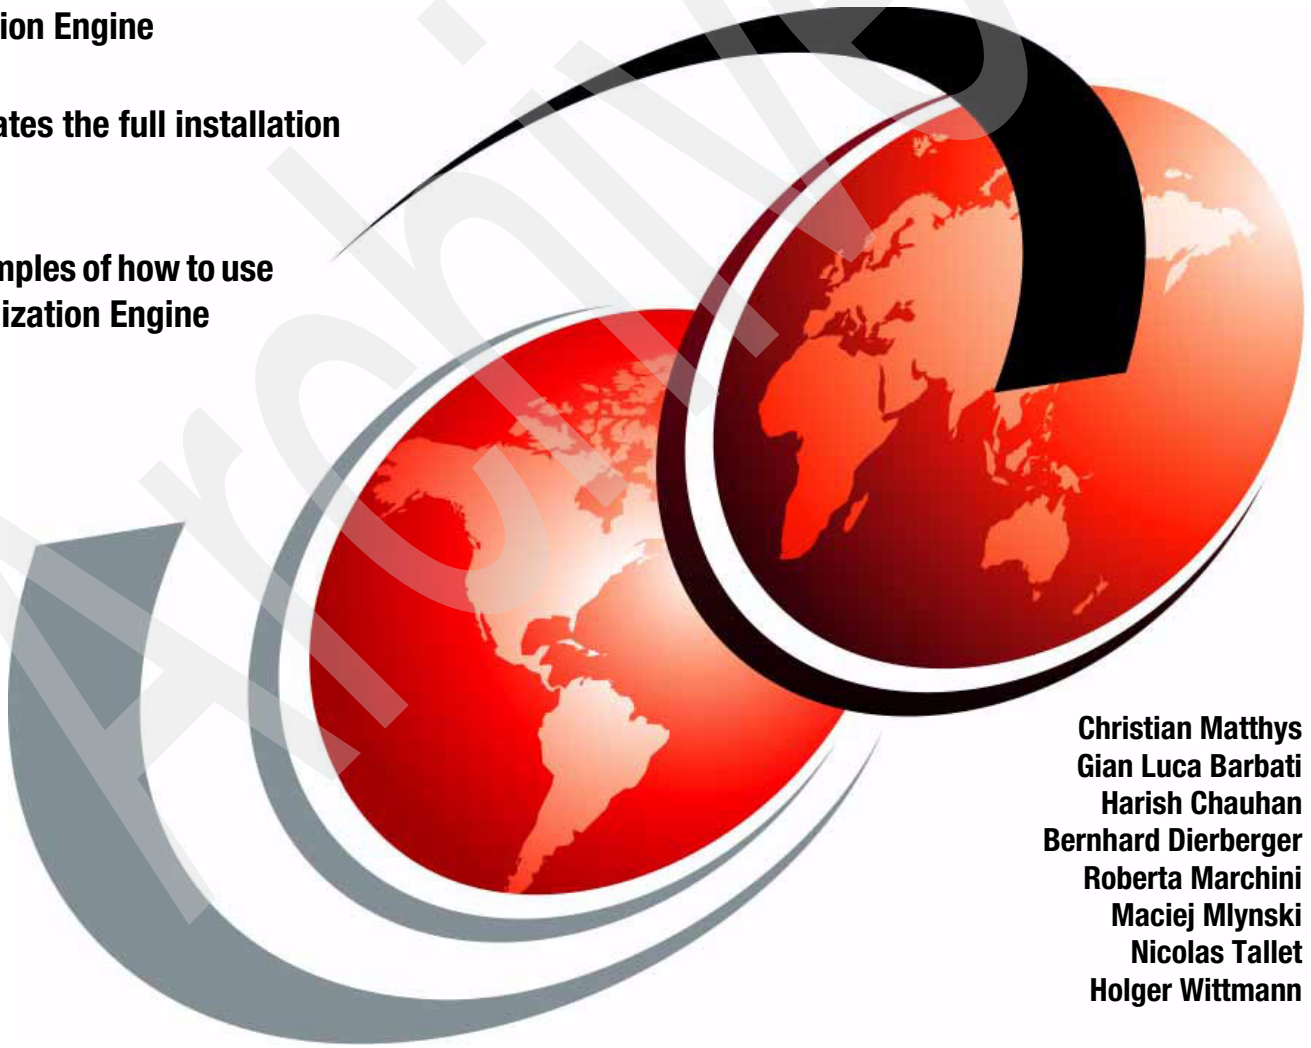
IBM

Planning, Installing, and Using the IBM Virtualization Engine Version 2.1

- Describes the components of the Virtualization Engine
- Demonstrates the full installation process
- Gives examples of how to use the Virtualization Engine



Christian Matthys
Gian Luca Barbati
Harish Chauhan
Bernhard Dierberger
Roberta Marchini
Maciej Mlynski
Nicolas Tallet
Holger Wittmann

Redbooks

ibm.com/redbooks

Archived

Archived



International Technical Support Organization

Planning, Installing, and Using the IBM Virtualization Engine Version 2.1

September 2006

Archived

SG24-7213-00

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

Archived

First Edition (September 2006)

This edition applies to the IBM Virtualization Engine Version 2.1 and its related components, namely the Virtualization Engine console, the IBM Director 5.10, Enterprise Workload Management V2.1, and Resource Dependency Service products (5648-F08,5655-EWM, 5733-DIR, 5733-EWM, and 5733-RDS).

© Copyright International Business Machines Corporation 2006. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices ix
Trademarks x

Preface xi
The team that wrote this IBM Redbook xi
Become a published author xiv
Comments welcome. xiv

Part 1. Planning and installing the Virtualization Engine Version 2.1 1

Chapter 1. Introducing the Virtualization Engine Version 2.1 3
1.1 The IBM Virtualization Engine. 4
1.1.1 The IBM Virtualization Engine platform. 4
1.1.2 New features in Version 2.1 8
1.2 Importance of standards 10
1.3 The vision: Service-oriented architecture 13
1.4 The Virtualization Engine roadmap 14
1.4.1 Services for managing the IT environment 15
1.4.2 The IBM Systems Agenda 15
1.5 The value of virtualization 16

Chapter 2. Planning 19
2.1 Terms used in the IBM Virtualization Engine 21
2.1.1 Components and common runtime components 21
2.1.2 Introducing the Virtualization Engine environment 25
2.1.3 The products. 25
2.1.4 The packaging 32
2.2 Planning tools 33
2.2.1 The Information Center. 34
2.2.2 The planning advisor. 35
2.2.3 Planning worksheets. 44
2.3 Prerequisites 50
2.3.1 AIX on IBM System p 51
2.3.2 Linux on POWER 54
2.3.3 Linux on IBM System x 56
2.3.4 Microsoft Windows on IBM System x / BladeCenter. 59

Chapter 3. Installing management servers. 63
3.1 The ITSO environments 65
3.2 Preinstallation tasks 67
3.2.1 Preparing the installation files. 67
3.2.2 Creating the operating system users and groups 69
3.2.3 General considerations. 71
3.2.4 Platform-specific considerations 72
3.2.5 Considerations before installing Resource Dependency Service 74
3.3 Installing the IBM Director. 77
3.3.1 Graphical installation process. 77
3.3.2 Silent installation process 85
3.4 Using the Virtualization Engine installer 88

© Copyright IBM Corp. 2006. All rights reserved.

3.4.1 Process overview . . . 88
3.4.2 The wizard steps . . . 89
3.5 Silent installation . . . 124
3.5.1 Virtualization Engine common components . . . 124
3.5.2 Virtualization Engine Services . . . 126
3.6 Post-installation checking tasks . . . 128
3.6.1 About the Directory Server . . . 128
3.6.2 EWLM Domain Manager . . . 129
3.6.3 IBM Director Server . . . 143
3.6.4 Resource Dependency Service . . . 145
3.6.5 Virtualization Engine console . . . 149
3.6.6 Bridges to Virtualization Engine console . . . 161

Chapter 4. Installing managed servers agents . . . 163
4.1 Preinstallation . . . 164
4.1.1 Enterprise Workload Manager agent . . . 164
4.1.2 IBM Director Agent . . . 164
4.2 Installation process . . . 165
4.2.1 Using the IBM Director software distribution . . . 165
4.2.2 Using the Remote Package Deployer (RPD) . . . 165
4.2.3 Manual installation . . . 167
4.3 Post-installation tasks . . . 182
4.3.1 EWLM Managed Server . . . 182
4.3.2 Director Agent . . . 186

Chapter 5. Applying fixes . . . 189
5.1 About fixes . . . 191
5.1.1 Terminology . . . 191
5.1.2 The Web site . . . 191
5.1.3 Verify the fix level of your environment . . . 194
5.2 About registration . . . 196
5.3 Installing and using the Update Installer tool . . . 199
5.4 Installing the updates . . . 200
5.4.1 Fix installation using the Virtualization Engine console . . . 200
5.4.2 Installing fixes without using the Virtualization Engine console . . . 201
5.5 Installing fixes for specific products . . . 205
5.5.1 Installing the DB2 fix . . . 205
5.5.2 Installing the IBM Director fix . . . 206
5.6 Uninstalling fixes . . . 207

Chapter 6. Lessons learned . . . 209

Part 2. Using the Virtualization Engine Version 2.1 . . . 213

Chapter 7. Using Resource Dependency Service to build manageable resources . . . 215
7.1 Resource Dependency Service concepts . . . 216
7.1.1 Resource Dependency Service components overview . . . 217
7.1.2 Understanding the Intelligent Device Discovery tool . . . 218
7.1.3 Understanding the Resource Model Services . . . 221
7.1.4 Understanding the Resource Modelling Framework . . . 222
7.2 Configuring Resource Dependency Service . . . 224
7.3 Planning and running IDD . . . 226
7.3.1 Planning the IDD campaign . . . 226
7.3.2 ITSO environment . . . 227

7.3.3 Running an IDD campaign 229
7.4 Planning and configuration of Resource Dependency Service data collectors 243
7.4.1 Where to look for the collectors. 244
7.4.2 Configuring a new collector. 245
7.4.3 Executing a collector. 251
7.4.4 Configuring and running the IBM Director collector 252
7.5 Managing Resource Dependency Service 257
7.5.1 Intelligent Device Discovery 257
7.5.2 Virtualization Engine management server 261
7.5.3 Integrated Solution Console 263
7.5.4 DB2 263
7.6 Troubleshooting 263

Chapter 8. Using the Virtualization Engine console to visualize the infrastructure environment 265
8.1 Accessing resource topology using Virtualization Engine console 266
8.2 The Virtualization Engine environment 270
8.3 The Managed Resources environment 275
8.3.1 What is a resource? 276
8.3.2 What is a filter? 278
8.3.3 Registering a new resource 280
8.3.4 What is a group? 282
8.3.5 What is a relationship? 285
8.4 Scenarios 287
8.4.1 Adding a new database in the system 287
8.4.2 Adding new resources in the environment 291
8.4.3 Using topology viewer for removing the resources 293
8.4.4 Interpreting the information shown in the topology viewer 297

Chapter 9. Using one single console to manage your environment 301
9.1 The environment 302
9.1.1 Management servers 302
9.1.2 Managed nodes 303
9.2 Configuring the environment 304
9.2.1 Configuring the management sources 304
9.2.2 Configuring the Hardware Management Console 309
9.2.3 Discovering systems with IBM Director. 312
9.2.4 Requesting access to systems in IBM Director. 315
9.2.5 Renaming BladeCenter. 317
9.2.6 Grouping resources 321
9.2.7 Changing the IBM Director Bridge configuration. 324
9.3 Diagnosing a hardware error. 330
9.3.1 Configuring resources in the Health Center 331
9.3.2 Seeing the notification and acting on the target system 334
9.4 Processing error 340
9.4.1 Process Monitor configuration in IBM Director 341
9.4.2 Adding the Process Monitor to the Health Center. 345
9.4.3 Adding the monitor to the watch list 347
9.4.4 Restarting the nfsd process 348
9.5 Monitoring the system CPU usage 350
9.5.1 Configuring IBM Director to monitor the CPU 351
9.5.2 Adding the resource monitor to the Health Center 353
9.5.3 Adding the CPU monitor to the dashboard 356
9.6 Managing a VMWare Virtual Machine. 358

Archived

9.6.1 Requesting access to the ESX host 360
9.6.2 Adding the resources to Health Center and the Watch List portlet 362
9.6.3 Configuring the Event Action plan 364
9.6.4 Visualizing in the Virtualization Engine console 370
9.7 Monitoring an LPAR 372
9.8 Starting and stopping services 375
9.8.1 Virtualization Engine console 375
9.8.2 IBM Director Bridge 375
9.8.3 IBM Director Server 377

Chapter 10. Using the Enterprise Workload Manager 379
10.1 Load Balancing with the CISCO Content Switching Module, SASP, and the IBM Enterprise Workload Manager 380
10.2 Performance monitoring using Enterprise Workload Manager 380
10.3 Monitoring, reporting, and partition management using Enterprise Workload Manager . 380

Part 3. Appendixes 383

Appendix A. Performance monitoring using Enterprise Workload Manager 385
Introduction 386
Prerequisites 387
Enabling ARM platform support 388
For AIX 388
For Linux 389
Creating an EWLM DM configuration ID 389
Using the configuration wizard 389
Using the command line 392
Troubleshooting 393
Creating an EWLM MS configuration ID 394
Using the configuration wizard 394
Using the command line 395
Troubleshooting 395
Creating an EWLM domain policy 395
Enabling ARM instrumentation in DB2 UDB 399
Troubleshooting 400
Enabling ARM instrumentation in Application Server 402
Troubleshooting 405
Enabling ARM instrumentation in IHS 406
Troubleshooting 407
Viewing transaction statistics using the Control Center 407
Conclusion 410

Abbreviations and acronyms 411

Glossary 415

Related publications 421
IBM Redbooks 421
Other publications 421
Online resources 421
How to get IBM Redbooks 423
Help from IBM 424

Index 425

Archived

Archived

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

© Copyright IBM Corp. 2006. All rights reserved.

Trademarks

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

AIX 5L™
AIX®
BladeCenter®
Cloudscape™
Collation®
CICS®
DB2 Universal Database™
DB2®
Enterprise Storage Server®
eServer™
HACMP™
i5/OS®
Intelligent Device Discovery®
IBM®
IMS™
iSeries™
Lotus Notes®
Lotus®
Micro-Partitioning™
MVS™
Notes®
OpenPower™
Processor Resource/Systems Manager™
POWER™
POWER Hypervisor™
POWER5™
PR/SM™
Redbooks™
Redbooks (logo) ™
RMF™
ServeRAID™
System i™
System p™
System p5™
System x™
System z™
Tivoli Management Environment®
Tivoli®
TotalStorage®
TME®
Virtualization Engine™
WebSphere®
z9™
z/OS®
z/VM®
zSeries®

The following terms are trademarks of other companies:

Enterprise JavaBeans, EJB, Java, JavaBeans, JavaServer, JavaServer Pages, JDBC, JMX, JRE, JSP, JVM, J2EE, J2SE, Power Management, Solaris, Sun, Sun Microsystems, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries 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.

Archived

Preface

The On Demand Operating Environment is the computing environment that is required for building and managing an On Demand Business. To become an On Demand Business, a company needs:

- Componentization of business processes to create the flexibility and innovation at the business level
- Applications modularity to support the business processes
- IT infrastructure simplification in order to support the changes the business needs

The concepts of automation and virtualization are fundamental to IT infrastructure simplification. *Virtualization* is the ability to separate the direct dependency of an application from a physical resource.

To address the virtualization aspects of the On Demand Operating Environment, the IBM® Virtualization Engine™ was announced in August 2004. Version 1 was then available. A new version, the Virtualization Engine Version 2.1, was announced and made available at the end of 2005. This new version enhances the functions provided in Version 1 and provides new products, new features, and new tools.

This IBM Redbook does not describe the components of the Virtualization Engine V2.1 in-depth. Read its companion IBM Redbook, *IBM Virtualization Engine Platform Version 2 Technical Presentation Guide*, SG24-7112, which has details about these components, before you read this book.

This IBM Redbook describes how to install and to use the Virtualization Engine V2.1 and has two parts:

- Part 1, “Planning and installing the Virtualization Engine Version 2.1” on page 1 helps you install and configure the Virtualization Engine V2.1 and all its products in different environments.
- Part 2, “Using the Virtualization Engine Version 2.1” on page 213 provides you with examples to help you understand how to use the products and the functions provided in this release.

This IBM Redbook is suitable for IT Specialists who want to install and to use the Virtualization Engine V2.1 products and components.

The team that wrote this IBM Redbook

This IBM Redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Christian Matthys has spent more than 25 years at IBM as a System Engineer, working with large mainframe-oriented clients. He spent three years as an ITSO project leader on assignment in Poughkeepsie, NY. In 2000, he joined the EMEA Design Center for On Demand Business, working with clients’ projects to make use of the leading edge technologies, in particular, the Autonomic Computing technologies. He was part of the launch team of client workshops to promote the IBM Virtualization Engine. He works as an ITSO project leader based in the Product and Solutions Support Center (PSSC) in Montpellier, France.

© Copyright IBM Corp. 2006. All rights reserved.

Gian Luca Barbati is a Certified IT Specialist based in Milan. He has been with IBM since 1989 when he joined the IBM System i™ support center. He then moved to the service area ITS to deliver services for the IBM System i platform. He has had a lot of experience in clients' projects and critical situation support outside Italy. Five years ago, he moved to the Linux® platform, and is currently architecturing and implementing Linux solutions in the IBM System x™ and IBM System i areas. He has been involved in many client projects and has extensive experience in IT management.

Harish Chauhan is a Senior Advisory IT Specialist with over 14 years of experience in IT, including eight years in IBM India. He has been working on Linux for the last four years. Prior to that, he has also worked on AIX®. His area of interests are Linux, clusters, high-performance computing, benchmarking, performance tuning, technical workshops, and solution implementation. He currently manages the IBM Linux Center of Competency based in Bangalore, India. He has published technical articles on IBM developerWorks. He is a member of the IBM worldwide black belt technical community for the IBM Virtualization Engine.

Bernhard Dierberger is a certified IT Specialist at the IBM Technical Marketing Competence Center (TMCC) in Boeblingen, Germany. He holds a PhD in physics from the University of Tuebingen, Germany. Before joining the IBM lab in Boeblingen, Bernhard was a technical consultant for IBM Global Services, Germany. From the beginning of his work at the IBM development lab, he has focused on Systems Management, J2EE™, and WebSphere® technologies. In 2001, he joined the TMCC Proof of Concept team. He is the technical focal point of Grid, On Demand, and Virtualization Engine client projects at the TMCC. He is a member of the IBM worldwide black belt team for the IBM Virtualization Engine.

Roberta Marchini is an IBM Certified IT Specialist working in the IBM System x Presales Technical Support team in Italy. She holds a degree in Electronic Engineering and has nine years of experience in Information Technology. She has been working in the Intel/AMD based server area for nine years, focusing on systems management since the beginning, both from the hardware and the software point of view. She has been a speaker on this subject at several IBM European events, such as the Systems and Technology Group University and the IBM System x Technical Symposium. She also holds the Microsoft® Certified Systems Engineer (MCSE) certification. In the last few years, she has been dealing with complex server consolidation projects involving both blade server based architectures and virtualization technologies, such as VMWare ESX and Microsoft Virtual Server.

Maciej Mlynski is an IBM Certified Advanced Technical Expert - IBM System p™ and AIX 5L™. He works as an IT Specialist for ASpartner Sp. z o.o., an IBM Premier Business Partner in Poland. He has nine years of experience in IBM System p and AIX. His areas of expertise include storage area networks, and highly-loaded and missing-critical computer systems. He holds a Master's Degree in Computer Science from the Technical University of Gliwice in Poland. He is presently involved in PhD research. One of his topics of interest is dynamic resources allocation.

Nicolas Tallet is an IT Specialist in Virtualization solutions at the Products & Solutions Support Center (PSSC) in Montpellier, France. He started his career in 1998 as a technical consultant within IBM Global Services, based in Paris, France. After six years of working on several projects dealing with SAP/Supply Chain Management (SCM) solutions implementation for various clients, he joined the European Virtualization Team hosted within the PSSC in December 2004. Throughout 2005, the team was involved in the Start Your Engines (SYE) education program, delivering Virtualization Engine technical workshop sessions throughout Europe to clients, business partners, and IBMers.

Holger Wittmann is a certified IT Architect, based in Stuttgart, Germany. He started working for IBM in 1991 in the scientific center in Heidelberg, Germany, building mobile radio software

solutions. In 1998, he joined IBM Global Service to build Java™ and Lotus® Notes® solutions in various client projects. Three years later, he took an opportunity in Sales & Distribution to support clients and major Independent Software Vendors (ISVs) on an architectural level in the arena of e-procurement and supply chain management solutions. In 2002, he joined the IBM Grid team to work with clients implementing grid solutions and then started working on Virtualization Engine. Now he is a member of the IBM worldwide black belt technical community to support virtualization in general. He is also a teacher and a developer of courseware for the technical community on On Demand Business, with an emphasis on business architecture comprising business transformation aspects with IT architecture.



Figure 1 The team who wrote this book, from left to right: Nicolas, Maciej, Gian Luca, Christian, Bernhard, Roberta, Holger and Harish

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

Jeremy Bauer	IBM USA, Rochester, MN
David Bennin	IBM USA, Poughkeepsie, NY
Boas Betzler	IBM USA, Poughkeepsie, NY
Amy Bildzok	IBM USA, Poughkeepsie, NY
Khai Q Bui	IBM USA, Raleigh, NC
Mike Buzzetti	IBM USA, Poughkeepsie, NY
Chekim Chhuor	IBM USA, Poughkeepsie, NY
Rich Conway	IBM USA, Poughkeepsie, NY
Jim Fall	IBM USA, Rochester, MN
Siavash Kazemian	IBM USA, Poughkeepsie, NY
Madeline Nick	IBM USA, Poughkeepsie, NY
Bob Padzieski	IBM USA, Rochester, MN
Eliuth Pomar	IBM USA, Poughkeepsie, NY
Susan Reuland	IBM USA, Rochester, MN
Daniel Riddering	IBM USA, Rochester, MN
Paul Sutera	IBM USA, Poughkeepsie, NY
Jeff Thelen	IBM USA, Rochester, MN
Scott Urness	IBM USA, Rochester, MN
Federico Vagnini	IBM Italy

Jeff Van Heuklon	IBM USA, Rochester, MN
Wade Wallace	IBM USA, Austin, TX
Joe Winkelbauer	IBM USA, Poughkeepsie, NY
Lynn Winkelbauer	IBM USA, Poughkeepsie, NY
Jim Zimmerman	IBM USA, Rochester, MN

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and clients.

Your efforts will help increase product acceptance and client satisfaction. As a bonus, you'll 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 IBM Redbooks™ to be as helpful as possible. Send us your comments about this or other IBM Redbooks in one of the following ways:

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbook@us.ibm.com

- Mail your comments to:

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

Archived

Part 1

Planning and installing the Virtualization Engine Version 2.1

In this part of this IBM Redbook, we describe the tasks and the steps you have to perform to successfully install the components of the Virtualization Engine Version 2.1. The installation is quite simple and more integrated than the Virtualization Engine Version 1. There are many options, and you have to decide which components you want and on which servers you want to install them before starting the installation process. Planning and installation involves:

- Planning the installation, which means understanding the Virtualization Engine components and its packaging, potentially using the tools provided by IBM and looking at the prerequisites in terms of product levels, performance, and security issues.
- Performing a few pre-installation tasks that may depend on the installation platform before starting the installation of the management servers or the managed servers agents. Post-installation tasks will depend on the product you have installed. An installation wizard is provided to facilitate the installation.
- Installing fixes, usually after the products installation. A new process with the Virtualization Engine V2.1 allows you to look at and install the fixes using the Virtualization Engine console.

© Copyright IBM Corp. 2006. All rights reserved.

Figure 2 shows the four major tasks for the complete installation process.

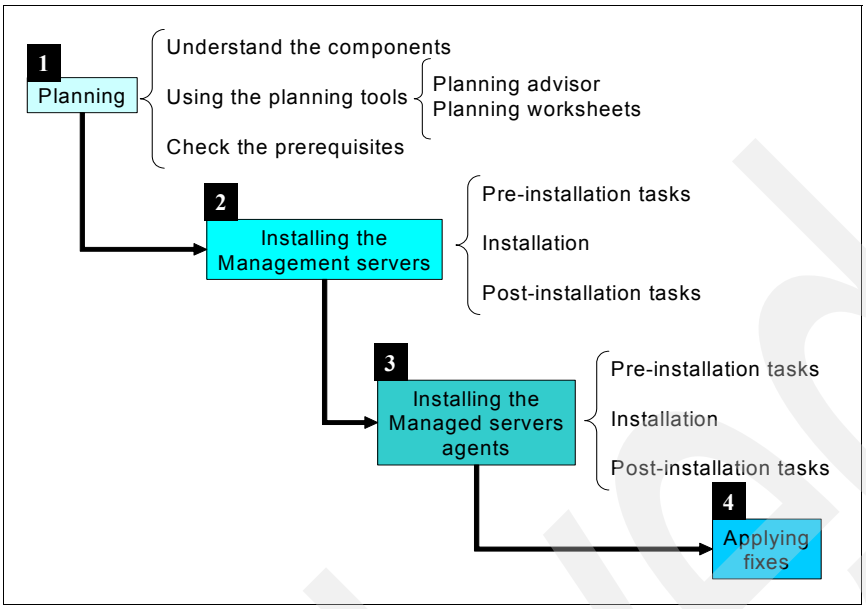


Figure 2 Installation tasks summary

The four main installation tasks are:

- The planning task, described in Chapter 2, “Planning” on page 19.
- The management servers installation tasks, described in Chapter 3, “Installing management servers” on page 63.
- The installation tasks of the agents on the managed servers, described in Chapter 4, “Installing managed servers agents” on page 163.
- The tasks to apply the fixes, described in Chapter 5, “Applying fixes” on page 189.

Before going through the tasks, Chapter 1, “Introducing the Virtualization Engine Version 2.1” on page 3 explains what the Virtualization Engine V2.1 is, and its value in managing a heterogeneous infrastructure environment.

The last chapter of this part of the book, Chapter 6, “Lessons learned” on page 209, summarizes hints and tips from our experience with the Virtualization Engine V2.1.

Archived

1

Introducing the Virtualization Engine Version 2.1

A first version of the IBM Virtualization Engine has been available since August 2004 and a new version, called the Virtualization Engine Version 2.1, has been available since the end of 2005 for most of its components.

This chapter summarizes what is new in version 2.1. For a more complete description, refer to the following:

- *IBM Virtualization Engine Platform Version 2 Technical Presentation Guide*, SG24-7112 available on the Web at:

  `http://www.redbooks.ibm.com/abstracts/sg247112.html?Open`

- The InfoCenter on the Web at:

  `http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp`

If you are not familiar with the Virtualization Engine Version 2.1, we recommend you read the IBM Redbook and the information in the InfoCenter. This book is about the usage of the products, not a description of the products.

As it is very important to understand the role of virtualization and how it fits in the IBM On Demand Operating Environment delivery, this chapter also describes the role of the standards in this arena and what we can expect from them, and presents virtualization in perspective. The Virtualization Engine will allow the Services-Oriented Infrastructure implementation, without which the benefits of the service-oriented architecture will be limited.

In this chapter:

- 1.1, “The IBM Virtualization Engine” on page 4 describes the IBM Virtualization Engine Version 2.1.
- The other four sections describe the key concepts to understanding the importance of the virtualization in the coming information technology (IT) environment.

© Copyright IBM Corp. 2006. All rights reserved.

1.1 The IBM Virtualization Engine

The traditional view of virtualization focuses on partitioning. A large box is used by multiple clients and each of them feels they are the only ones using the box. Until now, virtualization has been used to increase hardware utilization.

Now IBM expands virtualization capabilities to allow multiple systems to come together and perform together in a larger context and give you a common way to address component-level devices. Virtualization will play a major role in simplifying IT management. In particular, it will be extended from single servers to scale-out aggregations of servers, storage, and network components, making many small systems look like one single system from a management perspective.

Definition: A virtualized IT environment is a heterogeneous IT infrastructure, dynamically provisioned, managed by business rules, and where the work is monitored and scheduled as it flows through all the layers of the environment, from the network, through firewalls, Web servers, application servers, database servers, and back. These layers may be supported by different types of operating systems.

Virtualization allows data or computing resources to be logically represented and moved to where they are most needed and when they are most needed. This allows users to get better service and helps IT respond more effectively to business goals. Virtualization can be viewed as the creation of substitutes for real resources. These substitutes have the same functions and external interfaces as their real counterparts. However, they can differ in attributes, such as size or performance, and their users are usually unaware of the substitution.

The IBM Virtualization Engine Version 1 has been available since 2004. IBM is using this new version to take the next steps in virtualization by transforming the Virtualization Engine into an integrated platform of technologies that broadens and enhances the current capabilities. The IBM Virtualization Engine Version 2.1 has been available since December 2005.

1.1.1 The IBM Virtualization Engine platform

The IBM Virtualization Engine platform, as depicted in Figure 1-1 on page 5, is a comprehensive platform made up of IBM systems, servers and storage, with sophisticated virtualization technologies. In this platform view:

- These systems along with their virtualization technologies are referred to as *Virtual Resources*.
- *Virtual Access* and *Virtual Management* capabilities with the IBM TotalStorage® Productivity Center, the Virtualization Engine console, the IBM Director, Enterprise

Archive

Workload Manager, and Resource Dependency Service, offer the capability to view and manage multiple systems as one, with a consistent set of tools.

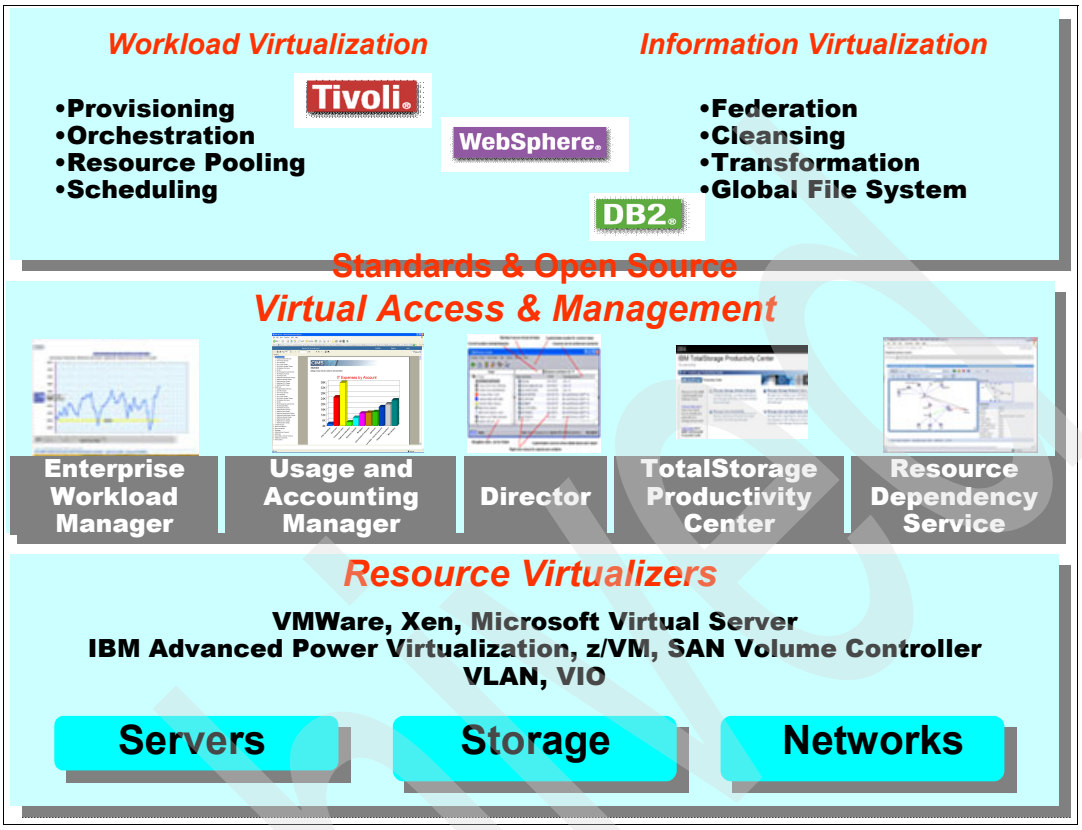


Figure 1-1 The Virtualization Engine platform

Virtual resources

The Virtualization Engine platform is based on virtual resources that are included in all IBM server and storage systems, as well as select server and storage systems from other vendors, such as Hewlett Packard (HP), Sun™, and EMC. These embedded technologies allow resource virtualizers to facilitate the virtualization of resources within a server or storage device. They prepare the virtual resources for the management and performance capabilities provided by the Virtualization Engine management collection.

They are the primary technologies to maximize infrastructure availability and to make better use of resource utilization. Usually these technologies are the starting point to consolidate the components of the infrastructure. They cover physical resource virtualization technologies as well as operating system virtualization technologies.

Physical resource virtualization

Physical resource virtualization technologies can be found at different infrastructure layers including the server layer, the storage layer, and the network layers:

- *Server virtualization* allows a single server to be used by multiple applications, middleware, and operating systems concurrently. Some of the technologies have been available for years though they continue to evolve. They include virtual memory, physical partitioning, logical partitioning, I/O virtualization, and so on. The Integrated Virtualization Manager (IVM), the Virtual I/O Server (VIOS), and the Hosting Partition Infrastructure are examples of the recent evolution of these technologies.

- *Storage virtualization* enables clients to make changes in the storage infrastructure with no or little impact on applications. This is particularly important with regard to the increase in storage needs. The storage virtualization technologies cover all the storage infrastructure levels. They include technologies at the disk controller level, at the block level, at the file level, and at the tape level.
- *Network virtualization* allows pooling and sharing of network components to make communication across IT infrastructures more efficient and more secure. The network virtualization technologies include features like the Open Systems Adapter Express (OSA) cards, virtual Ethernet, and Virtual IP Addressing (VIPA) facilities.

Operating system virtualization

Two virtualization technologies are mainly used, depending on the platform:

- The z/VM® system, available on the IBM System z™ for more than 30 years, is very dynamic. It allows real and virtual resources to be shared and switched dynamically without outages between virtual machines. Several hundred virtual machines can be hosted on a single system.

  For more information about the z/VM and Linux, refer to *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9*, SG24-6695.
- VMWare provides a set of tools, very similar to the z/VM facilities, for partitioning, consolidating, and managing IBM System x and IBM BladeCenters. Though the idea is the same as z/VM, given the difference in the size of the target box, the facilities are not totally equivalent. You can find more information about VMware on the Web at:

  `http://www.vmware.com/`

Virtual management

The Virtualization Engine management focus is providing workload and performance managers, and resource management tools in relation to the resource virtualizers. The goal is to provide common tools across the heterogeneous infrastructure.

These tools include the following products:

- The *IBM Director* provides a common, consistent solution for managing hardware and virtual partitions, not only for the IBM System x platform support, but also for the hardware configuration of systems that can be remotely discovered, monitored, and automated. It can be bridged with the Virtualization Engine console, offering one single console focal point for administrators.
- The *Resource Dependency Service* (*RDS*) provides a logical representation of hardware and software resources using a model based on open standards using Web services.
- *Enterprise Workload Manager* (*EWLM*) provides a consolidated end-to-end view of transactions performance and supports dynamic adjustments to CPU allocations across logical partitions (LPARs) for POWER5™ servers.
- *IBM Tivoli® Usage and Accounting Manager* (called *TUAM* or *UAM*) is a new product added in the Virtualization Engine components after IBM bought the CIMS Labs company at the end of January 2006. UAM is a cross-platform software system that determines the cost of providing information services. It manages IT costs by appropriating the costs of an organization's products services to business functions. UAM measures, analyzes, reports, and bills the utilization and costs of different computing resources.

The objective is to deliver all the functions on all the IBM eServer™ platforms[1]. Today, the following functions are available:

- For IBM System i:
  - EWLM domain manager and managed nodes

    EWLM is able to manage LPARs based on business goals.
  - IBM Director
  - Resource Dependency Service for Linux
  - The Virtualization Engine console
- For IBM System p:
  - EWLM domain manager and managed nodes

    EWLM is able to manage LPARs based on business goals.
  - IBM Director
  - Resource Dependency Service for Linux
  - The Integrated Virtualization Manager (IVM) on select systems
  - The Virtualization Engine console
- For IBM System x and BladeCenter®:
  - EWLM
  - IBM Director
  - Resource Dependency Service for Linux
  - Virtual Machine Manager
- For IBM System z:
  - EWLM domain manager on z/OS® and Linux

    EWLM manages nodes for z/OS and Linux.
  - IBM Director with the Virtualization Engine console for Linux
  - Resource Dependency Service for Linux

Virtual access

The IBM Virtualization Engine's virtual access provides:

- Virtual view

  The virtual view includes a new version of the Virtualization Engine console, which is a single Web-based interface for administering all virtual resources, such as resource health, monitoring, and integration access to planning tools. In other words, the Virtualization Engine console provides a single portal for administering the virtual resources; it allows a consolidated view of the performance, health, and topology across multiple heterogeneous operating systems.

  Though not provided as a component of the Virtualization Engine, IBM TotalStorage Productivity Center (TPC) can be put at this virtualization level. TPC is a suite of infrastructure management software that can centralize, automate, and simplify the management of complex and heterogeneous storage environments. It can help reduce the effort of managing complex storage infrastructures, improve storage capacity utilization, and improve administration efficiency. TPC allows you to respond to on demand storage

[1] IBM plans are subject to change without notice and this information represent goals and objective only. Such information is not intended as a definitive statement of a commitment to specific levels of function or delivery schedules.

Archived

needs and brings together, in a single point, the management of storage devices, fabric, and data.

- Programmatic interface

  The programmatic interface will offer the ability to access system resources as services; it is not externalized today. This allows access to the system infrastructure through a service-oriented architecture programming model. The key programmatic interfaces used in the Virtualization Engine platform today provide standards at the following levels:

  – Web services access to a resource
  – Mapping of Web services to the resource model
  – The resource model

Virtual planning tools

The virtual planning tools include tools that help you plan and build your virtualized environment. Examples of these tools include such tools as:

- The *IBM eServer Workload Estimator* provides sizing recommendations for a number of workloads for the IBM System i, IBM System p, and IBM System x. You can access the Workload Estimator on the Web at:

  `http://www.ibm.com/eserver/iseries/support/estimator`

- The *Virtualization Engine planning advisor* is an interactive tool that produces a customized plan based on your input. For more information about the planning advisor, refer to 2.2.2, “The planning advisor” on page 35.

1.1.2 New features in Version 2.1

The IBM Virtualization Engine Version 2 is a follow-up to the IBM Virtualization Engine Version 1, which has been available since August 2004.

Since August 2004, IBM had two offerings:

- The *IBM Virtualization Engine for Systems*, which focuses on utilizing the technologies integrated into the IBM eServer systems, such as Hypervisor, VLAN, Virtual I/O, and so on.
- The *IBM Virtualization Engine for Enterprise*, which provides a collection of multiplatform system tools to help simplify the management and utilization of IT resources, such as Enterprise Workload Manager, Tivoli Provisioning Manager, Director Multiplatform, Virtualization Engine console, and the IBM Grid Toolbox V3 for Multiplatforms.

In November 2005, IBM announced the IBM Virtualization Engine Version 2, which is an evolution of the two previous offerings (sometimes they were called *suites*). The following products are now delivered:

- Enterprise WorkLoad Manager (EWLM) Version 2.1

  This is a new version of EWLM that delivers new operating systems support and workload management on IBM System p systems.

- IBM Director Version 5.10

  This is a new version of IBM Director that integrates functions of IBM Director and IBM Director Multiplatform.

- Resource Dependency Service

  This is a new component of the IBM Virtualization Engine.

- IBMTivoli Usage and Accounting Manager (TUAM or UAM)

  This is a new product added in the Virtualization Engine components after IBM bought the CIMS Labs company at the end of January 2006.

Note that the following components, while not delivered with the IBM Virtualization Engine Version 2, deliver new virtualization facilities:

- Integrated Virtualization Manager (IVM)

  This provides LPAR support on IBM System p without a physical Hardware Management Console (HMC). With the IVM, IBM is simplifying the p5 system's virtualization deployment. Using a new wizard-driven, browser-based user interface, clients can create a micro-partition ready for installation on either the AIX 5L or Linux operating system. IVM comes packaged at no additional cost with the optional purchase of the Advanced POWER™ Virtualization (APV) feature on IBM System p5™ servers. IVM is designed to reduce the complexity and time involved to set up and configure multiple partition environments, using IBM Micro-Partitioning™ capabilities, in small environments.

- The TS7510 Virtualization Engine

  This is a high-performance, high-capacity open systems virtual tape product designed to augment the tape backup and restore process in large tape environments. The TS7510 is designed to aid in the tape backup and restore process by helping to provide clients significant operational and throughput efficiencies via tape virtualization. The Virtualization Engine TS7510 is a new member of the family of IBM Virtualization Engine for storage offerings, and is designed to help enhance Information Lifecycle Management by improving the utilization of tape resources. The TS7510 Virtualization Engine combines hardware and software into an integrated solution enabling tape virtualization for open systems servers connecting over Fibre Channel physical connections.

  The Virtualization Engine TS7510 Software V1.1 (5639-CC7) provides the software for the Virtualization Engine TS7510 and for a management console. The TS7510 utilizes IBM technology and software to emulate IBM tape libraries, drives, and media. This emulation is designed to help clients achieve the following operational and throughput efficiencies:

  – Backup window reduction
  – Restore time reduction
  – Data sharing and resource virtualization facilitation
  – Operational efficiencies
  – Low total cost of ownership (TCO)
  – Improved sharing of tape libraries across applications and servers

  For more information about TS7510, refer to *IBM Virtualization Engine TS7510: Tape Virtualization for Open Systems Servers*, SG24-7189.

- The IBM TotalStorage storage area network (SAN) Volume Controller (SVC)

  This is a key component in the IBM virtualization storage technologies. At the heart of the SAN Volume Controller's many benefits is its ability to consolidate devices from various suppliers, in different locations, into a single view of storage resource.

  For more information about SVC, refer to *IBM TotalStorage SAN Volume Controller*, SG24-6423.

IBM has a partnership with Network Appliance (NetApp). This partnership aims to promote the integration of Network Appliance's software applications with the IBM Tivoli Storage Manager product line, a backup and recovery package designed specifically for use with this partner's products. This package provides unified network-attached storage (NAS) and SAN access to data stored in IBM storage arrays and supports connectivity to IBM Enterprise Storage Server® (ESS) product line. You can find more information about NetApp on the Web at:

http://www.netapp.com/

Tivoli Provisioning Manager (TPM) and the IBM Grid Toolbox V3 for Multiplatforms are not delivered with the Virtualization Engine platform Version 2.1.

TPM, a key part of the On Demand Operating Environment (ODOE) architecture, will interact with the virtualization components. It is distributed through the IBM Software Tivoli brand.

You can download the IBM Grid Tool box from the Web at:

https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=grid

1.2 Importance of standards

One of the most important insights of the On Demand Operating Environment is that a *common pattern* supports both application services and infrastructure services, making it possible for applications to interact with the infrastructure components. For example, a provisioning application will be able to “dialog” with a hardware component to understand its state, its properties, how it works, and maybe to develop an application on it.

Standards are key to allow different types of components to understand each other in a heterogeneous environment. Standards are evolving with time and depending on needs. This is important to understand where we are and what we need.

Web services

Everything starts with Web services. A *Web service* is a software component designed to support interoperable machine-to-machine interaction over a network. Unlike typical applications, Web services describe themselves to the outside world by explaining what functions they perform, how they can be accessed, and what kinds of data they require. A Web service is *stateless*, meaning that it is a service whose implementation maintains no dynamic data, but which acts upon resources based on messages it sends and receives. The resources themselves may be *stateful*. A Web service will not maintain a dynamic state. This means that it will not maintain a state for which the service is responsible between message exchanges with its requestors. Operations on stateless Web services imply that the necessary parameters are passed on to each Web service request or invocation. A Web service is a program component deployed within some runtime environment that is responsible for executing the code (most often Java code) and for dispatching messages.

Web services are not standards. They *use* standards defined by independent companies. Web services are based today on the following standards:

- eXtensible Markup Language (XML)
- Simple Object Access Protocol (SOAP)
- Web Services Description Language (WSDL)
- Universal Description, Discovery, and Integration (UDDI)

The key point about Web services is that this piece of code is independent of the developer.

Providing stateful services through grid technologies

Grid technology was among the first to promote a virtual computing environment through the use of a set of services that represented the abstraction of heterogeneous resources. Grid technology needed a way to control the management of distributed and often *long-lived state* objects. The Open Grid Services Infrastructure (OGSI), developed by the Global Grid Forum (GGF), defines the conventions and specifications of the grid services, their protocols and interface behavior, properties, and attributes based on Web services specifications plus some extensions.

Note: The GGF is the community of users, developers, and vendors leading the global standardization effort for grid computing. For details about the GGF, refer to the following Web site:

`http://www.gridforum.org/`

OGSI is concerned primarily with creating, addressing, inspecting, and managing the lifetime of stateful grid services. One extension proposed by OGSI to the Web services community was the mechanism of *Service Data*, whose role is to expose the instance of a service to requestors, providing a *stateful* system. Stateful services allow service consumers to query and to manipulate service data that persist across service interactions.

Introducing message addressing

SOAP, the protocol used to access Web services, does not provide a standard way to specify where a message is going, how to return a response, or where to report an error. The WS-Addressing specification, defined in the WSDL 2.0 specification, defines a standard for incorporating message addressing information into Web services messages. This specification enables messaging systems to support message transmission through networks. In particular, this specification defines XML elements to identify Web service *endpoints*. A Web services endpoint is a resource where Web services messages can be targeted. Endpoint references convey the information needed to identify or reference a Web services endpoint. They can be used to provide addresses for individual messages sent to and from Web services. The only required element of an endpoint reference is the address.

Refactoring the interfaces

We were at a point where the users had too many separate but useful standards. The WS-Resource Framework (WS-RF) is a refactoring of some OGSI interfaces to produce a framework of independently useful Web services standards. WS-RF define terms, such as the *WS-Resource* approach to modeling and managing state in a Web services context. The WS-RF definition relies upon the WS-Addressing specification. The specifications of these interfaces retain all of the essential capabilities present in OGSI, while changing some of the syntax to exploit WS-Addressing and adopting a different terminology.

The WS-Resource construct is a means of expressing the relationship between stateful resources and Web services. The WS-RF goal allows stateful resources manipulation through Web services operations. WS-RF adopts the endpoint reference construct defined in the WS-Addressing specification as an XML syntax for identifying Web service endpoints. The creation of a WS-Resource in WS-RF terms is really the creation of a stateful Web service in OGSI terms.

Web Services Distributed Management: the state-of-the-art to manage Web services

Web Services Distributed Management (WSDM) is a recently approved Organization for the Advancement of Structured Information Standards (OASIS) standard (March 2005) to use Web services architecture and technology to manage distributed resource Web services (as well as non-Web services infrastructure).

Note: OASIS is “a not-for-profit, global consortium that drives the development, convergence, and adoption of e-business standards.” More information is available on the Web at:

`http://www.oasis-open.org`

WSDM allows a resource and all of its services to be manageable in a standard and interoperable manner. The WSDM specifications define:

- How to use Web services to expose manageable resources

  This is the Management Using Web Services (MUWS) specification.
- How to expose manageable Web service implementations

  This is the Management of Web Services (MOWS) specification.

To manage a Web services environment, the components (the Web services endpoints) have to be accessible and managed. A Web service endpoint provides access to a manageable resource. In order to discover the Web service endpoint providing access to a particular manageable resource, the requestor needs to obtain an *endpoint reference* (EPR) as defined by the WS-Addressing specification, and then obtain any other required descriptions, such as a WSDL document, an XML schema, or a policy document. A Web service endpoint providing access to some manageable resource is called a *manageability endpoint*. To exchange messages with a manageability endpoint, a manageability consumer needs to understand all of the required descriptions for the endpoint. The manageability consumer sends messages targeted to the manageable resource by using information contained in the EPR, for example, an address and some reference properties.

The WSDM specification defines how the manageability of Web service endpoints and resources exposed as Web services can be accessed through Web services.

Other standards involved

Other standards should be cited. These are:

- Open Services Gateway Initiative (OSGI)

  IBM is also involved in the OGSI through its Service Management Framework (SMF).

Note: The Open Services Gateway Initiative (OSGI) was founded in March 1999. Its mission is to create open specifications for the network delivery of managed services to local networks and devices. The OSGI organization is the leading standard for Internet services to homes, cars, small offices, and other environments. For more information, refer to the Web at:

`http://www.osgi.org`

SMF is an implementation of the OSGI Service Platform specification and provides a production-ready software management framework for network-delivered applications.

Stateful Web services (manageable resources) are hosted in a WebSphere runtime, using all the WS-RF specifications.

OSGI/SMF targets the environment where the size of the endpoint implementation needs to be limited and where a stateless Web service is enough to provide the functionality.

- The Common Information Model (CIM)

  The CIM class is considered the essential element of a resource model. The CIM classes determine:

  - The resources to monitor
  - The properties of each resource
  - The mechanism to collect the data

  CIM classes describe the status of each resource. A CIM class contains a list of properties relevant to a specific resource. CIM is built using a set of specifications and standards developed by the Distributed Management Task Force, Inc. (DMTF).

  Note: The DMTF is the industry organization leading the development of management standards and integration technology for enterprise and Internet environments. You can find more information about DMTF on the Web at:

  `http://www.dmtf.org/home`

  It provides common definitions of management information for IT infrastructure elements. CIM's common definitions enable the exchange of information between systems throughout the network. The use of CIM makes it possible to model IT infrastructure manageable resources, such as physical virtual servers, storage volumes, files, and so on, as stateful services according to the WS-RF specifications. In particular, the manageable resources provide a common interface in the form of the WS-Resource service, which is the base class of all resource services in a virtualized IT infrastructure environment. Each manageable resource has a *resource type* associated with it. Each manageable resource also has a *resource factory* associated with it to manage the creation and disposal of its instances. Additionally, the created manageable resource instances are registered in CIM-based repositories.

1.3 The vision: Service-oriented architecture

The On Demand Business is all about flexibility and integration. Therefore, we need to be able to quickly and easily assemble individual components of business and other functions into the precise application needed by the business at any particular time.

A *service* is any well-bounded, defined, and repeatable business task that can be invoked in a standard manner. *Service-oriented architecture* (SOA) is an IT architecture that supports service orientation based on open standards and that enables the assembly, the deployment, and the management of applications from reusable services that are independent of the applications and computing platforms on which they run. In this context, a business process will be treated as a flexible workflow composed of independent services across organizational boundaries and easily changed according to business needs.

SOA promotes the reuse of existing application components as services and the development of new services and their assembly in flexible ways.

SOA aims to develop applications that rely on functions, components, and objects. It is not a product or a platform. It defines a set of patterns and guidelines for creating loosely-coupled, business-aligned applications services. Because of the separation of concerns between description, implementation, and binding, this architecture provides unprecedented flexibility in responsiveness to new business opportunities.

SOA views every application or resource as a service, implementing a specific, identifiable set of functions. Each element of the architecture is a service that implements the operating environment capabilities together. Services can be invoked through defined communication protocols. SOA can be based on Web services. However, it may use other technologies instead.

SOA defines an interaction model between three primary parties:

- The *service provider* who publishes a service description and provides the implementation for the service.
- A *service requester* who can either use the Uniform Resource Identifier (URI) for the service description directly or can find the service description in a service registry and then bind and invoke the service. The service requester uses the contract to bind the client to the service.
- The *service broker* who provides and maintains the service registry. It gives the consumer directions on where to find the service and its service contract.

The conceptual model of SOA applies to the virtualization of both business functions and physical infrastructure. It spans the construction of applications as well as their deployment and management. Clients only see a collection of business services and are interested in their quality of service. The On Demand Operating Environment shields them from the details of application assembly and service delivery.

The On Demand Operating Environment consists of three main elements:

- Integration services, at the business and the application levels
- Infrastructure services, at the hardware and operating systems levels
- The enterprise service bus, that provides access to the services and so enables the building of SOA applications.

One of the most important insights of the On Demand Operating Environment is that a *common pattern* supports both application services and infrastructure services. Both kinds of events (at the integration services layer and at the infrastructure services layer) are encoded using a common event description format, and therefore, can be propagated through the infrastructure and correlated, or aggregated, independently of their source.

The virtualization capabilities are the foundation of the infrastructure services.

1.4 The Virtualization Engine roadmap

Two aspects can be highlighted to understand what the IBM Virtualization Engine wants and needs to provide:

- The request to make the IT infrastructure components very modular and easily accessible using standards.
- The necessity to break down the barriers between heterogeneous components.

Archived

1.4.1 Services for managing the IT environment

The concept of the SOA may apply to the IT environment, that is, the assembling of agile, automated business processes to support the business while reusing existing assets. To support this requirement, the use of standards is required:

- To allow system management disciplines to exchange information, because they may come from different vendors.
- To allow the system management discipline to exchange information with other types of resources.

Sharing and exposing configuration data is essential. This is why the Configuration Management Database (CMDB) is one primary component to develop. The CMDB is the primary and unique source of information to retrieve data about the IT configuration. To populate a CMDB, discovery tools are requested, and to visualize the configuration, a user interface is required.

The standard for resource management is WSDM, and the intent for the Virtualization Engine is to be completely WSDM-compliant in the near-future. Other IBM products, from Tivoli in particular, intend to link to a CMDB using WSDM standards and accessing the resources through Web services[2]. Today, RDS, one of the available Virtualization Engine products, has the tools for discovery and visualization, as described in Chapter 7, “Using Resource Dependency Service to build manageable resources” on page 215. IBM acquired Collation® at the end of 2005 and the product became available under the name of IBM Tivoli Application Dependency Discovery Manager (TADDM). TADDM is an agent-free discovery engine that enables you to create and visualize your application topology based on CIM models. We can expect some convergence in this area, driven by the WSDM standards.

CMDB will require the use of basic services, such as a Java runtime environment, a relational database, an LDAP directory, a message broker, and so on, for its role. All these services will build the foundation to implement IT management processes.

1.4.2 The IBM Systems Agenda

Virtualization is one of the three pillars of the IBM Systems Agenda, which provides a roadmap to enable an IT infrastructure for an On Demand Business. IBM is committed to breaking down the barriers between components and enabling them to be integrated. The three core design principles for the IBM Systems agenda are:

- Virtualization of everything
  - Continued enhancements to the virtualization facilities have to further simplify the infrastructure environment by integrating many of the features into the systems themselves. Virtualization functions need to simplify the platform management for small or midsize business with limited IT resources.
  - Systems need to better monitor and manage themselves, which will increase continuous availability, improve security, improve resource utilization, and improve quality of service levels. When virtualization capabilities are implemented, the automation concepts become more effective. The highest level of automation is achieved through autonomic computing technologies, another IBM initiative.

[2] IBM plans are subject to change without notice and this information represent goals and objective only. Such information is not intended as a definitive statement of a commitment to specific levels of function or delivery schedules.

Note: You can find more information about the IBM Autonomic Initiative on the Web at:

http://www-03.ibm.com/autonomic/

  - Each of these components will be part of the virtualization capabilities:
    - Storage
    - Processors
    - Networks
    - Applications
    - Data
    - Users

- Commitment to openness
  - IBM is committed to contributing to making open standards and openness part of its solutions. Openness reduces risk through freedom of choice.
  - IBM wants to offer systems that leverage technology innovation and fit within any IT environment, from small businesses that have limited resources, to very large multinational corporations, which have to deal with very complex heterogeneous environments.
  - IBM wants to offer systems that combine servers, operating systems, storage, middleware, networking, and applications that enable rapid deployment of a business solution based on industry standards.
  - For clients who want the flexibility to customize systems in their environment, IBM will deliver components that can be easily integrated because they are based on standards.
- Collaboration to innovate
  - Simplification across the enterprise infrastructure will provide more data security, more business flexibility, and more business opportunity.
  - Collaboration with business partners, systems integrators, and independent software vendors (ISVs) will offer complete and quicker business solutions.
  - To accelerate innovation, collaboration with the industry is fundamental.

1.5 The value of virtualization

Virtualization helps to improve IT system utilization and staff productivity with proven, industry-leading virtualization capabilities that are integrated throughout the full range of the IBM server and storage offerings. A virtualized infrastructure:

- Improves the total cost of ownership (TCO), by decreasing management cost, by increasing asset utilization, and by linking the infrastructure performance to business goals.
- Allows the access through a shared infrastructure, by leveraging common tools across many systems, by improving business resiliency and security, and by establishing the foundation for a Service Oriented Infrastructure.
- Increases the flexibility, by creating pools of system resources, by maintaining the freedom of choice with the use of open standards, and by masking the complexity.

Virtualization offers the following advantages:

- It eliminates the need to physically tie applications or data to the systems they use. It enables the consolidation of multiple systems, even large partitioned servers or storage products running different operating systems, with proven virtualization capabilities offered by IBM.
- It can improve the overall productivity and IT utilization for IBM and non-IBM platforms, as well as discover, monitor, and manage systems by consolidating and pooling IT resources, and consolidating tasks from multiple servers or storage products with IBM virtualization features that are included in the complete range of IBM systems.
- It logically combines and reallocates IT resources within a server or storage environment and manages system-level workload and performance across multiple systems with IBM virtualization features that support dynamic partitioning, virtual networks, virtual I/O, and virtual storage.
- By cascading proven virtualization technology from the mainframe and midrange systems throughout its server and storage offerings, IBM continues to lead the industry in overall virtualization capabilities.
- It improves manageability with a single view of your entire infrastructure using IBM systems that can apply virtualization capabilities synergistically to solutions from multiple vendors.
- It creates a single virtual storage pool that applications and employees across the enterprise can access by logically and securely combining servers and storage.
- It is able to move workloads to available resources to optimize performance by extending virtualization capabilities beyond servers and storage to internal and external resources, such as networks, applications and information, with technologies that support grid computing.
- Virtual access and management capabilities in the Virtualization Engine platform perform workload, platform, and systems management across a heterogeneous infrastructure. Technologies such as Enterprise Workload Manager, IBM Director, TotalStorage Productivity Center, and the Virtualization Engine console provide these capabilities.

Virtualization is the cornerstone for building the infrastructure necessary to implement a services-oriented architecture.

Archived

Archived

2

Planning

This chapter describes the planning task, which is the first task of the installation process. Proper planning will allow you to avoid many problems. Figure 2-1 describes the three main sub tasks for proper planning.

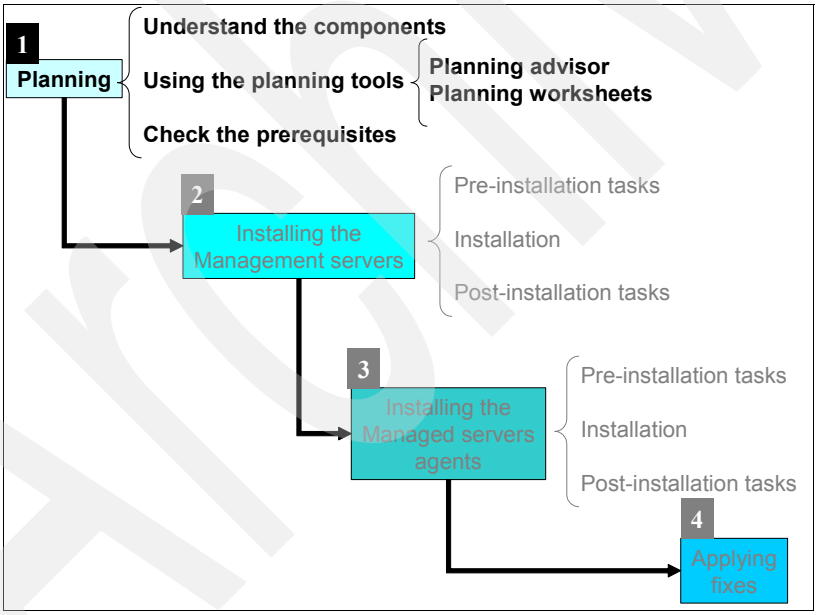


Figure 2-1 Task one: planning the installation

© Copyright IBM Corp. 2006. All rights reserved.

In this chapter:

- 2.1, “Terms used in the IBM Virtualization Engine” on page 21 helps you to understand how the products complement each other and how they are packaged.
- 2.2, “Planning tools” on page 33 describes the multiple tools that IBM provides you with to help you to decide how to implement the Virtualization Engine environment.
- As products, operating systems, and platforms usually have dependencies, 2.3, “Prerequisites” on page 50 describes the prerequisites to check when you have decided about your environment.

Archived

2.1 Terms used in the IBM Virtualization Engine

Some confusion could arise in the comprehension of the following terms, depending on the perspective in which they are looked at. In order to make things as clear as possible, we define these terms so that they are understood in the following way:

Components The *components* represent the functional or technical services provided by the IBM Virtualization Engine. They are sometimes called *features*, *services*, or *products* in the Virtualization Engine documentation. Some components can be made of other components: the installation process will install automatically dependent components as needed.

Common runtime components

Common runtime components are technical prerequisites for some of the components and they need to be installed: the installation process will install automatically the common runtime components as needed. They are sometimes called *common components* in the Virtualization Engine documentation.

Products *Products* are commercial assemblies of one or several components and they are named with product numbers. The products have the same logical set of Virtualization Engine common components. They can be obtained from IBM, either by purchase or free of charge through a hardware or software bundle. They usually have a product identification number.

Packaging The *packaging* defines the nature of the set of physical media containing the installation files of the components.

Virtualization Engine environment

The Virtualization Engine *environment* is a logical grouping of systems with Virtualization Engine software installed on them. The Virtualization Engine installation and configuration process gathers the information that defines the Virtualization Engine environment.

2.1.1 Components and common runtime components

The IBM Virtualization Engine management collection is made of the following basic *products*:

- IBM Enterprise Workload Manager (EWLM)
- Resource Dependency Service
- IBM Director
- IBM Usage and Account Manager (UAM)

The Virtualization Engine console is sometimes referred as a product, though there is no product identification number: the Virtualization Engine console is associated with one of these named products and is installed when chosen by the user or requested by one product; it would be more consistent to say that the Virtualization Engine console is a component.

The Virtualization Engine console, EWLM, Resource Dependency Service, and IBM Director are sometimes referred as *services,* in particular by the installer, as they do constitute the providers of the actual functional services. The products include the service that is the name of the product that can be selected for installation:

- For EWLM, it is called the Domain Manager.
- For IBM Director, it is called the Director Server.

Archived

These components can be made of other components with the following names that are externalized for the purpose of acquiring and installing fixes:

- IBM DB2® UDB Enterprise Server Edition
- IBM Integrated Solutions Console
- IBM WebSphere Application Server
- IBM Virtualization Engine console
- IBM Virtualization Engine base support
- IBM Virtualization Engine updates
- IBM Virtualization Engine security
- IBM Virtualization Engine shared WebSphere Application Server profile
- iLog JViews
- Cluster Systems Management bridge to IBM Virtualization Engine console
- Management Central bridge to IBM Virtualization Engine console
- IBM Director Bridge to IBM Virtualization Engine console

Beyond the console itself, the Virtualization Engine console also includes the following components:

- Director Bridge to Virtualization Engine console
- EWLM Bridge to Virtualization Engine console
- Cluster Systems Management bridge to the Virtualization Engine console
- Management Central Bridges to the Virtualization Engine console

The Virtualization Engine console and the bridges are selectable services at installation time.

The IBM Director and EWLM have endpoint code (agents) for which the installer presents a selection option (the installer does not install these endpoints: the images are loaded for a subsequent distributed installation):

- Director agent
- EWLM Managed Server

The Virtualization Engine common runtime includes core components that are not visible from the user's point of view. They are covertly installed during the installation wizard. They include the following:

- A J2EE server: JRE™ 1.4.2 is used.
- IBM Autonomic Deployment Engine framework: The IBM Autonomic Deployment Engine is a framework to centrally manage software changes during product and maintenance install and uninstall using a database of eXtensible Markup Language (XML) descriptors.
- iLog Jviews: This is a set of Java graphics products for building Web-based graphical user interface (GUI) displays, such as diagrams, charts, schedules, and so on. You can find more information about iLog Jviews on the Web at:

  `http://www.ilog.com/products/jviews/`
- Services to acquire or deploy modules for specific functions.

Two other components, namely the International Components for Unicode (ICU4J) and the IBM Common Logging module (ICL) are not visible to the user.

The common runtime components do not necessarily have to be installed in the same server on which the Virtualization Engine services (management servers) run.

Table 2-1 on page 23 describes the version used in the Virtualization Engine Version 2.1 for the common runtime components.

Table 2-1 Common runtime components

Common runtime component	Version
Integrated Solutions Console	6.0
WebSphere Application Server	6.0
IBM DB2 Universal Database™	8.2
IBM Tivoli Directory Server	6.0

Sometimes we use the term *Virtualization Engine base* to describe the components that need to be installed to install the products: understand that the *Virtualization Engine base* is not exposed at all to the user; this is a facility in regards to fixes installation.

Figure 2-2 depicts the relationship between the components.

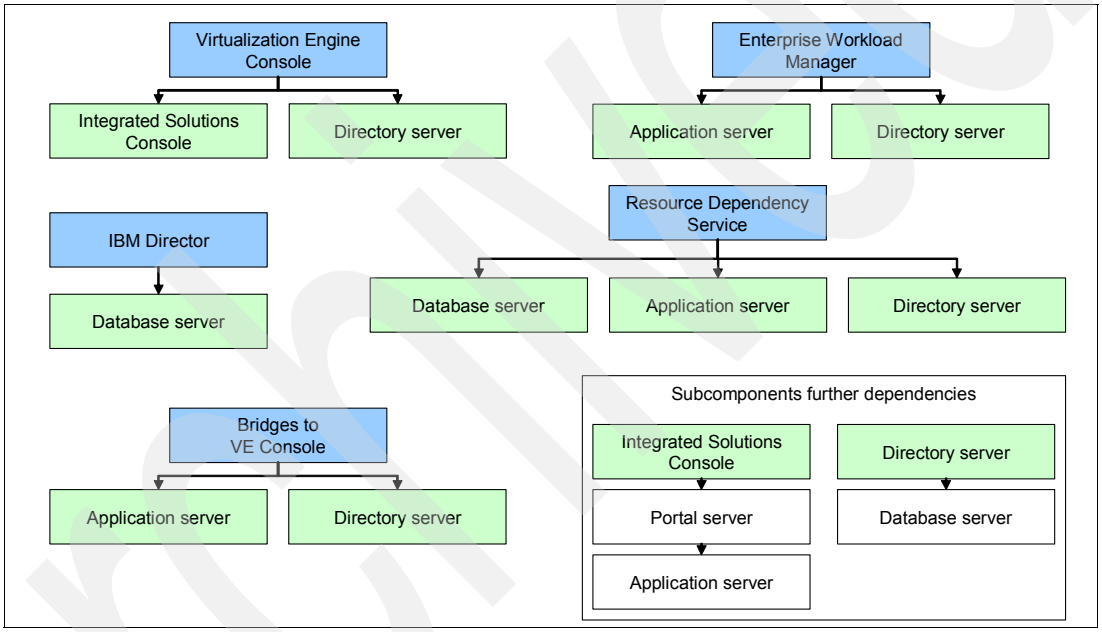


Figure 2-2 The Virtualization Engine components and their dependencies

The common runtime components do not necessarily have to be installed in the same server on which the Virtualization Engine services (management servers) run. The components have some "location" dependencies, as shown in Table 2-2 on page 24:

- "L" means that the common runtime component must be installed locally.
- "E" means that the common component can be installed on any server in the same Virtualization Engine environment. The Virtualization Engine environment is introduced in the next section (2.1.2, "Introducing the Virtualization Engine environment" on page 25). That means an ITDS/DB2 needs to be already installed.
- "R" means that the Virtualization Engine service can use a remote access to the common component. Note that Resource Dependency Service can access DB2 as a remote database.

Table 2-2 Virtualization Engine common runtime installation relationships

	EWLM domain manager	Resource Dependency Service	Virtualization Engine console	Console bridges
WebSphere	L	L	L	L
DB2	N/A	L/R/E	E	E
Directory Server	E	E	E	E
ISC	N/A	N/A	L	N/A
Resource modeling framework and type registry	N/A	L	N/A	N/A
Core components	L	L	L	L

Important: IBM Director does not have Virtualization Engine common runtime dependencies. The installation of IBM Director is a separate native process, except on i5/OS®. On i5/OS, the installation is also included in the Virtualization Engine installation wizard. IBM Director joins a Virtualization Engine environment when an IBM Director Bridge to the Virtualization Engine console is installed on the same system.

Note: It is important to understand that although those common runtime components are prerequisites for the Virtualization Engine components, they do not need to be ordered and installed separately, as they come embedded in the Virtualization Engine code and are automatically installed, whenever required, by the Virtualization Engine installer.

Naming conventions

Strict naming conventions apply to the different elements mentioned earlier. Table 2-3 shows a guidance summary that must be followed.

Table 2-3 Virtualization Engine naming conventions

Official name	Short name	Acronym
IBM Virtualization Engine console	Virtualization Engine console	
IBM Director 5.10	IBM Director	
IBM Enterprise Workload Manager	Enterprise Workload Manager	EWLM
IBM Virtualization Engine Resource Dependency Service	Resource Dependency Service	RDS
IBM Tivoli Usage and Account Manager		TUAM or UAM

2.1.2 Introducing the Virtualization Engine environment

The Virtualization Engine installation process creates a configuration repository in a Lightweight Directory Access Protocol (LDAP) Directory Server that contains all Virtualization Engine environment data, including some installation, configuration, and topology settings. This data is used for subsequent installations or services updates. The Virtualization Engine environment consists of the following type of data:

- Identity mapping repository location information.
- WebSphere Application Server security location information.
- The location of, and basic configuration information for, all installed Virtualization Engine products and components in the environment; for example, the location of Virtualization Engine console bridges, Enterprise Workload Manager (EWLM) domain manager, EWLM managed servers, IBM Director Server, IBM Director agents, and so forth. This information provides basic topology data for the Virtualization Engine console.
- Lightweight Third-Party Authentication (LTPA) keys and password.

All Virtualization Engine products and components must be part of the same environment. The Virtualization Engine environment is tied to the destination directory that you use for installing Virtualization Engine software. You establish the environment at installation time and you cannot change it until you remove all Virtualization Engine software from a given destination directory. Generally, a system has an association with a single environment; however, a system can belong to multiple environments if you install Virtualization Engine software into more than one destination directory. For example, you might choose to install a version of Virtualization Engine software to use for production and a second version to use for test. In this case, you install the applicable Virtualization Engine software into two different destination directories. The software in one destination directory is part of your production Virtualization Engine environment, and the software in the second destination directory is part of your test environment.

If the Virtualization Engine environment does not already exist on the directory server, the installation wizard creates a new suffix at the root of the Directory Server tree and stores the environment information under this suffix.

For all systems except i5/OS, the only compatible directory server is the IBM Tivoli Directory Server Version 5.1, 5.2, or 6.0. For i5/OS, the IBM Directory Server for i5/OS version V5R3, or later, is compatible.

2.1.3 The products

You can order the Virtualization Engine components through the following products:

- 5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on IBM System z
- 5655-EWM IBM Virtualization Engine Enterprise Workload Manager for z/OS
- 5733-DIR IBM Director with Virtualization Engine console
- 5733-EWM IBM Virtualization Engine Enterprise Workload Manager
- 5733-RDS Virtualization Engine Resource Dependency Service

Note: At the time of the writing of this IBM Redbook, the UAM product does not have any IBM product reference number. For more information, check with your IBM representative.

The following publications are available for ITUAM 5.1:

- *IBM Tivoli Usage and Accounting Manager Administrator's Guide*, GC32-1541, which provides information about installing and configuring the ITUAM system components.
- *IBM Tivoli Usage and Accounting Manager Web Reporting User's Guide*, GC32-1540, which provides information about how to use the Web-based reporting application, ITUAM Web Reporting.
- *IBM Tivoli Usage and Accounting Manager Data Collectors for Microsoft Windows User's Guide*, GC32-1543, which provides information about configuring and running ITUAM Data Collectors for Microsoft Windows®. This guide also includes information about configuring and running ITUAM Data Collector for Tivoli Decision Support for z/OSR (TDSz).
- *IBM Tivoli Usage and Accounting Manager Data Collectors for UNIX and Linux User's Guide*, GC32-1542, which provides information about configuring and running ITUAM Data Collectors for UNIX® and Linux.

The UAM product is not part of this IBM Redbook study.

They are also offered for free as bundles through the following products:

- 5733-DR1 Director with Virtualization Engine console for i5/OS
- 5733-EWA IBM Virtualization Engine Enterprise Workload Manager for i5/OS

The following sections provide the details of each product content.

5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and zSeries V2.1

Table 2-4 presents the functions provided by 5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on IBM System z9™ and zSeries® V2.1. It highlights the features that this product allows.

Table 2-4 Per platform components for product 5648-F08

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris™	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console						a a a			
Enterprise Workload Manager ► Domain manager ► Managed server	b	b	b	b	b	b	c	c	
Resource Dependency Service						d			

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris™	Microsoft Windows	z/OS
Virtualization Engine console						a d			
Bridges to Virtualization Engine console ► Director Bridge to Virtualization Engine console ► EWLM Bridge to Virtualization Engine console						 a a			

a. S01255B IBM Director with Virtualization Engine console V5.10
b. S01255D IBM Virtualization Engine Enterprise Workload Manager for Managing AIX, i5/OS, Linux, HP-UX servers V2.1
c. S01255G IBM Virtualization Engine Enterprise Workload Manager for Managing Windows or Solaris servers V2.1
d. S01255J IBM Virtualization Engine Resource Dependency Service V2.1

This product licensing policy is:

- Per managed processor: Features S01255B / S01255D / S01255G
- Per server: Feature S01255J

5655-EWM IBM Virtualization Engine Enterprise Workload Manager for z/OS V2.1

Table 2-5 presents the functions provided by the product (a “y” in the table cell means that the function is provided).

Table 2-5 Per platform components for product 5655-EWM

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console									
Enterprise Workload Manager ► Domain Manager ► Managed Server									 y y
Resource Dependency Service									
Virtualization Engine console									

Archive

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
Bridges to Virtualization Engine console ► Director Bridge to Virtualization Engine console ► EWLM Bridge to Virtualization Engine console									

This product licensing policy is per managed processor.

5733-DIR IBM Director with Virtualization Engine console V5.10

Table 2-6 presents the functions provided by the product. A “y” in the table cell means that the function is provided, and an “x” means that it is not.

Table 2-6 Per platform components for product 5733-DIR

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console	 y y		 y y y	 y y y	 y y y	 y		 y y y	
Enterprise Workload Manager ► Domain Manager ► Managed Server									
Resource Dependency Service									
Virtualization Engine console	y		y	x	y			y	
Bridges to Virtualization Engine console ► Director Bridge to Virtualization Engine console ► EWLM Bridge to Virtualization Engine console			y	y	y			y	

This product licensing policy is per managed processor.

5733-DR1 IBM Director with Virtualization Engine console for i5/OS V5.10

Table 2-7 on page 29 presents the function provided by the product. A “y” in the table cell means that the function is provided, and an “x” means that it is not.

Table 2-7 Per platform components for product 5733-DR1

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console			 x x x						
Enterprise Workload Manager ► Domain Manager ► Managed Server									
Resource Dependency Service									
Virtualization Engine console			x						
Bridges to Virtualization Engine console ► Director Bridge to Virtualization Engine console ► EWLM Bridge to Virtualization Engine console			x						

This product ships together with every i5/OS copy.

5733-EWA IBM Virtualization Engine Enterprise Workload Manager for i5/OS V2.1

Table 2-8 presents the functions provided by the product. A “y” in the table cell means that the function is provided. An “i” means that the component allows use of agents internal to the IBM System i machine only.

Table 2-8 Per platform components for product 5733-EWA

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console									
Enterprise Workload Manager ► Domain Manager ► Managed Server			 y y						
Resource Dependency Service									

Archived

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
Virtualization Engine console	i		i	i					
Bridges to Virtualization Engine Console ► Director Bridge to Virtualization Engine Console ► EWLM Bridge to Virtualization Engine Console			y						

This product is included in IBM System i Enterprise Edition.

5733-EWM IBM Virtualization Engine Enterprise Workload Manager V2.1

Table 2-9 presents the functions provided by the product. It highlights the features that this product allows.

Table 2-9 Per platform components for product 5733-EWM

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console									
Enterprise Workload Manager ► Domain Manager ► Managed Server	a[a] a	a	a	a a	a a		b[b]	a b	
Resource Dependency Service									
Virtualization Engine console	a	a		a	a			a	
Bridges to Virtualization Engine console ► Director Bridge to Virtualization Engine console ► EWLM Bridge to Virtualization Engine console	a	a		a	a			a	

a. 0001 IBM Virtualization Engine Enterprise Workload Manager V2.1
b. 0003 IBM Virtualization Engine Enterprise Workload Manager: Managed Server for Solaris, Windows Servers

This product licensing policy is per managed processor.

5733-RDS Virtualization Engine Resource Dependency Service V2.1

Table 2-10 presents the functions provided by the product (a “y” in the table cell means the function is provided).This product does not allow any optional features. This product licensing policy is per server.

Table 2-10 Per platform components for product 5733-RDS

Component	AIX	HP-UX	i5/OS	Linux on POWER	Linux on IBM System x	Linux on IBM System z	Solaris	Microsoft Windows	z/OS
IBM Director ► Director Server ► Director Agent ► Director Console									
Enterprise Workload Manager ► Domain Manager ► Managed Server									
Resource Dependency Service				y	y				
Virtualization Engine console				y	y				
Bridges to Virtualization Engine console ► Director Bridge to Virtualization Engine console ► EWLM Bridge to Virtualization Engine console									

Summary

To summarize:

- You can obtain the IBM Enterprise Workload Manager through the following products:
  - 5733-EWA IBM Virtualization Engine Enterprise Workload Manager for i5/OS V2.1
  - 5733-EWM IBM Virtualization Engine Enterprise Workload Manager V2.1
  - 5655-EWM IBM Virtualization Engine Enterprise Workload Manager for z/OS V2.1
  - 5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and z/Series
- You can obtain the Resource Dependency Service through the following products:
  - 5733-RDS IBM Virtualization Engine Resource Dependency Service V2.1
  - 5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and z/Series
- You can obtain the IBM Director through the following products:
  - 5733-DR1 IBM Director with Virtualization Engine console for i5/OS V5.10
  - 5733-DIR IBM Director with Virtualization Engine console V5.10
  - 5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and z/Series

Archive

– IBM Director is also available as part number under the following program numbers: 32R1224, 32R1223, 32R1227, 32R1225, 32R1230, 32R1226, and 32R1233.

2.1.4 The packaging

The IBM Virtualization Engine packaging adheres to the following general rules:

- Each Virtualization Engine component comes with its separate set of physical media. Note that the Virtualization Engine console is bundled as part of the Virtualization Engine common components.
- For packages that do not fit on a single CD-ROM because of their size, both the CD-ROM and the DVD formats are available. The client has a choice between the two formats.

Table 2-11 details the packaging for the IBM Virtualization Engine common components.

Table 2-11 Virtualization Engine common components packaging

Product	Packaging
IBM Virtualization Engine common components for AIX 5L	Seven CD-ROMs or one DVD
IBM Virtualization Engine common components for i5/OS	Five CD-ROMs or one DVD
IBM Virtualization Engine common components for Linux on POWER	Seven CD-ROMs or one DVD
IBM Virtualization Engine common components for Linux on IBM System x	Seven CD-ROMs or one DVD
IBM Virtualization Engine common components for Linux on IBM System z	Eight CD-ROMs or one DVD
IBM Virtualization Engine common components for Windows	Seven CD-ROMs or one DVD

Note: These packages always contain the Virtualization Engine console as part of the Virtualization Engine common components.

Each first CD-ROM of these sets is labeled START HERE in order to make clear that the installation process must start with the Virtualization Engine common components.

Table 2-12 details the packaging for the IBM Enterprise Workload Manager.

Table 2-12 Enterprise Workload Manager packaging

Product	Packaging
IBM Virtualization Engine Enterprise Workload Manager domain manager for Windows and Linux on IBM System x	One CD-ROM
IBM Virtualization Engine Enterprise Workload Manager domain manager for AIX 5L, i5/OS and Linux on POWER	One CD-ROM
IBM Virtualization Engine Enterprise Workload Manager domain manager for Linux on IBM System z	One CD-ROM
IBM Virtualization Engine Enterprise Workload Manager managed servers for AIX 5L, i5/OS, HP-UX, Linux on IBM System x, IBM System z and Linux on POWER	One CD-ROM

Product	Packaging
IBM Virtualization Engine Enterprise Workload Manager managed servers for Windows and Solaris	One CD-ROM

Table 2-13 details the packaging for the Resource Dependency Service.

Table 2-13 Resource Dependency Service packaging

Product	Packaging
IBM Virtualization Engine Resource Dependency Service for Linux on IBM System z, IBM System x, and POWER	One CD-ROM

Table 2-14 details the packaging for the IBM Director.

Table 2-14 IBM Director packaging:

Products	Packaging
IBM Director for IBM System p	One CD-ROM
IBM Director for IBM System i	One CD-ROM
IBM Director for IBM System x and BladeCenter	One CD-ROM
IBM Director for IBM System z	One CD-ROM
IBM Director Agents	One CD-ROM

Note: The IBM Director for IBM System p and IBM Director for IBM System x packages also contain all Director agents. The base directory CD for IBM System i and for IBM System z do not contain the agents; the packages include the Director agent CD.

2.2 Planning tools

This section describes the tools you need when you plan to install the Virtualization Engine V2.1.

These planning tools include:

- The Information Center
- The planning advisor
- Planning worksheets
  - Port assignment worksheet
  - User access and password planning worksheet
  - Installation planning worksheet
  - Security planning worksheet

Note: We have used these planning tools extensively in this ITSO project. All the installation examples in the following chapters are the result of in-depth investigation and analysis using these tools. We recommend that you use these facilities.

Archive

Although quite straightforward, deploying a Virtualization Engine environment within your enterprise has to be well planned to avoid time-consuming mistakes. Complex factors include which products or components you want to use, how many platforms and machines you have to manage, and how wide the network is.

Proper planning in advance is essential. The tools will help you in researching and evaluating which combination of the Virtualization Engine products you are interested in, and then deciding on which servers type and operating system to store them. They will also help you in identifying the minimum required prerequisites either for the hardware (processor, disk, and memory size) or the software (operating system levels). The planning advisor and the worksheets are both able to go in depth into technical details because they contain links that redirect you to the Information Center.

2.2.1 The Information Center

The Information Center on the Web is a single location of the repository for the Virtualization Engine documentation. This is a well-structured Web site where you can find important up-to-date information about the Virtualization Engine. You can find information easily using the tree index or the powerful search button.

You can access the Information Center on the Web at:

`http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp`

Note: In this book, the terms “Information Center” and “InfoCenter” have the same meaning and refer to the IBM System Software Information Center.

Figure 2-3 shows the Information Center Web page.

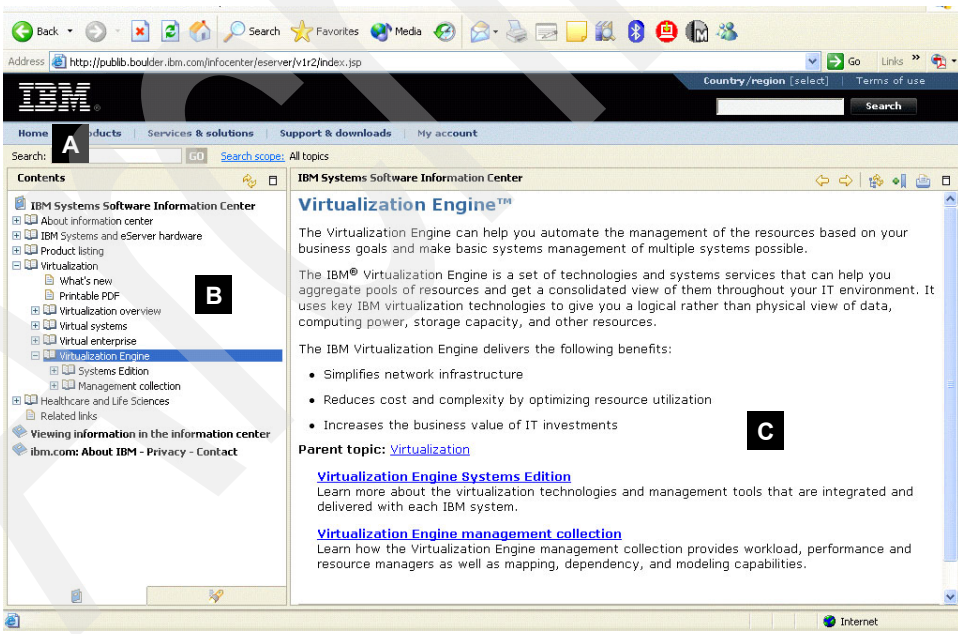


Figure 2-3 Information Center main page

As shown in Figure 2-3, the Information Center Web page has three areas:

- A: The search button is at the top of the page. You can make searches using words or phrases. Search results are listed on the left part of the window, sorted by matching percentage. Select one result to display the related text on the right part of the page.

- B: The left side organizes the topics into a tree-based hierarchy list. You can use this hierarchy to access printable PDF files as well as to access Virtualization Engine information quickly. Click **Systems Edition** or **Management collection** to get a list of topics that you can access. The selected topic is displayed in the right side of the window.
- C: This is the part of the window that displays the selection from the left part (B) of the screen. You can adjust the dividing bar between part B and part C of the window to help you read specific topics in the window better.

The Information Center is only available on the Internet. It cannot be accessed offline. From the Information Center, you can print selected topics in PDF format. You can also order some of them from the IBM Publications Center on the Web at:

http://www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi

The Information Center organizes all individual topics into topic collections. Each topic collection has a printable PDF. You can see this in the first part of the Web page that contains the same topic information as the product's online help information.

For the Virtualization Engine V2.1, your browser settings determine the language version of the Information Center.

2.2.2 The planning advisor

The Virtualization Engine planning advisor is an interactive tool that produces customized planning output based on your needs. The planning advisor provides you with the specific information about which products to order, how to deploy the various components, and how to install the system services on your environment.

This section describes how to use this planning advisor tool to help you plan your installation. Use the planning advisor to:

- Identify what products you want to install.
- Produce a design of your Virtualization Engine topology.
- Obtain detailed information about the product to be ordered
- Generate customized installation flow charts.
- Print and record your customized output before beginning the installation process.

The tool is continuously updated and provides you with the latest technical information, such as operating system support, hardware requirements, documentations, and so on. However, it does not perform an installation or configuration directly on any systems in your enterprise.

Accessing the planning advisor

You can start the planning advisor tool directly from the Information Center on the Web at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/veinfo/planner.htm

You can also start the planning advisor tool directly from the Virtualization Engine Web site under the planning advisor link on the Web at:

http://www-03.ibm.com/servers/eserver/about/virtualization/getstarted.html

When you click **Planning advisor**, the first window that appears is a disclaimer. The Virtualization Engine planning advisor is under the specific terms of service. You must click **Accept** before you can continue using this tool.

Archived

Figure 2-4 shows the planning advisor starting page.

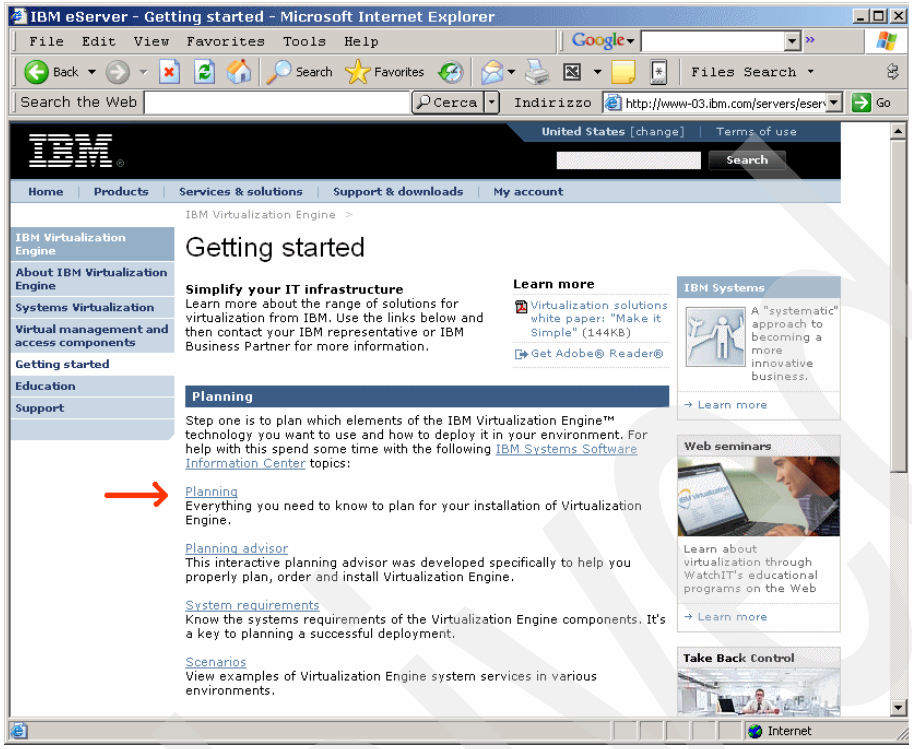


Figure 2-4 Planning advisor starting page

Configuring your Virtualization environment

To configure your Virtualization Engine environment, perform the following steps:

1. A new flash-based tool opens in a separate browser window, as shown in Figure 2-5 on page 37. This is the presentation window. Click **Start**.

Archive

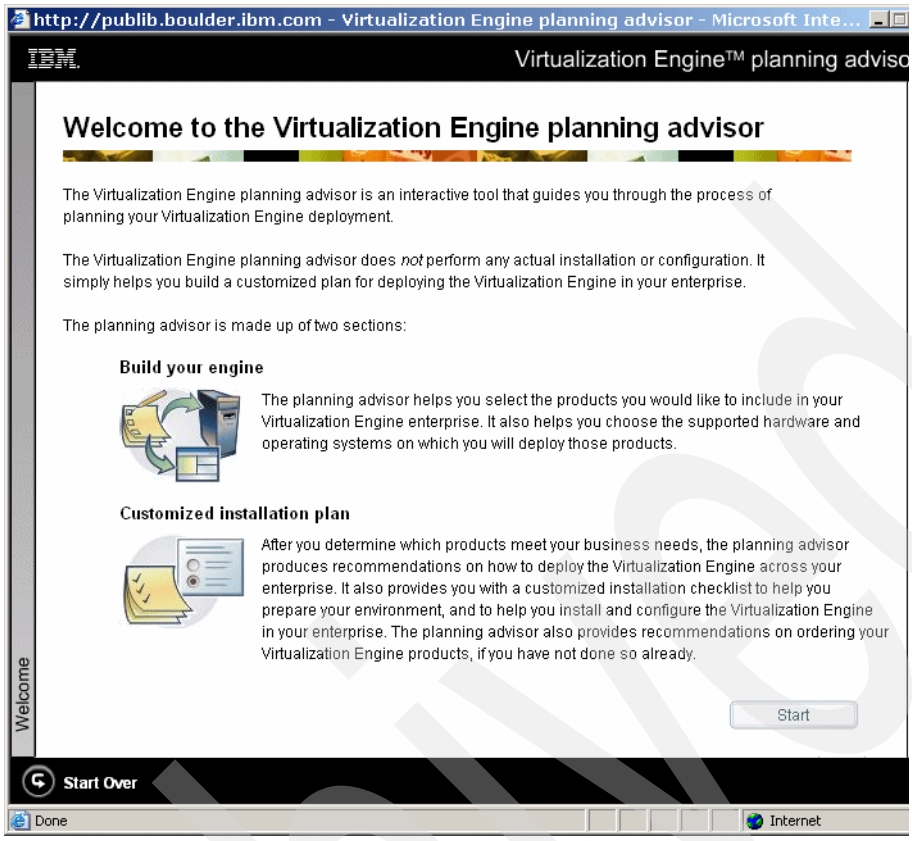


Figure 2-5 The first window of the planning advisor

2. The products selection window appears, as shown in Figure 2-6.

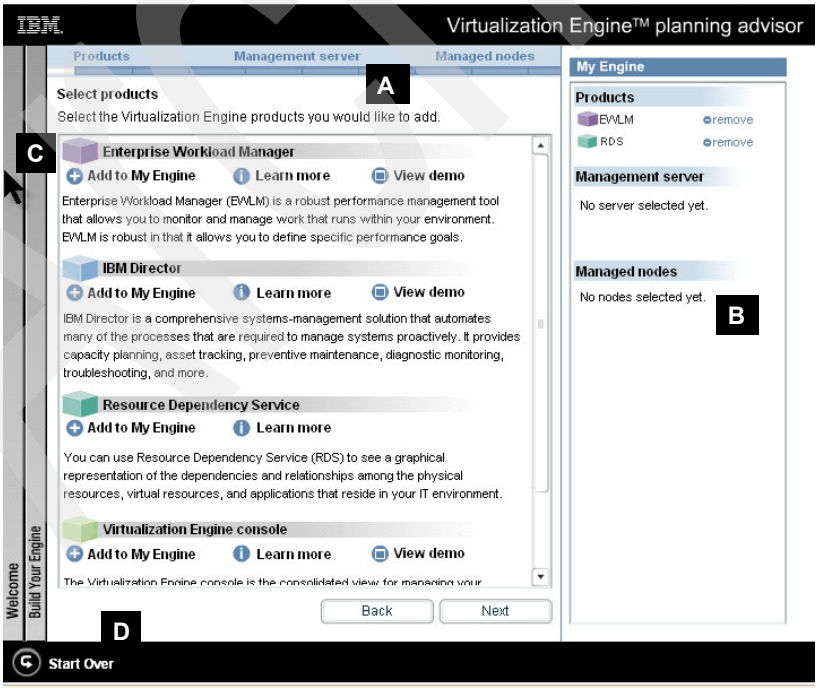


Figure 2-6 Products selection

There are some important considerations in order to take the best advantage of the planning advisor tool.The products selection window shown in Figure 2-6 on page 37 has the following panes:

– A: Note the progress bar at the top of the pane. This bar shows you the progress of the planning. The three main planning steps you have to follow are listed above this progress bar:

  • Products

    This lists the Virtualization Engine products you can add in your enterprise.

  • Management server

    This defines the Management Servers on which to deploy the selected products.

  • Managed nodes

    This selects the types of environments that will be managed by the Virtualization Engine.

    Note: If you select only Resource Dependency Service, there will be no managed nodes configuration, as Resource Dependency Service products do not expect any agent code to be distributed on the client machines.

– B: On the left side of the window is a report of the selections you made in the windows during the usage of the planning advisor. You can review or edit your previous selections from this pane at any time while using the tool.

  Note: You can select the Virtualization Engine console at any time, as it is always requested by other products.

– C: Each product has one specific color assigned. Each of them has a brief explanation and a flash demo to show the main characteristics and functionality of the products. If you click the **Learn more** link for each product, you will see a brief explanation about what the product is and a detailed list of software and hardware requirements sorted by platform and operating system.

– D: The Start Over button at the very bottom of the window allows you to restart the tool any time. It will reset any kind of selection that was made earlier in the session.

Press **Next**.

3. The window that appears has about information about what a management server is and its function in a Virtualization Engine enterprise. Once you have read the recommendations listed here, click **Next**. The window in which you should define your management nodes appears.

   Note: We recommend that you install all the management server components on a single system to simplify the installation process and to avoid potential error and extra post-installation steps. However, other considerations may prevail, such as your security environment, the level of availability and the independence of the components you want, and the skills of your staff.

   Complete the selection and the definitions of one or more management servers for the services. The planning advisor helps you during the selection. It presents only the suitable alternatives, and will take care of the software requisites of the products. For example, at the time of the writing of this IBM Redbook, Resource Dependency Service can be installed only on Linux.

As shown in Figure 2-7, we defined the installation of all products (except the IBM Director) on a single machine, that is, an IBM System x with the Linux SUSE Enterprise Server 9 operating system.

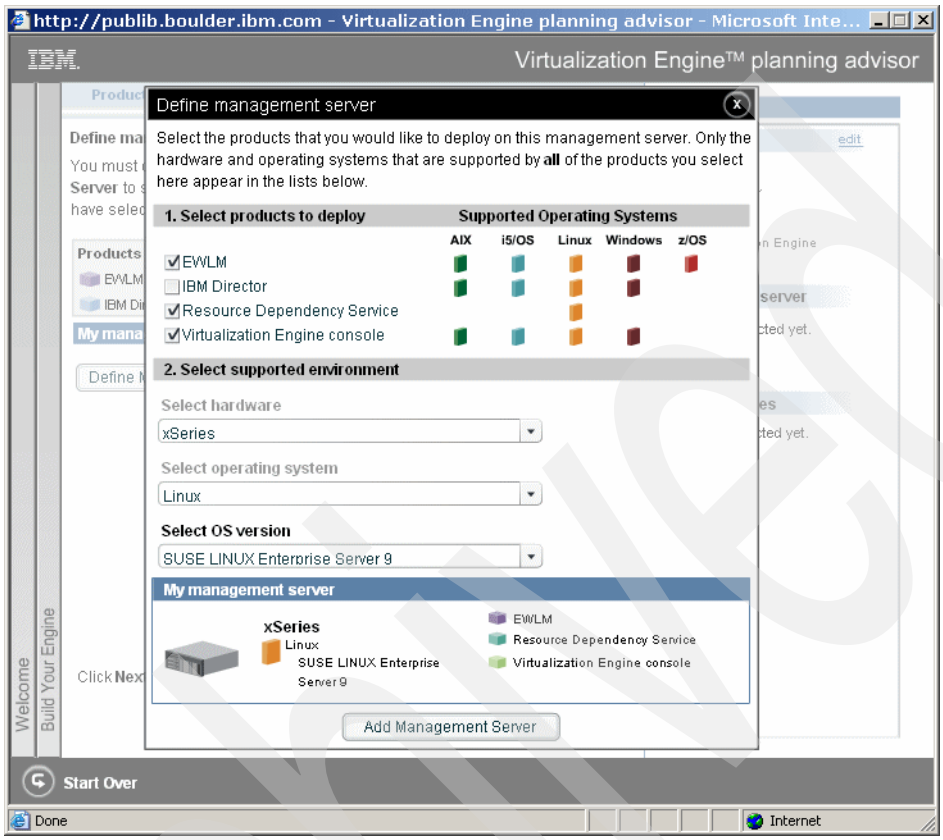


Figure 2-7 Example of management servers definitions example

4. After you have completed the management node configuration, you should provide the directory server requirements and the managed nodes information.

   For the managed nodes, select the different hardware platforms. For each of them, select the operating system that will be a managed node. You should define each unique

Archive

hardware platform and the operating system combination that will be on your managed nodes, as shown in Figure 2-8.

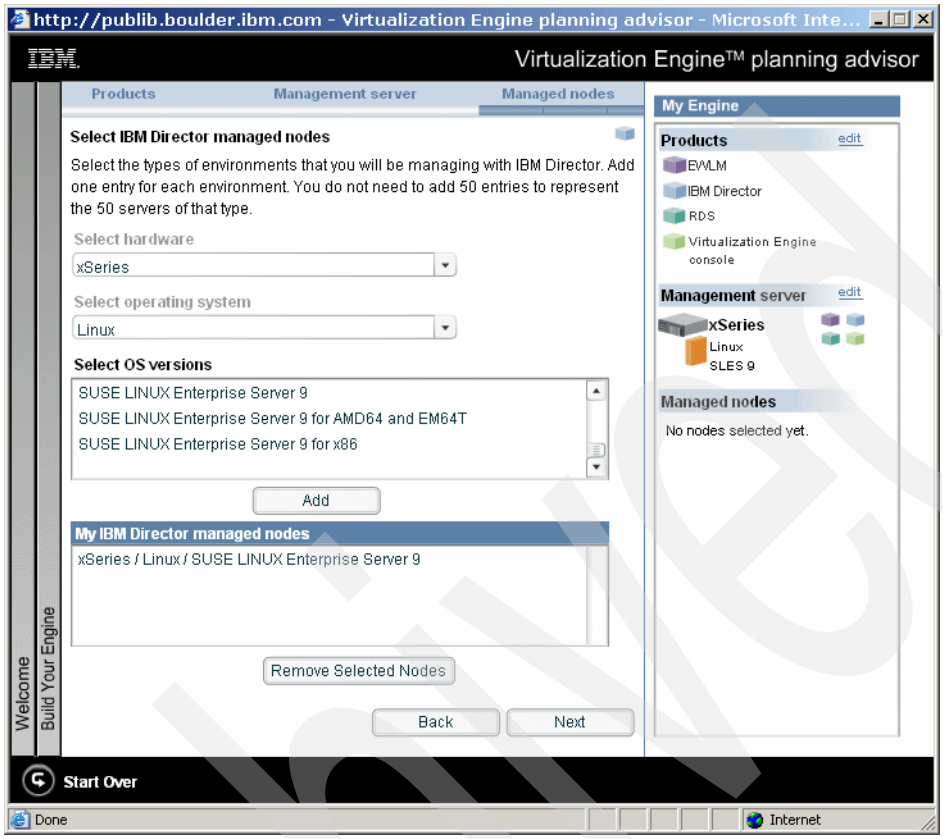


Figure 2-8 Managed node definitions example

Note: The management node concept can be applied only for IBM Director and EWLM. It is not applicable to Resource Dependency Service or the Virtualization Engine console.

At the time of the writing of this IBM Redbook, in the Linux operating system, for each type of platform, we could install the managed server EWLM code. However, we could not create the managed server domain. The error message `FEW239E "This operating system doesn't support the EWLM manager server"` was then displayed.

Although the managed nodes task in the planning advisor lets you configure any kind of managed node using the Linux operating system, this support will be made available later through Virtualization Engine Fix Packs.

Results and findings

Once you complete the managed nodes configuration, the results window appears. This window is a report of the Virtualization Engine environment as configured by you. You can

Archived

use this plan to install and deploy the Virtualization Engine in your enterprise. This window has the following three tabs:

- The My Engine tab

  This tab shows a collection and review of what you have selected in the previous windows. It contains the recommended hardware, the list of the engagement nodes, and the type of managed nodes selected, as shown in Figure 2-9.

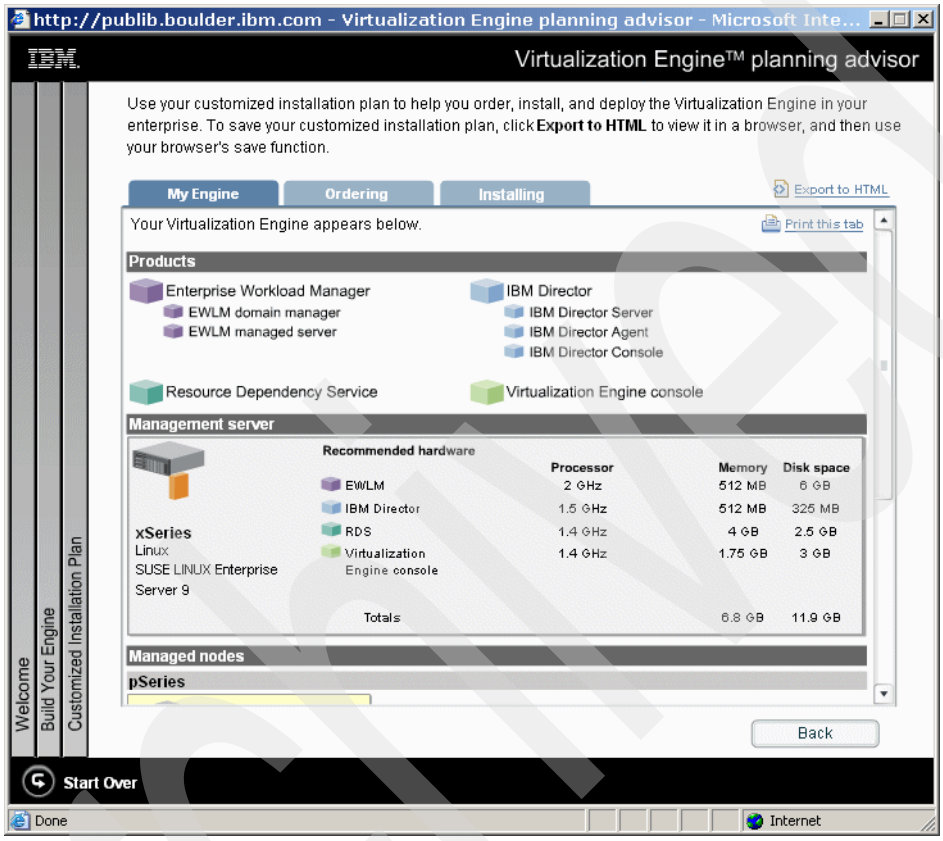


Figure 2-9 Results part 1: the My Engine tab

Archive

- The Ordering tab

  The Ordering tab, as shown in Figure 2-10, is based on data collection and details all the necessary software products that you have to order to implement the designed environment.

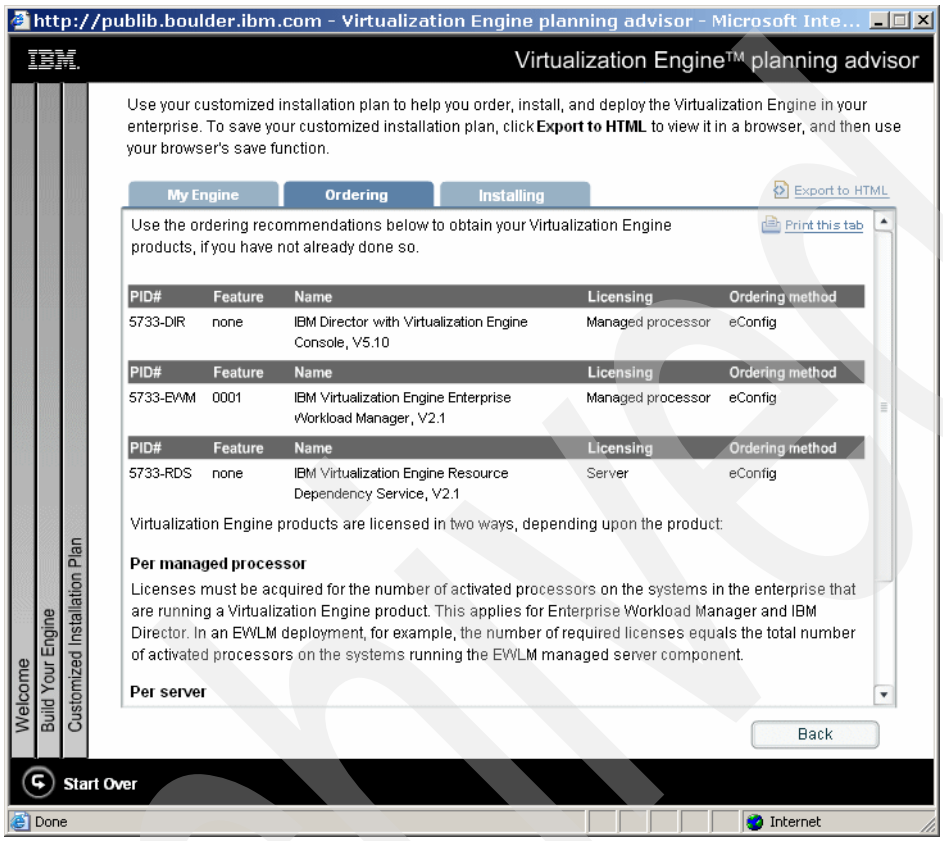


Figure 2-10 Results part 2: the Ordering tab

- The Installing tab

  The Installing tab, as shown in Figure 2-11 on page 43, contains the installation steps in the right order based on the products, management servers, and managed nodes you have previously described.

  The results are displayed as a checklist that you can use to for the step-by-step installation of the Virtualization Engine environment. There are special icons to define the level of

Archived

difficulty of each task. They are in different colors depending on the level of expertise and time required for each task.

To maximize the usability of the tool, there are links to specific topics in the Information Center, as Figure 2-11 shows.



Figure 2-11 Results part 3: the Installing tab

You can either export or save your results in Hypertext Markup Language (HTML) format. This can prove useful later when you log out. This is a useful option that allows you to save your configuration and compare it with others if required.

The HTML-exported results page, as shown in Figure 2-12, displays the customized installation plan. It includes details about the engine, the ordering, and the installation instructions. In this HTML-exported version, there are also links to the Information Center that you can select.

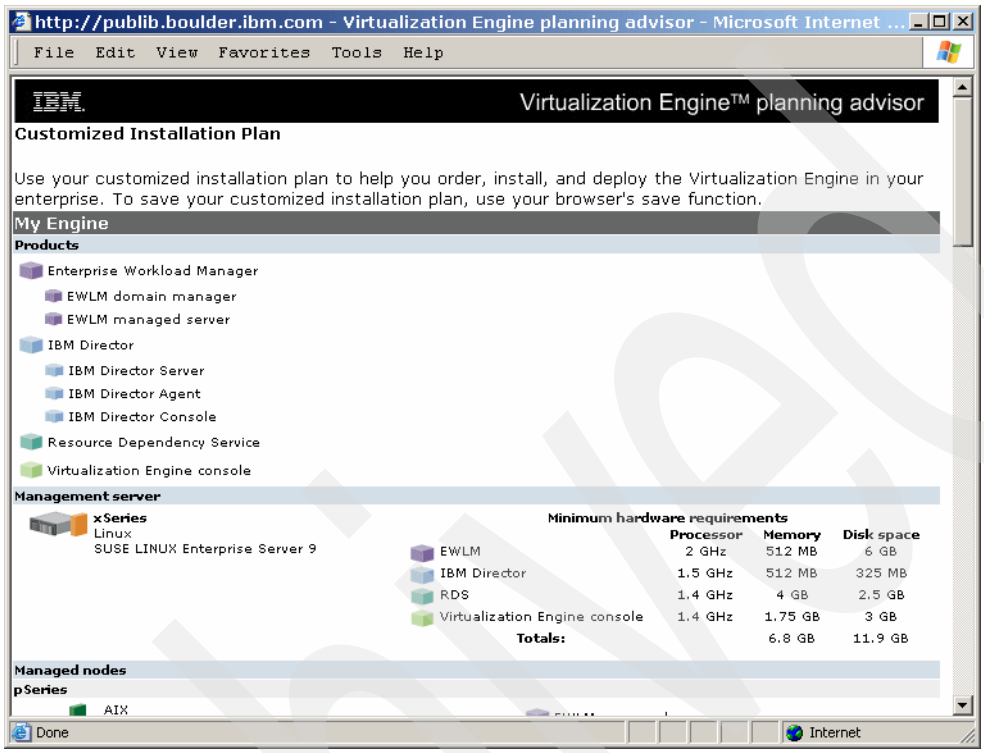


Figure 2-12 Plan exported in HTML format

2.2.3 Planning worksheets

The planning worksheets contain the tasks and considerations that you must address before you begin the installation of the Virtualization Engine products.

These worksheets help you retrieve and record the information that you will need during the installation and configuration of the Virtualization Engine products. Print each page, record your answers, and use the completed worksheets for installing and configuring the Virtualization Engine environment.

In this section, we describe the worksheets we used in our project. They are presented in the same sequence in which they should be filled.

You can also find details about these planning worksheets in the Information Center. We always used these worksheets during the setup of our laboratory's environments and found them very useful. We recommend that you use them.

Archive

Port assignment worksheet

Use this worksheet, described in Table 2-15, to record the ports you assign to all the Virtualization Engine components, subcomponents, and products.

We recommend that you accept the default port assignment values during the installation process as long as the default ports are not currently in use on your system. Table 2-15 lists the default ports used as well.

Table 2-15 Port assignment worksheet

Application	Description	Default port values	Your port value
IBM Tivoli Directory Server: ► Main server port ► Admin daemon port	► Directory server port for common communication with other applications ► Directory server port for administrator connections with the directory server	► 389 ► 3538	
Shared WebSphere Application Server profile	The shared WebSphere Application Server profile requires the use of a block, or set, of 15 ports for internal communications.	18551 (default starting port in block) through 18565 (default ending port in block).	
Integrated Solutions Console (ISC)	The Integrated Solutions Console (ISC) requires the use of a set, or block, of 30 ports for internal communications. Two of the most useful ports include the ports used to log into the Virtualization Engine console and the HTTP port that the ISC uses. Other ports in the block of consecutive ports enable other types of internal communications between the Virtualization Engine console and other applications.	► 8421 (default starting port in block) through 8450 (default ending port) ► 8421 (all operating systems) port used to log on to the Virtualization Engine console ► 8425 HTTP port (i5/OS) or 8426 HTTP port (other operating systems)	

Note: If you are installing the Tivoli Directory Server on a Linux system, you need to consider the possibility that there might be an existing OpenLDAP configuration on the system.

Archived

The Virtualization Engine installation wizard, detailed in 3.4, “Using the Virtualization Engine installer” on page 88, will store information about the port assignments in specific files on your system. You can check the following files after installation to verify the Virtualization Engine console port assignments. On AIX, Linux, and Windows, use the following command:

```
<Install_location>/p/VEISC/config/cells/VEISCNode01Cell/nodes/VEISCNode01/serverindex.xml
```

In this command, <*Install_location*> represents the Virtualization Engine installation directory.

Users access and password planning worksheet

Use this worksheet, described in Table 2-16, to record the user IDs and passwords you want to use for your Virtualization Engine environment.

There are a number of user IDs and passwords associated with the products of the Virtualization Engine. You should set up many of these user IDs during the installation process. You should set up other user IDs for individual Virtualization Engine products, such as EWLM, during the post-installation and configuration processes for each individual product.

For convenience, the default used values are also reported.

Table 2-16 User access and password planning worksheet

Value	Description	Default Value	Answer
Virtualization Engine administrator	Used to sign on to the Virtualization Engine console and administer the WebSphere Application Servers. The Virtualization Engine administrator user ID is the same as the following user IDs: ► Virtualization Engine console administrator ► WebSphere Application Server shared profile administrator ► WebSphere Application Server administrative console administrator	veadmin	
Directory server administrator	Used to manage the directory server. Note: If you are using multiple directory servers for the Virtualization Engine, you may be using different administrator distinguished names (DNs) or passwords for each one. In that case, ensure that you record all of the DNs and passwords here.	cn=Admin	
Director server instance owner (operating system user ID)	Used as the owner of all of the directory server installation objects. A system user ID appropriate for use as the directory server instance owner must exist prior to beginning installation of the directory server.	Not applicable	

Archived

Value	Description	Default Value	Answer
DB2 instance owner (operating system user ID)	Used to manage the DB2 server. It controls all DB2 processes and owns all file systems and devices used by the databases contained within the instance. For example, you use this user ID to start and stop the DB2 server.	vedb2ur	
DB2 fenced user (operating system user ID)	Used to run user-defined functions (UDFs) and stored procedures outside of the address space that the DB2 database uses.	vedb2fc	vedb2fc You cannot change the default value for this user ID.
DB2 administration server user (operating system user ID)	Used to run the DB2 administration server on the system. Also used to sign on to DB2 administration tools.	vedb2ad	vedb2ad You cannot change the default value for this user ID.
Lightweight Third Party Authentication (LTPA) key password	Used to protect the file that contains the WebSphere Application Server LTPA keys.	Not applicable	

Important: This worksheet may contain sensitive information. Once compiled, keep it safely stored in a secure place.

Installation setup planning worksheet

Use the installation setup planning worksheet described in Table 2-17 to record the values you need to start the installation process.

This worksheet needs to be used during the Virtualization Engine Common Component installation.

Table 2-17 Installation setup planning worksheet

Information	Description	Answer
Destination directory	The directory in which you want to save the installation images.	
Temporary directory	(This pertains to an i5/OS installation only.) The directory in which you store the installation images that the wizard needs to copy from the distribution media. This location is a temporary directory on the Windows workstation that you are using to perform your i5/OS installation.	
Log file directory	The directory in which you want to save the log files.	

Information	Description	Answer
Host name of directory server	The fully-qualified host name of the directory server that you want to use to host the Virtualization Engine environment information.	
Server port	The non-secure port number that the directory server uses for communication with other applications. The default port number for the server port is 389.	
Admin daemon port	The port number that the directory server uses for administrator connections with the directory server. The default port number for the Admin port is 3538.	
Directory server administrator	The directory server administrator user ID in the form of a distinguished name.	
Name for the Virtualization Engine environment	The unique name that you specify to identify the Virtualization Engine environment.	

Directory server planning worksheet

Use the directory server setup planning worksheet described in Table 2-18 to record the values to install the directory server. They are used during the Virtualization Engine Common Component installation.

Table 2-18 Directory server setup planning worksheet

Information	Description	Answer
Host name of directory server	The fully qualified host name of the directory server that you want to use to host the Virtualization Engine environment information.	
DB2 instance owner (operating system user ID)	The operating system user ID you specify to be the DB2 instance owner.	
Server port	The port number that the directory server uses for communications with other applications. The default port number for the server port is 389.	
Admin daemon port	The port number that the directory server uses for administrator connections with the directory server. For example, this port number is used in commands to start and stop the directory server. The default port number for the Admin port is 3538.	
Encryption seed	A string of characters for the directory server to use as an encryption seed.	

Archived

Information	Description	Answer
Director server instance owner (operating system user ID)	The operating system user ID you specify to be the directory server instance owner.	
Directory server administrator	The directory server administrator user ID in the form of a distinguished name.	

Security planning worksheet

Use the Virtualization Engine security planning worksheet described in Table 2-19 to record the values to install the Virtualization Engine products that require input for the Virtualization Engine security.

Use this worksheet after the Virtualization Engine Common Component installation and before the deployment of the services. For example, you will need it before you start the installation of the Virtualization Engine console. This information pertains to the Enterprise Identity Mapping (EIM) security and the shared WebSphere Application server installation.

Note: If the only product you plan to install in your environment is the IBM Director, you do not need to complete the following worksheet.

Table 2-19 Security planning worksheet

Information	Description	Answer
Virtualization Engine administrator	The User ID used to sign on to the Virtualization Engine console and to administer the WebSphere Application Servers.	
LTPA key password	A password that secures the LTPA cryptographic keys.	
LTPA time-out	The time-out interval for the LTPA bind. The value is the number of minutes that an LTPA token passed from one machine to another machine is valid. If your value is set to 120, after 120 minutes, that token expires. If you are concerned about your network speed, you should increase this number to avoid communication errors. If the token has expired when it reaches a bridge application, the bridge application shows up with a status of unknown in the Virtualization Engine console.	
LTPA single sign-on domain name	The single sign-on domain name of the HTTP server that the Integrated Solutions Console uses. The value of this parameter will be used as the domain name of the LTPA token cookie, and LTPA tokens will only be accepted when they come from this domain. The value is not the same as the directory server host name.	

Archived

There are two other parts of this table that are not described in this book. You can find them in the Information Center. We did not include them here because we did not use them in our project. The information included in these parts pertains to the user registry settings within the directory server and the Identity Mapping Repository.

Note: If you are not familiar with Lightweight Directory Access Protocol (LDAP) usage, we recommend that you use the IBM default values. By using the recommended default values provided by the installation wizard, you can avoid potential configurations errors.

More planning worksheets

You can find a complete repository of all the planning worksheets available in the Information Center on the Web at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/esmcinfo/eicacinstallmslinux.htm

We have more worksheets for each service component, such as EWLM, IBM Director, and Resource Dependency Service. They relate to prerequisites, user account, domain topology, and managed server planning. All of them contains tasks and considerations that must be addressed before beginning the installation of the service.

All the worksheets are not detailed in this IBM Redbook. You will find all the detailed worksheets in the Information Center.

2.3 Prerequisites

This section discusses the prerequisites the Virtualization Engine installation depends on, both from a hardware and a software point of view.

The Virtualization Engine planning advisor provides the first level of hardware and software prerequisites as a first step to help plan and deploy the environment. You must study these elements attentively in order to avoid any missing prerequisite issues during the installation phase itself.

Note: We strongly recommend that you consult the Information Center for the latest detailed and current information.

Note: Precise information about IBM Virtualization Engine requirements is available on the Web at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/veicinfo/eicarsystemservicereqs.htm

The up-to-date, comprehensive list of operating system fixes required by the IBM Virtualization Products is available on the Web at:

http://techsupport.services.ibm.com/server/VirtualizationEngine/osDownloads.html

2.3.1 AIX on IBM System p

The following sections discuss the hardware and software requirements for an AIX environment on System p.

Hardware requirements

Table 2-20 describes the hardware requirements for management servers running on AIX on IBM System p systems.

Table 2-20 Hardware requirements for AIX on IBM System p

Component	Processor	Memory	Disk Space	Comment
Virtualization Engine console	1.4 GHz	1.75 GB	3 GB	The paging space should be double the physical memory.
Enterprise Workload Manager	1.4 GHz	512 MB	6 GB	
IBM Director Console	1.5 GHz	512 MB	170 MB /opt 2 MB /usr	
IBM Director Server	1.5 GHz	1024 MB	325 MB /opt 2 MB /usr	
IBM Director Agent	1.5 GHz	512 MB	75 MB /opt 1 MB /usr	

The hardware requirement depends on the environment you have to manage. A large number of managed servers would require stronger performances from the management servers.

Note: From a performance view, the best method is to install every component of the Virtualization Engine on a different server. Alternatively, for easy management, we recommend that you install all Virtualization Engine components in one AIX box.

Both the installation process and the execution process require fast storage. It is better to place the file system /opt on a very fast external storage. It is also possible to install the Virtualization Engine console on different file systems.

Software requirements

This section details the software requirements, including the supported operating systems, operating system fixes, and recommended filesets for AIX on IBM System p.

Operating system

The least supported level for Virtualization Engine code is AIX 5L V5.3 ML03. The Virtualization Engine introduces a lot of new features to manage an IT infrastructure. Therefore, we recommend that you install the latest AIX fixes. The environment that is discussed in this chapter uses Technology Level 04.

Table 2-21 describes the operating systems that are supported by the Virtualization Engine components.

Table 2-21 Supported operating systems for AIX on IBM System p

Component	AIX on POWER
Enterprise Workload Manager	
Domain Manager	AIX 5L V5.3 *
Managed Server	AIX 5L V5.2 ** AIX 5L V5.3 *
IBM Director	
Director Server	AIX 5L V5.3 *
Director Agent	AIX 5L V5.2 AIX 5L V5.3
Director Console	AIX 5L V5.3
Resource Dependency Service	
Resource Dependency Service	N/A
Virtualization Engine console	
Virtualization Engine console	AIX 5L V5.3 *
Director Bridge	AIX 5L V5.3
EWLM Bridge	AIX 5L V5.3

* Requires 5300-03 recommended maintenance package or later.
** Requires 5200-03 recommended maintenance package or later.

The minimum recommended technology level for EWLM agents is an AIX 5L V5.2 ML03, and AIX 5L V5.2 ML04 for IBM Director. As AIX 5L V5.3 has a lot of useful features for virtualization, such as micro-partitioning, virtual I/O, Virtual Small Computer System Interface (SCSI), Virtual Ethernet, Shared Ethernet and SMT, we therefore also recommend that you use AIX 5L V5.3 for the Virtualization Engine agents and install the latest operating system fixes.

The required operating system fixes are as follows:

- The fixes for the Virtualization Engine components are described in Table 2-22.

Table 2-22 Required operating system fixes for Virtualization Engine

APAR	Description
IY78923	C++ Streams affected by setlocale() routine
IY79636	EWLM Fix AeWLM MS may hang when ewlm_get_activedata returns 1 continuously
IY76140	Child process unable to load module
IY78114	EWLM Fix ewlm_query() reports zero for phisical_memory_size and paging_space_faults
IY78228	These values are inaccurate at best and wrong on shared processor logical partitions (LPARs)

APAR	Description
IY80573	EWLM Fix lpar_get_info() reports wrong max pcu value
IY80619	EWLM Fix VP was disabled by kernel

Note: You can get the AIX authorized program authorization reports (APARs) on the Web at:

http://www-912.ibm.com/eserver/support/fixes/fixcentral

- Table 2-23 describes the fixes for the IBM Director.

Table 2-23 Required operating system fixes for IBM Director

APAR	File name
N/A	vacpp.70.aix.nov2005.ptf.tar.Z

Others

Table 2-24 describes the recommendation for filesets.

Table 2-24 Required filesets on AIX

Component	Fileset	Source
Enterprise Workload Manager	bos.net.ewlm.rte	AIX CDs
IBM Director	sysmgt.pegasus.cimserver sysmgt.pegasus.osbaseproviders xlC.aix5.0 xlC.rte openssl-0.9.6k-2.aix5.1.ppc.rpm	AIX CDs AIX CDs AIX CDs AIX CDs Linux Toolbox
Virtualization Engine console	-	

Note: The OpenSSL file set is present on the AIX Linux Toolbox CD, but it is also available on the Web at:

http://www-03.ibm.com/servers/aix/products/aixos/linux/download.html

The installation of the required file sets must be done as follows:

- To install the OpenSSL file set:
  a. Download the OpenSSL package from the Linux Toolbox and copy it to the /tmp directory.
  b. Run the following command:

```
rpm -i openssl-0.9.6k-2.aix5.1.ppc.rpm
```

- To install the Pegasus CIM server:
  a. Put the AIX 5L V5.3 Expansion Pack CD into the CD-ROM drive.
  b. Run the following command:

```
smitty install_latest
```

Archived

c. Enter the software location directory, and select the following two filesets:
   - sysmgt.pegasus.cimserver
   - sysmgt.pegasus.osbaseproviders

d. Accept the license agreement.

You can check the existence of the required filesets by using the following commands:

```
lslpp -l sysmgt.sguide*
lslpp -l sysmgt.websm*
lslpp -l xlC.aix50*
lslpp -l xlC.rte*
```

2.3.2 Linux on POWER

This section describes the hardware and software requirements for Linux on POWER.

Hardware requirements

Table 2-25 describes the recommended hardware for the management servers running on Linux on POWER.

Table 2-25 Hardware requirements for Linux on POWER

Component	Processor	Memory	Disk Space	Swap space
Enterprise Workload Manager	1.4 GHz	512 MB	6 GB	
IBM Director	1.5 GHz	512 MB	325 MB	
Resource Dependency Service	POWER5	4 GB	2.5 GB	2x memory
Virtualization Engine console	1.4 GHz	1.75 GB	3 GB	

Archived

Software prerequisites

Table 2-26 lists the required Red Hat Package Managers (RPMs) needed in order to install the Virtualization Engine components.

Table 2-26 Required rpms for Linux on POWER

Component	Required RPM
IBM Director	compat-libstdc++-296
	compat-libstdc++-33
	librtas-1.2.4-1.ppc64.rpm*
	ppc64-utils-2.5.3-1.ppc64.rpm*
	lspvd-0.14.1-1.ppc.rpm*
Tivoli Directory Server	compat-libstdc++-296
	compat-libstdc++-33
	compat-gcc
	compat-gcc-c++
	libgcc
	glibc-kernheaders
	glibc-headers
	glibc-devel

Note: For Director, you can find the marked RPMs on the Web at:

`https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html`

There are additional RPMs listed on the site that provide added functionality for Director, but are not required for the installation of Virtualization Engine.

All other RPMs can be found on the Red Hat installation CDs. Install both the ppc and ppc64 versions of the RPMs, if both exist.

Operating system

Table 2-27 lists the supported operating systems.

Table 2-27 Supported operating systems for Linux on POWER

Component	Linux
Enterprise Workload Manager	
Domain Manager	RHEL AS 4.0 SLES 9
Managed Server	Refer to [a]
IBM Director	
Director Server	RHEL AS 4.0 SLES 9

Component	Linux
Director Agent	RHEL AS 3.0 RHEL AS 4.0 SLES 8 SLES 9
Director Console	RHEL AS 4.0 SLES 9
Resource Dependency Service	
Resource Dependency Service	RHEL AS 4.0 SLES 9
Virtualization Engine console	
Virtualization Engine console	RHEL AS 4.0 SLES 9
Director Bridge	RHEL AS 4.0 SLES 9
EWLM Bridge	RHEL AS 4.0 SLES 9

a. The EWLM Agent (or managed server) was not supported on Linux operating systems in the standard release of theVersion 2.1 of the product at the time of the writing of this IBM Redbook. This support has be made available through the Virtualization Engine Fix Packs 2.1.20.

Note: The version of Linux on POWER used in the context of this book was Red Hat Enterprise Linux AS 4.0.

2.3.3 Linux on IBM System x

This section describes the hardware and software requirements for Linux on IBM System x.

Hardware requirements

Table 2-28 lists the hardware requirements for management servers running on Linux on IBM System x.

Table 2-28 Hardware requirements for Linux on IBM System x

Component	Processor	Memory	Disk space	Temporary directory space
Enterprise Workload Manager	2 GHz	512 MB	6 GB	An additional 2 GB in the /tmp directory
IBM Director	1.5 GHz	1024 MB	500 MB	n/a
Resource Dependency Service	1.4 Ghz	4 GB	2.5 GB	An additional 2 GB in the /tmp directory
Virtualization Engine console	1.4 Ghz	1.75 GB	3 GB	An additional 2 GB in the /tmp directory

Archive

Software prerequisites

This section describes the software requirements for Linux on IBM System x.

Operating system

Table 2-29 lists the supported operating systems.

Table 2-29 Supported operating systems for Linux on IBM System x

Component	Linux (32-bit)	Linux (64-bit)
Enterprise Workload Manager		
Domain Manager	RHEL AS 4.0 SLES 9	-
Managed Server	Refer to [a]	-
IBM Director		
Director Server	RHEL AS 3.0 RHEL AS 4.0 RHEL ES 3.0 RHEL ES 4.0 SLES 8 SLES 9	RHEL AS 4.0 RHEL ES 4.0
Director Agent	RHEL WS 3.0 RHEL WS 4.0 RHEL AS 3.0 RHEL AS 4.0 RHEL ES 3.0 RHEL ES 4.0 SLES 8 SLES 9	RHEL WS 3.0 RHEL WS 4.0 RHEL AS 3.0 RHEL AS 4.0 RHEL ES 3.0 RHEL ES 4.0 SLES 9
Director Console	RHEL AS 3.0 RHEL AS 4.0 RHEL ES 3.0 RHEL ES 4.0 SLES 8 SLES 9	RHEL AS 4.0 RHEL ES 4.0 SLES 9
Resource Dependency Service		
Resource Dependency Service	RHEL AS 4.0 SLES 9	-
Virtualization Engine console		
Virtualization Engine console	RHEL AS 4.0 SLES 9	-
Director Bridge	RHEL AS 4.0 SLES 9	-
EWLM Bridge	RHEL AS 4.0 SLES 9	-

a. The EWLM Agent is not supported on Linux operating systems in the standard release of the version V2.1 of the product. This support will be made available later on through Virtualization Engine Fix Packs.

Archive

You have to ensure that you have applied all the most recent operating system fix packages and RPMs. You can find the required RPMs and recommended fixes on the Web at:

http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/osDownloads.html

The Web site is always up to date and you should visit it before any new installation.

Note: The version of Linux on IBM System x used in the context of this book was SUSE Linux Enterprise Server 9 SP3.

IBM Director Server on IBM System x

Before you install IBM Director, you may need to install the following two drivers:

- LM78 device driver
- SMBus device driver

These two drivers ensure that certain IBM Director tasks and functions work properly.

Those drivers are provided with the installation files in the following location:

<*Installation Files*>/drivers/linux

Refer to the following README files for installation details:

- IBMSMB-README.TXT
- IBMLM78-README.TXT

Note: You must install the kernel sources and compilers on the system in order to perform the drivers installation.

Resource Dependency Service

At the time of writing, a specific requirement for the Resource Dependency Service must be added. Resource Dependency Service requires a connection to the Internet in order to be able to retrieve XML schemas from two external repositories to be used by the Web services engine. This connectivity is used both during installation and the bridge startup time. Access is not used again once the machine is up and running, until the next reboot or Bridge restart.

The technical parameters of this connection are as follows:

- Remote Address: schemas.xmlsoap.org
- Remote Port: 80 (HTTP)

Note: The RDS system does not need a direct connection to the Internet. It can be situated behind a router or firewall providing network address translation (NAT) services or behind a proxy server.

If no route to the Internet exists on the Resource Dependency Service system, the impact on the product features will be the following:

- Data collection will be fully operational.
- Managed resources visualization using the graph view will not be operational.

This requirement will be removed in a future release of the Resource Dependency Service. At the time of the writing of this IBM Redbook, no due date is available as yet. The proposed workaround is detailed in 3.2.5, “Considerations before installing Resource Dependency Service” on page 74.

2.3.4 Microsoft Windows on IBM System x / BladeCenter

This section details the hardware and software requirements for Microsoft Windows on IBM System x or BladeCenter.

Hardware requirements

Table 2-30 lists the hardware requirements for management servers running on Microsoft Windows Server® 2003 Enterprise Edition.

Table 2-30 Hardware requirements for Microsoft Windows on IBM System x / BladeCenter

Component	Processor	Memory	Disk Space
Enterprise Workload Manager	2.0 GHz	512 MB	6 GB
IBM Director	1.5 GHz	512 MB	325 MB
Virtualization Engine console	1.4 GHz	1.75 GB	3 GB

Note: Based on these hardware requirements, a single mid-range IBM System x or Blade system may host all the Virtualization Engine components (excluding Resource Dependency Service, not available on Microsoft Windows).

In the context of this IBM Redbook, given the fact that performance was not to be taken into account, installations were conducted on the following hardware: Blade HS20 / 2 CPU 3.20 GHz Intel® Xeon® / 4 GB RAM.

Software prerequisites

This section details the software prerequisites for Microsoft Windows on IBM System x or BladeCenter.

Operating system

Table 2-31 lists the supported operating systems.

Table 2-31 Supported operating systems for Microsoft Windows on IBM System x / BladeCenter

Component	Microsoft Windows (32-bit)	Microsoft Windows (64-bit)
Enterprise Workload Manager		
Domain Manager	Server 2003 Standard Ed. Server 2003 Enterprise Ed.	-
Managed Server	2000 Server 2000 Advanced Server Server 2003 Standard Ed. Server 2003 Enterprise Ed.	-
IBM Director		
Director Server	2000 Server 2000 Advanced Server Server 2003 Web Ed. Server 2003 Standard Ed. Server 2003 Enterprise Ed.	Server 2003 Web Ed. Server 2003 Standard Ed. Server 2003 Enterprise Ed.

Component	Microsoft Windows (32-bit)	Microsoft Windows (64-bit)
Director Agent	2000 Server 2000 Advanced Server Server 2003 Web Ed. Server 2003 Standard Ed. Server 2003 Enterprise Ed.	Server 2003 Web Ed. Server 2003 Standard Ed. Server 2003 Enterprise Ed.
Director Console	2000 Professional 2000 Server 2000 Advanced Server XP Professional Server 2003 Web Edition Server 2003 Standard Ed. Server 2003 Enterprise Ed.	Server 2003 Web Edition Server 2003 Standard Ed. Server 2003 Enterprise Ed.
Resource Dependency Service		
Resource Dependency Service	N/A	N/A
Virtualization Engine console		
Virtualization Engine console	Server 2003 Standard Ed. Server 2003 Enterprise Ed.	-
Director Bridge	Server 2003 Standard Ed. Server 2003 Enterprise Ed.	-
EWLM Bridge	Server 2003 Standard Ed. Server 2003 Enterprise Ed.	-

For all of these operating systems, we recommend that you use the latest level of service pack and fixes.

Note: The version of Microsoft Windows used in the context of this book was Microsoft Windows 2003 Server Standard Edition SP1.

Others

The Virtualization Engine installer has no specific other software prerequisite.

The IBM Director has no mandatory prerequisite. However, if the default database server (Apache Derby) that comes with the product is not appropriate, you must install another database server first.

Table 2-32 lists the supported database servers.

Table 2-32 Supported database servers for IBM Director 5.10

Product	Version	Fix Pack
IBM DB2 Universal Database	8.1[a]	Fix Pack 9a
Microsoft Structured Query Language (SQL) Server	2000	Service Pack 3a
Oracle Server	9.2 / 10g	
PostgreSQL	7.3 / 7.4	

a. Since the book has been written, DB2 8.2 provided with the Virtualization Engine is now supported with Fix Pack 9a.

Note: The database on which the Director Server relies can be changed at any time after the installation. Therefore, it is possible to:

- Keep the default Apache Derby database at installation time.
- Install one of the database servers listed in Table 2-32 and switch from Derby at a later time.

Archived

Archived

3

Installing management servers

This chapter describes the installation of the management servers, the main task of the installation process. This is the second task of the installation process, as shown in Figure 3-1.

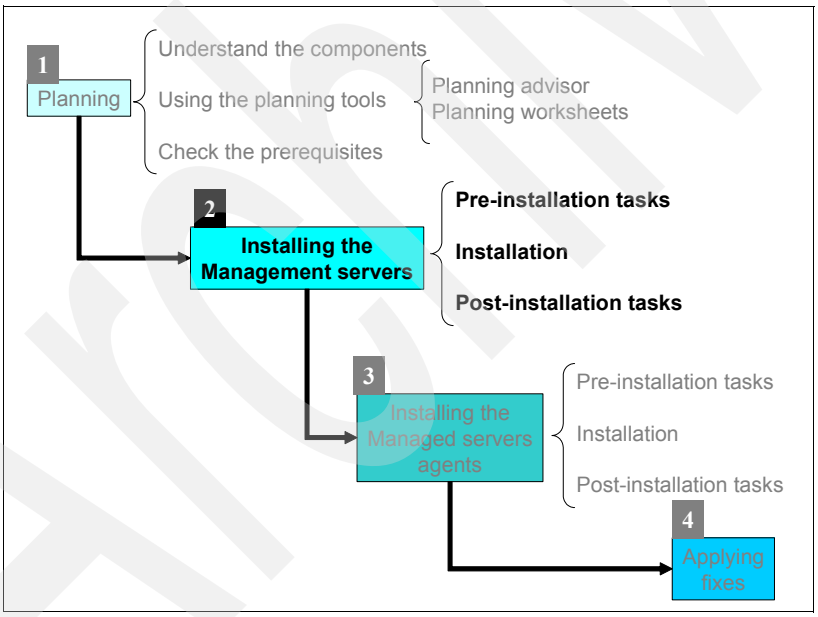


Figure 3-1 Task two: installing the management servers

With Virtualization Engine V2.1, all necessary components are federated under a single installer and do not need to be installed separately as prerequisites. This makes the installation process much easier to understand and much easier to perform than the Virtualization Engine Version 1.

© Copyright IBM Corp. 2006. All rights reserved.

The installation can be divided into three sub-tasks:

- Preinstallation tasks

  For example, the installation files may need some preparation for ease-of-use, user IDs and passwords need to be created, and so on. The installation of Resource Dependency Service requests a specific preinstallation task. This is discussed in 3.2, “Preinstallation tasks” on page 67.

- The installation tasks

  3.4, “Using the Virtualization Engine installer” on page 88 describes the steps of the Virtualization Engine installer. Except on i5/OS, IBM Director has its own independent installation process and you install it independently of the installer, as described in 3.3, “Installing the IBM Director” on page 77.

- Post-installation tasks

  We provide information about checking whether the installation was successful. 3.6, “Post-installation checking tasks” on page 128 provides this information.

Before describing all the tasks, sub-tasks, and steps of the installation, 3.1, “The ITSO environments” on page 65 describes the International Technical Support Organization (ITSO) environments we built to test it for this documentation.

Archived

3.1 The ITSO environments

Multiple environments were set up and tested during our ITSO project. The objectives were to test the installation with different operating systems using different configurations.

Note: The options are numerous and we did not build all the possible configurations. We did not use any IBM System i or IBM System z servers due to the constraints of time and the participants' skills. However, the IBM System i and the IBM System z support the Virtualization Engine V2.1, as explained in 2.1.3, "The products" on page 25.

Other documentation covers these platforms: they are not published at the time of the writing of this IBM Redbook and we recommend you visit the IBM Redbook Web site to get the latest available information at:

`http://www.redbooks.ibm.com/`

In our project, we used the following configurations:

- The first configuration, as depicted in Figure 3-2, used one single BladeCenter with all the Virtualization Engine products and Resource Dependency Service requiring Linux. The operating system was SUSE Linux Enterprise Server 9 SP3. The BladeCenter was an HS20 with two control programs (CPs), 4 GB of memory, and 40 GB of disk space.

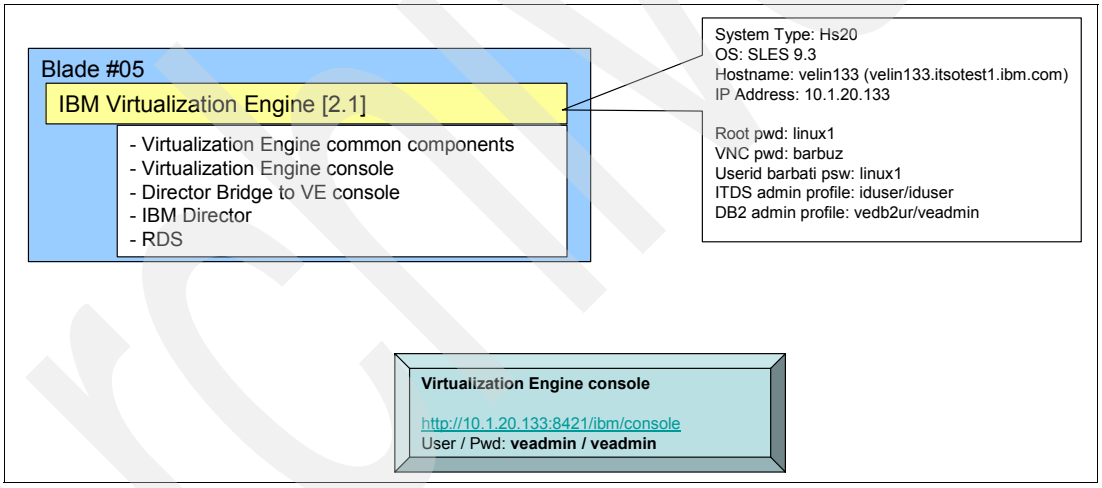


Figure 3-2 Configuration one: one BladeCenter with Linux

- The second configuration, as depicted in Figure 3-3 on page 66, used two boxes. These included one BladeCenter running Windows with the IBM Director, the Enterprise Workload Manager (EWLM) and the Virtualization Engine console, and another

BladeCenter to run Resource Dependency Service, running SUSE Linux Enterprise Server 9.

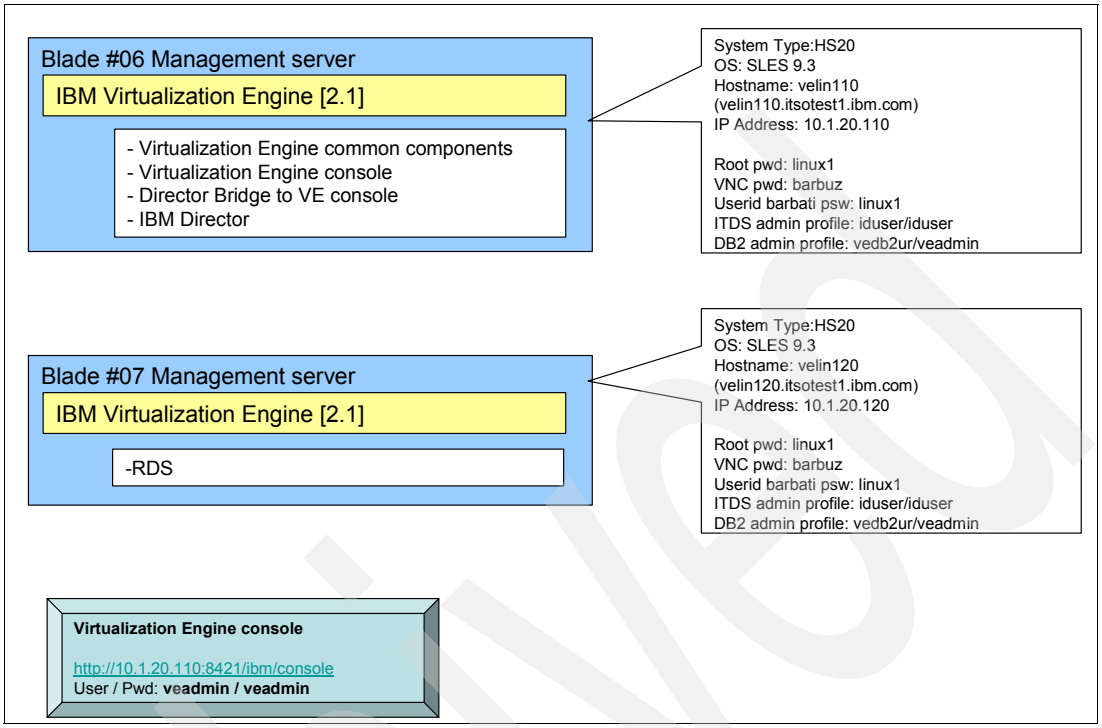


Figure 3-3 Configuration two: one BladeCenter with Windows plus one BladeCenter with Linux

- Using a System p (one p550 with 4 CPs and 16 GB of memory) and multiple partitions, as depicted in Figure 3-4, we set up multiple Virtualization Engine environments.

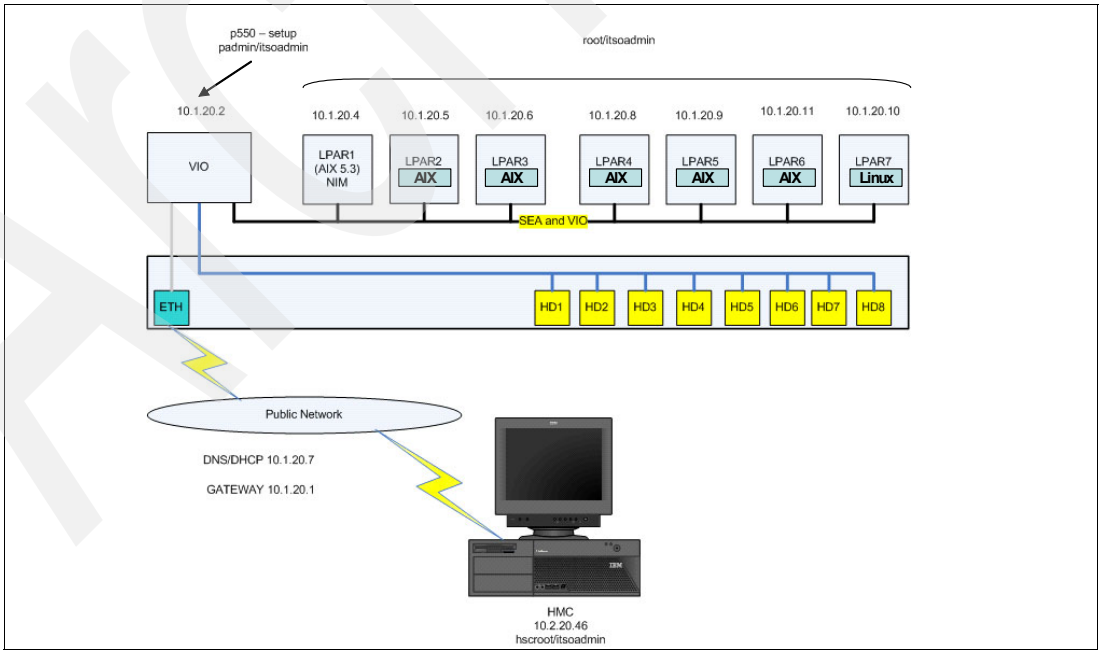


Figure 3-4 The System p configuration

Figure 3-5 describes how the System p was used. Two Virtualization Engine environments were installed:

– In one AIX 5L V5.3 partition, VEAIX2, we installed the IBM Director, EWLM, and the Virtualization Engine console.

– In one Linux partition, pLinux7, we installed the IBM Director, EWLM, the Virtualization Engine console, and Resource Dependency Service. We used Red Hat Version 4.

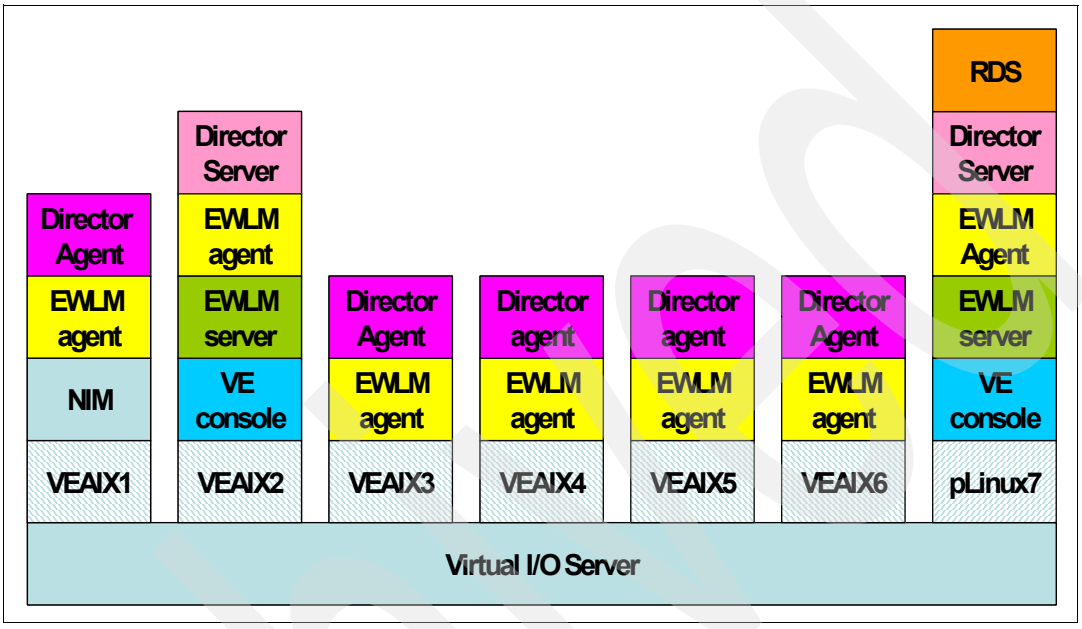


Figure 3-5 Configurations three and four on System p: VEAIX2 and pLInux7

The VEAIX3, VEAIX4, VEAIX5, and VEAIX6 partitions were managed servers and they are not considered during the management servers installation process.

Important: Although there are always a few differences in the installation depending on the operating systems and the options you use, the installation process is almost the same because of the Virtualization Engine V2.1 packaging and the Virtualization Engine installer. In the following sections, we describe the generic process. When operating systems or options need specific attention, the details are provided in these sections.

3.2 Preinstallation tasks

This section describes the preinstallation tasks to be performed. It is important that you understand these tasks before starting the installation process itself.

3.2.1 Preparing the installation files

Most of the time, using the physical media in order to perform the installation is not an option, for various reasons. These reasons could be:

- The access to the physical machine may be difficult because of security constraints or geographical distance.
- No CD-ROM drive or DVD drive is available on the system.

We recommend that you make the installation files contained on the physical media available for the installation in the following way:

1. Choose a system that can host the installation files.

   This system, referred to as the *file server*, must:

   – Reside on the same network as the systems on which the installation is to be performed. You must consider the bandwidth size, as it may badly impact the performance of the installation.

   – Provide a sufficient disk space capacity to host the whole content of all the physical installation media. In our scenario, we used 40 GB for the complete set of files for all the platforms.

   Note: During the installation process itself, some files get unpacked, leading to extra disk space consumption. It is, therefore, necessary to plan to provide twice as much disk space as the capacity required by the installation files gathered from physical media.

   The amount of disk space you need on the target system depends on the services that you select to install. Since unpacking is done one CD image at a time, it should never need more than 1 GB additional storage during the installation than the disk space provided in “Hardware requirements” on page 54.

2. Copy the installation files to the file server.

   You must dump the entire content of the physical installation media onto the file server.

   Note: As the IBM Director provides its own distinct installation process, we recommend that its installation files remain totally separated from the Virtualization Engine installation files.

