


# IBM Systems Director Navigator for i

Learn about the new Web console for  
IBM i

Explore the functionality of the  
new Web console

Leverage this new powerful  
tool to manage IBM i



Brandon Schulz  
Giancarlo Omati  
Morten Buur Rasmussen  
Johnnie Talamantes  
Claudio Villalobos  
Brian Younger

**Redbooks**





International Technical Support Organization

**IBM Systems Director Navigator for i**

December 2009

Archived

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

Archived

**First Edition (December 2009)**

This edition applies to IBM i 6.1 (product number 5761-SS1).

© Copyright International Business Machines Corporation 2009. 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

<b>Notices</b> .....	vii
Trademarks .....	viii
<b>Preface</b> .....	ix
The team who wrote this book .....	ix
Become a published author .....	xi
Comments welcome .....	xi
<b>Chapter 1. Overview of IBM Systems Director Navigator for i</b> .....	1
1.1 Strategic positioning .....	2
1.2 What is included in Systems Director Navigator .....	4
<b>Chapter 2. Installation, configuration, and support tips</b> .....	7
2.1 IBM i server .....	8
2.2 Required jobs .....	8
2.2.1 IBM i required products .....	9
2.2.2 Recommended PTF levels .....	9
2.2.3 Client-side requirements .....	9
2.2.4 Browser security warning .....	10
2.2.5 Access the console .....	12
2.3 Support tips .....	14
<b>Chapter 3. The Web console</b> .....	17
3.1 Welcome page .....	18
3.2 Navigation within the interface .....	19
3.3 Settings .....	25
<b>Chapter 4. Out of the box</b> .....	27
4.1 New controls in IBM Systems Director Navigator for i .....	28
4.2 Overview of the new user interface .....	28
4.2.1 System task .....	29
4.2.2 Basic Operations page .....	30
4.2.3 Work Management page .....	30
4.2.4 Configuration and service functions .....	31
4.2.5 Network .....	31
4.2.6 Integrated server administration .....	32
4.2.7 Security .....	33
4.2.8 Users and groups .....	34
4.2.9 Databases .....	35
4.2.10 Journal management .....	36
4.2.11 Performance .....	37
4.2.12 File Systems page .....	37
4.2.13 Internet Configurations page .....	38
<b>Chapter 5. Network</b> .....	39
5.1 TCP/IP configuration .....	41
5.2 Remote Access Services .....	49
5.3 Servers .....	53
5.4 IP policies .....	61

5.5 Internet . . . . .	66
5.6 i5/OS NetServer . . . . .	68
<b>Chapter 6. Integrated Server Administration . . . . .</b>	<b>71</b>
6.1 Server Management functions . . . . .	73
6.2 Disk management functions . . . . .	75
6.3 User enrollment functions . . . . .	77
6.4 iSCSI: Network Server Host Adapter functions . . . . .	77
6.5 iSCSI: remote system functions . . . . .	78
6.6 iSCSI: service processor functions . . . . .	79
6.7 iSCSI: connection security functions . . . . .	79
<b>Chapter 7. Security, users, and groups . . . . .</b>	<b>81</b>
7.1 Security . . . . .	82
7.1.1 Authorization list . . . . .	82
7.1.2 Create an authorization list . . . . .	84
7.1.3 Change object permissions . . . . .	84
7.1.4 Cryptographic Services Key Management . . . . .	84
7.1.5 Cryptographic Coprocessor configuration . . . . .	87
7.1.6 Intrusion detection . . . . .	88
7.2 Users and Groups page . . . . .	92
7.2.1 Users . . . . .	93
7.2.2 Users not in a Group task . . . . .	94
7.2.3 Create user . . . . .	95
7.2.4 Change user . . . . .	95
7.2.5 Groups . . . . .	96
7.2.6 Create Group task . . . . .	97
<b>Chapter 8. Database and journal management . . . . .</b>	<b>101</b>
8.1 Working with database tasks . . . . .	102
8.2 Select which schemas to work with . . . . .	104
8.3 Generate SQL for objects . . . . .	109
8.4 Change query attributes . . . . .	113
8.5 SQL Performance Monitor . . . . .	117
8.5.1 Create a new SQL Performance Monitor . . . . .	117
8.5.2 Work with SQL performance monitors . . . . .	119
8.5.3 Detailed database monitor . . . . .	119
8.5.4 Example of database monitor analysis . . . . .	120
8.6 SQL Plan Cache . . . . .	130
8.7 Show SQL details for job . . . . .	135
8.8 Index advisor . . . . .	136
8.9 Health center . . . . .	140
8.10 Journal management . . . . .	145
<b>Chapter 9. Performance . . . . .</b>	<b>147</b>
9.1 Performance Data Investigator . . . . .	148
9.2 Collection Services . . . . .	148
9.2.1 How Collection Services works . . . . .	148
9.2.2 Configuring Collection Services . . . . .	149
9.3 Job Watcher . . . . .	152
9.3.1 How Job Watcher works . . . . .	152
9.3.2 Configuring Job Watcher . . . . .	153
9.4 Disk Watcher . . . . .	155
9.4.1 How Disk Watcher works . . . . .	156

9.4.2 Configuring Disk Watcher .....	156
9.5 Performance Data Investigator .....	160
9.6 Health indicators .....	164
<b>Chapter 10. File systems</b> .....	<b>167</b>
10.1 Integrated File System .....	168
10.1.1 What the Integrated file System is .....	168
10.1.2 Root file system .....	169
10.1.3 QopenSys file system .....	171
10.1.4 QDLS file system .....	173
10.1.5 QSYS.LIB file system .....	175
10.1.6 QFileSvr.400 file system .....	177
10.1.7 QOPT file system .....	178
10.1.8 QNTC file system .....	179
10.2 File shares .....	180
10.3 Create file share .....	182
<b>Chapter 11. Advanced Job Scheduler for     IBM i</b> .....	<b>185</b>
11.1 Accessing AJS from Systems Director Navigator .....	186
11.2 Scheduling a job .....	186
11.2.1 The work .....	187
11.2.2 The schedule .....	188
11.2.3 The notification .....	189
11.3 Job groups .....	189
11.4 Job dependencies .....	191
11.5 Notification and report distribution .....	192
11.6 Remote command jobs .....	192
11.7 Multiple scheduling environments .....	193
<b>Chapter 12. Plug-ins and add-ons</b> .....	<b>195</b>
12.1 Backup, Recovery, and Media Services .....	196
12.2 High-availability solution .....	198
<b>Related publications</b> .....	<b>201</b>
IBM Redbooks publications .....	201
Online resources .....	201
How to get Redbooks publications .....	201
Help from IBM .....	202
<b>Index</b> .....	<b>203</b>

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.

# Trademarks

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

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

AIX®	IBM Systems Director Active Energy	System i®
AS/400®	Manager™	System Storage™
BladeCenter®	IBM®	System x®
DB2®	iSeries®	System z®
Domino®	OS/400®	TotalStorage®
Electronic Service Agent™	PowerHA™	WebSphere®
i5/OS®	Redbooks®	xSeries®
	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

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

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

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

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

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

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

# Preface

In this IBM® Redbooks® publication we discuss IBM Systems Director Navigator for i, which is a Web console interface for IBM i administration where you can work with the Web-enabled tasks of System i® Navigator. IBM Systems Director Navigator for i includes a number of welcome pages that allow you to quickly find the task that you want to perform.

The IBM Systems Director Navigator for i interface is not just a set of URL addressable tasks, but is a robust Web console from which you can manage your IBM i system. However, the System i Navigator Tasks on the Web, which are a set of URL-addressable tasks, can be accessed by using the URL or from within the IBM Systems Director Navigator for i interface.

The IBM Systems Director Navigator for i console contains the following task categories for IBM i management. These categories are similar to those that you see within the System i Navigator client application.

- ▶ System: system status, system operator messages, history log, disk status, change password, and more
- ▶ Basic operations: messages, printers, printer output, and more
- ▶ Work management: active jobs, server jobs, output queues, active subsystems, and more
- ▶ Configuration and service: system values, time management, disk units, disk pools, and more
- ▶ Network: TCP/IP servers, lines, DNS servers, and more
- ▶ Integrated server administration: all virtual disks, remote systems, domains, and more
- ▶ Security: authorization lists, cryptographic services key management, intrusion detection, and more
- ▶ Users and groups: users, create users, groups, and more
- ▶ Databases: performance monitors, health center, databases, and more
- ▶ Journal management: journals, create a journal, and more
- ▶ Performance: investigate data and collections, and more
- ▶ File systems: integrated file system, file shares, create file share, and more
- ▶ Internet configurations: digital certificate manager, System i Navigator Tasks on the Web, and more
- ▶ Cluster resource services: display the list of nodes, delete the cluster, and more

The information in this book is intended to help you start using the Web-based console, IBM Systems Director Navigator for i, by providing you with a look at the new interface as well as tips for working with various parts of the new console.

## The team who wrote this book

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

**Brandon Schulz** is currently working for the IBM STG Lab Services and Training group in Rochester, Minnesota, focusing on multi-core, high-performance computing and cluster

application development. He previously worked on IBM System Director products including IBM System Director 6.1 and IBM System Director Virtualization Manager. Prior to working on IBM System Director, he worked on IBM Toolbox for Java™ and IBM iSeries® Access for Web.

**Giancarlo Omati** has been working in System i software technical support for the past 13 years. His work has focused on communication protocols and applications and high availability during that time. His skills include OS/400® system administration, System i Navigator, database, journaling, system maintenance and tuning, and general OS/400 system knowledge. He completed the SNA and TCP roadmap training as part of the IBM Rochester Software Support Internship program in November 2006 and co-wrote IBM Redbooks publication *IBM i 6.1 Technical Overview*, SG24-7713, in November 2008.

**Morten Buur Rasmussen** is a Senior IT Specialist at the IBM Systems and Technology Group Lab Services (former CTC) in La Gaude, France. He covers Europe and the Middle East in client-facing areas of IBM System i, database, and WebSphere® performance. Before joining the Lab Services, he worked for IBM Denmark and different European banks. He has 22 years of experience in the computing field. His areas of expertise are database technology, application development, and work management for IBM System i.

**Johnnie Talamantes** is currently a subject matter expert for Systems Director Navigator at the IBM System i Support Center in Rochester, MN. He has extensive experience with the System i Access family of products, including System i Navigator. He has written many technical presentations, team notes, and knowledge-base documents as a subject matter expert for this product and others.

**Claudio Villalobos** is an IBM IT Specialist for System i working at IBM Argentina. He has 18 years of experience in systems management, performance and tuning, software and hardware configuration, logical partitions (LPARs), backup and recovery, System i installation, and Backup Recovery Media Services (BRMS) installation and configuration. Claudio has experience on System/36, AS/400®, and System i platforms. In 2004 he joined IBM, where he currently provides technical support and performs system management. He has participated in the preparation of four IBM Redbooks publications, including *Going Green with IBM Systems Director Active Energy Manager 3.1*, REDP-4361; *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789-03; *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946; and *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

**Brian Younger** is a Product Manager for Pinnacle Business Systems, Inc. in Edmond, Oklahoma. He is the lead designer and a developer for the IBM Advanced Job Scheduler for IBM i. Prior to working on the Advanced Job Scheduler, he was a developer for the IBM Backup, Recovery, and Media Services for i product. His articles written for the IBM Systems Magazine include *Batch Job Dependencies with Advanced Job Scheduler*, *How to use Advanced Job Scheduler Enhancements*, *New CL Features in Advanced Job Scheduler*, *Seven Time-Saving Tips for AJS*, *AJS Features Outmatch Base Scheduler*, *Passing Parameters in Scheduled Jobs*, *In the Know with Manual Automation*, *Multiple Scheduling Environments Create a Buzz*, and *Report Distribution Using AJS and iSeries Navigator*.

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

Hernando Bedoya  
STG Lab Services and Training, Rochester

Scott Vetter  
International Technical Support Organization, Austin

Robert Bestgen  
Tim Rowe  
Diane Olson  
Tim Kramer  
Bunny Chaney  
Margaret Fenlon  
Charles Emig  
Jim Flanagan  
Jeff Meaden  
Dawn May  
Adam Clark  
IBM Development, Rochester

Gang SG Shi  
Cheng Min Chi  
IBM China

## Become a published author

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

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

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

[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

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

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

[ibm.com/redbooks](http://ibm.com/redbooks)

- ▶ Send your comments in an e-mail to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

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

Archived



# Overview of IBM Systems Director Navigator for i

Introduced in V6R1, the IBM Systems Director Navigator for i is a full-featured, Web-based console that provides a single, easy-to-use view of the IBM i system. Systems Director Navigator is the first Web-based system management console that is integrated into IBM i. There is nothing to purchase or install. It is included in every V6R1 IBM i operating system.

IBM Systems Director Navigator provides a strategic tool for managing a specific IBM i partition. It can be used in conjunction with IBM Systems Director products to provide a multi-partition or heterogeneous (or both) cross-system management solution.

IBM Systems Director Navigator can be extended by business partners or IBM products using plug-ins. These plug-ins have the same look and feel as the rest of IBM Systems Director Navigator, thereby providing a seamless perspective into these areas.

## 1.1 Strategic positioning

IBM Systems Director Navigator is a merger of the best from existing client interfaces and the power and flexibility of Web-based solutions. The functions, groupings, and tasks will be very familiar to those used to the client-based System i Navigator interface. This rich set of functions is integrated into an Integrated Solutions Console (ISC) framework. This same framework is used by IBM Systems Director, thus giving the Director Navigator interface the same rich user experience as is available with the Systems Director products. It provides a single console for administering your IBM i products and applications by using single sign-on, thereby supporting access to all functions on the system within the Web interface.

IBM Systems Director Navigator's strategic design allows IBM i administrators to take advantage of a portlet-based infrastructure. The power of portlets allows you to work with a wide variety of data and tasks within a single *window*. This allows for multi-tasking and switching between active tasks. Portlets allow for partial refresh, which means that you have the option to refresh only the selected portlet or even a specific area within a portlet. The power of the ISC framework allows you to highly customize your environment to make it yours. You can define the pages important to you so that when you enter Systems Director Navigator, your startup pages immediately display with data relevant to you.

### **How do the Navigator interfaces compare**

Traditionally, IBM i systems management tasks have been administered from several user interfaces provided in the System i Access client application. With the evolution of Systems Director Navigator, we can compare each of these solutions to see why IBM Systems Director Navigator is well positioned to be the future interface for IBM i systems management.

#### ***System i Navigator client application***

The Navigator client interface is the graphical interface that has been available for many years. This interface is only available to users who have installed it on their PCs. It is a traditional client/server application. Also:

- ▶ It is a Microsoft® Windows® technology-based application.
- ▶ It must be installed on a PC or workstation.
- ▶ The System i Access for Windows client product is required to obtain this interface.
- ▶ It includes a Management Central feature to:
  - Manage multiple systems.
  - Monitor the systems in your environment.
  - Send PTFs to systems in your environment.
- ▶ Some functions require the 5761XW1 product license (per workstation).
- ▶ It provides graphical charting functions within database and disk management tasks.

#### ***System i Access PC5250 emulator client application***

The PC5250 emulator is the traditional *green screen* command-line interface. It allows you to access the system using IBM i CL commands. Also:

- ▶ The emulator is a Microsoft Windows technology-based application.
- ▶ It must be installed on a PC or workstation.
- ▶ The System i Access for Windows client product provides this interface.
- ▶ It is a non-graphical user interface.
- ▶ It requires the 5761XW1 product license (per workstation).



### ***IBM Systems Director Navigator for i***

This is the new, strategic Web-based graphical interface. It allows users to access and manage jobs and resources using a standard Web browser. There is no additional software to install or configure. Also:

- ▶ You can use the existing Web browser installed on your workstation.
- ▶ It is a Web-based portlet application.
- ▶ It is integrated into the IBM i base operating system (5761SS1, Option 3).
- ▶ It includes over 300 Web-enabled tasks similar to System i Navigator.
- ▶ It is the only location for new IBM i management plug-ins, such as Power HA and IBM Performance Tools.
- ▶ It can work with many different tasks at once (portletlets within the console's work area).
- ▶ All Web-enabled tasks are still URL addressable and can be accessed individually as servlets or as portlets within the new Web console.

IBM Systems Director Navigator contains a very rich set of Web-enabled functions and tasks, and this set continues to increase rapidly. However, there are certain functions currently available in the client navigator applications that are not yet available through the Director Navigator Web interface. These functions are:

- ▶ Management Central functions
- ▶ Database graphics and charting functions
  - Visual Explain
  - Database navigator
- ▶ Disk Management graphics and charting functions
  - Disk unit/ASP capacity chart
  - Disk unit physical location

If you use these functions, you must continue to utilize the System i Navigator client application for V6R1.

## 1.2 What is included in Systems Director Navigator

The IBM Systems Director Navigator for i console is a Web-based interface for managing your IBM i system. All that is required on your PC or workstation is a Web browser. The Systems Director Navigator interface is installed and available to use out-of-the-box in V6R1. Since it is a complete Web-based system management tool, there is no need to remember specific URLs for each task. Each task is available via an easy-to-use, menu-driven interface. Instead, you only must remember one URL, `http://<systemname>:2001`, where `<systemname>` is the name of your IBM i system. This directs you to the main console window, as shown in Figure 1-1.

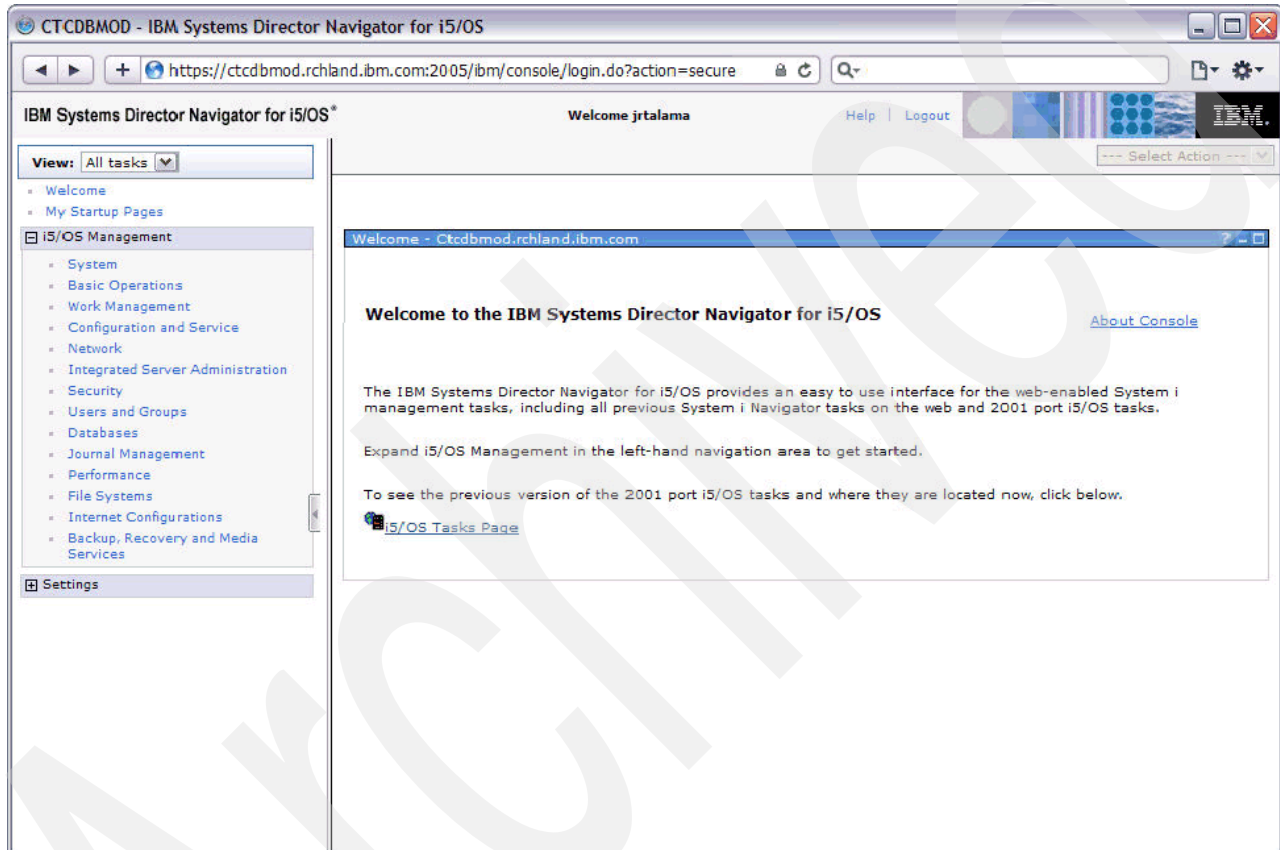


Figure 1-1 IBM Systems Director Navigator for i

**Note:** In previous releases, the port 2001 URL for your system would have taken you to the IBM i Task Page, which contained links to Web-based IBM i administration interfaces. This task page has been replaced with the Systems Director Navigator interface. If you still wish to access these Web interfaces, you can click the **i5/OS Tasks Page** link located on the Welcome page.

Systems Director Navigator has over 300 tasks available through the Web interface. These tasks are grouped into 13 categories. For those users who have been used to managing their systems with the Navigator for i client interface, they should feel comfortable with the Web-based interface, as the tasks have been grouped and organized in a similar manner. The categories appear in the task navigation when you open the i5/OS® Management, as shown in Figure 1-2.



Figure 1-2 Task categories

The system management tasks are divided into 13 logical groupings. Click each group to view the Welcome page for that group. The welcome page contains the most commonly used tasks for that group, such as:

- ▶ **System:** commonly used system tasks, including viewing system status, messages, disk status, and passwords. Click **Application Administration** to control which users are able to use specific functions within this interface.
  - **Basic Operations:** basic tasks for an individual user to easily view messages, printer output, printers, and their user jobs.
  - **Work Management:** Control the work performed on the system including support for interactive and batch work. Distribute resources to ensure that applications run well.
- ▶ **Configuration and Service:** functions to help you manage both hardware and software on your system, including system values, time management, and disk units.
- ▶ **Network:** Control TCP/IP configuration, remote access services, networking servers, and IP policies. Work with and manage your DNS servers.
- ▶ **Integrated Server Administration:** Manage servers hosted by IBM i including virtual disks, network server host adapters, remote systems, service processors, enrolling and managing users and groups, and domains.
- ▶ **Security:** Functions allow you to plan and implement security on your system, authorization lists, object permissions, cryptographic functions, and intrusion detection.
- ▶ **Users and Groups:** Display and manage a complete list of the users and groups. Create, delete, and modify users and groups on this system.
- ▶ **Databases:** graphical interface to perform many common administration database operations, access and modify objects, work with performance monitors, and access the Health Center.
- ▶ **Journal Management:** Work with journals and journal receivers.

- ▶ Performance: Tools to help manage the collection of performance data on the system. View the collection data using a powerful graphical charting tool to investigate potential performance issues.
- ▶ Files Systems: links to the most commonly used file systems tasks, Integrated File System, i5/OS Netserver, and Network File System.
- ▶ Internet Configurations: IBM i Internet configuration tasks, HTTP Server, IBM Web Administration GUI, Digital Certificates, IBM IPP, Web-based Help Server, and Internet Setup Wizard.

### IBM licensed program products plug-ins for Systems Director Navigator

The IBM Systems Director Navigator interface is a pluggable interface that allows other products to be added into the Director Navigator console. IBM has several licensed program products (LPPs) that are plug-ins to the console. This makes these functions seamlessly available via the Web-based console. When one of these LPPs is installed (and the HTTP Admin server has been restarted), the Web console automatically detects this new function and makes it available in the console, as shown in Figure 1-3.



Figure 1-3 i5/OS Management: BRMS plug-in

The following IBM LPPs have Systems Director Navigator Web interfaces. When these LPPs are installed, they are made available in the console:

- ▶ Backup Recovery and Media Services (BRMS) 5761-BR1  
<http://www.ibm.com/systems/i/support/brms>
- ▶ System i High Availability Solutions Manager (iHASM) 5761-HAS
  - High Availability Solutions Manager (solution-based) GUI
  - Cluster Resource Services (task-based) GUI
  - <http://www.ibm.com/systems/power/software/availability/i5os.html>
- ▶ IBM Performance Tools for i5/OS DiskWatcher 5761-PT1 (option 1)  
<http://www-03.ibm.com/systems/i/advantages/perfmgmt/ptools.html>
- ▶ JobWatcher 5761-PT1 (option 3)  
[https://www-912.ibm.com/i\\_dir/idoctor.nsf](https://www-912.ibm.com/i_dir/idoctor.nsf)



## Installation, configuration, and support tips

This chapter provides details on what product options will be needed to run IBM Systems Director Navigator for i, what is needed to gain access to the Web-based console, what server jobs are required, additional LPPs that can plug in to the Web-based console, support tips, and much more. This chapter describes the following topics:

- ▶ IBM i server
- ▶ Required jobs
- ▶ Required products
- ▶ Recommended PTFs and levels
- ▶ IBM license program products plug-ins for System Director Navigator
- ▶ Client-side requirements
- ▶ Web browser information
- ▶ How to access the console
- ▶ Support tips

## 2.1 IBM i server

As mentioned previously, the IBM Systems Director Navigator for i Web-based console in V6R1 is installed and available *out-of-the-box*. The product options mentioned in this chapter should already be installed on your IBM i server. These products are integrated into the base operating system. However, you should ensure that they are installed and up to date on PTF levels described.

## 2.2 Required jobs

At a minimum, the HTTP Admin Server must be running to use the console. The required jobs should be running under the QHTTSPVR subsystem, as seen in Figure 2-1.

```
Work with Active Jobs                                LP02F8F
                                                    09/02/09 11:08:42
CPU %:      .0      Elapsed time: 00:00:00      Active jobs: 218
Current
Opt  Subsystem/Job  User      Type  CPU %  Function      Status
-----
QHTTSPVR  QSYS            SBS      .0     DEQW
ADMIN     QLWISVR         BCI      .0     JVM-com.ibm.lw THDW
ADMIN     QTMHHTTP        BCH      .0     PGM-QZHBMAIN  SIGW
ADMIN     QTMHHTTP        BCI      .0     PGM-QZSRLOG   SIGW
ADMIN     QTMHHTTP        BCI      .0     PGM-QZSRHTTP  SIGW
ADMIN2    QLWISVR         BCI      .0     JVM-com.ibm.lw THDW
ADMIN3    QLWISVR         BCI      .0     JVM-com.ibm.lw THDW
ADMIN4    QLWISVR         BCI      .0     JVM-com.ibm.lw THDW

====>
F21=Display instructions/keys

Bottom
```

Figure 2-1 QHTTSPVR subsystem and admin server jobs

To start the HTTP Admin server, type the following at the IBM i command line:

```
STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)
```

## 2.2.1 IBM i required products

The following i products (which should be included in the default V6R1 system configuration) must be installed on your system for the console to operate and provide the 13 categories generally available. The product options:

- ▶ 5761SS1 option 3 - Extended Base Directory Support
- ▶ 5761SS1 option 12 - Host Servers
- ▶ 5761SS1 option 30 - QShell
- ▶ 5761SS1 option 33 - Portable Applications Solutions Environment
- ▶ 5761DG1 - IBM HTTP Server for i5/OS
- ▶ 5761JV1 - IBM Developer Kit for Java
- ▶ 5761JV1 option 8 - J2SE 5.0 32 bit
- ▶ 5761JC1 - IBM Toolbox for Java

## 2.2.2 Recommended PTF levels

The recommended PTF levels are suggested prior to working with the Web-based console. This should ensure that you will have a successful experience with the product and prevent any issues when initially accessing the interface for the first time.

**Note:** The following individual and group PTFs were the latest at the time that this document was written. We recommend checking what the latest PTF levels are on the Web site referenced.

### PTFs

At IBM i command line: DSPPTF LICPGM(5761SS1) SELECT(SI34446)

- ▶ Systems Director Navigator PTF:  
Product: 5761SS1, PTF# SI34446

At IBM i command line: WRKPTFGRP

- ▶ Group PTF: DB2® for System i  
SF99601 - Level: 10 (installed status, preferred)
- ▶ Group PTF: Java  
SF99562 - Level: 10 (installed status, preferred)
- ▶ Group PTF: IBM HTTP server for i  
SF99115 - Level: 9 (installed status, preferred)

IBM i Support Group PTFs by Release:

[http://www-912.ibm.com/s\\_dir/sline003.nsf/GroupPTFs?OpenView&view=GroupPTFs](http://www-912.ibm.com/s_dir/sline003.nsf/GroupPTFs?OpenView&view=GroupPTFs)

## 2.2.3 Client-side requirements

There is nothing to install on your workstation. All that is needed is a Web browser.

## How to access

Simply enter the following into your browser address:

`http://<systemname>:2001`

Where `<systemname>` is the name or TCP/IP address of your server.

**Note:** Your browser will be redirected to a secure path, similar to the following:

`https://<systemname>:2005/ibm/console/logon.jsp`

## 2.2.4 Browser security warning

When the user's browser application is used to access a server/Web site that provides a Secure Sockets Layer (SSL) connection, the server presents a certificate to the user's client software as proof of its identity. The client software must then validate the server's certificate before the server can establish the session. To validate the server certificate, the client software must have access to a locally stored copy of the certificate for the certificate authority (CA) that issued the server certificate.

If the server presents a certificate from a public Internet CA, the user's browser or other client software must already have a copy of the CA certificate. The server presents a certificate from a private local CA. Each user has the option to take the appropriate steps per browser instructions to install a copy of the local CA certificate.

When a client accesses `http://<systemname>:2001` the browser will automatically redirect to `https://<systemname>:2005/ibm/console`. The new console is secured with a local certificate that was shipped with the ADMIN LWI server. It is issued for a URL of `www.ibm.com` and has a Certificate Authority of `www.ibm.com`. This may cause the browser to display two warnings:

- ▶ The browser will not recognize the certificate authority as a trusted authority.  
This can be resolved by configuring the browser to trust the certificate authority.
- ▶ The URL that is typed in the browser does not match the URL that the certificate for which it was issued. Since we will never type `www.ibm.com` to access this console this presents a problem and may require a new certificate to be created.



Figure 2-2 shows a problem with the security certificate.

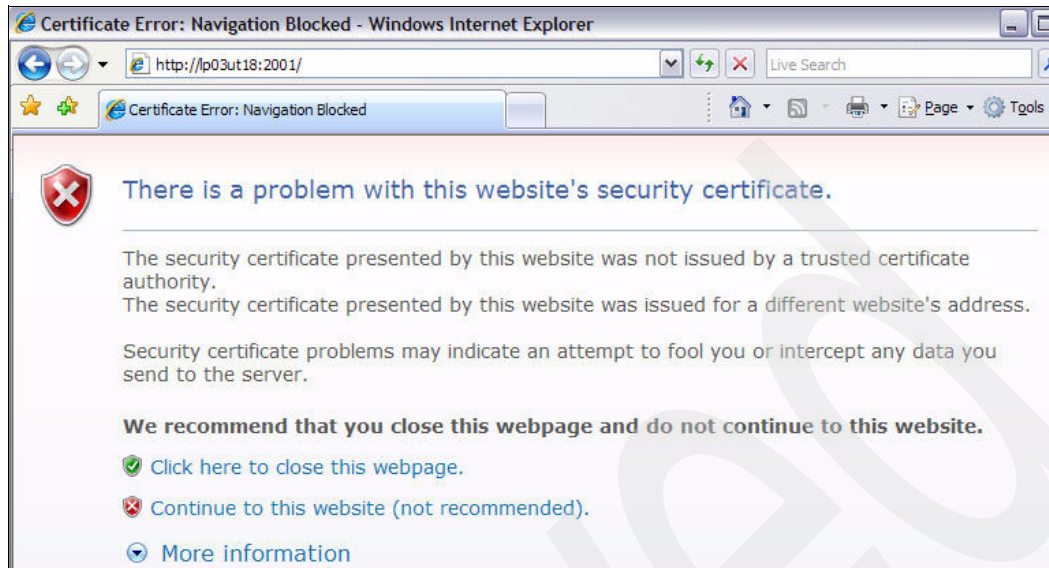


Figure 2-2 Problem with security certificate



Figure 2-3 Unknown authority



Figure 2-4 Security error

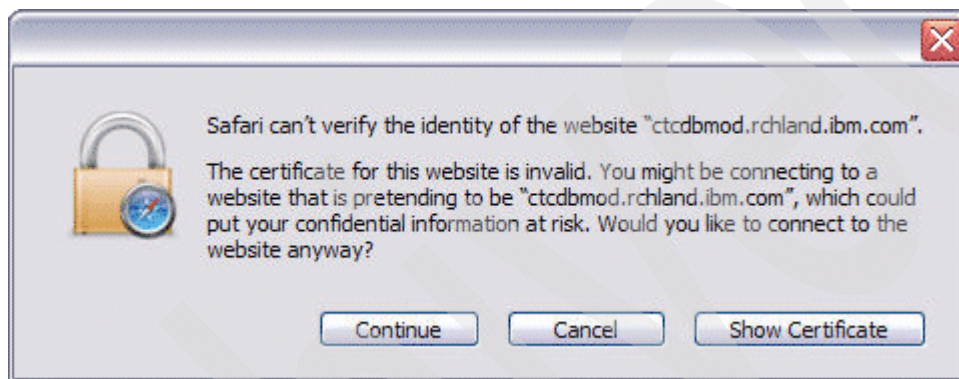


Figure 2-5 Safari verification

In review, if you prefer to not be presented with these warning messages, we recommend installing the certificate authority according to your browser specifications.

## 2.2.5 Access the console

As mentioned, to access the console enter the following URL in your browser address:

`http://<systemname>:2001`

You will be redirected to the IBM Systems Director Navigator for i logon prompt, as seen in Figure 2-6.

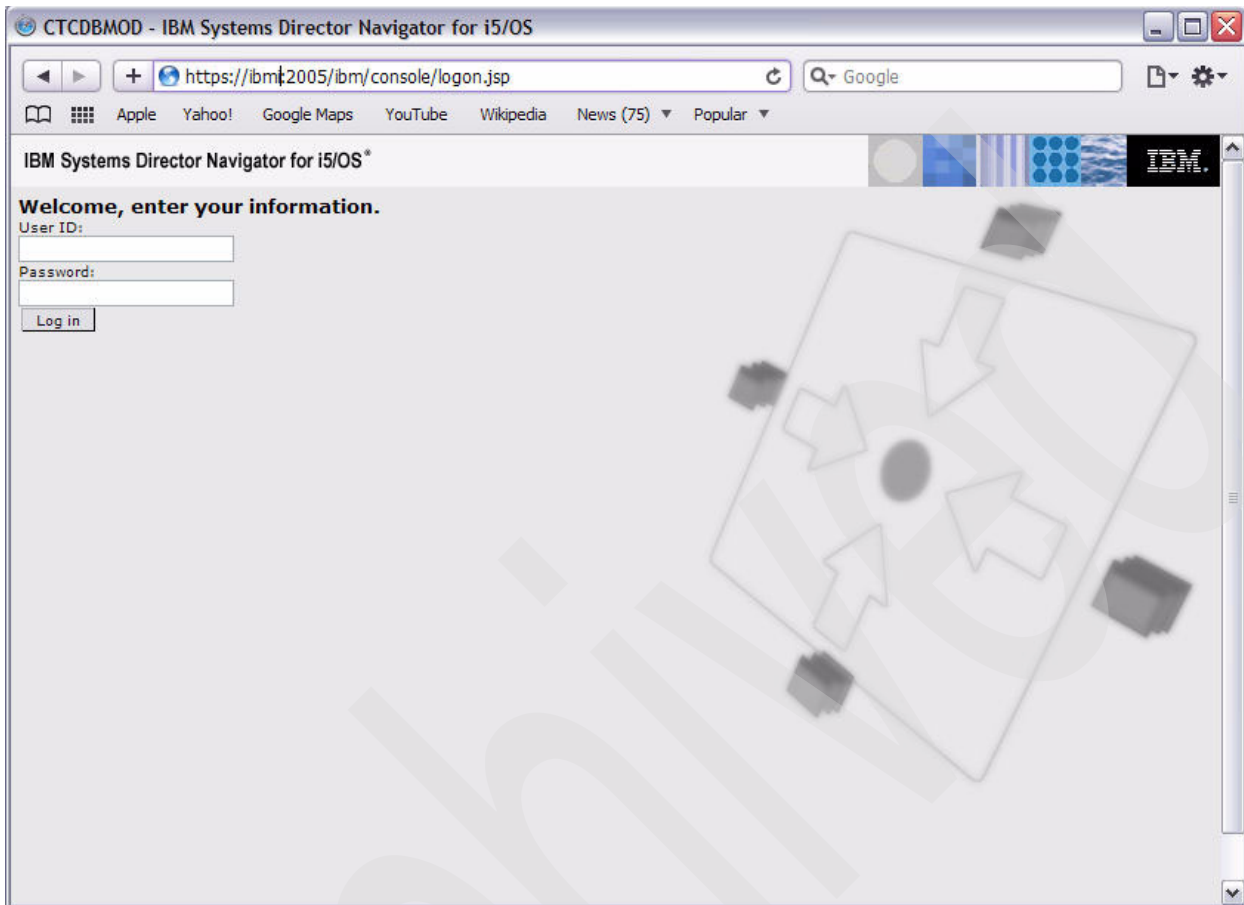


Figure 2-6 IBM Systems Director Navigator for i logon page

Enter your IBM i user ID and password and you will see the IBM Systems Director Navigator for i startup page, as shown in Figure 2-7.

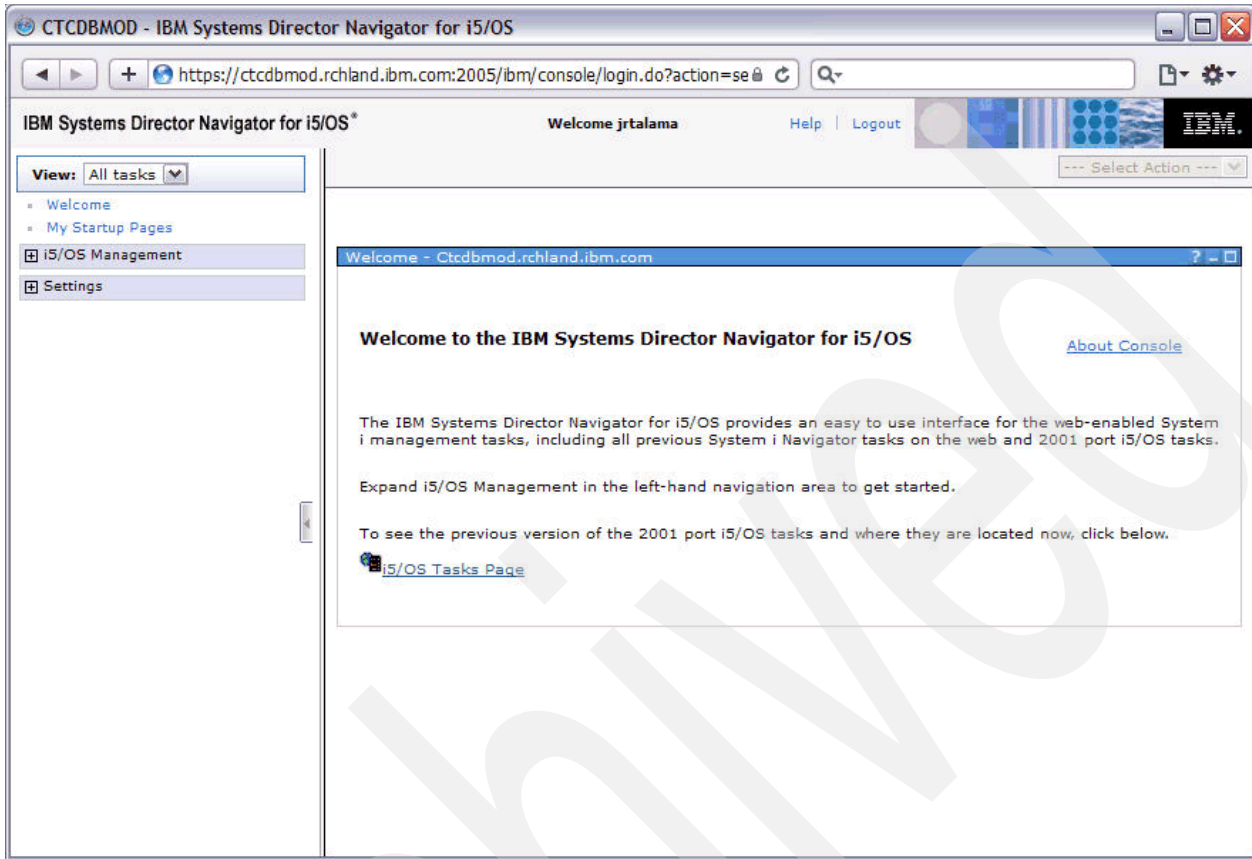


Figure 2-7 IBM Systems Director Navigator for i startup page

## 2.3 Support tips

In this section we review a few technical tips provided by our IBM i Rochester Support Center to assist in preventing any known or common issues when initially attempting to connect to the Web console.

- ▶ HTTP 404: page not found when trying to access the URL.

In most cases, the HTTP Admin server is not running. Refer to 2.2, “Required jobs” on page 8, and Figure 2-1 on page 8.

- ▶ HTTP 500.

In most cases, the required PTFs are not up to date or are not applied. Refer to 2.2.2, “Recommended PTF levels” on page 9, to ensure that the recommended PTFs are applied and installed.

- ▶ HTTP 403: forbidden.

In most cases, the required PTFs are not up to date or are not applied. Refer to 2.2.2, “Recommended PTF levels” on page 9, and ensure that the recommended PTFs are applied and installed.

- ▶ Connection is successful to IBM Systems Director Navigator for i and you are able to log on. However, the Navigation panel is empty.
  - a. IBM Systems Director Navigator for i makes use of the /tmp directory. In this incident, we determined that the profile \*Public did not have the proper data authority to use the /tmp directory in the IFS.
  - b. Make the appropriate changes.
- ▶ Connection is successful to IBM Systems Director Navigator for i and you are able to log on. However, when selecting a task in the Navigation panel to execute, the task has slow performance.
  - a. This can be caused by invalid DNS entries on their IBM i TCP/IP configuration. IBM Systems Director Navigator for i uses J2SE 5.0 32 bit, which impacts DNS behavior. J2SE 5.0 32 bit runs in the PASE environment. When IPv6 is enabled, which is the default, each DNS entry is checked to see whether it is IPv6 capable. When PASE hits a slow or unresponsive DNS entry, it waits for a reply or a time-out.
  - b. There are a couple things that can be done to resolve this issue:
    - i. Disable IPv6; this would require a reset of the TCP/IP stack and every connection to the system would be terminated. This would cause significant down time.
    - ii. Remove the bad DNS entry. This can be done without affecting connections to the system and is dynamic, meaning that changes take effect immediately.

We encourage you to research these and additional issues not mentioned, at our online knowledge base at the following Web site. Additionally, address any undocumented issues to the IBM i Rochester Support Center team for further investigation.

<https://www-947.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000027>

Archived



## The Web console

This chapter discusses the usage of the IBM Systems Director Navigator for i interface, where you can work with the Web-enabled tasks of System i Navigator.

## 3.1 Welcome page

Once logged in, you will see the IBM Systems Director Navigator welcome page. The Web tasks that used to be displayed on the i5/OS Tasks page on operating system releases prior to 6.1 are now accessible by clicking the **i5/OS Tasks Page** link at the bottom of the welcome page (Figure 3-1).

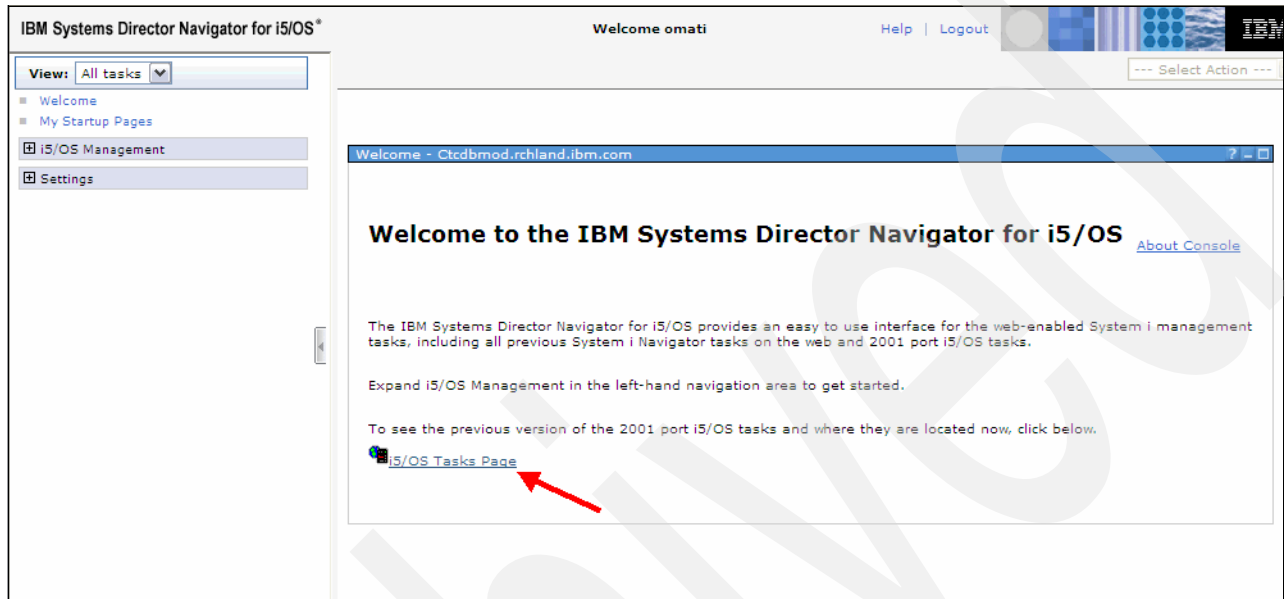


Figure 3-1 Welcome page

On the top of the page is shown the logged-in user, a help link, and a logout link:

- ▶ Help link: You will access the IBM Systems Director Navigator for i Help (Figure 3-2).

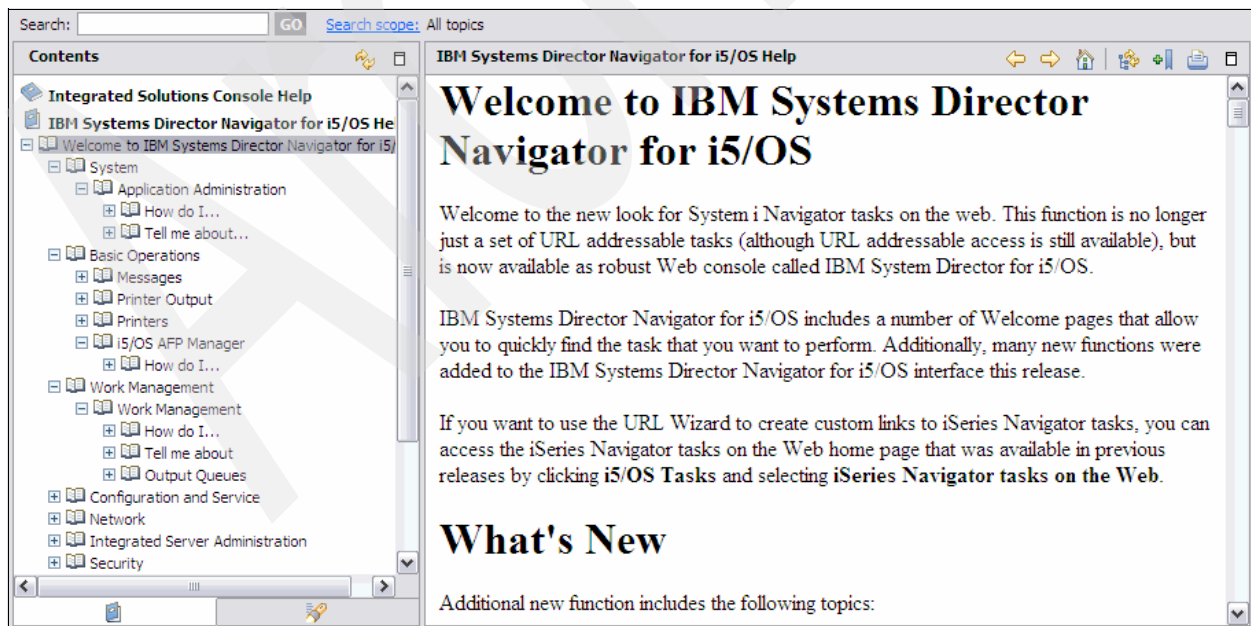


Figure 3-2 Help

- ▶ Logout: Disconnect the session.



## 3.2 Navigation within the interface

In the left navigation area, under i5/OS Management, is shown the main navigation view containing the list of all the components that you can access. Each of the task categories has a Welcome page that has links to the most common tasks. See Figure 3-3.

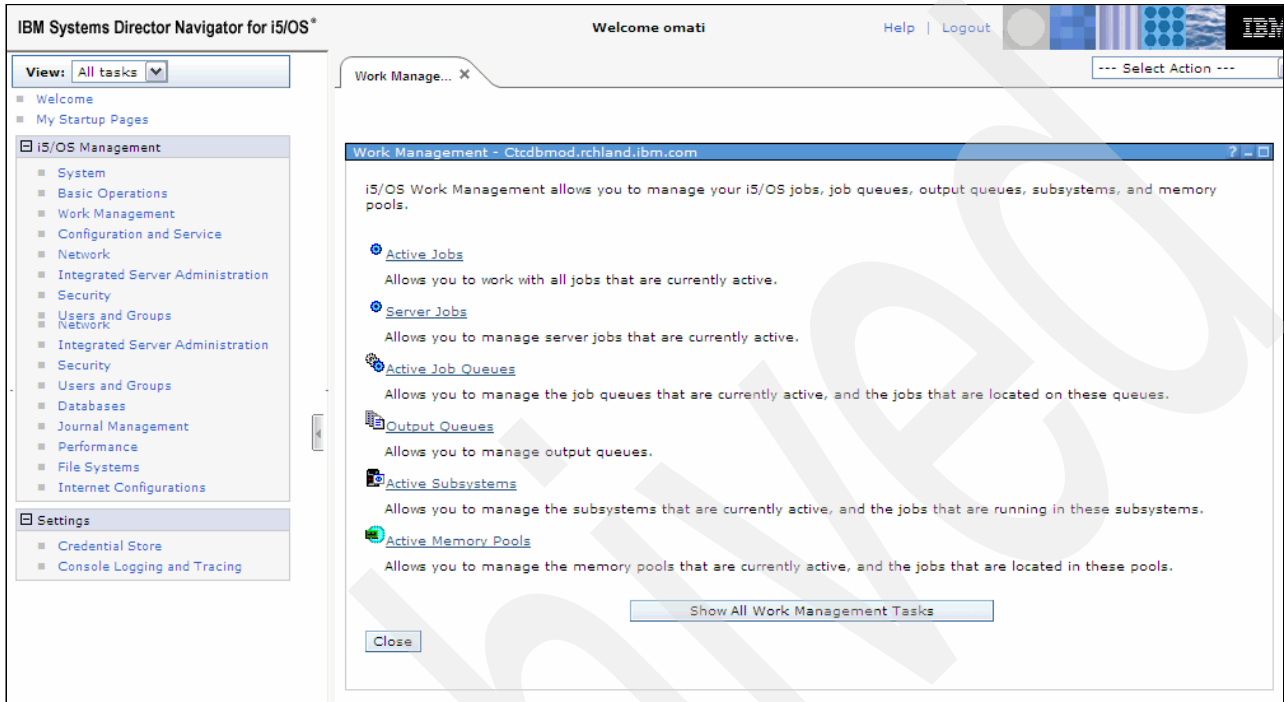


Figure 3-3 Work Management task

As you become familiar with the IBM Systems Director Navigator for i you will find that you can customize it to meet your needs and preferences. You can have a different welcome page to quickly find the task that you want to perform. Use the tabs to navigate between open pages or to close a specific page. If you open more than five pages, you can scroll to access other tabs. Each tab can be chosen to be in the preferred startup pages list (Figure 3-4) and then you can choose which one of them will be shown as the default page at the next login (Figure 3-5 on page 21). The original welcome page is always available by simply clicking the **Welcome** link on top of the left navigation panel. An overview of all the components and functions is available in Chapter 4, “Out of the box” on page 27.

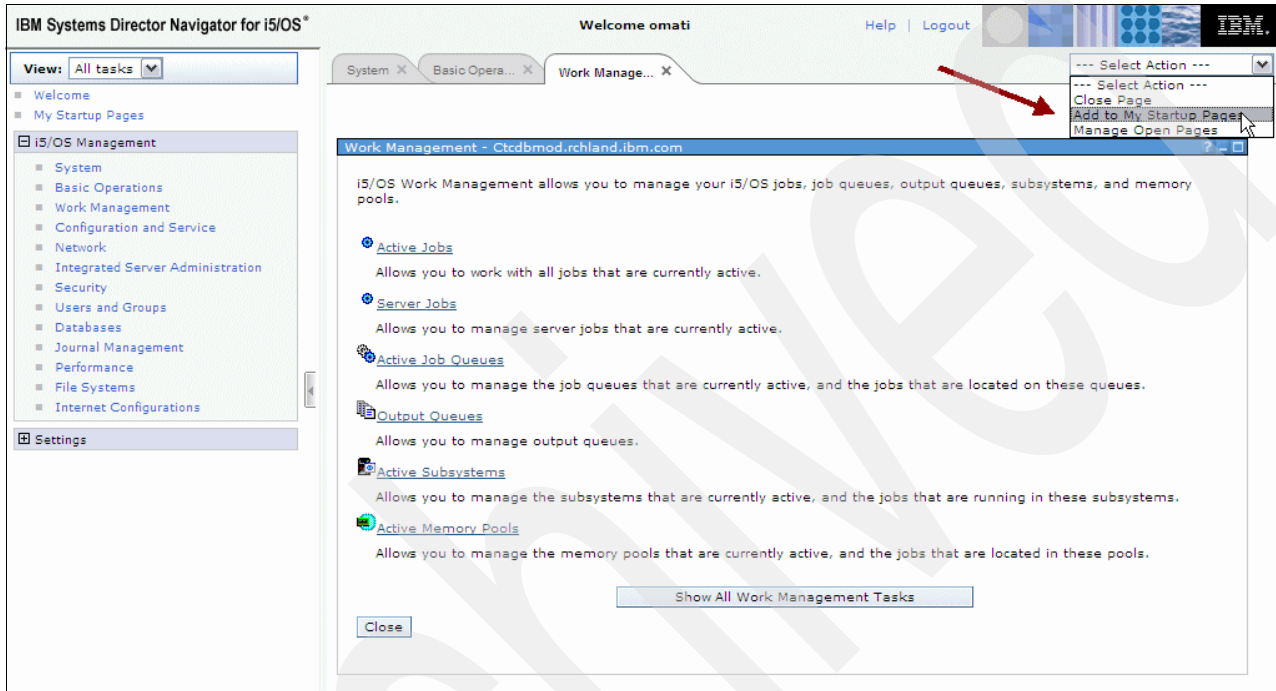


Figure 3-4 Choosing the current tab (Work Management) as the preferred startup page

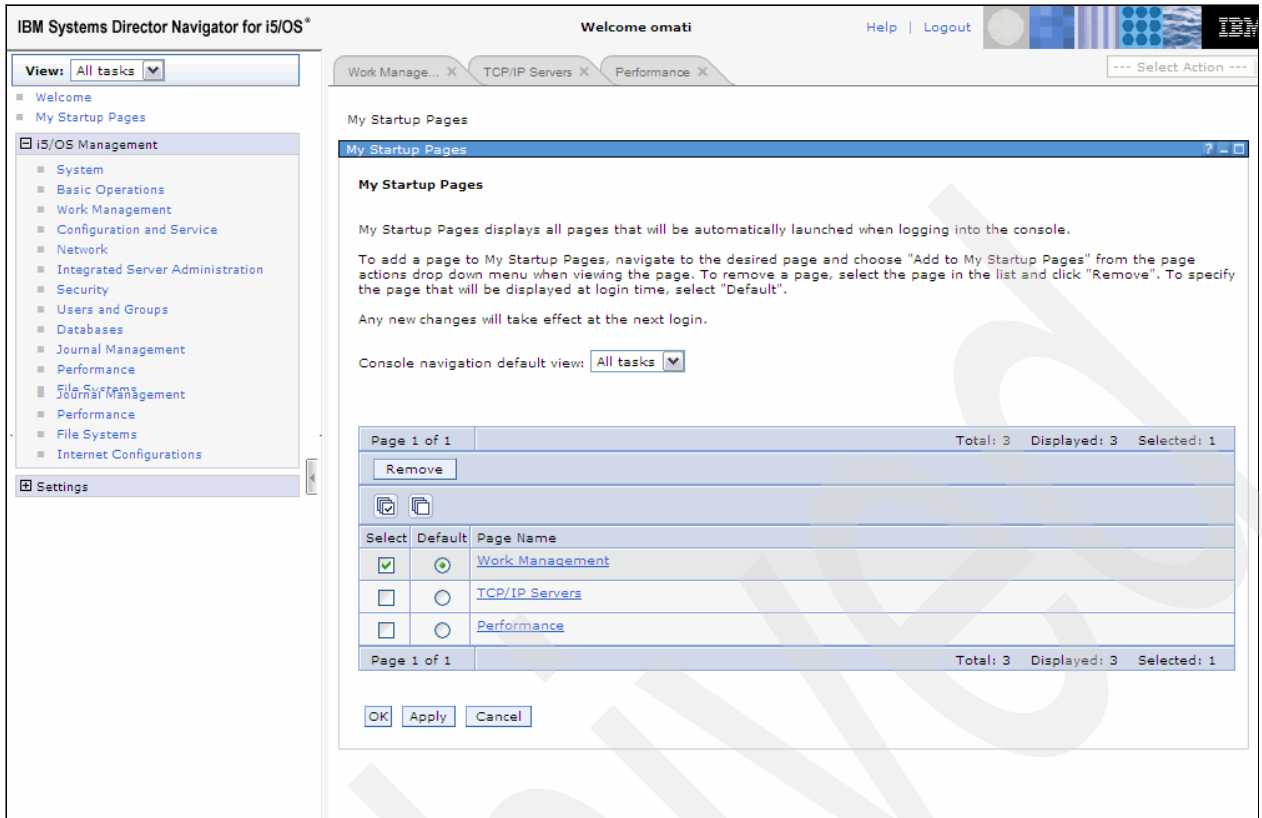


Figure 3-5 Default startup page selection

If multiple tabs are open and you must close one or all of them, you can manage the open pages as shown in Figure 3-6 and select those to be closed. See Figure 3-7 on page 22.

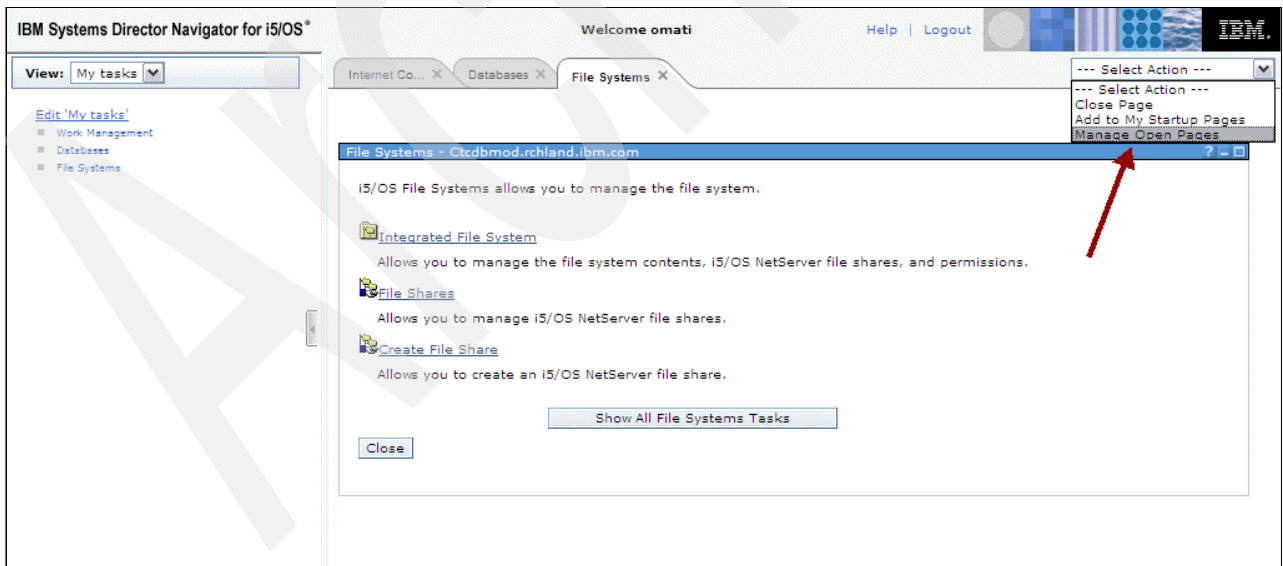


Figure 3-6 Manage open pages

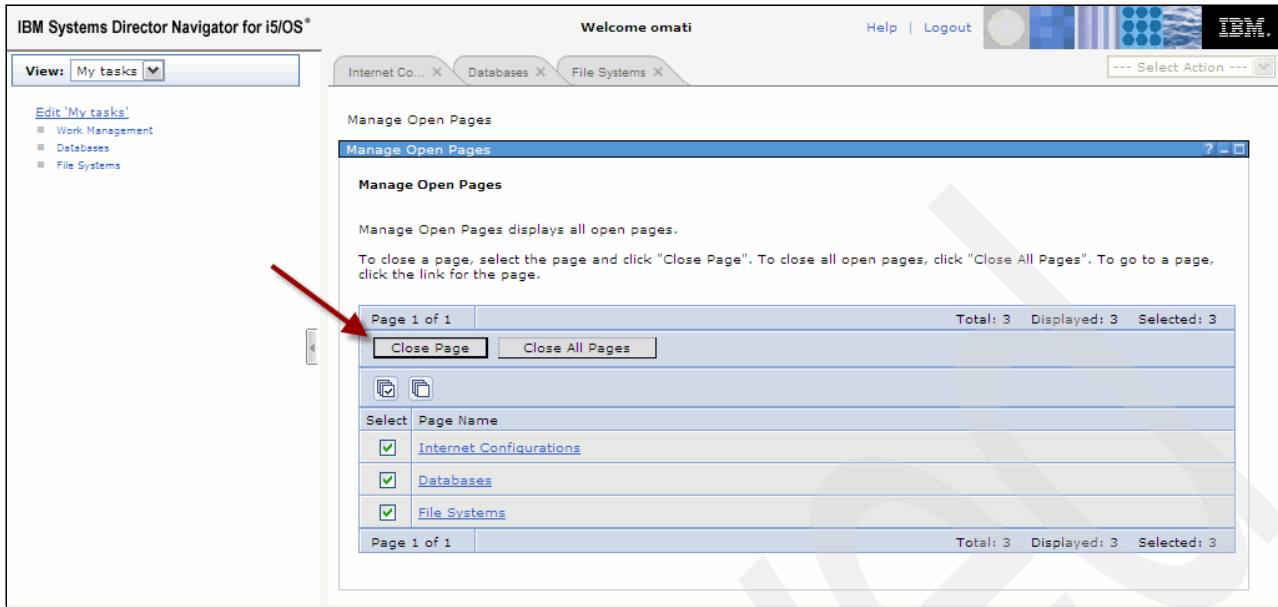


Figure 3-7 Select pages to close

On the top of the left navigation panel you can see a view bar with a drop-down list showing two options:

- ▶ All tasks: the entire list of the available tasks.
- ▶ My tasks: a customization of the default view. You can choose to create a view of some of the tasks, typically those that are highly used. The selected tasks will be listed in the left panel (Figure 3-8).

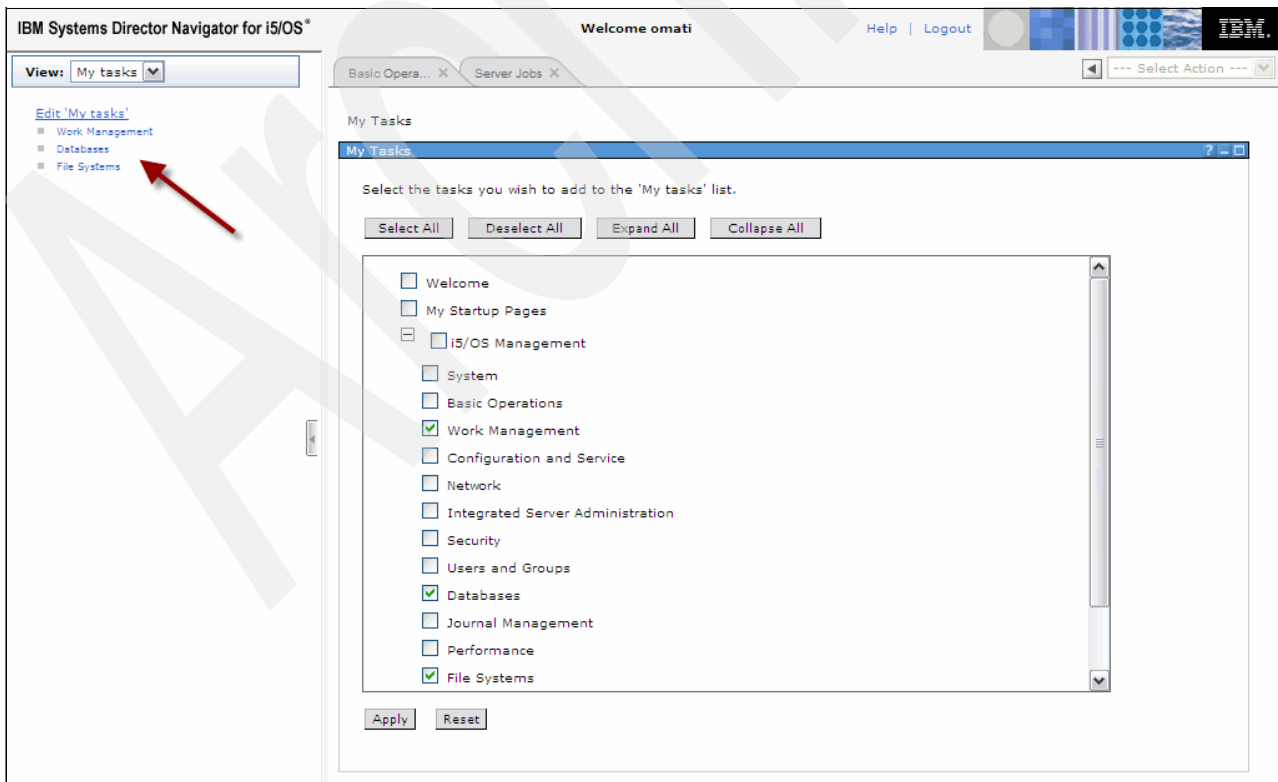


Figure 3-8 My Tasks setup

While navigating the IBM Systems Director Navigator for i, many of the IBM i resources are displayed in a Web table (for example, see the Server Jobs task shown below). By default, 15 items are displayed per page. The total number of items and pages are shown at the bottom of the table. To change the default table size, click the edit mode icon, which is a pen icon displayed in the upper right corner of the portlet, as shown in Figure 3-9, and choose how many rows per page you want (Figure 3-10).

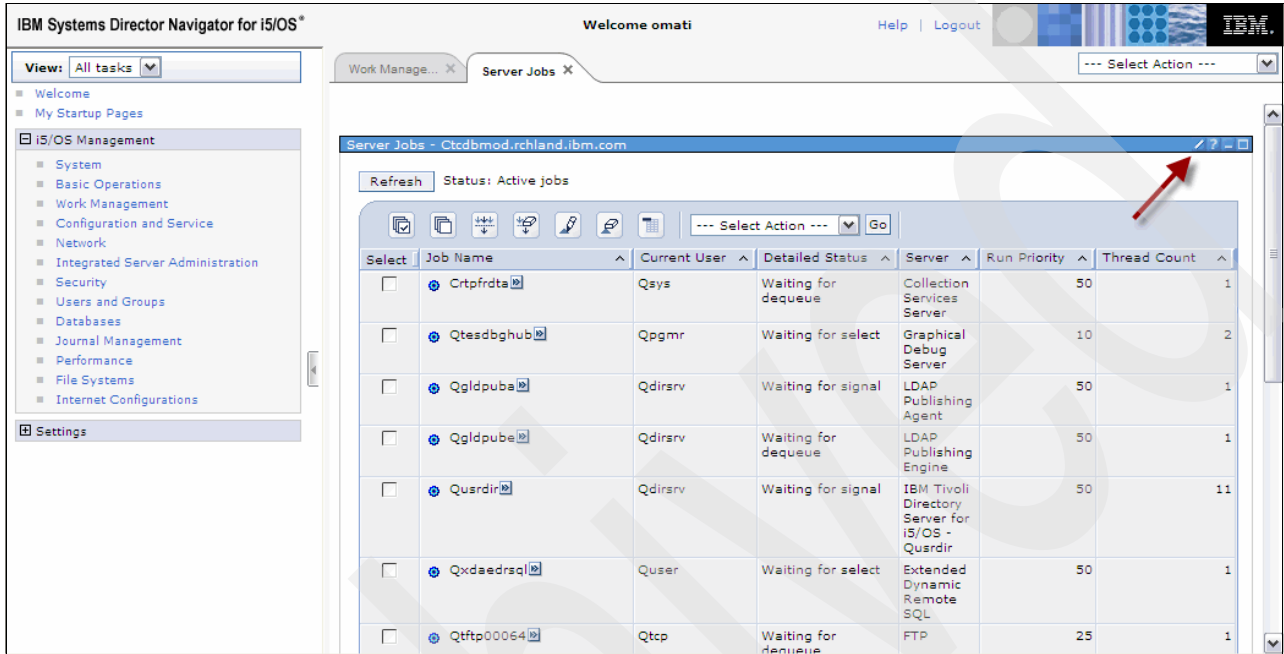


Figure 3-9 The pen icon

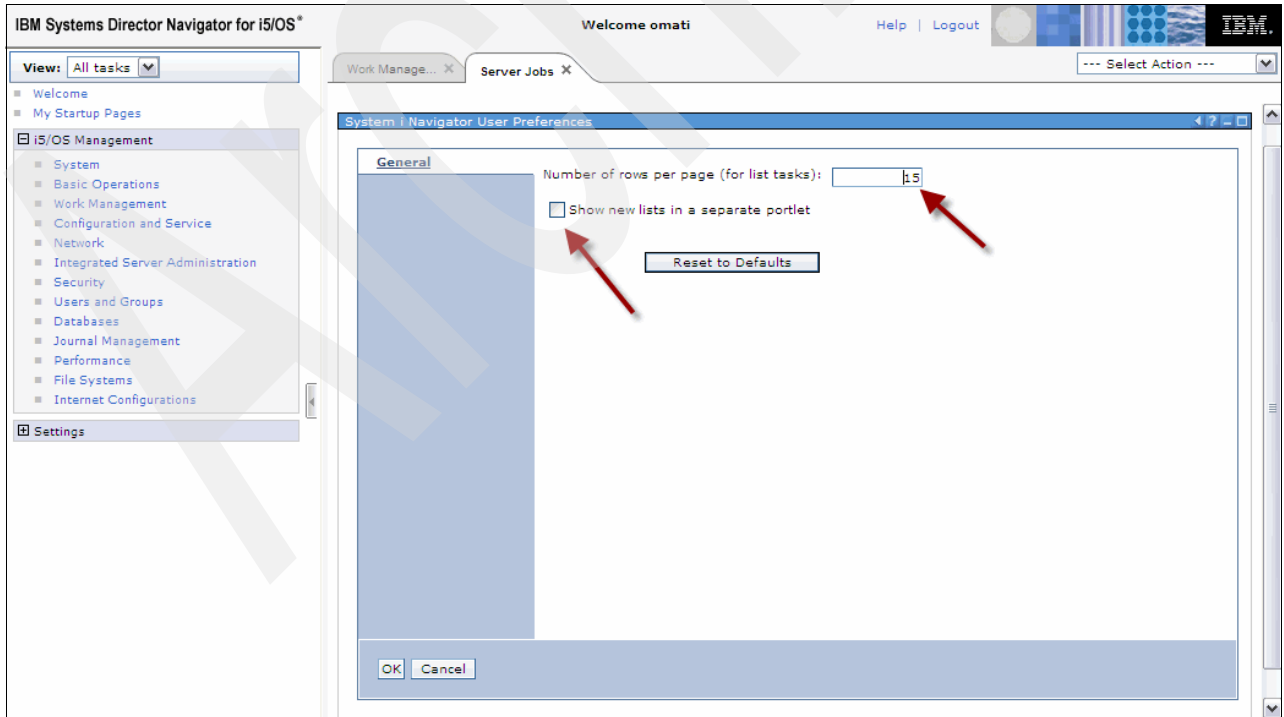


Figure 3-10 Select the number of rows

You can also add, remove, and sort columns by using the icons on top of the table or with the drop-down list (Figure 3-11).

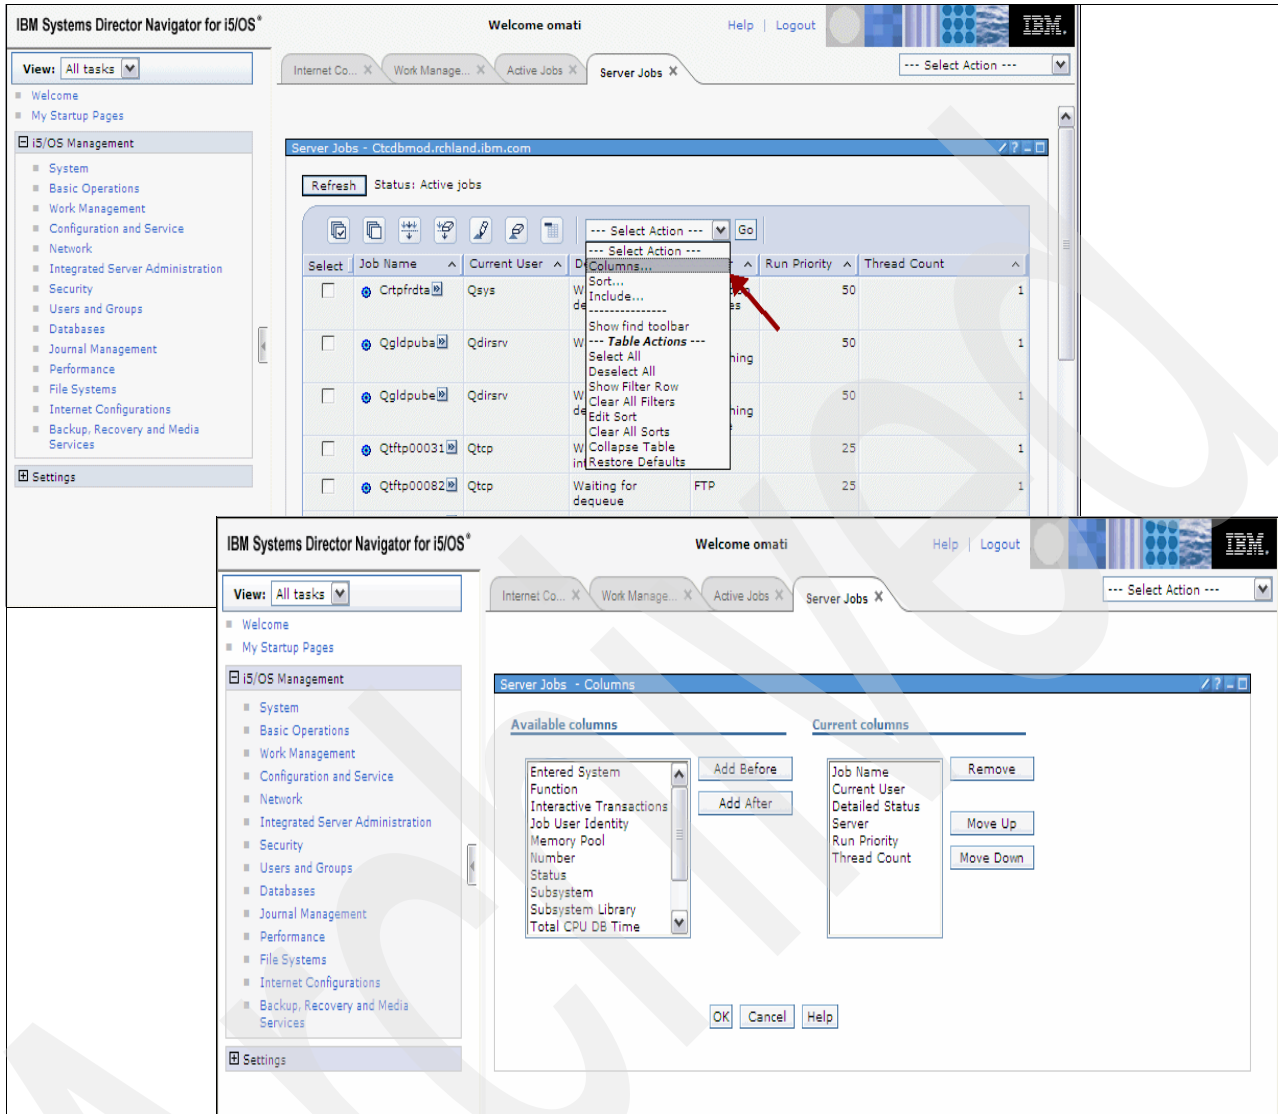


Figure 3-11 Sort columns

It is also possible to hide the left navigation frame by clicking **Hide Navigation**, as shown in Figure 3-12.

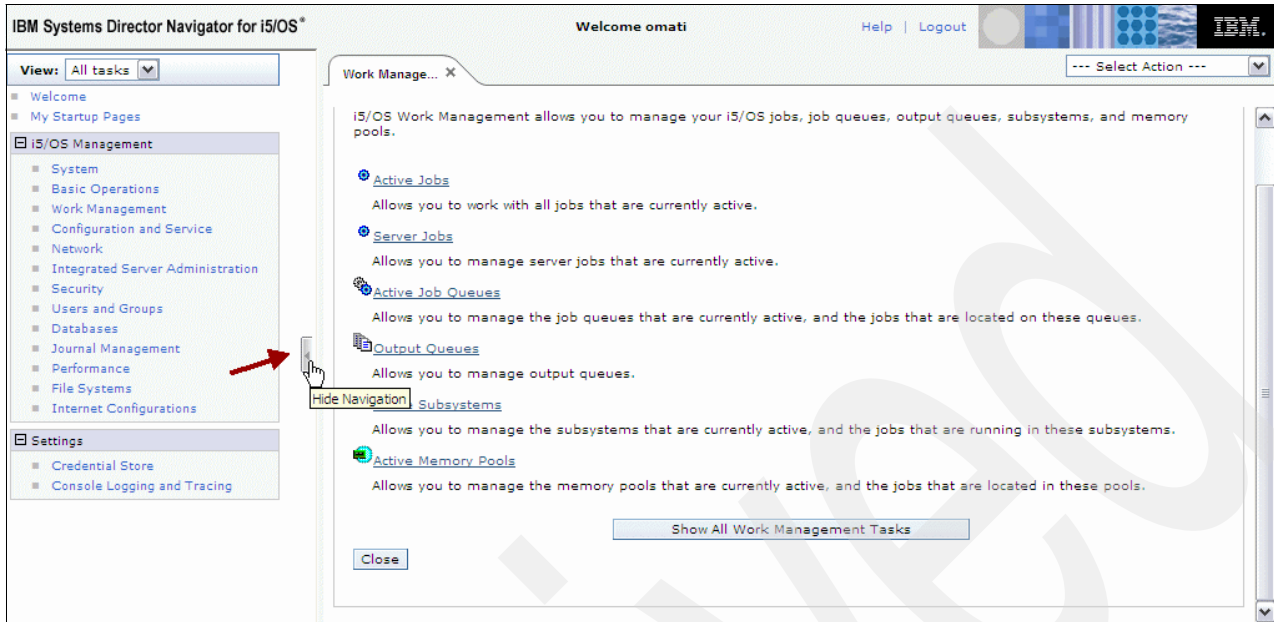


Figure 3-12 Hide Navigation

### 3.3 Settings

In the bottom part of the left navigation frame you can find the settings task, which contains the following options:

- ▶ **Credential Store**  
Save sensitive information in encrypted form.
- ▶ **Console Logging and Tracing**  
Used to control the settings for recording and maintaining runtime information for the console.

For more detailed information about the settings refer to Chapter 2, “Installation, configuration, and support tips” on page 7.

Archived





## Out of the box

In this chapter we introduce the new user interface and controls that are provided in IBM System Director Navigator for i.

## 4.1 New controls in IBM Systems Director Navigator for i

When you open IBM Systems Director Navigator for i you will see the same task categories that you see in the System i Navigator client application, as Figure 4-1 shows. The link to the Web applications that used to be displayed on the i5/OS tasks page, which previously appeared at the 2001 port, is at the bottom of the Welcome page (Figure 4-1).

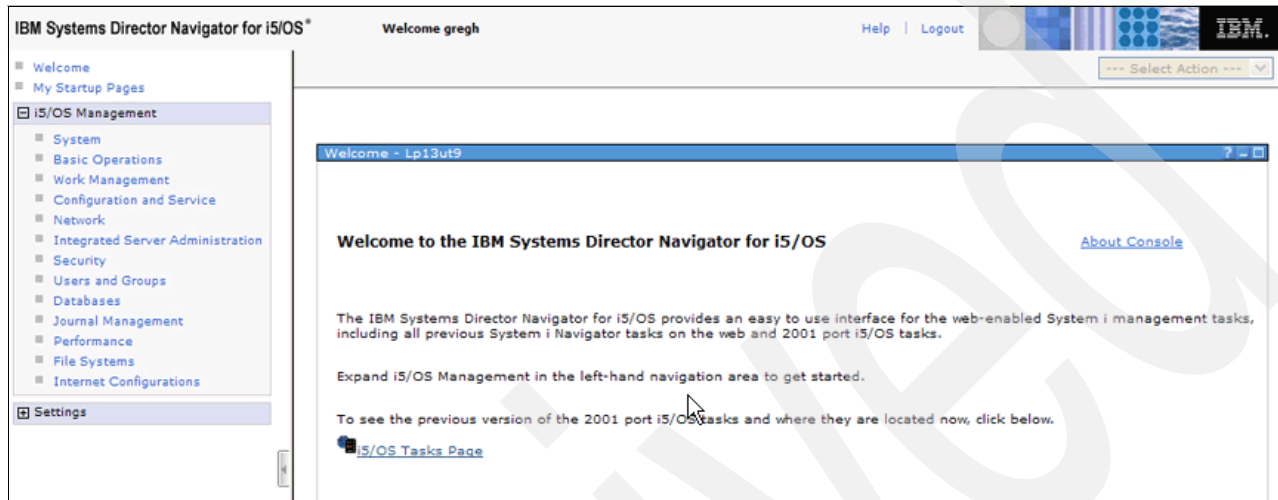


Figure 4-1 Welcome menu for the IBM System Director Navigator for i

## 4.2 Overview of the new user interface

In the left navigation area (Figure 4-1) you will find a list of the new tasks. When you select one of them with a click, a summary of sub-categories is displayed.

## 4.2.1 System task

Commonly used system tasks include viewing system status, operator messages, disk status, and passwords (Figure 4-2).

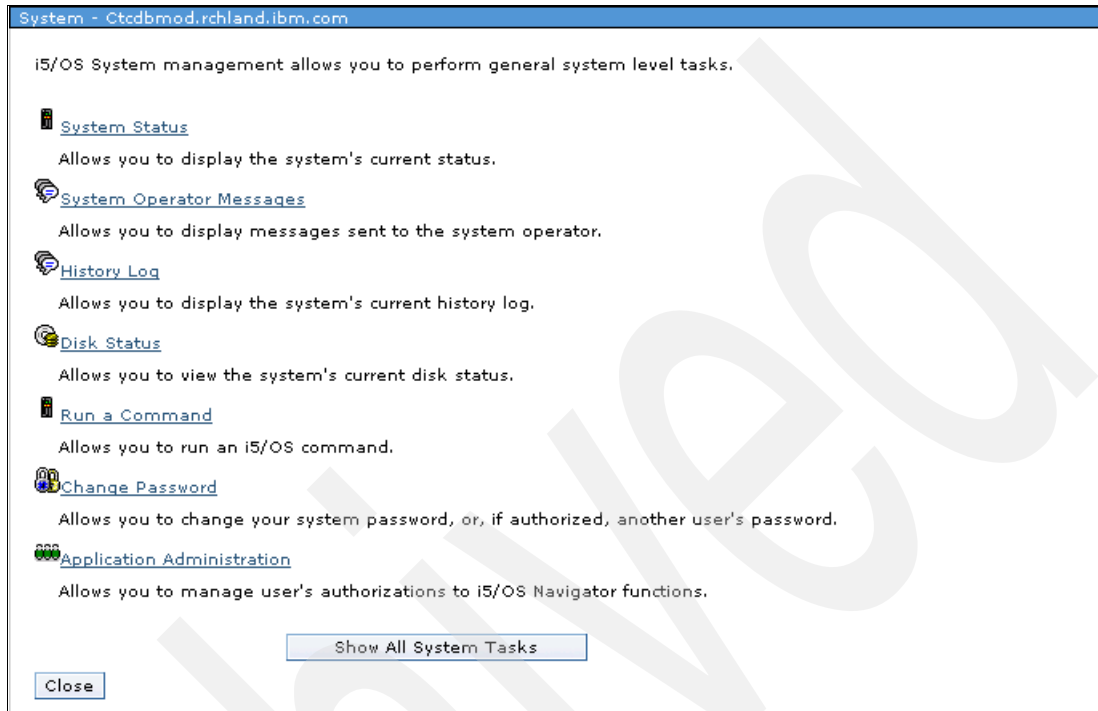


Figure 4-2 Task system functions

## 4.2.2 Basic Operations page

The Basic Operations page contains basic tasks to easily manage lists of messages, printer output, and printers (Figure 4-3).

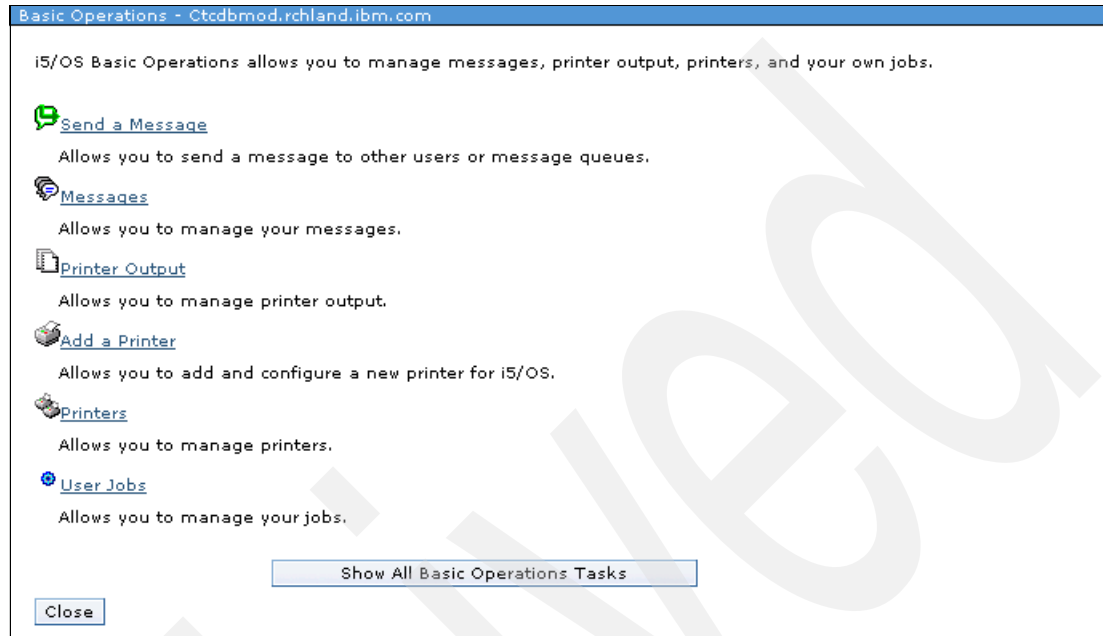


Figure 4-3 Basic Operations functions

## 4.2.3 Work Management page

The Work Management page includes tasks that allow you to control the work performed on the system, including support for interactive and batch work. These tasks also allow you to distribute resources to ensure that applications run well (Figure 4-4).

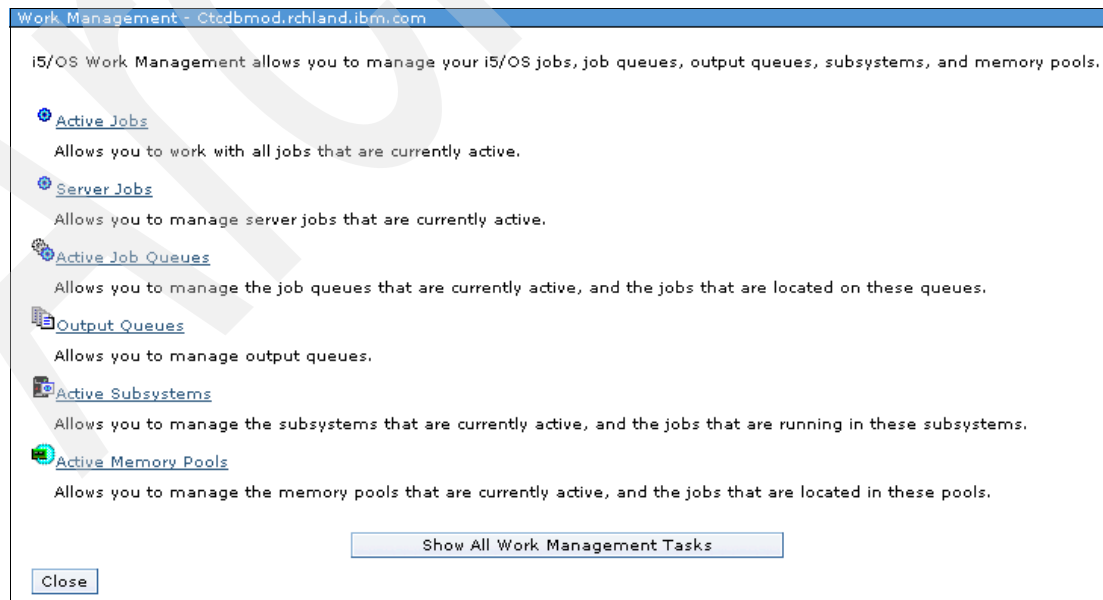


Figure 4-4 Work management functions

## 4.2.4 Configuration and service functions

Configuration and Service functions help you to manage both hardware and software on your system, including system values, time management, and disk units (Figure 4-5).