   It may be suitable to create two different directories on the file server that would host the installation files of each of the installers. In our case, we used the following:

   - IbmDirector
   - IbmVirtualizationEngine

   The name “IbmVirtualizationEngine” was not a good choice and we do not recommend you use a long name because the Virtualization Engine install root is limited to 28 characters. You can only use basic Latin characters. Spaces are not allowed on AIX and Linux. Windows has a 256-character limit on file names and any path names beyond this are truncated.

   Because the directory structure is identical on all the physical Virtualization Engine installation media, the file copy process may issue warnings regarding possible file overwriting. You can safely discard the message or replace the names if requested (we recommend you use the “replace all” option).

3. Share the installation files with the target system.

   The way of sharing the installation files depends on the platform types:

   – If both systems (file server and management server) use Microsoft Windows, we recommend that you use Microsoft File Sharing.

   – If the file server runs under AIX or Linux and the target system runs under Microsoft Windows, we recommend that you use Samba.

– If both servers run under AIX or Linux, we recommend that you use NFS.

Note: You can consider the option of copying the installation files directly onto the target system if technical reasons prevent the use of a file server as described earlier. However, this is less flexible for the following reasons:

- The files will have to be copied on *each* target system.
- The installation files consume extra disk space capacity on the target systems.
- Unpacking the files takes time (though understand there is no unpacking during the copy from the media; the files are only unpacked when running the installer).

The Virtualization Engine installation files are packaged under the Solution Install format. Solution Install packages are made of the following elements:

- Installable unit (IU) of software

  The IU is self-extracting and self-installing (for the installation media, the IU format is used, but it is not packaged into a jar file and it is not self installing).

- Installable unit deployment descriptor (IUDD) eXtensible Markup Language (XML) file

  The IUDD file describes the package content, its technical requirements, and its dependencies. The IUDD file is always named <*Component*>_packagedIU.xml.

- Artifact XML file

  The artifact file describes the installation instructions of the package (not used for installing the services; used for the remote deployment of Director agents and EWLM managed servers).

The installation files are placed in the following directory structure:

- FILES: This directory contains the installable units corresponding to all the technical components constituting the IBM Virtualization Engine.
- META-INF: This directory contains the IUDD files and artifact files linked to the IUs present in the FILES directory.

3.2.2 Creating the operating system users and groups

The following users are required during the installation process. Consider using the appropriate planning worksheets to better plan your environment before the installation. Filling in all these profiles in the user planning worksheet will help you during the installation.

- The IBM Director Server user

  This user is the owner of the IBM Director Server process or service. On Microsoft Windows systems, the user must belong to the administrator group. On Linux, the user root is the default and there is no need to create a specific user.

Important: This user gets the Director Server administration privileges. You should use this to connect to the Director Console after the installation process.

By default, the local operating system account used for connecting to the system will be proposed by the Director Console.

Archivo

- DB2 instance owner:

  This user is the owner of the DB2 database instance hosting the Virtualization Engine Directory Server. If requested to do so, the Virtualization Engine installer automatically creates the user when required. This depends on the operating system:

  - On Microsoft Windows systems, the user must belong to the administrator group.
  - On AIX and Linux systems, the user must belong to a specific group. The Virtualization Engine installer may automatically create the group if requested to do so.

  The naming proposal of the Virtualization Engine installer is the following:

  - User: vedb2ur
  - Group (if applicable): vedb2grp

- Directory Server instance owner:

  This user is the owner of the Directory Server instance hosting the various Virtualization Engine repositories. This depends on the operating systems:

  - On Microsoft Windows systems, the user must belong to the administrator group.
  - On AIX and Linux systems, the user must belong to a specific group.

  In order to remain consistent with the way the DB2 instance owner is named by the installer, we recommend that you use the following names:

  - User: vetdsur
  - Group (if applicable): vetdsgrp

- WebSphere Application Server Security administrator:

  This user is the administrator of the WebSphere Application Server security.

  In order to remain consistent with the way that the DB2 instance owner is named by the installer, we recommend that you use the following names:

  - User: vewasur
  - Group (If applicable): vewasgrp

  A veadmin user is created by the installation wizard on the directory server; this ID is used as the WebSphere admin ID and is named veadmin; the name of this ID cannot be changed.

- EWLM Control Center users:

  EWLM Control Center users are operating system users. EWLM Control Center users are assigned one of the following roles:

  - Administrator
  - Operator
  - Monitor

  It is necessary to create one user for the administrator role as a starting point.

Attention: The installer does not create the ITDS instance owner and group: the user must create it (see 3.2.4, “Platform-specific considerations” on page 72). If the user and group for the ITDS instance owner have exactly the same name, this will cause the ITDS installation to fail.

Archived

3.2.3 General considerations

This section provides information that is applicable to all types of platforms. The following considerations are applicable to all types of platforms:

- Before the installation, ensure that your systems meet the minimum hardware and software requirements. The planning advisor tool can help you understand your needs.
- Installing all Virtualization Engine products on a single system and operating system can simplify the installation process and can avoid potential errors and the need to perform extra post-installation configuration tasks. However, before installing all the products on a single system, ensure that it meets the minimum hardware requirements. Other considerations may prevail, like performance, availability, skills, and so on.
- Virtualization Engine components must be able to resolve host names consistently for systems within the Virtualization Engine environment. If host names do not resolve correctly, the installation will fail. Implementing a Domain Name System (DNS) is a solution. Even the name resolution via local host tables (/etc/hosts file) is acceptable. Single sign on relies on a topology with a consistent host name resolution.

  Note: We recommend that you use the fully qualified name during the installation.

  Note: A tool is available on the Virtualization Engine fix Web site to assist with host name resolution. Refer to:

  `http://www14.software.ibm.com/webapp/set2/sas/VirtualizationEngine/DNSConsiderations.html`

- If you decide to install the Virtualization Engine products on multiple systems, you must consider the order in which you will install the products. If you are installing any combination of the management collection products, in addition to the IBM Director, follow this order when you install the products:
  a. Install the directory server (ITDS) and DB2 first (other common components are installed as needed when services are selected).
  b. Install the IBM Director, which may use the DB2 installed in the previous step.
  c. Install the other Virtualization Engine services that include the Virtualization Engine console plus the console bridges, EWLM, and Resource Dependency Service.
- We recommend that you install a console bridge and its associated management source on the same system. When you install both the console bridge and its associated management source on the same system, you can simplify the installation process, avoid potential errors, and avoid the need to perform extra post-installation manual configuration.
- To perform the installation, you must use the system administrator user ID.
- The Virtualization Engine uses specific Internet protocol (IP) network ports by default (for example, the Directory servers by default listen on transmission control port (TCP) 389). If you choose to use the default port numbers, you must ensure that they are available for use.
- We recommend that you make a backup (for example, using the **`mksysb`** command on an AIX environment) before the installation starts. The installation may take some time. If the installation fails, it may be easier to restore the operating system image than to run the uninstaller and start the installation again.

- Ensure that all the hardware prerequisites and the authorized program analysis reports (APARs) available have been installed. For example, on AIX, you can check if the selected APARs have been installed correctly by using the following command:

  `instfix -ivk <apar number>`

- Before starting the installation, check if the host resolving works correctly. Host name resolution is a frequent problem in the Virtualization Engine environment. If you are not using an external DNS server:

  a. Open the /etc/netsvc.conf file, and add `hosts=local` at the end. The result will be that the operating system will resolve the host names to IP addresses by using the /etc/host file.

  b. Define /etc/hosts for the domain manager server and management server, for example:

     - `10.1.20.6 veaix3.itsotest1.ibm.com veaix3`
     - `10.1.20.5 veaix2.itsotest1.ibm.com veaix2`

3.2.4 Platform-specific considerations

This section provides specific platform information when planning the installation of the Virtualization Engine (faced in our ITSO environment).

Windows systems prerequisite considerations

There is no further prerequisite information about a Windows operating system before deploying the Virtualization Engine.

AIX systems prerequisite considerations

There is no further prerequisite information for an AIX operating system before deploying the Virtualization Engine.

Linux on IBM System x prerequisite considerations

The following considerations apply to a Linux environment running on an System x or a BladeCenter:

- Many Virtualization Engine components use a directory server (ITDS) to store authentication information. Before initiating the installation, you should create one user profile to be used as the directory server instance owner. This user ID is used as the owner of all of the directory server installation objects. The user ID has the following considerations:
  - The user ID must exist before starting the installation of the directory server.
  - The user ID must have the user root as a member of the user's primary group.
  - The user ID must have a home directory that the primary group owns and the primary group must have read, write, and execute access to the home directory.

  Example 3-1 shows the commands used to create a new user ID named vetdsur with the same password. The home directory is created and named vetdsur. RWX authorities are given to the user and groups.

Example 3-1 Create user ID for directory server

```
# groupadd vetdsgrp
# useradd -d /home/vetdsur -m -g vetdsgrp -s /bin/bash vetdsur
# passwd vetdsur
# groupmod -A root vetdsgrp
```

```
# chmod -R 775 /home/vetdsur
```

Note: The Virtualization Engine Common Component installation process will also automatically create other system user profiles. Please be aware that it can potentially represent a security issue if it is not managed according to the enterprise rules.

- You can determine whether you have an existing directory server, such as OpenLDAP, installed on your system. To determine if you have an existing directory server, run the following command:

  ```
  # rpm -qa | grep -i ldap
  ```

  This identifies any installed directory server that includes “ldap” in its application name.

  Then determine whether the existing directory server is using port 389 with the following command:

  ```
  # netstat -an
  ```

  This identifies which ports are actively in use on your system. If the existing directory server is not currently running, check the configuration properties for the directory server to determine whether the directory server is configured to use port 389.

- During the installation, there is the possibility that the online Help windows will not be displayed.

  Online Help windows are Web pages in JavaServer™ Pages™ (JSP™) and an appropriate browser is required to display them. If the installation is performed on a brand-new installed machine and no additional browser is installed, and there is only the basic in distribution, there will be no possibility of reading the Help windows. The following browsers are enabled:

  – Mozilla-based browser 1.7 or higher
  – Firefox 1.0 or higher

  Accessing and browsing inside the Virtualization Engine console requires a browser with the same characteristics. Otherwise, you will not be able to use all the console functions.

- Once the installation is initiated from a shared mounted directory, you should modify the *autorun* script, as shown in Example 3-2; otherwise the installation fails to start because it does not find the correct path for the product’s files and cannot continue.

Example 3-2 script autorun

```
#!/bin/sh
#
# IBM Virtualization Engine Setup and Install autorun script - Linux for IBM
System x
#source="`mount | grep udf | awk '(NR==1){print$3}'`"
source="/VE2code"
if [ "$source" == "" ]; then
   source="`mount | grep iso9660 | awk '(NR==1){print$3}'`"
   if [ "$source" == "" ]; then
      source="`mount | grep subfs | grep cd | awk '(NR==1){print$3}'`"
   fi
fi
```

A variable source must contain the starting path where the installation file is located. By default, it is set to execute installation by mounting the CD-ROM or the iso image. If you have to perform a different installation, you must change this value. As shown in

Example 3-2 on page 73, in our ITSO environment, we had to put the variable named `source` as the value `/VE2code`. This, in our case, was the mount point of the NFS shared directory.

Important: IBM does not recommend changing the autorun file. The InfoCenter provides instructions on how to run multiple installations.

- Ensure that all the hardware prerequisites and that the AIX APARs described before in this IBM Redbook have been installed. You can check if the selected APARs have been installed correctly by using the following command:

  ```
  instfix -ivk <apar number>
  ```

3.2.5 Considerations before installing Resource Dependency Service

At the time of the writing of this IBM Redbook, Resource Dependency Service requires a working connection to the Internet to download schema files, when you are installing Resource Dependency Service and when Resource Dependency Service is started, from the Web at:

`http://www.w3.org/`
`http://schemas.xmlsoap.org/`

This may cause a problem for clients who are behind a secure firewall, or who do not have access to the Internet. The following solutions are available:

- Set up the firewall so that the access to these specific Web sites can be permitted.
- Allow the Resource Dependency Service machine to download the schema files that are required to a local Web server.

Both solutions have been tested in the ITSO environment, and they are both working well. In the following section, we detail the latter solution.

Workaround steps

This solution setup requires the following steps:

1. Install Apache 2 on a local server.
   - On an SLES 9 system, install the following RPMs:
     - "apache2-2.0.49-27.8
     - "apache2-worker-2.0.49-27.8
     - "libapr0-2.0.49-27.8
   - On an RHEL 4 system, install the following RPMs:
     - "httpd-suexec-2.0.52-19.ent
     - "httpd-2.0.52-19.ent
     - "httpd-manual-2.0.52-19.ent
2. Download the schemas.

   On a Linux machine that has access to the Internet, run the following commands:

   ```
   wget -m --no-host-directories http://www.w3.org/2001/xml.xsd
   wget -m --no-host-directories
   http://schemas.xmlsoap.org/ws/2004/08/addressing/
   ```

   Now that you have downloaded the schemas, copy them to the “Apache 2” machine, in the default directory for Apache2. SLES 9 creates the default server in the /srv/www/htdocs

directory. This is where you should place the schema files that Resource Dependency Service needs. For RHEL, replace /srv/www/htdocs with /var/www/html.

Now you have the schema files downloaded with the following directory structure:

– /srv/www/htdocs/2001/xml.xsd

– /srv/www/htdocs/ws/2004/08/addressing/index.html

On SLES 9, Apache runs as user wwwrun, so change the file ownership to that user using the following command:

```
chown -R /srv/www/htdocs wwwrun.
```

On RHEL, Apache runs as user apache, so change the file ownership to that user using the following command:

```
chown -R /var/www/html apache.
```

3. Configure and start Apache 2 with the following steps:

   a. Edit the Apache 2 configuration file in the following locations:

      - On SLES 9, configure the default server file that is located in /etc/apache2/default-server.conf.
      - On RHEL, the Apache 2 configuration file is located in /etc/httpd/conf/httpd.conf.

   b. Because the schema files are XML-based, ensure that the Multipurpose Internet Mail Extensions (MIME) types tell the Resource Dependency Service that you are sending text or XML. Add the configuration shown in Example 3-3 inside the directory directive (the comments have been removed for brevity).

Example 3-3 Apache 2 configuration

```
<Directory "/srv/www/htdocs">  Options None
        AllowOverride None
        Order allow,deny
        Allow from all
AddType text/xml .html .xsd
</Directory>
```

   c. Start Apache as follows:

      - On SLES, run the following command:

        ```
        /etc/init.d/apache2 start
        ```

      - On RHEL, run the following command:

        ```
        /etc/init.d/httpd start
        ```

4. Configure the RDS Machine

   On the machine that hosts the Resource Dependency Service, edit /etc/hosts to add the following:

   – *vv.xx.yy.zz*. www.w3.org
   – *vv.xx.yy.zz* schemas.xmlsoap.org

   Replace *vv.xx.yy.zz* with the IP of the machine that Apache 2 was installed on.

Summary

At this point, you can install and start Resource Dependency Service. Since Resource Dependency Service will pick up the schema information from the local Web server, it will install and start without problem.

The Apache 2 access log (that is, /var/log/apache2/access_log on SLES and /var/log/httpd/access_log on RHEL) should have messages similar those shown in Example 3-4. The IP address making the request should match the machine that Resource Dependency Service has been installed on.

Example 3-4 Messages when Resource Dependency Service is started

```
129.40.150.48 - - [21/Mar/2006:06:13:42 -0500] "GET /2001/xml.xsd HTTP/1.1" 200
5829 "-" "Java/1.4.2"
129.40.150.48 - - [21/Mar/2006:06:13:42 -0500] "GET /ws/2004/08/addressing
HTTP/1.1" 301 348 "-" "Java/1.4.2"
```

Troubleshooting

If Resource Dependency Service does not start, or if it takes a very long time to start, it is possible that it cannot find the files. There are a few things to check:

- Ensure that the Resource Dependency Service machine can resolve the IP addresses of the Apache 2 machine for `www.w3.org` and `schemas.xmlsoap.org` using the **`ping`** command, as shown in Example 3-5.

Example 3-5 Resolving the IP addresses

```
#ping   www.w3.org
-> vv.xx.yy.zz
#ping schemas.xmlsoap.org
-> vv.xx.yy.zz
```

- Make sure Apache is running as follows:
  - On SLES, run the following command:

    ```
    # ps -def |grep apache
    root  7225 1 0 06:14 ?  00:00:00 /usr/sbin/httpd2-worker -f
    /etc/apache2/httpd.conf wwwrun 7226  7225 0 06:14 ? 00:00:00
    /usr/sbin/httpd2-worker -f /etc/apache2/httpd.conf
    wwwrun 7227  7225 0 06:14 ? 00:00:00 /usr/sbin/httpd2-worker -f
    /etc/apache2/httpd.conf
    wwwrun 7228 7225 0 06:14 ? 00:00:00 /usr/sbin/httpd2-worker -f
    /etc/apache2/httpd.conf
    ```

  - On RHEL, run the following command:

    ```
    # ps -def |grep httpd
    root       986     1 36 15:54 ?        00:00:03 /usr/sbin/httpd
    apache     991   986  3 15:55 ?        00:00:00 /usr/sbin/httpd
    apache     992   986  3 15:55 ?        00:00:00 /usr/sbin/httpd
    apache     993   986  0 15:55 ?        00:00:00 /usr/sbin/httpd
    ```

- Ensure that the files are accessible as follows:
  - Open a Web browser and go to:

    `http://`*`vv.xx.yy.zz`*` /2001/xml.xsd`
    `http://`*`vv.xx.yy.zz`*` /ws/2004/08/addressing/`

    (Replace *vv.xx.yy.zz* with the IP address of the Apache machine.)

    This will load the schemas in the browser.

  - Open a Web browser on the Resource Dependency Service machine and go to:

    `http://www.w3.org/2001/xml.xsd`
    `http://schemas.xmlsoap.org/ws/2004/08/addressing/`

This will load the schemas in the browser.

Additional configuration consideration

Resource Dependency Service only supports running on a Linux operating system. Therefore, it is possible to install the Apache server on the same machine. As an added security measure, you can lock down the Apache server so that it honors requests only from the local machine.

You can do this by setting the Listen directive in the Apache 2 configuration file to `Listen 127.0.0.1:80`. This tells Apache to only allow connections to port 80 from the 127.0.0.1 IP.

On an SLES 9 machine, the Listen directive is located in /etc/apache2/listen.conf.

On an RHEL machine, the Listen directive is located in /etc/httpd/conf/httpd.conf.

Replace the line that contains `Listen 80` with `Listen 127.0.0.1:80`.

Once this is done, you can follow steps two and three, as described in “Workaround steps” on page 74. Then step four will need to be changed, as shown in Example 3-6.

Example 3-6 RDS configuration changes

```
127.0.0.1  www.w3.org
127.0.0.1  schemas.xmlsoap.org
```

3.3 Installing the IBM Director

The IBM Director has its own independent installation process. If you intend to install it and other components of the Virtualization Engine as well, we recommend that you install the IBM Director first. This section describes this IBM Director installation process.

The Director Server installation can be performed in two ways:

- Through a graphical installation wizard

  This mode is available on Microsoft Windows systems only.
- In silent mode

  This mode is valid for all types of systems, including Microsoft Windows systems.

3.3.1 Graphical installation process

This section describes the installation of the IBM Director using the graphical wizard.

> **Reminder:** The installation through the graphical wizard is available on Microsoft Windows systems only.

The wizard is started with the following command:

`<Installation Files Directory>\setup.exe`

Archive

In this command, <*Installation Files Directory*> is the directory where the code of the IBM Director has been installed. The wizard goes through the following installation steps:

1. The IBM Director is made of several components. The first window, shown in Figure 3-6, allows you to select the components you want to install. Select **Install IBM Director Server**. This is the first component you should install, if nothing has been installed before.

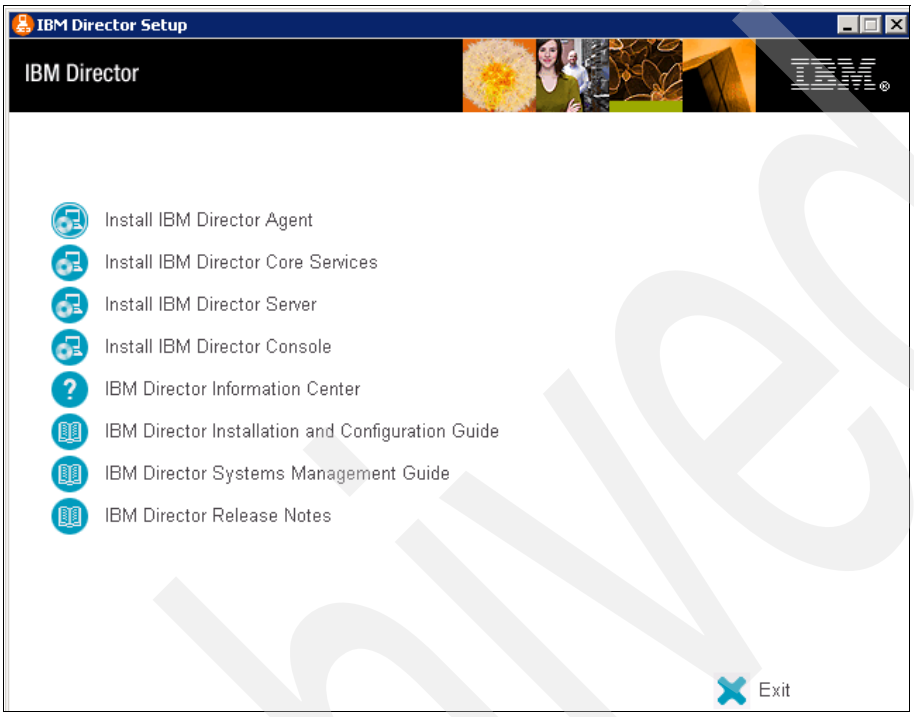


Figure 3-6 IBM Director setup

2. Once you have chosen the IBM Director Server, the actual installation process starts. A welcome window appears. Click **Next**.
3. The licences agreement window appears. You should accept it and click **Next**.
4. The features and installation directory selection window appears.

   By default, the following components are selected:

   – IBM Director Server
   – Level 2: IBM Director Agent
   – Level 1: IBM Director Core Services
   – IBM Director Console
   – IBM System x support and tasks

   Note: The Director Server always comes with the Director Agent (including Director Core Services) and the Director Console. These components cannot be deselected.

   Additional optional components may be selected if necessary. They will extend the base features of the IBM Director Server:

   – IBM Director Remote Control Agent
   – BladeCenter Management Extension
   – Rack Manager

   For details about the nature of each feature, refer to *Implementing IBM Director 5.10*, SG24-6188.

Archived

In our case, as shown in Figure 3-7, we did not choose any additional components. Once you have made your selection, click **Next**.

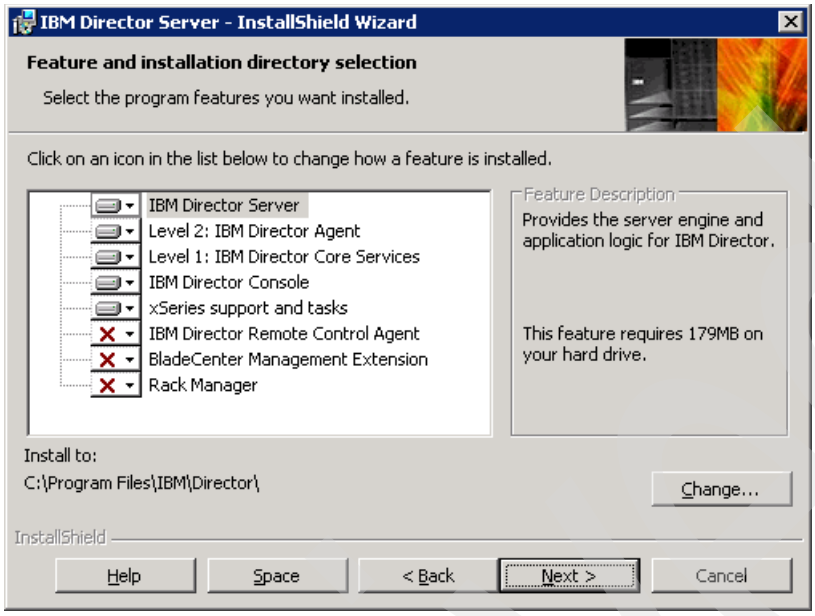


Figure 3-7 Feature and installation directory selection

5. The IBM Director service account information window appears.

   This window shown in Figure 3-8 gathers the elements related to the local account under which the IBM Director Server service will run after installation.

   You should have created the user during the preinstallation phase. (Refer to 3.2.2, “Creating the operating system users and groups” on page 69).

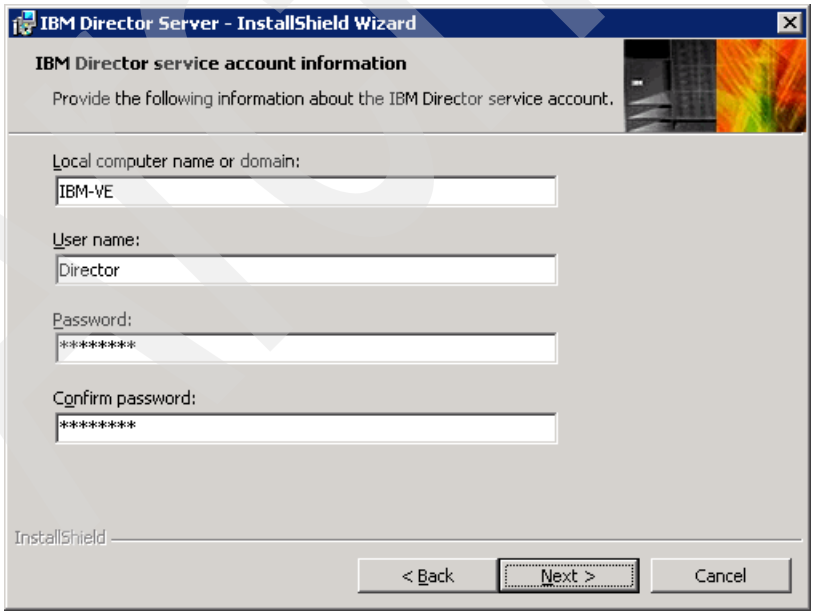


Figure 3-8 IBM Director service account information

6. The encryption settings window appears.

   Data transmissions between the IBM Director Server and the IBM Director agents can optionally be encrypted. This window shown in Figure 3-9 allows you to give the encryption algorithm you want to use, should you want to use one. The default values are acceptable.

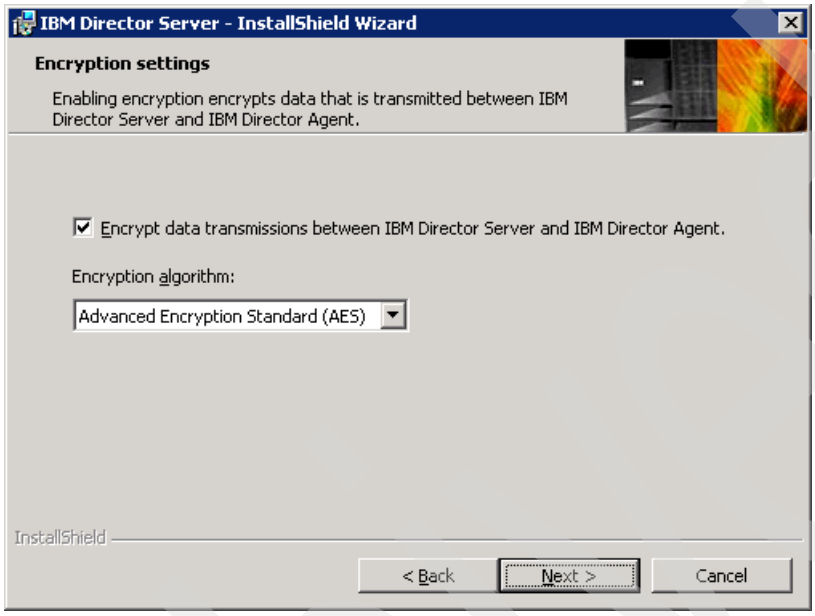


Figure 3-9 Encryption settings

7. The Software Distribution settings window, shown in Figure 3-10, defines the directories where the software packages, created and distributed with the software distribution feature, will be stored. The default values are acceptable. Click **Next**.

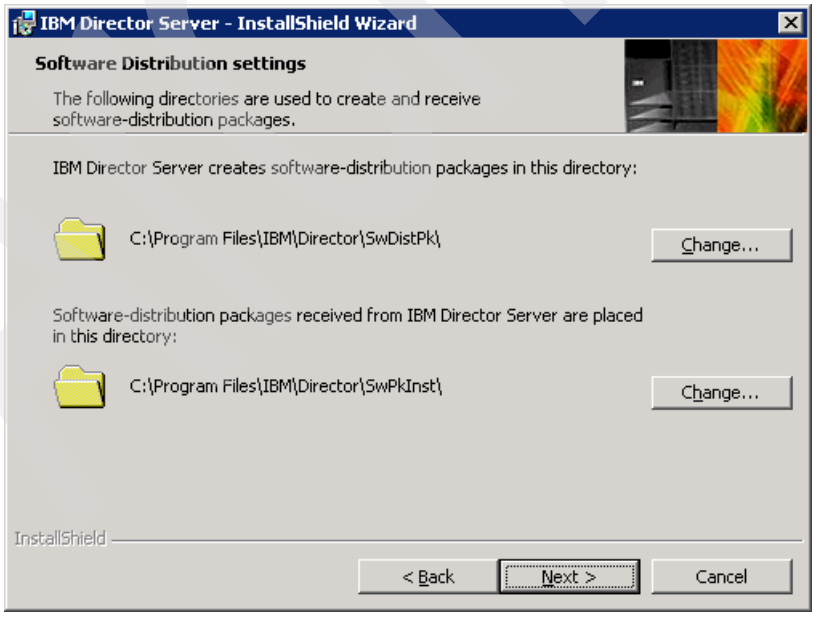


Figure 3-10 Software distribution settings

8. The next window that appears, shown in Figure 3-11, is the last chance to move backward and to modify the selection made before the installation starts. When you are ready, click **Install**.

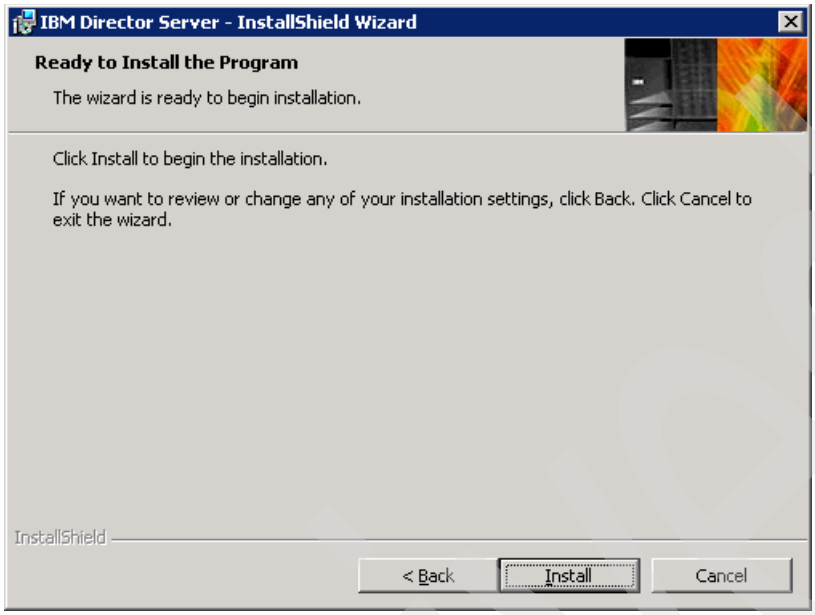


Figure 3-11 Ready to Install the Program

9. The next window, shown in Figure 3-12, informs you of the installation progress.

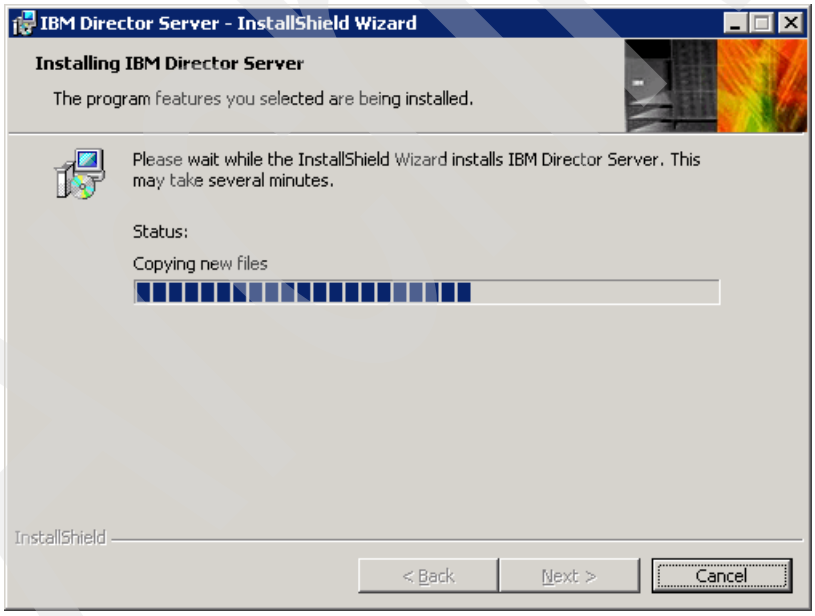


Figure 3-12 Installing IBM Director Server

10. The Network driver configuration window displayed in Figure 3-13 on page 82 allows you to specify the configuration of the IBM Director agent network connectivity, and you have to select the network adapters on which the IBM Director agent should listen. By default, the IBM Director agent listens on all the network adapters of the system. However, you

can modify this configuration in a more restrictive way. When you have made your choice, click **OK**.

Note: This setting particularly makes sense when the IBM Director Server resides on several networks, with one of them being dedicated to system administration and distinct from the actual *production* network. In that case, we recommend restricting the frame of the IBM Director to the administration network only.

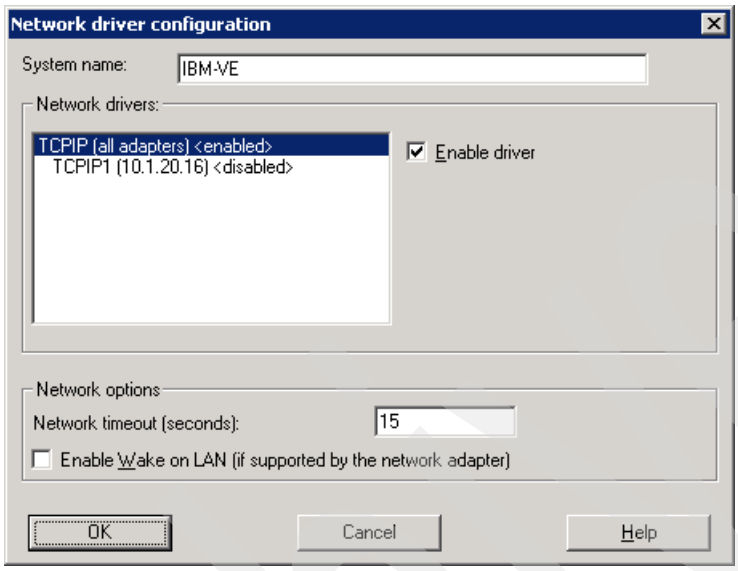


Figure 3-13 Network driver configuration

11. In the IBM Director database configuration window, you have to select a database for the IBM Director and you can choose between Apache Derby, the IBM DB2 UDB, the Microsoft SQL server, or Oracle. Apache Derby is the default database of the IBM Director Server. It is included in the IBM Director Server code and is not a prerequisite to be installed. After making your selection, as shown in Figure 3-14 on page 83, click **Next** to configure your database.

Note: The database on which the IBM Director Server relies on can be changed at any time after the installation. Therefore, keeping the default selection does not prevent you from switching to any of the other supported database servers at a later time.

To reinvoke the IBM Director database configuration wizard after the installation, use the **`cfgdb`** command.

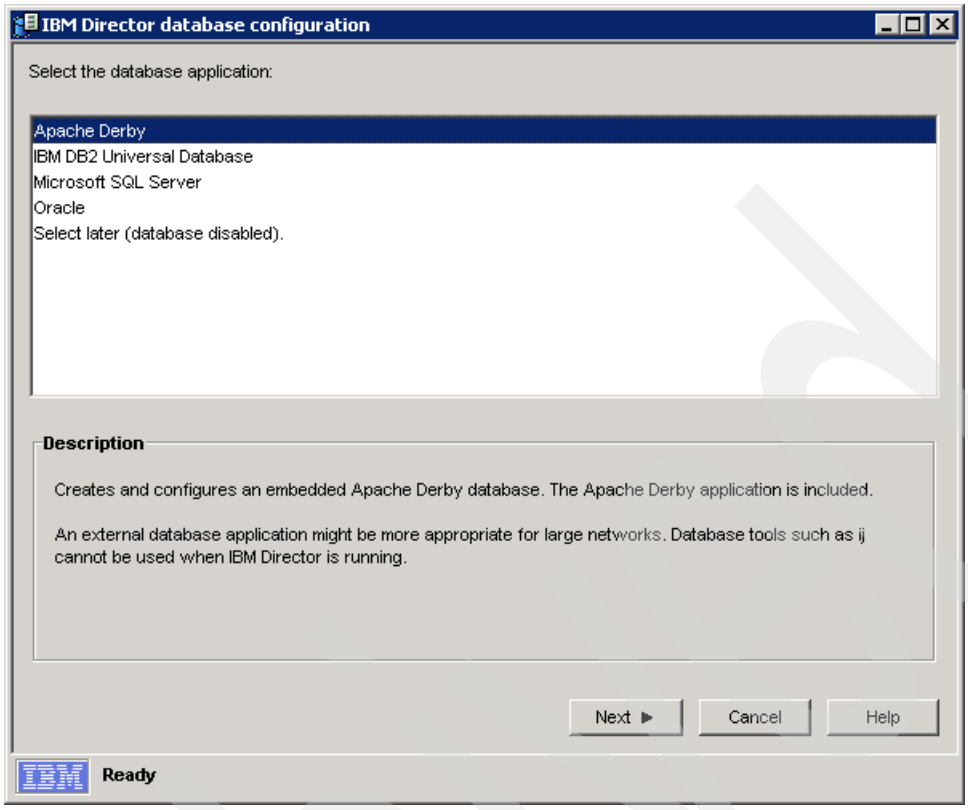


Figure 3-14 IBM Director database configuration - Select the database application

Archived

12. In the database configuration window that appears, the Apache Derby, the choice shown in Figure 3-15, does not require any user input for its configuration. The database name is hardcoded to IBM Director.

Note: If you have selected another database application in the previous step, the installer requires more information at this point, in order to be able to connect to the chosen database. This information normally involves:

- Server name
- TCP/IP port
- Database name
- User ID or password

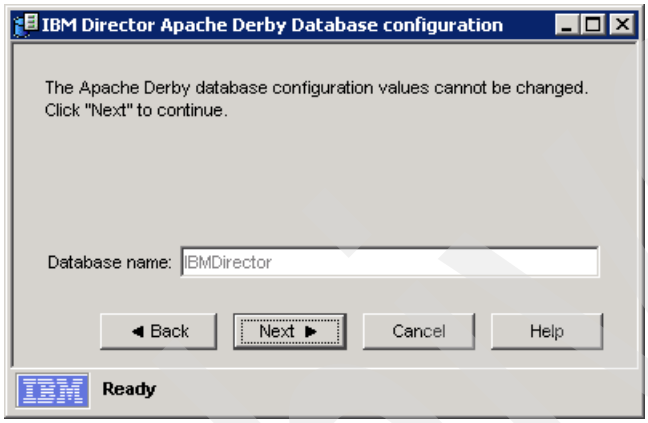


Figure 3-15 IBM Director Apache Derby database configuration

Archived

13. When the installation process is complete, the window shown in Figure 3-16 appears. Click **Finish**.

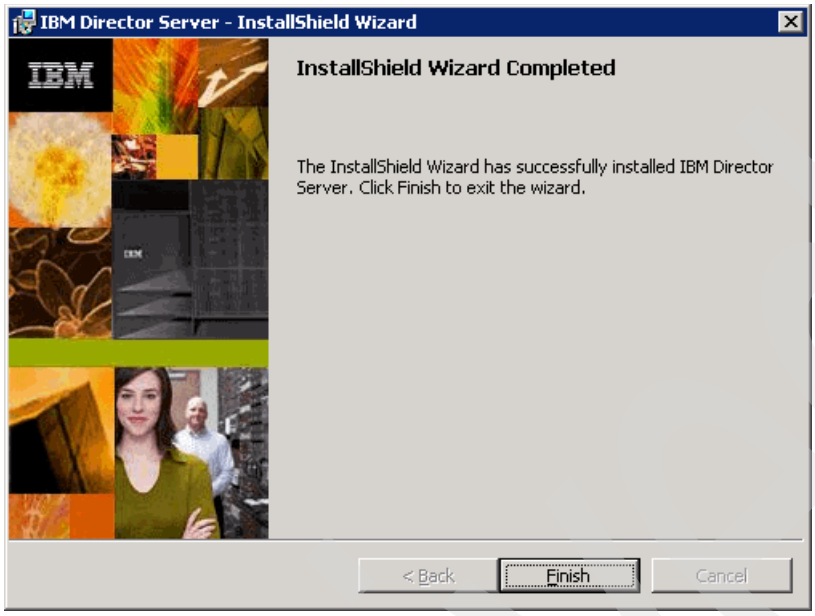


Figure 3-16 InstallShield wizard completed

14. You have to reboot the system at the end of the installation process. You can decide to postpone the reboot. In order to restart the system immediately, select **Yes** in the last window that appears, as shown in Figure 3-17.

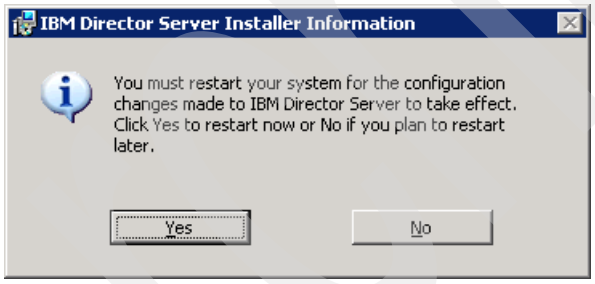


Figure 3-17 IBM Director Server installer information

3.3.2 Silent installation process

This is the ability to run the installation without any user interaction, using a response file to drive the user interactions. You should customize the parameters before being able to start the installation process itself.

Passing the parameters

The silent installation relies on the following parameters:

- An executable file that launches the installation process.

  This file is named:

  - For AIX: dirinstall
  - For Linux: dirinstall
  - For Microsoft Windows: ibmsetup.exe

- A text response file that provides the installer with the appropriate user settings.

  This file is named dirserv.rsp.

Depending on the platform, those two files are located in the following directories:

- For AIX: <*Installation Files Directory*>/director/server/aix/
- For Linux on POWER: <*Installation Files Directory*>/director/server/linux/ppc/
- For Linux on System x: <*Installation Files Directory*>/director/server/linux/i386/
- For Windows: <*Installation Files Directory*>\director\server\windows\i386\

Example 3-7 and Example 3-8 are examples of the files we used in our environment.

Example 3-7 Executable and response file on Linux on System x systems

```
/mnt/source/IbmDirector/server/linux/i386/dirserv.rsp
/mnt/source/IbmDirector/server/linux/i386/dirinstall
```

Example 3-8 Executable and response file on Microsoft Windows systems

```
M:\IbmDirector\director\server\windows\i386\dirserv.rsp
M:\IbmDirector\director\server\windows\i386\ibmsetup.exe
```

These response files must be edited in the following way:

- For Microsoft Windows, the domain, a user ID, and a password should be configured as follows:
  - Domain = <*Director Server host name*>
  - User ID = <*Director Server user name*>
  - Password = <*Director Server user password*>
- For AIX and Linux, specify the database method with the following command:

  `DbmsConfigMethod=Rspfile`
- Whatever the environment, you should choose and configure a database in the following way:

  `DbmsApplication=[ ApacheDerby | DB2 | MicrosoftSQLServer | Oracle | PostgreSQL | noDatabase ]`

  Warning: Exact value spelling differs on Microsoft Windows systems.

  If you choose `DbmsApplication <> ApacheDerby` and `DbmsApplication <> noDatabase`, you have to edit the following parameters:

  - DbmsTcpIpListenerPort=<*TCP/IP DB connection port*>
  - DbmsServerName=<*DB server host name*>
  - DbmsDatabaseName=<*DB name*>
  - DbmsUserId=<*User ID*>
  - DbmsPassword=<*Password*>

Archived

Example 3-9 and Example 3-10 show our installation parameters.

Example 3-9 Parameter values on AIX / Linux systems using DB2

```
DbmsConfigMethod=Rspfile
DbmsApplication=DB2
DbmsTcpIpListenerPort=50000
DbmsServerName=ibm-ve
DbmsDatabaseName=DIRECTOR
DbmsUserId=director
DbmsPassword=director
```

Example 3-10 Parameter values on Microsoft Windows systems using Apache Derby

```
Domain = ibm-ve
UserID = Director
Password = Director
DbmsApplication = Apache Derby
```

Note: For Microsoft Windows systems, do not specify a fully-qualified domain name as a value. Use the short host name for the Domain instead.

Starting the process

In order to launch the silent installation process, you should pass the following parameters to the executable:

- For AIX or Linux, use the following command:

  ```
  -r <Response file filename>
  ```

- For Microsoft Windows, use the following command:

  ```
  silent RSP="<Response file filename>"
  ```

Note: On AIX and Linux systems, the response file parameter is optional. If you do not specify it, the installation will be executed using the default installation settings.

Example 3-11, Example 3-12, and Example 3-13 show the commands we used for our environments.

Example 3-11 Launching the silent installation on AIX systems

```
/mnt/IbmDirector/server/aix/dirinstall -r /mnt/IbmDirector/server/aix/dirserv.rsp
```

Example 3-12 Launching the silent installation on Linux on System x systems

```
/mnt/IbmDirector/server/linux/i386/dirinstall
                                  -r
                                  /mnt/IbmDirector/server/linux/i386/dirserv.rsp
```

Example 3-13 Launching the silent installation on Microsoft Windows systems

```
M:\IbmDirector\director\server\windows\i386\ibmsetup.exe
                silent
                RSP="M:\IbmDirector\director\server\windows\i386\dirserv.rsp"
```

Archived

Note: On Microsoft Windows systems, the system, by default, automatically reboots at the end of the installation process. You can change the default in the dirserv.rsp file if you want.

3.4 Using the Virtualization Engine installer

The Virtualization Engine installation process is a collection of modular pieces assembled into a wizard. The objective of this section is to:

- Describe the different modules.
- Define their technical content.
- Represent the various paths the installer goes through, depending on the components that have been selected for the installation.

3.4.1 Process overview

The installation process is made of six steps preceded by a *prologue* phase and followed by an *epilogue* phase, as described in Figure 3-18.

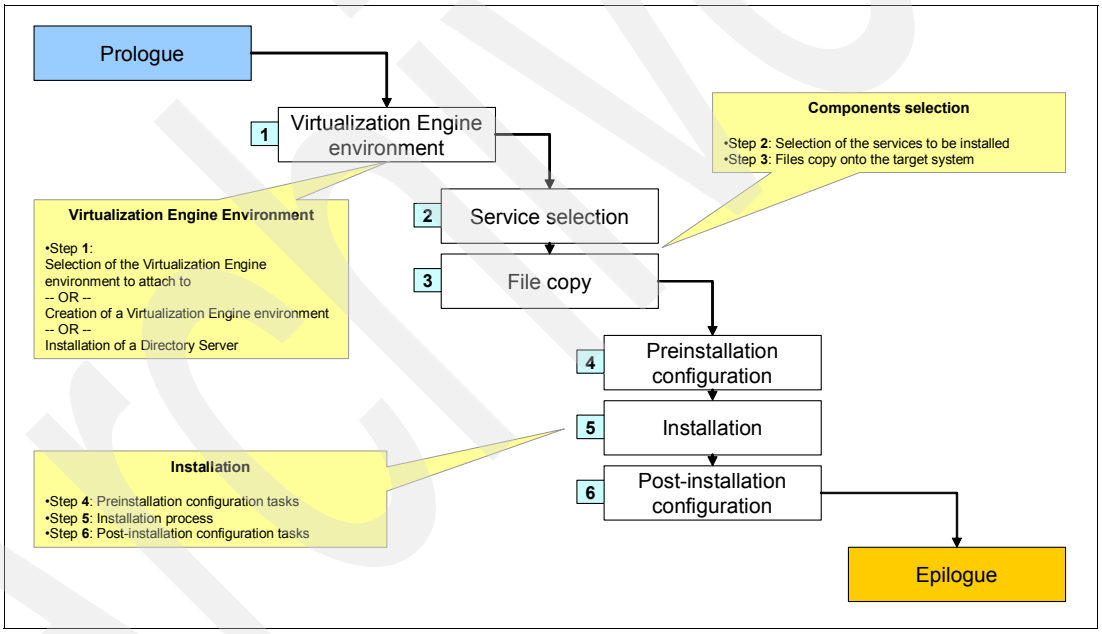


Figure 3-18 The Virtualization Engine installation process

The process can be summarized in the following way:

- The *prologue* task, which involves the welcome and the license agreements primitive steps. This phase does a preliminary check of the environment to detect whether some Virtualization Engine components are already installed.
- The Virtualization Engine Environment phase, shown as step 1 in Figure 3-18, *Establishing the Virtualization Engine environment*, in the Virtualization Engine installation wizard, allows you to:
  - Choose a Virtualization Engine Environment to link to if it does exist.
  - Create a Virtualization Engine Environment if a Directory Server exists and has been selected by you.

Archived

  - Install a Directory Server if none exists, or if you want to install one.

- The components selection, defined as steps 2 and 3 in the wizard in Figure 3-18 on page 88, group:
  - The selection of the components to be installed, named *Selecting the Virtualization Engine services to install* in the wizard.
  - The copy of the installation files to the target system, named *Copying the installation media* in the wizard.
- The installation itself, defined as steps 4 through 6, in the wizard in Figure 3-18 on page 88, can be divided into:
  - A preinstallation configuration step that is the entry of the parameters for each of the selected components, named *Performing preinstallation configuration* in the wizard.
  - The installation process itself, named *Installing the selected Virtualization Engine services* in the wizard.
  - A post-configuration step for each of the components that has been installed, named *Performing post-installation configuration* in the wizard.
- The *epilogue* task displays a final summary of the installation.

Note: The exact content of some of these steps varies, depending on:

- The state of the system in terms of the operating system configuration.
- The state of the system in terms of already existing Virtualization Engine components.
- The components selected for the installation.

3.4.2 The wizard steps

This section describes the steps in the installation process using the installation wizard.

Important: When you use the wizard, the installation process is almost the same, whatever the environment on which you install. In this section, we describe the installation on a Windows environment. If at any step, the installation procedure is different for a Linux or an AIX environment, notes are provided.

Though tests on System i are not described in this IBM Redbook, understand that the installer for an i5/OS system is Windows based and is *very* similar. The primary difference is that the installation is initiated from a personal computer. Thus, all installation files are transferred by the installer from the PC to i5/OS. Then the installer initiates the installation locally on the i5/OS system. From a user's perspective, the flow of windows is very similar to the installation for the other operating systems

The installation is much more integrated than in the Virtualization Engine Version 1 and is quite straightforward.

Prologue

1. When you insert the initial CD, an initial window appears, as shown in Figure 3-19 on page 90, that allows you to start the installation process.

Note: There are multiple "first" CDs in the full media set; the "initial" CD is identified in your pack as "IBM Virtualization Engine common components START HERE - Installation (CD 1 of *n*) for *OS*".

Or you can start the graphical wizard by changing the current directory to the <Installation Files Directory>/FILES and then use the following commands (when files have been copied, for example):

- For AIX: installVEAix.bin
- For Linux on POWER: installVELinPPC.bin
- For Linux on System x: installVELinX.bin
- For Windows: installVEWin.exe

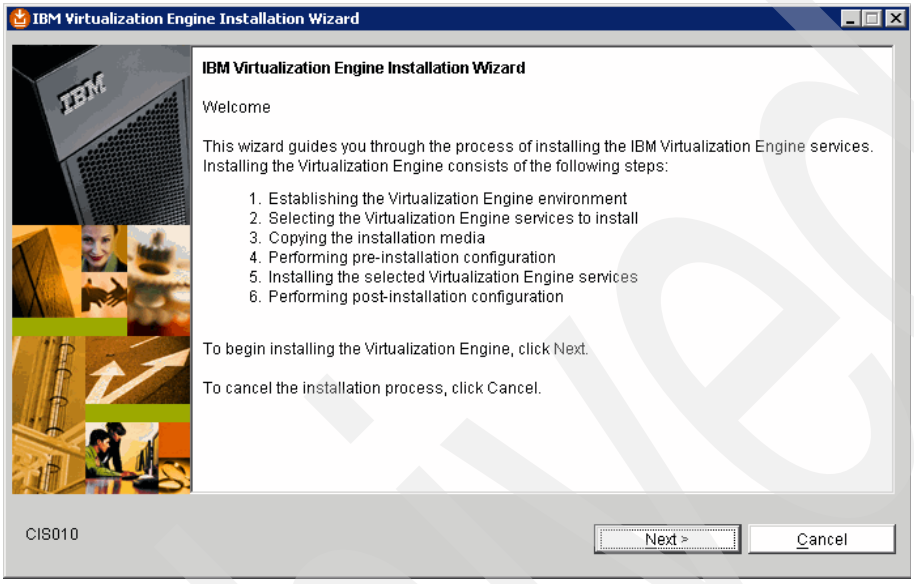


Figure 3-19 The very first Virtualization Engine installation wizard window

2. Press **Next** and the License window is displayed. Each Virtualization Engine service (the IBM Tivoli Directory Server, the WebSphere Application server, the IBM Director with the Virtualization Engine console, the Enterprise Workload Manager Domain manager, and Resource Dependency Service) is under specific licensing. You should accept the terms for each product you intend to install before you go into the installation steps.

3. As shown in Figure 3-20 on page 91, you should decide the installation path. All the selected services will be installed under this path. This is very important information and you should remember your choice. This is something we recommend you write down in the Installation setup planning worksheet. The default path for Linux and AIX is /opt/IBM/VE2 and the default path for Windows is \Program Files\IBM\VE2. You can change this to the value you want (remember the 28-character limit).

Archived

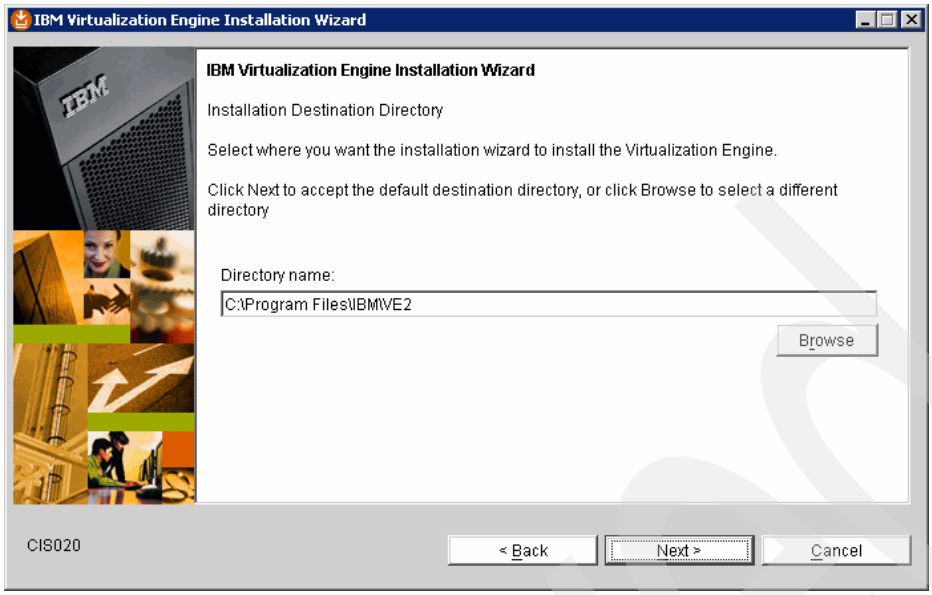


Figure 3-20 The destination directory (Windows example)

4. The Log Definition window that appears, shown in Figure 3-21, allows you to define where the installation and configuration logs will be created. The default location, for AIX and Linux, is the /var/log directory. The default directory location for Windows is \Program Files\IBM\VE2\Logs. This value could be really important in the problem determination process. We recommend you write down the value that you choose in the installation setup worksheet.

   During the installation, the wizard will register all activities under the file installDebug.txt. Specific error messages will be sent to the installLog.txt file. These logs are important. The installation information will be posted and you may need to read the files should you have any problems with during the installation. On AIX and Linux, the **`tail -f/var/log/installDebug.txt`** command is very convenient to watch how the log is growing during the installation process.

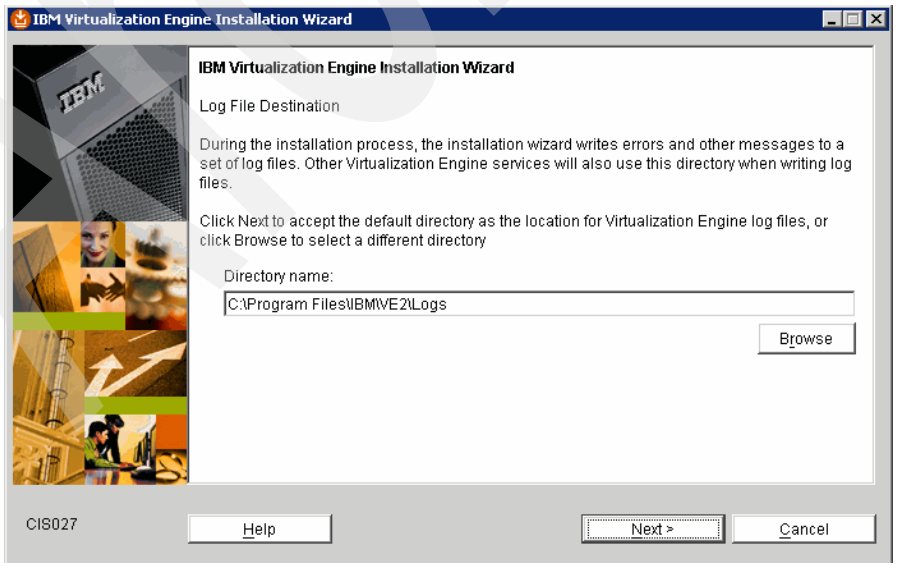


Figure 3-21 Log file destination (Windows example)

5. The last window of this phase is displayed in Figure 3-22. Before starting the installation process, the wizard:
   a. Installs the minimum requirements to run the installation tasks, like a runtime environment, the default directories, the solution installer components, and so on.
   b. Checks your environment to verify if some Virtualization Engine components already exist.
   c. Verifies that your operating system is at least at the minimum supported level.

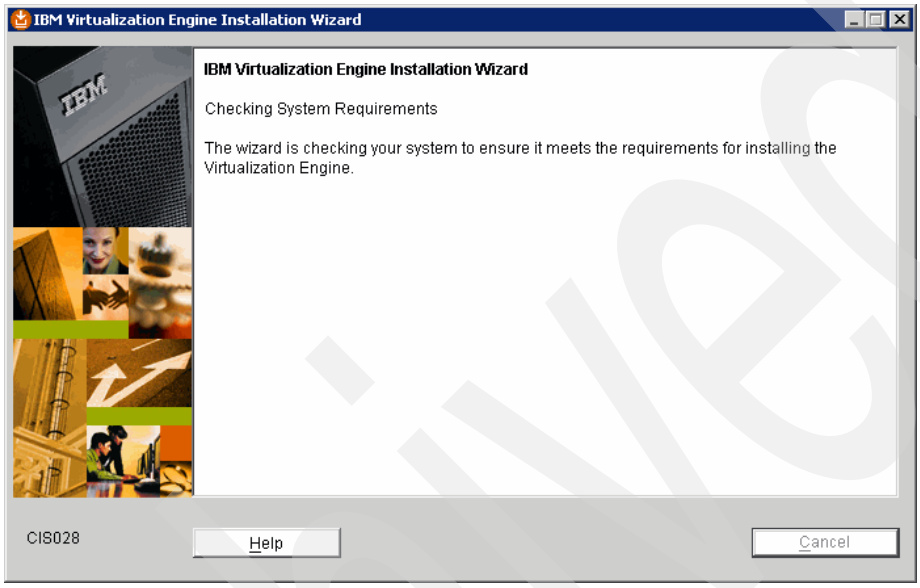


Figure 3-22 Preliminary check by the wizard

The next window that will be displayed by the wizard depends on the status of this check, as explained in the step "Step 1: Establishing the Virtualization Engine environment" on page 92 of ""Installation steps" on page 92".

Installation steps

Note: The core installation steps are made of six steps that are detailed in this section:

- "Step 1: Establishing the Virtualization Engine environment" on page 92 establishes the Virtualization Engine environment.
- "Step 2: Service selection" on page 106 is the service selection.
- "Step 3: File copy" on page 108 copies the code files.
- "Step 4: Preinstallation configuration" on page 109 prepares the installation.
- "Step 5: Installation" on page 118 installs the code according to the user choice and the prerequisites.
- "Step 6: Post-installation configuration" on page 119 provides information about the installation process.

Each step itself can be divided into multiple sub-steps.

Step 1: Establishing the Virtualization Engine environment

A *Virtualization Engine environment* is a logical grouping of systems with Virtualization Engine products installed on them. The Virtualization Engine environment is hosted by a

directory server that must be an IBM Tivoli Directory Server (ITDS). A Virtualization Engine requires the use of an LDAP server (a directory server) and a directory server is linked to a Virtualization Engine environment: a Virtualization Engine environment contains a repository of environment-global information that is stored on an ITDS; a given installation of Virtualization Engine software retains information about the ITDS and the Virtualization Engine environment on that ITDS so that this information repository can be referenced at runtime by the Virtualization Engine management servers. If an existing Virtualization Engine environment already exists on the directory server, the installation can be included in that environment. If a Virtualization Engine environment does not exist, you must create it.

This first installation step offers different choices that depend on your environment. Figure 3-23 summarizes the wizard algorithm of this first step.

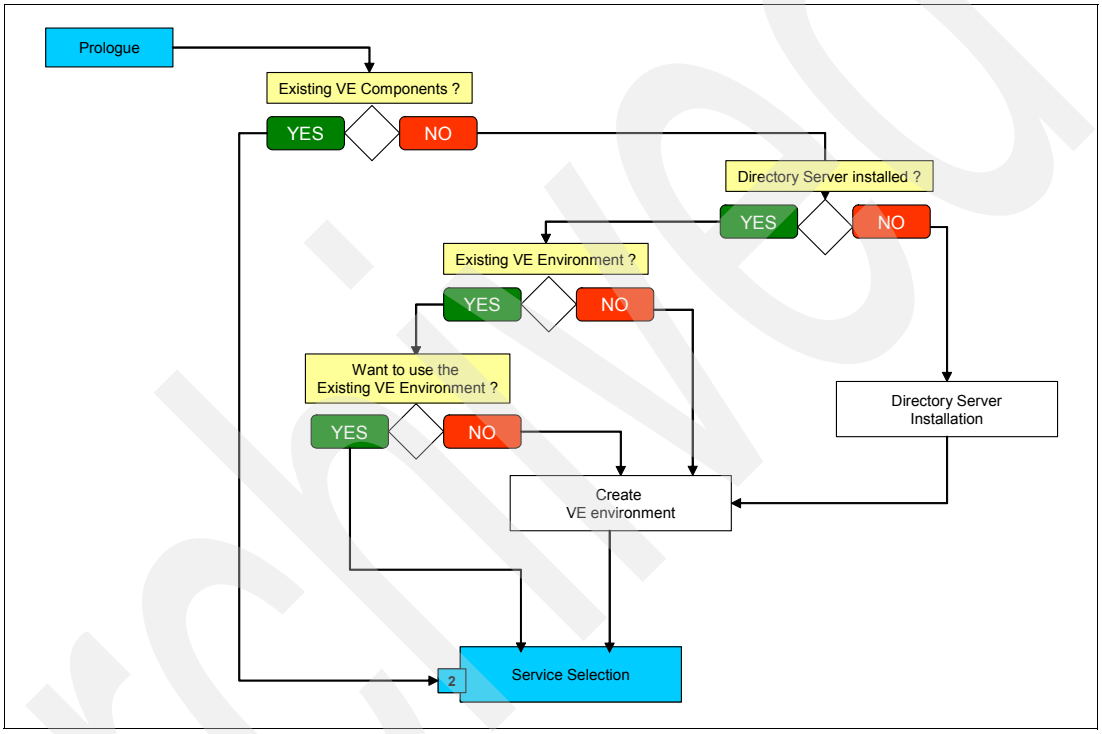


Figure 3-23 Establishing the Virtualization Engine environment

1. The wizard first checks to see if there is already any Virtualization Engine products on the system. If there are, then the window shown in Figure 3-24 on page 94 appears.

Archive

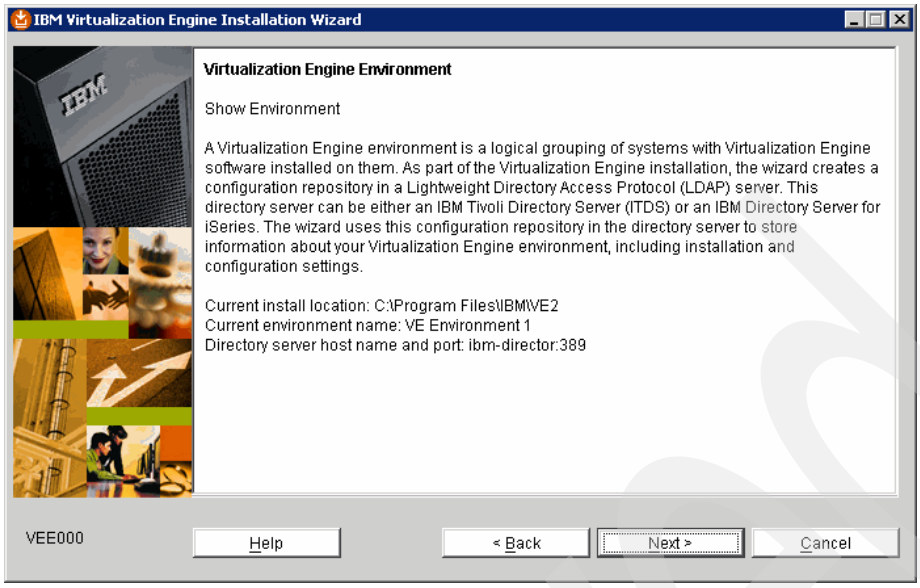


Figure 3-24 The wizard found an existing Virtualization Engine environment

The wizard finds the associated Virtualization Engine environment. You cannot change the properties of the already existing Virtualization Engine environment. You have to use it, or you can decide to restart the installation process on another installation directory (you can install more than one instance of the Virtualization Engine on the same machine using different installation directories, but you should not attempt to run services from different directories at the same time). If you decide to continue, press **Next** and the window shown in Figure 3-46 on page 106 will appear. You can then start Step 2 of the installation, as detailed in "Step 2: Service selection" on page 106.

2. If the wizard does not detect any Virtualization Engine product, it will ask you if you want to use an existing directory server or if you want to install a new one. A window, as shown in Figure 3-25, appears in which you have to enter your choice. The Virtualization Engine requires the use of a Directory Server. It is used to store the environment information that includes the installation and configuration settings.

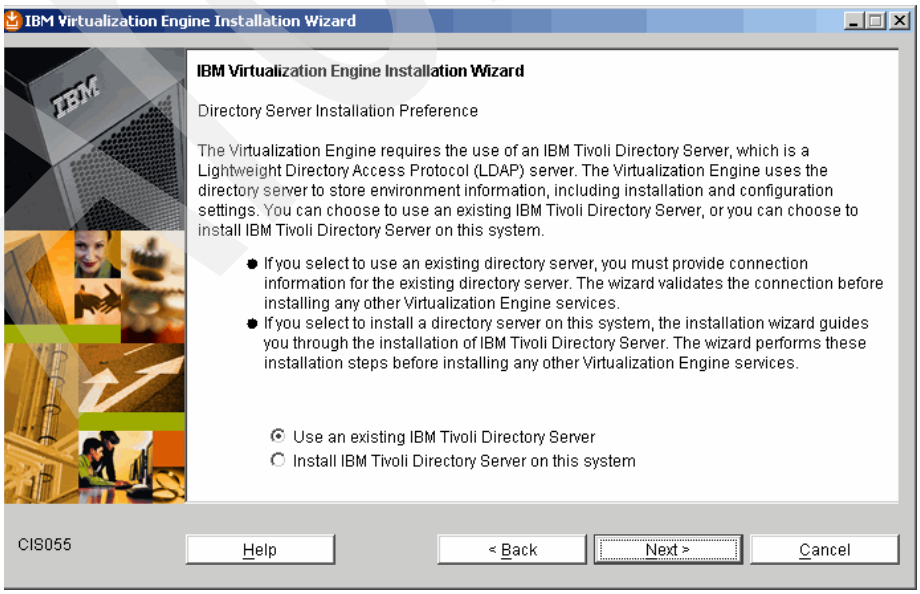


Figure 3-25 Decision to use an existing directory server or create a new one

You can choose to either use an existing directory server or create a new one:

– If you decide to use an existing directory server:

  i. You will be asked to provide the parameters of the existing directory server to the wizard, as shown in Figure 3-26, and the installation wizard checks if additional media is needed. The existing directory server must be either IBM Tivoli Directory Server 5.1, 5.2, or 6.0 or IBM Directory Server for i5/OS.

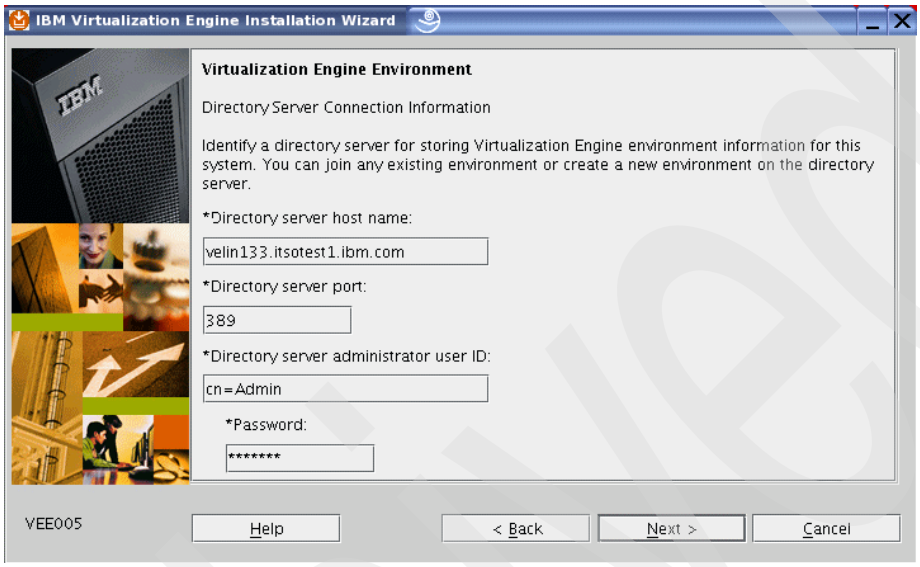


Figure 3-26 Using an existing Directory Server

  ii. You will be asked about the Virtualization Engine environment, as described in "Selecting the Virtualization Engine environment" on page 103, and then you continue with the Step 2 of the installation process, as described in "Step 2: Service selection" on page 106.

– If you decide to install a new directory server, as described in "Directory server installation process" on page 96, the wizard will help you install the Tivoli Directory server that requires a DB2 database. For DB2 and the Tivoli Directory server, you have to provide some parameters, as explained in "Directory server installation process" on page 96. This process by itself is a *full installation process in which there is no service selection* (EWLM, IBM Director, or Resource Dependency Service) proposed.

If you are using Windows, once your new directory server is installed, Figure 3-39 on page 102 is displayed and you have to re-start the wizard to install the Virtualization Engine components; on other operating systems, Figure 3-38 on page 102 is displayed and you can exit or continue with the Virtualization Engine environment setup and service installation. You will then be in a situation where no Virtualization Engine service is detected. You should use an installed directory server, possibly the one you just created.

Directory server installation process

You should perform the following steps to install the directory server. The process starts when you decide to create the directory server, as shown in the window shown in Figure 3-27.

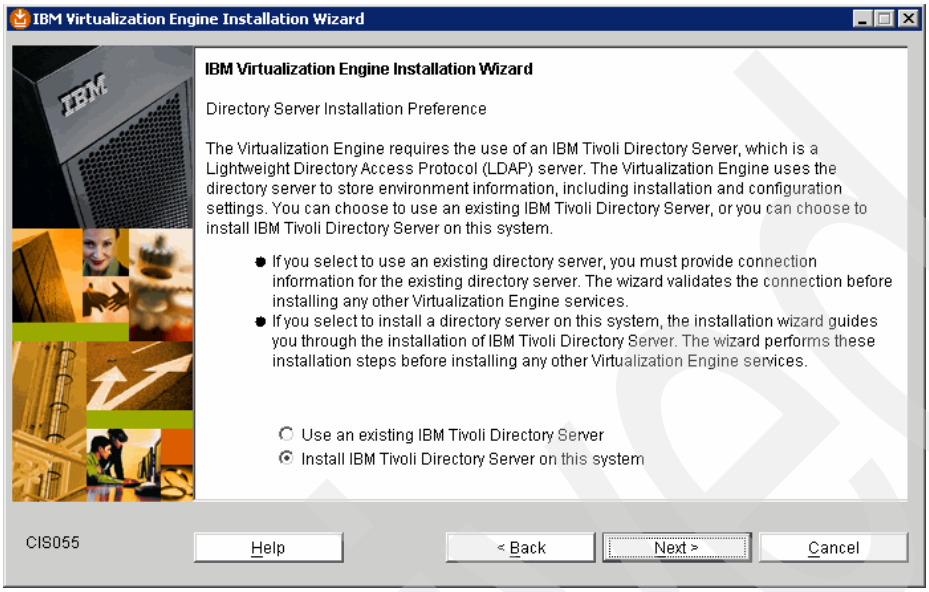


Figure 3-27 Choosing a directory server

1. The wizard checks the system requirements and then copies all the necessary files for the installation into the destination directory. Then the wizard unpacks all the necessary files. This can take some time.
2. The directory server needs a database associated with it. The wizard prompts you to either create a new operating system user ID for DB2 database, or to use an existing one, as shown in Figure 3-28. We chose **Create a new operating systems user ID for DB2 instance owner**.

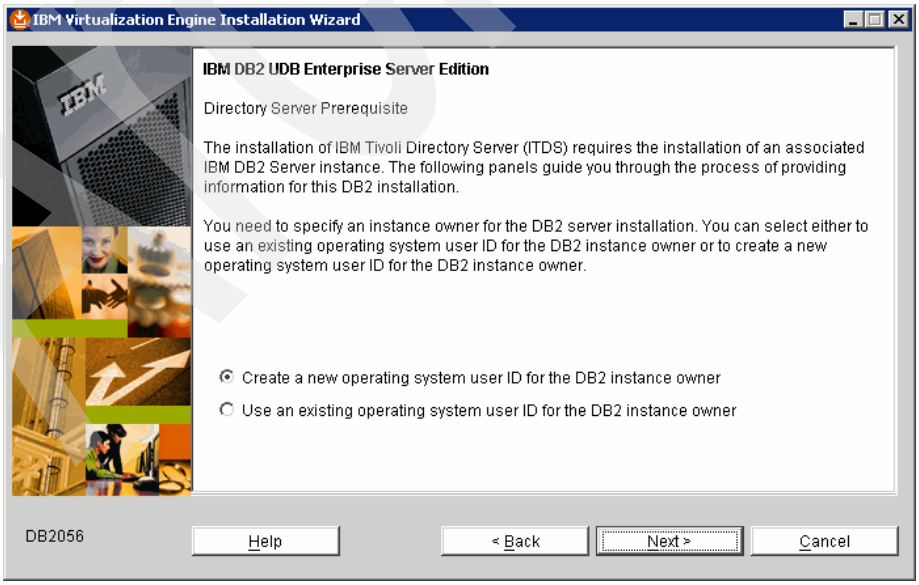


Figure 3-28 Directory server parameters

3. A window appears asking for the DB2 user and password to be entered, as shown in Figure 3-29.

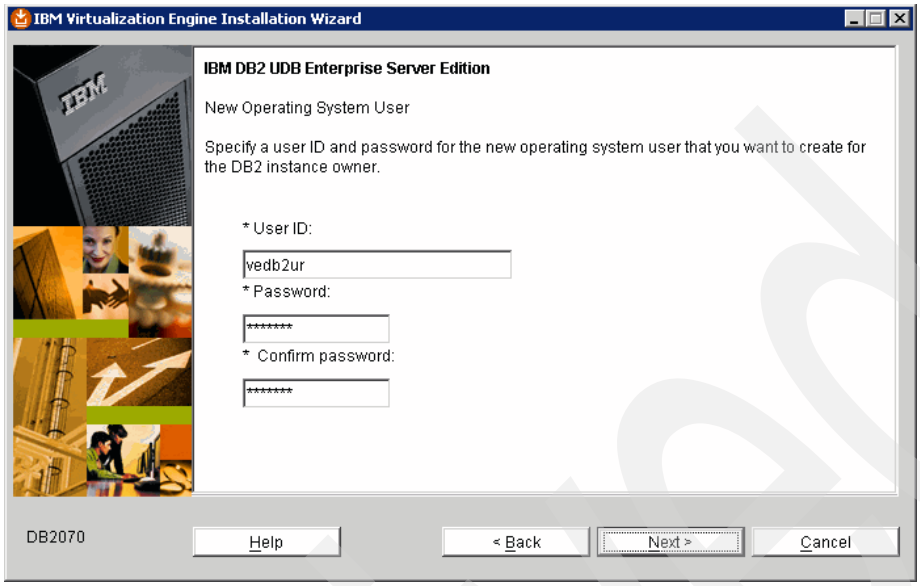


Figure 3-29 Specifying the database parameters

AIX and Linux considerations: The home directory and the user group notations are different between Windows and AIX or Linux. The displayed windows may vary slightly depending on the operating system on which you are installing.

Figure 3-30 is the window displayed in a Windows environment. Figure 3-30 is the window displayed in a Linux on System x environment.

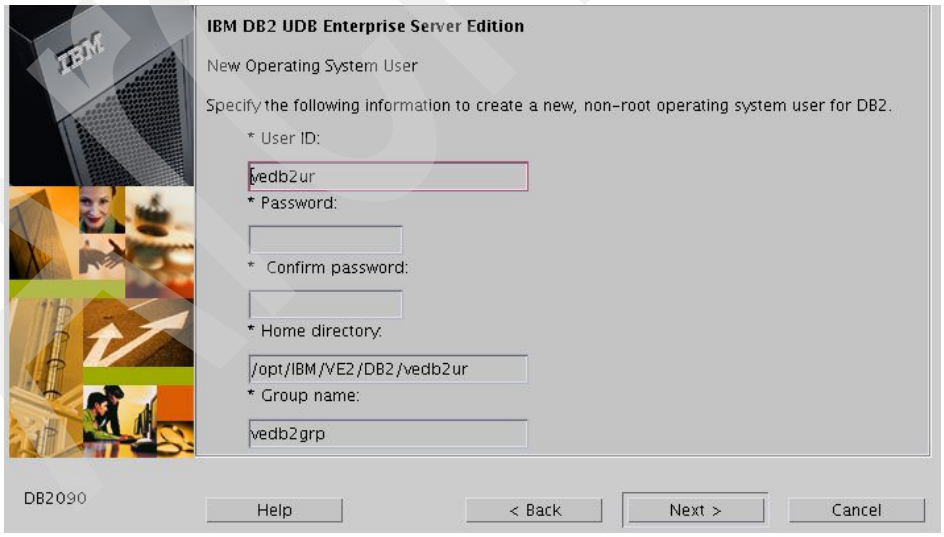


Figure 3-30 Database parameters in a Linux environment

4. The default directory server ports are then displayed, as shown in Figure 3-31. This is 389 for the server, and 3538 for the Administrator port. The port numbers can be modified. We recommend that you fill in the appropriate worksheet for your records.

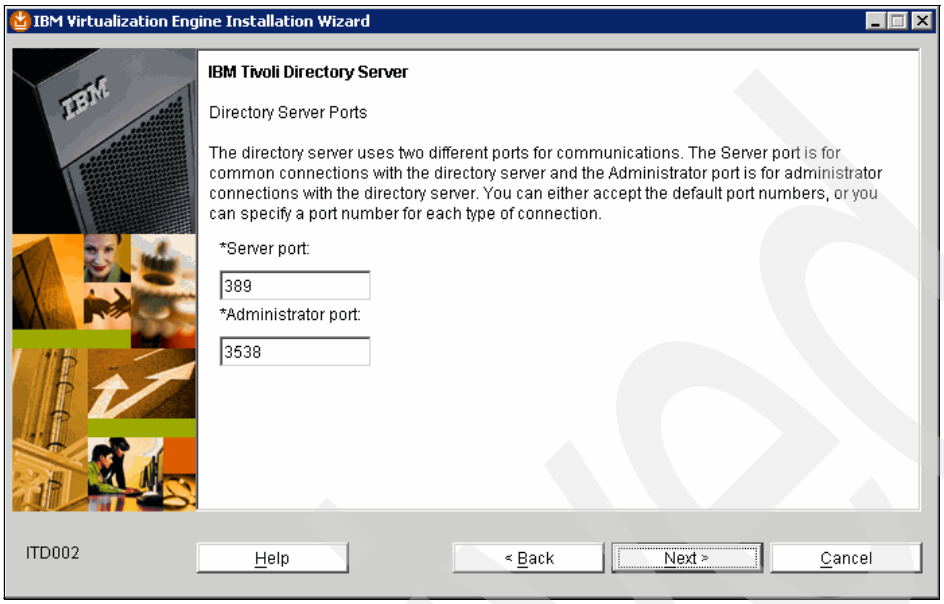


Figure 3-31 The directory server ports

5. In the next window that appears, shown in Figure 3-32, you must specify a character string that the Virtualization Engine will use as an encryption seed to generate a cryptographic key. The key is used to secure passwords in the directory. A directory seed string cannot contain spaces or special characters. Once you enter this, press **Next**.

   A seed value is necessary in case you have to export directory server data to another LDAP server. We recommend that you store the directory server seed in the directory server planning worksheet. This represents a sensitive security information. We recommend that you store it in a secure place.

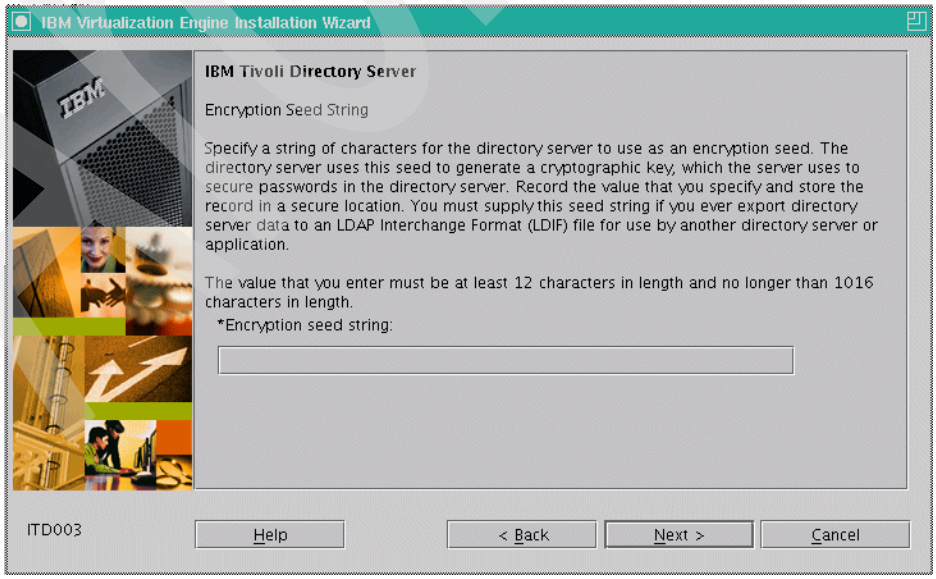


Figure 3-32 The encryption seed string

6. The directory server and database instance owner is the user ID you created before the installation, as described in 2.3, “Prerequisites” on page 50. Another window (not shown) asks for the DB Instance Owner ID and password. Enter the user name and password. Press **Next** to continue.
7. You are then asked to enter the directory server administrator parameters, as shown in Figure 3-33.

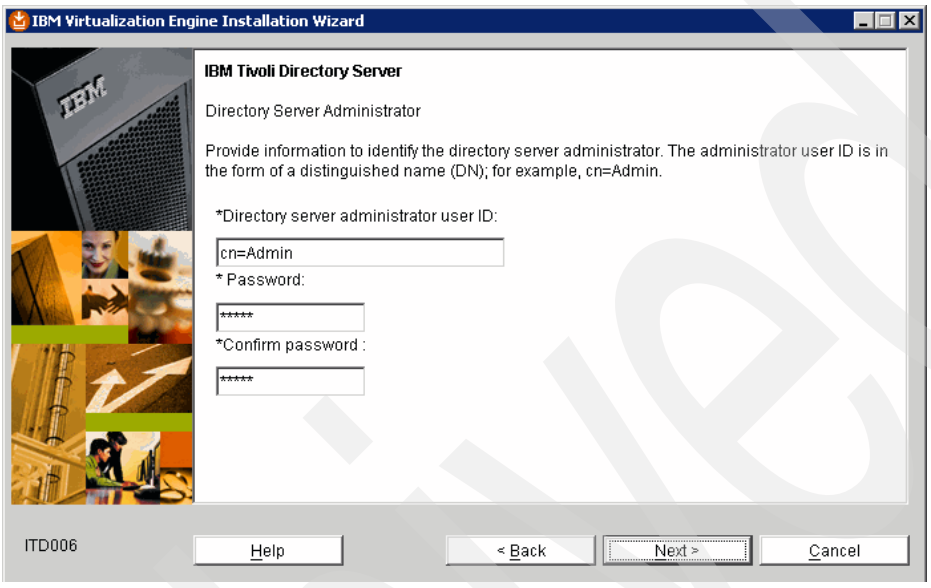


Figure 3-33 Directory server administrator

The directory server administrator does not reflect an operating system user ID. It is the administrator of the LDAP directory server database. A default value (DN syntax) is presented. In this window, you should provide a password for the administrator. This is important and sensitive security information that you have to manage according to your own security rules. We recommend that you record the value that you specify and store it in a secure place, and fill in the user access and password planning worksheet.

8. The first part of this configuration is now complete and the wizard will execute all the installation steps for the directory server installation only and other subcomponents it finds as a prerequisite. Going through this directory server installation process does not allow you to install other components (or services) of the Virtualization Engine. As explained in the previous section, you will have to restart the wizard for installing the Virtualization Engine components if you are on Windows.

   The wizard will now display a summary of the common runtime components it will install, as shown in Figure 3-34 on page 100. The installation will start when you click **Install**.

Archived

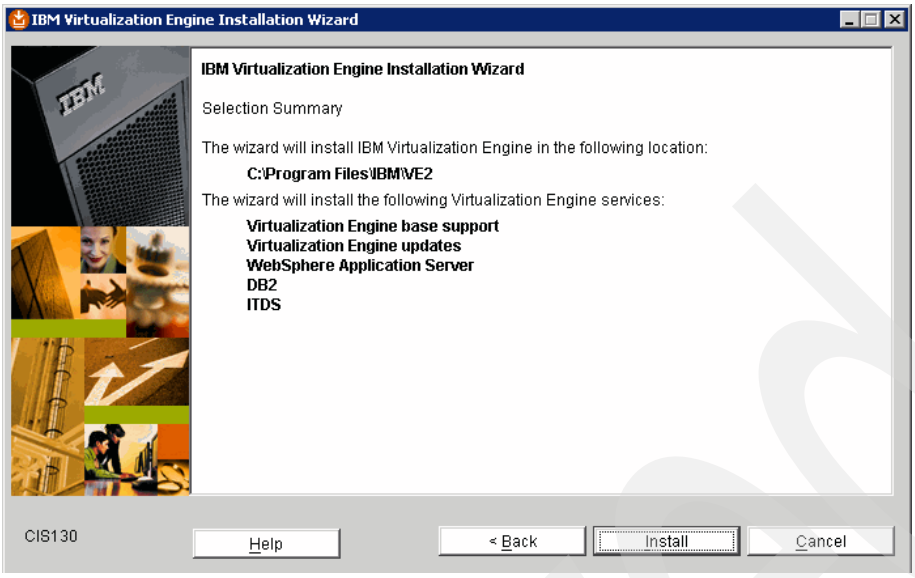


Figure 3-34 Summary of the Virtualization Engine services installation

The installation may take some time. A progress bar will indicate the progress, as shown in Figure 3-35.

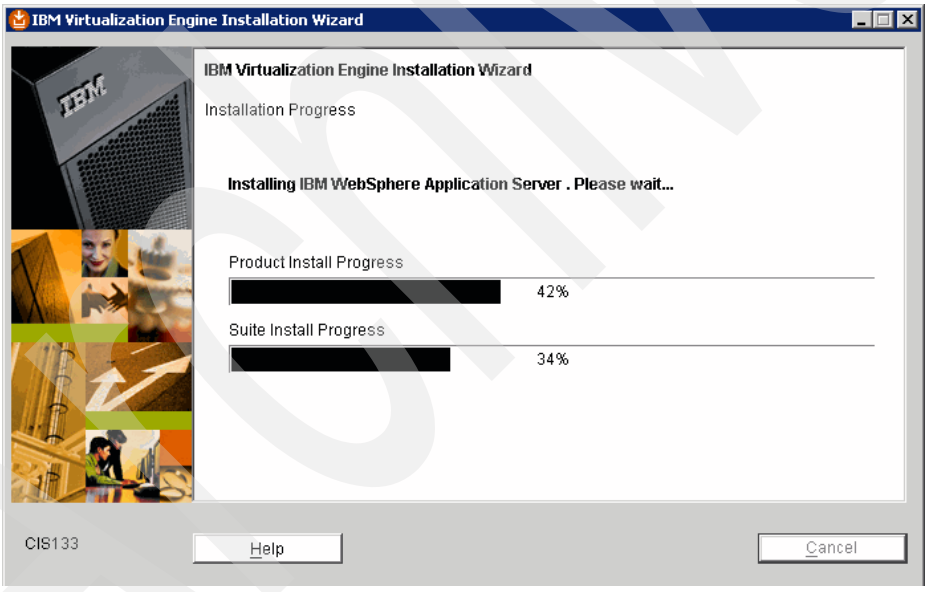


Figure 3-35 Installation progress information

9. Some post-installation steps may be required for each of the sub components. The wizard will prompt you with an initial window in which you should click **Next** to start the post-installation process, as shown in Figure 3-36 on page 101.

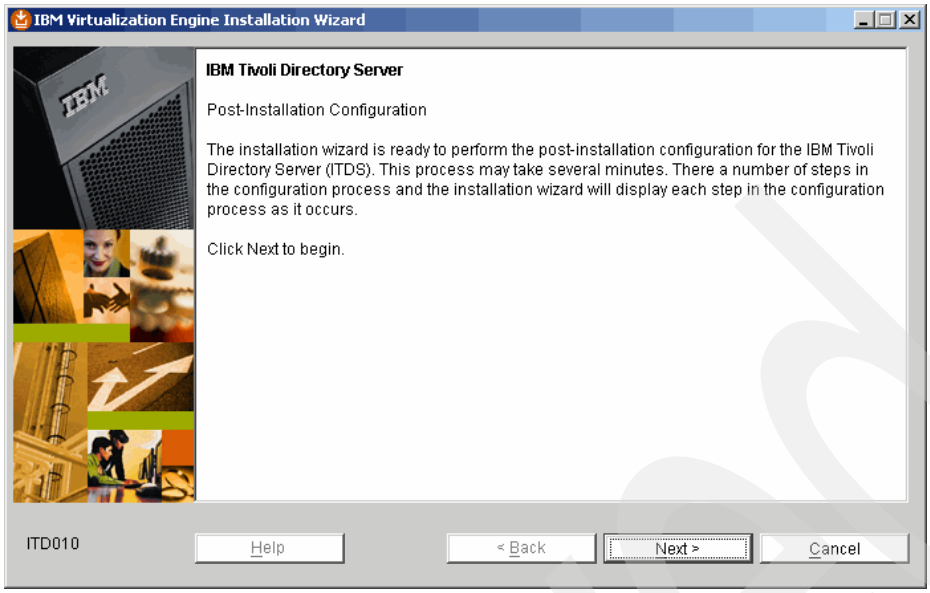


Figure 3-36 Post-installation process

Multiple windows will be displayed to inform you about the progress of the post-installation tasks. Other summary windows may be displayed to show the parameters that have been selected, depending on your entry in the previous steps, asking for your approval before the process continues, as shown in Figure 3-37.

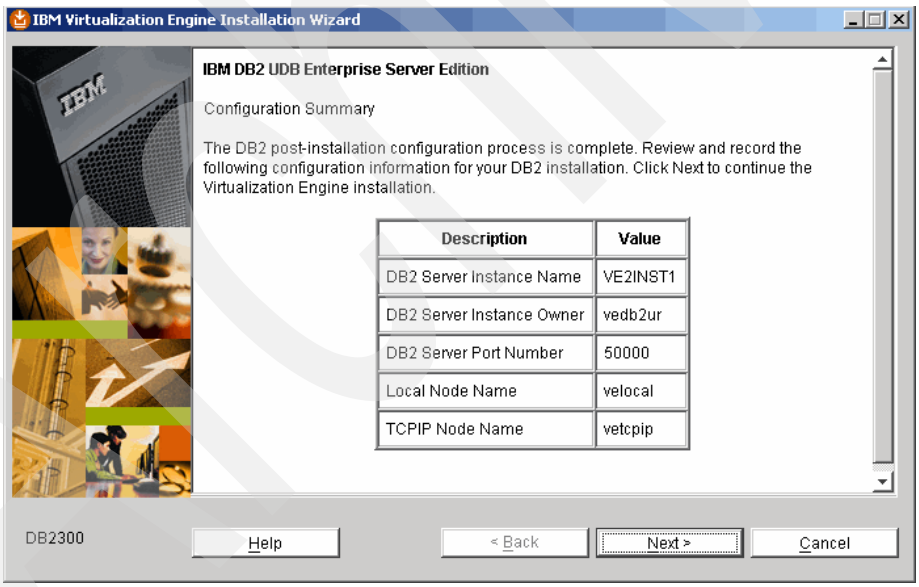


Figure 3-37 DB2 configuration summary

10. At the end, the wizard prompts a new window, displayed in Figure 3-38, where you have to decide to either exit or to install the Virtualization Engine components now.

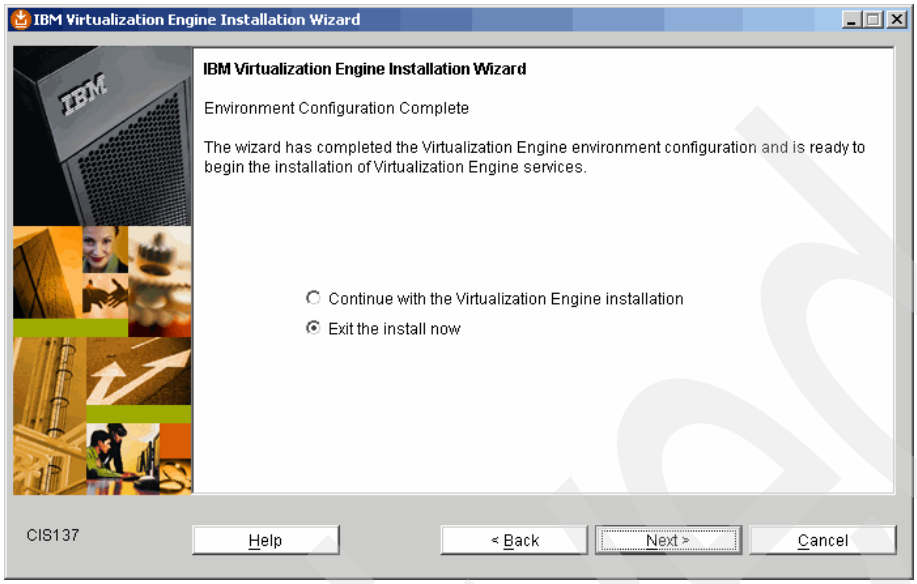


Figure 3-38 Environment configuration final window

Important: On Microsoft Windows, the reboot is mandatory (no reboot is required for AIX or Linux platforms).

With the Microsoft Windows operating systems, a window, as shown in Figure 3-39, appears. You should click **Finish** first to exit the process. If you want to install Virtualization Engine products, you must reboot your machine and start the wizard again.

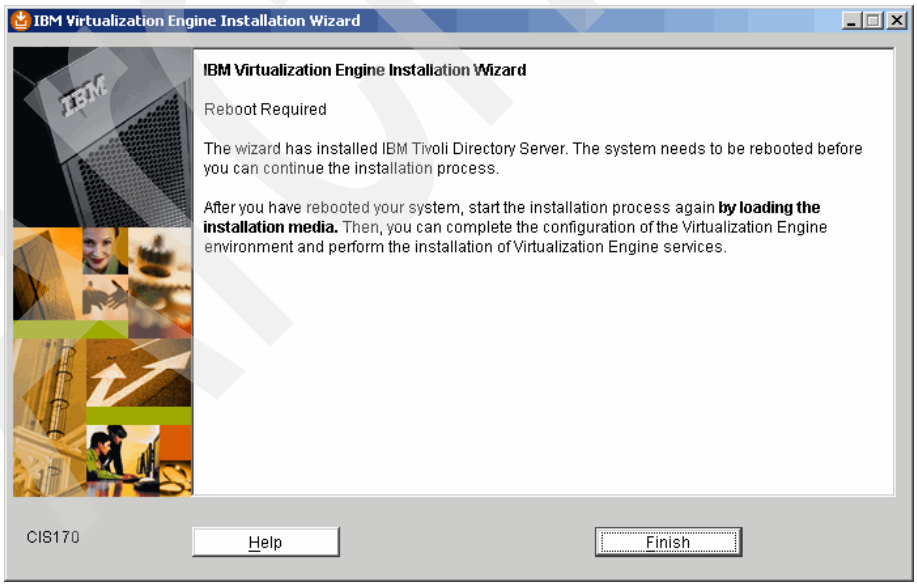


Figure 3-39 A reboot is required for Microsoft Windows operating system

If you choose **Continue with the Virtualization Engine installation**, click **Next**, as shown in Figure 3-37 on page 101, the wizard goes to the window for the Directory Server Connection Information (Figure 3-26 on page 95) so that the user can set up the

Virtualization Engine environment using the ITDS the user just installed; this window appears if a Virtualization Engine environment has not been established for this installation directory yet, regardless of whether or not any services have been installed.

Figure 3-40 shows the installation summary window that appears after the wizard has successfully installed the Virtualization Engine. Select **Finish** to exit the wizard.

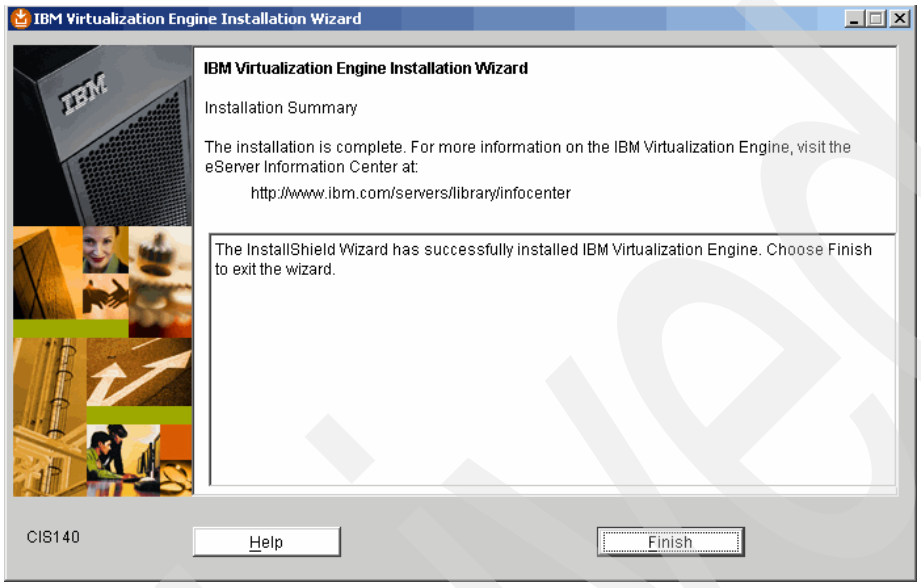


Figure 3-40 Installation summary

Selecting the Virtualization Engine environment

This task is started when the wizard does not detect any existing Virtualization Engine component on the machine and if you stated that you want to use an existing directory server, as shown in Figure 3-41.

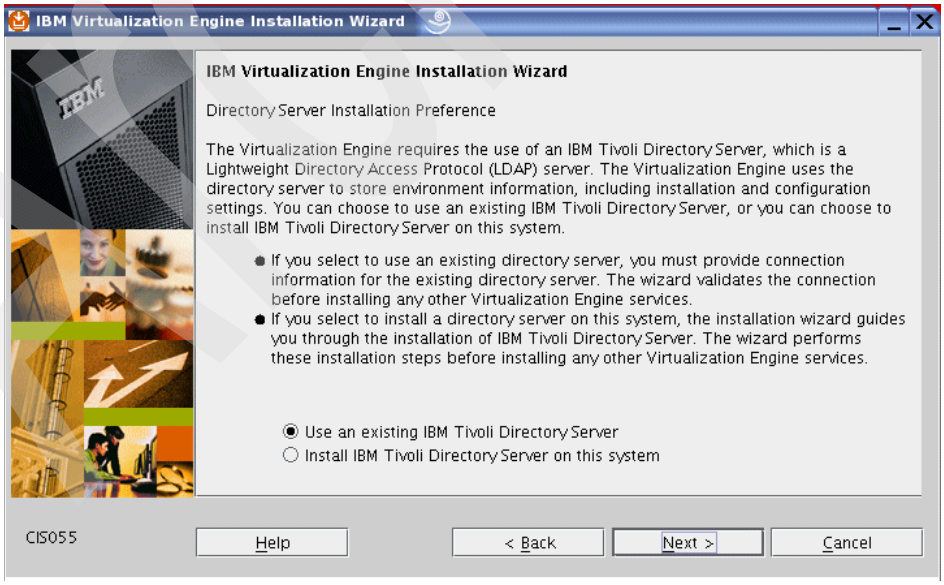


Figure 3-41 Choosing to use an existing directory server

To select the Virtualization Engine environment, perform the following tasks:

1. In the beginning, a window with information about the Virtualization Engine environment appears, as shown in Figure 3-42.

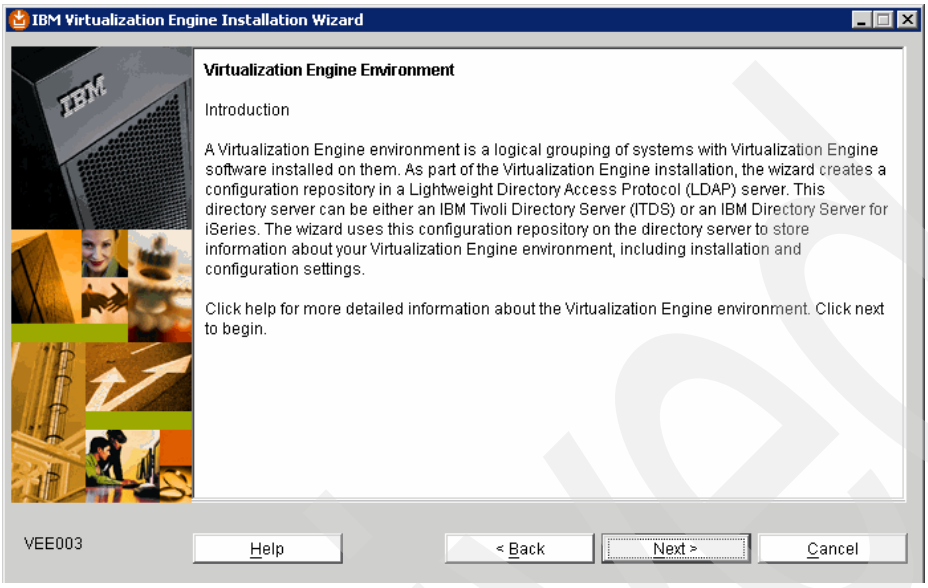


Figure 3-42 The Virtualization Engine environment information window

2. The wizard then asks you to provide information about the directory server you want to use, as shown in Figure 3-43.

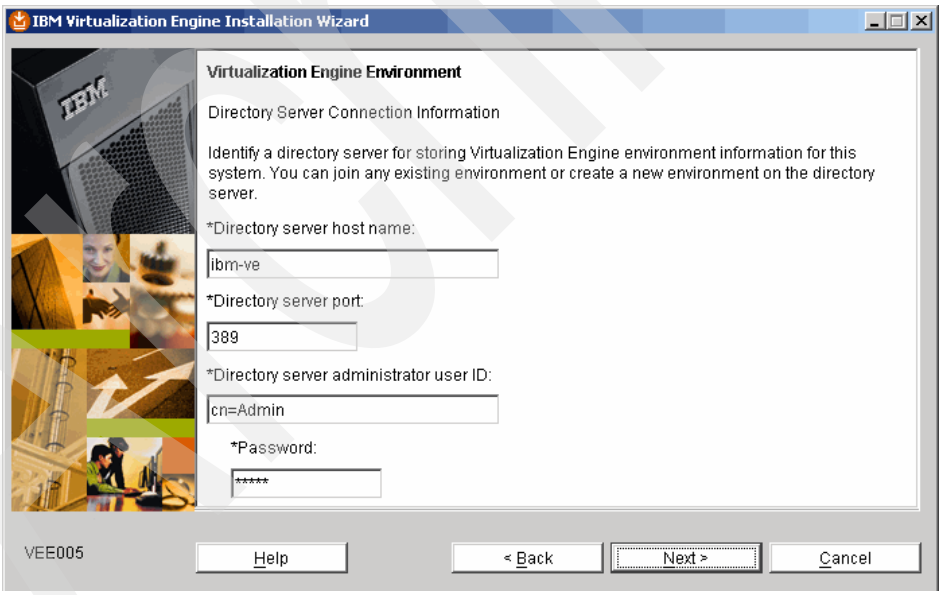


Figure 3-43 Directory server properties

The wizard checks the parameters, and makes a connection between your current machine and the directory server machine.

3. The wizard then checks whether an existing Virtualization Engine environment name exists or not. This results in two possibilities:

   a. If the wizard finds an existing Virtualization Engine environment, the window, shown in Figure 3-44 appears, and you can choose between using the existing environment or creating a new one.

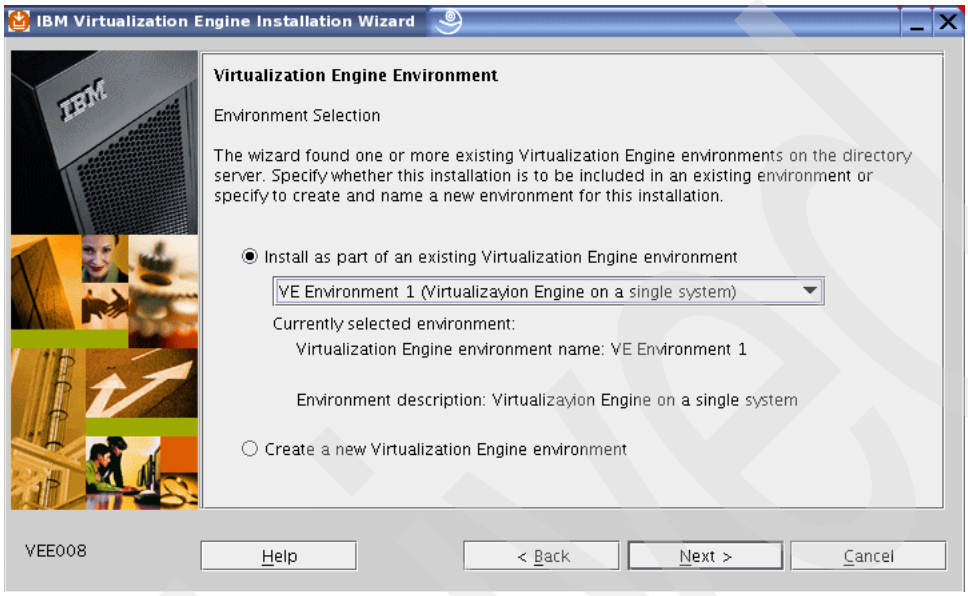


Figure 3-44 Virtualization Engine environment selection

Press **Next**, and Step 2 of the installation process will start, as explained in Step “Step 2: Service selection” on page 106.

   b. If the wizards does not find an existing environment, you are asked to provide some information for its creation, as shown in Figure 3-45.

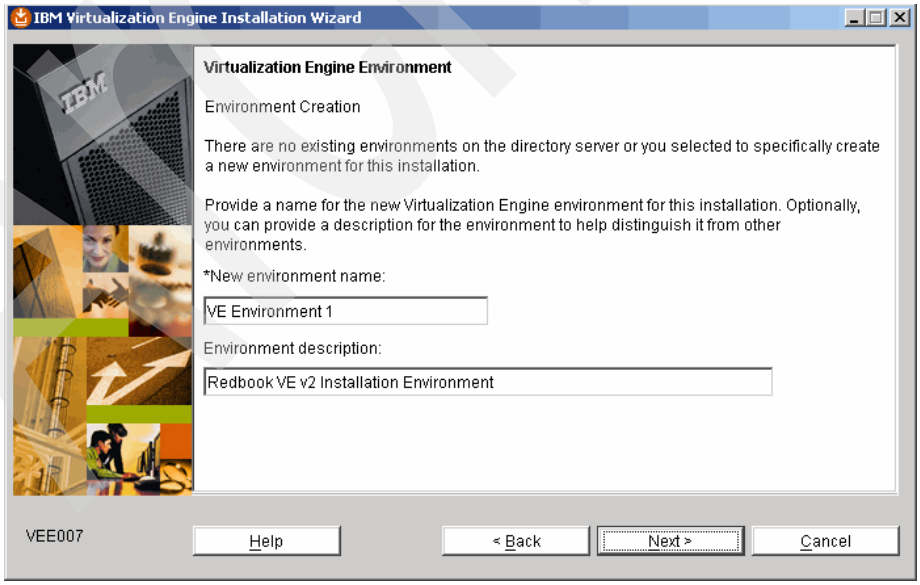


Figure 3-45 Providing the Virtualization Engine environment name

Step 2: Service selection

1. The wizard will now ask you to select the Virtualization Engine components you want to install, as shown in Figure 3-46. This window appears when you are not installing under Linux. The Resource Dependency Service component is not shown.

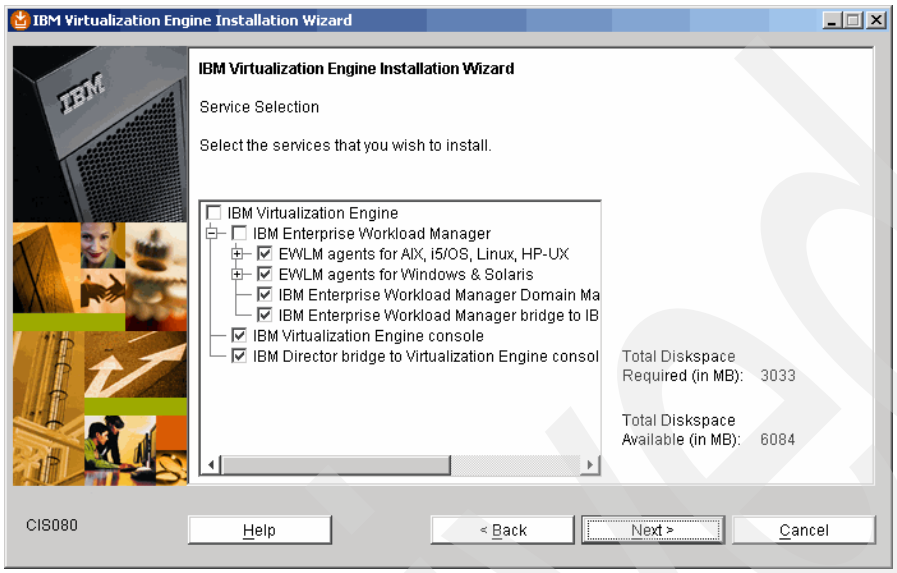


Figure 3-46 Selection of the Virtualization Engine products to install (installation under Windows)

Figure 3-47 shows the window that appears when you are installing on a Linux system. Resource Dependency Service is proposed in the list.

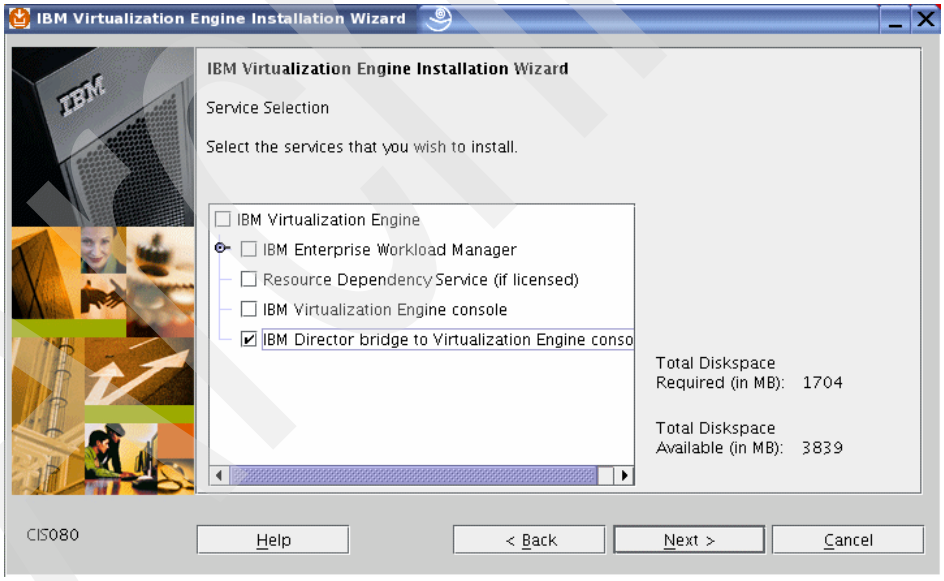


Figure 3-47 Selection of the Virtualization Engine products to install (installation under Linux)

The key icon shown next to the IBM Enterprise Workload Manager means that you have to choose other subcomponents. Click the icon to view the other possible subcomponents.

If a product is already installed, the wizard will inform you. However, the wizard cannot check whether a product has been installed on *another* machine. This is something you have to check yourself.

On the right side of the window, the amount of available space on disk and how much disk space is required by the selected services is displayed. This information is important for you before you start the installation. Click **Next**.

Note: For the EWLM agents, the Virtualization Engine installer does not actually install the agent code on any managed server. It simply prepares the agent code for further distribution through the Remote Package Deployer.

2. A selection summary window appears, as shown in Figure 3-48, which names all the products, components, and subcomponents that will be installed when you click the **Next** button. It also displays the components the wizard will install as prerequisites of some of the products.

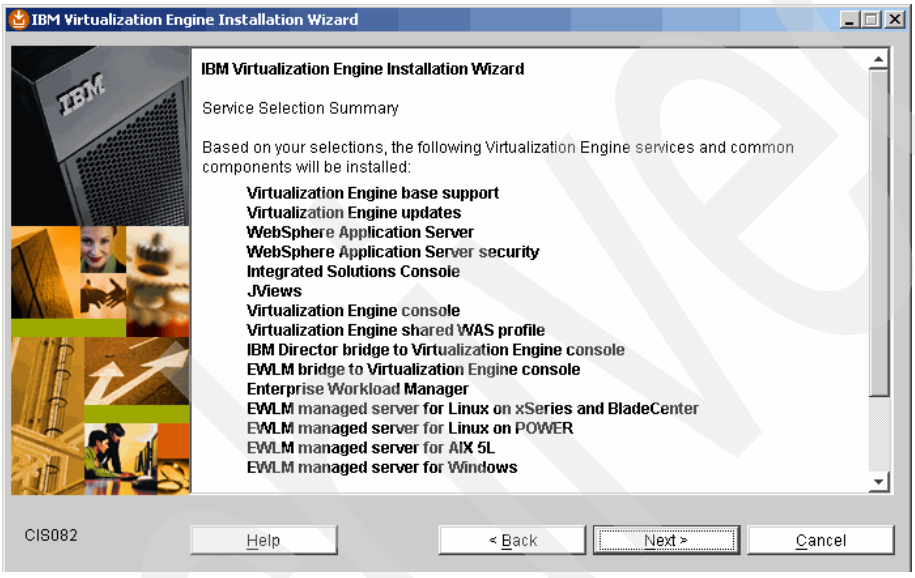


Figure 3-48 Selection summary example (with no RDS installation)

The content of this window depends on the options you choose for your installation. In our example, shown in Figure 3-48, Resource Dependency Service has not been chosen. Figure 3-49 on page 108 shows another example, where Resource Dependency Service has been chosen. All the components that will be installed to run Resource Dependency Service are shown. Note that DB2, in particular, is in the list.

Archived

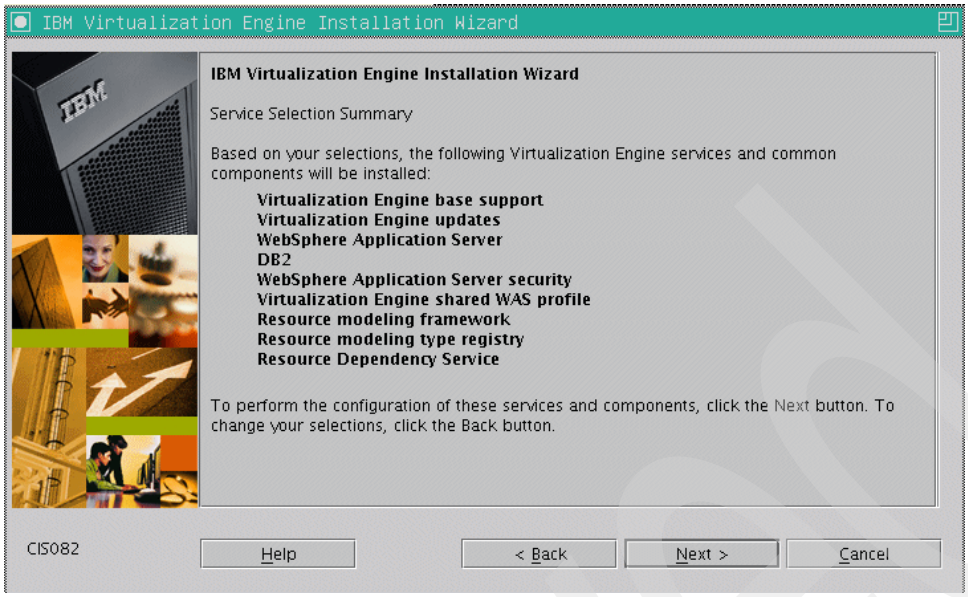


Figure 3-49 Selection summary example with Resource Dependency Service installation

Step 3: File copy

1. The wizard first checks if additional media needs to be copied. It then informs you and you should click **Next** to start the copy, as shown in Figure 3-50. No further media copy is done when you are doing the installation from already copied files on a server.

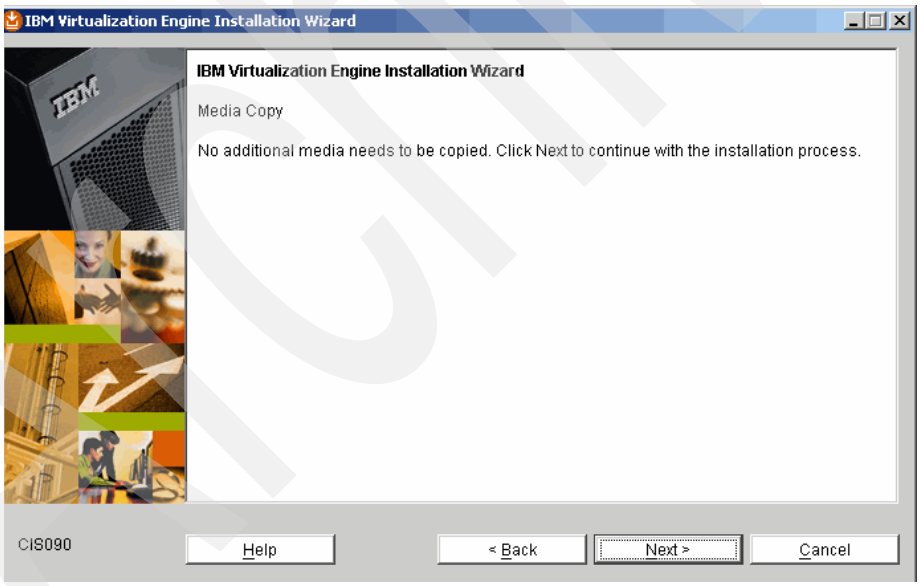


Figure 3-50 Media copy

2. The wizard will inform you about what it is doing, for example, unpacking the files as shown in Figure 3-51 on page 109. There is no need during this process to enter any information or click any button.

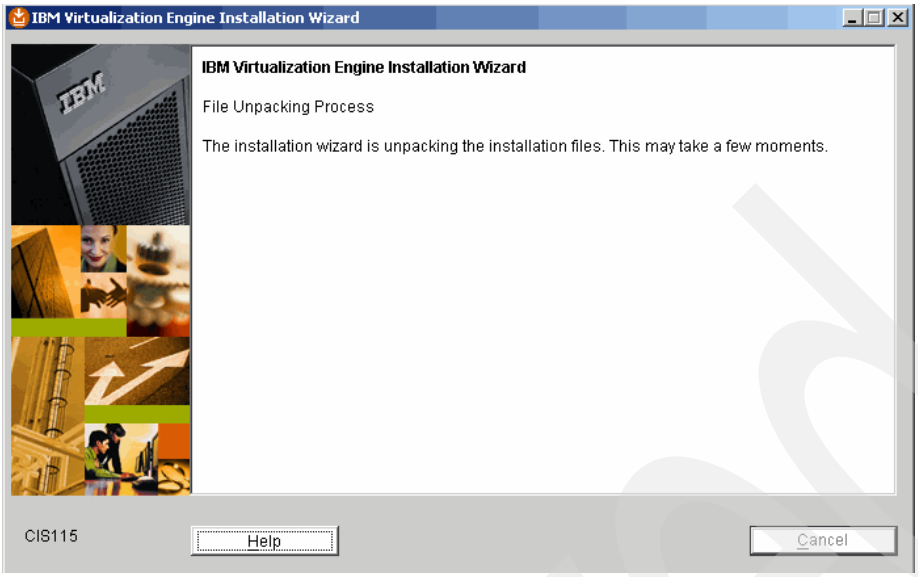


Figure 3-51 Information window

Step 4: Preinstallation configuration

1. The first task of the preinstallation process is to customize the WebSphere application server security environment according to your own security rules. The window shown in Figure 3-52 will start the process.

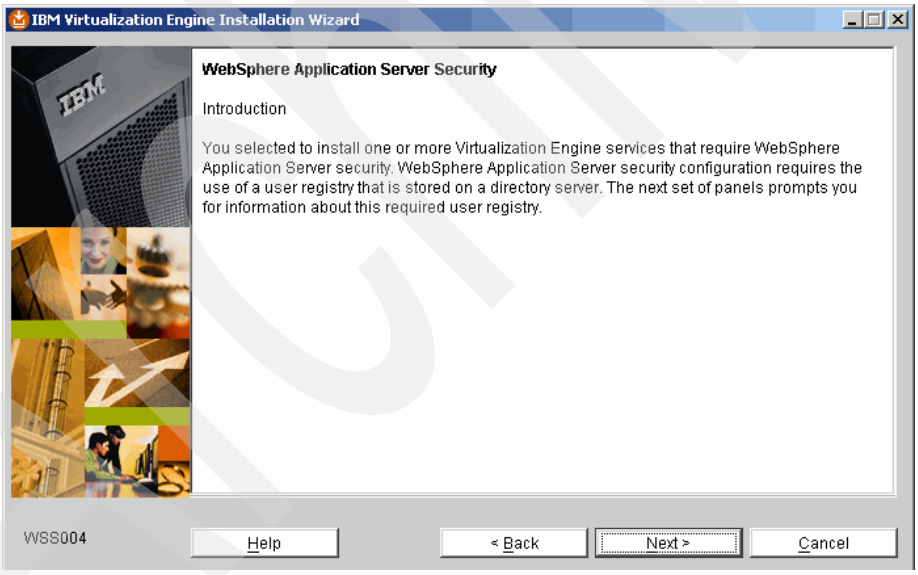


Figure 3-52 WebSphere Application Server security start task

   a. You should provide your directory server preference, as shown in Figure 3-53 on page 110. Click **Next**.

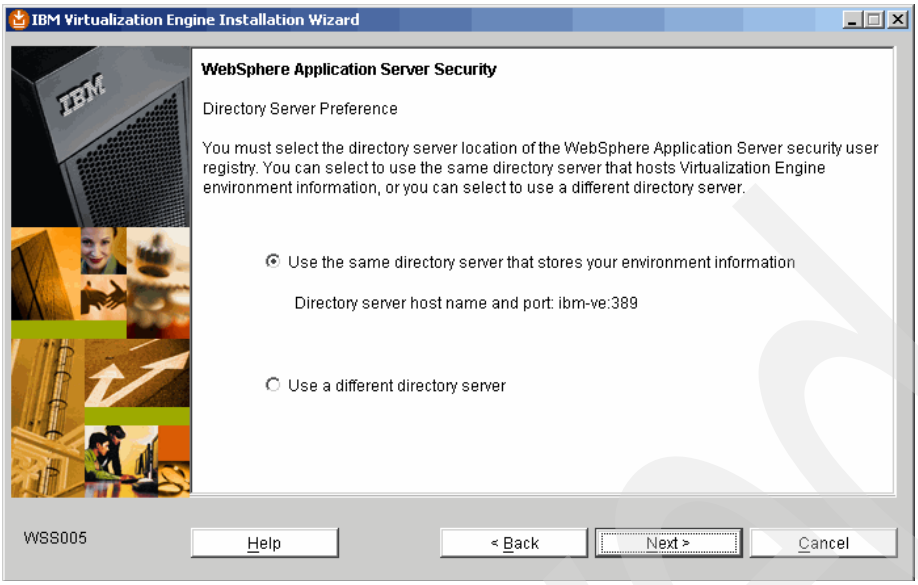


Figure 3-53 WebSphere Application server security: directory server preferences

b. You should provide the directory server configuration preference, as shown in Figure 3-54. Click **Next**.

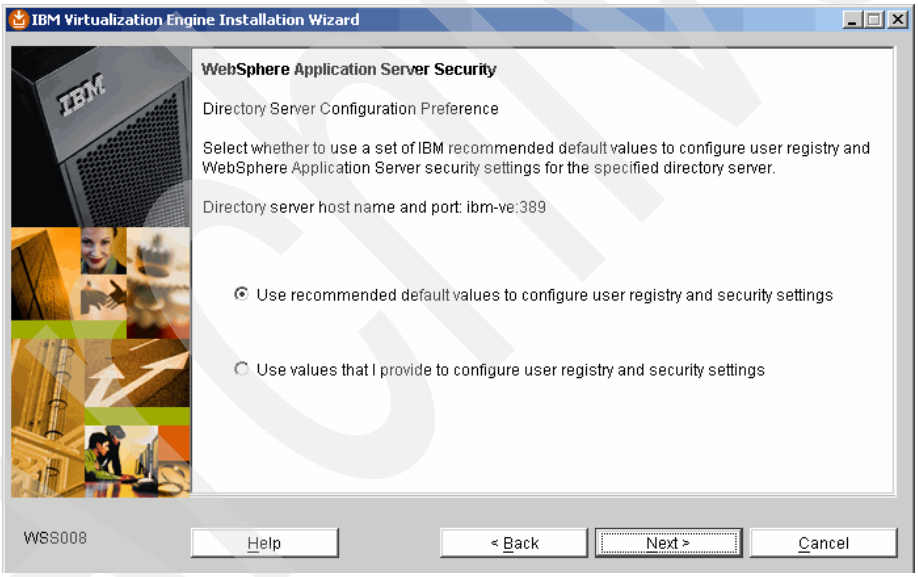


Figure 3-54 WebSphere Application server security: directory server configuration preference

Note: If the Virtualization Engine installer is not able to automatically determine the network domain of the system to use as a Lightweight Third-Party Authentication (LTPA) single sign-on domain, you will have to specify the WebSphere Application Server security settings manually using the second radio button “Use values that I provide...”; you will go through a series of windows until you encounter the window containing the LTPA Single Sign On Domain in which you have to provide the network domain.

c. The wizard then displays a summary of your parameters, as shown in Figure 3-55 on page 111. Click **Next**.

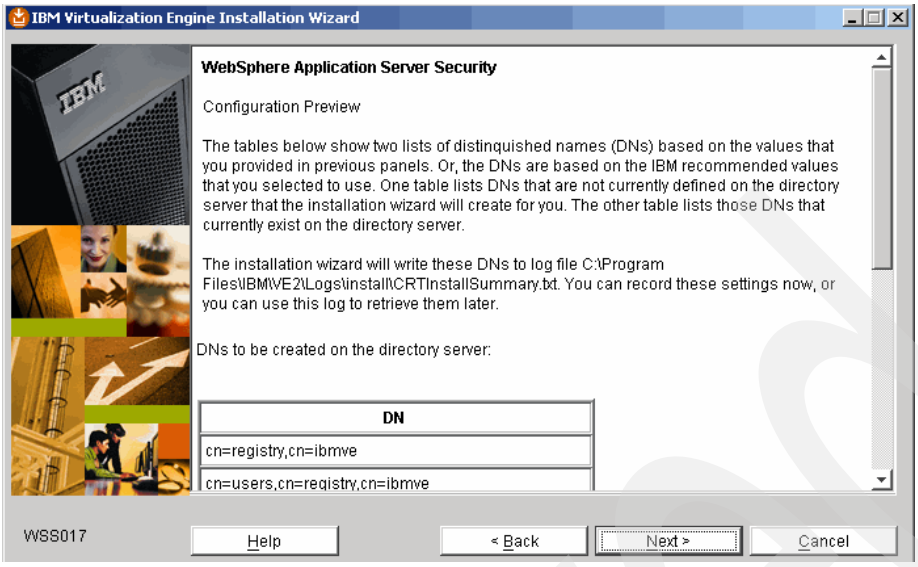


Figure 3-55 WebSphere application server security: preferences summary

2. The Virtualization Engine services may require an identity mapping repository. The installation wizard assists you in creating one when needed.

   a. The wizard introduces the identity mapping repository, as shown in Figure 3-56.

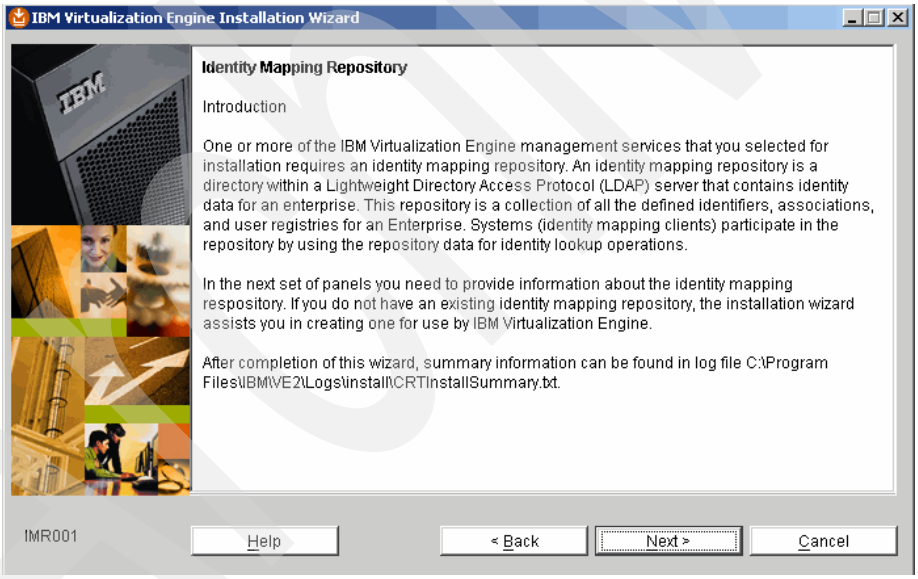


Figure 3-56 Identity mapping repository

   b. Then the wizard asks you to select your preferences, as shown in Figure 3-57 on page 112 and Figure 3-58 on page 112.

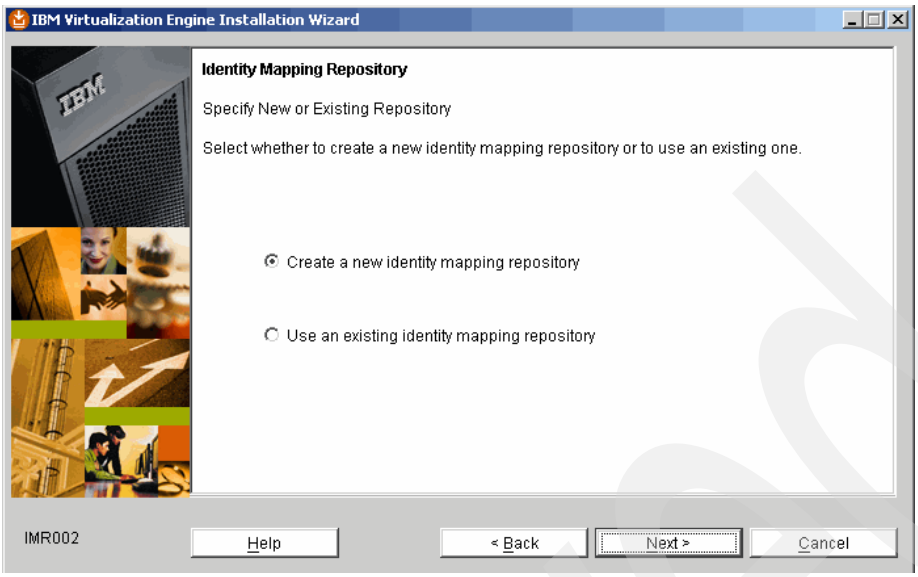


Figure 3-57 Selecting the identity mapping repository

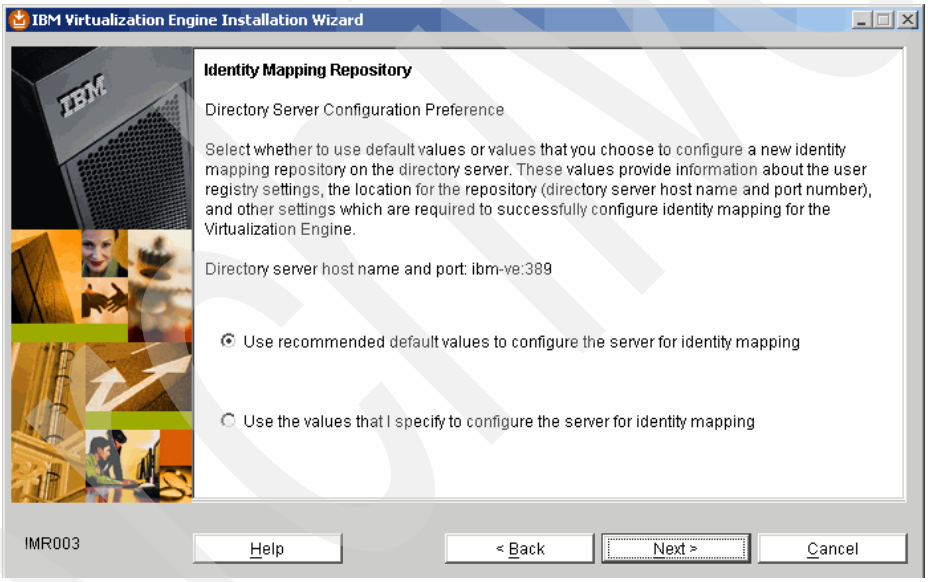


Figure 3-58 Identity mapping repository configuration preference

c. Then the wizard displays the summary of the settings of your installation, as shown in Figure 3-59 on page 113.

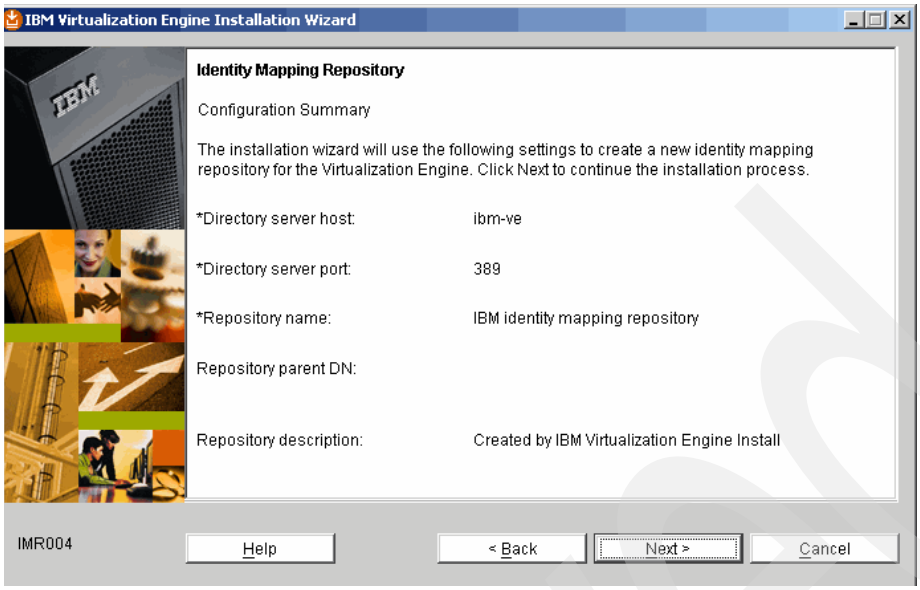


Figure 3-59 Identity mapping repository: configuration summary

d. The Virtualization Engine console needs to run the Integrated Solutions Console (ISC), which requires a range of 30 TCP ports. The wizard will ask you to provide the starting port number, as shown in Figure 3-60.

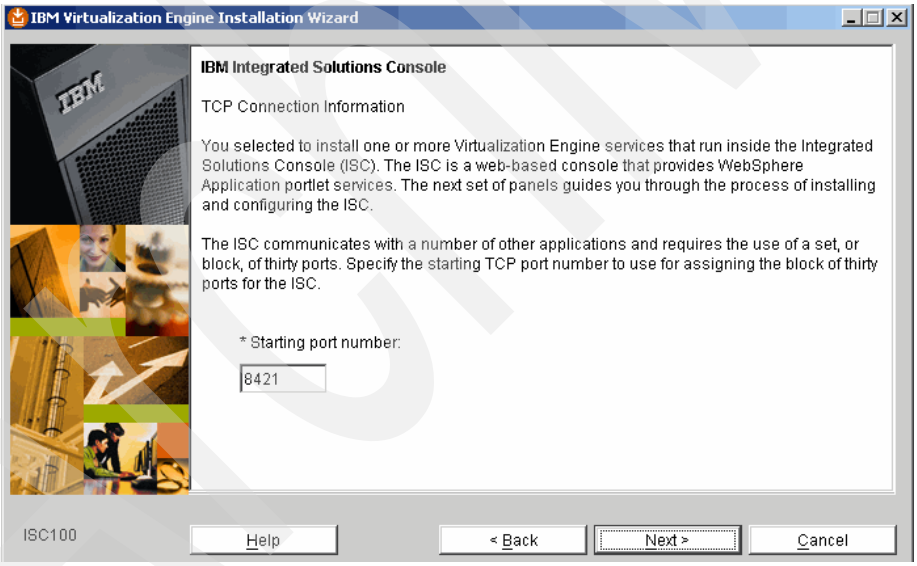


Figure 3-60 ISC TCP port number

Then the wizard displays the port numbers it will use for the ISC and the WebSphere port used to run the ISC, as shown in Figure 3-61 on page 114. We suggest you verify, using the **`netstat -an`** command, that the full block of ports specified are not in use.

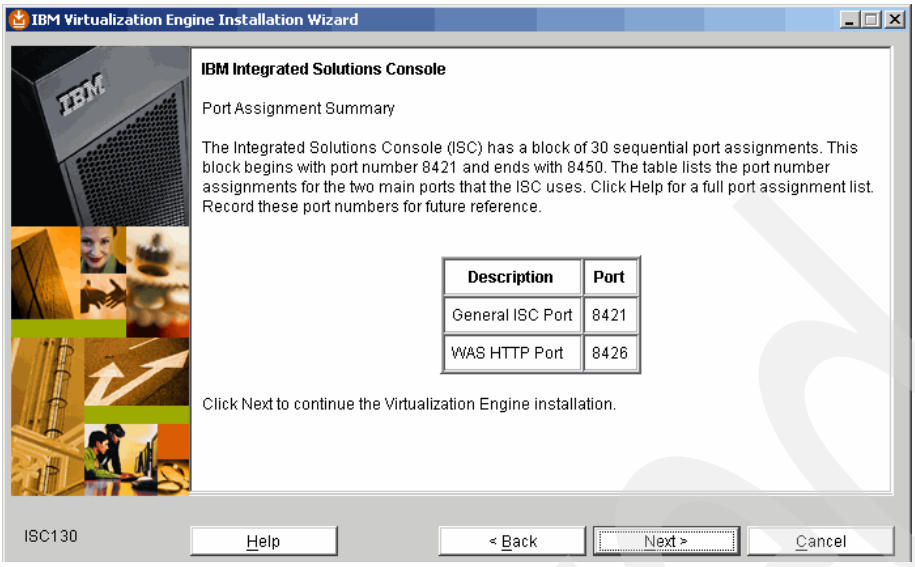


Figure 3-61 ISC port assignment summary

3. You should decide whether the Virtualization Engine should start automatically, as shown in Figure 3-62.

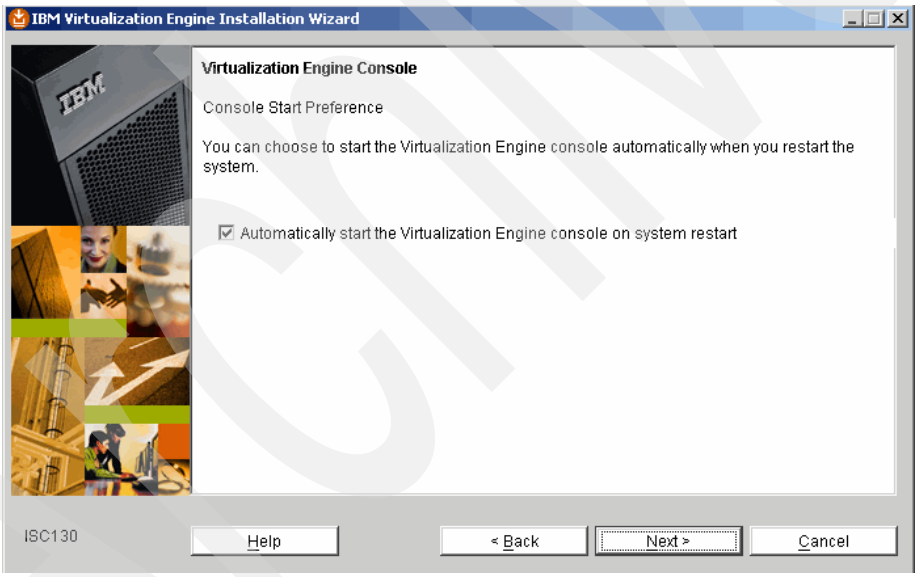


Figure 3-62 Virtualization Engine console start preference

4. If you install Resource Dependency Service, DB2 needs to be installed and you will be asked whether you want to use an existing DB2 server or install a new one, as shown in Figure 3-63 on page 115.

Archived

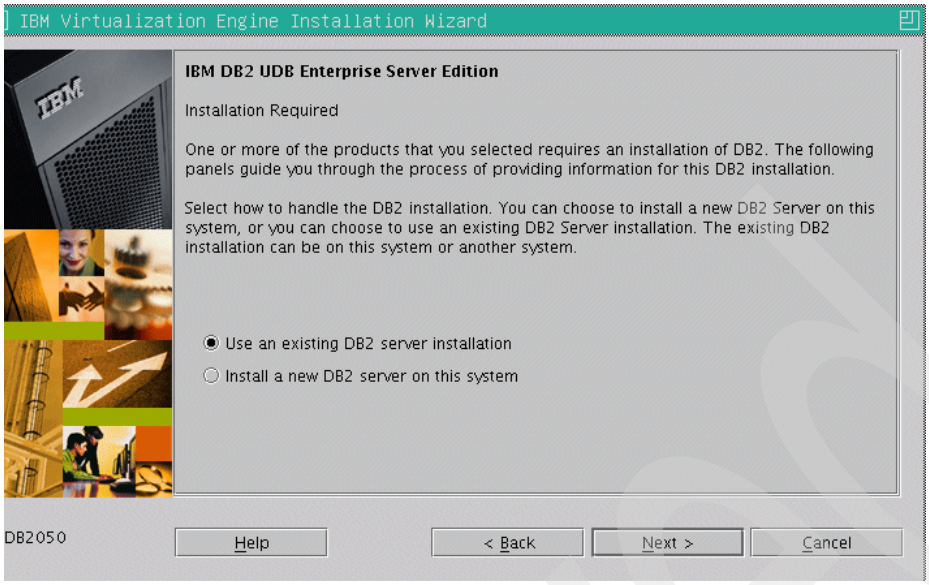


Figure 3-63 DB2 use

If you are using a DB2 server instance, you should provide the host name, as shown in Figure 3-64, and the wizard will make the connection.

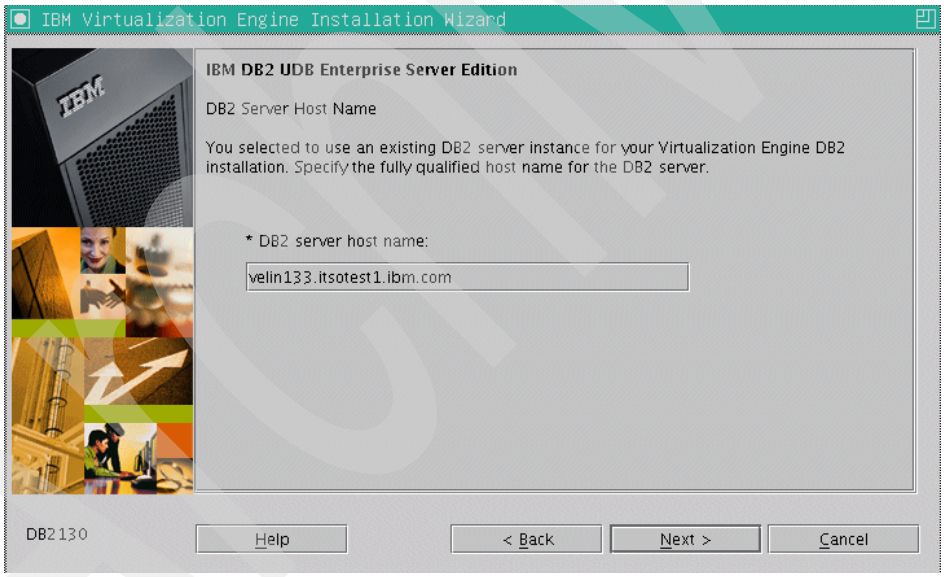


Figure 3-64 Providing the parameters of an existing DB2

Otherwise, you will have to provide some information, such as an operating system user ID and password, as shown in Figure 3-65 on page 116, and information about the TCP connection, as shown in Figure 3-66 on page 116.

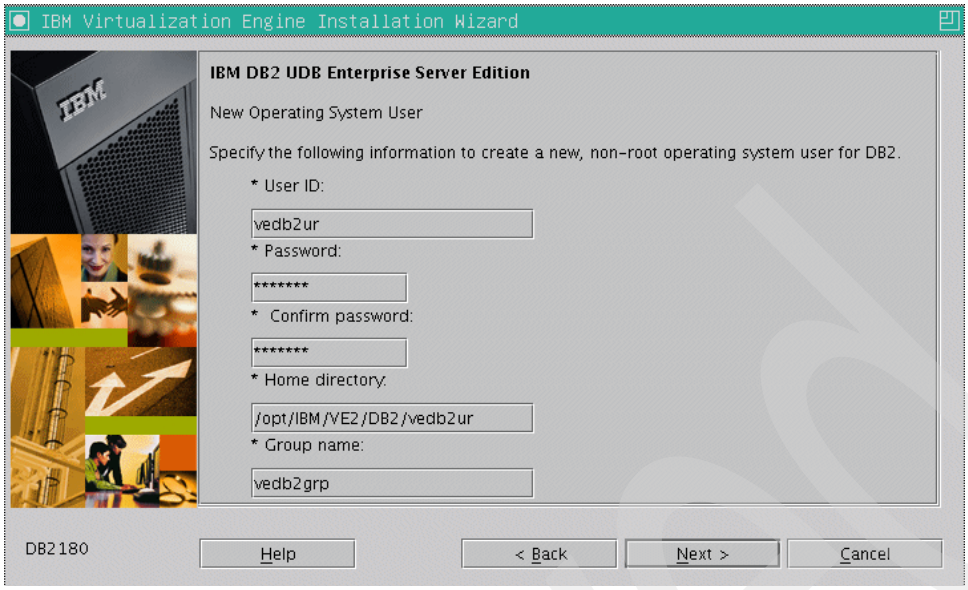


Figure 3-65 Providing the parameters to create a new DB2 environment for Resource Dependency Service (1/2)

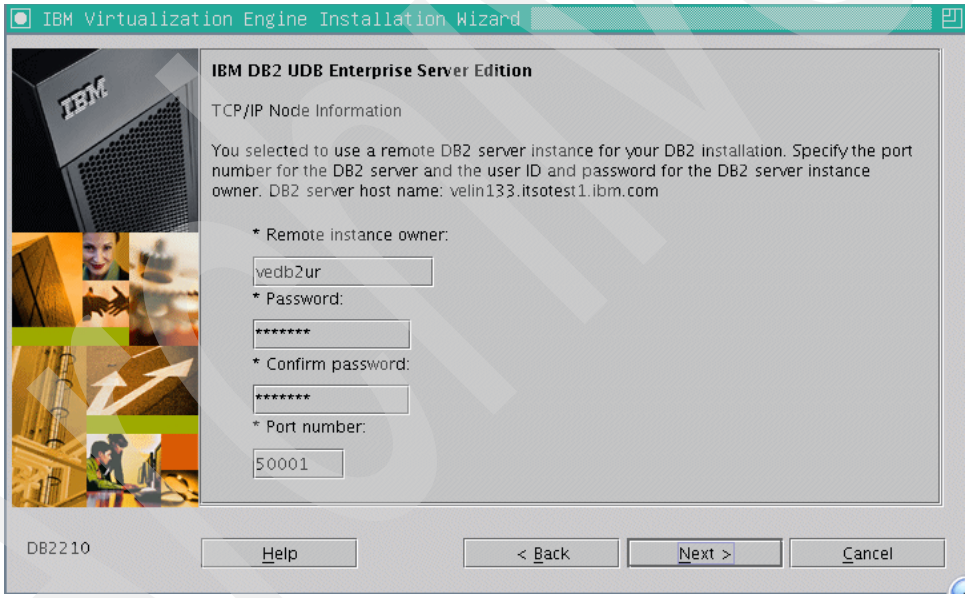


Figure 3-66 Providing the parameters to create a new DB2 environment for RDS (2/2)

The default operating system user profile name for the local DB2 administrator name is vedb2ur. You should record this information in the planning worksheet.

This information is used to link the Virtualization Engine to the remote database server. The information is made of the remote user profile owner of the instance (the default is vedb2ur) and the remote database port number for which the default value is 50001. You can check on the machine hosting the DB2 server in the file named /etc/services and look for the service name db2c_vedb2ur (assuming that the DB2 instance owner is vedb2ur). Using the same DB2 that is used for ITDS is fine: there should be no performance or size issues in most cases.