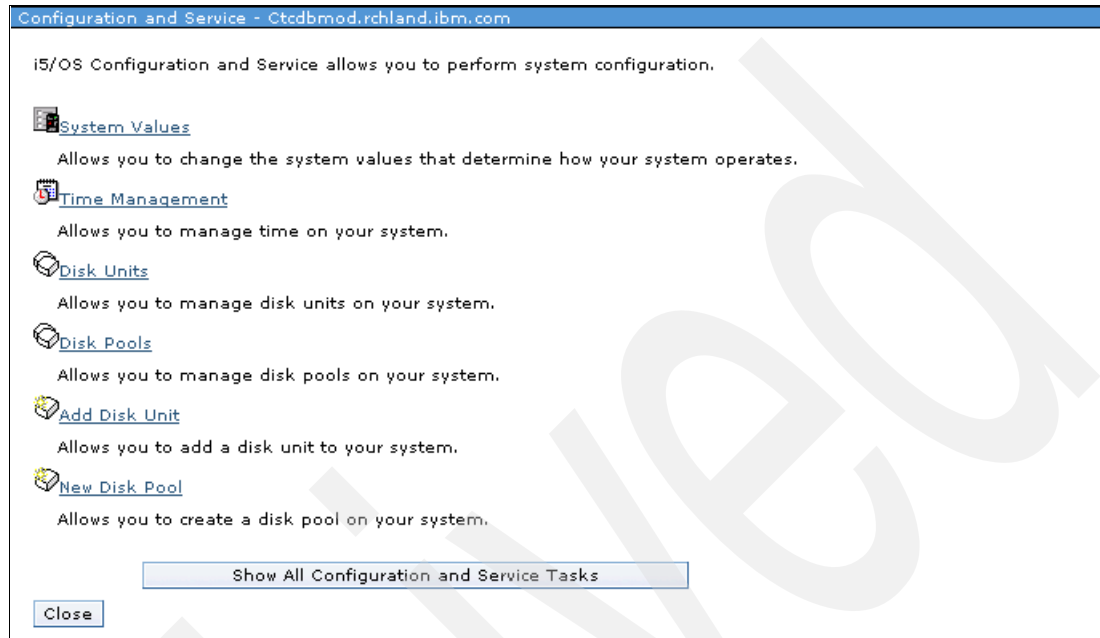


Figure 4-5 Configuration and service functions

## 4.2.5 Network

The Network page tasks allow you to control TCP/IP configuration, remote access services, networking servers, and IP policies.

### TCP/IP configuration

Use TCP/IP configuration to manage and monitor your TCP/IP network. TCP/IP configuration provides wizards for configuring both IPv4 and IPv6 lines, interfaces, and routes. TCP/IP configuration allows you to view and configure TCP/IP properties, such as domain and host name settings, and to start or stop TCP/IP interfaces. In addition, TCP/IP configuration provides you with utilities for troubleshooting your network, such as traceroute and ping.

## Domain Name System

The Domain Name System (DNS) is a distributed database of all of the IP addresses and their associated domain names. Each DNS server is responsible for knowing only a small fraction of all the addresses that exist. The servers work together, referring queries that they cannot answer to other servers that are authoritative. See Figure 4-6.

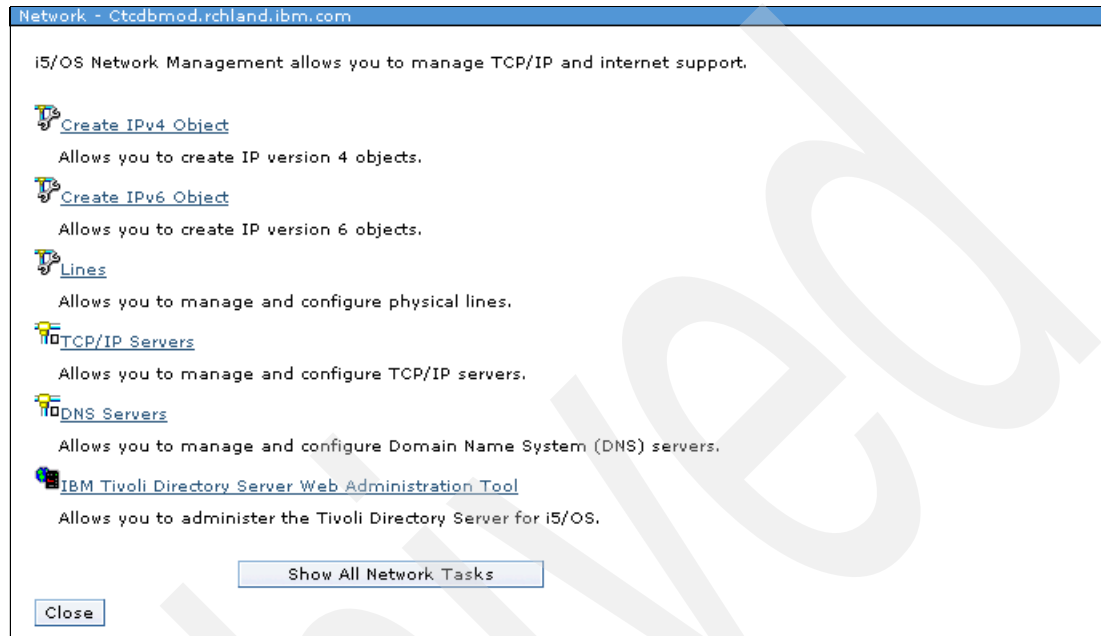


Figure 4-6 Network functions

### 4.2.6 Integrated server administration

Integrated server administration allows you to perform management tasks on servers that are hosted by i5/OS, including the following:

- ▶ Manage integrated BladeCenter® blade and System x® servers that use virtual disks provided by i5/OS.
- ▶ Manage Power server logical partitions that use virtual disks provided by i5/OS.
- ▶ Manage i5/OS virtual disks.
- ▶ Manage i5/OS configuration objects for iSCSI adapters.
- ▶ Manage i5/OS configuration objects for iSCSI attached servers.
- ▶ Manage enrollment of i5/OS users and groups to integrated Windows servers.

See Figure 4-7.

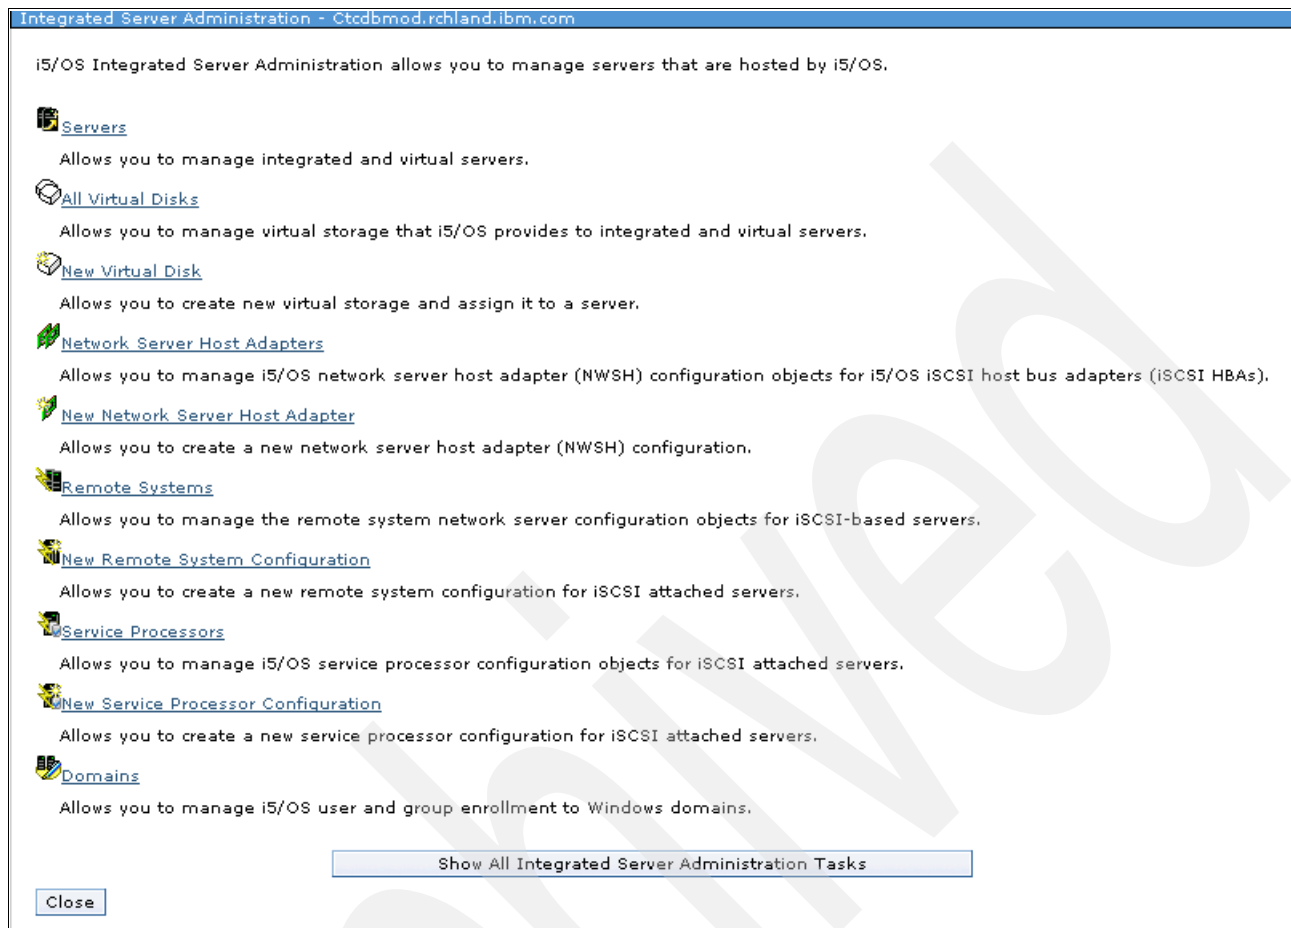


Figure 4-7 Integrated server administration functions

## 4.2.7 Security

The security interface functions allow you to plan and implement security on your system, including authorization lists, object permissions, and cryptographic functions.

### Authorization lists

An authorization list groups objects with similar security requirements together. An authorization list contains a list of users and groups and the authority that each has to the objects secured by the list.

### Object permissions

Permissions includes all of a user's authority (or the type of access) allowed to an object.

### Cryptographic Services Key Management

From IBM System Director Navigator for i, Cryptographic Services Key Management allows you to store and manage master keys and keystores. Since you will be exchanging sensitive data to manage master keys and keystores, we recommend that you use a secure session.

## Intrusion detection system

The intrusion detection system (IDS) notifies you of attempts to hack into, disrupt, or deny service to the system. These potential intrusions are logged as journal entries in the security audit journal and displayed in the new Intrusion Detection System graphical user interface (GUI). IDS also monitors for potential extrusions, where your system might be used as the source of the attack. See Figure 4-8.

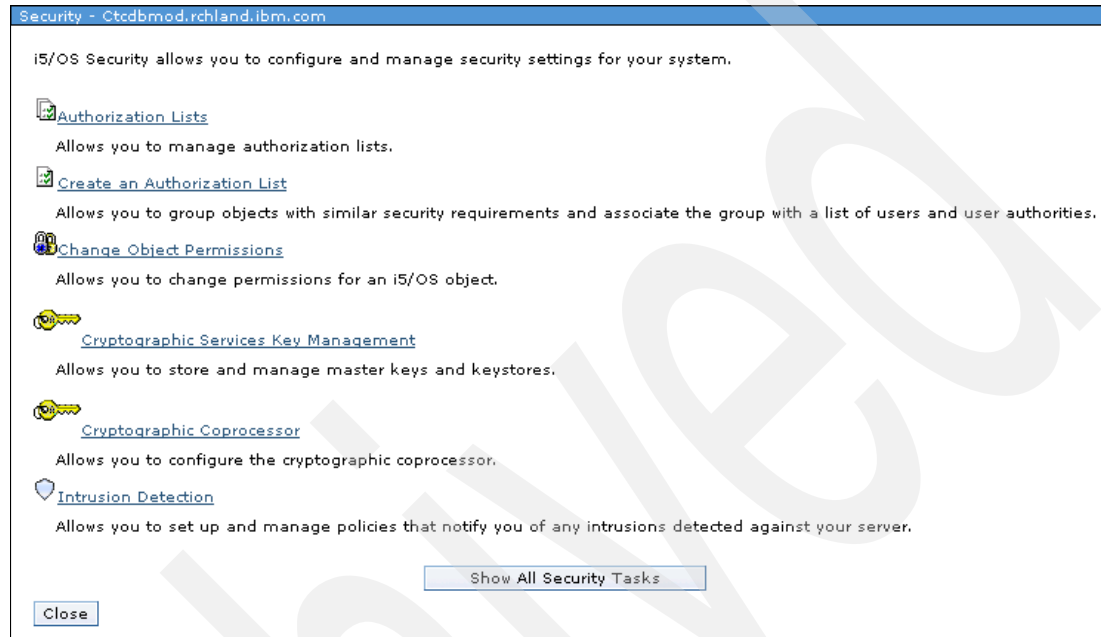


Figure 4-8 Security functions

## 4.2.8 Users and groups

The Users and Groups page allows you to display and manage a complete list of the users and groups, as well as create, delete, and modify users and groups. See Figure 4-9.

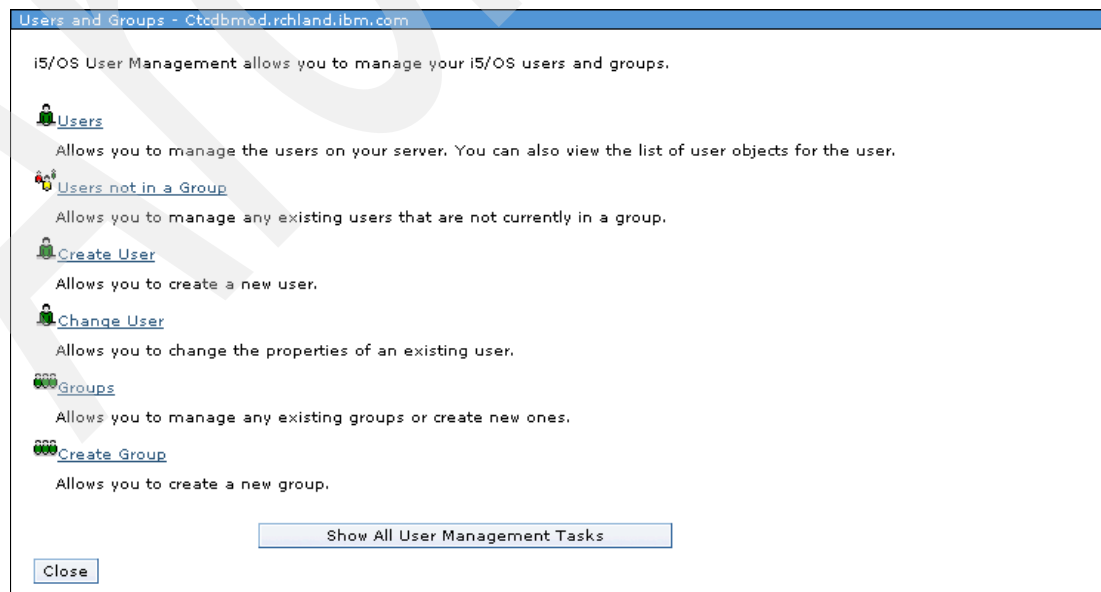


Figure 4-9 User and groups functions



## 4.2.9 Databases

Graphical interface to perform many common admin database operations, access and modify objects, work with performance monitors, and access the health center. The following functions are supported:

- ▶ On demand performance service
  - Customization of columns for analysis reports.
  - Save analysis reports to spreadsheet.
  - SQL plan cache event monitors.
  - Show active jobs that have run a specific SQL statement.
  - Show list of users that have run a specific SQL statement.
  - Visually explain an SQL statement while it is running.
  - Reset last used date and statistics for indexes.
  - Show indexes and MQTs for all tables in a schema.
- ▶ Health center
  - Display environmental limits for jobs.
  - Display file activity categories.
  - Journals and journal receivers added to overview.
  - Package and procedure added to size limits.
- ▶ Run SQL statements
  - Check statements for SQL portability.
  - Save script to source member on i5/OS.
  - Copy result window to spreadsheet.
  - Print the result window to a printer.

- Launch directly from .sql IFS objects.
- Provide GUI front-end to the WRKOBJLCK CL command to database objects, work with performance monitors, and access the health center (Figure 4-10).

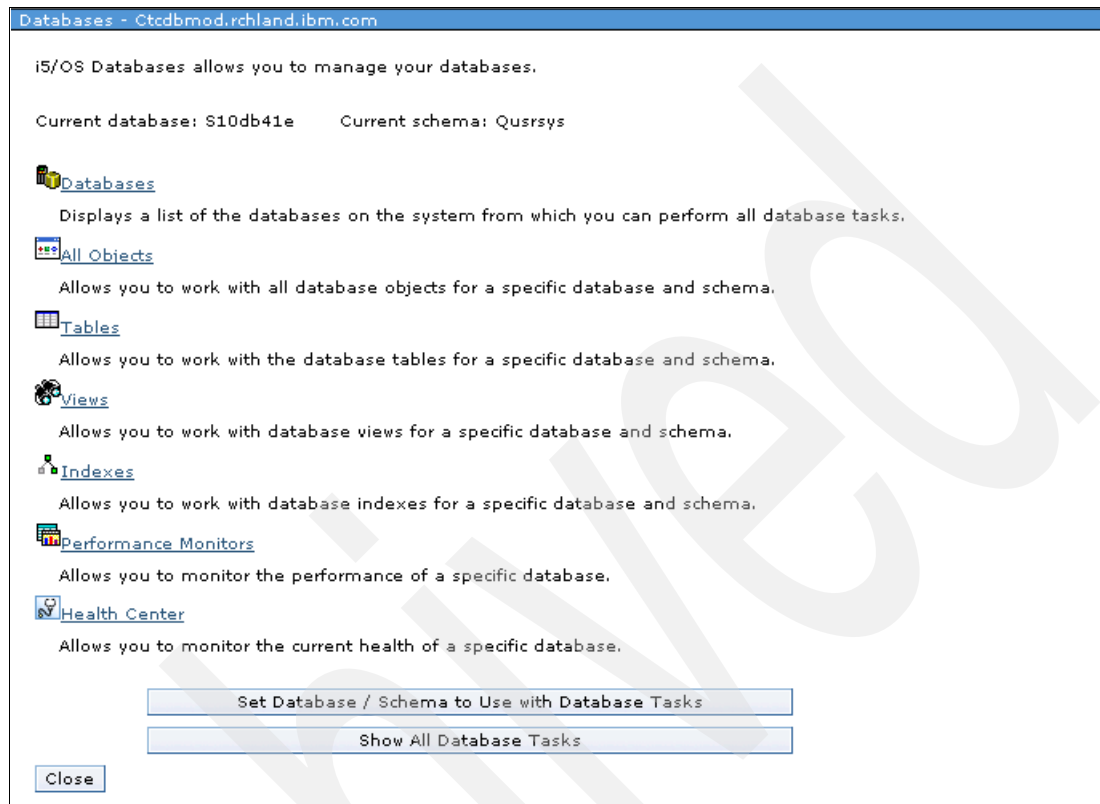


Figure 4-10 Databases functions

## 4.2.10 Journal management

Work with journals and journals receivers (Figure 4-11).

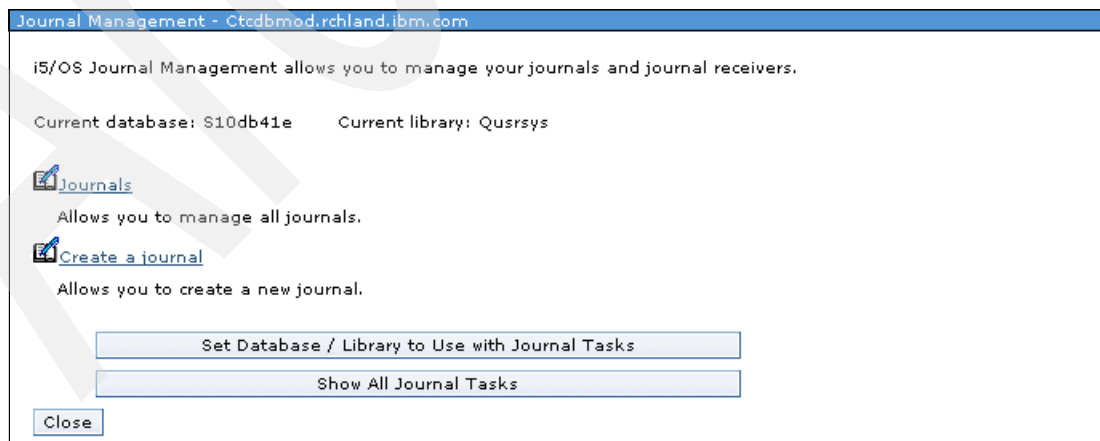


Figure 4-11 Journal management functions

## 4.2.11 Performance

The performance section provides tasks to collect performance data and view the collections to investigate potential performance issues (Figure 4-12).

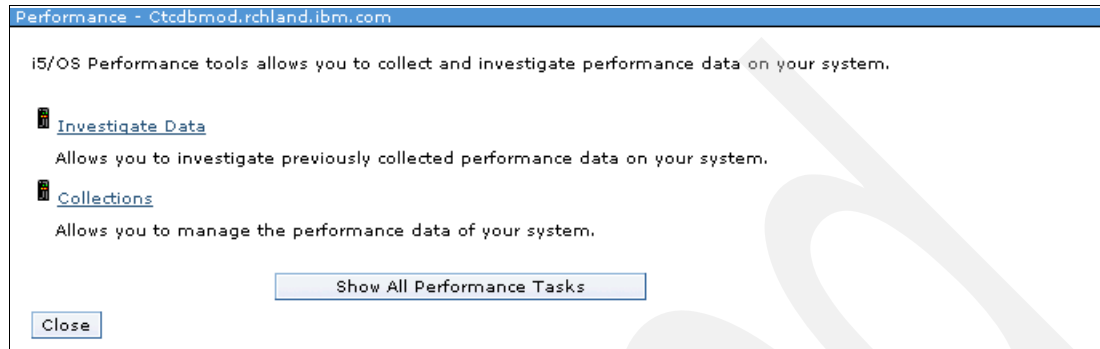


Figure 4-12 Performance functions

## 4.2.12 File Systems page

The File Systems page includes links to the most commonly used file system tasks, Integrated File System, i5/OS NetServer, and Network File System (NFS).

Dynamic mount information has been added to Integrated File System. It shows the list of currently mounted file systems. You can view the name, the mount type, and the mounted directory of the file system. You can also unmount file systems at the same time or view the properties of a file system. Mounting a Network File System is now available from the System Director Navigator for i. It is no longer necessary to go to the green screen to mount the NFS. You can now select a folder and then collect, display, and analyze attribute information from the folder and its contents, create filters, sort the information, or obtain a report of the objects that are consuming the most system storage space, for instance. You will be able to check out or check in all the contents of a folder and its subfolders with only a few clicks (Figure 4-13).

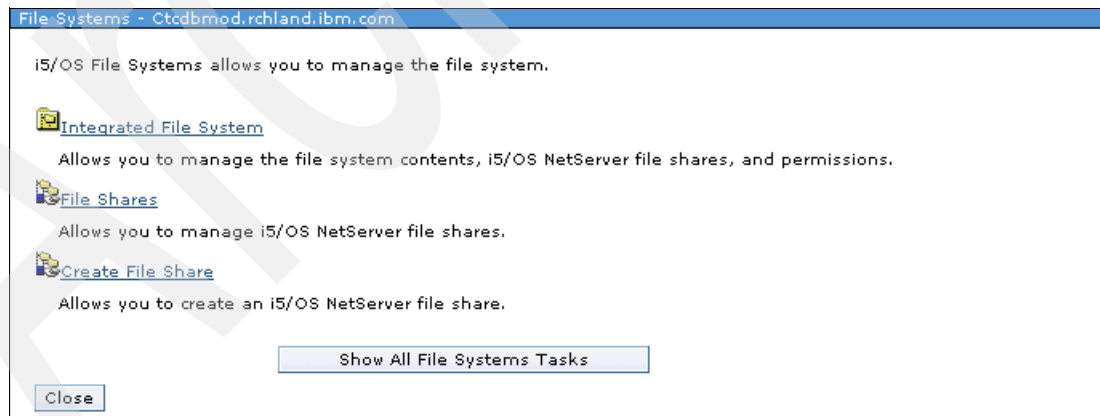


Figure 4-13 File system functions

## 4.2.13 Internet Configurations page

The Internet Configurations page contains i5/OS Internet configuration tasks, the HTTP server, digital certificates, IBM IPP, the Web-based help server, and the Internet setup wizard (Figure 4-14).

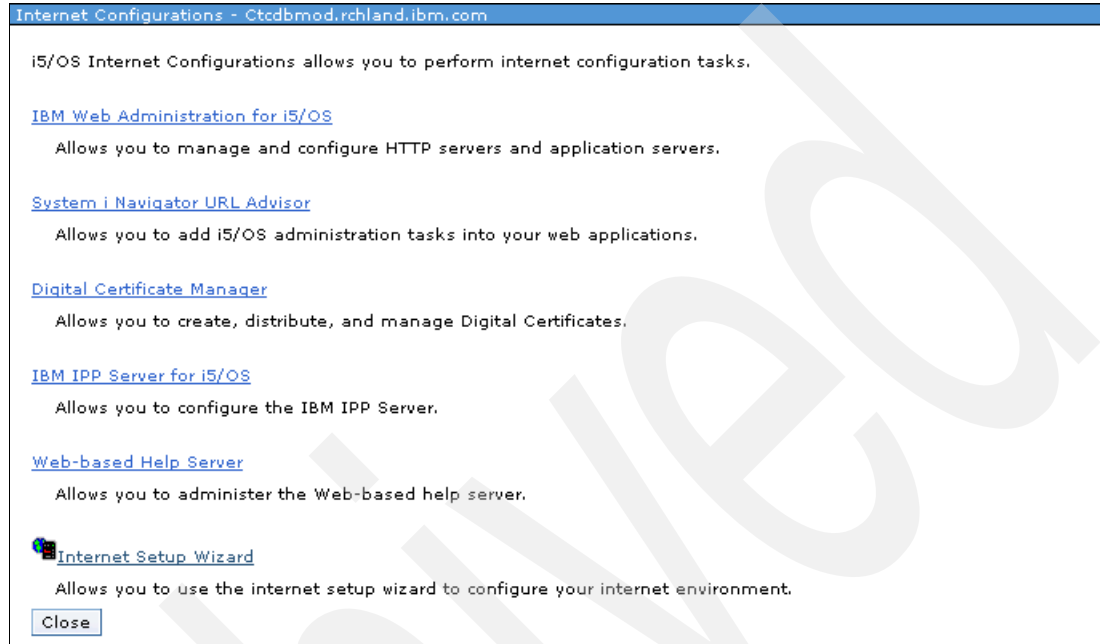


Figure 4-14 Internet configurations functions

From Internet Configurations page you can work with IBM i Internet configuration tasks. Configure your HTTP servers, access System Director Navigator for i Tasks on the Web, work with digital certificates, and more.

For further information about this task and other tasks referenced in this chapter, see:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp?topic=/rzahg/icmain.htm>

# Network

This chapter provides an overview of links to the most commonly used tasks available in the network section of the IBM Systems Director Navigator for i, as shown in Figure 5-1.

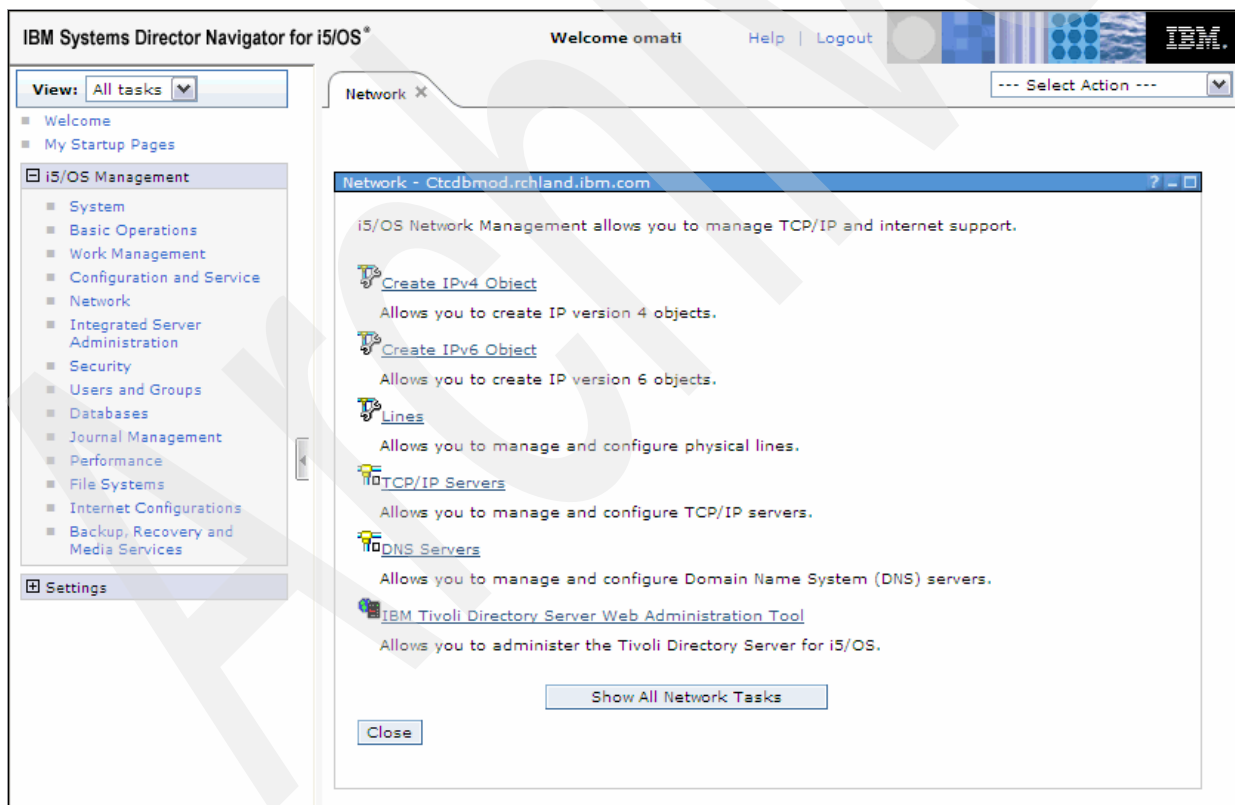


Figure 5-1 Network welcome page

Click **Show All Network Tasks** to view the complete list of all the network functions available (Figure 5-2).

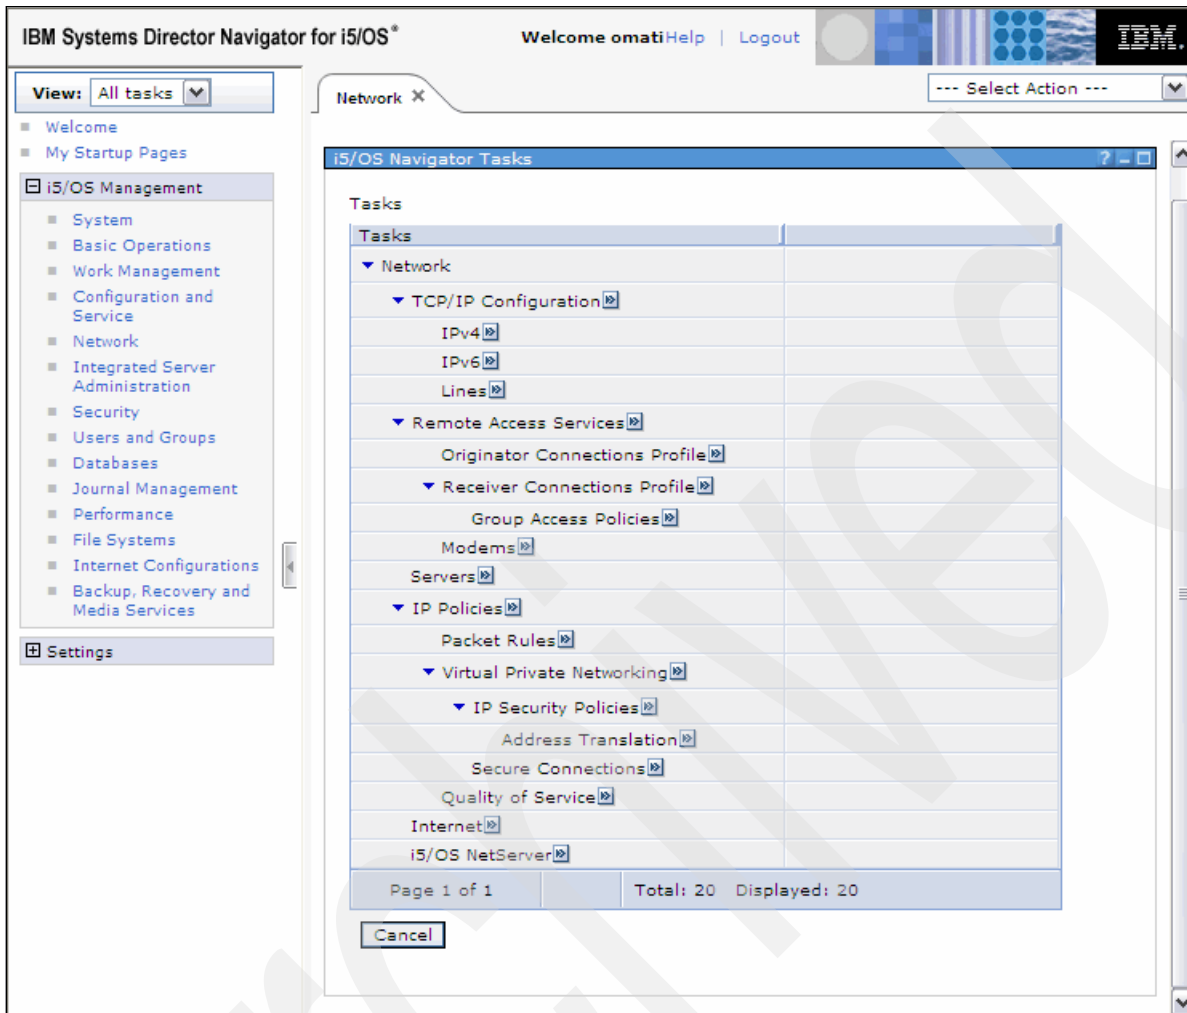


Figure 5-2 All the network tasks

## 5.1 TCP/IP configuration

Use TCP/IP configuration to manage and monitor your TCP/IP network. TCP/IP configuration provides wizards for configuring both IPv4 and IPv6 lines, interfaces, and routes. Figure 5-3 shows all the options of this task.

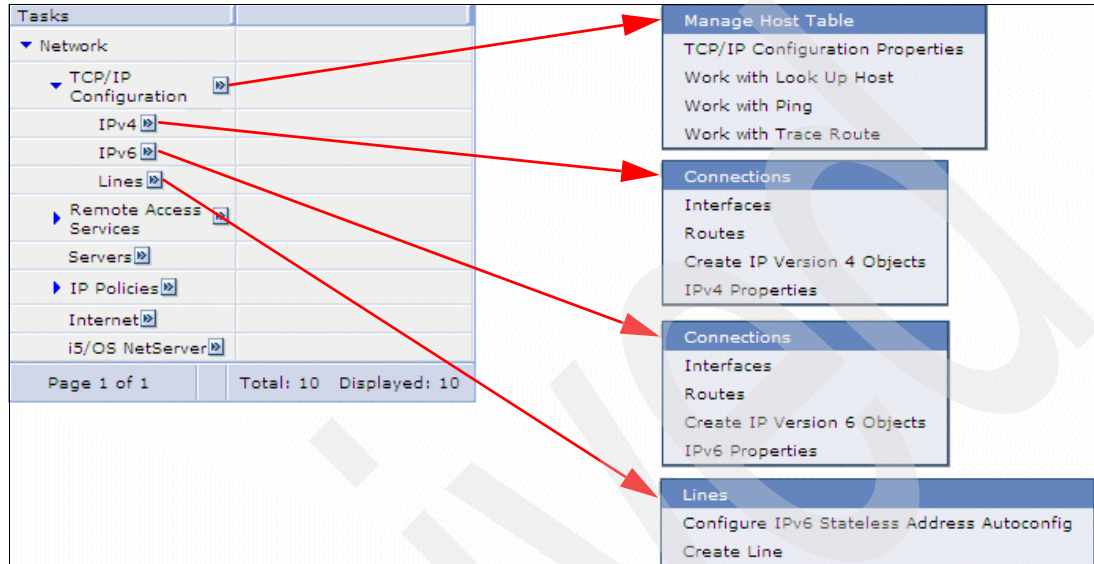


Figure 5-3 All of the TCP/IP tasks options

## TCP/IP properties

TCP/IP configuration allows you to view and configure TCP/IP properties (Figure 5-4), such as domain and host name settings, as shown in Figure 5-5 on page 43, and to start or stop TCP/IP interfaces.

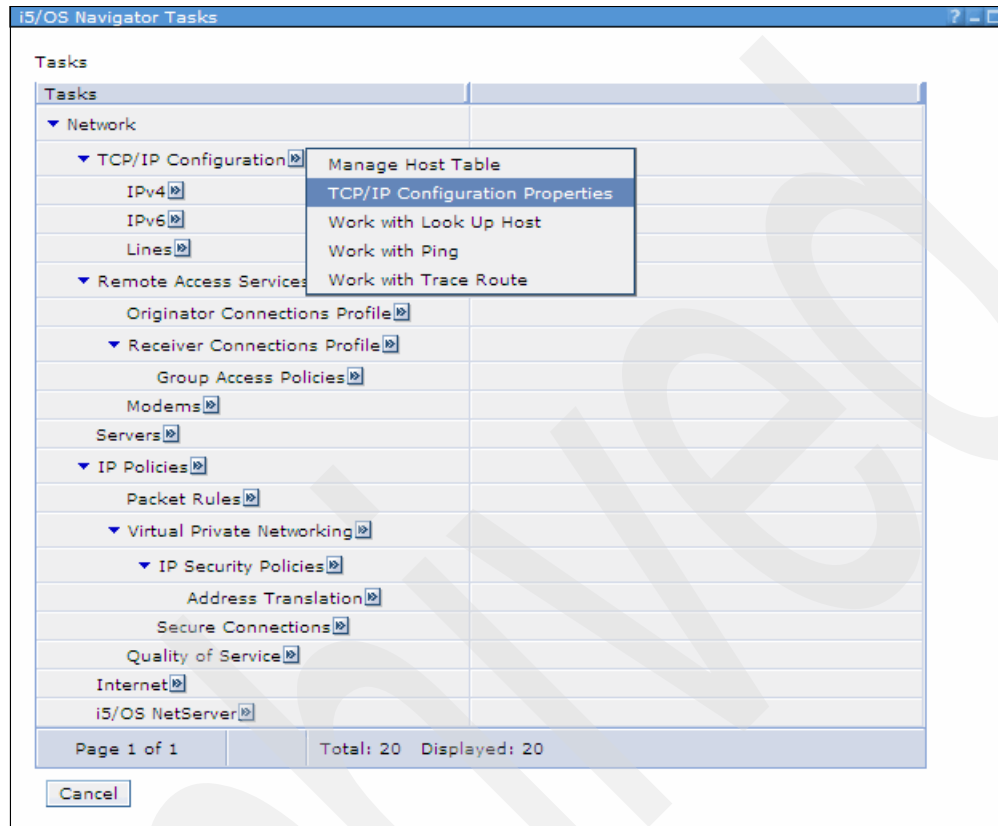


Figure 5-4 TCP/IP configuration properties



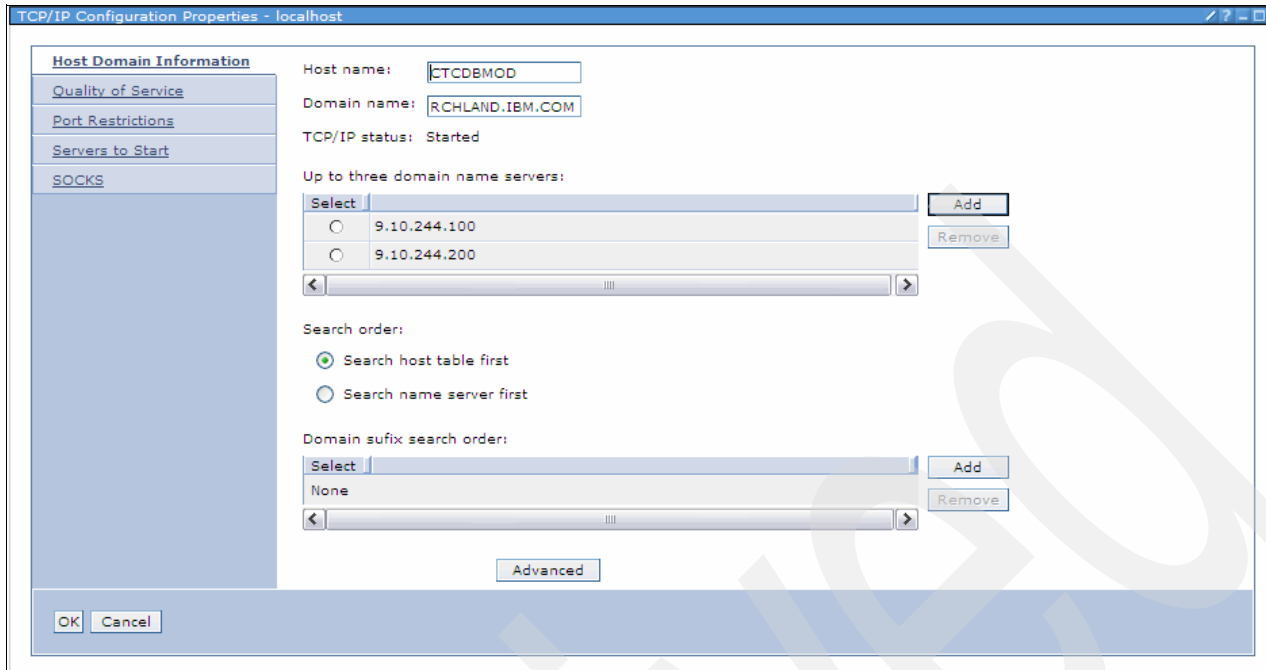


Figure 5-5 Host domain information

## Troubleshoot

In addition, TCP/IP configuration provides you with utilities to troubleshoot your network, such as traceroute (Figure 5-6) and ping (Figure 5-7 on page 44).

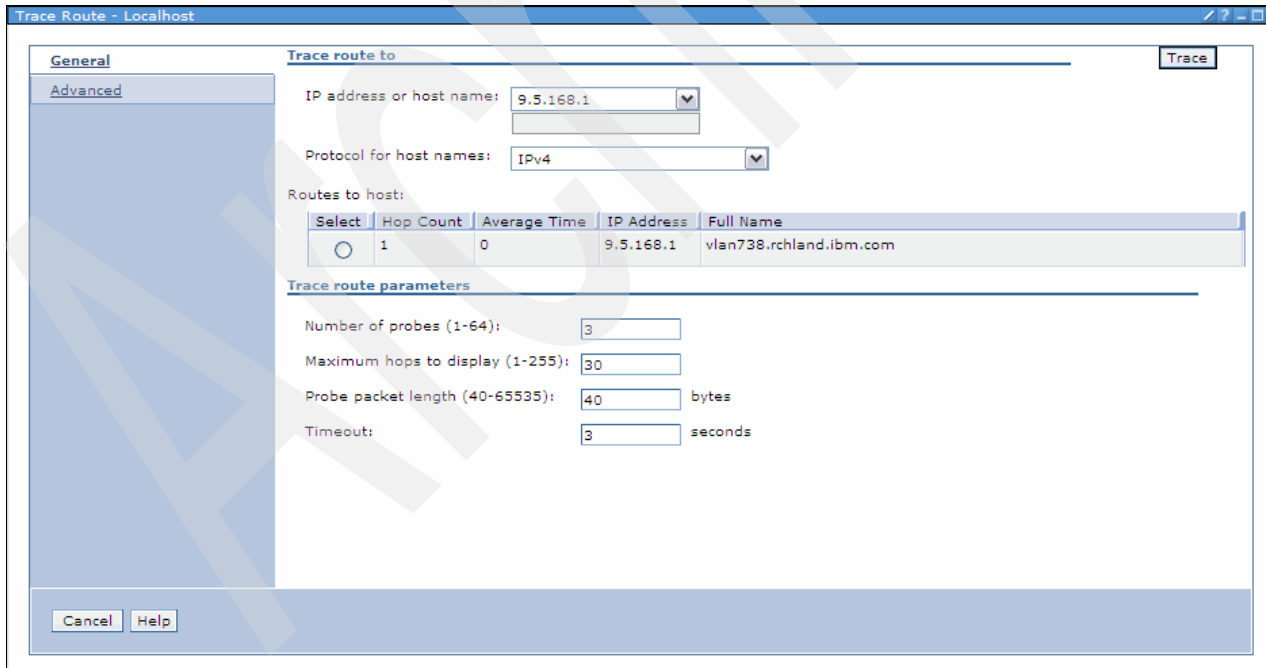


Figure 5-6 Trace route

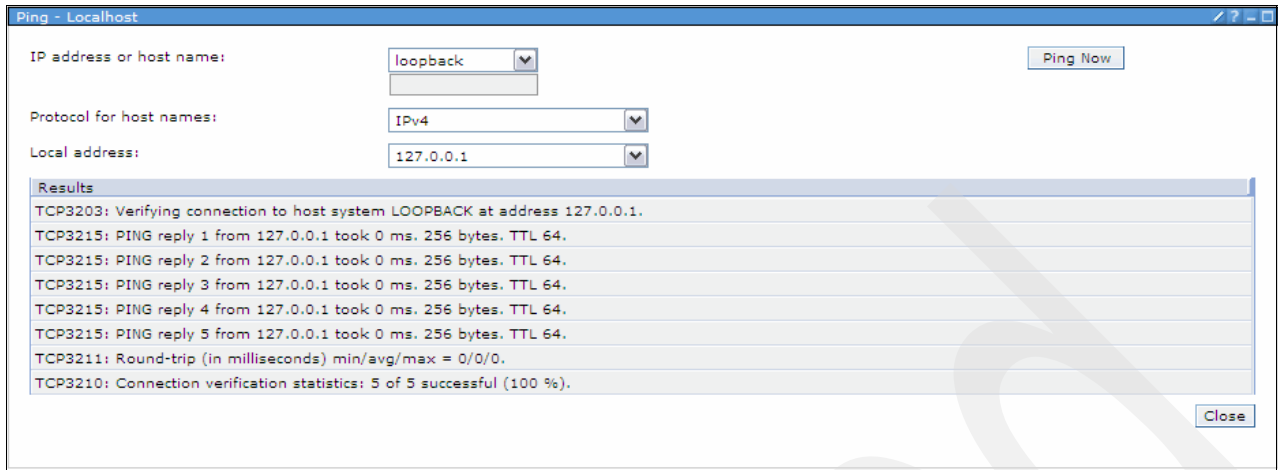


Figure 5-7 Ping

## Servers and connections

On the TCP/IP Configuration properties panel, you can choose which servers you want to be automatically started when TCP/IP starts (Figure 5-8). In Figure 5-9 on page 46 the panel is shown where you can view all the listening ports and active connections of the system.

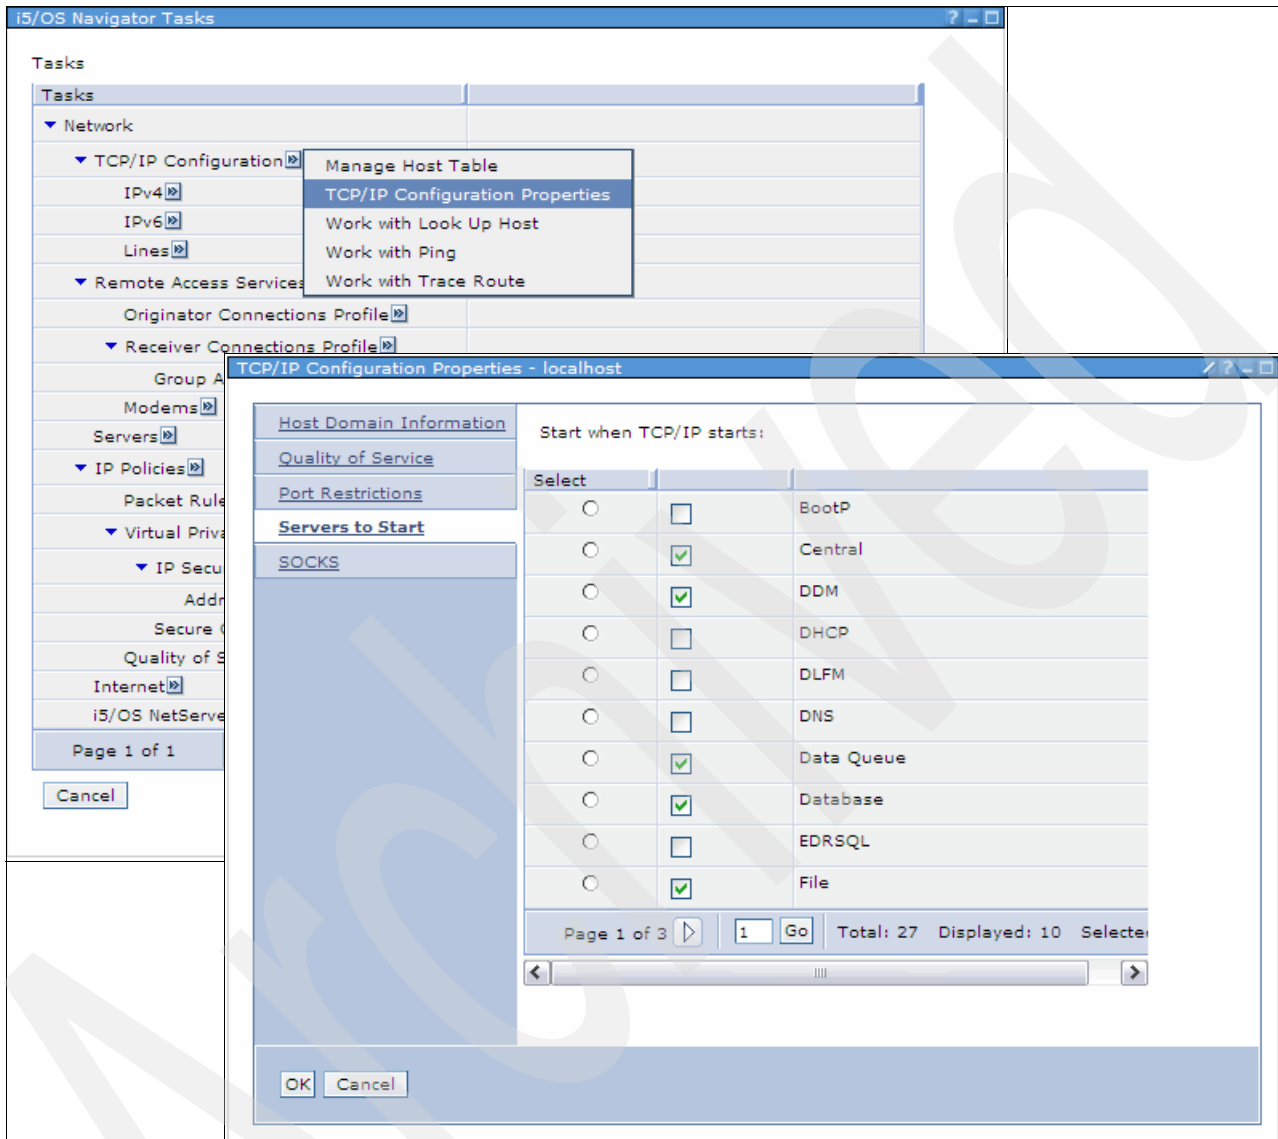


Figure 5-8 Servers to start

Select	Remote Address	Remote Port	Local Address	Local Port	State	Idle Time	Bytes Received	Bytes Sent	Opening Type	Associated User Profile
<input type="checkbox"/>	*	*	*	59078	Listen	1d 00:24:54	0	0	Passive	Qibmhelp
<input type="checkbox"/>	9.49.187.181	3039	9.5.168.161	23	Established	00:02:09	3861	483234	Passive	Qtcp
<input type="checkbox"/>	9.57.139.124	1188	9.5.168.161	23	Established	00:00:14	428	8718	Passive	Qtcp
<input type="checkbox"/>	*	*	9.5.168.161	427	Listen	1d 00:25:58	0	0	Passive	Qsys
<input type="checkbox"/>	9.57.139.109	1252	9.5.168.161	445	Established	00:09:15	6640236	12392631	Passive	Qsys
<input type="checkbox"/>	9.57.139.124	2219	9.5.168.161	445	Established	01:29:12	513	387	Passive	Qsys
<input type="checkbox"/>	9.57.139.115	2914	9.5.168.161	2005	Time-wait	00:00:45	1352	205478	Passive	Qlwavr
<input type="checkbox"/>	9.57.139.115	2917	9.5.168.161	2005	Time-wait	00:00:40	3086	539	Passive	Qlwavr
<input type="checkbox"/>	9.57.139.115	2918	9.5.168.161	2005	Time-wait	00:00:39	1085	171359	Passive	Qlwavr
<input type="checkbox"/>	9.57.139.115	2925	9.5.168.161	2005	Time-wait	00:00:16	3459	539	Passive	Qlwavr
<input type="checkbox"/>	9.57.139.115	2926	9.5.168.161	2005	Time-wait	00:00:15	1079	205480	Passive	Qlwavr

Figure 5-9 Connections list

Figure 5-10 shows how to manage the TCP/IP servers.

Select	Name	Status	Description
<input type="checkbox"/>	ASFTomcat	Started	ASFTomcat Basic Servlet and JSP Engine for the Apache Web Server
<input type="checkbox"/>	BootP	Stopped	BootP Server
<input type="checkbox"/>	BootP DHCP Relay Agent	Stopped	BootP DHCP Relay Agent
<input type="checkbox"/>	DDM	Started	DRDA DDM Server TCP/IP
<input type="checkbox"/>	DHCP	Stopped	Dynamic Host Configuration Protocol
<input type="checkbox"/>	DLFM	Stopped	Datalink File Manager
<input type="checkbox"/>	EDRSQL	Stopped	Extended Dynamic Remote SQL
<input type="checkbox"/>	FTP	Started	File Transfer Protocol
<input type="checkbox"/>	HTTP Administration	Stopped	HTTP Administration Server
<input type="checkbox"/>	i5/OS NetServer	Stopped	i5/OS NetServer for i5/OS Network Neighborhood
<input type="checkbox"/>	IBM Tivoli Directory Server for i5/OS	Stopped	IBM Tivoli Directory Server for i5/OS (LDAP)
<input type="checkbox"/>	INETD	Stopped	Internet Super Server
<input type="checkbox"/>	LPD	Started	Line Printer Daemon
<input type="checkbox"/>	Management central	Started	Management Central
<input type="checkbox"/>	NFS	Stopped	NFS Server

Figure 5-10 TCP/IP servers

For more information about the TCP/IP configuration refer to the IBM i Information Center at the following link and choose **Networking** → **TCP/IP setup**:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>

## Create a line description

This task allows you to create a new line description using the provided wizard. Figure 5-11 to Figure 5-15 on page 49 are the panels that you will see while using the configuration wizard.

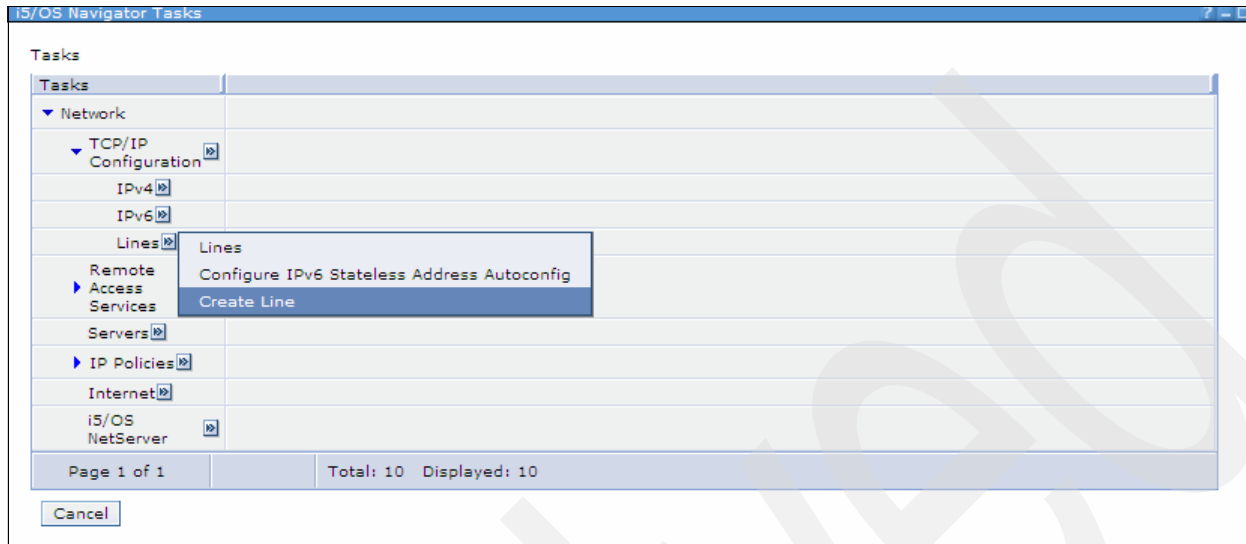


Figure 5-11 Create a line description

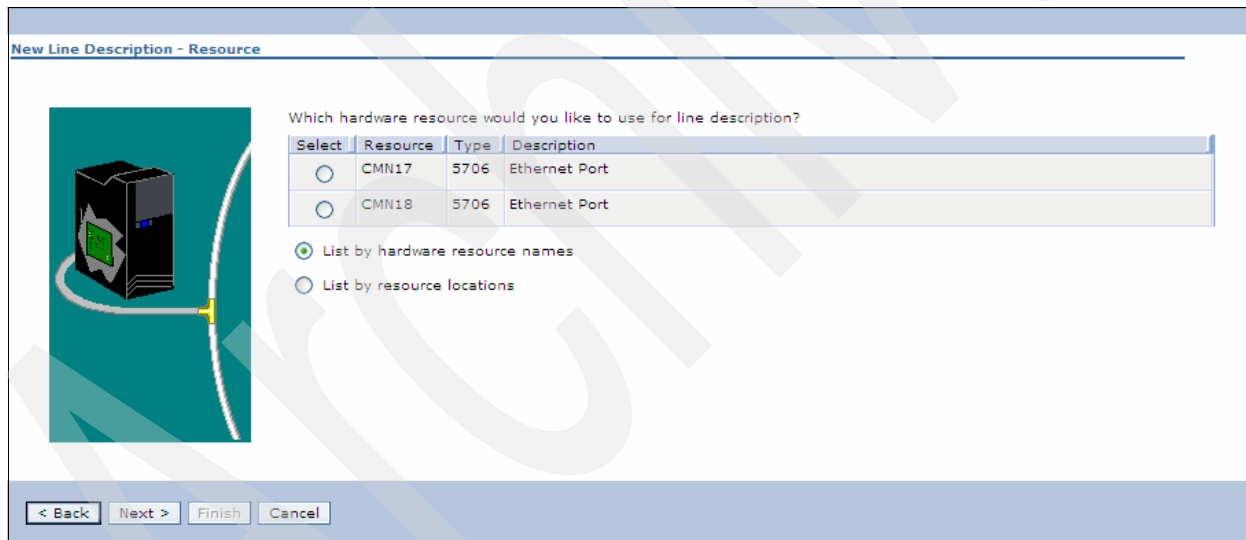


Figure 5-12 Choose resource

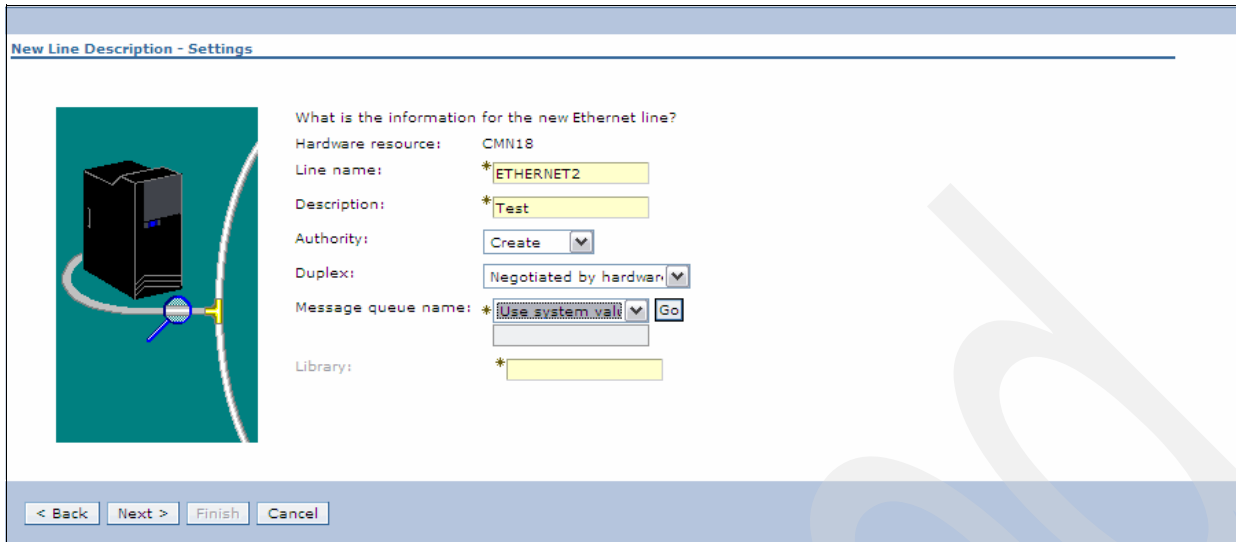


Figure 5-13 New line description settings

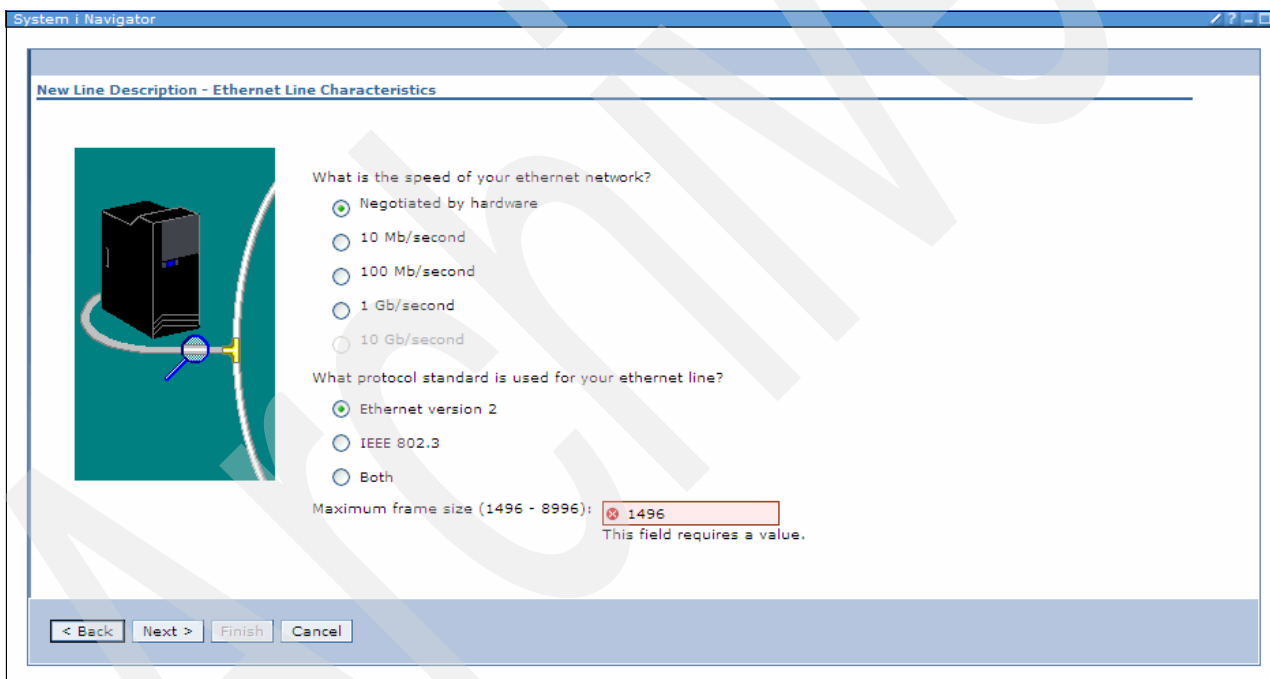


Figure 5-14 New line characteristics

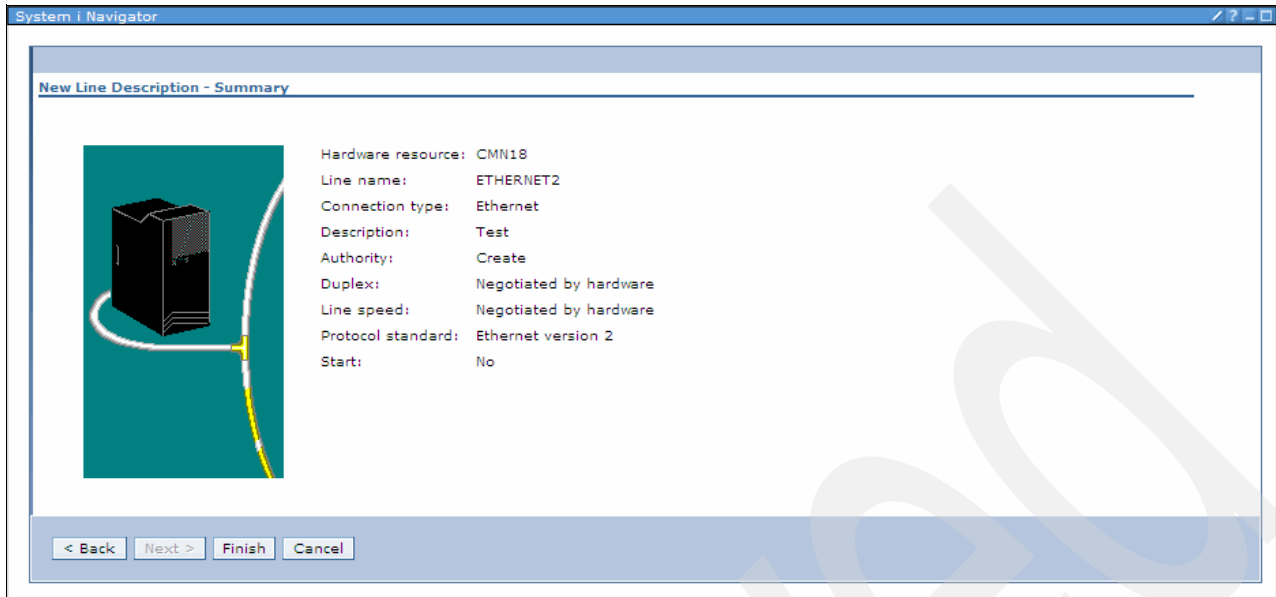


Figure 5-15 New line description summary

## 5.2 Remote Access Services

Point-to-point connection profiles define a set of parameters and resources for Remote Access Services (RAS) connections. You can activate profiles to use the predefined settings to either dial out (originate) or listen for (receive) RAS connections. See all the available options in Figure 5-16.

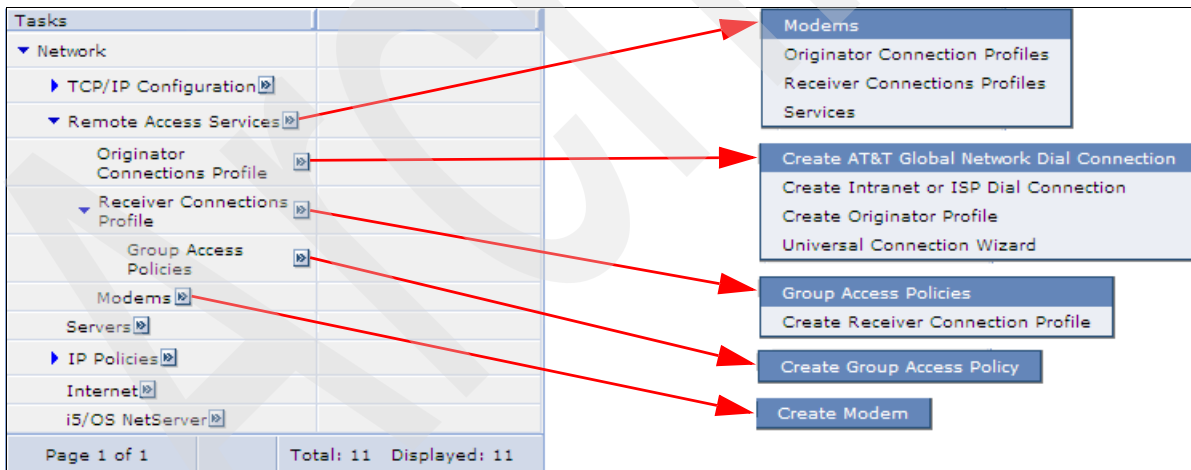


Figure 5-16 All of the Remote Access options

### Protocol type

You must specify the protocol that you want this profile to use when a remote system establishes a point-to-point connection:

- ▶ PPP
- ▶ SLIP

## Mode selections

Specifies the operating mode to be used by this connection profile when this system establishes a point-to-point connection:

- ▶ Connection type
  - Switched line
  - Leased line
  - L2TP
  - PPPoE
- ▶ Operating mode

The operating mode choices depend on the type of profile that you are creating:

- Originator profile
- Receiver profile

Figure 5-17 shows the window that appears when you choose to create a PPP answer profile on a switched line.

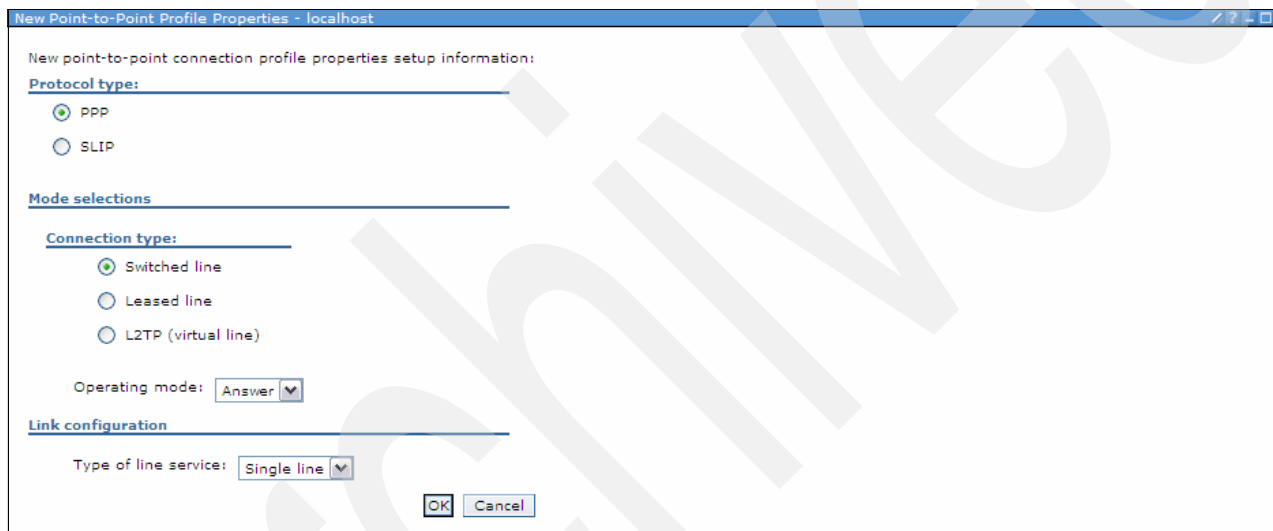


Figure 5-17 New point-to-point profile

## Universal Connection Wizard

Remote Access Services let you configure a connection between your system and IBM through the Universal Connection wizard. It enables electronic support-related functions such as reporting problems, receiving PTFs, sending service information, or receiving documentation updates. IBM uses this configuration when you allow the following applications to provide support, service, and system documentation:

- ▶ Electronic Customer Support
- ▶ Electronic Service Agent™



Figure 5-18 shows how to start it. Then you will be asked to fill in your company information (Figure 5-19).

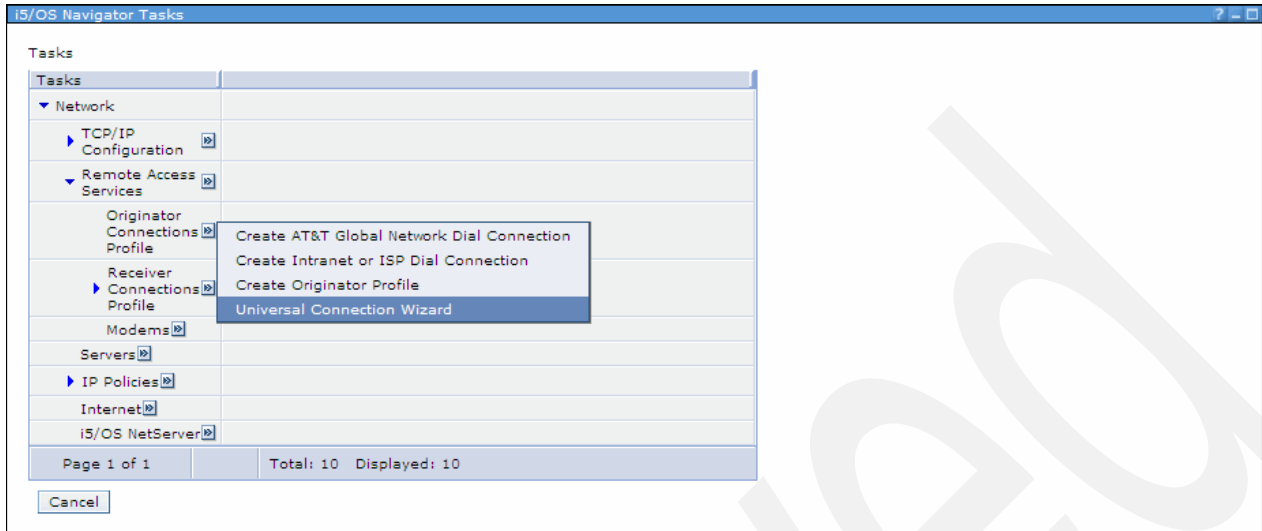


Figure 5-18 Universal Connection Wizard

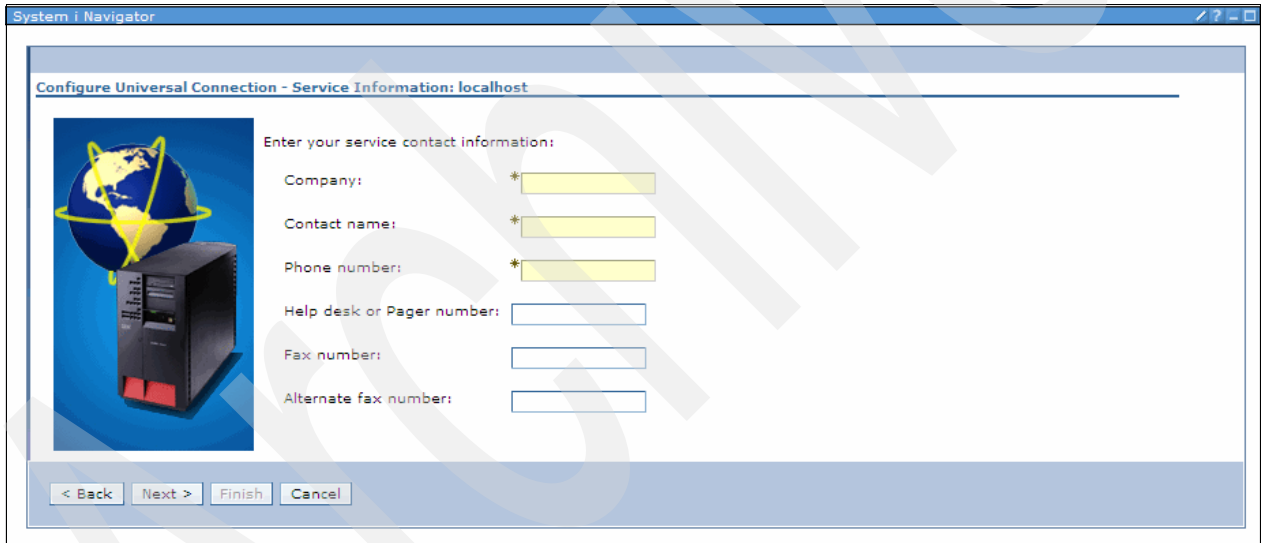


Figure 5-19 Service information

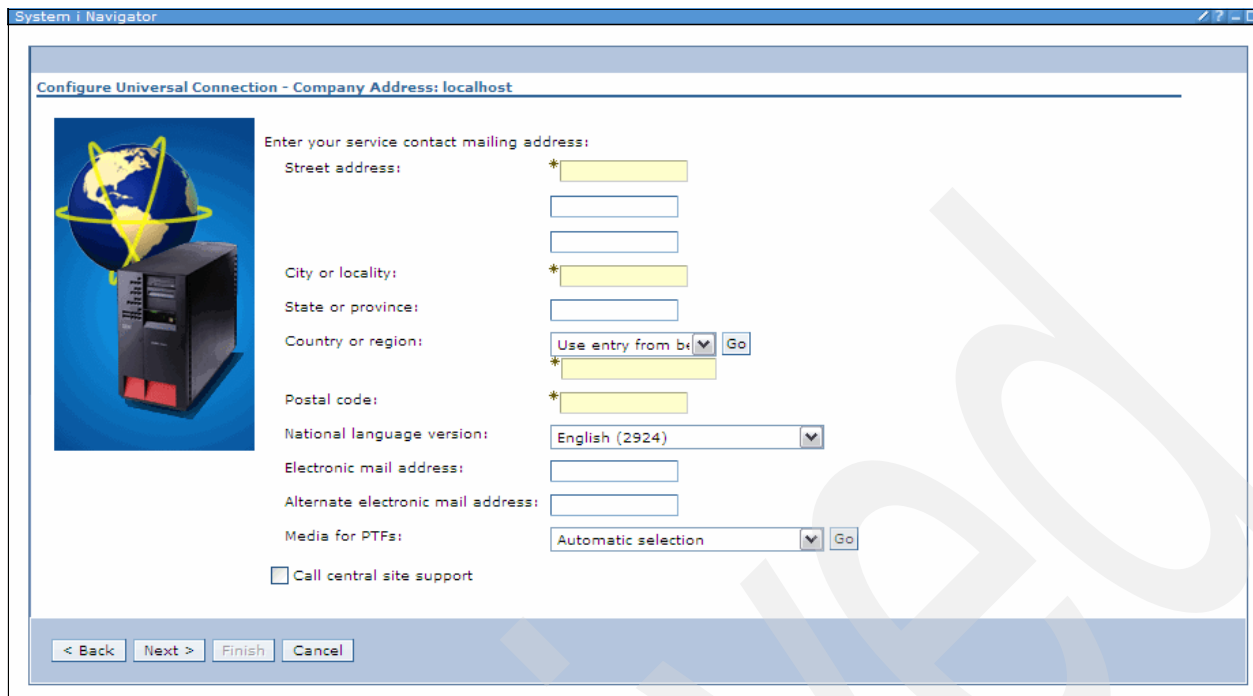


Figure 5-20 Company address

After inserting company details, such as company name, contact name, country, and region, you will choose the connection type, as shown in Figure 5-21. When finished, the wizard will show you a summary of all the information provided and ask you to check the connection.

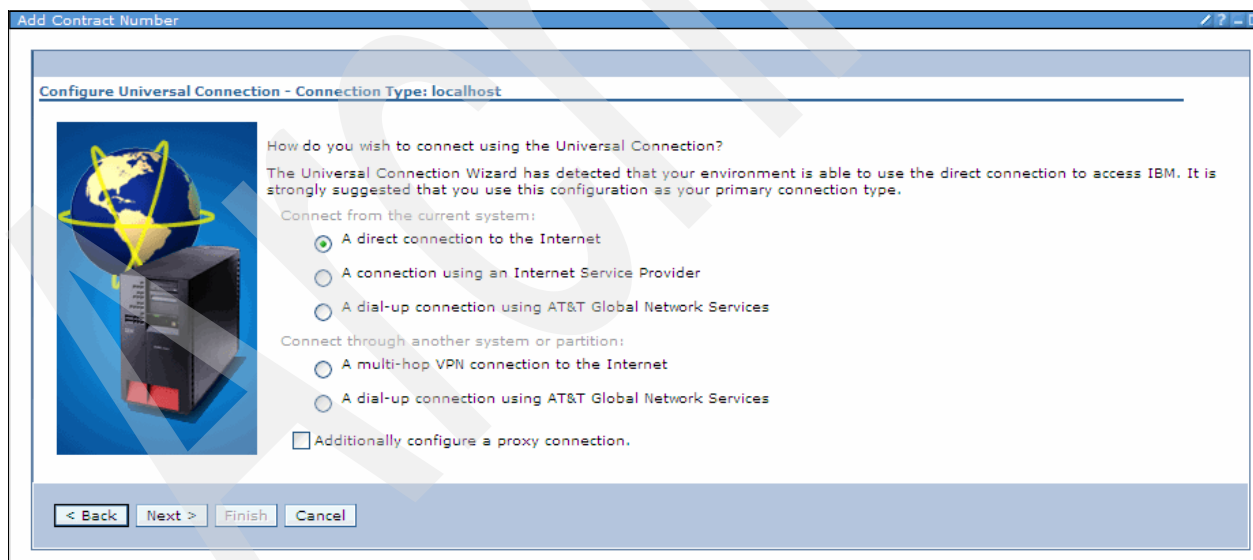


Figure 5-21 Connection type

For detailed information about the Remote Access Services refer to the IBM i Information Center at the following link and choose **Networking** → **TCP/IP applications, protocols, and services** → **Remote Access Services: PPP connections**:

<http://publib.boulder.ibm.com/infocenter/iseres/v6r1m0/index.jsp>

## 5.3 Servers

The Servers section in the Network task lets you check the status and manage the following servers jobs:

- ▶ DNS
- ▶ Domino® Servers.
- ▶ System i Access.
- ▶ TCP/IP Servers.
- ▶ User-defined.
- ▶ Create DNS server.
- ▶ Create user-defined server.
- ▶ Manage DNS dynamic update keys.
- ▶ Servers properties.

### DNS

Shows you the status of the Domain Name System (DNS) if configured. Choose the Create DNS Server task if you need to create one.

### Domino Servers page

This page lets you see the status of the Domino Servers. Using the pop-up menu or the select action drop-down list, you can access all the actions that can be taken for the selected server. See Figure 5-22.

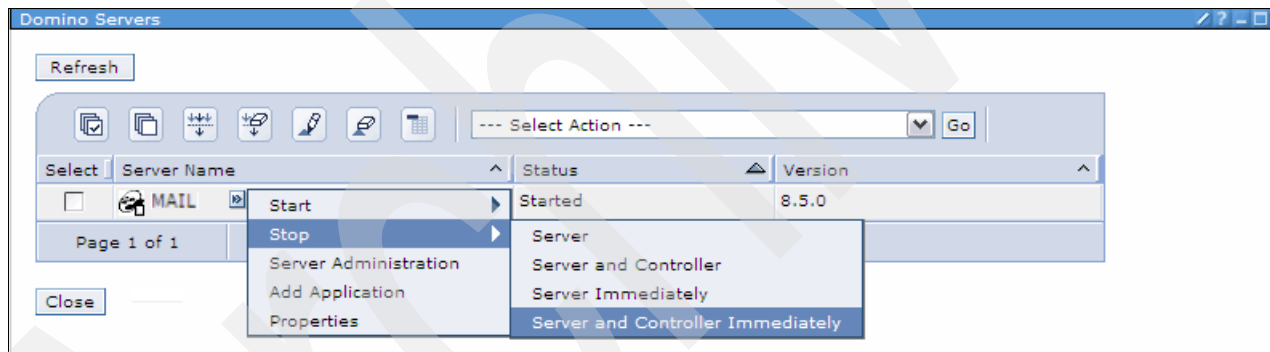


Figure 5-22 Manage Domino Servers

## System i Access

This option displays a list of all the System i Access servers, as shown in Figure 5-23, and allows you to list the server jobs running on the system for the selected server (Figure 5-24).

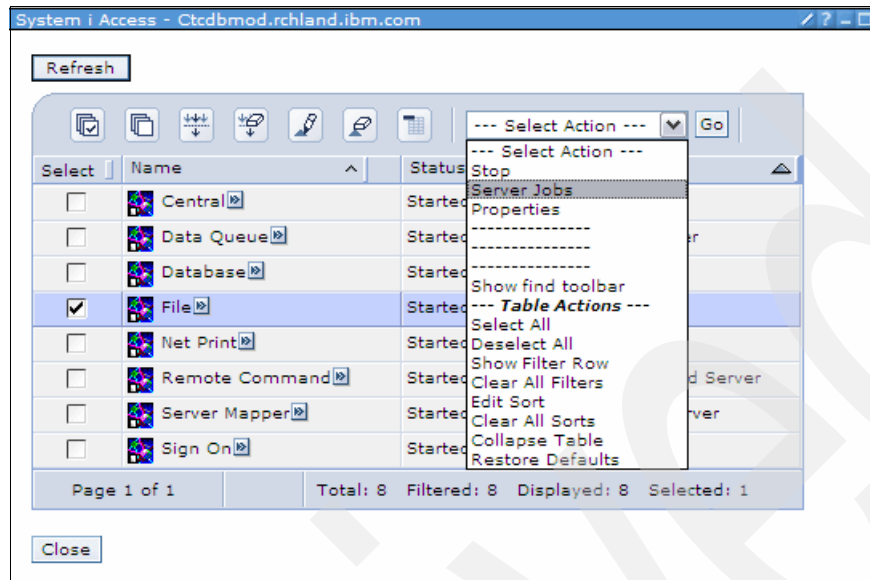


Figure 5-23 System i servers

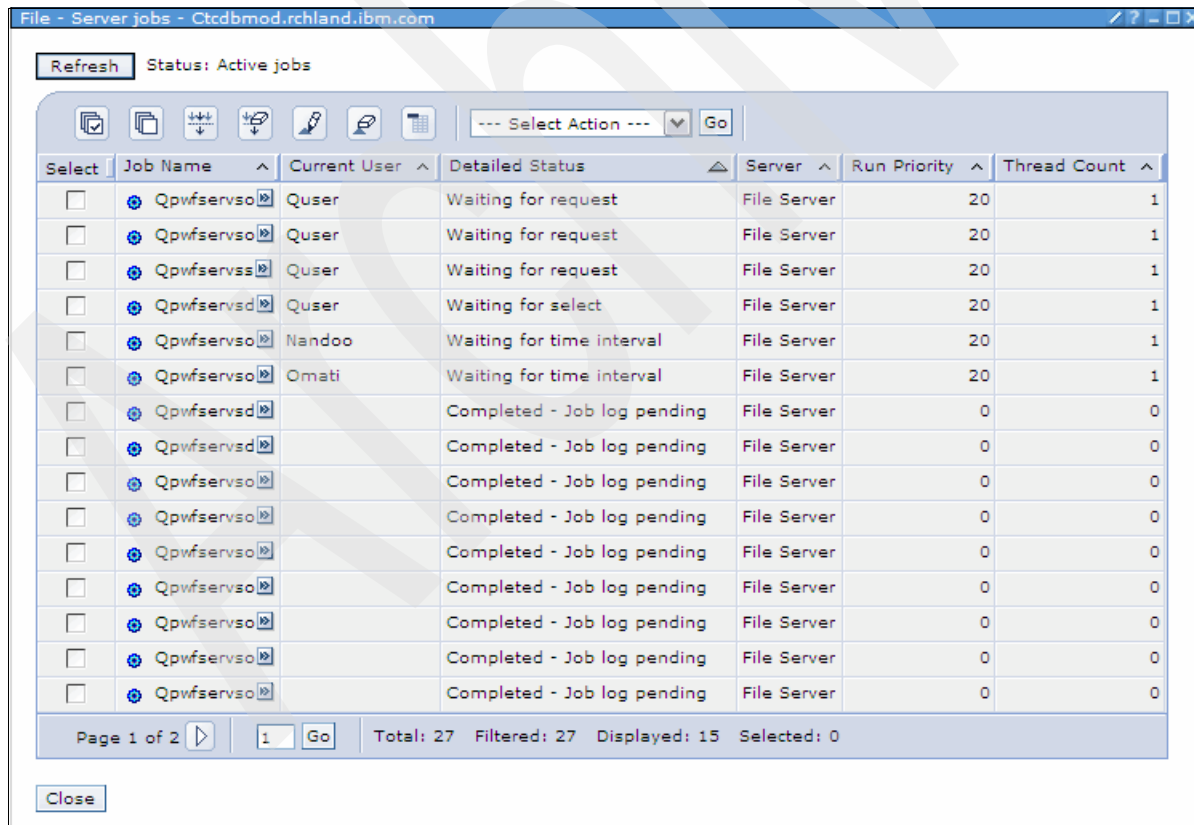


Figure 5-24 File: server jobs

## TCP/IP Servers page

This page lists all the TCP/IP servers available on the system (Figure 5-25). In our example we show how to configure an NFS Mount. See Figure 5-26 on page 56 and Figure 5-27 on page 56.