Click **Next**.

5. An introduction window is displayed.You should enter the port numbers for the shared WebSphere Application Server profile, as shown in Figure 3-67.

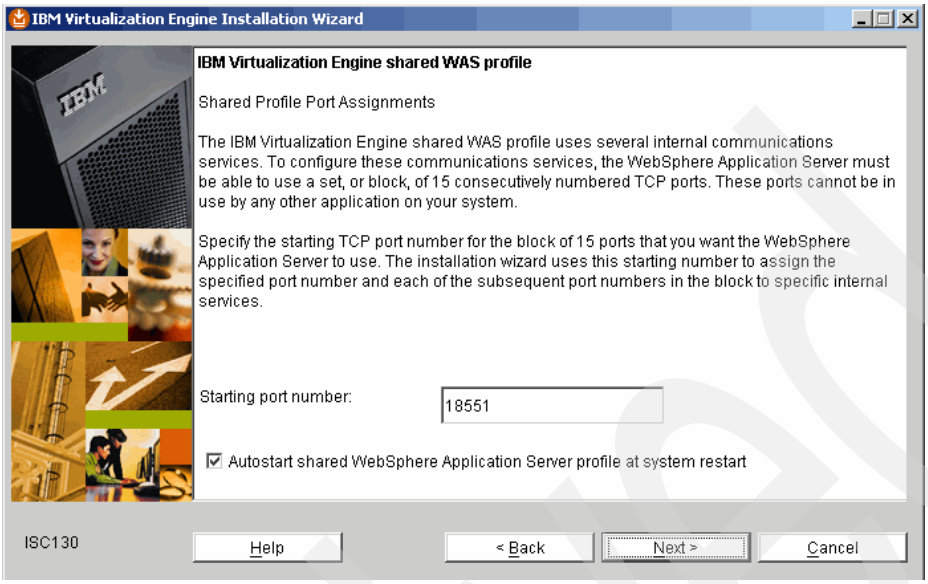


Figure 3-67 Shared WebSphere Application Server profile ports assignment

6. After the introduction window appears, the wizard will ask you to provide the Director Server information, such as the host name, as shown in Figure 3-68.

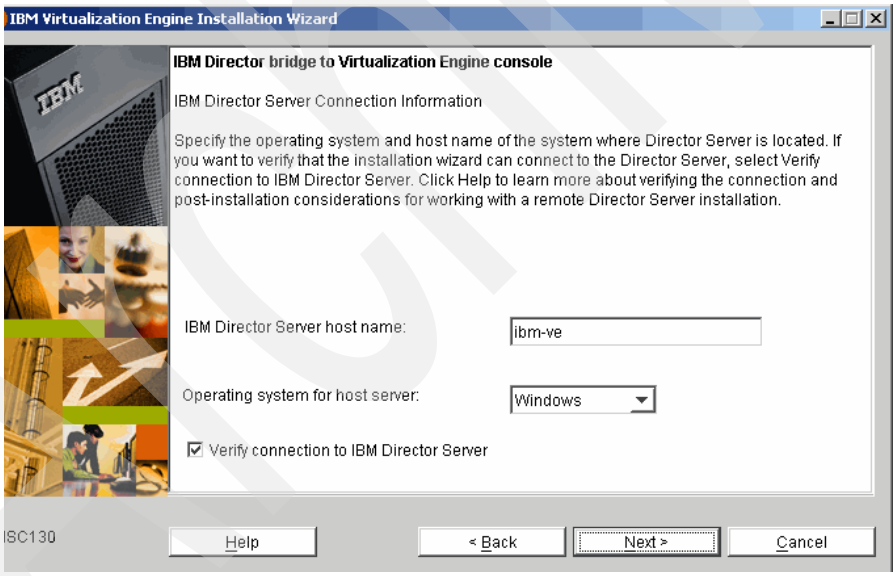


Figure 3-68 IBM Director Bridge connection information

Note: We recommend that you test the connectivity between the Director Server and the Director Bridge now, especially in the case where the two components are not installed on the same system.

7. For EWLM, you should provide the path information in which the installation will store the EWLM domain manager information, as shown in Figure 3-69 on page 118.

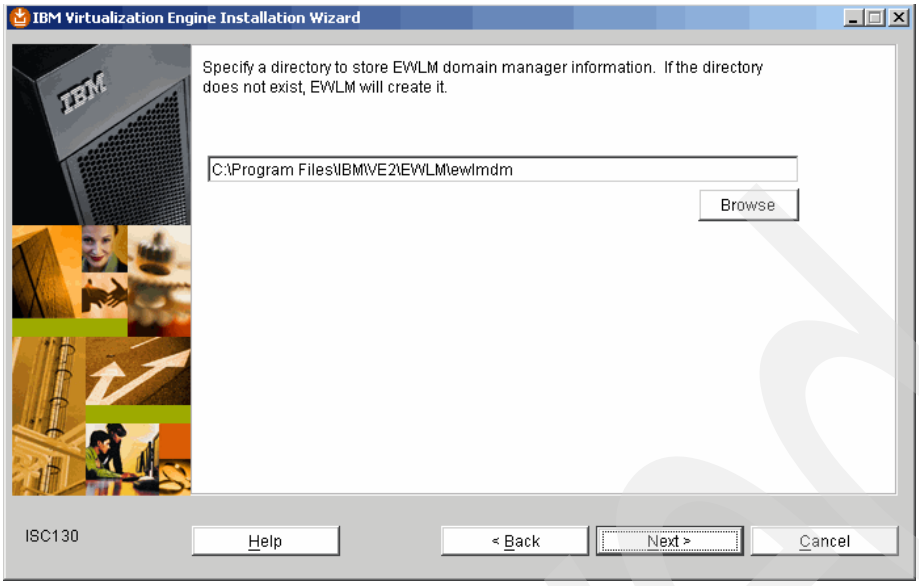


Figure 3-69 EWLM Domain Manager preference

Step 5: Installation

Once all these parameters have been provided, the installation process itself can begin. First, a summary window appears from which you can go back to change any parameters if necessary. To start the installation, click **Install**, as shown in Figure 3-70.

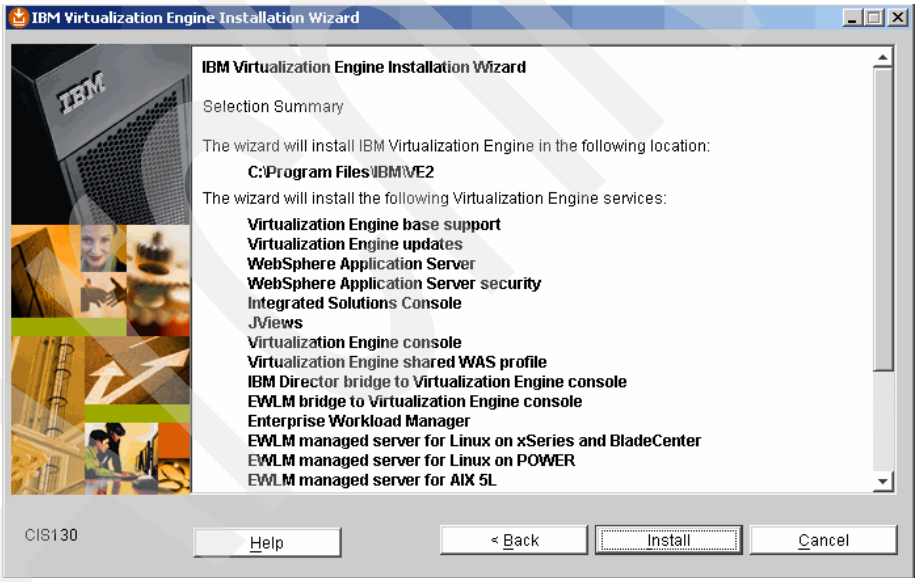


Figure 3-70 Selection summary for the installation

An installation progress bar window, as shown in Figure 3-71 on page 119, provides you with some information during the installation that can take one or two hours, depending on the system on which you install.

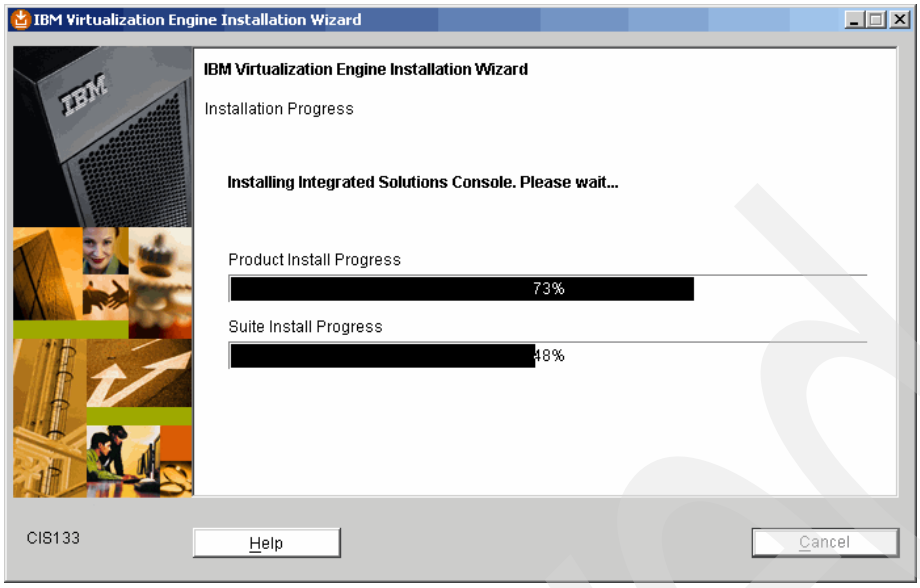


Figure 3-71 Installation progress

Step 6: Post-installation configuration

1. The post-installation tasks will inform you about the results of the installation, as shown in Figure 3-72, or show you the results for each of the components or their status, as shown in Figure 3-73 on page 120. You can also see a summary, as shown in Figure 3-74 on page 120.

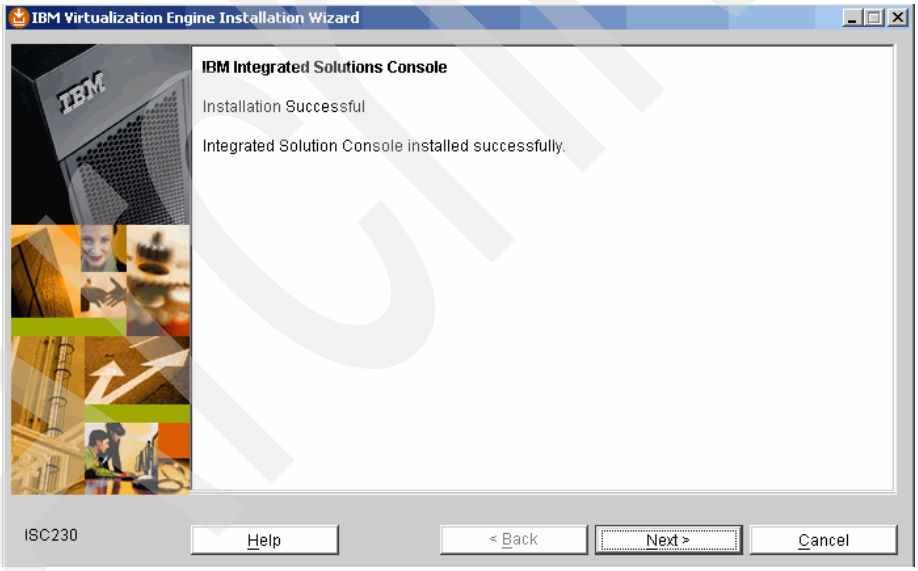


Figure 3-72 Installation result by component

Archived

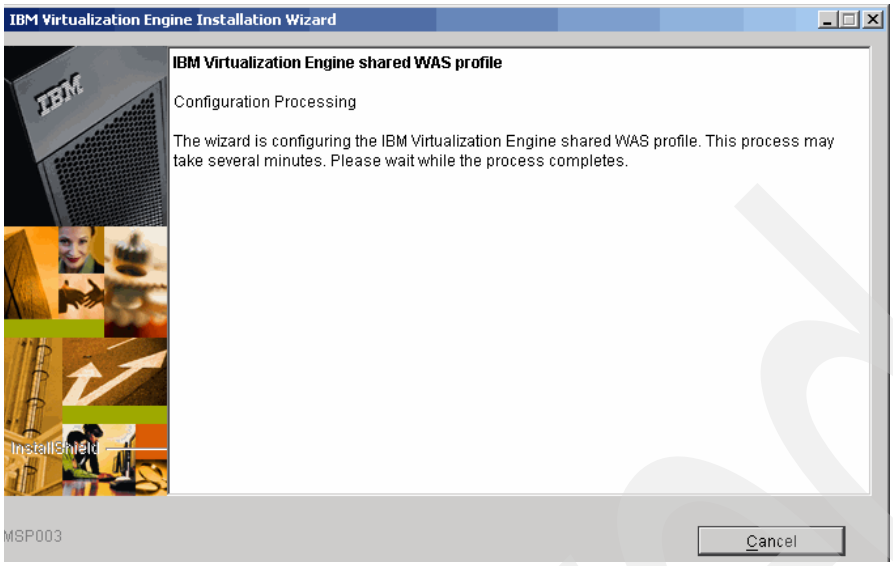


Figure 3-73 Configuration processing

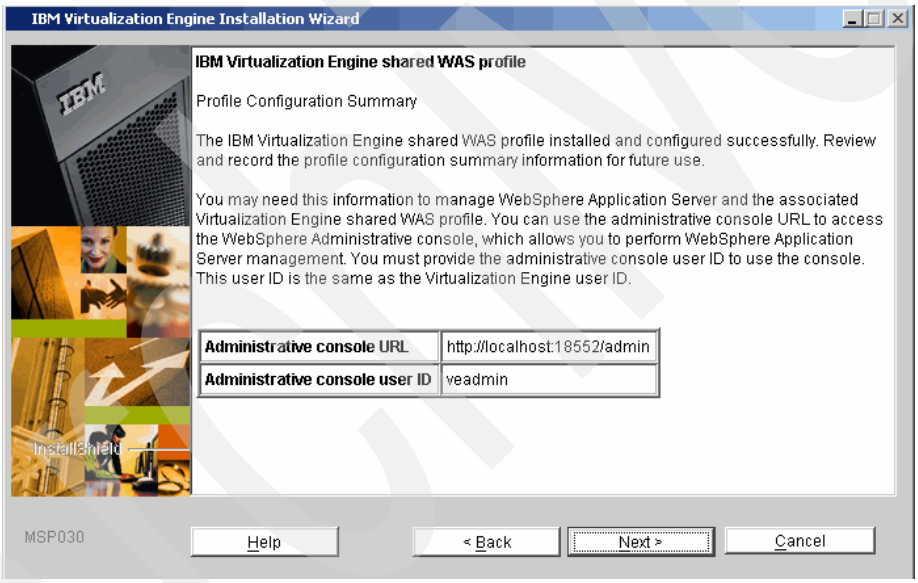


Figure 3-74 Configuration summary

2. For EWLM, you have the choice to configure the domain manager immediately or later as shown in Figure 3-75 on page 121.

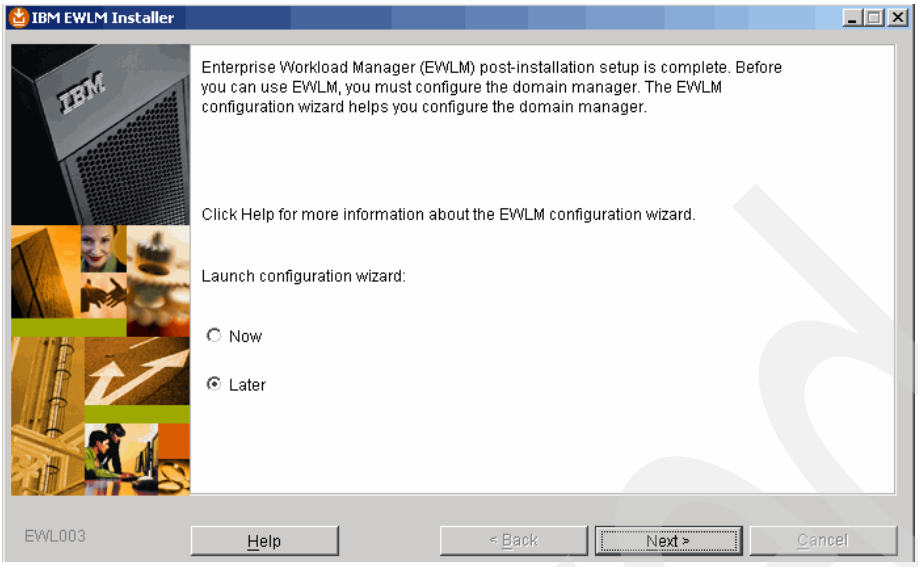


Figure 3-75 EWLM domain manager configuration choice

3. The last component that is installed is the Virtualization Engine console. The wizard will inform you about the progress of this installation, as shown in Figure 3-76.

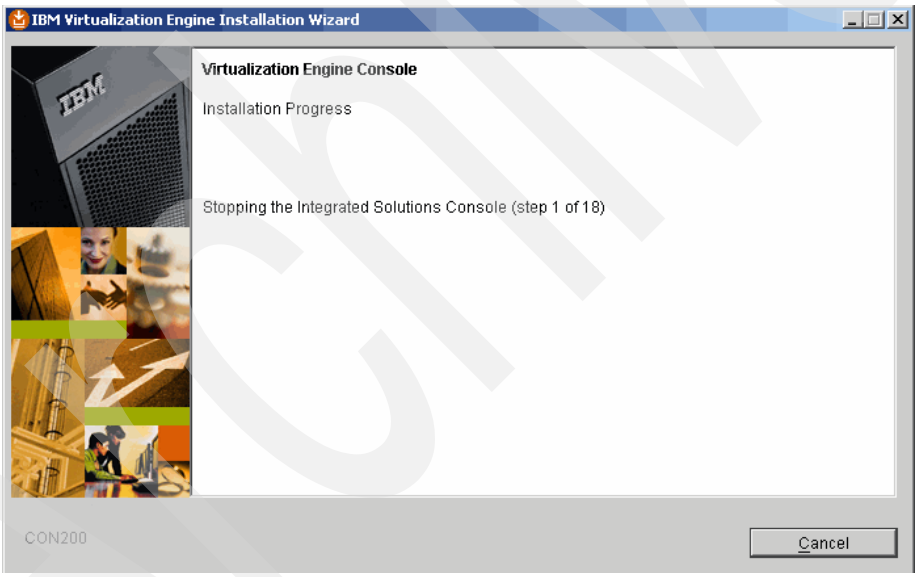


Figure 3-76 Virtualization Engine console installation steps

The components installation ends with the window displayed in Figure 3-77 on page 122, with a message that the Virtualization Engine installed successfully. This window allows you to launch the Virtualization Engine console to demonstrate that the installation was successful.

Archived

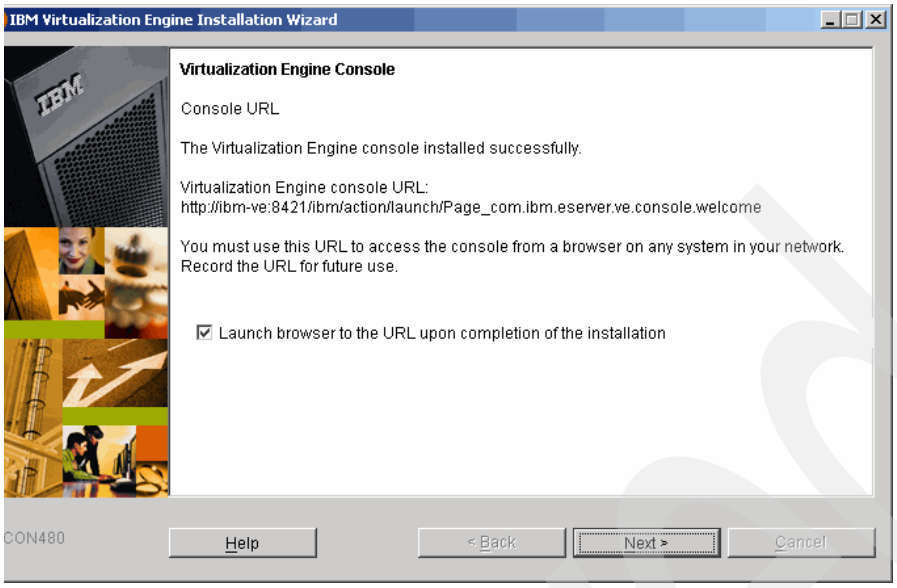


Figure 3-77 Virtualization Engine final installation window

At this point, you can log in to the console. In our ITSO environment, we used the following URL:

`http://hostname:8421/ibm/defaultconsole`

This displayed the window shown in Figure 3-78.

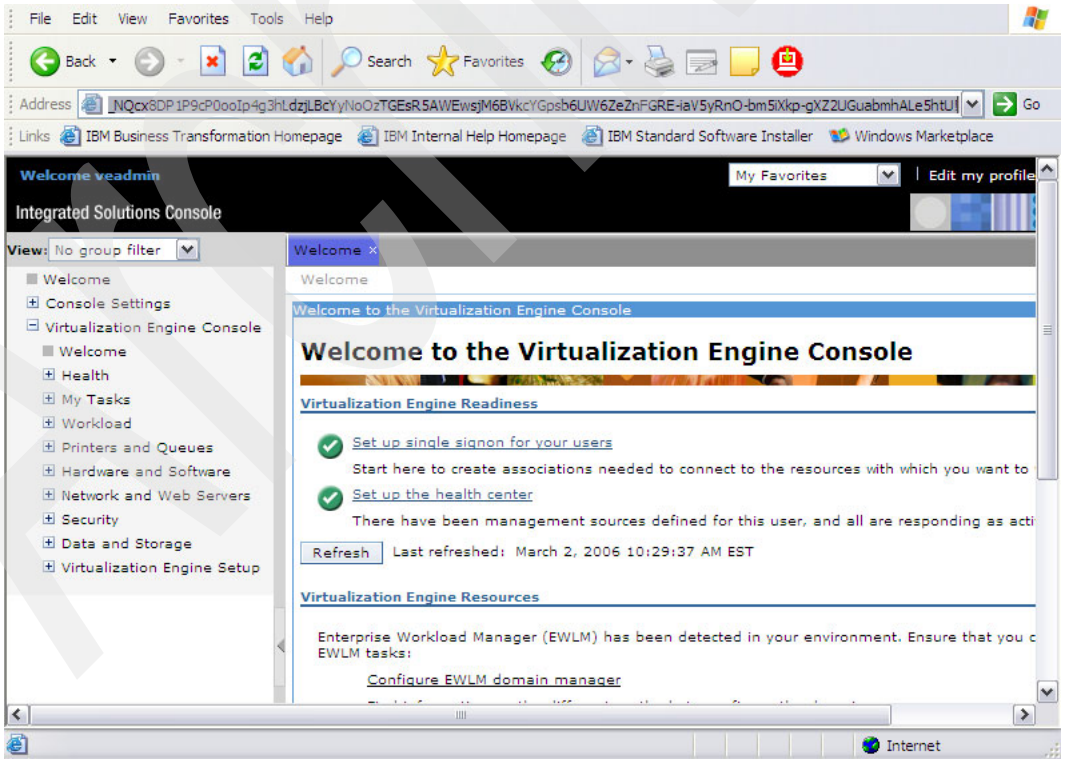


Figure 3-78 Welcome window to the Virtualization Engine console

Epilogue

Two tasks still have to be performed:

- The first one registers the installed components and a window appears, as shown in Figure 3-79.

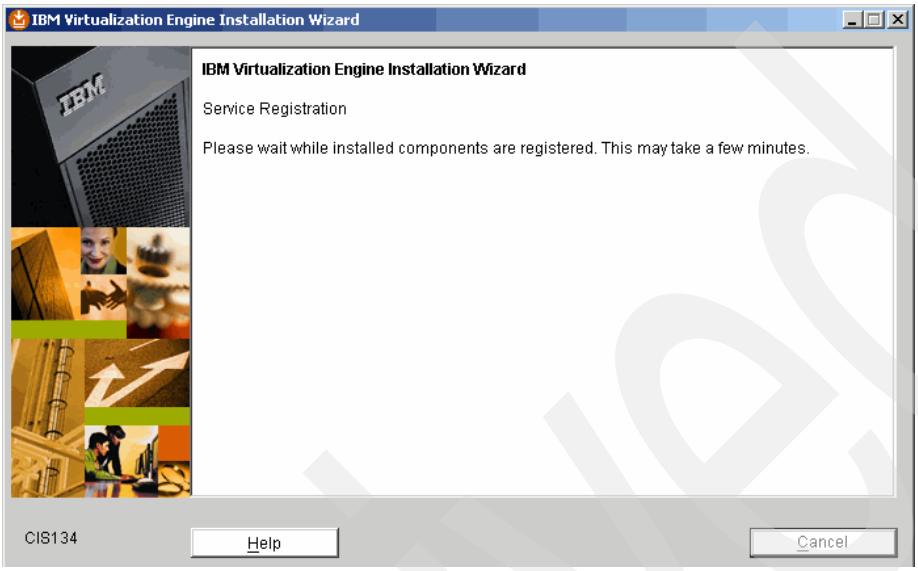


Figure 3-79 Service registration

- The other one wants to inform you that the installation completed successfully. A window is displayed, as shown in Figure 3-80, from which you can exit the Virtualization Engine wizard by clicking **Finish**.

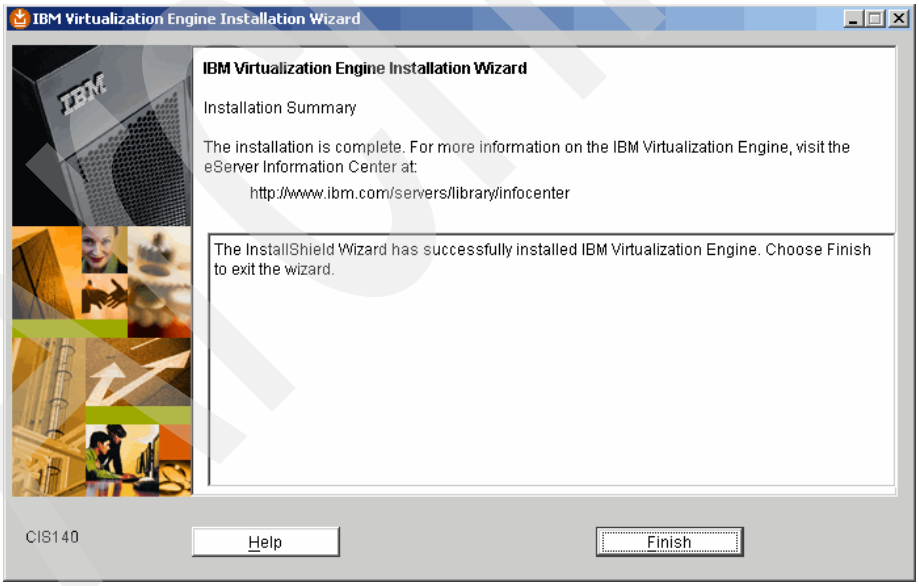


Figure 3-80 The final Virtualization Engine installation window

3.5 Silent installation

The silent installation allows you to run an unattended installation of any component of the IBM Virtualization Engine.

Silent installation is based on the customization of a few response files.

Note: We recommend that you always back up a response file before editing it, for example, by making a copy named <*filename*>.original.rsp.

This customization implies:

- Uncommenting lines
- Providing parameter values

Note: Providing values for passwords in the response files implies the use of the IBM Virtualization Engine HashPassword dialog that translates the clear text password into a hashed stream. Usage of the tool is documented on top of each response file.

As we choose to install the Director Server on the local system rather than using an existing Directory Server, the installation process must be performed in two steps:

- Installation of the Virtualization Engine Common Components (including Directory Server)
- Installation of the Virtualization Engine services

This is the same process as with the graphical interface.

3.5.1 Virtualization Engine common components

In a scenario where ITDS and DB2 need to be installed, you must edit the following response files:

- <*Virtualization Engine Source Directory*>/FILES/VEinstall.rsp

  Uncomment the following lines, as shown in Example 3-14:

  - `-P itds.active=true`

    This requests Directory Server installation.

  - `-P db2.active=true`

    This requests database server installation.

  - `-G licenseAccepted=true`

    This specifies the agreement to all licenses that are part of the installation process.

Example 3-14 The VEinstall.rsp file

```
-P itds.active=true
-P db2.active=true
-G licenseAccepted=true
```

- <*Virtualization Engine Source Directory*>/FILES/Common/DB2.rsp

  Uncomment the following lines and modify the parameter values appropriately, as shown in Example 3-15 on page 125:

  - `-G DB2_INSTALL_LOCALLY=true`

    This specifies the local installation of the database server.

– `-G DB2_NEW_USER=true`

  This requests automatic creation by the installer of the DB2 instance owner.

– `-G DB2_INSTANCE_OWNER_PASSWORD=<Password>`

  This specifies the password of the DB2 instance owner. The parameter value must be a hashed password.

Example 3-15 The DB2.rsp file

```
-G DB2_INSTALL_LOCALLY=true
-G DB2_NEW_USER=true
-G DB2_INSTANCE_OWNER_PASSWORD=oGOJ+dcv3dbrNGtqm8skrLl8lutBjB/Y4l0dbFGVnjk=
```

- <*Virtualization Engine Source Directory*>/FILES/Common/ITDS.rsp

  Uncomment the following lines and modify the parameter values, as shown in Example 3-16:

  – `-G ITDS_DIRECTORY_SERVER_ADMINISTRATOR_USER_ID_PASSWORD=<Password>`

    This specifies the password for the Directory Server administrator. The parameter value must be a hashed password.

  – `-G ITDS_DIRECTORY_SERVER_INSTANCE_OWNER_USER_ID=<Username>`

    This specifies the user name of the Directory Server instance owner.

  – `-G ITDS_DIRECTORY_SERVER_INSTANCE_OWNER_USER_ID_PASSWORD=<Password>`

    This specifies the password of the Directory Server instance owner. The parameter value must be a hashed password.

  – `-G ITDS_ENCRYPTION_SEED_STRING=<Encryption Seed String>`

    This specifies an encryption seed to be used by the Directory Server. The parameter value must be a hashed password.

Example 3-16 The ITDS.rsp file

```
-G
ITDS_DIRECTORY_SERVER_ADMINISTRATOR_USER_ID_PASSWORD=sJFexRnyeYQKKFzOfi7IfR1S/GTXd
O9TcjnZKGL3luE=
-G ITDS_DIRECTORY_SERVER_INSTANCE_OWNER_USER_ID=vetdsur
 -G
ITDS_DIRECTORY_SERVER_INSTANCE_OWNER_USER_ID_PASSWORD=HyZMDG5LX8fGCLvbDsgPnL44ypfA
EZiEhIhUZ2oH94A=
-G ITDS_ENCRYPTION_SEED_STRING=mFRkVWNzrrkZkvwy3f7IuqtOI/aQnFDLKTyYnJRYBcc=
```

The following command launches the installation process:

```
<Virtualization Engine Source Directory>/FILES/InstallVEWin.exe -optionsve -silent
```

The installation process is logged in the following files:

- <*Virtualization Engine Installation Directory*>/Logs/installLog.txt
- <*Virtualization Engine Installation Directory*>/Logs/installDebug.txt

Archived

3.5.2 Virtualization Engine Services

You must edit the following response files:

- <*Virtualization Engine Source Directory*>/FILES/VEinstall.rsp

  Uncomment the following lines, as shown in Example 3-17:

  - `-P ewlm<Platform>.active=true`

    This requests the EWLM installation.

  - `-P vec.active=true`

    This requests the Virtualization Engine Console installation.

  - `-P vecewlm.active=true`

    This requests the EWLM Bridge to Virtualization Engine Console installation.

  - `-P vecdir.active=true`

    This requests the Director Bridge to Virtualization Engine Console installation.

  - `-G licenseAccepted=true`

    This specifies the agreement to all licenses that are part of the installation process.

Example 3-17 The VEinstall.rsp file

```
-P ewlmwin.active=true
-P vec.active=true
-G licenseAccepted=true
```

- <*Virtualization Engine Source Directory*>/FILES/Common/VEEnv.rsp

  Uncomment the following lines and modify the parameter values appropriately, as shown in Example 3-18 on page 127:

  - `-G VEENV_INSTALL_INTO_EXISTING_ENVIRONMENT=false`

    This requests the creation of a new Virtualization Engine Environment.

  - `-G VEENV_DIRECTORY_SERVER_HOST_NAME=<Hostname>`

    This specifies the host name of the Directory Server to host the Virtualization Engine Environment.

  - `-G VEENV_DIRECTORY_SERVER_PORT=389`

    This specifies the Directory Server port. 389 is the default value.

  - `-G VEENV_DIRECTORY_SERVER_ADMIN_USER_ID=cn=Admin`

    This specifies the Directory Server Administrator User. cn=Admin is the default value.

  - `-G VEENV_DIRECTORY_SERVER_ADMIN_USER_PASSWORD=<Password>`

    This specifies the password of the Directory Server administrator user. The parameter value must be a hashed password.

  - `-G VEENV_NEW_ENVIRONMENT_NAME=<Virtualization Engine Environment>`

    This specifies a name for the Virtualization Engine Environment.

  - `-G VEENV_NEW_ENVIRONMENT_DESCRIPTION=<Description>`

    This specifies a description of the Virtualization Engine Environment.

Archived

Example 3-18 The VEEnv.rsp file

```
-G VEENV_INSTALL_INTO_EXISTING_ENVIRONMENT=false
-G VEENV_DIRECTORY_SERVER_HOST_NAME=ibm-ve
-G VEENV_DIRECTORY_SERVER_PORT=389
-G VEENV_DIRECTORY_SERVER_ADMIN_USER_ID=cn=Admin
-G
VEENV_DIRECTORY_SERVER_ADMIN_USER_PASSWORD=sJFexRnyeYQKKFzOfi7IfR1S/GTXdO9TcjnZKGL
3luE=
-G VEENV_NEW_ENVIRONMENT_NAME=virtualization
-G VEENV_NEW_ENVIRONMENT_DESCRIPTION=Virtualization Environment
```

Warning: The set of parameter values mentioned above leads to the creation of a new Virtualization Engine Environment inside the Directory Server. If the installation process is re-run later (for adding new components, for example), you must modify the parameter values in order to request the use of the existing Virtualization Engine Environment.

- <*Virtualization Engine Source Directory*>/FILES/Common/WASSecurity.rsp

  Uncomment the following lines and modify the parameter values appropriately, as shown in Example 3-19:

  - `-G WASSEC_VIRTUALIZATION_ENGINE_ADMINISTRATOR_USER_PASSWORD=<Password>`

    This specifies a password for the Virtualization Engine administrator (veadmin). The parameter value must be a hashed password.

  - `-G WASSEC_LTPA_PASSWORD=<Password>`

    This specifies a password for the LTPA. The parameter value must be a hashed password.

Example 3-19 The WASSecurity.rsp file

```
-G
WASSEC_VIRTUALIZATION_ENGINE_ADMINISTRATOR_USER_PASSWORD=TKIISfyEV7gxraOI6G66CEhf9
IzFj5hUlYNb/9MJOpI=
-G WASSEC_LTPA_PASSWORD=CusjeiZ3sS+fQYIjbGNT9OfWbscrBO67Gg41VEH36xw=
```

- <*Virtualization Engine Source Directory*>/FILES/VEC/VECDir.rsp

  Uncomment the following lines and modify the parameter values appropriately, as shown in Example 3-20:

  - `-G DirServerPlatform=<Platform name>`

    This specifies the platform type on which the Director Server is installed.

Example 3-20 The VECDir.rsp file

```
-G DirServerPlatform=WINDOWS
```

The following command launches the installation process:

`<Virtualization Engine Source Directory>/FILES/InstallVEWin.exe -optionsve -silent`

The installation process is logged in the following files:

- <*Virtualization Engine Installation Directory*>/Logs/installLog.txt
- <*Virtualization Engine Installation Directory*>/Logs/installDebug.txt

You can get more information about the Virtualization Engine installation from “Virtualization Engine management collection Version 2 Release 1”, which is found in the Information Center (described in Example 2.2.1 on page 34).

3.6 Post-installation checking tasks

These post-installation sections will help you check your installation.

3.6.1 About the Directory Server

This section provides some information to be sure the Directory Server is correctly installed.

Post-installation configuration

The Directory Server does not require any post-installation configuration.

Startup

You can start the Directory Server either in manual mode or in automatic mode.

Manual mode

The command to start the Directory Server is:

- For AIX**:/opt/ibm/ldap/V6.0/sbin/ibmslapd -I <*Instance Owner*>**
- For Linux: **/opt/ibm/ldap/V6.0/sbin/ibmslapd -I <*Instance Owner*>**
- For Microsoft Windows; **C:\idsslapd-<*Instance Owner*>\etc\ibmslapdService.cmd <Instance Owner>** or using the **idsslapd-<*Instance Owner*>** service.

Example 3-21 shows an example in an AIX or Linux environment.

Example 3-21 Starting the Directory Server on AIX / Linux systems

```
/opt/ibm/ldap/V6.0/sbin/ibmslapd -I vedsur
```

Note: The DB2 database instance created during the installation of the Directory Server must run before the Directory Server can start successfully. The startup of the DB2 instance is configured for an automatic start by default on all the platforms.

The command to start the Directory Server Web Administration Tool is:

- For AIX:

  ```
  <Virtualization Engine Installation
  Directory>/WebSphere/AppServer/profiles/IDSWebApp/bin/startServer.sh server1
  -profileName IDSWebApp
  ```

- For Linux:

  ```
  <Virtualization Engine Installation
  Directory>/WebSphere/AppServer/profiles/IDSWebApp/bin/startServer.sh server1
  -profileName IDSWebApp
  ```

- For Microsoft Windows:

  ```
  <Virtualization Engine Installation
  Directory>\WebSphere\AppServer\profiles\IDSWebApp\bin\startServer.sh server1
  -profileName IDSWebApp
  ```

Archived

Example 3-22 shows an example for an AIX or Linux environment.

Example 3-22 Starting the Directory Server Web Administration Tool on AIX / Linux systems

```
/opt/IBM/VE2/WebSphere/AppServer/profiles/IDSWebApp/bin/startServer.sh server1
-profileName IDSWebApp
```

Note: If required, you can get the connection to the Directory Server Web Administration Tool at:

`http://<hostname>:12100/IDSWebApp/IDSjsp/Login.jsp`

Automatic mode

On AIX or Linux systems, the installer automatically adds an entry in the inittab, as shown in Example 3-23. The startup is thus automated at boot time.

Example 3-23 Directory Server entry in inittab for AIX / Linux systems

```
veit:2345:once:/opt/ibm/ldap/V6.0/sbin/ibmslapd -I idsuser > /dev/console 2>&1
```

On Microsoft Windows systems, the idsslapd-<*Instance Owner*> Directory Server service is configured by default with a manual startup type. You should change this to automatic.

Post-installation checkup

One eventual functional checkup you should perform is opening a connection to the Directory Server administration console and browsing through the registries it contains.

3.6.2 EWLM Domain Manager

This section provides some information to ensure that the EWLM Domain Manager is well installed.

Post-installation configuration

Once you have installed the EWLM Domain Manager, you should configure it. This step also implies the creation of the Control Center, which is the Web user interface to access the EWLM Domain Manager.

This configuration can be performed in two ways, either using a graphical configuration wizard (this requires a graphical user interface available on the system) or using commands manually. The following section describes the first option.

Graphical configuration wizard

The graphical configuration wizard is launched by using the following commands:

- For AIX: **<*Installation Directory*>/EWLM/bin/configWizardDM.sh**
- For Linux: **<*Installation Directory*>/EWLM/bin/configWizardDM.sh**
- For Microsoft Windows: **<*Installation Directory*>\EWLM\bin\configWizardDM.bat**

Figure 3-24 shows an example of the commands in an AIX environment.

Example 3-24 Launching EWLM Domain Manager configuration wizard on AIX systems

```
/opt/IBM/VE2/EWLM/bin/configWizardDM.sh
```

Example 3-25 shows an example of the commands on a Windows Microsoft environment.

Example 3-25 Launching EWLM Domain Manager configuration wizard on Microsoft Windows systems

```
C:\Program Files\IBM\VE2\EWLM\bin\configWizardDM.bat
```

To configure the EWLM domain manager using the graphical configuration wizard, perform the following steps:

1. Create the domain manager and Control Center.

   The window shown in Figure 3-81 allows you to select the action to perform.

   Note: The configuration wizard can be used later on for modifying the Domain Manager configuration.

   Select **Create domain manager and Control Center**, as shown in Figure 3-81, and click **Next**.

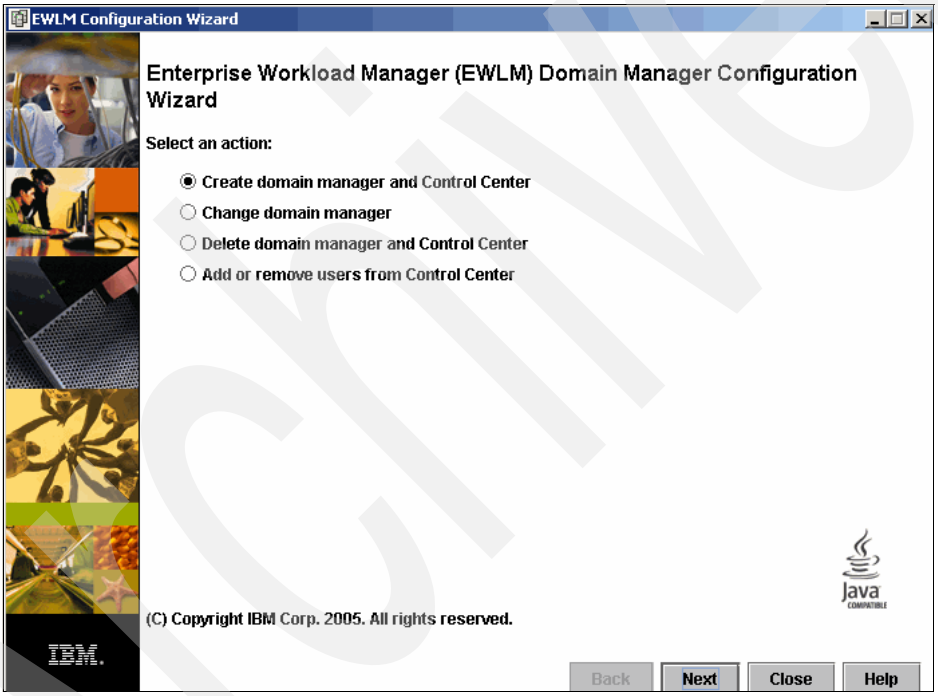


Figure 3-81 Create domain manager and Control Center

2. Specify the configuration ID.

   The configuration ID, as displayed in Figure 3-82 on page 131, constitutes the unique identifier of the EWLM domain manager configuration settings. Click **Next**.

Archived

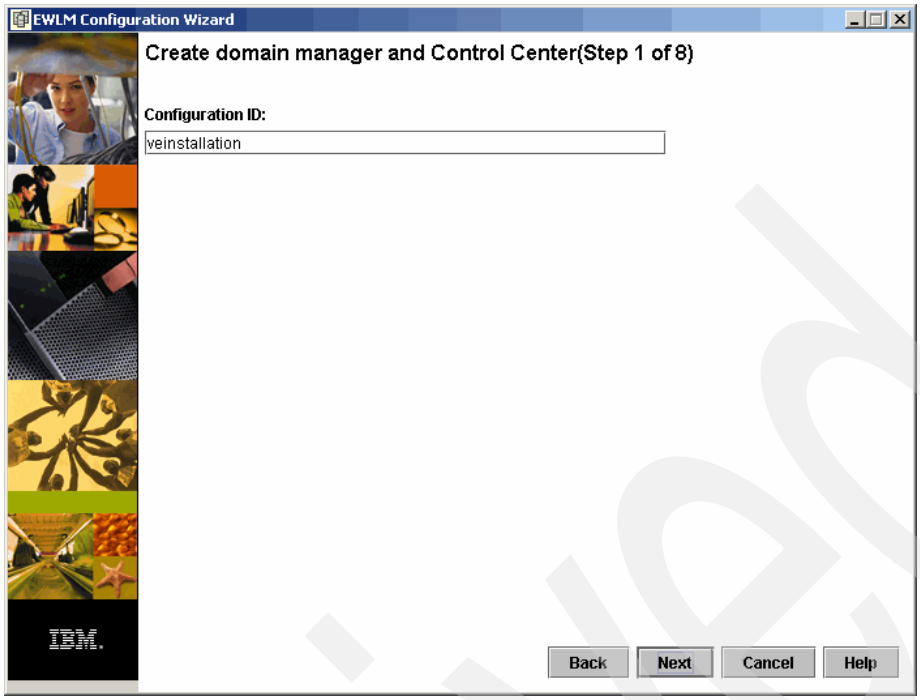


Figure 3-82 Configuration ID

3. Specify the domain manager settings.

   This window, as shown in Figure 3-83 on page 132, allows you to specify the following main settings of the domain manager:

   – The Domain name: A name representing the nature of the domain to manage from an EWLM point of view.

   – The Domain Manager IP address: The IP address or the host name of the domain manager.

   – The default port values are acceptable and should be kept as is.

   Click **Next**.

Archived

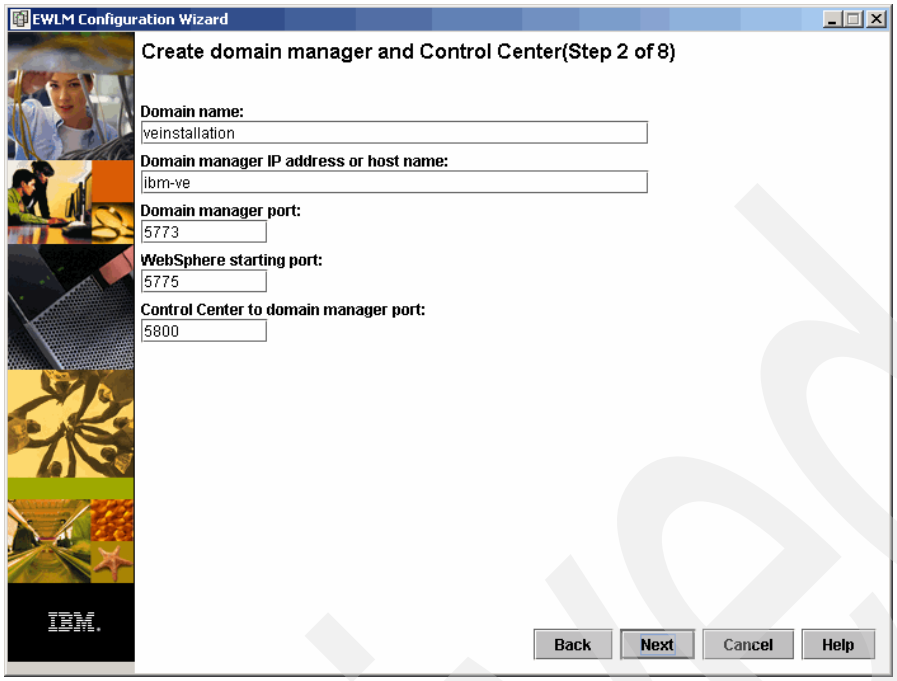


Figure 3-83 Domain manager settings

4. Set the security level.

   Select the appropriate level of security for your environment in terms of communication flows between the EWLM Domain Manager and the EWLM Managed Servers, as shown in Figure 3-84. Click **Next**.

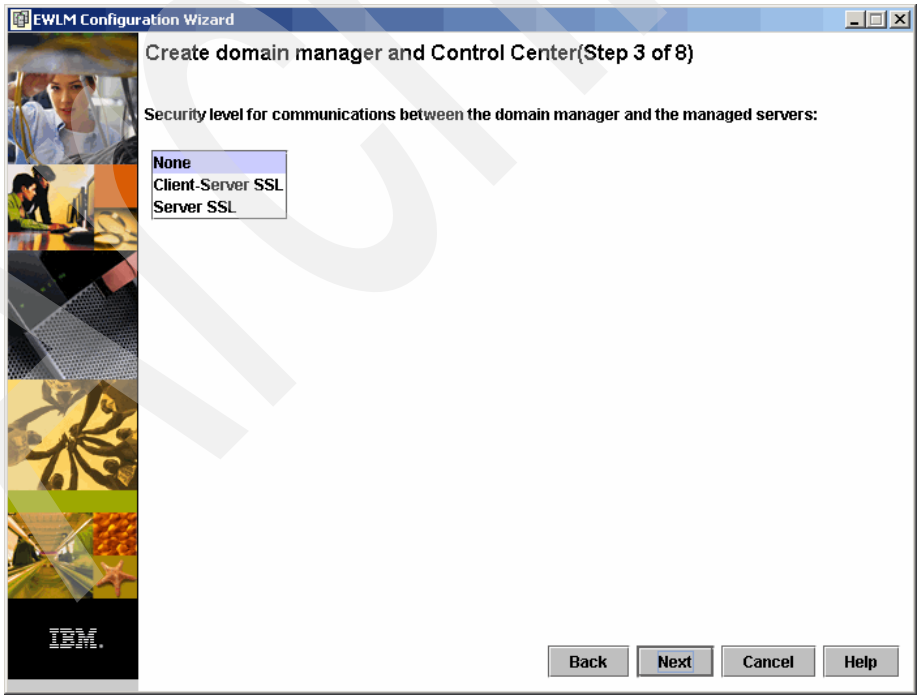


Figure 3-84 Security level

5. Define the Secure Sockets Layer (SSL) certificates.

The SSL certificates secure the communication between the EWLM Domain Manager and the EWLM Control Center. Unless you want to define the SSL certificates by yourself, keep the default selection, as shown in Figure 3-85. Click **Next**.

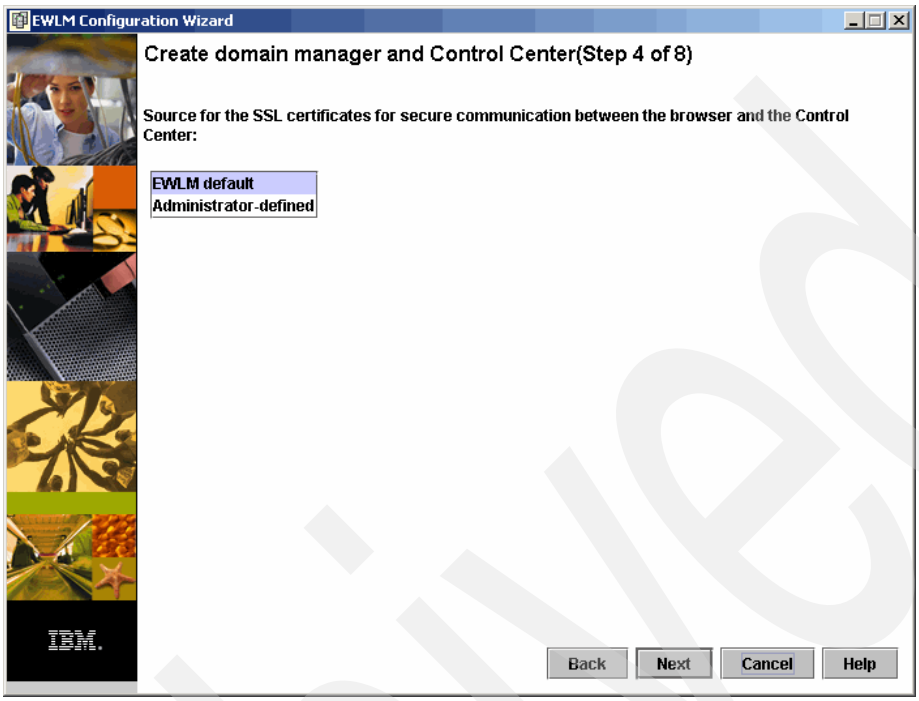


Figure 3-85 SSL certificates

6. Configure the firewall broker.

   The firewall broker is a feature that enables the communication between the EWLM Domain Manager and the EWLM Managed Servers through a firewall. If this is the case in your environment, you must specify the configuration of the firewall broker here, as shown in Figure 3-86 on page 134. Click **Next**.

Archived

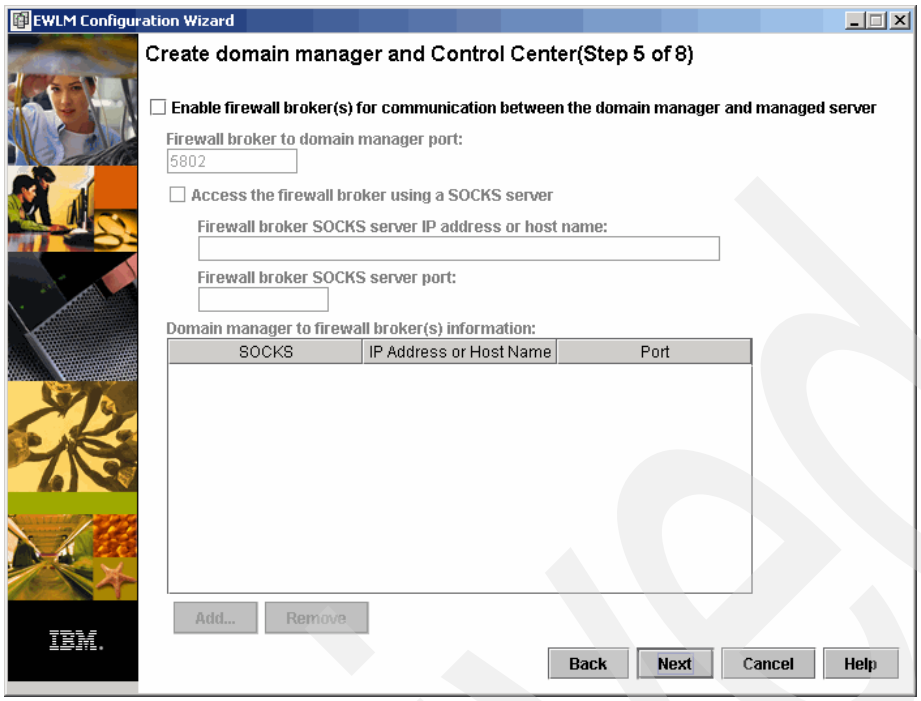


Figure 3-86 Firewall Broker

7. Set the WebSphere administrator user ID and password.

   In this window, as shown in Figure 3-87, you should provide the WebSphere administrator user ID and password. You should have created the WebSphere administrator user as a prerequisite to the installation. Click **Next**.

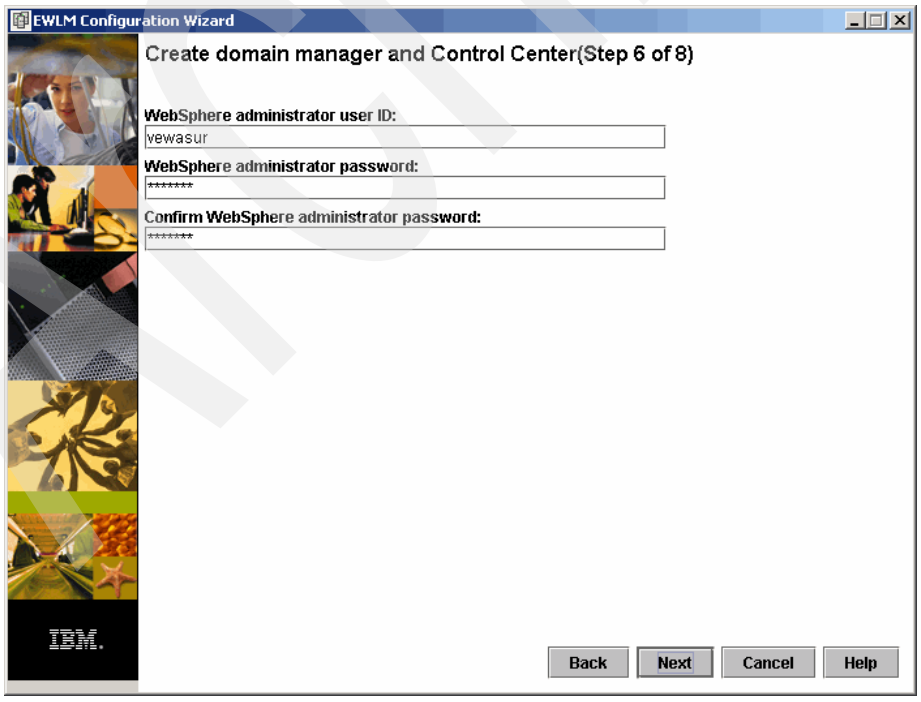


Figure 3-87 WebSphere administrator

8. Assign the Control Center users.

   This window, as shown in Figure 3-88, helps to define the users and groups that are allowed to connect to the EWLM Control Center, and the level or rights they will get once connected.

   Note: All users and groups that you specify in this window must match existing local user or group accounts.

   You can assign EWLM Control Center roles to either individual users in the Add Users tab, or to a group of users in the Add Groups tab. They provide the following level of authorizations:

   – Administrator: Full rights on the EWLM Control Center, including configuration of the EWLM Domain Manager.
   – Operator: Monitoring rights extended with management of service policies.
   – Monitor: Monitoring rights only.

   Click **Next** when you have completed this task.

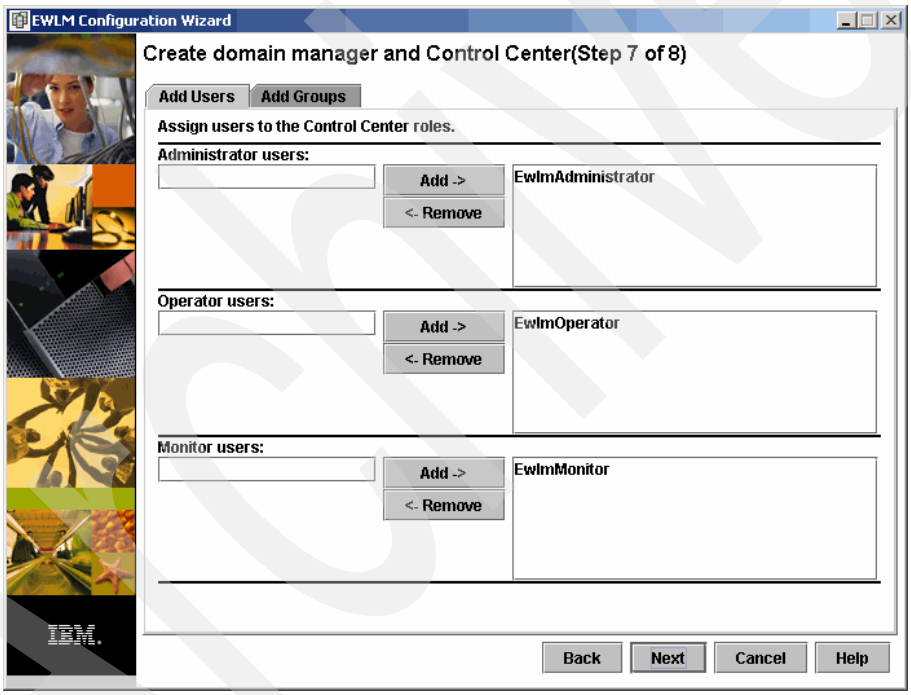


Figure 3-88 Control Center users

9. Start the parameters for the domain manager and the Control Center.

   The wizard can start the EWLM Domain Manager and the EWLM Control Center automatically if you select the check box, as shown in Figure 3-89 on page 136. Then click **Finish** to start the configuration.

Archived

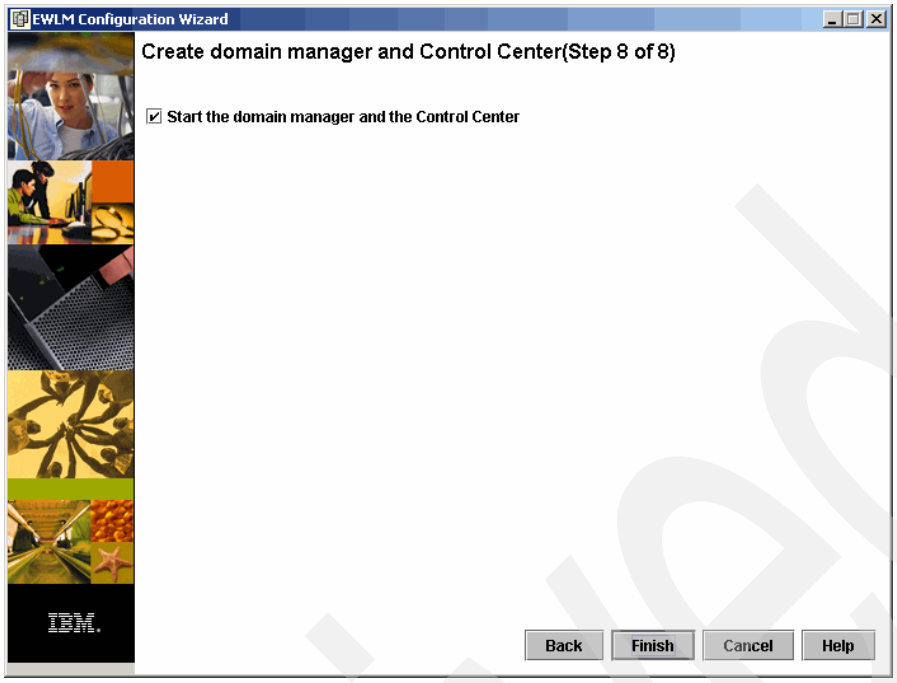


Figure 3-89 Starting the domain manager and Control Center

The configuration then starts the creation of the domain manager with the parameters that you have specified in the wizard and will display some information, as shown in Figure 3-90 and Figure 3-91 on page 137. Note that this operation is time consuming.

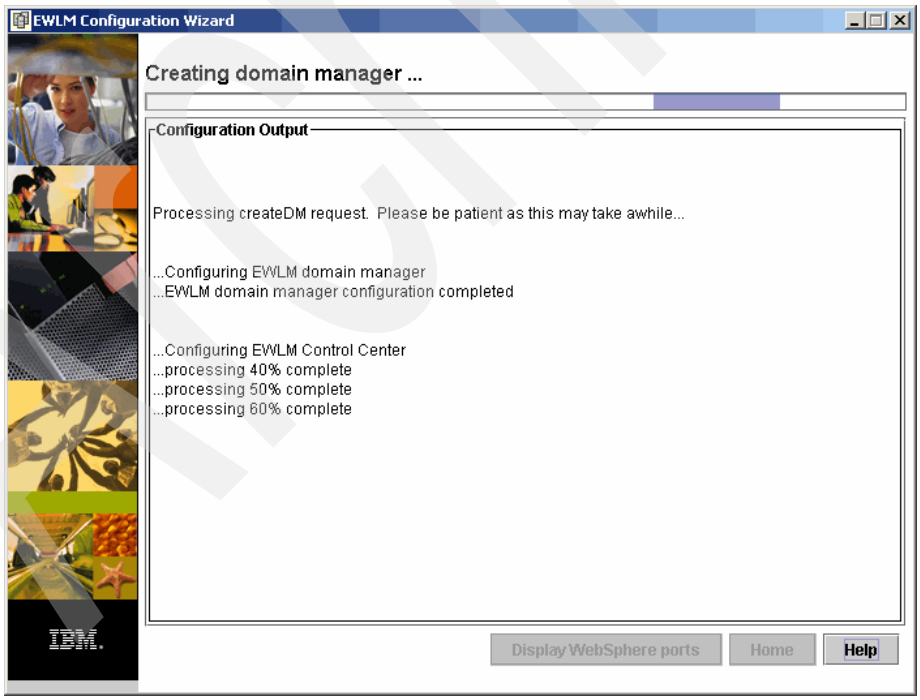


Figure 3-90 Domain manager - output configuration

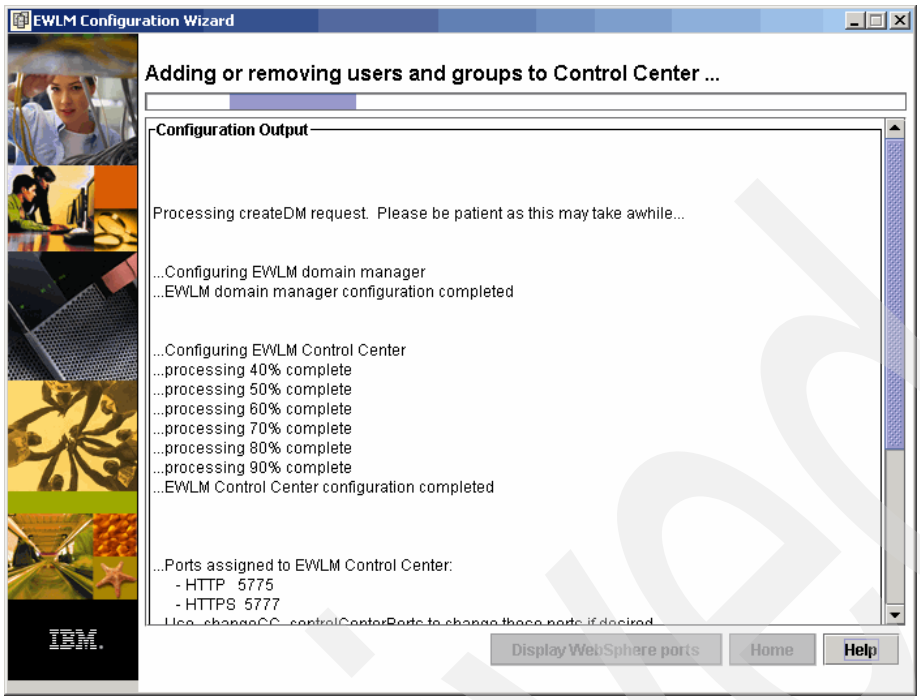


Figure 3-91 Adding users

When the domain manager configuration is complete, the information shown in Figure 3-92 is displayed. Click **OK**.

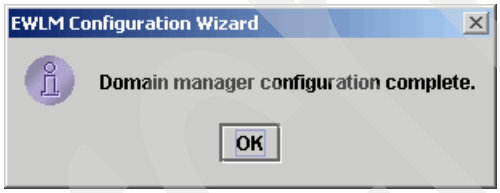


Figure 3-92 Domain manager configuration complete

Considerations for IBM System p EWLM Partition Management configuration

EWLM V 2.1 provides the capability to automatically adjust the amount of processing units among partitions defined in a partition workload group for an IBM System p. By comparing the actual performance of each partition to the performance defined in the domain policy, EWLM adjusts the resource assignment (cp) in order ensure that the goals defined in the domain policy are met.

Note: These configuration tasks only apply to IBM System p with POWER5 systems.

You must complete the following steps in order to enable the logical partition management:

1. Configure the partition profile to specify a partition workload group.

   To monitor a partition, EWLM requires that the partition profile indicate that it belongs to a *partition workload group*. EWLM can adjust the resources of partitions that belong to a partition workload group. EWLM automatically identifies the partitions to manage if the partition is defined a partition workload group.

Archived

To define the partition in a partition workload group, you must complete the following steps:

a. Log in to the Hardware Management console.

b. Select **Server and Partition** in the navigation area.

c. Select **Server Management**.

d. In the contents area, open the server on which you want to create or to change a logical partition profile.

e. If you want to create a new partition, complete the following steps:

   i. Right-click **Partitions**, and select **Create** → **Logical Partition**.

   ii. Complete the Create Logical Partition wizard to create a logical partition. On the Create Logical Partition - Workload groups page, specify a Partition Workload Group. Figure 3-93 indicates where to define the logical partition workload group.

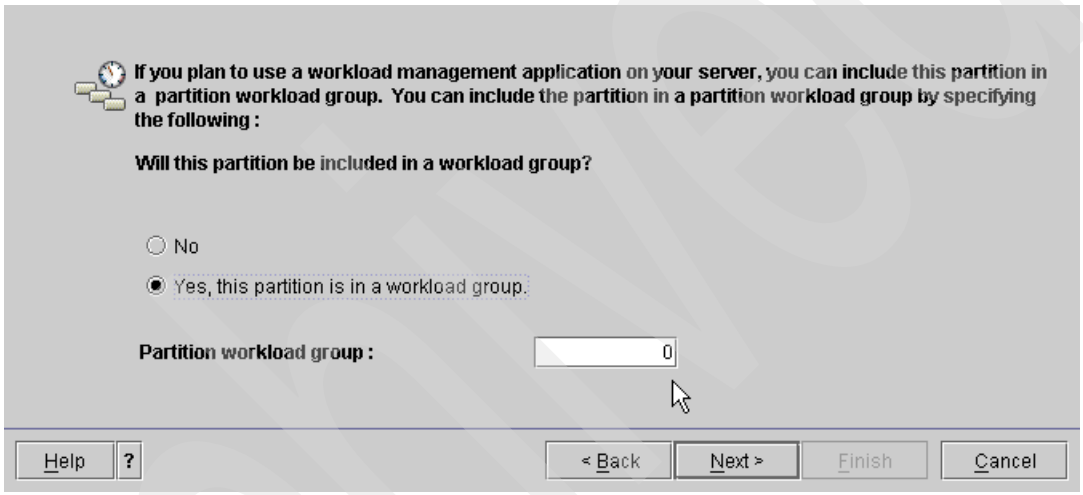


Figure 3-93 Setting the partition workload group during the LPAR creation process

   iii. On the Create Logical Partition Profile - Processors page, select **Shared**. This enables the EWLM to move processing power from the shared processor pool to this partition, if necessary.

   iv. Right-click the logical partition profile and select **Properties**, then select **Hardware** → **Processors and Memory**. Ensure that the **Allow shared processor pool utilization authority** is selected. By selecting this option, EWLM has the proper authority to make resources adjustment, if necessary.

> **Note:** This change only applies to the logical partition properties. You must make a similar change to the Partition Profile if you want to apply this change to all subsequent profile activation and deactivation requests.

If you want to change the existing partition properties, complete the following steps:

i. Right-click the logical partition profile **Properties**.

ii. On the other page, ensure that a Partition workload group is specified. Figure 3-94 shows the Partition Properties page that indicates to which partition workload group this logical partition belongs.

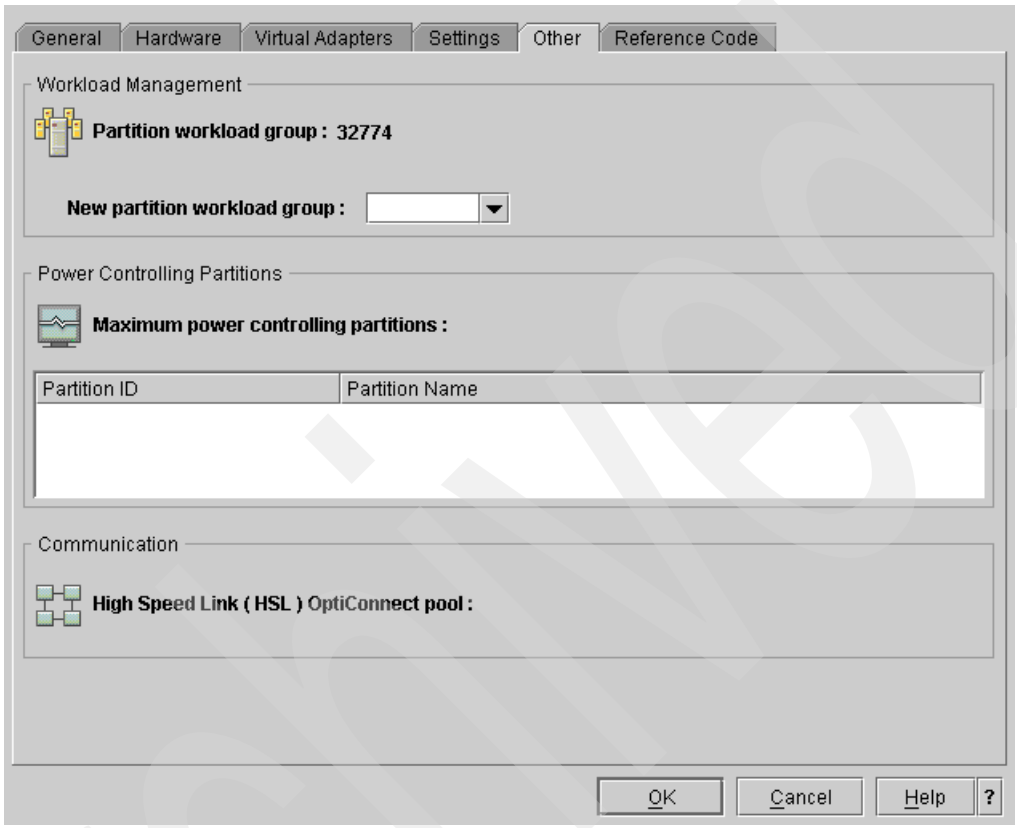


Figure 3-94 Partition properties

iii. Set the EWLM partition workload group.

iv. Restart the EWLM managed server process on the partition.

Repeat this process for each logical partition that you want to define into the Partition workload group.

2. Change the CPU and memory limits on LPAR.

   If you want to change the amount of resources automatically by using an EWLM partition group manager, you should define the maximum and minimum resources initially. After this, the EWLM will move resources between the defined minimum and maximum limits.

   To change the parameters:

   a. Log in to the Hardware Management Console (as hscroot).

   b. Select **Server and Partition** → **Server Management** → **Partitions**.

   c. Right-click the partition profile, and select **Properties**.

   d. Change the minimum and maximum capability values as desired.

   e. Click **OK** to accept the changes.

Figure 3-95 shows the possible options for memory resources. When the amount of resources reaches Minimum, EWLM will not be able to remove resources for this partition. Maximum means that EWLM will not add more resources. The Desired option means that this is the default, and these resources are allocated to the partition after a reboot. With EWLM V2.1, only the CPU resources can be moved between partitions. Moving memory and I/O channels is not possible.

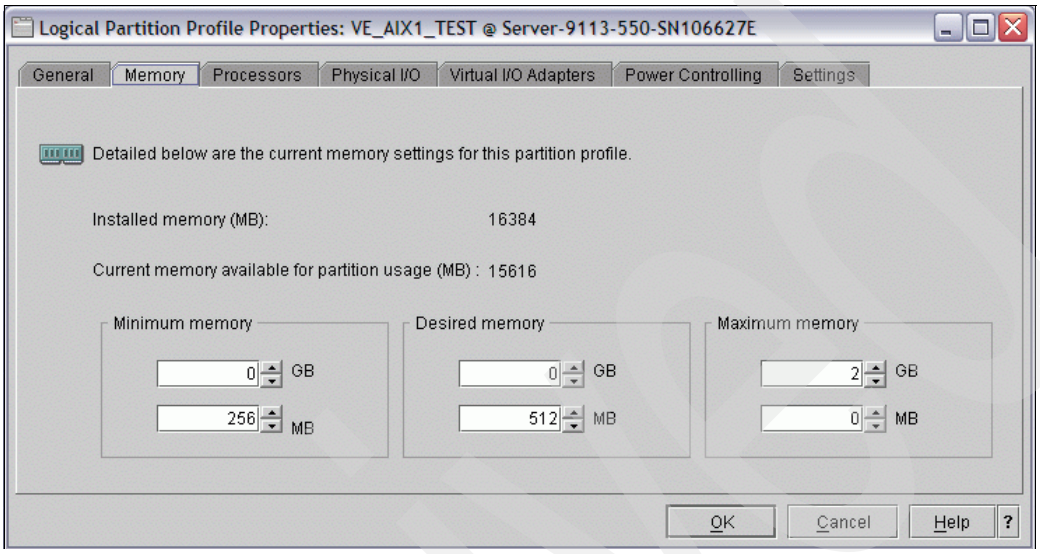


Figure 3-95 Memory properties on the LPAR profile

Figure 3-96 on page 141 shows the parameters concerning processors. We recommend that you use **Shared processing mode** and **Uncapped sharing mode**. This means that this partition is able to use the resources if another partition does not use it.

Archive

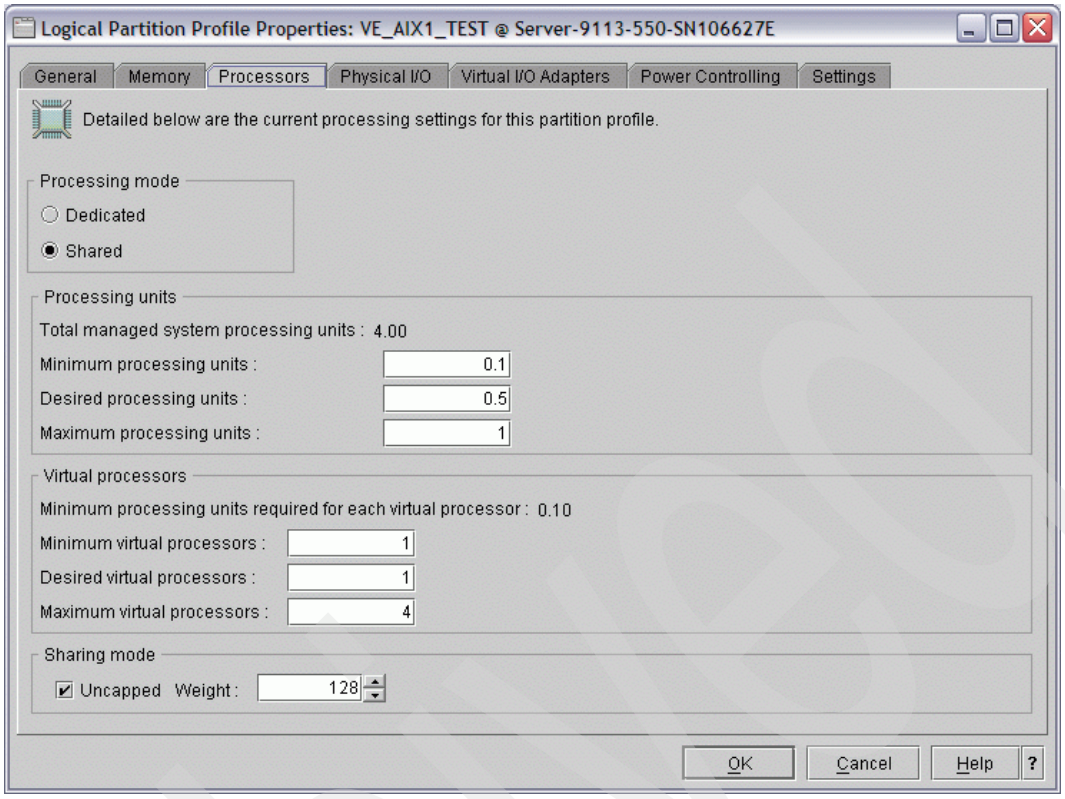


Figure 3-96 Processors properties on the LPAR profile

When you have finished configuring the logical partitions, you can enable the partition workload management using the EWLM Control Center.

3. Enable the partition workload group management.

   Before using the partition management in the EWLM console, you must enable it. Complete the following steps to enable the partition management:

   a. Log in to EWLM Control Center as Administrator.

   b. Navigate to **Domain Settings** → **EWLM management** → **Enable partition workload group management**.

   c. Select **Enable partition workload group management** and click **OK** to accept the changes, as shown in Figure 3-97 on page 142.

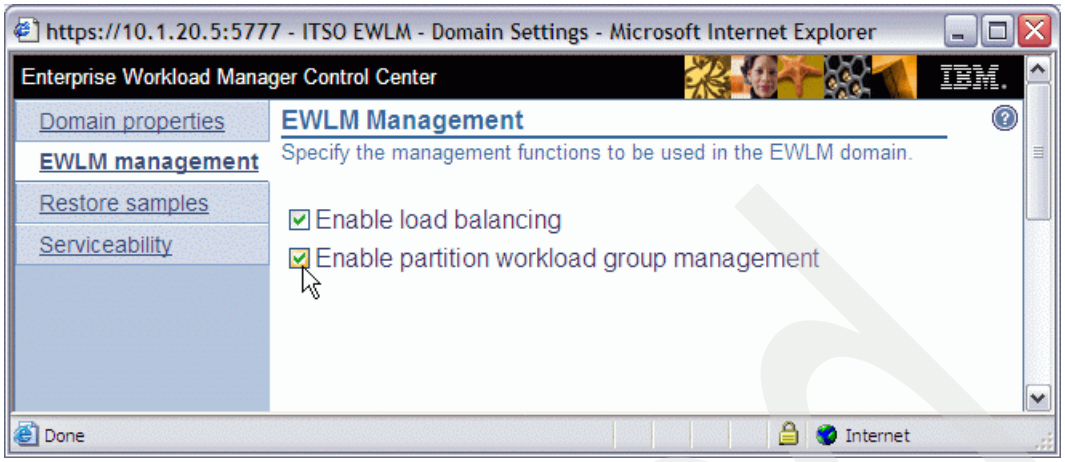


Figure 3-97 Enabling partition group management

Partition management is enabled on the domain. EWLM will make resource adjustments based on the goals defined in the active service policy.

Startup

Two modes are available to start the EWLM domain manager:

- Manual mode
- Automatic mode

Manual mode

The commands to start the EWLM Domain Manager are:

- For AIX: **<*Installation Directory*>/EWLM/bin/startDM.sh <*Config. ID*>**
- For Linux: **<*Installation Directory*>/EWLM/bin/startDM.sh <*Config. ID*>**
- For Microsoft Windows: **<*Installation Directory*>\EWLM\bin\startDM.bat <*Config.ID*>** or using the IBM EWLM Domain Manager Service 1 service.

The commands to start the EWLM Control Center are:

- For AIX: **<*Installation Directory*>/EWLM/bin/startWAS.sh <Config.ID>**
- For Linux: **<*Installation Directory*>/EWLM/bin/startWAS.sh <Config.ID>**
- For Microsoft Windows: **<*Installation Directory*>\EWLM\bin\startWAS.bat <*Config.ID*>** or using the IBMWAS6Service - IBM EWLM Control Center - <EWLM Domain Manager Name> service.

Example 3-26 shows the commands used to start the EWLM domain manager and the control center in an AIX and Linux environment.

Example 3-26 Starting the EWLM Domain Manager and Control Center on AIX / Linux systems

```
/opt/IBM/VE2/EWLM/bin/startDM.sh redbook
/opt/IBM/VE2/EWLM/bin/startWAS.sh redbook
```

Automatic mode

Automating the startup of the service depends on the platform type:

- For AIX, an entry should be added to the inittab, as described in Example 3-27 on page 143.

Example 3-27 Adding an entry in inittab for EWLM Domain Manager and EWLM Control Center on AIX systems

```
mkitab "ewlmdm:2:wait:/opt/IBM/VE2/EWLM/bin/startDM.sh redbook"
mkitab "ewlmcc:2:wait:/opt/IBM/VE2/EWLM/bin/startWAS.sh redbook"
```

- For Linux, the installer does not create a startup script for the two components. You should, therefore, add an entry to the inittab, as described in Example 3-28.

Example 3-28 Entries in inittab for EWLM Domain Manager and EWLM Control Center for Linux systems

```
ewlmdm:35:once:/opt/IBM/VE2/EWLM/bin/startDM.sh redbook > /dev/console 2>&1 # EWLM
DM
ewlmcc:35:once:/opt/IBM/VE2/EWLM/bin/startDM.sh redbook > /dev/console 2>&1 # EWLM
CC
```

- For Microsoft Windows, by default, the two services are configured with the manual startup type. You can change this setup to automatic if necessary.

Post-installation checkup

Here we discuss the post-installation checkup.

Functional checkup

A simple connection to the EWLM Control Center is a fast way to check that both the Control Center and the Domain Manager are up and running.

Technical checkup

Open the TCP 5777 TCP/IP port for connection. This is the connection port to the EWLM Control Center.

3.6.3 IBM Director Server

This section provides information to ensure that the IBM Director Server is well installed.

Post-installation configuration

The IBM Director Server needs the following post-configuration task if you intend to use the security facility.

Single Sign-on (SSO) environment for Linux systems

This applies only when the IBM Director Server is installed on a Linux system.

If the IBM Director Server is supposed to be part of a Virtualization Engine environment, with a link to the Virtualization Engine console, the following operations should be carried out in order to be able to properly configure the SSO later:

- The cURL library must be linked with the /usr/lib/libcurl.so base.

  This is not always the case, especially for SLES9 systems. If /usr/lib/libcurl.so does not exist and if the installed cURL library is /usr/lib/libcurl.so.2, then you should create a symbolic link by running the following command:

  ```
  # ln -s /usr/lib/libcurl.so.2 /usr/lib/libcurl.so.
  ```

  Different Linux distributions may include different versions of the installed cURL library, such as /usr/lib/libcurl.so.3. You still need a link to the base cURL library.

Archived

- You must install the following extra package:

  <*Installation Files*>/director/server/linux/i386/VE_SingleSignOn-1.00-1.i386.rpm

  Install this package with the following commands:

```
# cd /mnt/IbmDirector/director/server/linux/i386
# rpm -i VE_SingleSignOn-1.00-1.i386.rpm
```

Startup

The IBM Director Server can be started either in manual or automatic mode.

Manual mode

The command to start the Director Server is:

- For AIX: **<*Installation Directory*>/bin/twgstart**
- For Linux: **<*Installation Directory*>/bin/twgstart** or **/etc/init.d/twgserver start**
- For Microsoft Windows: **<*Installation Directory*>\bin\twgservw.exe** or using the TWGSERVER service.

The command to start the Director Console is:

- For AIX: **<*Installation Directory*>/bin/twgcon**
- For Linux: **<*Installation Directory*>/bin/twgcon**
- For Microsoft Windows: **<*Installation Directory*>\bin\twgconw.exe**

Example 3-29 shows the command to start the IBM Director Console in an AIX or a Linux environment.

Example 3-29 Starting the Director Console on AIX / Linux systems

```
/opt/ibm/director/bin/twgcon
```

Automatic mode

Automating the startup of the service depends on the platform type:

- For AIX, an entry should be added to the inittab, as shown in Example 3-30.

Example 3-30 Automatic startup of the IBM Director Console in an AIX environment

```
mkitab "dirserv:2:wait:/opt/ibm/director/bin/twgstart"
```

- For Linux, by default, the startup script /etc/init.d/twgserver is active for Runlevels 3 and 5.
- For Microsoft Windows, by default, the TWGSERVER service is configured with an automatic startup type.

Post-installation checkup

Here we discuss the post-installation checkup.

Functional checkup

A fast and easy first functional checkup can be done in the following ways:

1. Connecting to the IBM Director Server through the Director Console.

   This validates that the Director Server is up and running.

Archived

2. Running a discovery.

   By default, the IBM Director discovery feature is configured to use TCP/IP broadcast. Thus, running a discovery should make all devices present on the same subnet as the Director Server appear in the Director Console interface.

   The way to configure and execute a discovery in IBM Director is described in *Implementing IBM Director 5.10*, SG24-6188.

   This validates that the Director Server is able to establish network connections with the Managed Systems.

Technical checkup

The TCP 389 TCP/IP port must be opened for connection. This is the IBM Director Server connection port.

3.6.4 Resource Dependency Service

This section provides information to ensure that Resource Dependency Service is well installed.

Post-installation configuration

No specific post-installation configuration is required for the Resource Dependency Service. However, after its installation, the Virtualization Engine Console must be restarted using the following commands:

- For AIX:

```
<Virtualization Engine Install Dir>/ISC/PortalServer/bin/stopISC.sh ISC_Portal
<userid> <password>
<Virtualization Engine Install Dir>/ISC/PortalServer/bin/stopISC.sh ISC_Portal
```

- For Linux:

```
<Virtualization Engine Install Dir>/ISC/PortalServer/bin/stopISC.sh ISC_Portal
<userid> <password>
<Virtualization Engine Install Dir>/ISC/PortalServer/bin/stopISC.sh ISC_Portal
```

- For Microsoft Windows:

```
<Virtualization Engine Install Dir>\ISC\PortalServer\bin\stopISC.bat ISC_Portal
<userid> <password>
<Virtualization Engine Install Dir>\ISC\PortalServer\bin\stopISC.bat ISC_Portal
```

In these commands, the user ID or password required correspond to the administrator for the Integrated Solutions Console.

Once restarted, the Welcome page of the Virtualization Engine Console should reflect the presence of the Resource Dependency Service in the environment, as shown in Figure 3-98 on page 146.

Archived

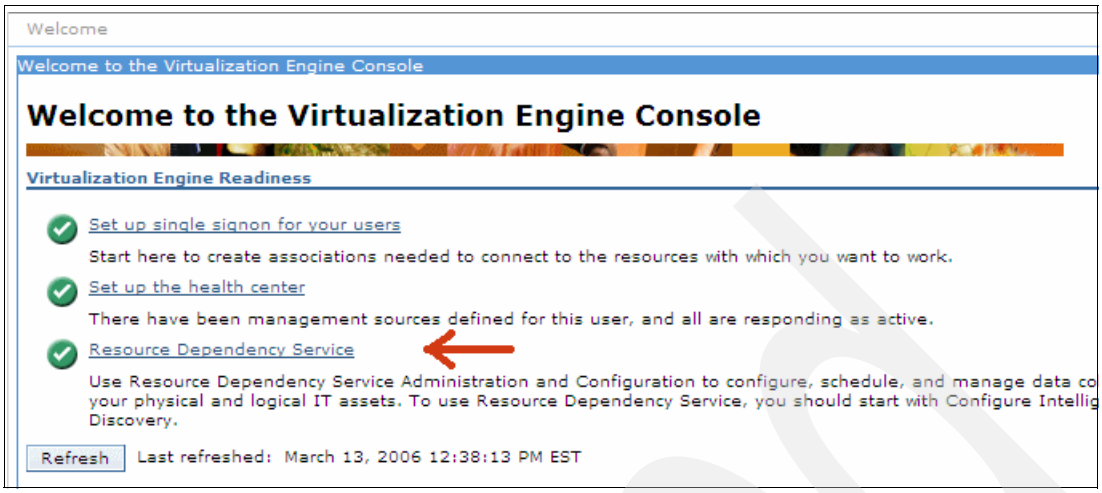


Figure 3-98 Virtualization Engine Console Welcome page with Resource Dependency Service

Shared memory and DB2 Database server on Linux on POWER processor

The following instructions only apply to systems running Linux on a POWER processor.

If there are messages from the DB2 database server or other applications indicating that the shared memory parameters are too small, the Linux kernel parameters of the system hosting the DB2 database server must be configured to match the workload.

This situation can eventually lead to the inability for Resource Dependency Service to connect to its database. The message `Unable to connect to RDS database` (appearing inside the Virtualization Engine Console) may be the consequence of such a memory allocation error.

To configure the Linux kernel parameters, you should complete the three following steps:

1. Determine the current value of the Maximum Shared Memory by running the **`sysctl -a | grep shmmax`** command.

   Based on the current value, choose a new value for the parameter that is larger than the existing value but still appropriate for the system.

   > **Note:** The amount of shared memory that Resource Dependency Service requires depends on the number of database instances running on the Resource Dependency Service server. For example, if IBM Tivoli Directory Server is installed on the Resource Dependency Service server and is using a DB2 database on the server, the shared memory parameter should be at least 2 GB. If the Resource Dependency Service server is only hosting the database required by Resource Dependency Service, 1 GB of shared memory may be sufficient.

2. Reset the Maximum Shared Memory value for the current session by issuing the **`sysctl -w kernel.shmmax=<`*`New Value`*`>`** command, where *<New Value>* is the value that has to be chosen from the previous step. This command changes the shmmax parameter until the next time you restart the server.
3. If this modification seems to fix the problem, the change should be made permanent. To make it permanent:

   a. Edit the /etc/sysctl.conf file. If the file does not exist, it must be created.

b. Add the following line into the file:

```
kernel.shmmax = <New Value>
```

c. Save the file.

This makes the change to the shmmax parameter permanent.

Recommended additional software

Before running an Intelligent Device Discovery® (IDD) campaign, we recommend that you install the following software applications:

- Nmap

  Nmap is an open source tool for network exploration and is also a security scanner. Installing Nmap on the Resource Dependency Service server will greatly enhance the scanning capabilities of IDD. The tool is available for download at:

  `http://www.insecure.org`

  To install Nmap, complete the two following steps:

  a. Check to see what the latest version of Nmap is from the Nmap Web site and download it on the Resource Dependency Service server. Nmap can be downloaded as an RPM or source code. To install the RPM, perform the following steps. For the source code, refer to the instruction on the Web site:

     i. Open a command line as root on the local host.

     ii. Change the path to the directory you have downloaded the RPM of Nmap to and install it with the following command:

     ```
     rpm -vhU <Nmap version>-1.i386.rpm
     ```

     In this command, <*Nmap version*> is the version of the RPM you have downloaded from the Web site.

  b. After the command is completed, check the installation by typing the **`rpm -q nmap`** command.

- Windows Intelligent Device Discovery (IDD) Helper

  Windows Intelligent Device Discovery Helper allows IDD to discover more information from Windows targets, both servers and workstations, that are across the Windows Management Instrumentation (WMI).

  The Windows IDD helper is part of the provided Resource Dependency Service code and is downloaded from the server that hosts Resource Dependency Service. One IDD Helper is needed for each Resource Dependency server.

  To install the Windows IDD helper, complete the three following steps:

  a. In the Virtualization Engine console, select **Virtualization Engine Console** → **Virtualization Engine Setup** → **Resource Dependency Service Configuration.**

  b. In the Resource Dependency Service welcome page, click **Download Windows Intelligent Device Discovery Helper**. Then save the iddWin.exe. file.

  Note: The Windows IDD Helper installation executable file size is around 95 MB.

  c. On the target Microsoft Windows server:

     i. Uninstall any previously existing version of the Windows IDD helper. Complete the uninstallation by deleting the C:\idd.project directory.

ii. Run the executable iddWin.exe code. Follow the instructions of the installation wizard. When prompted, specify the host name of the system running the Resource Dependency Service.

iii. Launch the Windows IDD Helper service. Navigate to the operating system services window. Check that a service called IDD Server is already present. If not, execute the following program:

```
C:\IDD.project\bin\install\installService.bat
```

This registers the service. Once the service is registered, launch it or ensure that it is running.

The installation is complete.

Troubleshooting

At the time of the writing of this IBM Redbook, there is a known issue that may impact the Resource Topology feature of Resource Dependency Service. The issue is that when you go into the Virtualization Engine console and select **Hardware and Software → Resource Topology**, any of the following behavior can be observed in the Resource Topology tab:

- The resources are displayed by default in a list format rather than in the graphical topology view.
- The graphical topology view button is disabled.
- When switching from Virtualization Engine Environment to Managed Resources view using the combo box, the following error message may appear:

```
AVE0408I - The graph view only supports a total of -1 resources, but there are
5 resources in the graph. Because it supports more resources, you have been
switched to the table view of resources.
```

The following actions help to fix the problem and to restore the topology view of the environment:

1. Stop the VEMS server by issuing the following command, as shown in Example 3-31:

   ```
   <Virtualization Engine Install Dir>/VEMS/bin/stopServer VEMS -username <User>
   -password <Pwd>
   ```

Example 3-31 Stopping the VEMS server

```
/opt/IBM/VE2/VEMS/bin/stopServer.sh VEMS -username veadmin -password veadmin
```

2. On the Virtualization Engine Console system, go to the <*Virtualization Engine Install Dir*>/VEConsole/ESM/data/persistence directory.
3. In this directory, delete all files that have the following *suffix* as part of the name, as shown in Example 3-32:

   ```
   topology_preference.xml
   ```

Example 3-32 Removing existing "Topology Preference" files

```
rm -f *_topology_preference.xml
```

4. Restart the VEMS server with the following command:

   ```
   <Virtualization Engine Install Dir>/VEMS/bin/startServer.sh VEMS
   ```

Note: This issue is under investigation at the time of the writing of this IBM Redbook and a fix should be included in a future Fix Pack for the Resource Dependency Service component.

Archived

Startup

Resource Dependency Service can be started in manual or automatic mode.

Manual mode

The command to start the Resource Dependency Service, as shown in Example 3-33, is as follows:

`<Installation Directory>/VEMS/bin/startServer.sh VEMS or /etc/init.d/ibmstartvems`

Example 3-33 Starting the Resource Dependency Service on Linux systems

```
/opt/IBM/VE2/VEMS/bin/startServer.sh VEMS
```

Note: The service requires the Directory Server to be up and running. If not, the service will not start.

Automatic mode

By default, the startup script created by the installer is set as active at boot time for runlevels 2, 3, and 5.

However, if the bridge runs on the same system as the Directory server, this service fails to start at boot time. The reason is that this service requires the Directory server to be up and running in order to start properly, as it needs to retrieve the user ID and password from the repository defined inside the Shared WebSphere Application profile. As the Directory Server is configured to start automatically through an entry in the inittab, its startup command is always executed *after* the startup script of the bridges service.

The suggested workaround for this issue is to perform the following:

- Disable the automatic mode for the startup script using the **chkconfig** command, as shown in Example 3-34.

Example 3-34 Disabling automatic mode

```
chkconfig ibmstartvems off
```

- Add an entry at *the very end* of the inittab, making sure that this entry comes *after* the entry related to the Directory Server, as shown in Example 3-35.

Example 3-35 Adding an entry into the inittab file

```
vems:35:once:/opt/IBM/VE2/VEMS/bin/startServer.sh VEMS > /dev/console 2>&1
```

3.6.5 Virtualization Engine console

This section provides information to ensure that the Virtualization Engine console is well installed.

Post-installation configuration

Some specific configuration may be needed for two aspects: the SSO environment and the management resources environment.

- Single Sign-on

  The Virtualization Engine Console features an SSO mechanism. It allows the Virtualization Engine console users to connect to other Virtualization Engine management servers without the need to authenticate again. In particular, this SSO mechanism applies

to the communication between the Virtualization Engine console and the following services:

- EWLM Domain Manager (management source)
- IBM Director Server (management source)
- Resource Dependency Service

Note: The SSO configuration of the EWLM Domain Manager is automatically set up by the installer. No further configuration task is required.

You must, therefore, perform the SSO configuration for both the IBM Director and Resource Dependency Service (if present in the environment). It is achieved through the Virtualization Engine console.

- Management sources

  The management sources represent the providers of information on which the Virtualization Engine console relies. The Management Sources may be:

  - Components of the IBM Virtualization Engine itself:
    - EWLM Domain Manager
    - IBM Director Server
  - External providers:
    - Cluster System Management (CSM)
    - Management Central

The following sections describe these customization issues.

Single Sign-on configuration

Perform the eight following steps:

1. Connect to the Virtualization Engine console.

   The connection URL is `http://<hostname>:8421/ibm/console`.

   The browser should display the Integrated Solutions Console login page, as shown in Figure 3-99, where the expected user is the Virtualization Engine administrator (the default value is veadmin).

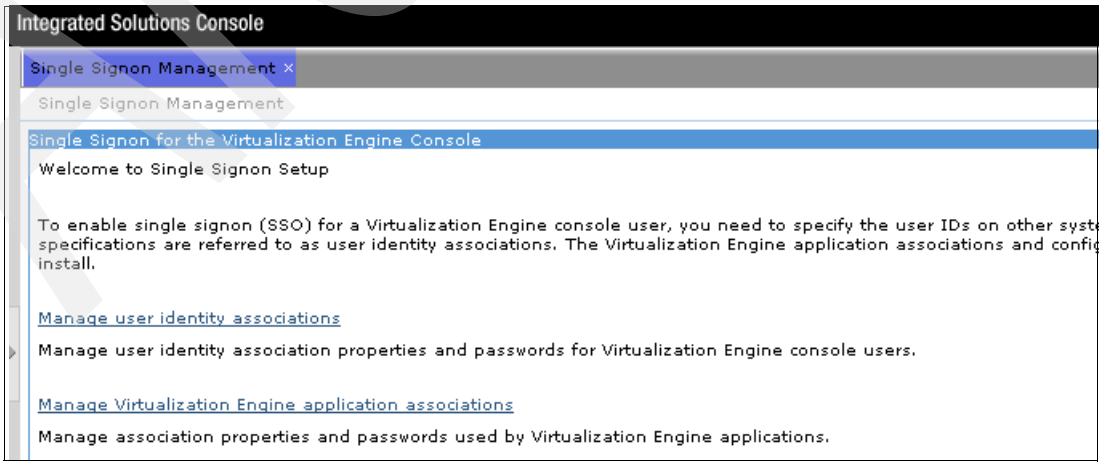


Figure 3-99 Log on to Virtualization Engine Console

2. Set up the SSO management environment as follows:
   a. Select **Virtualization Engine Console** → **Virtualization Engine Setup** → **Single Signon Management**.
   b. Click **Manage user identity associations**, as shown in Figure 3-100.

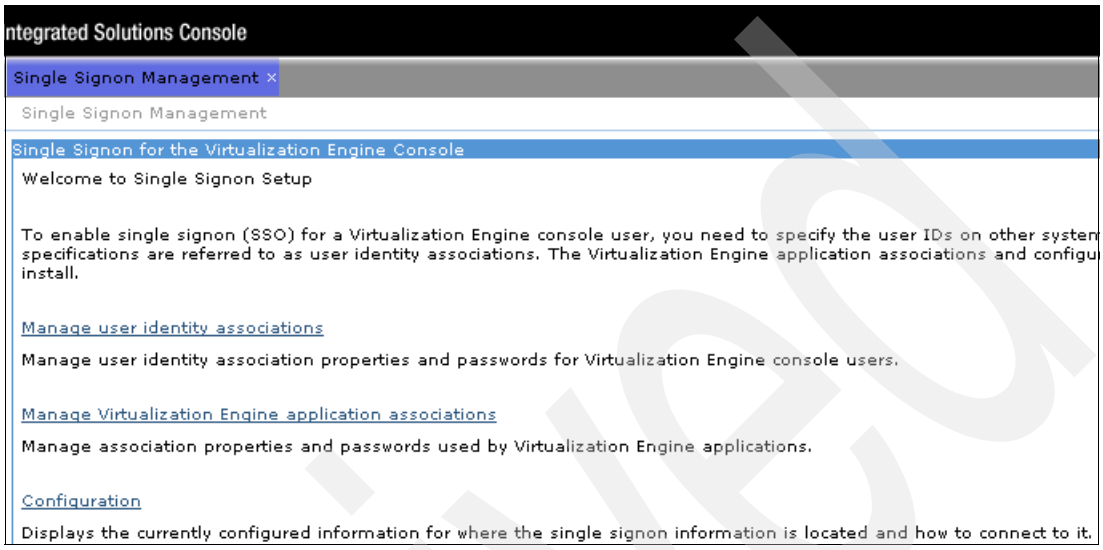


Figure 3-100 Single Signon Management

3. Manage User Identity Associations by selecting **New...** and clicking **Go**, as shown in Figure 3-101.

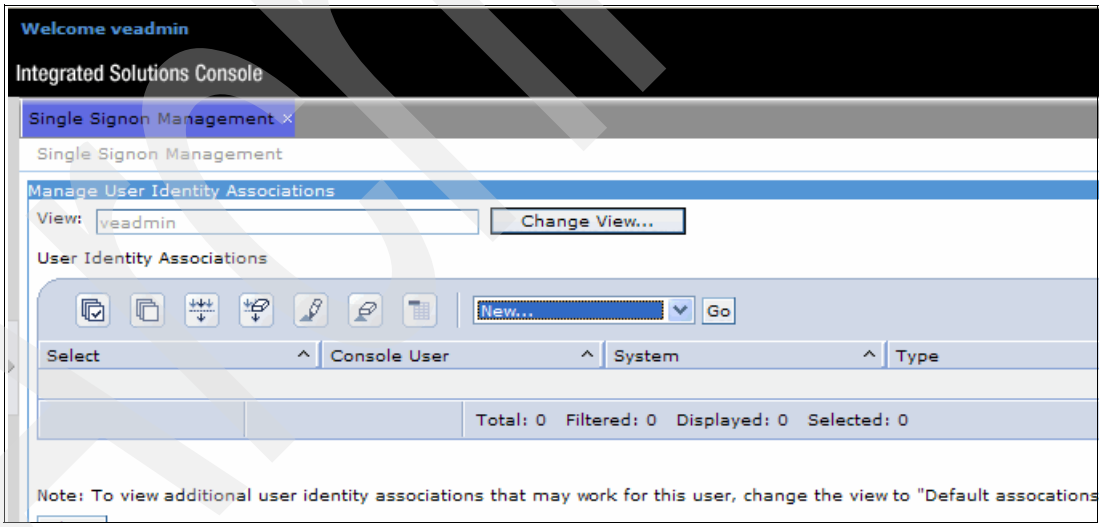


Figure 3-101 User Identity Associations

4. Add the New User Identity Association by clicking **Select**, as shown in Figure 3-102 on page 152.

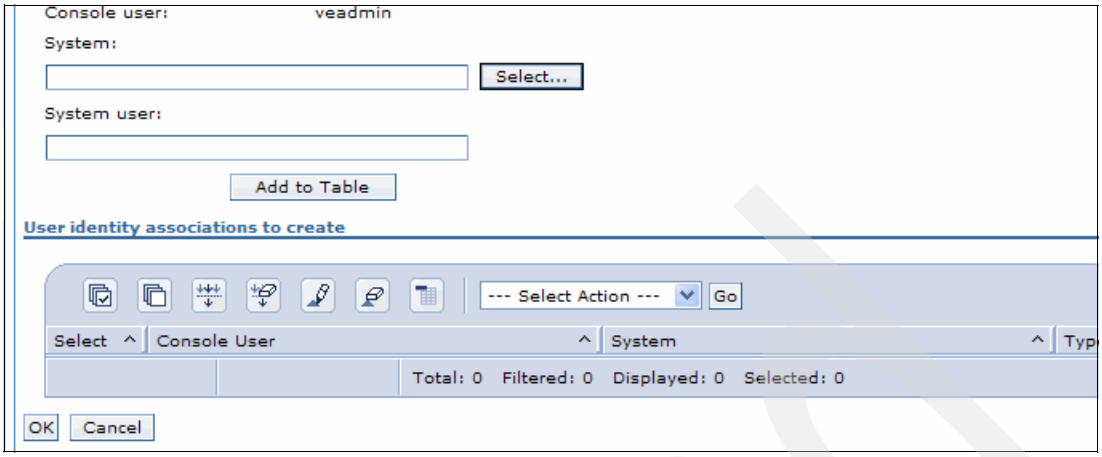


Figure 3-102 New user identity association

5. Select **System** and click **OK.** As shown in Figure 3-103, the system lists <hostname>:5800, which refers to the EWLM Bridge to Virtualization Engine console.

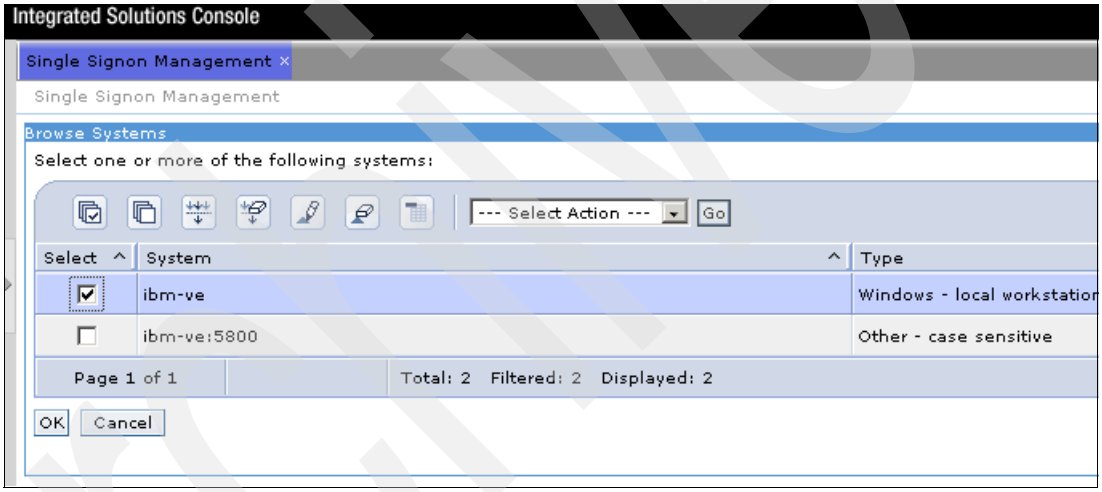


Figure 3-103 SSO Management Browse Systems

6. You should now associate the user identity to the system, as shown in Figure 3-104 on page 153. Click the **Add to Table** button.

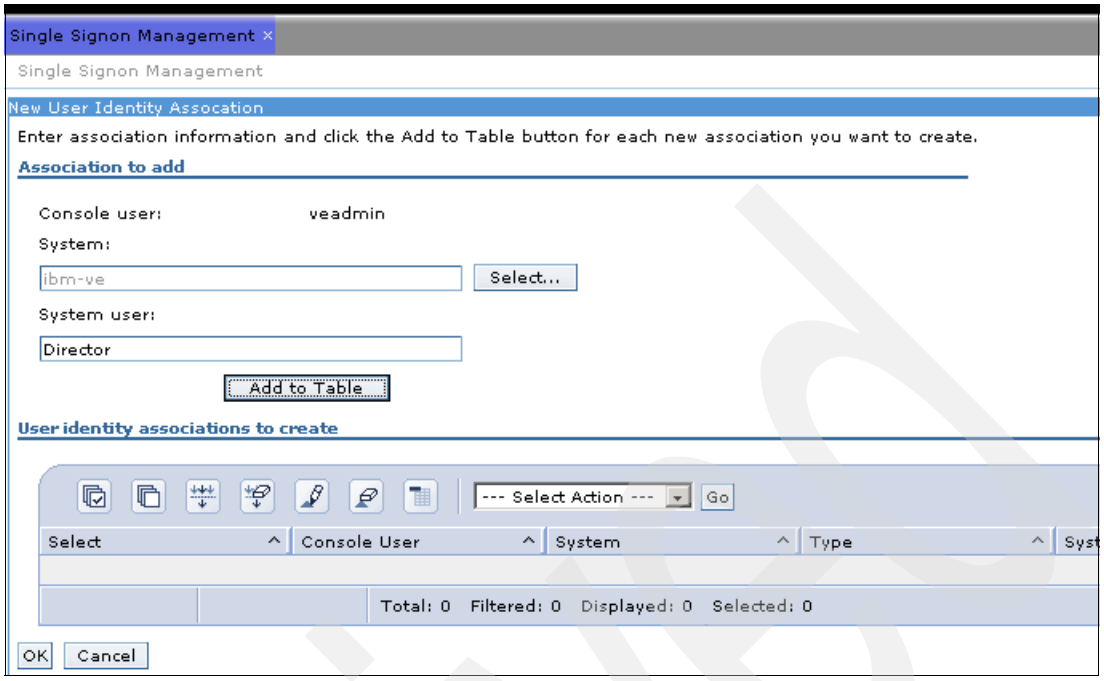


Figure 3-104 SSO Management - New User Identity Association

7. Add the selected system to the User Identity Association to create a New User Identity Association, as shown in Figure 3-105.

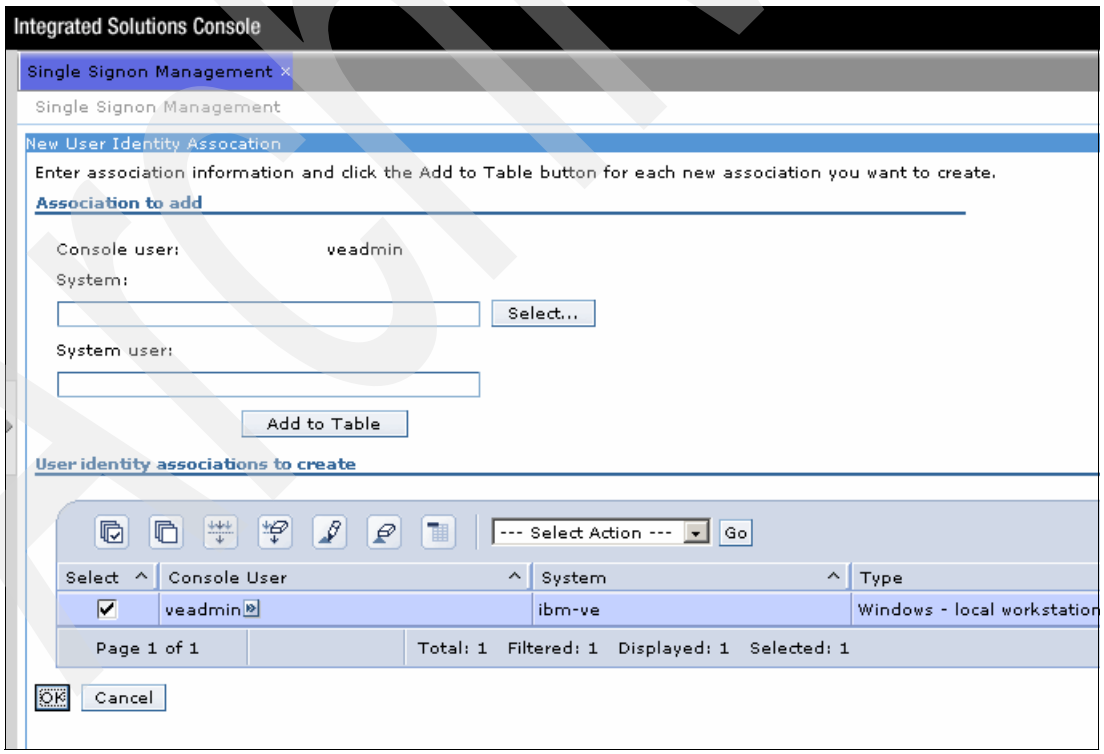


Figure 3-105 Create User Identity Association

8. Set the password by displaying the contextual menu and selecting **Set Password**, as shown in Figure 3-106 on page 154.

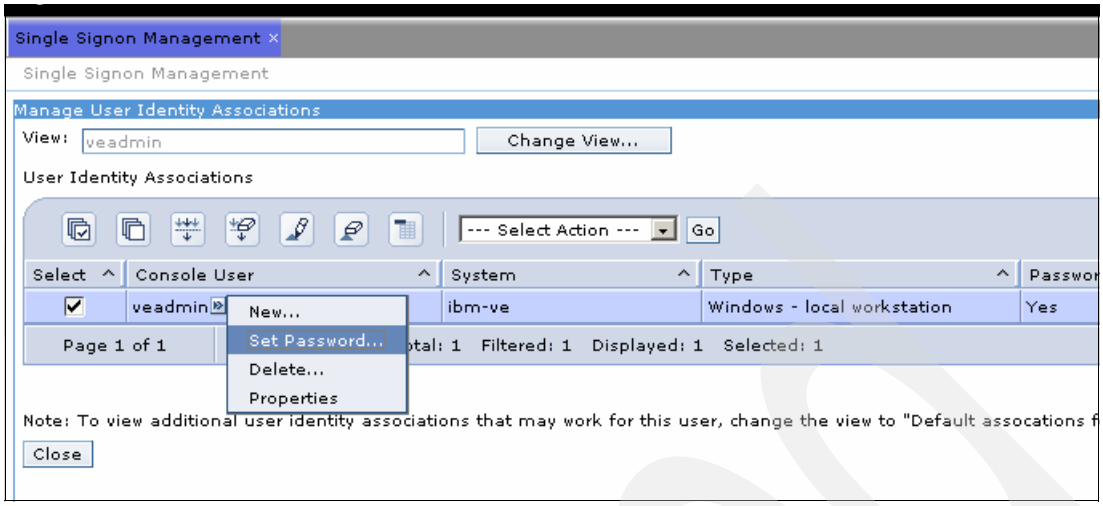


Figure 3-106 Set password option

Enter the password for the system user and click **OK**, as shown in Figure 3-107.

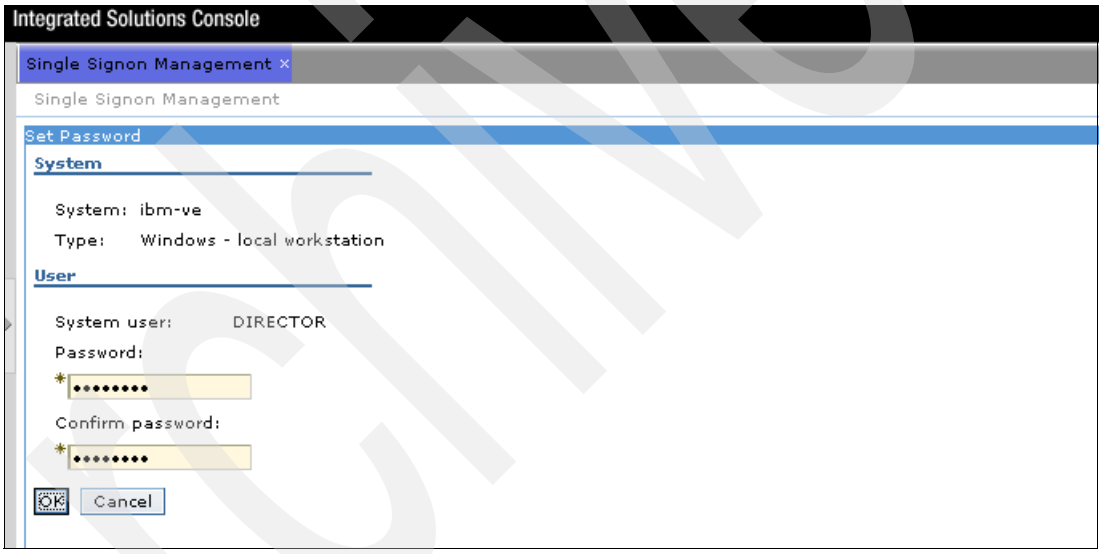


Figure 3-107 Set password

The User Identity Association is created, as shown in Figure 3-108 on page 155.

Archive

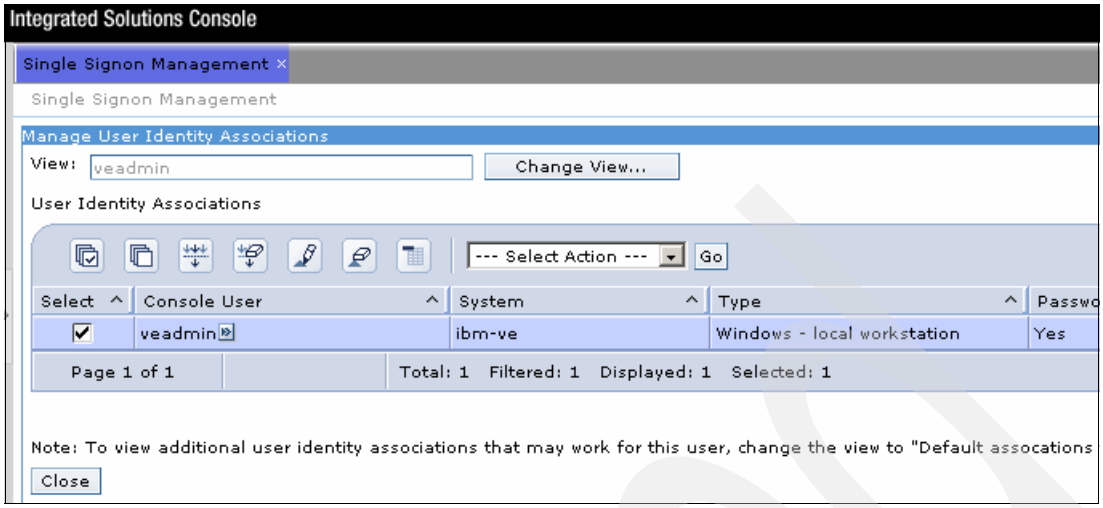


Figure 3-108 User identity associations

For Resource Dependency Service, an extra step is required. Perform the configuration as follows:

1. Single Sign-on Management

   Select **Virtualization Engine Console** → **Virtualization Engine Setup** → **Single Signon Management**. The window in Figure 3-107 on page 154 should appear. Click **Manage user identity associations**.

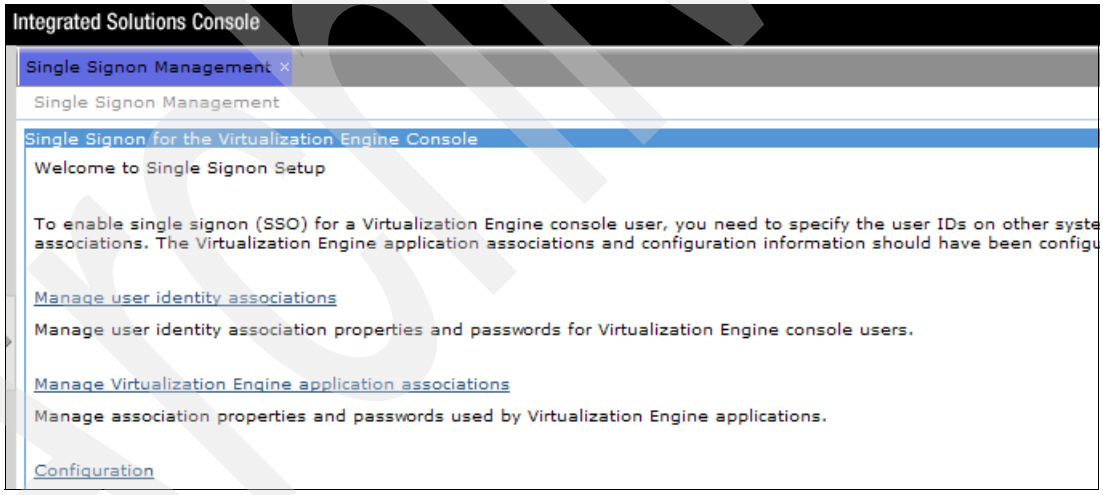


Figure 3-109 Single Signon Management

2. SSO management: manage Virtualization Engine application associations

   Display the contextual menu for the Resource Dependency Service record and select **Set Password**, as shown in Figure 3-110 on page 156.

Archive

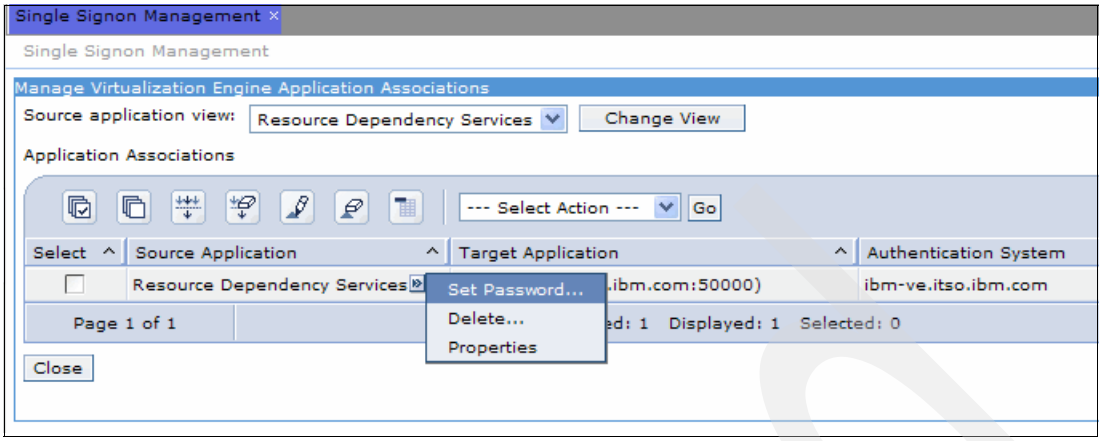


Figure 3-110 SSO Management - Manage Virtualization Engine application associations

3. SSO Management: set password

   Input the password for the DB2 instance owner and click **OK**, as shown in Figure 3-111.

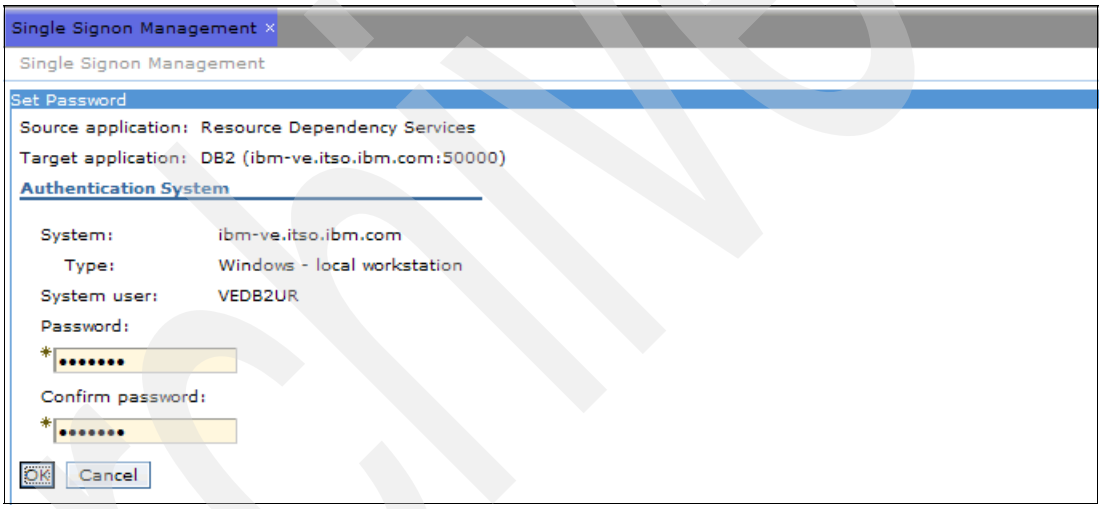


Figure 3-111 SSO Management - Set Password

4. SSO Management: manage Virtualization Engine Applications Association

   The Application Association is created, as shown in Figure 3-112.

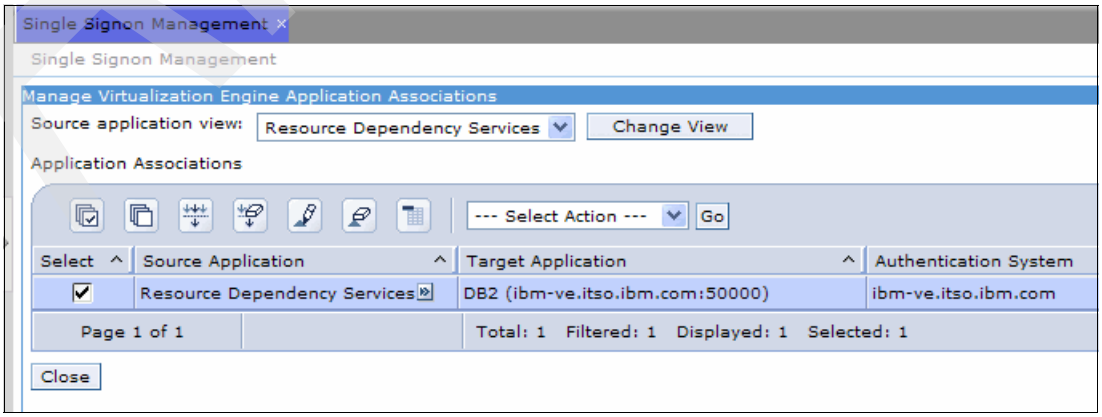


Figure 3-112 SSO Management - Manage Virtualization Engine Applications Association

Adding management sources

The management sources represent the providers of information on which the Virtualization Engine console relies. The management sources may be:

- Components of the IBM Virtualization Engine itself:
  - EWLM Domain Manager
  - Director server
- External providers:
  - Cluster System Management (CSM)
  - Management Central

Perform the following steps to define the way to add management sources to the Virtualization Engine console:

1. In the Health Center wizard Welcome window, select **Virtualization Engine Console** → **Virtualization Engine Setup** → **Health Center Setup**.

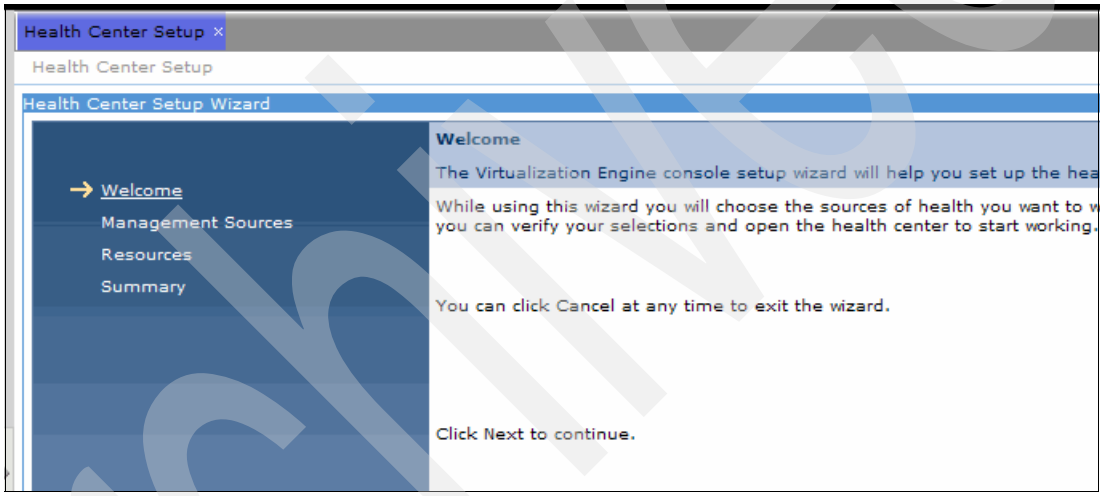


Figure 3-113 Virtualization Engine Console - Health Center Setup

2. In the management sources selection window, you can select management sources from the list of management sources that have been the automatically detected, as shown in Figure 3-114 on page 158. Another option is that you can rename these Management Sources using the **Rename** button.

Archived

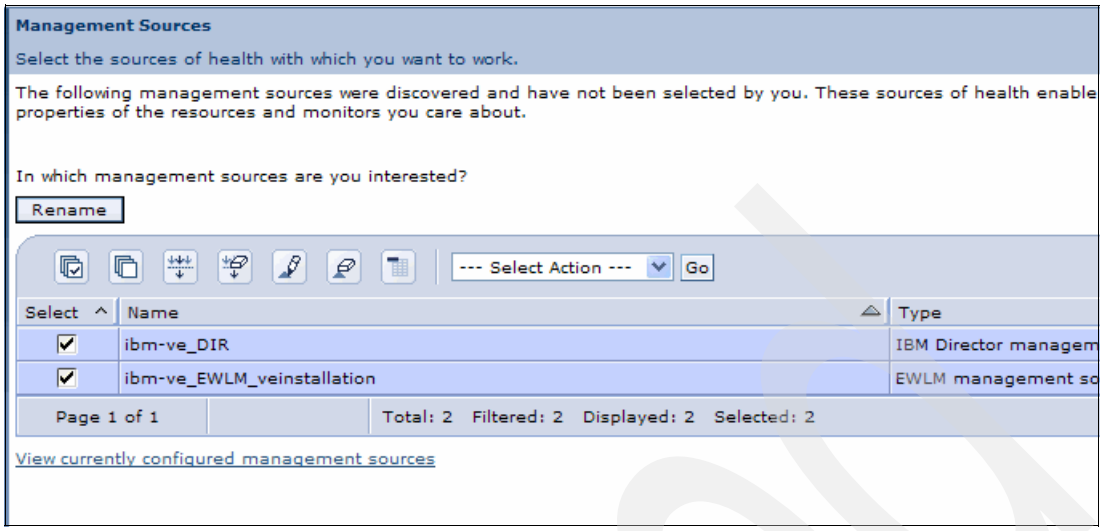


Figure 3-114 Virtualization Engine Console - detecting of new resources

3. In the Health Center wizard, in the Resources window, accept the default selection of resources to be retrieved.

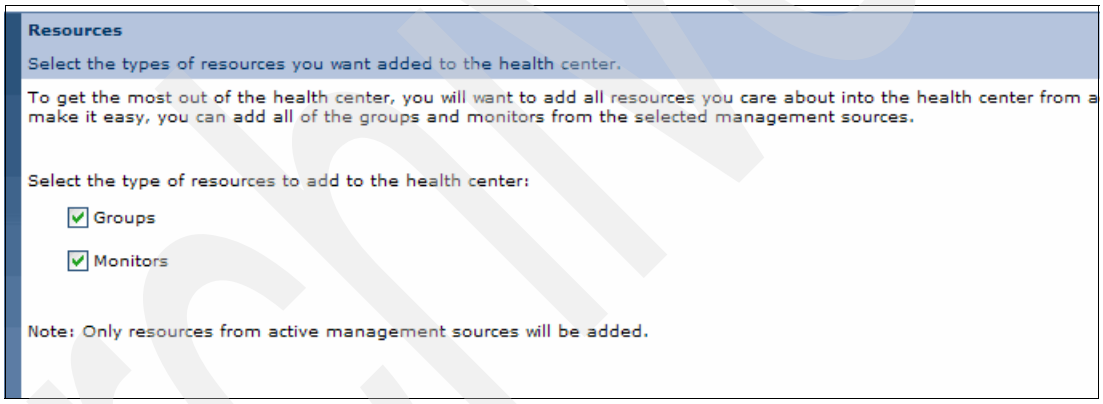


Figure 3-115 Virtualization Engine Console - retrieve new resources

4. In the Summary window of the Health Center wizard, click **Finish** and exit the wizard.

Archive

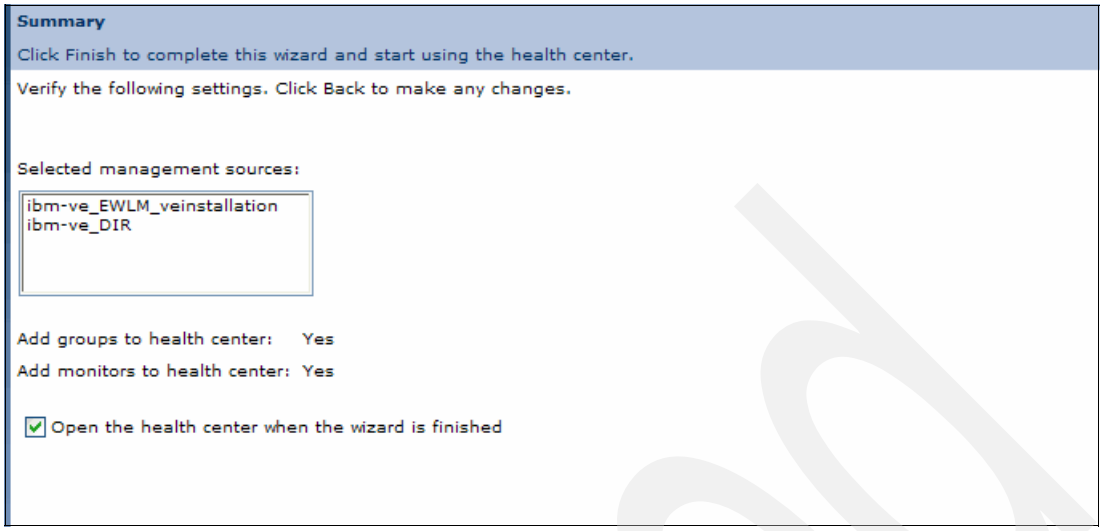


Figure 3-116 Virtualization Engine Console - exit from Health Center

5. To check the status, from the menu on the left of the window, select **Virtualization Engine Console → Virtualization Engine Setup → Management Sources**. The management sources have just been added to the Virtualization Engine console and should be reported as an active status, as shown in Figure 3-117.

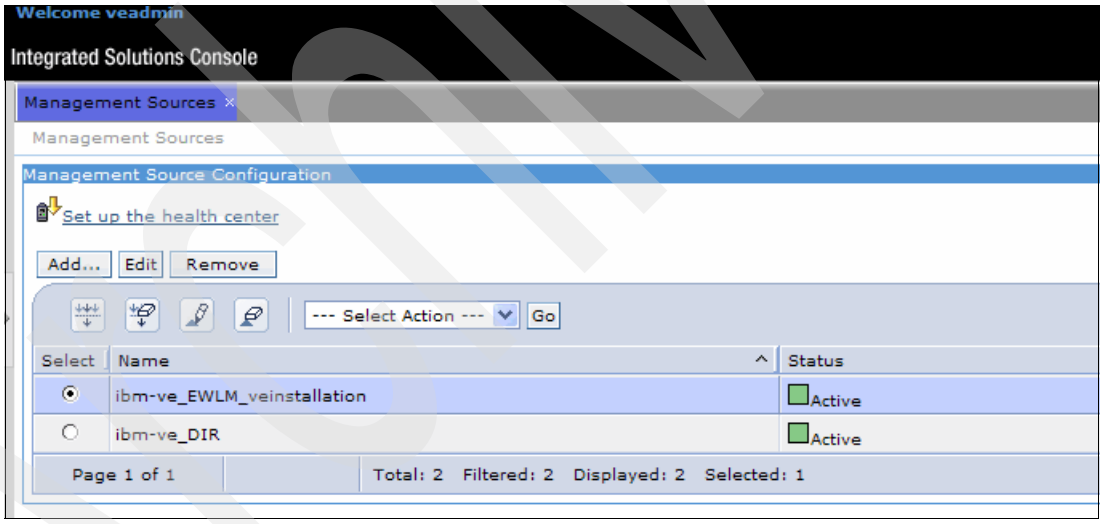


Figure 3-117 Virtualization Engine console: management sources

Post-installation checkup

Here we discuss the post-installation checkup.

Functional checkup

You can perform a fast and easy first functional checkup by doing the following:

- The definition of a simple resource monitor inside IBM Director

  This resource monitor can be, for example, the CPU usage of the Director Server itself.
- The configuration of the related Monitor inside the Virtualization Engine Console
- The visualization of the monitor values by using a graph in the Virtualization Engine Console

This validates that:

- The Virtualization Engine console is up and running.
- The Director Server is up and running.
- The Director Bridge to Virtualization Engine console is up and running.
- The communication between the Virtualization Engine console and the Director Bridge is active.

Startup

The Virtualization Engine console can be started manually or automatically.

Manual mode

To start the Virtualization Engine console, use the following commands:

- For AIX: **<*Installation Directory*>/ISC/PortalServer/bin/startISC.sh ISC_Portal**
- For Linux: **<*Installation Directory*>/ISC/PortalServer/bin/startISC.sh ISC_Portal**
- For Microsoft Windows:**<*Installation Directory*>\ISC\PortalServer\bin\startISC.bat ISC_Portal** or using the IBMWAS6Service - WebSphere for Virtualization Engine Console service.

Example 3-36 shows the command to start the Virtualization Engine console in an AIX and a Linux environment.

Example 3-36 Starting the Virtualization Engine Console on AIX / Linux systems

```
/opt/IBM/VE2/ISC/PortalServer/bin/startISC.sh ISC_Portal
```

Note: In order for the startup of the Virtualization Engine Console to complete successfully, the Directory Server service must already be running. If not, the startup fails.

Automatic mode

The automatic startup of the Virtualization Engine console depends on the operating system type:

- For AIX, an entry should be added to the inittab. If the Virtualization Engine console runs on the same system as the Directory server, this entry must come *after* the entry related to the Directory Server in order to respect the dependency mentioned earlier. Example 3-37 shows the entry to add to the inittab file.

Example 3-37 Adding an inittab entry for Virtualization Engine console on AIX systems

```
mkitab "vec:2:wait:/opt/IBM/VE2/ISC/PortalServer/bin/startISC.sh ISC_Portal"
```

- For Linux, the installer does not create a startup script for this service. You should, therefore, add an entry to the inittab. If the Virtualization Engine console runs on the same system as the Directory server, this entry must come *after* the entry related to the Directory server in order to respect the dependency mentioned above. Example 3-38 shows the entry to add in the inittab file.

Example 3-38 Inittab entry for the Virtualization Engine console on Linux systems

```
vec:35:once:/opt/IBM/VE2/ISC/PortalServer/bin/startISC.sh ISC_Portal >
/dev/console 2>&1
```

- For Microsoft Windows, by default, the Virtualization Engine console service is configured with an automatic startup type. However, if the Virtualization Engine console runs on the

same system as the Directory server, it fails to start at boot time. The reason is that by default, the startup type of the Directory server service is set to Manual. Therefore, when the Virtualization Engine Console service tries initializing, it cannot access the Directory server as required. We could not find any valid workaround for this issue.

Technical checkup

Open the TCP 8421 TCP/IP ports for connection. This is the connection port to the Integrated Solutions Console.

3.6.6 Bridges to Virtualization Engine console

This section provides information to ensure that the bridges to the Virtualization Engine console are well installed.

Post-installation configuration

The bridges to the Virtualization Engine console do not require any post-installation configuration at the level of the bridges themselves. The configuration is actually done directly through the Virtualization Engine console.

Startup

The bridges to the Virtualization Engine console may be started manually or automatically.

Manual mode

To start the bridges to the Virtualization Engine console, issue the following command:

- For AIX: **<*Installation Directory*>/VEMS/bin/startServer.sh VEMS**
- For Linux: **<*Installation Directory*>/VEMS/bin/startServer.sh VEMS** or use a startup script: **/etc/init.d/ibmstartvems**
- For Microsoft Windows: **<*Installation Directory*>\VEMS\bin\startServer.bat VEMS** or using the IBMWAS6Service - IBMVEManagementServices service

Example 3-39 shows an example of the command for an AIX or for a Linux environment.

Example 3-39 Starting the Bridge on AIX / Linux systems

```
/opt/IBM/VE2/VEMS/bin/startServer.sh VEMS
```

Note: In order for the startup of the bridge to complete successfully, the Directory Server service must already be running. If not, the startup fails.

Automatic mode

To automate the startup of the service depends on the platform type:

- For AIX, an entry should be added to the inittab. If the Bridge runs on the same system as the Directory server, this entry must come *after* the entry related to the Directory Server in order to respect the dependency mentioned earlier. Example 3-40 shows the inittab entry.

Example 3-40 Adding an inittab entry for Bridge service on AIX systems

```
mkitab "vems:2:wait:/opt/IBM/VE2/VEMS/bin/startServer.sh VEMS"
```

- For Linux, if the autostart flag was selected during the installation, a startup script is created by the installer and is turned on by default for runlevels 3 and 5.

  However, if the bridge runs on the same system as the Directory Server, this service fails to start at boot time. The reason is that this service requires the Directory server to be up and running in order to start properly, as it needs to retrieve the user ID and password defined inside the Shared WebSphere Application profile from the repository. As the Directory Server is configured to start automatically through an entry in the inittab, its startup command is always executed *after* the startup script of the bridges service.

  We recommend the following fix for this issue:

  – Disable the automatic mode for the startup script with the following command:

    ```
    chkconfig ibmstartvems off
    ```

  – Add an entry at *the very end* of the inittab, making sure that this entry comes after the entry related to the Directory server as follows:

    ```
    vems:35:once:/opt/IBM/VE2/VEMS/bin/startServer.sh VEMS > /dev/console 2>&1
    ```

  Note: This issue will be solved in IBM Virtualization Engine V3, with *all* components relying on startup scripts, instead of inittab for some of them.

- For Microsoft Windows, by default, the bridge service is configured with an automatic startup type.

  However, if the Bridge runs on the same system as the Directory server, it fails to start at boot time. The reason is that by default, the startup type of the Directory Server service is set to Manual. Therefore, when the Bridge service tries initializing, it cannot access the Directory Server as required.

  We could not find any valid workaround for this issue.

Post-installation checkup

Here we discuss the post-installation checkup.

Technical checkup

Open the TCP 18551 TCP/IP port for connection. It is the connection port to the bridge.

Archived

4

Installing managed servers agents

After the installation of the management servers, you have to define which servers you want to manage. This next to the last task of the installation process, as depicted in Figure 4-1, is to deploy the agents on the managed servers. This is what is discussed in this chapter.

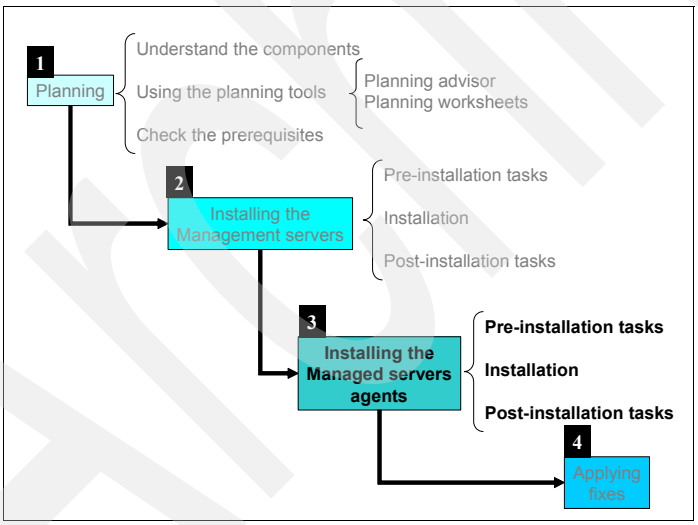


Figure 4-1 Task three: installing the agents on the managed servers

This task, similar to installing the management servers, can be subdivided into three sub-tasks:

1. Preinstallation tasks, described in 4.1, “Preinstallation” on page 164. For example, looking at the Application Response Measurement (ARM) enablement.
2. The installation process itself, where different methods can be used, as described in 4.2, “Installation process” on page 165.
3. Post-installation tasks, described in 4.3, “Post-installation tasks” on page 182. For example, the configuration of the managed servers.

© Copyright IBM Corp. 2006. All rights reserved.

4.1 Preinstallation

This section provides a summary of the elements to check on the target systems before starting the installation of the agents.

4.1.1 Enterprise Workload Manager agent

The following list summarizes the items you should look at before starting the installation of the Enterprise Workload Manager (EWLM) agent:

- Make sure that the hardware and software prerequisites are met (for the latest information, visit the Information Center).
- Check that the host name resolution works fine, testing, in particular, the connectivity between the EWLM Managed Server system and the EWLM Domain Manager system (or the firewall broker, depending on the environment topology).

  If you are using local hosts file for name resolution, ensure that both the EWLM Managed Server and the EWLM Domain Manager or firewall broker have corresponding entries in the file.
- For AIX operating systems, activate the ARM at the operating system level. To do this:

  a. Check if ARM is enabled as follows:

     i. Run the **`lsarm -a`** command.
     ii. Run the **`smitty ewlm`** command.
     iii. Select Change/Show status of the EWLM services.
     iv. Select Show the status of the EWLM services.

  b. If ARM is reported as *not* enabled, then perform the following steps:

     i. Move back to the previous screen by pressing F3.
     ii. Select Enable the EWLM services.
     iii. Display the parameter value list by pressing F4.
     iv. Choose Both.
     v. Validate the selection by pressing Enter.
     vi. Apply the change to the system by pressing Enter.
     vii. Exit SMITTY.

4.1.2 IBM Director Agent

The following list summarizes the items to look at before starting the installation of the IBM Director Agent:

- Make sure the hardware and software prerequisites are met (visit the Information Center for the latest information).
- Check that the host name resolution works fine, testing, in particular, the connectivity between the Director Agent system and the Director Server system.

  If you use the local hosts file for name resolution, ensure that both the Director Server and the Director Agent have corresponding entries in the file.

Archived

4.2 Installation process

There are three different methods to install the agents on the managed endpoints:

1. Using the Software Distribution feature of IBM Director.
2. Using the Remote Package Deployer (RPD) tool that comes with the IBM Virtualization Engine.
3. Manually and directly on the target systems.

The following sections describe each of these methods.

Note: The EWLM Agent and the Director Agent installation files are located in the following places:

- EWLM Agent: <Virtualization Engine Installation Files Directory>/FILES/EWLM/

  The directory contains one zip archive per platform.
- Director Agent: <IBM Director Installation Files Directory>/director/agent/<platform>/

4.2.1 Using the IBM Director software distribution

The software distribution feature that comes with IBM Director allows the automated deployment of software packages onto target systems. It is a very convenient way of dealing with the managed server installation task. (Note that the XML files are provided and can be imported in IBM Director.)

Detailed usage of the IBM Director software distribution feature is documented in *Implementing IBM Director 5.10*, SG24-6188.

4.2.2 Using the Remote Package Deployer (RPD)

The Remote Package Deployer allows distribution of the agent code onto the managed system from the management server through the command line. To use the RPD, perform the following steps:

1. Place the agent code on the management server.

   Note: The Virtualization Engine installer automatically takes care of this operation for the EWLM agents that are selected for installation in the Service Selection step of the Virtualization Engine installer. In this case, the agent code is placed in the <*Virtualization Engine Installation Directory*>/ManagedNodes/EWLM/<Platform> directory.

   When manually created, the directory structure hosting the agent code should comply with the Virtualization Engine installer standard, which implies:

   – Creating a subfolder by component, for example, for Director Agent:

     <Virtualization Engine Installation Directory>/ManagedNodes/Director

   – Creating a subfolder by platform inside the component folder, for example:

     - <*Virtualization Engine Installation Directory*>/ManagedNodes/Director/Aix
     - <*Virtualization Engine Installation Directory*>/ManagedNodes/Director/Win

   The agent code must be copied directly from the installation files under the appropriate folder, depending on the product and the platform. The agents are packaged with the

Archived

Solutions Install format. They are made of the following components, as shown in Example 4-1:

- A FILES directory containing the Installable Unit (IU):
  - Installation executable file
  - Response file
- A META-INF directory containing the Installable Unit Deployment Descriptor (IUDD) eXtensible Markup Language (XML) file and the Artifact XML file.

Example 4-1 Installation files for EWLM Agent on Microsoft Windows systems

```
<...>\FILES\EWLMMSWin.exe
<...>\FILES\EWLMMSWin.rsp
<...>\META-INF\EWLMMSWinRPD_packagedIU.xml
<...>\META-INF\installArtifact.xml
```

2. Edit the response files. You can specify the desired installation settings by editing the response files corresponding to the agents to be deployed. Refer to “Silent installation” on page 181 for each of the two agents for information related to customizing the response files.
3. Launch the RPD script. The RPD script is located in the following directory (it depends on the operating system):
   - For AIX and Linux, it is located in <*Virtualization Engine Installation Directory*>/rpd/RPDx.sh.
   - For Microsoft Windows, it is located in <*Virtualization Engine Installation Directory*>\rpd\RPDWin.bat.

   You must specify the following parameters:

   - -system <*system*> (Host name/IP address of the target system)
   - -userid <*username*> (User name on the target system)
   - -password <*pwd*> (Password of the user on the target system)
   - -p <*filename*> (Path of the Package Description file)

   Example 4-2 and Example 4-3 show an RPD command in different operating system environments.

Example 4-2 Launching RPD command for EWLM Agent on AIX / Linux

```
/opt/IBM/VE2/rpd/RPDx.sh -system ewlm-ms -userid root -password ibm
                -p
                /opt/IBM/VE2/ManagedNodes/EWLM/Win/META-INF/EWLMMS_packagedIU.xm
                l
```

Example 4-3 Launching RPD command for EWLM Agent on Microsoft Windows

```
C:\Program Files\IBM\VE2\rpd\RPDWin.bat -system ewlm-ms -userid Administrator
-password ibm -p "C:\Program
Files\IBM\VE2\ManagedNodes\EWLM\Win\META-INF\EWLMMS_packagedIU.xml"
```

The installation process is silent and there will be no indication of progress, either on the target system or on the management server. The completion of the script command corresponds to the completion of the installation on the target system.

4.2.3 Manual installation

The manual installation of the agents remains an option for a reduced environment where the number of managed systems is limited.

It requires getting the adequate agent code copied onto the managed system before the installation by any appropriate means, such as File Transfer Protocol (FTP) transfer, copy from installation physical media, and so on.

EWLM agent

The EWLM agent comes packaged as zip archives that need to be copied and unzipped on the managed system:

- For AIX: EWLMMSAixRPD.zip
- For Linux on POWER:EWLMMSLinPPCRPD.zip
- For Linux on IBM System x: EWLMMSLinXRPD.zip
- For Microsoft Windows: EWLMMSWinRPD.zip

Inside the zip archives, the agent code is packaged according to the Solution Install format.

You can install the EWLM Agent in one of the following ways:

- Graphical installation wizard

  Note: This requires the presence of a graphical environment on AIX and Linux systems.

- Console mode

  Console mode is a text-based, user-interactive installation mode. It does not require any graphical environment to be installed.