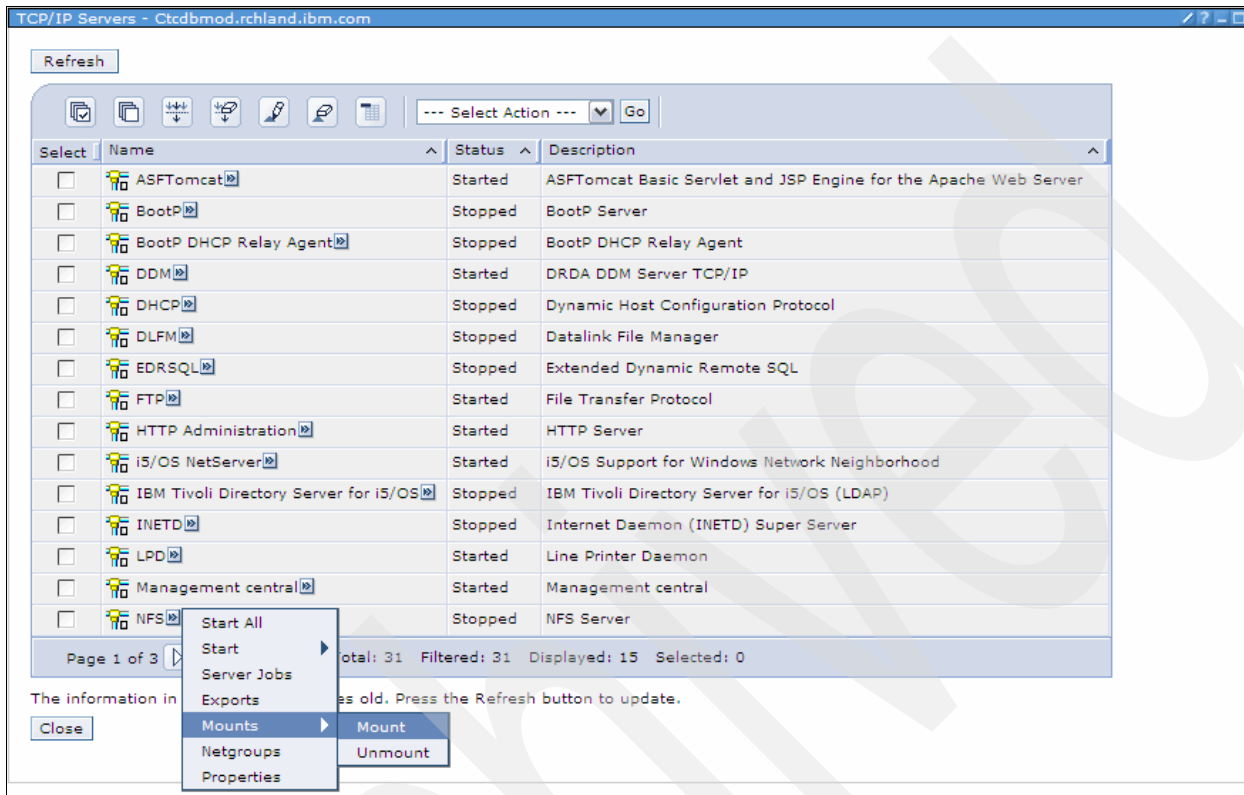


Figure 5-25 TCP/IP Servers

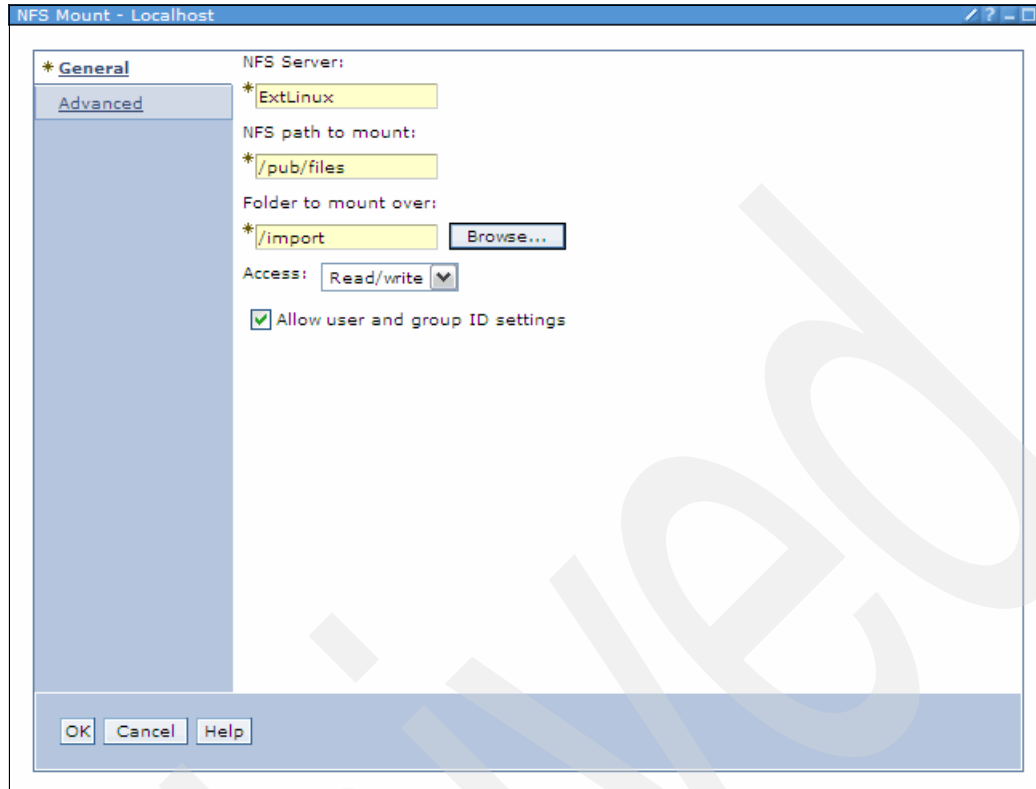


Figure 5-26 NFS Mount

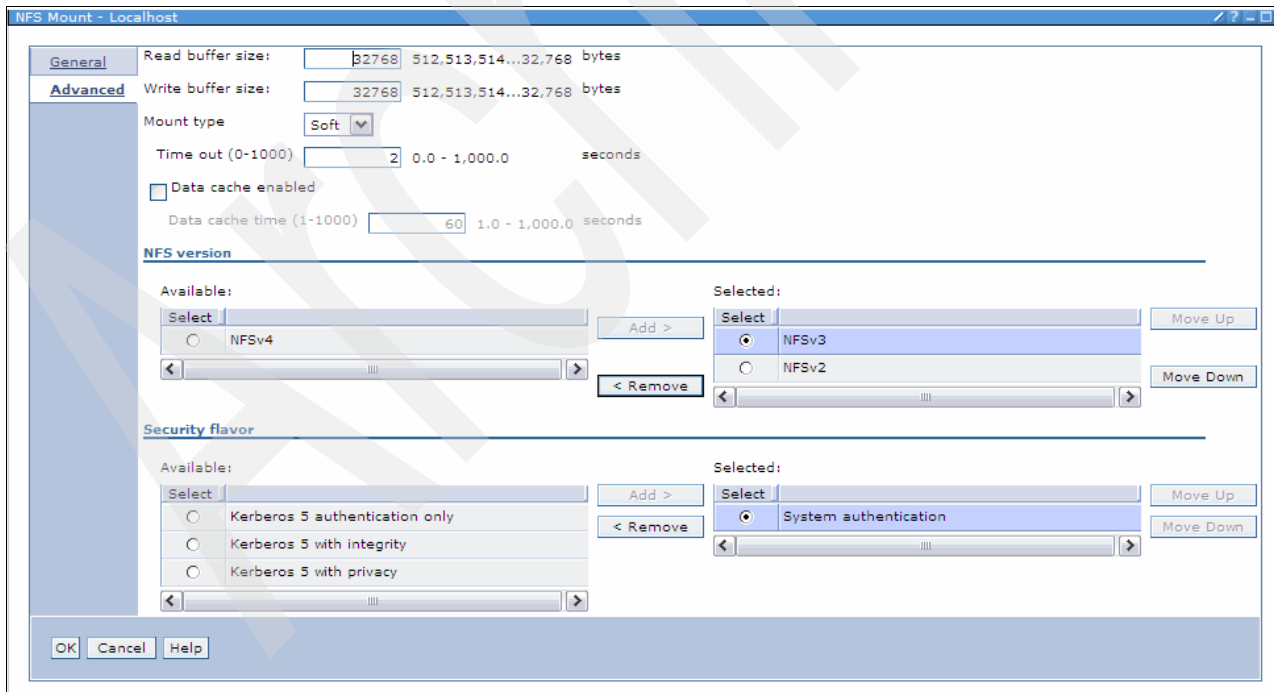


Figure 5-27 NFS Mount: advanced options

### User defined

The System i platform supports user-defined servers. With this support, you can add custom servers and perform administrative tasks for these custom servers. For example, you can

stop, start, or monitor a custom server in the same way in which you administer servers that are included with the operating system.

When you start or end a user-defined server, the data structure that is passed to the server exit program includes the following fields:

- ▶ Action
- ▶ Instance name
- ▶ Instance startup values
- ▶ Length of instance startup values

For further information about user-defined servers refer to the IBM i Information Center at the following link and choose **Networking** → **TCP/IP applications, protocols, and services** → **User-defined servers**:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>

### Create DNS Server task

This task provide a wizard to create your own DNS Server. Domain Name System (DNS) offers a variety of solutions. Before you configure DNS, it is important to plan how it works within your network. For details about the requirements, planning, configuring, and managing of this server refer to the IBM i Information Center at the following link and choose **Networking** → **TCP/IP applications, protocols, and services** → **Domain Name System**:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>

### Create User-Defined Server wizard

Use the Create User-Defined Server wizard (Figure 5-28) to add servers to the user-defined servers list.

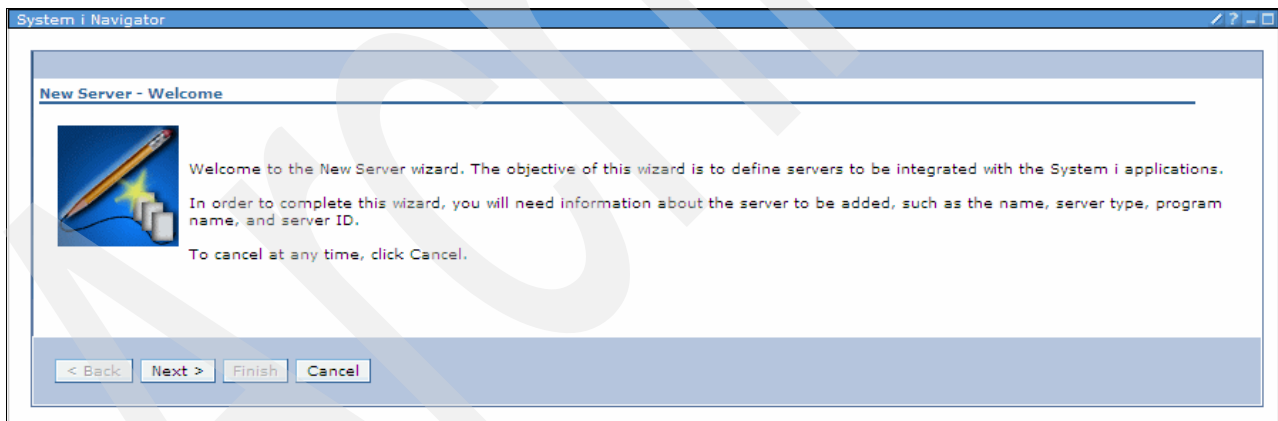


Figure 5-28 New Server - Welcome

Once the wizard complete, the Summary page appears displaying the information specified on the previous panels. To make changes, click **Back** to return to the page that you want to change. Click **Finish** to create the server. See Figure 5-29.

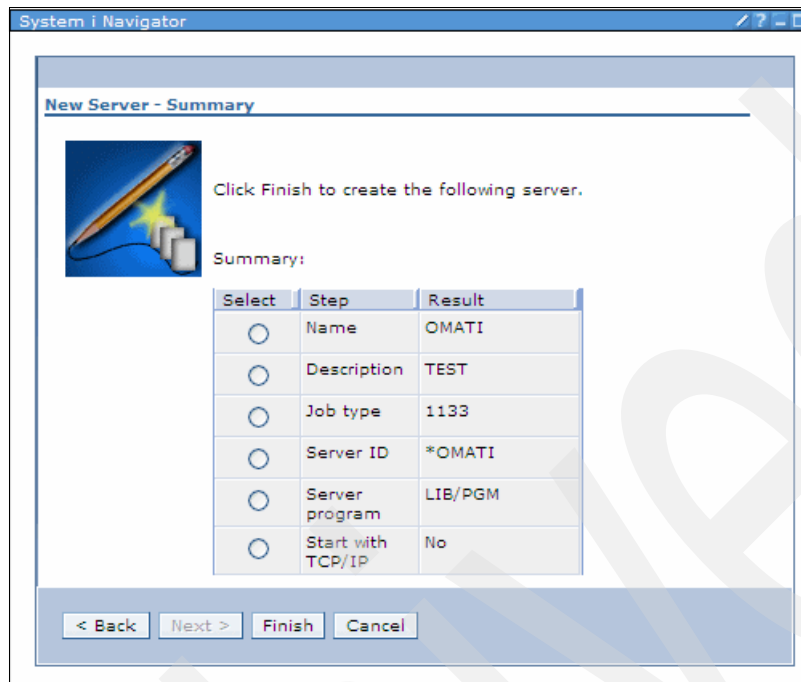


Figure 5-29 New Server - Summary

Also see “User defined” on page 56.

### Manage DNS Dynamic Update Keys dialog

Use the Manage Dynamic Update Keys dialog to manage the dynamic update key files. The key files contain the private keys used by the DNS dynamic update program that is running on this System i product to perform secure dynamic updates. The table displays each key by name, along with the domain names of the zones for which the key can be used. One dynamic update key file is created for each zone domain name specified in the table. Use the appropriate buttons to manage the dynamic update keys table:

- ▶ Add: Add a new key to the table.
- ▶ Edit: Open the selected key in the table so that you can change its properties.
- ▶ Remove: Delete the selected key from the table.

**Note:** Do not remove a key that is named in the allow-update statement of the DNS configuration file (named.conf).

- ▶ View Key: Display the value of the selected key.



## Servers properties

From the Subsystems page you can specify which subsystem you want this server's job to run in. Specific servers use this subsystem function. These include the Distributed Data Management (DDM) server, the central server, the data queue server, the net print server, the remote command server, and the sign on server. See Figure 5-30.

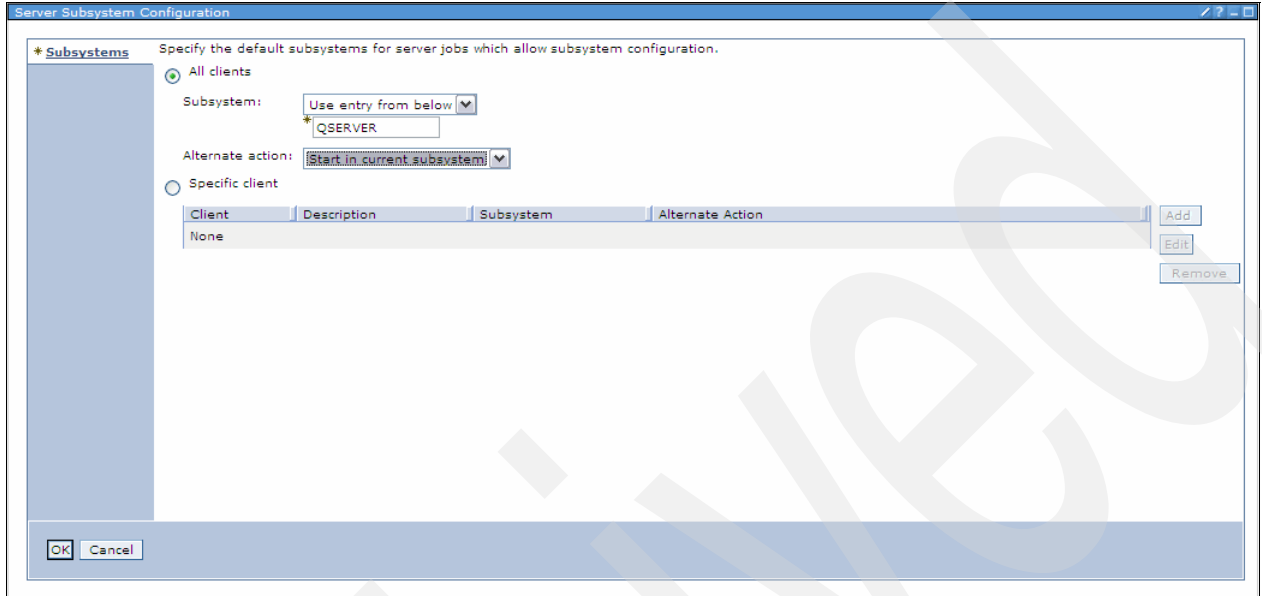


Figure 5-30 Server Subsystem Configuration

Two configuration options are shown in the subsystem configuration:

► All Clients

- Subsystem: Specifies the subsystem that you want this server's job to run in. If None is selected, the server jobs will perform the alternate action.
- Alternate action: Specifies what to do if the server jobs cannot run in the specified subsystem. Possible values are Reject and Start in current subsystem. When Reject is selected, the job will be rejected if it cannot run in the specified subsystem. When Start in current subsystem is selected, the job will attempt to run in the same subsystem that the server daemon is running if it cannot run in the specified subsystem.

► Specific Client

This specifies that you want to set up a unique subsystem configuration for specific clients. When you add a specific client configuration, a <Public> entry will be added to the end of the list. The <Public> entry applies to all clients not included by the specific client entries that you have added. See Figure 5-31 and Figure 5-32 on page 61.

The screenshot shows the 'Add Client' dialog box. It is titled 'Add Client' and has a 'Client information' section. There are two radio buttons: 'IPv4 Client' (selected) and 'IPv6 Client'. Below them is a 'Description:' text box. Under the 'Client:' label, there are two radio buttons: 'IP address:' (selected) and 'IP address range:'. The 'IP address:' section includes a text box for the IP address, a 'Subnet mask: \*' text box, and another 'Subnet mask: \*' text box. The 'IP address range:' section includes two text boxes separated by '-- \*'. Below these are 'Subsystem:' and 'Alternate action:' dropdown menus. The 'Subsystem:' dropdown is set to 'Use entry from' and 'Qusrwrk'. The 'Alternate action:' dropdown is set to 'Start in current'. At the bottom are 'OK' and 'Cancel' buttons.

Figure 5-31 Add Client

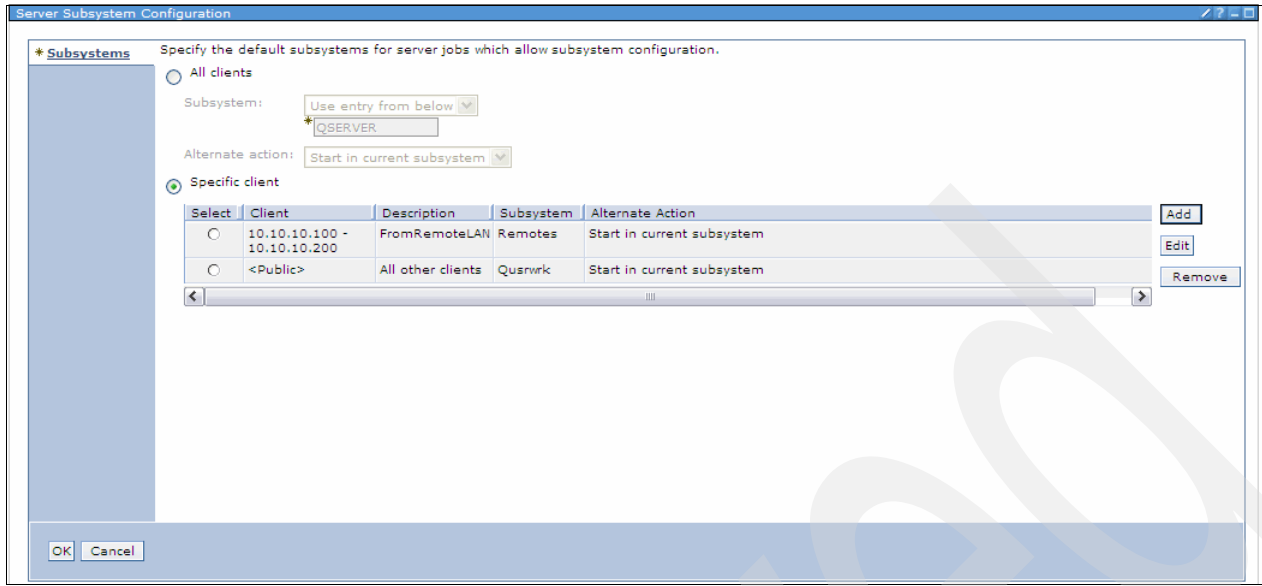


Figure 5-32 <Public> entry automatically added

## 5.4 IP policies

In this section we discuss IP policies.

## Packet rules

Use the Packet Rules Editor to get started creating packet rules on your system. You can create a new file, edit an existing one, or work with the sample files provided on the system. IP packet filtering lets you control what IP traffic you allow in your network. It protects your network by filtering packets according to rules that you specify. See the entire list of available options in Figure 5-33 and a sample of a configuration file in Figure 5-34 on page 63.

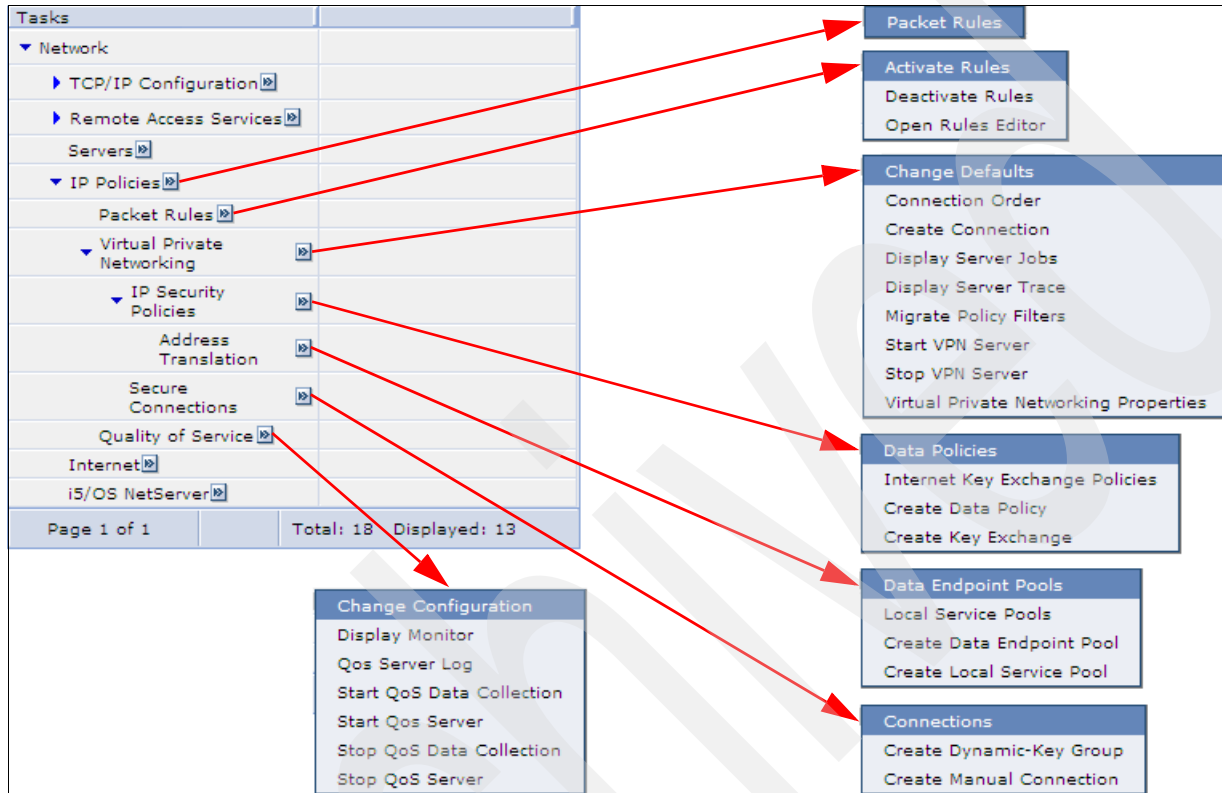


Figure 5-33 All the IP policies tasks

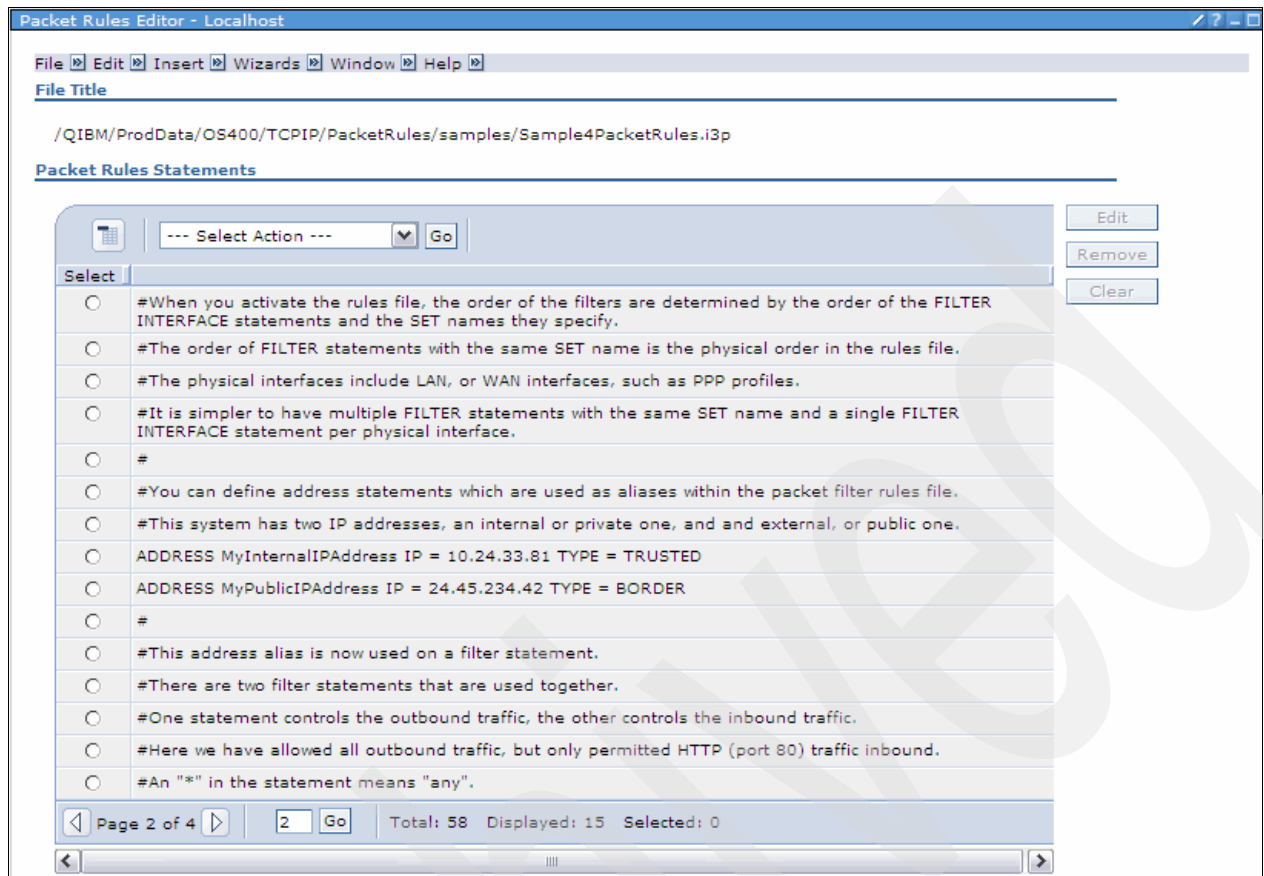


Figure 5-34 Packet rules sample file

For details about the packet rules refer to the IBM i Information Center at the following link and choose **Security** → **IP filtering and network address translation**:

<http://publib.boulder.ibm.com/infocenter/iseres/v6r1m0/index.jsp>

## Virtual private networking

It is important that you have at least a basic knowledge of standard VPN technologies before you implement a VPN connection. A virtual private network (VPN) allows your company to securely extend its private intranet over the existing framework of a public network, such as the Internet. With VPN, your company can control network traffic while providing important security features such as authentication and data privacy. It allows you to create a secure end-to-end path between any combination of host and gateway. VPN uses authentication methods, encryption algorithms, and other precautions to ensure that data sent between the two endpoints of its connection remains secure. Figure 5-35 shows the VPN general defaults.

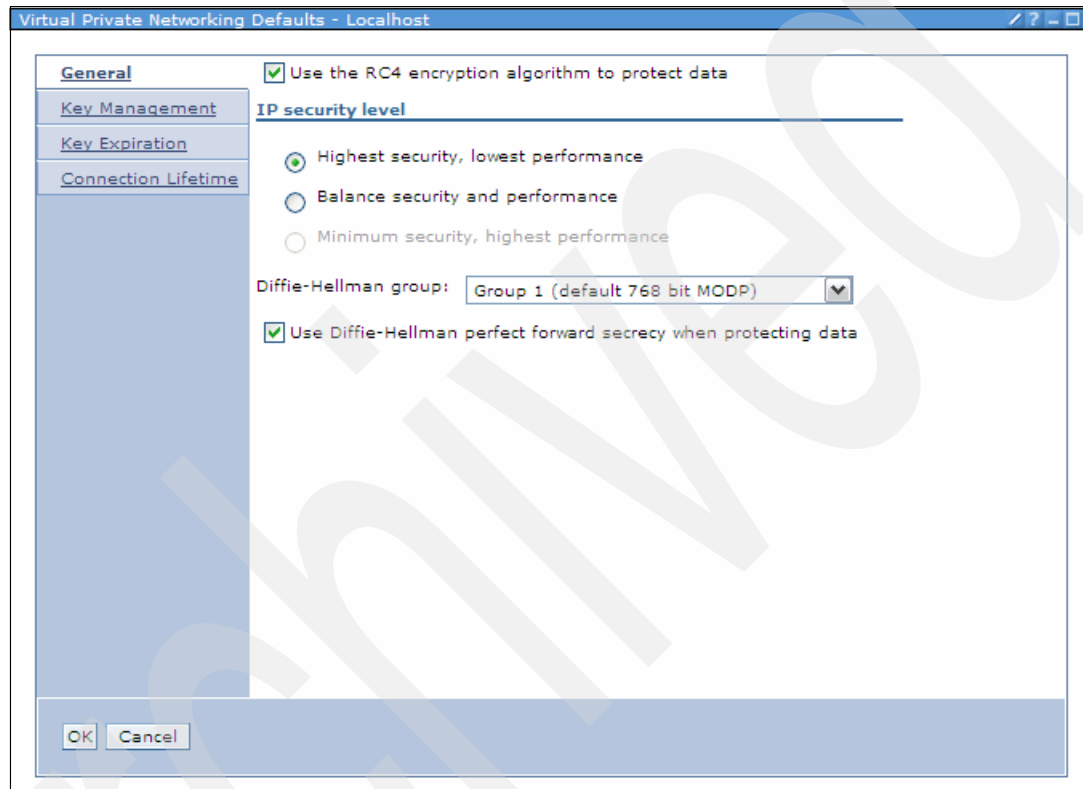


Figure 5-35 VPN defaults

To better understand how any VPN connection works, familiarize yourself with these protocols and concepts and how VPN uses them. See the IBM i Information Center at the following link and choose **Security** → **Virtual Private Networking**:

<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/index.jsp>

## Quality of service

The quality of service (QoS) solution enables the policies to request network priority and bandwidth for TCP/IP applications throughout the network. QoS policies can manage packet priority and can also limit data leaving your system, manage connection requests, and control system load. The QoS server must be running to activate the intrusion detection policy. Use the New Quality of Service Configuration wizard (Figure 5-37) to begin policy configuration. This wizard asks you to define start-up instructions for the different QoS parameters, as shown in Figure 5-36. You must use this wizard in the following circumstances:

- ▶ You are upgrading your server to a new release. You will have the option to configure a directory server in which to export information. No data is lost during this conversion.
- ▶ This is the first time that you are using the QoS graphical user interface (GUI) on this system.
- ▶ You want to manually remove any previous configuration information and start over.

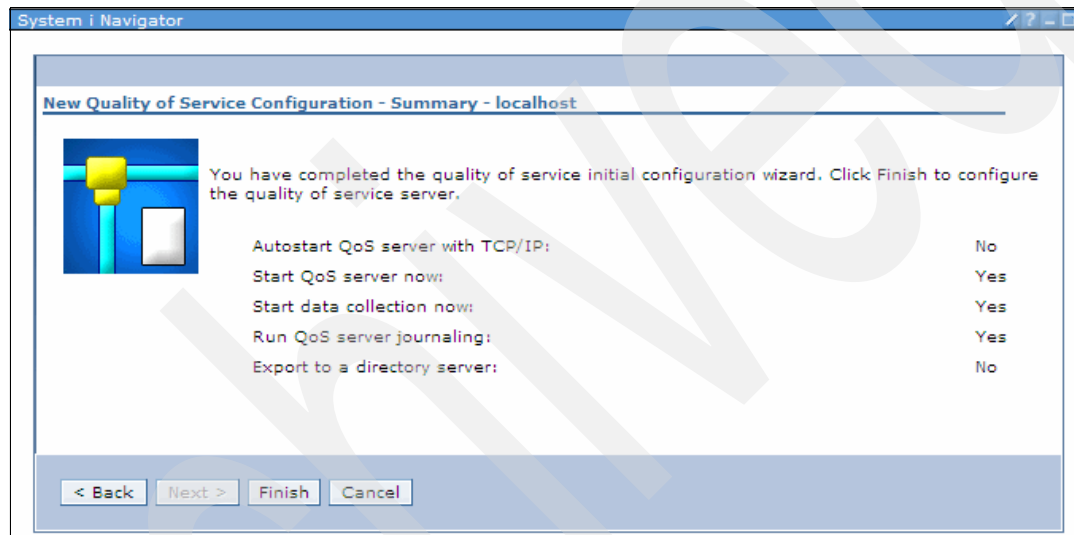


Figure 5-36 QoS configuration summary

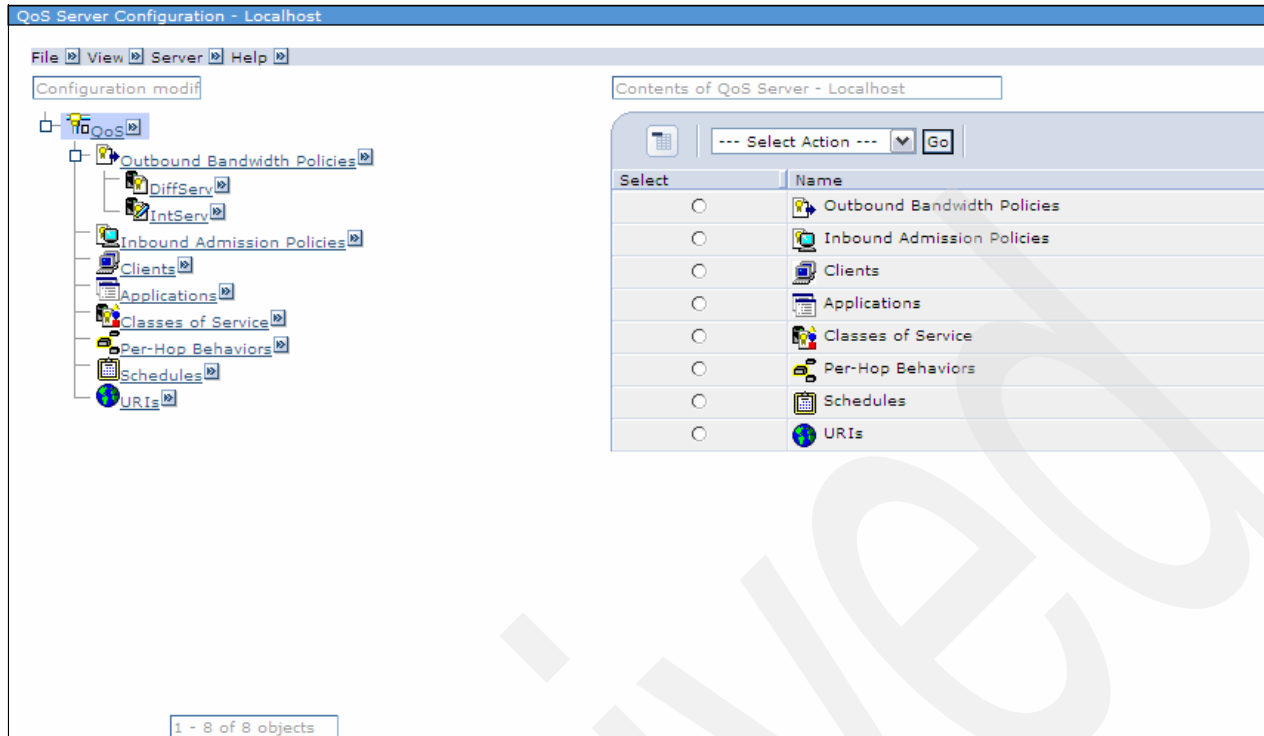


Figure 5-37 QoS Configuration wizard

For details about the quality of service refer to the IBM i Information Center at the following link and choose **Networking** → **TCP/IP applications, protocols, and services** → **Quality of service**:

<http://publib.boulder.ibm.com/infocenter/iseres/v6r1m0/index.jsp>

## 5.5 Internet

The Internet Setup Wizard will help you connect your system to a network and select the type of network that your system is connected to (Figure 5-38 on page 68). Options include the following:

- ▶ To a protected network

Your system is connected to a public network that anyone may access. You protect your system by setting up filtering rules and starting only services that you want people to be able to access. This is also known as a DMZ configuration.

The wizard helps you to do the following:

- Configure a connection to the network.
- Configure routes to your firewall and private network.
- Protect your system using IP packet filtering.
- Allow your system to provide Web serving.
- Allow your system to act as an HTTP proxy server.

Features of this option include the following:

- TCP/IP network connection
- A default route to your router
- A route to your intranet back through your firewall



- IP packet filtering
- ▶ To a private network

Your system is connected to a private network behind a firewall. Though the network is private, you may allow systems on the network to access external systems, such as for Web browsing.

The wizard helps you to do the following:

- ▶ Configure a connection to the network.
- ▶ Configure routes to your firewall and optionally to other networks.
- ▶ Protect your system with IP packet filtering.
- ▶ Configure a public IP address for the system (virtual IP).
- ▶ Allow your system to provide Web serving.
- ▶ Allow your system to act as an HTTP proxy server.

Features of this option include the following:

- ▶ TCP/IP network connection on your private network
  - A default route to your firewall
  - A network route to your subnet router
  - IP packet filtering for traffic from firewall
- ▶ Company/office to ISP

Your system connects to the Internet through a dial-up (modem) connection to an Internet service provider (ISP). You may also choose this option if you are working in a branch office of a company that has a private network that uses dial-up connections.

The wizard helps you to configure a dial-up (modem) connection to your ISP or private network. Features of this option include the following:

- ▶ A PPP profile that uses a dynamic IP address from your ISP
- ▶ A route for traffic to go to your ISP and the Internet when not resolved internally (It will be the default route for the PPP interface.)
- ▶ Network Address Translation (also known as *full masquerading*)
- ▶ IP packet filtering (Deny all incoming connections that are not a response from NAT.)

Refer to Figure 5-38.

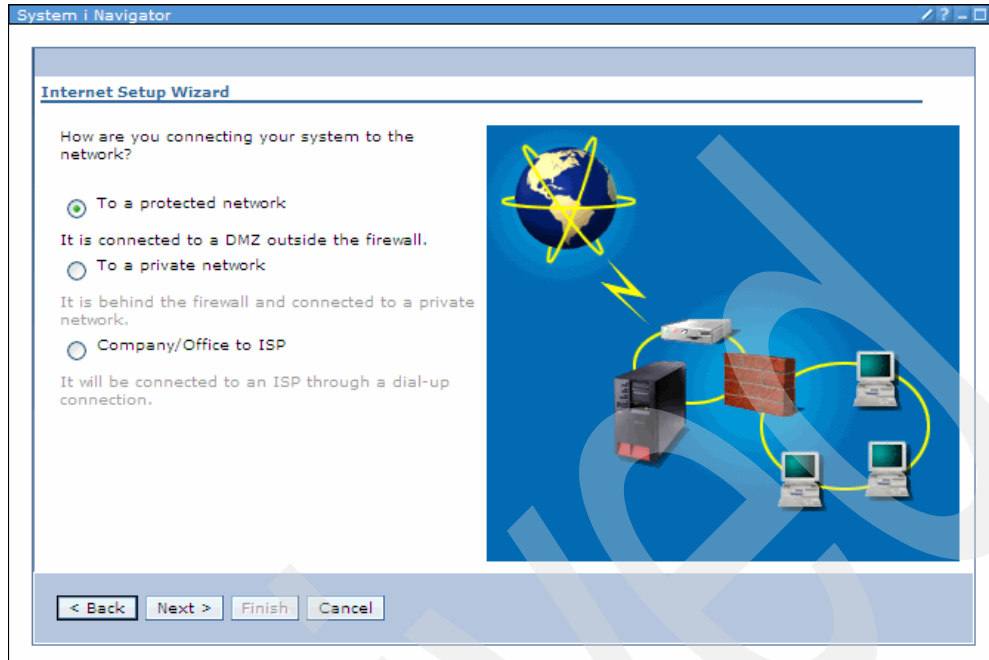


Figure 5-38 Internet Setup Wizard

## 5.6 i5/OS NetServer

In this section you can manage the following i5/OS NetServer functions:

- ▶ NetServer Sessions

This page lists all the active sessions on the system (Figure 5-39).

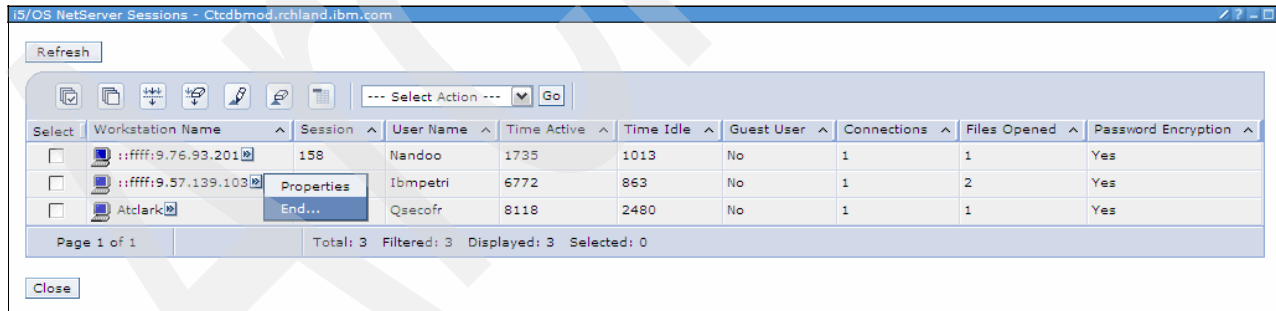


Figure 5-39 NetServer Sessions

- ▶ NetServer Disabled User IDs

If you are using Windows, an i5/OS server user ID may become disabled if the password for the Windows user ID and the i5/OS server user ID do not match or after a number of unsuccessful attempts to connect to i5/OS NetServer. Select the user ID that you want to enable and click **Enable user ID**.

► NetServer Properties

The General page displays the current i5/OS NetServer general properties. Changes can be made to current properties (Figure 5-40).

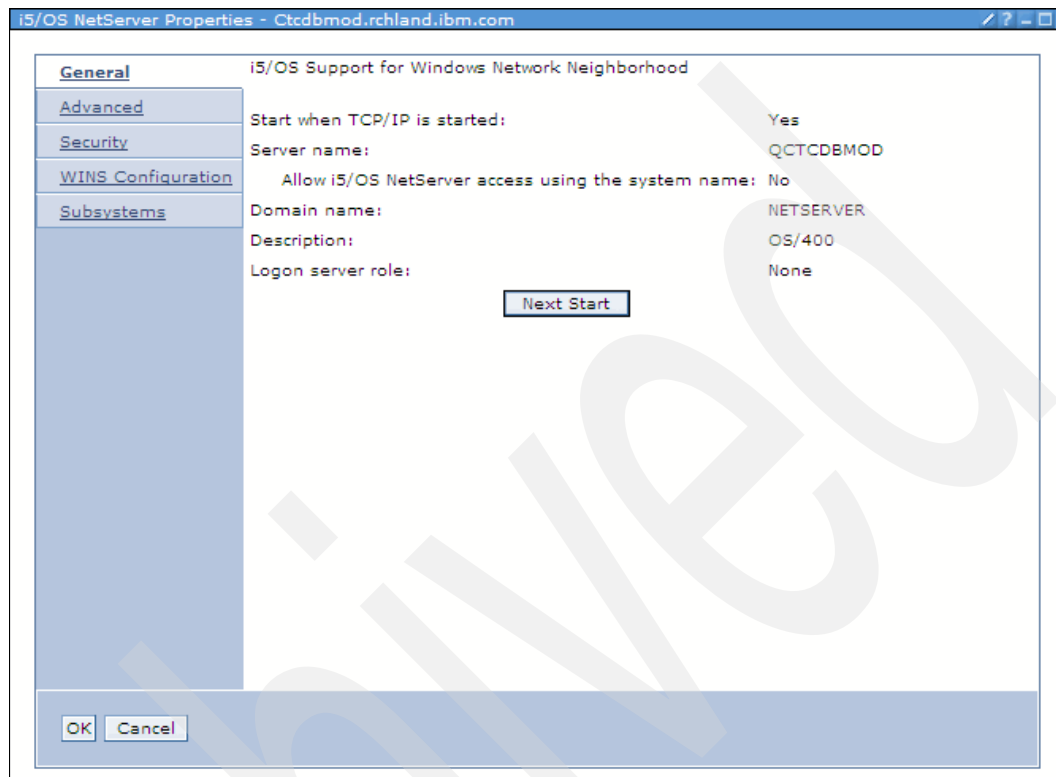


Figure 5-40 Netserver Properties

► NetServer Status

The Status dialog displays the current statistics for i5/OS, as shown in Figure 5-41 on page 70:

- Started  
Specifies the date and time that the server was started.
- Status since  
Specifies the time and date that the current set of statistics began recording.
- File opens  
Specifies the number of server files that have been opened.
- Print jobs queued  
Specifies the number of print jobs that have been spooled to the server.
- Sessions started  
Specifies the number of server sessions started.
- Sessions disconnected automatically  
Specifies the number of sessions that disconnected automatically.
- Sessions ended  
Specifies the number of sessions that disconnected normally.