- Silent mode

  For the EWLM Agent, the silent mode is available in two variants:

  - Using a response file
  - Using command-line parameters

Graphical wizard

The graphical installation wizard is launched through the following executable programs:

- For AIX: <*Source Directory*>/Aix/FILES/EWLMMSAix.bin
- For Linux on POWER: <*Source Directory*>/LinPPC/FILES/EWLMMSLinPPC.bin
- For Linux on IBM System x: <*Source Directory*>/LinX/FILES/EWLMMSLinX.bin
- For Microsoft Windows: <*Source Directory*>\FILES\EWLMMSWin.exe

Archived

Perform the following steps using the installation wizard:

1. Start the InstallShield wizard. The welcome window, as shown in Figure 4-2, is displayed.



Figure 4-2 Welcome to the InstallShield Wizard

2. You should specify the installation directory, as shown in Figure 4-3. The default value is acceptable.

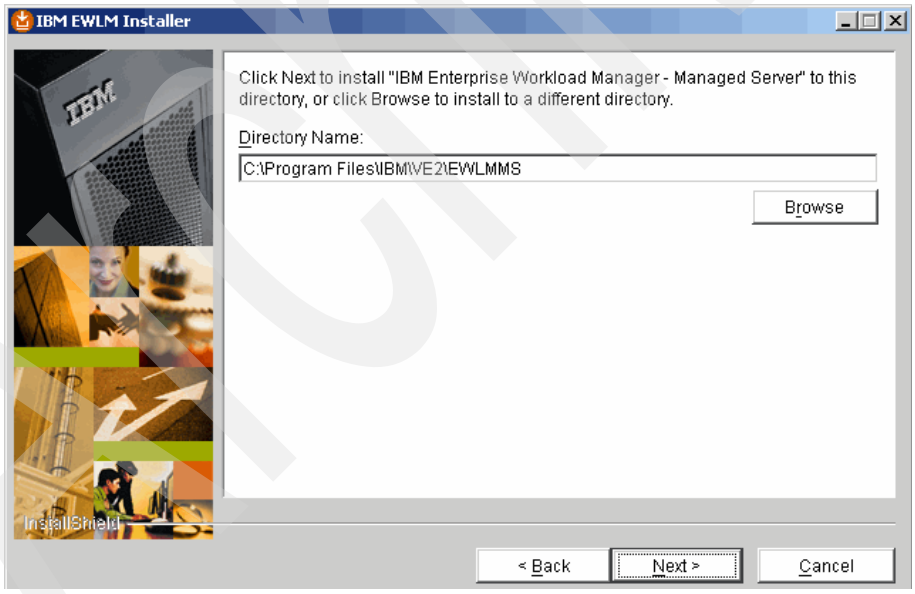


Figure 4-3 Installation Directory

3. You should provide a data directory, as displayed in Figure 4-4 on page 169. A data directory is used to store the data files related to the EWLM Managed Server configuration. The default value is acceptable.

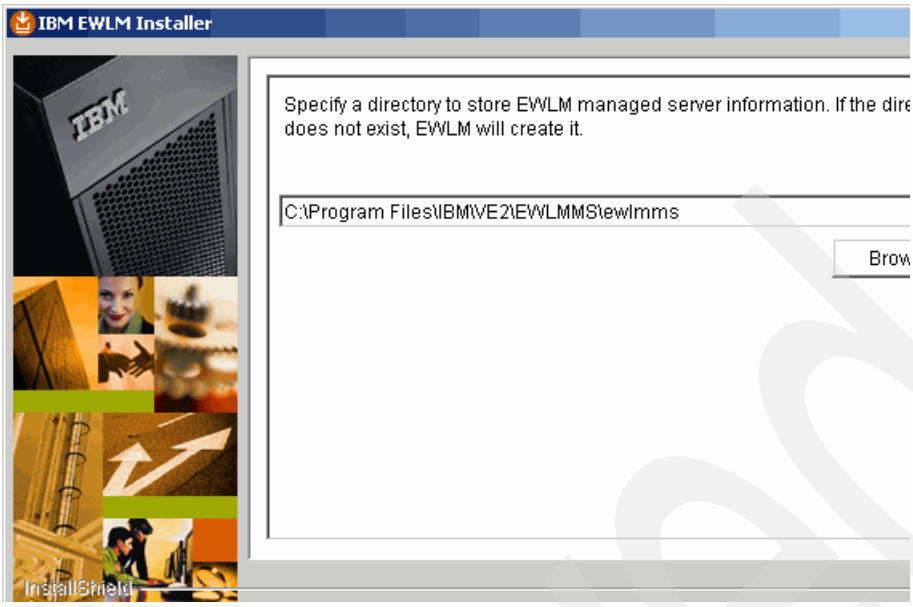


Figure 4-4 Data Directory

4. Select the features you want to install. The window, as shown in Figure 4-5, presents the following choices:
   - Managed Server represents the EWLM Agent and is mandatory.
   - Firewall Broker allows the communication between the EWLM Domain Manager and the EWLM Managed Server to go through a firewall. The counterpart component must be installed on the EWLM Domain Manager side.

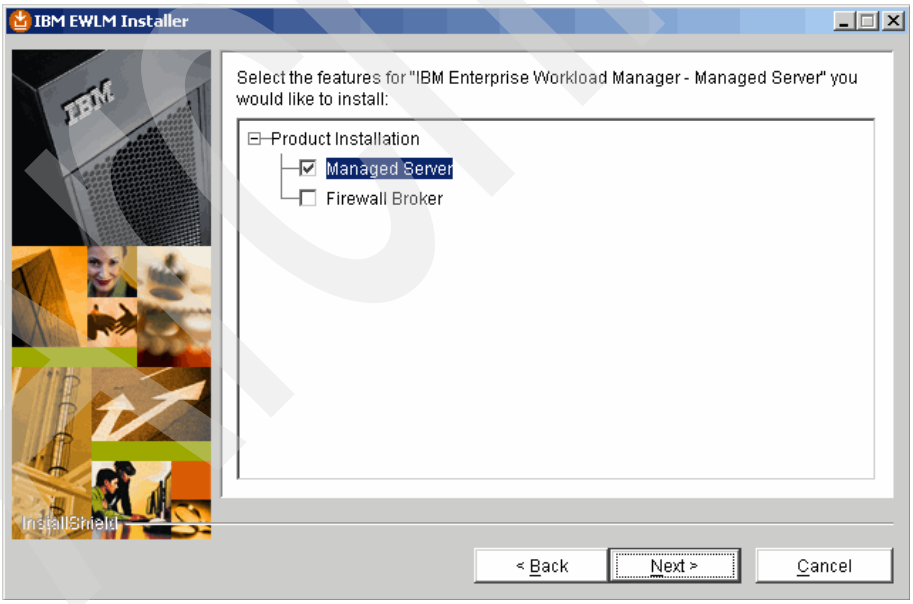


Figure 4-5 Feature Selection

5. Before the installation, a summary window, as shown in Figure 4-6 on page 170, is displayed. This is your last chance to move backwards and modify the selection before the actual installation starts. Click **Next** to start the installation.

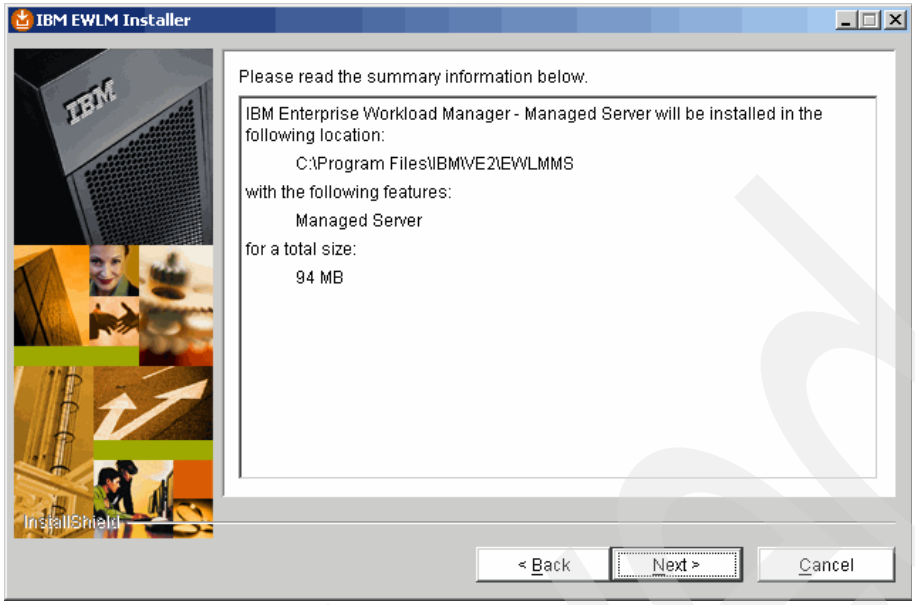


Figure 4-6 Summary window

The installation process informs you about the progress of the process.

6. At the end of the process, the wizard displays a window, as shown in Figure 4-7, from which you can directly start the post-installation configuration. If this is your choice, click **Next**.

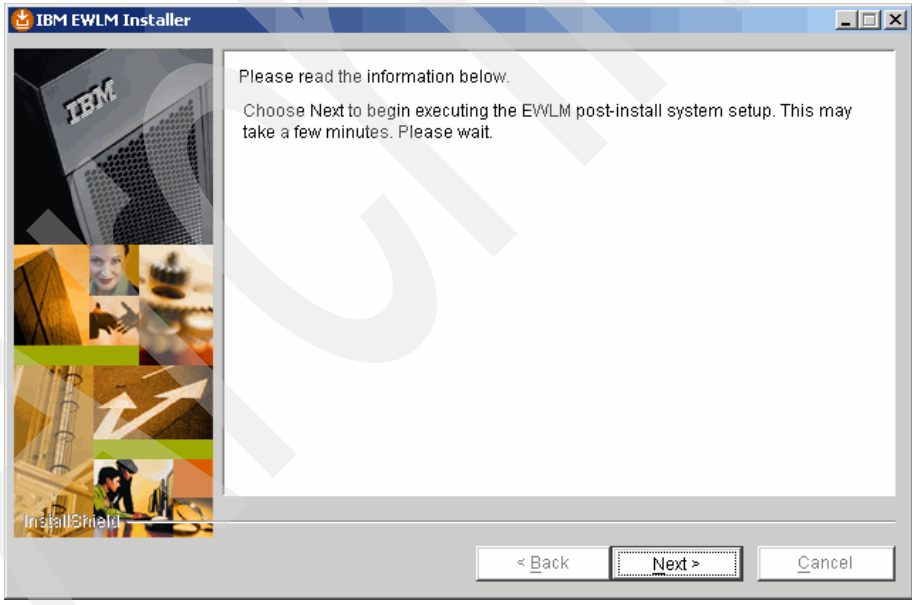


Figure 4-7 EWLM post-install setup

7. Select **Now** in order to launch the configuration wizard and start with the post-installation configuration process. Otherwise, select **Later** to exit the wizard.

Note: The post-installation configuration of the EWLM Managed Server is described in 3.6.3, “IBM Director Server” on page 143.

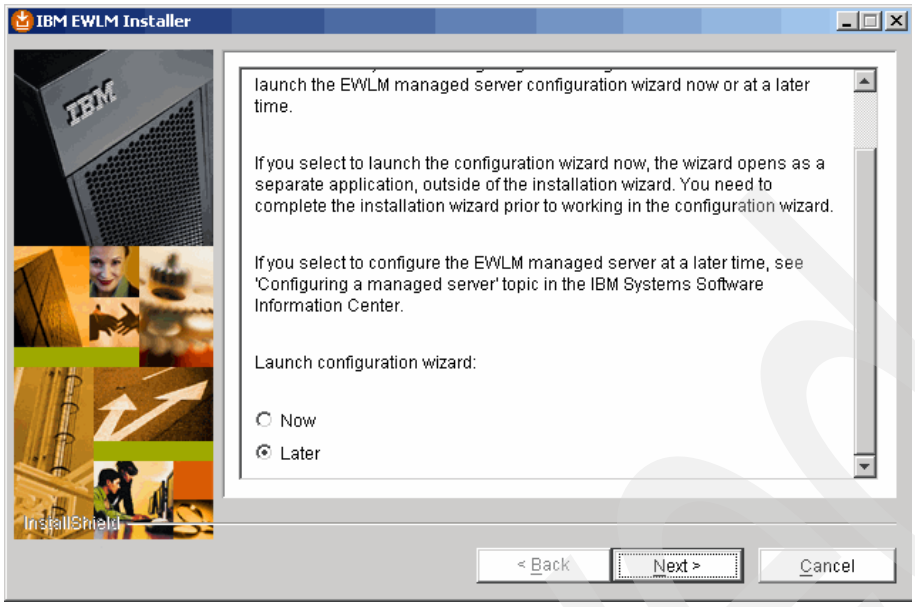


Figure 4-8 Launch Configuration Wizard

8. On Microsoft Windows systems, the EWLM Agent installation requires a reboot and the wizard displays a window, as shown in Figure 4-9. Select **Reboot now** and click **Next** to reboot.

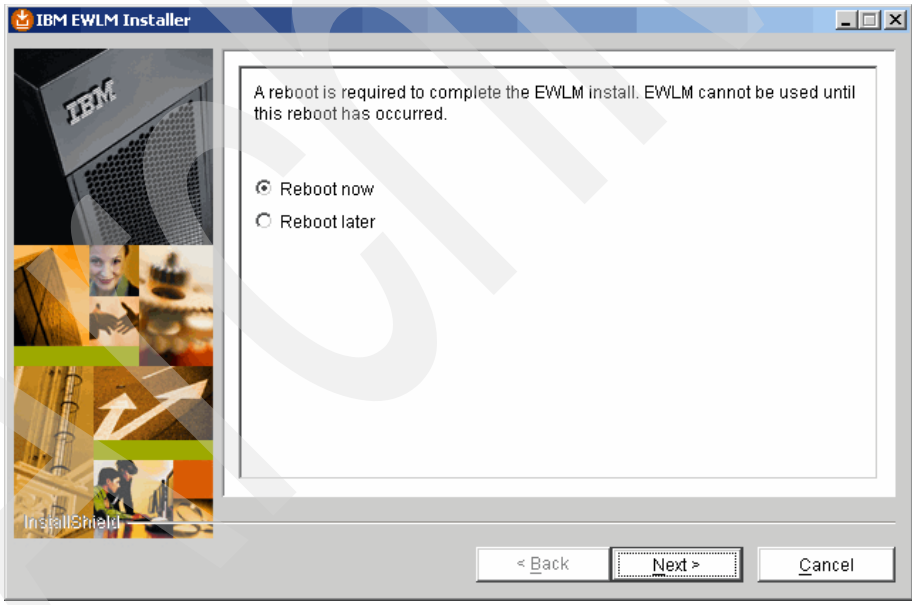


Figure 4-9 Reboot now

9. A window appears informing you about the success of the installation, as shown in Figure 4-10 on page 172. Click **Finish** to leave the process.

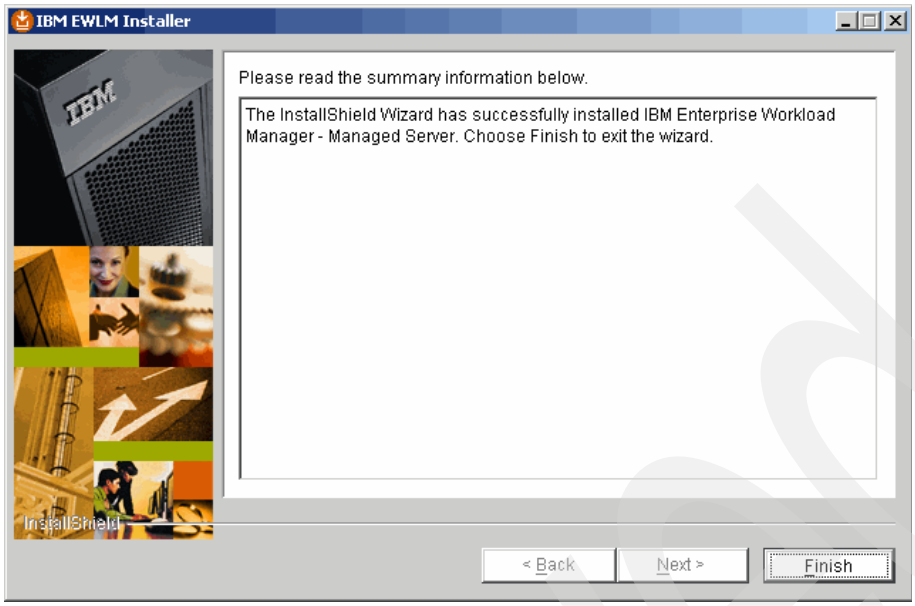


Figure 4-10 Summary information

Console mode

The Console mode is launched by passing the -console parameter to the installation executable program, depending on the operating system:

- For AIX: <*Source Directory*>/Aix/FILES/EWLMMSAix.bin -console
- For Linux on POWER: <*Source Directory*>/LinPPC/FILES/EWLMMSLinPPC.bin -console
- For Linux on IBM System x: <*Source Directory*>/LinX/FILES/EWLMMSLinX.bin -console
- For Microsoft Windows: <*Source Directory*>\FILES\EWLMMSWin.exe -console

This triggers a new shell terminal in which the installation process takes place.

The installation steps are identical to the installation graphical wizard. Refer to the “Graphical wizard” on page 167 for a detailed description of each step’s content. Example 4-4 is an example of the EWLM agent on an AIX system.

Example 4-4 EWLM Agent installation in console mode on AIX systems

```
# /tmp/EWLM/Aix/FILES/EWLMMSAix.bin -console

            Initializing InstallShield Wizard........
            Launching InstallShield Wizard........
-------------------------------------------------------------------------------
Welcome to the InstallShield Wizard for IBM Enterprise Workload Manager -
Managed Server
The InstallShield Wizard will install IBM Enterprise Workload Manager - Managed
Server on your computer.
To continue, choose Next.
IBM Enterprise Workload Manager - Managed Server
IBM
http://www.ibm.com

Press 1 for Next, 3 to Cancel or 5 to Redisplay [1]
-------------------------------------------------------------------------------
IBM Enterprise Workload Manager - Managed Server Install Location
Please specify a directory or press Enter to accept the default directory.
```

```
Directory Name: [/opt/IBM/VE2/EWLMMS]
Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1]
-------------------------------------------------------------------------------
Specify a directory to store EWLM managed server information. If the directory
does not exist, EWLM will create it.

Enter a file: [/ewlmData/ewlmms]

Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1]
-------------------------------------------------------------------------------
Select the features for "IBM Enterprise Workload Manager - Managed Server" you
would like to install:
IBM Enterprise Workload Manager - Managed Server
To select/deselect a feature or to view its children, type its number:
1.  [ ] Managed Server
2.  [ ] Firewall Broker

   Other options:
0. Continue installing

   Enter command [0] 1

Select the features for "IBM Enterprise Workload Manager - Managed Server" you
would like to install:

   IBM Enterprise Workload Manager - Managed Server
To select/deselect a feature or to view its children, type its number:

     1.  [x] Managed Server
     2.  [ ] Firewall Broker

   Other options:

     0. Continue installing

   Enter command [0]

Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1]
-------------------------------------------------------------------------------
IBM Enterprise Workload Manager - Managed Server will be installed in the
following location:
/opt/IBM/VE2/EWLMMS
with the following features:
Managed Server
for a total size:
89.4 MB
Press 1 for Next, 2 for Previous, 3 to Cancel or 5 to Redisplay [1]
Installing IBM Enterprise Workload Manager - Managed Server. Please wait...
|-----------|-----------|-----------|------------|
0%         25%         50%         75%        100%
||||||||||||||||||||||||||||||||||||||||||||||||||

Creating uninstaller...
Updating the inventory ...
Creating uninstaller...
```

Archived

```
--------------------------------------------------------------------------------
Please read the information below.
Choose Next to begin executing the EWLM post-install system setup. This may
take a few minutes. Please wait.
Press 1 for Next or 5 to Redisplay [1]
--------------------------------------------------------------------------------
The InstallShield Wizard has successfully installed IBM Enterprise Workload
Manager - Managed Server. Choose Finish to exit the wizard.

Press 3 to Finish or 5 to Redisplay [3]
```

Silent installation using the response file

Silent installation is based on the parameters contained in the response file provided as part of the EWLM agent installation package, and depends on the operating system as follows:

- For AIX: <*Source Directory*>/Aix/FILES/EWLMMSAix.rsp
- For Linux on POWER: <*Source Directory*>/LinPPC/FILES/EWLMMSLinPPC.rsp
- For Linux on IBM System x: <*Source Directory*>/LinX/FILES/EWLMMSLinX.rsp
- For Microsoft Windows: <*Source Directory*>\FILES\EWLMMSWin.rsp

All parameters already have default values that may be acceptable. However, you can consider the following changes:

- -P firewallBrokerBean.active=true

  By default, the firewall broker feature is installed. This feature is not required in most environments.

- -W reboot.option="true"

  This parameter is specific to the Microsoft Windows version of the EWLM agent. As the reboot is mandatory at the end of the installation, we recommend that you replace the default value by true.

Example 4-5 presents the parameters of the response file in a Windows environment.

Example 4-5 Response file parameters for Microsoft Windows systems

```
-P installLocation="C:\Program Files\IBM\VE2\EWLMMS"
-W ewlmUserData.location="C:\Program Files\IBM\VE2\EWLMMS\ewlmms"
-P managedServerBean.active=true
-P firewallBrokerBean.active=false
-W configWizard.launchOption="false"
-W reboot.option="true"
```

You can launch the silent installation by passing the following parameters:

- -silent
- -options <*Response file filename*>

Pass these parameters to the installation executable that depending on the operating system as follows:

- For AIX: <*Source Directory*>/Aix/FILES/EWLMMSAix.bin -silent -options <*Source Directory*>/Aix/FILES/EWLMMSAix.rsp
- For Linux on POWER: <*Source Directory*>/LinPPC/FILES/EWLMMSLinPPC.bin -silent -options <*Source Directory*>/LinPPC/FILES/EWLMMSLinPPC.rsp

- For Linux on IBM System x: *<Source Directory>*/LinX/FILES/EWLMMSLinX.bin -silent -options *<Source Directory>*/LinX/FILES/EWLMMSLinX.rsp
- For Microsoft Windows: *<Source Directory>*\FILES\EWLMMSWin.exe -silent -options "*<Source Directory>*\FILES\EWLMMSWin.rsp"

Silent installation using command line parameters

This mode is very similar to the previous one except that rather than picking the parameter values from the response file, the installer expects them passed as command line parameters.

The parameter list is as follows:

- -P managedServerBean.active=[true | false]

  Specify `true` to install the agent. The default value (if omitted) is false.

- -P firewallBrokerBean.active=[true | false]

  Specify `true` to install the firewall broker feature. The default value (if omitted) is false.

- -P productBean.installLocation=<Installation directory>

  This is to specify the installation directory.

Example 4-6 is an example of using the command line in an AIX environment.

Example 4-6 Silent installation using command line parameters on AIX

```
/tmp/EWLMMSAAix.bin -silent -P managedServerBean.active=true
                                        -P firewallBrokerBean.active=false
```

Director Agent

Three levels of Director Agent exist, providing distinct footprints on the system and distinct levels of functionality:

- Level-0 ("Agentless"): No Director Agent installed, limited functionality.
- Level-1 ("Core Services"): Basic features available only, with a reduced footprint on the managed system.
- Level-2 ("IBM Director Agent"): Full functionality available, with a higher footprint on the managed system.

You can find details about the differences of the three levels of functionality in *Implementing IBM Director 5.10*, SG24-6188.

Note: The installation process for the Level-1 and Level-2 agents are very similar and will be described together, with explicit mentions of the slight differences.

As for the Director Server, the Director Agent installation can be performed in two ways:

1. Through a graphical installation wizard

   This mode is available on Microsoft Windows systems only.

2. In silent mode

   This mode is valid for all types of systems, including Microsoft Windows ones.

Archived

Graphical installation wizard

Reminder: The graphical wizard installation is available on Microsoft Windows systems only.

Start the wizard with the following command:

```
<Installation Files Directory>\setup.exe
```

Perform the following steps:

1. Set up the IBM Director.

   This window, as shown in Figure 4-11, offers the choice between the components to install. Choose either **Install IBM Director Agent** (Level-2) or **Install IBM Director Core Services** (Level-1.)

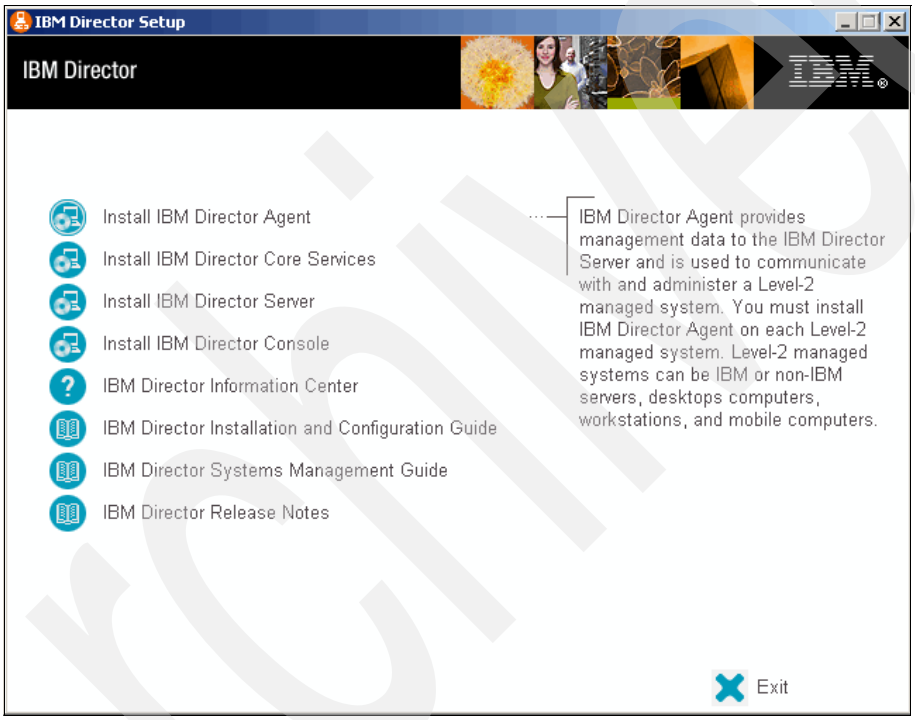


Figure 4-11 IBM Director Setup welcome screen

The installation wizard starts and presents a welcome window. Accept the license agreement.

If the wizard finds services installed already, a window, as shown in Figure 4-12, mentions the existence of a previous installation.

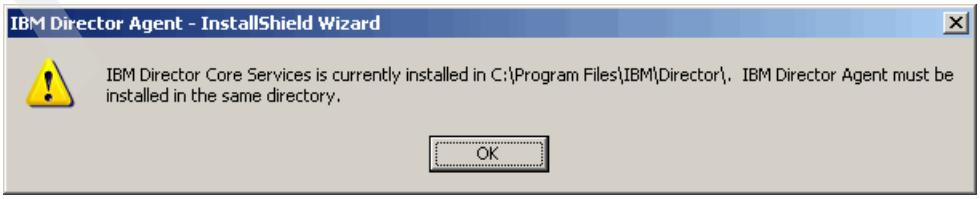


Figure 4-12 InstallShield wizard warning: Core Services already installed

Archived

2. As shown in Figure 4-13, you can choose the optional features to install for a Level-2 Director Agent. You can also specify the installation directory of the agent, as shown in Figure 4-14.

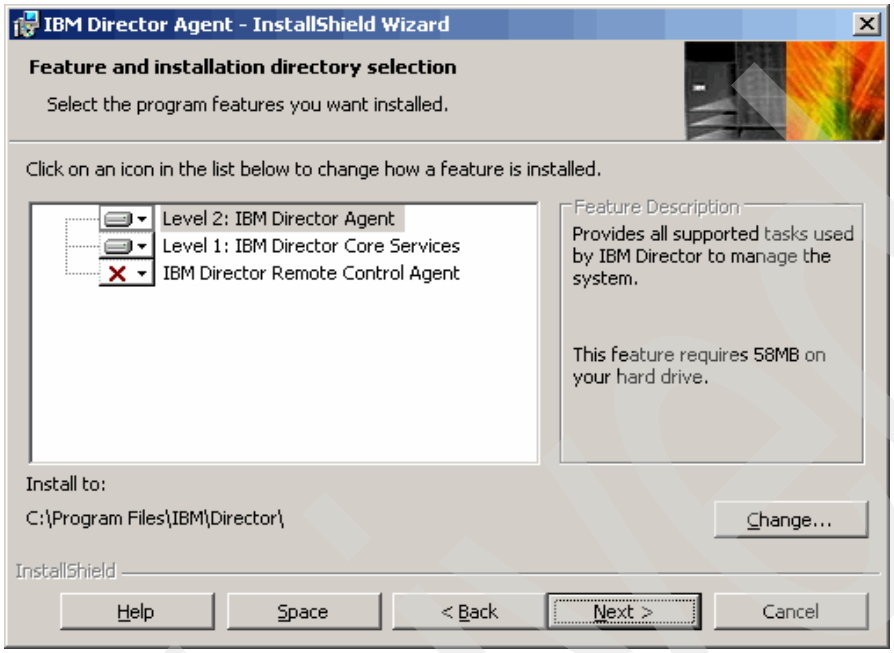


Figure 4-13 Feature and installation directory selection (Level-2 agent installation only)

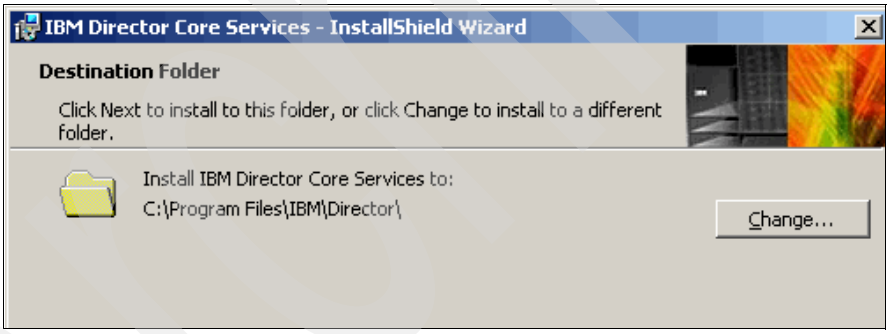


Figure 4-14 Destination folder (Level-1 agent installation only)

3. The next step allows you to configure, as shown in Figure 4-15 on page 178, the security settings applicable to the Director Agent (Level-2 only), including:
   - Encrypting the data transmission between the Director Server and the Director Agent
   - Forcing the Director Server to request access to the managed system
   - Declaring the Director Server to which the Director Agent should automatically register after installation

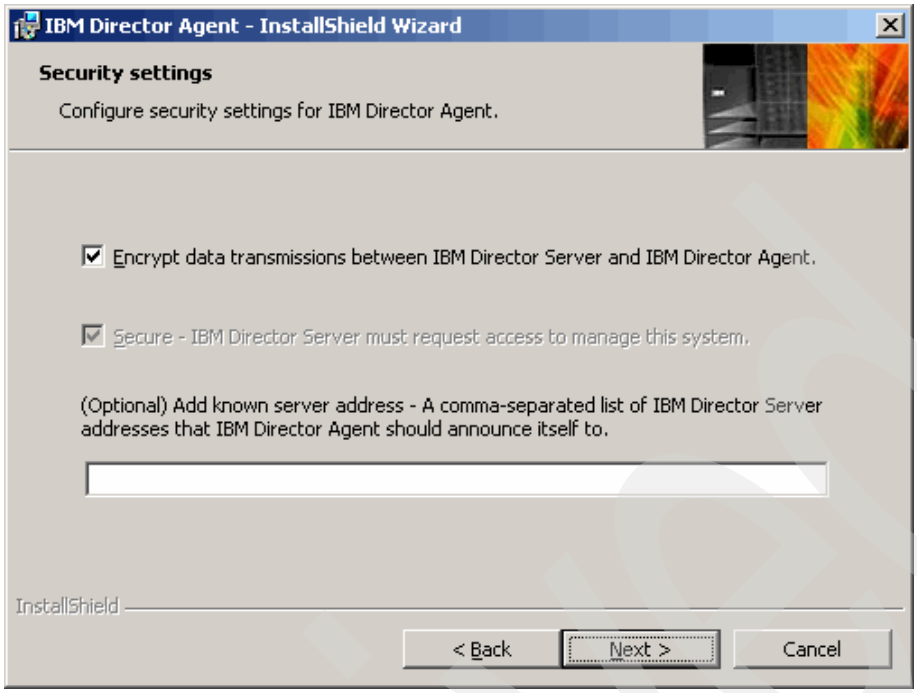


Figure 4-15 Security Settings (Level-2 agent installation only)

4. You can specify the location for the software package distributed to the managed system. The default value is acceptable (Figure 4-16).

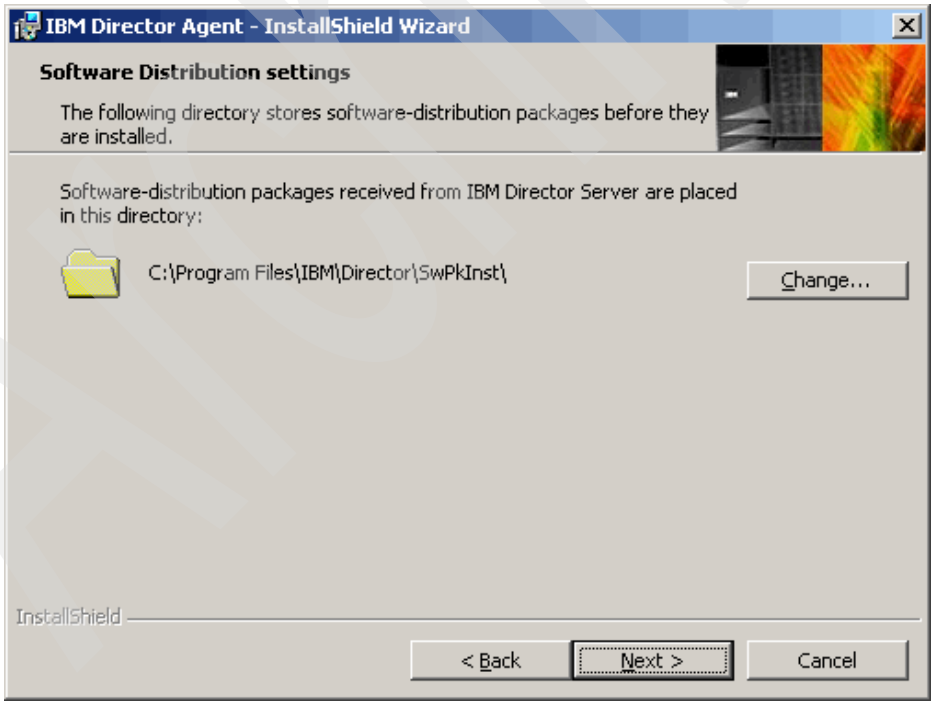


Figure 4-16 Software Distribution settings (Level-2 agent installation only)

5. This step is the last chance to come back to previous selection windows before the actual installation starts. A window appears, as shown in Figure 4-17, where you can click **Install** or go back to the previous steps.

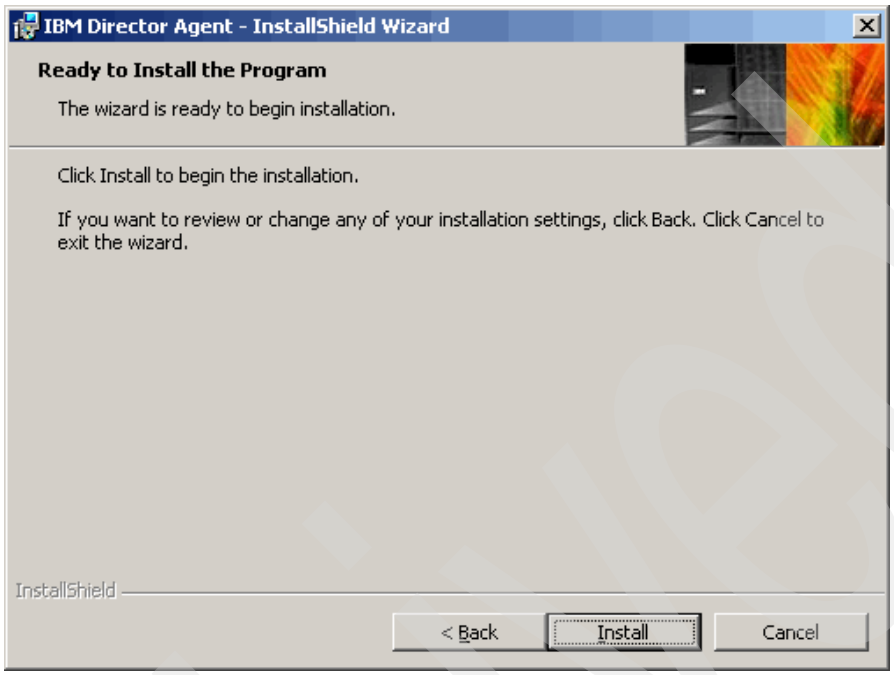


Figure 4-17 Ready to Install the Program

6. The installer displays the installation progress, as shown in Figure 4-18.



Figure 4-18 Installing the IBM Director Agent

7. The next step allows you to configure the network driver (Level-2 agent installation only). This window, as shown in Figure 4-19 on page 180, configures the Director Agent network connectivity. You should select the network adapters on which the Director Agent should listen. By default, the Director Agent listens on all the network adapters of the system. However, you can modify this configuration in a more restrictive way.

> **Note:** This setting is significant whenever the Director server resides on a network dedicated to the system administration, and is distinct from the actual “production” network. In this case, we recommend that you restrict the frame of the Director Agent to the administration network only.

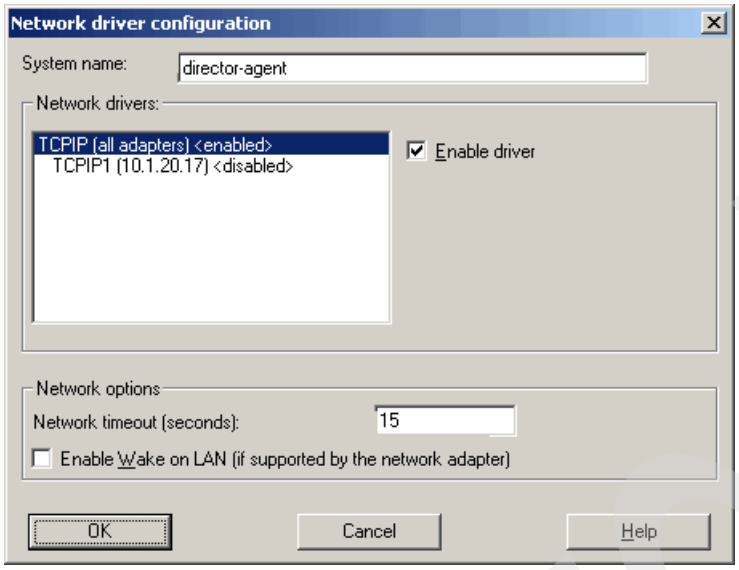


Figure 4-19 Network driver configuration

8. When the installation is complete, the window shown in Figure 4-20 appears. To exit the installation wizard, click **Finish**.

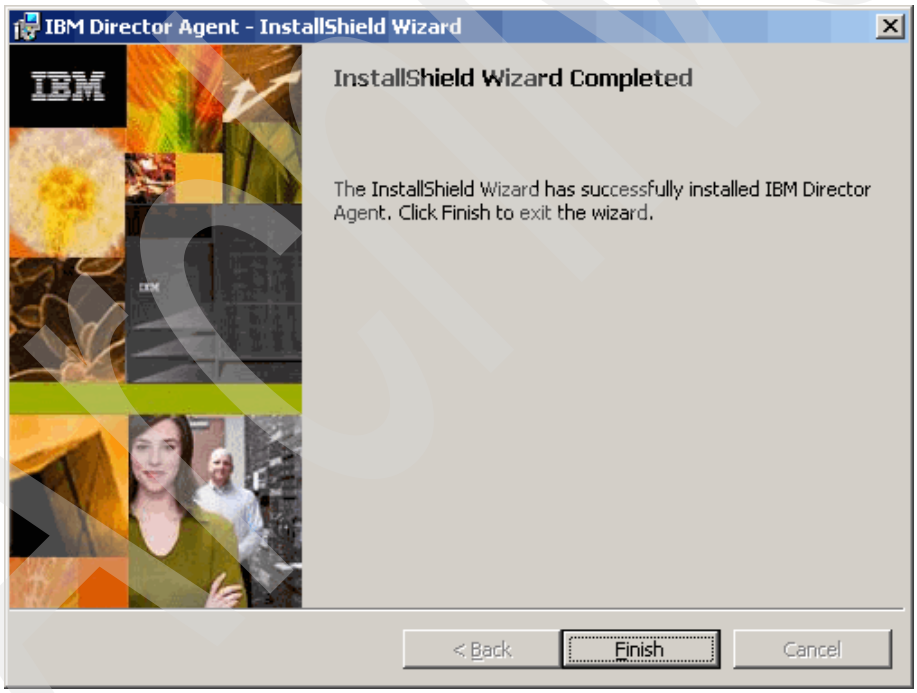


Figure 4-20 InstallShield Wizard Completed

Silent installation

The silent installation relies on two files:

- An executable file that launches the installation process. This file name depends on the operating system:
  - For AIX: dir5.10_agent_aix.sh
  - For Linux: dir5.10_agent_linux.sh
  - For Microsoft Windows: dir5.10_agent_windows.exe
- A text response file that provides the installer with the appropriate user settings. This file is named dirserv.rsp.

Depending on the operating system, these two files are located in different directories:

- For AIX: <*Installation Files Directory*>/director/agent/aix/FILES
- For Linux on POWER:<*Installation Files Directory*>/director/agent/linux/ppc/FILES
- For Linux on System x: <*Installation Files Directory*>/director/agent/linux/i386/FILES
- For Windows: <*Installation Files Directory*>\director\agent\windows\i386\FILES

Example 4-7 and Example 4-8 are the files we used in our environment.

Example 4-7 Executable and response file on Linux on System x

```
/mnt/source/IbmDirector/director/agent/linux/i386/FILES/dirserv.rsp
/mnt/source/IbmDirector/director/agent/linux/i386/FILES/dir5.10_agent_linux.sh
```

Example 4-8 Executable and response file on Microsoft Windows systems

```
M:\IbmDirector\director\agent\windows\i386\FILES\dirserv.rsp
M:\IbmDirector\director\agent\windows\i386\FILES\dir5.10_agent_windows.exe
```

The default values of the response files are acceptable for a standard installation. For specific customization, refer to the parameter comments inside the response file itself.

Note: The number of parameters contained in the response file is much greater on the Microsoft Windows platform than on the other ones (AIX, Linux). This is due to the history of the product, which is a former IBM System x / BladeCenter management tool. Still, even on Microsoft Windows systems, the default values remain perfectly acceptable.

In order to launch the silent installation process, pass the following parameters to the executable:

- For AIX or Linux: -r <*Response file file name*>
- For Microsoft Windows: silent RSP="<*Response file file name*>"

Note: On all the systems, the response file parameter is optional. If not specified, the installation will be executed using the default installation settings.

Archived

Example 4-9, Example 4-10, and Example 4-11 shows the commands we used in our environments.

Example 4-9 Launching the silent installation on AIX systems

```
/mnt/source/IbmDirector/director/agent/aix/FILES/dir5.10_agent_aix.sh
                      -r
                      /mnt/source/IbmDirector/director/agent/aix/FILES/dirserv.rs
                      p
```

Example 4-10 Launching the silent installation on Linux on System x

```
/mnt/source/IbmDirector/director/agent/linux/i386/FILES/dir5.10_agent_linux.sh
                -r
                /mnt/source/IbmDirector/director/agent/linux/i386/FILES/dirserv.rs
                p
```

Example 4-11 Launching the silent installation on Microsoft Windows systems

```
M:\IbmDirector\director\agent\windows\i386\FILES\dir5.10_agent_windows.exe
                silent
                RSP="M:\IbmDirector\director\agent\windows\i386\FILES\dirserv.rsp"
```

Note: On Microsoft Windows systems, the system *automatically* reboots at the end of the installation process.

4.3 Post-installation tasks

There are a few configuration tasks you should perform to make the agents function as you want.

4.3.1 EWLM Managed Server

Similar to the EWLM Domain Manager, each EWLM Managed Server needs to be configured before you start it after the installation.

Configuration

You can perform this configuration in two ways:

- Graphical configuration wizard

  This requires a graphical use interface available on the system.

- Command line mode

Graphical configuration wizard

Launch the graphical configuration wizard with the following commands, depending on the operating systems:

- For AIX: **<*Installation Directory*>/EWLMMS/bin/configWizardMS.sh**
- For Linux: **<*Installation Directory*>/EWLMMS/bin/configWizardMS.sh**
- For Microsoft Windows: **<*Installation Directory*>\EWLMMS\bin\configWizardMS.bat**

Example 4-12 on page 183 and Example 4-13 on page 183 show the commands we used in our environment.

Example 4-12 Launching the graphical configuration wizard on Linux systems

```
/opt/IBM/VE2/EWLMMS/bin/configWizardMS.sh
```

Example 4-13 Launching the graphical configuration wizard on Microsoft Windows systems

```
C:\Program Files\IBM\VE2\EWLMMS\bin\configWizardMS.bat
```

For the configuration, perform the following steps:

1. First, select the action to perform, as shown in Figure 4-21.

   Note: The configuration wizard can be used later on for modifying the Managed Server configuration.

   Select **Create managed server** and click **Next**.

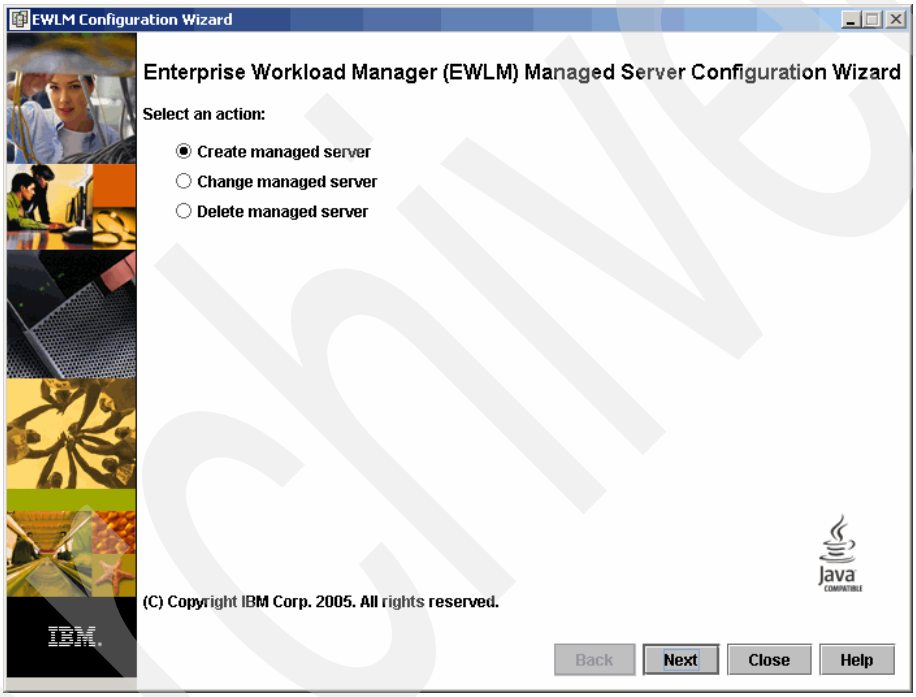


Figure 4-21 Select an action

2. Provide the configuration ID, as shown in Figure 4-22 on page 184. This is the unique identifier of the EWLM Managed Server configuration settings. You may consider using the domain name set at the Domain Manager level. Click **Next**.

Archived

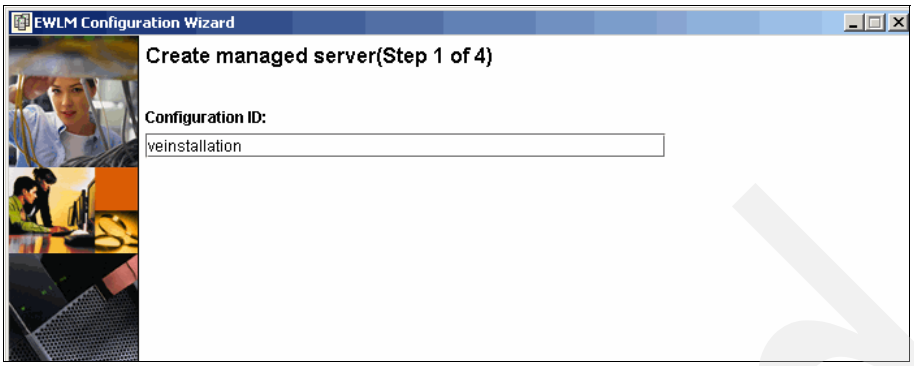


Figure 4-22 Configuration ID

3. Define the address and the TCP/IP port to connect to, as shown in Figure 4-23. It may be either one of the EWLM Domain Manager or the one of the firewall broker if this feature has been implemented. The address may be either the IP address or the host name (short or fully-qualified) of the peer. The port must match the one specified during the EWLM Domain Manager configuration. Click **Next**.

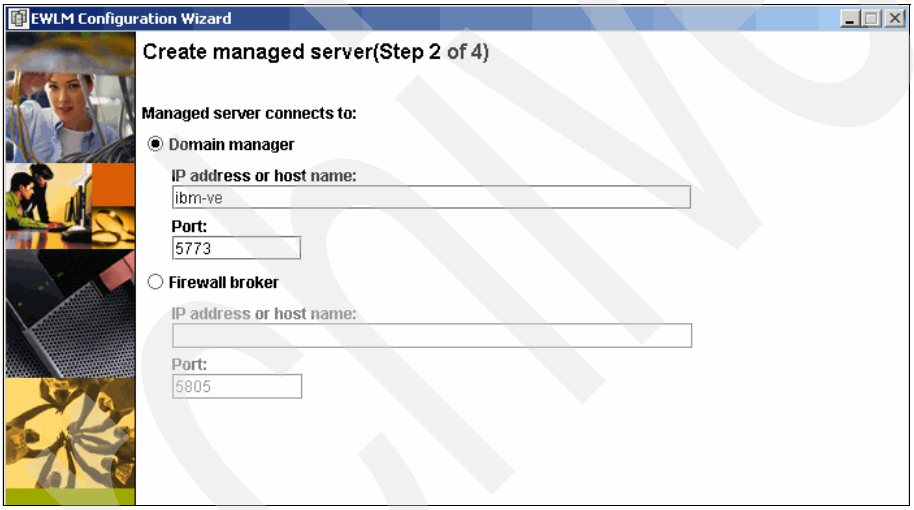


Figure 4-23 Domain Manager

4. It may be necessary to use an HTTP proxy server to establish the communication between the EWLM Managed Server and the EWLM Domain Manager. If this is the case, specify the proxy parameters, as shown in Figure 4-24. Click **Next**.

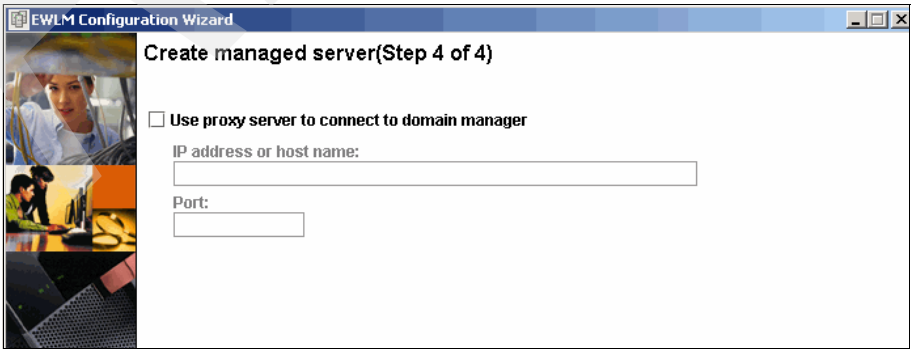


Figure 4-24 Use proxy server

5. The managed server is then created and the configuration progress is displayed as shown in Figure 4-25.

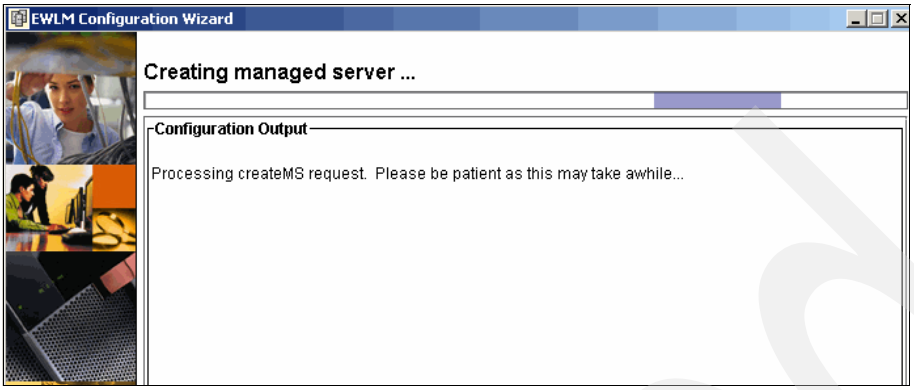


Figure 4-25 Creating managed server

When the configuration process is complete, a window appears as shown in Figure 4-26. Click **OK** and exit the wizard.

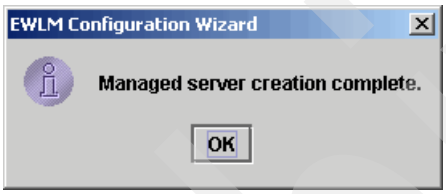


Figure 4-26 Managed server creation complete

Command-line mode

The command-line mode allows the configuration of the EWLM Managed Server using the command line.

The command to launch depends on the operating system:

- For AIX or Linux: **<*Installation Directory*>/EWLMMS/bin/createMS.sh**
- For Microsoft Windows: **<*Installation Directory*>\EWLMMS\bin\createMS.bat**

The parameters you should provide are as follows:

- `<configID>`, which is the configuration ID.
- `-ma <address>`, which is the address of the EWLM Domain Manager or the firewall broker.
- `-mp <port>`, which is the port of the EWLM Domain Manager or firewall broker.
- `-auth [None | ServerSSL | ClientServerSSL]`, which is the security level for the connection between the EWLM Domain Manager and the EWLM Managed Server.

Example 4-14 is an example of the command we used in our AIX and Linux environment.

Example 4-14 EWLM Managed Server command line configuration on AIX / Linux

```
/opt/IBM/VE2/EWLMMS/bin/createMS.sh redbook -ma ewlm-dm -mp 5773 -auth None
```

Archived

For an AIX platform: If you plan to use EWLM on a High-Availability Cluster Multi-Processing (HACMP™) environment, you should install one agent per management node. The EWLM should be associated with a server, not with an HACMP resource group. Therefore, you should be careful about the assignment rules. The rules specify that you should classify resources, independent of a server box. If you are using a user name for classified EWLM resources, then the same user name should exist on all HACMP nodes. If you are using a path name where the program are located, then the program should be placed in the same path name on all HACMP nodes.

Startup

When the configuration is complete, you should launch the Managed Server, if it is not already running.

The command to invoke this action depends on the operating system:

- For AIX or Linux: **<*Installation Directory*>/EWLMMS/bin/startMS.sh <configID>**
- For Microsoft Windows: **<*Installation Directory*>\EWLMMS\bin\startMS.bat <configID>**

Example 4-15 and Example 4-16 are the commands we used in our environments.

Example 4-15 EWLM Managed Server startup command on AIX / Linux systems

```
/opt/IBM/VE2/EWLMMS/bin/startMS.sh redbook
```

Example 4-16 EWLM Managed Server startup command on Microsoft Windows systems

```
C:\Program Files\IBM\VE2\EWLMMS\bin\startMS.bat redbook
```

Note: Microsoft Windows systems also have an operating system service that is automatically created by the EWLM agent installer. You can use it to start and stop the EWLM Managed Server.

On Linux or Microsoft Windows systems, the startup of the EWLM Managed Server at boot time is automatic. On AIX systems, you can consider adding an entry in the inittab, as shown in Example 4-17.

Example 4-17 Adding EWLM Managed Server into inittab on AIX systems

```
mkitab "ewlmms:2:wait:/opt/IBM/VE2/EWLMMS/bin/startMS.sh redbook"
```

4.3.2 Director Agent

There is no post-installation configuration task for the Director Agent.

Startup

The Director agent is automatically started in Linux. You can start the Director Agent with the following commands, depending on the operating systems:

- For AIX or Linux: **/opt/ibm/director/bin/twgstart.sh**
- For Microsoft Windows: **C:\Program Files\IBM\Director\bin\twgipcsv.exe**

Note: Microsoft Windows systems also feature an operating system service that is automatically created by the Director Agent installer. You can use it to start and stop the Director Agent.

Archived

Archived

5

Applying fixes

This chapter describes the fixes and update process for a Virtualization Engine environment. Maintaining up-to-date Virtualization Engine components and system services is very important for correct functioning of the entire environment. This is important for functionality and for performance.

Fixes and patches may exist for all the components, services, or the Common Component. Fixes are divided by services and operating system. A given update can apply to either the management server or agent managed node. This means relative fixes can be installed on any managed or management nodes in the enterprise.

New tools for Virtualization Engine Version 2.1 are available to make finding the fixes and installing the updates easier. The following tools are available:

- You can now use the Virtualization Engine console to install an update. The installation is managed by the *Update Installer*. When using the Virtualization Engine console to install an update, the Virtualization Engine console uses this utility. If you have to install an update without using the Virtualization Engine console, you can use this utility directly. In either case, you *must download* and install the Update Installer before you can install any updates. The Update Installer is an alternative method for determining what products are installed.
- The *Virtualization Engine Maintenance Environment Assistant* is a program that checks the environment for Virtualization Engine products that failed to register during the automated installation process and attempts to register the products.
- A Web site can show you the latest available fixes.
- For EWLM updates, you should also review some product-specific information about installing updates.

© Copyright IBM Corp. 2006. All rights reserved.

This task is the last task of the installation process, as depicted in Figure 5-1.

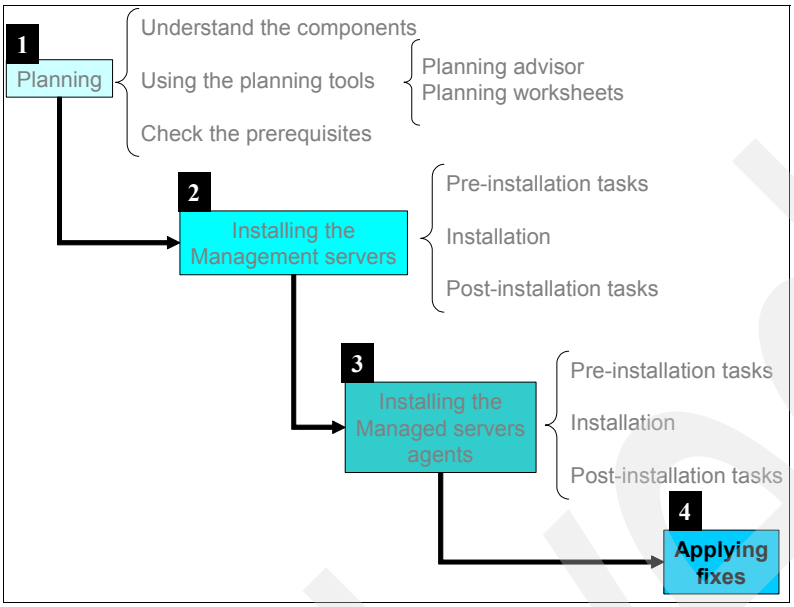


Figure 5-1 Task four: applying fixes

Archived

5.1 About fixes

As with every product, the Virtualization Engine may have defaults and bugs: they are corrected through fixes. All Virtualization Engine products share a common set of processes for acquiring and managing fixes and updates. New in this release are several features that make it easier to find the fix or update you need and to then install and manage updates to your Virtualization Engine products.

5.1.1 Terminology

The Virtualization Engine updates are currently serviced using a *FixPeriod* model. A FixPeriod provides a logical grouping for a set of fixes. Any fixes that belong to the same fix level are part of the same FixPeriod. You cannot download and install a FixPeriod. You must download and install the individual *Fix Packs* that comprise the FixPeriod.

A *Fix Pack* is one single installable unit for multiple individual problems. Fix Packs are the smallest entity that you can install. An operating system specific Fix Pack is provided for each product component when that product component needs to make fixes on a particular operating system. Each Fix Package for a FixPeriod contains a *readme file* with information specific to that package. The readme file might contain any special instructions for installing the update and describing important changes delivered by the update.

Attention: Updates for different platforms may have the same number (level), however, they may include different fixes.

Updates are cumulative. One update will include all the fixes of all updates that have been made available until then for this release.

The FixPeriod name follows the V.R.F format, where V is the version, R the release, and F the fix period level. For example, at the time of writing this book, three FixPeriods are available for the Virtualization Engine version 2.1. These are 2.1.21, 2.1.20, and 2.1.10.

5.1.2 The Web site

You can refer to or download the fixes from the following Web site:

`http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/`

This Web page allows you to download *critical fixes* for Version 2 and Version 1 of the Virtualization Engine, as shown in Figure 5-2 on page 192.

If this is your first installation, this is a good source of information regarding the latest fixes and we recommend you go through the provided list.

Archived



Figure 5-2 IBM Virtualization Engine fix Web page

You cannot chose the FixPeriod in the window shown in Figure 5-2. To select a product of the Virtualization Engine, when you select a product from the product name list, another window is displayed, where you can choose the FixPeriod, as shown in Figure 5-3.

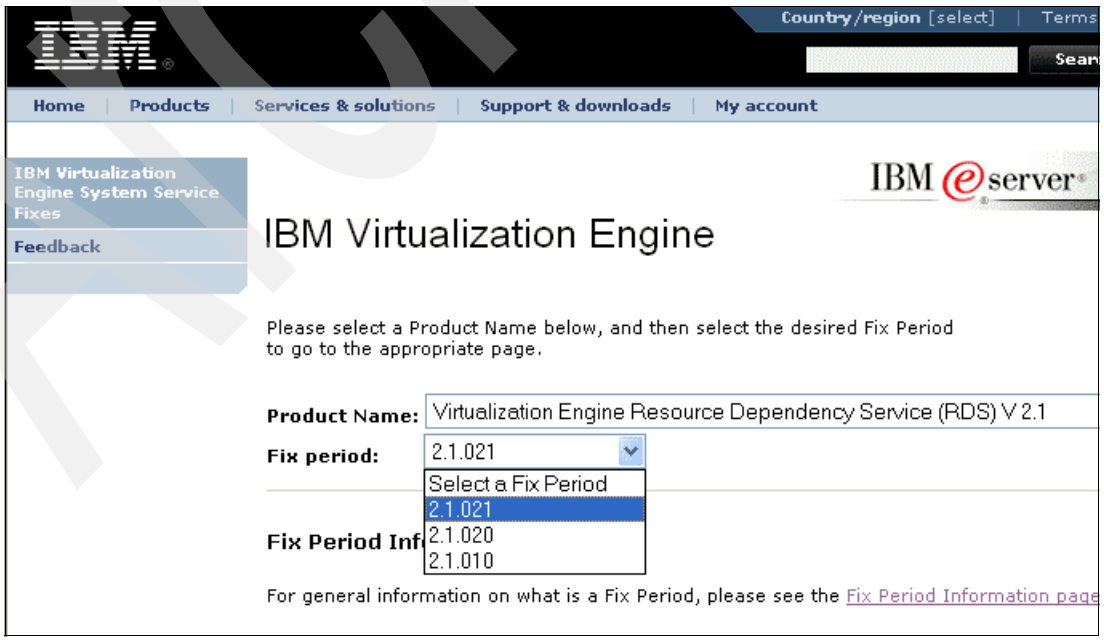


Figure 5-3 Get fixes list by product name

When you choose the FixPeriod, a third window appears, as shown in Figure 5-4, in which you can get the list of the Fix Packs by operating system or by product component.

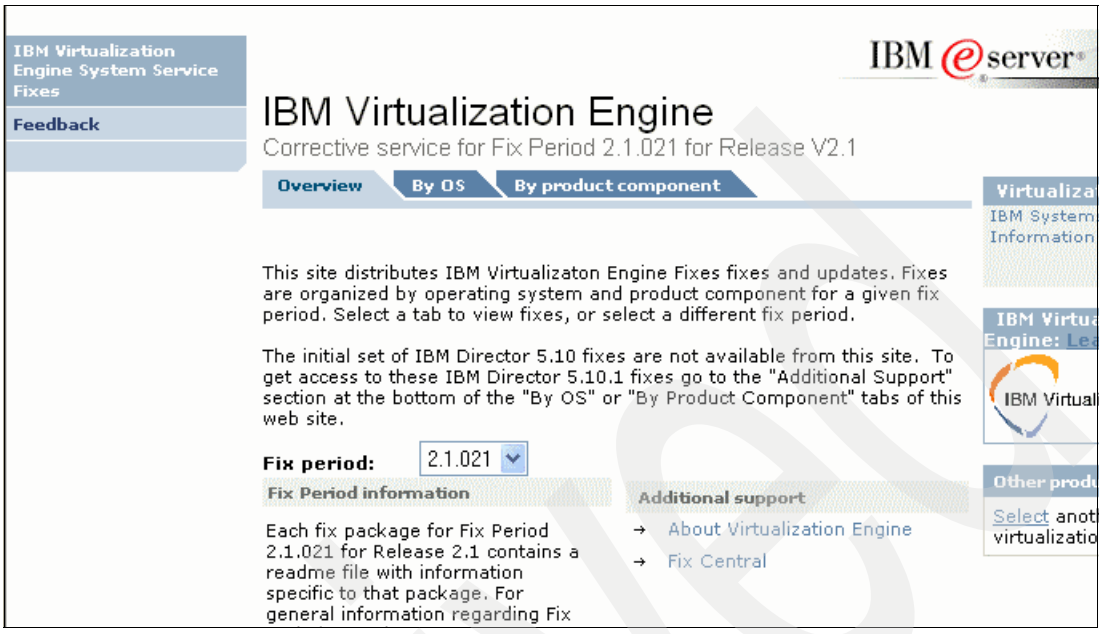


Figure 5-4 Fix Web site download page

If you choose the **By product component** tab, the list of the packages is provided. From this window, you can download the required fix, as shown in Figure 5-5.



Figure 5-5 Package list

5.1.3 Verify the fix level of your environment

It may be important to identify which Virtualization Engine components or subcomponents have been installed on a particular machine and which updates have been already installed.

The Virtualization Engine provides two methods to verify which patches or fix levels, if any, are installed and on which products in your environment:

- The first method uses the Virtualization Engine console.

  Log in to the console, expand **Virtualization Engine console**, open **Hardware and software**, and then select **Virtual Engine Updates**, as shown in Figure 5-6.

Note: The Virtualization Engine console will display only inventory information for management servers running i5/OS, AIX, and Linux on IBM System x. Software installed on other operating systems will not be displayed.

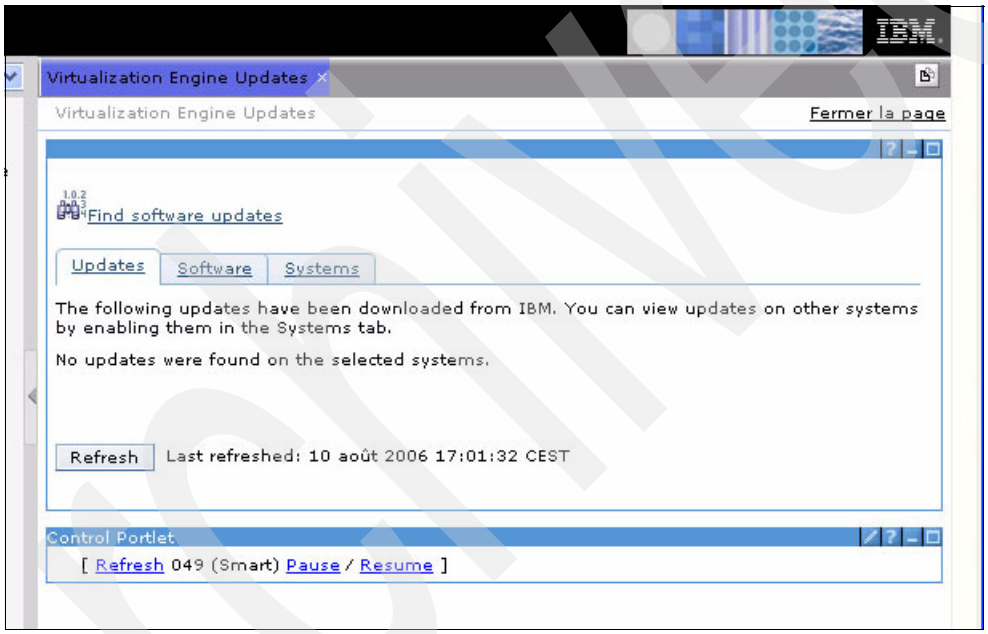


Figure 5-6 How to use the Virtualization Engine console to verify the fix level

- The second method is to do it manually using the command line.

  You may check the patch level of the components using the following commands, which depend on the operating system:

  – On AIX and Linux: `# ./vefixlevel -<service ID>`

  – On Windows: `> vefixlevel -<service ID>`

  You will find the command in the /bin folder of the Virtualization Engine installation directory.

  In this command, <*service ID*> is the ID of the product for which you would like to display the information. Running the command without the <*service ID*> will display the fix for all installed Virtualization Engine products, as shown in Figure 5-7 on page 195.

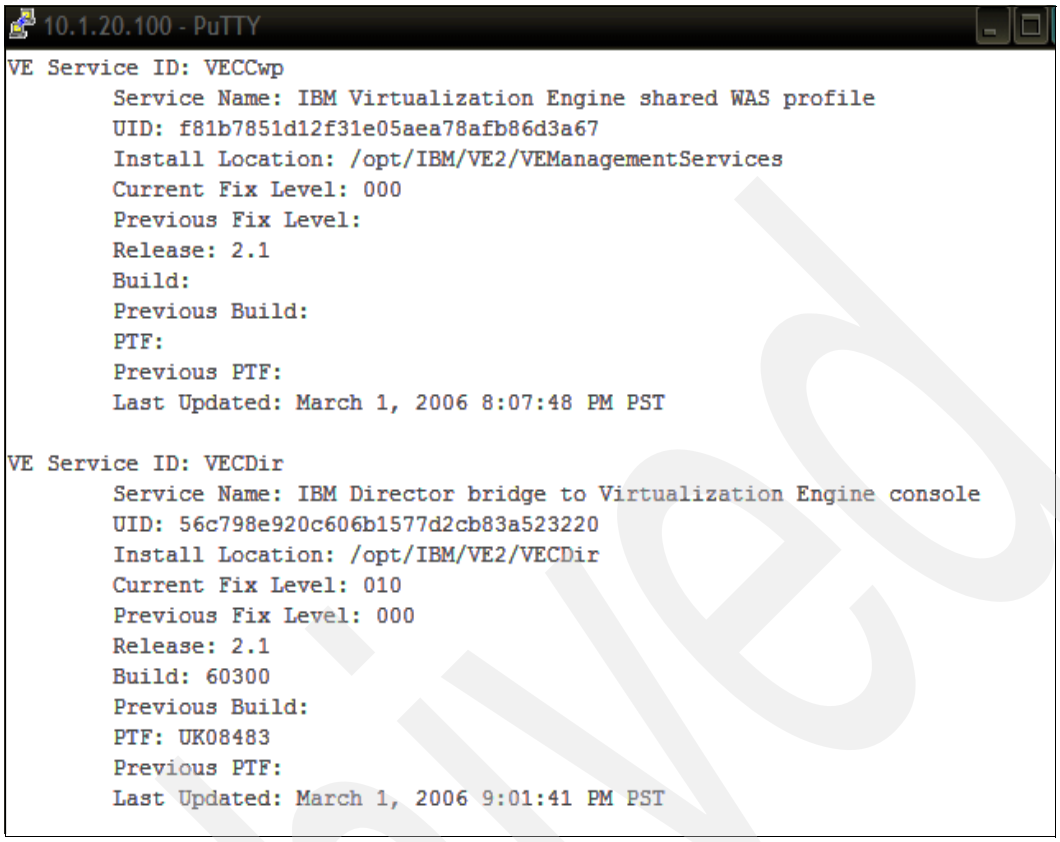


Figure 5-7 Result of vefixlevel command

For each component, there is a great deal of information, such as the product unique identifier (UID) and the installation location, the release, and the current fix level. In Figure 5-7, we can see that:

- The VECCwp (the shared WebSphere Application Server profile) does not have any fixes installed.
- The VECDir (IBM Director Bridge to Virtualization Engine console) has a build level of 60300 and a program temporary fix (PTF) number UK08483 installed.

The field “Previous PTF” records previous patches installed before the patch currently applied. You can also remove patches if any functionality problems occur.

Archived

Table 5-1 describes the Virtualization Engine service IDs you can use in this command for a specific component or subcomponent.

Table 5-1 Product key description

Product key	Service ID	Description
VEC	VEC	Virtualization Engine console
EWLM	EWLM	Enterprise Workload Manager
Director	n/a	IBM Director
RDS	RDS	Resource Dependency Service
WAS	WAS	WebSphere Application Server
ITDS	ITDSServer	IBM Tivoli Directory Server
DB2	DB2	IBM DB2
ISC	ISC	Integrated Security Console
JView	JVIEWS	iLog JViews
VECMc	VECMc	Management Central Bridge to Virtualization Engine console
VECDir	VECDri	IBM Director Bridge to Virtualization Engine console
VECEWLM	VECEWLM	EWLM bridge to Virtualization Engine console
VECCsm	VECCsm	CSM bridge to Virtualization Engine console
ODITypeReg	ODITypeReg	Resource Modeling type registry
ODIAppSrv	ODIAppSrv	Resource Modeling framework
VEMS	VECCwp	Virtualization Engine shared WebSphere Application Server profile
VESS	VESS	Virtualization Engine security
VEService	VEService	Virtualization Engine Updates
VECore	VECore	Virtualization Engine base support

5.2 About registration

The Virtualization Engine products need to be registered successfully with the IBM Autonomic Deployment Engine to ensure that they can be identified. The Autonomic Deployment Engine is a component of the Virtualization Engine that is installed when you install the Virtualization Engine. All Virtualization Engine products need to be registered with the Autonomic Deployment Engine to ensure that they can be identified and that they can receive and apply the fixes.

All products should be registered successfully during the installation process. However, problems can occur.

There are two types of messages that indicate that one or more products failed to register:

- A window, like the one shown in Figure 5-8, may appear when you run the Virtualization Engine installation wizard.

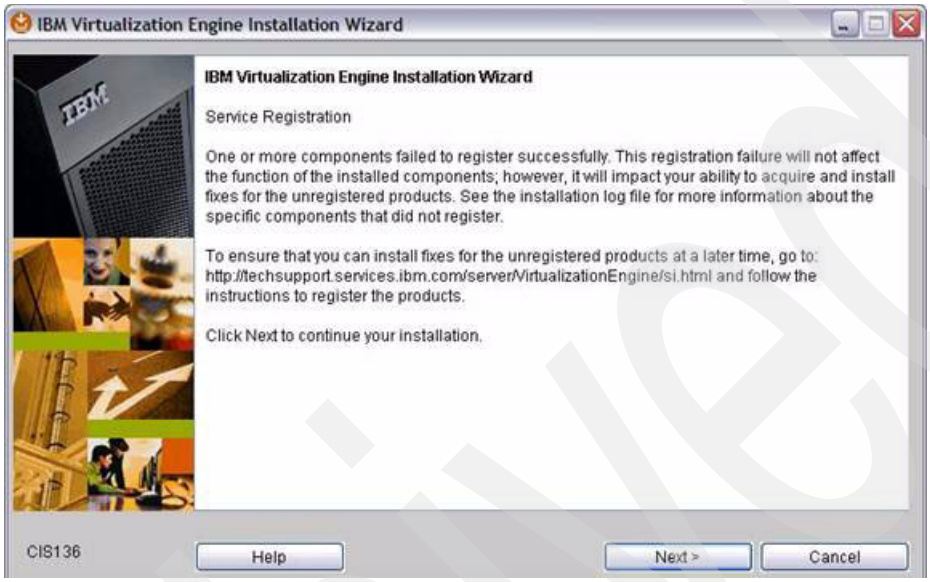


Figure 5-8 Service registration failure window

- You may see a message in the installation log, like the one shown in Example 5-2 on page 204.

Example 5-1 Service registration failure message

```
AVEINS006e An error occurred registering the product with SI
```

If you do not receive the window shown in Figure 5-8 or the AVEINS006e message in the installation log, then all the products registered successfully and you do not need to take any further action.

If you do get an error message, then you have to solve this problem before installing any fixes. You can register the products that failed to register during the installation by running the *Virtualization Engine Maintenance Environment Assistant* tool. You can download this tool from the Web at:

http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/si.html#_Download_and_Run

You have to download a package and unpack it into the directory where the Virtualization Engine has been installed. A new directory named srt will be created. To run the tool, you should go into the directory and run the **/srt.sh** command for AIX or Linux, or the **srt.bat** command for Windows. The results that are displayed are shown in Figure 5-9 on page 198.

Archived

```
Product Key: DB2
Registered: true
Generating response file...
Response file generated.
------------------------------------------------
Product Key: VECORE
Registered: true
Generating response file...
Response file generated.
------------------------------------------------
Product Key: VEMS
Registered: true
Generating response file...
Response file generated.
------------------------------------------------
Product Key: VECDir
Registered: true
Generating response file...
Response file generated.
------------------------------------------------
Product Key: ISC
Registered: true
Generating response file...
Response file generated.
velin94:/opt/IBM/VE2/srt/bin #
```

Figure 5-9 Results of the Virtualization Engine Maintenance Environment Assistant

For products that are already registered, the **Registered** value will be true. For products that need registration, this value will be false. In this case, the Virtualization Engine Maintenance Environment Assistant will attempt to register the non-registered products. If the product registration is successful, a message will be displayed. However, if the registration fails, an error message with further information about the cause of the problem is displayed. Table 5-1 on page 196 explains the content of the Product Key field.

If some products still fail to register, you should verify that the Autonomic Deployment Engine is active. If the Autonomic Deployment Engine is unavailable, you need to restart it using the **net start acsisrv** command (for a Windows environment). If the Virtualization Engine Maintenance Environment Assistant still fails to register after starting the Autonomic Deployment Engine, you may need to restart the Cloudscape™ server.

You can manually register a product as follows:

- On Windows, run the following command:

  `manageIUD -o create -r <discriminant> -i <responseFile> -force <forceFile> -p <packagedIUFile>`

- On AIX or Linux, run the following command:

  `manageIUD.sh -o create -r <discriminant> -i <responseFile> -force <forceFile> -p <packagedIUFile>`

In these commands:

- <*discriminant*> is the path name of the Virtualization Engine installation directory.
- <*responseFile*> is the path name of the response file to use for registration.
- <*forceFile*> is the path name of the IBM Virtualization Engine installation forcefile. This file is created in the <*Install_Location*>/META-INF directory after the Virtualization Engine Maintenance Environment Assistant runs and is named forceFile.properties.

- <*packagedIUFile*> is the path name of the packaged installable unit file for the product to register. This file is located in the <*Install_Location*>/META-INF directory, and has a name format similar to <*product-key*>_packagedIU.xml.

5.3 Installing and using the Update Installer tool

Before installing any updates, you must install the *Update Installer* tool. The Update Installer will help you install updates for the Virtualization Engine products and is necessary for whatever method you use to install the fixes, either through the Virtualization Engine console or using the command-line mode. A version of the Update Installer is installed when the Virtualization Engine is installed. However, we recommend that you download the latest version of this tool available on the Web at:

`http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/gmi.html`

At the time of the writing of this IBM Redbook, the current version of the Update Installer is 1.2.1.1.

The tool is available in zip format (for Windows) or in tar format (for AIX and Linux). The installation is very simple. Both formats can be used in an i5/OS environment. To install the tool, you should download the code and unarchive the package to the <*VE2 Install Path*>/gmi directory. The existing files belonging to an old version of the tool will be replaced and overridden.

To execute the tool, from <*VE2 Install Path directory*>/gmi, run the following command:

- On Linux or AIX: **# ./gmi_inst -javaHome <*VE2 Install Path*>/jre142/jre**
- On Windows: **> gmi_inst -javaHome <*VE2 Install Path*>\jre142\jre**

Figure 5-10 is an example of the results of the tool installation.



Figure 5-10 Installation of the Update Installer tool

Archived

Note: We recommend that you use the latest version of the tool. On Microsoft Windows systems, in our tests we experienced a bug with the version delivered with the Virtualization Engine. This bug does not show up when using the latest version available from the Web site.

5.4 Installing the updates

Fixes (or updates) can be downloaded in two ways:

- The first method is to download fixes manually, as explained already in 5.1.2, “The Web site” on page 191, from the fix Web site at:

  https://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/home2.html

  We recommend that you create a specific directory to receive the fixes, for example, /var/IBM/VE2/Updates.

- The second method is to download fixes using the Virtualization Engine console. This means the management servers, where the fixes have to be installed, must be able to connect to an IBM service provider to search for and download updates. A direct connection to the Internet may be unacceptable for some environments. You may have to configure a Hypertext Transfer Protocol (HTTP) proxy environment to reroute the traffic to the external service provider. You can do an HTTP proxy server configuration through the Virtualization Engine console, as detailed in the Information Center under the “Connection requirements” index.

5.4.1 Fix installation using the Virtualization Engine console

Once the update has been found and downloaded to the management server system, you can also use the Virtualization Engine console to install the update on the management server, or (for agent systems) launch IBM Director to distribute and install the update on the systems where agents, requiring an update, are installed.

You can only use the Virtualization Engine console to install a management server update on the server used to download the update. Remote installations are not supported. However, you can use the Virtualization Engine console to download an update, then manually copy the update to another management server, and finally install it using the Update Installer.

To access the Virtualization Engine updates user interface, complete the following steps:

1. Open the Virtualization Engine console.
2. Click **Hardware and software** in the navigation tree.
3. Click **Virtualization Engine updates** in the navigation tree.

Important: An update for the WebSphere Application Server or the Virtualization Engine console itself requires that the console is stopped and restarted. For these updates, the only way to be installed is to use the Update Installer tool. Updates to DB2 or IBM Tivoli Directory Server cannot be installed using the Virtualization Engine console.

5.4.2 Installing fixes without using the Virtualization Engine console

In some cases, it is not possible to use the Virtualization Engine console to install updates.

In this case, you can use the Update Installer tool instead. It can be used in three ways:

- Using Update Installer in graphical mode
- Using Update installer in console mode
- Using Update Installer in silent mode

Using the Update Installer in graphical mode

This procedure assumes that the updates already reside on the target management server. You should log in with an administator user profile, with root authorization for Linux, or administrator authorization for Windows. You should perform the following steps:

1. Open a command line window and navigate to the installation location for the Update Installer. For example:
   - On AIX and Linux: /usr/ibm/common/ci/gmi/bin
   - On Windows: C:\Program Files\IBM\Common\ci\gmi\bin

   Launch the Update Installer by entering the `gmi` command.The Update Installer displays a welcome screen.
2. Select **Install maintenance package**, such as fixes, Fix Packs, or refresh parts, and then click **Next** to proceed. The wizard will determine which Virtualization Engine products are installed on this system.
3. Select the component or subomponent to which the update applies, as shown in Figure 5-11, and then click **Next** to proceed.

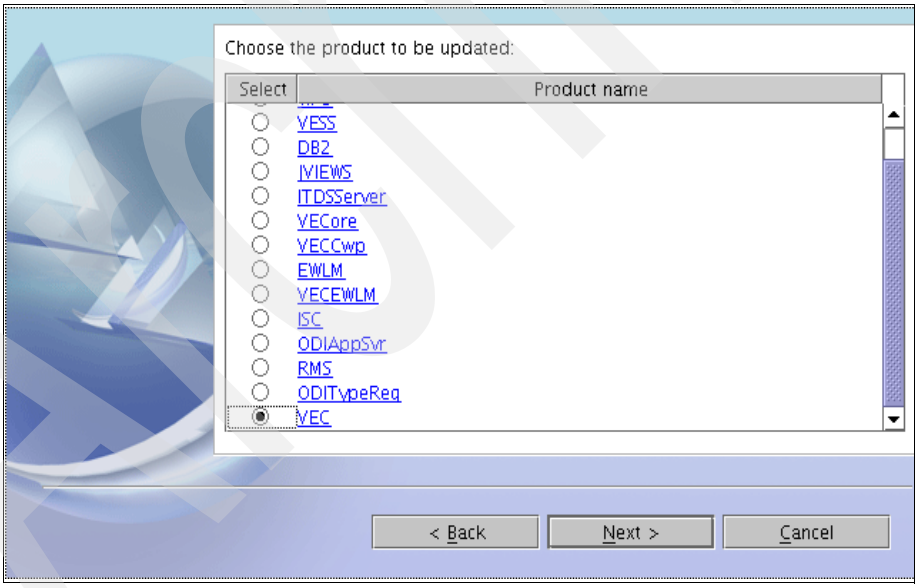


Figure 5-11 List of installed products

4. Complete the path to the update by adding the directory containing the fixes. Select the directory and click **Add**.
5. When the path to the update is displayed in the Paths to search box, click **Next** to proceed. The installation wizard searches the listed directories for any updates, and will display the results, as shown in Figure 5-12 on page 202.

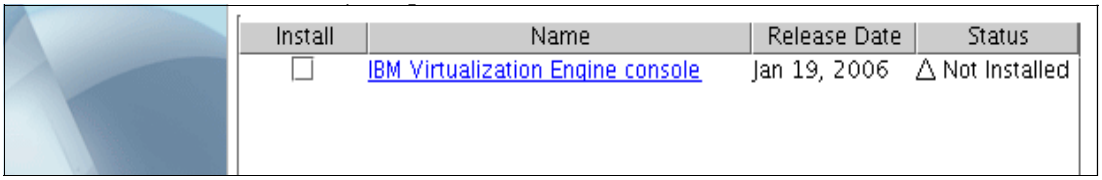


Figure 5-12 Search results for updates in the listed directories

Click the **Install** check box for the update you would like to install, and click **Next** to proceed.

You must perform this step for each of the components or subcomponents you want to update. Ensure that **Install maintenance on your computer** is selected and click **Next** to proceed. You can create a response file to allow additional installations to be performed silently. To do this, check **Save your settings in a response file** and specify a path for the new file.

6. Review the information in the installation summary window to ensure that all the settings are correct and click **Install**. The wizard begins to install the update.

   When the process is complete, the tool displays a summary, as shown in Figure 5-13. Click **Finish** to exit the wizard, or repeat the steps to install other updates.

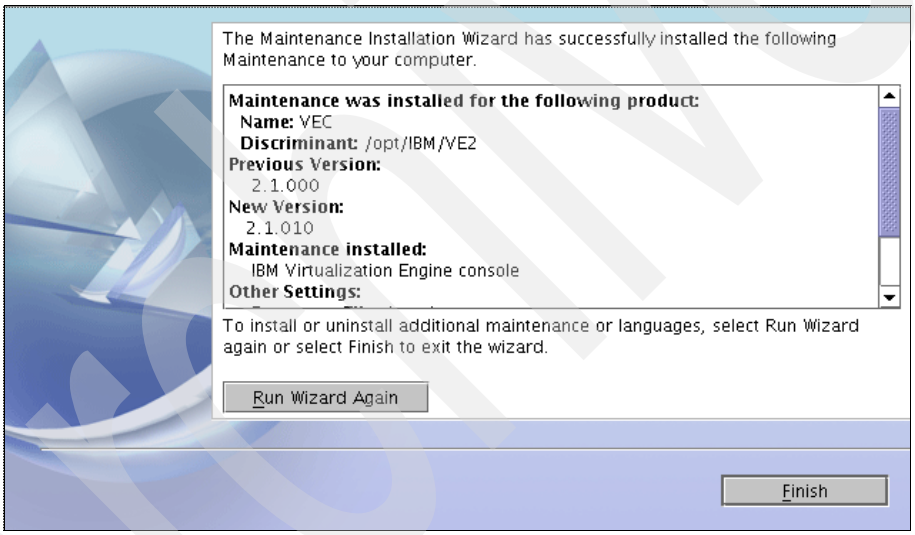


Figure 5-13 Update installation summary

Note: Installing the updates could take some time (several hours in our environment), so be patient and do not disturb the process.

For a Linux on System x platform, if after applying fixes for the Integrated Solutions Console (ISC) the service does not restart automatically, we suggest that you look in the /opt/IBM/VE2/p/VEISC/logs/ISC_Portal/ISC_Portal.pid file to check whether there is a PID for the ISC process. Otherwise, you can start the service manually with the ***<VE2 install path>*/ISC/PortalServer/startISC.sh ISC_Portal** command.

Using the Update Installer in console mode

This procedure assumes that the update currently resides on the target management server, and that the user has administrative authority on that system. You should perform the following steps:

1. Open a command-line window and navigate to the installation location for the Update Installer. For example:
   - On AIX and Linux: /usr/ibm/common/ci/gmi/bin
   - On Windows: C:\Program Files\IBM\Common\ci\gmi\bin

   Launch the Update Installer in console mode by entering the **`# gmi -console`** command.
2. Select **Next** (option 1) and press Enter at the Welcome window.
3. Select **Install maintenance package**, such as fixes, Fix Packs or refresh parts (option 1), and then enter 0 when finished. Select **Next** (option 1) and press Enter to proceed. The Update Installer V1.2.1 will determine which Virtualization Engine products are installed on your system.
4. Select the product to which the update applies by entering the corresponding number, then enter 0 when finished, and press Enter to proceed.
5. Select whether you would like to display details about the installed product, then select **Next** (option 1), and press Enter to proceed.

   If the path to the update is not currently displayed in the Paths to search list, complete the following steps to add it:
   a. Select **Add a path to this list** (option 1) and press Enter.
   b. Enter the path to the update and press Enter.
   c. Select **Yes** or **No** depending on whether you want to search the subdirectories of this path.
6. Confirm that the list of paths to search contains the path to the update by selecting **Use this list** (option 3). Then select **Next** (option 1) and press Enter to proceed. The Update Installer will search for any update in the specified locations, and will prompt you to select which one to install.
7. Select the update you want to install by entering the corresponding number, and press Enter.
8. Select whether you would like to display details about the update, then select **Next**, and press Enter to proceed.
9. Ensure that **Install maintenance on this computer** is selected and enter 0 when you have finished. Select **Next** (option 1) and press Enter to proceed. You can create a response file to allow additional installations to be performed silently. To do this, also select **Save your settings in a response file**.
10. Review the information about the installation summary window to ensure that all the settings are correct. Select **Next** (option 1) and press Enter to start the installation.
11. After the installation is complete, a summary will be displayed. You can indicate whether you would like to run the installation wizard again. Select **Yes** to install another update, or **No** to exit.

Using the Update Installer in silent mode

You can install updates on a management server silently (without requiring any user interaction during the installation). You can do this either by creating a response file, or by specifying any required arguments on the command to launch the installation.

Example 5-2 is a response file to install an update for the Virtualization Engine console; the update file is named UK10951 and was downloaded by the IBM Virtualization Engine Fix Web site. The path where the file is located is /var/IBM/VE2/Updates.

Example 5-2 Response file to install an update for the Virtualization Engine console

```
################################################################################
# Update Installer, version 1.2.1.1
# Response file created Mar 2, 2006 6:28:53 PM.
# For more information, see the Update Installer User's Guide.
################################################################################
# -uuid <uuid>
# Identifies the 32-character UUID of the offering that will be updated by Update
Installer.
# If more than one instance of the offering is installed, the -discriminant
# option is required.
# -discriminant <discriminant>
# Identifies the discriminant that uniquely identifies an installed instance
# of the offering to Solution Install.
#
# uuid 603A294ADE8D38CD6D44FCE5D1601EF2 = VEC (2.1.000)
# Offering Details:
#
#  Product:
#    VEC
# Version:
#    2.1.000
# Release Date:
#    Oct 10, 2005
# Discriminant:
#    /opt/IBM/VE2
# Installation Date:
#    Mar 1, 2006
# Features Installed:
#    VEC_F01
################################################################################
-uuid 603A294ADE8D38CD6D44FCE5D1601EF2
-discriminant "/opt/IBM/VE2"
################################################################################
# -mdvFile <mdv-file-list>
# Identifies one or more maintenance package, separated by
# spaces, to be installed or uninstalled by Update Installer.
#
# -mdvFile /var/IBM/VE2/Updates/UK10951.jar
# Maintenance Details:
#
#  Maintenance Package Name:
#    IBM Virtualization Engine console
#  Version:
#    2.1.010
#  Release Date:
#    Jan 19, 2006
#  Type:
#    FixPack
#  Applicable Products:
#    VEC
```

```
# ############################################################################
-mdvFile "/var/IBM/VE2/Updates/UK10951.jar"
```

The response file can be modified with the following parameters:

- *uuid* is unique identifier of the application.
- *discriminant* should be set to the installation path for the Virtualization Engine.
- *mdvFile* is the path to the archived update.

Once the response file is updated, you can install the update by running the Update Installer silently. Perform the following steps:

1. Open a command line and navigate to the installation location of the Update Installer. For example:
   - On AIX and Linux: /usr/ibm/common/ci/gmi/bin
   - On Windows: C:\Program Files\IBM\Common\ci\gmi\bin
2. Run the Update Installer with the -silent option of the **gmi** command. For example:

   `gmi -silent -options <path to response file>,`

   In this command, `<Path to response file>` is the full path to the response file.

You can also install the update silently without using a response file. If you do not want to create a response file, use the following command:

`gmi -discriminant <discrimantName> -mdvFile <mdvFileName> -silent.`

5.5 Installing fixes for specific products

Two of the products of the Virtualization Engine environment get fixes in a different way: DB2 and the IBM Director Server. Their specific processes are described in this section.

5.5.1 Installing the DB2 fix

If you have DB2 installed, then you may need to install the IBM DB2 security fix. Some DB2 Universal Database for Linux, UNIX, and Windows Version 8 (DB2 UDB) products contain a security vulnerability. DB2 V8.2 is provided with the Virtualization Engine. At the time of the writing of this IBM Redbook, Fix Pack 11 was available for DB2 UDB.

You can download the fixes and get instructions on how to apply them on the Web at:

http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/DB29a.html

The Web site instructions are very clear and are divided by platform. Before you install the DB2 updates, it is necessary to stop the directory server and the DB2 instance.

Note: If you have to remove the installed DB2 Fix Pack, *do not* execute the db2_deinstall script which is included in the unpacked directory. This script *removes* the entire database from the system.

Archived

Note: At the time of the writing of this IBM Redbook, we faced a problem after the DB2 Fix Pack installation finished. The ITDS service could not start anymore. To solve this problem, you should perform the following steps. Check that `[instance name]`, `[DBname]`, and `[adminID]` have the same value if the default was used during installation.

For Linux and AIX:

1. Stop the ITDS server, log into the system as a root user, and run the following commands:
   - **`# idsilist`** (This returns a list of all instances on your system.)
   - **`# ibmdiradm -I [instance name] -k`** (Repeat for all instances.)
   - **`# ibmslapd -I [instance name] -k`** (Repeat for all instances.)
2. De-configure the database with the **`# idsucfgdb -I [instance name]`** command.

   Use option 1 to erase the configuration, but leave the data. *Do not* choose option 2, which deletes the entire database.
3. Re-configure the database with the following command:

   ```
   # idscfgdb -I [instance name] -a [ITDS adminID] -w [ITDS adminID password]
   -t [DBname] -l [adminID home directory]
   ```
4. Restart the ITDS server with the following commands:
   - `# ibmdiradm -I [instance name]`
   - `# ibmslapd -I [instance name]`

For Windows:

1. Log in as a user with Administrator privileges.
2. Ensure that all ITDS services are stopped. Use the services Control Panel applet (select **Control Panel** → **Administrative Tools** → **Services**) and stop all services starting with “IBM Tivoli Directory Server Instance V6.0”.
3. Open a disk operating system (DOS) command line and run the **`> idsucfgdb -I [instance name]`** command.

   Select option 1 to erase the configuration but leave the data. *Do not* choose option 2, which deletes the entire database.

   The following command reconfigures the database:

   ```
   > idscfgdb -I [instance name] -a [ITDS adminID] -w [ITDS adminID password]
   -t [DBname] -l [DBlocation]
   ```

   (Here, `DBlocation` must be a drive letter.)
4. Start all ITDS services using the services Control Panel applet (select **Control Panel** → **Administrative Tools** → **Services**) and start all services starting with “IBM Tivoli Directory Server Instance V6.0”.

5.5.2 Installing the IBM Director fix

The procedure for the IBM Director Server and Agent updates is independent of the Virtualization Engine. You can find fixes for both IBM Director Server and Agent on the Web at:

https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=dmp

To access this Web site, you should enroll and use a valid user ID and password.

For details about how to install updates for the IBM Director, refer to *Implementing IBM Director 5.10*, SG24-6188.

5.6 Uninstalling fixes

To uninstall fixes:

- You can uninstall Virtualization Engine updates. This restores the component or subcomponent to the fix level it had before. You can use only the Update Installer to remove the most recently applied update.
- You can uninstall an update from the Virtualization Engine console by clicking the action button on an application listed in the Software tab. For more information, refer to the online help.
- Updates to IBM Director cannot be uninstalled.
- You cannot use the Virtualization Engine console to uninstall an update for the following products on the management server that is running the affected Virtualization Engine console:
  - WebSphere Application Server
  - Integrated Solutions Console
  - Virtualization Engine console
  - Any Virtualization Engine bridges

You can either use a Virtualization Engine console on a remote management server, or you can use the Update Installer. You can find technical details about how to remove a fix in the Information Center.

Archived

Archived

6

Lessons learned

Virtualization Engine is a powerful suite of products, all joined together and accessible through the Virtualization Engine console. To achieve these results, many components have to be installed and coexist on the same environment. The lessons learned from our laboratory experiences are listed in this chapter.

Importance of planning

- A detailed plan at the beginning is very important. All hardware and software prerequisites must be controlled and satisfied.
- Not having the minimum hardware requisites may dramatically affect the performance of the whole Virtualization Engine environment.
- Single sign-on preparation has to be defined.
- Both preinstallation and post-installation requisites must be understood and prepared carefully.

Retrieving information

- The Information Center is the primary repository of information. This is a very well-structured Web site where you can find very useful information. It is easy to browse through by using the index or the Search button.
- Read through the Information Center any time before and after any Virtualization Engine installation. It provides the latest technical informations about products and services, and includes indications about services usage.
- Use the planning advisor and the planning worksheets to design the environment before you start the installation. You cannot ignore hardware and software requirements.

Choosing a platform

- It does not matter on which platform the environment is deployed. Installation and processes are almost the same for all operating systems or platforms. You can choose a mix of operating systems and hardware.
- Platform and operating systems depend on various criteria, such as number of managed servers, availability of the environment, security framework, skills, and so on. Most importantly, you must respect hardware and software requirements. For example, if you

© Copyright IBM Corp. 2006. All rights reserved.

have to use Resource Dependency Service in your environment, you must install it on Linux.

- Clients can choose the operating platforms with which they are more confident and have the best expertise. Again, it is not possible to ignore the minimum hardware and software requisites.

Installation considerations

- If you want to simplify the installation process and avoid potential error and extra post-installation steps, we recommend that you install all the management server components on a single system.
- Installation of products can take some time (several hours in our case), but is fairly easy. The common components may take two or more hours to be installed as well. The same applies for the installation of updates.
- The time required is almost the same for all the platforms and depends mostly on how many hardware resources in terms of memory and CPU are on the servers and the number of servers to be managed.
- You can run the silent installation, using response files, as many times as necessary without consuming too much user time.
- Post-installation requires particular attention and should be well detailed.

Performance and usability considerations

- Before you install the Virtualization Engine components, think about performances, stability, and recovery solutions.
- From a performance point of view, you should have a great deal of physical memory and processing power. We recommend that you add more than the planning advisor recommends. You should remember that the paging space should be twice the size of the physical memory size. We recommend that you place the Virtualization Engine code on a fast external storage. It will be helpful during the installation process.
- From a disaster recovery point of view, you must remember to create a backup before installing the updates. On an AIX system, we recommend that you install the operating system from a Network Installation Management (NIM) server. It may help you to restore a clean operating system image. You should also create an mksysb before installing the Virtualization Engine components.
- We recommend that you change some operating system parameters as described in (refer to 3.6, “Post-installation checking tasks” on page 128), where you will find specific changes suggested for each operating system platform.
- Be careful about the operating system user IDs, both automatically created during installation or manually added. There are many of them, so it is important to have all them clearly listed with the password secured in a safe place.

Maintenance of the environment

- You must maintain the Virtualization Engine environment carefully. A Virtualization Engine environment joins together many products, and the maintenance of those products can be done separately.
- Use the Virtualization Engine console for the maintenance. It allows you to have a direct view of all the installed components in the environment.
- Create more than one operator profile to manage single tasks.
- If any problems arise, have a look at the log files. There are many log files and messages may be stored in the system log or in any specific log for each service or bridge.

- We recommend that you always be at the last patch level. These patches are delivered to address known problems, and may also introduce new functionality on the products. The readme file with that come with each patch contains all the information about what a patch modifies.

Archived

Archived

Part 2

Using the Virtualization Engine Version 2.1

In this part, we describe the following scenarios for using the Virtualization Engine products.

- Chapter 7, “Using Resource Dependency Service to build manageable resources” on page 215 discusses how to build manageable resources using Resource Dependency Service.
- Chapter 8, “Using the Virtualization Engine console to visualize the infrastructure environment” on page 265 discusses how to visualize the manageable resources using the Virtualization Engine console.
- Chapter 9, “Using one single console to manage your environment” on page 301 discusses how the IBM Director and the Virtualization Engine console complement each other.
- Chapter 10, “Using the Enterprise Workload Manager” on page 379 discusses several situations in which the Enterprise Workload Manager automates the infrastructure management.

© Copyright IBM Corp. 2006. All rights reserved.

Archived

7

Using Resource Dependency Service to build manageable resources

Manageable resources are the key component to build a service-oriented infrastructure. The Resource Dependency Service product allows the discovery and collection of the resource properties to build manageable resources.

This chapter:

- Introduces the key Resource Dependency Service concepts in 7.1, “Resource Dependency Service concepts” on page 216.
- Describes how to configure Resource Dependency Service in 7.2, “Configuring Resource Dependency Service” on page 224.
- Describes how to use the Intelligent Device Recovery in 7.3, “Planning and running IDD” on page 226.
- Describes how to configure and use Resource Dependency Service data collectors in 7.4, “Planning and configuration of Resource Dependency Service data collectors” on page 243.
- Provides some information about how to manage Resource Dependency Service in 7.5, “Managing Resource Dependency Service” on page 257 and 7.6, “Troubleshooting” on page 263.

© Copyright IBM Corp. 2006. All rights reserved.

7.1 Resource Dependency Service concepts

Resource Dependency Service focuses on the discovery and mining of information from existing applications, tools, and IT systems of your heterogeneous IT environment. It then displays a view of these assets through the Virtualization Engine console. With Resource Dependency Service, you can examine your current IT environment for assets and correlate them to their business functions. The overall goal of Resource Dependency Service is to create an environment where the elements are abstracted into virtualized resources, which can be managed and accessed.

Resource Dependency Service is the first stage of developing a representation of the resources and the relationships between components so that they can become on demand services.

Resource Dependency Service is a key component on the roadmap to using WSDM for managing systems. Automatic discovery of the resources and the relationships between them is crucial in large and heterogeneous environments where it is hard or nearly impossible to have an overview of all resources. The representation of this information has to be in a consistent and standardized way.

A virtualized environment has the strong requirement to understand and manage resources as well as interactions and dependencies among resources in a uniform way. The uniformity allows components in a virtualized environment to be plugged in and to be used according to the uniform semantics, without having to understand specialized semantics for each function.

Resource models provide the theoretical basis for modelling the components and dependencies among them. A resource model is an architecture intended to provide the means by which IT resources and the management approaches applied to them can be expressed in a common and standardized way.

Resource models today are tied closely to our major stovepipe applications. Existing hardware and software products have inconsistent provisioning or deprovisioning procedures and proprietary and incompatible management (monitoring, configuration, and so on) application programming interfaces (APIs). Dependencies among products are often implied and difficult to trace. As a result, integrated solutions combining multiple products are difficult and often impractical to build and maintain, and impossible to create automatically.

Standardization of the resource models is key in getting a virtualized environment. Services that have to be managed must be a Manageable Resource, which is a WS-Resource Framework (WS-RF) compliant component. A manageable resource has manageability interfaces, which are sets of properties, operations, events, and meta-information described by Web Services Description Language (WSDL), Resource Properties, Manageable Resource Prototype, and Policy documents.

The basic steps towards the creation of a manageable resource is depicted in Figure 7-1 on page 217. Several discovery technologies for IT assets, applications, and database federation techniques provide the resource information, which Resource Dependency Service analyzes, correlates, and transforms into the resource model. The resource data and its respective metadata are stored in a Resource Dependency Service datastore. In a next step, the resources are exposed as manageable resource through Web Services and are available for management application. One of the exploiters of these exposed Web Services is the topology view with the Virtualization Engine console.

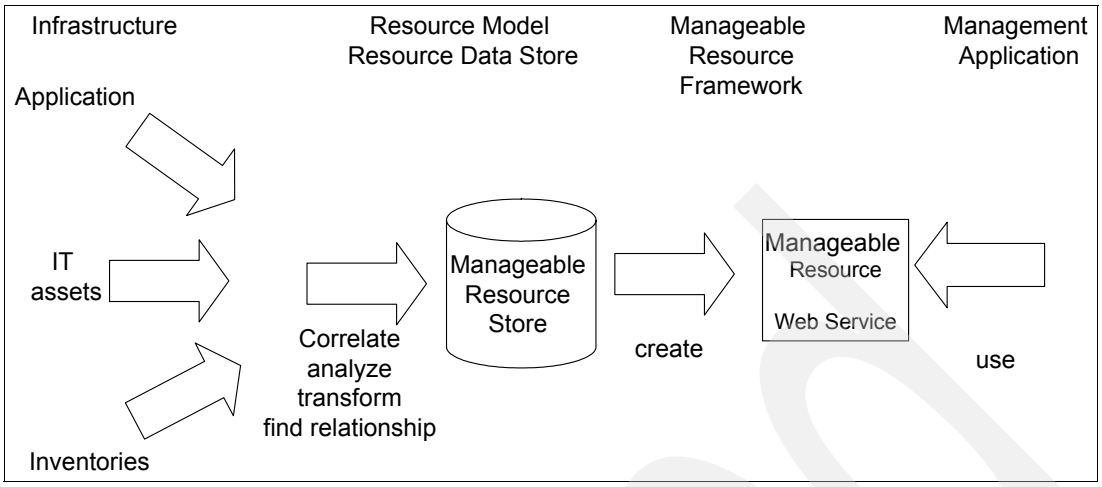


Figure 7-1 Steps toward a manageable resource

Resource Dependency Service helps you to assess your current environment and the connections to the business functions. With this information an impact analysis of planned and unplanned of resource outages is possible. In addition, the link to the business functions allows you to assess the potential financial costs associated with the resource outage. In this sense, Resource Dependency Service is a major building block for a number of management disciplines. Management disciplines leveraging Resource Dependency Service include business resilience, workload management, capacity planning, or provisioning.

7.1.1 Resource Dependency Service components overview

Resource Dependency Service consists of several functional units that cover the discovery, collection, mining, federation, correlation, and analysis of resource information within your environment. Figure 7-2 on page 218 depicts the basic components of Resource Dependency Service.

- The Resource Dependency Service server component is a WebSphere application. It provides a common set of components developed to integrate and represent existing IT resources and discover existing relationships among them. The application runs within the VEMS shared profile of the WebSphere installation.

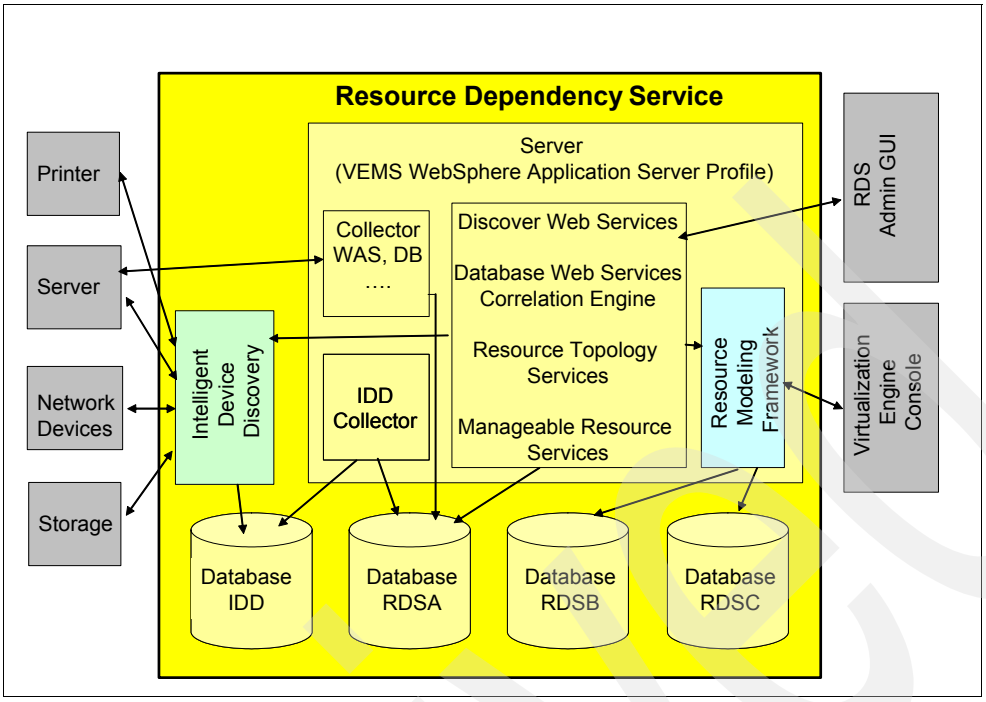


Figure 7-2 Resource Dependency Service component overview

- The collected data is stored within the datastore, which is a DB2 database called RDSA. The Resource Dependency Service configuration graphical user interface (GUI) allows a view of the contents of this datastore.
- The Intelligent Device Discovery (IDD) is an IBM developed tool used to discover IT resources in your environment. IBM Strategic Outsourcing uses this tool in a stand-alone version in projects to analyze client environments to better understand the environment. The IDD server runs as a stand-alone server that starts at boot time. IDD utilizes WebSphere applications that run in the default WebSphere server profile.
- The resource visualization within the Virtualization Engine console utilizes manageable resource services in conjunction with the Resource Modelling Framework (RMF™). Currently, the RMF APIs are not exposed to external clients. The Resource Modelling Framework has a separate database (RDSB and RDSC) and contains all WSDL documents of the manageable resources and its properties.

The following sections provide a more detailed view about these Resource Dependency Service components.

7.1.2 Understanding the Intelligent Device Discovery tool

Here we discuss the Intelligent Device Discovery tool.

Intelligent Device Discovery

In today's fast changing IT environments, the tracking of computing devices as assets and their usage becomes more and more critical. This gets more important in virtualized environments.

The IDD tool gathers information about network-attached devices and their categorization. It integrates various established network discovery and fingerprinting programs and integrates their collected information into a single database.

The IDD tool supports:

- The discovery of Internet Protocol (IP) interfaces and attached systems
- The discovery of hardware information for servers and workstations
- The discovery of operation system information
- The discovery of installed applications

An overview of the IDD integration into Resource Dependency Service is given in Figure 7-3. IDD runs in its own server environment with a database. The IDD database is a DB2 instance, called VEDB2UR in this case (IDD and Resource Dependency Service use the same DB2 instance, which is created as part of the Virtualization Engine installation process). The IDD collector within the Resource Dependency Service server performs the transformation into the Resource Dependency Service data model, correlation, and identification of the resource relationships. This process is covered in the “Description of Resource Dependency Service Collectors” on page 221.

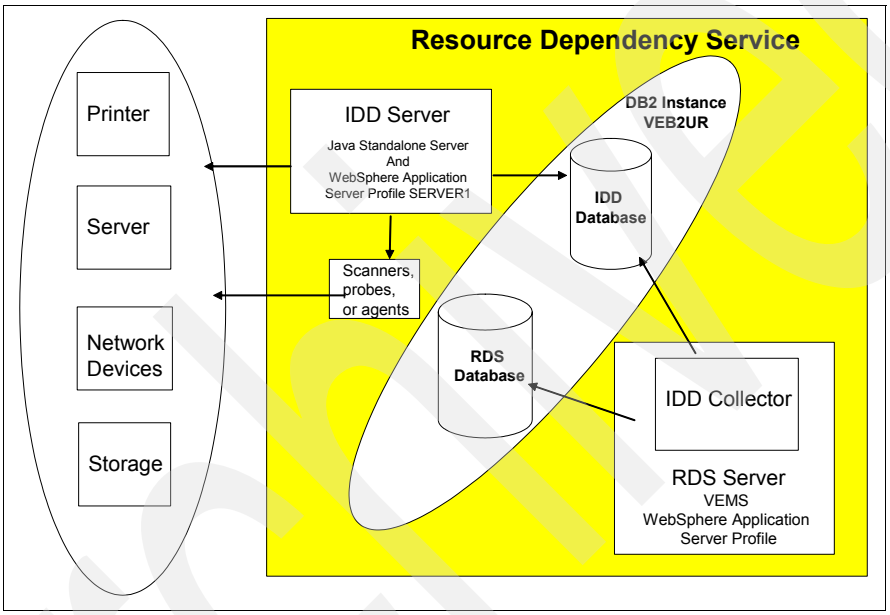


Figure 7-3 IDD Resource Dependency Service integration

IDD technologies

There are generally three methods for collecting data in a network environment:

- Discovering with scanners, without agents installed on the resource. Discovery with scanners is the least intrusive technique. No code is installed or executed on the end points. The scanners can passively listen and map the network, walk through the network, or interact and fingerprint individual hosts and applications.
- Discovering with probes means that software is installed at runtime on the end points and only at the time of the discovery. This software gathers data and pushes it out of the endpoint and then deletes itself.
- Discovering with agents is the most intrusive discovery technique. Agents run permanently on the end points. They can be installed and configured either manually or automatically. The agents gather and push the information to the management server.

IDD offers configurable levels of resource scanning processes with increasingly detailed discovery.

- The *Level 1* scanning is an agentless, non-intrusive scanning technique without using a credential or using IP, port, and wireless intelligent network (WIN) scanners. This scan discovers or identifies all Transmission Control Protocol/Internet Protocol (TCP/IP) attached devices and is able to indicate some general information about the system, such as IP, Domain Name System (DNS) host name (if resolved), Mac address (primarily Windows), computer name (primarily Windows), operating system, service bindings that exist, and active ports. The scan results are improved if an NMAP[1] program is installed on the IDD server. For this scan, specify a scan range and a list of servers that you want to exclude from the scan.
- The *Level 1+* scanning is typically used to enhance what is found in a Level 1 scan. A Level 1+ collection can run as a stand-alone discovery. The scanning techniques include Remote Windows Management Instrumentation (WMI) queries, Remote Registry Reads, and Remote SSH scripting on UNIX. Detailed hardware, software, application, and configuration information may be collected. WMI must be enabled to gather detailed hardware information about Windows. Data collected with this level includes HW configuration information, model, processors, speed, memory, hard disks, and installed software packages.
- The *Level 2* collection technique can be used as an alternative to the Level 1+ detailed data collection method. Currently, IDD Level 2 is not supported within Resource Dependency Service. It is only a feature of the IBM Global Services IDD offering. A script/executable is distributed to target systems via a client-approved method (login script, push technology, IDD push technology, e-mail agents, and so on). Script executed locally produces eXtensible Markup Language (XML) based data, which is transferred to a remote endpoint in the network (IDD system or a specified share point). Data is processed by the IDD system and the database is populated.

The task of collecting data that RDS can use is called a *campaign*.

Note: Scanning of networks and devices is not completely without risk. IDD has scanned more than 150,000,000 IPs since 2002. Only 21 problems are reported. All problems reported to date have been related to defective code or unmaintained code on the devices that caused them problems when they were being scanned.

If you have *Mission Critical* devices or subnets, you can exclude these devices from the scanning.

IDD database

IDD uses a relational database to store the data. The IDD database schema contains a general configuration and history data. The discovered and mined data are stored in tables belonging to another database schema with the name of the campaign. For each new campaign, a new database schema is created.

The IDD discovery process has two steps:

1. The scanned resource data is stored in database tables.
2. A combiner process combines the data and transfers the results in tables.

[1] Nmap ("Network Mapper") is a free open source utility for network exploration or security auditing that can downloaded from http://www.insecure.org/nmap/download.html.

7.1.3 Understanding the Resource Model Services

Here we discuss the Resource Model Services.

Resource Model Services

Resource Model Services is the server component within Resource Dependency Service, which hosts particular collectors and control functions for the resource discovery. Mining and correlation of the data is also part of Resource Dependency Service. The datastore stores all resource information in the main storage for all correlated retrieved information. This database contains the relationships between data, which is necessary for further processing of the Resource Modelling framework.

Description of Resource Dependency Service Collectors

Resource Dependency Service Data Collectors provide a mechanism to get data from systems and workloads coming from multiple data sources. These sources include IDD, Tivoli Business Systems Management (TBSM), Tivoli Configuration Manager (TCM), IBM Director, Enterprise Workload Manager (EWLM), WebSphere Application Server, WebSphere MQ, DB2 UDB, IBM System i, and IBM System z. The collected data is stored in the RDSA database. The collectors also attempt to make relationships between resources, when they are placed into the datastore. The collector process flow involves the following steps:

1. Connect to external repository using JDBC™ or application-specific interfaces.
2. Discover Resource Dependency Service specific information.
3. Analyze, transform, and correlate the data.
4. Store the appropriate information into the database.
5. Start the correlation engine to check if data is redundant.
6. Start the creation of a manageable resource if sufficient information is available.

Every time a collector has finished, the correlation engine starts. The correlation engine combines rows in the database and synchronizes the data between Resource Dependency Service and the Resource Modelling Framework database.

The correlation engine uses rules read from two XML files:

- One XML file keeps the set of previously correlated resources so that when an object is discovered a second time, it can be correlated more quickly.
- The second XML file is used to maintain the correlation values for each model attribute in the database. When two attributes match, the corresponding correlation value from the XML file is used. When a specified threshold is met, the two objects can be combined. The attribute correlation values and the threshold for each resource type still has to be determined.

When two objects can be combined, the correlation engine updates one object with attributes from the second and the second object is removed from the datastore. The correlation engine saves a new rule into the first XML file mentioned previously.

Description of the database

The Resource Dependency Service database stores all the collected resource data from all the Resource Dependency Service data collectors. Each resource type has its own table. In addition, there are tables for the resource model metadata. Resource Dependency Service specific application data is also stored in this database.

Archive

7.1.4 Understanding the Resource Modelling Framework

Here we discuss the Resource Modelling Framework.

Resource Modelling Framework

In a typical virtualized environment, a large number of different kinds of resources are managed. For example, the domain is comprised of computer resources, operating system resources, or storage resources, which in turn can be categorized into different types. Thus, we recommend an open resource model and open standard technologies, such as WS-RF, WS-Addressing, WSDM, and CIM to manage the heterogeneous environment. A common framework simplifies the building of virtualized systems management components. Within Virtualization Engine Version 2, the Resource Modelling Framework provides this framework.

The framework is comprised of many features. It includes services that may be used as helper classes. These may be extended or used as utilities within specific domains (for example, helpers for generating identity or testing equality), reusable components that do not stand on their own and have to be incorporated as part of a larger solution (for example, connectors to CIM Servers), as well as patterns for exploitation of the various framework components. In this framework, all manageable resources are modelled as stateful services according the WS-RF standard.

The Virtualization Engine console topology view shown in Figure 7-4 is an application that uses the Resource Modelling Framework. More details about how to use the Topology Service is described in Chapter 8, “Using the Virtualization Engine console to visualize the infrastructure environment” on page 265. Figure 7-4 presents a list of all the available resource types. The resources of each type are manageable resources, which must conform to the specification of the underlying resource model.

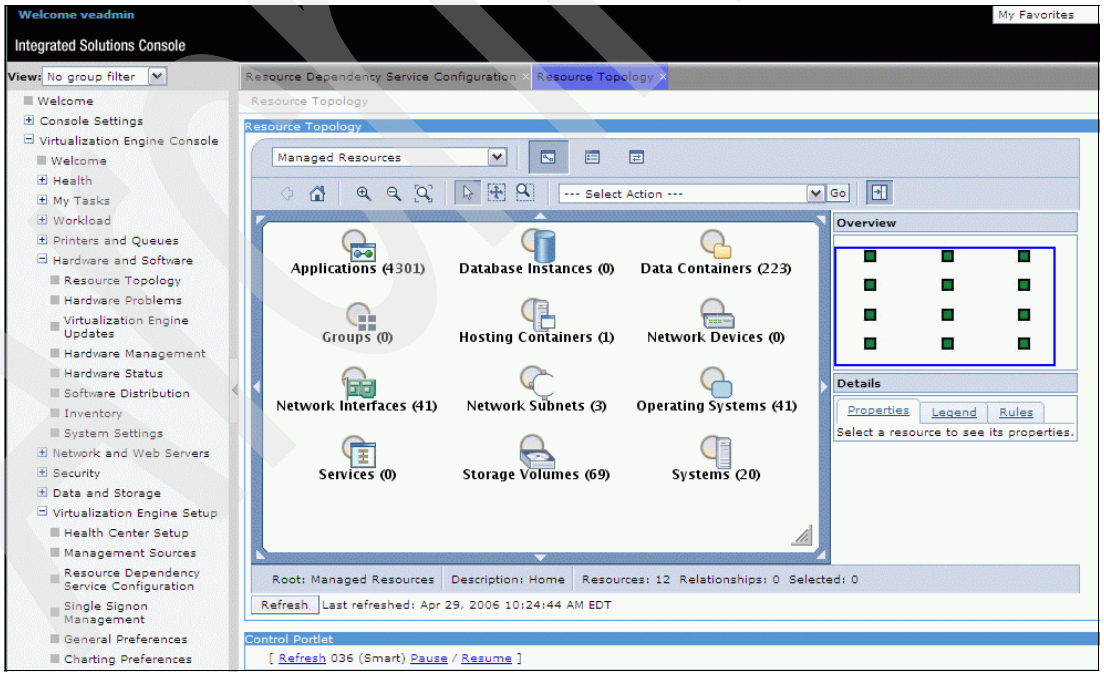


Figure 7-4 Virtualization Engine topology view

A resource in this context has a management interface with the operations that control the state and function of the resource. Every managed resource must implement the ResourceIdentification portType, which is a read-only WS-Resource Property called *resourceID*. The resourceID is globally unique, unique in time, consistent across and within

management endpoints, and persistent and opaque to the client. The resourceIDs are used to test the equality of two WS-Resources. If two resourceIDs are equal, then the two manageability endpoints represent the same resource. As an example for a system resource, the resource model requires a complete and valid set of values for the *UUID* or *MachineType*, the *Model*, the *SerialNumber* and, optionally, the *LPARid*. If any information is missing, the resource creation will fail.

Because a managed resource is a WS-Resource, its state is expressed through resource properties. Web service operations may be invoked upon the WS-Resource to monitor the state, effect the state change, or to subscribe for a state change notification. Therefore, a managed resource must implement at least the GetResourceProperty interface.

For each resource, there is a group of three Web Services: Resource Factory, Resource Registry, and Manageable Resource, as shown in Figure 7-5.

- The Resource Factory service is used to create new resources of that type based on given creation input parameters. As a result, the client gets the End Point Reference (EPR) for the newly created resource. The client must provide all the information that is necessary to create the resource. The resource factory stores the resource data in the Resource Dependency Service datastores.
- The Resource Registry service is used to query for existing manageable resources in the Resource Dependency Service datastores. Search criteria can be passed in to narrow down the resources you are searching for. The EPRs of the matching resources is returned.
- There is a Manageable Resource Service for each type. This service can be used to query and set any property or relationship of the resource.

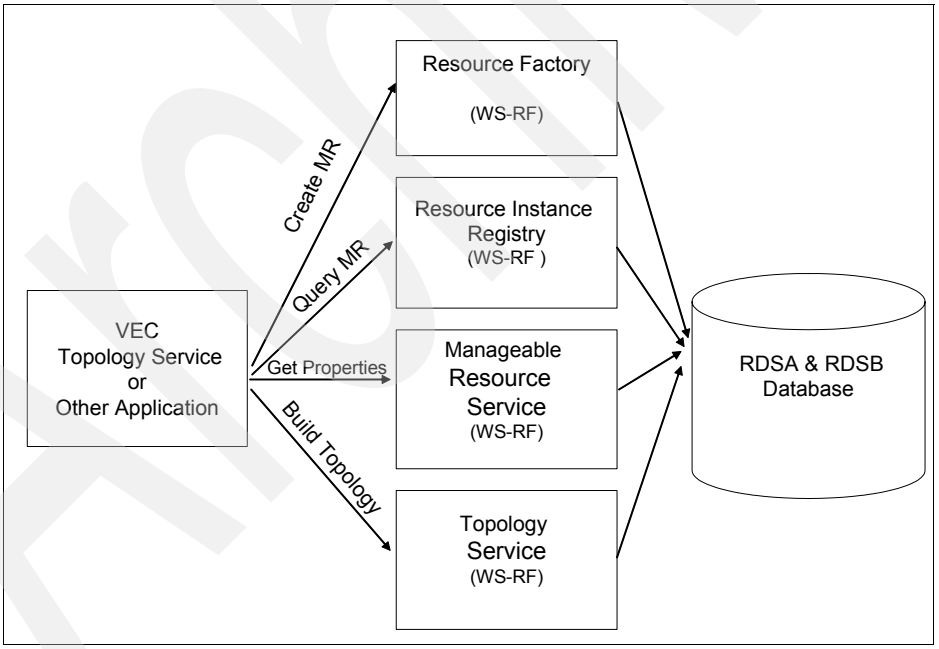


Figure 7-5 Services involved in the creation of a manageable resource

The Topology Services are a separate component. Resource topology services are based on the WS-RF stateful service component of the framework. Clients can access topology information using the standard WS-RF Web services. A topology service client can use the topology service by asking for the topology related to a resource. That is, clients call a static Web service method, providing a starting End Point Reference (EPR). When created, a topology service builds a representation of the topology for the resource and returns an EPR

to the client. Clients can build a graphical representation of the topology by getting the topology information from the topology resource.

Description of the database

The RDSB database stores the resource information and their relationship information for all manageable resources. The tables can be large in case of large environments.

Note: All three databases (IDD, RDSA, and RDSB) are used by different components of Resource Dependency Service, which gather the data with different tools and methods. Therefore, some of the data is stored in all of these databases, for example, the resource properties. Only the data in the ODIRM tables is visible in the topology view.

7.2 Configuring Resource Dependency Service

Before using Resource Dependency Service, you have to configure it. To select the configuration task, in the Virtualization Engine console navigation window, select **Virtualization Engine Setup** → **Resource Dependency Services Configuration**. Figure 7-6 shows the window that appears.

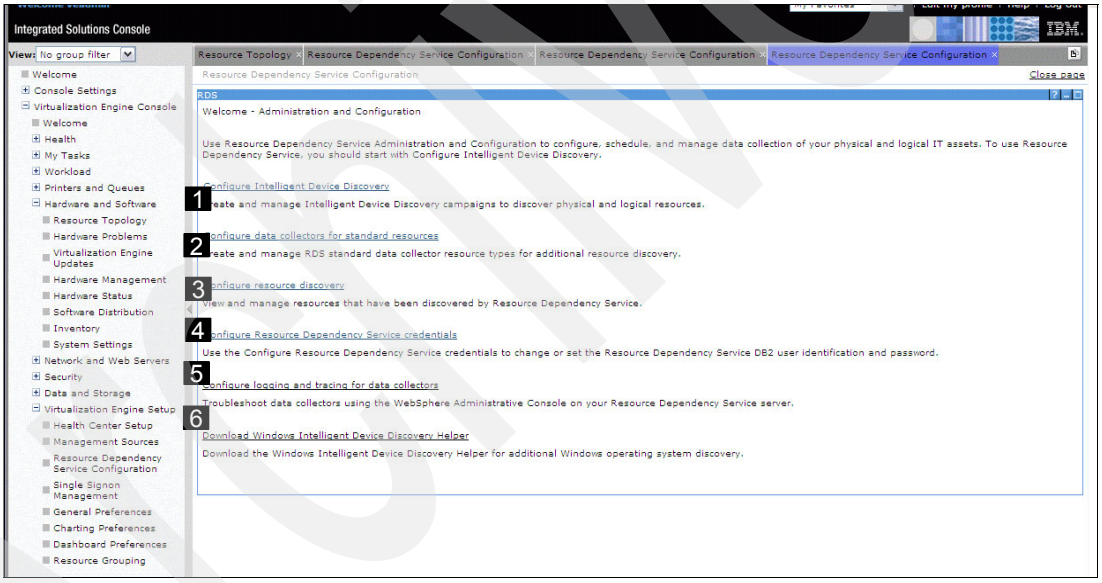


Figure 7-6 Resource Dependency Service configuration overview screen

You can select from the following configuration tasks:

1. The Intelligent Device Discovery configuration is described in “Configuring an IDD campaign without credentials” on page 229 and “IDD campaign configuration with credentials” on page 237.
2. The data collectors for standard resources configuration is described in “Configuring the IBM Director collector” on page 253.
3. The resource discovery configuration is described in 7.4.2, “Configuring a new collector” on page 245.
4. You can change, if you want, the user ID/password defined at the installation process after the Resource Dependency Service installation.
   a. Select **Configure Resource Dependency Service credentials** (see Figure 7-6).

b. The window shown in Figure 7-7 opens. Specify the Resource Dependency Service user ID (see 1) and password for the Resource Dependency Service database.

   The field Password set (2) shows status Unknown when no password is provided. To provide a password, click **Password** (3).

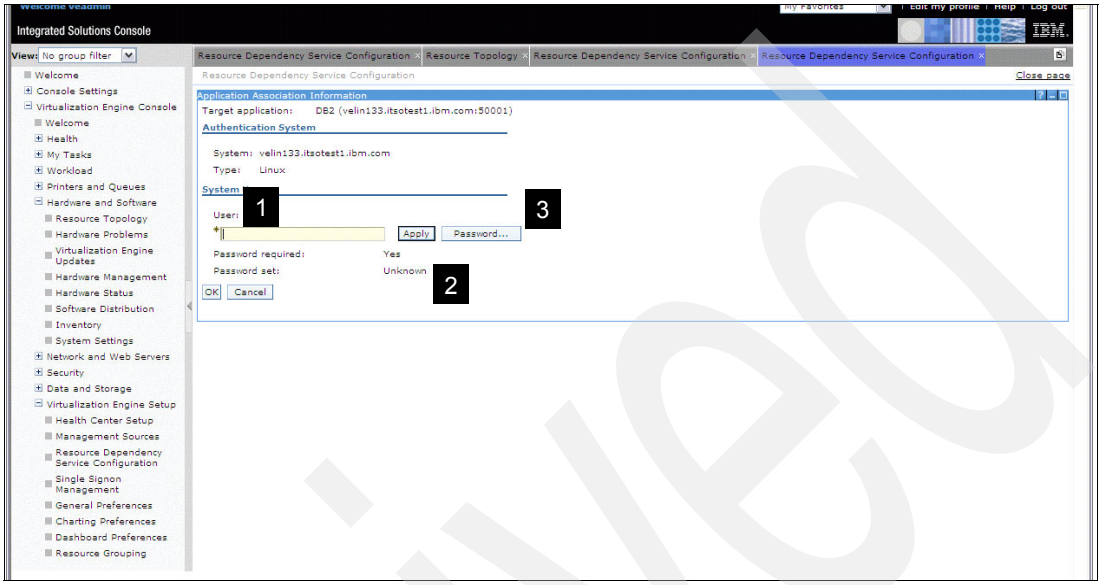


Figure 7-7 Configuring the Resource Dependency Service credentials

The Resource Dependency Service set password window opens, as shown in Figure 7-8. In the **Password** field, type in the password, and confirm it in the Confirm password field. Click **OK** to finish the task.

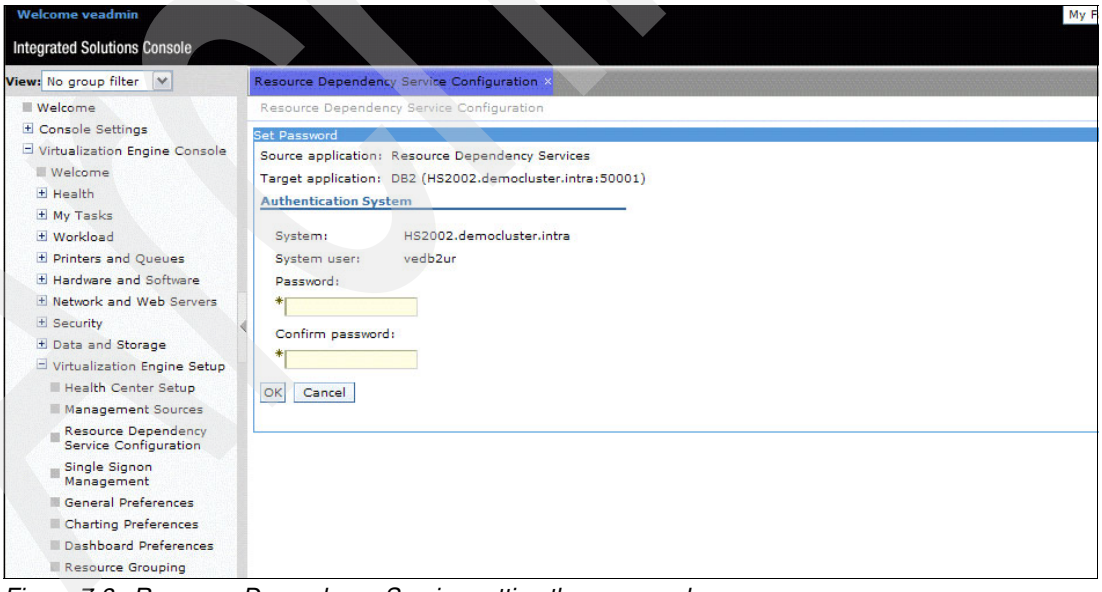


Figure 7-8 Resource Dependency Service setting the password

5. Select **Configure logging and tracing data collector** to start the WebSphere Administration Console for the VEMS profile. In the navigation pane in the left of the administration console, select **Troubleshooting** → **Logs and Trace** and the Logging and Tracing window is displayed, as shown in Figure 7-9 on page 226. Select the server **VEMS**.

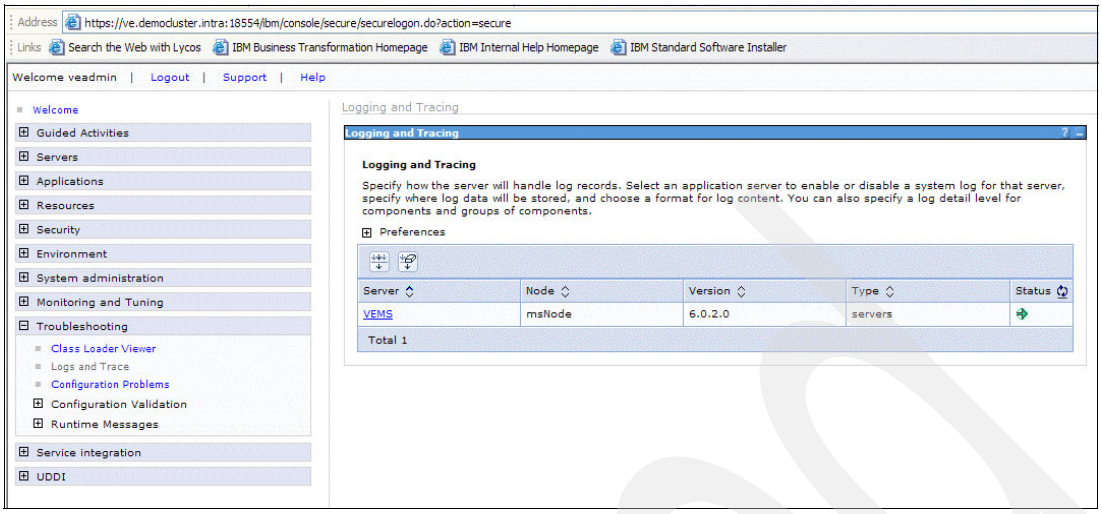


Figure 7-9 WebSphere Administration Console Logging and Tracing

Select **Change Log Detail Level** (see Figure 7-10) to specify the log level of the VEMS components. For more details about how to configure the log levels, refer to the WebSphere online manual.

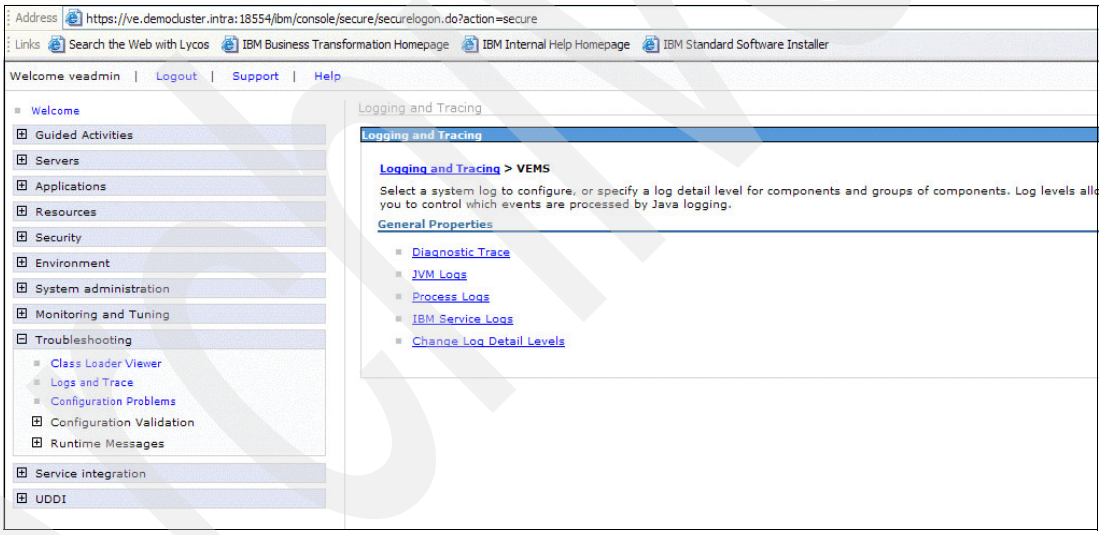


Figure 7-10 Logging and Tracing for VEMS

6. Download Windows Intelligent Device Discovery Helper.

7.3 Planning and running IDD

IDD gathers information about network-attached devices. The gathering of the information, called an IDD campaign, needs to be customized.

7.3.1 Planning the IDD campaign

Before running an IDD campaign in your environment, a proper planning of the campaign is necessary.

Before starting the first test campaigns, we recommend that you contact your network security team: the IDD tool uses several techniques for discovery, some of which may be noticed by the network security teams. While these techniques are harmless and non-intrusive, unwarned security teams may perceive IDD's probes as attempting to hack or intrude into their systems.

Be aware that you may see the following activities during an IDD campaign, though they are part of a normal operation:

- Repeated failed login attempt: An IDD discovery module attempts a *Telnet Login* into each address in the IP ranges. No user ID or password is provided, and there is no attempt to login. The standard *Telnet Prompt Response* is captured by IDD, which attempts to extract the host name and operating system.
- IDD discovers information about a device by checking specific TCP/IP ports. The port usage is different on each platform and operating system. An empty network packet may respond with some information that is used to decode attributes of the system or device. This technique is harmless and is in no way intruding into the device, but intrusion detection techniques and network monitoring systems can detect it.

The planning of IDD campaigns is comprised of the following steps:

- Review the network topology.
  - Network ranges and subnet masks
  - Exclusions from scanning such as firewalls and routers
- Identify secured network segments.
- Determine optimal time and duration.
- Obtain information for scans.
  - For Level 1+ scans
    - Primary NT Domain user account with local administrator authority
    - Simple Network Management Protocol (SNMP) (read-only) community strings
    - Secure SHell (SSH) credentials
  - Optional configuration or enabling SSH, SNMP, and WMI on endpoints, if access to endpoints is required for detailed endpoint information.

Typically the first campaign is a Level 1 scan, that is, a scan without specified credentials (see "Configuring an IDD campaign without credentials" on page 229). This scan is the basis for further detailed Level 1+ campaigns.

7.3.2 ITSO environment

The ITSO environment is made of a BladeCenter with HS20 and JS20 Blades, a p550 with eight logical partitions (LPARS) and branch office desktop computers.

Figure 7-11 on page 228 shows all the available systems in the ITSO project environment. The IBM Director discovery scan, without using SNMP, Common Information Model (CIM), or the Director specific agents, provide all the IP resources within the environment.

Archived

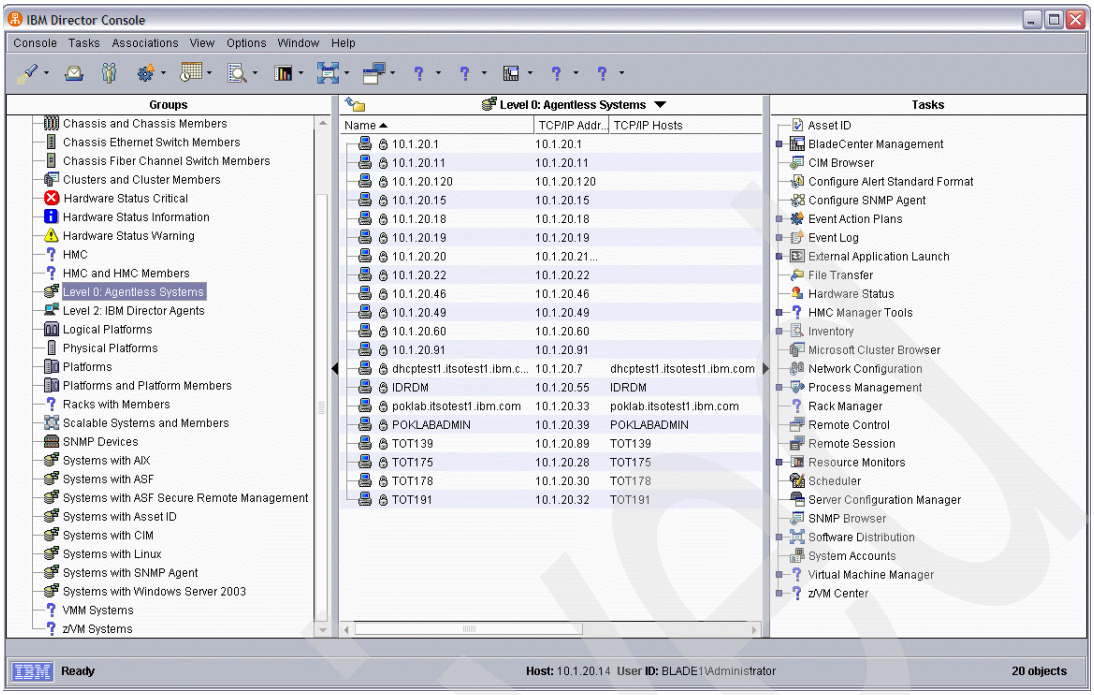


Figure 7-11 ITSO environment IBM Director Level 0 Agentless Systems

Figure 7-12 shows more details about the system in the project environment. The IBM Director scan with the Director agents shows that we use AIX, Linux, and Windows as operating systems on the servers.

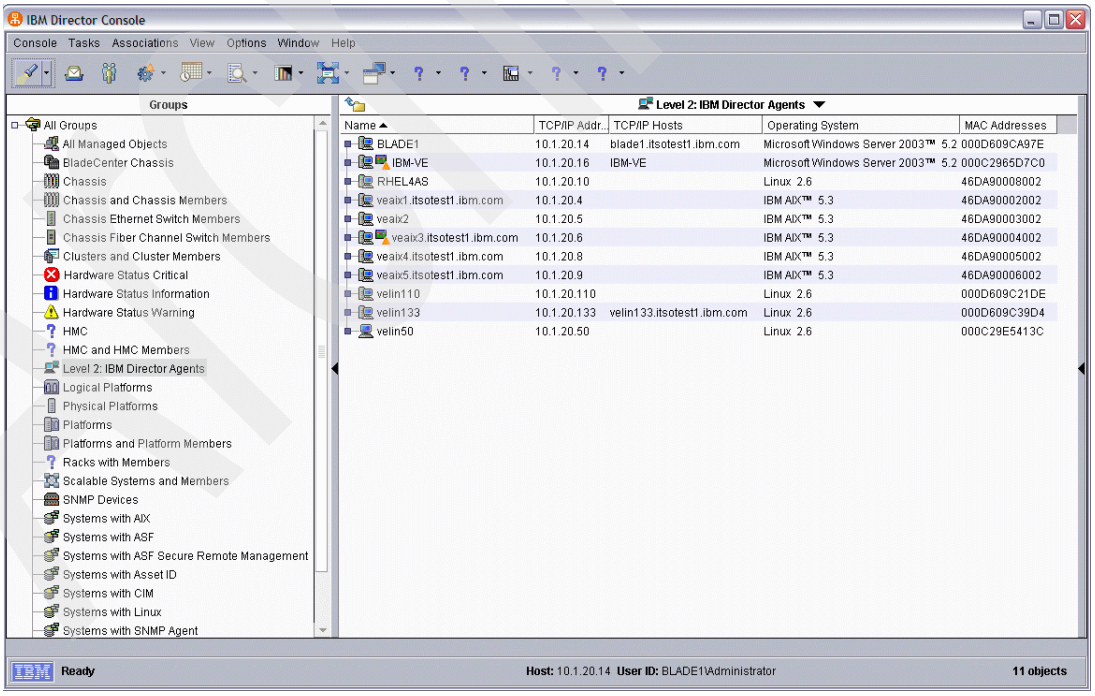


Figure 7-12 ITSO environment IBM Director Level 2 Director Agents

7.3.3 Running an IDD campaign

In this section, we describe two types of IDD campaigns and we compare the potential results:

- An IDD campaign without credentials
- An IDD campaign with credentials

Configuring an IDD campaign without credentials

An IDD scan requires the IP addresses of the devices or the address range of the network as input. Hosts that should not be discovered are specified in an exclude list. To configure the IDD campaign, from the Virtualization Engine console navigation pane, select **Resource Dependency Service Configuration → Configure Intelligent Device Discover Configuration**. A list of already configured campaigns opens.

The installation procedure creates the BASELINE (A in Figure 7-13) campaign. To create a new campaign, select button B and the drop-down list with the possible actions appears. Select the action **New** (C) and click **Go** (D).

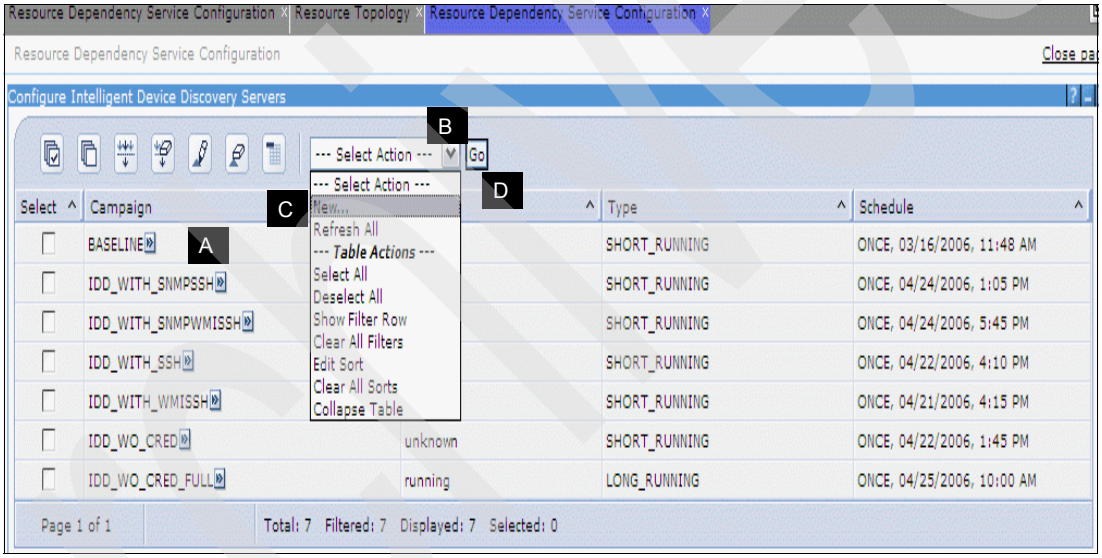


Figure 7-13 IDD create campaign

Another option to create a new campaign is to select the action beside the campaign (A in Figure 7-14). Select **New Based On...** (B).

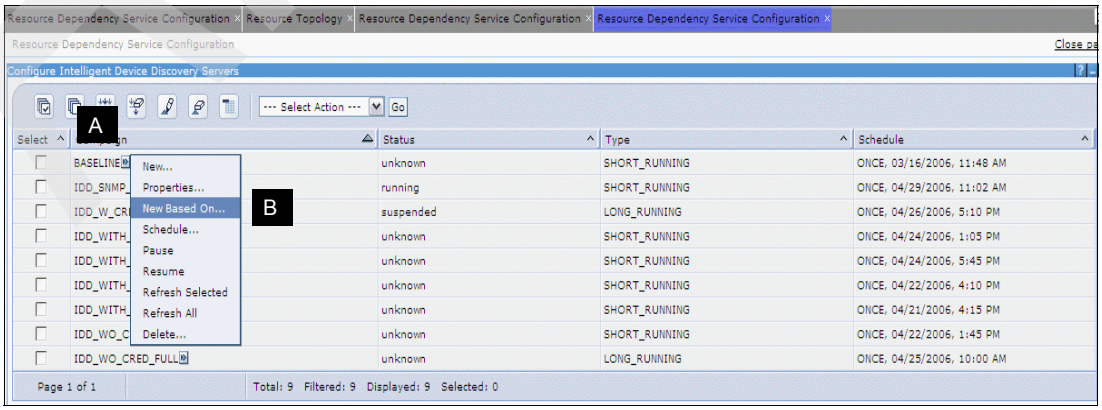


Figure 7-14 IDD campaign creation based on existing campaign

In the IDD campaign configuration window (shown in Figure 7-15), the list on the right area displays the configuration steps:

1. In the first task (General), specify the Campaign name (A) and the campaign type (C). The default value for the campaign type is Short running, which means a scan period of one hour. For campaigns scanning a large network, you must specify Long running as the type of the campaign.

   In the field Base referenced campaign (B), specify a referenced campaign name. This is optional, the entry none means that you do not use a referenced campaign.

   If you want to select a referenced campaign, click the button close to the name field and from the drop-down list select the name of the existing campaign you want as a referenced campaign. If a referenced campaign is defined, the network information is automatically transferred.

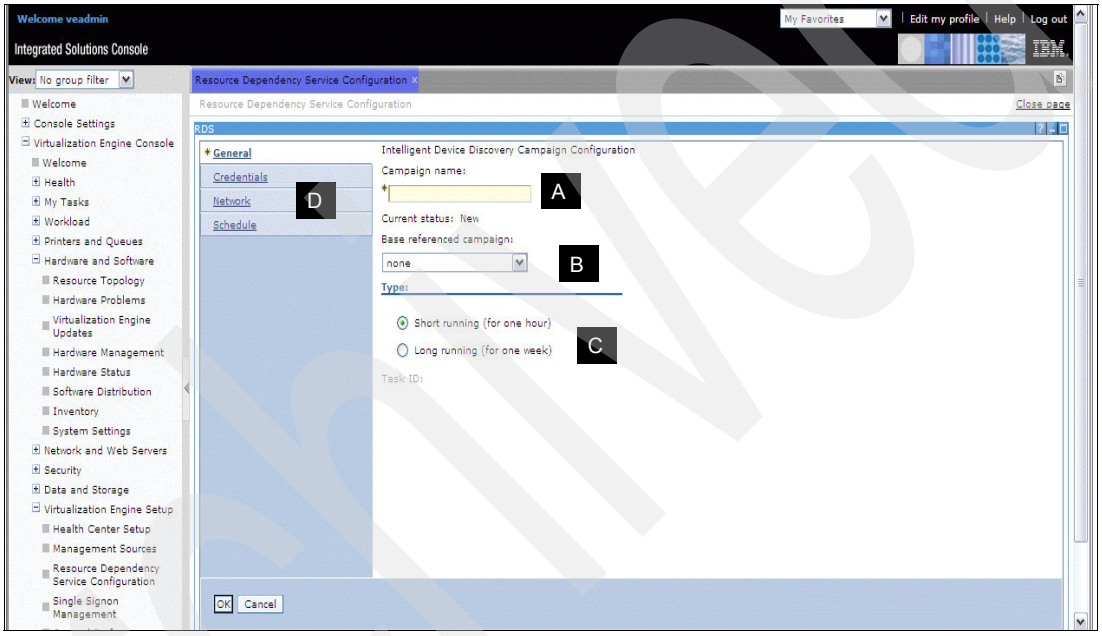


Figure 7-15 IDD general configuration

2. We want to configure a campaign without credentials, that is, a Level 1 scan. Therefore, we skip the credential task (Credentials) and we select **Network** (D).
3. In the network configuration window (shown in Figure 7-16 on page 231), specify a list of hosts, an IP range for the network (**A**), or both, and optionally a list of hosts that should be excluded (**B**).

Archived

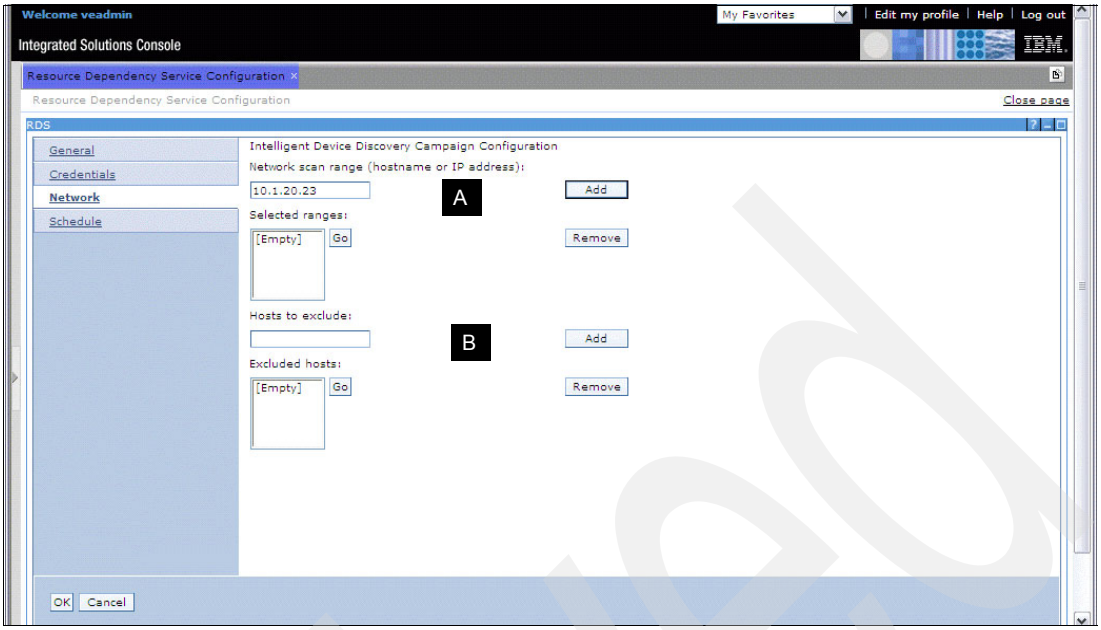


Figure 7-16 IDD network configuration

4. The last step in the campaign configuration is the schedule parameter specification. Select the task **Schedule** on the Resource Dependency Configuration Campaign pane and the schedule configuration window (as shown in Figure 7-17 on page 232) is displayed.

   In the date field (A), the default value is the current date for a new campaign. You can overwrite the default value according to your requirements. You have the option of selecting the schedule date with a graphical calendar interface (B).

   There are four choices (C) for how often a data collection should run:

   – Once: A data collection occurs one time. Provide Start Date and Time.
   – Daily: The collection starts every day at the specified Start time.
   – Weekly: The collection occurs every week.
   – Every two weeks: The collection runs every two weeks.

   Specify the schedule time (D) and click **OK**.

   Note: The schedule time must be in the future. The creation of the campaign requirements sometimes depends on the power of the management server. Therefore, if you want to start the collection at once, do not specify a time that is close to the current time.

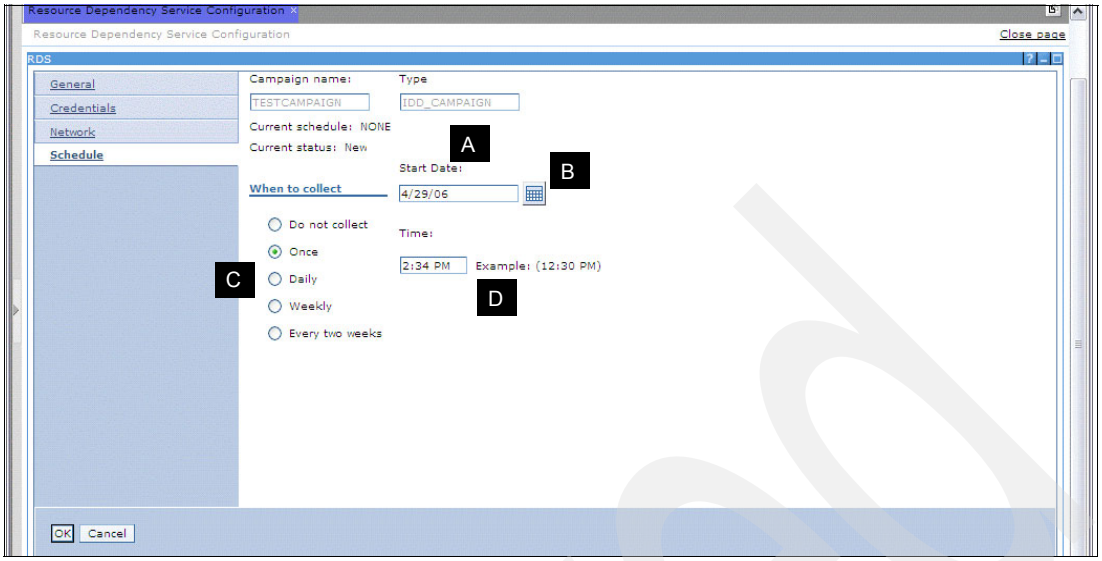


Figure 7-17 IDD campaign schedule configuration

The campaign overview window (as shown in Figure 7-18) shows the status of the campaign. In our case, the campaign is only created and not started. The scheduler starts the campaign according to the specified start time and date.

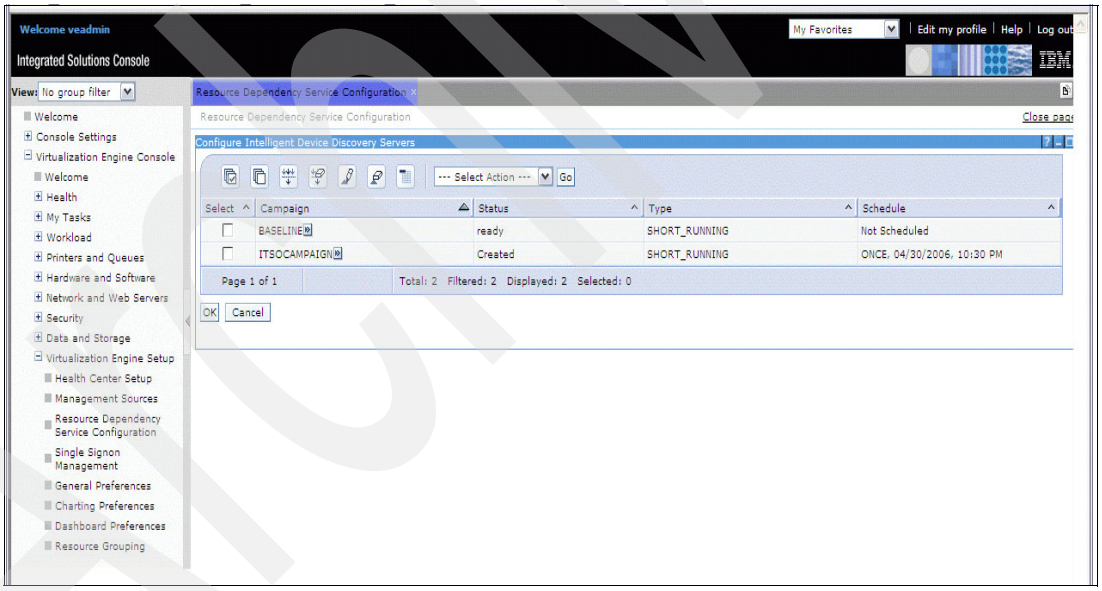


Figure 7-18 Status overview of IDD campaigns

There are two ways within the Virtualization Engine console to show or analyze the discovered resources:

- The topology view
- The discovered resource

The Virtualization Engine console topology view

The Virtualization Engine console topology view provides an overview of all created manageable resources with its relationships. To get the topology view of the resources, in the Virtualization Engine console navigation window, select **Hardware and Software** → **Resource Topology**. Figure 7-19 on page 233 presents the overview of all the MR created in

our ITSO environment after the IDD campaign ran without credentials. The specified IP range comprises the whole environment.

The IDD campaign without credential specification creates one system MR. You can select the resource type and get more information about it.

Note: This overview does not contain all the discovered information; it presents only the manageable resources that have been created from the discovery.

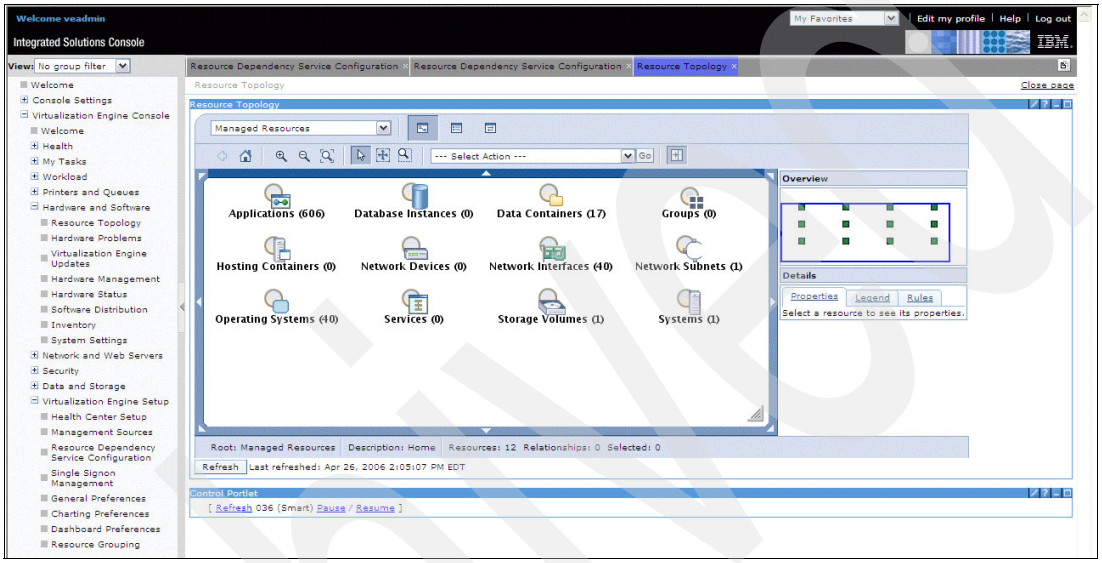


Figure 7-19 Virtualization Engine console topology overview: IDD campaign without credentials

Configure the discovered resource

In the Resource Dependency Service configuration menu, select the third task, **Configure Resource Discovery**, and the Configure Resource Discovery window opens, as shown in Figure 7-20. This window provides the entry point to the tables view of the discovered resource. This window offers a link to the table view for each of the ten manageable resource types. Another way to look at the tables is to use standard DB2 tools.

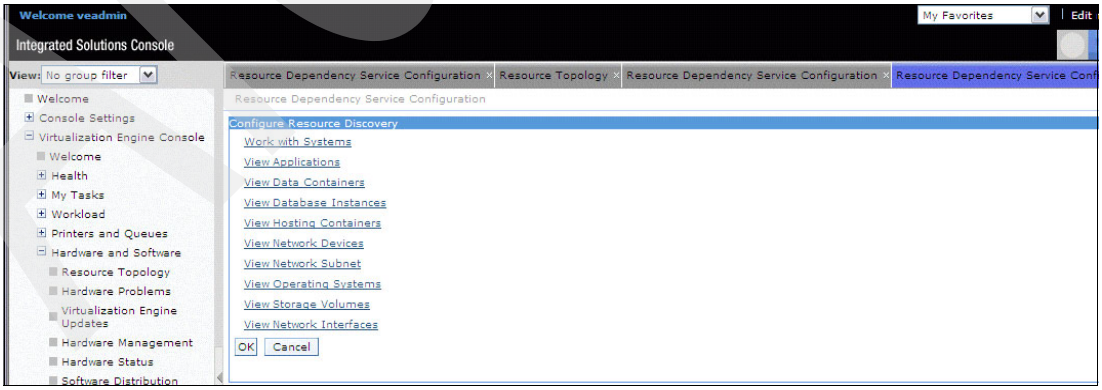


Figure 7-20 Configuring the discovery resources

The ten manageable resource types are the following:

- Systems
- Applications
- Data Containers
- Database Instances
- Hosting Containers
- Network Devices
- Network Subnet
- Operating Systems
- Storage Volumes
- Network Interfaces

These types correspond to the Resource Dependency Service tables in the Resource Dependency Service database. We display a few of them in the following images.

- Figure 7-21 shows the Work with Systems view. For comparison with the Topology view (Figure 7-19 on page 233), we used results from the same campaign.

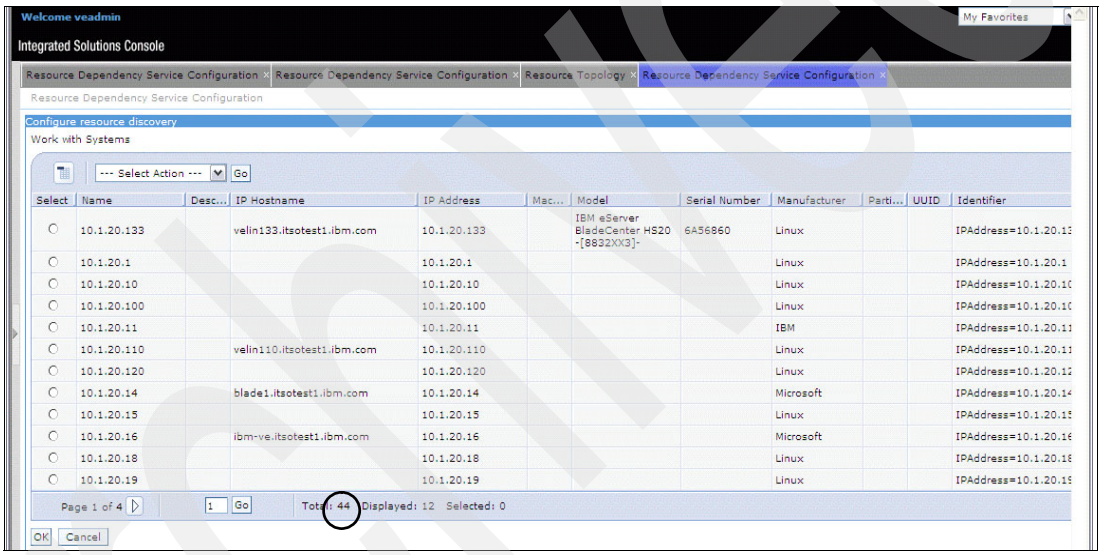


Figure 7-21 Work with System view for IDD campaign without credentials

We can see that 44 systems are discovered. The IDD campaign discovers a Model and a Serial Number only for the first system (velin133.itsotest1.ibm.com). This is the server on which our Virtualization Engine runs. For this run, we installed all the Virtualization Engine components on a Linux system running on a blade. For all other systems, this information is not found. This machine is the only one that is visible in the Topology view (Figure 7-20 on page 233). The model and serial number information are required for the generation of a unique ResourceID. This ID is a mandatory property for each manageable resource system type. Not having a ResourceID prevents the Resource Modeling Framework from creating a manageable resource.

- The Operating Systems view, shown in Figure 7-22 on page 235, shows that the required information to create an operating system type of a manageable resource is discovered. The Y in the field Viewable indicates that you can view the manageable resource from the topology view. Figure 7-22 on page 235 shows that 43 operation systems have been discovered and 40 can be viewed on the topology view: three of them do not have the required information to become a manageable resource.

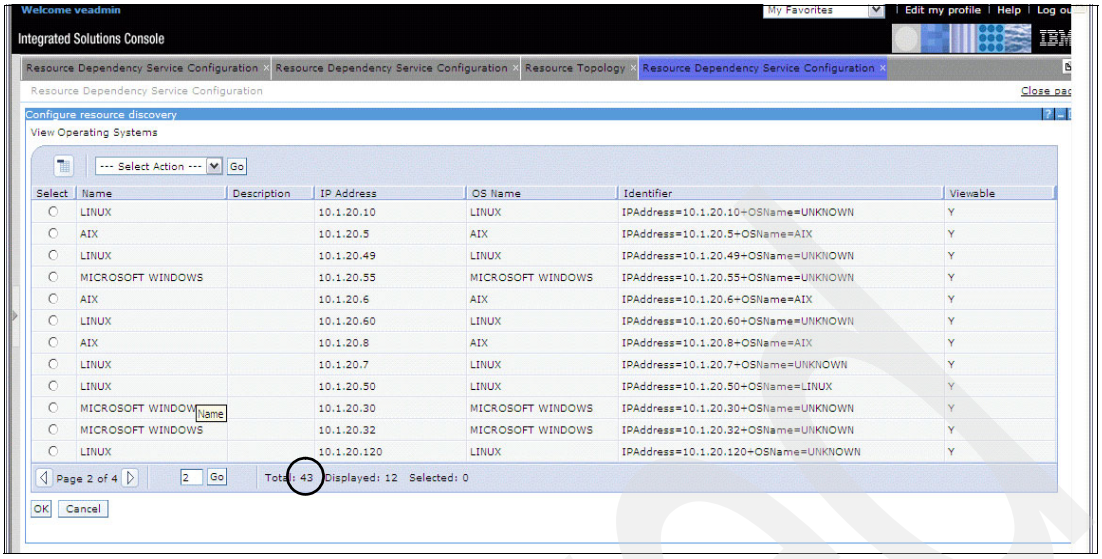


Figure 7-22 Operating Systems view for IDD campaign without credentials

- Figure 7-23 shows the discovered Network Interfaces. Only one management server (10.1.20.133) is described with its MAC address. Eighty-six network interfaces have been discovered, of which 40 can be defined as manageable resources (displayed in the topology view).

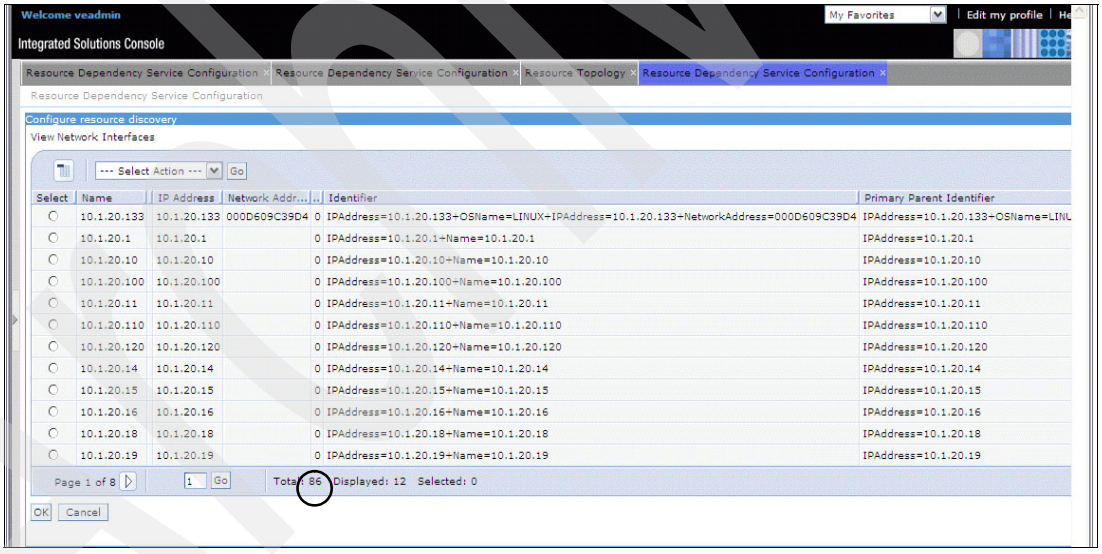


Figure 7-23 Network Interfaces view for IDD campaign without credentials

- A list of the discovered Data Containers is shown in Figure 7-24. All the discovered data containers belong to the Virtualization Engine management server, as shown in the Identifier field. The Identifier field contains the concatenated list of resource attributes, which identifies the resource.

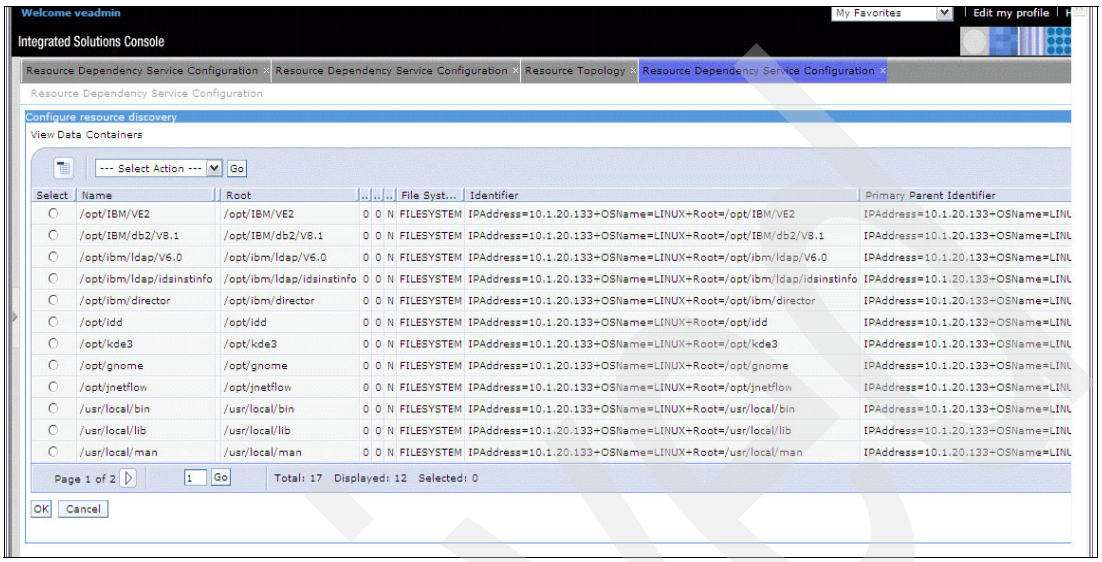


Figure 7-24 Data Containers view for IDD campaign without credentials

The IDD scan without credentials does not find any data containers on the other systems because credentials are required for this information.

- Figure 7-25 is an example of the Applications discovered resource types. All the discovered applications belong to the Virtualization Engine management server. As with the Data Containers, due to the lack of credentials, IDD is not able to get application information from other systems.

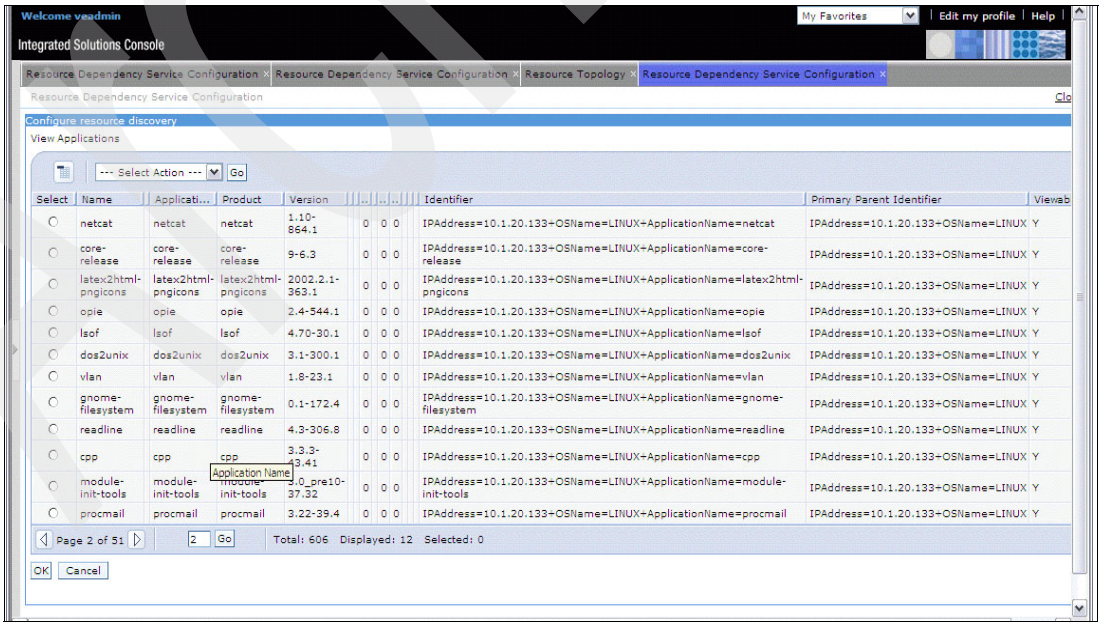


Figure 7-25 Application view for IDD campaign without credentials

- The Storage Volumes view (Figure 7-26) shows the storage device of the Virtualization Engine management server.

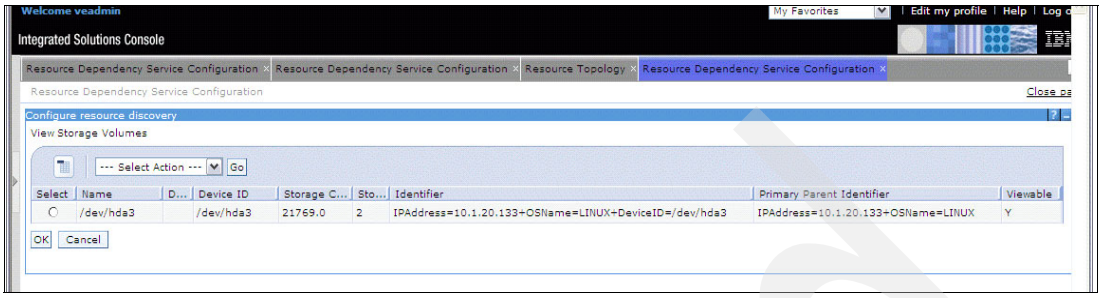


Figure 7-26 Storage Volumes view for IDD campaign without credentials

IDD campaign configuration with credentials

You can select this configuration task during the creation of a new campaign or you can change an existing campaign. To create a new campaign with credential usage, you have to perform four steps:

1. Specify general campaign information.
2. Define campaign credentials.
3. Define the network range.
4. Schedule information.

Steps 1, 2, and 4 are already described in “Configuring an IDD campaign without credentials” on page 229. These steps are also applicable for the campaign configuration with credentials. Only the defined credential step is described here.

To add the credentials to a campaign, select the **Credential** task, as shown in Figure 7-27. Select **Selection Action** → **New**.

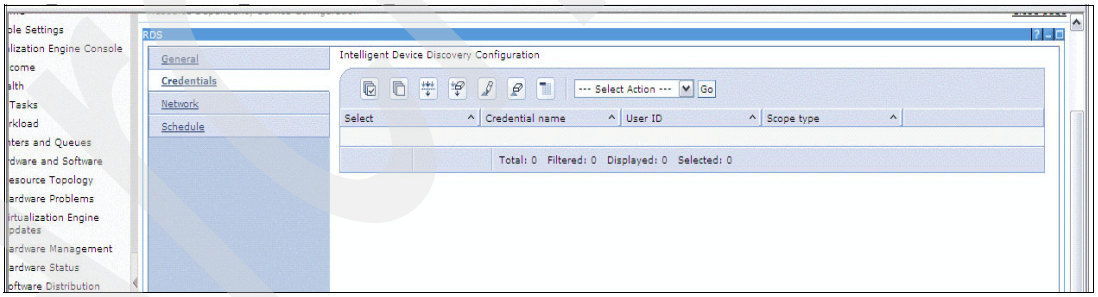


Figure 7-27 IDD configuration credentials

In the credential configuration window, you have to provide the necessary information. The information you have to provide depends on the credential type: SNMP, WMI, or SSH.

SNMP

If you choose SNMP, as shown in Figure 7-28 on page 238, you have to provide the Credential scope, which is a name for the SNMP credential scope. The User ID field is not active. In the Password field, specify the community string. In the Confirm password field, you have to confirm the community string. If you have more than one community string for your devices and resources, you have to repeat the configuration for each community string.

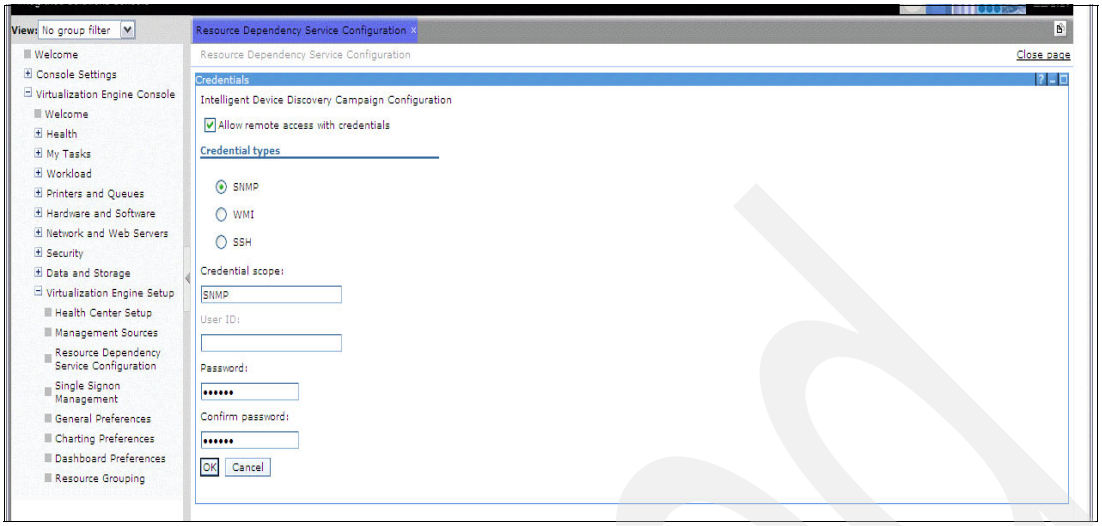


Figure 7-28 IDD campaign credential configuration: SNMP

WMI

If you choose WMI, as shown in Figure 7-29, you have to provide the Credential scope, the User ID or Domain\User ID, and the Password of the WMI credential. You have to repeat this configuration for each WMI credential in your environment.

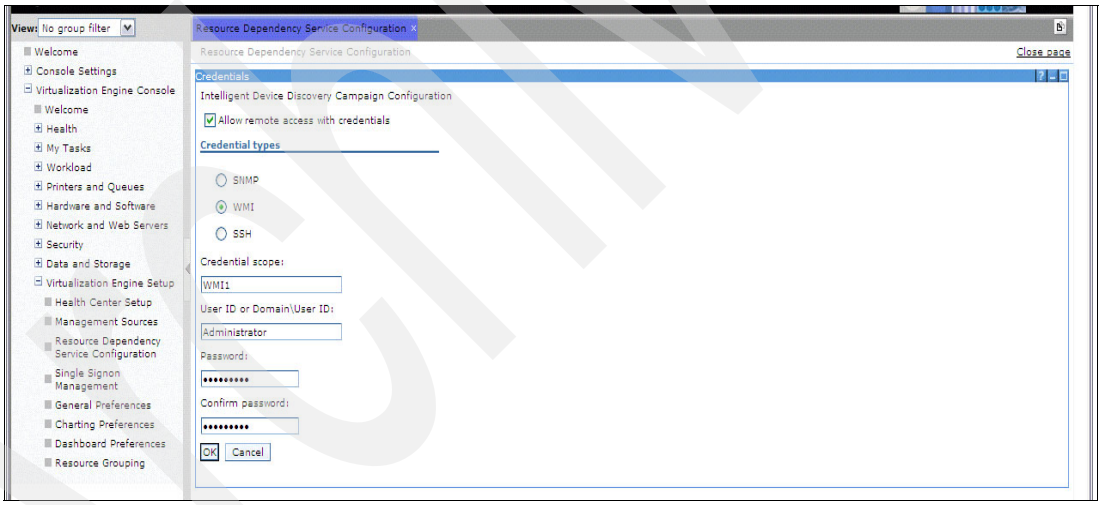


Figure 7-29 IDD campaign credential configuration: WMI

Archived

SSH

For SSH, you have to provide the Credential scope, which is the name of the credential, the User ID and the Password**,** as shown in Figure 7-30**.** For each combination of User ID and Password in your environment, you have to configure the campaign credentials.

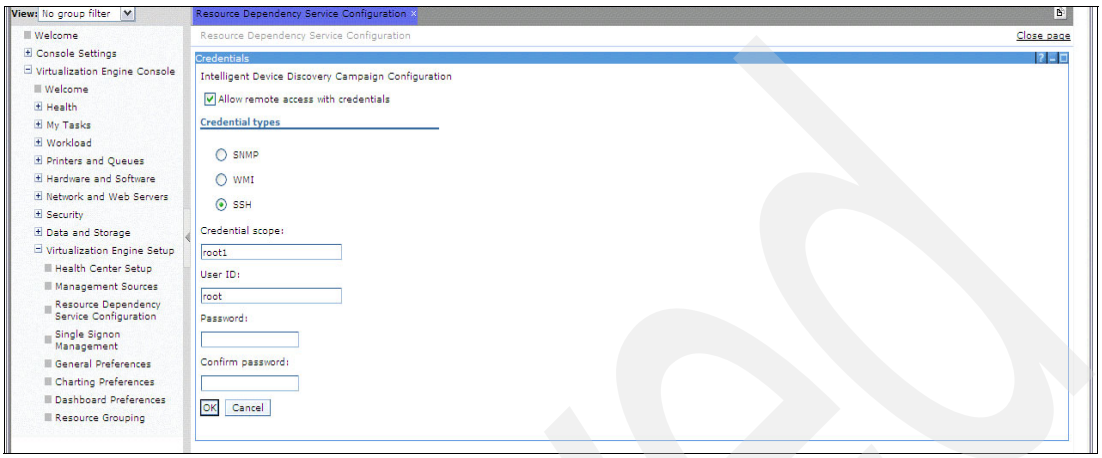


Figure 7-30 IDD campaign credential configuration: SSH

Update credential

When you click the **OK** button after the credential configuration step, the list of all configured credentials appears, as shown in Figure 7-31. In our example, we configured one SNMP credential, two WMI credentials, and four SSH credentials.

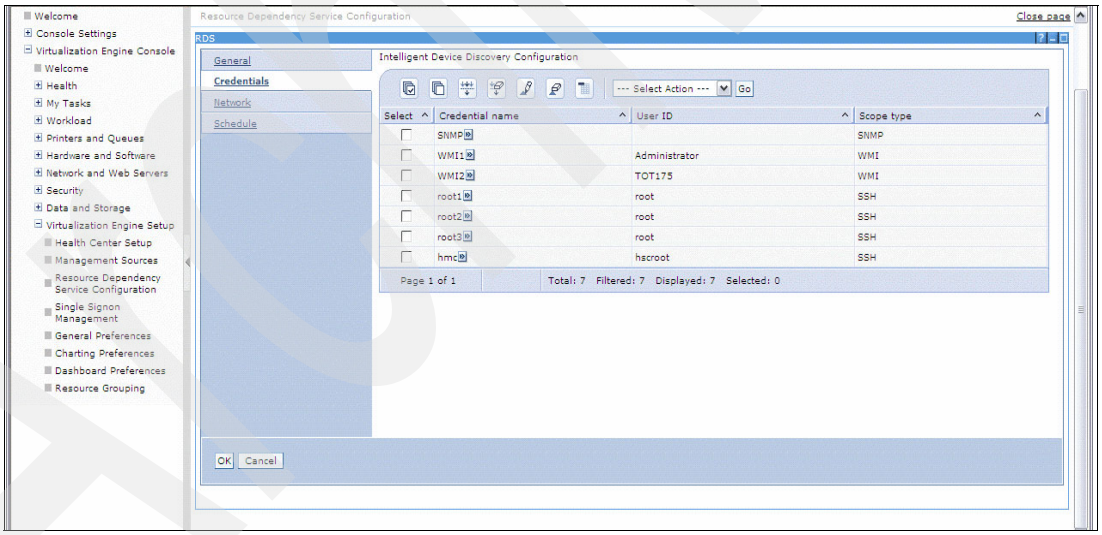


Figure 7-31 IDD campaign configured credential overview

If you want to change the password of the credentials, select the credential name and click the action button beside the name, as shown in Figure 7-32 on page 240. You can select three possible actions:

- New to create a new credential type.
- Update to change the password.
- Delete to remove the credential from the list.

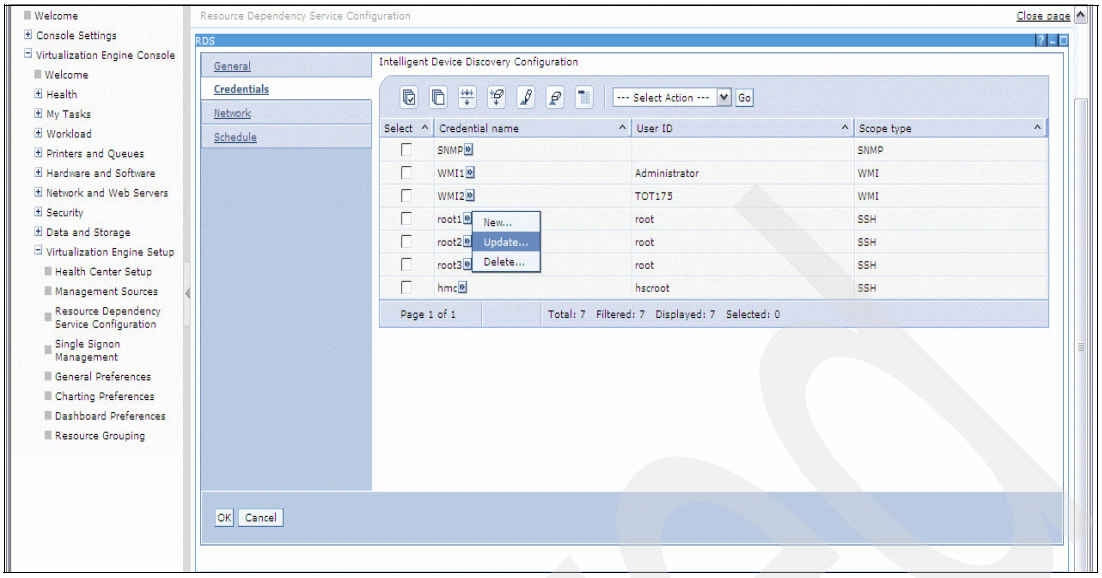


Figure 7-32 IDD campaign credential update

Note: Be aware than configuring more than three credentials per campaign can result in user ID lock-out situations if three invalid credentials are tried for a system.

The IDD scan with credentials retrieves more information from the systems than the scanning without credentials. Therefore, the ODIRM component can create more manageable resources. The total number of detected systems is not increased, but the information per system is more complete, as shown in Figure 7-33 on page 241.

For a credential campaign, the following scanners provide the information:

- idduxscan
- wmiscan
- hostscanner
- snmpwalk
- httpscan
- nsa

Archive

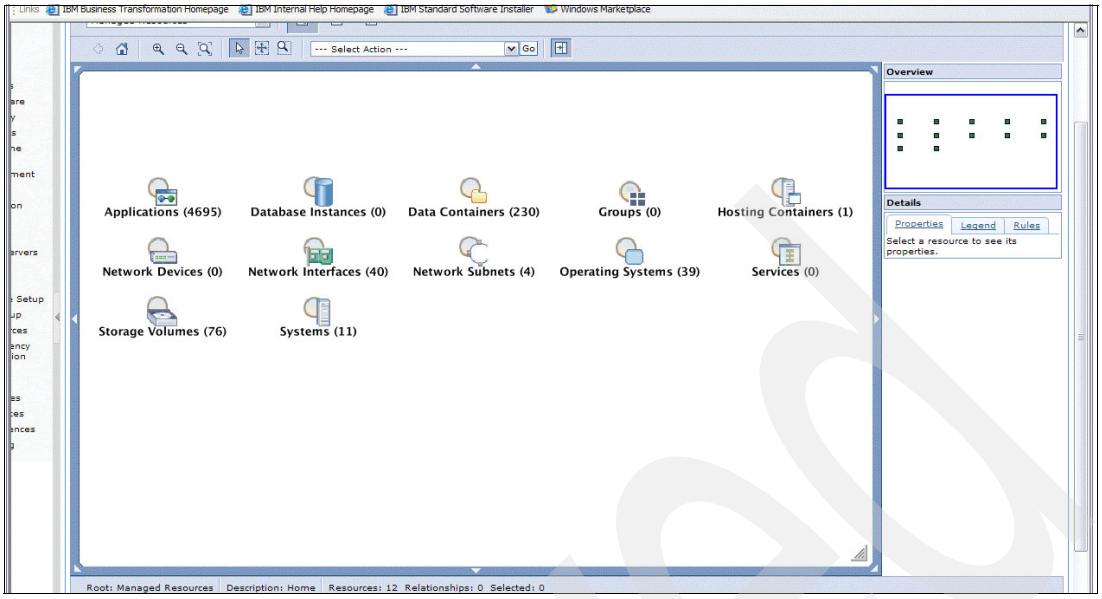


Figure 7-33 Topology View of IDD campaign with credential

IDD creates iSeries™, z/OS, and WebSphere data collectors when it detects these as part of its scanning. For example, if, during an IDD scan, IDD identifies a particular IP as an iSeries system, it will then create an iSeries data collector for that IP and attempt to run it if credentials exist for the iSeries. If not, the collector will be created and the user can add credentials to it from the “Configure data collectors for standard resources” option of the main Resource Dependency Service window.

Figure 7-34 shows a list of the automatically generated data collectors. The name starts with IDD followed by a W or Z depending on the type of collector. W indicates a WebSphere collector and Z indicates a z/OS collector. The end of the name forms the IP address or the host name. For example, the collector name IDDW_velin133 stands for the WebSphere data collector for host velin133.

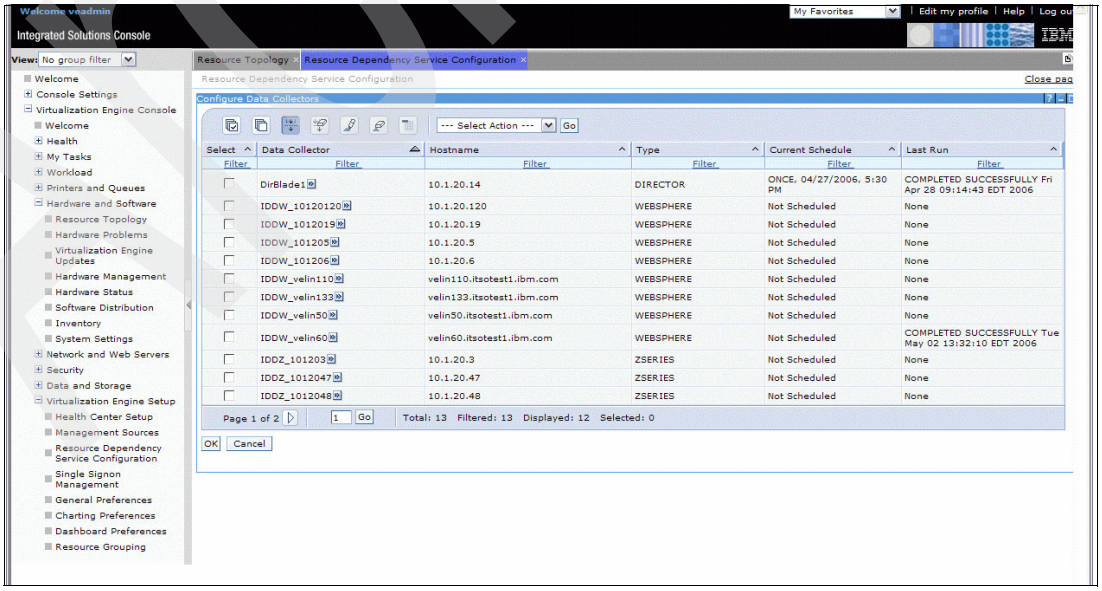


Figure 7-34 IDD generated data collectors

The result of this collector is shown in Figure 7-35. The data collector needs user credentials for retrieving the data. If the credentials are configured during the single sign-on (SSO) configuration, the automatic data collector process can use these credentials and the collector gathers data. If no credentials are specified, the collector finishes with errors. In an additional step, the credentials must be specified.

In our example, the server velin133.Itsotest hosts a standard WebSphere installation. Thus, only the default applications with the default server1 are shown.

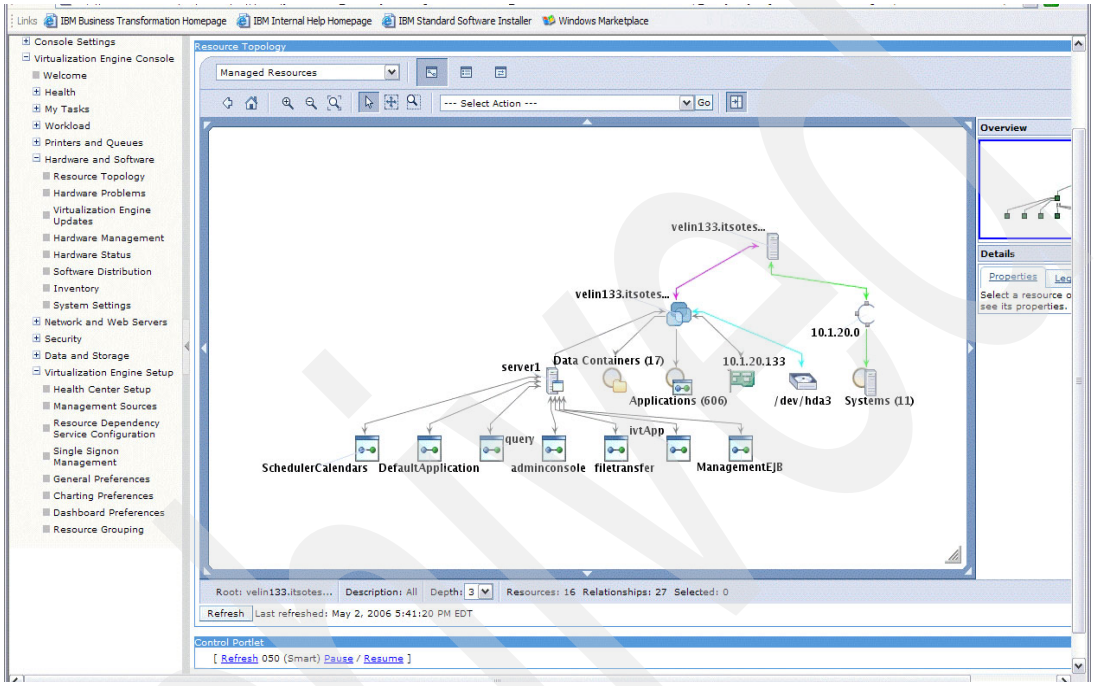


Figure 7-35 WebSphere topology as discovered in IDD campaign

The results of the IDD campaign can be displayed within the Resource Dependency Service configuration windows, as described in 7.3.1, “Planning the IDD campaign” on page 226. Figure 7-36 on page 243 shows the systems view. The table view of the discovered systems displays more details with the credential campaign for the same environment; more Model and Serial Number properties have been found.

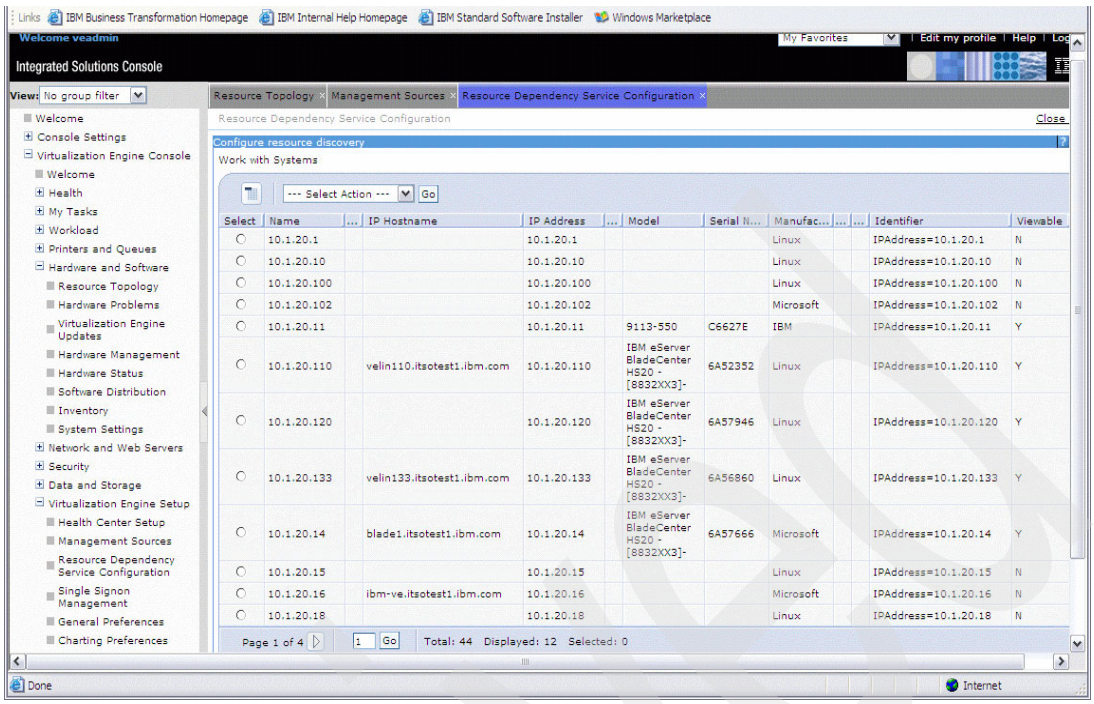


Figure 7-36 IDD campaign with credentials results: Systems view

7.4 Planning and configuration of Resource Dependency Service data collectors

There are two types of scanners bundled with Resource Dependency Service: the IDD tool and collectors. The main difference between the two is in the way the discovery of resources is done. Collectors are specific to a machine, while IDD can take a range of IP addresses. Resource Dependency Service collectors add information that is specific to particular applications or systems management tools into the Resource Dependency Service databases. The collectors need access and connection information for the specific data collector. Before using the collectors, you have to gather the necessary access information.

At the time of the writing of this IBM Redbook, Resource Dependency Service provides the following collectors:

- Tivoli Business Systems Management (TBSM)
- Tivoli Configuration Manager (TCM)
- IBM Director
- Enterprise Workload Manager (EWLM)
- IBM System i
- DB2 Universal Database
- WebSphere MQ
- WebSphere Application Server
- IBM System z z/OS

7.4.1 Where to look for the collectors

Before you can access and configure collectors, you have to log in to the Virtualization Engine console using your Virtualization Engine credentials.

1. Open your Web browser and enter the following:

   `http://<hostname or IP Address> :8421/ibm/console`

   In this example, 8421 is the default port number for the Virtualization Engine console.

2. The login screen opens, as shown in Figure 7-37. Enter the User ID and password. Click **Log in**.

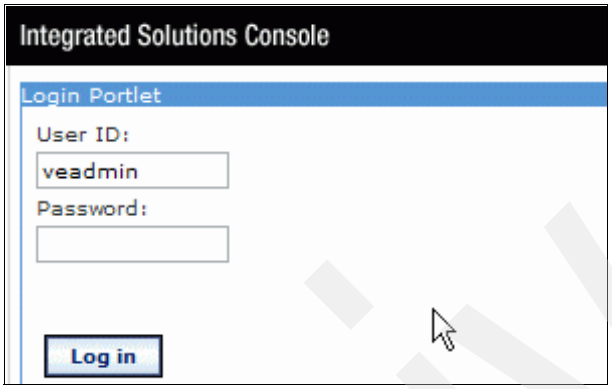


Figure 7-37 Login window

3. The navigation window opens, as shown in Figure 7-38. Expand **A Virtualization Engine Console** → **B Virtualization Engine Setup** → **C Resource Dependency Service Configuration**. Click **D Configure data collectors for standard resources**.

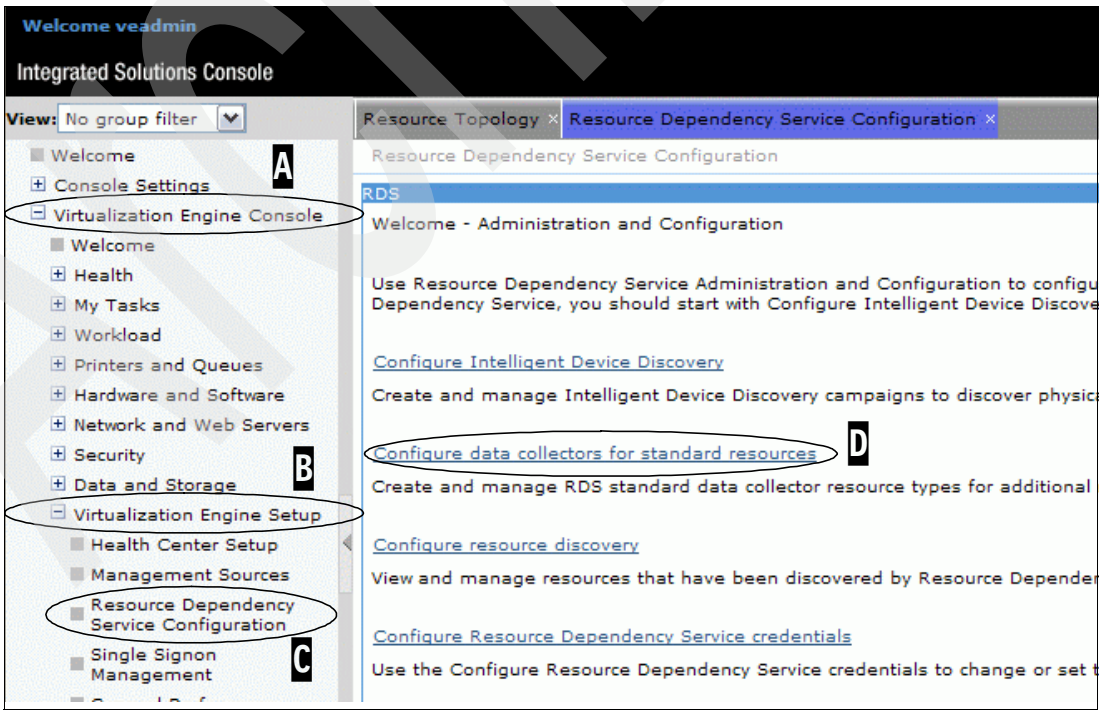


Figure 7-38 Navigation window

Archived

4. The **Configure Data Collector** window opens, as shown in Figure 7-39. You can register new collectors in this window.

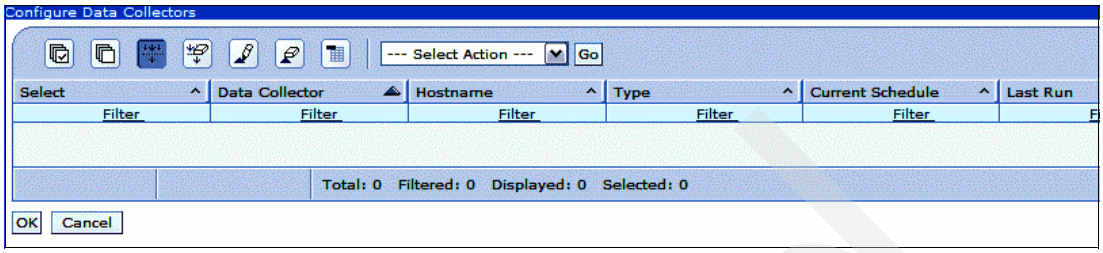


Figure 7-39 Configuring the data collectors

7.4.2 Configuring a new collector

In this section, we explain how to configure the following collectors:

- A collector for the machine hosting DB2 Universal Database, installed on 10.1.20.50
- A collector for the machine hosting WebSphere Application Server, installed on 10.1.20.60

DB2 Universal Database

To define a DB2 Universal Database collector:

1. In the **Select Action** menu, choose the **New** option, as shown in Figure 7-40. Click **OK**.

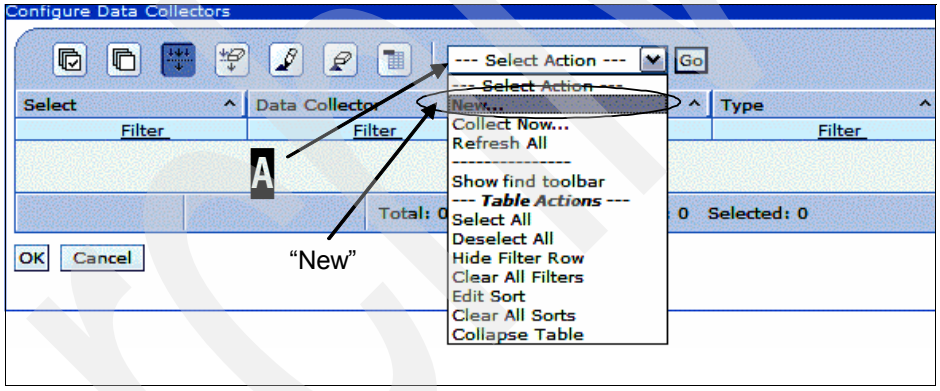


Figure 7-40 Registering a new collector

2. A New Data Collector window is displayed. From the drop-down list, select **DB2 Universal Database**, as shown in Figure 7-41. Click **OK**.

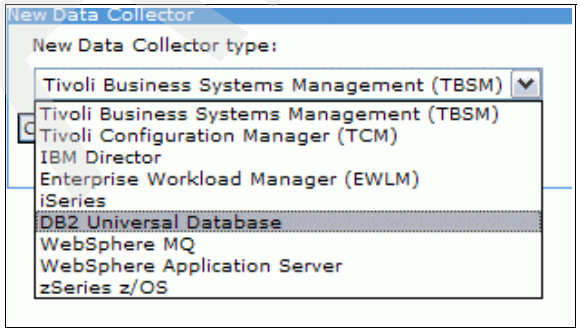


Figure 7-41 New data collector: DB2 Universal Database

3. A new window is displayed with three main tags: General, Instances, and Schedule.

   In the General section, specify the collector name A (you can enter any name) and the IP address B, as shown in Figure 7-42.

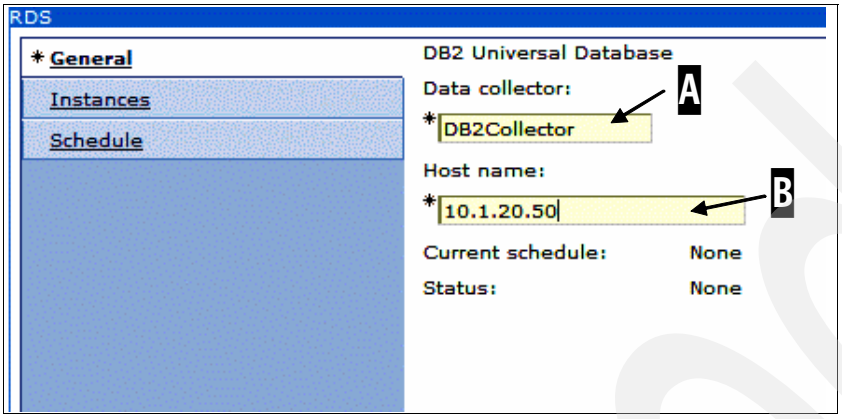


Figure 7-42 General description

4. In the Instances section, from the drop-down list A, select **New**. Click **Go**, as shown in Figure 7-43.

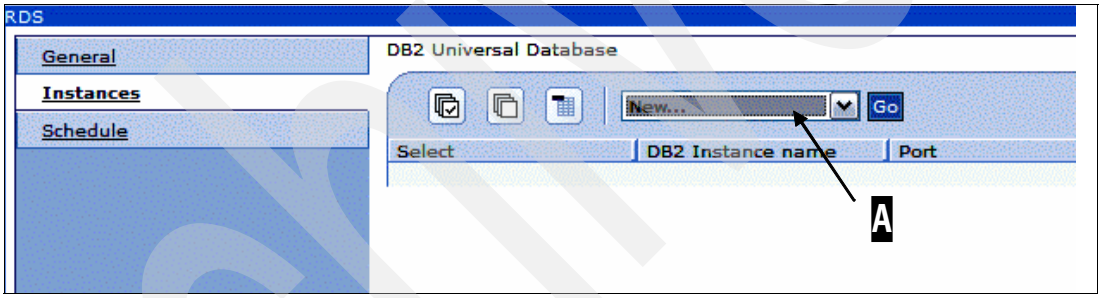


Figure 7-43 New instance

5. In the new window that is displayed, specify the DB2 instance details, such as name and port details A. Click **Credentials** B, as shown in Figure 7-44 on page 247.

Archived

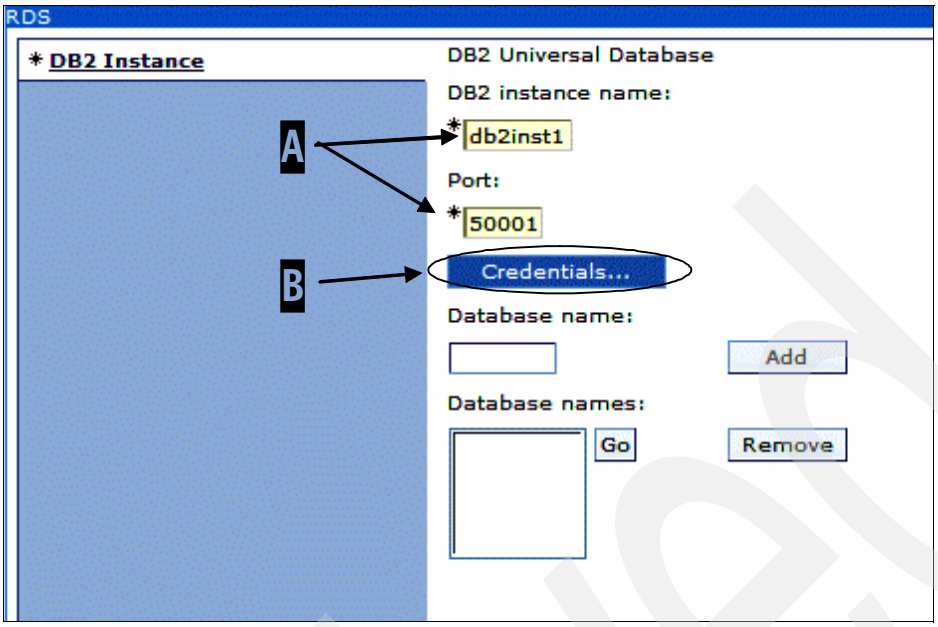


Figure 7-44 Instance details

Note: The instance details are usually available with the database administrator who manages the particular server.

6. The Application Association Information window is displayed (Figure 7-45). Specify **Password** details, click **Apply**, and then click **OK**.

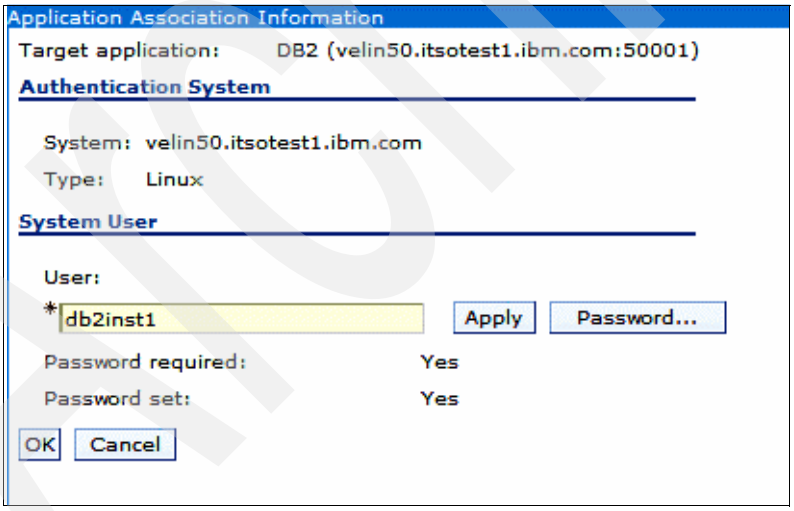


Figure 7-45 Password details

7. This takes you back to the previous window where you specified the instance details (see Figure 7-44). In the **C** field, specify any valid database name that exists on the machine 10.1.20.50. Click **Add** (**D**), as shown in Figure 7-46 on page 248.

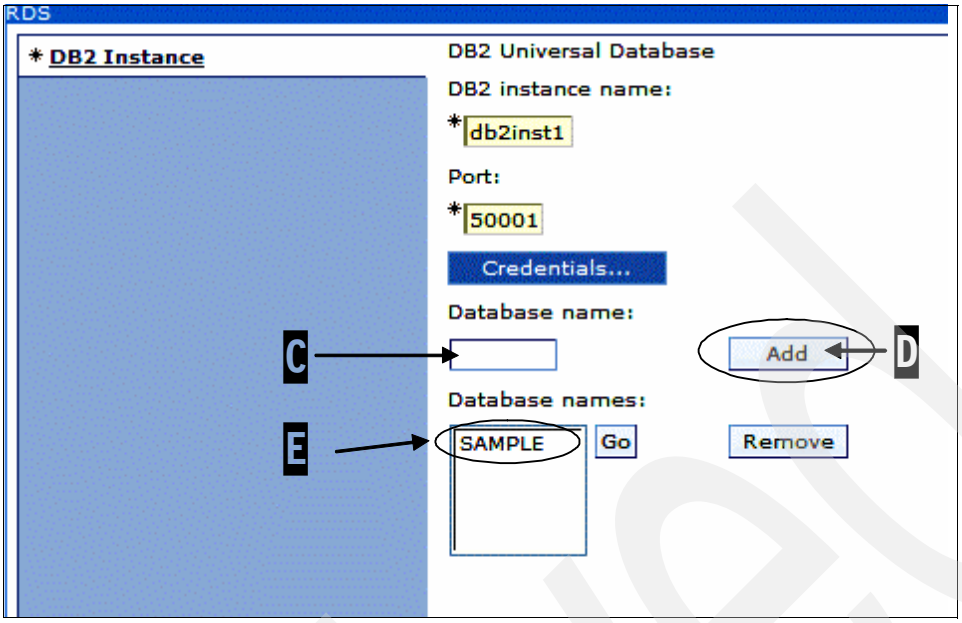


Figure 7-46 Database details

The added database names are reflected in the field **E**. You can specify more than one database name. When you click **OK**, the newly created DB2 instance is registered, as shown in Figure 7-47.

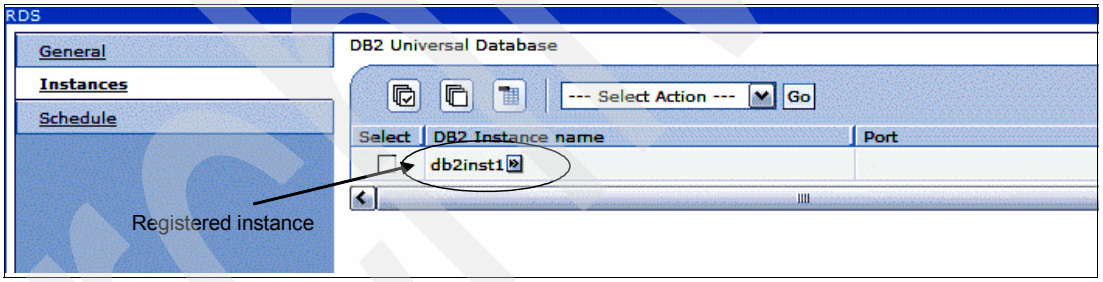


Figure 7-47 db2inst1 registered

8. In the Schedule section, you see the window shown in Figure 7-48 on page 249. Define the scheduling of the collection. Depending upon your requirement, you can choose from the various choices available. Click **OK**.

Archived

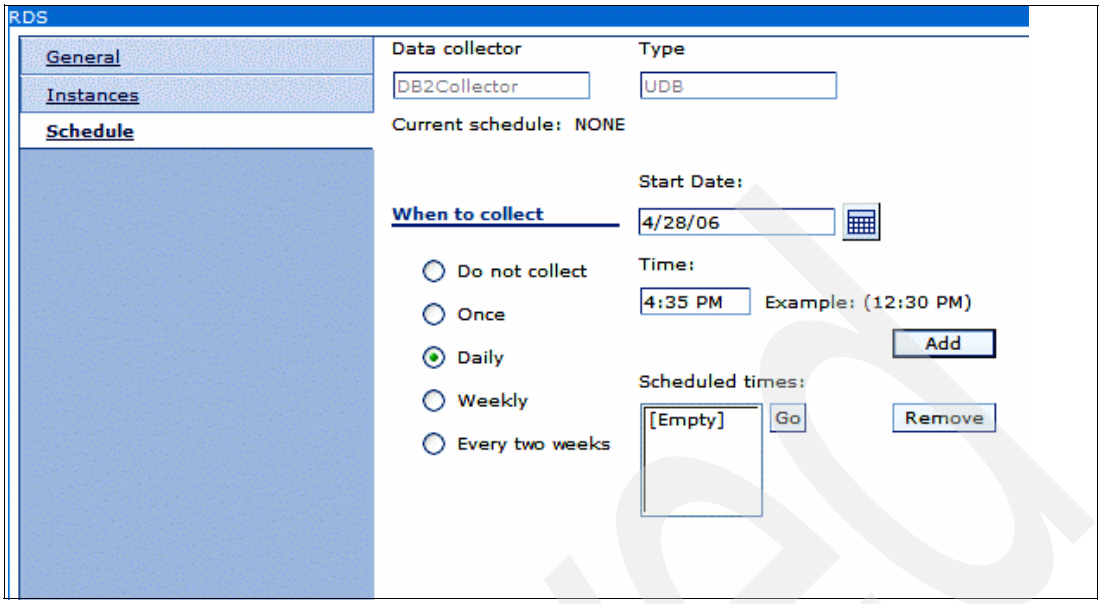


Figure 7-48 Scheduling the collector

This registers the newly defined collector, the DB2Collector in Resource Dependency Service, as shown in Figure 7-49.

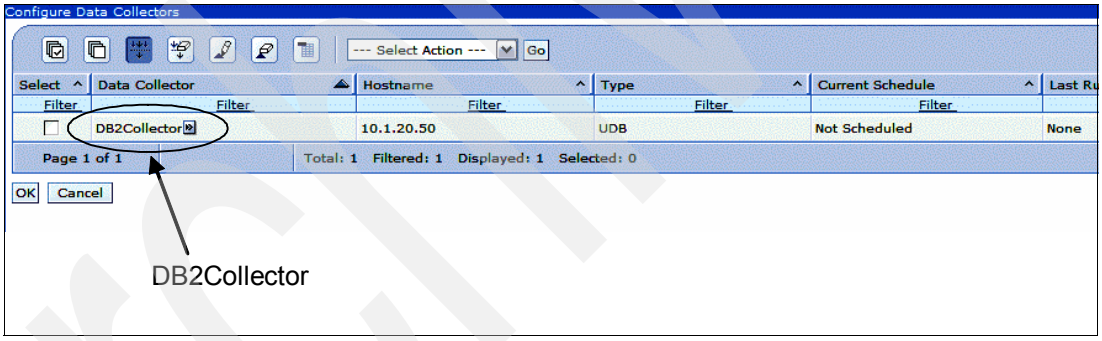


Figure 7-49 Registered DB2Collector

WebSphere Application Server

To define a WebSphere Application Server collector:

1. Repeat the step shown in Figure 7-40 on page 245, that is, click **New**.
2. From the drop-down list, select **WebSphere Application Server**, as shown in Figure 7-50. A window is displayed with two main tags, General and Schedule.

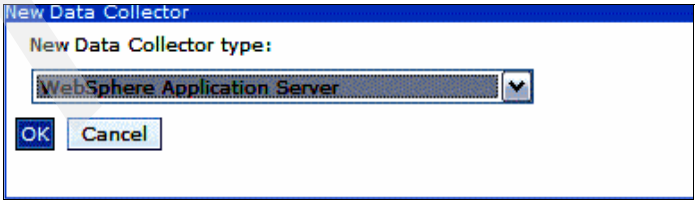


Figure 7-50 New data collector: WebSphere Application Server

Archived

3. In the General section, specify a collector name A (any name is possible), the IP address, and the port details (B). To define the credentials, click **Credentials** (C), as shown in Figure 7-51.

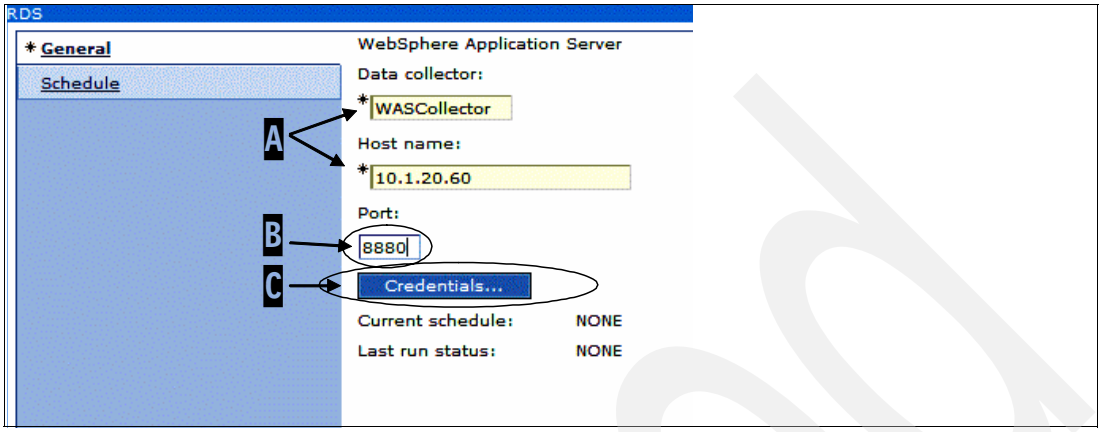


Figure 7-51 WebSphere Application Server collector details

> **Note:** You can check the port details in the WebSphere Application Server Administration window using the following URL:
>
> `http://<hostname>:9060/admin`
>
> Select **Servers** → **Application servers**. Select your server name, that is, **server1**. Select **Communications**. Expand **Ports** and look for the port information associated with SOAP_CONNECTOR_ADDRESS.

4. When you click **Credentials**, the Application Association Information window opens, as shown in Figure 7-52. Specify the **System User** details (the user name and password of the machine hosting WebSphere Application Server). Click **Apply**, and then click **OK**.

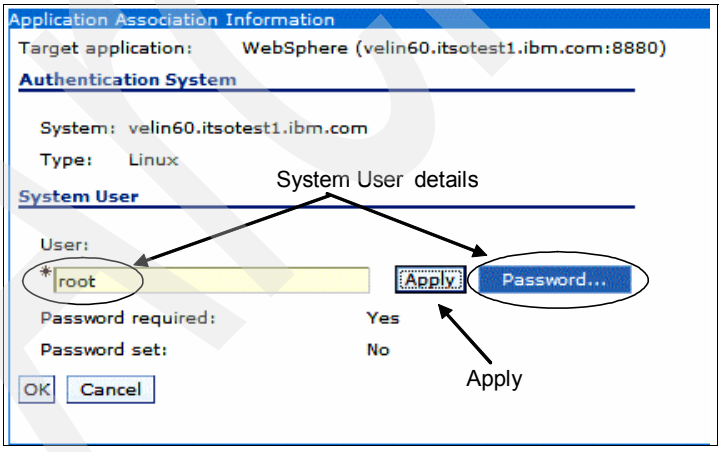


Figure 7-52 Application association information

5. This takes you back to the previous window, as shown in Figure 7-51. In the Schedule tab, schedule the execution of this collector, as shown in Figure 7-53 on page 251. Click **OK**.

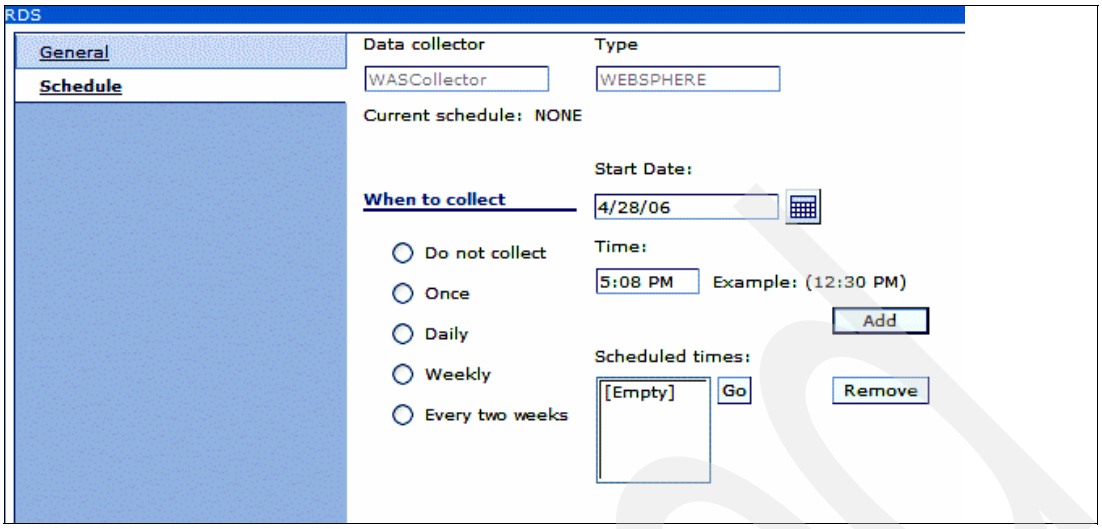


Figure 7-53 Scheduler for WASCollector

6. Your collector is defined in the Resource Dependency Service, as shown in Figure 7-54.

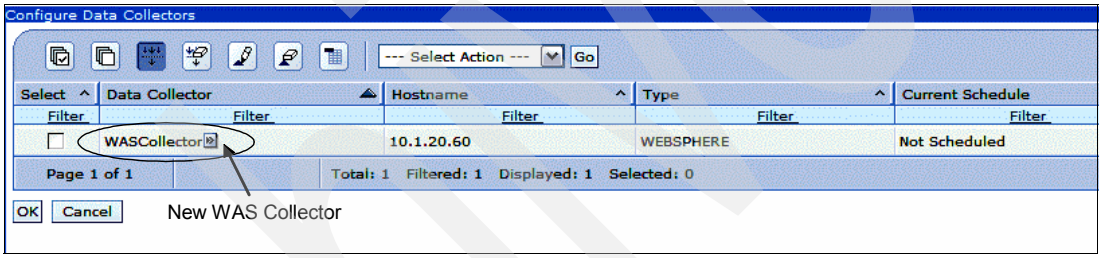


Figure 7-54 WASCollector

7.4.3 Executing a collector

A collector can collect the information from the resource in two ways:

1. Automatically from the scheduling specified while configuring a collector, as shown in Figure 7-48 on page 249. The following options are available:

   a. Once
   b. Daily
   c. Weekly
   d. Every two weeks

   For each of the options, you can set the time of execution. At a specified time, the collector is executed and information is gathered. The collected information is stored in the Resource Dependency Service database.

2. Manually forcing the collector to execute. This method can be used to fetch the information any time. Perform the following steps to execute the collector manually.

   a. Select the collector.
   b. From the drop-down list, select **Collect Now**, as shown in Figure 7-55 on page 252. Click **Go**.

Archive

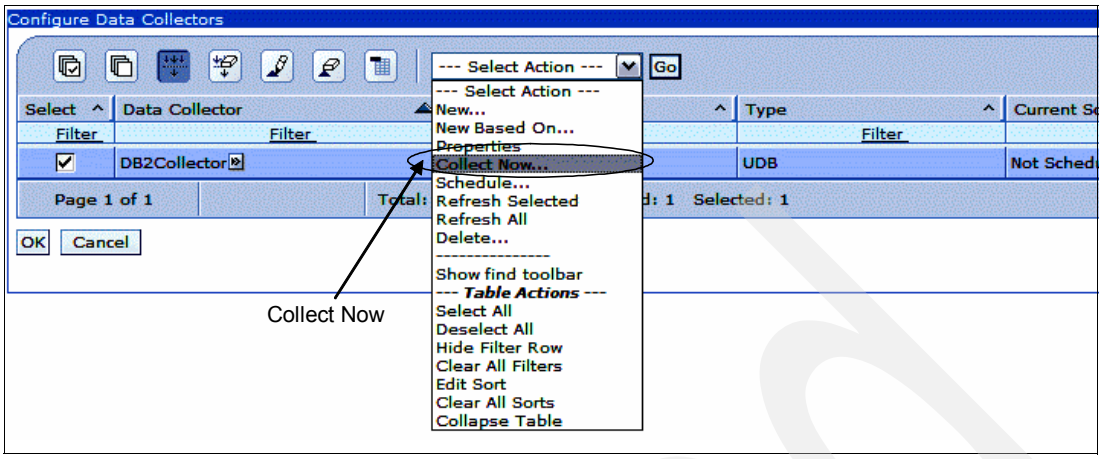


Figure 7-55 Manually executing the collector

c. The next window prompts you to specify all the collectors or a selected one, as shown in Figure 7-56. Click **Collect**.

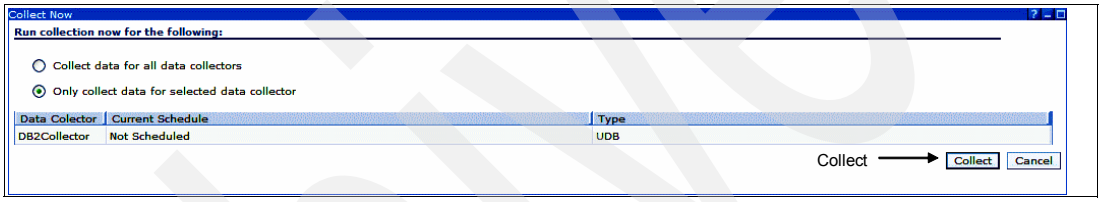


Figure 7-56 Collect Now window

On successful execution, the collector shows the COMPLETED SUCCESSFULLY message, as shown in Figure 7-57.

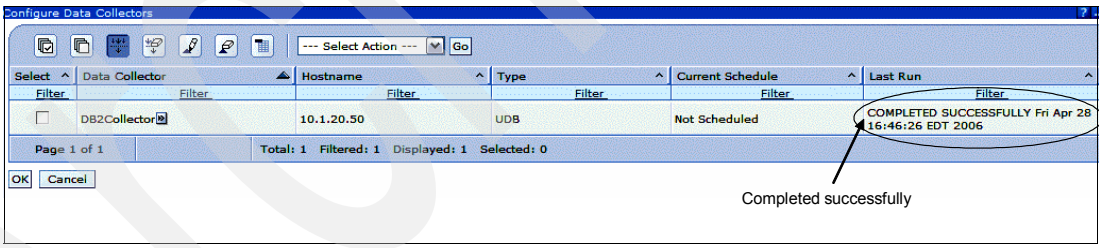


Figure 7-57 Information captured status

7.4.4 Configuring and running the IBM Director collector

The **New Data Collector** window (Figure 7-58 on page 253) allows you to select the collector type that you want to configure. In our case, we select the IBM Director collector.

Archive

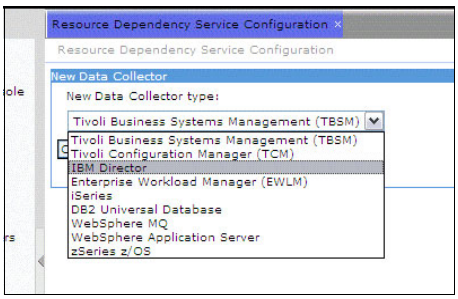


Figure 7-58 Collector selection list

Configuring the IBM Director collector

The configuration of the IBM Director collector is made of two steps: provide the parameters and schedule the collection.

1. In the first step, in the General tab, specify the name of the collector (A), the host name, or the IP address of the IBM Director Server (B) and the credentials (C), as shown in Figure 7-59. The name of the collector is arbitrary and you can specify any meaningful name.

   For the input of the credentials, click **Credentials** (C), and the Application Association Information window is displayed.

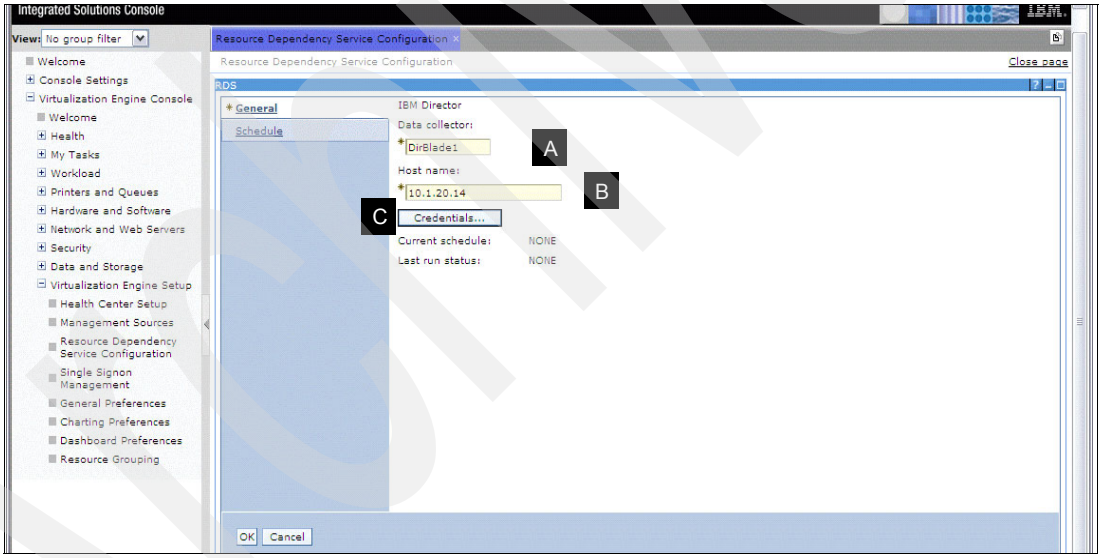


Figure 7-59 Resource Dependency Service IBM Director collector general configuration

2. In the Application Association Information window, shown in Figure 7-60 on page 254, specify the credential data. If no IBM Director application association is already available, you have to select the type of the system. Click the drop-down list (A) and select the appropriate type.

   Type the user name in the User field (B) and click **Apply** (C). To provide the password, click **Password** (D).

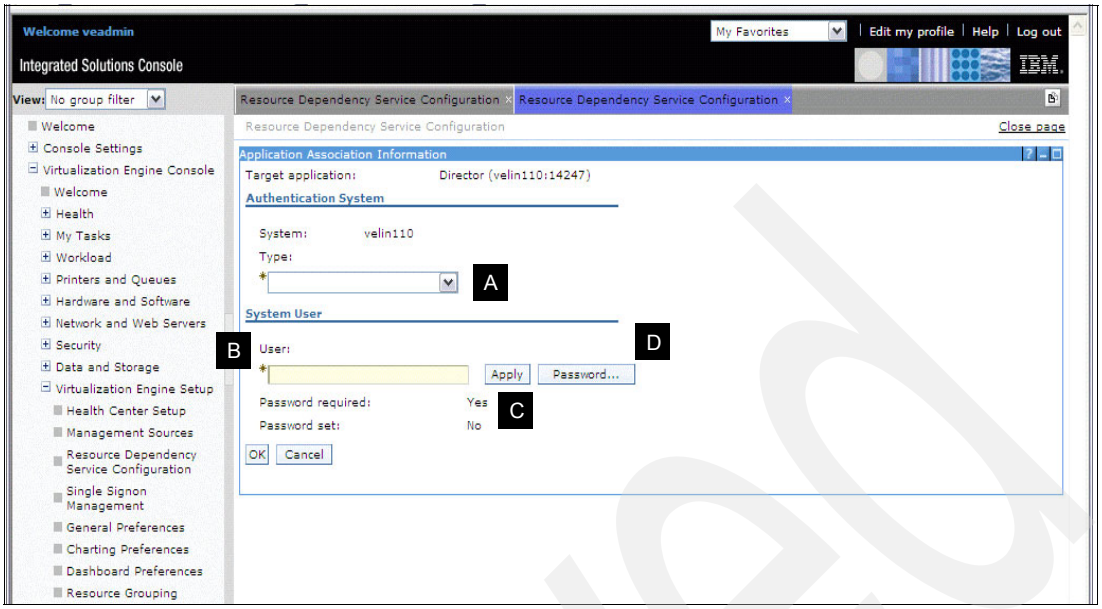


Figure 7-60 Resource Dependency Service IBM Director collector configuration credentials

3. In the Set Password window, shown in Figure 7-61, type the password and confirm it. Click **OK**.

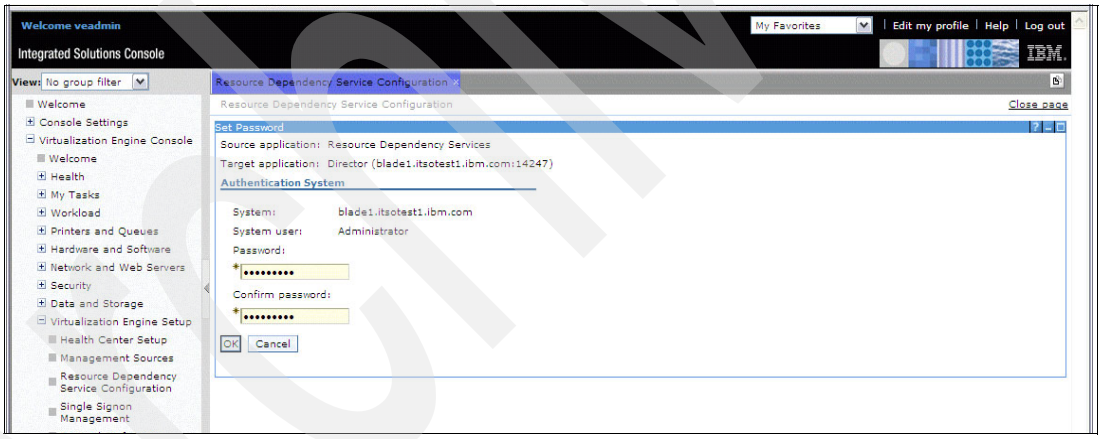


Figure 7-61 Resource Dependency Service IBM Director collector configuration credentials: Setting the password

4. The schedule configuration is the same as for the other collectors, as shown in Figure 7-62 on page 255.

Archived

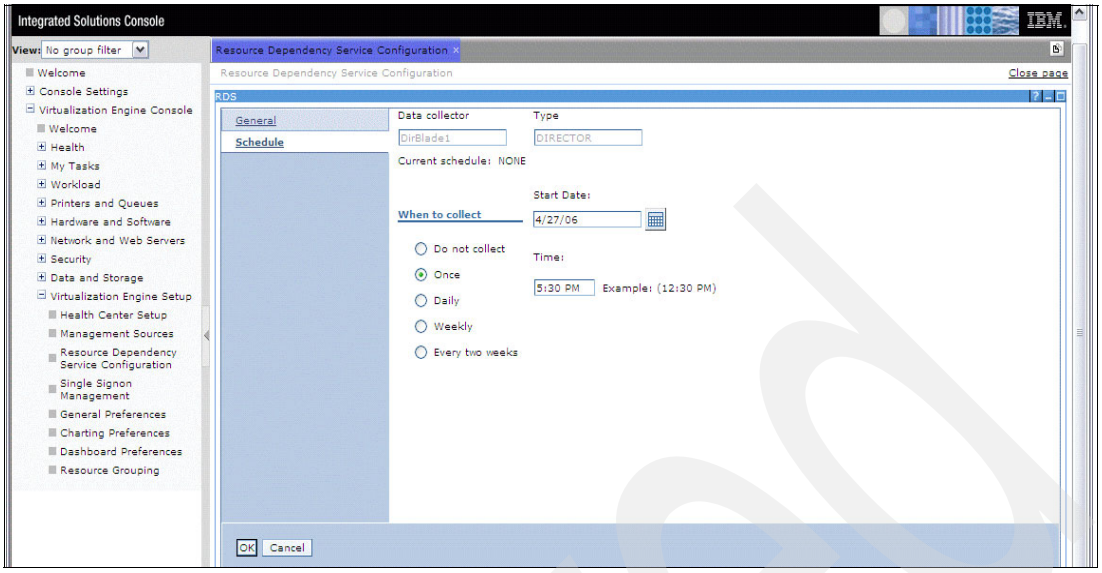


Figure 7-62 Resource Dependency Service IBM Director Collector configuration schedule

After configuring the data collector, the Configure Data Collector view (Figure 7-63) shows the list of already configured and newly created collectors. In this view, the Current Schedule and Last Run field provide run status information. To know whether the collector run is active or not, you must look into the VEMS logs.

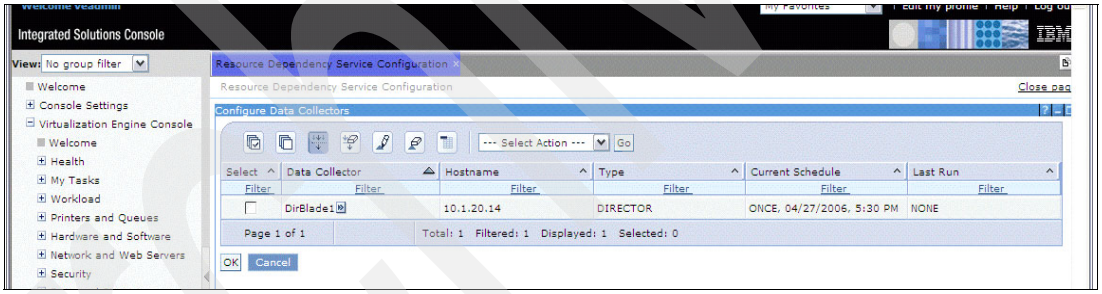


Figure 7-63 Resource Dependency Service data collectors list

Starting the IBM Director collector

To start a collector run:

1. In the Configure Data Collector view, shown in Figure 7-64, click the button next to the Data Collector name. Select **Collect Now**.

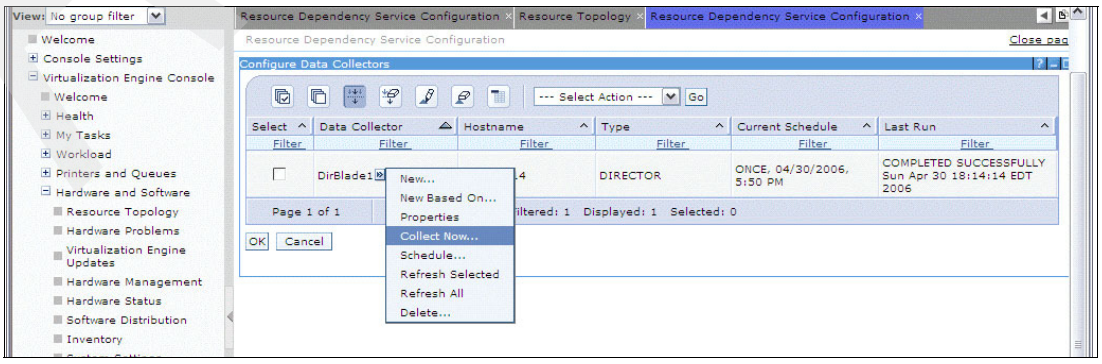


Figure 7-64 Resource Dependency Service data collector: Collect Now task

2. Specify if you want to start all data collectors or only the selected data collector, as shown in Figure 7-65. Click **Collect** and the run starts.

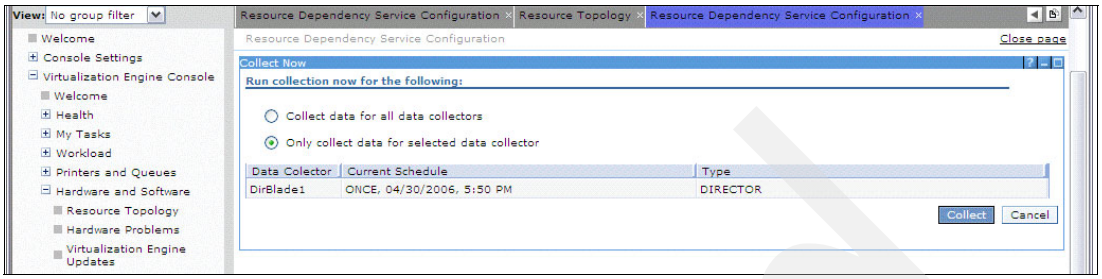


Figure 7-65 Resource Dependency Service data collector: Collection Now

Results from the IBM Director collection

The topology view of the run with the IBM Director collector now shows an increased number of systems compared to the view obtained with the IDD scan (see Figure 7-66). The IBM Director collector provides the serial number and model information for more systems. Thus, the resource ID can be generated from ODIRM.

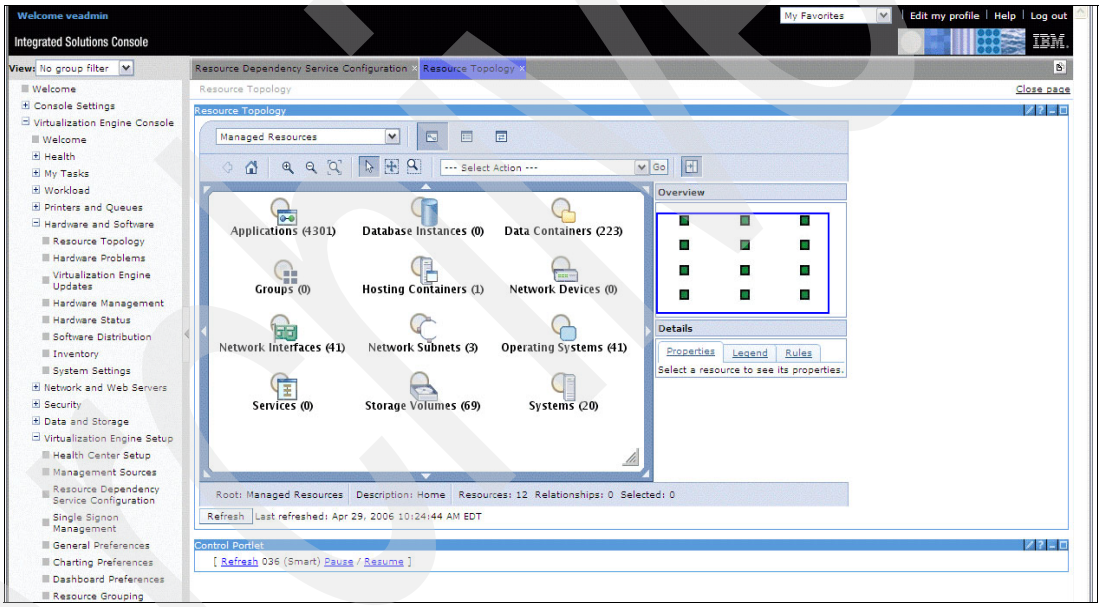


Figure 7-66 Collector run result

When we compare the number of discover systems within the IBM Director GUI (Figure 7-12 on page 228), we see that all level two systems are now visible in the Resource Dependency Service topology view. The table view of the systems (Figure 7-67 on page 257) contains 64 systems. Some of the available information for the systems are not sufficient to generate a unique resource ID. The correlation engine identified some of the entries as duplicated and merges these entries into the ODIRM database. Therefore, only 20 systems are visible in the topology view.

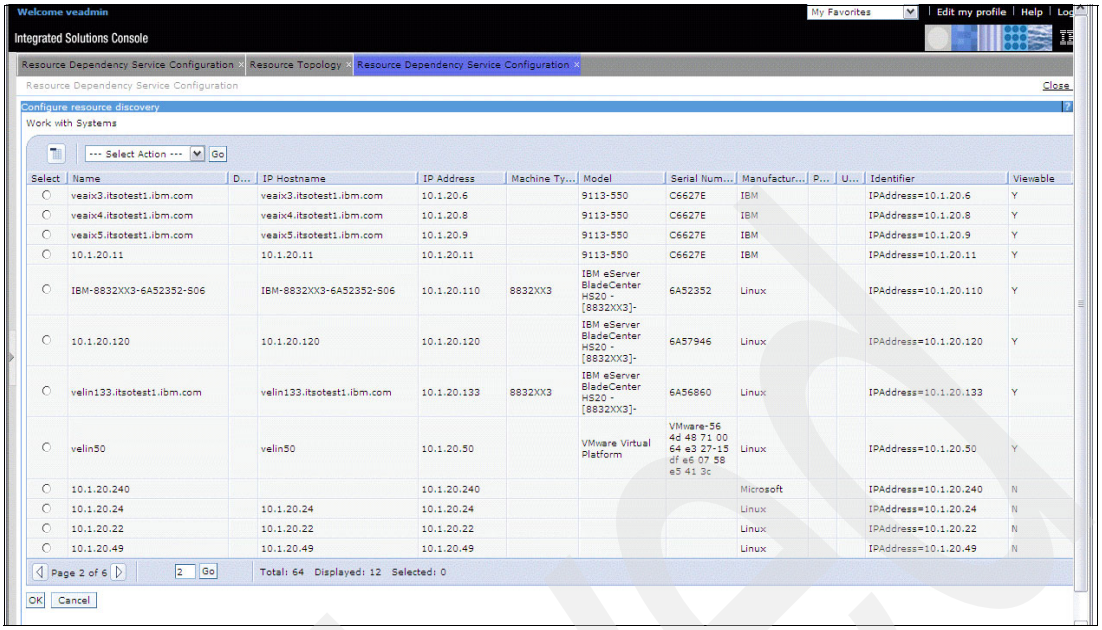


Figure 7-67 Combined IDD campaign and IBM Director collector results

7.5 Managing Resource Dependency Service

Resource Dependency Service consists of the IDD and the VEMS component, which are WebSphere applications or use a WebSphere Application Server infrastructure. The topology view of Resource Dependency Service resides within the Virtualization Engine console, which is an Integrated Solution Console (ISC) application. The ISC is a WebSphere portal application and resides in a WebSphere Application Server. This means that we have to manage three WebSphere Application Server profiles. The administration tools and the log files of these components are in different locations. The following section give an short overview of the administrative tools and commands for managing the Resource Dependency Service installation.

7.5.1 Intelligent Device Discovery

The IDD component is installed in the /opt/idd directory. The bin directory contains the administrative tools of IDD. The iddAdmin tool offers basic administration functions. Example 7-1 shows the list of the available functions.

Example 7-1 IDD administration tool

```
velin120:/opt/idd/bin # iddAdmin
iddAdmin: unknown command ''
IDD admin usage : iddAdmin <command>
Command is one of the following:
 help             : Show this message and exit.
 status           : Display status of IDD services.
 start            : Start IDD services.
 stop             : Stop IDD services.
 restart          : Restart IDD services.
 update           : Update IDD software.
 update-nsa       : Update NSA config files.
 log <level>      : Set logging level to: SEVERE | WARNING | INFO | FINE
```

```
 trace <level>    : Set tracing level to: ON | OFF
 flush            : Flush Workspace internal state and cancel all pending tasks.
 clean            : Flush, then delete all runtime logs and temporary files.
 workfor <servers>: Update the workgroup servers this machine works for.
```

Starting IDD servers

You can start the IDD processes with the **/opt/idd/bin/iddAdmin start** command. The output of this command is shown in Example 7-2.

Example 7-2 IDD admin start

```
velin120:/opt/idd/bin # ./iddAdmin start
IDD services status:
  * workspace                                                    ready
  * server                                                       ready
  * was                                                          ready
```

Another way to start IDD is by executing the idd script with the **/etc/init.d/idd start** command.

These procedures start the IDD server, the IDD workspace process, and the WebSphere server profile SERVER1.

Resetting IDD

During the testing phase, you may want to reset and cancel all IDD tasks. To do this, on the command line, enter the **/opt/idd/iddAdmin flush** command. IDD stores all history and log data in a table. Therefore, you have to clean the history table. The **/opt/idd/iddAdmin clean** command deletes all the log entries.

Note: The cleaning process of IDD requires some time. **iddAdmin clean** should only be used in extreme cases where a user needs to wipe out the history, tasks, and log tables. Usually, an **iddAdmin flush** is all that is required to reset and clean up IDD.

IDD command-line tool

The /opt/idd/bin directory also contains the command-line interface iddCli of IDD, which provides functions on the campaign and task level. Example 7-3 shows the available options of the iddCli tool. Resource Dependency Service uses this interface to create, drop, or change the campaigns.

Example 7-3 IDD command-line tool

```
velin120:/opt/idd/bin # iddCli
03 May 2006 11:00:13,509 [main] INFO    DDKIDD000134I: iddCli (com.ibm.idd.AUDIT)
Usage: <command> [<arguments>] [<options>]
Supported commands:
 help [<command>]
 campaign create <campaign> [-b <basecampaign> [-A|-assets all|confirmed]
[-i|-inventory] [-a|-address] [-c|-cabling] [-p|-people]] [-t
LongRunning|Default|ShortRunning] [-u <user>]
 campaign drop <campaign> [all|tables|files] [-u <user>]
 campaign defs <campaign> <category> {get <key> [<val>] | set <key> <val> | delete
<key> | list} [-u <user>]
 campaign exclude <campaign> <iprange> [-u <user>]
```

```
 certificate import <password> <filename> [<alias>]
 certificate export <password> [<filename>]
 certificate exchange <password> <hostname> <ssh_user> <ssh_passwd>
<remote_storepwd> [--remote-alias <alias>]
 certificate list <password> [all|own|trusted]
 database statement <sql>
 database query <sql> [--with-metadata]
 server info [all|version|interfaces|dir]
 server credentials <accesspwd> {get <key> [<val>] | set <key> <val> | delete
<key> | list}
 task submit <svc> <campaign> [-u <user>] ['parameter=value; ...']
 task pause|continue|stop <taskid> [-u <user>]
 task status <taskid>
```

IDD server information

Using the IDD command-line interface, you can get a summary of the IDD services and the location of the Windows Helper system. Example 7-5 on page 260 shows the server information output of our Resource Dependency Service installation.

Example 7-4 IDD command-line tool server information

```
velin120:/opt/idd/bin # iddCli server info all
03 May 2006 11:04:51,622 [main] INFO    DDKIDD000134I: iddCli 'server' 'info'
'all' (com.ibm.idd.AUDIT)
IDD Version 3.2 (Build 12022) / Linux/2.6.5-7.244-bigsmp/x86 / J2RE 1.4.2 IBM
build cxia32142-20050609 (JIT enabled: jitc)

IP addresses : 10.1.20.120, [fe80:0:0:0:20d:60ff:fe4e:834a]

Hostname     : velin120.itsotest1.ibm.com
Localname    : velin120
Domains      : itsotest1.ibm.com
DNS Servers  :

Address         Aliases, Version
-----------------------------------------------------
10.1.20.120, fe80:0:0:0:20d:60ff:fe4e:834a      velin120.itsotest1.ibm.com,
velin120, 10.1.20.120, fe80:0:0:0:20d:60ff:fe4e:834a         working for:
127.0.0.1, 127.0.0.1      IDD version: 3.2 (Build 12022)
                services: combiner.baddress, combiner.bannerscans,
combiner.bcabling, combiner.binventory, combiner.bpeople, combiner.cimscans,
combiner.dnswalker, combiner.hostscans, combiner.httpscans, combiner.iprange,
combiner.netflowscans, combiner.ntdomains, combiner.ospfscans,
combiner.pcharscans, combiner.snmpscans, combiner.sqlrules, combiner.toposcans,
combiner.vulnascans, combiner.winscans, db2db, egatherer.cleaner,
egatherer.xml.cleaner, exporter, importer, input, logdrainer

10.1.20.14      blade1.itsotest1.ibm.com, blade1, 10.1.20.14            working
for:    IDD version: 3.2 (Build 12022)
                services: wmiscan

10.1.20.120, fe80:0:0:0:20d:60ff:fe4e:834a      velin120.itsotest1.ibm.com,
velin120, 10.1.20.120, fe80:0:0:0:20d:60ff:fe4e:834a         working for:
127.0.0.1  IDD version: 3.2 (Build 12022)
```

```
              services: dnswalker, hostscanner, httpscan, ipscan, ldap, netflow,
nmap, nsa, nsavulna, ospffeeder.collector, slpscan, snmpwalk, ssh.idduxscan,
sysscan
```

Checking IDD task status

During testing, you may want to know the status of an IDD task, which is not visible in the Resource Dependency Service GUI. To get the task information, perform the following steps:

1. Log in as the IDD admin user. We use the default user vedb2ur.
2. Connect to the IDD database by typing **db2 connect to IDD** on the command line.
3. Submit the following SQL statement to get the task ID:

   ```
   db2 "select taskid, services, campaign from idd.tasks where ustate='running'"
   ```

4. Mark the task ID and copy it into the command **/opt/idd/bin/iddCli status *taskid***.

   Example 7-5 presents a command line interface task status output for a combiner task.

Example 7-5 IDD task status

```
velin133:~ # iddCli task status 20060503140540031-1.2.5
03 May 2006 14:26:51,070 [main] INFO    DDKIDD0O0134I: iddCli 'task' 'status'
'20060503140540031-1.2.5' (com.ibm.idd.AUDIT)
03 May 2006 14:26:55,047 [main] INFO    DDKIDD0O0240I: Connected to Database
jdbc:db2://velin133.itsotest1.ibm.com:50001/IDD:retrieveMessagesFromServerOnGetMes
sage=true; DB2/LINUX SQL08022 using IBM DB2                              JDBC
Universal Driver Architecture 2.6.80 (com.ibm.idd.db.ConnectionInformation)
TaskId:      20060503140540031-1.2.5
State:       running
Campaign:    UPDATE_CAMPAIGN
User:        RMS
Service:     combiner.httpscans
Description: Combiner of Httpscans
Schedule:    start=May03 14:06:43; stop=May03 14:56:43; maxIters=40;
period=00:01:00; isFixedDelay=true
Concurrency: 1
Burstsize:   10
Parameters:  {destServerList=*, paramRange=SELECT h.* FROM {campaign}.HTTPSCANS h
where h.id NOT IN (SELECT inputid FROM {campaign}.COMBINERLOG WHERE inputtable =
'HTTPSCANS') ORDER BY h.recorded, param
Type=mapArray}
Statistics:  inserts=0 ; error = false ; done = 45 %
```

In campaign table ASSETLOG, a summary of the scans or tasks is given. Use this to analyze the campaign actions in the network.

Archived

Getting IDD Program Identification

The Program Identification (PID) files are in the /var/run directory. Using the **cat** command, you can list the PID of the IDD server and workspace processes, as shown in Example 7-6.

Example 7-6 Getting IDD PID

```
velin120:/var/run # cat /var/run/idd-server.pid
32386
velin120:/var/run # cat /var/run/idd-workspace.pid
32365
```

7.5.2 Virtualization Engine management server

Virtualization Engine management server (VEMS) is the WebSphere Application Server profile of Resource Dependency Service. The HOME directory of the profile is /opt/IBM/VE2/.

Checking server status

You check the server status with the statusServer.sh script, which is located in the bin directory of the profile home directory. The script requires the user ID and password of the administrator user as an input parameter. Example 7-7 shows an example of the command result.

Example 7-7 Checking the VEMS server

```
velin120:/opt/IBM/VE2/VEMS/bin # ./serverStatus.sh VEMS -username veadmin
-password password
ADMU0116I: Tool information is being logged in file
           /opt/IBM/VE2/VEMS/logs/VEMS/serverStatus.log
ADMU0128I: Starting tool with the VEMS profile
ADMU0500I: Retrieving server status for VEMS
ADMU0508I: The Application Server "VEMS" is STARTED
```

Starting VEMS

To start the VEMS server, invoke the /opt/IBM/VE2/VEMS/startServer.sh VEMS script. The output of the start command is shown in Example 7-8.

Example 7-8 Starting the VEMS server

```
velin120:/opt/IBM/VE2/VEMS/bin # ./startServer.sh VEMS
ADMU0116I: Tool information is being logged in file
           /opt/IBM/VE2/VEMS/logs/VEMS/startServer.log
ADMU0128I: Starting tool with the VEMS profile
ADMU3100I: Reading configuration for server: VEMS
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server VEMS open for e-business; process id is 6123
```

Stopping VEMS

To stop the VEMS server, use the **/opt/IBM/VE2/VEMS/bin/stopServer.sh VEMS -username user -password password** command. Example 7-9 on page 262 shows the output of the command.

Example 7-9 Stopping the VEMS server

```
velin120:/opt/IBM/VE2/VEMS/bin # ./stopServer.sh VEMS -username veadmin -password
password
ADMU0116I: Tool information is being logged in file
           /opt/IBM/VE2/VEMS/logs/VEMS/stopServer.log
ADMU0128I: Starting tool with the VEMS profile
ADMU3100I: Reading configuration for server: VEMS
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server VEMS stop completed.
```

VEMS administration console

To launch the administration console, open a browser and enter the following URL:

`http://`*`yourhost`*`:18552/ibm/console`

Within the administration console, under Troubleshooting → Runtime Messages, you can find the error, warning, and information messages. Figure 7-68 shows a list of the runtime information messages. You can click the message text to get further information about the message.

If you do not want to use the WebSphere administration GUI, you can directly look into the log files located in the /var/log/VEServices/VEMS directory. In this directory, you can find the PID file of VEMS.

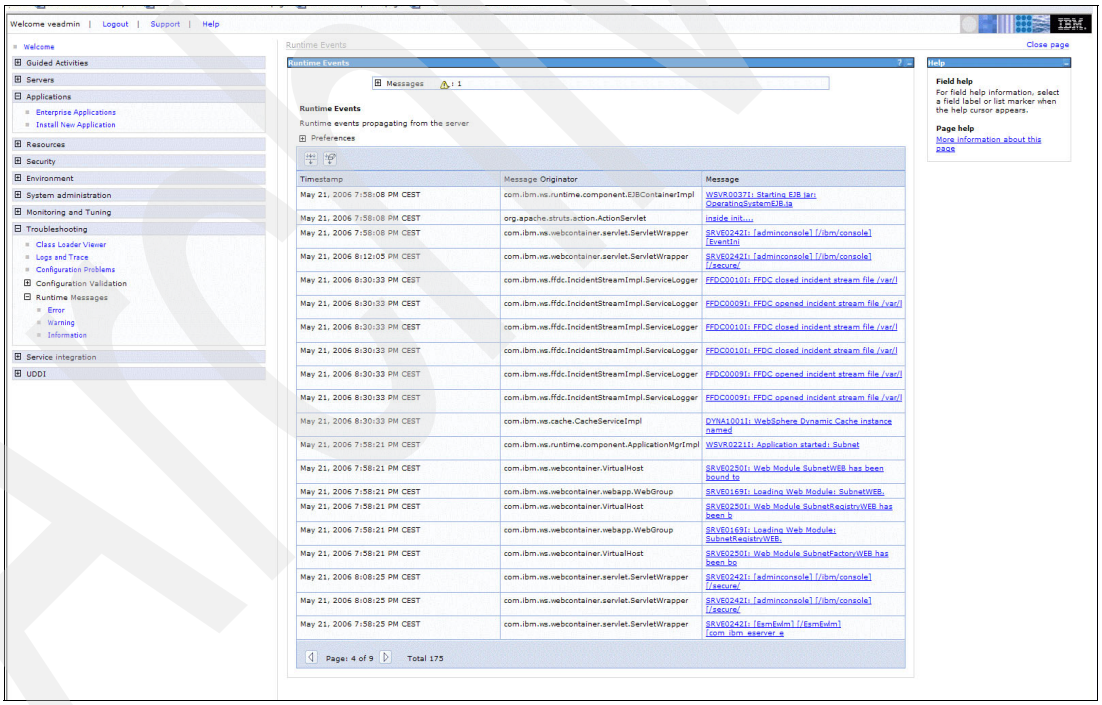


Figure 7-68 VEMS administration console

7.5.3 Integrated Solution Console

Here we discuss the Integrated Solution Console.

Stopping the Virtualization Engine console

The topology view of Resource Dependency Service runs within the Virtualization Engine console, which is in turn an ISC application. If you want to stop the ISC, enter the following command:

```
/opt/IBM/VE2/ISC/PortalServer/bin/stopISC.sh ISC_Portal userid password
```

Note: Specify the full qualified stop ***file*** **stopISC.sh** command. Otherwise, the command will fail with the error: `/stopISC: line 44: /opt/IBM/VE2/ISC/PortalServer.sh/bin/stopServer: No such file or directory`

Starting the Virtualization Engine console

The start command for ISC is **/opt/IBM/VE2/ISC/PortalServer/bin/startISC.sh ISC_Portal**. The PID file is /opt/IBM/VE2/p/VEISC/logs/ISC_Portal.pid.

Virtualization Engine console administration tool

The Virtualization Engine console is an application within the ISC and runs in the VEISC WebSphere profile. The administrative console of this profile is not started by default. To start it, enter the following command:

```
/opt/IBM/VE2/WebSphere/Appserever/bin/startServer.sh server1 -profileName VEISC
-username userid -password -yourpw
```

After starting the administrative console, open a browser and enter the following URL:

`http://`*`hostname`*`:8437/ibm/console`

Log in with the user ID and password of your installation.

7.5.4 DB2

The DB2 database log files is in the /var/log/db2 directory. You can use the DB2 standard tools to manage DB2.

7.6 Troubleshooting

- Within the Resource Dependency Service configuration windows, no log or progress information about the running IDD scan is available. Only the status of the campaign, such as running, stopped, and so on, is presented. During the configuration of new campaigns and in the testing phase, it may be worthwhile to get more information about the running campaigns.

  The IDD logs are stored in /opt/idd/logs.

  The IDD server log file contain the scan messages. Use this file to understand the IDD actions. It also helps in debugging error situations. Example 7-10 on page 264 shows the log entry of IDD for an SSH scan with wrong credentials. IDD wants to connect to server 172.30.40.7 with a credential configured within Resource Dependency Service with the name 172.30.20.22. The IDD added the type of the credential, in this case ssh, in front of the credential name as configured in Resource Dependency Service.

Example 7-10 Extract of IDD server log file

```
29 May 2006 20:10:25,717
[ssh.idduxscan[2/10]-[wms://localhost/idd.workspace/workspace?service=ssh.idduxsca
n&jmsid=0a84e9f4fb3b485decee6c73110a134f0000000000000001]] INFO    DDKIDD014016I:
Could not connect to 172.30.40.7 with given credential ssh_172.30.20.22.
(com.ibm.idd.AUDIT)
```

- Configure only one or a small part of your network for the first scans, because a large scan may run for a long time. Check the log files for errors and look into the IDD results to get more experience in how IDD works.
- Inconsistencies of the IDD task status occur in the case of IDD crash situations. To resolve this, clean the IDD task table according to the procedure described in 7.5, “Managing Resource Dependency Service” on page 257.

  If you clean the IDD tables with the iddAdmin tool, Resource Dependency Service cannot determine the status of a campaign. Therefore, in the Resource Dependency Service configuration window, the campaign status of unknown is displayed. As soon as the campaign is restarted by configuring a new schedule time, the status information will be available within the GUI.

  During the cleanup of the IDD Database, we got DB2 errors indicating that the transaction log is full. If this error situations occurs often, consider changing the LOGSIZ parameter of the IDD database.
- The Resource Dependency Service table CONFIGURATION_INFORMATION contains the parameter of the Resource Dependency Service collector runs and the scheduler status. During our tests, we experienced a problem where Resource Dependency Service was unable to start a new data collector run. This was because an old collector run was not finished correctly and the run status within the table was not updated. In this case, we manually removed the status entry of this run from the table.
- In case of database connection problems, check the Resource Dependency Service SSO configuration. We experienced connection problems after the Resource Dependency Service installation due to an improperly configured Resource Dependency Service SSO. Another source of database connection problems can be wrong DB2 configuration parameters. Check whether the recommended DB2 parameter are configured according to the latest information provided in the InfoCenter.

Archived

8

Using the Virtualization Engine console to visualize the infrastructure environment

The main advantage provided by Virtualization Engine console, visualizing the complete infrastructure, is meant for system administrators. It helps them take the necessary action as and when required. For example, system administrators should know where resources are physically being shared. If a shared media drive is failing, they have to locate the DVD drive in order to replace it. For example, in a homogeneous environment, IBM System p administrators probably view resources through the Hardware Management Console (HMC). If the administrators use a variety of operating systems, they may also have iSeries Navigator or a Linux management tool in the environment. As resources become virtual across multiple operating systems, administrators benefit from using one interface to help simplify the management.

The Topology Viewer is an integral part of the Virtualization Engine console. It helps visualize the resources along with their dependencies and relationship with other resources. What you see in the Topology Viewer is basically the data captured through various scanning processes, such as the Intelligent Device Discovery/Collectors and stored in the DB2 databases. The most important databases are ODIRMDB and RMS1.

With Topology Viewer, you can do the following:

- Display the resources in a Graphical User Interface (GUI).
- Understand the side effects associated with the change in the resource status.
- The discovery of resources ultimately populates the resources and relationships shown in the topology view.

© Copyright IBM Corp. 2006. All rights reserved.

8.1 Accessing resource topology using Virtualization Engine console

To access the resource topology, perform the following tasks:

1. Open your Web browser and specify the following:

   `http://<hostname or IP address>:8421/ibm/console`

   The login window opens as shown in Figure 8-1, where 8421 is the default port number for the Virtualization Engine console. Enter the user ID and password. Click **Log in**.

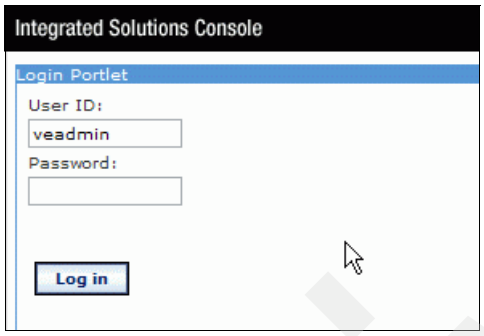


Figure 8-1 Login window

2. In the new window that opens, in the navigation window on the left, expand **Virtualization Engine Console** → **Hardware and Software** → **Resource Topology**, as shown in Figure 8-2. Click **Resource Topology**.



Figure 8-2 Navigation window on the left

Archived

3. The Virtualization Engine console displays the window shown in Figure 8-3.

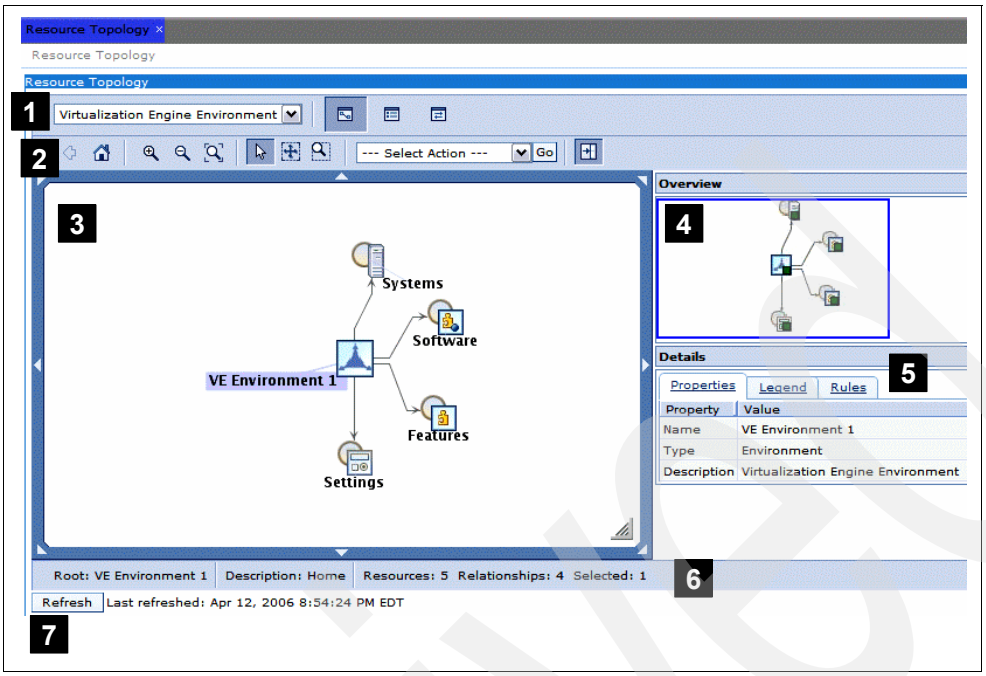


Figure 8-3 Virtualization Engine environment

This window is made of multiple fields:

– 1 is a drop-down menu to select between Virtualization Engine Environment or Managed Resources.

> **Note:** You must have Resource Dependency Service installed and configured before beginning to use the Managed Resource data source.

– 2 is the navigation toolbar. This helps you navigate within a particular environment, for example, Zoom Out, Zoom In, Toggle, and so on, as shown in Figure 8-4.

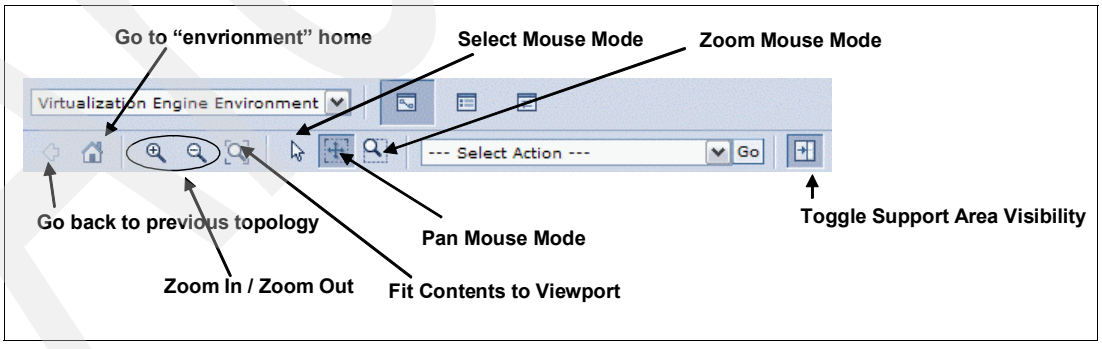


Figure 8-4 Navigation toolbar

– 3 is the topology window. This displays the resources and their dependencies and relationship with other resources. This window can exhibit two views: a Radial view or a Tree view. To switch between the views, click anywhere in the topology window, right-click, and select **Layout** → **Radial/Tree**, as shown in Figure 8-5.

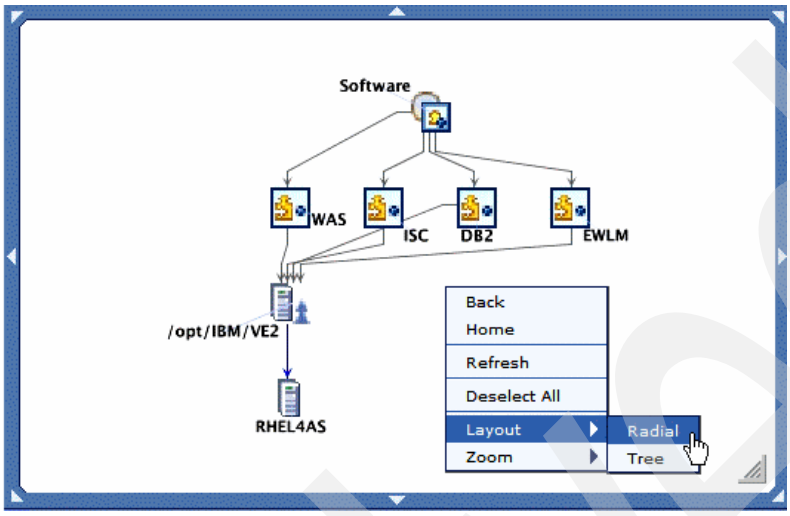


Figure 8-5 Topology window showing resources in tree view

The radial view shows the resources shown in Figure 8-6.

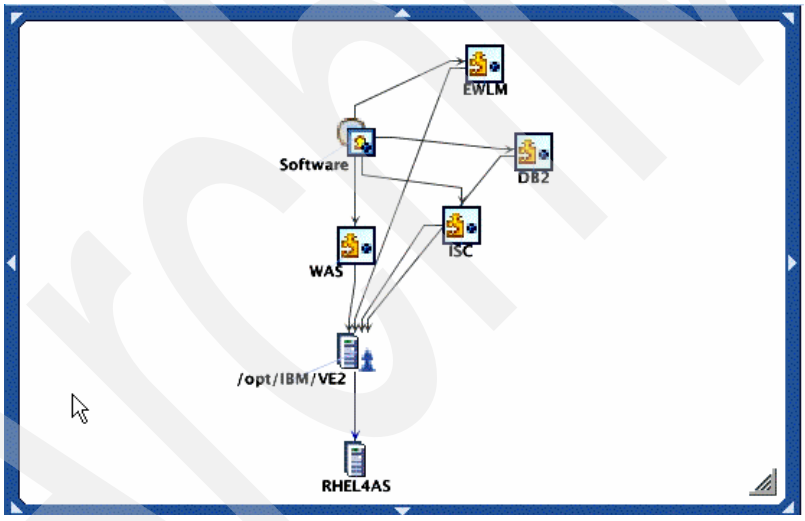


Figure 8-6 Topology window showing resources in radial view

The information displayed in the topology view varies depending on the type of environment you select, either the Virtualization Engine environment or the Managed Resources. Table 8-1 explains the components displayed when you select the Virtualization Engine environment.

Table 8-1 Components displayed when Virtualization Engine environment is selected

Component	Description
Systems	The computer systems where Virtualization Engine features and components have been installed.
Features	Virtualization Engine offerings, for example, Enterprise Workload Manager (EWLM), Virtualization Engine console, that comprise multiple components.

Component	Description
Software	Pieces of software, for example, Integrated Solution Console (ISC), WebSphere Application Server, and so on, that come together to form features.
Settings	Data shared by Virtualization Engine applications, for example, security information.

Table 8-2 explains the components you view when you select Managed Resources.

Table 8-2 Components displayed when Managed Resources is selected

Component	Description
Applications	Programs that run within a hosting container or an operating system interface.
Database Containers	Similar to file systems, data containers contain a collection of files and attributes associated with these files. They allow applications to create and manage files on the storage devices. Data containers are contained on a system or remotely mounted from a file server.
Data Instances	A collection of interrelated or independent data items stored together to serve one or more applications.
Groups	Use groups to organize unrelated resources, creating personal views of specific resources from your environment.
Hosting Containers	Programs or applications that host other applications, for example, application servers host applications that are served up on the Internet.
Network Devices	Special purpose, dedicated computers, such as routers and load balancers, that help manage network traffic.
Network Interface	A network interface (also known as Network Interface Card or NIC) provides the interface control between a system and an external high-speed network.
Network Subnet	A subnet is an interconnected cut independent segment of an entire network that is identified by its Internet Protocol (IP) address.
Operating Systems	A collection of system programs that control the overall operation of a computer system: Windows XP, Linux, i5/OS, and so on.
Services	A process or subprocess used to automate common tasks, such as verifying security credentials or calibrating system clocks. Services contain resources, including other services, used by the process.
Storage Volumes	A storage volume represents a medium used for storage of data on a computer. A storage volume can be an entire physical hard disk or portions of a hard disk.
Systems	A system represents physical computers or virtual computers, such as logical or physical partitions.

Archived

– 4 is the Overview Area. It shows the miniaturized view A of the partial topology B, as shown in Figure 8-7.

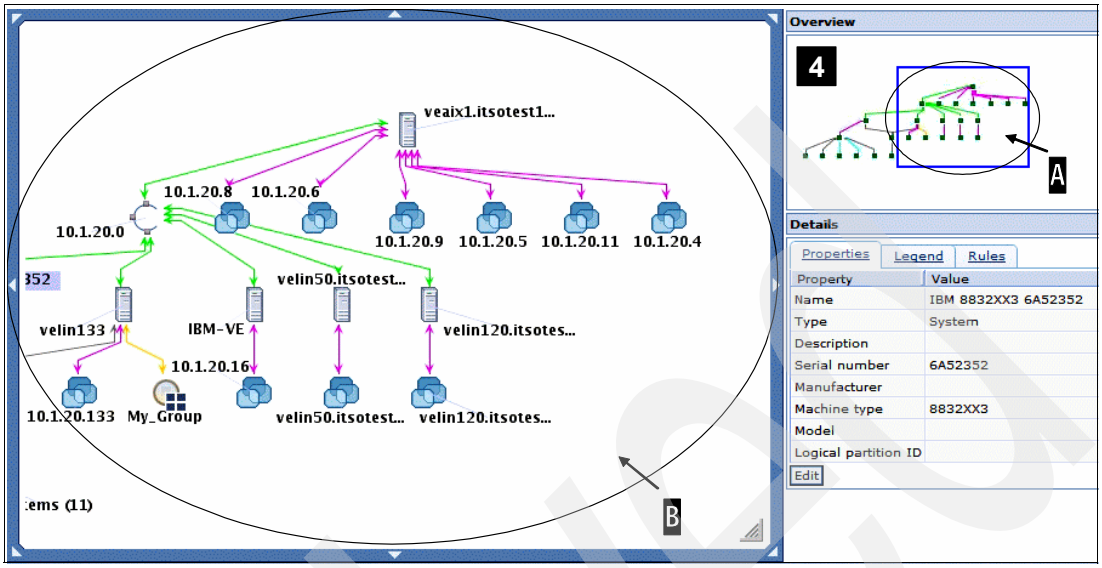


Figure 8-7 Overview

– 5 is the details window and contains three tabs:
  - Properties: When you select a resource or relationship in the topology, all its properties appear in the properties pane, which is editable. You can change the properties of a resource without opening another work page.
  - Legend: Graphical representations of relationships are displayed by colored lines. The description of each type of colored line listed in the legend can be used to differentiate relationship types.
  - Rules: Relationship rules describe all the combinations of node and relationship types that can possibly appear in the current topology.
– 6 is the status line. It displays the name of the resource that is currently the root node. It shows the count of resources or relationships currently shown and the number selected. When you select the Managed Resource data source, the status line also displays a selectable depth menu that allows you to display the number of relationships from the root node to the furthest object away from that node.
– 7 is the Refresh button. Click **Refresh** to update the current view of the topology. Any changes made to the underlying data source of the topology is reflected in the updated view when you click this button.

8.2 The Virtualization Engine environment

The Virtualization Engine environment data source allows you to explore the relationship between Virtualization Engine resource types, such as computer systems, settings, software, and features. All the data displayed in the Virtualization Engine environment is data that is stored in a common repository during the installation of the Virtualization Engine. A Virtualization Engine environment is a logical grouping of systems or partitions with the Virtualization Engine software installed on them.

As part of the Virtualization Engine installation, the installation wizard creates a configuration repository in a directory server, which is a Lightweight Directory Access Protocol (LDAP). The

wizard uses this configuration repository in the directory server to store information about your Virtualization Engine environment, including installation, configuration, and topology settings. Data from this environment is used to populate the resource topology viewer. The resource topology allows the visualization of this information but not the modification of the data.

To display the Virtualization Engine environment, perform the following six tasks:

1. Log in to Virtualization Engine console using Virtualization Engine credentials, as shown in Figure 8-1 on page 266.
2. Click **Resource Topology** in the navigation window, as shown in Figure 8-2 on page 266.
3. Select **Virtualization Engine Environment** from the drop-down list 1, as shown in Figure 8-3 on page 267. The Virtualization Engine environment shown in Figure 8-8 is displayed.

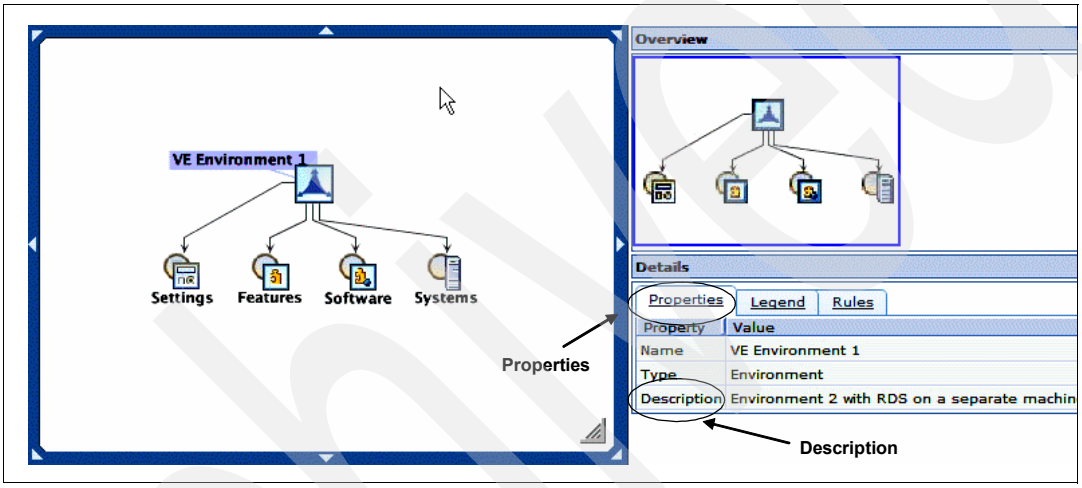


Figure 8-8 Virtualization Engine environment

Archived

Note: In our environment, since we had Resource Dependency Service installed separately, you see the Properties (Description).

For more details, select any resource, right-click, and select **Open** → **Next Level/All Levels**.

Note: Next Level takes you one level down in the hierarchy, while All Level shows you the expanded form, as shown in Figure 8-9.

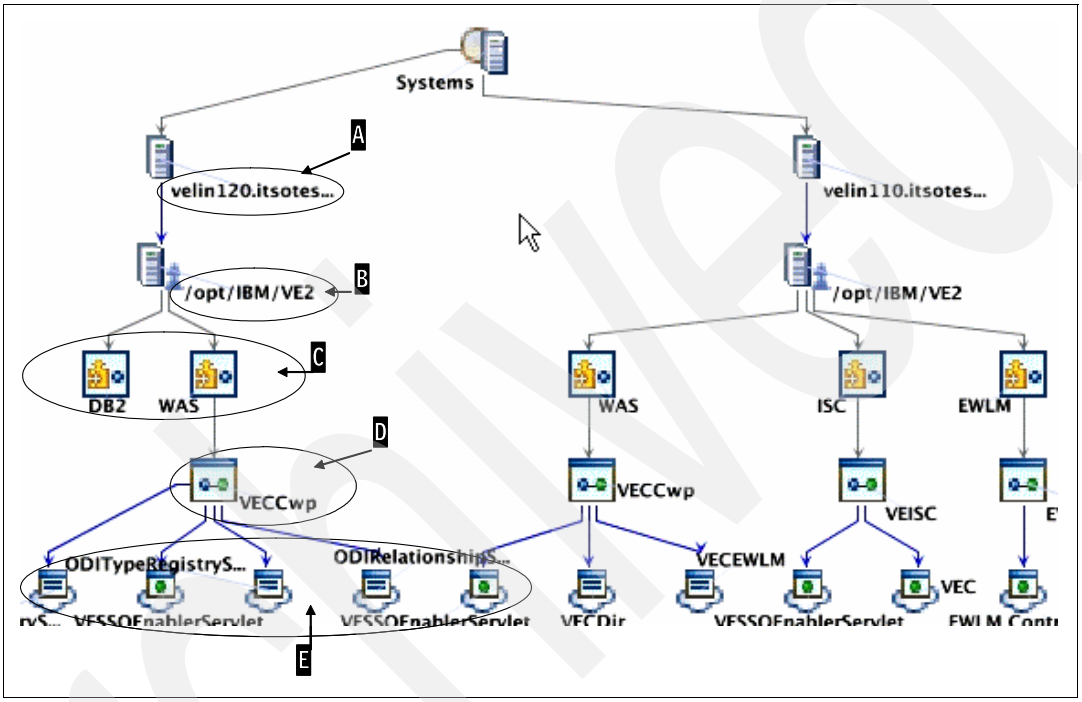


Figure 8-9 All level view of a systems resource

Figure 8-10 on page 273 shows that:

- There are two systems (velin110 and velin120) participating in the Virtualization Engine environment.
- Resource Dependency Service is installed on velin120 (A).
- The Virtualization Engine environment is installed in path /opt/IBM/VE2 (B).
- System velin120 has IBM DB2, IBM WebSphere Application Server, and Resource Dependency Service installed (C).
- System velin110 has IBM WebSphere Application Server, EWLM, and ISC installed.
- WebSphere Application Server installed on velin120 contains the VECCwp profile (D).
- VECCwp contains various services, such as ODITypeRegistryService, ODIRelationshipStoreService, and so on (E).

Archive

4. To check the *relationship* between velin120 and Systems, click **A** → **Read Information in the Properties pane** **B**, as shown in Figure 8-10.

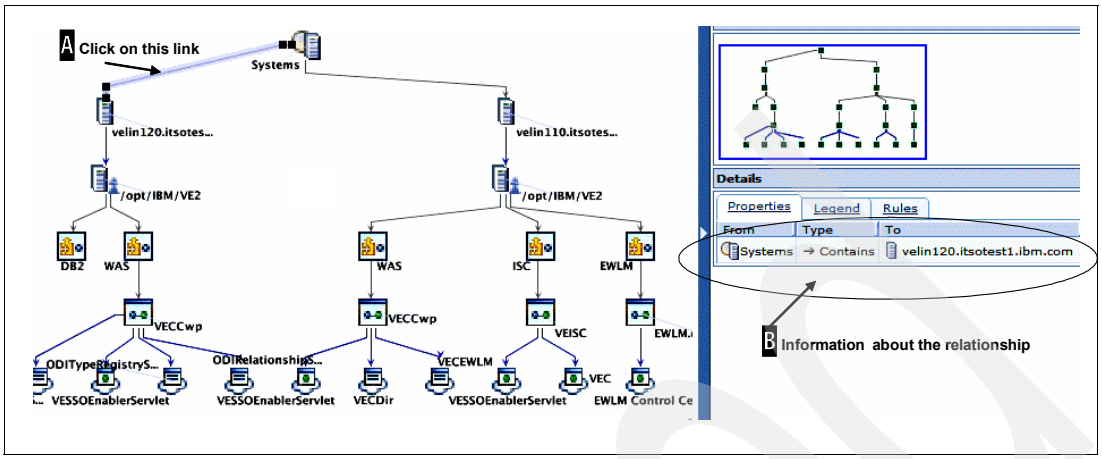


Figure 8-10 Relationship between resources

Note: It is important to see the Properties pane whenever you make any link selection. It provides you more descriptive information about the resource.

5. You can switch between the Systems View or Features View, as shown in Figure 8-11. Select **WAS**, right-click, and select **Open** → **Systems View**.

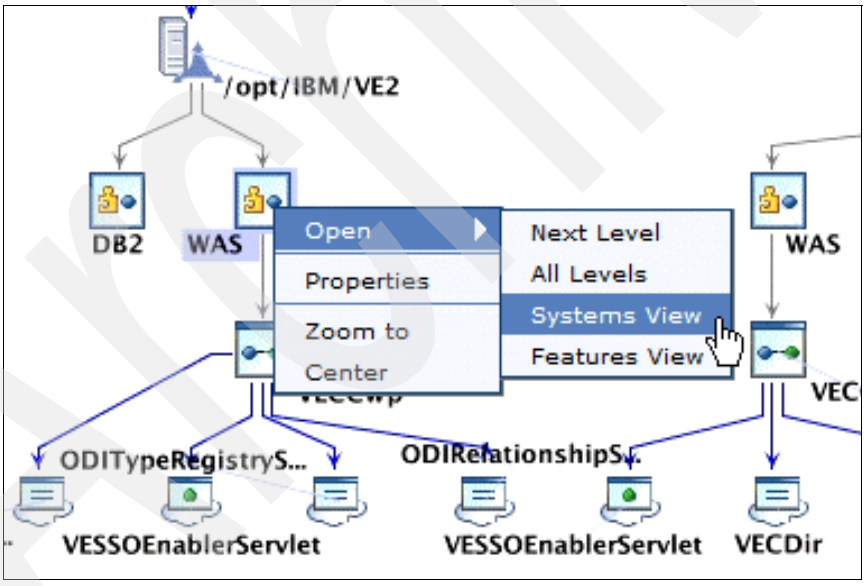


Figure 8-11 Choosing a systems view

Archive

The system opens, as shown in Figure 8-12.

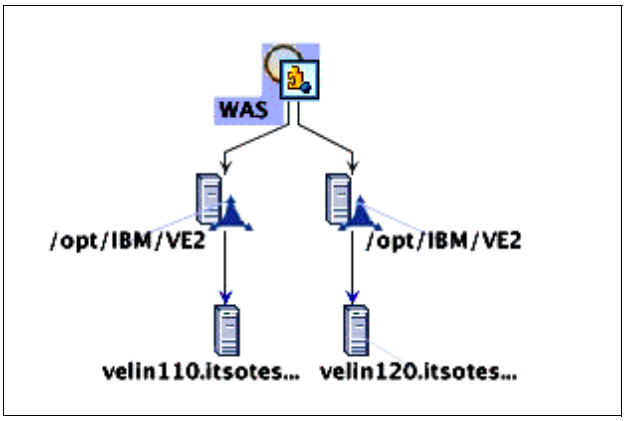


Figure 8-12 System view of the selected resource

6. Another important capability of the Virtualization Engine Management Environment is its ability to display port information. To see port information, select A and read the information in the Properties pane B, as shown in Figure 8-13.

Note: The property description field provides port information for easy management.

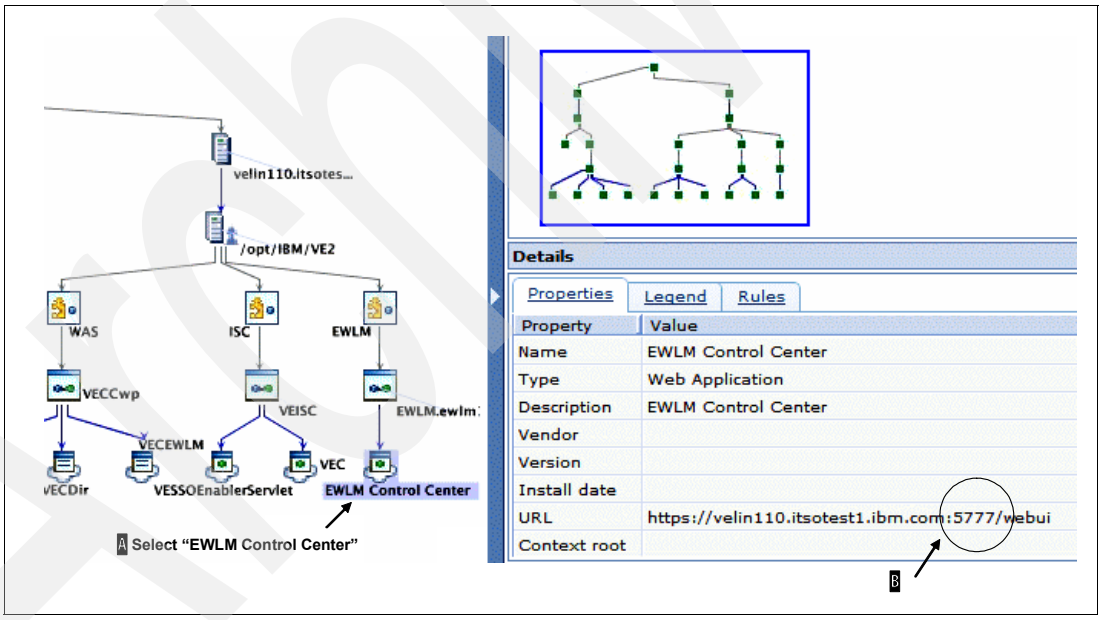


Figure 8-13 Port information about the selected resource

8.3 The Managed Resources environment

The Managed Resources environment displays the resources discovered by the Resource Dependency Service, such as network interfaces, services, applications, host containers, and the relationships between them. In addition, you can create managed resources, delete managed resources, create relationships, delete relationships, and edit resource properties.

The Resource Dependency Service uses a combination of data collection and data discovery to gather information about the resources associated with a business service and the relationships between these resources. This data allows you to quickly find and address issues associated with your resources, contributing to fast problem resolution.

To display the Managed Resources environment, perform the following tasks:

1. Log in to Virtualization Engine console using the Virtualization Engine credentials, as shown in Figure 8-1 on page 266.
2. Expand **Virtualization Engine Console** → **Hardware and Software** → **Resource Topology**, as shown in Figure 8-2 on page 266.
3. Select **Managed Resources** from the drop-down list 1, as shown in Figure 8-3 on page 267.

   You will see a collection of manageable resources, such as applications, operating systems, database instances, systems, services, and so on, in the topology window, as shown in Figure 8-14. A description for each of these resources is provided in Table 8-2 on page 269.

   What you see is the information captured by various scanners, such as Intelligent Device Discovery (IDD) and Collectors. In our scenario, there are 11 systems in the network, 11 operating systems installed, and eight applications running in the environment. Expand these resources further to get more details. For more information, refer to 8.3.1, “What is a resource?” on page 276. It is also possible *not* to see anything for any resource type. This can happen if nothing is discovered or if there is not enough information for that resource to show up in the topology pane.

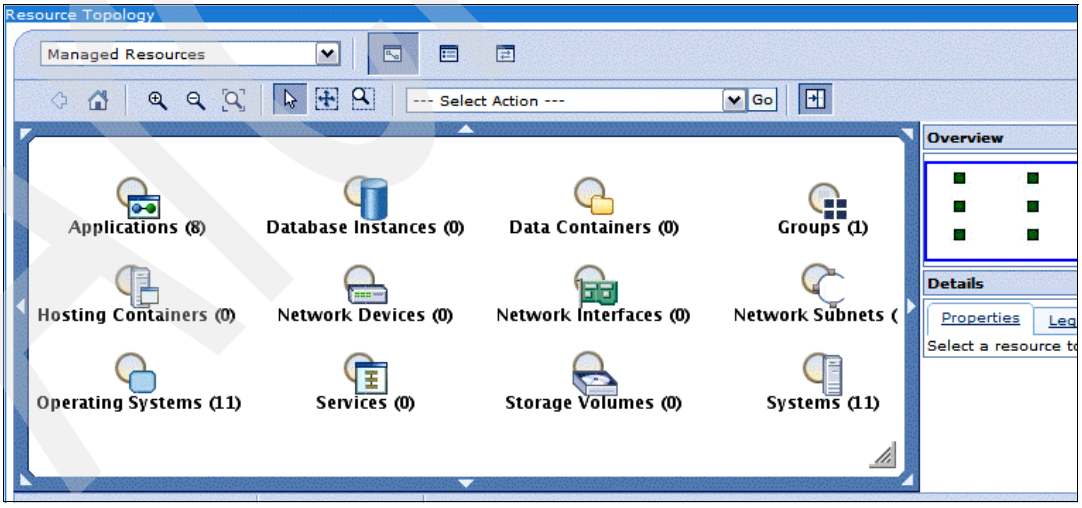


Figure 8-14 Managed resources environment

Note: The resource types are predefined, that is, applications, services, database containers, and so on. You cannot add a new resource type, but you can add new resources in a specific resource type.

When you select Managed Resources for the first time in the topology viewer, you may get an error, as shown in Figure 8-15. This error tells you that IDD scans or Collectors have not yet been run, and therefore there is nothing available to display.

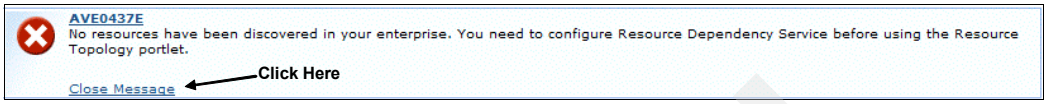


Figure 8-15 Error message

You can close the error message by clicking **Close Message**.