- Unknown user attempts  
Specifies the number of unknown users who have tried to establish a connection to the server or the number of guest users that have established a connection to the server. This field is enabled if a guest user profile is unavailable.
- Password violations  
Specifies the number of server password violations.
- Average response time  
Specifies the average server response time in milliseconds.
- Bytes sent  
Specifies the number of server bytes sent to the network.
- Bytes received  
Specifies the number of server bytes received from the network.
- Refresh Now  
Retrieves the current statistics for the server and displays them in the dialog fields.
- Timed Refresh  
Opens a dialog box that allows you to set the time between refresh requests. Once set, timed refresh values are saved, so you do not have to refresh each time that the i5/OS NetServer Status dialog is opened.

**Note:** Timed refresh values are saved for each system, not for each user.

- Reset Status  
Resets all status values to zero and prompts the server to begin recounting the statistics from zero.

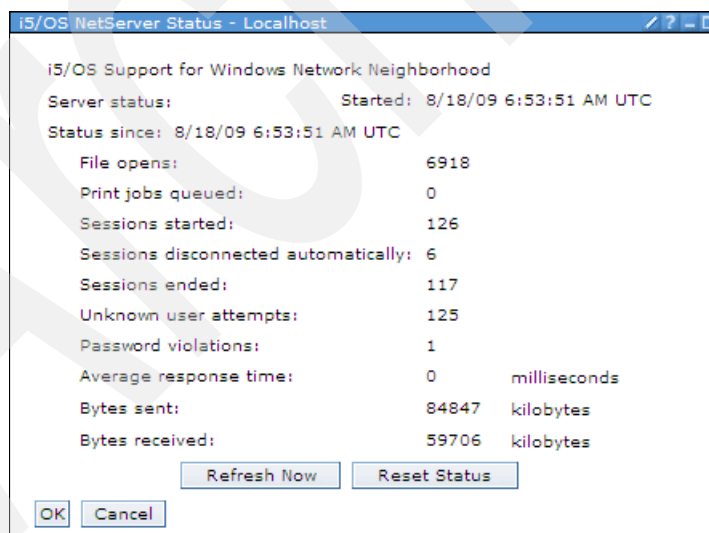


Figure 5-41 NetServer Status

# Integrated Server Administration

This chapter presents the tasks provided in the Integrated Server Administration GUI portlet. Over 60 tasks/panels are available to manage the integrated and virtual servers that the system hosts. Figure 6-1 shows the main functions view.

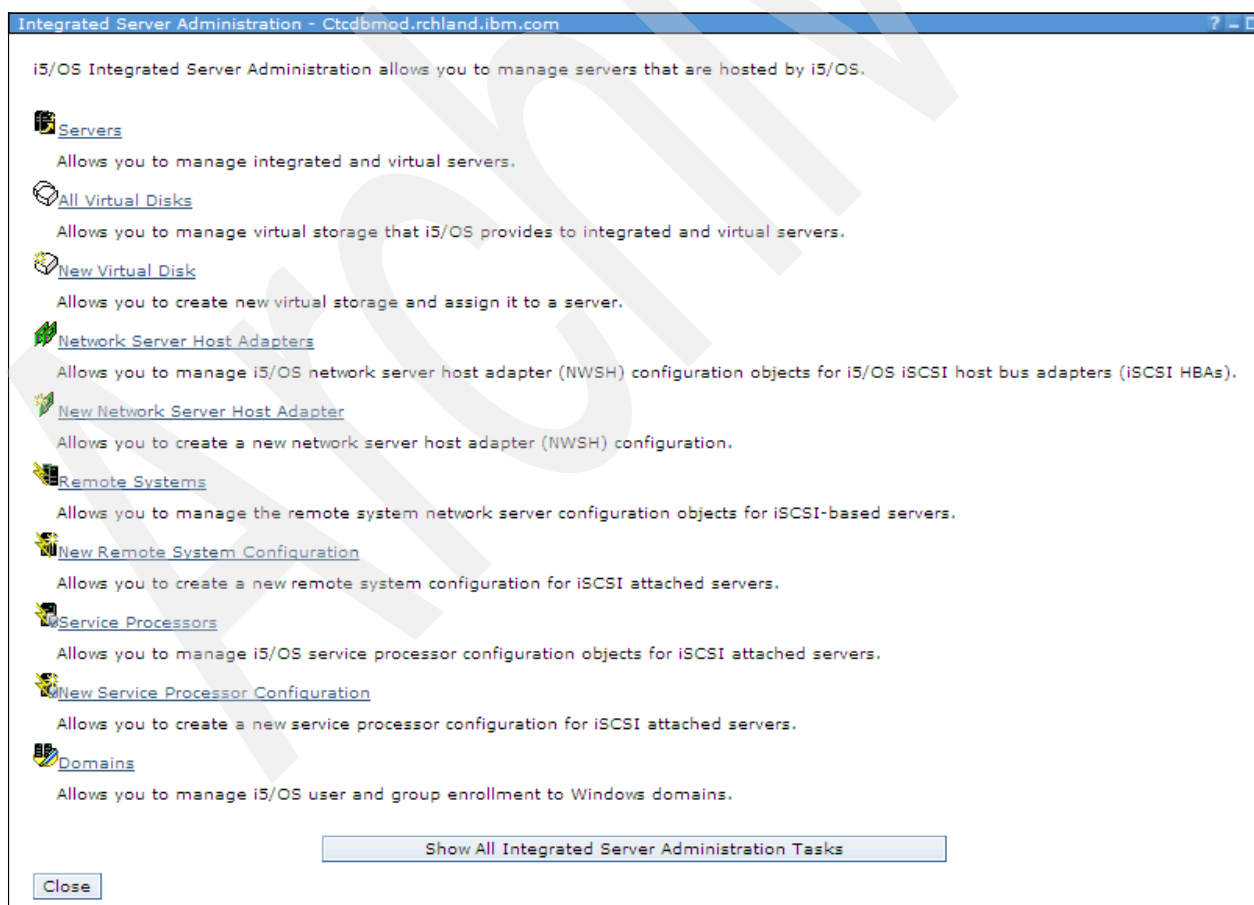


Figure 6-1 Integrated Server Administration main view

You can expand all the main functions by clicking **Show All Integrated Server Administration Tasks** and having the view of the sub-tasks available. Click the arrow next to the task that you are going to manage to get the list of the actions that you can take as shown in Figure 6-2.

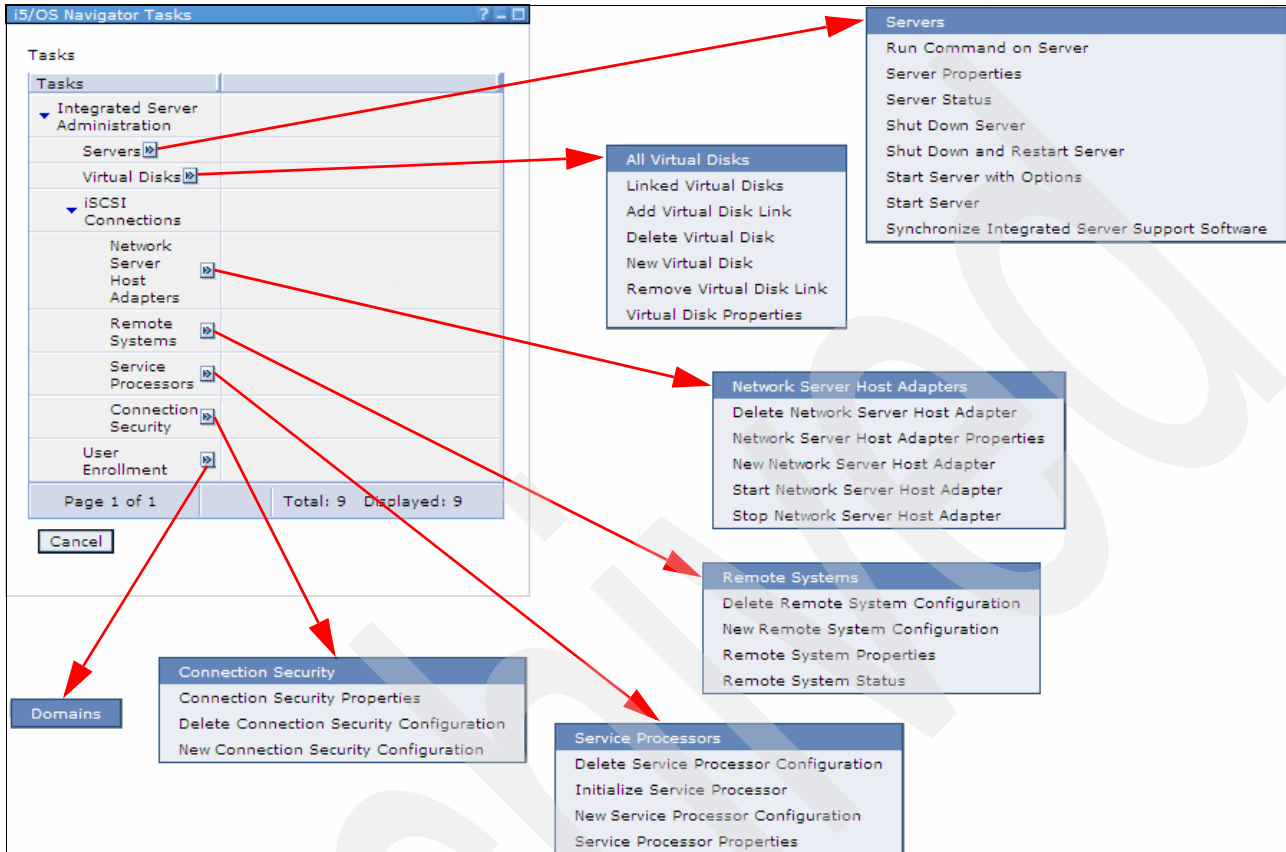


Figure 6-2 All of the available tasks

## 6.1 Server Management functions

This function lets you manage the integrated servers. These could be an Integrated xSeries® Server (IXS), an Integrated xSeries Adapter (IXA), or an iSCSI Integrated Server. The available options are:

- ▶ Show list of all integrated and virtual servers (All types: Windows, Linux®, VMware, AIX® and i5/OS).
- ▶ Start server.
- ▶ Start all Servers.
- ▶ Shut down server.
- ▶ Shut down all Servers.
- ▶ Restart server.
- ▶ Display server status.
- ▶ Submit remote Windows commands & view command output.
- ▶ Synchronize i5/OS integration code from i5/OS to Windows.
- ▶ Display/change server properties Hardware to run on, NWSHs to use (for iSCSI), S/W version installed, message logging, and so on.
- ▶ Hot spare server hardware (via server properties).

Also see Figure 6-3 and Figure 6-4 on page 74.

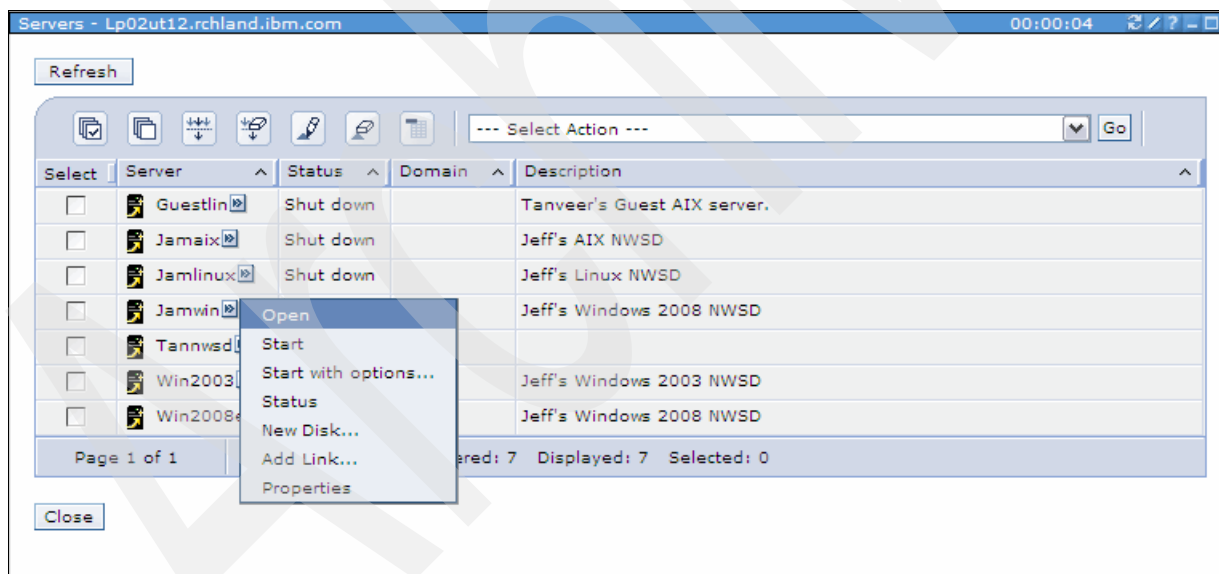


Figure 6-3 Servers pop-up menu

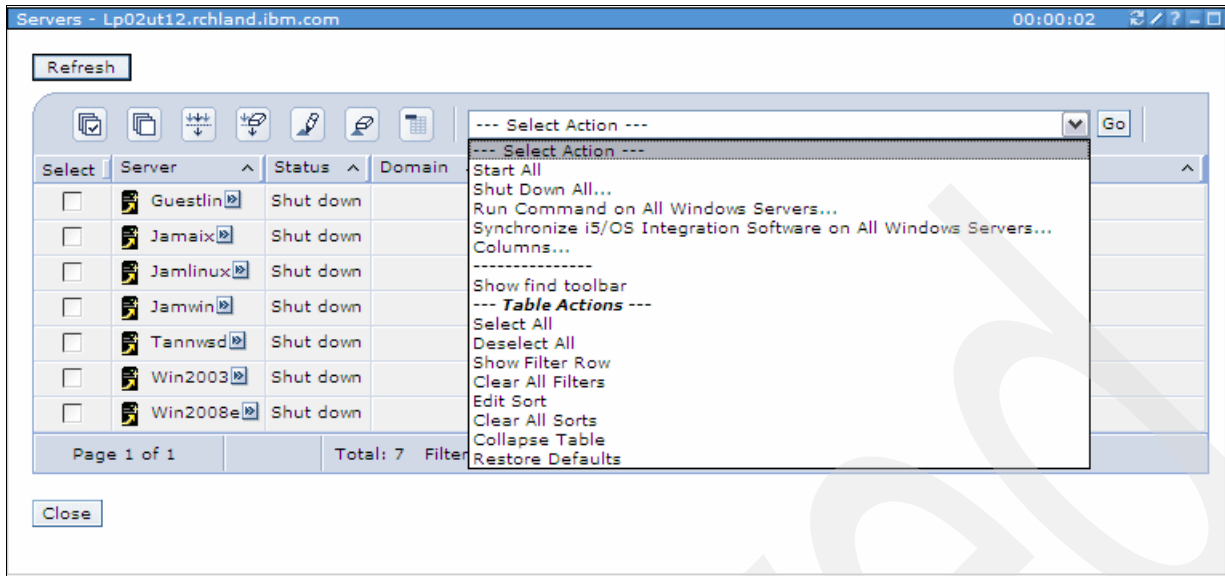


Figure 6-4 Servers drop-down list



## 6.2 Disk management functions

Use disk management functions to manage all the virtual disks that you created, which can be linked or not. The available options are:

- ▶ Create disk drive.
- ▶ Copy (New Based On).
- ▶ Grow disk drive (via properties, see Figure 6-6 on page 76).
- ▶ Link disk drive to server.
- ▶ Unlink disk drive from server.
- ▶ Can link/unlink dynamically to/from an active server.
- ▶ Delete disk drive.
- ▶ Display disk drive properties (Figure 6-7 on page 76).

See also Figure 6-5.

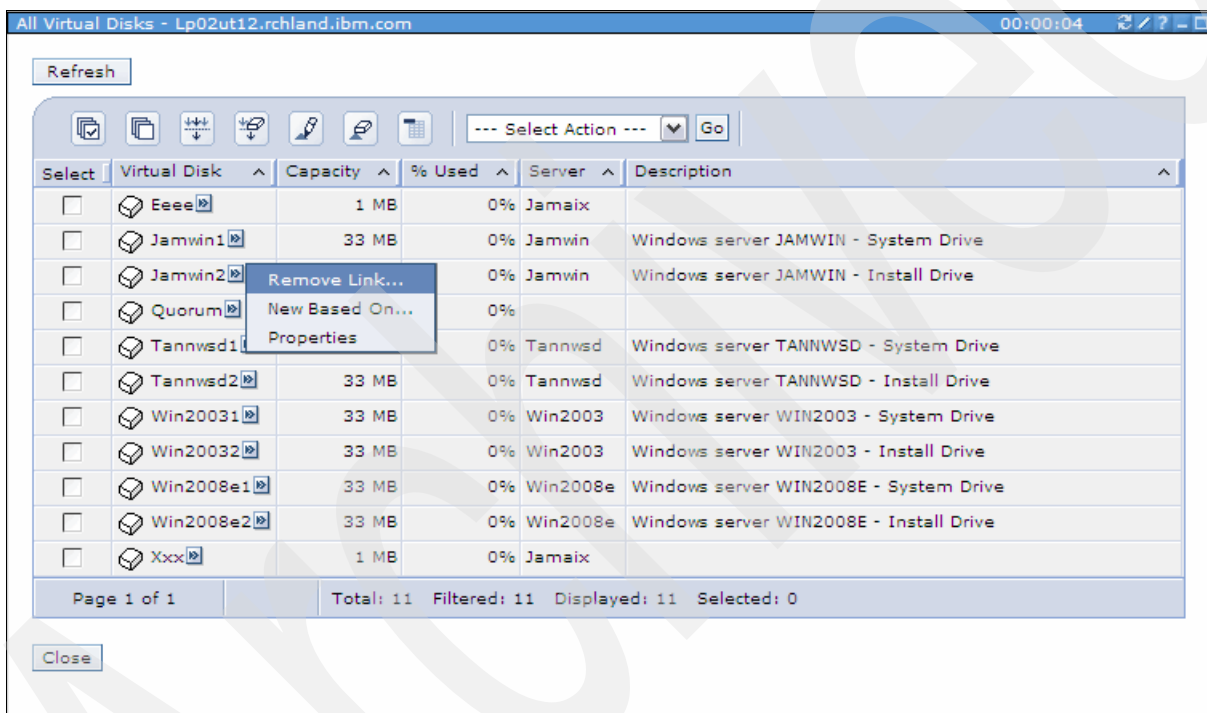


Figure 6-5 All Virtual Disks

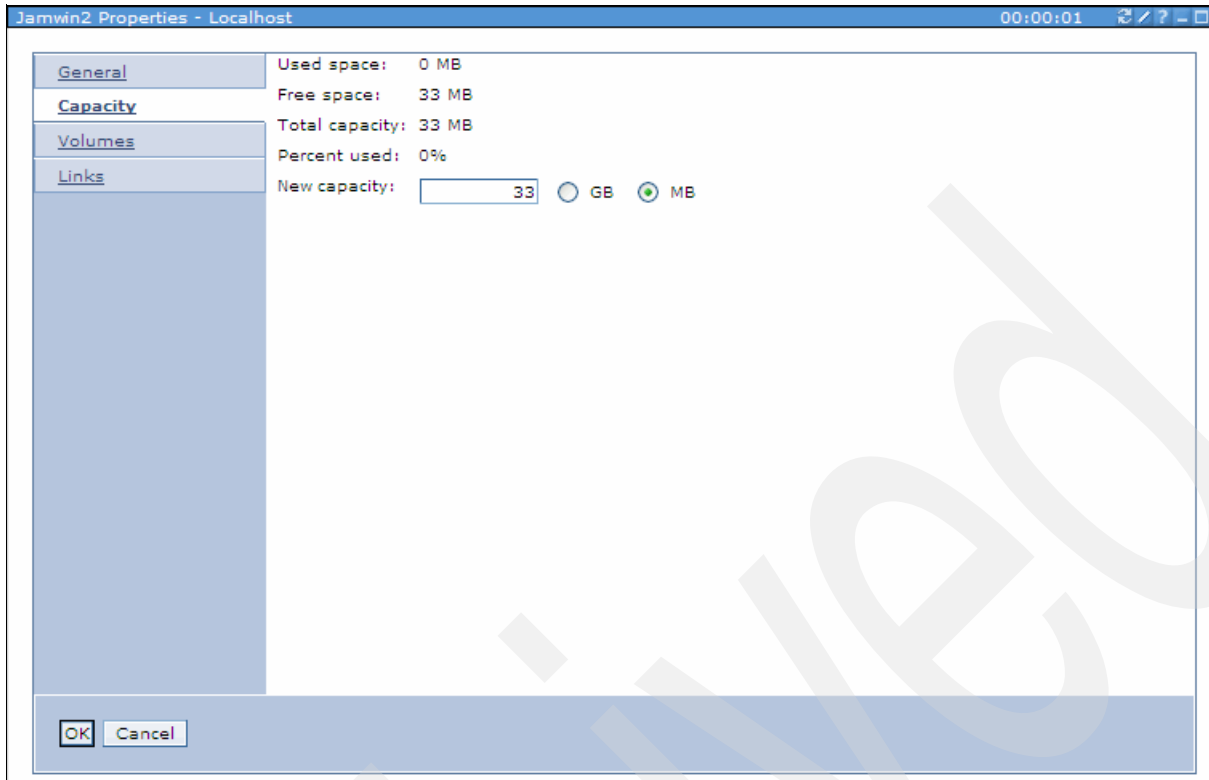


Figure 6-6 Disk properties: Capacity




Figure 6-7 Disk properties: General

## 6.3 User enrollment functions

Use these tasks to manage enrollment of i5/OS users and groups to your Windows domains and servers. User enrollment tasks are available from the domains listed in the Domains folder and from the servers listed in the Servers folder. The available options are:

- ▶ Show list, including status.
- ▶ Enroll users and groups to Windows domains and servers.
- ▶ Typically, create new Windows users via a template user.
- ▶ One-way attribute and password synchronization.
- ▶ Can manage Windows password in Windows (no i5/OS password).
- ▶ Display enrollment error details.
- ▶ Retry failed enrollments.
- ▶ Display enrollment properties.
- ▶ Unenroll users and groups from Windows domains and servers.
- ▶ Take advantage of EIM to enable single sign-on.

## 6.4 iSCSI: Network Server Host Adapter functions

The Network Server Host Adapters (NWSH) are used to configure the System i target iSCSI Host Bus Adapter (HBA). The available options are:

- ▶ Show list, including status (Figure 6-8).
- ▶ Start and stop (vary on/off).
- ▶ Create.
- ▶ Copy (New Based On).
- ▶ Delete.
- ▶ Display/change properties: network interfaces, message queue, online at IPL, and so on (Figure 6-9 on page 78).
- ▶ Hot spare iSCSI adapter hardware.

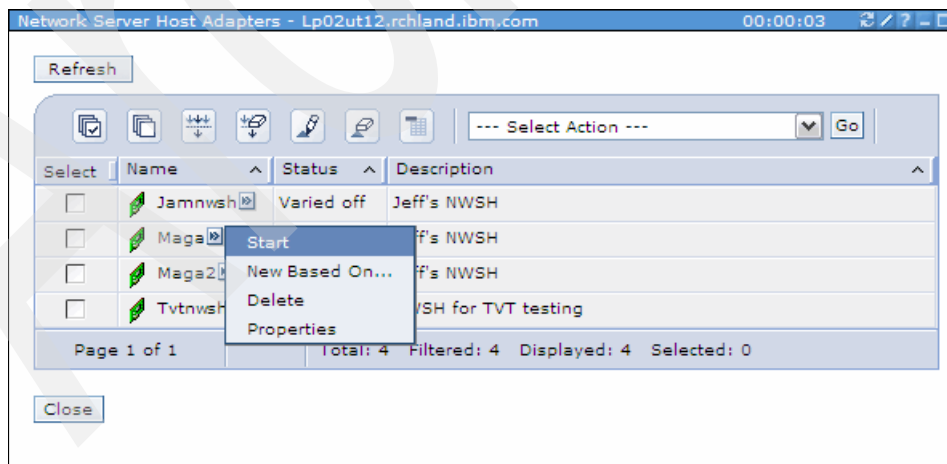


Figure 6-8 Network Server Host Adapter list

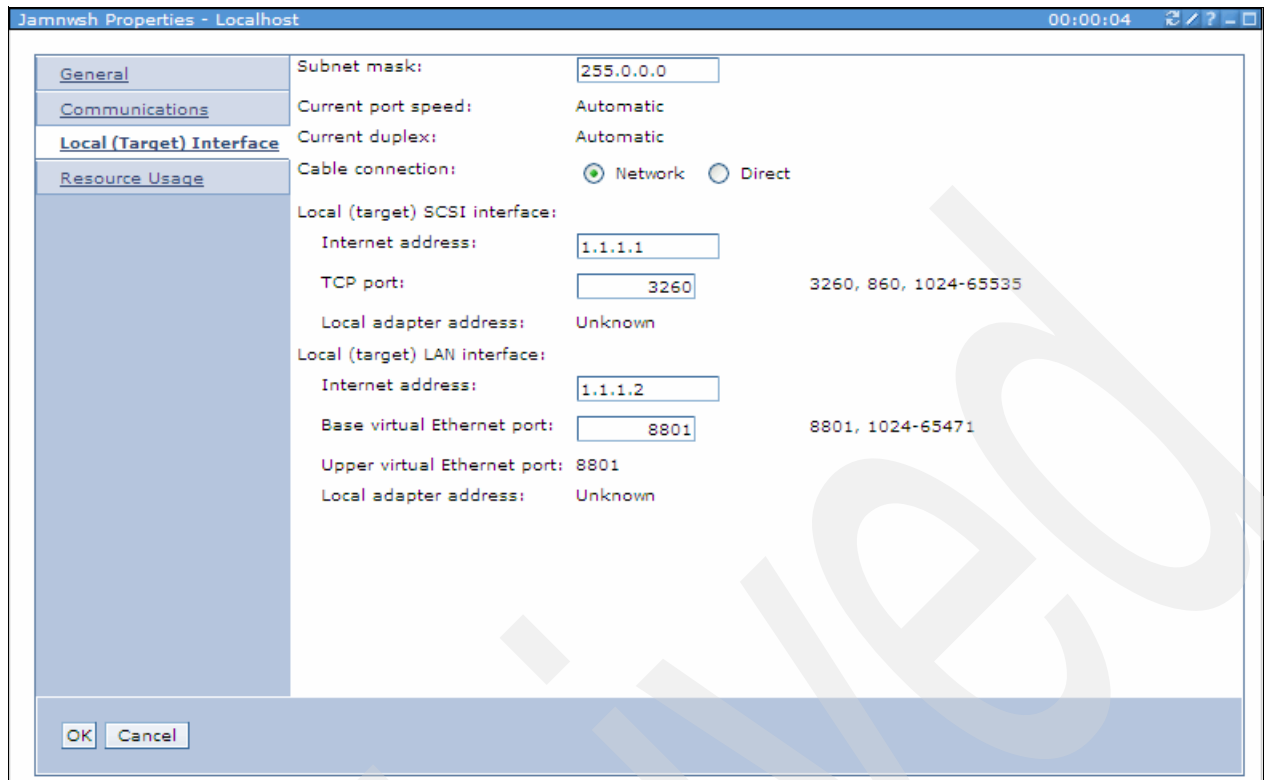


Figure 6-9 Network Server Host Adapter: Properties - local Interface

## 6.5 iSCSI: remote system functions

This is used to manage an iSCSI-attached System x or Blade within an IBM BladeCenter chassis. The available options are:

- ▶ Show list.
- ▶ Create.
- ▶ Copy (New Based On).
- ▶ Delete.
- ▶ Display/change properties: system identity, network interfaces, boot options, and so on.
- ▶ Show status of remote System x or blade server hardware: powered off/on, booting, OS running, and so on.

Also see Figure 6-10.

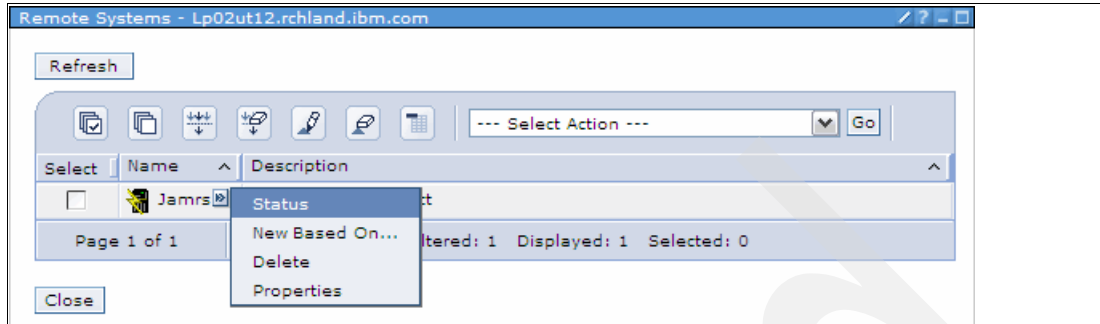


Figure 6-10 Remote Systems list

## 6.6 iSCSI: service processor functions

This is used to manage the service processor (SP) of a System x or of the Management Module (MM) of a IBM BladeCenter server. The available options are:

- ▶ Show list (Figure 6-11).
- ▶ Create.
- ▶ Copy (New Based On).
- ▶ Delete.
- ▶ Display/change properties: enclosure identity, authentication/security, and so on.
- ▶ Initialize System x or BladeCenter service processor (user name/password and so on).

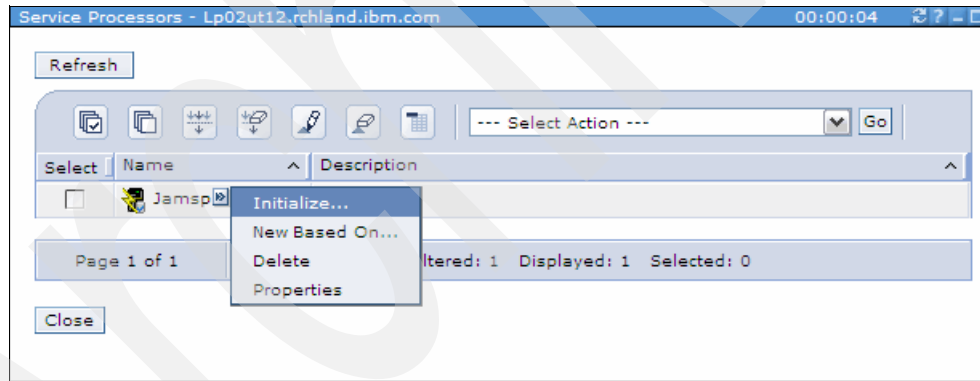


Figure 6-11 Service Processors list

## 6.7 iSCSI: connection security functions

The available options are:

- ▶ Show list.
- ▶ Create.
- ▶ Copy (New Based On).
- ▶ Delete.
- ▶ Display/Change properties.

For more information about the Integrated Server Administration refer to IBM i Information Center at the following link and choose **System i Integration**:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>

Archived



## Security, users, and groups

This chapter provides an overview of working with security, users, and groups in IBM Systems Director Navigator for i. An overview of IBM i authority terminology corresponding to IBM Systems Director Navigator for i is also included.

## 7.1 Security

IBM Systems Director Navigator for i has new security functions and interfaces that do not exist within the System i Navigator client application. Figure 7-1 shows the Security Menu tasks.

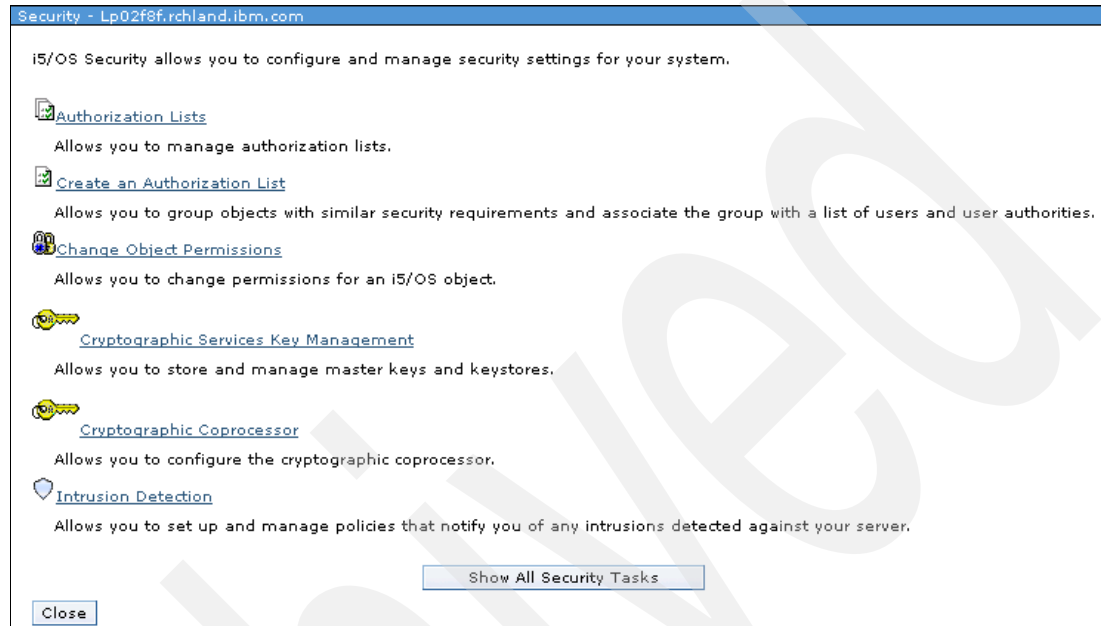


Figure 7-1 Security menu

The functions are described below.

### 7.1.1 Authorization list

An authorization list groups together objects with similar security requirements. An authorization list contains a list of users and groups and the authority that each has to the objects secured by the list. When you select this option, a new page appears (Figure 7-2).

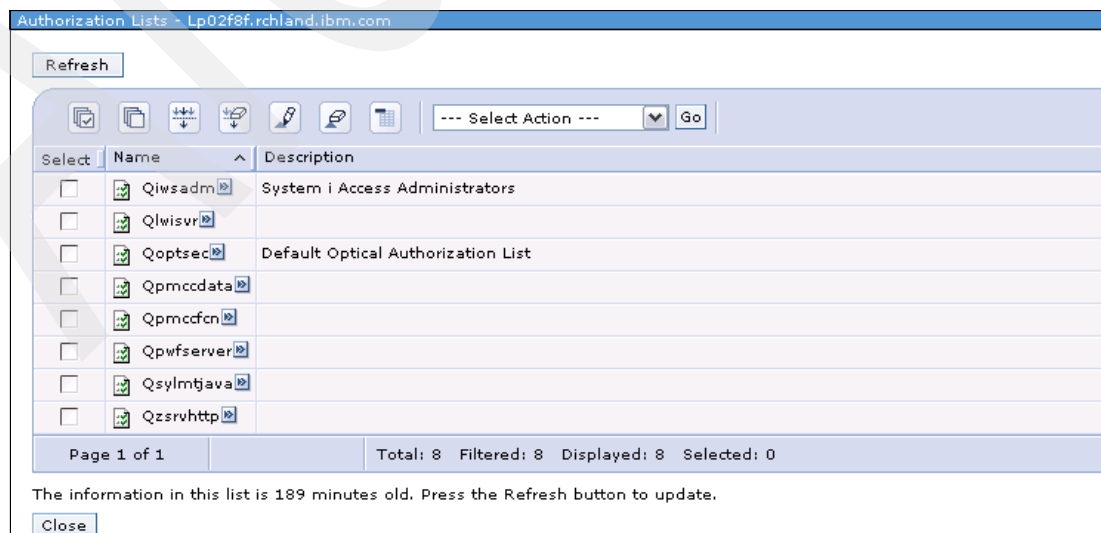


Figure 7-2 Authorization list functions



You can select any authorization list to perform different functions. Figure 7-3 shows the available options.

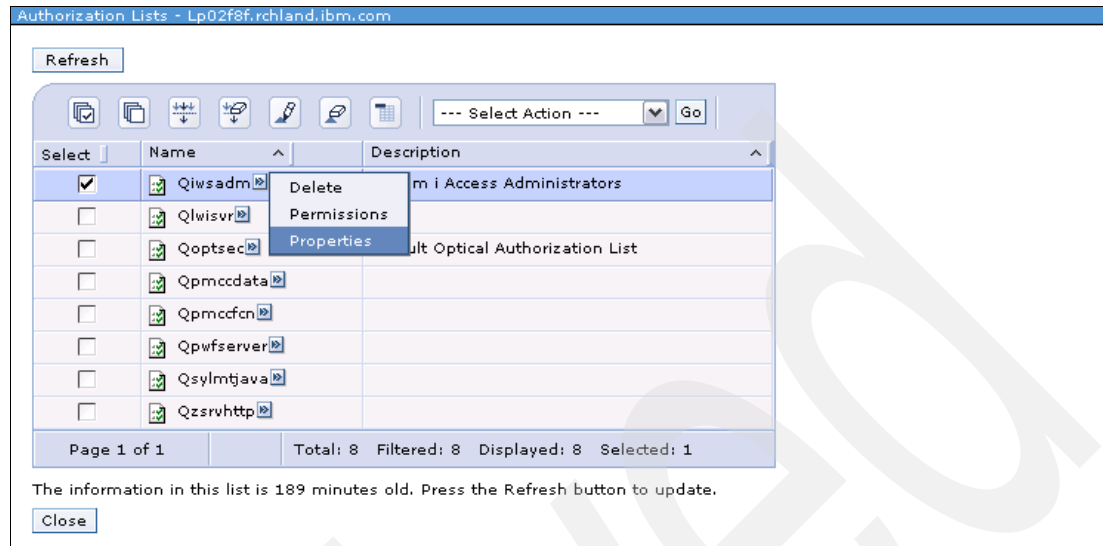


Figure 7-3 Selection of the options

As shown In Figure 7-3, you can select:

- ▶ Delete the authorization list.
- ▶ Permissions of the authorization list.
- ▶ Properties of the authorization list.

You can also select **Select Action** from the drop-down list, as shown in Figure 7-4.

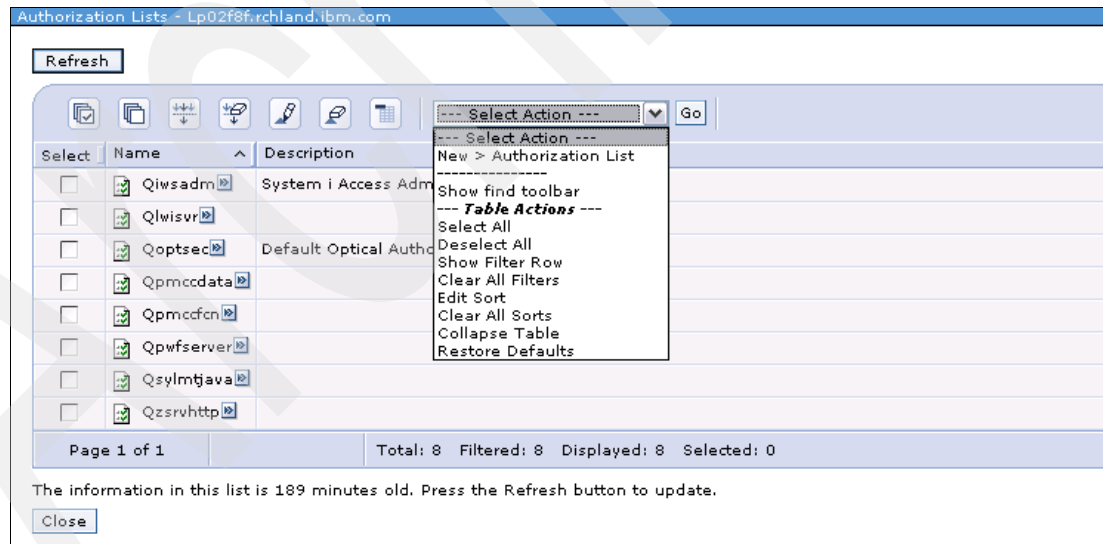


Figure 7-4 Selecting options

Below Select Action is the option New Authorization List. Select it if you want to create a new list. See 7.1.2, “Create an authorization list” on page 84.

## 7.1.2 Create an authorization list

As shown in Figure 7-5, you have options to the public authority:

- ▶ **USE:** The user can perform basic operations on the object, such as running a program or reading a file. The user cannot change the object. Use authority provides object operational authority, read authority, and execute authority.
- ▶ **CHANGE:** Change authority allows the user to change and perform basic functions on the object. The user can perform all operations on the object except those limited to the owner or controlled by object existence authority and object management authority.
- ▶ **ALL:** The user can perform all operations except those limited to the owner or controlled by authorization list management authority. The user can control the object's existence, specify the security for the object, change the object, and perform basic functions on the object. The user also can change ownership of the object.
- ▶ **EXCLUDE:** No user other than the owner can access the object.

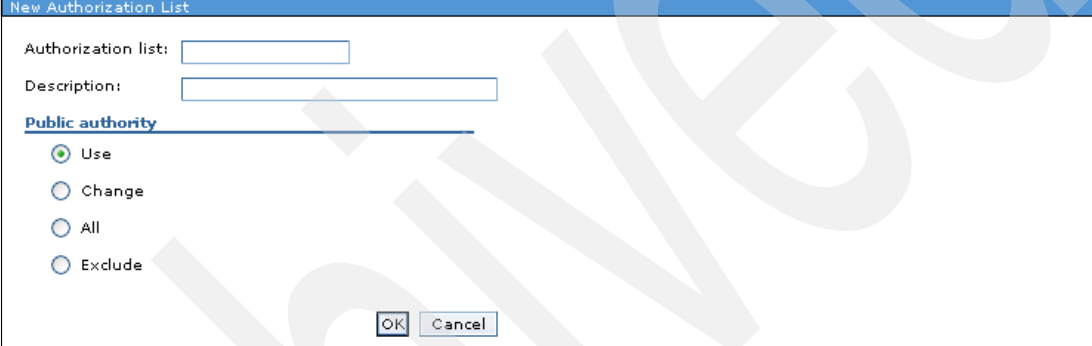


Figure 7-5 Create a new authorization list

## 7.1.3 Change object permissions

When you select this option, the page shown in Figure 7-6 appears. You need to write the correct path to modify the permissions on the object selected.

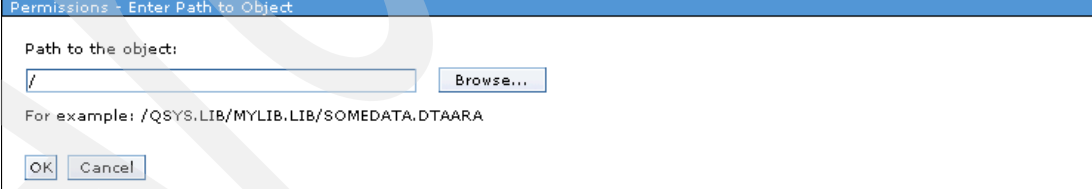


Figure 7-6 Enter path to the object

## 7.1.4 Cryptographic Services Key Management

From IBM Systems Director Navigator for i, Cryptographic Services Key Management allows you to store and manage master keys and keystores. Since you will be exchanging sensitive data to manage master keys and keystores, we recommend that you use a secure session.

Cryptographic Services supports a hierarchical key system. At the top of the hierarchy is a set of master keys. These keys are the only key values stored in the clear (unencrypted). Cryptographic Services securely stores the master keys within the i5/OS Licensed Internal Code (LIC).

Eight general-purpose master keys are used to encrypt other keys that can be stored in keystore files. Keystore files are database files. Any type of key supported by cryptographic services can be stored in a keystore file, for example, AES, RC2, RSA, and SHA1-HMAC.

In addition to the eight general-purpose master keys, Cryptographic Services supports two special-purpose master keys. The ASP master key is used for protecting data in the Independent Auxiliary Storage Pool (in the Disk Management GUI, this is known as an independent disk pool). The save/restore master key is used to encrypt the other master keys when they are saved to media using a save system (SAVSYS) operation.

Figure 7-7 shows the functions.