To get more details for any specific resource, select the resource, right-click, and select **Open → Search "resource type selected"** or **All "resource type selected"**

Sometimes it is important to refresh the topology window because you do not have the latest update on the viewer. To refresh the view, right-click in the topology pane and select **Refresh**. Figure 8-16 shows the refresh view.

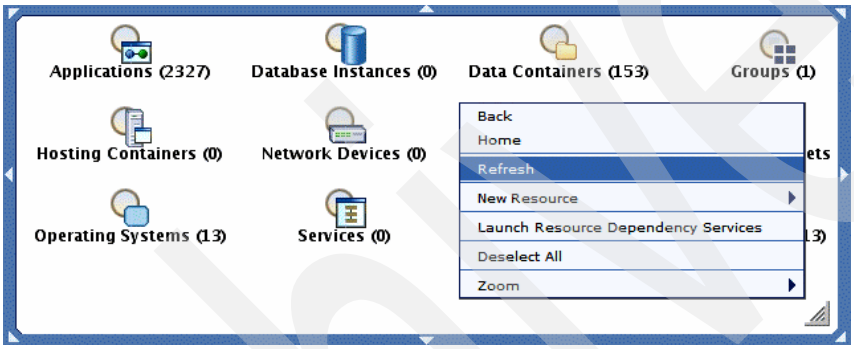


Figure 8-16 Refreshing the view

8.3.1 What is a resource?

A resource is a container that holds information of a specific type, for example, services, applications, storage volumes, and so on. Refer to Table 8-2 on page 269 for more details. For example, if you select an operating system, right-click it, and select **Open → All Operating Systems,** you will see all the operating systems discovered in the environment, as shown in Figure 8-17 on page 277.

An operating system is identified by an IP address. If you want to know what operating system has been installed on a specific machine, for example, 10.1.20.11, select A and read the information in the Details window, as indicated by B in Figure 8-17.

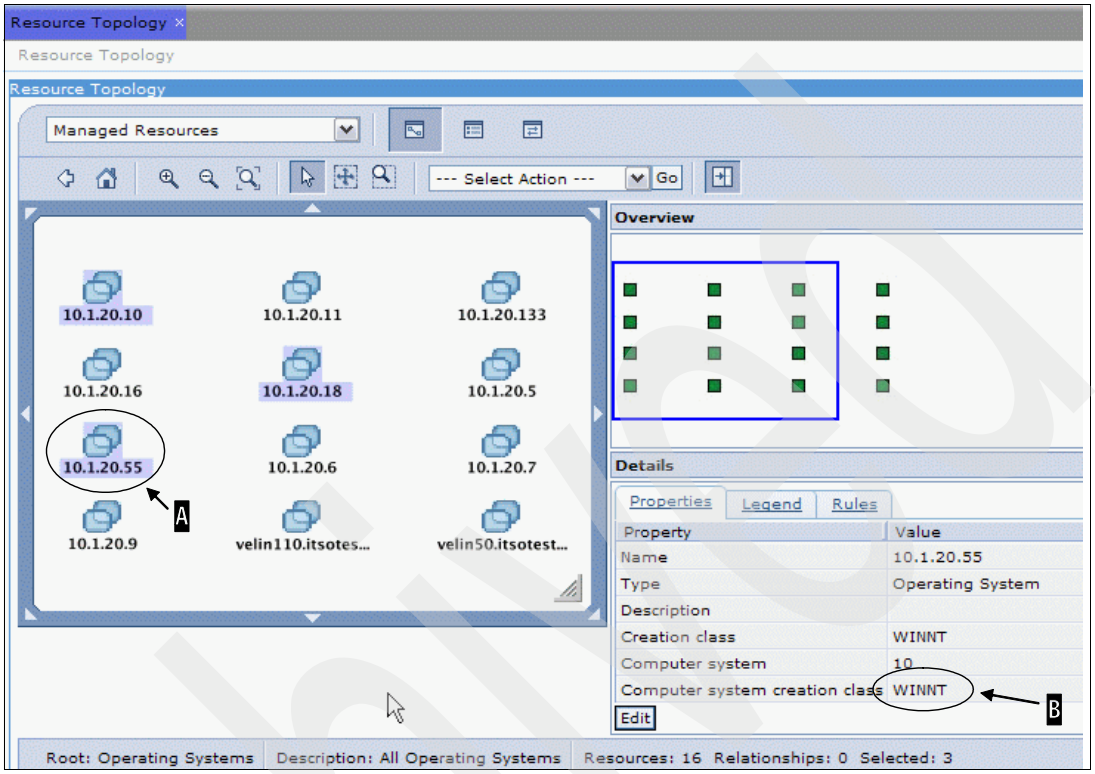


Figure 8-17 Operating system details

You can see that 10.1.20.11 has AIX installed.

In the same window, if you want to add more descriptive information, click the **Edit** button. A window opens, as shown in Figure 8-18. Type in more descriptive information here and click **OK**.

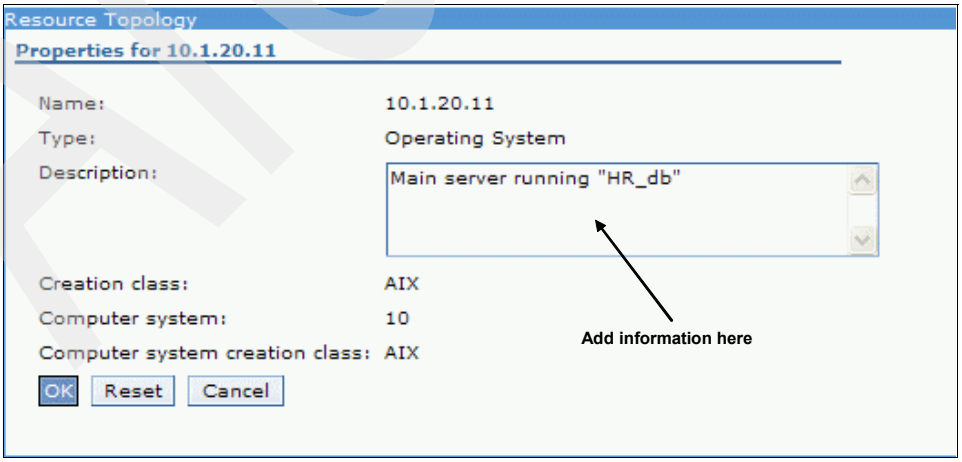


Figure 8-18 Edit screen

You see that the updated information is now visible on the Details window, as shown in Figure 8-19.

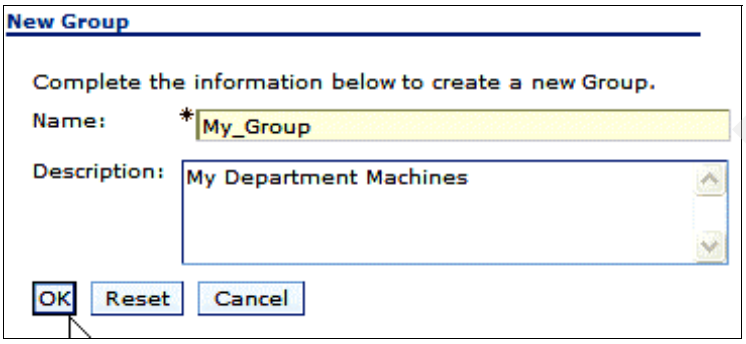


Figure 8-19 Updated information in the details window

This additional information helps you recollect that you have other tools or applications installed on the same machine.

Note: The editable contents vary from resource to resource, that is, the operating system only allows you to add description, while network interface allows you to add description, network address, switch port, and so on.

8.3.2 What is a filter?

Many times, when a particular machine has hundreds of applications installed, you may only be interested in knowing whether a particular software has been installed on it or not. In such a situation, filters are very useful. They allow you to find only the required information. For example, Figure 8-20 shows that 10.1.20.18 has 190 applications.

To define a filter, perform the following tasks:

1. Select **A**, right-click, and select **Applications**.

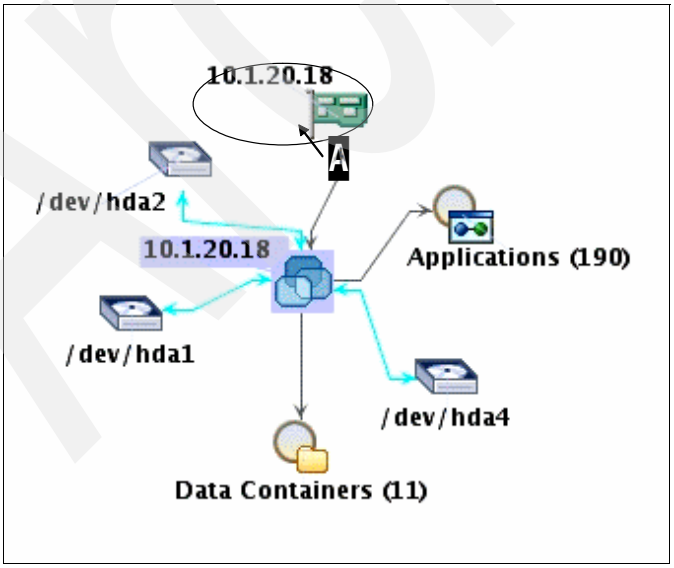


Figure 8-20 Discovered applications on a particular machine

2. Switch from **Show topology view** to **Show resources in table view**, as shown in Figure 8-21.

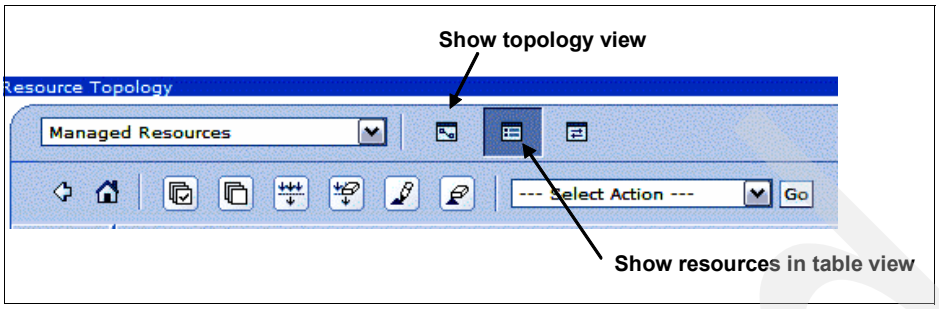


Figure 8-21 Switching to table view

3. After switching, you will see the resources in the table view, as shown in Figure 8-22. You can see that there are 10 pages of information available. In this case, a filter will help drastically. Click **B** and you will find that an additional tab, **C**, appears.

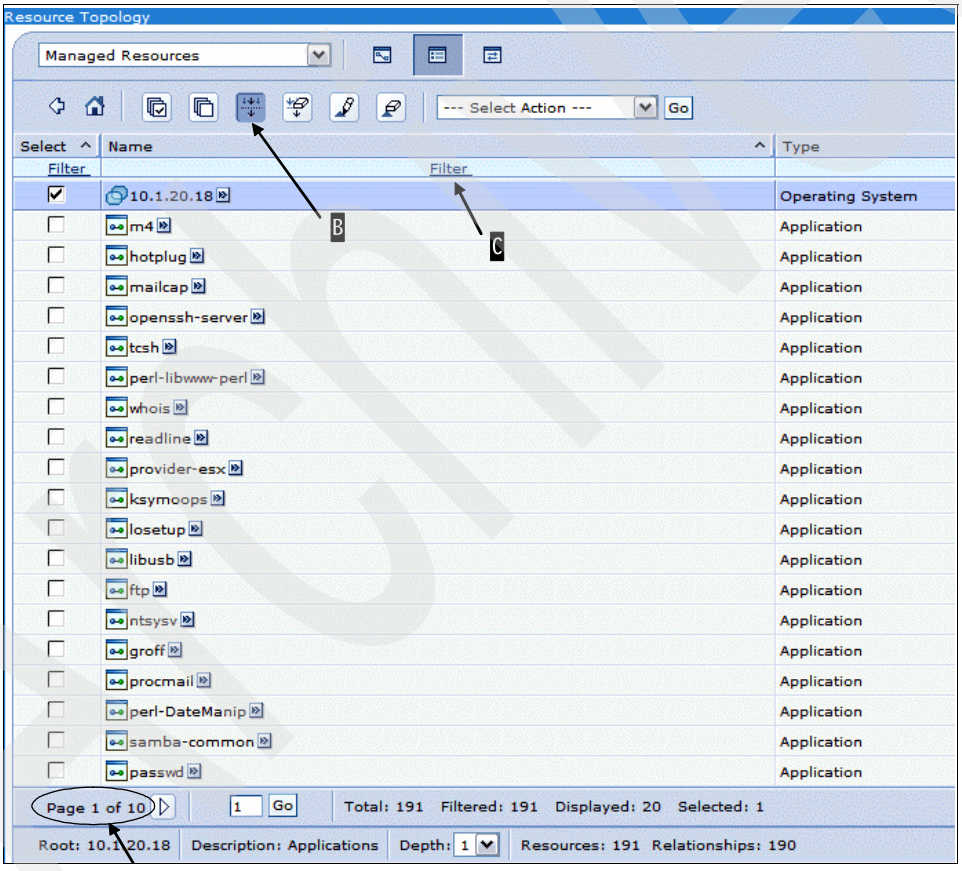


Figure 8-22 Resources in the table view

4. After clicking C, a window appears, as shown in Figure 8-23. Here you can define the filter according to your requirement. For example, you may want to check whether an ftp application is installed on this machine or not. Type ftp in the Text field and click **OK**.

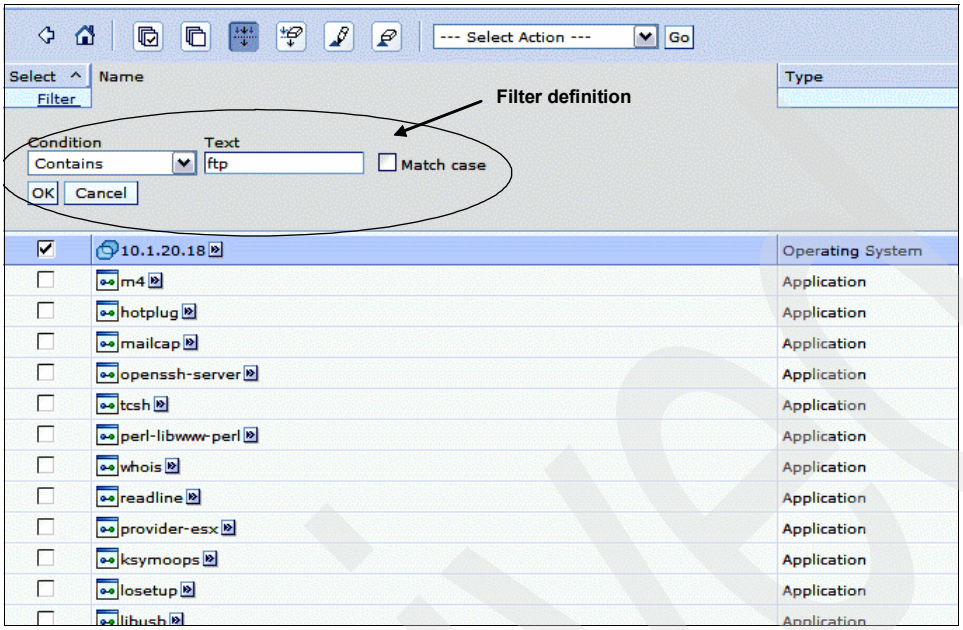


Figure 8-23 Defining a filter

The returned results are shown in Figure 8-24. It clearly shows that the ftp RPM is installed.

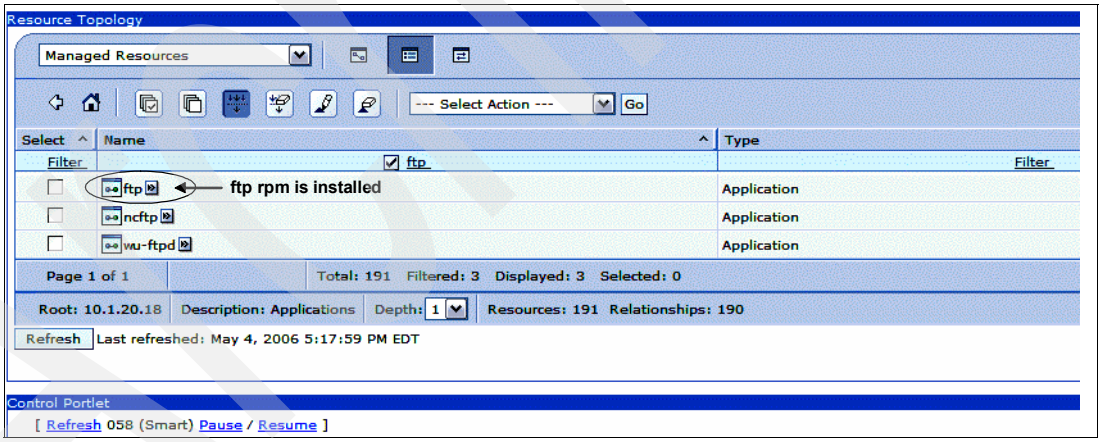


Figure 8-24 Filter results

Note: Filters can be defined in a variety of ways. You can choose a condition from the filter. You can also remove the defined filters from the navigation windows.

8.3.3 Registering a new resource

Sometimes, a particular system is not detected by the scanner, or, to segregate the view based on the systems in your department, you may want to add a new resource.

Note: New resources can only be added in the Managed Resources environment.

To add a new resource in your Managed Resource, perform the following tasks:

1. Right-click anywhere in your topology pane and select **New Resource** → **Group**, as shown in Figure 8-25.

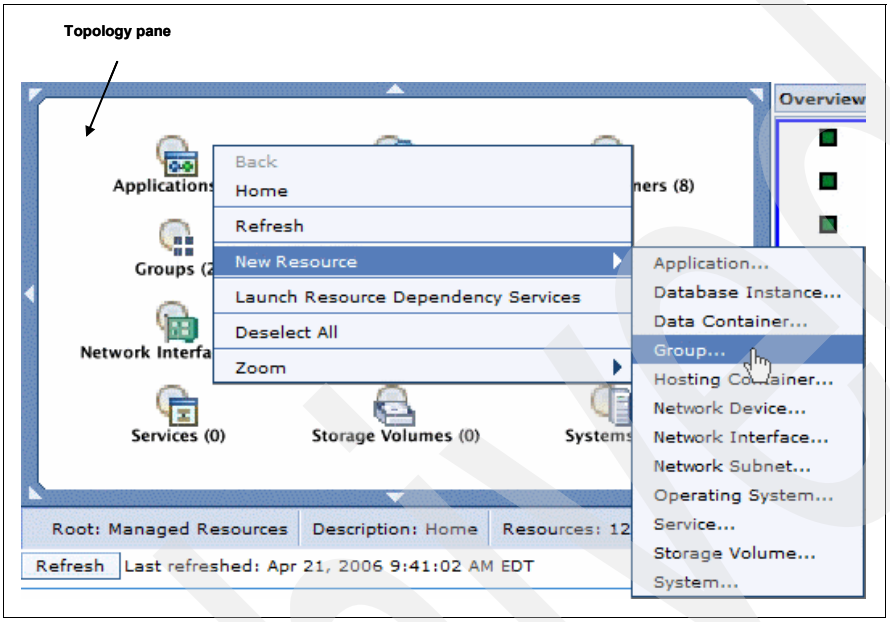


Figure 8-25 Adding a new resource

2. Enter the details shown in Figure 8-26. You can give it a logical name so that you can easily identify the group in the topology viewer, for example, you can select the name My_Group and give a description of My Department Machines.

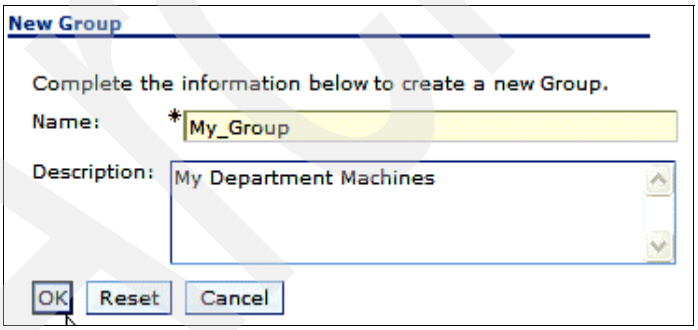


Figure 8-26 New group details

A new group is registered in your groups resource, as shown in Figure 8-27 on page 282.

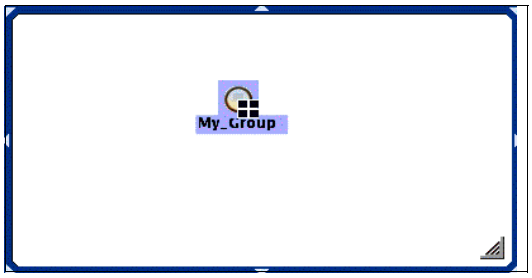


Figure 8-27 My_Group registered

8.3.4 What is a group?

A group is a collection of unrelated resources. You can create your own groups depending on your requirements, for example, machines in your department, machines hosting a specific application or database, and so on.

Adding resources in your group

To add resources in your group:

1. Create a group using the steps explained in 8.3.2, “What is a filter?” on page 278.
2. From the topology pane, select a system, right-click, and select **Open** → **All Systems**. You will see all the systems, as shown in Figure 8-28.

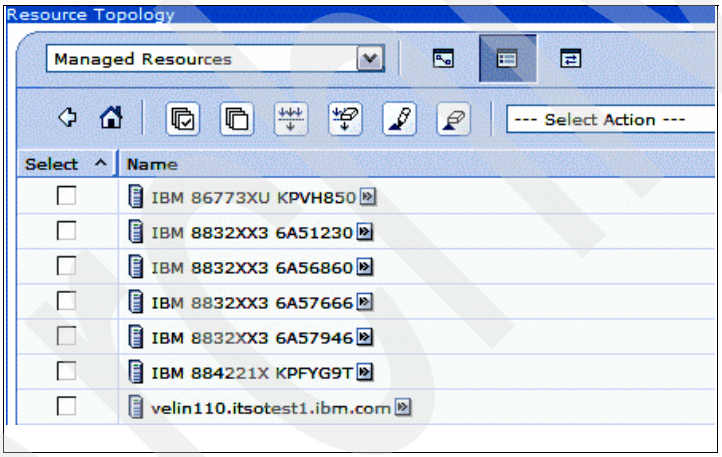


Figure 8-28 Complete system listing

Archived

3. Choose systems (A), click the corresponding buttons (B), and select **Add to Group**, as shown in Figure 8-29.

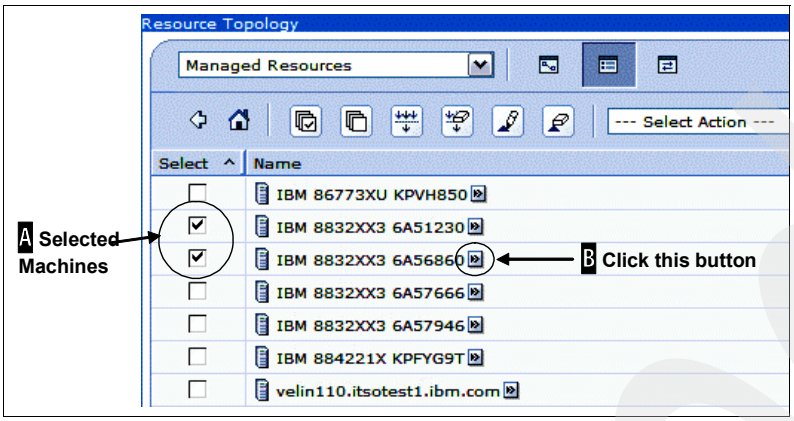


Figure 8-29 Choosing the systems

4. Select the group created, that is, My_Group. Alternately, you can create a new one. Select **My_Group** and click **OK**, as shown in Figure 8-30.

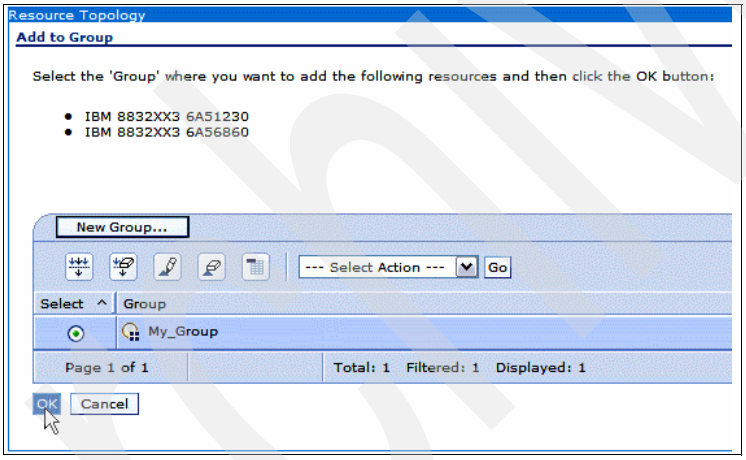


Figure 8-30 Selecting a group

You see the contents of the newly created group in the topology viewer, as shown in Figure 8-31 on page 284.

Archive

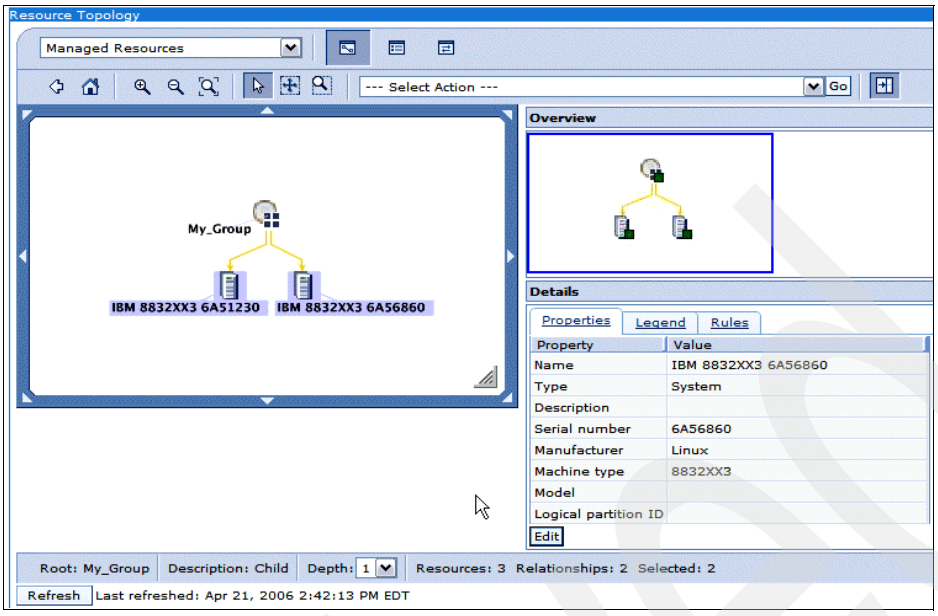


Figure 8-31 Group contents

Note: Resources participating in a group are always linked using the *yellow* color.

You can visualize the complete topology of My_Group along with other resources, as shown in Figure 8-32. Select a resource, right-click, and select **Open** → **All**.

Figure 8-32 shows the following listed dependencies and relationships:

- IBM Blade Chassis IBM 86773XU KPVH850 holds the two blades selected in this group, that is, IBM 8832XX3 6A56860 and IBM 8832XX3 6A51230.
- 10.1.20.133 has been installed on IBM 8832XX3 6A56860.
- 10.1.20.133 has IBM Director installed on it.

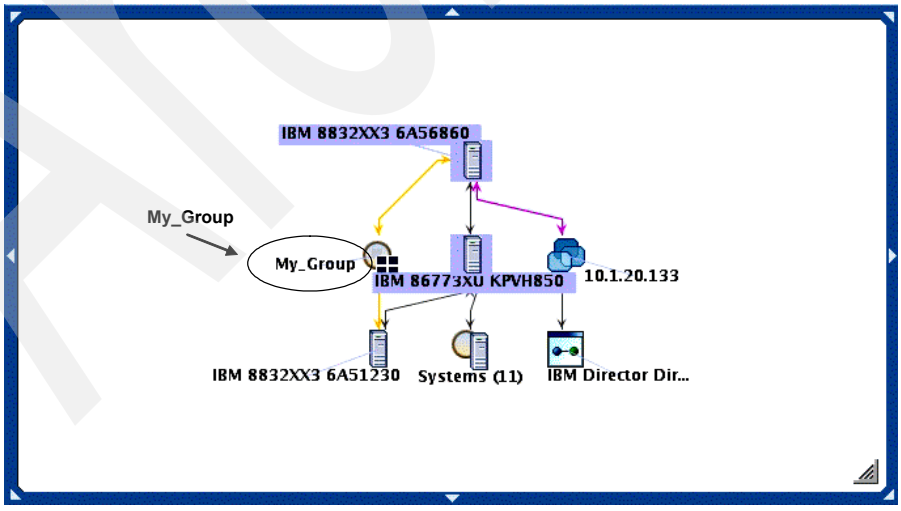


Figure 8-32 Complete topology of My_Group

8.3.5 What is a relationship?

A relationship is a connection defined between two resources. It is always defined between two or more resources. Define relationships manually only in two cases: when scans are not able to detect the relationship or you explicitly want to create a new relationship between the resources.

Note: You cannot create a relationship between dissimilar resources.

Creating a new relationship

Let us take a case where we want to create a new relationship between the operating system 10.1.20.133 and the system IBM 8832XX3 6A51230.

1. Select A, right-click, and select **New Relationship** (B), as shown in Figure 8-33.

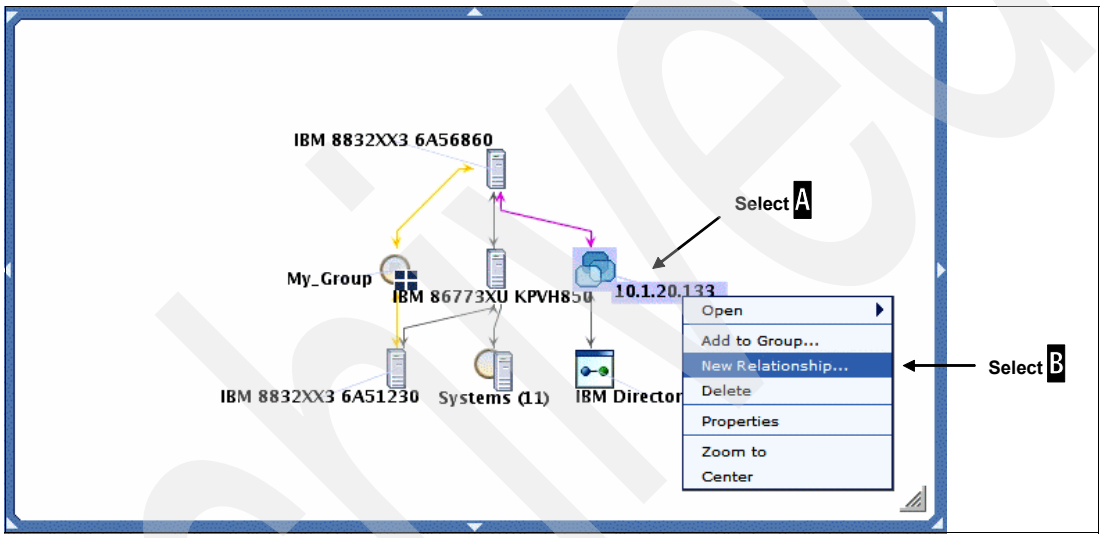


Figure 8-33 Creating a new relationship

Note: You have to be in the topology pane to define a new relationship.

2. In the next window, provide the necessary inputs, as shown in Figure 8-34 on page 286.
   a. A Specify the system identifier that has to be linked with this operating system, that is, IBM 8832XX3 6A51230.
   b. B Select **System** from the Resource Type drop-down list.
   c. C Click **Search**.
   d. D The system name appears in the field.
   e. E Click **Next**.

Archived

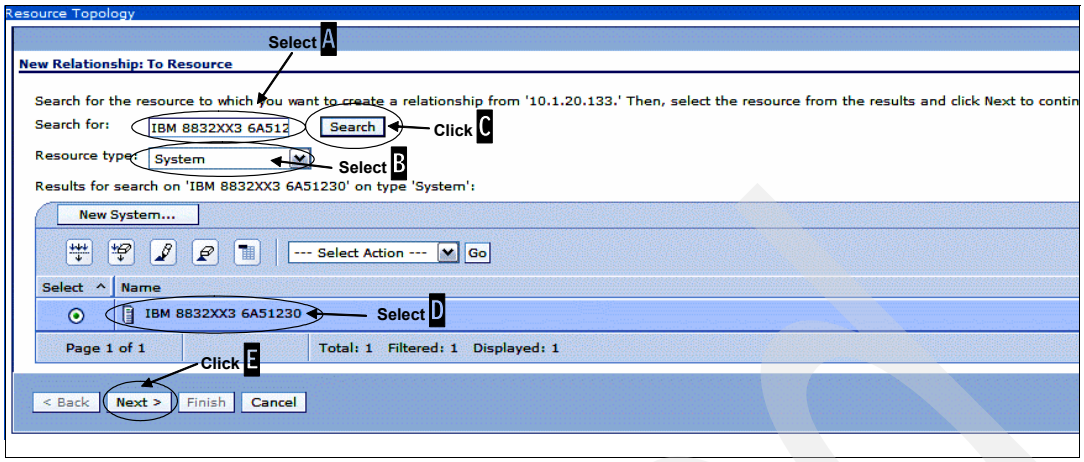


Figure 8-34 Inputs for creating a new relationship

3. Figure 8-35 shows the relationship details. Click **Finish**.

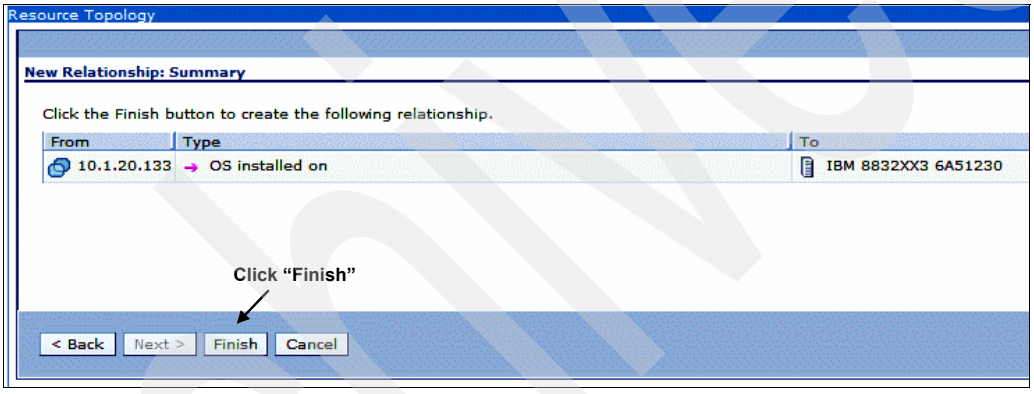


Figure 8-35 Relationship details

Figure 8-36 shows that a new relationship has been defined between operating system 10.1.20.133 and system IBM 8832XX3 6A51230.

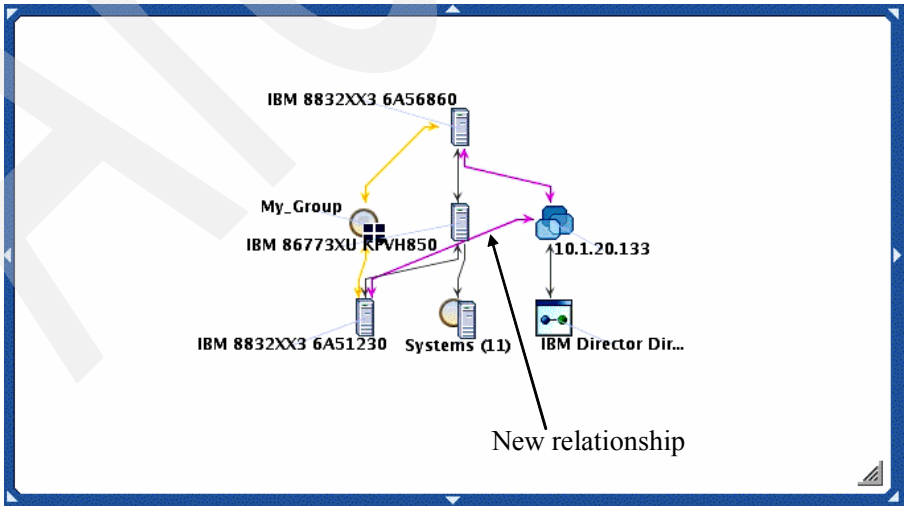


Figure 8-36 New relationship registered

8.4 Scenarios

In the following sections, we describe how the topology viewer reflects the changes of the infrastructure in three cases:

- When we add a database
- When we add a resource
- When we remove a resource

8.4.1 Adding a new database in the system

Consider this scenario: A human resources department requires a new application to be deployed for its employees. This application requires a new database. The system administrator has identified the machine velin50 for creating this new database. We demonstrate how this new database appears in the topology viewer.

1. Check the Managed Resources status in the topology viewer before registering any new database on velin50. Figure 8-37 shows that there are no database instances.

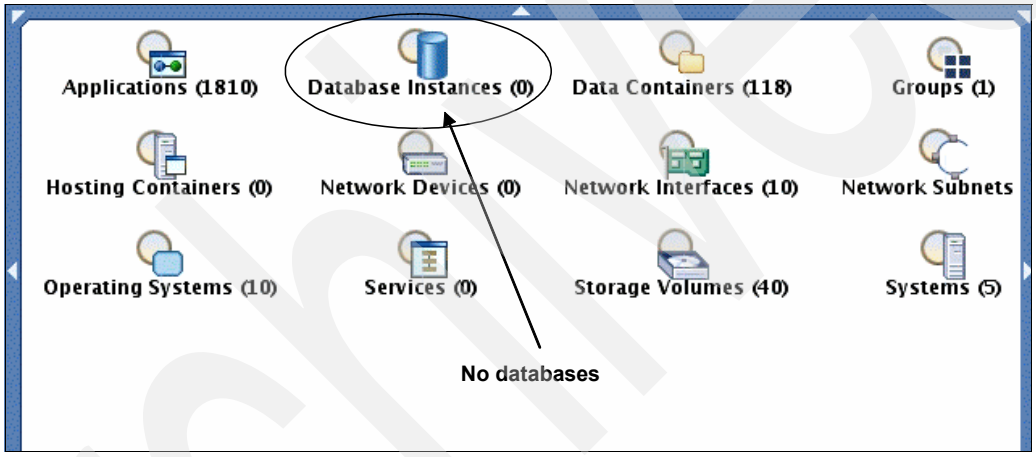


Figure 8-37 No databases registered under Managed Resources

Archived

2. Run the DB2 Collector on the identified machine velin50 / 10.1.20.50.
3. Select A from the navigation window and click B, as shown in Figure 8-38.

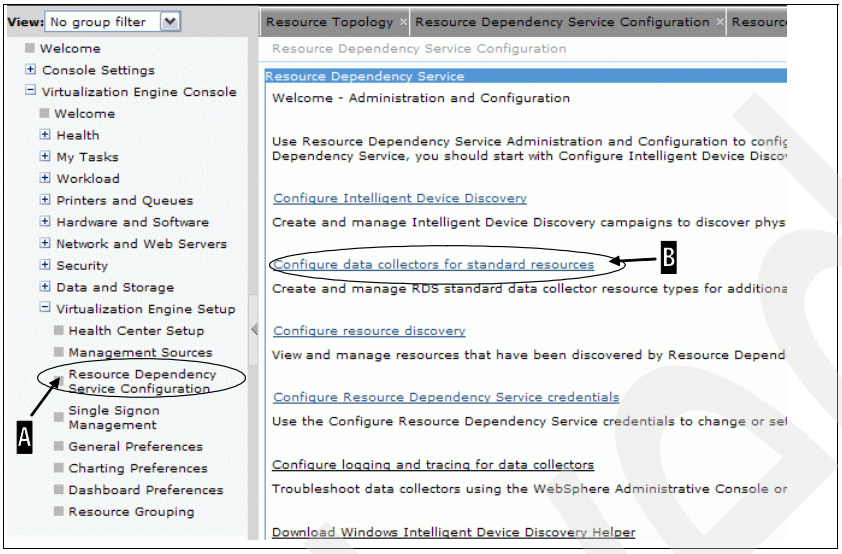


Figure 8-38 Resource Dependency Service Configuration

4. Select **DB2Collector**, as shown in Figure 8-39.

> **Note:** DB2Collector_HR has been specifically defined for the machine hosting the IBM DB2 UDB machine.

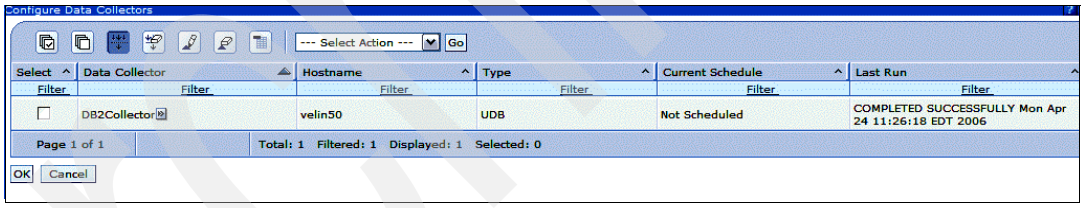


Figure 8-39 Defined DB2Collector for velin50

5. Open the Topology Viewer. From the navigation window, select and expand **Virtualization Engine Console** → **Hardware and Software** → **Resource Topology**.

Archive

6. Select **Managed Resources** from the drop-down list 1, as shown in Figure 8-3 on page 267.
7. You can now see two database instances discovered by the DB2Collector, as shown in Figure 8-40. Select **Database Instances**, right-click, and select **Open** → **All Database Instances**.

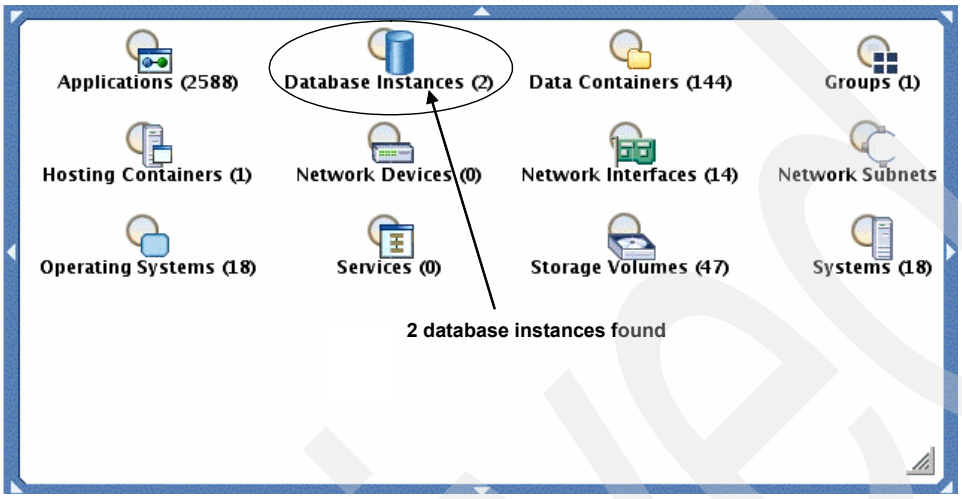


Figure 8-40 Database instances registered under Managed Resources

You can see the database instance details, as shown in Figure 8-41.

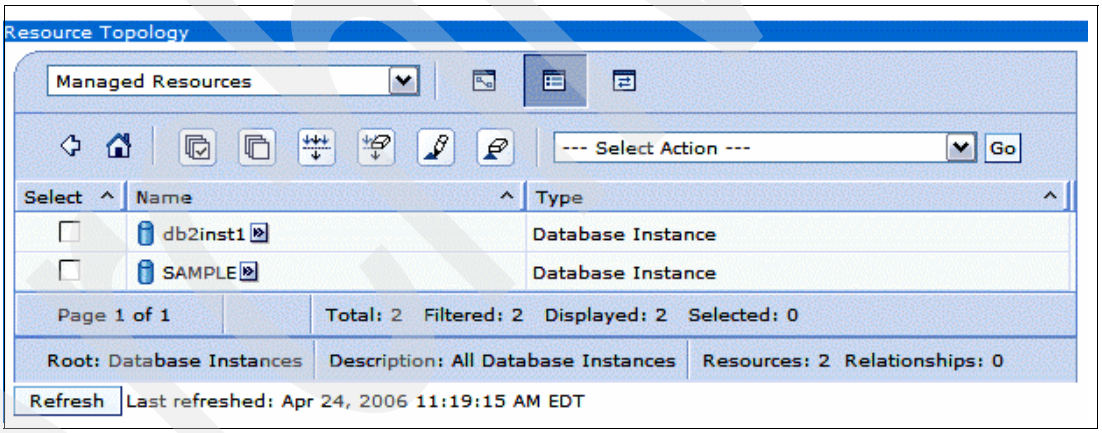


Figure 8-41 Database instances details

8. Create a new database HR_db on the velin50 machine.
9. Re-run DB2Collector_HR on velin50.
10. From the navigation window, select A and click B, as shown in Figure 8-38 on page 288.
11. Select **DB2Collector_HR**, as shown in Figure 8-42 on page 290.

> **Note:** DB2Collector_HR has been specifically defined for the machine hosting IBM DB2 UDB.

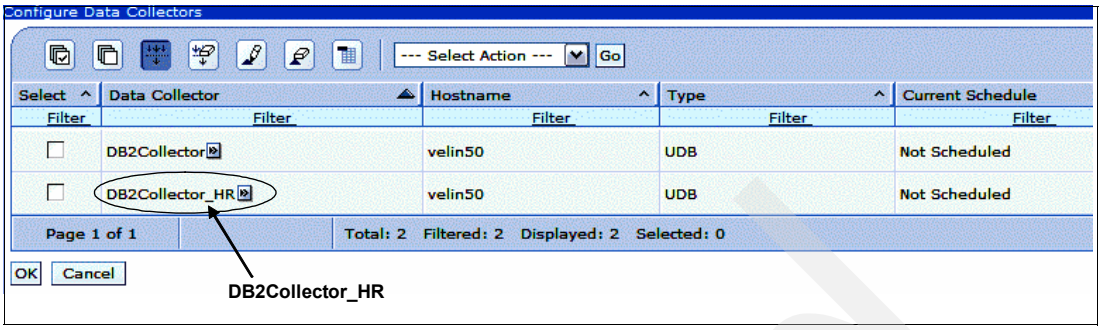


Figure 8-42 DB2Collector_HR

12. Open the Topology Viewer again. From the navigation window, select and expand **Virtualization Engine Console** → **Hardware and Software** → **Resource Topology**.
13. Select **Managed Resources** from the drop-down list 1, as shown in Figure 8-3 on page 267.
14. From the topology pane, select **Database Instances**, right-click, and select **Open** → **All Database Instances**.

    You see that the HR_db database is registered under Database Instances, as shown in Figure 8-43.

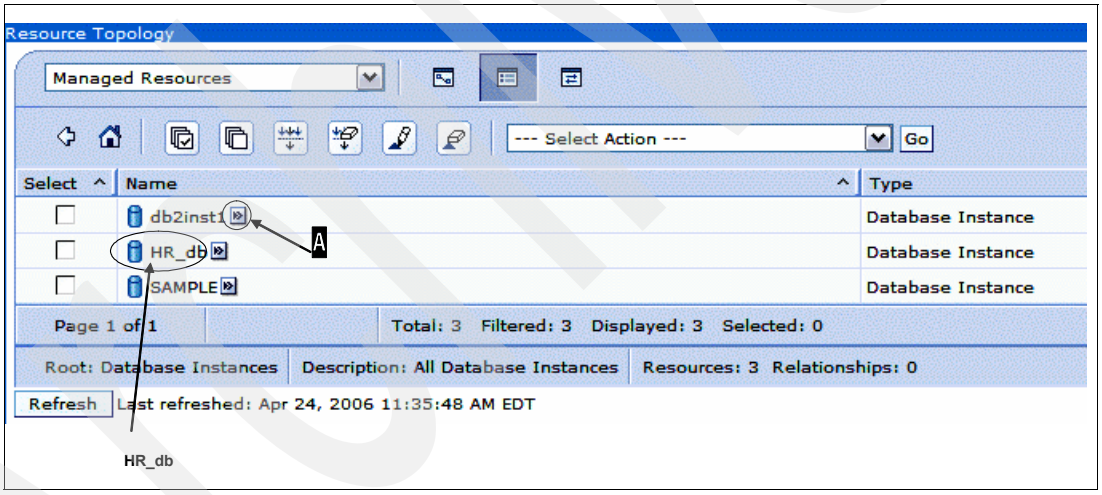


Figure 8-43 New database registered

To get more information about the instance name db2inst1, click A and select **Open** → **Child**. You will see the details in the topology window, as shown in Figure 8-44.

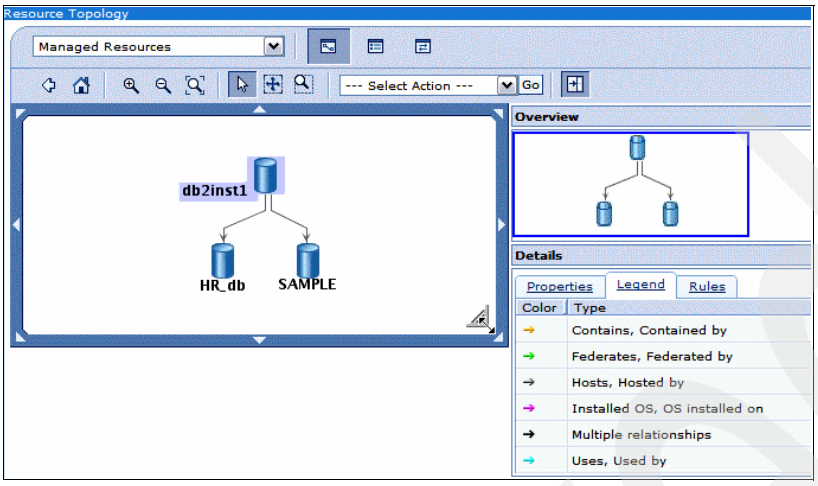


Figure 8-44 Database instance db2inst1 details

You can see the complete topology with all the resource dependencies, as shown in Figure 8-45. Select any resource, right-click, and select **Open** → **All**.

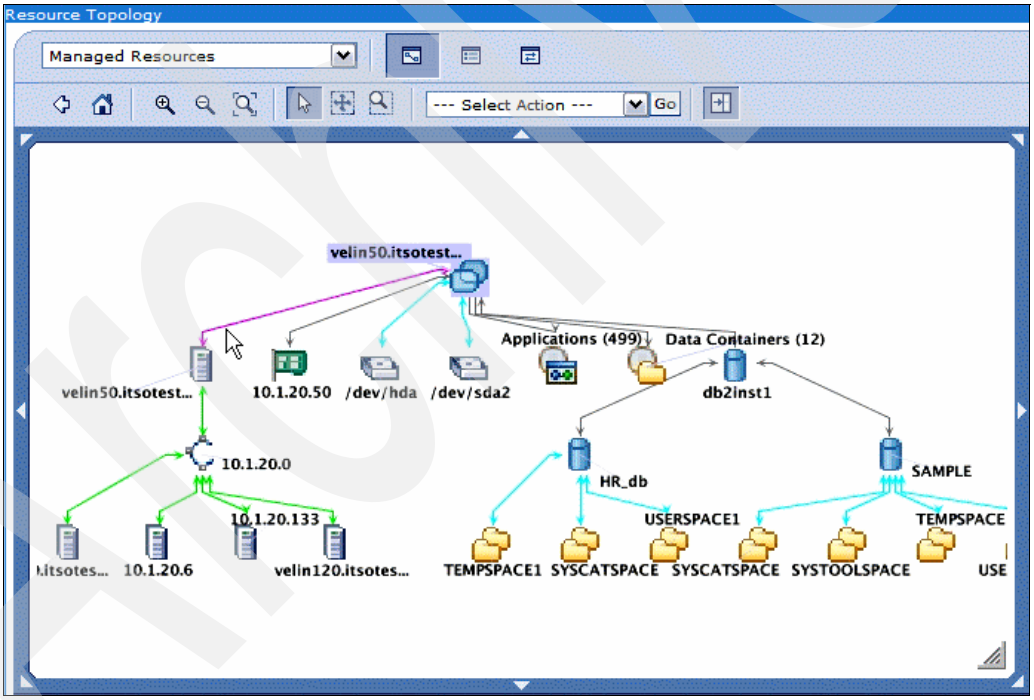


Figure 8-45 Complete topology details of velin50

8.4.2 Adding new resources in the environment

In this scenario, we see how the addition of new resources reflects the changes in the topology viewer.

Note that the IBM Director and the Virtualization Engine console are two separate components. The Virtualization Engine talks to the IBM Director using a bridge. Refer to Chapter 9, “Using one single console to manage your environment” on page 301 for more

details. The IBM Director captures resource information from time to time based on its schedule. This information is then fetched from the IBM Director using the IBM Director collector defined on the Virtualization Engine console. For more details about collectors, refer 7.4.1, “Where to look for the collectors” on page 244.

To visualize the infrastructure change, perform the following tasks:

1. Run the IBM Director collector to get the present status before adding any new machine:
   a. From the navigation window, select **A** and click **B**, as shown in Figure 8-38 on page 288.
   b. Select **IDCollector**, as shown in Figure 8-46.

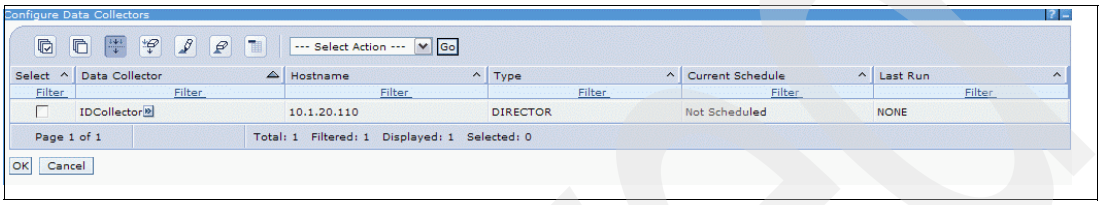


Figure 8-46 IBM Director collector definition

   c. After the collector has run successfully, you can see what the topology viewer looks like. Select **Systems** from Managed Resources environment, right-click, and select **Open** → **All Systems**. Figure 8-47 shows the present status.

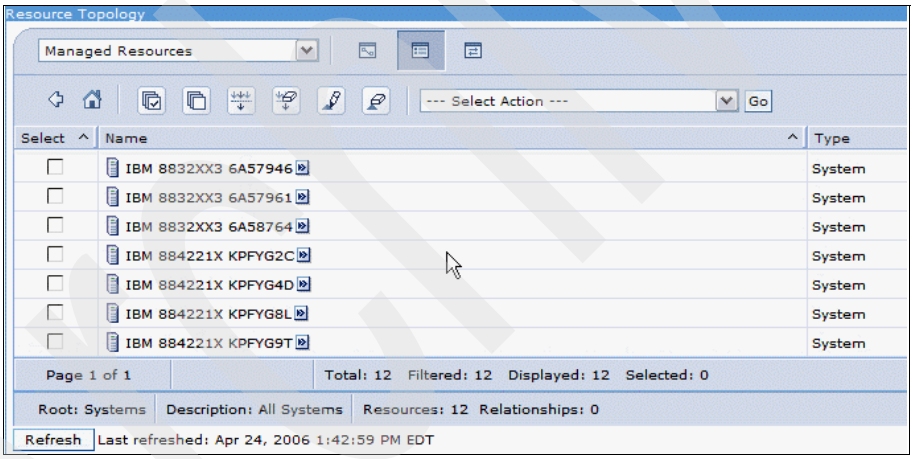


Figure 8-47 System list

Now add new resources, such as computer systems, printers, and so on in your environment.

2. Re-run the IDCollector.
   a. From the navigation window, select **A** and click **B** as shown in Figure 8-38 on page 288.
   b. Select **IDCollector**, as shown in Figure 8-46.
3. Open Managed Resources in the topology pane. Select **Systems** from the Managed Resources environment, right-click, and select **Open** → **All Systems**.

You can now find more resources in your topology viewer, as shown in Figure 8-48.

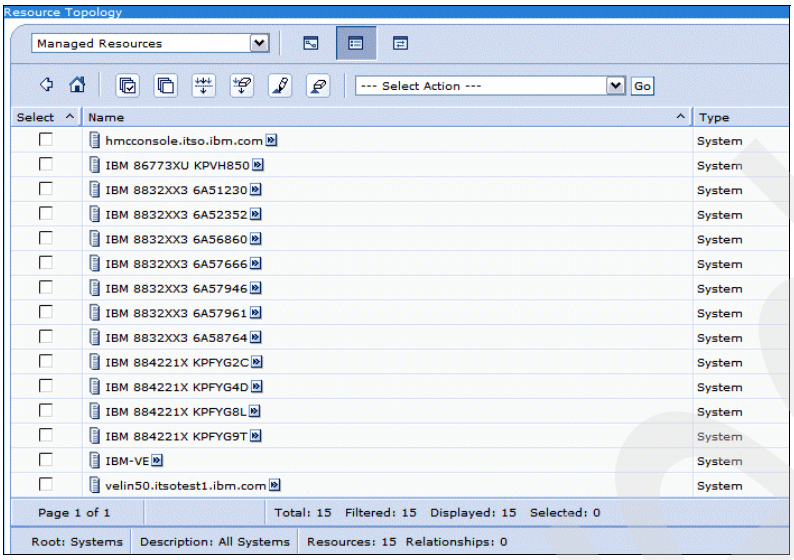


Figure 8-48 More resources discovered by IBM Director

8.4.3 Using topology viewer for removing the resources

In this scenario, we discuss how to remove resources from the topology viewer. This helps in keeping the view up to date. For example, you have a machine hosting a service. However, after some time, you do not require it and you want to remove the resource from the viewer.

1. Log in to the Virtualization Engine console using the Virtualization Engine credentials, as shown in Figure 8-1 on page 266.
2. Expand **Virtualization Engine Console** → **Hardware and Software** → **Resource Topology**, as shown in Figure 8-2 on page 266.
3. Select **Managed Resources** from the drop-down list 1, as shown in Figure 8-3 on page 267.
4. Select **Systems** from the Managed Resources environment, right-click, and select **Open** → **All Systems**. The systems are displayed, as shown in Figure 8-47 on page 292.
5. Select A, click B, and select **Open** → **All**, as shown in Figure 8-49.

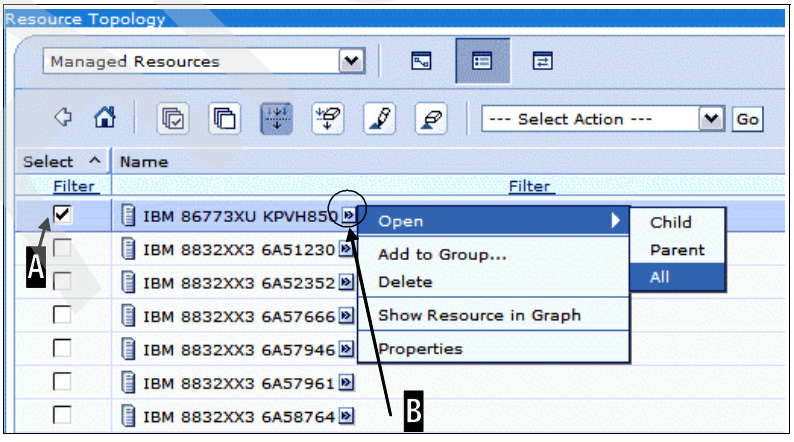


Figure 8-49 Selected system

The selected resource is displayed, as shown in Figure 8-50.

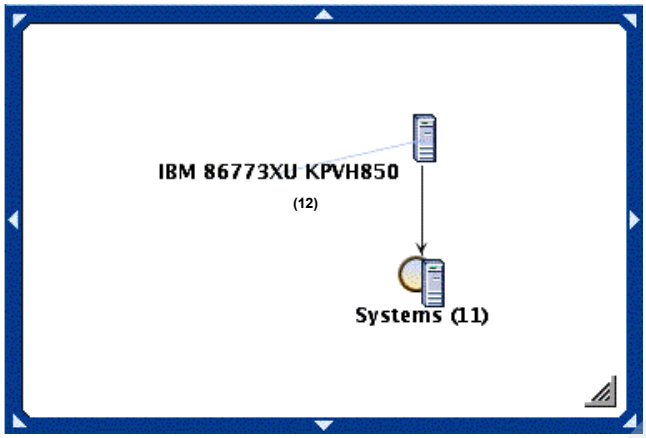


Figure 8-50 Resource in a topology view

6. Select **Systems(12)**, right-click, and select **Open** → **Systems**. The complete system view is displayed, as shown in Figure 8-51.

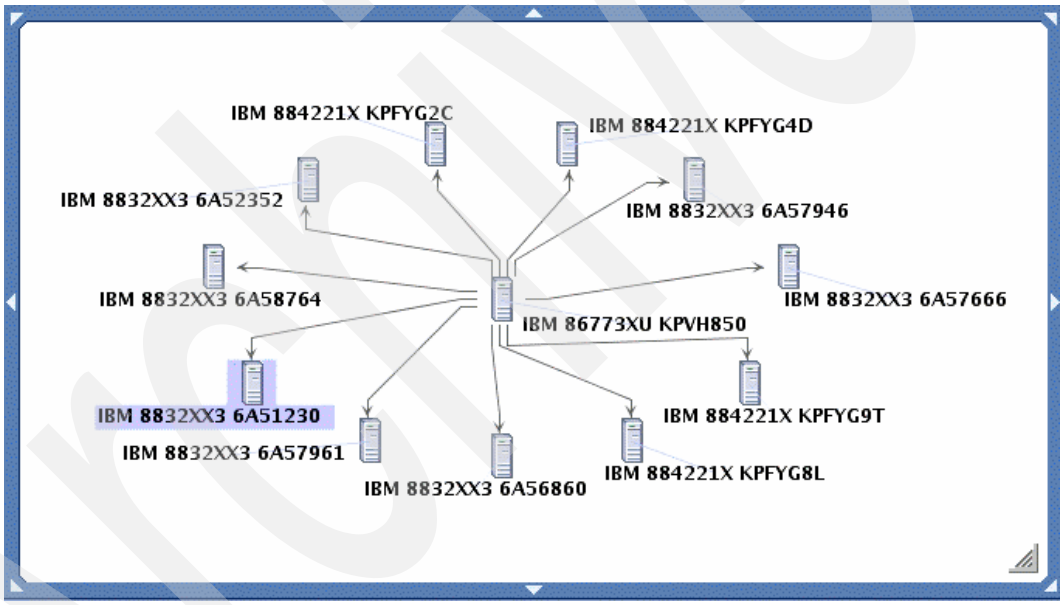


Figure 8-51 All systems in the topology window

Archived

7. Select the machine you want to remove, in this example, **IBM 8832XX3 6A51230**, right-click, and select **Open** → **All**. Figure 8-52 shows the resource dependencies with the other resources.

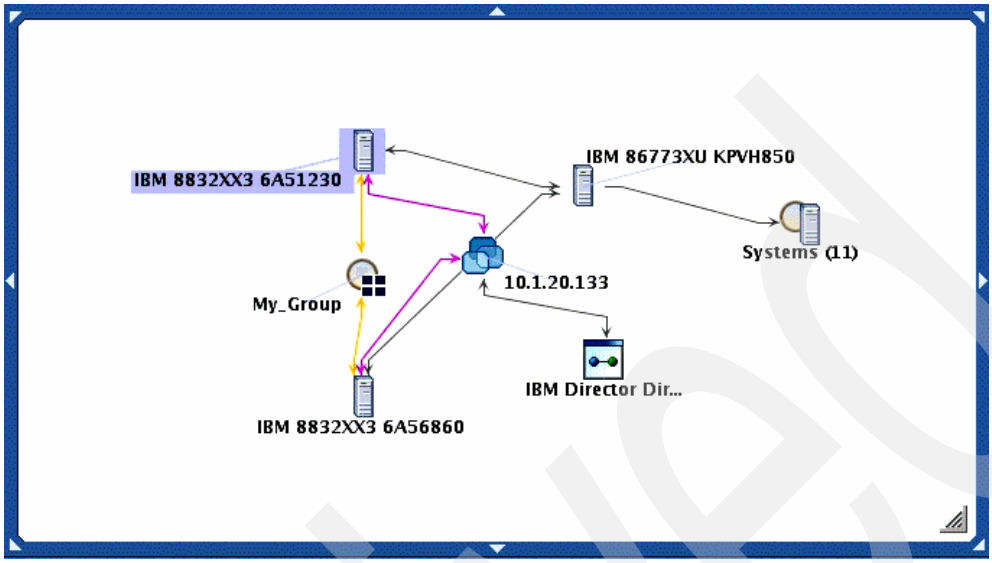


Figure 8-52 Selected resource dependencies

8. You can now remove the resource that is no longer required. Click **IBM 8832XX3 6A56860**, right-click, and select **Delete**, as shown in Figure 8-53.

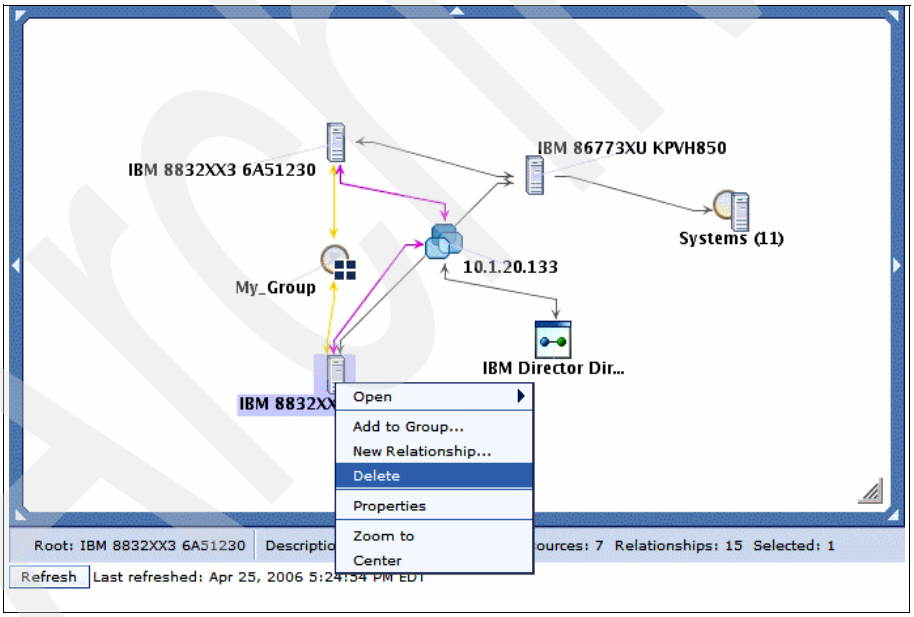


Figure 8-53 Removing the resource

Archived

9. When you select **Delete**, a confirmation window is displayed, as shown in Figure 8-54. Click **Yes**.

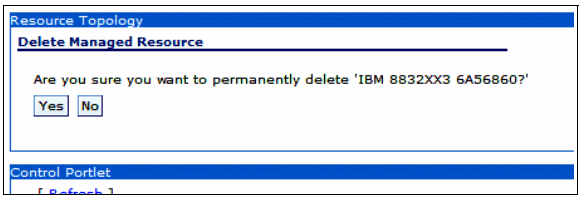


Figure 8-54 Confirmation window for the removal

The selected resource gets removed from the topology view, as shown in Figure 8-55.

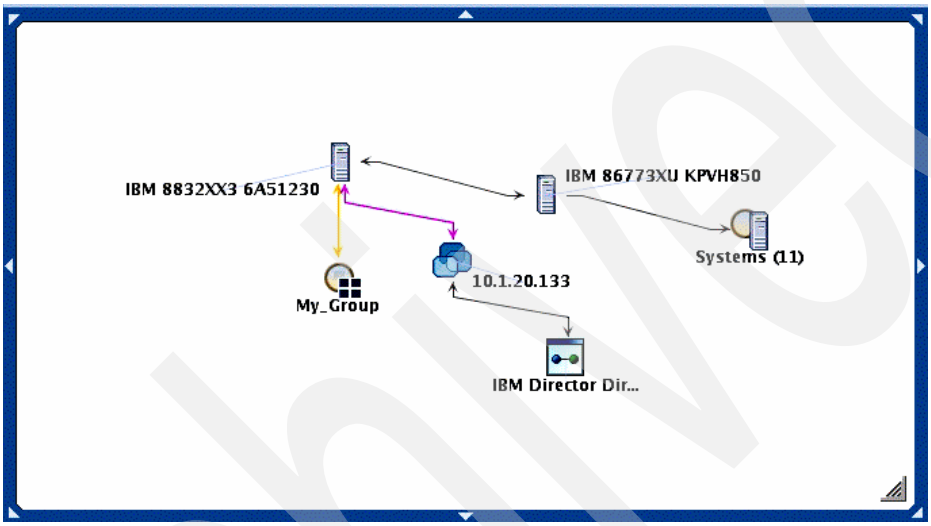


Figure 8-55 Selected resource removed

In Figure 8-55, note that when you delete the resource, all its dependencies and relationships with other resources are automatically removed, that is, the relationship with My_Group, with 10.1.20.133, and with Blade Chassis.

Archived

To verify that the resource has been removed, select **System(11)**, right-click, and select **Open** → **Systems**. Figure 8-56 shows the output.

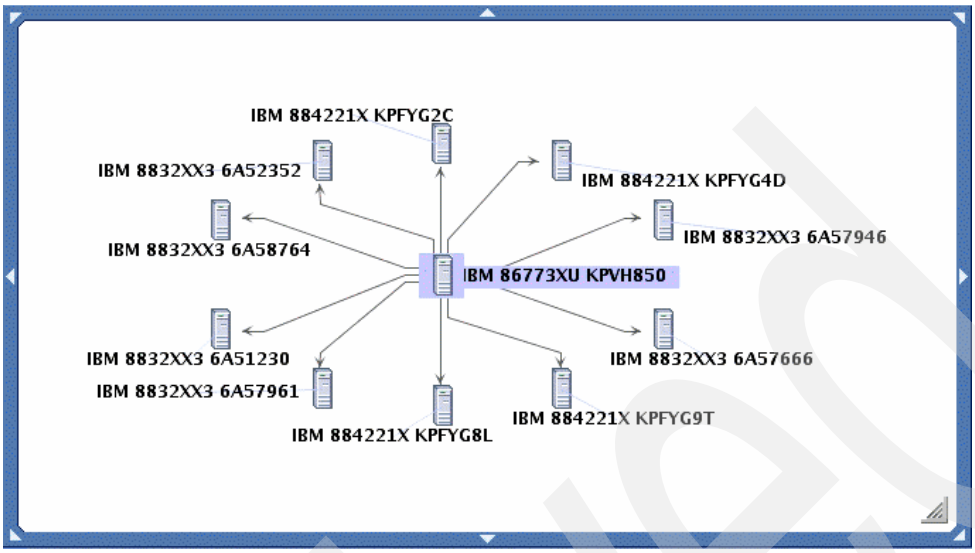


Figure 8-56 Topology window after resource has been removed

8.4.4 Interpreting the information shown in the topology viewer

In this section, we explain how to understand the topology viewer in two cases.

- With an IBM System p machine in the environment
- With an IBM System x Blade Center in the environment

With an IBM System p machine in the environment

In this section, we explain how to interpret the information shown up in the topology viewer. Refer to Figure 8-57 for understanding the process.

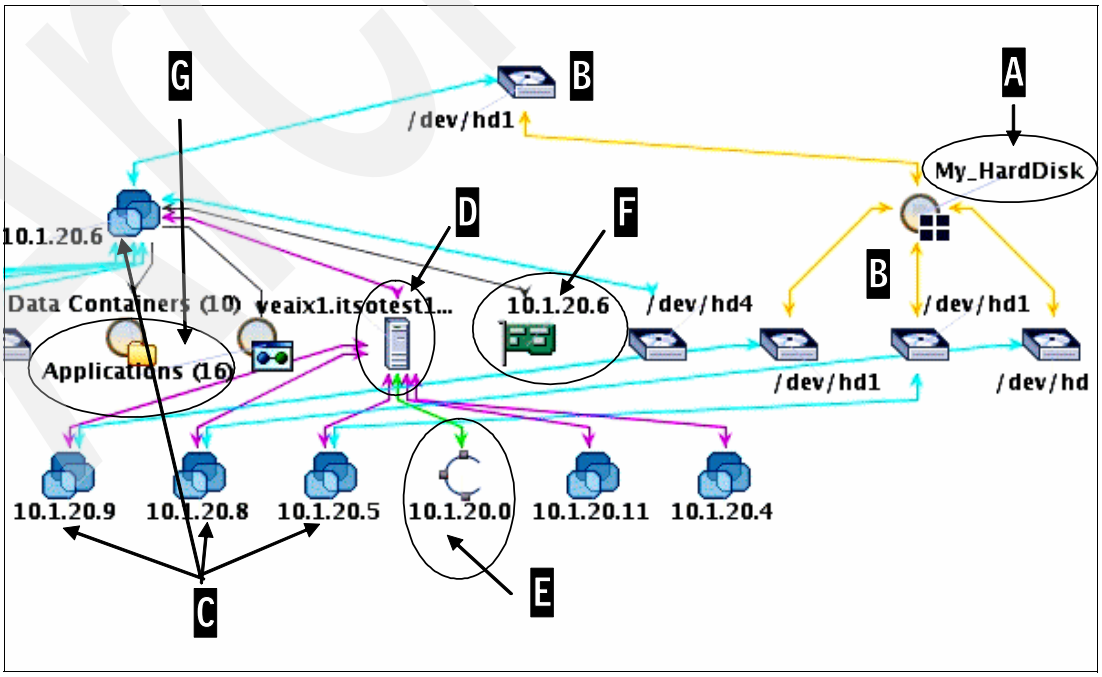


Figure 8-57 Sample topology

Note: Figure 8-57 shows only a part of the overall information captured in the topology viewer.

We can observe the following resources, their dependencies, and relationships:

- A is a manually created group called My_HardDisk.
- A has four group members named /dev/hd1 identified by B.
- Four operating systems, identified by C, have been installed on each of these hard disk partitions.
- All these operating systems are hosted on the same machine D. Therefore, it identifies the machine with logical partitioning (LPAR) capability. (We had IBM System p 550 in our test environment.)
- Machine D is on a 10.1.20 subnet, as pointed out by E.
- IP address F 10.1.20.6 is linked to one of the operating system installed on machine D.
- There are 16 applications G installed on 10.1.20.6.

With an IBM System x Blade Center in the environment

Let us look at another example, as shown in Figure 8-58.

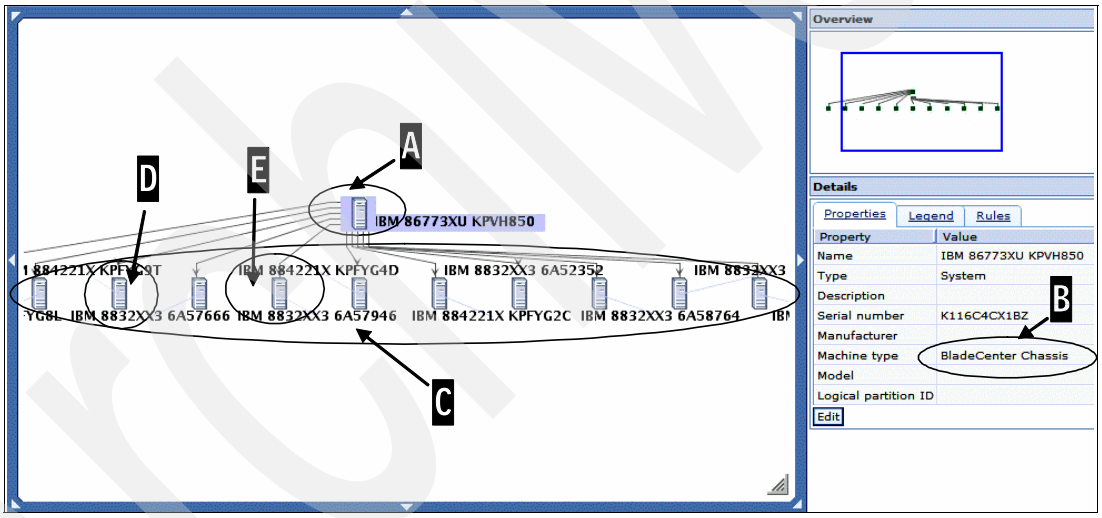


Figure 8-58 Blade chassis with 10 blades

Note: Figure 8-58 shows the Systems view in the topology window.

- There is one BladeCenter chassis, A, detected in the environment, as pointed out in the Properties window in B.
- There are 10 blades, C, installed in the BladeCenter chassis A.
- There are two blade models D and E installed in the blade chassis: 8832 and 8842.

- For each blade, you can find more details in the Properties window B, that is, serial number and model details, as shown in Figure 8-59.

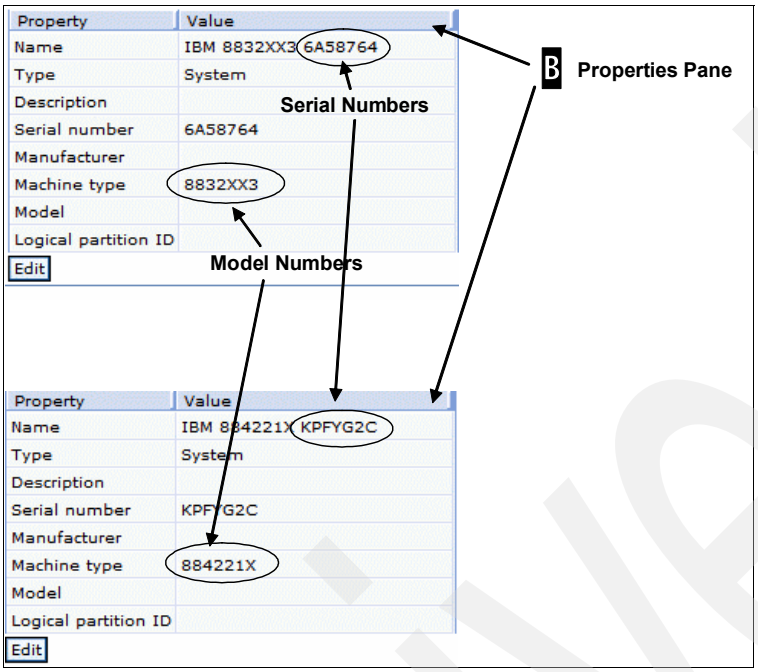


Figure 8-59 Properties window showing details of each blade

Archived

Archived

9

Using one single console to manage your environment

One of the purposes of the Virtualization Engine console is to provide a single portal for administering IT resources that hides the complexity of the underlying technologies.

In order to manage a heterogeneous environment, you usually need several consoles. For example, the IBM Director Console allows you to manage several IBM eServer platforms. However, in order to perform specific actions, you could need other consoles. For example, Web-based Systems Manager (WebSM) to connect to a Hardware Management Console (HMC) for IBM System p or IBM System i management, or a Web browser to connect directly to a BladeCenter Management Module, and so on.

This chapter gives you an overview of how you can use the Virtualization Engine console as a single point of access for management tasks, without needing to know which tool you are actually using under the cover. The tools we are using are IBM Director and the WebSM. We assume that you are familiar with the main features of the Virtualization Engine, including the Console, IBM Director, and WebSM. For more information, refer to *IBM Virtualization Engine platform Version 2, Technical Presentation Guide*, SG24-7112.

The beginning of this chapter, 9.1, “The environment” on page 302, describes the environment we used to run the scenarios. 9.2, “Configuring the environment” on page 304 describes the basic configurations we performed on the management environment and on the managed systems. Finally, this chapter describes each scenario in detail, comprising setup, configuration, and the results for each of them. We chose basic scenarios that you can combine to create more complex situations.

These are the following scenarios:

- 9.3, “Diagnosing a hardware error” on page 330
- 9.4, “Processing error” on page 340
- 9.5, “Monitoring the system CPU usage” on page 350
- 9.6, “Managing a VMWare Virtual Machine” on page 358
- 9.7, “Monitoring an LPAR” on page 372

© Copyright IBM Corp. 2006. All rights reserved.

9.1 The environment

In this section, we describe the environment we set up at the International Technical Support Organization (ITSO) facilities to run the scenarios.

9.1.1 Management servers

The management servers are two HS20 blade servers, one running SUSE Linux Enterprise Server (SLES) 9.3, and one running Windows 2003 Standard Edition Service Pack (SP) 1. Figure 9-1 shows the distribution of the Virtualization Engine components. All the components are running on Linux, except for the IBM Director Server, which is on a Windows-based system. The reason for this is that the Windows-based installation is richer in functions and extensions. For example, Virtual Machine Manager is not available on non-Windows installations.

Our IBM Director Server database is a DB2 database configured in the same DB2 instance used by the other Virtualization Engine components (running on the Linux server).

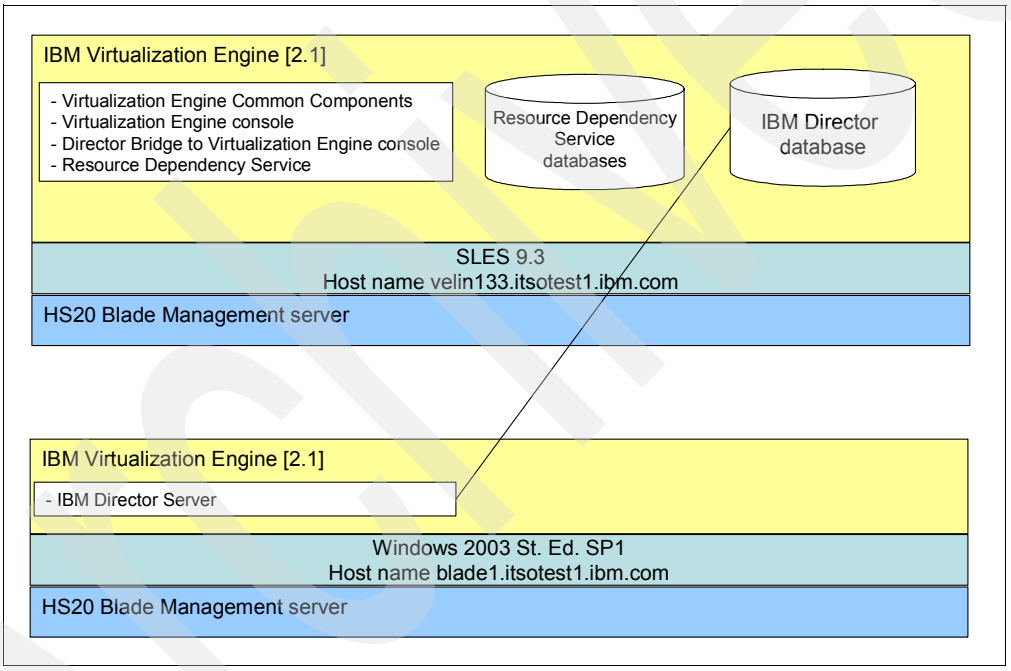


Figure 9-1 Management servers

IBM Director is comprised of a base framework, whose functions can be enhanced by IBM Director extensions.

The IBM Director Server default installation does not include two tasks:

- BladeCenter Management
- Rack Manager

We installed both features on our management server. If you did not select these components during the installation, you can add them at a later time. On Windows, you can run the installation wizard again, while on UNIX-like systems (AIX, i5/OS or Linux), you can run the installation shell script changing the response file to include the new components. Refer to

the *IBM Director Installation and Configuration Guide* for more details. This is available on the Web at:

http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-61788

In addition, we upgraded to IBM Director 5.10.1 and installed the following extensions:

- Virtual Machine Manager (VMM) V2.1
- Hardware Management Console V5.10.1

For more details on IBM Director features, available extensions, and installation tips, refer to *Implementing IBM Director 5.10*, SG24-6188.

Important: We did not install the ServeRAID™ Manager extension in our environment because version 8.25, the latest available at the time of the writing of this IBM Redbook, causes the communication between the IBM Director Bridge and the IBM Director Server to fail.

When launching the IBM Director Console, either from the Virtualization Engine Console or independently, you should ensure that the proper console extensions are installed. You will see question marks close to the tasks installed on the server that have the corresponding console components missing. Refer to *Implementing IBM Director 5.10*, SG24-6188 for more details on the IBM Director extensions.

9.1.2 Managed nodes

The managed nodes in our environment, as shown in Figure 9-2 on page 304, are:

- One BladeCenter Chassis containing several blade servers, some of which are hosting the Virtualization Engine components.
- One IBM System p HMC V5.2, connected to an IBM System p550, configured with seven logical partitions (LPARs), as described in Figure 9-2 on page 304.

For our scenarios, we used the following:

- One HS20 server running VMWare ESX 2.5.2 (Blade #04 in Figure 9-2 on page 304)
- One VMWare ESX 2.5.2 virtual machine running Windows 2003 Standard Server Edition and the IBM Director L2 Agent
- The p550 LPAR named veaix3 running AIX 5L V5.3
- The p550 LPAR named veaix4 running AIX 5L V5.3
- The BladeCenter Management Module
- The HMC

Note: We used ESX 2.5.2 with IBM Director 5.10.1, even if it is not officially supported. IBM Director 5.10.2 fully supports ESX 2.5.2.

For more information about supported operating systems, refer to the *IBM Director Hardware and Software Compatibility* document for the IBM Director version that you are using. You can find this on the Web at:

http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-61788

You must also check which Virtualization Machine Manager (VMM) version is supported with your release of IBM Director.

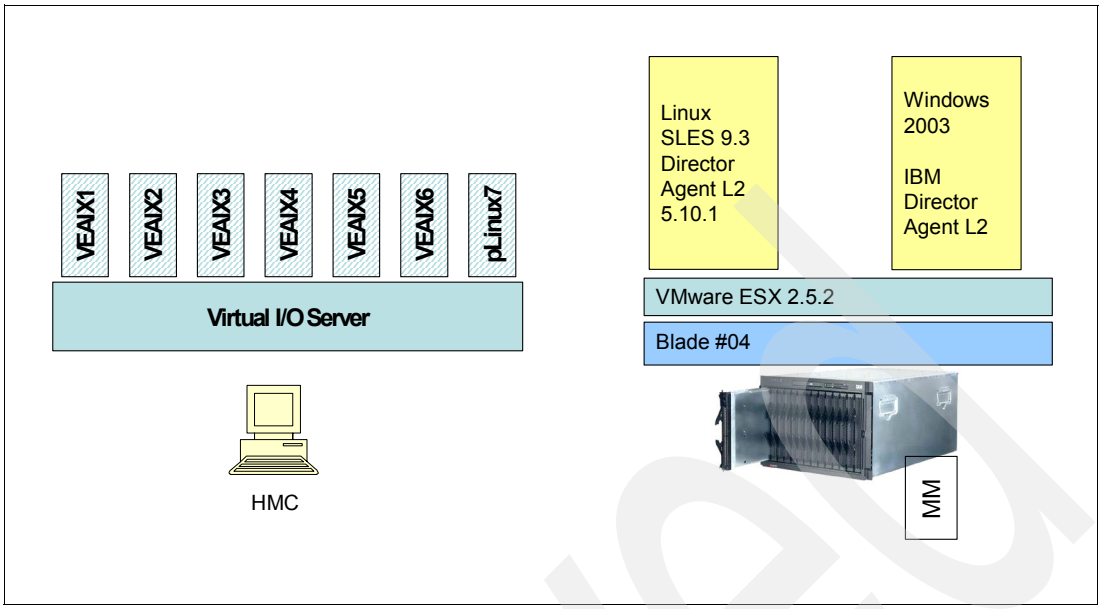


Figure 9-2 Managed nodes

9.2 Configuring the environment

Before performing actions on the managed systems from the Virtualization Engine console, you should configure the management sources to select which systems you want to monitor from the IBM Director Server through Web services. Depending on the type of the managed system, you might need to do additional configurations on these as well. In this section, we describe both the management sources and the managed systems configurations.

9.2.1 Configuring the management sources

In order to manage IBM Director resources from the Health Center (a function of the Virtualization Engine console), you should configure the management sources corresponding to the Web service that communicates with the IBM Director Server. You can do this using the *Health Center setup wizard*.

Archive

There are four ways of accessing the Health Center setup wizard:

- In the left navigation window of the window shown in Figure 9-3, expand **Virtualization Engine Console** and click **Welcome**. In the Welcome to the Virtualization Engine Console window, click **Setup the Health Center**.

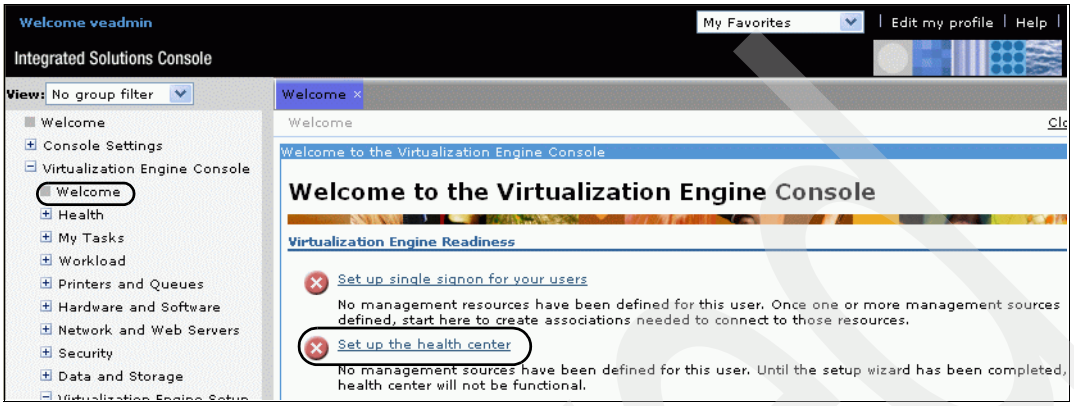


Figure 9-3 Virtualization Engine console welcome page

> **Tip:** There are two welcome pages:
>
> - The first one is related to the Integrated Solutions Console, the framework on which the Virtualization Engine console is based. This welcome page is the first displayed after you successfully log into the Virtualization Engine console.
> - The second one is related to the Virtualization Engine console. Ensure that you go to the Virtualization Engine console welcome page.

- In the left navigation window, expand **Virtualization Engine Console**, expand **Health**, and click **Health Center**. In the Workspace portlet, click **Add Management Sources**, as shown in Figure 9-4.

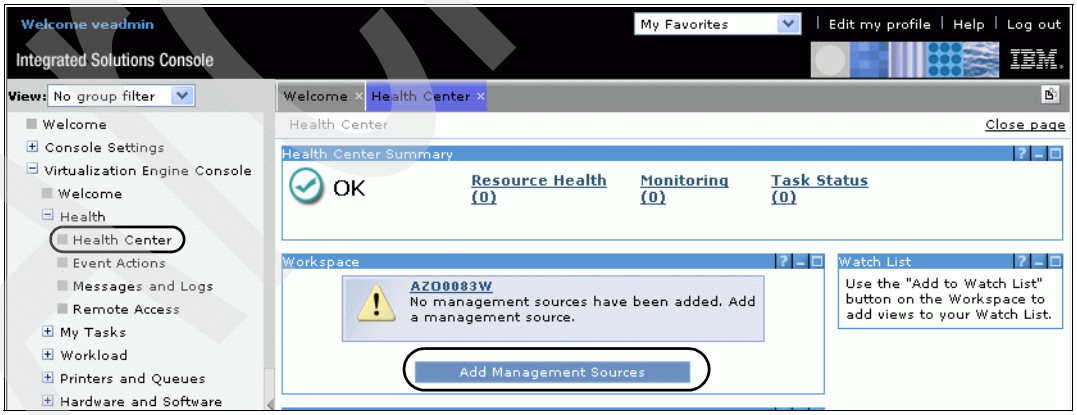


Figure 9-4 Virtualization Engine console Health Center

- In the left navigation window, expand **Virtualization Engine Setup** and click **Health Center Setup**, as shown in Figure 9-5.



Figure 9-5 Health Center setup

- In the left navigation window, expand **Virtualization Engine Setup** and click **Management Sources**. In the Management Sources Configuration window, click **Setup the Health Center**, as shown in Figure 9-6.

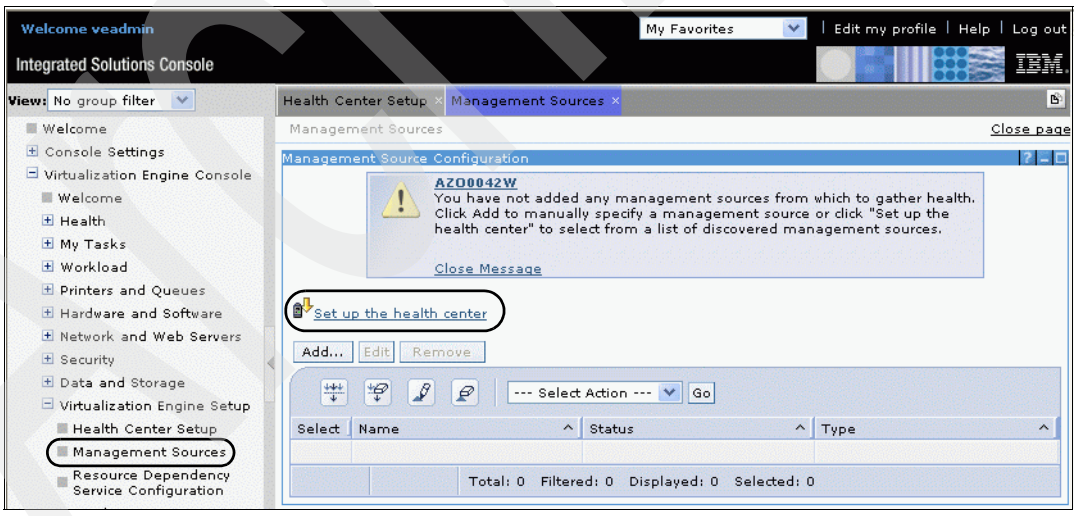


Figure 9-6 Management Sources configuration portlet

Once you are in the Health Center setup wizard, you are guided through the creation of everything you need to manage your resources from the Virtualization Engine console in the following steps:

1. In the Welcome window (shown in Figure 9-5), click **Next**. This brings you to the management sources window shown in Figure 9-7 on page 307. In this window, you should select the management sources you require from the list. This list comes from the

Web services registered in the application server. You can change the names by checking the box next to the management source name and clicking the **Rename** button.

2. Select the check box corresponding to the IBM Director management source and click **Next**. This opens the Resources tab, as shown in Figure 9-8 on page 308.

> **Note:** The Management Source default name refers to the server where the IBM Director Bridge is installed, and also includes an extension to identify which type of resource it is (_DIR, in our case). It does not refer to the server where the IBM Director Server runs. This is why in our environment we see the name of the Virtualization Engine console server and not the name of the IBM Director Server itself.

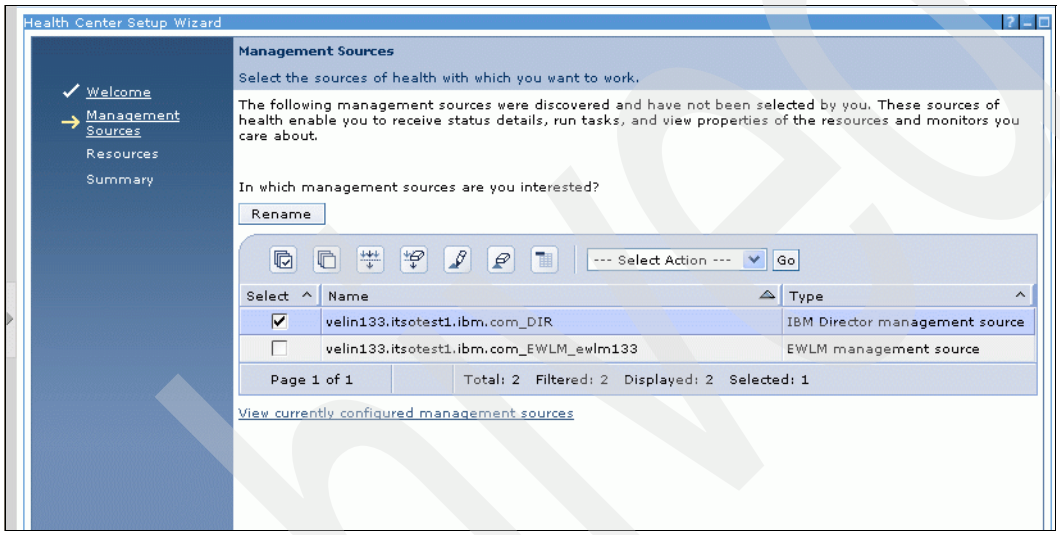


Figure 9-7 Health Center setup wizard: selecting the management sources

In the Resources tab, you can add all the resources and groups registered in the IBM Director to the Health Center, as shown in Figure 9-8 on page 308. We recommend that you *uncheck* both Groups and Resources and configure what you want to see in the Health Center at a later time. The IBM Director environment can be made of many systems, groups of systems, and alerts configured. Therefore, if you choose to add all groups and resources, you will have a heavily populated Health Center window that will be less readable.

Note: If you have no systems registered in the IBM Director Server, checking Groups and Resources has no effect.

3. Click **Next** and then click **Finish**.

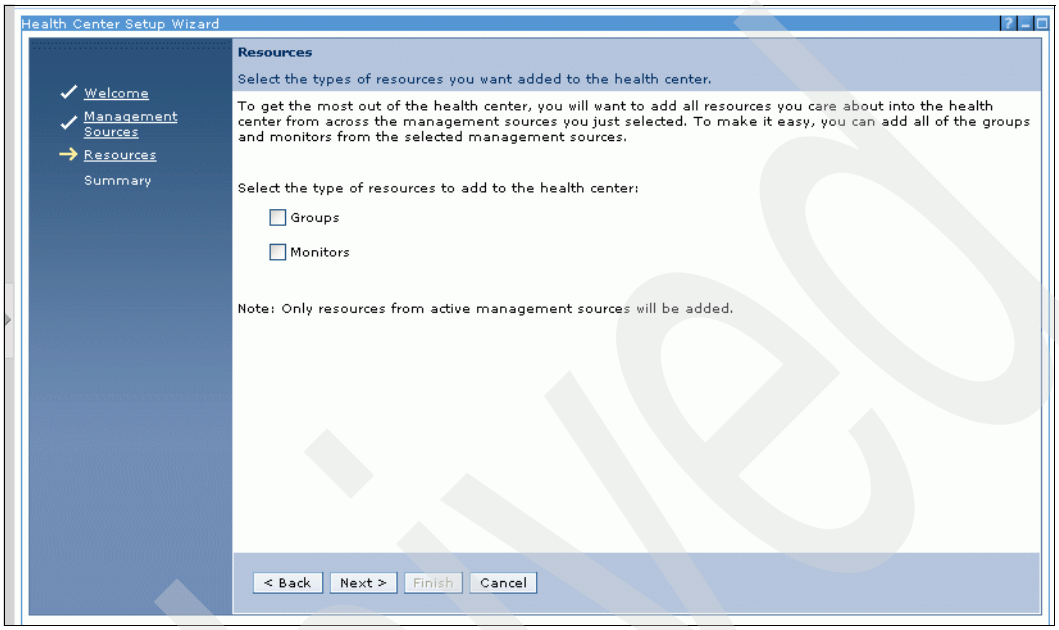


Figure 9-8 Health Center setup wizard: selecting the types of resources

4. Once the management source is configured, verify that it has been added in the list of available management sources. In the left navigation window, select **Virtualization Engine Console** → **Virtualization Engine Setup** and click **Management Sources**. You will see a window like the one shown in Figure 9-9 on page 309. The green square next to the management source name indicates that the communication has been established properly.

Note: You can configure the management sources without using the Health Center wizard from the Management Source Configuration portlet, by clicking **Add** and following the instructions. However, with this method, you will not have the ability to add resources automatically to the Resource Health.

Archived

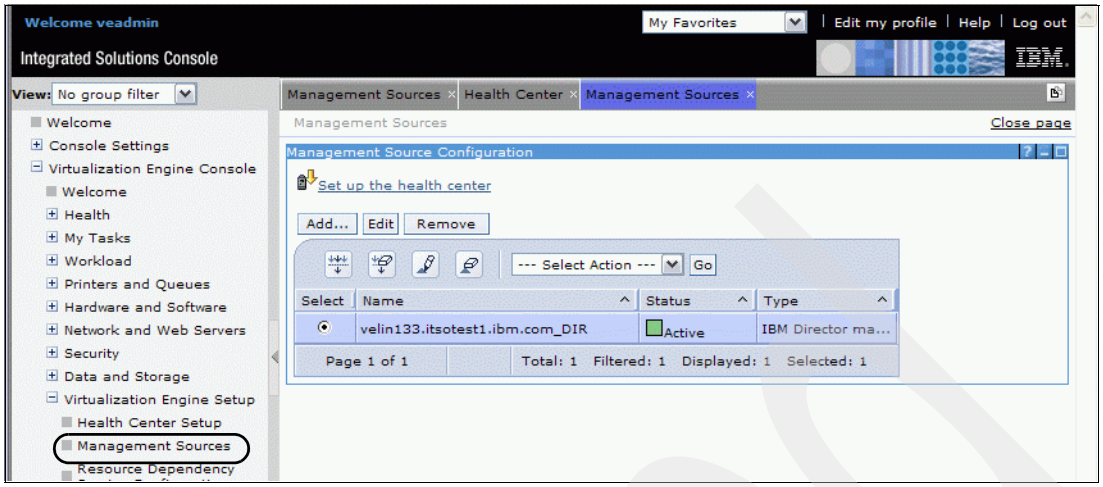


Figure 9-9 Configured management sources

9.2.2 Configuring the Hardware Management Console

IBM Director is able to automatically discover an HMC through the Service Location Protocol (SLP) agent that is installed on the HMC itself. The discovery of the HMC is followed by an attempt to access the Pegasus Common Information Model Object Manager (CIMOM) so that information regarding Central Electronic Complexes (CECs) and LPARs are gathered and represented as IBM Director objects.

In order to manage an HMC using IBM Director, there are three prerequisites:

1. The IBM Director HMC plug-in must be installed on the server and on the console.
2. The HMC must be Version 5.2 or higher. You can get the latest HMC code from the IBM support Web site at:

   `http://www.ibm.com/support`
3. You must open the Open Pegasus port (TCP 5989) in the HMC firewall.

   To open the necessary ports of the HMC firewall, open the WebSM and connect to the target HMC providing the correct credentials. Select **HMC Management** → **HMC Configuration** and click **Customize Network Settings** (encircled in Figure 9-11 on page 310).

> **Note:** WebSM is a Java application required to connect to the HMC. You can install it from `http://<HMC>/remote_client.html`, where *<HMC>* is the HMC IP address or host name.
>
> You can also install it from the Virtualization Engine console by selecting **Virtualization Engine Console** → **My tasks** → **Launchpad**. Click **Web-based System Manager** (as shown in Figure 9-10) to install the application for you if it is not already installed.
>
> You can configure the settings from the HMC physical console as well.

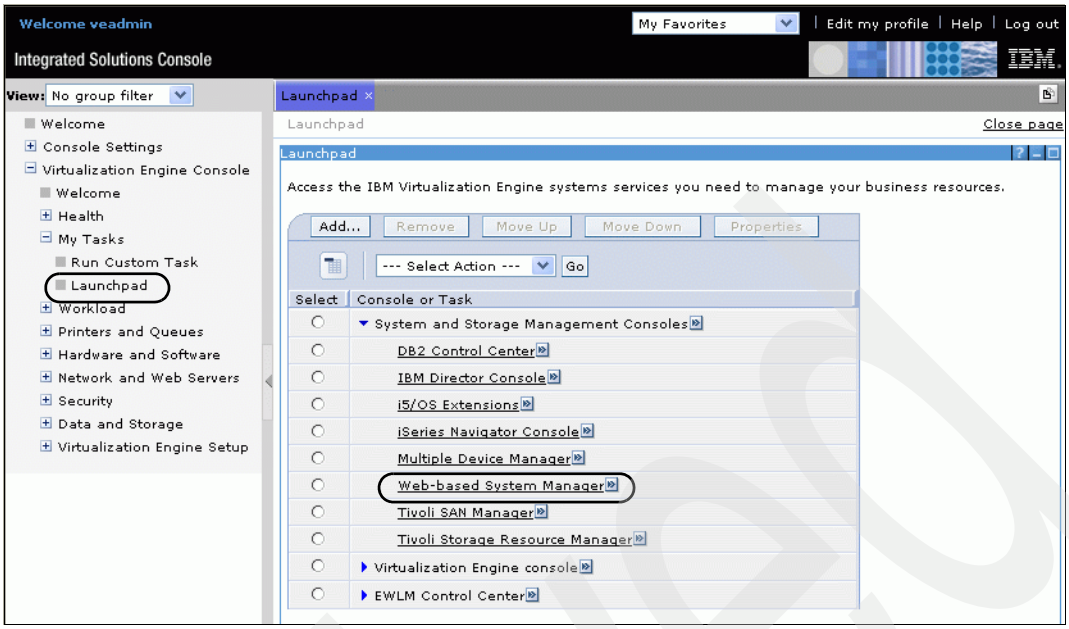


Figure 9-10 Virtualization Engine console launchpad

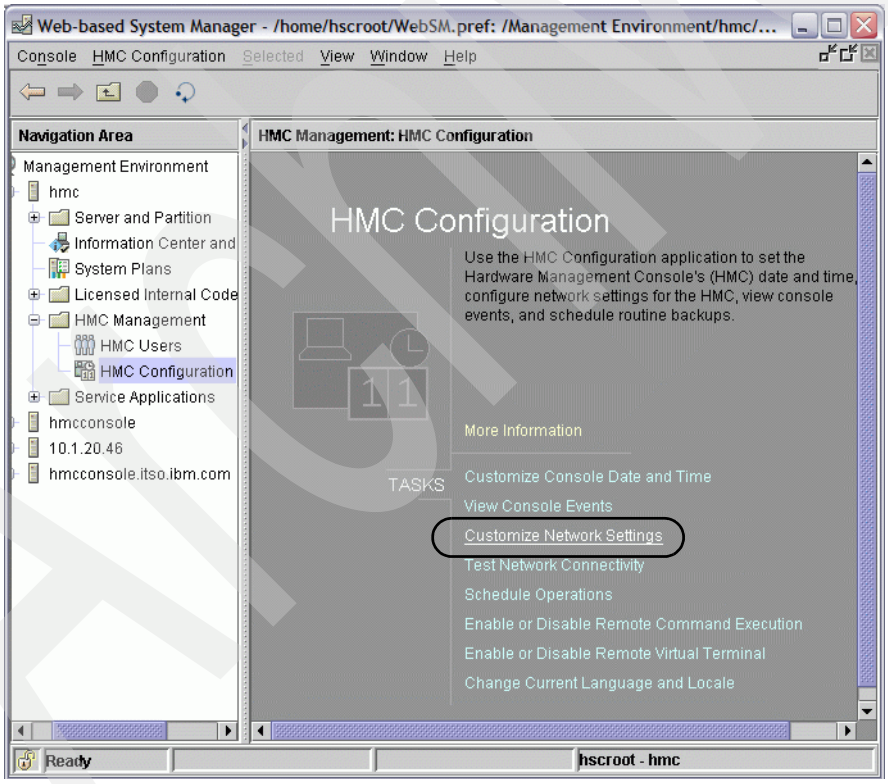


Figure 9-11 HMC network configuration

4. In the **Customize Network Settings** window shown in Figure 9-12 on page 311, select the Ethernet card connected to the management environment network in the LAN Adapters tab. One adapter is always connected to a private network segment for Management Processor communication. Click **Details**.

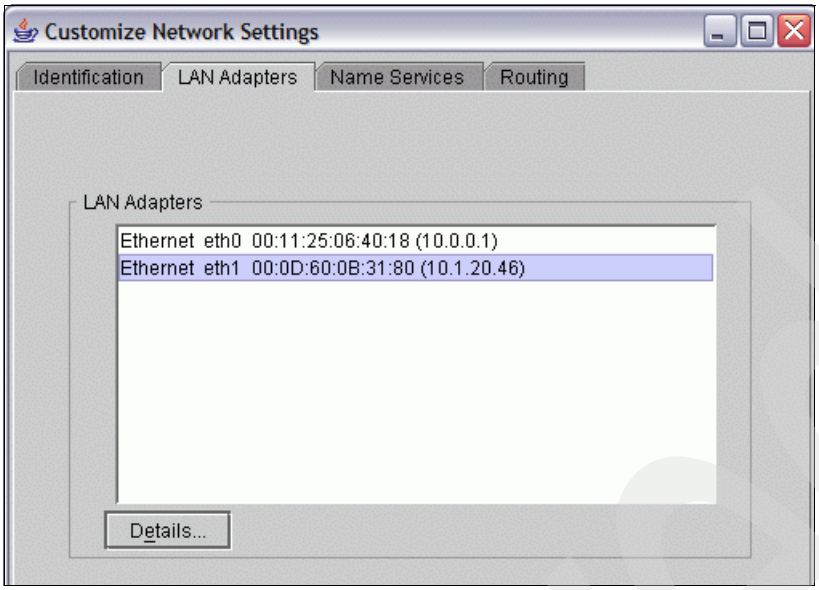


Figure 9-12 HMC LAN Adapters configuration

5. In the LAN Adapter Details window that opens, as shown in Figure 9-13, select the **Firewall** tab. The upper pane contains the applications for which communication is allowed and the lower pane contains the applications for which communication is forbidden. Select the application called **Open Pegasus** in the lower pane and click **Allow Incoming**. This will move the Open Pegasus application into the upper pane. Click **OK**.

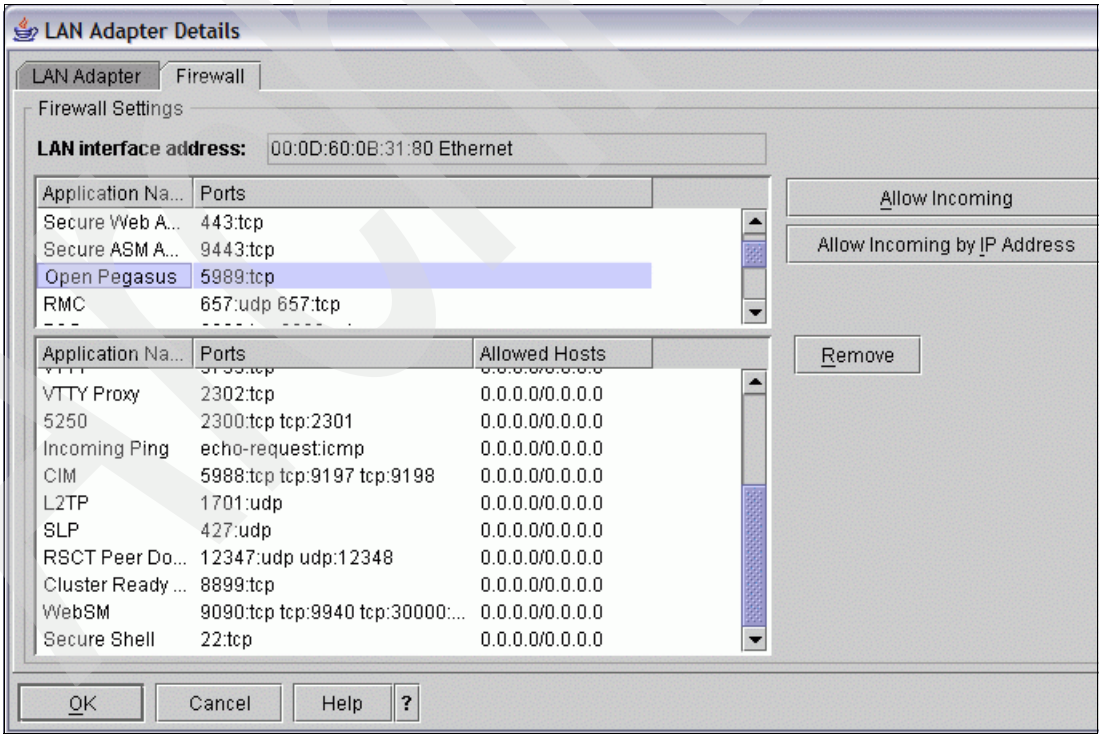


Figure 9-13 HMC firewall configuration

9.2.3 Discovering systems with IBM Director

IBM Director is able to perform an automatic discovery of the systems in your network. To discover your systems, you have two choices:

- Open the IBM Director Console and launch the Discovery task.

  To do this, open the IBM Director Console from your workstation and click the discovery icon in the upper taskbar, as shown in Figure 9-14. Refer to *Implementing IBM Director 5.10*, SG24-6188 for more details.

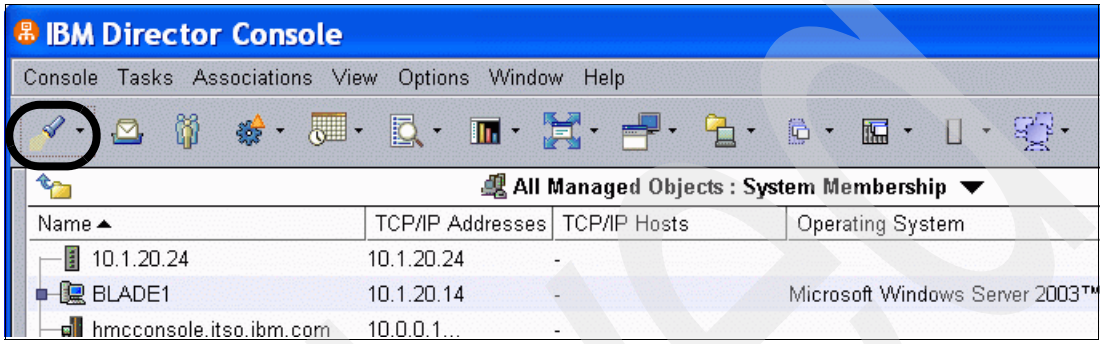


Figure 9-14 IBM Director Discovery task

- Launch the discovery from the Virtualization Engine console directly.

  To do this, select **Virtualization Engine Console** → **My Tasks** in the navigation pane, and select **Run Custom Task**. The list of **Available Actions** appears, as shown in Figure 9-15. Click **Discover Systems**.

> **Note:** This method works only if the IBM Director Bridge is properly configured.

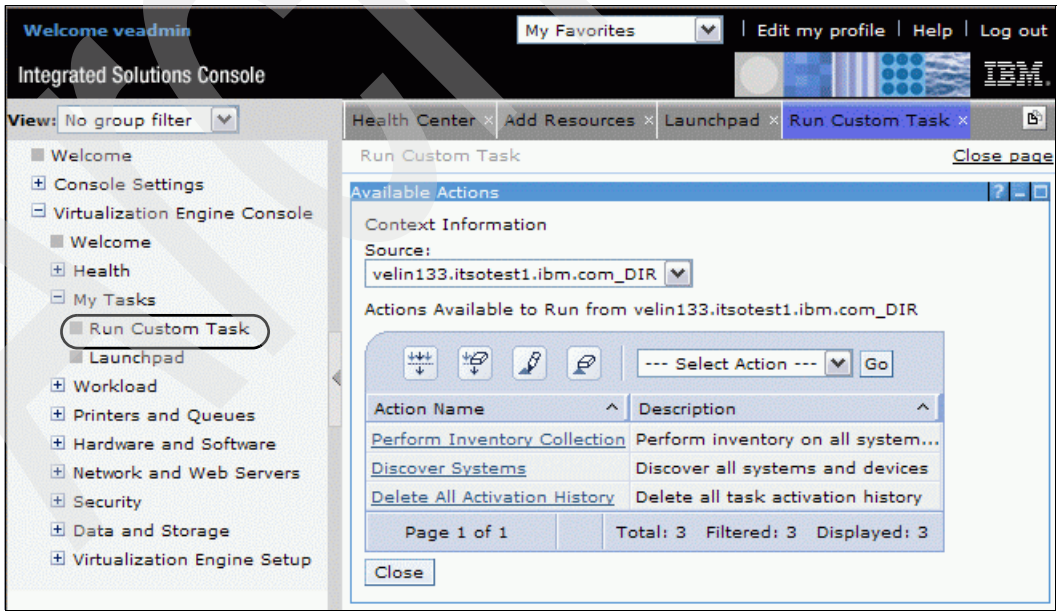


Figure 9-15 Run Custom Task actions

The **Selected Action to Run** portlet appears, showing you the optional input parameters you can select. In our case, there is nothing else to configure, so click **Next** and then click **Run**, as shown in Figure 9-16.

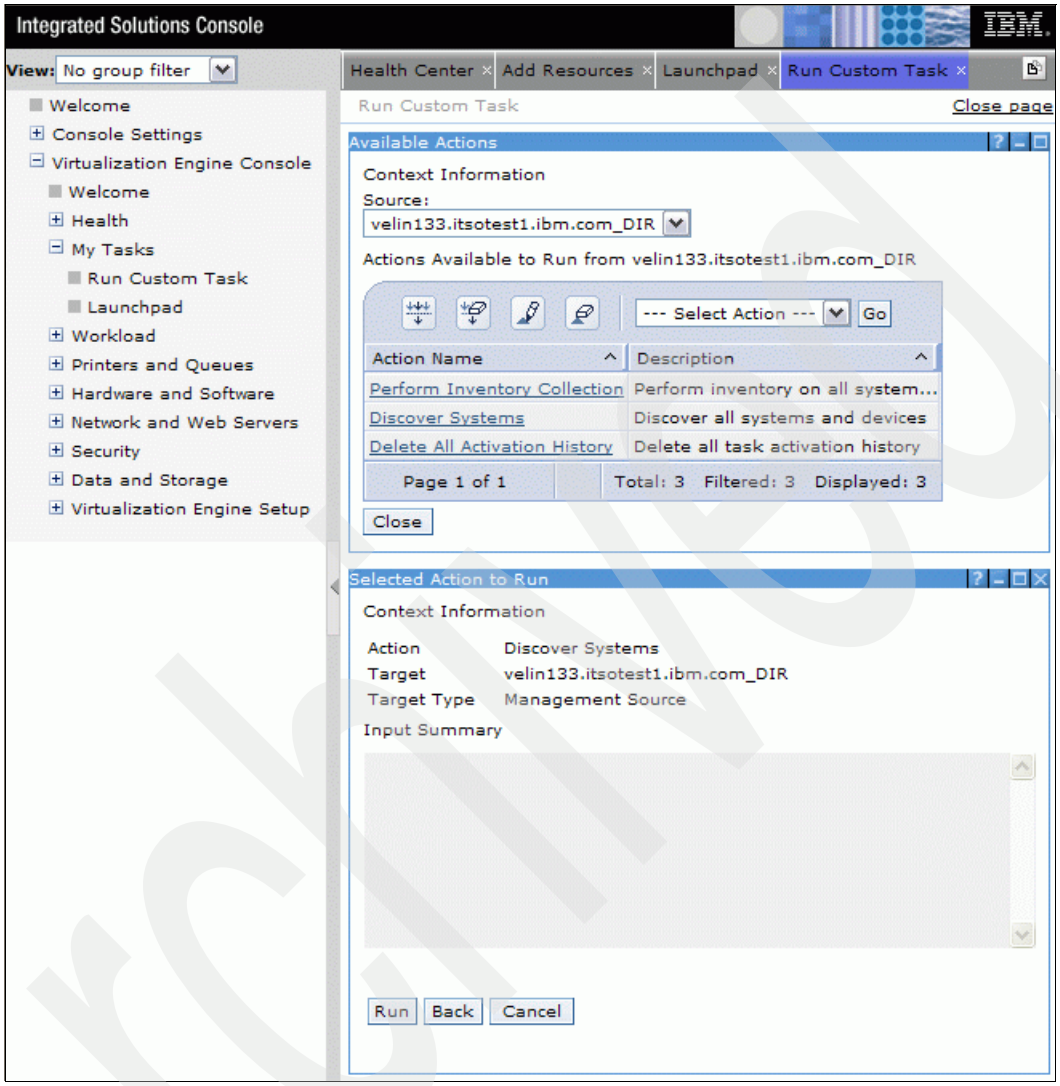


Figure 9-16 Perform Discovery action details

You will get a notification that the action has been submitted, as shown in Figure 9-17 on page 314.

During this process, the Virtualization Engine console has been instructed to launch the discovery command to the IBM Director Server through Web services.

Archive

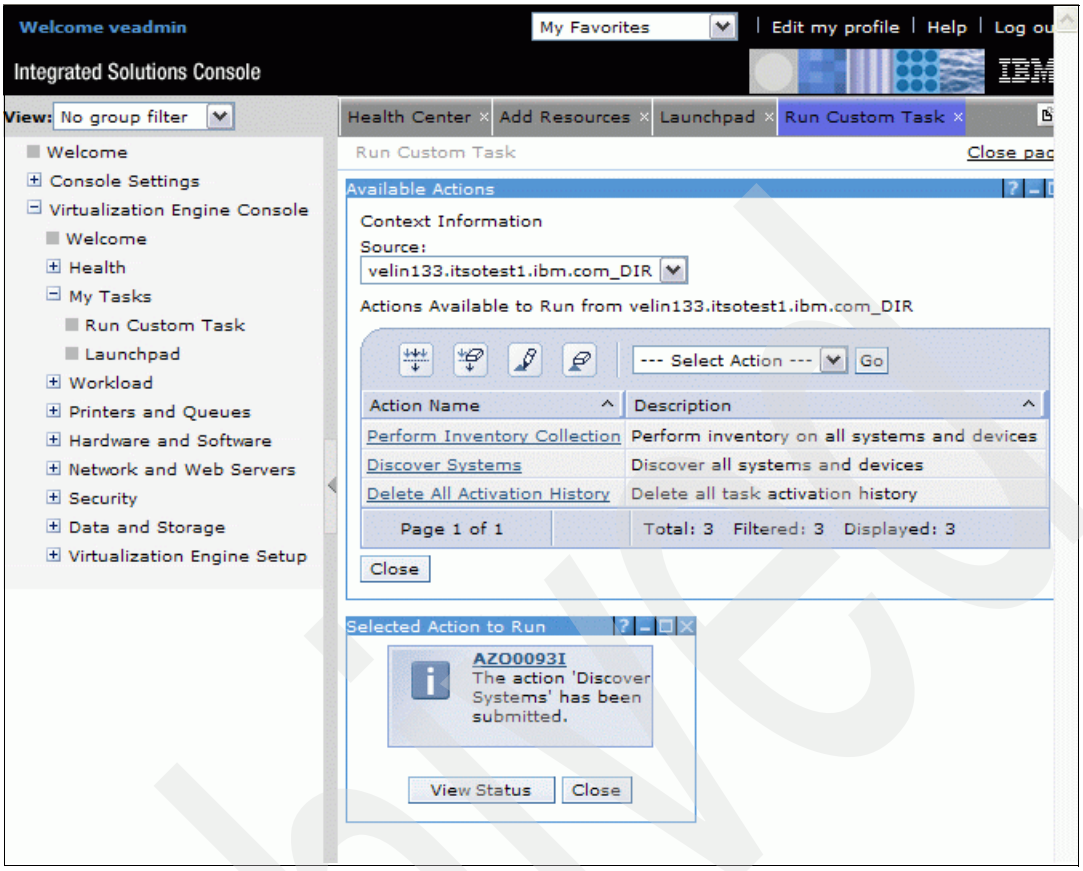


Figure 9-17 The Discovery Action has been submitted

To check the status of your action, click **View Status** in the Selected Action to Run portlet. This opens the Task Status portlet, which you can also open from the Health Center by clicking **Task Status** (encircled in Figure 9-18 on page 315). Here, you see the status of the action. If you click the arrow next to the name of the action launched and select **Properties**, you will see the details log for the launched action.

Note: If your environment comprises several networks connected by routers or firewalls, then you might have to configure additional parameters in the IBM Director in order for the discovery to work. You can find more details on discovery in Chapter 5, “IBM Director configuration“, of *Implementing IBM Director 5.10*, SG24-6188.

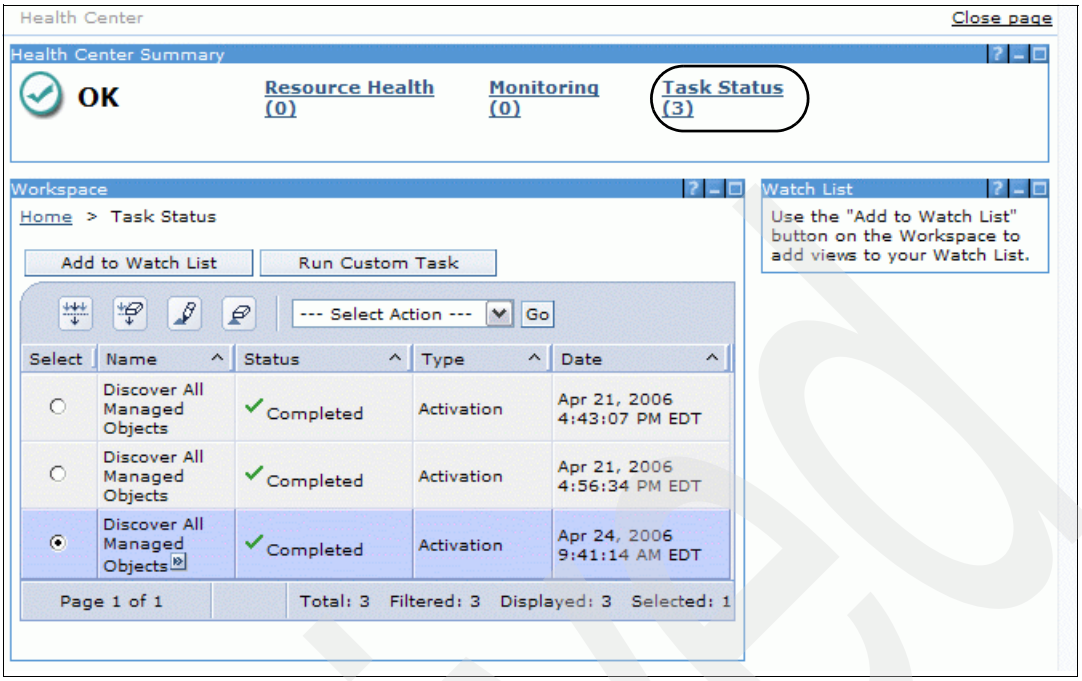


Figure 9-18 Task status portlet

9.2.4 Requesting access to systems in IBM Director

The discovery process in IBM Director registers all the systems that IBM Director is able to manage. However, this does not mean that it is automatically able to access them. In fact, for security reasons, systems may require you to insert the proper credentials to manage them, depending on the object type and the options you selected when you installed the management agents. IBM Director security is out of the scope of this book. For more details, refer to Chapter 3 in *Implementing IBM Director 5.10*, SG24-6188.

There is no way to gain access to the systems from the Virtualization Engine console, so you should open the IBM Director Console. There are two methods by which you can do this:

- Launch it from your workstation if you have it installed. Refer to *IBM Director Installation and Configuration Guide* for more details on IBM Director Console installation.
- Use the Virtualization Engine console launchpad. To do this, open your browser and connect to the Virtualization Engine console portal. In the navigation pane, expand **Virtualization Engine Console** → **My Tasks** and click **Launchpad**. The Launchpad portlet shown in Figure 9-19 on page 316 appears. Click **IBM Director Console** to launch it. If you have the console installed on your workstation, then this launches the console; otherwise, you will be asked if you want to install the console on your system.

> **Important:** The IBM Director Console code is stored in a .jar file on the system running the Virtualization Engine console portal. This code is Version 5.10 and does not include any extensions, nor the BladeCenter or Rack Manager options. If you choose this installation method, you should add patches and extensions to your console manually.

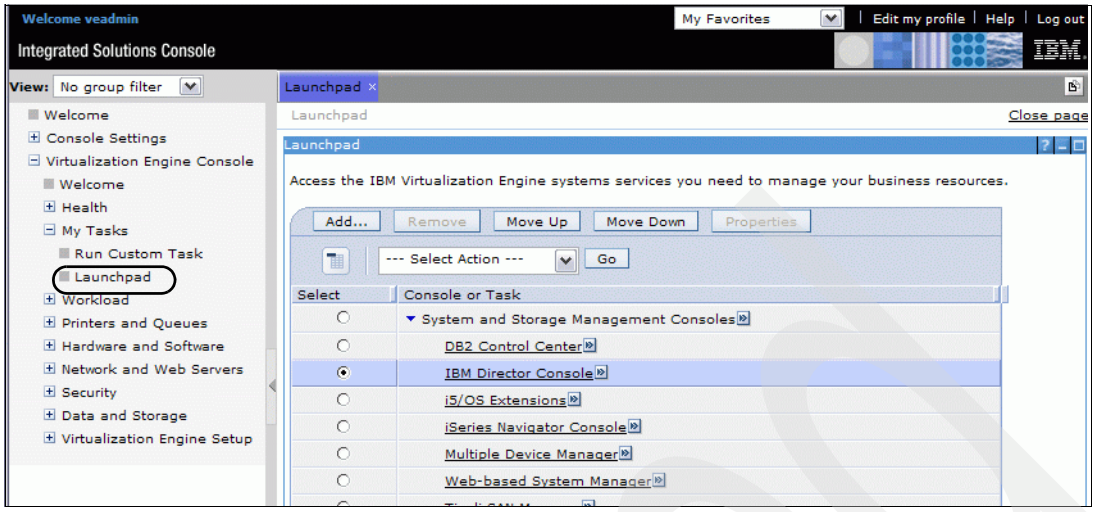


Figure 9-19 Virtualization Engine console launchpad

Once you are in the IBM Director Console, as shown in Figure 9-20, you can identify the locked systems (those with a symbol next to the system icon). These are the systems which the IBM Director is not allowed to perform actions on because of missing credentials. Right-click one or multiple systems, and select **Request Access** from the menu, as shown in Figure 9-20. Enter the correct credentials and click **OK**.

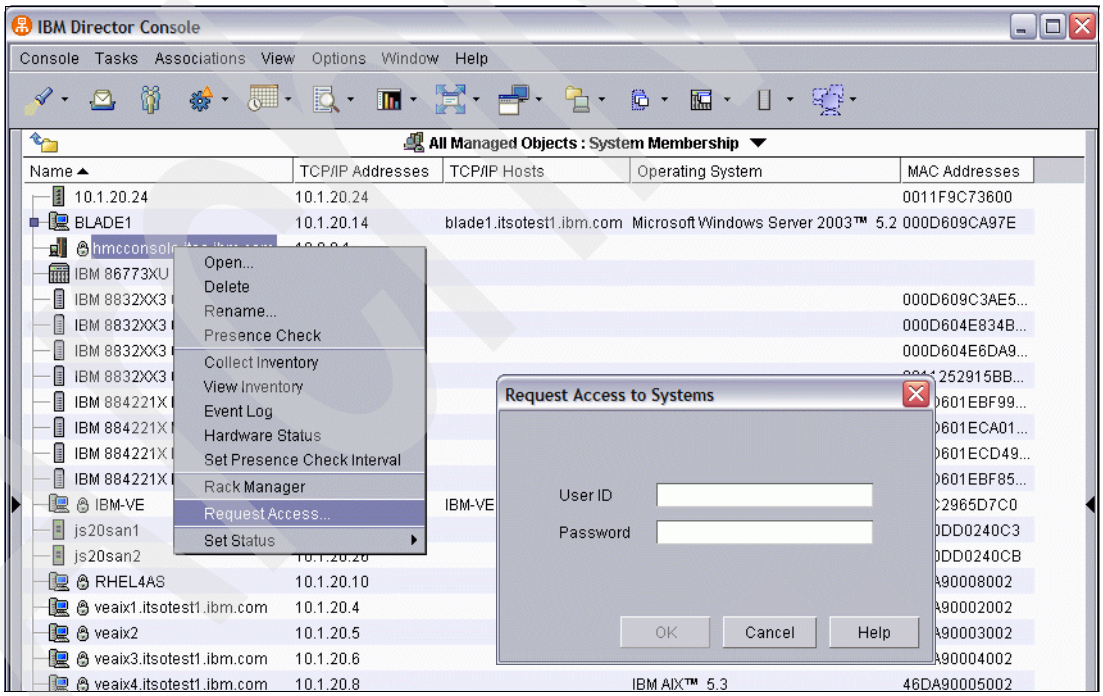


Figure 9-20 Request access to a locked object in the IBM Director Console

Tip: If you have many locked systems requiring the same credentials, you can select more than one by keeping the Ctrl button pressed and selecting the systems that you want to request access to. For the Request Access item to appear in the menu, select locked objects only.

An alternative way of requesting access to systems at one time is to use the IBM Director command-line interface. It is an executable program installed with every IBM Director component. From a command prompt on the IBM Director Server host (or from a remote shell), run the following command:

```
dircli lsmo --groups "NAME" | dircli accessmo –u USERID –p PASSWORD -f –
```

In this command, `NAME` is the group name containing the systems you want to request access to (brackets are required only if the group name includes spaces). `USERID` is the user name, `PASSWORD` is the password.

This command instructs **`dircli`** to list all the managed objects in the group called `NAME` and use them as argument for the **`dircli accessmo`** command.

If you do not include the –u and –p switches, you are asked for the USERID and PASSWORD at run time.

9.2.5 Renaming BladeCenter

When IBM Director discovers systems, it creates different objects based on the system type. For example, when a BladeCenter is discovered, IBM Director communicates with the Management Module and queries it for the objects inside the BladeCenter itself. It then creates a chassis object representing the management module and physical platform objects to represent blade servers installed into the chassis itself. The default names assigned to these objects are based on the machine type and serial number of the system. However, this is not the best way of identifying the systems at first sight in the console. In this section, we show you how to automatically rename the systems so that the name contains the slot number where the blades are installed. You also have the ability to rename these systems manually according to any criteria you find useful for systems identification, for example, the application the server is hosting.

The naming convention we use here is simply to append the slot number the server is installed into to the default name. Therefore, the name of each blade server is in the *IBM-MachineType-SerialNumber-SNumber* format, where MachineType is the four-digit machine type of the blade server, *SerialNumber* is its seven-digit serial number, and *Number* is the slot number where the blade server is installed.

To change the naming convention, perform the following steps:

1. Open the IBM Director Console (either directly or from the Virtualization Engine console) and select **Options** → **Discovery Preferences**, as shown in Figure 9-21.

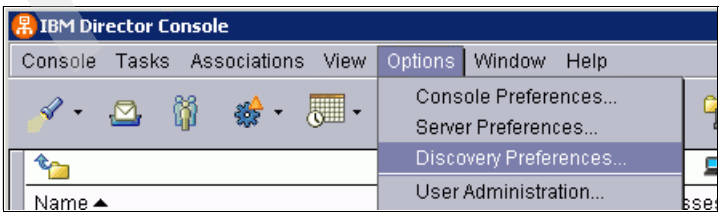


Figure 9-21 Options menu in IBM Director Console

2. The Discovery Preferences window opens, as shown in Figure 9-22.

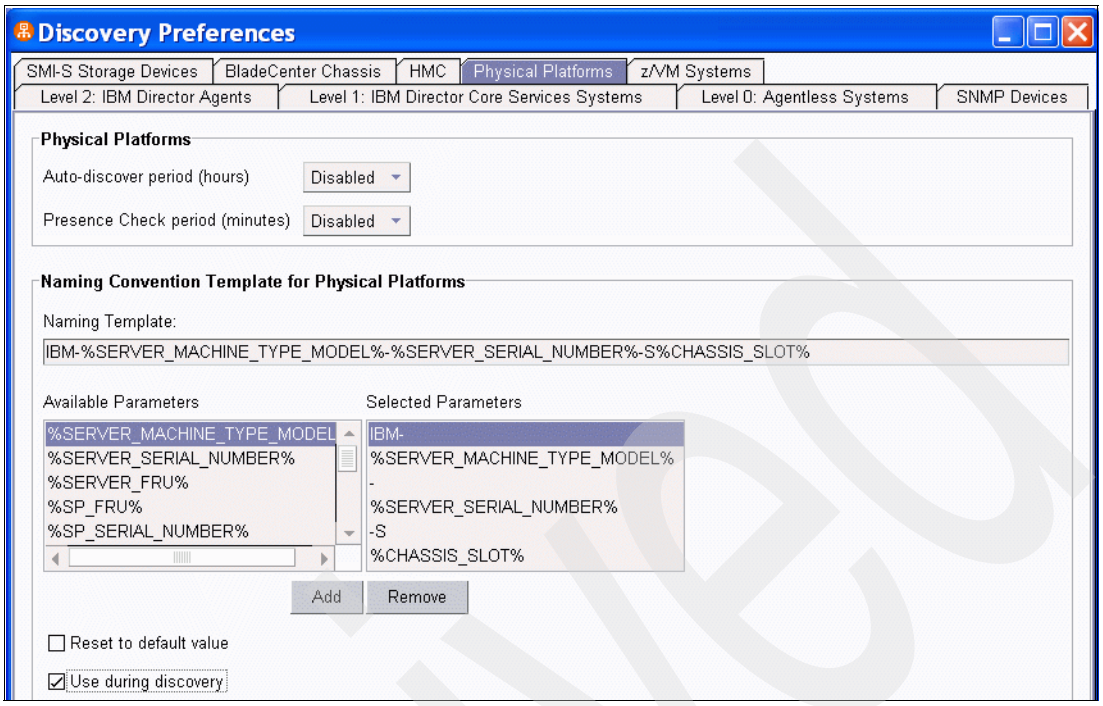


Figure 9-22 Discovery Preferences window

Select the **Physical Platforms** tab. This is where you should build the name for a Physical Platform object.

In the Naming Template field, you have the currently used template for naming the objects. It is a combination of text strings and variables (included between the % symbol). This field is not editable.

In the Available Parameters column, you have the list of the available parameters. These are variables whose values change according to the object properties. For example, %SERVER_MACHINE_TYPE_MODEL% is the four-digit machine type, and so on.

In the Selected Parameters column, you see the parameters currently selected. In this column:

- To add a parameter:
  i. Select the chosen parameter in the Available Parameters column.
  ii. In the Selected Parameters column, select the parameter right after which you want the new one to be added.
  iii. Click **Add**.
- To remove a parameter:
  i. Select the parameter in the Selected Parameters column.
  ii. Click **Remove**.
- To change any line in the Selected Parameters column, you must remove it and add the new parameter you want.

To add text not bound to any variable, select **Insert additional text** in the Available Parameters column and click **Add**. The Input Dialog box shown in Figure 9-23 on page 319 appears, asking you to type the text you want to add. Type the text and click **OK**.

Archived

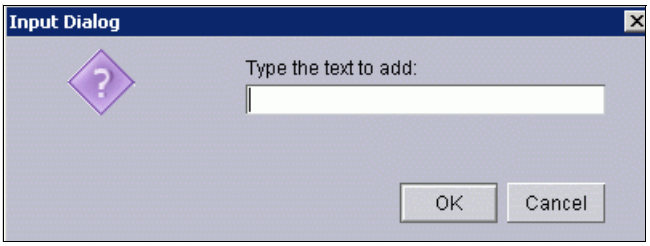


Figure 9-23 Input dialog box to type the text to add

With respect to the default values, change the spaces to hyphens for clarity and append the text “-S” (S stands for Slot), followed by the %CHASSIS_SLOT% parameter to easily identify the location of the blade server.

If you check the box next to **Use during discovery**, the new naming convention is applied during the discovery process for all the physical platform objects. This is not the best option when you have other servers in your environment that are not blades, as the CHASSIS_SLOT parameter will be blank. In this case, we recommend that you do not check this box. Click **OK** to save your changes.

Note: This operation does not change the names of the currently discovered systems.

3. The next step is to rename the systems in the console.

   Select the systems you want to rename. In our case, we selected the group **Chassis and Chassis Members** (encircled in Figure 9-24). We then selected all the Blade Servers objects. Right-click one of the selected objects and select **Auto Rename**, as shown in Figure 9-24.

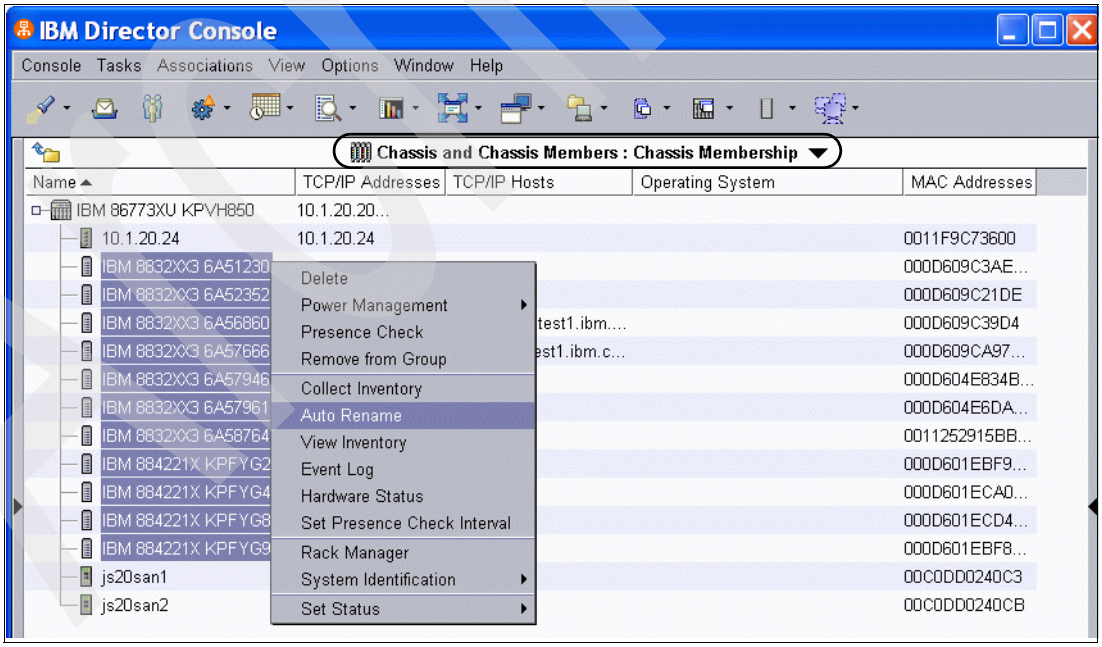


Figure 9-24 Auto rename the blade servers

4. A window appears asking whether you want to **Schedule** or **Execute Now**, as shown in Figure 9-25 on page 320.

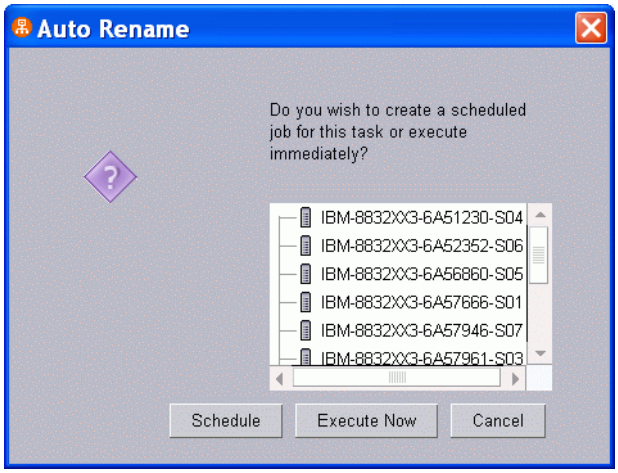


Figure 9-25 Auto Rename task confirmation window

Click **Execute Now**.

5. The Execution History window appears, as shown in Figure 9-26

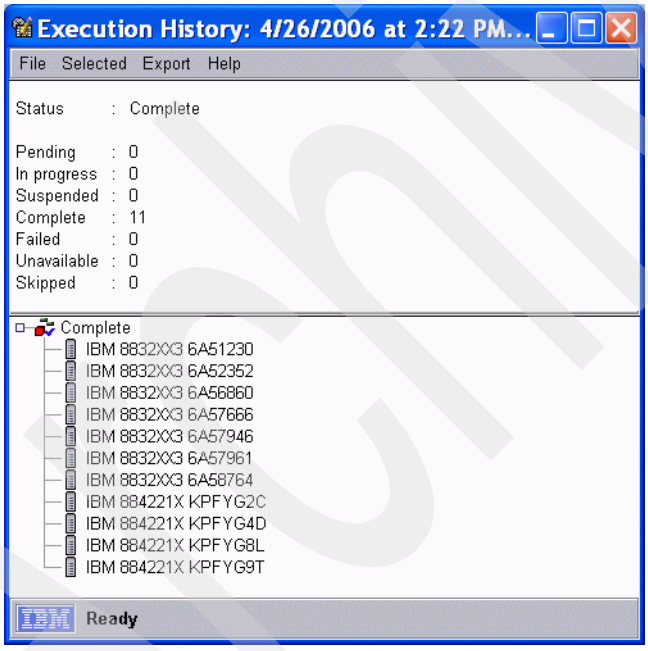


Figure 9-26 Execution History window

You can close this window at any time.

Archived

When the action is complete, you will see that the blade servers are renamed, as shown in Figure 9-27.

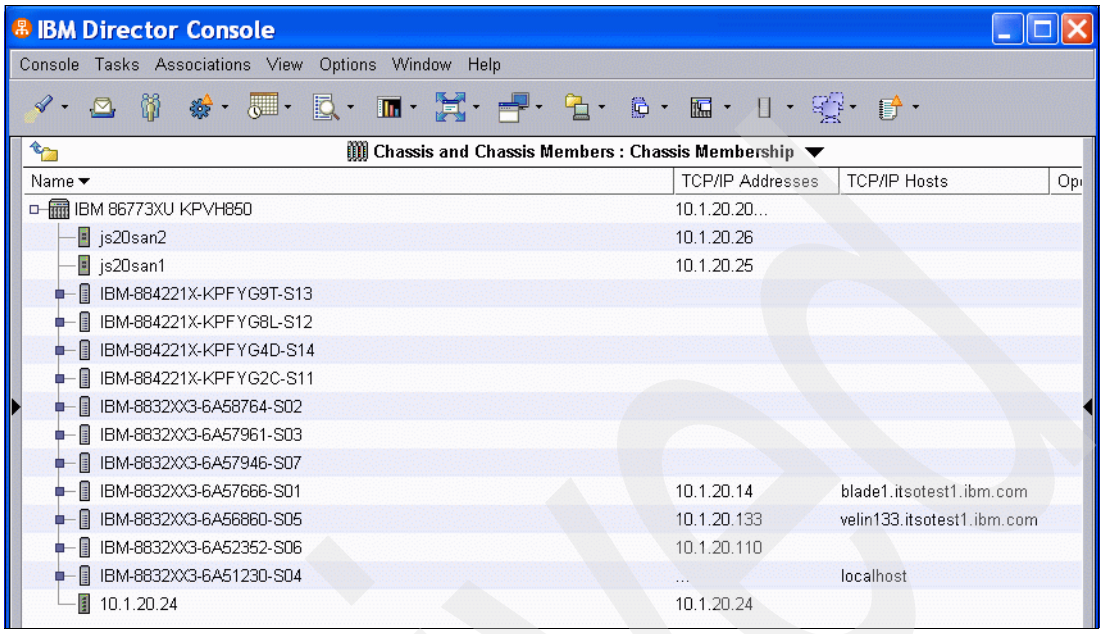


Figure 9-27 IBM Director Console showing the renamed blade servers objects

This renaming is reflected in the Virtualization Engine console as well.

Tip: You can automatically rename BladeCenters as well, using the same method, by selecting the **BladeCenter Chassis** tab in the Discovery Preferences window, as shown in Figure 9-22 on page 318.

9.2.6 Grouping resources

Beyond the default groups available in IBM Director, you can also create your own group definition, for example, to include the systems that are part of the same business application for easy access. You can do this using the IBM Director Console. This section does not describe this method. For details, refer to *Implementing IBM Director 5.10*, SG24-6188. You can also create your own group definition from the Virtualization Engine console by performing the following steps:

1. Expand **Virtualization Engine Setup** → **Resource Grouping**, as shown in Figure 9-28 on page 322.

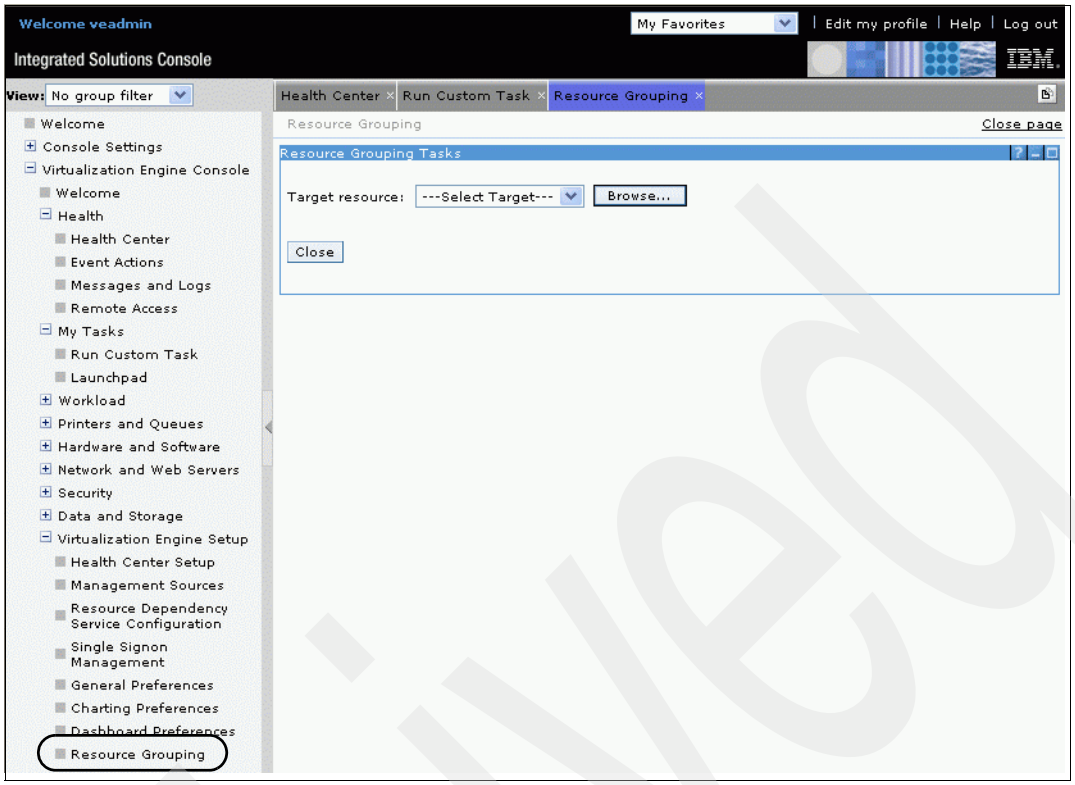


Figure 9-28 Resource Grouping configuration portlet

2. The Virtualization Engine console covertly launches the IBM Director task to create new groups. You will be asked to select a **Target Resource** in which context the IBM Director task is to be launched. This task does not need to be launched against a resource, so you can select any target. Click **Browse** to see the list of selectable systems, as shown in Figure 9-29 on page 323, and select the first one. Then click **OK**.

> **Note:** You are asked for the target resource only the first time you open this window. When you successfully select a **Target Resource**, it is kept for future usage.

Archived

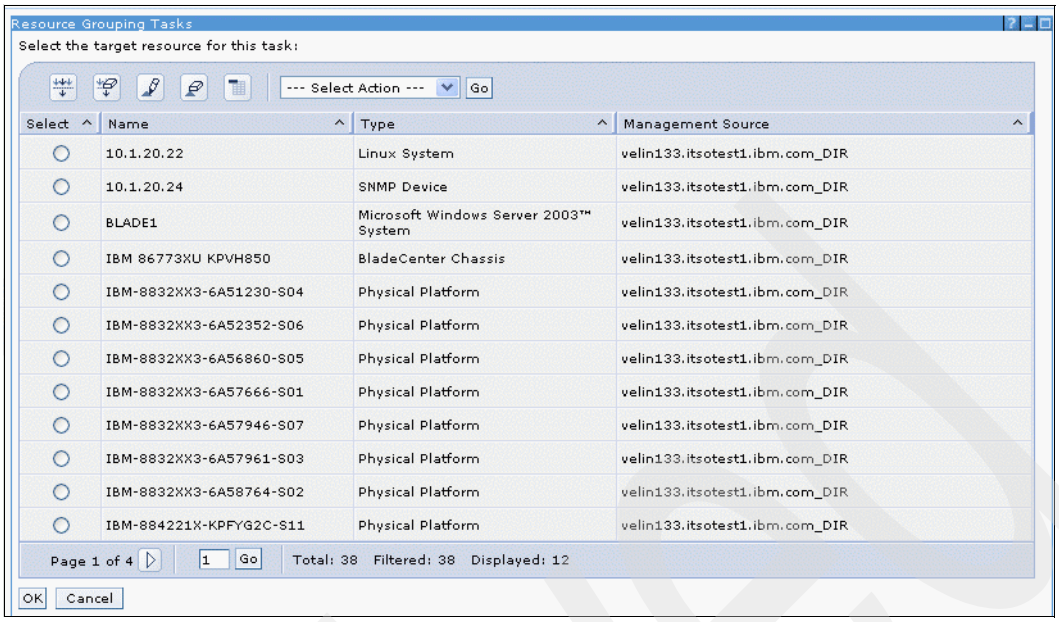


Figure 9-29 Resource Grouping Tasks: target selection

You see now the list of Resource Grouping Tasks available, as shown in Figure 9-30.

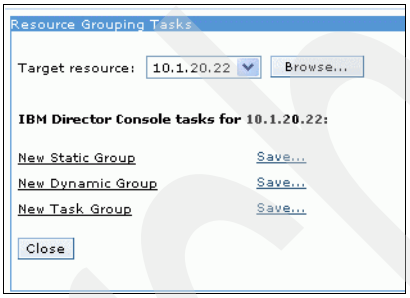


Figure 9-30 Resource Grouping Tasks available

3. Click **New Static Group**. This opens the IBM Director Static Group Editor window, as shown in Figure 9-31 on page 324.

Archive

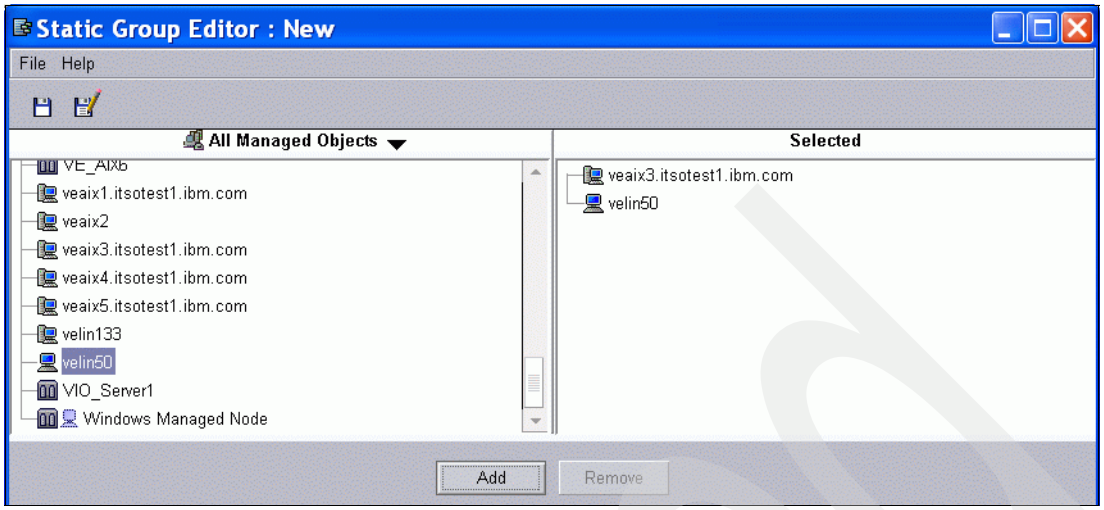


Figure 9-31 Static Group Editor: new window

Here you select the list of systems you want to add in the **Static Group Editor** by selecting the objects in the left column and clicking **Add**. In this example, we want to create a group containing the Web server and application server for our intranet. Select **File** → **Save** and insert a descriptive name for the group in the Save As input window shown in Figure 9-32. Click **OK** to create the group.



Figure 9-32 Save As input box

To add the new group to the list of Health Center managed resources, refer to 9.3.1, “Configuring resources in the Health Center” on page 331.

Refer to *Implementing IBM Director 5.10*, SG24-6188 for more details on grouping resources in IBM Director.

9.2.7 Changing the IBM Director Bridge configuration

You may need to change which IBM Director Server the bridge is communicating with, either because you changed the server name, or because your environment has grown and for performance reasons you need to switch from a shared installation (with all the other Virtualization Engine components) to a dedicated installation.

The IBM Director Bridge is a software running as an application in the WebSphere Application Server. Hence, to change its configuration, you must use the WebSphere Application Administration console. You have to go through the following steps to change the IBM Director Server host name:

1. Open a Web browser and go to the following Web site:

   `http://hostname:portnumber/admin`

Here, `hostname` is the Virtualization Engine console server host name and `portnumber` is the port used by the WebSphere Application Administration console. 18552 is the default value of the port number default.

> **Tip:** To find the WebSphere Application Administration console port number, open the VEInstall.properties file in the Virtualization Engine installation root directory. The value of the *ibm.ve2.VEMS.adminPort* variable is the running session administration port.
>
> To ease the navigation in the Virtualization Engine configuration files, the Virtualization Engine installation process defines an environment variable named *IBM_VE2_LOCATION*, which contains the Virtualization Engine installation root directory.

2. Log in with the Virtualization Engine administrator credentials.
3. In the left navigation pane, expand **Applications** and select **Enterprise Applications**. This opens the Enterprise Applications portlet shown in Figure 9-33.

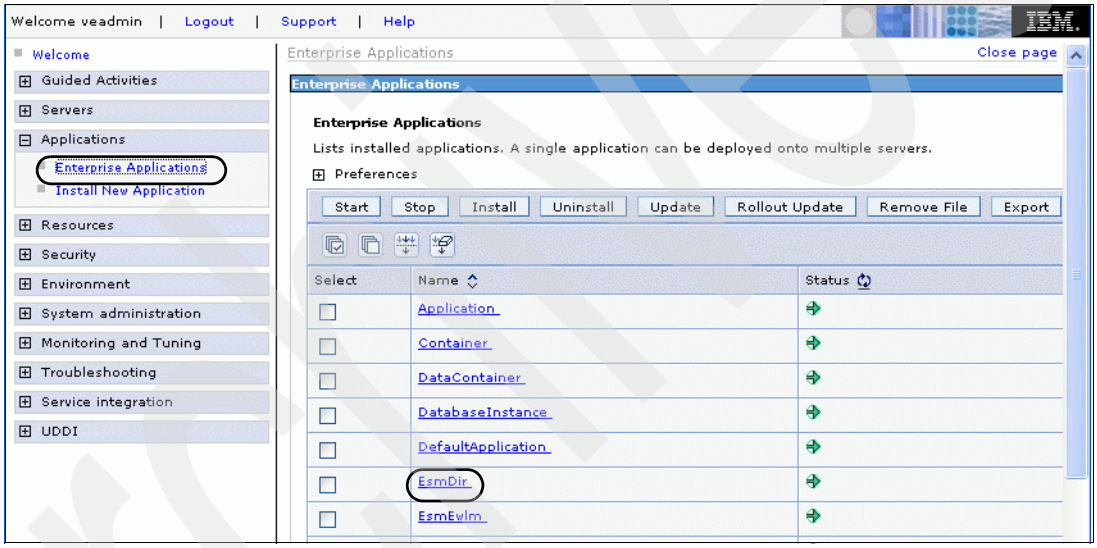


Figure 9-33 WebSphere Application administration console

4. Click **EsmDir** (encircled in Figure 9-33).
5. Click **Connector Modules**.
6. Click **DirectorJca.rar**.
7. Click **Resource Adapter**.

8. The configuration window appears. Go to the bottom of the page and click **Cancel**, as shown in Figure 9-34.

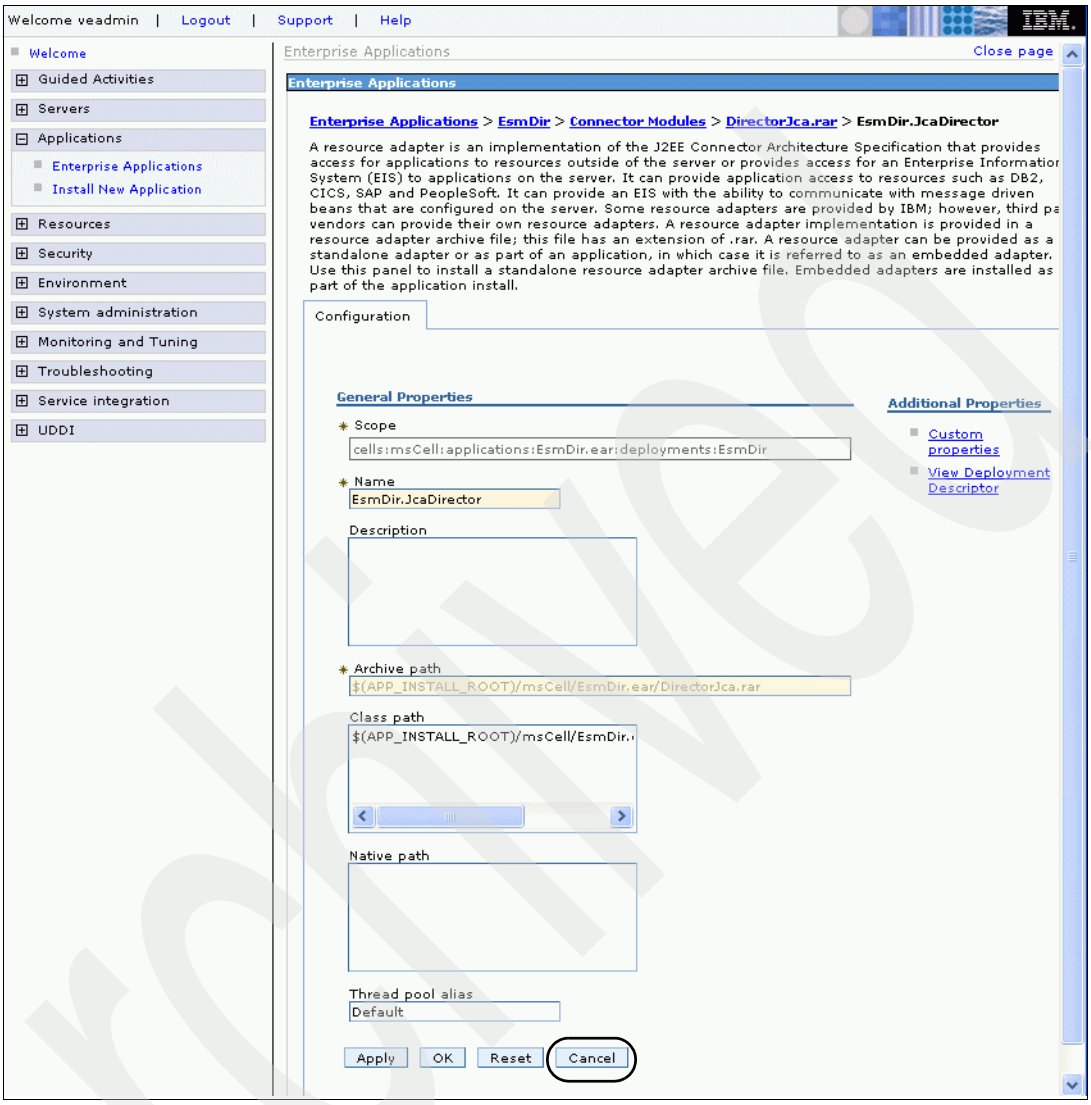


Figure 9-34 Configuring the DirectorJca.rar

9. Click **Resource Adapter**.
10. Click **J2C connector factories**.
11. Click **Virtualization Engine consolejca_EsmDir**.
12. Click **Custom properties**.
13. You will see a list of properties, as shown in Figure 9-35 on page 327. Click **DirectorHost** and modify the host name accordingly.

> **Note:** The IBM Director Server host name must be resolved by the host where the IBM Director Bridge is running. Either configure the Domain Name System (DNS) server properly or add the host name and IP address to the hosts file of the server running the IBM Director Bridge.

Archive

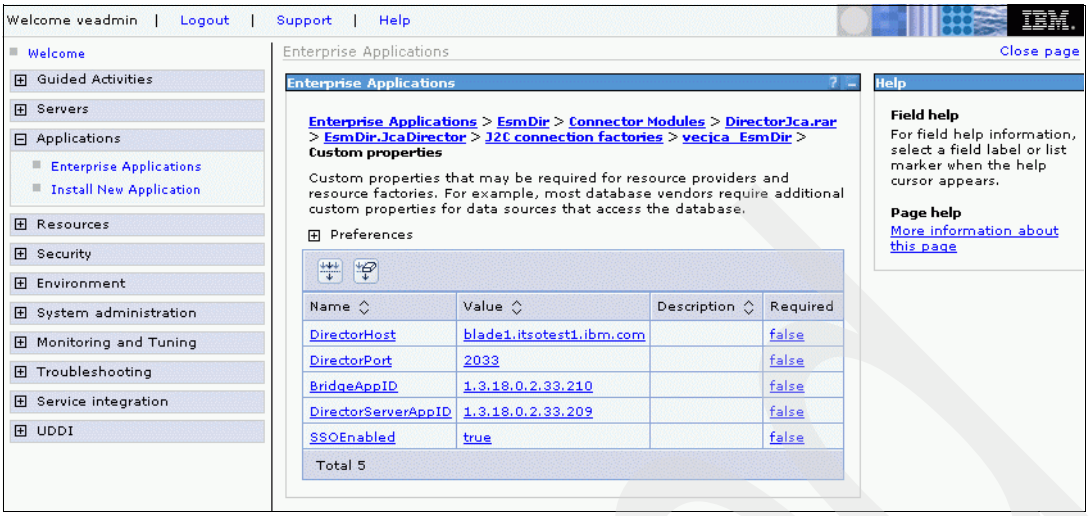


Figure 9-35 IBM Director Bridge properties window

14. Click **OK,** then click **Save**. Click **Save** once again.

    During the saving process, a warning message appears stating that you may need to restart the application server. This is not necessary. However, if you want to restart it, refer to 9.8, “Starting and stopping services” on page 375.

15. Now that the host name has been changed, you should configure the proper credentials that the bridge uses to access the new IBM Director Server. You should do this from the Single Sign-on configuration menu. However, before configuring the proper credentials, you should run a few commands. On the server running the IBM Director Bridge, open a command prompt and run the following command:

    ```
    IBM_VE2_LOCATION/bin/vessoconfig -action setAppInstance -appid
    "1.3.18.0.2.33.209" -host ibmdirectorserverhostname
    ```

    In this command, `ibmdirectorserverhostname` is the host name of the new IBM Director Server.

    This command adds the new IBM Director Server system host name to the list of hosts for which the application IBM Director (whose ID is 1.3.18.0.2.33.209) is configured. To view the complete list, enter the following command:

    ```
    IBM_VE2_LOCATION/bin/vessoconfig -action viewAppInstance -appid
    "1.3.18.0.2.33.209"
    ```

    Tip: Type the **`IBM_VE2_LOCATION/bin/vessoconfig`** command for a complete list of commands available.

16. On the server running the bridge, enter the following command from a command prompt:

    ```
    IBM_VE2_LOCATION/bin/vessoconfig -action cfgSysRegistry  -host
    ibmdirectorserverhostname -mappingRegType OS_TYPE
    ```

    In this command, `ibmdirectorserverhostname` is the host name of the new IBM Director Server, and OS_TYPE is its operating system type. The list of valid operating system types are:

    – WINDOWS
    – OS400
    – AIX
    – LINUX

This command adds the new IBM Director Server system to the list of systems for which you are allowed to configure credentials to log in to, in the Single Sign-on Management configuration portlet (as we do in step 19).

17. If the IBM Director Server and the bridge are running on the same host, you should run the following command on the Virtualization Engine console server (which is also hosting the bridge):

    ```
    IBM_VE2_LOCATION/bin/vessoserverconfig
    ```

    There is no other step you have to do. You can skip the remaining items in this list.

    If the IBM Director Server and the bridge are running on different hosts (such as in our case), then run the following command:

    ```
    IBM_VE2_LOCATION/bin/vessoserverconfig ibmdirectorserverhostname
    ```

    In this command, `ibmdirectorserverhostname` is the host name of the new IBM Director Server. This generates the vessoclientinfo.cfg in the current directory. This file contains connection information for the IBM Director Server to accept connections from the SSO servlet.

18. Copy the vessoclientinfo.cfg file created in step 17 into the IBM Director Server host in the following directory, if you are using a Windows system:

    ```
    C:\Documents and Settings\All Users\Application
    Data\IBM\VE\var\ve\vessoclientinfo.cfg
    ```

    If your system is an AIX, i5/OS, or Linux system, copy the vessoclientinfo.cfg file into the following directory:

    ```
    /var/VE/vessoclientinfo.cfg
    ```

19. The last step configures the credentials that the bridge uses to access the IBM Director Server. To do so, in the Virtualization Engine console navigation window:

    a. Expand **Virtualization Engine Setup** and click **Single Sign-on Management**, as shown in Figure 9-36.

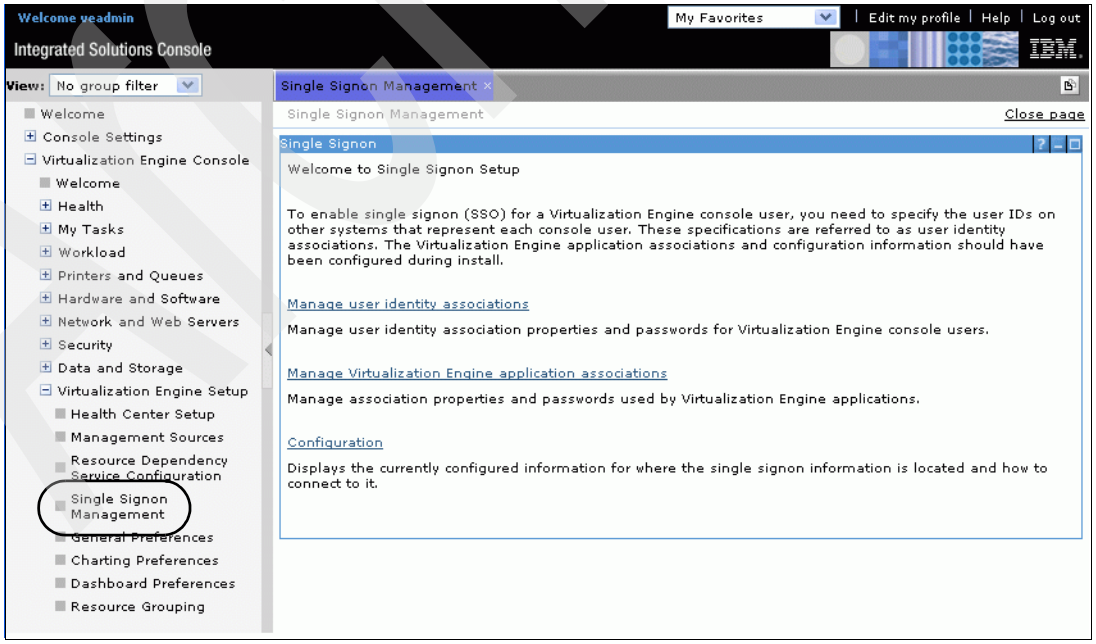


Figure 9-36 Single Sign-on Management

b. Click **Manage user identity associations**. This displays a list of Virtualization Engine console users mappings when requesting access to other systems, as shown in Figure 9-37.

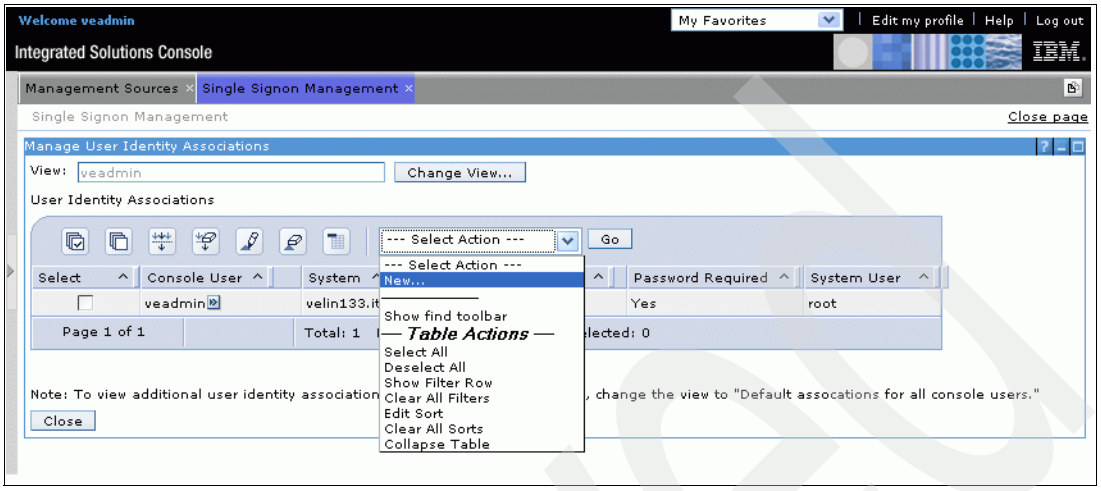


Figure 9-37 Manage user identity associations window

c. Select **New** from the drop-down list. This opens the New User Identity Association portlet, as shown in Figure 9-38.

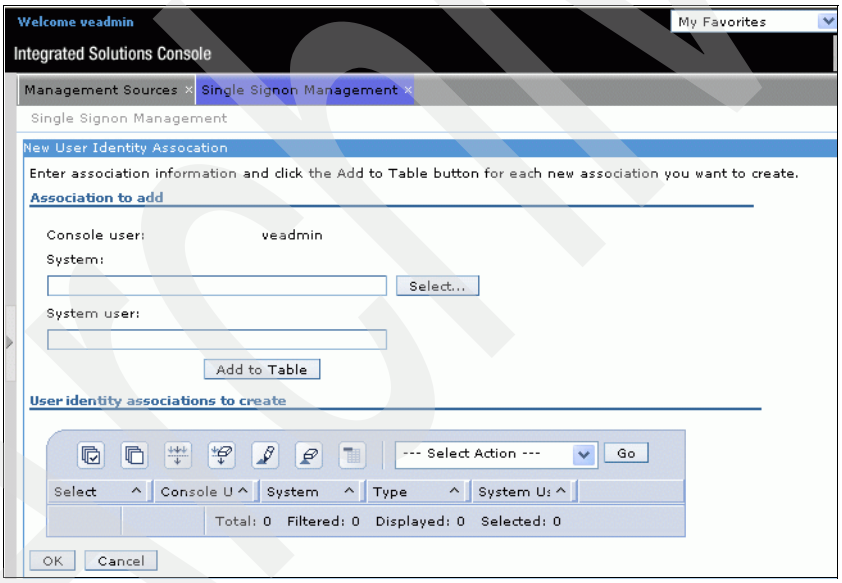


Figure 9-38 New User Identity Association

d. Click **Select** to see a list of systems for which you can configure credentials. Search for the new IBM Director Server system in the list. It is listed because you ran the commands listed in step 16 on page 327. Select the new IBM Director Server host name and click **OK**.

e. Once the System field is filled, insert a user name that is an IBM Director administrator in the System user field (Administrator in our case) and click **Add to Table**. Then click **OK**.

f. You have now added the new system in the list of User Identity Associations, as shown in Figure 9-39 on page 330. Select the check box next to this new line and click the

arrow next to the user name in the Console User column. Select **Set password**. Type the password corresponding to the user you specified for the new IBM Director Server twice, then click **OK**. This operation associates the Virtualization Engine user named veadmin (the one we are logged in with) with the Administrator user in windows.

Note: When the bridge and the IBM Director Server run on the same server, the user name association is not necessary, because the veadmin user is already mapped to the appropriate user.

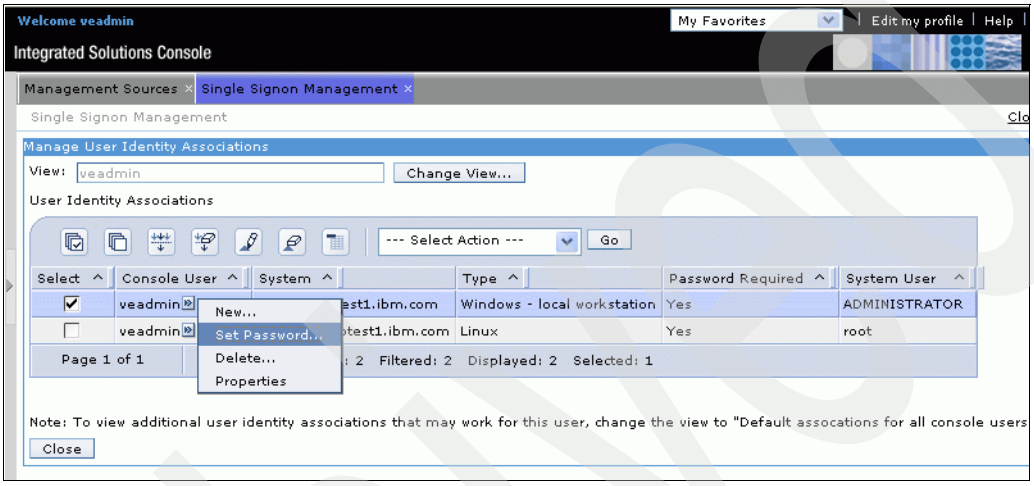


Figure 9-39 Set the password on the new host

9.3 Diagnosing a hardware error

One of the functions of the Virtualization Engine console is to receive a notification of a hardware failure and try to diagnose the error. In our example, we will cause a power supply failure on the BladeCenter chassis by physically removing it from the chassis itself.

When this happens, the BladeCenter Management module detects it and sends a notification to the IBM Director Server (field 1 in Figure 9-40 on page 331). This is polled by the bridge (field 2 in Figure 9-40 on page 331) that populates the Health Center portlet (field 3 in Figure 9-40 on page 331).

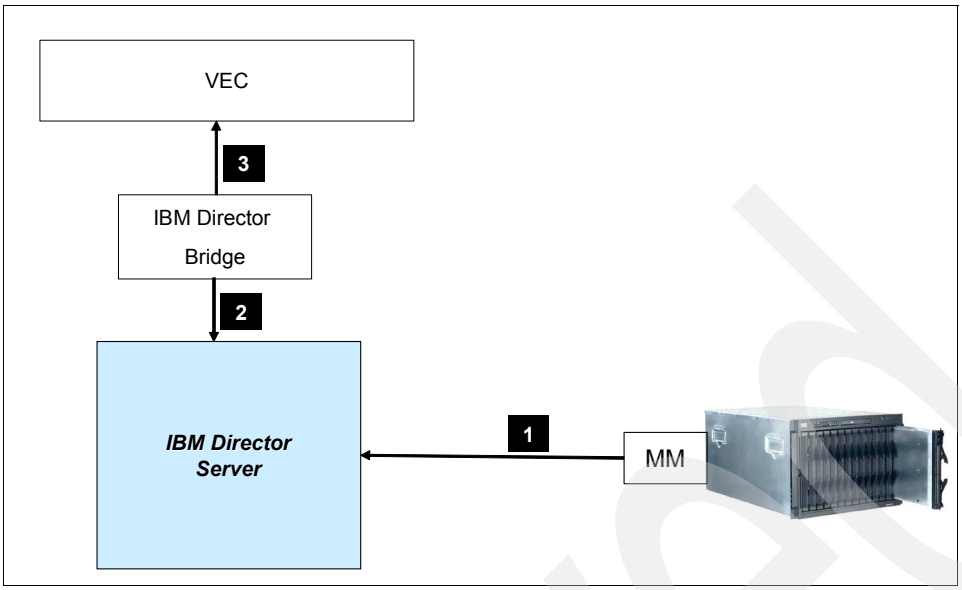


Figure 9-40 Event flow to the Virtualization Engine console

The following section describes how to configure the Health Center to get this notification. We assume that:

- The IBM Director Bridge has been properly configured. Refer to 9.2.1, “Configuring the management sources” on page 304 for more details.
- You already discovered your BladeCenter Chassis in IBM Director and requested access with the proper credentials. Refer to 9.2.3, “Discovering systems with IBM Director” on page 312 for more details.

Note: You need to insert the proper credentials to access a BladeCenter Chassis only if you changed the default user ID and password pair.

9.3.1 Configuring resources in the Health Center

9.2.1, “Configuring the management sources” on page 304 describes how to configure the management sources. In order to receive notifications about the managed systems, you should configure the Health Center to show the resources you are interested in.

Open a Web browser and log on to Virtualization Engine console. To do this, perform the following steps:

1. In the left navigation pane, expand **Virtualization Engine Console** → **Health** and select **Health Center**. This opens the Virtualization Engine console Health Center, as shown in Figure 9-41 on page 332. The Health Center summary is empty because we did not select to automatically add all groups and resources in the Health Center configuration wizard.

Tip: When you do not need the navigation bar on the left side, you can hide it by clicking the arrow highlighted in Figure 9-41. This expands the portlets area.

Archived

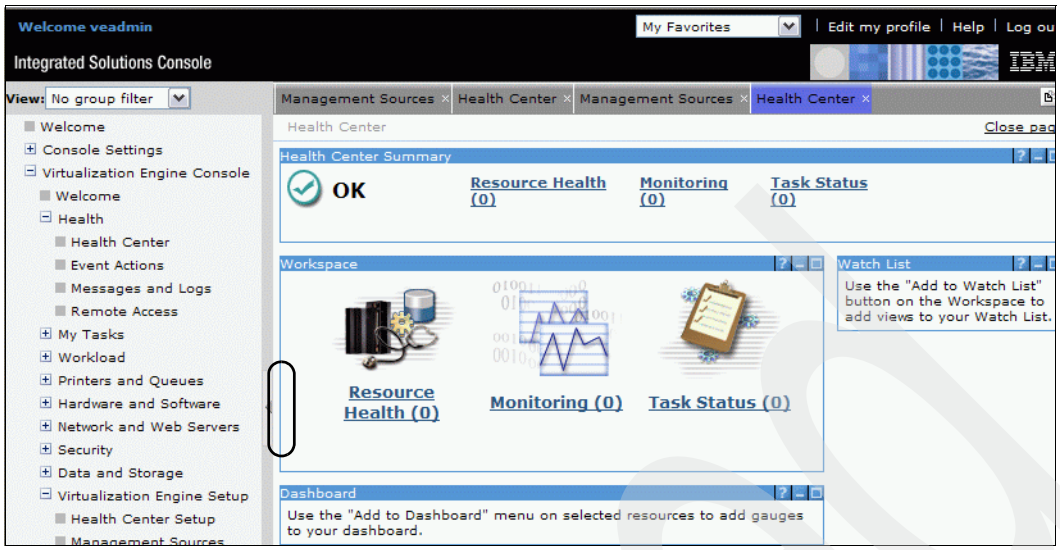


Figure 9-41 Health Center

2. Click **Resource Health**, either from the Health Center summary or from the Workspace portlet to see the list of resources in the workspace, as shown in Figure 9-42. In our case, we have a warning message, since no resources have been configured yet.

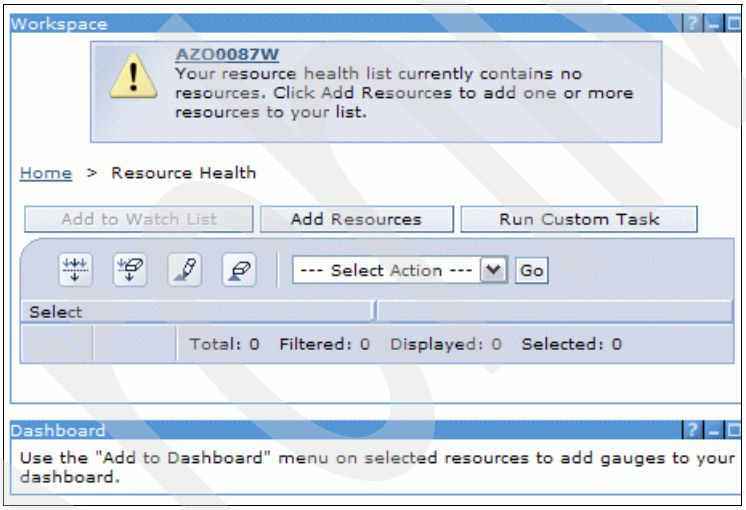


Figure 9-42 Resource health

3. Click the **Add Resources** tab. You are presented with a list of the groups and systems managed by the selected management source, IBM Director in our case, as shown in Figure 9-43 on page 333.

> **Note:** The display of the available resources normally takes some time, since the IBM Director Bridge needs to query the IBM Director Server for the available resources.

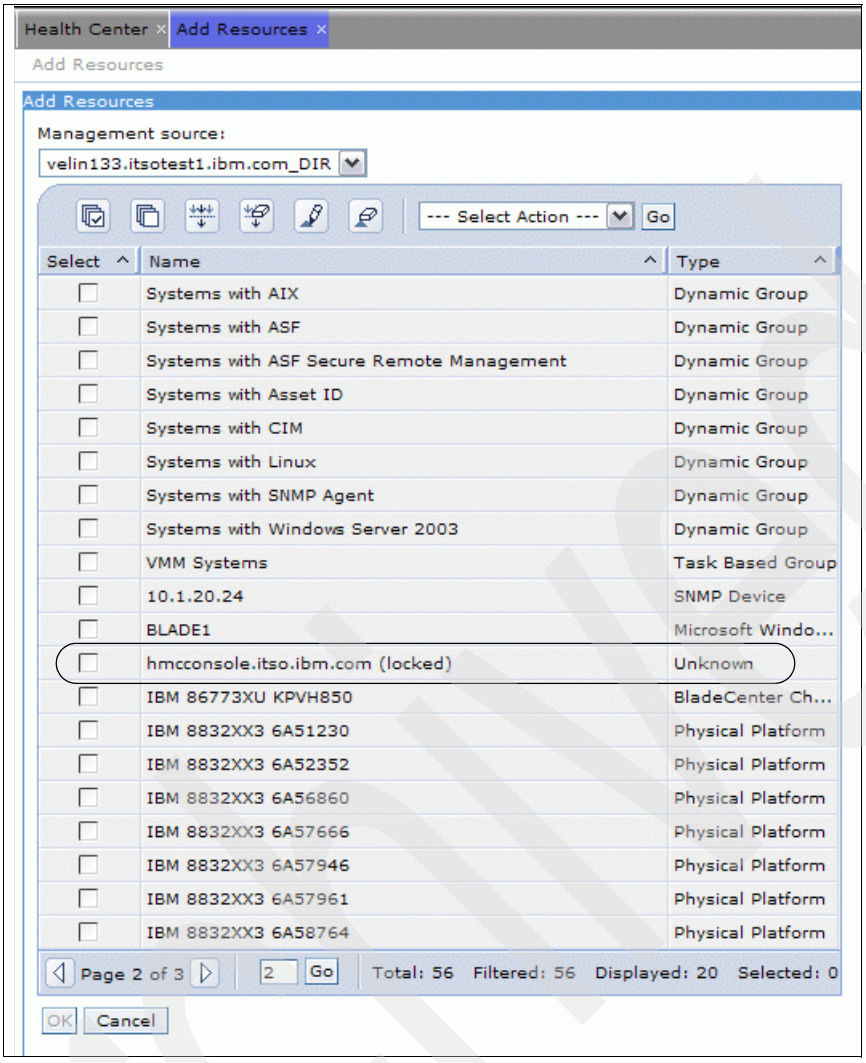


Figure 9-43 Resource selection window

If you have no objects in the IBM Director database, you will have only the default groups defined the time of installation in the IBM Director in this list. To discover systems with IBM Director, refer to 9.2.3, “Discovering systems with IBM Director” on page 312.

Note: The system circled in Figure 9-43 is locked. This means that your IBM Director Server has no valid credentials to manage it. Refer to 9.2.4, “Requesting access to systems in IBM Director” on page 315 for more information about requesting access to managed systems. The system type is unknown because it is not one of the native IBM Director object types (HMC management functions are added by an IBM Director plug-in). This is not related to the fact that the system is locked.

The choice of which groups or resources to add is crucial. We recommend you plan carefully what to do since if your Health Center is too busy, you will end up missing notifications. Given our list of managed objects and the scenarios we plan to run, we add the following groups to the list of managed resources:

– *Chassis and Chassis Members*, which contains the BladeCenter Chassis and all the components installed into it, such as blade servers and I/O modules.

– *HMC and HMC Members*, which contains the HMC and all the systems it is managing, such as the physical IBM System p server and configured LPAR. This group is needed for the scenario described in 9.7, “Monitoring an LPAR” on page 372.

– *Level 2: IBM Director Agents*, which contains all the systems with the full IBM Director Agent installed.

To add the systems to the Resource Health list, select the check box next to the groups of systems you want to add and click **OK**. We now have the groups listed in the Resource Health. The summary shows that we are currently monitoring three resources, as shown in Figure 9-44.

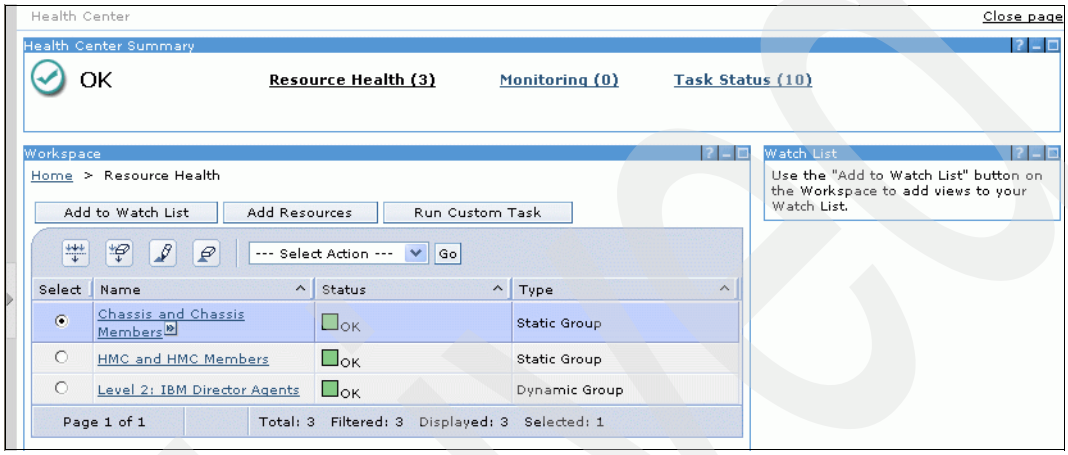


Figure 9-44 Health Center containing resources

Note: You can also create your own group definitions that contain systems that are part of the same business application. You can do so from the Virtualization Engine console launch in context.

Refer to *Implementing IBM Director 5.10*, SG24-6188 and *IBM Director Systems Management Guide* for more information about how to create groups in the IBM Director. *IBM Director Systems Management Guide* is available on the Web at:

`http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-61788`

9.3.2 Seeing the notification and acting on the target system

Now that you have added the resources to the Resource Health, there is nothing else you need to configure to have a hardware failure notification in the Health Center. Figure 9-45 on page 335 shows the Health Center after the power supply failure we generated in our test environment. The Health Center summary shows a red exclamation mark meaning that something is wrong. The red exclamation mark next to the Resource Health link shows that the problem is in one of the resources you are monitoring there. If you click **Resource Health** in the summary portlet, you will see the details in the workspace portlet.

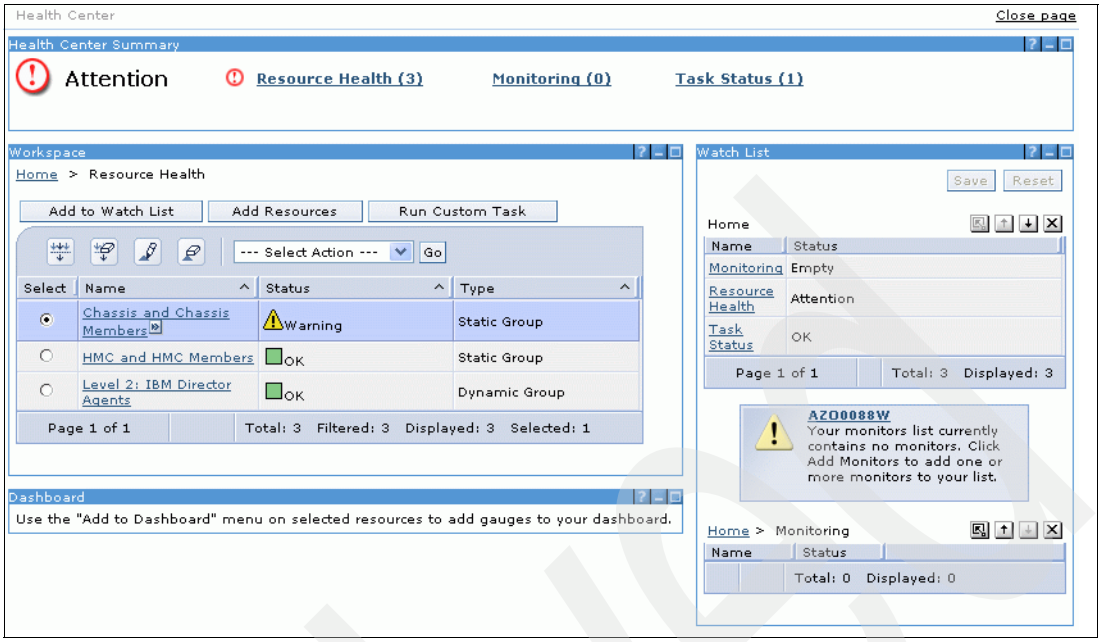


Figure 9-45 Health Center: indication of a hardware problem

From this view, you can get more details about what is happening to the systems. To do so, click the group name, **Chassis and Chassis Members,** in our example. Here, you see the details about the group, as shown in Figure 9-46.

Tip: When you navigate further to check resources, you can check at which level you are at the top of the workspace portlet, as shown in Figure 9-46.

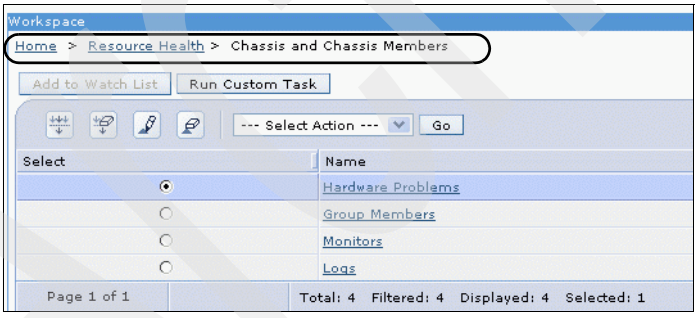


Figure 9-46 Group chassis and chassis members details

In this case, we want to see what is causing the warning notification, so we click **Hardware Problems**, which brings us one level down, as shown in Figure 9-47 on page 336.

Note: The Hardware Problems portlet corresponds to the Hardware Status task in the IBM Director. If the notification is not caused by a hardware problem, the Hardware Problems portlet is empty.

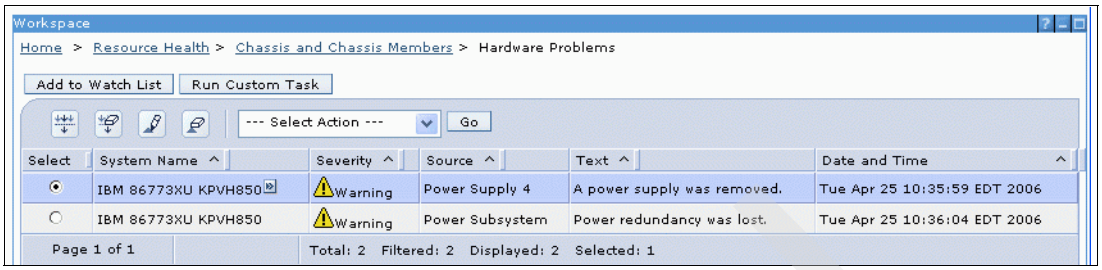


Figure 9-47 Hardware Problems details

The Hardware Problems portlet has more details about which system generated the alert. In our example, we see that the machine that generated the problem is the BladeCenter Chassis. Here, it is identified with the name given by IBM Director when the management module is automatically discovered, that is, the string IBM followed by the machine type and serial number. Refer to 9.2.5, “Renaming BladeCenter” on page 317 for options to name physical platforms differently.

We have two indications:

- The first one saying that the power supply has been removed.
- The second one related to the loss of redundancy on power due to the power supply failure.

Select one of the entries and click the arrow next to the system name. For more details about the event itself, click **Properties** from the menu, as shown in Figure 9-48.

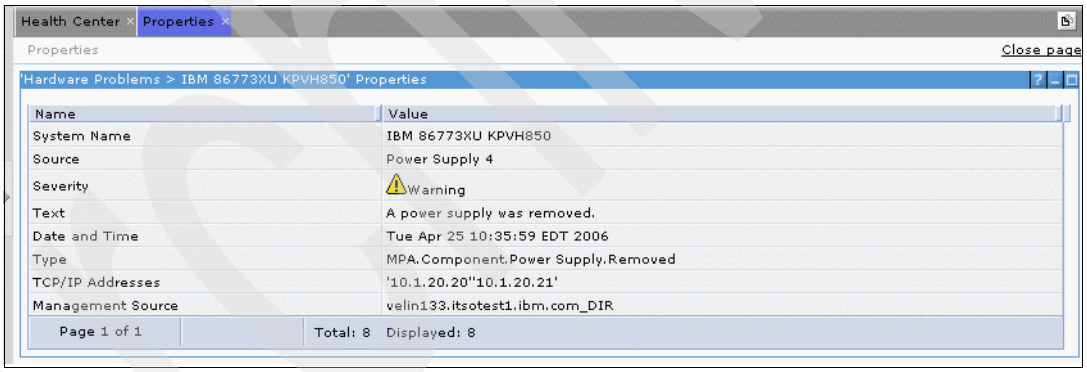


Figure 9-48 Error message details

If you close this page, you return to the Hardware Problem details window. Once you know that the problem is a hardware failure, you have two choices:

- Replace the failed part. This automatically resets the hardware status on the corresponding machine and group.
- Acknowledge the status so that the warning indication is deleted. To do so, from the Hardware Problem Details window, select the event entry that you want to clear, click the arrow next to the system name, and select **Status Event Action**. You have the choice to Ignore or to Clear the status, as shown in Figure 9-49 on page 337.

  If you choose to ignore the event, this means that the event is not available any more and the status is reset to normal. Future events received by the selected system, category, or component type will be logged; however, the hardware status will not be set to warning.

  Clear means that the status is reset to normal; however, future events will set the status to warning. In both cases, the error code is kept in the log.

Note: Clearing or Ignoring an event in Virtualization Engine console corresponds to the same action in the IBM Director hardware status task.

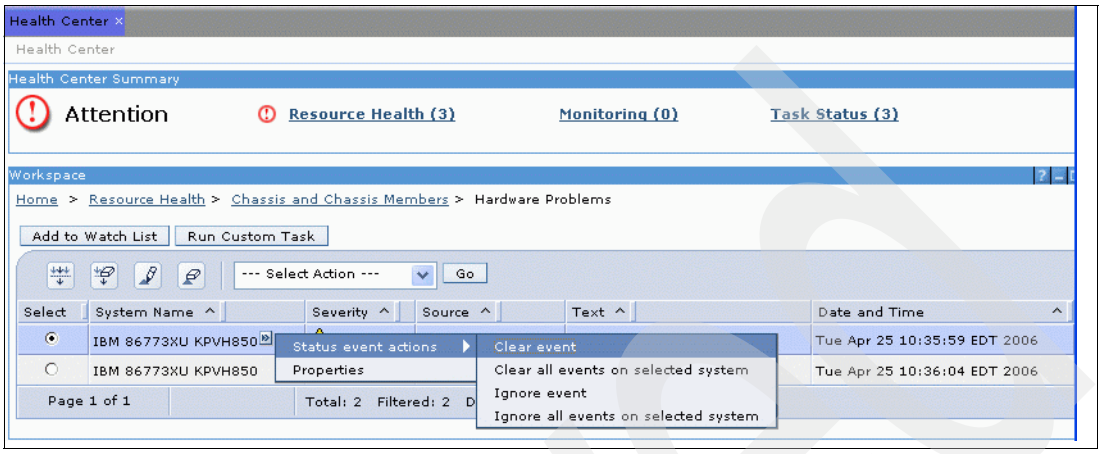


Figure 9-49 Clearing the status on a failed system

Once you clear all the entries in the Hardware Problems list, the Health Center Summary portlet in the top portion of the portal returns to normal, as shown in Figure 9-50.



Figure 9-50 Health Center summary back to normal

You can access the event properties even if the hardware problem has been cleared by clicking **Logs** from the group details window, as shown in Figure 9-46 on page 335. In the Logs portlet (shown in Figure 9-51) that appears, click **EventLog**, then select which logs you want to see, as shown in Figure 9-52 on page 338. The list of available filters is the same defined in the IBM Director Event Action Plan configuration. Click **All Events** to see the list af all the events pertaining to the object (the whole group in this case). Figure 9-53 on page 338 highlights the events generated by the hardware failure.

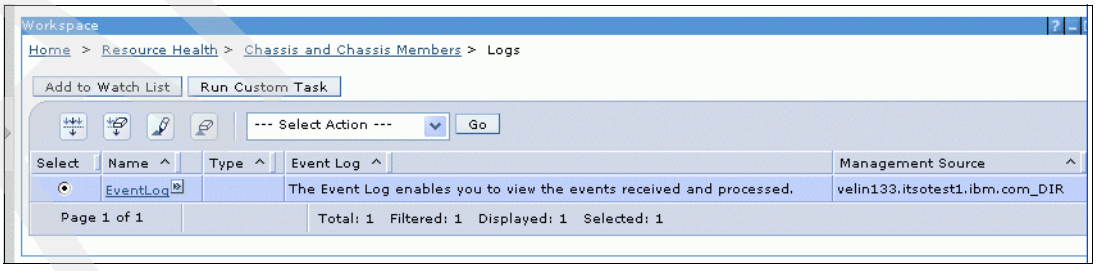


Figure 9-51 Group details: logs

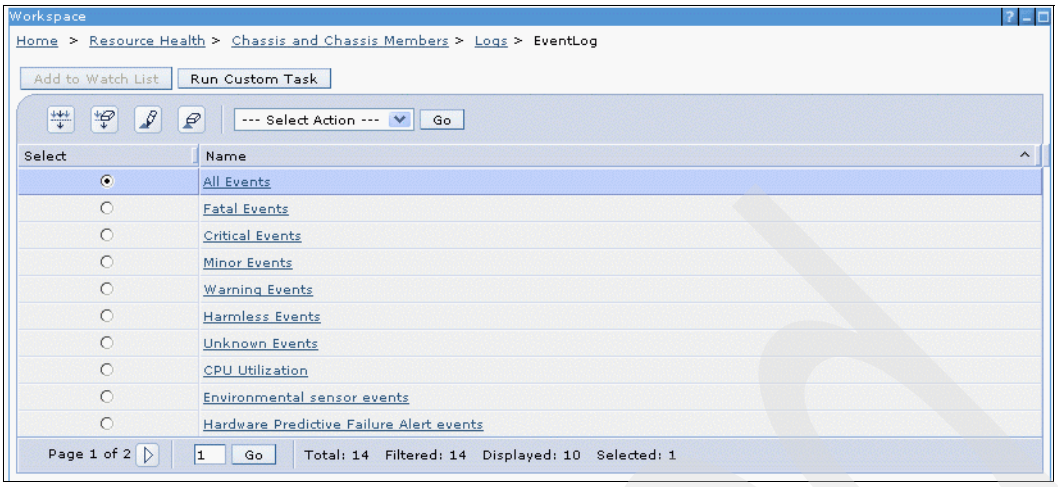


Figure 9-52 List of logs by filter type

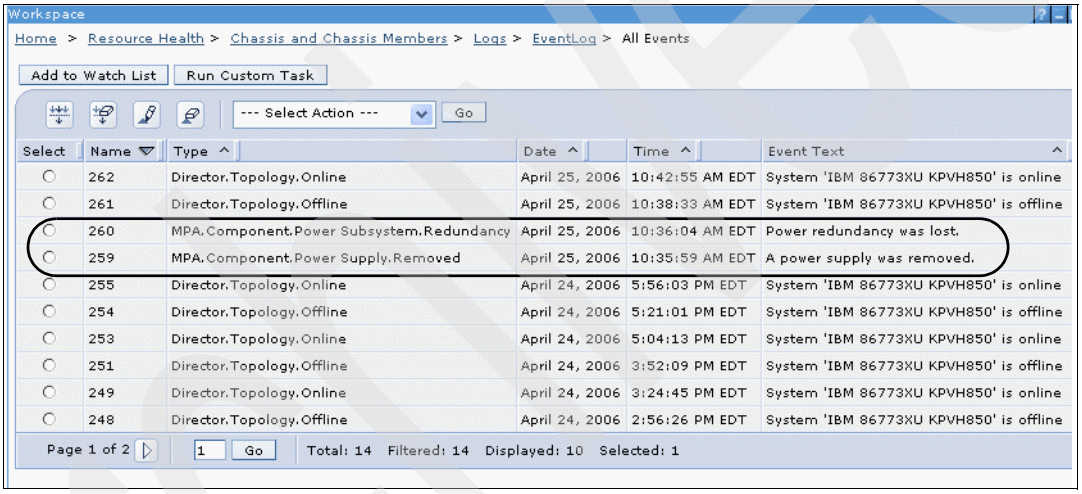


Figure 9-53 List of event log entries for the group

Building an alert on the BladeCenter unit

You can make the Indication LED blink on the BladeCenter unit that needs the hardware replacement to make it easy for the people in charge of hardware maintenance to identify. There is an easy way to do this from the Virtualization Engine console:

1. Open the Health Center and select **Resource Health** → **Chassis and Chassis Members** → **Group Members**.
2. Select the BladeCenter Chassis object, as shown in Figure 9-54 on page 339, and click the arrow next to the chassis name.
3. From the launch-in-context menu, select **System Identification** → **LED Strobe**. To turn it off, select **LED Off**.

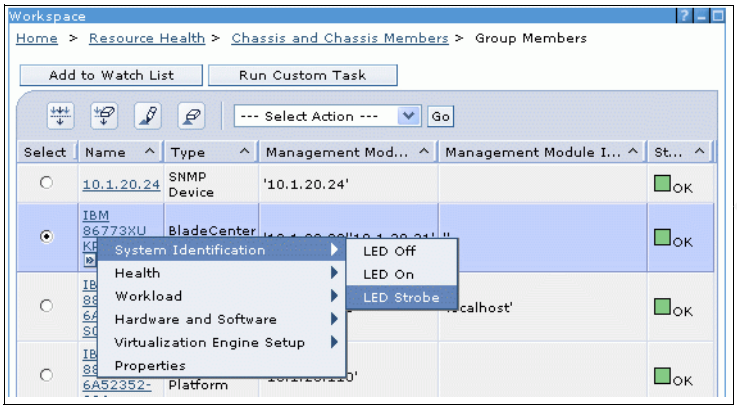


Figure 9-54 Launch-in-context menu on a BladeCenter

Checking the power status

If you want to check the power status of a blade server, you have the ability to do so from the BladeCenter management module Web interface. You can open it from the launch-in-context menu on the BladeCenter Chassis object. To check the power status of a blade server, perform the following steps:

1. From the Health Center home, select **Resource Health** → **Chassis and Chassis Members** → **Group Members**.
2. Select the BladeCenter Chassis object, click the arrow next to the name and select **Hardware and Software** → **Hardware Management** → **Launch BladeCenter Management Module**, as shown in Figure 9-55.



Figure 9-55 Launch the BladeCenter Web Interface from the Virtualization Engine console

This opens a Web browser window pointing to the BladeCenter management module IP address.

3. Insert the proper user ID and password and select the timeout desired.
4. To access the power control menu, expand **Blade Tasks**, then click **Power/Restart**, as shown in Figure 9-56 on page 340.

   Here you see the power status of each blade server. If you want to start or stop one, select the check box next to the target blade and click **Power On Blade** or **Power Off Blade** accordingly.

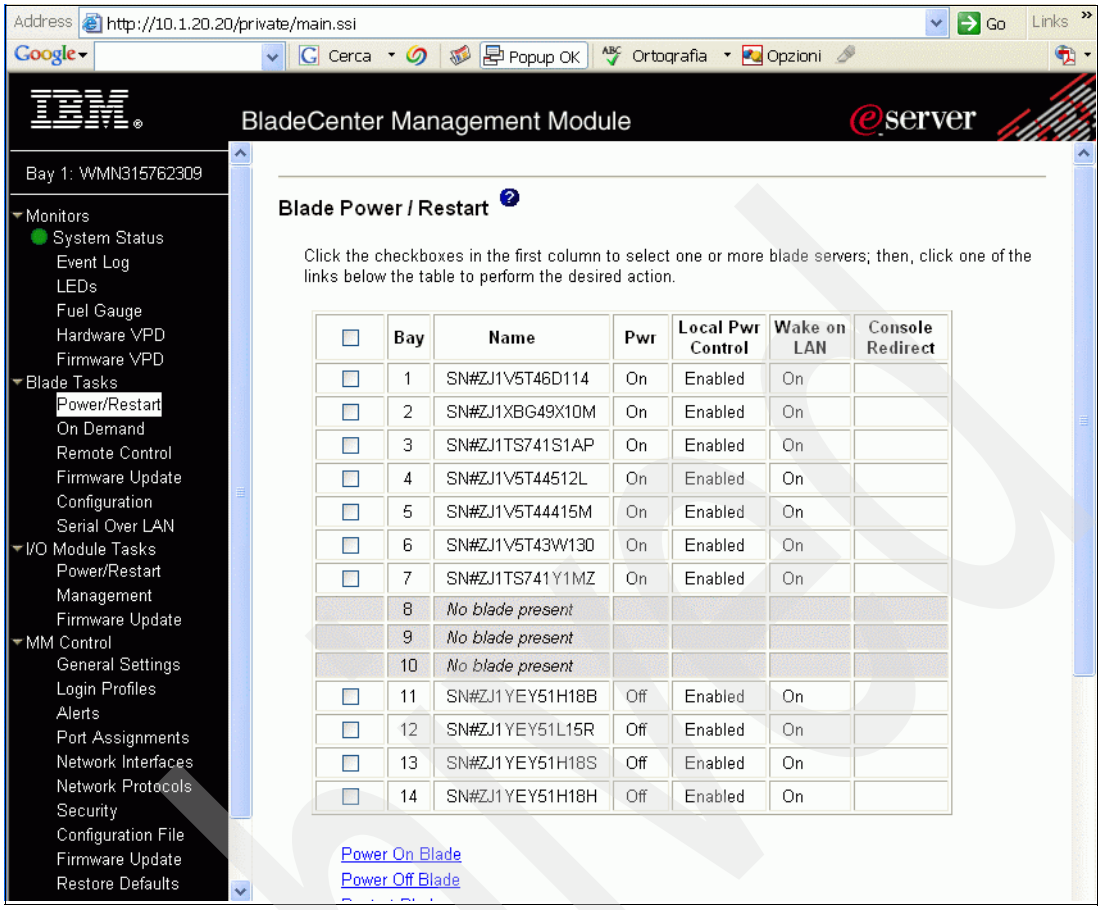


Figure 9-56 BladeCenter management module Web interface

9.4 Processing error

In this scenario, we want to monitor a critical process running on an IBM System p LPAR running AIX 5L V5.3. The process we chose for the example is the NFS daemon. The event flow, illustrated in Figure 9-57 on page 341, is as follows:

1. The IBM Director L2 Agent monitors the process (field 1 shown in Figure 9-57 on page 341).
2. If the process stops, the IBM Director Agent sends a notification to field 2, the IBM Director Server.
3. The notification is polled by field 3, the IBM Director Bridge.
4. The IBM Director Bridge updates field 4, the Virtualization Engine console Health Center.

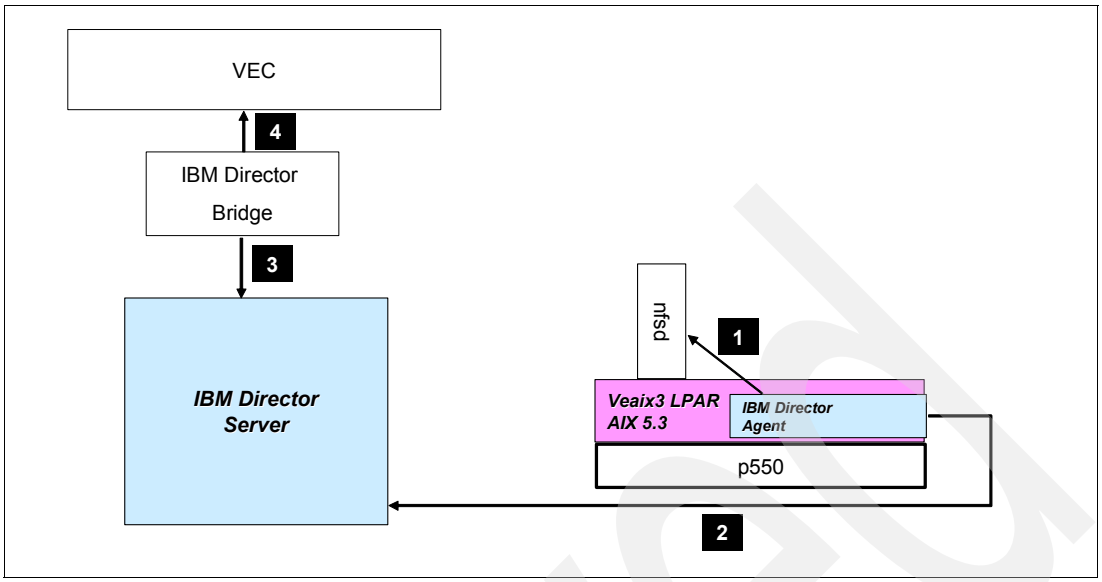


Figure 9-57 Event flow

Once we receive the event, we configure the IBM Director to restart the process.

Note: You can automate the process restart using the IBM Director Event Action Plans.

For this scenario, we assume that:

- The IBM Director Server and bridge are communicating properly.
- IBM Director L2 Agent is installed on the AIX 5L V5.3 LPAR.

9.4.1 Process Monitor configuration in IBM Director

As we said, the IBM Director agent is in charge of monitoring the process. This must be configured in the IBM Director using the process management task. You can open this from the Virtualization Engine console using either of the following methods:

- Using the left navigation pane

  To use this method, perform the following tasks:

  a. From the left navigation pane, expand **Virtualization Engine Console**, expand **Workload**, and then select **Workload Management**. This opens the Workload Management Tasks portlet, as shown in Figure 9-58 on page 342.

Archived

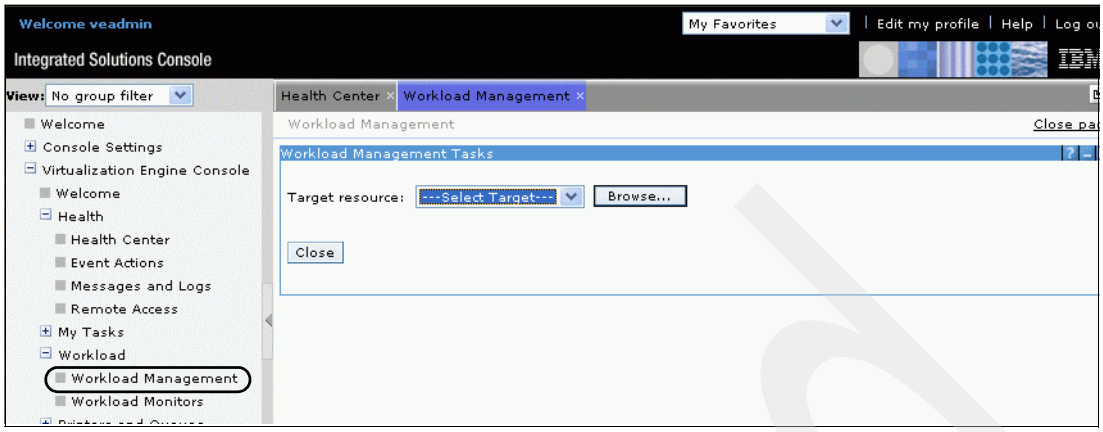


Figure 9-58 Workload Management Tasks portlet

b. Select the target system from the Target resource list.

   If it is not listed, then it means that none of these functions was ever launched against the system itself.

   To add it, click **Browse** and select your target system. Then click **OK**. We selected the veaix3 machine, which is the system with the IBM Director L2 Agent installed.

   Note: When you click **Browse**, the list of available target resources contains only the systems that IBM Director Server has registered with the proper credentials. The locked systems are not presented in the list, since no management action can be performed against them from the IBM Director Console.

c. Once you have selected the target system, a list of the available tasks appears, as shown in Figure 9-59. Click **Process Management**.

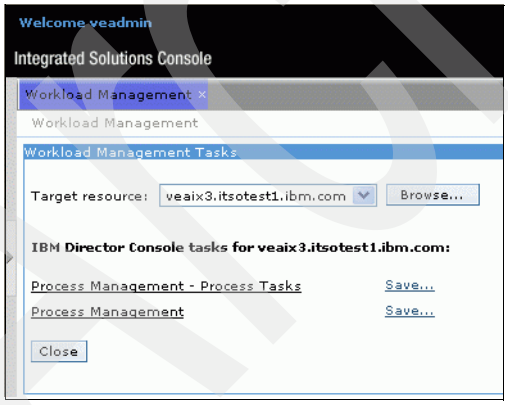


Figure 9-59 Workload management tasks selection

This launches the IBM Director Console and opens the process management task against the selected system, as shown in Figure 9-60 on page 343.

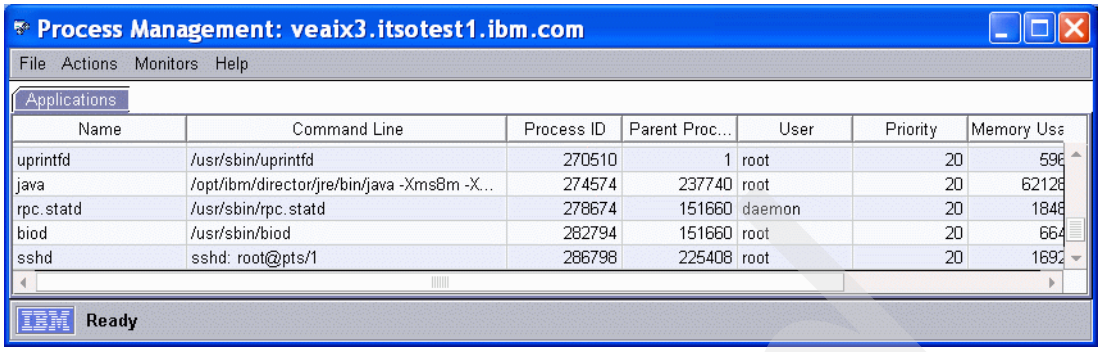


Figure 9-60 IBM Director Process Management task

- Using the launch-in-context

  To use this method, perform the following steps:

  a. From the Health Center portlet, select **Resource Health**, then click **Level 2: IBM Director Agents**. Next, click **Group Members**. You see the list of systems contained in your group, as shown in Figure 9-61.

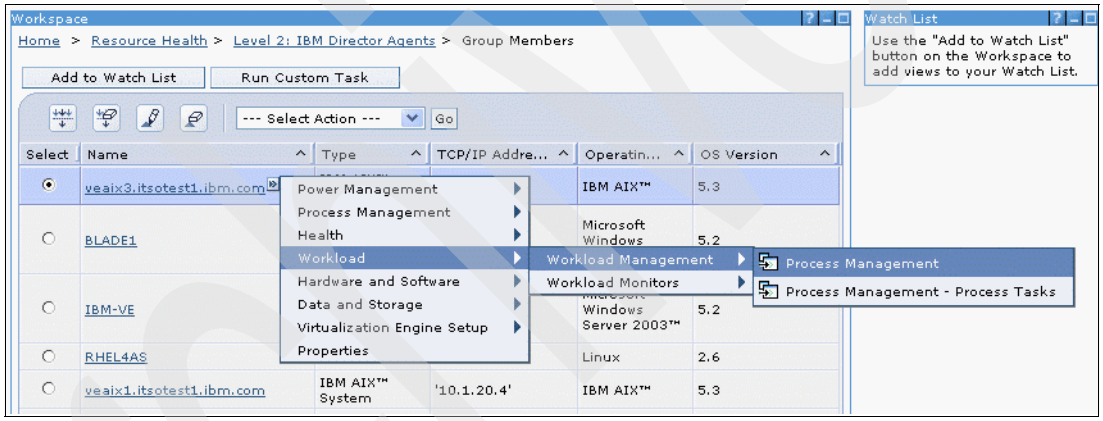


Figure 9-61 Health Center workspace portlet showing the systems in the Level 2: IBM Director Agents group

Select the target system (veaix3.itsotest1.ibm.com in our example), and click the arrow next to the system name. This lists the tasks available on that specific system, as shown in Figure 9-61.

b. Select **Workload** → **Workload Management** → **Process Management**. This instructs the Virtualization Engine console to launch the IBM Director Console. You see the window in Figure 9-62 that notifies you that an external application is being launched.



Figure 9-62 Indication of an external application launch from the Virtualization Engine console

The IBM Director Console process management task in the context of the selected system (Figure 9-60) appears.

Note: For the launch-in-context method only, your browser must have software to block pop-up ad windows disabled for the Virtualization Engine console Web site. The reason for this is that a new browser window is launched covertly. This is not the case when you launch the same task from the navigation window (described in page 341).

c. In the Process Management task window for the selected system, scroll to find the nfsd process. Right-click it and select **Add to Monitors** (as shown in Figure 9-63). This adds the process to the list of those you want to monitor.

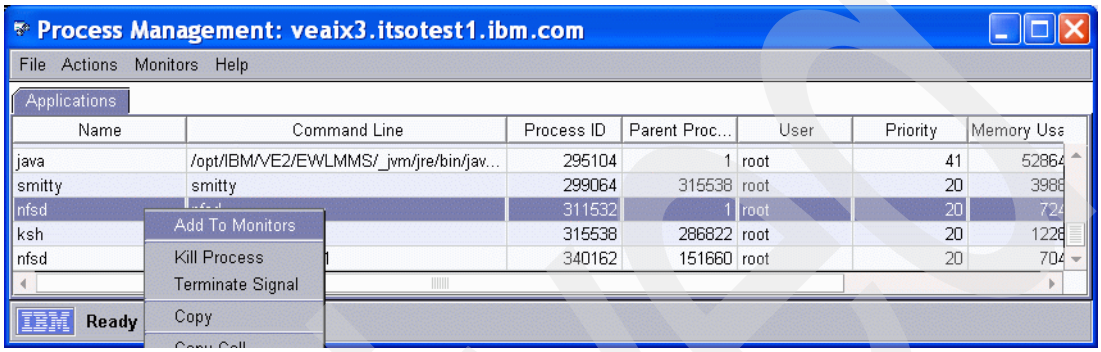


Figure 9-63 Adding the nfsd process monitor

d. Select **Monitors** from the top menu, and click **Process Monitors** (as shown in Figure 9-64).

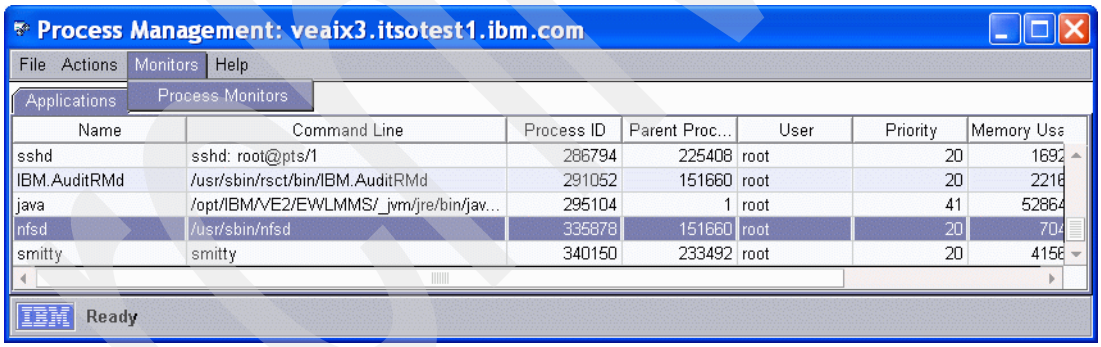


Figure 9-64 Open the Process Monitors window

The Process Monitors window, shown in Figure 9-65, appears. In the line corresponding to nfsd, you must select which conditions you want to be triggered, in our example, we select the **Stop** and **Fail** check boxes.

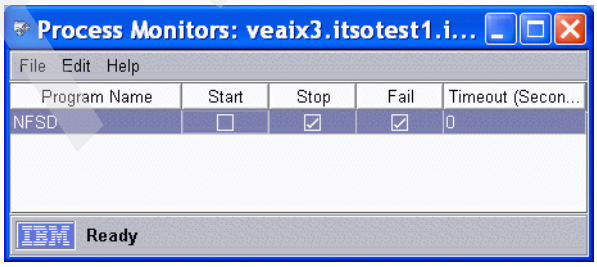


Figure 9-65 Monitored Processes for the target host

Note: If you select the check box next to Start, you will get a warning notification when the process starts. If the normal operations correspond to a started status for the process, then we recommend you do *not* select Start.

e. Select **File** → **Save** to save your settings, as shown in Figure 9-66.

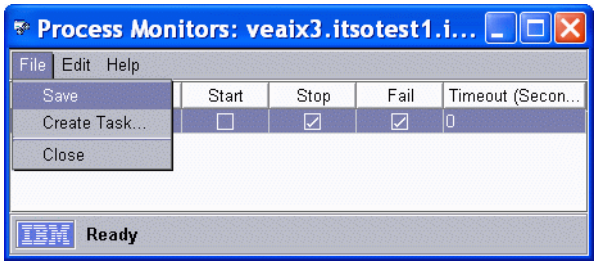


Figure 9-66 Saving the Process Monitor settings

Note: You can also launch the Process Monitors window for a system directly by selecting **Workload** → **Workload Monitors** → **Process Management - Process Monitors** from the object launch-in-context menu, as shown in Figure 9-67.

However, when you create monitors directly from the Process Monitors window, you must create the process monitor row. This means that you must know the exact name of the process. We recommend you follow the steps we described to avoid errors in the process name.

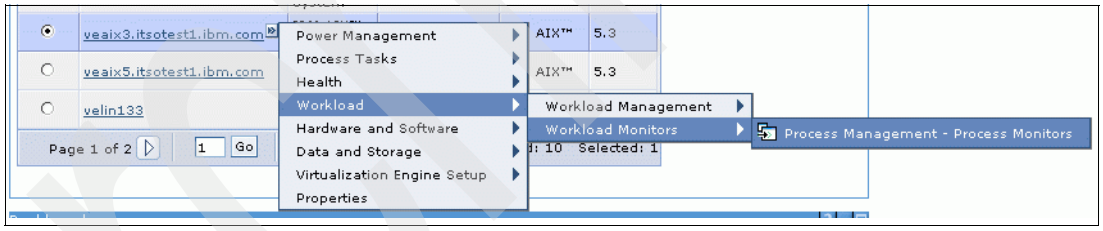


Figure 9-67 Launching the Process Monitors window directly from Virtualization Engine console

9.4.2 Adding the Process Monitor to the Health Center

Until now, we have configured the IBM Director agent to monitor the selected process so that the IBM Director Server receives a notification of the event. To get the corresponding notification in Virtualization Engine console, you must perform some additional configurations.

The first thing to do is to add the process monitor we just configured in IBM Director to the list of monitors in the Health Center. To do this, open the Health Center and perform the following steps:

1. Select **Monitoring**. In the Monitoring portlet, click the **Add Monitors** tab, as shown in Figure 9-68.

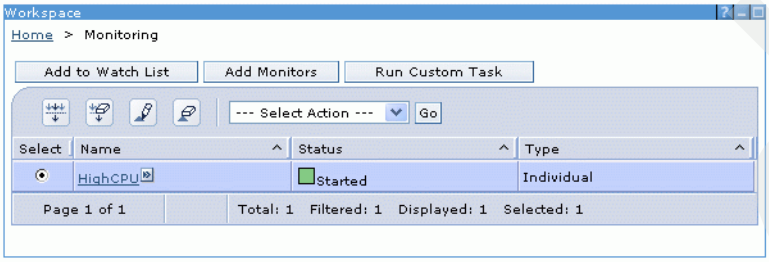


Figure 9-68 Health Center Monitoring portlet

This opens the **Add Monitors** portlet, as shown in Figure 9-69. All the currently configured monitors in IBM Director are listed here by name.

2. Find the name of the process monitor just created, in our case PMON: NFSFD Stop, and select the check box next to its name. Then click **OK**. This adds the selected monitor to the list of currently monitored parameters, as shown in Figure 9-70 on page 347.

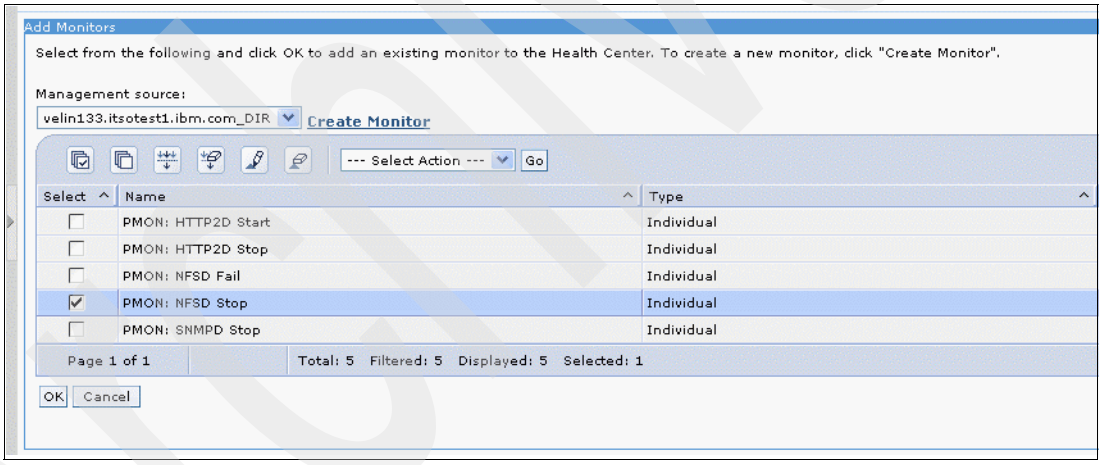


Figure 9-69 Add Monitors portlet: list of available monitors

Note: The list of monitors does not include the system name to which the monitor applies. If you have the same process monitored on several systems, you will have more entries with the same description. To find the one pertaining to a particular system, click it to see the details, as shown in Figure 9-71.

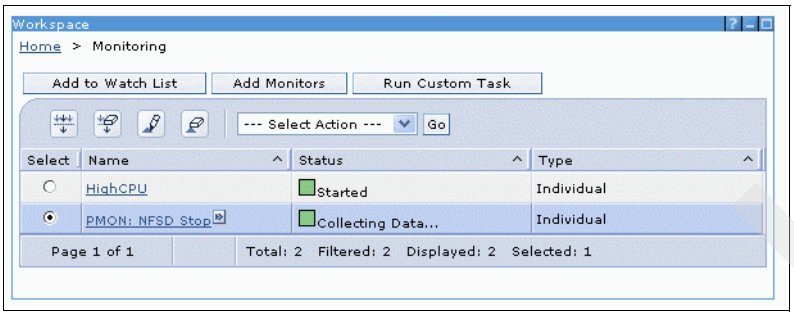


Figure 9-70 Workspace portlet: new monitor added to the current list

9.4.3 Adding the monitor to the watch list

One of the functions of the Health Center is to provide a bird's eye view of the resource status of the complete environment. One way of doing this is to add resources to the watch list. The aim of the watch list is to contain a list of critical resources for which the status is immediately available. To do this, perform the following steps:

1. Open the Health Center and click **Monitoring**. You see the list of currently configured monitors, as shown in Figure 9-70. If you select the resource you want to add to the watch list from the list of monitored resources, you have to add the *status* of the resource monitor collection, which is not what we are interested in. To have the *value* of the resource monitored on the selected system, you should go down one level, by clicking **PMON:NFSD Stop**. This opens the monitor properties, as shown in Figure 9-71.

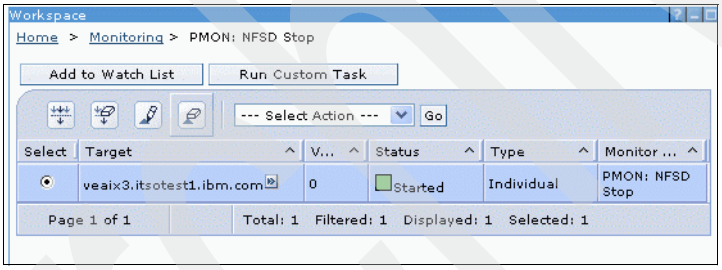


Figure 9-71 PSMON:NFSD Stop monitor properties

2. Now that you have the value of the monitor, click **Add to Watch List**. This way, the Health Center is composed of the portlets shown in Figure 9-72 on page 348. The Watch List shows clearly that the nfs daemon is stopped. This status is reflected in the Health Center Summary and in the Workspace portlets with the red exclamation mark. However, you do not need to drill down to the resources to have the details of the problem reported.

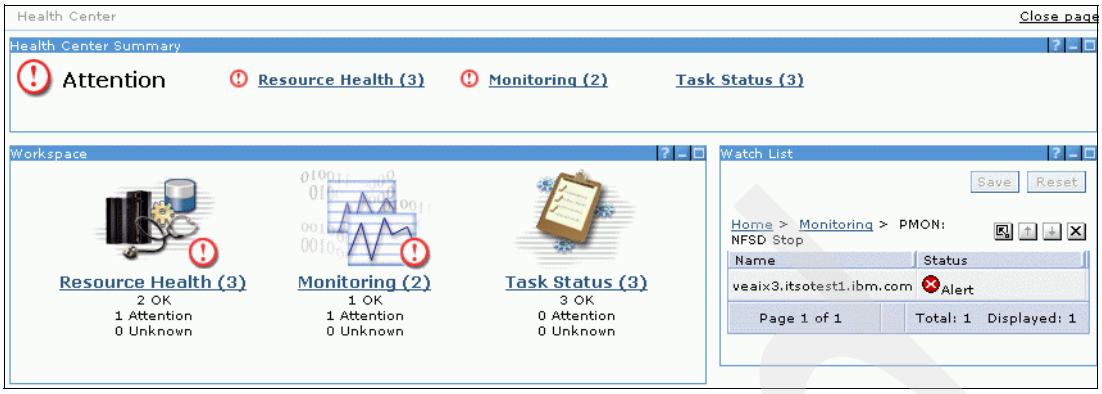


Figure 9-72 Health Center with watch list configured

9.4.4 Restarting the nfsd process

Once you get a notification that the nfsd process is stopped, you should restart the process.

To do this, open the Health Center and do the following steps:

1. Select **Resource Health** → **Level2: IBM Director Agents** → **Group Members**.
2. Select the **veaix3.itsotest1.ibm.com** host and click the arrow next to the name to display the launch-in-context menu. Select **Workload** → **Workload Management** → **Process Management - Process Tasks**, as shown in Figure 9-73.

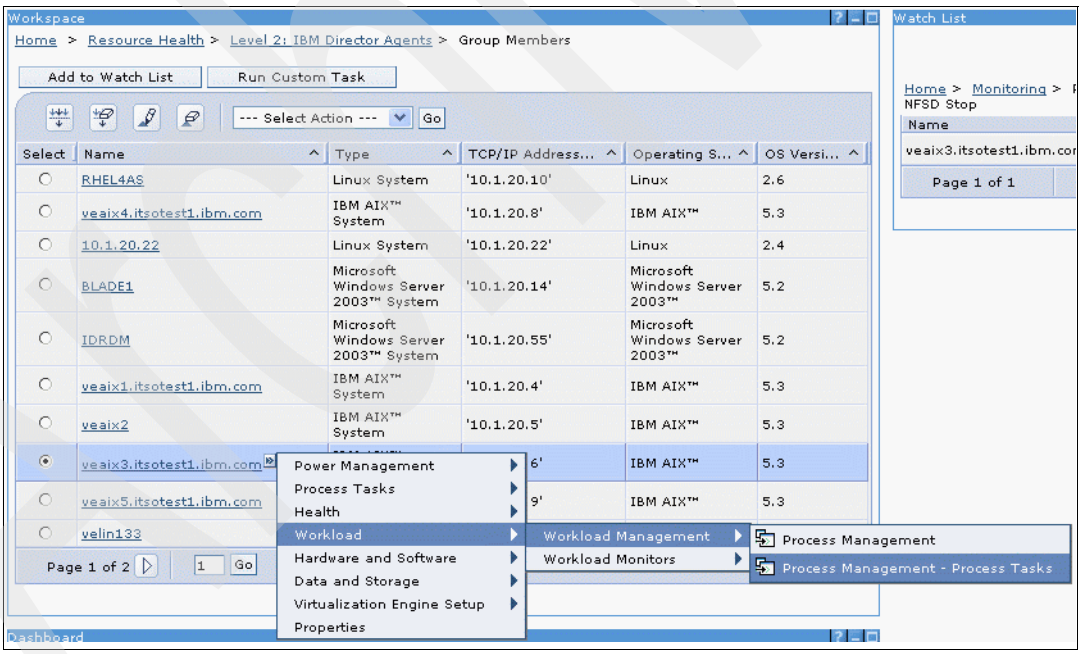


Figure 9-73 Launch the Process Tasks from Virtualization Engine console

This opens the IBM Director Process Task window, as shown in Figure 9-74 on page 349.

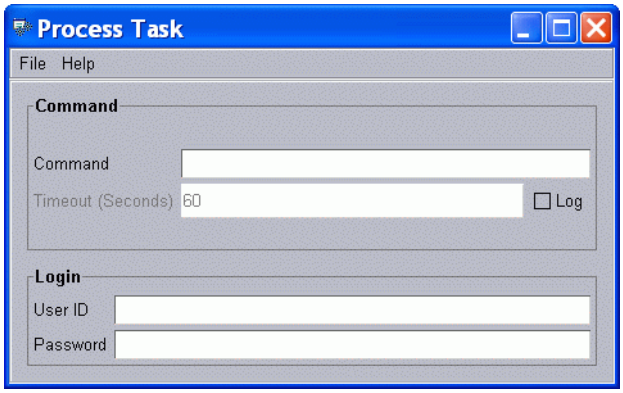


Figure 9-74 IBM Director Process Task window

Note: The Process Task window is not launched in the context of the target system since this is a general task. You can launch it from any system launch-in-context menu.

3. In the Command field, insert the name of the command you want to execute. In our case we run the **`startsrc -s nfsd`** command that starts the nfs daemon on a system running AIX.

   The login field overrides the default user ID, which is the one that launches the IBM Director agent on the target machine (root, in our case). Therefore, leave this field blank.

4. Select **File** → **Save As**, as shown in Figure 9-75.

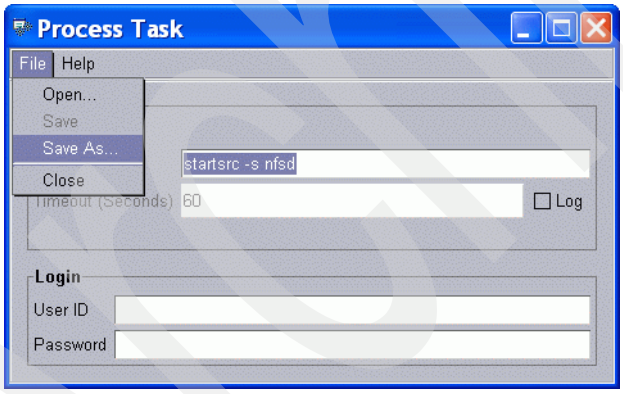


Figure 9-75 Save the Process Task

The Process Task: Save As window appears, asking you for a name for the new process task, as shown in Figure 9-76 on page 350.

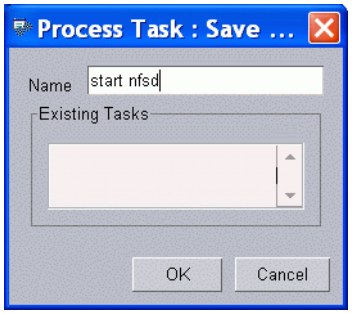


Figure 9-76 Process Tasks: Save As window

5. Type start nfsd in the Name field. This saves the tasks so that you can execute the specified command on a system without typing it or logging on to the system itself.
6. To restart the nfs daemon on our target machine, from the launch-in-context menu, select **Process Tasks → start nfsd**, which you just created, as shown in Figure 9-77. This action restarts the nfsd process and clears the error indication and the status in the watch list.

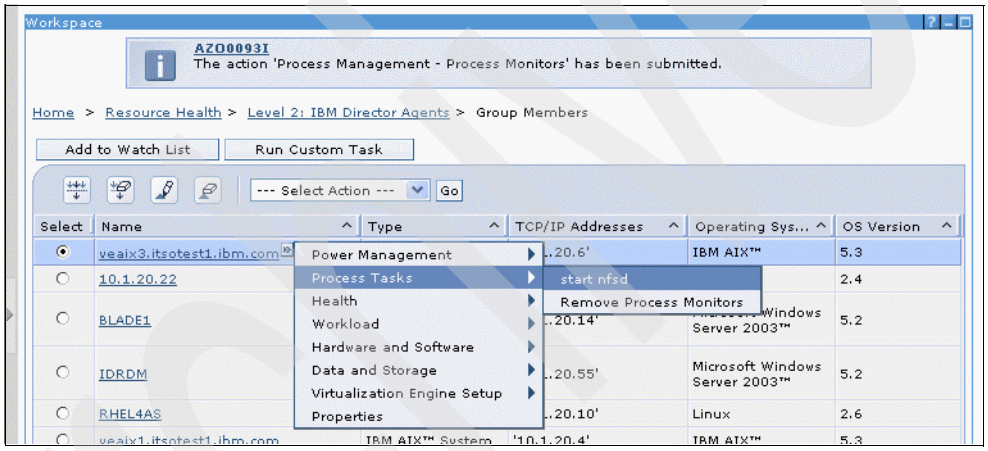


Figure 9-77 Launch the start nfsd process task from Virtualization Engine console

Note: The process restart can be automated using IBM Director Event Action Plans (EAPs). We give an example of how to use EAPs in 9.6, “Managing a VMWare Virtual Machine” on page 358. See *Implementing IBM Director 5.10*, SG24-6188 for more details.

9.5 Monitoring the system CPU usage

In this scenario, we want to monitor the CPU usage on one of the IBM System p LPARs running AIX 5L V5.3. To do this, we create a Resource Threshold in IBM Director using the Virtualization Engine console and add it to the Monitoring tab. The prerequisites and the event flow for this scenario are the same as those described in 9.4, “Processing error” on page 340. The main difference is that we must use a different task in IBM Director to configure the CPU monitor.

9.5.1 Configuring IBM Director to monitor the CPU

To configure the IBM Director to monitor the CPU, perform the following steps:

1. From Health Center, click **Monitoring**, then click **Add Monitors**. This opens the Add Monitors portlet, as shown in Figure 9-78.

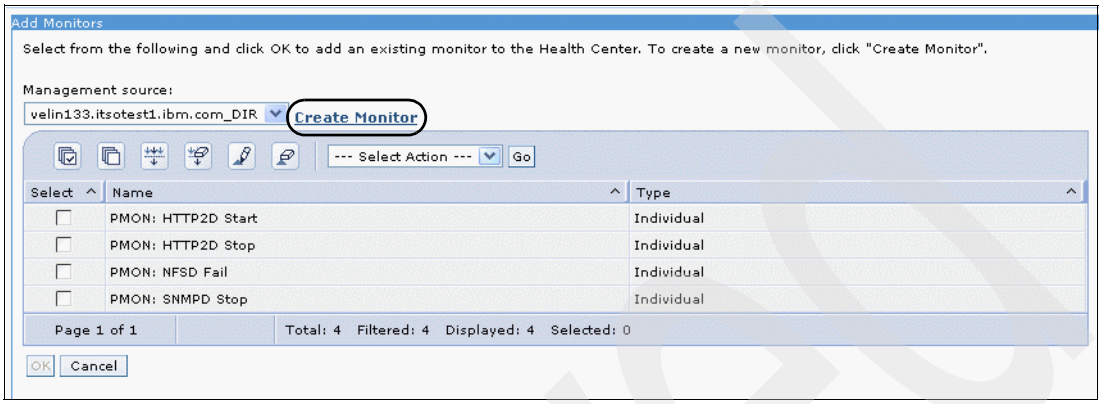


Figure 9-78 Add monitors portlet

2. Click **Create Monitor**. This opens the Create Monitor Tasks portlet as shown in Figure 9-79. Select the **Target resource** from the list, or click **Browse** if it is not selected. Once you have selected the correct target, click **Resource Monitors**.

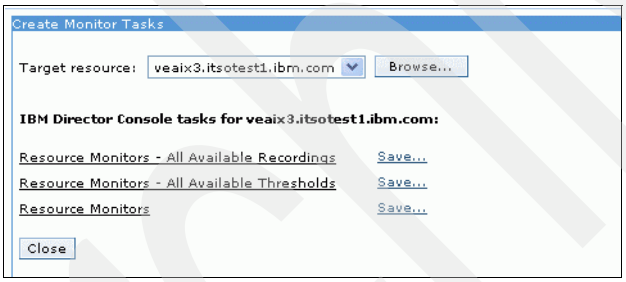


Figure 9-79 Create Monitor Tasks portlet

This launches the IBM Director Console Resource Monitor task against the selected system, as shown in Figure 9-80.

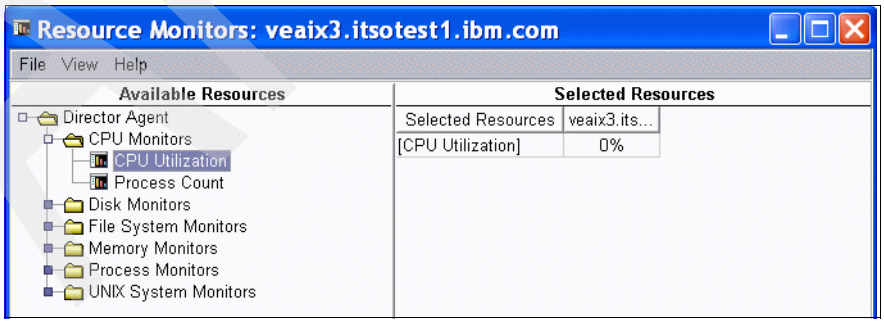


Figure 9-80 CPU utilization selection in the Resource Monitor task

3. Select **Director Agent** → **CPU Monitors**, and double-click **CPU Utilization**. This adds the CPU utilization to the Selected Resources pane. Right-click the Resource Value, as shown in Figure 9-81 on page 352, and select **Individual Threshold...**.

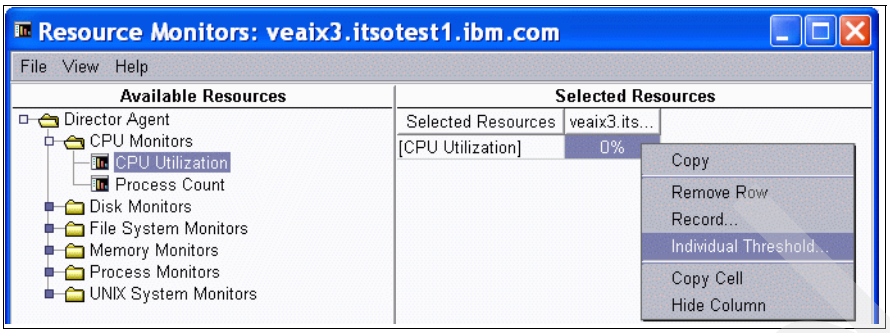


Figure 9-81 Set individual threshold on a resource

This opens the System Threshold veaix3.itsotest1.ibm.com, as shown in Figure 9-82.

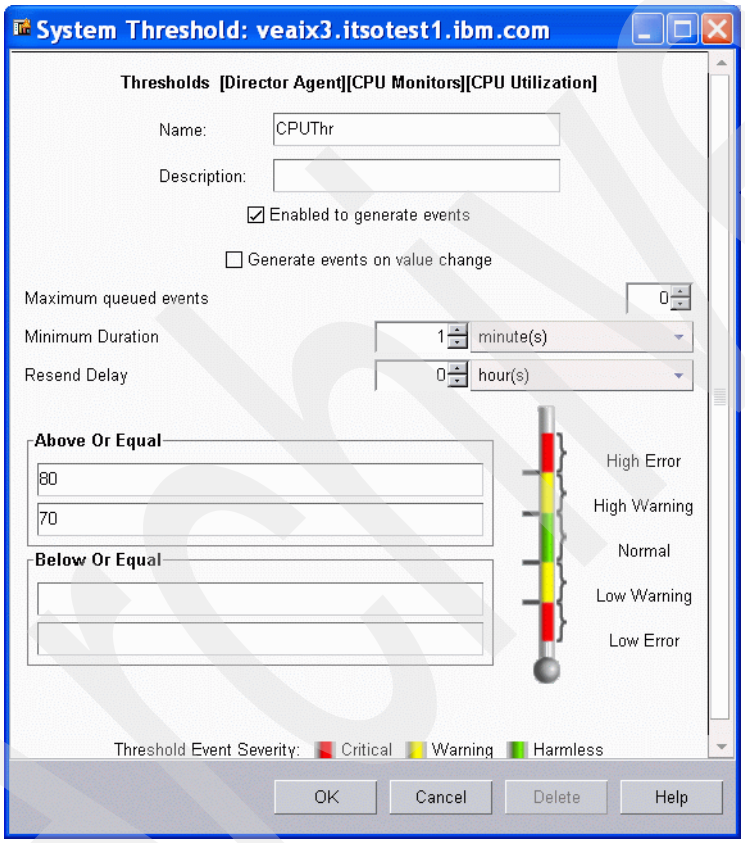


Figure 9-82 System Threshold configuration window

This is where you configure the notification you want to receive in the following steps:

a. First, you have to specify a name for the threshold. We chose CPUThr. Then, specify an optional description. Check the box **Enabled to generate events**. This means that when the monitored parameters meets the conditions specified below, it generates an event in IBM Director.

   Do *not* check the Generate events on value change. If you do so, the IBM Director generates an event with a severity level of warning when the monitor returns to the normal state.

b. The Maximum queued events field specifies how many events are held by the IBM Director Agent (that generated the event) should the IBM Director Server be unavailable for delivery. We left this field at 0.

c. In the Minimum Duration field, we set the value to one minute. This is the amount of time the condition must persist prior to generating the alert.

d. The Resend Delay field indicates for how long the condition generating the alert must persist prior to sending another event to the IBM Director Server. Since we only wanted to receive one event, we left this field at 0.

e. Finally, specify the values generating the threshold. In our example, we wanted to generate an error when the CPU utilization was above 80%. Hence, we wrote 80 in the High Error field. We wanted to generate a warning if the CPU utilization was above 70%. Therefore, we wrote 70 in the High Warning field. We did not want to generate errors or warnings if the CPU utilization is below any value, so we left the Low Warning and Low Error fields blank.

   You can find more information about the meaning of each field by clicking **Help**.

   Click **OK** to save your settings.

4. You are back in the Resource Monitor window, as shown in Figure 9-83. Since a threshold is set on the selected resource, you see the icon next to the detected resource value.

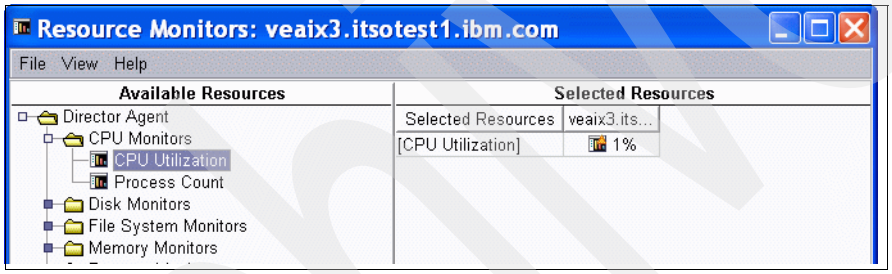


Figure 9-83 Resource Monitors with threshold set

The CPU threshold is set and working. Close the IBM Director Console and return to the Add Monitors portlet.

Note: The CPU utilization monitor reports the *logical CPUs* utilization. This means that if you have two logical CPUs, one of which is used at 100% and the second one at 0%, the CPU Utilization value will show 50%.

9.5.2 Adding the resource monitor to the Health Center

Now that you have configured the IBM Director to monitor the CPU status, you should add it to the resources monitored by the Health Center. To do this, open the Health Center and perform the following steps:

1. Go to the **Add Monitors** portlet, as shown in Figure 9-84 on page 354. Click **Monitoring**, and then click **Add Monitors**.

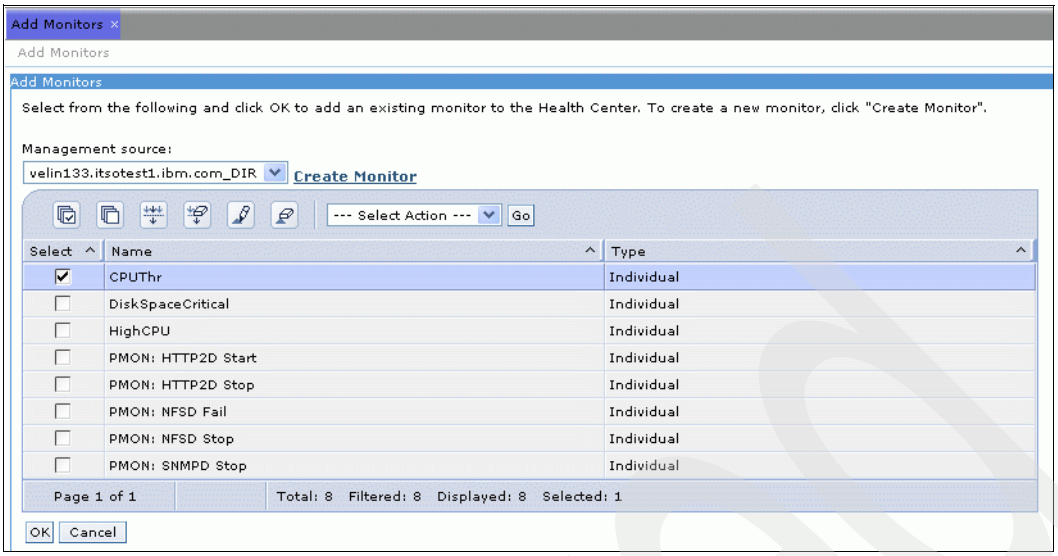


Figure 9-84 Add monitors portlet

Note: If you perform the steps in the order we present, when you close the IBM Director Console, the Add Monitors portlet will still be open in the Virtualization Engine console. In this portlet, there is no Refresh button. Therefore, you should close the tab (encircled in Figure 9-85) in the Integrated Solution Console, and re-open the Add Monitors portlet to see the threshold just created.

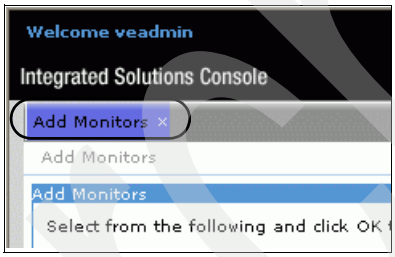


Figure 9-85 Closing the Add Monitors portlet

2. Now we have just created the threshold listed in the available monitors, as shown in Figure 9-85. Select the check box next to it and click **OK**. This adds the selected CPU monitor to the monitoring list, as shown in Figure 9-86.

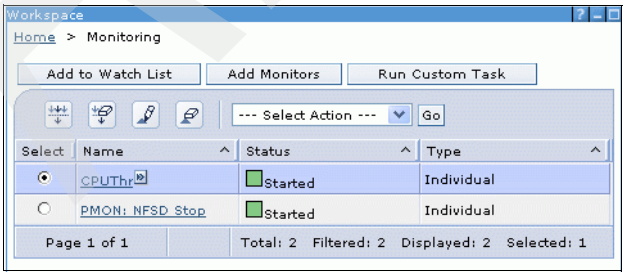


Figure 9-86 Workspace Portlet showing the newly added monitor

3. From this view, you can see what the status of the monitor is. It is “Started” in our case. To see the monitor value, you must drill down one level by clicking **CPUThr.** This shows the CPUThr monitor details (Figure 9-87).

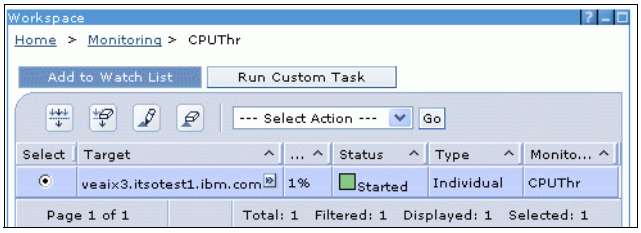


Figure 9-87 Threshold details

Note: IBM Director provides a method to configure thresholds on one system and replicate all its settings to other systems by using the *threshold plan task*. This creates a task that you can drag onto systems or groups to which you want to apply the same thresholds. For more details about threshold plans, refer to the *IBM Director Systems Management Guide*.

At the time of the writing of this IBM Redbook, the list of monitors in Virtualization Engine console is populated based on the content of the IBM Director *All Available Thresholds* task launched against all the systems with an L2 IBM Director agent installed. Thresholds applied to a system as threshold plan tasks are not listed in the *All Available Thresholds* of the system itself. You can see which threshold plans are applied to a system or group of systems in the IBM Director Console by selecting **Resource Monitors** in the Associations menu, as shown in Figure 9-88. The Threshold Plans are listed under the Resource Monitors directory that appears below the group L2 IBM Director agents.

This means that thresholds set on systems using IBM Director threshold plan tasks are not included in the list of monitors that can be added to the Health Center. Therefore, you must configure thresholds system-by-system, in order to monitor them in the Health Center.

The same concept of threshold plan tasks applies to both monitors on Windows services and monitors on processes, that is, you can create threshold tasks to propagate service monitors and process monitors. In the first case, the resulting service monitors behave as resource monitors. Therefore, they cannot be used in the Health Center Monitors portlet. On the other hand, Process Monitors are seen in the All Available Thresholds task on the selected system. Therefore, you can use these in the Virtualization Engine console Health Center.

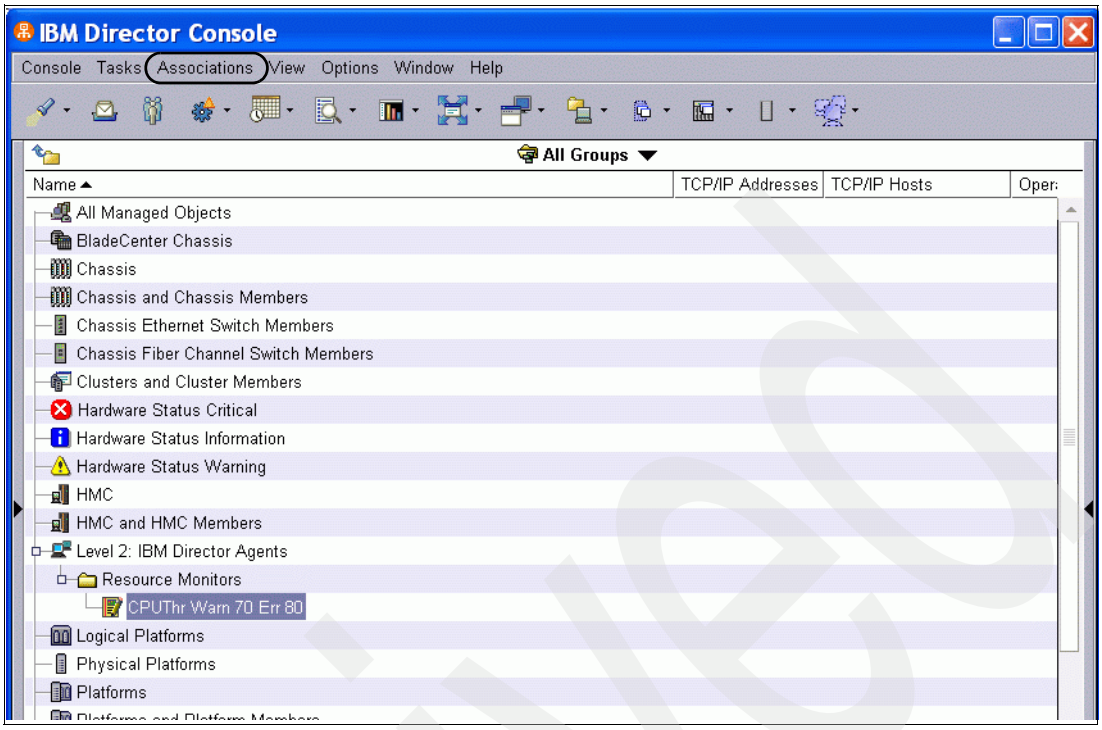


Figure 9-88 Threshold plans applied to a group in the IBM Director Console

9.5.3 Adding the CPU monitor to the dashboard

In addition to adding the CPU monitor to the list of monitored resources in the Health Center, you should also be able to see the value of the monitor in real time. For this, add the monitor to the Dashboard with the following steps:

1. Select **Health Center** → **Monitoring** → **CPUThr**. Click the arrow next to the system name and select **Add to Dashboard**, as shown in Figure 9-89.

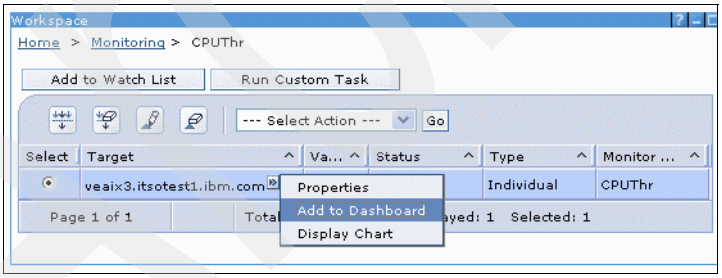


Figure 9-89 Select to add the monitor to the Dashboard.

2. In the Dashboard (as shown in Figure 9-90 on page 357), you should insert the parameters needed to create the desired graph. Insert the name and an optional description in the corresponding fields. Select **Dial** as Chart Type. Then, specify the lower and higher values for the graph. We insert 0 as low, 100 as high. Click **OK**.

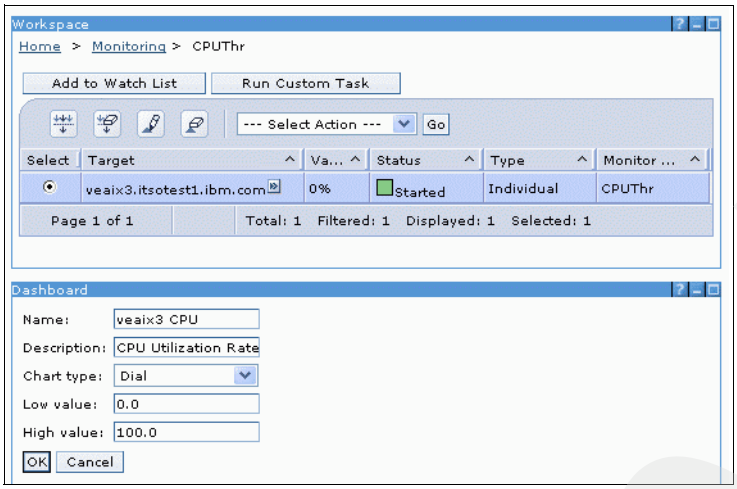


Figure 9-90 Dashboard configuration portlet

3. As you see in Figure 9-91, the Dashboard portlet now contains the value of the CPU utilization for the veaix3 system. If you move the mouse over the gauge, you will see the detailed value of the counter.

> **Note:** To have the flyover text feature in the Dashboard portlet, you should install the Scalable Vector Graphics (SVG) plug-in for your browser. If you do not have it installed, you will see a notification, as shown in Figure 9-92. Not all browsers support SVG. You can find more information and downloads on the Web at:
>
> `http://www.adobe.com/svg/viewer/install/main.html`

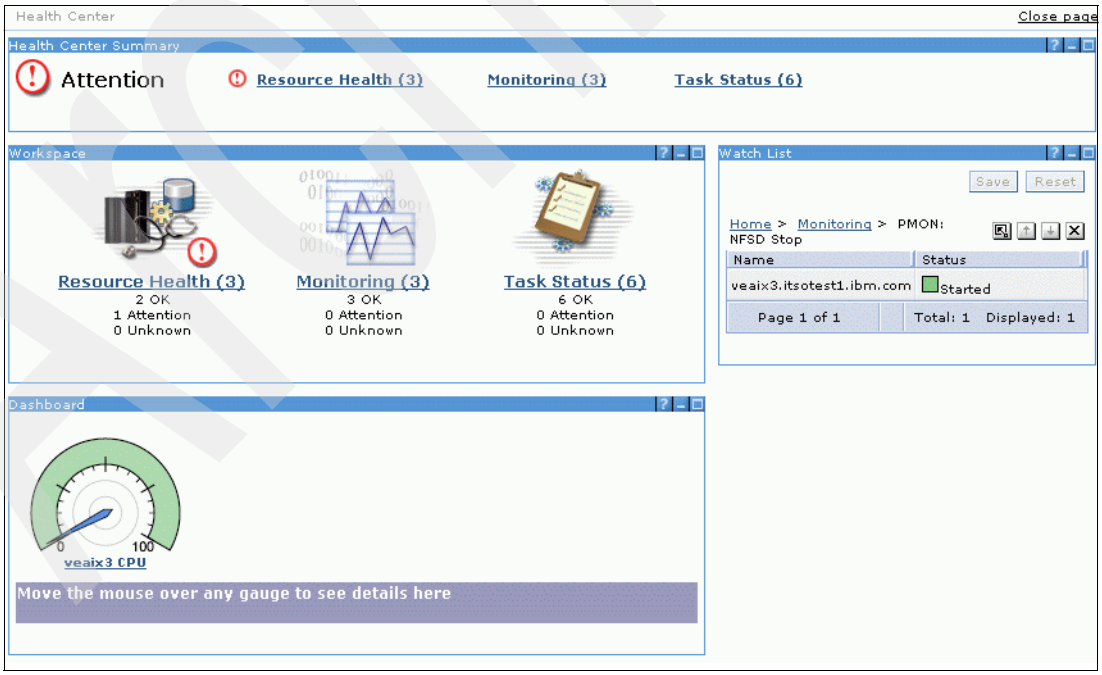


Figure 9-91 The Dashboard portlet contains the CPU utilization real time value for the selected system

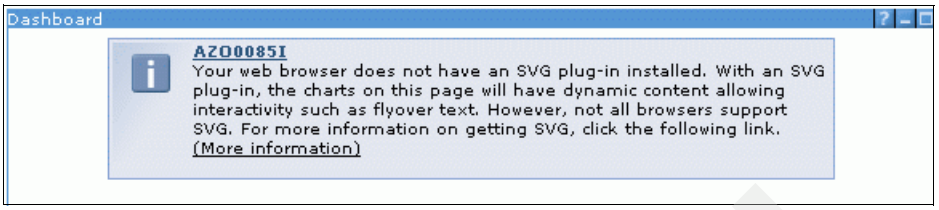


Figure 9-92 Missing SVG plug-in notification

When the CPU utilization goes beyond 80%, as shown in Figure 9-93, you see the Health Center summary showing that some monitors are in an error status (indicated with a red exclamation mark). You also have a notification in the Resource Health. As shown in Figure 9-93, we drilled down to the groups in the workspace portlet to show that the notification pertains to the Level 2: IBM Director Agent group, which contains the veaix3 system on which we caused the event. You now see the dashboard portlet showing the exact value of the CPU utilization monitor.

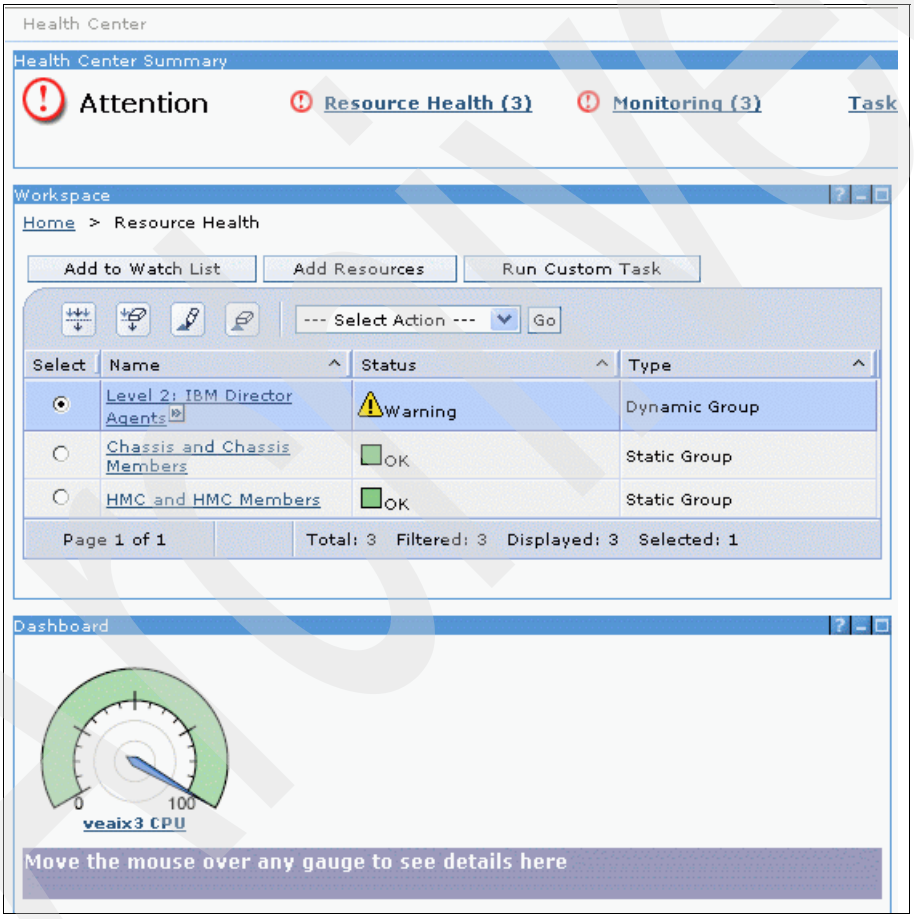


Figure 9-93 Health Center showing the CPU utilization to 100%

9.6 Managing a VMWare Virtual Machine

The use of virtualization technologies is becoming widespread. The term *virtualization* can mean several different things according to the context in which it is used. Here, by virtualization technologies, we mean the ability to run multiple operating systems on a single

piece of hardware. In the x86-based server space, there are several products available. However, there are three main products used by enterprises at a production level:

- VMWare ESX
- VMWare GSX
- Microsoft Virtual Server

A detailed description of these technologies is out of the scope of this book. You can find more information in *Virtualization on the IBM System x3950 Server*, SG24-7190.

In this scenario, we outline a possible usage of the Virtualization Engine console and the IBM Director Console to manage a virtual machine based on VMWare ESX. To do this, we exploit the Virtual Machine Manager (VMM), one of the IBM Director extensions.

As the IBM Director, VMM has three components:

- Console
- Server
- Agent

There are many possible ways of using VMM in a virtualized infrastructure. Refer to Chapter 5, “Management with IBM Director“, in *Virtualization on the IBM System x3950 Server*, SG24-7190. In this scenario, we are not using VMWare add-ons like VMWare Virtual Center to manage the VMWare infrastructure. We only have a VMWare ESX server where we installed the IBM Director L2 Agent and the VMM plug-in agent. The IBM Director Server has also been extended with the VMM plug-in. Figure 9-94 shows you a schematic view of this scenario.

Note: In order to visualize correctly the VMM added functions, you must install the VMM Console component on the system running the IBM Director Console. This is not included in the installation performed automatically by the Virtualization Engine console.

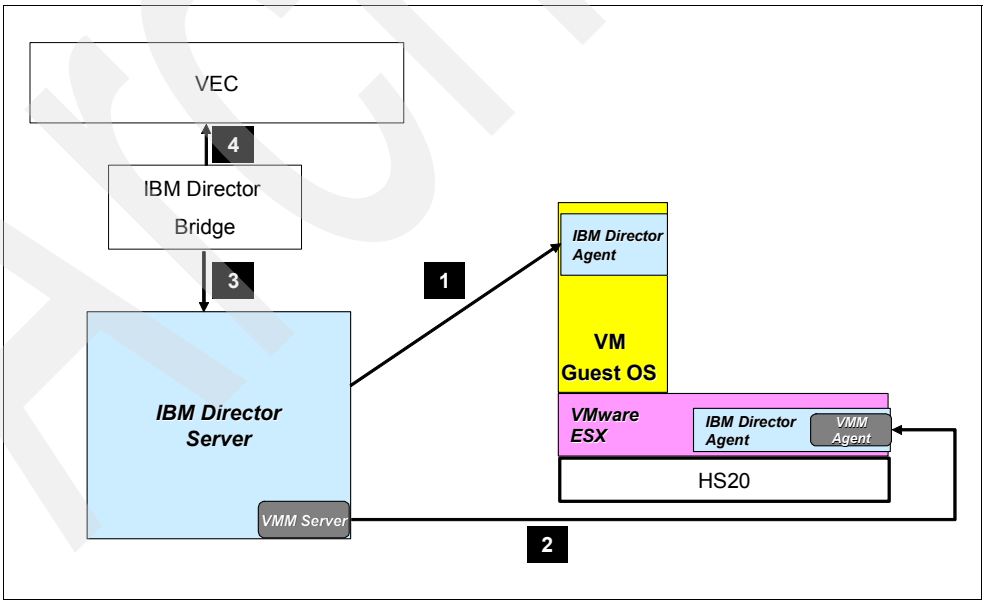


Figure 9-94 VMM environment

To summarize, perform the following steps. The IBM Director Server polls the IBM Director agent (presence check operation) periodically running on the Virtual Machine guest operating system (OS):

1. When the agent is detected as offline, an event is generated on the IBM Director Server (as shown in field 1 in Figure 9-94 on page 359).
2. We create an automatic action on the IBM Director Server that instructs VMWare ESX to restart the Virtual Machine through the VMM plug-in (as shown in field 2 in Figure 9-94 on page 359).
3. We then visualize the indications of what happened in Virtualization Engine console since the IBM Director Bridge polls the IBM Director Server (as shown in field 3 in Figure 9-94 on page 359), and populates the Virtualization Engine console Health Center (as shown in field 4 in Figure 9-94 on page 359).

Note: What we describe in this scenario is only a very limited subset of the functions implemented in VMM 2.1. For a detailed description, refer to *IBM Virtual Machine Manager Installation and Users Guide Version 2.1*, downloadable from the Web at:

`http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56955`

9.6.1 Requesting access to the ESX host

The VMWare ESX host is discovered by the IBM Director Server. For details on how to discover systems, refer to 9.2.3, “Discovering systems with IBM Director” on page 312. By default, all Linux hosts are secured. This means that you need to insert the credentials to request access to the managed resource. To do this, perform the following steps:

1. Open the IBM Director Console.
2. Locate the VMWare ESX host in the IBM Director Console, right-click it, and select **Request Access**. Enter the root user credentials.

   In our case, we located the VMWare ESX host by its IP address, which is 10.1.20.22, as shown in Figure 9-95. For more details on requesting access to systems, refer to 9.2.4, “Requesting access to systems in IBM Director” on page 315.

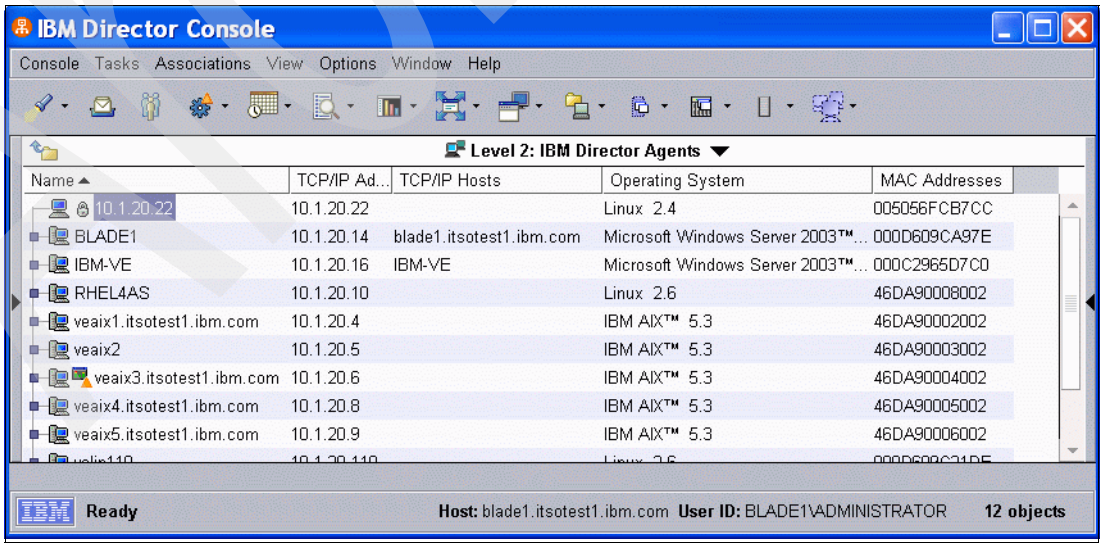


Figure 9-95 VMWare ESX object in IBM Director

Note: The VMware ESX host operating system detected by IBM Director is Linux 2.4. This is because VMWare console OS in ESX 2.5.2 is based on a Linux 2.4 kernel.

3. Once the ESX host object is unlocked, a full inventory scan is automatically performed. This detects the presence of the VMM plug-in on the host itself. The new icon , next to the host icon, identifies the object as a VMM object, as shown in Figure 9-96.

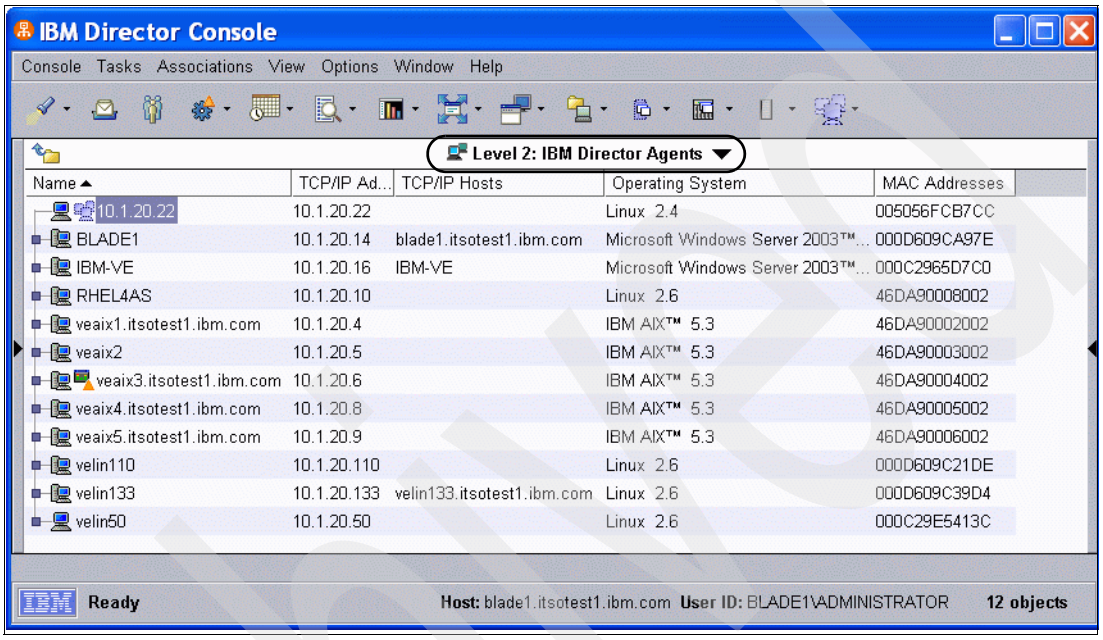


Figure 9-96 VMWare ESX Object unlocked in IBM Director Console

4. From the Group Selection list shown in Figure 9-96, select the group named **VMM Systems**. From the Associations menu, select **VMM Systems Membership**, as shown in Figure 9-97.

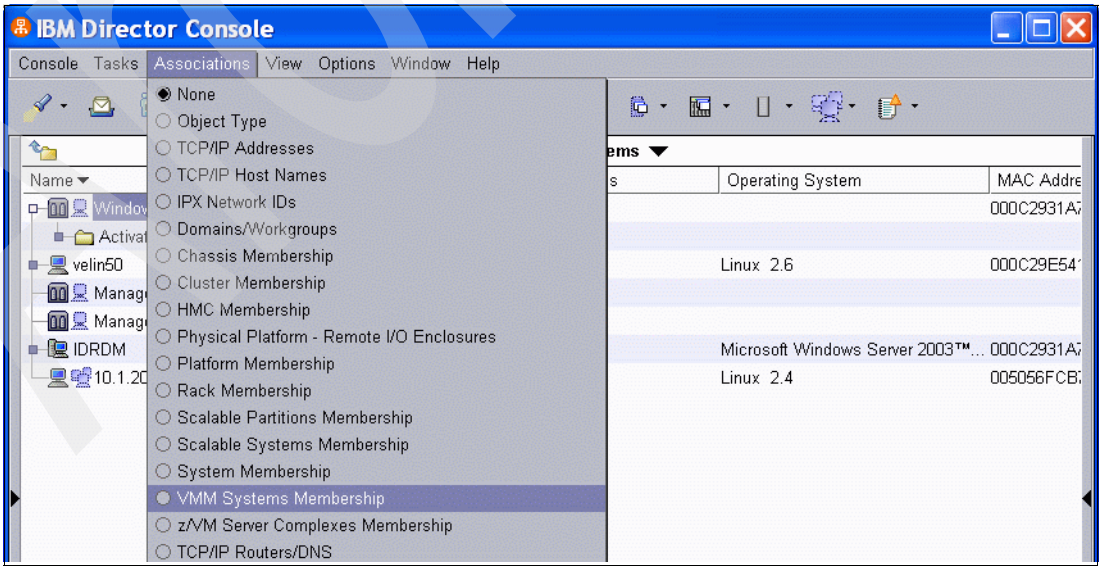


Figure 9-97 Select the VMM Systems Membership Association

This visualizes the VMM resources as a tree, as shown in Figure 9-98 on page 362.

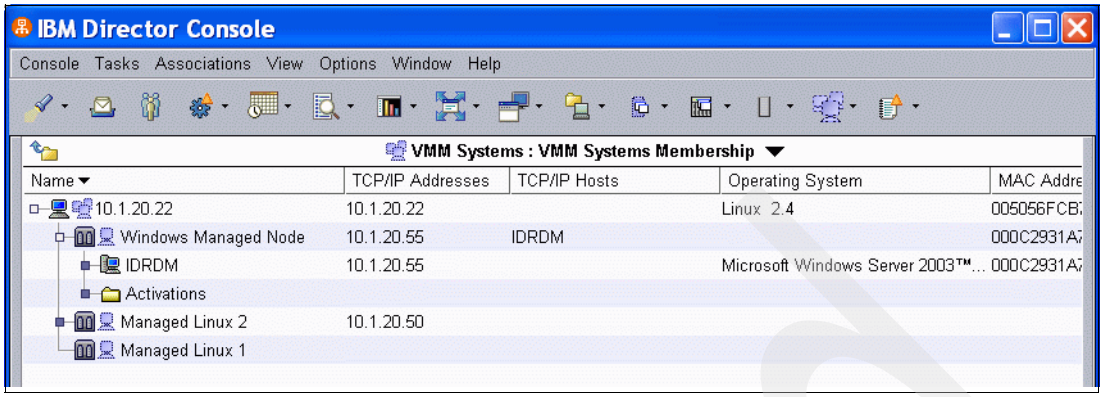


Figure 9-98 VMM group list

In the VMM Systems group, we see the following objects:

- The ESX host, named 10.1.20.22, which is the root of the tree.
- The list of Virtual Machines registered in the ESX host, identified by this icon , next to the name of the object.

Under each Virtual Machine, the corresponding Level2 IBM Director agent object, if any, is installed in the Virtual Machine guest OS.

9.6.2 Adding the resources to Health Center and the Watch List portlet

In order to manage the systems from the Virtualization Engine console, you should add them to the list of resources in the Health Center. To do this, open the Virtualization Engine console, open the Health Center, and perform the following steps:

1. Click **Resource Health**. In the Workspace, click **Add Resources**, as shown in Figure 9-99.

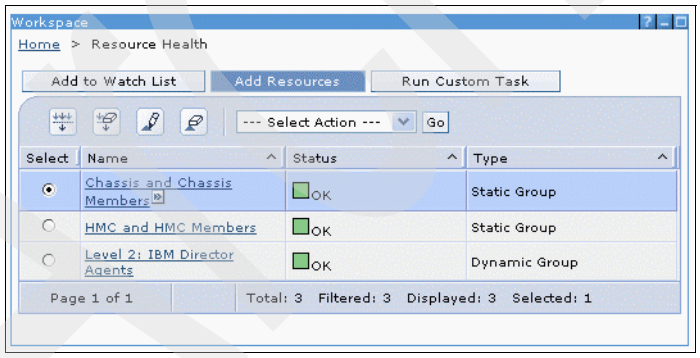


Figure 9-99 Adding the VMM Group to the Resource list

2. This brings you to the Add Resources portlet. Select the check box next to the VMM Systems group (browsing through the pages if necessary), and then click **OK**, as shown in Figure 9-100 on page 363.

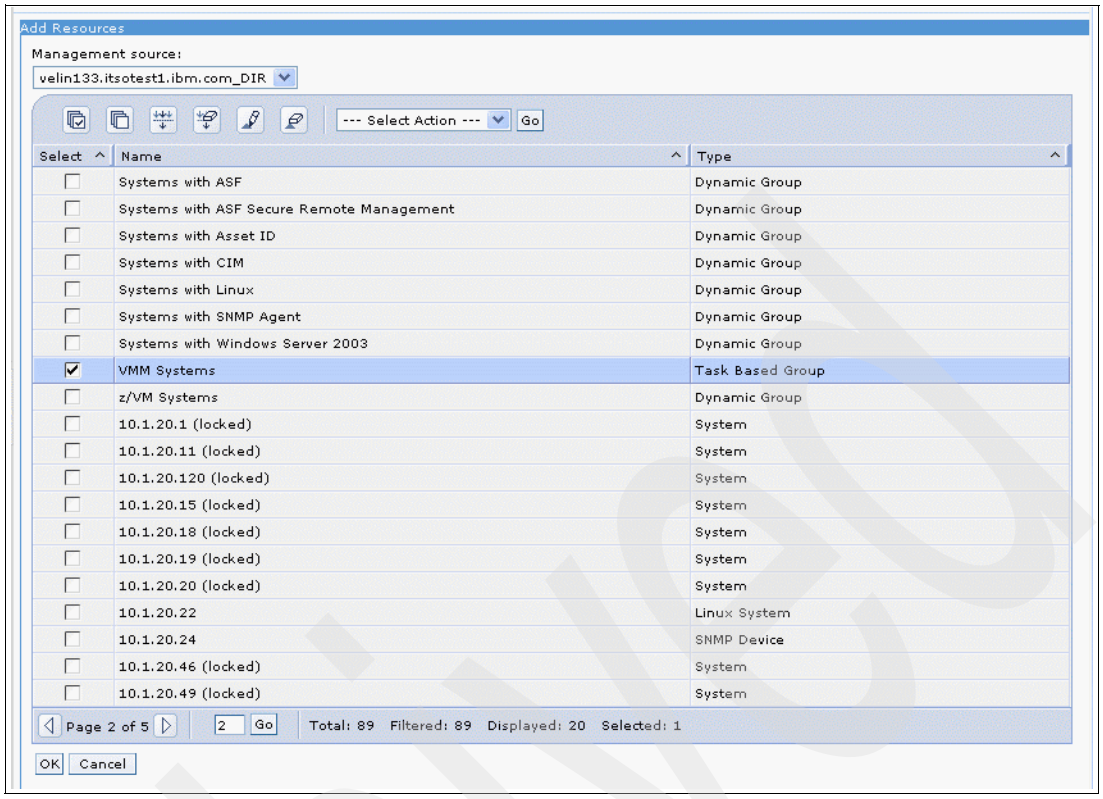


Figure 9-100 Selecting the VMM Systems group in the Add Resource portlet

3. In order to be able to see the VMM Systems group members status more easily, add them to the Watch List. To do this (as shown in Figure 9-101), from the Workspace portlet, select **Resource Health** → **VMM Systems** → **Group Members**. Click **Add to Watch List**.

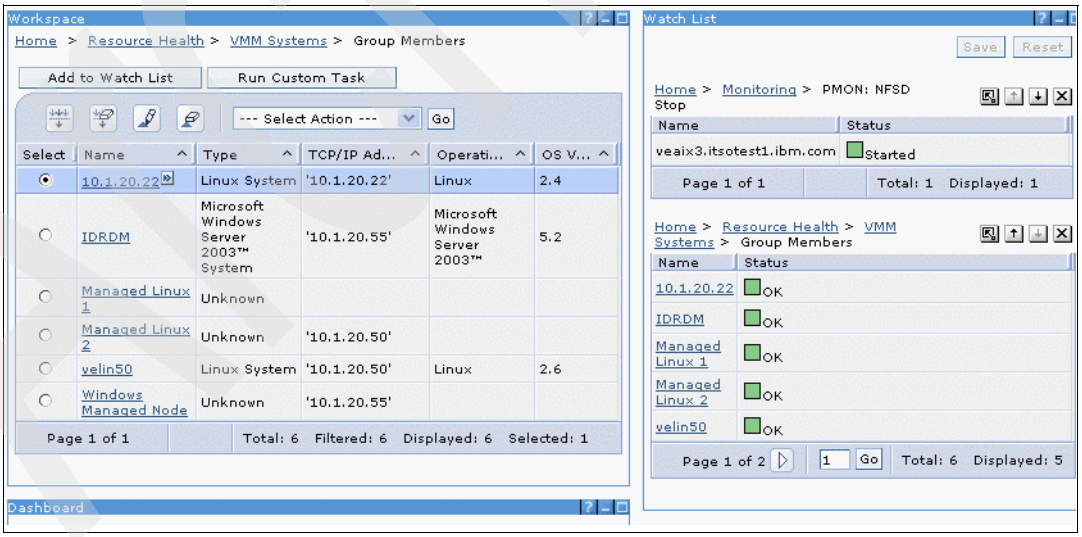


Figure 9-101 Adding the VMM Systems group to the Watch List

Note: The Group Members list in the Virtualization Engine console does not show the same tree view that you have in the IBM Director Console. Furthermore, the Virtual Machines have the system type set to unknown. We recommend that you use the IBM Director Console if you want to be able to understand which Virtual Machine is registered in which ESX host.

9.6.3 Configuring the Event Action plan

The next step is to configure an IBM Director Event Action plan to restart the Virtual Machine as soon as the IBM Director agent is detected as offline. To do this, open the Event Action Plan Builder in the IBM Director Console. You can open it from the IBM Director Console directly, or use the Virtualization Engine console to launch it by opening the Virtualization Engine console, expanding **Virtualization Engine Console** → **Health**, and selecting **Event Actions**, as shown in Figure 9-102.

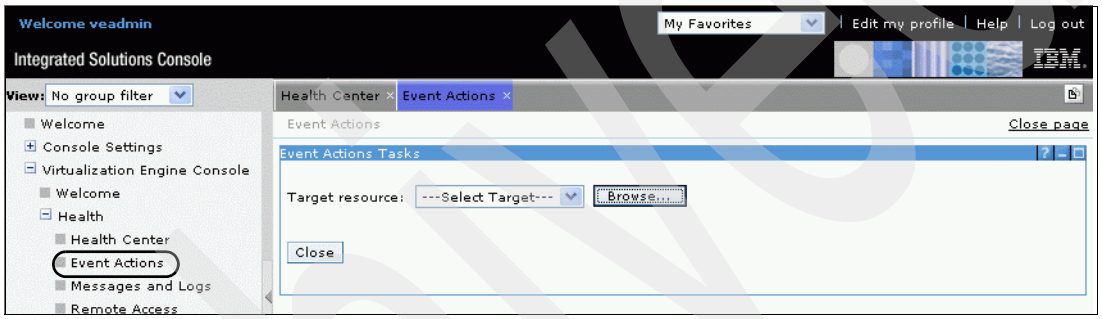


Figure 9-102 Launch the Event Action Plan Builder

Once you have opened the Event Action Plan Builder, perform the following steps:

1. Select the target resource from the list. If this is not populated, click **Browse** to select the resource in the list of available ones, as shown in Figure 9-103 on page 365. In our case, we select IDRDM, which is the Level 2 Agent associated with the VM we want to monitor.

Note: The action that we want to launch, the Event Action Plan Builder, does not need to be run against a target resource. Therefore, it does not matter which object you select as a Target resource in Figure 9-102. This applies to the Event Action Plan wizard action as well (as shown in Figure 9-104).

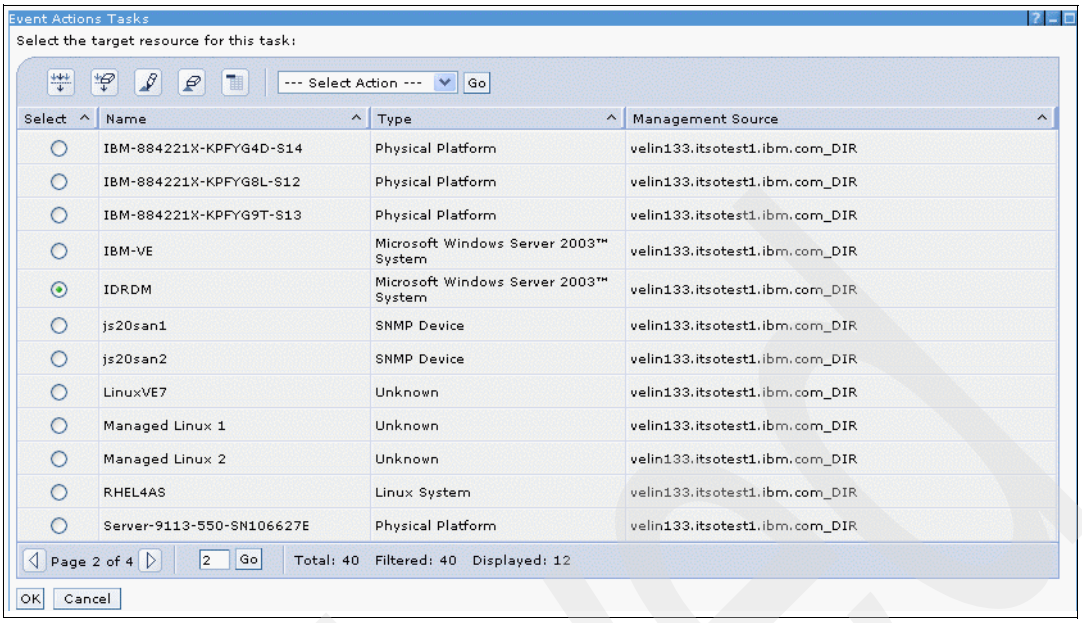


Figure 9-103 Selecting the object to launch the Event Action Plan Builder

2. After you select a target resource, click **OK**. The Event Actions Tasks portlet appears, as shown in Figure 9-104.

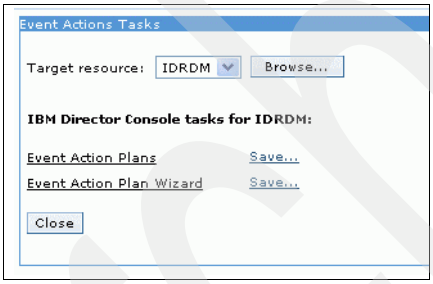


Figure 9-104 Launching the Event Action Plans Builder

3. Click **Event Action Plans**. This opens the IBM Director Console and launches the Event Action Plan Builder window, as shown in Figure 9-105 on page 366.

> **Note:** The other option is to launch the Event Action Plan wizard. This is aimed at performing basic configurations using a list of predefined actions. The action that we want to configure is not in this list, so you *must* use the Event Action Plan Builder.

Archive

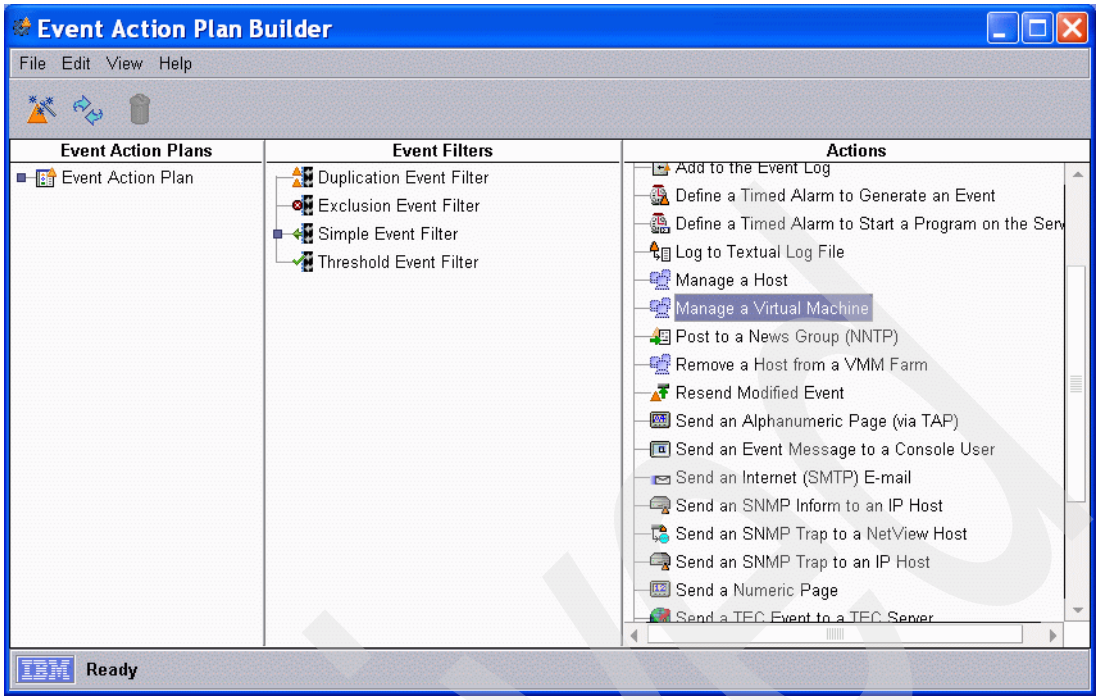


Figure 9-105 Event Action Plan Builder window

The Event Action Plan Builder window is composed of three panes, from right to left:

- *Actions*, where you select which action you want to perform
- *Event Filters*, where you select which event triggers the action
- *Event Action Plans*, where you bind actions to filters.

For a complete description of the Event Action Plan Builder, refer to *Implementing IBM Director 5.10*, SG24-6188.

4. Restart the Virtual Machine named Windows Managed Node, if the system is not responding. If the IBM Director agent has an offline status, this means that the Virtual Machine is hanging.

> **Note:** Identifying a condition to be sure that a system is not responding and needs to be restarted is not simple. In fact, the IBM Director Agent offline status is detected every time the IBM Director Agent is stopped or it cannot be reached, for example, because the network connection is not available. In a real life environment, we recommend also including other indications to check that the system is hanging prior to automating the restart of a virtual machine.

5. In the Actions pane, right-click **Manage a Virtual Machine** and select **Customize**, as shown in Figure 9-106 on page 367.

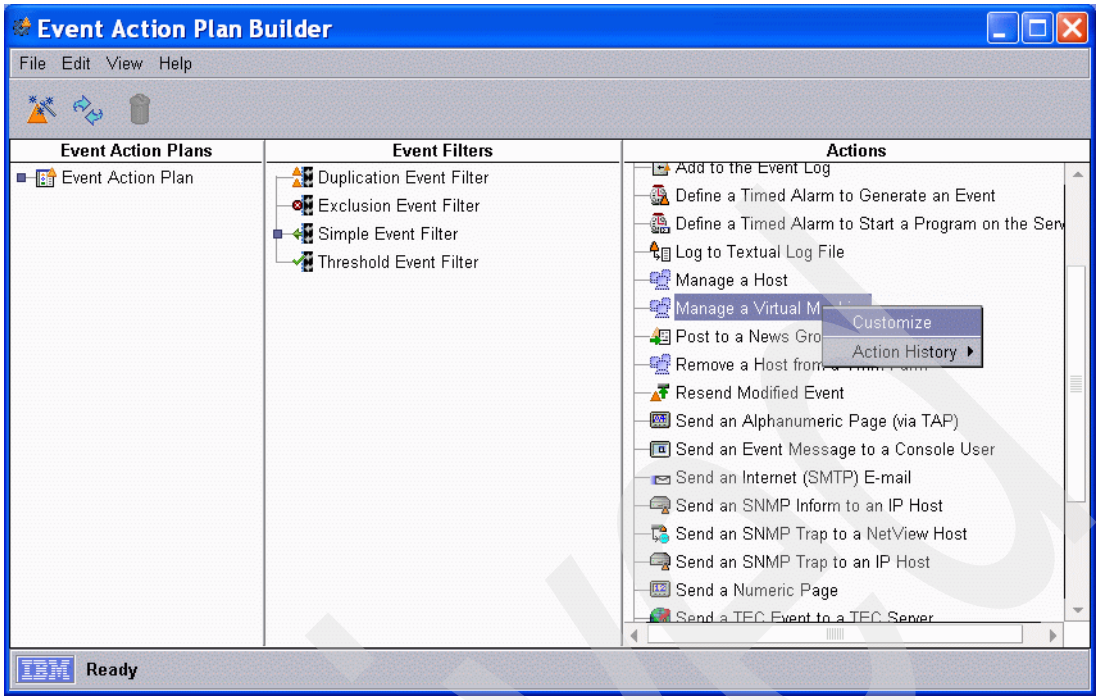


Figure 9-106 Customizing the Manage a Virtual Machine action

This opens the Customize Action window, as shown in Figure 9-107.

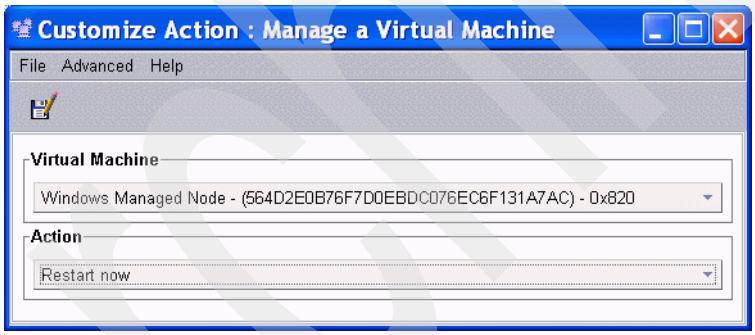


Figure 9-107 Customize Action: Manage a Virtual Machine window

6. Select the target Virtual Machine (Windows Managed Node, in our example). Select the type of action you want to perform (Restart now, in our example). Click the **Save As** icon and save the action with a name. In our scenario, we chose the name RestartWindowsVM. The saved action is listed under the Manage a Virtual Machine category of actions, as shown in Figure 9-108 on page 368.