**Note:** The Cryptographic Services Key Management function is available in IBM i 6.1

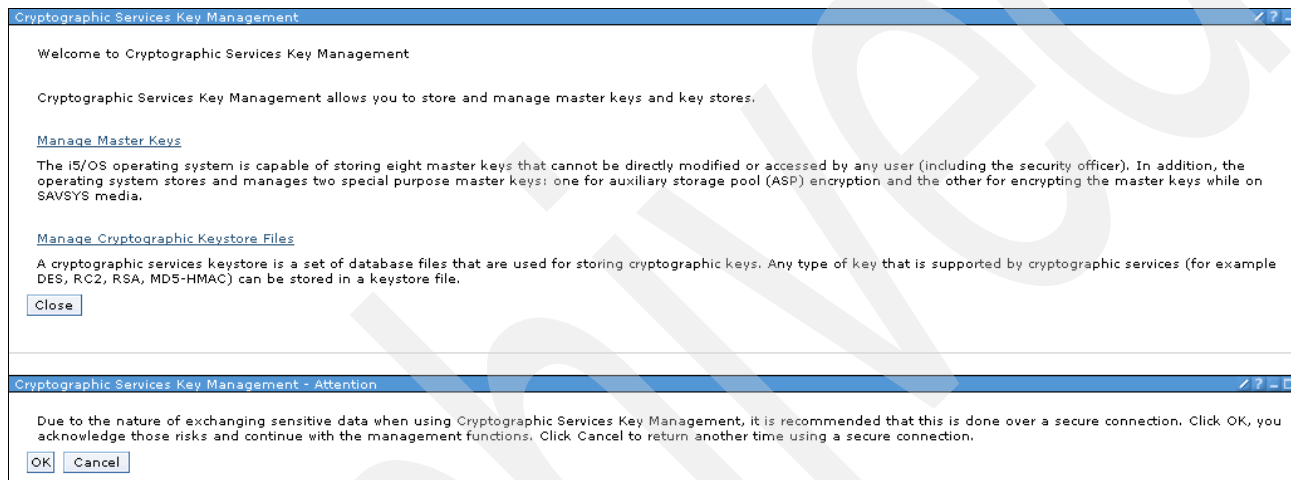


Figure 7-7 Cryptographic Services Key Management

As shown the Figure 7-7, Cryptographic Services Key Management allows you to store and manage master key and key stores.

### Managing master keys

Master keys are used to encrypt other keys. You can load, set, and test master keys. You can clear a master key only after you have set it.

Cryptographic Services allows you to set up eight general-purpose master keys and two-special purpose master keys that cannot be directly modified or accessed by the user (including the security officer). The two special-purpose master keys are the save/restore master key used for encrypting the master keys while on SAVSYS media and the auxiliary storage pool (ASP) master key used for ASP encryption. Cryptographic Services master keys are 256-bit AES keys that are securely stored within the i5/OS LIC.

Master keys are used to encrypt other keys. If a master key is lost, all keys encrypted under that master key, and consequently all data encrypted under those keys, are lost. It is important that you back up the master keys both by saving the passphrases and by using a SAVSYS operation. To protect the master keys while on the save media, they are encrypted with the save/restore master key.

**Note:** You should use Secure Sockets Layer (SSL) to reduce the risk of exposing key values while performing key management functions.

Each master key is composed of four 32-byte values, called versions. The versions are new, current, old, and pending:

- ▶ The new master key version contains the value of the master key while it is being loaded.
- ▶ The current master key version contains the active master key value. This is the value that will be used when a master key is specified on a cryptographic operation (unless specifically stated otherwise).
- ▶ The old master key version contains the previous current master key version. It is used to prevent the loss of data and keys when the master key is changed.
- ▶ The pending master key version holds a master key value that has been restored to the system but cannot be correctly decrypted.

Each version of a master key has a key verification value (KVV). The KVV is a 20-byte hash of the key value. It is used to determine whether a master key has changed or what version of a master key was used in an encryption operation.

The following section describes master key operations. All master key operations will create a Cryptographic Configuration (CY) audit record.

### **Managing cryptographic keystore files**

You can create keystore files and add, generate, delete, import, export, and retrieve attributes for key records.

A keystore is a set of database files that are used for storing cryptographic keys. Any type of key that is supported by cryptographic services can be stored in a keystore file. Some examples of the types of keys supported by cryptographic services are AES, RC2, RSA, and MD5-HMAC. You can create as many keystore files as you want and add as many key records as you want into a keystore file. Since each keystore file is a separate system object, you can authorize different users to each file. You can save and restore each keystore file at different times. This depends on how often key records are added to the keystore file and how often the master key for the keystore file is changed.


You can manage keystore files from the System i Navigator or the IBM System Director Navigator for i interfaces, or use the Cryptographic Services APIs or control language (CL) commands.

**Note:** You should use Secure Sockets Layer to reduce the risk of exposing key values while performing key management functions.

## 7.1.5 Cryptographic Coprocessor configuration

When you select this function you see the page shown in Figure 7-8.

### Cryptographic Coprocessor Configuration



You must run the Cryptographic Coprocessor configuration function using secure sockets layer (SSL) encryption.

Currently SSL encryption is: OFF

Do you wish to switch to a secure session, or return to the i5/OS tasks page?

**Notes:**

1. If you are configuring a coprocessor for an SSL session, you **must** select either a 4758-023 or a 4764 coprocessor.
2. You **must** navigate within this function using the navigation keys provided. The browser **Back** key will not be operational for security reasons.

Figure 7-8 Cryptographics Coprocessor Configuration menu

### The Cryptographic Coprocessor can be used to augment system

IBM offers a Cryptographic Coprocessor, which is available on a variety of system models. Cryptographic Coprocessors contain hardware engines, which perform cryptographic operations used by i5/OS application programs and i5/OS SSL transactions.

You can use a Cryptographic Coprocessor to implement a broad range of i5/OS-based applications. Examples are applications for performing financial PIN transactions, bank-to-clearing-house transactions, EMV transactions for integrated circuit (chip)-based credit cards, and basic SET block processing. To do this, you or an applications provider must write an application program using a security programming interface (SAPI) to access the security services of your Cryptographic Coprocessor. The SAPI for the Cryptographic Coprocessor conforms to the IBM Common Cryptographic Architecture (CCA). The SAPI is contained in the CCA Cryptographic Service Provider (CCA CSP), which is delivered as i5/OS Option 35.

To meet capacity and availability requirements, an application can control up to eight coprocessors. The application must control access to individual coprocessors by using the Cryptographic\_Resource\_Allocate (CSUACRA) and Cryptographic\_Resource\_Deallocate (CSUACRD) CCA APIs.

- ▶ You can use a Cryptographic Coprocessor along with DCM to generate and store private keys associated with SSL digital certificates. A Cryptographic Coprocessor provides a performance assist enhancement by handling SSL private key processing during SSL session establishment.
- ▶ When using multiple coprocessors, DCM configuration gives you the following options for using hardware to generate and store the private key associated with a digital certificate.
- ▶ The private key is generated in hardware and stored (retained) in hardware. With this option the private key never leaves the coprocessor, and thus the private key cannot be used or shared with another coprocessor. This means that you and your application must manage multiple private keys and certificates.

Although the private key is generated in hardware, it can be stored in a keystore file. This option allows a single private key to be shared among multiple coprocessors. A requirement is that each coprocessor must share the same master key. You can use the clone master keys page to set up your coprocessors to have the same master key. The private key is generated in one of the coprocessors and is then saved in the keystore file, encrypted under the master key of that coprocessor. Any coprocessor with an identical master key can use that private key.

**Note:** For further information go to:

<http://publib.boulder.ibm.com/infocenter/iseres/v6r1m0/index.jsp?topic=/rzamv/rzamvwhatnew.htm>

## 7.1.6 Intrusion detection

The term intrusion detection is used two ways in i5/OS documentation. In the first sense, intrusion detection refers to the prevention and detection of security exposures. For example, a hacker might be trying to break into the system using a user ID that is not valid, or an inexperienced user with too much authority might be altering important objects in system libraries. In the second sense, intrusion detection refers to the intrusion detection function that uses policies to monitor suspicious traffic on the system.

The intrusion detection and prevention system (IDS) notifies you of attempts to hack into, disrupt, or deny service to the system. IDS also monitors for potential extrusions, where your system might be used as the source of the attack. These potential intrusions and extrusions are logged as intrusion monitor audit records in the security audit journal and displayed as intrusion events in the Intrusion Detection System graphical user interface (GUI). You can configure IDS to prevent intrusions and extrusions from occurring.

Intrusion detection involves gathering information about attacks arriving over the TCP/IP network. Intrusions encompass many undesirable activities, such as information theft and denial of service attacks. The objective of an intrusion might be to acquire information that a person is not authorized to have (information theft). The objective might be to cause a business harm by rendering a network, system, or application unusable (denial of service), or it might be to gain unauthorized use of a system as a means for further intrusions elsewhere. Most intrusions follow a pattern of information gathering, attempted access, and then destructive attacks. Some attacks can be detected and neutralized by the target system. Other attacks cannot be effectively neutralized by the target system. Most of the attacks also make use of spoofed packets, which are not easily traceable to their true origin. Many attacks make use of unwitting accomplices, which are machines or networks that are used without authorization to hide the identity of the attacker. For these reasons, a vital part of intrusion detection is gathering information and detecting and preventing system attacks.

As shown in Figure 7-9, the Intrusion Detection Management window has four functions, which we discuss in the following sections.

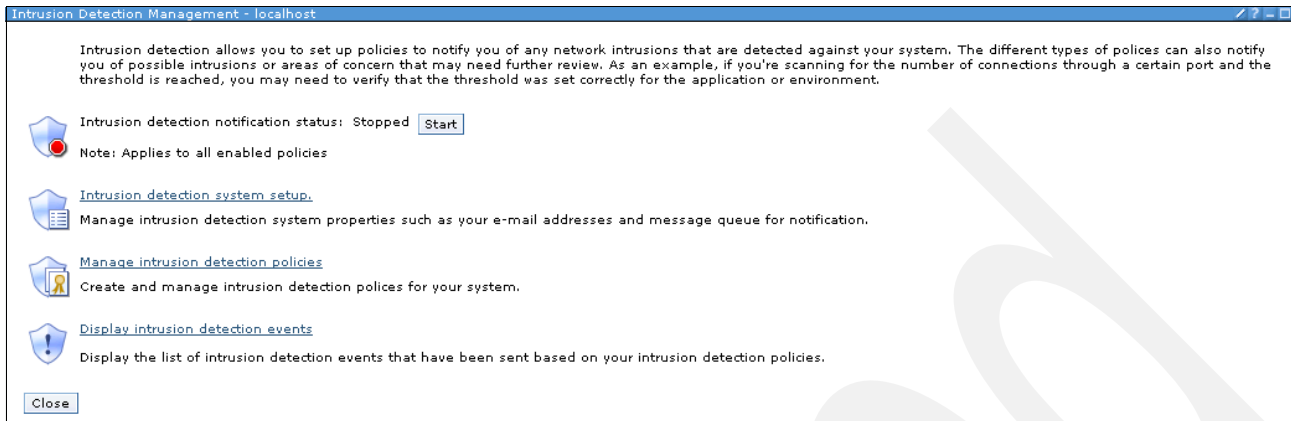


Figure 7-9 Intrusion Detection Management

### Intrusion Detection System Setup

IDS is a notification system. You can optionally configure IDS to send real-time intrusion notifications to a message queue and to specific e-mail addresses. That way, you can alert systems administrators about specific types of intrusions and extrusions so that they can take actions to stop further intrusions from occurring. You can enable or disable IDS e-mail and message notification per policy. See Figure 7-10.

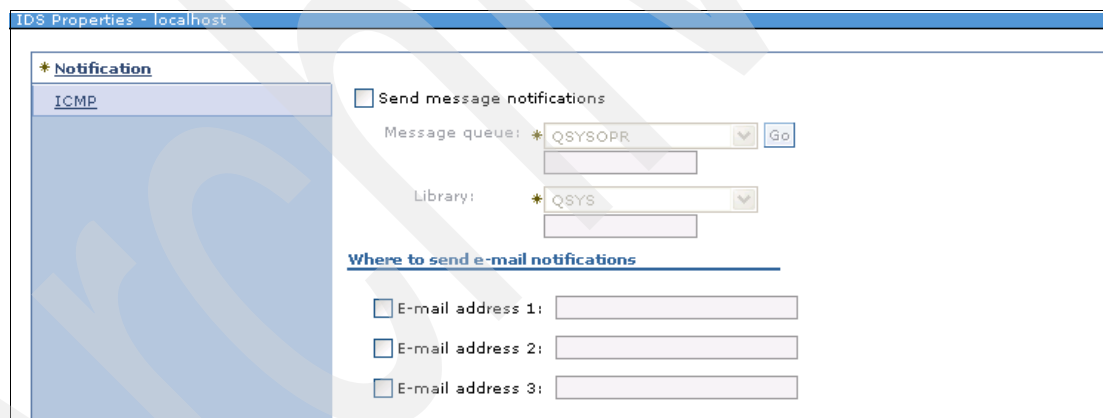


Figure 7-10 Intrusion detection setup properties

In the IDS Properties page, select the **Notification** tab. To send intrusion messages to a message queue, select the **Send message notifications** check box and specify the name of the message queue and library. (If the check box remains cleared, IDS does not send notifications to a message queue.) To send intrusion messages to an e-mail address, select the **E-mail address** check box and enter the e-mail address. You can send intrusion messages to up to three e-mail addresses. (If the check box remains cleared, IDS does not send notifications to an e-mail address.) To allow Internet Control Message Protocol (ICMP) redirect messages, click the **ICMP** tab and select the check box. (If the check box remains cleared, IDS does not notify you of ICMP redirect messages.) See Figure 7-11.

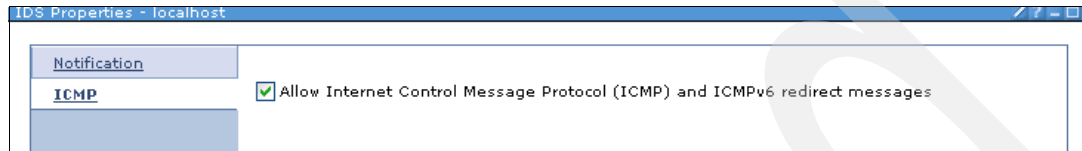


Figure 7-11 Internet control message protocol allowed

ICMP redirect messages are used to inform a host of a more optimal route to a destination. However, a hacker could send an ICMP redirect message to a host to have future traffic directed to the hacker's system.

**Note:** You must have \*ALLOBJ and \*IOSYSCFG authority to be able to display or change the IDS properties. To use IDS e-mail notification, System i SMTP must be configured and running.

### Manage intrusion detection policies

You can create a set of default intrusion detection policies that will monitor for all types of intrusions or extrusions for the entire system. You can also create specific attack, scan, and traffic regulation policies. See Figure 7-12.

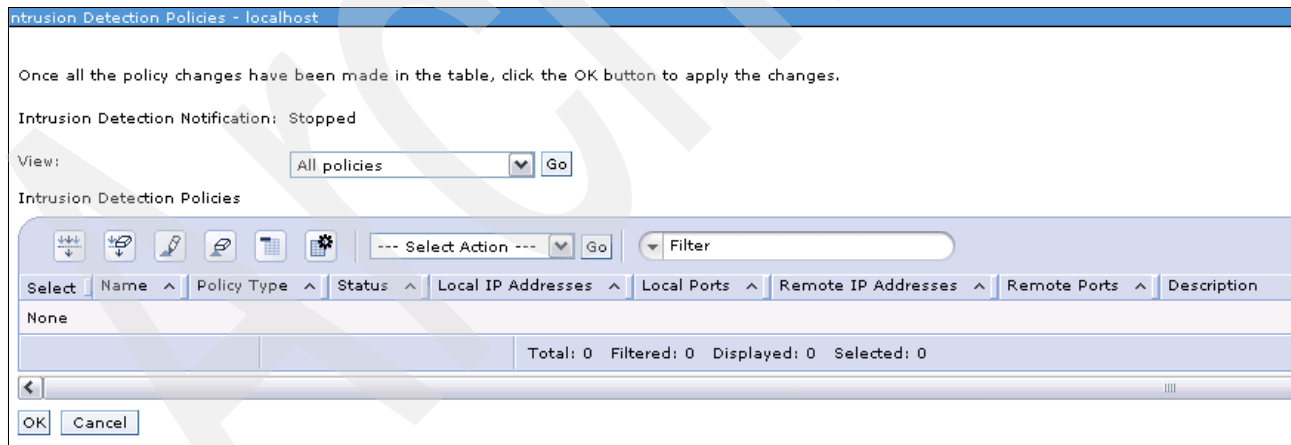


Figure 7-12 Intrusion Detection Policies

On the Intrusion Detection Policies page, select **New** from the Actions menu. The New intrusion detection policy wizard is displayed.



On the Select Policy to Create page, select **Create a set of default intrusion detection policies**. (This function is disabled if the default policies already exist.) See Figure 7-13.

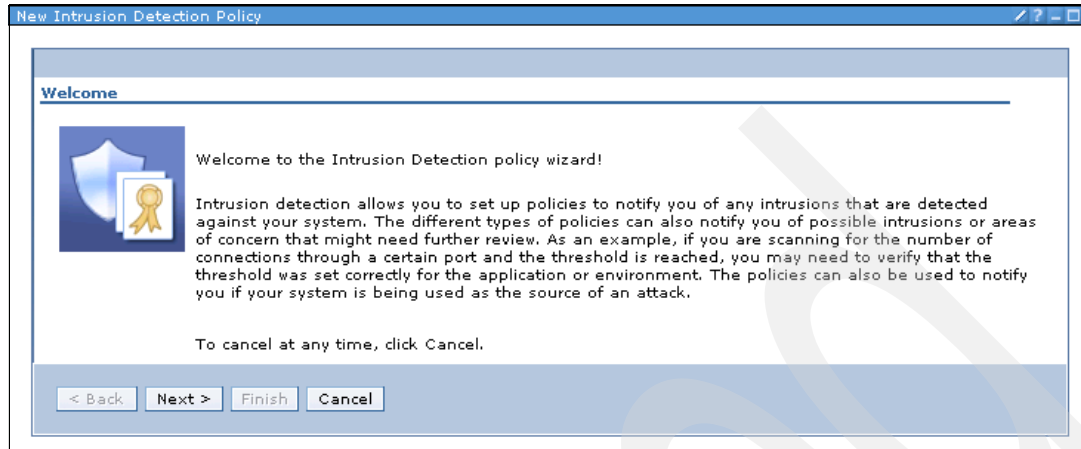


Figure 7-13 New Intrusion Detection Policy

Follow the instructions in the wizard to create the policies. Click **OK** on the Intrusion detection policies page to apply the changes. Now your system is ready to catch suspicious events coming in through the TCP/IP network.

### **Intrusion detection events**

**Note:** You must have \*ALLOBJ and \*IOSYSCFG authority to work with intrusion detection policy.

Use the Intrusion Detection System GUI to display a list of potential intrusion events as well as detailed information about each event. If you are using System i Navigator, you also can refresh the intrusion events at specified intervals. See Figure 7-14.

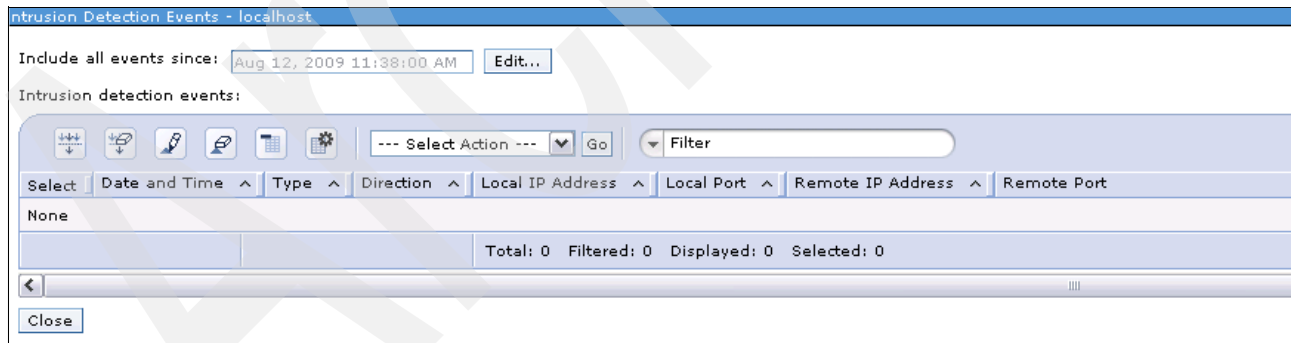


Figure 7-14 Intrusion Detection Events

By default, the Intrusion Detection Events page lists events that have occurred in the previous 24 hours. Perform any of the following tasks.

To refresh the intrusion detection events immediately, select **Refresh Now** from the Actions menu. To refresh the intrusion detection events at intervals, select **Timed Refresh** from the Actions menu. On the Timed Refresh page, check to turn on the timed refresh and enter the time in minutes between refreshes.

To display event details, select the event and select **Details** from the context or Actions menu. You also can find these event details in the intrusion monitor audit record.

If the QAUDJRN is no configured into the QSYS library before you active this function, the panel shown in Figure 7-15 appears.

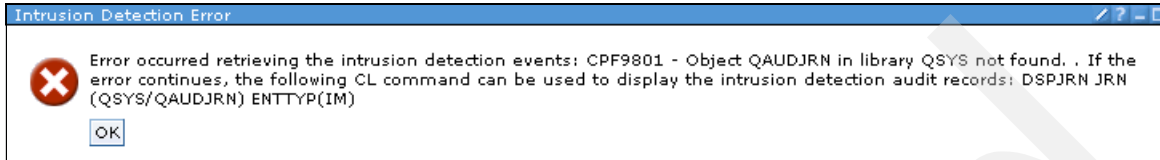


Figure 7-15 Intrusion Detection Error

To continue, click **OK** and you will see the window shown in Figure 7-14 on page 91.

## 7.2 Users and Groups page

The Users and Groups page displays a complete list of the users and groups configured on the system. You can use this functional area to manage users and groups. Figure 7-16 illustrates the Welcome page for the Users and Groups function.

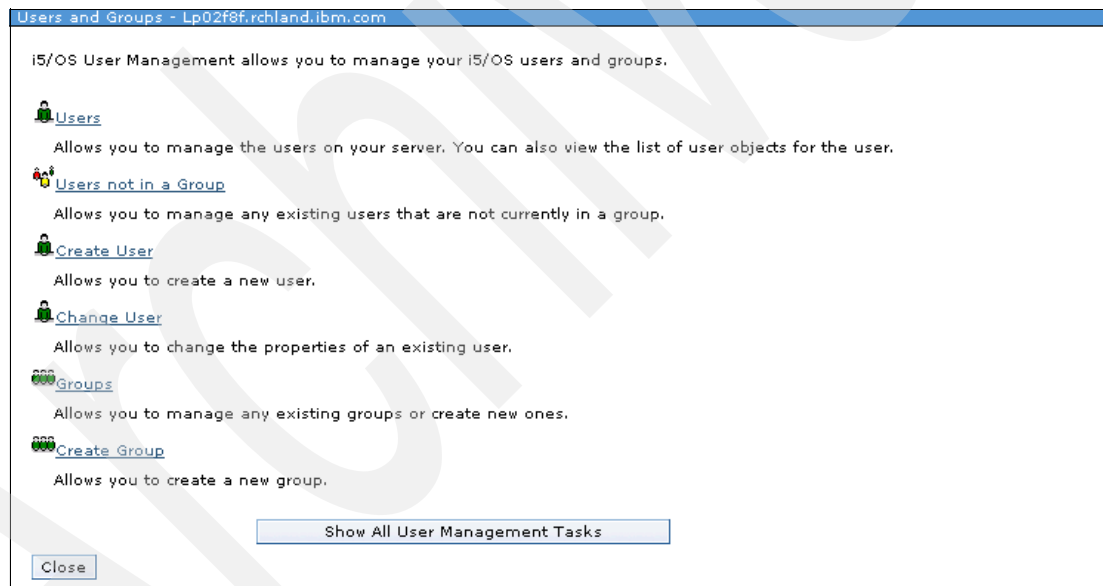


Figure 7-16 Users and Groups

## 7.2.1 Users

As shown in Figure 7-16 on page 92, the first option, Users, allows you to manage the users on your system as well as view the list of user objects for each user. You can also create a new user based on a specific user profile, delete a user profile, send a message to a user, view printer output, and so on. Figure 7-17 shows you the specific tasks mentioned.

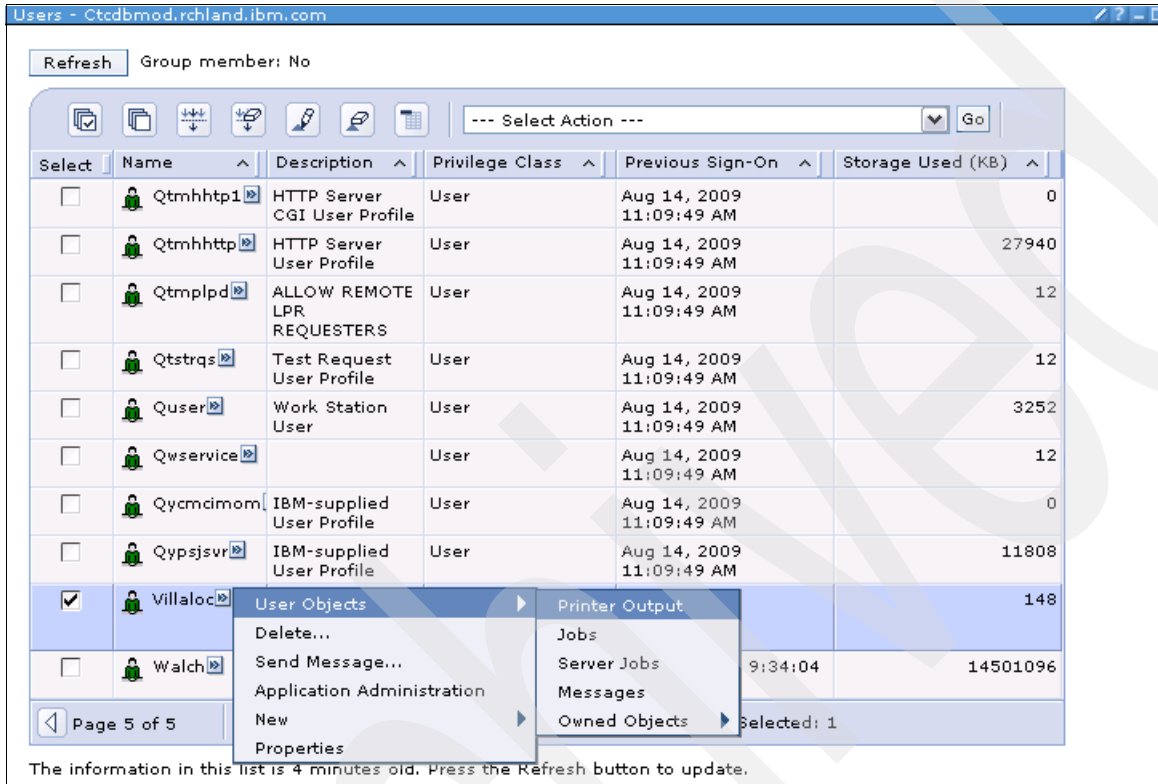


Figure 7-17 Users tasks

As shown in Figure 7-17, when you select a user, a drop-down list appears showing you several available tasks.

If you select New from this list, as shown in Figure 7-18, you can create a new user ID based on the ID previously selected. You are copying the same characteristics from the selected user ID to the new ID.

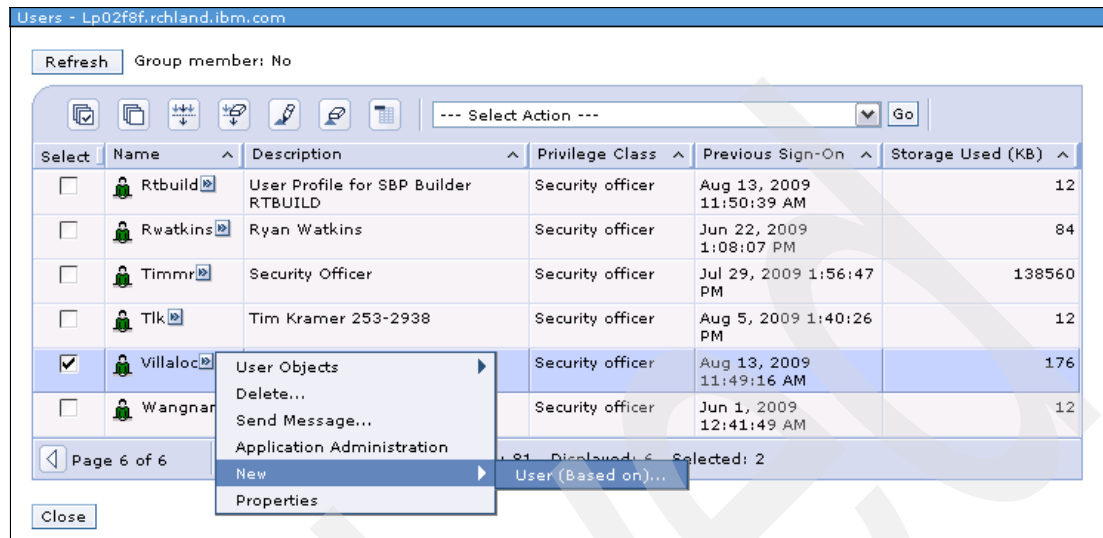


Figure 7-18 New user

After that, you see the window shown in Figure 7-19.

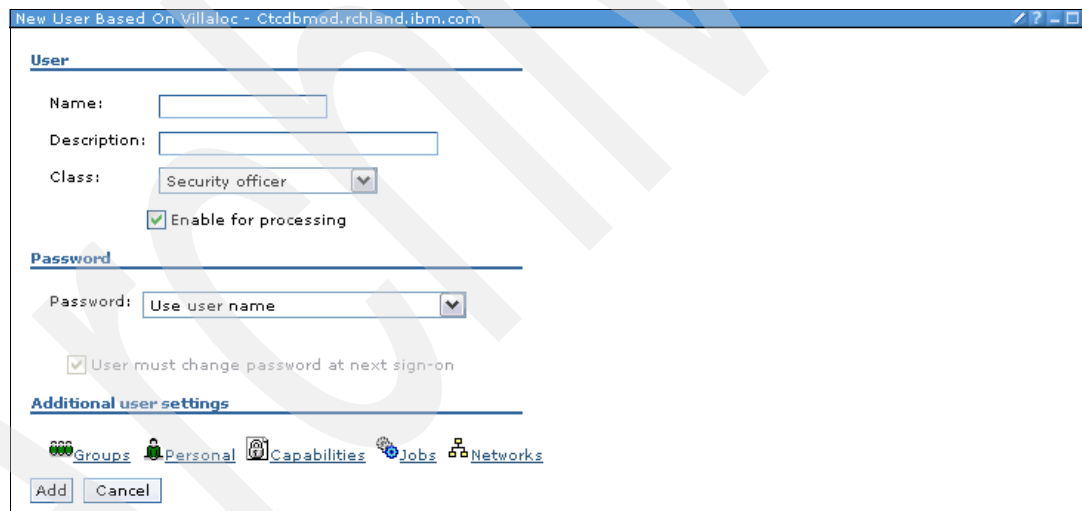


Figure 7-19 New user based on existing user

## 7.2.2 Users not in a Group task

The Users not in a Group task has the same basic function as the Users task previously described. This option shows only the user IDs that are not in a group.

## 7.2.3 Create user

The Create User task allows you create a new user ID. You have the option to create a new user ID based on an existing ID. See Figure 7-20.

The screenshot shows a web browser window titled "New User - Ctcdbmod.rchland.ibm.com". The main content area is titled "User" and contains the following fields and options:

- Name:** A text input field.
- Description:** A text input field.
- Class:** A dropdown menu with "User" selected.
- Enable for processing

The next section is titled "Password" and contains:

- Password:** A dropdown menu with "Use user name" selected.
- User must change password at next sign-on

The final section is titled "Additional user settings" and features a row of icons for "Groups", "Personal", "Capabilities", "Jobs", and "Networks". At the bottom of the dialog are "Add" and "Cancel" buttons.

Figure 7-20 Creating a new user

You must complete the fields to create the user ID.

## 7.2.4 Change user

The Change User task allows you to change characteristics of the selected user ID, as shown in Figure 7-21.

The screenshot shows a dialog box titled "User Information". The text inside reads "Please provide the user information for your task:". Below this text is a "User name:" label followed by a text input field containing the text "villalod" and a "Browse..." button. At the bottom of the dialog are "OK" and "Cancel" buttons.

Figure 7-21 Change selected user

After you enter the name of the user ID to change, click the **OK** button. A page like that shown in Figure 7-22 will appear.

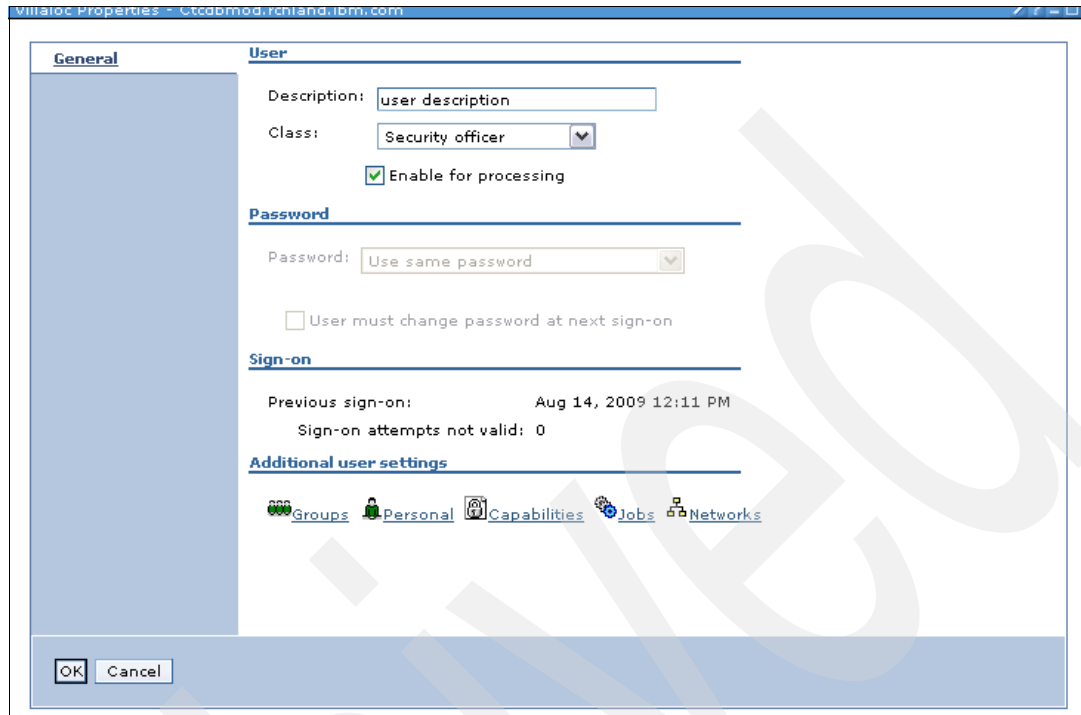


Figure 7-22 User properties

You can change more specific characteristics by selecting the options in the Additional user settings section.

## 7.2.5 Groups

The Groups task allows you see the groups configured on the system, as shown in Figure 7-23.

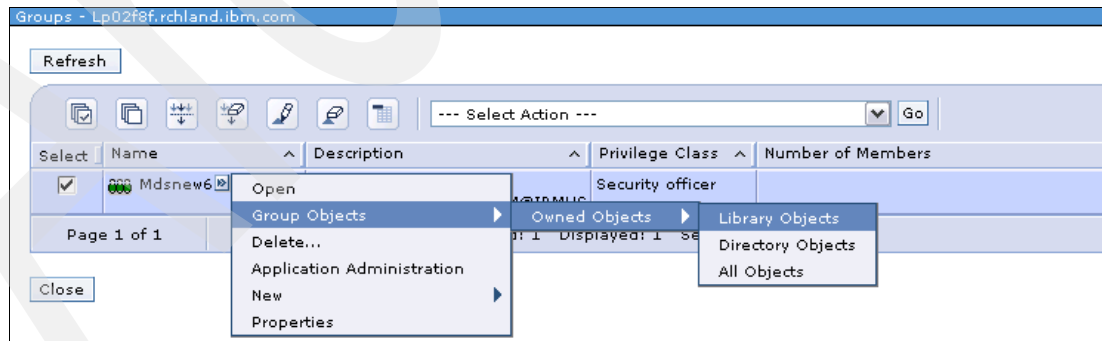


Figure 7-23 Groups

When you select a group (as in Figure 7-23 on page 96) and expand its menu, you can select several options for the group. Figure 7-24 shows that you can create a new user as a member of the selected group, as well as create a new group based on the selected group. For further information about user ID creation, see 7.2.3, “Create user” on page 95.

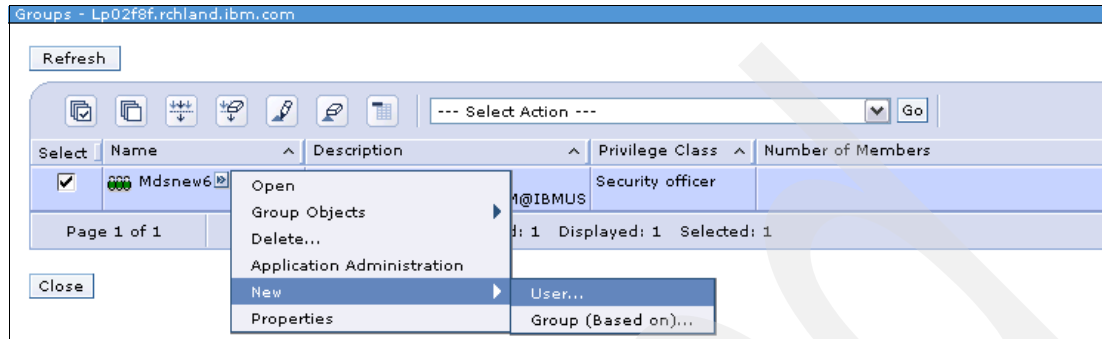


Figure 7-24 New User and Group

You can also select the drop-down list to see additional options, as shown in Figure 7-25.

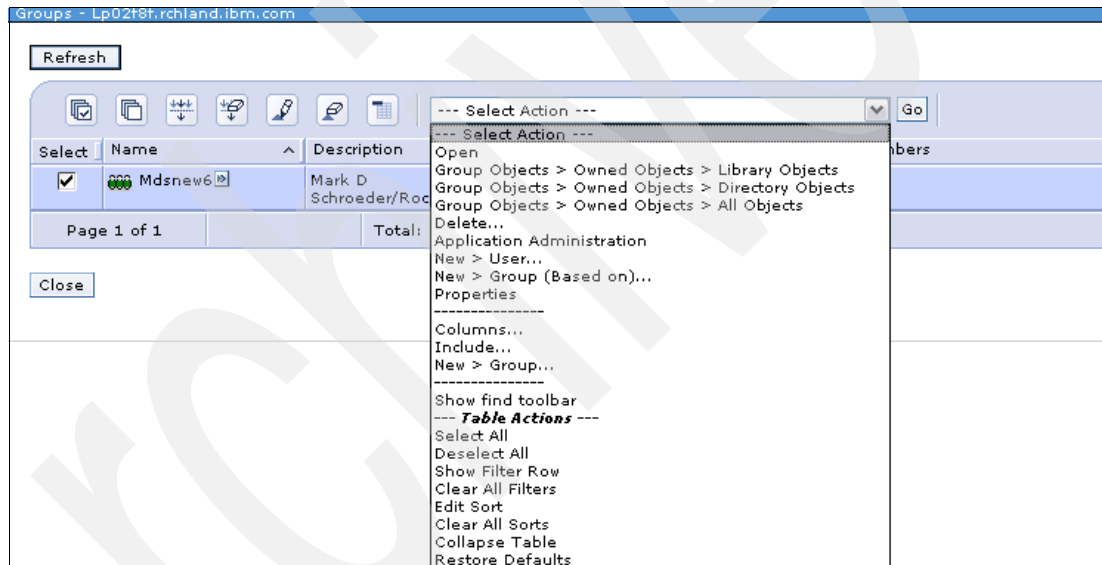


Figure 7-25 Additional options

## 7.2.6 Create Group task

The Create Group task allows you create a new group, as shown in Figure 7-26.



Figure 7-26 Create a new group

You will be asked whether the new group should be based on an existing group, as shown in Figure 7-26 on page 97. The default option is No. Click **OK** and the next window appears (Figure 7-27).

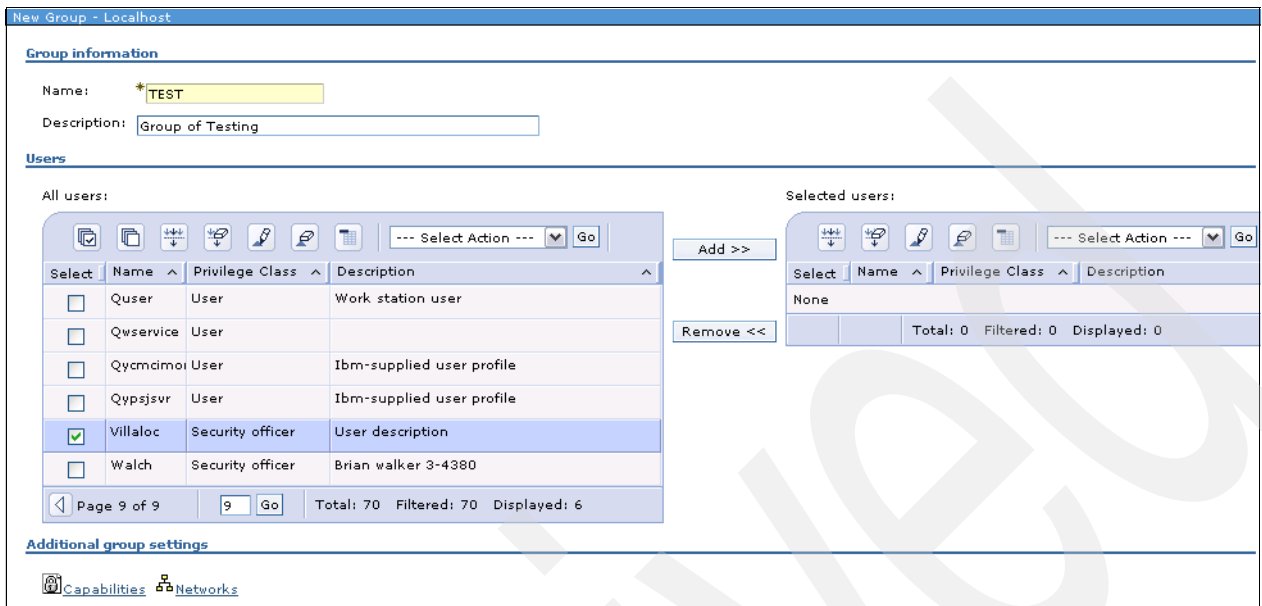


Figure 7-27 Creating a new group

As shown in Figure 7-27, when you create a new group, a list of user IDs appears. You can select user IDs and add them to your new group by clicking **Add**. The added user IDs will be moved to the Selected users list, as shown in Figure 7-28.

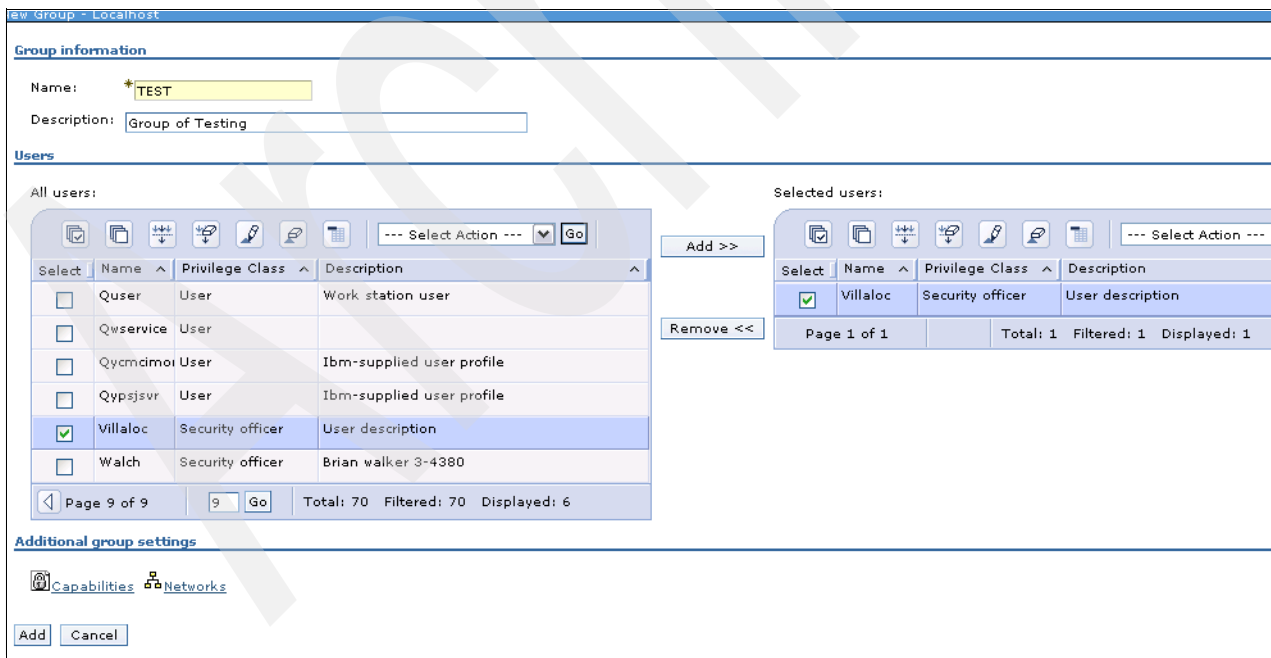



Figure 7-28 Adding user to group



When you finish adding users, click **Add** at the bottom of the window. A confirmation window will be displayed, as shown in Figure 7-29.