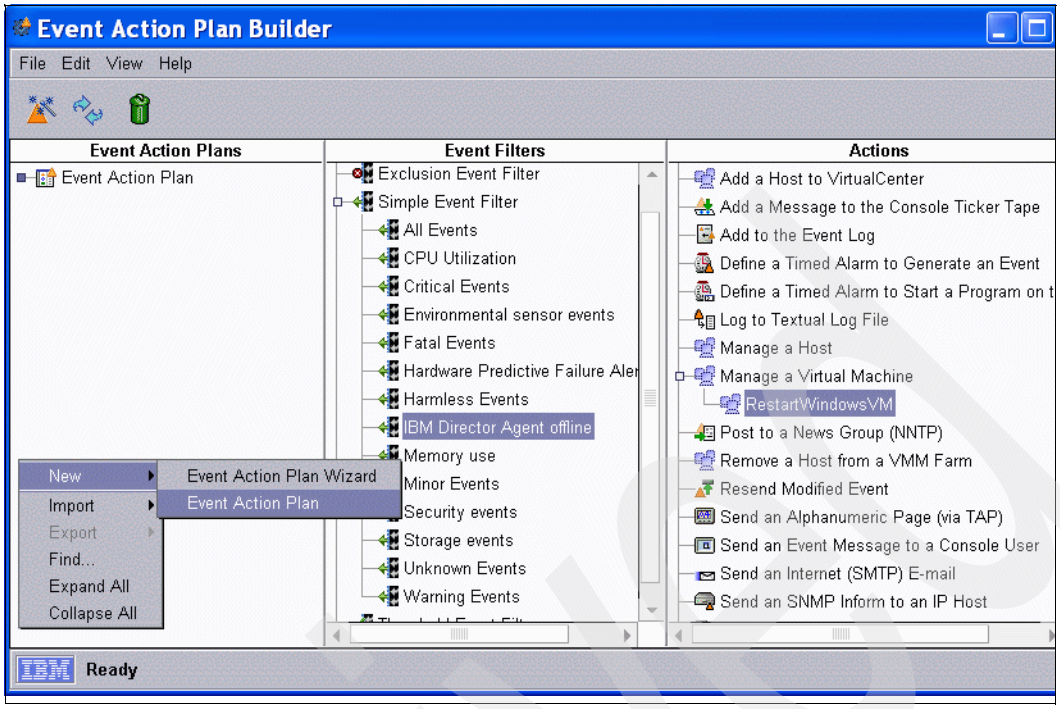


Figure 9-108 New Event Action Plan

7. As a filter, we use the predefined IBM Director Agent: offline, which is listed under the Simple Event Filter category in the Event Filters pane (as shown in Figure 9-108). Now create the association between the Event Filter and the Action, by creating the Event Action Plan. Right-click anywhere in the Event Action Plans pane and select **New** → **Event Action Plan**. This opens the Create Event Action Plan window, as shown in Figure 9-109.

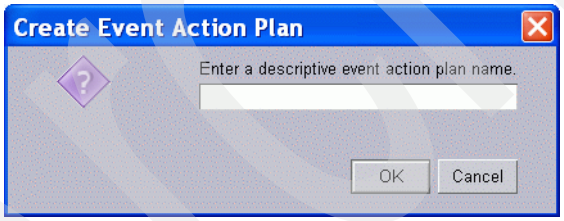


Figure 9-109 Create Event Action Plan window

8. Enter a name for the new Event Action Plan. In our case, we named it RestartWindowsVM. Click **OK**. You now have a new Event Action Plan, as shown in Figure 9-110 on page 369.

Archive

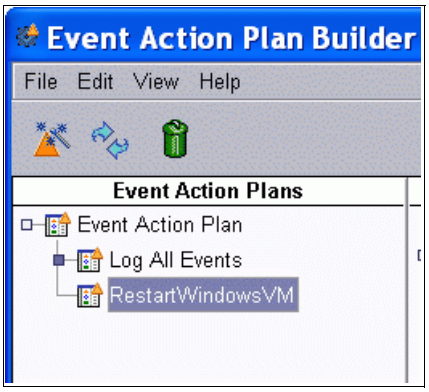


Figure 9-110 A new empty Event Action Plan

9. Drag the IBM Director Agent offline Event Filter onto the RestartWindowsVM Event Action Plan (as shown in step 1 in Figure 9-111). Drag the RestartWindowsVM Action on the IBM Director Agent offline Filter in the Event Action Plan pane (not in the Event Filter pane) (as shown in step 2 in Figure 9-111). This binds the action to the filter.

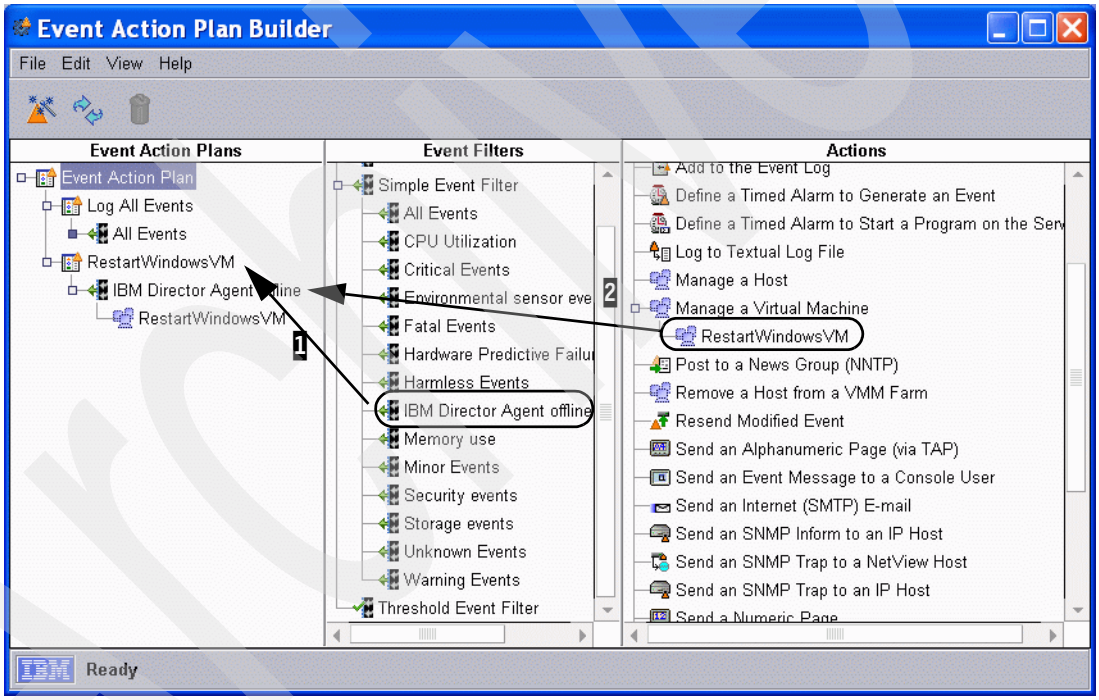


Figure 9-111 Populating the newly-created Event Action Plan

10. Up to now, you have only created a rule that says that every time a certain kind of event is received, the IBM Director must perform a certain action. You selected the target of the action, but not which machine is generating the event. You can do this by applying the Event Action Plan to the machine you want to filter the event from. To do this, perform the following steps:

    a. In the IBM Director Console, select the Level2 IBM Director Agent object (IDRDM, in our case), as in Figure 9-112 on page 370.

    b. From the Navigation pane, select the Event Action Plans icon. Select **RestartWindowsVM**, and then select **RestartWindowsVM:IDRDM**.

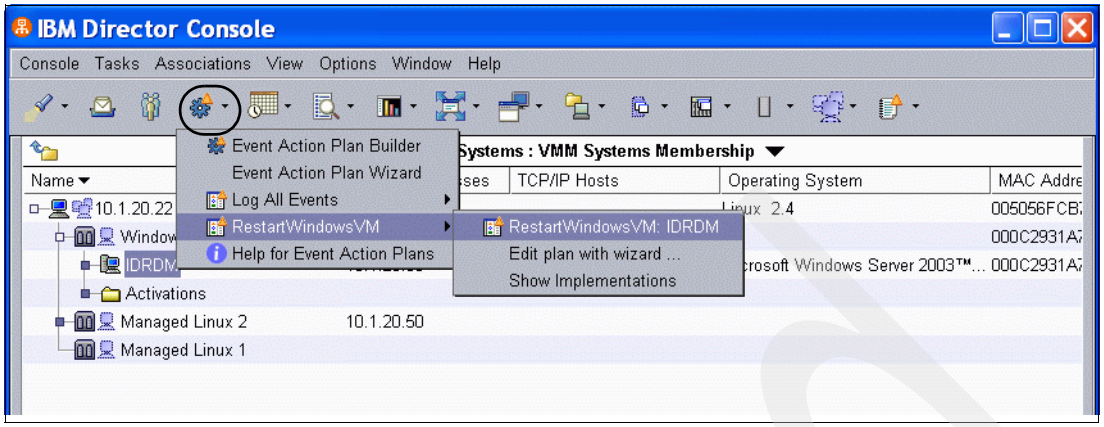


Figure 9-112 Applying the Event Action Plan to the target system

An information window pops up stating that the Event Action Plan has been successfully applied, as shown in Figure 9-113.

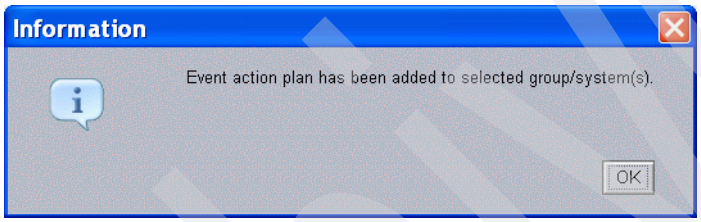


Figure 9-113 Information window

From now, everytime the IBM Director Server receives an IBM Director agent offline event on the system named IDRDM, the corresponding Virtual Machine is restarted.

9.6.4 Visualizing in the Virtualization Engine console

As we configured the Virtualization Engine console to show the VMM Systems group content in the Watch List in 9.6.2, “Adding the resources to Health Center and the Watch List portlet” on page 362, when the IBM Director Agent on the guest OS (IDRDM, in our example) of the Virtual Machine is offline, the warning message shown in Figure 9-114 on page 371 appears.

Archived

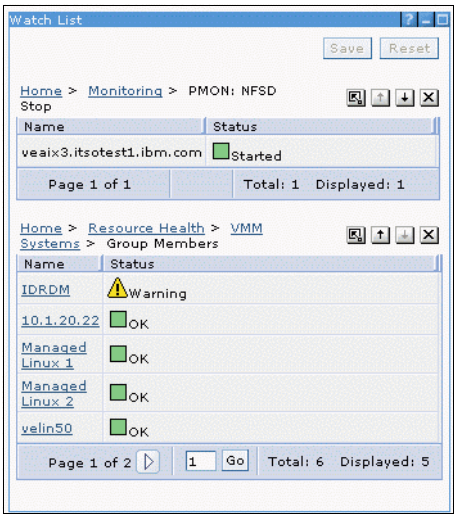


Figure 9-114 IDRDM is offline in the Watch List

You can see the details of the event from the Health Center workspace portlet. Select **Resource Health → VMM Systems → Logs → EventLog → All Events**, as shown in Figure 9-115. You will also see an entry stating that the Virtual Machine has been restarted because of the Event Action Plan that has been configured.

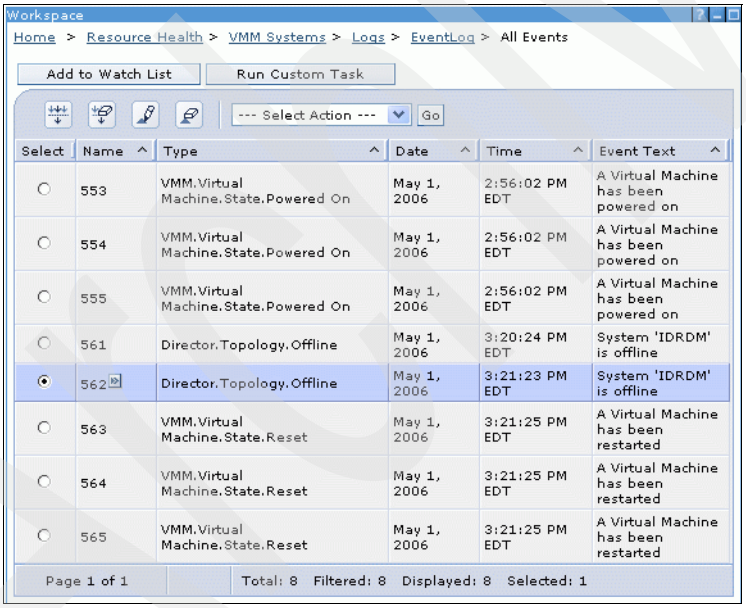


Figure 9-115 Workspace showing that the Virtual Machine has been restarted

Archived

Note: The Virtual Machine restarted notification is registered three times. This is because three different objects in the IBM Director are generating an event following the restart of the Virtual Machine:

- The IBM Director Agent object IDRDM
- The Virtual Machine object Windows Managed Node
- The ESX object 10.1.20.22

The difference between the entries is more clear in the IBM Director Event Log (shown in Figure 9-116), where you see the object that generated the event in the System Name field.

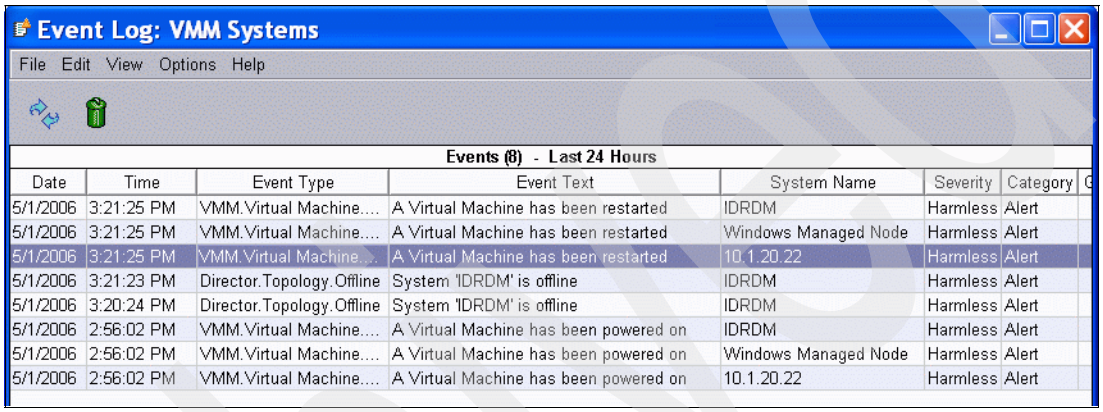


Figure 9-116 IBM Director Event Log containing the three indications that the Virtual Machine has been restarted.

9.7 Monitoring an LPAR

IBM Director is able to automatically discover an HMC through the Service Location Protocol (SLP) Agent that is installed on the HMC itself. After discovering an HMC, you must request access to the system with the proper credentials, as described in 9.2.4, “Requesting access to systems in IBM Director” on page 315. Once access has been successfully granted, the IBM Director attempts to access the Pegasus Common Information Model Object Manager (CIMOM) so that information regarding Central Electronic Complexes (CECs) and LPARs are gathered and represented as IBM Director objects.

All the actions that the IBM Director is able to launch on the CECs and LPARs managed by the HMC are performed through the CIM instrumentation on the HMC itself.

In this scenario, we describe how to receive events from the HMC managing the p550 system in our environment. Figure 9-117 on page 373 schematically describes the event flow. We power off the LPAR named VELPAR4. This is detected by the HMC (1 in Figure 9-117 on page 373). This in turn communicates this event to the IBM Director Server (2 in Figure 9-117 on page 373) that is polled by the Bridge (3) to update the Virtualization Engine console (4), as shown in Figure 9-117 on page 373. Once we see the event in the Health Center, we restart the LPAR manually.

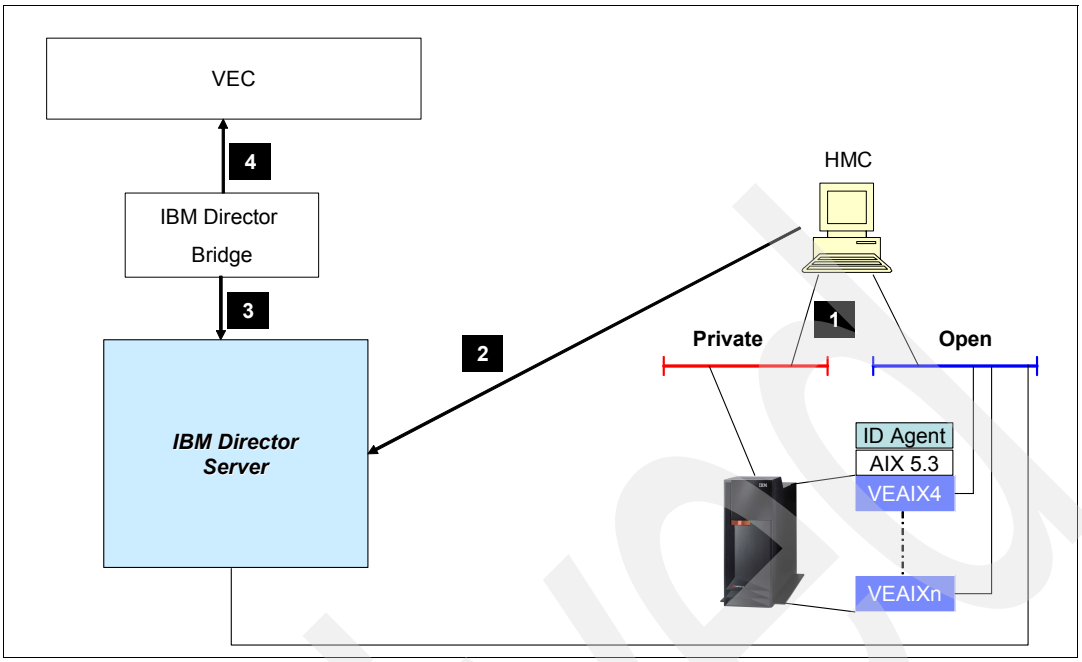


Figure 9-117 Event flow from HMC

Note: Though the IBM Director Agent is installed on the AIX 5L V5.3 instance running in the VEAIX4 LPAR, it is not used in this scenario. All the actions described also work when the agent is not installed.

Since we already added the HMC and HMC Members group to the Resource Health, as described in 9.3, “Diagnosing a hardware error” on page 330, when the LPAR is turned off, an error message appears in the Virtualization Engine console. To view it, from the Health Center, select **Resource Health** → **HMC and HMC Members** → **Logs** → **Event Log** → **All Events** as shown in Figure 9-118. In the Event Text field, you will see a notification that the LPAR has been powered off.

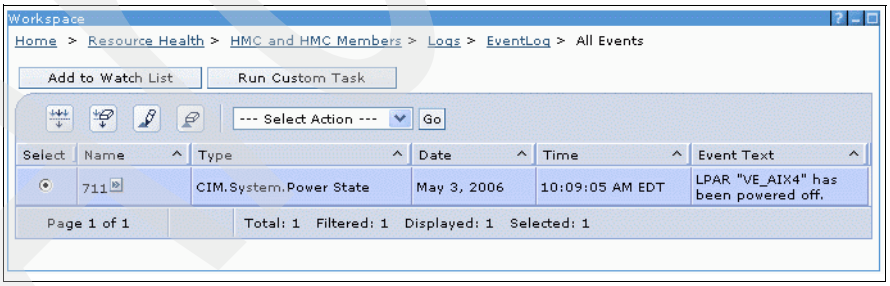


Figure 9-118 Event Log showing that an LPAR has been powered off

If you want to restart the LPAR, perform the following steps:

1. Select the LPAR resource in the Health Center by opening **Resource Health** → **HMC and HMC Members** → **Group Members**, as shown in Figure 9-119 on page 374.

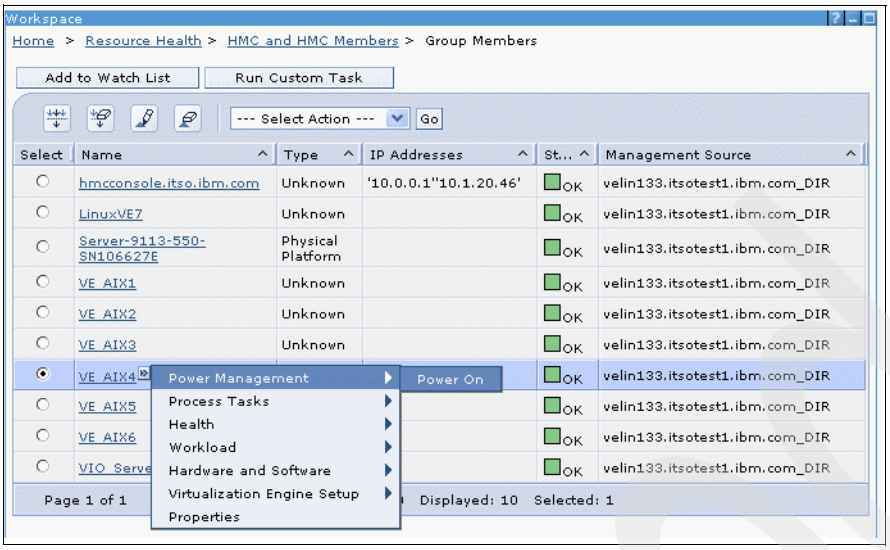


Figure 9-119 Launch-in-context for an LPAR

2. Click the arrow next to the name of the system and select **Power Management™ → Power On**.

> **Note:** The launch-in-context Power Management menu content is updated according to the LPAR status. If it is powered on, you have the option to Power Off Now or Restart Now.

If you want to perform more advanced tasks on the LPAR, such as changing its performance parameters, then you must use WebSM. To do this, launch the specific task in WebSM from the HMC object itself. For example, to launch Server Management, select the HMC object, click the arrow next to the name of the system, as shown in Figure 9-120, and select **Hardware and Software → Hardware Management → Server and Partition Management: Server Management**.

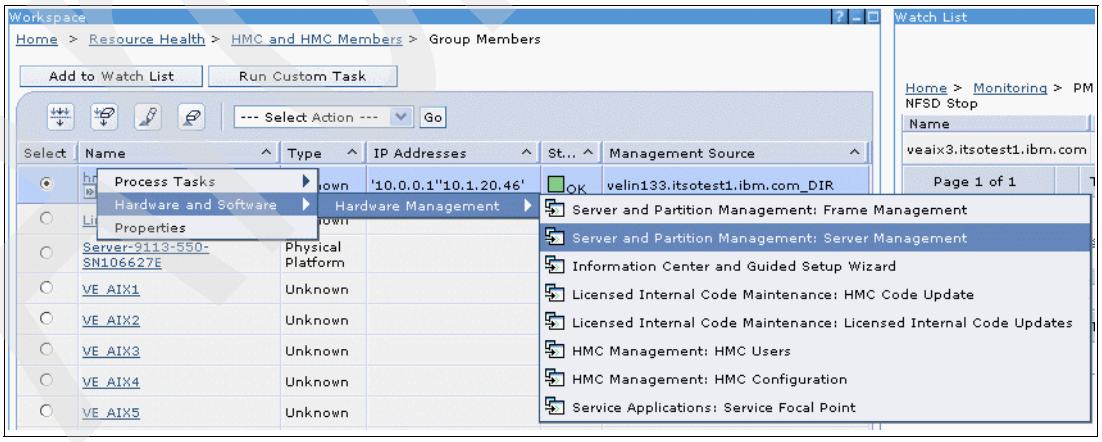


Figure 9-120 WebSM launch-in-context tasks

This opens the Server Management tasks in the WebSM application.

Note: The HMC host name must be resolved by the workstation where you are running the Virtualization Engine console. You also need to enter the credentials to connect to the WebSM application.

9.8 Starting and stopping services

In this section, we describe how to start and stop the services we are using.

These services are:

- Virtualization Engine console
- IBM Director Bridge
- IBM Director Server

9.8.1 Virtualization Engine console

To start the Virtualization Engine console, type the following command from a command prompt on the host running the Virtualization Engine console (both for Windows and Linux):

```
IBM_VE2_LOCATION/ISC/PortalServer/bin/startISC.sh ISC_Portal
```

In this command, *IBM_VE2_LOCATION* is the Virtualization Engine installation root.

To stop Virtualization Engine console, type the following command:

```
IBM_VE2_LOCATION/ISC/PortalServer/bin/stopISC.sh ISC_Portal username password
```

In this command, *IBM_VE2_LOCATION* is the Virtualization Engine installation root, *username* is the Virtualization Engine administrator user name, and *password* is the corresponding password.

9.8.2 IBM Director Bridge

As we said, the IBM Director Bridge is software running as an application inside a WebSphere Application Server. There are two methods to start and stop the IBM Director Bridge:

- Act on the application itself, through the WebSphere Application Administration Console.

  To do this, perform the following steps:

  a. Open a Web browser and enter the following Web address:

     ```
     http://hostname:portnumber/admin
     ```

     In this Web address, *hostname* is your Virtualization Engine console server host name, and *portnumber* is the port used by the WebSphere Application Administration Console. The *portnumber* default value is 18552.

Tip: To find the WebSphere Application Administration Console port number, open the VEInstall.properties file in the Virtualization Engine Installation root directory. The value of the variable named *ibm.ve2.VEMS.adminPort* is the running session administration port.

Archivo

b. Expand **Application** and click **Enterprise Applications**. The list of applications currently running appears, as shown in Figure 9-121. Check the box next to EsmDir, and click **Stop**.

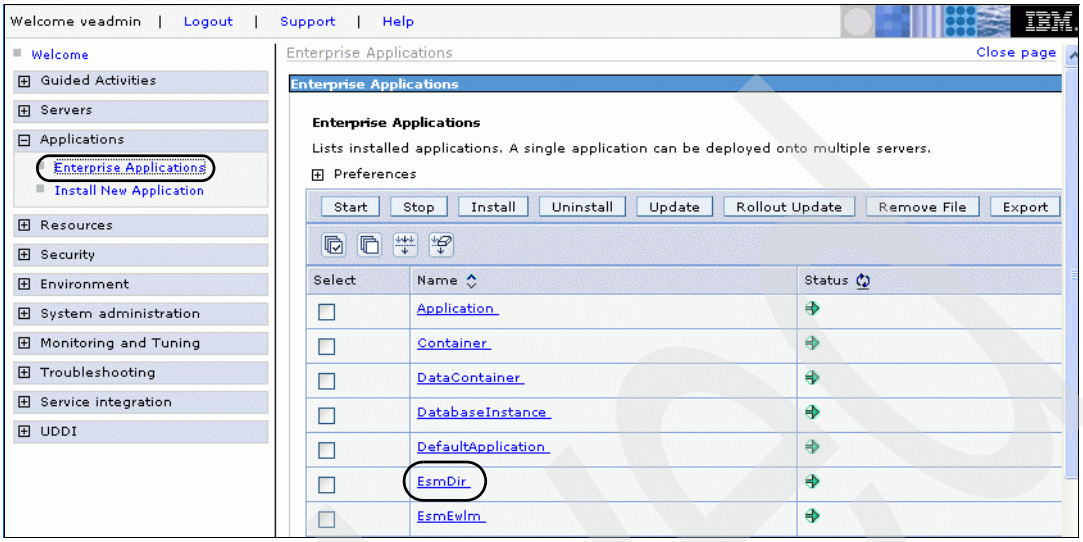


Figure 9-121 Enterprise application list in the WebSphere Application Administration Console

When the service is stopped, the ✖ icon appears in the Status column.

To start the application, check the box next to EsmDir, and then click **Start**.

- Act on the WebSphere Application Server as a whole.

  To start it, open a command prompt on the host running the Bridge and type:

  `IBM_VE2_LOCATION/VEMS/bin/startServer.sh VEMS`

  In this command, *IBM_VE2_LOCATION* is the Virtualization Engine installation root.

  When you start the WebSphere Application Server, you find the corresponding process identifier in the following text file:

  `/var/log/VEServices/VEMS/VEMS.pid`

  To stop the WebSphere Application Server, open a command prompt on the host running the Bridge and type:

  `IBM_VE2_LOCATION/VEMS/bin/stopServer.sh VEMS -username yourveadmin -password yourpwd`

  In this command, *IBM_VE2_LOCATION* is the Virtualization Engine installation root, *yourveadmin* is the Virtualization Engine administrator user name, and *yourpwd* is the corresponding password.

  Note: This operation stops the WebSphere Application Server and all the applications it is running. For example, if you are running Resource Dependency Service on the same machine, this operation will stop it too.

Archive

Note: The IBM Director Bridge logs are in the following directory:

`IBM_VE2_LOCATION/VEMS/logs/VEMS`

The IBM Director Bridge log files are:

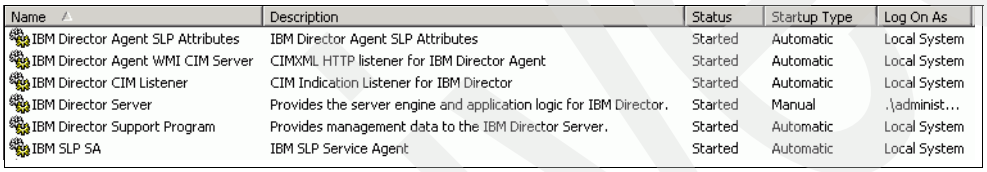
- VEBridge_err.log
- VEBridge_out.log

9.8.3 IBM Director Server

The IBM Director Server component runs on either a Windows system or a UNIX-like system.

Windows systems

On a Windows system, IBM Director is managed through Windows Services. Several services are registered after the IBM Director Server installation, as you can see from the output of the Services Microsoft Management Console (MMC) in Figure 9-122.



Name	Description	Status	Startup Type	Log On As
IBM Director Agent SLP Attributes	IBM Director Agent SLP Attributes	Started	Automatic	Local System
IBM Director Agent WMI CIM Server	CIMXML HTTP listener for IBM Director Agent	Started	Automatic	Local System
IBM Director CIM Listener	CIM Indication Listener for IBM Director	Started	Automatic	Local System
IBM Director Server	Provides the server engine and application logic for IBM Director.	Started	Manual	.\administ...
IBM Director Support Program	Provides management data to the IBM Director Server.	Started	Automatic	Local System
IBM SLP SA	IBM SLP Service Agent	Started	Automatic	Local System

Figure 9-122 IBM Director Services in a Windows host

These services are:

- IBM Director Agent SLP attributes
- IBM Director Agent WMI CIM server
- IBM Director CIM Listener
- IBM Director Server
- IBM Director Support Program
- IBM SLP SA

All these are related to the IBM Director Agent component that is part of the server installation except for the *IBM Director Server* service, which is the only one related to the server component.

Important: The Startup Type for all these services is set to Automatic except for the IBM Director Server. The reason for this is that the IBM Director Support Program controls the IBM Director Server service too. That is, when the IBM Director Support Program service starts, the IBM Director Server service starts too. This also happens when you stop the services.

You have two ways of starting IBM Director services:

- Through the services MMC:

  Select the service you want to start, right-click it, and select **Start**.

- Through the command line:

  To start the services, type the following command in a command prompt:

  `net start` *`ServiceName`*

  In this command, *ServiceName* is the name of the service you want to start.

Note: You can also type the short name for the service. For example, you can start the IBM Director Support Program with the following command:

```
net start twgipc
```

There are two methods to stop the services:

- Through the services MMC:

  Select the service you want to stop, right-click it, and select **Stop**.

- Through the command line:

  Type the following command in a command prompt:

  ```
  net stop ServiceName
  ```

  In this command, *ServiceName* is the name of the service you want to stop.

UNIX-like systems

To start the IBM Director Server on a UNIX-like system, open a command prompt and type:

```
/opt/ibm/director/bin/twgstart
```

To stop the IBM Director, type:

```
/opt/ibm/director/bin/twgstop
```

To check its status, type:

```
/opt/ibm/director/bin/twgstat
```

Tip: The IBM Director installation process adds the /opt/ibm/director/bin directory to the PATH environment variable for both Windows and UNIX-like systems. Therefore, you do not need to type the full path to the commands to execute them.

Error logs

The error generated by the IBM Director Bridge when ServeRAID Manager 8.25 is installed on the IBM Director Server is in the Bridge Error log located in the following directory:

```
IBM_VE2_LOCATION/VEMS/logs/VEMS
```

The file is named VEBridge_err.log. The ServeRAID error appears as shown in Example 9-1.

Example 9-1 ServeRAID error

```
- AVE4143E Communication to Director server failed.
com.ibm.sysmgtlib.director.task.TaskActivationServiceException: Communication to
Director server  failed.
Reason code: 1
com.ibm.sysmgt.sdk.remote.ServerCommunicationException: Exception received:
com.ibm.sysmgt.sdk.remote.TaskActivation; local class incompatible:streamclassdesc
serialVersionUID = -5211320047777772187, local class serialVersionUID =
-1106630560579716477
```

Archived

10

Using the Enterprise Workload Manager

This IBM Redbook does not focus on how to use the Enterprise Workload Manager (EWLM). Many books and papers are already available that demonstrate the value of EWLM and describe how to use it.

EWLM is part of the Virtualization Engine and provides many functions that can be used to facilitate or automate the infrastructure management. In this chapter, we refer to the three main functions of the EWLM:

- Workload balancing
- Application Response Measurement (ARM) instrumentation
- Logical partition (LPAR) management

These are already described in books or papers to which we refer.

© Copyright IBM Corp. 2006. All rights reserved.

10.1 Load Balancing with the CISCO Content Switching Module, SASP, and the IBM Enterprise Workload Manager

This chapter describes how to use a Content Switching Module and the IBM Enterprise Workload Manager for efficient load balancing. It briefly explains the functionality of each component and details the configuration necessary for both. Readers of this chapter include technologists seeking to discover more about the interaction between these two products, and administrators attempting to configure this interaction for use in their environments.

The paper can be found in the following Web site:

http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlmload1/

10.2 Performance monitoring using Enterprise Workload Manager

To meet real-world business commitments, gathering detailed statistics of the transactions in your environment can be critical. Learn how you can set up end-to-end performance monitoring in a functional IBM Enterprise Workload Manager (EWLM) domain and see how to turn on ARM instrumentation for IBM HTTP Server, WebSphere Application Server (Application Server), and DB2 UDB on AIX and Linux. This enables you to gather detailed statistics. This article also demonstrates how to troubleshoot this implementation.

The article can be found at:

http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-arm/

You can find a copy of the full paper in Appendix A, “Performance monitoring using Enterprise Workload Manager” on page 385.

10.3 Monitoring, reporting, and partition management using Enterprise Workload Manager

Two following sources of information provide you with information about how to manage a partition with EWLM.

- *IBM Enterprise Workload Manager Version 2.1*, SG24-6785 is an IBM Redbook dedicated to EWLM. This IBM Redbook describes multiple scenarios about how to use EWLM for monitoring and reporting, and for LPAR management on an IBM System p.

  In particular, this IBM Redbook details four scenarios:

  - The first three scenarios describe POWER partition management functionality of Enterprise Workload Manager V2.1.
  - The fourth scenario includes z/OS hardware and discusses the EWLM and WLM monitoring functionalities for z/OS.
- Additional scenarios that describe how EWLM can dynamically shift processing capacity to manage workloads using an IBM eServer i5 system running i5/OS logical partitions can be found at:

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/veicinfo/eicar_i5p5_scenario.htm

The scenarios include information and steps on the assessment, planning, and setup of an EWLM Domain Policy, as well as information about how EWLM performs partition management.

We recommend that you read *IBM Enterprise Workload Manager Version 2.1*, SG24-6785 to get the full information. You can download this IBM Redbook at:

http://www.redbooks.ibm.com/redbooks/pdfs/sg246785.pdf

- On the Developer Works IBM web site, you can find multiple articles about “Enabling autonomic partition management on System p5 servers with the Enterprise Workload Manager”. Refer to:

  http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-lpar1/
  http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-lpar2/

Archived

Archived

Part 3

Appendixes

Archived

© Copyright IBM Corp. 2006. All rights reserved.

Archived

A

Performance monitoring using Enterprise Workload Manager

To meet real-world business commitments, gathering detailed statistics of the transactions in your environment can be critical. This appendix shows you how you can set up end-to-end performance monitoring in a functional IBM Enterprise Workload Manager (EWLM) domain, and how to turn on Application Response Measurement (ARM) instrumentation for IBM HTTP Server, WebSphere Application Server (Application Server), and DB2 UDB on AIX and Linux, which enables you to gather detailed statistics. This appendix also demonstrates how to troubleshoot this implementation.

The author of this appendix is CheKim Chhuor, IBM.

You can find this article on the Web at:

http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-arm/

Archived

© Copyright IBM Corp. 2006. All rights reserved.

Introduction

Today, multi-tier Web applications and clusters are very common in IT environments. System administrators need to know how their Web applications are performing to satisfy established business goals or service level agreements (SLAs). There are several tools to let you monitor systems and network capacity. Most middleware has some means of providing internal statistics. Other tools and service providers can tell you how your applications are performing from different viewpoints around the world.

However, very few tools can give you an end-to-end view of every transaction coming across your middleware, from the moment the request hits the front-end Web server all the way to the back-end database, across multiple hops and multiple platforms. This is the unique capability of the EWLM for Multiplatforms, a component of the IBM Virtualization Engine Server Suite.

For any given transaction, EWLM can tell you the path that was taken, the time it has spent at each middleware hop, and other system statistics. With this data, EWLM can draw a topology of the middleware hosting your application, which helps you identify a potential bottleneck server among a cluster, if it occurs.

Middleware can mean different things in different contexts. As define in Wikipedia (), “Middleware is now used to describe database management systems, Web servers, application servers, content management systems, and similar tools that support the application development and delivery process.” In this article, middleware refers to the bundle of IBM HTTP Server, Application Server, and DB2 UDB together.

A lot of IBM middleware has been enhanced in the past few years to exploit the ARM open standard as a way to provide internal transaction statistics to any external tool supporting this standard. Application Server (Express, Base, ND, XD) and DB2 UDB have been ARM-instrumented, as well as the WebSphere plug-in for IBM HTTP Server (IHS) and Microsoft Internet Information Server (IIS). Using this middleware, you can monitor your Web application's performance from end to end. However, ARM instrumentation is disabled by default. This chapter will explain how you can enable ARM instrumentation.

You need the EWLM Managed Server (MS), an agent that runs inside each server to collect the ARM statistics and send them to the EWLM Domain Manager (DM). The DM aggregates all the data and presents it through the EWLM Control Center, which you can access from any Web browser. Figure A-1 on page 387 illustrates the high-level architecture.

Archive

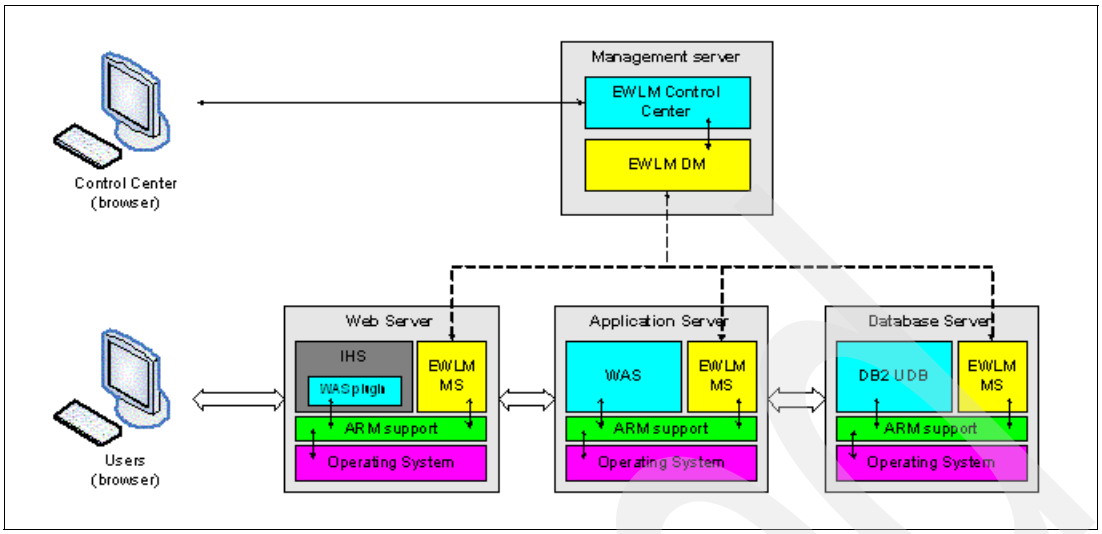


Figure A-1 High-level architecture of EWLM and middleware support

The overall process includes the following steps:

1. Installing EWLM DM on the management server.
2. Installing EWLM MS on each machine running the middleware.
3. Enabling ARM platform support in the operating system of each machine.
4. Creating an EWLM DM configuration ID.
5. Creating an EWLM MS configuration ID.
6. Creating an EWLM domain policy.
7. Enabling ARM instrumentation in DB2 UDB.
8. Enabling ARM instrumentation in the Application Server.
9. Enabling ARM instrumentation in IHS.
10. Using the EWLM Control Center to see the transactions statistics.

Prerequisites

EWLM supports IHS, Application Server, and DB2 on many platforms. Refer to the EWLM Information Center (refer to Resources) for the latest information. ARM instrumentation of the middleware is identical on most platforms (except i5/OS and z/OS), but the ARM platform support layer and related security procedures are specific to each operating system. For convenience, this appendix mainly discusses AIX, with some mention of Linux-specific items when necessary.

The software discussed in this article includes:

- AIX 5L V5.3 ML1
- SLES 9.0 SP3
- IBM Enterprise Workload Manager V2R1
- IBM HTTP Server V6.0.2
- WebSphere Application Server V6.0 Network Deployment
- DB2 UDB V8.2 Enterprise Edition

You should have practical knowledge of the middleware and the operating systems. You should have already installed the EWLM DM on the management server and EWLM MS on each node where the middleware is running. Both could run in the same system, provided you have enough processing power and memory.

Archive

The test environment consists of four logical partitions on an IBM System p 570 box. Each partition is configured with the desired capacity of one processor (POWER5 Micro-Partitioning enabled) and 2 GB of memory. You can do this exercise on any IBM System p stand-alone box or IBM System i partition, as long as you can run AIX 5L V5.3 on them. In our scenario, we have AIX and Linux running in 64-bit mode, and the same for DB2 UDB. EWLM is installed as part of the IBM Virtualization Engine CD stack, which we will not discuss in this chapter.

Enabling ARM platform support

This section explains how to enable ARM platform support for AIX and Linux.

For AIX

As shown in Figure A-1 on page 387, neither the middleware or any ARM-instrumented application really interacts with EWLM MS directly. They send data to the ARM platform support and MS, in turn, fetches the data from the ARM platform support and sends it over to DM. Therefore, you must start with enabling ARM support in AIX.

Ensure that the ARM support package, named bos.net.ewlm.rte, is installed. Use the **# lslpp -l bos.net.ewlm.rte** command to query if it is installed. If it is not already installed, you will have to get it from the AIX installation CD. You can use System Management Interface Tool (SMIT) to install the package.

After confirming that the package is installed, you can enable ARM support as follows:

1. Run the **smitty ewlm** command.
2. Change or show the status of the EWLM services.
3. Enable the EWLM services.
4. When prompted to enable EWLM services now, at next boot, or both, press F4 to choose both. Press Enter to confirm.

Run the **ewlmcfg -q** command to confirm that EWLM support is enabled. You should get the message shown in Example A-1.

Example: A-1 EWLM enablement message

```
EWLM is enabled.
No process is using the EWLM services.
```

ARM support in AIX should now be enabled. When starting EWLM MS, if you get a message saying that the platform support is not enabled, run the **ewlmcfg -q** command again to verify. Look for the following line in /etc/inittab:

```
rcnetwlm:23456789:wait:/etc/rc.netwlm start> /dev/console 2>&1 # Start netwlm
```

Make sure this line is not commented out and that it is being invoked properly. We have seen problems where AIX was initializing another service before it and that service never exited. Therefore, AIX never got to the point of running this service. If in doubt, you can try to initialize it manually. This does not cause any harm if it was already initialized. Run the **/etc/rc.netwlm start** command, which may yield the message shown in Example A-2 on page 389.

Example: A-2 Message when problems initializing the service

```
/usr/lib/drivers/netwlm already loaded
Error loading kernel extension: cfgnetwlm failed
```

For Linux

Before running EWLM MS on Linux, refer to Chapter 5, “Applying fixes” on page 189 for additional patches to apply to the operating system. If you are compiling your own kernel, ensure that you do not disable the Class-Based Kernel Resource Management (CKRM) feature in the kernel. The ARM platform support relies on CKRM to get statistics from the operating system. You also need to install Fix Pack 020 on top of the Linux MS. Linux MS was deactivated in the GA version.

The ARM platform support in Linux is shipped in the EWLM MS package and is installed as part of MS installation. Normally, you do not have to do anything special for that. The ARM daemon should be started at system startup. You can verify that it started using the **`ps aux | grep ewlmarmd`** command. If it is not running, verify the existence of the /etc/init.d/ewlm file and make sure the daemon is not being disabled during the startup sequence. Use your favorite platform tool for verification, such as YaST2 or **`chkconfig`**.

On the Linux platform, EWLM MS uses a file system mounted at /mnt/relayfs, so be careful not to mount anything at /mnt. Create a subdirectory under /mnt and use that instead.

Another important directory that you should not delete is /tmp/ewlmarmd. Every time there is an application making ARM registration, some related files are created in it. If it is deleted by accident, you can simply restart the ARM daemon, which will recreate it. However, you will have to restart all the ARM applications as well because all the registrations will be lost. Use the **`/etc/init.d/ewlm`** restart to restart.

Creating an EWLM DM configuration ID

Before you can use EWLM DM, you must create a configuration ID similar to an Application Server profile or a DB2 instance. You can use the command line or the graphical user interface (GUI) wizard to do so. Creating the configuration ID creates a working directory for this configuration ID, assigns several TCP/IP ports to the DM, and creates an Application Server profile to run the Control Center application.

Using the configuration wizard

Make sure you have a working X Window System environment. You should do this, as root, using the **`# /opt/IBM/VE2/EWLM/bin/configWizardDM.sh`** command.

When the screen opens, select **Create domain manager**, then select **Control Center** and click **Next**. Then do the following steps:

1. At step 1 of 8 of the configuration, the wizard asks for the configuration ID name. Do not use any special characters. We used chekimDM. The wizard creates the /opt/IBM/VE2/EWLM/ewlmdm_instances/servers/chekimDM directory.
2. In the next step, the domain name is only for display purposes. You can use long words and spaces. The IP address could be your actual address (you could put 0.0.0.0 as a way to tell DM to listen to all IP addresses on the system). You can keep the default port assignments or set your own, as shown in Figure A-2 on page 390.

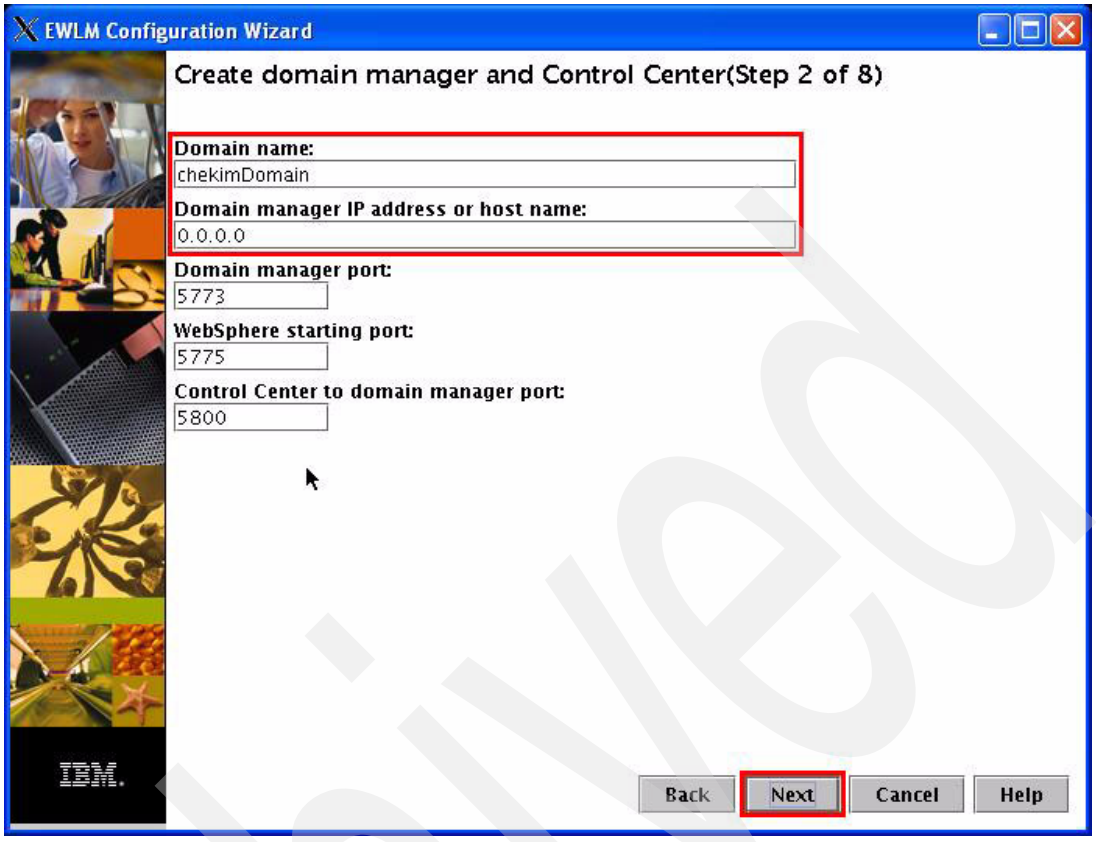


Figure A-2 DM configuration wizard, IP address, and listening ports

3. In the next step, it is not necessary to enable Secure Sockets Layer (SSL) encryption between DM and MS for an intranet environment.
4. Step 4 asks if you want to replace the default SSL certificate provided by WebSphere, which is running the Control Center application, with Global Security enabled. Again, in a controlled environment, it is not really necessary. Leave the EWLM default choice and select **Next**.
5. In step 5, do not enable the firewall broker unless you are completely sure it is necessary.
6. Step 6 is for the user ID that will be used to stop and manage WebSphere for the Control Center. It has to be root, because Global Security is using local operating system (OS) authentication.
7. For step 7, you can assign operating system users or groups to different roles for the Control Center operations, as shown in Figure A-3 on page 391. You have to assign at least one user or group to the Administrator role to use EWLM. However, you can add or remove any assignment later on if you want.

Archived

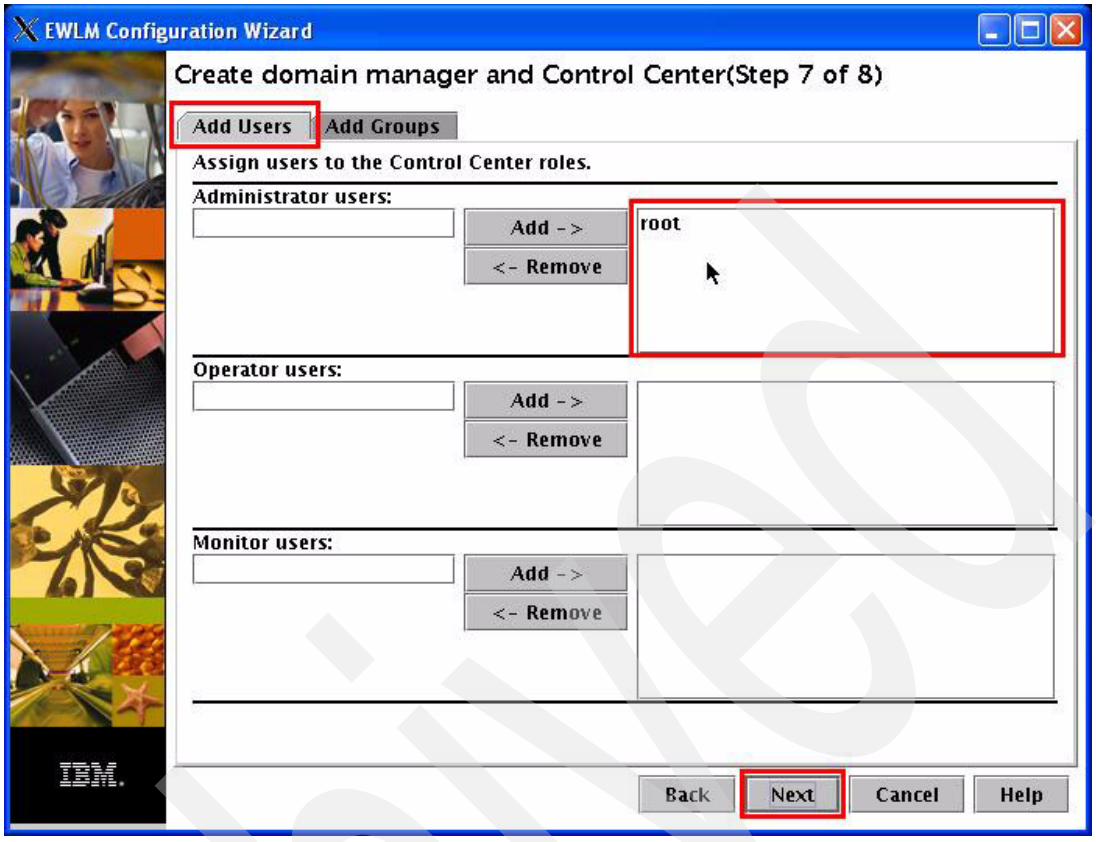


Figure A-3 DM configuration wizard, adding role for the Control Center

8. At step 8, leave the box unchecked. Creating the DM configuration ID is a lengthy process, so we prefer to get it done properly, and then start DM manually afterward.

After about 10 minutes, the configuration is complete, as shown in Figure A-4 on page 392. Notice the Control Center we address at the bottom of the window. You should substitute your actual IP address or host name in this Web address.

Archived

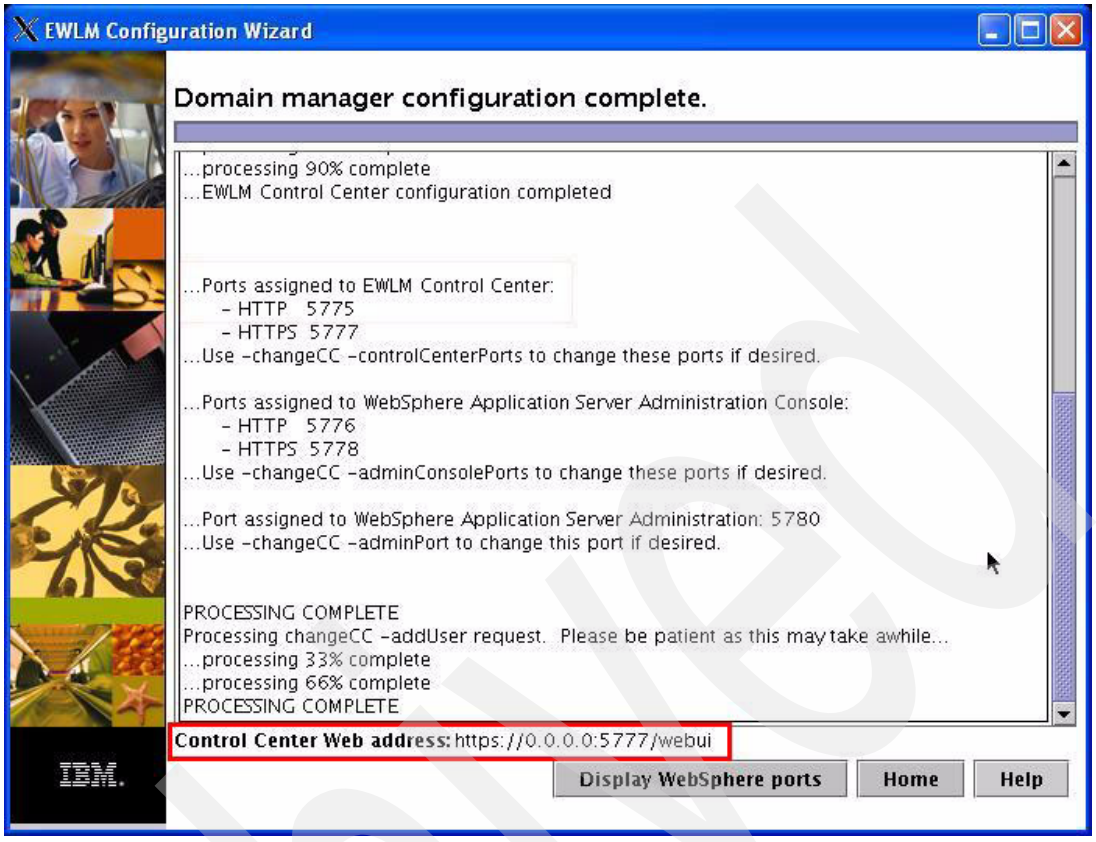


Figure A-4 DM configuration wizard, creation process completed

Using the command line

You can perform this task with the command line as well. This is especially useful if you do not have root authority and you have to provide the command to be run by another person who has root authority. These commands do exactly what the configuration wizard did.

Example A-3 is the command-line version of configuration wizard tasks.

Example: A-3 Command-line versions of configuration wizard tasks

```
# cd /opt/IBM/VE2/EWLM/bin
# ./createMS.sh chekimDM -ma 0.0.0.0 -mp 5775 -wasPorts 5775 -jp 5800 -dn
chekimDomain
-auth None -adminUser root -adminPW password

# ./changeCC.sh -addUser chekimDM -adminUser root -adminPW password -roleUser root
-role
Administrator
```

To start the DM and Control Center, use the commands listed in Example A-4.

Example: A-4 Command-line versions to start DM and Control Center

```
# cd /opt/IBM/VE2/EWLM/bin
# ./startDM.sh chekimDM
# ./startWAS.sh
chekimDM
```

Now, log in to the Control Center to see if everything is working properly. Open the Internet Explorer® and type the following Web address:

```
https://dm_hostname:5777/webui
```

Log in with the user ID and password that you provided as the Administrator role. After a successful login, you should see the window shown in Figure A-5. Make sure you allow pop-ups for this Web page.



Figure A-5 Control Center home page

Troubleshooting

If you have errors and the wizard does not complete, you need to know where and why it is failing, and you can find this information in the execution logs. The following three log files are created during the process in two directories:

- /var/adm/ewlm/ewlmdm/Diagnostics/createDM_timestamp.log
- /var/adm/ewlm/ewlmdm/Diagnostics/wizardDM_timestamp.log
- /var/adm/ewlm/ewlmdm/ConfigWizardDiagnostics/ConfigurationWizard.log

If there are errors during the WebSphere profile creation process, you should be able to determine the issue from the createDM log file. If you have port conflict issues, try to stop all the running Application Server instances and ensure that the ports are available for use.

If the failure occurs after 90 percent of the process, you are most likely having a problem with the Lightweight Directory Access Protocol (LDAP) server. By default, after the WebSphere profile is created and the Control Center application deployed, it tries to register your

newly-created DM configuration ID into the LDAP server that was specified during Virtualization Engine installation. If a problem occurs during this process, your newly-created DM is dropped so that you can recreate it again. It only keeps the log files to tell you what went wrong.

The LDAP registration is only used by the Virtualization Engine console. If you do not use the Virtualization Engine console in your environment and thought that LDAP was not required anymore, you do require the LDAP server for creating and deleting DM. If for some reason you can no longer get your LDAP server back, there is a workaround to skip the LDAP registration during DM creation and removal. Edit the /opt/IBM/VE2/EWLM/bin/ewlmdm_environment.sh file and comment out the two lines, as shown in Example A-5.

Example: A-5 Commenting out definitions

```
# VE_BASE=/opt/IBM/VE2
# VE_LIB=/opt/IBM/VE2/lib
```

Rerun the wizard or command line. It will not try to do any LDAP registration. (Do not do this if you are planning to use the Virtualization Engine console in the future or you will not be able to hook up to this DM from the Virtualization Engine console bridge.)

If you want to delete a DM configuration ID, you must use the **deleteDM.sh** command or the wizard. Do not simply delete the working directory. There is lots of cleanup work that needs to be done.

Creating an EWLM MS configuration ID

You must create a configuration ID on each EWLM MS node using either the wizard or the command line. This process is much shorter than creating a DM configuration ID.

We ran the MS process as root. You could run MS as a non-root user. Refer to the Resources section in the Information Center for information about how you can do this. This section discusses the configuration wizard, command line, and troubleshooting.

Using the configuration wizard

Run the configuration wizard as root using the # **/opt/IBM/VE2/EWLMMS/bin/configWizardMS.sh** command.

At the startup window, select **Create managed server** and click **Next**. Then perform the following steps:

1. For step 1 of 4, the configuration ID should not contain any special characters. Our configuration ID is chekimMS. A working directory is created in /ewlmData/ewlmms/servers/chekimMS.
2. At step 2, you should provide the DM's host name or IP address. Do not forget to change the port if you assigned a different one earlier during the DM creation.
3. Step 3 is for SSL encryption and must match the setting in the DM. Select **None** and then click **Next**.
4. At step 4, if you have no direct communication between DM and MS, you may use a standard HTTP proxy to bridge them, making sure the CONNECT method is enabled. Otherwise, just select **Next** to start the MS creation.

Ensure that you wait until the message `PROCESSING COMPLETE` appears.

Using the command line

To perform all the wizard steps with one command, run the following command on all the MS nodes:

```
# /opt/IBM/VE2/EWLMMS/bin/createMS.sh chekimMS -ma hci246 -mp 5773 -auth None
```

To start the MS, enter the following command:

```
# /opt/IBM/VE2/EWLMMS/bin/startMS.sh chekimMS
```

If you go to the Control Center and select **Manage** → **Managed servers**, you will see the new MS joining the domain. If you do not see it right away, click the refresh icon at the top right of the window a few times.

Troubleshooting

The MS creation process is quite straightforward. However, you might experience some problems when you start the MS. It may not start successfully if the platform support is not set up properly, in which case refer to , “Enabling ARM platform support” on page 388.

Another potential problem you might have eventually is with the file system space. By default, the MS configuration ID is created at the root file system under the /ewlmData directory. AIX allocates very little space to this file system by default. If your MS is running into problems and starts dumping XML diagnostics files, it might quickly fill up the root file system. As a remedy, you could increase the root file system size or move the /ewlmData's content to a larger file system and create a symlink at /ewlmData pointing to the new place.

Creating an EWLM domain policy

EWLM is shipped with sample policies to get you started. However, we will create a new one in this section.You should specify how to classify incoming requests into transaction classes that are subsets of a service class in the policy, and specify a performance objective for the service class. We are going to use a basic Web application called *Trade* that simulates a simple trading Web application in which users can buy or sell stocks and query their portfolio. If you do not have access to this application, you can use any other sample Web application for this process and perform the following steps:

1. To begin the domain policy creation, log in to the Control Center again, select **Domain policies**, and then select **New**.

   The domain policy wizard appears. Give the policy a name and a short description.

2. At the next step, you can keep the default EWLM workload and select **Next**.

3. In the Service classes step, create a new service class by clicking **New**.

   You will see the details of the service class's performance objective. All the fields are self-explanatory, as shown in Figure A-6 on page 396. After filling in all the values, select **OK** and then **Next**.

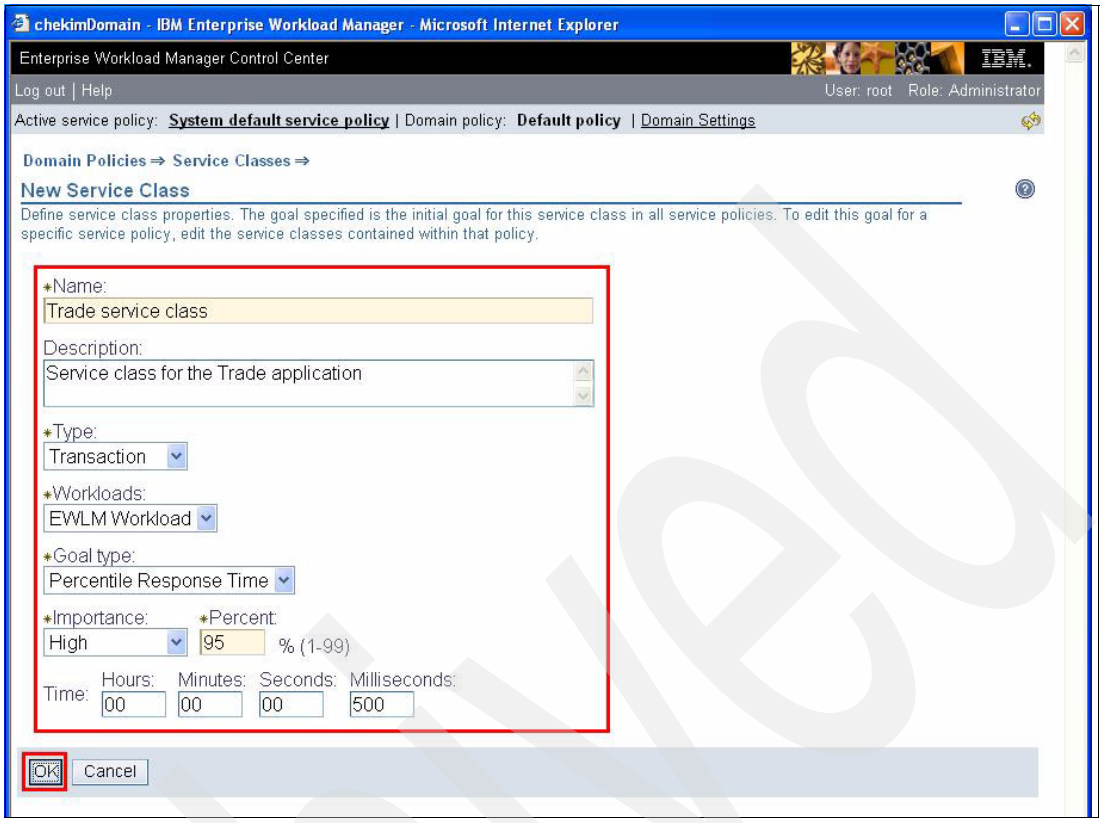


Figure A-6 Domain policy wizard: creating a service class

4. Add an application to the policy to classify requests coming into this application by selecting **Add** in the Applications step. Remember, only the entry application matters. This is where each incoming request is classified and the classification is carried on to all subsequent middleware hops. Each application here has a certain number of predefined filters that are pertinent to it.

   For the Trade application we used in our setup, all the incoming requests hit the IHS first, so we added the IBM Webserving Plugin application to our policy.

5. To create a transaction class, make sure the list is on the IBM Webserving Plugin entry and select **New**. Give it a name and assign it to position 1 of Trade service class, as shown in Figure A-7 on page 397. Remember that the order is very important in the classification rules. You should always have the most specific rule first, followed by the more generic ones.

Archived

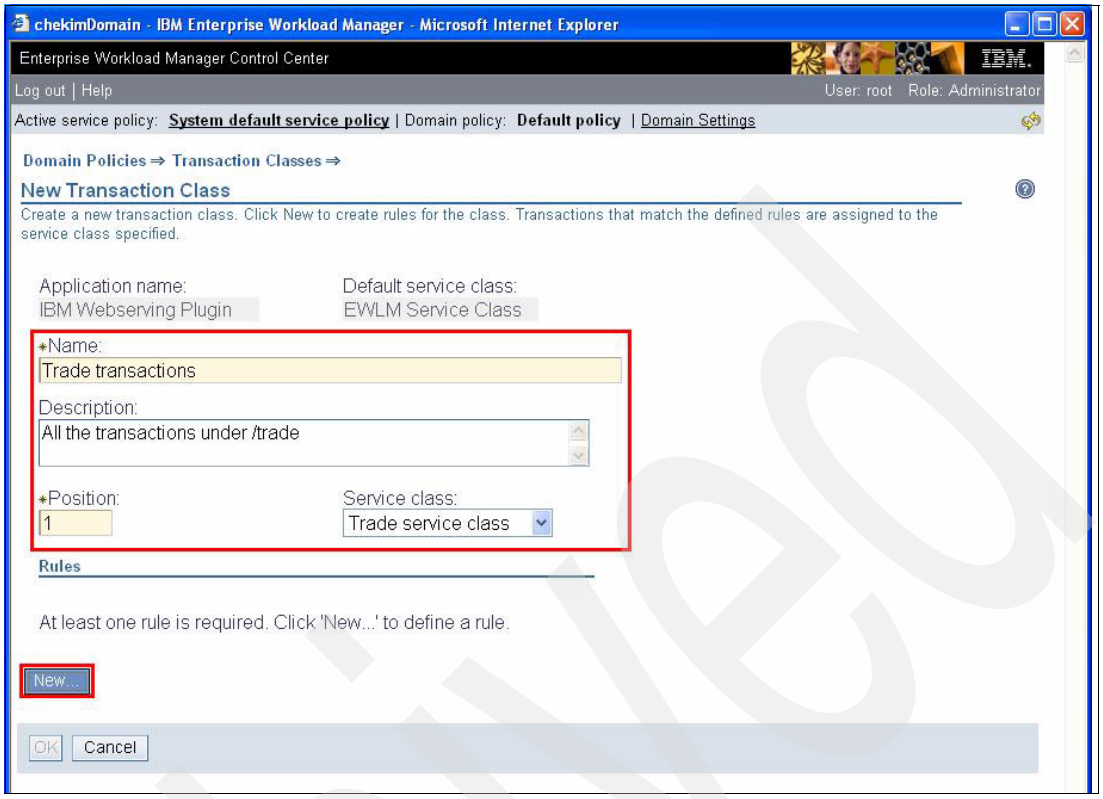


Figure A-7 Domain policy wizard: creating a transaction class

6. You should add a rule to this transaction class by selecting **New**. Our classification rule is based on the EWLM:URI filter and the value is /trade(*) without any quotes, as shown in Figure A-8 on page 398. Any incoming request starting with /trade is classified in this transaction class. You can create very complex rules by grouping several individual rules with logical operators. However, in this scenario, we keep it very simple.

Archived

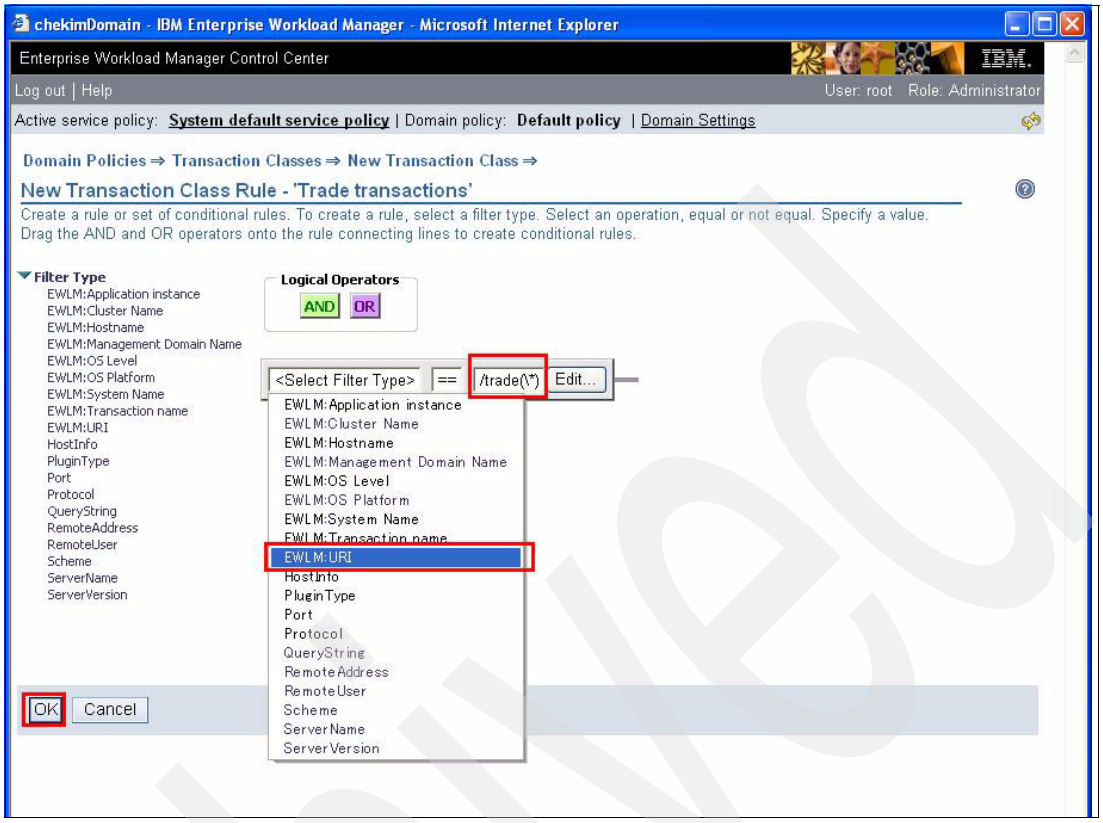


Figure A-8 Domain policy wizard, creating a rule

7. When you have finished adding the rule, make sure that the new transaction class is positioned before the default one that will catch everything else coming into this application. Select **Finish** to save the policy.
8. After the policy is saved, you must deploy it to actually make the policy effective. Select the newly-created policy, and then select **Deploy** in the list. Then select **Go**.
9. Next, you will be prompted to select the service policy to activate. Because you did not create any new service policies, there is only one that you can pick to activate. Select it and then select **Deploy**. You can also actually deploy a domain policy without activating any service policy. However, this is not necessary.

Make sure the service policy is successfully activated on all the managed servers, as shown in Figure A-9 on page 399.

Archived

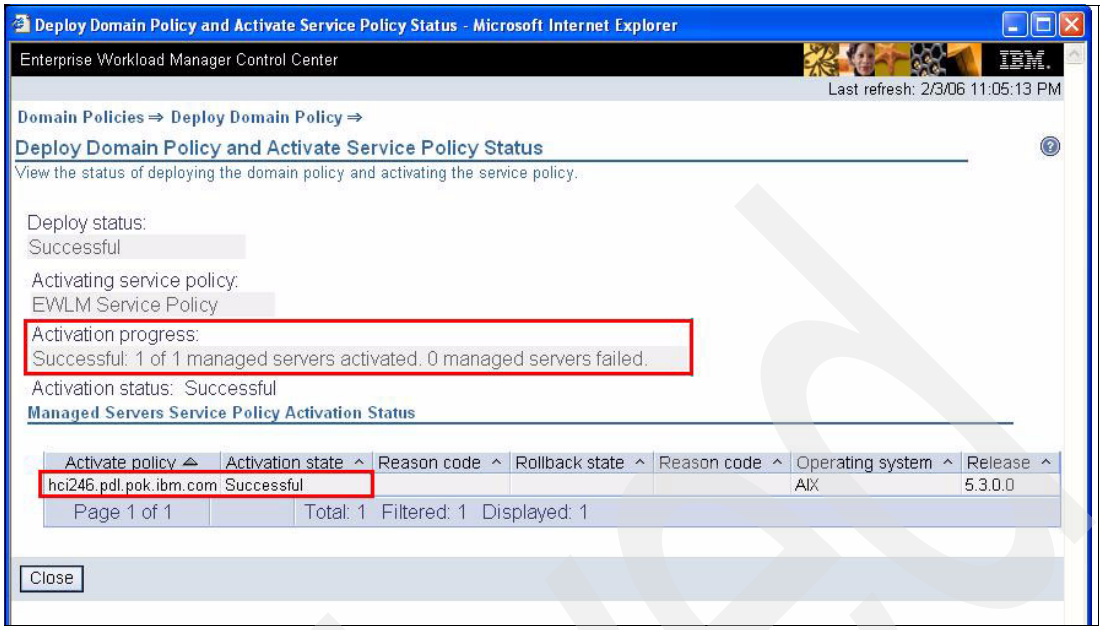


Figure A-9 Status of service policy activation

At any time, you can tell which policy is active by looking at the top of the window. You can also see the service policy that is active on any managed server. Normally, all the managed servers should have the same active service policy. However, sometimes a server might be down or disconnected while a domain policy is being deployed. Therefore, there is a discrepancy. As soon as the server is connected back to the domain manager, it attempts to activate the current policy.

You are now ready to receive transaction statistics from the middleware. However, you will need to enable the ARM instrumentation in your middleware first.

Enabling ARM instrumentation in DB2 UDB

Beginning with DB2 UDB V8.2, ARM instrumentation is already turned on by default. However, you have to tell it where to look for the ARM library libarm4.a (or libarm4.so for Linux). For the db2inst1 instance, use the following commands:

```
# su - db2inst1
$ db2set DB2LIBPATH=/usr/lib
```

Before you start DB2, you should give the db2inst1 user some special privileges on AIX to register itself to ARM platform support. You have to be root to do this task and use the following command:

```
# chuser "capabilities=CAP_ARM_APPLICATION,CAP_PROPAGATE" db2inst1
```

On Linux, you simply have to add the db2inst1 user to the ewlm_armusers group (created during MS installation). In both cases, you may have to log out and log back in to make the change effective. Then you can start the DB2 instance.

For Linux operating systems, only DB2 for Linux on Intel x86 platforms has ARM instrumentation support. Linux on all other hardware platforms does not have ARM support yet, so it is not useful to do this. Future versions of DB2 UDB will have ARM support on all Linux hardware platforms.

Now you can start the DB2 instance. Example A-6 shows if the DB2 instance has successfully registered to ARM.

Example: A-6 Starting DB2 instance

```
# lsarm -a
APPL: IBM DB2 Universal Database
```

On Linux, this command is found at /opt/IBM/VE2/EWLMMS/usr/bin/lsarm.

Troubleshooting

The most common DB2 ARM instrumentation problem is caused by unauthorized registration, either because it lacks the required privilege or the privilege was not yet effective. To double check that on AIX, run the **# lsuser db2inst1** command, as shown in Example A-7.

Example: A-7 Checking DB2 ARM instrumentation registration

```
# lsuser db2inst1
db2inst1 id=108 pgrp=db2grp1 groups=db2grp1,staff,dasadm1 home=/home/db2inst1
shell=/usr/bin/ksh login=true su=true rlogin=true daemon=true admin=false
sugroups=ALL
admgroups= tpath=nosak ttys=ALL expires=0 auth1=SYSTEM auth2=NONE umask=22
registry=files SYSTEM=compat logintimes= loginretries=0 pwdwarntime=0
account_locked=
false minage=0 maxage=0 maxexpired=-1 minalpha=0 minother=0 mindiff=0 maxrepeats=8
minlen=0 histexpire=0 histsize=0 pwdchecks= dictionlist= capabilities=
CAP_ARM_APPLICATION,CAP_PROPAGATE fsize=-1 cpu=-1 data=491519 stack=32767
core=-1 rss=-1 nofiles=2000 time_last_login=1139331346 tty_last_login=
host_last_login=
hci246 unsuccessful_login_count=0 roles=
```

On Linux, run the **id db2inst1** command and make sure it is in the right group. Another possible problem on Linux is that, if for any reason the /tmp/ewlmarm4 directory is not accessible to **db2inst1** because it has the wrong directory permission or may have been deleted by accident during cleanup activities, then lsarm will not show any application.

A more serious problem could be that DB2 somehow cannot find or load the ARM library. In that case, you should capture some traces with the db2trc tool. If you are not familiar with this tool, refer the DB2 Command Reference at:

http://publib.boulder.ibm.com/infocenter/db2luw/v8//index.jsp

Example A-8 on page 401 illustrates a quick trick that shows you how far DB2 has gone in loading the ARM library.

Archived

Example: A-8 Determining how far DB2 has gone in loading the ARM library

```
# su - db2inst1
$ db2trc on -f 1.trc                    (turn on tracing)
$ db2start                              (start DB2)
$ db2trc off                            (turn off tracing)
$ db2trc format 1.trc 1.txt             (format binary trace file to text)

$ cd sqllib/db2dump
$ grep -i arm 1.txt
rc = 0x800F00FC = -2146500356 = SQLO_INV_PARM
2948    entry DB2 UDB base sys utilities sqleLoadArmFunctions fnc (1.3.5.1132.0)
2949    data DB2 UDB base sys utilities sqleLoadArmFunctions fnc (3.3.5.1132.0.1)
        libarm4.a
6C69 6261 726D 342E 61                    libarm4.a
        6172 6D5F 7265 6769 7374 6572 5F61 7070    arm_register_app
        6172 6D5F 7374 6172 745F 6170 706C 6963    arm_start_applic
        6172 6D5F 7265 6769 7374 6572 5F74 7261    arm_register_tra
        6172 6D5F 7374 6172 745F 7472 616E 7361    arm_start_transa
        6172 6D5F 6269 6E64 5F74 6872 6561 64      arm_bind_thread
        6172 6D5F 626C 6F63 6B5F 7472 616E 7361    arm_block_transa
        6172 6D5F 756E 626C 6F63 6B5F 7472 616E    arm_unblock_tran
        6172 6D5F 756E 6269 6E64 5F74 6872 6561    arm_unbind_threa
        6172 6D5F 7374 6F70 5F74 7261 6E73 6163    arm_stop_transac
        6172 6D5F 6465 7374 726F 795F 6170 706C    arm_destroy_appl
        6172 6D5F 6765 745F 6572 726F 725F 6D65    arm_get_error_me
        6172 6D5F 7374 6F70 5F61 7070 6C69 6361    arm_stop_applica
        6172 6D5F 6765 745F 636F 7272 656C 6174    arm_get_correlat
2999    data DB2 UDB base sys utilities sqleLoadArmFunctions fnc (3.3.5.1132.0.10)
3000    exit DB2 UDB base sys utilities sqleLoadArmFunctions fnc (2.3.5.1132.0)
11569   entry DB2 UDB base sys utilities sqleArmRegisterAgent fnc (1.3.5.1135.0)
11578   data DB2 UDB base sys utilities sqleArmRegisterAgent fnc (3.3.5.1135.0.10)
11579   data DB2 UDB base sys utilities sqleArmRegisterAgent fnc (3.3.5.1135.0.20)
11580   data DB2 UDB base sys utilities sqleArmRegisterAgent fnc (3.3.5.1135.0.30)
11581   exit DB2 UDB base sys utilities sqleArmRegisterAgent fnc (2.3.5.1135.0)
11880   entry DB2 UDB base sys utilities sqleArmRegisterAgent fnc (1.3.5.1135.0)
11883   data DB2 UDB base sys utilities sqleArmRegisterAgent fnc (3.3.5.1135.0.10)
11884   data DB2 UDB base sys utilities sqleArmRegisterAgent fnc (3.3.5.1135.0.20)
11885   data DB2 UDB base sys utilities sqleArmRegisterAgent fnc (3.3.5.1135.0.30)
11886   exit DB2 UDB base sys utilities sqleArmRegisterAgent fnc (2.3.5.1135.0)
```

The result of **grep** in Example A-8 is when DB2 has successfully loaded all the ARM functions from libarm4.a. If you see this in your trace file but your DB2 instance still does not show up in lsarm, then you know that it is not related to ARM library loading. Most likely it is a user-privilege problem or ARM-platform-support problem.

If DB2 could not find libarm4.a, you will only see the first five rows from Listing 5 and not the other ARM functions in the trace file. If you have to call technical support for this kind of issue, providing them with such a trace file is very helpful.

Enabling ARM instrumentation in Application Server

Enabling ARM in Application Server is a two-step process:

1. First, you should enable Request Metrics globally.
2. Then, inside each AppServer's setting, you should tell it where to look for the ARM implementation class and Java Native Interface (JNI) library.

You can, therefore, selectively enable ARM instrumentation in each AppServer of your WebSphere cell. With WebSphere Application Server Express, Base, and Network Deployment editions, ARM instrumentation is applicable to AppServers only. With WebSphere Application Server XD edition, you can apply that to Proxy Servers and On Demand Routers as well.

To enable Request Metrics at the cell scope, perform the following steps:

1. Log in to the WebSphere Administrative Console, go to the Request Metrics window, and use the following settings (shown in Figure A-10 on page 403):
   - Check Enable Request metrics.
   - Components to be instrumented = ALL.
   - Trace level = Hops.
   - Check Application Response Measurement(ARM) agent.
   - Agent Type = ARM4.
   - ARM transaction factory implementation class name = com.ibm.wlm.arm40SDK.transaction.Arm40TransactionFactory.

Archived

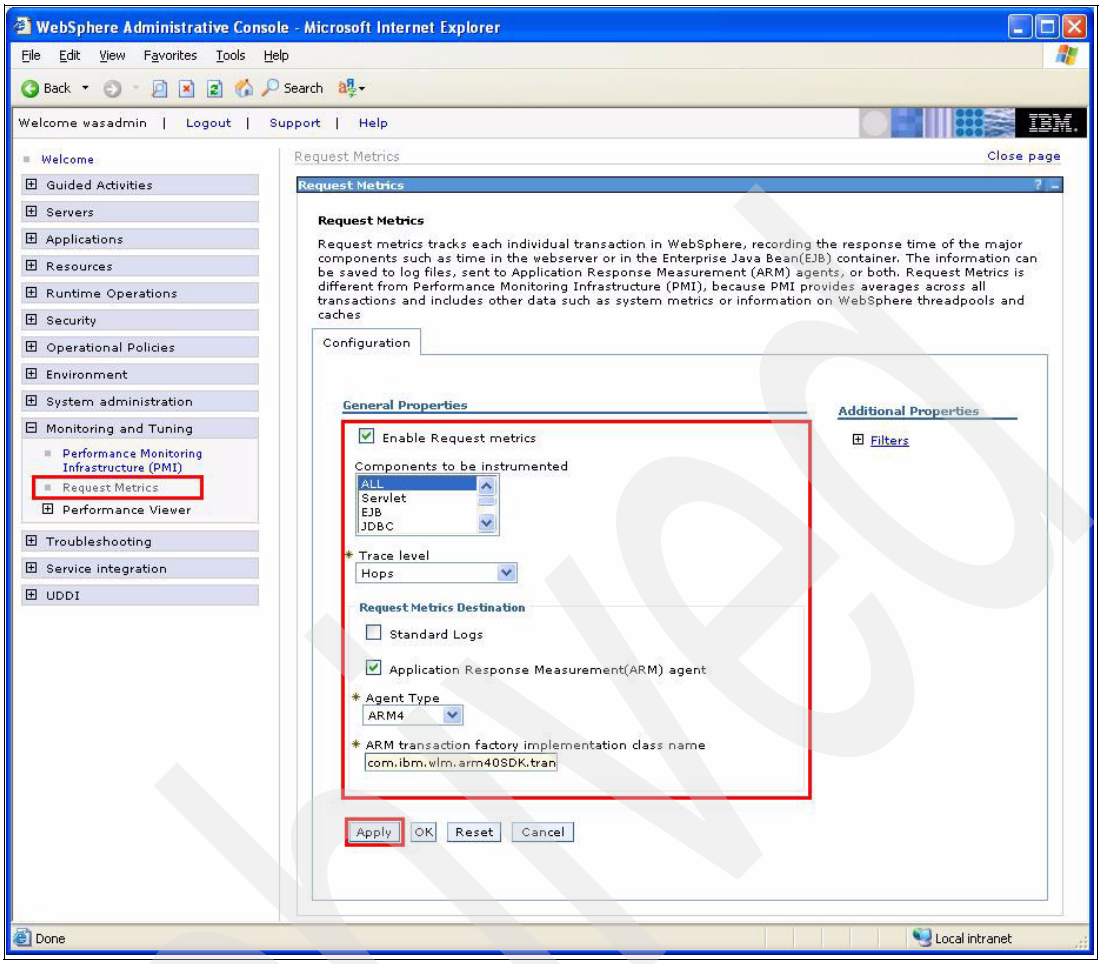


Figure A-10 Enabling Request Metrics in WebSphere AdminConsole

2. Go into each AppServer (or other types of server in the case of WebSphere Application Server XD) to add custom properties to the Java Virtual Machine (JVM™). As shown in Figure A-11 on page 404, select **Application servers** → **yourAppServer** → **Process Definition** → **Java Virtual Machine** → **Custom Properties**, and create the two following properties:
   - ws.ext.dirs = /opt/IBM/VE2/EWLMMS/classes/ARM
   - com.ibm.websphere.pmi.reqmetrics.PassCorrelatorToDB = true
3. The PassCorrelatorToDB property should be set to true only if you have an ARM-instrumented DB2 instance at the next hop. Otherwise, it is preferable to set it to false or just not set it to avoid some minor overhead. If your WebSphere application is connecting to other types of database servers, or if your back-end DB2 instance is an older version, you can skip this property. Note that only the new DB2 Universal Java Database Connectivity (JDBC) Driver (db2jcc) supports flowing ARM correlator to downstream hop. If you are using another JDBC provider, the ARM correlator will not flow to the next hop regardless of the database server.
4. You only need to set the java.library.path property (java.library.path = /usr/lpp/ewlm/java/lib_64.) if your Application Server is a 64-bit edition. Otherwise, Application Server is able to find the EWLM JNI library in /usr/lib by itself. On Linux, the 64-bit JNI library is in the /usr/lib64 directory.

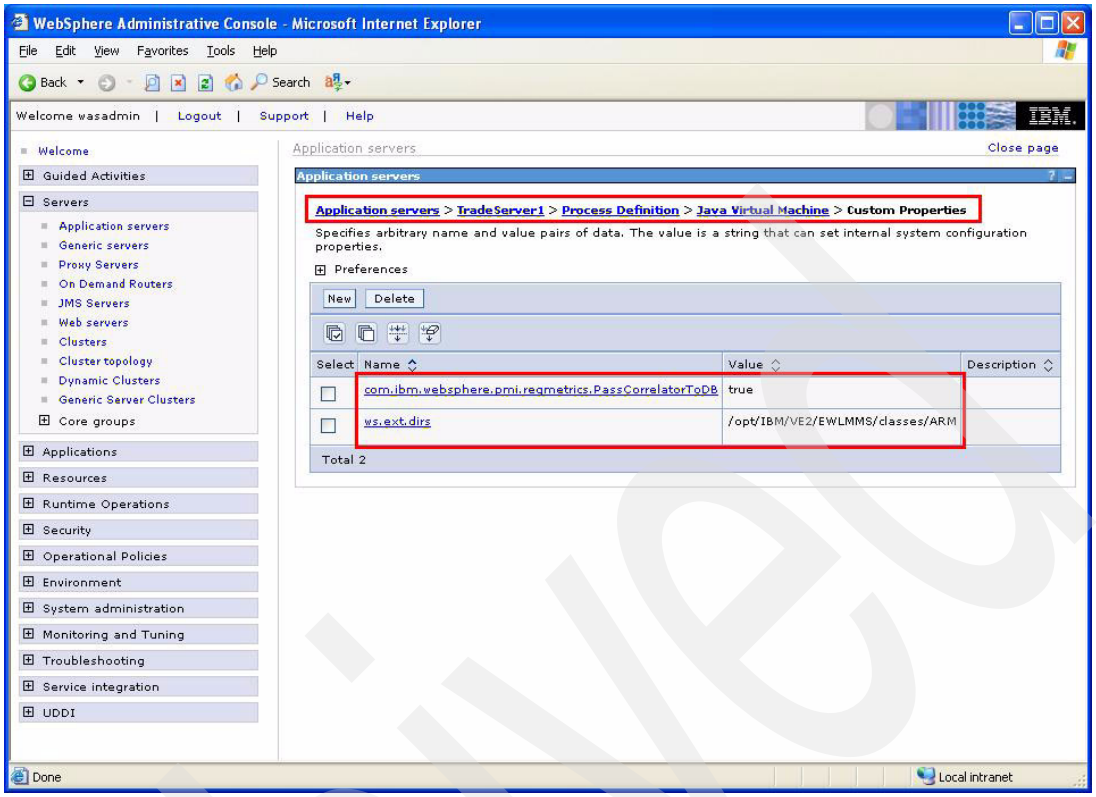


Figure A-11 JVM custom properties for ARM instrumentation

5. Save the changes and restart the AppServers to make the ARM instrumentation effective. Verify if ARM registration is successful by using the **`# lsarm -a`** command. You should see the following message:

   ```
   APPL: WebSphere:APPLICATION_SERVER
   ```

Although this procedure is fairly straightforward for a handful of AppServers, it could become tedious if you have a few dozen on different platforms to ARM instrument, because each platform might need a different directory path. For this purpose, refer to the wsadmin script called was_arm.jacl to automate the process. This is available at on the Web:

`http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-arm/#download`

The script will interactively loop through a list of servers (AppServer, ODR, and so on) in a WebSphere ND or XD cell to add JVM custom properties for ARM instrumentation. It has the default directory paths for all the EWLM supported platforms except i5/OS and z/OS. You are free to use or modify it to fit your own environment. We recommend that you become familiar with this procedure before using the script. Otherwise, you may not know exactly what has been done to your WebSphere environment.

If you are running Application Server as a non-root user, then the same security rule applies as with DB2. On AIX, you should add some privileges to the user as follows:

```
# chuser "capabilities=CAP_ARM_APPLICATION,CAP_PROPAGATE" userID
```

On Linux, you should add that user into the ewlm_armusers group and log in again to make it effective.

If, for any reason, you want to temporarily disable ARM instrumentation in your entire WebSphere cell, you can simply uncheck Enable Request metrics on the Request Metrics

window and leave those JVM custom properties alone. If you want to disable ARM permanently, it is better to remove those properties.

Troubleshooting

If you followed all the steps and do not see a WebSphere in lsarm, perform the following steps:

1. Check and see if the user running WebSphere has the required privilege to register to ARM.
2. Check whether WebSphere has been able to load the ARM implementation class provided by EWLM in the arm4.jarfile. You determine this from the SystemOut.log during Application Server startup. If there is no error, then the ARM implementation class was found and loaded. If there are any errors, you should double-check the ws.ext.dirs Java virtual machine (JVM) custom property.

   When all is well, you should see the results in the SystemOut.log, as shown in Example A-9.

Example: A-9 SystemOut.log results when ARM implementation class is found

```
0000000a PmiRmArmWrapp I   PMRM0001I: ARM implementation is:
com.ibm.ws.pmi.reqmetrics.PmiRmArmXfcArm4Impl
0000000a PmiRmArmWrapp I   PMRM0004I: PmiRmInit
servername=6.0.0.0/hcp121Cell01/hcp121Node02/TradeServer1/APPLICATION_SERVER
```

Example A-10 shows the results when ARM implementation class is not found.

Example: A-10 SystemOut.log results when ARM implementation class is not found

```
0000000a PmiRmArmWrapp I   PMRM0001I: ARM implementation is:
com.ibm.ws.pmi.reqmetrics.PmiRmArmXfcArm4Impl
0000000a PmiRmArmWrapp W   PMRM0002E: Failed to instantiate:
com.ibm.ws.pmi.reqmetrics.PmiRmArmXfcArm4Impl
```

You will also see the exception java.lang.NoClassDefFoundError in the SystemErr.log.

Sometimes the problem is less obvious to detect. For example, if the libewljarm4.a JNI library is not found, you might not see any error message in the logs or you might see the exception java.lang.UnsatisfiedLinkError in the SystemErr.log, depending on the circumstance of the failure. For those cases where no obvious error is seen, you should turn on PMI tracing with pmi.rm=all, as shown in Figure A-12 on page 406.

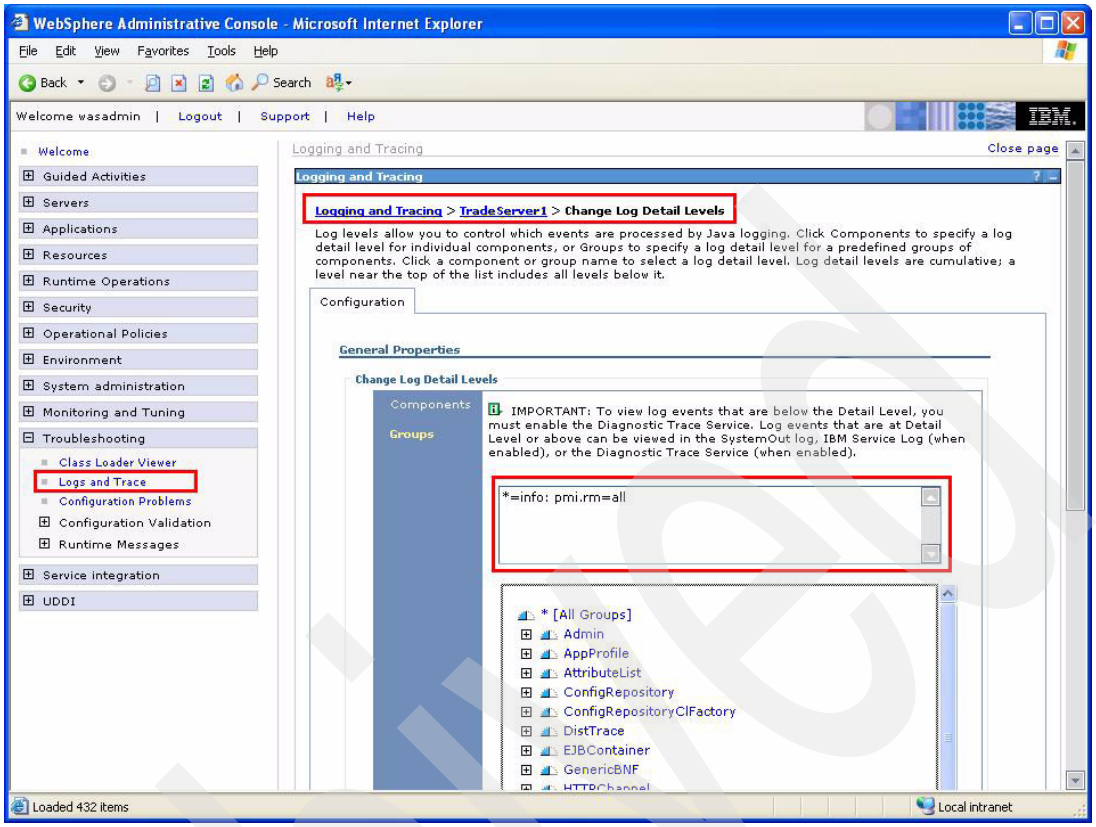


Figure A-12 Using pmi.rm=all to trace debugging for ARM enablement

If WebSphere is not able to find or load the JNI library, you should see the following message in the trace.log, as shown in Example A-11.

Example: A-11 WebSphere cannot find or load the JNI library

```
0000000a PmiRmArmXfcAr 3   init: creating newArmApplication object with parms:
WebSphere:APPLICATION_SERVER, cell=hcp121Cell01, node=hcp121Node02.TradeServer1,
server=TradeServer1
0000000a PmiRmArmXfcAr 3   init: newArmApplicationDefinition error_code=-15001,
error_msg=-15001
0000000a PmiRmArmXfcAr 3   init: newArmApplication error_code=-15001,
error_msg=-15001
```

The error code -15001 is returned for any problems while loading or calling the JNI library. You would see this error code if the JNI library was not found, or if registration to ARM platform support failed. You should then double-check the java.library.path property. To read about all the ARM return code descriptions, search for “ARM return codes” in the IBM EWLM Information Center.

Enabling ARM instrumentation in IHS

Technically, you are going to enable ARM in the Application Server plug-in running inside IHS, not in IHS itself. After you enable request metrics in WebSphere, you should regenerate the

plug-in configuration file to enable ARM instrumentation to the plug-in. Look for the following lines in your new plugin-cfg.xml:

```
<RequestMetrics armEnabled="true" loggingEnabled="false" rmEnabled="true"
traceLevel="HOPS">
```

Remembering the user privilege, IHS by default, is running as the user nobody. But this time on AIX, you do not have to give nobody any additional capabilities. It is inherited when you start IHS with root, which is normally required for IHS to listen on port 80 or 443. If you have a different setting in your environment, you should give the same capabilities as usual. On Linux, you always have to add nobody in the ewlm_armusers group.

After you restart IHS with the new plugin-cfg.xml using the `# lsarm -a` command, you should see it registered to ARM with the following message:

```
APPL: IBM Webserving Plugin
```

Troubleshooting

If there is an error while the Application Server plug-in is trying to initialize ARM or making ARM registration, you would see it in the plug-in log file, normally in /opt/IBM/WebSphere/Plugins/logs/http_plugin.log.

The following are some major differences between IHS and the previous middleware:

- With DB2 and Application Server, if ARM support is not initialized successfully, everything is still fully functional. However, no ARM data gets transmitted.
- With the Application Server plug-in, if ARM support fails to initialize, the plug-in will be totally stuck. IHS will still serve static content. However, it will not be able to forward any requests to WebSphere.

If you are debugging a strange case of Application Server plug-in failure, keep in mind that ARM support might make it stuck sometimes. If in doubt, change `armEnabled=false` in the plug-in configuration file and try again.

Viewing transaction statistics using the Control Center

Now that you have turned on instrumentation everywhere, run some transactions and see what statistics you can get out of EWLM. For this exercise, we set up a cluster of three Application Server nodes front-ended by an IHS node with a DB2 database server at the back end. One of our servers acted as Application Server and IHS at the same time.

After accessing the Trade application several times from our browser, the Control Center informed us about how the cluster is performing with regard to the performance objective previously set in the domain policy. As shown in Figure A-13 on page 408, our goal was to have 95 percent of the transactions belonging to the Trade service class respond within half a second (500 ms). We got to 99 percent in that range. EWLM gave us a performance index of 0.5 for the transactions occurring during that time interval. A performance index of less than 1.0 means we exceeded the performance objective. If this goal was the SLA for this application, we would know that we exceeded the SLA and our clients would be satisfied.

Archived

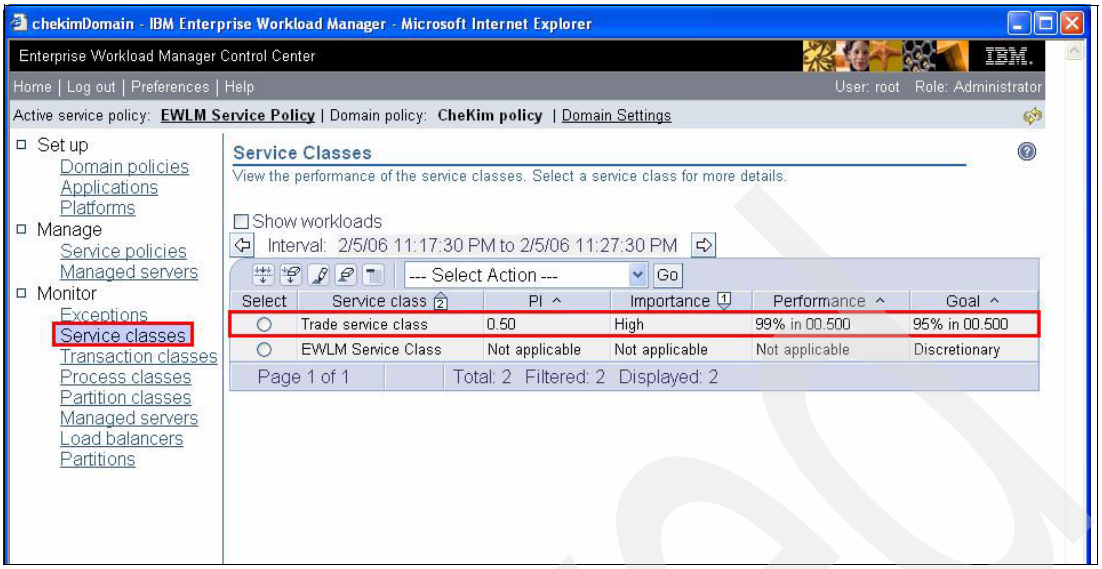


Figure A-13 EWLM Control Center: service classes view

In the Service classes view, perform the following steps:

1. Select **Trade service class**.
2. Select **Server topology** in the Action list, and then select **Go**. A window appears showing the system topology that EWLM has discovered on-the-fly while transactions are being run. As shown in Figure A-14 on page 409, you also see that there were 75 requests being sent to the hci245 Application Server node while 28 and 29 requests were sent to the two other nodes. At this point, EWLM is only capturing statistics about those transactions. It does not influence the outcome in any way.

Archive

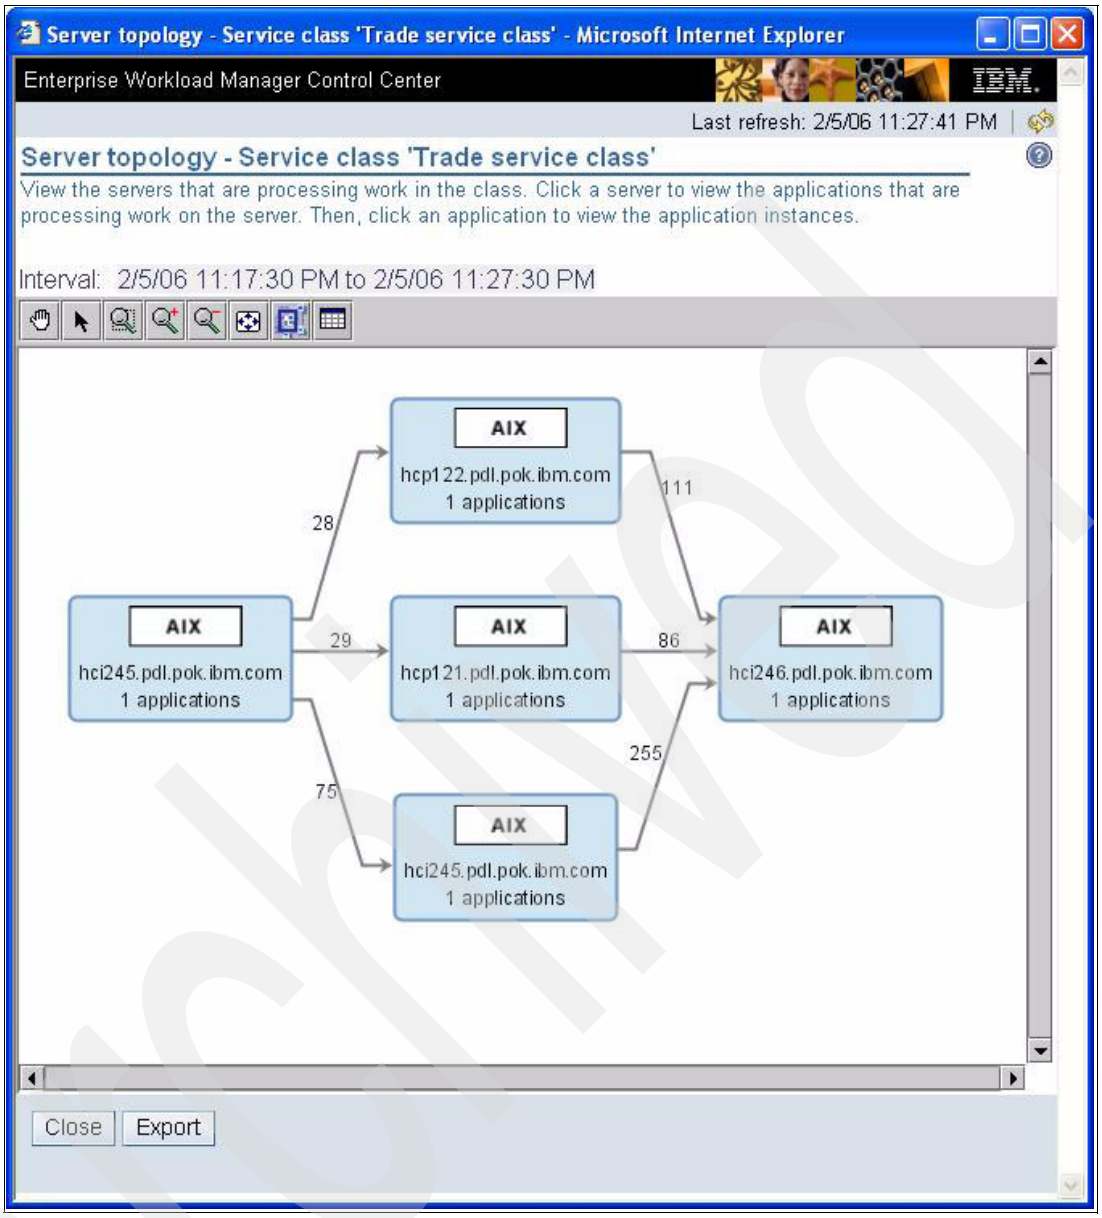


Figure A-14 EWLM Control Center: server topology graph

3. Select the **Trade service class** again by choosing **Application topology,** and then selecting **Go**. A window appears showing the same topology, but in the application perspective (as shown in Figure A-15 on page 410).

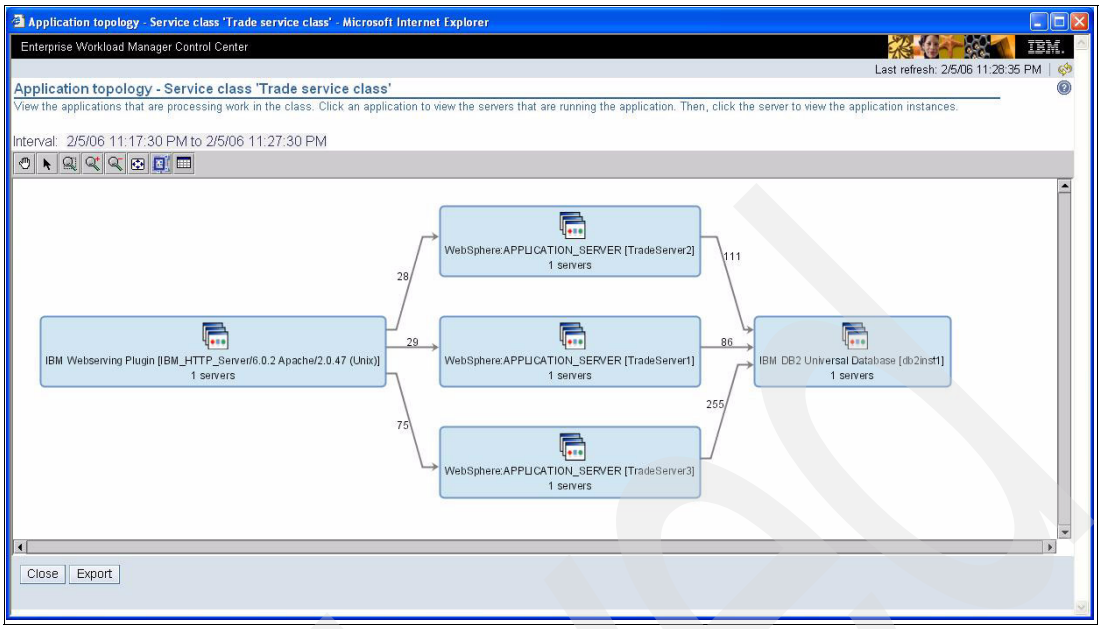


Figure A-15 EWLM Control Center, Application topology graph

In both topology graphs, you can click each box to drill down to a more detailed level, down to the individual process with its respective statistics. You can position the mouse cursor on the box that you want and the statistics table will appear.

Conclusion

We have just scratched the surface in describing what the EWLM Control Center is and the types of statistics it can present to you. In this appendix, we showed you how to enable ARM instrumentation in various middleware. This appendix showed you the summarized and granular statistics that you can get from the EWLM Control Center, as well as the topology graphs showing the transaction paths. It demonstrated how you can get those statistics without touching the application code, using off-the-shelf IBM middleware.

As IBM continues to ARM-instrument more and more middleware, you will see the benefit of your EWLM investment growing. It is only a matter of time before you can have an end-to-end view of your Web transactions all the way back to your earlier applications running on the mainframe, or your corporate enterprise resource planning (ERP). In the meantime, if you want to get the same benefits from your home-made applications, you can ARM-instrument them yourself. Search for “EWLM APIs” in the Information Center or visit The Open Group.

You can apply these lessons to a WebSphere Portal Server or WebSphere Commerce Server environment as well, provided that it is running on Application Server 6.0 or higher.

Abbreviations and acronyms

AAS	Advanced Administrative System
AMC	Autonomic Monitoring Engine
API	Application Programming Interface
ARM	Application Response Measurement
ARP	Address Resolution Protocol
AUIML	Abstract User Interface Markup Language
BPEL4WS	Business Process Execution Language for Web Services
CBE	Common Base Event
CBM	Component Business Modeling
CCI	Common Client Interface
CEC	Central Electronic Complex
CICS®	Customer Information Control System
CIM	Common Information Model
CIM Provider	Common Information Model Provider
CIMOM	Common Information Model Object Manager
CMDB	Common Model Database Schema
CMR	Converged ID Mapping Repository
CP	Control Processor
CRT	Common Runtime
CSM	Cluster System Manager
DASD	Direct Access Storage Device
DB2	DATABASE 2
DBCTL	Data Base Control Subsystem
DBRC	Data Base Recovery Control
DCCTL	Data Communication Control
DCM	Data Center Model
Director MP	Director Multiplatform
DLPAR	Dynamic Logical Partitioning
DMTF	Distributed Management Task Force
DN	Domain Name
DNS	Domain name System
EAP	Event Action Plan
EIM	Enterprise Identity Mapping
EJB™	Enterprise Java Beans
EPR	Endpoint Reference
ERD	Enterprise Resource Director
ESB	Enterprise Service Bus
ESS	Enterprise Storage Server
EWLM	Enterprise Workload Manager
FTP	File Transfer Program
GCR	Global Configuration Repository
GUI	Graphical User Interface
HLD	High Level Design
HMC	Hardware Management Console
HPI	Hosting Partition Infrastructure
HSL	High Speed Link
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IBM	International Business Machines Corporation
ICL	IBM Common Logging
ICMP	Internet Control Message Processing
IDD	Intelligent Device Discovery
IGT	IBM Grid Toolbox
IMR	Identity Mapping Repository
IMS™	Information Management System
IP	Internet Protocol
IRD	Intelligent Resource Director
ISC	Integrated Solution Console
ISV	Independent Software Vendor
IT	Information Technology
ITDS	IBM Tivoli Directory Services
ITSO	International Technical Support Organization
IUDD	Installable Unit Deployment Descriptor
IVM	Integrated Virtualization Manager
iXA	Integrated IBM System x Adapter
J2EE	Java 2 Platform Enterprise Edition
J2SE™	Java 2 Platform Standard Edition
JAAS	Java Authentication and Authorization Service
JCA	Java Connector Architecture
JCL	Job Control Language

© Copyright IBM Corp. 2006. All rights reserved.

JDBC	Java DataBase Connector
JMX™	Java Management Extensions
JNI	Java Native Interface
JNLP	Java Network Launching Protocol
JRE	Java Runtime Environment
JVM	Java Virtual Machine
JWS	Java Web Start
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LDW	LPAR Deployment Wizard
LIC	Launch in Context
LMB	Logical Memory Block
LPAR	Logical Partition
LPW	LPAR Planning Wizard
LTA	Log and Trace Analyzer
LTPA	Lightweight Third-Party Authentication
LVT	LPAR Validation Tool
MAC	Media Access Control
MC	Management Central
MIB	Management Information Base
MOM	Microsoft Operation Manager
MOWS	Management of Web Services
MP	Multiplatform
MQAI	Message Queuing Administration Interface
MUWS	Management Using Web Services
NCP	Network Control Program
NRM	Network Resource Model
OASIS	Organization for the Advancement of Structured Information Standards
ODI	On Demand Infrastructure
ODI-RM	On Demand Infrastructure Resource Modeling
ODOE	On Demand Operating Environment
OGSA	Open Grid Services Architecture
OGSI	Open Grid Services Infrastructure
OSA	Open Systems Adapter
OSGi	Open Services Gateway Initiative
PAM	Pluggable Authentication Modules
PAP	Provisioning Automation Pack
PICS	Platform Installation and Configuration Services
PLM	Partition Load Manager
PMAC	Policy Management for Autonomic Computing
PR/SM™	Processor Resource System Manager
RDBMS	Relational Database Management System
RDM	Remote Deployment Manager
RDS	Resource Dependency Services
REXX	Restructured EXtended eXecutor language
RMC	Resource Monitoring and Control
RMC	Remote Monitoring and Control
RMS	Resource Mapping Services
RMTR	Resource Modeling Type Registry
RTDS	Resource Topology and Discovery Services
SAN	Storage Area Network
SASP	Server Application State Protocol
SDK	Software Development Kit
SLP	Service Location Protocol
SMASH	Systems Management Architecture for Server Hardware
SMB	Server Message Block
SMF	System Management Framework
SMI-S	Storage Management Initiative Specification
SNIA	Storage Network Industry Association
SNMP	Simple Network Management Protocol
SOA	Service-oriented architecture
SOAP	Simple Object Access Protocol
SPX	IBM Software Platform Extension
SSO	Single Sign On
SVC	SAN Volume Controller
TBSM	Tivoli Business Service Manager
TCM	Tivoli Configuration Manager
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol /Internet Protocol
TIM	Tivoli Monitoring
TIO	Tivoli Intelligent Orchestrator
TM	Transaction Manager
TME®	Tivoli Management Environment®
TPC	Total Storage Productivity Center
TPF	Transaction Processing Facility

Archive

TPM	Tivoli Provisioning Manager
UCD	User Centered Design
UDDI	Universal Description, Discovery, and Integration
UDP	User Datagram Protocol
URI	Unviversal Resource Identifier
URL	Uniform resource Locator
VA	Virtualization Agent
VIOS	Virtual I/O Server
VIPA	Virtual IP Addressing
VM	Virtualization Manager
VMM	Virtualization Machine Manager
VSC	Virtual Server Complex
VSD	Virtual Server Deployment
VTS	Virtual Tape Server
WBEM	Web Based Enterprise Management
WII	WebSphere Information Integrator
WLM	Workload Manager
WMI	Windows Management Instrumentation
WSA	Web Services Addressing
WSDL	Web Service Definition Language
WSDM	Web Services Distributed Management
WSN	Web Service Notification
WS-RF	Web Services Resource Framework
WSRL	Web Resource Lifetime
WUWS	Management Using Web Services
XML	eXtensible Markup Language

Archived

Archived

Glossary

Abstract User Interface Markup Language (AUIML). An XML implementation that provides a platform and technology-neutral method of representing windows, wizards, property sheets, and other user interface elements. It defines the intent of the user interface, such that it can be described once and rendered to the user in multiple environments and on various devices. AUIML is being developed and defined by the team who is developing the IBM Autonomic Computing Common Systems Administration core capability.

Advanced Administration System (AAS). The AAS Order Fulfillment system provides the order data, and the associated invoice data for an order.

Application programming interface (API). An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

Application Response Measurement (ARM). A standard of Open Group. The application can use the ARM API set to report the information to calculate the response time, or monitor the status of work processed. In Virtualization Engine R1, an application must be instrumented using ARM 4.0 standard to be monitored by Enterprise Workload Manager. In Virtualization Engine R2, non-ARM instrumented clients are supported by Enterprise Workload Manager in a different way.

Application Response Measurement API. The ARM API is a simple API that applications can use to pass vital information about a transaction to an agent. All the applications call the ARM API just before a transaction (or a subtransaction) starts and then again just after it ends. The API is supported by an agent, which measures and monitors the transactions, and makes the information available to management applications.

The ARM calls identify the application, the transaction, and (optionally) the user, and provide the status of each transaction when it completes. This is sufficient information for the management solution to answer vital questions. Eliminating this guesswork helps both the users of the applications and the administrators responsible to make sure that they are meeting the requirements of the business.

Autonomic computing. Autonomic computing is an approach to self-managed computing systems with a minimum of human interference. The term derives from the body's autonomic nervous system, which controls key functions without conscious awareness or involvement.

Common Information Model (CIM) server. A CIM server provides access to instrumented CIM resources. Pegasus is an open source implementation of a CIM server that has been ported onto all IBM servers.

Common Base Event (CBE). The Common Base Event model is a standard that defines a common representation of events. This standard, developed by the IBM Autonomic Computing Architecture Board, supports the encoding of logging, tracing, management, and business events using a common XML-based format. Using this format, you can correlate different types of events that originate from different applications. The Common Base Event model is part of the IBM Autonomic Computing Toolkit. For more information, see http://www.ibm.com/autonomic.

Common Client Interface (CCI). Defined by the J2EE Connector architecture specification, the CCI defines a set of interfaces and classes whose methods allow a client to perform typical data access operations. A client or application component that uses the CCI to interact with an underlying service does so in a prescribed manner. A client application that relies on a CCI resource adapter is very much like any other J2EE client that uses enterprise bean methods.

Common Information Model (CIM). The CIM is a conceptual information model for describing management properties that are not bound to a particular implementation. This allows for the interchange of management information between management systems and applications through the Common Information Model object manager (CIMOM).

Common Information Model object manager (CIMOM). CIMOM is an object management engine that exists between the managed system and the management application. Management applications can be local or remote, but currently CIMOM and the programs that collect management data, or "providers, must be located on the machine that is being instrumented. See also: Pegasus

Common Information Model (CIM) provider. A program that corrects the management data of CIM and reports it to CIMOM.

Common Model Database schema (CMDB). The resource repository of RMS (RMS database) is based on CMDB.

© Copyright IBM Corp. 2006. All rights reserved.

Data Base Control (DBCTL) subsystem. An interface between CICS Transaction Server and IMS or ESA that allows access to IMS DL/I full-function databases and to data entry databases (DEDBs) from one or more CICS systems without the need for data sharing. It also provides release independence, virtual storage constraint relief, operational flexibility, and failure isolation.

Database Recovery Control (DBRC). A feature of the IMS Database Manager that facilitates easier recovery of IMS databases. DBRC maintains information required for database recoveries, generates recovery control statements, verifies recovery input, maintains a separate change log for database data sets, and supports sharing of IMS databases and areas by multiple IMS subsystems.

Data Communications Control (DCCTL). A subsystem that allows IMS TM to act as a stand-alone, full-function transaction manager that can connect to DB2 or other external subsystems.

Distributed Management Task Force (DMTF). A standardization organization that designed CIM and WBEM.

Domain Name System (DNS). The distributed database system used to map domain names to IP addresses.

Endpoint reference (EPR). The combination of the address of a Web service that implements the manageability interface for a managed resource type and the properties that allow that Web service to identify the managed resource.

Enterprise Identity Mapping (EIM). A mechanism for associating a person or entity to the appropriate user identities in various user registries throughout an enterprise. EIM provides application programming interfaces (APIs) for creating and managing these identity mapping relationships, as well as APIs that applications use to query this information.

Enterprise JavaBeans™ (EJB). A component architecture defined by Sun Microsystems™ for the development and deployment of object-oriented, distributed, enterprise-level application.

Enterprise Workload Manager Resource Director. A function of Enterprise Workload Manager that monitors the processor resource of the logical partition (LPAR) of an Enterprise Workload Manager managed server.

eXtensible Markup Language (XML). A text-based tag language used for document processing and for publishing information on the Web.

Hardware Management Console (HMC). A system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information to IBM for analysis.

Hypertext Markup Language (HTML). A markup language that conforms to the SGML standard and was designed primarily to support the online display of textual and graphical information that includes hypertext links.

Hypertext Transfer Protocol (HTTP). An Internet protocol that is used to transfer and display hypertext and XML documents on the Web.

Hypertext Transfer Protocol Secure (HTTPS). A Web protocol for secure transactions that encrypts and decrypts user page requests and pages returned by the Web server.

IBM Grid Toolbox (IGT). IGT in Virtualization Engine Rel.1 is factored and integrated as "IBM Grid Services" and "Virtualization Engine Common Services and Components" in Virtualization Engine Rel.2.

IBM Software Platform Extension (SPX). IBM Software Platform Extension Software Componentization Architecture. A packaged software brand product to install and support a solution, such as Virtualization Engine, easily. SPX is the collection of architectures and specifications that define how the IBM software componentization strategy is implemented, and guide the changes to IBM software development processes, management systems, and culture.

Integrated Solutions Console (ISC). ISC framework is a common console technology for systems administration based on WebSphere Portal technology. It is also used for Virtualization Engine console.

Integrated Virtualization Manager (IVM). IVM is a browser-based system management interface that you can use to manage a single managed system without a Hardware Management Console (HMC). The VIOS in a managed system can use IVM to create logical partitions on a single managed system, manage the virtual storage and virtual Ethernet on the managed system, and view service information related to the managed system.

Integrated IBM System x adapter. An external card located in an IBM System x server that provides direct, high-speed attachment to an IBM System i server.

Intelligent Device Discovery (IDD). A network-based IT asset discovery and categorization tool. Data discovery in Resource Dependency Service is performed through IDD.

Internet Control Message Protocol (ICMP). An Internet protocol that is used by a gateway to communicate with a source host, for example, to report an error in a datagram.

Internet Protocol (IP). A protocol that routes data through a network or interconnected networks. Internet Protocol (IP) acts as an intermediary between the higher protocol layers and the physical network.

Java 2 Platform Enterprise Edition (J2EE). An environment for developing and deploying enterprise applications, defined by Sun Microsystems Inc. The J2EE platform consists of a set of services, application programming interfaces (APIs), and protocols that provide the functionality for developing multitiered, Web-based applications. (Sun).

Java 2 Platform Standard Edition (J2SE). The core Java technology platform.

Java Authentication and Authorization Service (JAAS). In J2EE technology, a standard API for performing security-based operations. Through JAAS, services can authenticate and authorize users while enabling the applications to remain independent from underlying technologies.

J2EE Connector architecture (JCA). Java Connector architecture is a standard architecture for connecting the J2EE platform to heterogeneous enterprise information systems (EIS). A resource adapter compliant with this architecture can plug into a J2EE server and provide connectivity between the application server, the EIS, and the enterprise application.

Java Database Connectivity (JDBC). An industry standard for database-independent connectivity between the Java platform and a wide range of databases. The JDBC interface provides a call-level API for SQL-based database access.

Java Management Extensions (JMX). A means of doing management of and through Java technology. JMX was developed through the Java Community ProcessSM program, by Sun Microsystems, Inc. and some leading companies in the management field. JMX is a universal, open extension of the Java programming language for management that can be deployed across all industries, wherever management is needed.

Java Native Interface (JNI). A programming interface that allows Java code to interoperate with functions that are written in other programming languages.

Java Network Launching Protocol (JNLP). To enable your Java application to execute using Java Web start, you must create a Java Networking Language Protocol (JNLP) file for the application. The JNLP is an XML file that provides the basic elements and descriptions for your application. The JNLP conforms to the Java Community Process (JCP), under specification number JSR 000056 (see Resources). The purpose of the JNLP file is as follows:

- To specify JAR files to download from the Web server
- To provide other possible package requirements
- To specify system properties
- To allow for any necessary runtime parameters
- To specify which version of the Java 2 platform to use

Java Runtime Environment (JRE). JRE provides the libraries, the Java Virtual Machine, and other components to run applets and applications written in the Java programming language. In addition, two key deployment technologies are part of the JRE: Java Plug-in, which enables applets to run in popular browsers; and Java Web Start, which deploys stand-alone applications over a network.

Java Virtual Management (JVM). A software implementation of a central processing unit that runs compiled Java code (applets and applications).

Lightweight Directory Access Protocol (LDAP). An open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 directory access protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory.

Lightweight Third Party Authentication (LTPA). A protocol that uses cryptography to support security in a distributed environment.

Local area network (LAN). A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

Logical partition (LPAR). A subset of a single system that contains resources (processors, memory, and input/output devices). A logical partition operates as an independent system. If hardware requirements are met, multiple logical partitions can exist within a system.

Management of Web Services (MOWS). Management of Web services (MOWS) is an application of Management using Web services (MUWS) to the resources that are elements of the Web Services Architecture [WS-Arch]. This WSDM specification defines how the manageability of Web service endpoints and resources exposed as Web services can be accessed via Web services. In order to achieve this goal, MOWS is based on the MUWS specifications.

Management using Web Services (MUWS). MUWS defines how an Information Technology resource connected to a network provides manageability interfaces such that the IT resource can be managed locally and from remote locations using Web services technologies.

Message Queuing Administration Interface (MQAI). A programming interface that performs administration tasks on a WebSphere MQ queue manager through the use of data bags. Data bags allow the user to handle properties (or parameters) of WebSphere MQ objects.

Network Control Program (NCP). A program that controls the operation of a communication controller.

On Demand Infrastructure Resource Modeling (ODI:RM). ODI-RM (ODI:RM) is an architecture intended to facilitate an integration of the IBM products relevant to On Demand. ODI:RM is converged to RMS.

Open Grid Services Infrastructure (OGSI). The set of interfaces that enable developers to create, discover, manage, and exchange information among grid services.

Pegasus. Pegasus is one of the open source implementations of the CIMOM that adheres to the Distributed Management Task Force (DMTF) CIM and Web-based Enterprise Management (WBEM) standards. Pegasus is designed to be inherently portable and builds and runs on the AIX, Linux, and Windows operating systems.

Processor Resource/Systems Manager™ (PR/SM). A feature that allows a processor to use several MVS™ images simultaneously and provides logical partitioning capability.

Resource Mapping Service (RMS). The architecture for Resource Mapping Service focuses on the discovery and mining of information from existing applications, tools, and IT systems, and presenting a static view of those assets on the RMS database.

Server Application State Protocol (SASP). The communication protocol between Enterprise Workload Manager domain manager and a load balancer. The load balancer obtains the weight of each server or the way to distribute requests to the managed servers from the Enterprise Workload Manager domain manager.

Service-Oriented architecture (SOA). A conceptual description of the structure of a software system in terms of its components and the services they provide, without regard for the underlying implementation of these components, services, and connections between components.

Simple Object Access Protocol (SOAP) SOAP provides an open, extensible way for applications to communicate using XML-based messages over the Web, regardless of what operating system, object model, or language they use. A lightweight, XML-based protocol for exchanging information in a decentralized, distributed environment, SOAP can be used to query and return information and invoke services across the Internet. The protocol consists of a framework for describing what is in a message and how to process it, a set of encoding rules for expressing application-defined data types, and a convention for representing remote procedure calls and responses.

Software development kit (SDK). A set of tools, APIs, and documentation to assist with the development of software in a specific computer language or for a particular operating environment.

Storage area network (SAN). A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

Transmission Control Protocol/Internet Protocol (TCP/IP). An industry-standard non-proprietary set of communication protocols that provide reliable end-to-end connections between applications over interconnected networks of different types.

Transmission Control Protocol. A communication protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Universal Description, Discovery, and Integration (UDDI) UDDI is the service discovery protocol for Web services through which Web service providers and endpoints can be located. A set of standard-based specifications that enables companies and applications to quickly and easily find and use Web services over the Internet.

Uniform Resource Locator (URL). The unique address of a file that is accessible in a network, such as the Internet. The URL includes the abbreviated name of the protocol used to access information resource and the information used by the protocol to locate the information resource.

Universal Resource Identifier (URI). A unique address that is used to identify content on the Web, such as a page of text, a video or sound clip, a still or animated image, or a program. The most common form of URI is the Web page address, which is a particular form or subset of URI called a Uniform Resource Locator. A URI typically describes how to access the resource, the computer that contains the resource, and the name of the resource (a file name) on the computer.

Virtual I/O Server (VIOS). VIOS is part of the Advanced POWER Virtualization hardware feature on System p5 and IBM eServer p5 systems and part of the IBM POWER Hypervisor™ and VIOS feature on OpenPower™ systems. It is a single-function appliance that resides in an IBM POWER5 partition and facilitates the sharing of physical I/O resources between client partitions (AIX or Linux) within the server. The VIOS provides virtual SCSI target and shared Ethernet adapter (SEA) virtual I/O function to client LPARs.

Virtualization Agent (VA). A component of Enterprise Workload Manager that runs in each IBM System p or i LPAR partition, and handles interactions between Enterprise Workload Manager domain manager and the IBM System p or i Hypervisor.

Virtualization Manager. A component of Enterprise Workload Manager that runs on Enterprise Workload Manager domain manager; the resource arbitrator for virtual LPAR groups.

Web Based Enterprise Management (WBEM). A set of Distributed Management Task Force (DMTF) standards that defines the protocols used to communicate with a particular Common Information Model (CIM) implementation that uses CIM servers. This architecture provides the management design framework that consists of devices, device providers, the object manager, and the messaging protocol for the communication between client applications and the object manager.

Web Service Description Language (WSDL). The standard format for describing a Web service. Expressed in XML, a WSDL definition describes how to access a Web service and what operations it will perform.

Web service. A software system identified by an URI whose public interfaces and bindings are defined and described using XML. Its definition can be discovered by other software systems. These systems may then interact with the Web service using XML based messages conveyed by Internet protocols.

Web Services Distributed Management (WSDM). WSDM is a specification and specification set, MUWS and MOWS, for managing devices as well as Web services using Web services, which are inherently dependent upon a service-oriented architecture foundation.

Web Services Resource Framework (WS-RF) WS-RF provides a common pattern for accessing resources using stateful Web services.

Windows Management Instrumentation (WMI). An application programming interface (API) in the Windows operating system that enables devices and systems in a network to be configured and managed. WMI uses the CIM to enable network administrators to access and share management information.

Archived

Archived

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this IBM Redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 423. Note that some of the documents referenced here may be available in softcopy only.

- *IBM Enterprise Workload Manager Version 2.1*, SG24-6785
- *IBM TotalStorage SAN Volume Controller*, SG24-6423
- *IBM Virtualization Engine Platform Version 2 Technical Presentation Guide*, SG24-7112
- *IBM Virtualization Engine TS7510: Tape Virtualization for Open Systems Servers*, SG24-7189
- *Implementing IBM Director 5.10*, SG24-6188
- *Virtualization on the IBM System x3950 Server*, SG24-7190
- *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9*, SG24-6695

Other publications

These publications are also relevant as further information sources:

- *IBM Tivoli Usage and Accounting Manager Administrator’s Guide*, GC32-1541
- *IBM Tivoli Usage and Accounting Manager Data Collectors for Microsoft Windows User’s Guide*, GC32-1543
- *IBM Tivoli Usage and Accounting Manager Data Collectors for UNIX and Linux User’s Guide*, GC32-1542
- *IBM Tivoli Usage and Accounting Manager Web Reporting User’s Guide*, GC32-1540

Online resources

These Web sites and URLs are also relevant as further information sources:

- Adobe - SVG - Manual Download

  `http://www.adobe.com/svg/viewer/install/main.html`
- Autonomic load balancing, Part 1: Cisco Content Switching Module

  `http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlmload1/`
- Distributed Management Task Force, Inc. (DMTF)

  `http://www.dmtf.org/home`
- Download the Free Nmap Security Scanner for Linux/MAC/UNIX or Windows

  `http://www.insecure.org/nmap/download.html`

© Copyright IBM Corp. 2006. All rights reserved.

- Global Grid Forum

  http://www.gridforum.org/
- IANA Port Number Assignments

  http://www.iana.org/assignments/port-numbers
- IBM AIX Toolbox Download Page

  http://www-03.ibm.com/servers/aix/products/aixos/linux/download.html
- IBM Autonomic Computing

  http://www-03.ibm.com/autonomic/
- IBM Director 5.10 Update 2

  https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=dmp
- IBM eServer - Getting started

  http://www-03.ibm.com/servers/eserver/about/virtualization/getstarted.html
- IBM Grid Toolbox

  https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=grid
- IBM Personal computing support - IBM Director v5.10 - Publications

  http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-61788
- IBM Personal computing support - IBM Virtual Machine Manager Installation and User's Guide - IB

  http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56955
- IBM Publications Center

  http://www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi
- IBM - Service and productivity tools for Linux on POWER

  http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html
- IBM Support and downloads

  http://www.ibm.com/support
- IBM Support: Fix Central

  http://www-912.ibm.com/eserver/support/fixes/fixcentral
- IBM Systems Software Information Center

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp
- IBM Systems Workload Estimator

  http://www.ibm.com/eserver/iseries/support/estimator
- IBM Virtualization Engine

  http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/home2.html
- IBM Virtualization Engine - Operating System Fixes

  http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/osDownloads.html
- IBM Virtualization Engine Services fixes and updates

  http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/gmi.html

Archived

- ILOG JViews: Java graphics for better Web GUIs

  http://www.ilog.com/products/jviews/
- Insecure.Org

  http://www.insecure.org
- Installing bridges to supported Linux management sources

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/esmcinfo/eicacinstallmslinux.htm
- Installing DB2 Fix Pack 9a for IBM Virtualization Engine clients

  http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/DB29a.html
- Network Appliance

  http://www.netapp.com/
- Open Services Gateway Initiative (OSGI)

  http://www.osgi.org
- Organization for the Advancement of Structured Information Standards (OASIS)

  http://www.oasis-open.org
- Performance monitoring with Enterprise Workload Manager

  http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-arm/
- Scenario: Dynamically shifting processing capacity to manage workloads

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/veicinfo/eicar_i5p5_scenario.htm
- Scenario: Enable autonomic partition management on IBM System p5 servers with the Enterprise Workload Manager

  http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-lpar1/
  http://www-128.ibm.com/developerworks/autonomic/library/ac-ewlm-lpar2/
- Virtualization Engine planning advisor

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/veinfo/planner.htm
- Virtualization Engine product requirements

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/veicinfo/eicarsystemservicereqs.htm
- Virtualization Engine Registration Assistance

  http://www14.software.ibm.com/webapp/set2/sas/f/VirtualizationEngine/si.html
- VMware - Virtualization Software

  http://www.vmware.com/

How to get IBM Redbooks

You can search for, view, or download IBM Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy IBM Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Archived

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Archived

Index

Numerics

5648-F08 IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and zSeries, V2.1 26
5655-EWM IBM Virtualization Engine Enterprise Workload Manager for z/OS, V2.1 27
5733-DIR IBM Director with Virtualization Engine console, V5.10 28
5733-DR1 IBM Director with Virtualization Engine console for i5/OS, V5.10 28
5733-EWA IBM Virtualization Engine Enterprise Workload Manager for i5/OS, V2.1 29
5733-EWM IBM Virtualization Engine Enterprise Workload Manager, V2.1 30
5733-RDS Virtualization Engine Resource Dependency Service, V2.1 31

A

adding
  a new database in the system 287
  monitor to watch list 347
  new resource 280
  new resources in the environment 291
  Process Monitor to Health Center 345
  resource monitor to Health Center 353
  resources to Health Center and Watch List portlet 362
Adding resources in your group 282
Advanced Power Virtualization (APV) 9
AIX 9, 26, 46, 67, 86, 102, 128, 142, 160, 164, 181, 194, 228, 277, 302, 328, 373, 380, 385, 387
  on IBM System p 51
  systems prerequisite considerations 72
Application Response Measurement 379
Application Response Measurement.See also ARM
Applications resource types 236
applying fixes 189
APV (Advanced Power Virtualization) 9
ARM 385
  instrumentation in Application Server 402
  instrumentation in DB2 UDB 399
  instrumentation in IHS 406
  platform support 388
ARM. See also Automated Response Measurement
artifact XML file 69
autonomic 15, 380, 385, 404
Autonomic Deployment Engine 22, 196

B

BladeCenter 317, 330
  Management 302
bridge 291, 324, 330, 375
Building an alert 338
business
  functions 216
  goals 386
  rules 4
business goals 4
business rules 4

C

CIM 13, 227
  CIMOM 309
  repository 13
  servers 222
CIM. See also Common Information Model
CIMS Labs 6
CKRM 389
Class-Based Kernel Resource Management (CKRM) 389
CMDB 15
Collation 15
collectors 221, 244, 275
  configure 245
command 194
  cat 261
  cfgdb 82
  chkconfig 149, 162
  chown 75
  configWizardDM 129
  configWizardMS 182
  createMS 185
  DbmsConfigMethod 86
  dircli 317
  ewlm 388
  gmi 201
  ibmdiradm 206
  ibmslapd 128, 206
  iddAdmin 258
  idscfgdb 206
  idsilist 206
  idsucfgdb 206
  InstallVEWin 125, 127
  instfix 72, 74
  jre 199
  lsarm 164
  lslpp 54
  manageIUD 198
  mksysb 71
  netstat 73
  RPD 166
  rpm 53, 147
  RSP 181
  serverStatus 261
  setup 77, 176
  srt 197
  start 75
  startDM 142

© Copyright IBM Corp. 2006. All rights reserved.

startISC 160, 202, 263, 375
startMS 186
startServer 128, 161–162, 261
startsrc 349
startWAS 142
stopISC 145, 263, 375
stopServer 148, 261
sysctl 146
tail 91
twgcon 144
twgipcsv 186
twgstart 144
vefixlevel 194
vessoconfig 327
vessoserverconfig 328
wget 74
wgstart 186
command line 175, 185
command-line mode 182, 199
Common Information Model 13
Common Information Model. See also CIM
components 9, 21
common runtime components 21
Configuration Management Database 15
configuring
collector 245
Event Action plan 364
Hardware Management Console 309
IBM Director to monitor CPU 351
management sources 304
resources in Health Center 331
console mode 167, 172, 201
console topology view 232
Control Center 386
correlation engine 221, 256
creating
a new relationship 285
EWLM DM configuration ID 389
EWLM domain policy 395
EWLM MS configuration ID 394
operating system users and groups 69
credential 269, 293, 309, 325, 333, 342, 360, 372
credentials 224, 229, 317

D

dashboard 356
data collectors 241
Data Containers 236
DB2 fix, installing 205
DB2 instance owner 70
DB2 Universal Database 245
DB2 Universal Database collector 245
Defining a filter 278
diagnosing a hardware error 330
Director Agent 175
Director Server 21, 46, 69, 78, 124, 143
Directory Server instance owner 70
Directory Server planning worksheet 48
discovered resource 233
discovering systems with IBM Director 312
Discovering with agents 219
Discovering with probes 219
Discovering with scanners 219
Distributed Management Task Force (DMTF) 13
DMTF (Distributed Management Task Force) 13
domain manager.See EWLM

E

encryption 177
endpoint reference. See EPR
enterprise service bus 14
Enterprise Service Bus (ESB) 14
Enterprise Workload Manager 6, 379
EPR 12
ESB (Enterprise Service Bus) 14
Event Action plan 364
EWLM 25, 164
5655-EWM 25
5733-EWA 26
5733-EWM 25
agent 164, 167
Control Center users 70
domain manager 7, 21, 25, 72, 90, 117
Managed Server 182
EWLM DM configuration ID, creating 389
EWLM domain policy, creating 395
EWLM MS configuration ID, creating 394

F

filter 278, 337, 366, 396
Firewall Broker 133, 164, 169
fix installation using the Virtualization Engine console 200
Fix Pack 191
fixes 189
critical 191
installing 201
fixPeriod 191

G

Global Grid Forum (GGF) 11
graphical configuration wizard 129, 182
graphical installation process 77
graphical installation wizard 77, 167, 175
graphical mode 201
Grid Toolbox 10
group 282, 298, 321, 364
adding resources in a group 282
grouping resources 321

H

hardware error 330
Hardware Management Console (HMC) 9, 303
Hardware Management Console. See also HMC
Health Center 305, 347, 362
HMC 265, 301, 309, 334, 372
HMC plug-in 309
HTTP proxy server 184

Archived

I

I/T infrastructure
  element 13
  manageable resource 13
i5/OS 24, 47, 64, 89, 95, 194, 302, 328, 380, 387, 404
  5733-DR1 26
  5733-EWA 26
IBM Common Logging (ICL) 22
IBM Director 6, 21, 25, 28
  5733-DIR 25
  5733-DR1 26
  Agent 164
  Server 143
  server on IBM System x 58
  Server user 69
IBM Director collector 252
IBM Director extensions 302
IBM Director fix, installing 206
IBM Director software distribution 165
IBM Director to monitor CPU, configuring 351
IBM Director, installing 77
IBM eServer
  pSeries 8
IBM Tivoli Usage and Accounting Manager 6
IBM TotalStorage
  component 9
IBM Virtualization Engine
  new component 8
  packaging 32
  platform 4
  Version 1 8
  Virtual Access 7
ICL (IBM Common Logging) 22
IDD 218, 275
IDD campaign 226
IDD campaign configuration with credentials 237
IDD campaign with credentials 237
IDD campaign without credentials 229
IDD database 220
IDD technologies 219
iddAdmin tool 257
Installable unit (IU) 69
Installable unit deployment descriptor (IUDD) 69
installation setup planning worksheet 47
installation wizard 201
installing
  DB2 fix 205
  fixes 201
  IBM Director 77
  IBM Director fix 206
  managed servers agents 163
  management servers 63
Integrated Solution Console. See also ISC
Integrated Solutions Console 202, 305
Integrated Virtualization Manager 9
Integrated Virtualization Manager (IVM) 5
Intelligent Device Discovery. See IDD
Interpreting the information shown in the topology viewer 297
ISC 202, 257, 263
iSeries 8
ITDS 72
ITSO environment 227
ITSO environments 65
IU (Installable unit) 69
IUDD (Installable unit deployment descriptor) 69
IVM (Integrated Virtualization Manager) 5

L

launch-in-context 343
launchpad 309, 315
LDAP 25, 270
Level 1 scanning 220
Level 1+ scanning 220
Level 2 334, 343
Level 2 scanning 220
Level-0 175
Level-1 175
Level-2 175, 179
Lightweight Directory Access Protocol 25
Linux 25–26, 53, 55, 65, 74, 128, 142, 160, 172, 203, 228, 234, 327, 361, 375, 380
  5648-F08 25
Linux on IBM System x 56
Linux on IBM System x prerequisite considerations 72
Linux on POWER 54
Linux Toolbox 53
list of the packages 193
LM78 device driver 58
logical partition 6
  CPU allocations 6
  management 379

M

manageability endpoint 12
manageable resource (MR) 12
Manageable Resource Service 223
managed nodes 303
Managed Resources 267, 269
Managed Resources environment 275
managed servers 163
managed servers agents, installing 163
Management of Web Services (MOWS) 12
management servers 63, 302
management servers, installing 63
Management Using Web Services 12
Management Using Web Services (MUWS) 12
manager 90
managing a VMWare Virtual Machine 358
Manual installation 167
Micro-Partitioning 388
Microsoft Virtual Server 359
Microsoft Windows on IBM System x / BladeCenter 59
monitoring
  LPAR 372
  system CPU usage 350
MR (manageable resource) 12

Archive

N

naming conventions 25
network driver 179
Network Installation Management (NIM) 210
Network Interfaces view 235
Network virtualization 6
new features in Version 2.1 8

O

ODIRMDB database 224
Open Grid Services Infrastructure (OGSI) 11
Open Pegasus 311
Open Services Gateway Initiative (OSGI) 12
open standards 16
Open Systems Adapter (OSA) 6
operating system users and groups, creating 69
operating system virtualization 6
Operating Systems view 234
Organization for the Advancement of Structured Information Standards (OASIS) 12

P

Packaging 21
partition management 137, 141, 380, 423
parts of IBM Virtualization Engine 21
patches 189
performance monitoring using Enterprise Workload Manager 380
PID 261
planning
  IDD campaign 226
planning advisor 33, 38, 50, 71
planning worksheets 33
port assignment worksheet 45
portType 222
post-installation 64, 163, 170
post-installation checking tasks 128
POWER5 388
Preinstallation 64, 163
prerequisites 20, 55
Process Monitor to Health Center, adding 345
Products 21
PTF 195

R

Rack Manager 302
RDS. See Resource Dependency Service
reboot 85, 102, 140, 171, 182
Red Hat 67
Redbooks Web site 423
registering a new resource 280
registration 196
relationship 285
Remote Package Deployer 165
Remote Package Deployer (RPD) 165
renaming BladeCenter 317
requesting access
  systems in IBM Director 315
Resource Dependency Service 6, 25, 74, 106, 215, 272
  5733-RDS 25
Resource Dependency Service Collectors 221
resource factory 13
resource factory service 223
Resource Health 334
Resource model 216
resource model services 221
resource modelling framework 222
resource monitor to Health Center, adding 353
resource registry service 223
resource topology 266
resource type 13
resource types 234
resource virtualization
  facilitation 9
resourceID 222
resources in Health Center, configuring 331
resources to Health Center and Watch List portlet, adding 362
restarting the nfsd process 348
RMS data collectors 221
RPD 165
RPD script 166

S

SAN Volume Controller (SVC) 9
Schedule 231, 248
Secure Sockets Layer 390
security 177
security planning worksheet 49
Server virtualization 5
Service Location Protocol 372
Service Management Framework (SMF) 12
service-oriented architecture 13
service-oriented architecture.See also SOA
services 21
Shared memory and DB2 Database server on Linux on POWER processor 146
Shared processing mode 140
Silent installation 174
silent installation process 85
silent installation using command line parameters 175
silent mode 167, 175, 201
Single Object Access Protocol (SOAP) 10
Single Sign-on 328
Single Sign-on (SSO) environment for Linux systems 143
SLES 55, 74, 143, 387
SLP 372
SMBus device driver 58
SNMP 237
SOA 13
SOA. See also service-oriented architecture
SSH 239
SSL 132, 390
Static Group 323
Storage virtualization 6
Storage Volumes view 237
SUSE 39, 58, 65, 302

Archived

SUSE Linux Enterprise Server (SLES) 302
SUSE Linux Enterprise Server.See also SLES
SVC 9
System i 7, 29, 65, 89, 243, 301
System p 7, 33, 51, 66, 137, 265, 297, 301, 334, 340, 350, 381, 388
System x 6, 26, 39, 56, 78, 167, 181, 297
System z 6, 25, 65, 221, 243
  5648-F08 25
systems in IBM Director, requesting access 315
systems view 242

T

TADDM 15
TBSM 221
TCM 221
the monitor to watch list, adding 347
Tivoli Provisioning Manager 10
Tivoli Provisioning Manager (TPM) 10
Topology Service 222
Topology Viewer 265
TS7510 Virtualization Engine 9
TUAM 6

U

UAM 6
Uncapped sharing mode 140
Uninstalling fixes 207
Update Installer 189, 199, 203
Updates 191
Usage and Account Manager (UAM) 21
users access and password planning worksheet 46
Using topology viewer for removing the resources 293

V

VEMS 261
  administration console 262
viewing transaction statistics using Control Center 407
virtual access 4, 7, 17
Virtual Engine Updates 194
virtual Ethernet 6
Virtual Machine Manager 7, 303
virtual management 4, 6
virtual planning tools 8
virtual resources 4–5
virtual storage pool 17
Virtualization 4
virtualization capability 16
virtualization component 10
Virtualization Engine
  console 149
  IBM Director 7
  Maintenance Environment Assistant 189
  planning advisor 8
  Version 2.1 3
  visualizing 370
Virtualization Engine components 194
Virtualization Engine console 1, 17, 21, 25, 31, 40, 51, 65, 90, 113, 222, 244, 257, 263, 265, 301
  5733-DIR 25
  5733-DR1 26
  Accessing resource topology 266
Virtualization Engine environment 21, 25, 36, 66, 88, 267–268, 270
Virtualization Engine installer 24, 64, 88, 110, 165
Virtualization Engine management server 261
Virtualization Engine subcomponents 194
virtualized 4
virtualized environment 8
virtualized resources 216
VMM 303, 359, 370
VMM plug-in 359
VMWare 6, 303, 358
  ESX 360
VMWare ESX 359
VMWare GSX 359

W

Watch List 362
Web Service 12, 216, 223, 313
Web Services Description Language (WSDL) 10
Web Services Distributed Management 12
Web Services Distributed Management (WSDM) 12
Web-based System Manager 309
WebSM 309, 374
WebSphere 217
WebSphere Application Server 249
  security administrator 70
WebSphere Application Server collector 249
Windows 2003 60, 302
Windows systems prerequisite considerations 72
Windows XP 269
WMI 238
Work with Systems view 234
workload balancing 379
Workload Estimator 8
WS-Addressing specification 11
WSDL 216
WSDM 12, 15, 216
WSDM specification 12
WS-Resource 223
WS-Resource Framework (WS-RF) 11
WS-Resource Property 222
WS-RF 11, 216

Z

z/OS 7, 25, 380, 387, 404
  5655-EWM 25
z/VM 6

Archive

Archived





Redbooks

Planning, Installing, and Using the IBM Virtualization Engine Version 2.1

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM CONFIDENTIAL

Archived

Archived

Planning, Installing, and Using the IBM Virtualization Engine Version 2.1



Describes the components of the Virtualization Engine

Demonstrates the full installation process

Gives examples of how to use the Virtualization Engine

The concepts of automation and virtualization are fundamental to IT infrastructure simplification. *Virtualization* is the ability to separate the direct dependency of an application from a physical resource.

To address the virtualization aspects of the On Demand Operating Environment, the IBM® Virtualization Engine™ was announced in August 2004. Version 1 was then available. A new version, the Virtualization Engine Version 2.1, was announced and made available at the end of 2005. This new version enhances the functions provided in Version 1 and provides new products, new features, and new tools.

This IBM Redbook describes how to install and to use the Virtualization Engine Version 2.1. This IBM Redbook is the companion of *IBM Virtualization Engine Platform Version 2 Technical Presentation Guide*, SG24-7112, which describes the functions of the components. This IBM Redbook is suitable for IT Specialists who want to install and use the Virtualization Engine Version 2.1 products and components.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7213-00 ISBN 0738494895