*Figure 7-29 New group created*

Archived



## Database and journal management

This chapter describes the database and journal tasks that you can access using the IBM Systems Director Navigator for i. This chapter discusses the following topics:

- ▶ Working with and creation of objects in a database schema
- ▶ Working with and changing database settings
- ▶ Creation of and working with database monitors
- ▶ Working with journals

## 8.1 Working with database tasks

IBM Systems Director Navigator for i has a new Web interface compared to System i Navigator. This chapter provides an overview of links to the most commonly used tasks available in the database section of the Systems Director Navigator, as shown in Figure 8-1.

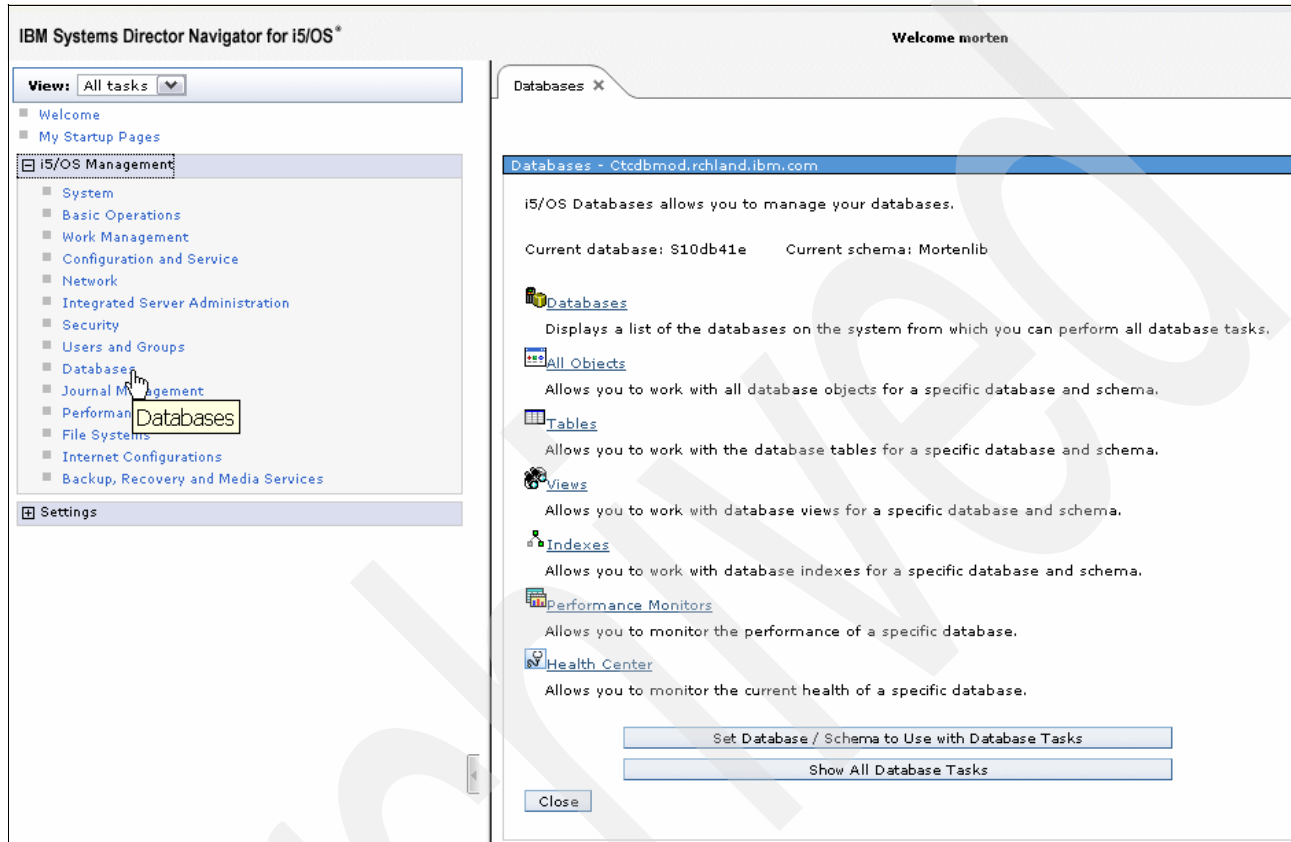


Figure 8-1 Database tasks in System Director Navigator

Figure 8-1 shows you some of the database tasks. You will find additional tasks when you click **Show All Database Tasks**, as shown in Figure 8-2.

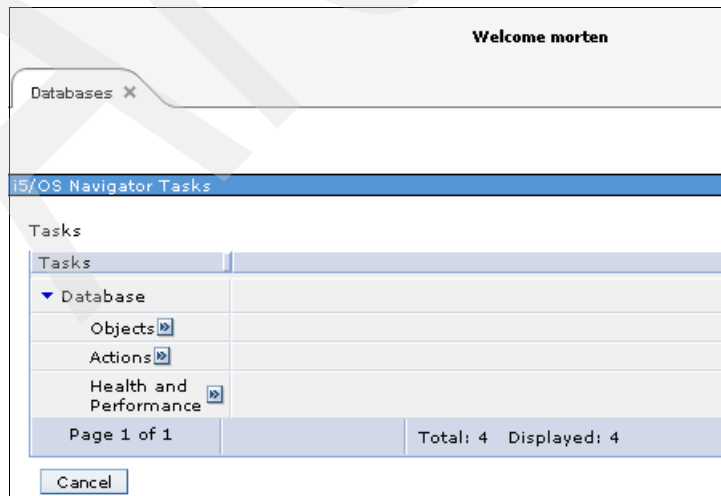


Figure 8-2 Show all database tasks

You can expand the menu items by clicking the arrow to the right of the task name, as shown in Figure 8-3.

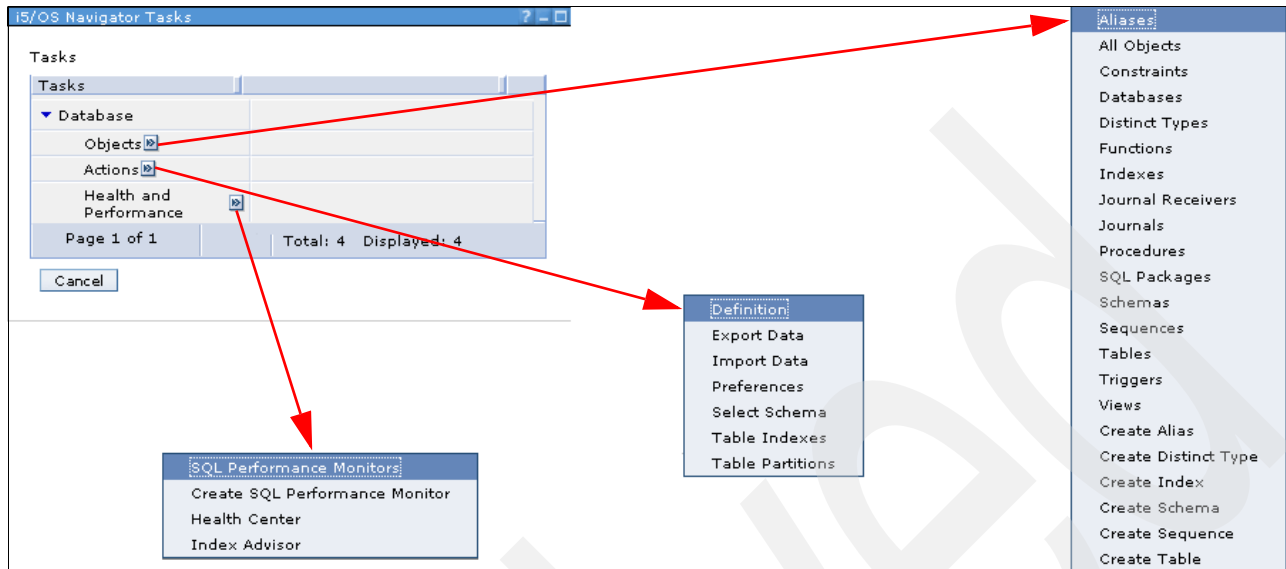


Figure 8-3 Overview of database tasks

In the following sections we will go through some of the different tasks that you can perform with IBM Systems Director Navigator for i.

## 8.2 Select which schemas to work with

When you work with the Systems Director Navigator, you should set the library that you work with. The default library is QUSRSYS, and unless you are using this library, you can select to set the library here in Figure 8-4.

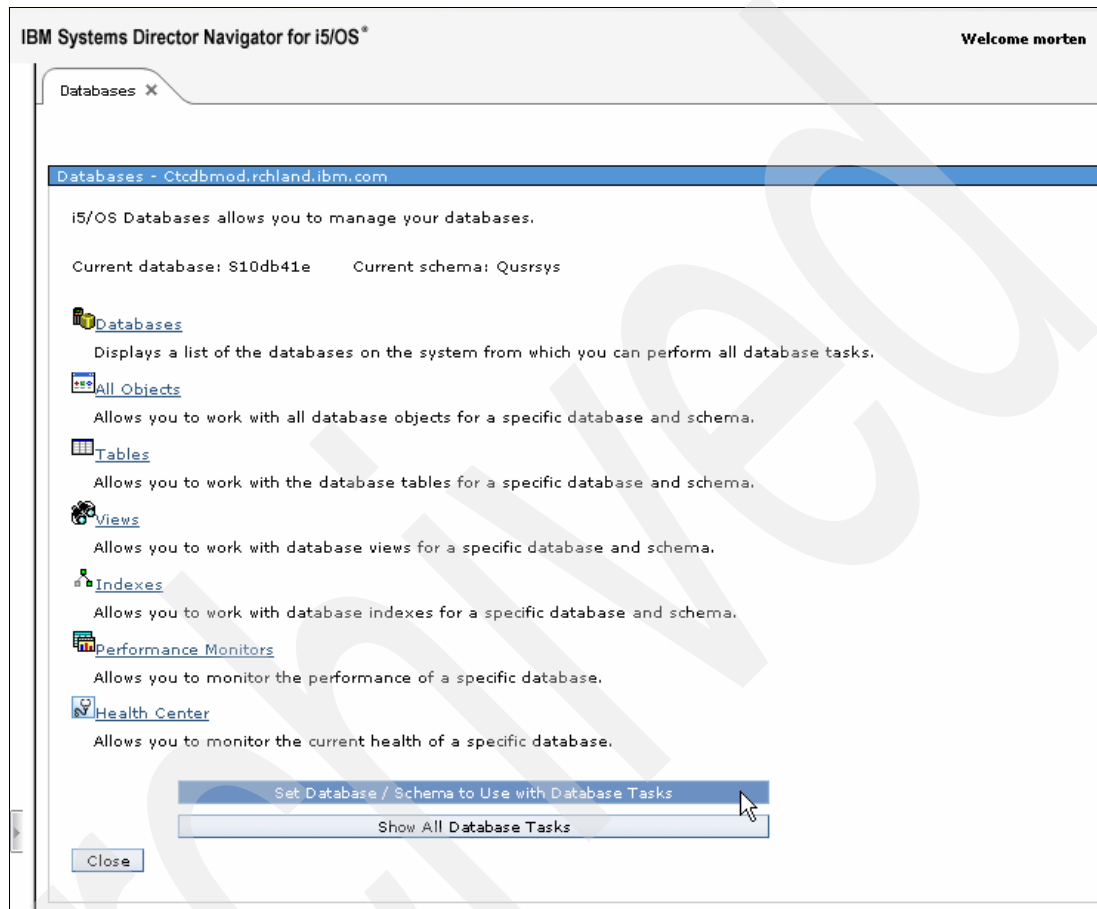


Figure 8-4 Select database/schema to use with database tasks

You can set the current library as shown in Figure 8-5. You can select your database from the system. By default, the local database will be selected, and this will in many cases be the only one on your system. You can take a look at other databases by clicking **Browse** to the right of the Database field. For the schema you can type in the schema that you want to work with or you can simply browse all available schemas for the database by clicking **Browse** to the right of the schema name. After selecting your database and schema click **OK**.

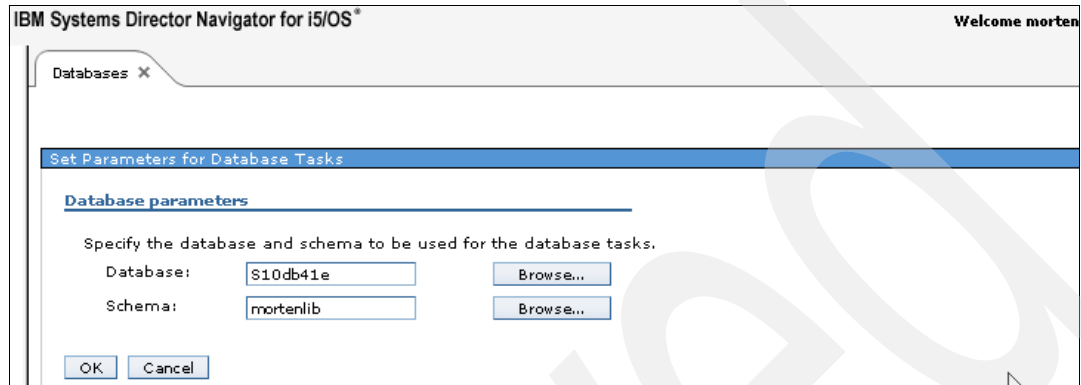


Figure 8-5 Set database/schema to use with database tasks

After setting the library you are ready to work with the different tasks shown in Figure 8-4 on page 104.

A number of functions are under the Databases task, as you can see in Figure 8-6.

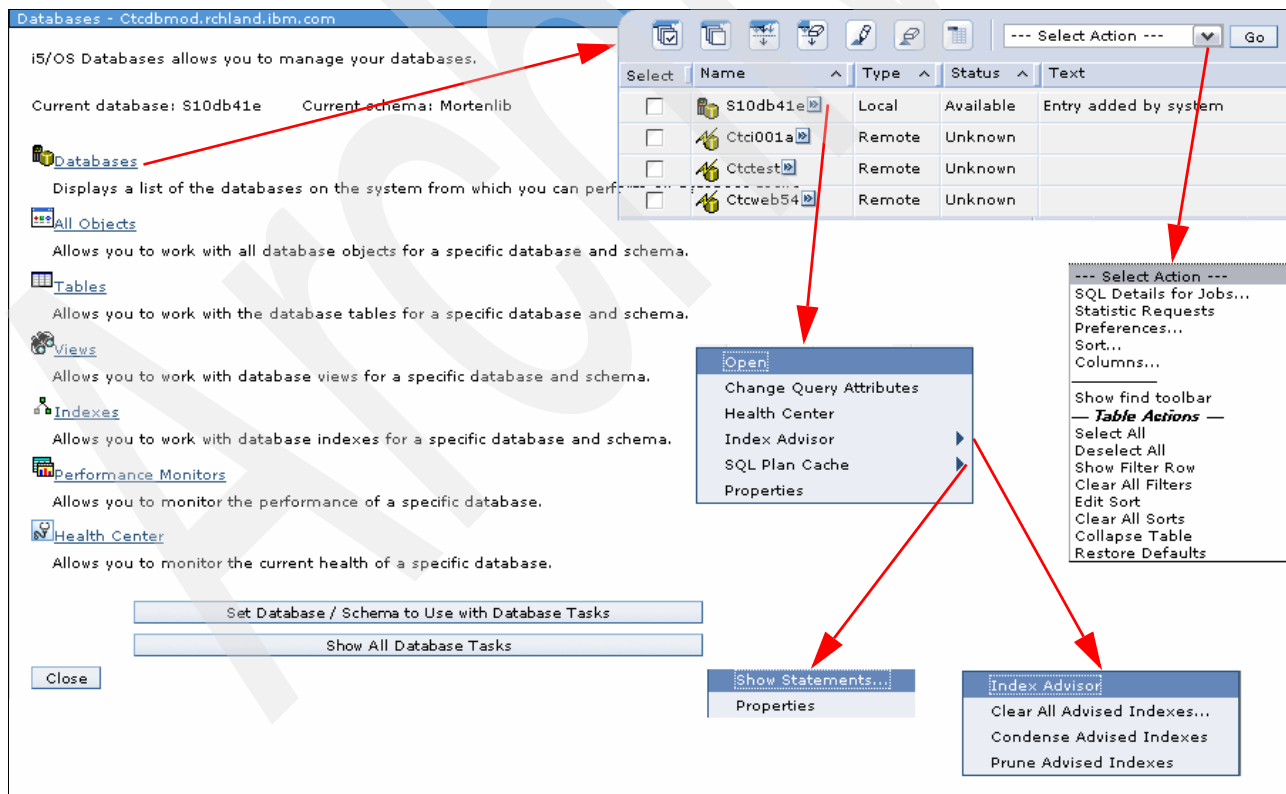


Figure 8-6 Expansion of the Databases menu item

Here you can work with a list of databases on the system. You can have a local database and a number of remote databases. For the local database, you can click **Open** and see further options and functions, as shown in Figure 8-7.

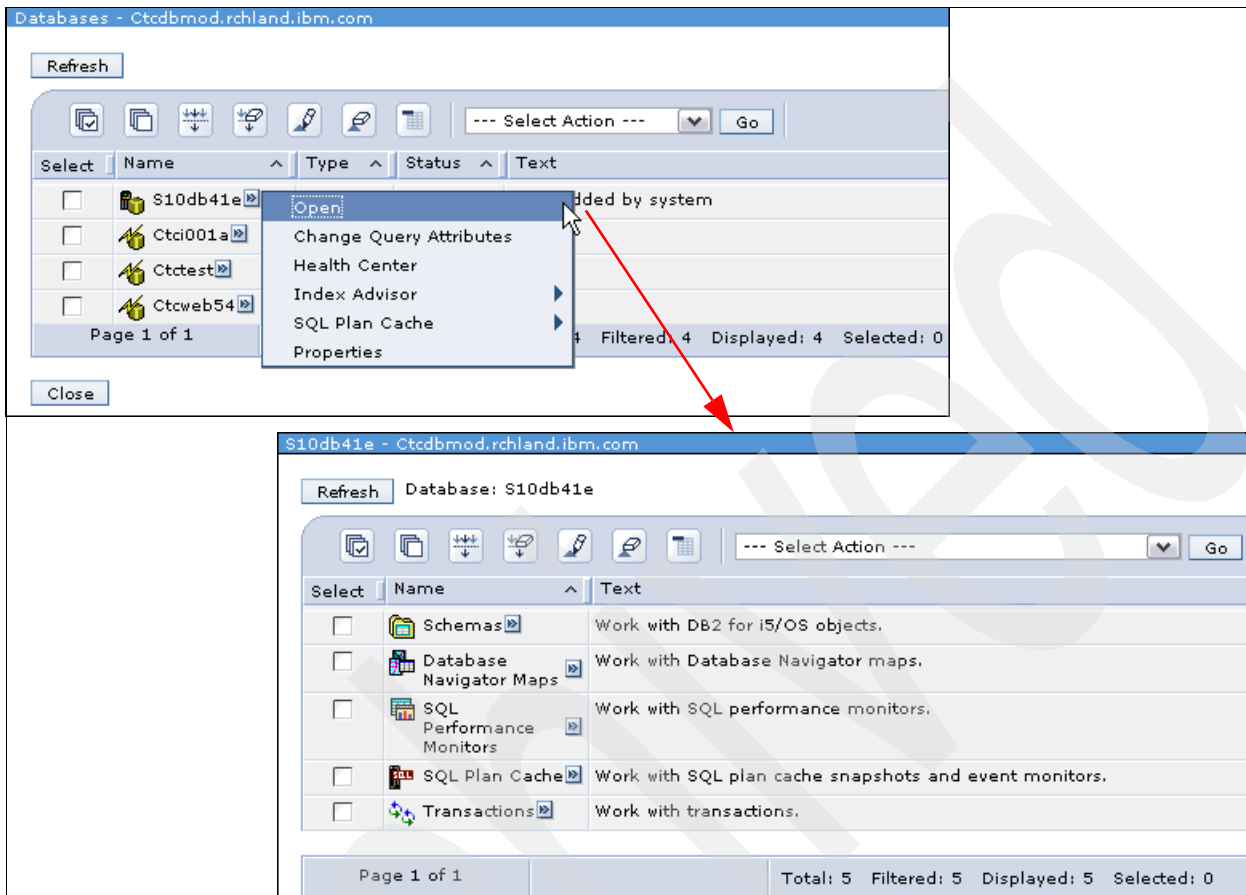


Figure 8-7 Database open gives you some good functionality

Here you can:

- ▶ Work with and drill down to the libraries in your library list. The list is initially your library list, but then can be whatever the user wants to have for his list by managing it with select schemas to display.
- ▶ Show an overview over existing database navigator maps.
- ▶ Show SQL statements from the plan cache.
- ▶ View global and database transactions.

For the rest of the options on the main database page (shown in Figure 8-6 on page 105), a number of options are available on the database system level, such as:

- ▶ Change Query Attributes
- ▶ Health Center
- ▶ Index Advisor
- ▶ SQL Plan Cache
- ▶ Properties

For the menu items All Objects, Tables, Views, and Indexes, you will work with the library that you have chosen by setting the database and schema to work with. Here you have a number of options. If you choose the All Objects menu you will see all the objects in the library and all



the drop-down menus will be different, depending on the object type. When you choose the Tables menu item, you will have a number of different options in your menus, as shown in Figure 8-8.

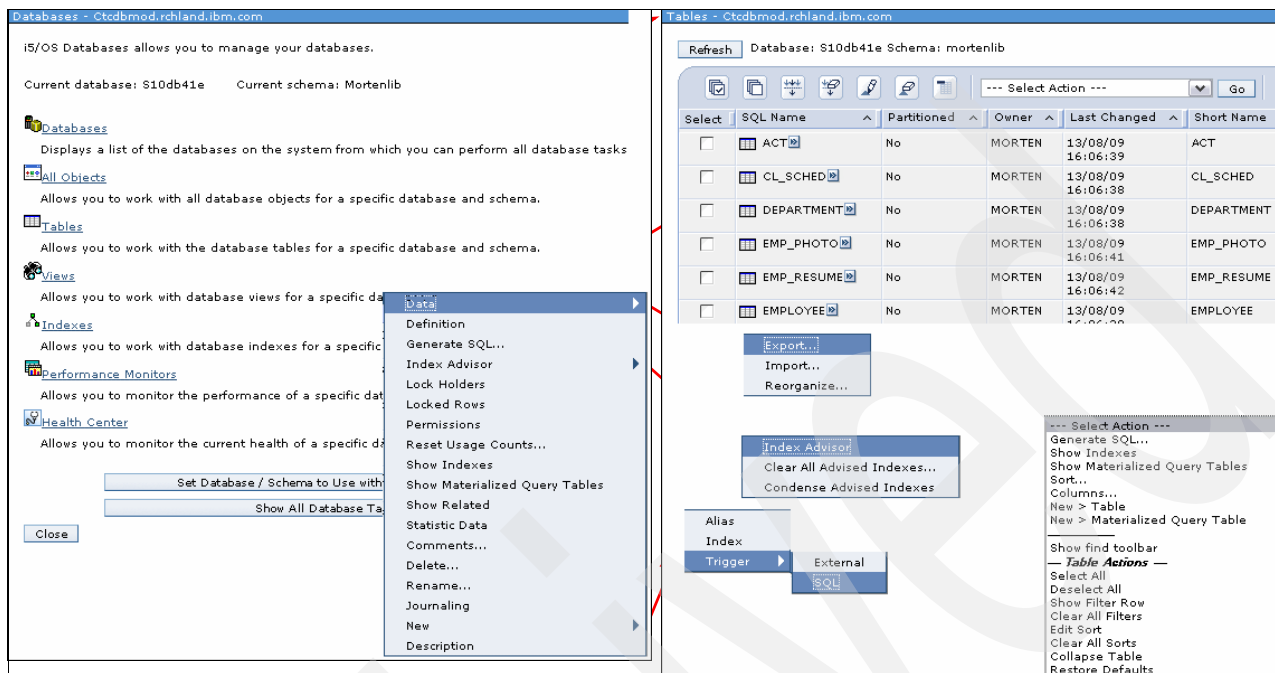


Figure 8-8 Expansion of the Tables menu item

There are a number of options for the objects in the schema, including:

- ▶ Work with all objects in a schema.
- ▶ Work with aliases in a schema.
- ▶ Work with constraints in a schema.
- ▶ Work with distinct types in a schema.
- ▶ Work with functions in a schema.
- ▶ Work with indexes in a schema.
- ▶ Work with journals in a schema.
- ▶ Work with journal receivers in a schema.
- ▶ Work with SQL procedures in a schema.
- ▶ Work with sequences in a schema.
- ▶ Work with SQL packages in a schema.
- ▶ Work with tables in a schema.
- ▶ Work with triggers in a schema.
- ▶ Work with views in a schema.

You can also create different objects in the schema, including:

- ▶ Create table.
- ▶ Create alias.
- ▶ Create distinct type.
- ▶ Create index.
- ▶ Create schema.
- ▶ Create sequence.

You can also maintain data and transfer data. This is done under the following tasks:

- ▶ Reorganize table.
- ▶ Import data into a table from a text file.
- ▶ Export data from a table or view to a text file.

### Table creation

Here we will provide an example of creating a table. You can start from the main database menu and select tables, as shown in Figure 8-9.

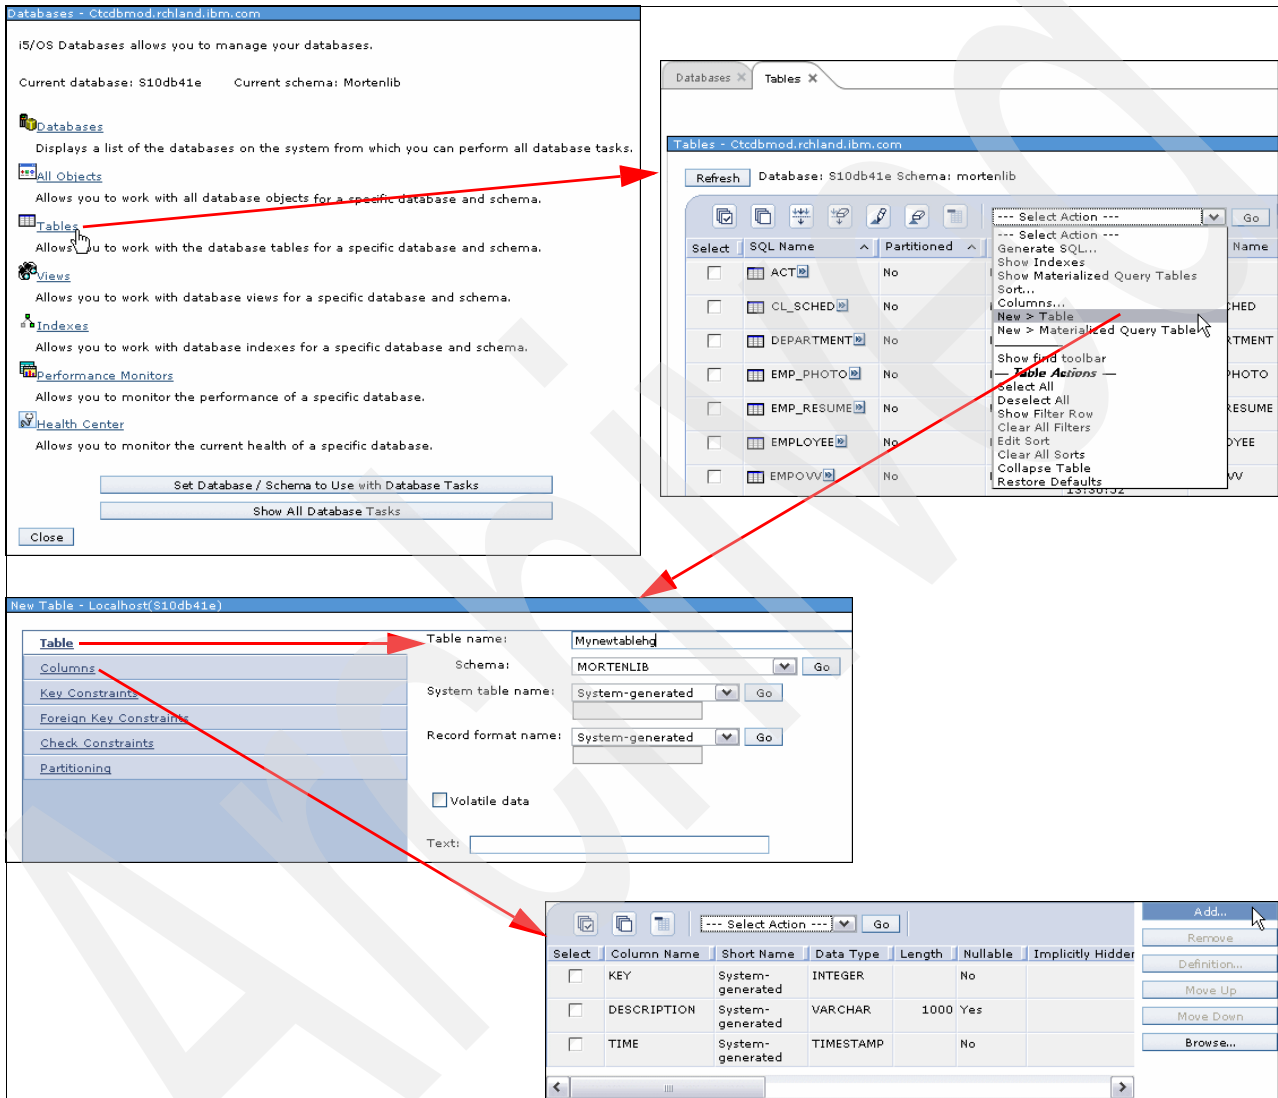


Figure 8-9 Creation of table

Then select **New** **Table** in the pull-down menu and click **Go**. Still in the page shown in Figure 8-9, then select a table name and description on the Table tab. On the Columns tab, enter the columns for the table.

Next go to the Key Constraints tab and enter primary key information if you want that for your table, as shown in Figure 8-10.

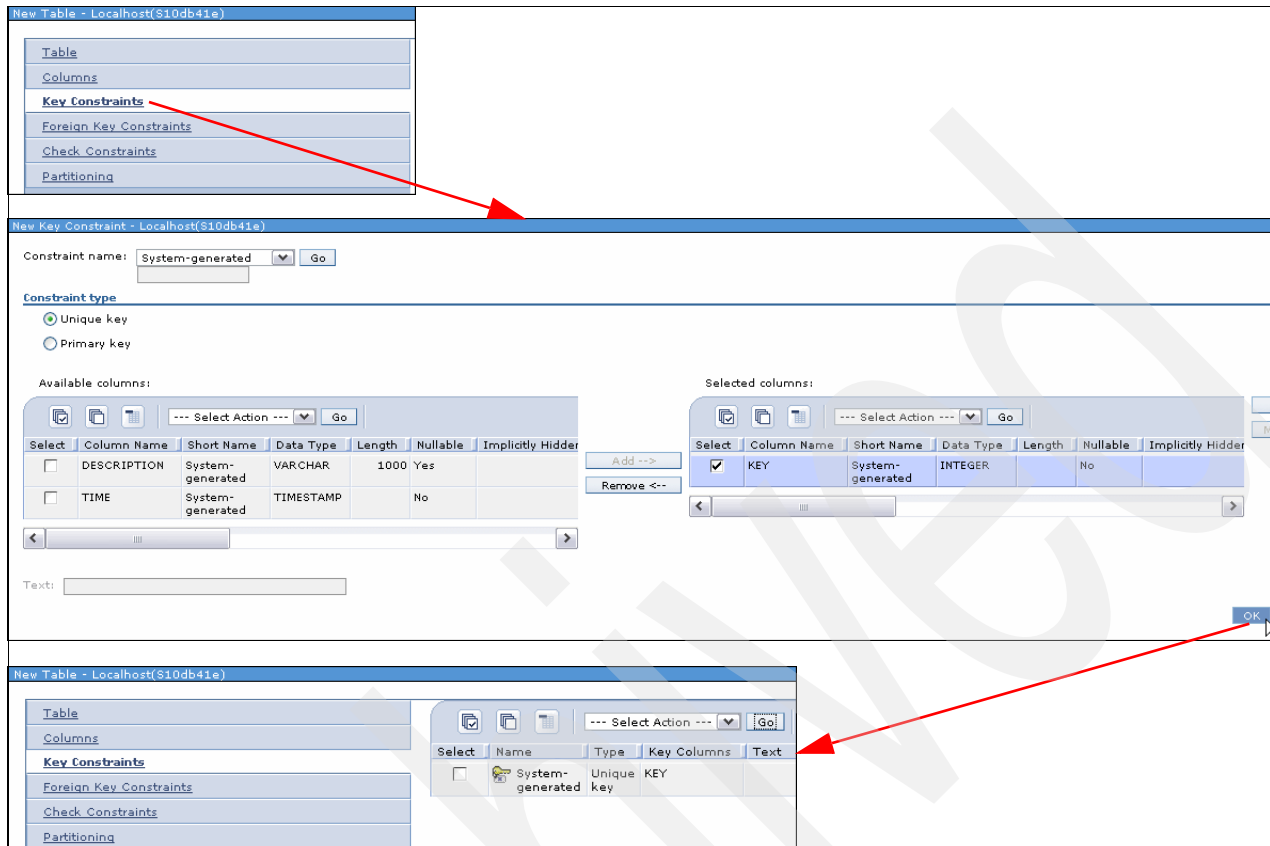


Figure 8-10 Table creation: more options

You can continue adding the needed parameters for your table. When finished, finalize by clicking **OK** on the lower right side of the page and your new table is created.

## 8.3 Generate SQL for objects

One useful feature for working with database objects is being able to construct the SQL that was used to create existing database objects.

You can generate SQL for most database objects regardless of whether they have been created by SQL or by native objects via DDS. Additionally, if you generate SQL for a table that has constraints or triggers associated with it, the SQL is generated for those constraints or triggers as well. You can generate SQL for one object or many objects at a time. To generate SQL for an object, click the small double arrow icon to the right of the object so that you see the pop-up menu. Here you can select **Generate SQL**, as shown in Figure 8-11.

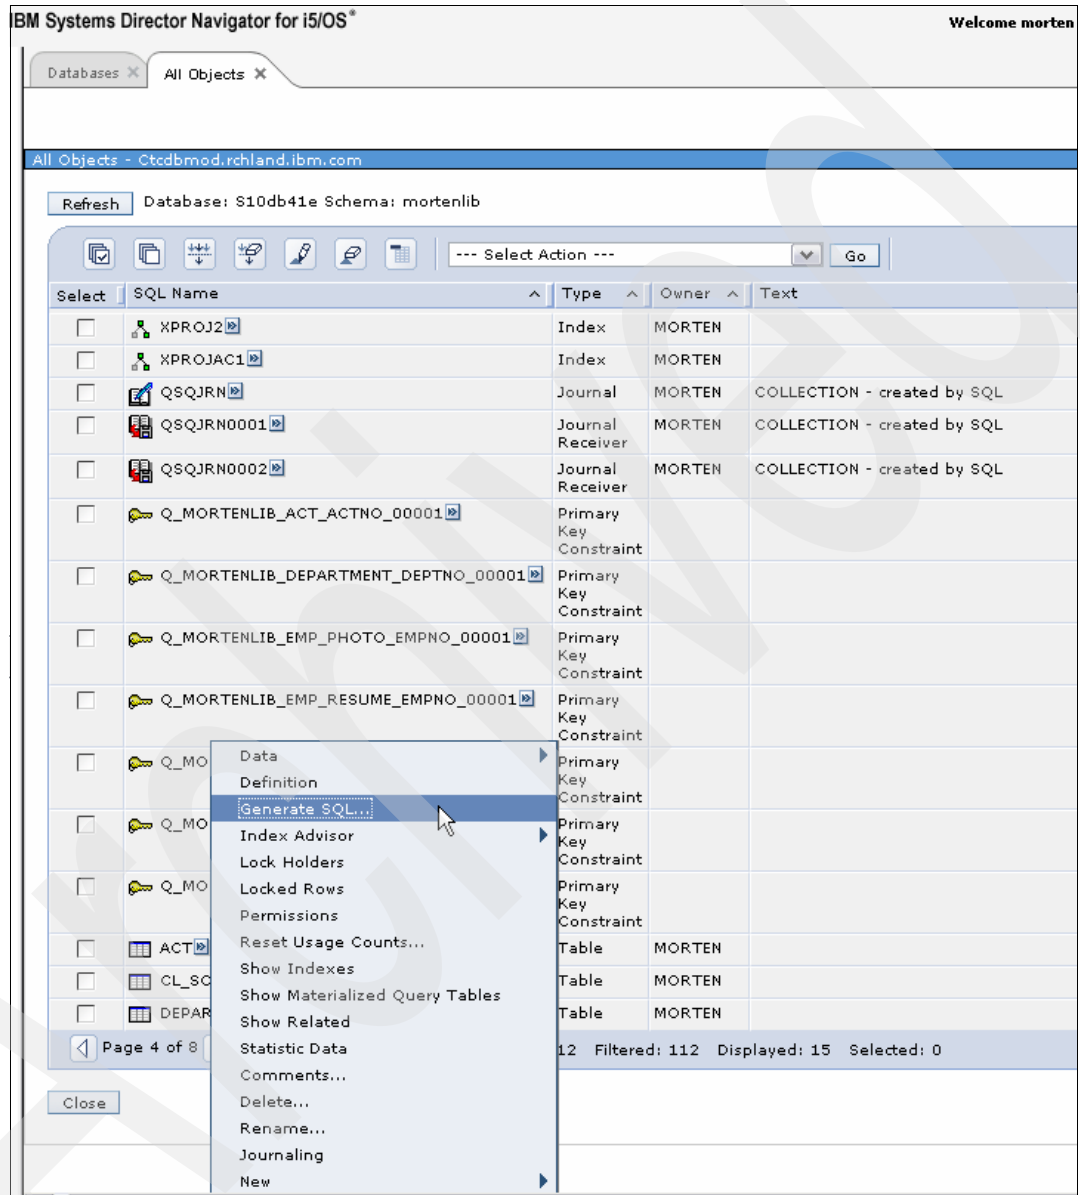


Figure 8-11 Generate SQL for an object

Then a new window will appear with additional instructions about how you want the SQL to be generated. On the first tab called Output you can choose between having the source generated in a file or just showing the SQL source in your browser (Figure 8-12). If you choose Write to file, the SQL source will be written to the specified file. You can also select to add the source to an existing file by checking the Append box.

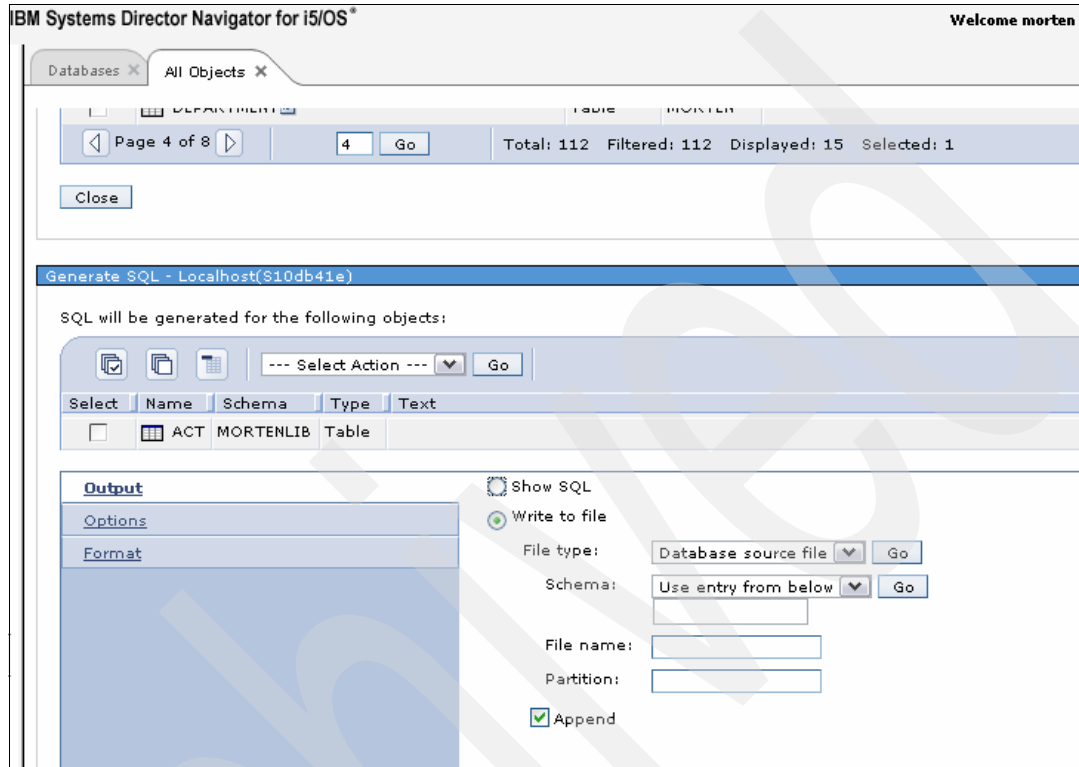


Figure 8-12 Generate SQL: output options

When you choose to just Show SQL in the browser, the output file options cannot be logically edited, as shown in Figure 8-13.

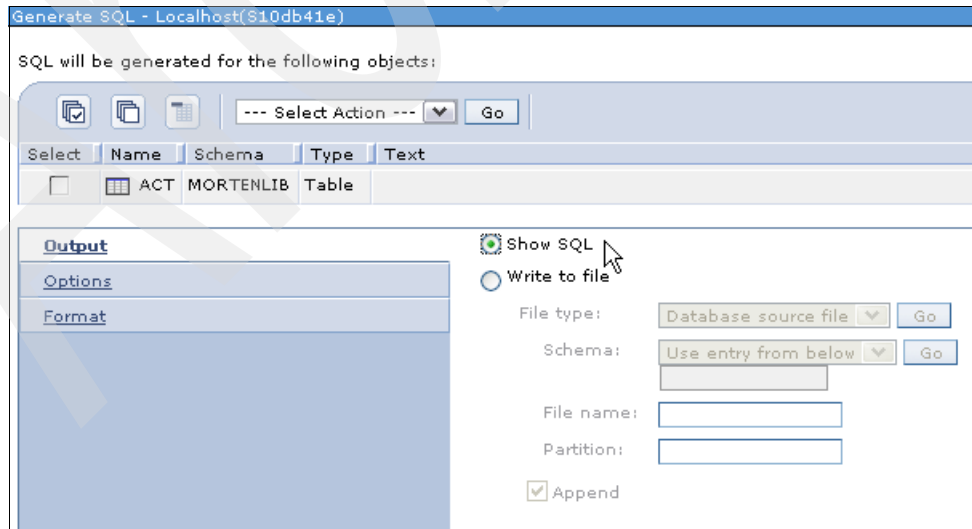


Figure 8-13 Generate SQL: output to browser

Under the Options tab, you can select to have additional details for the SQL that you want to generate, like if you want to include triggers and constraints for a table for instance. You can also specify whether you want to include the SQL statement for dropping the object before it is generated. In Figure 8-14 you can see that we choose to include as much as possible in the generation.

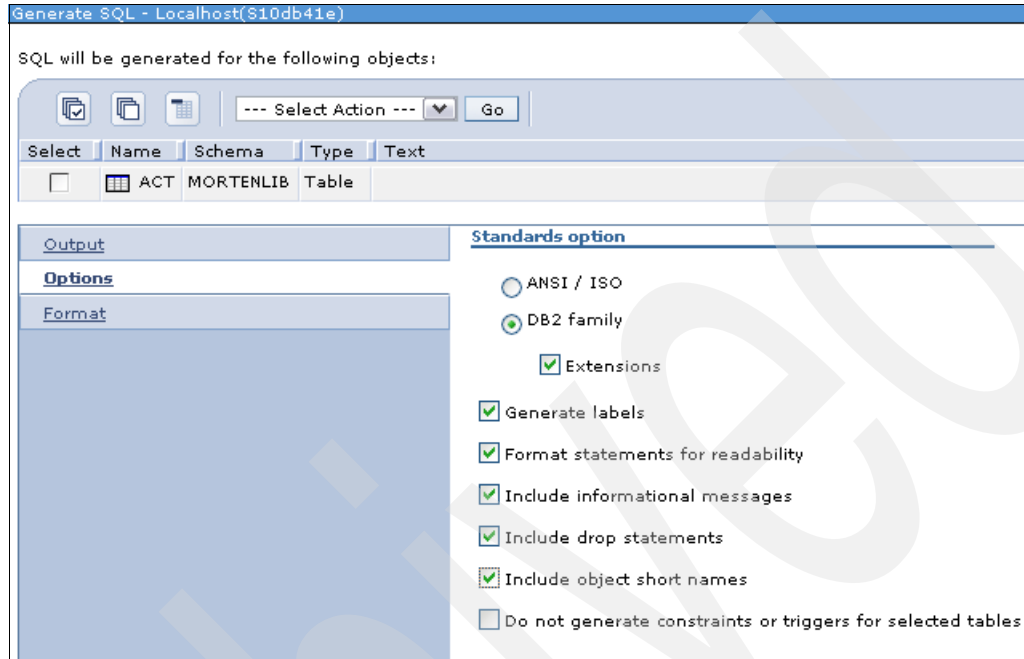


Figure 8-14 Generate SQL options

Under the Format tab, you can select SQL or system naming format, as well as time and date format, as shown in Figure 8-15.

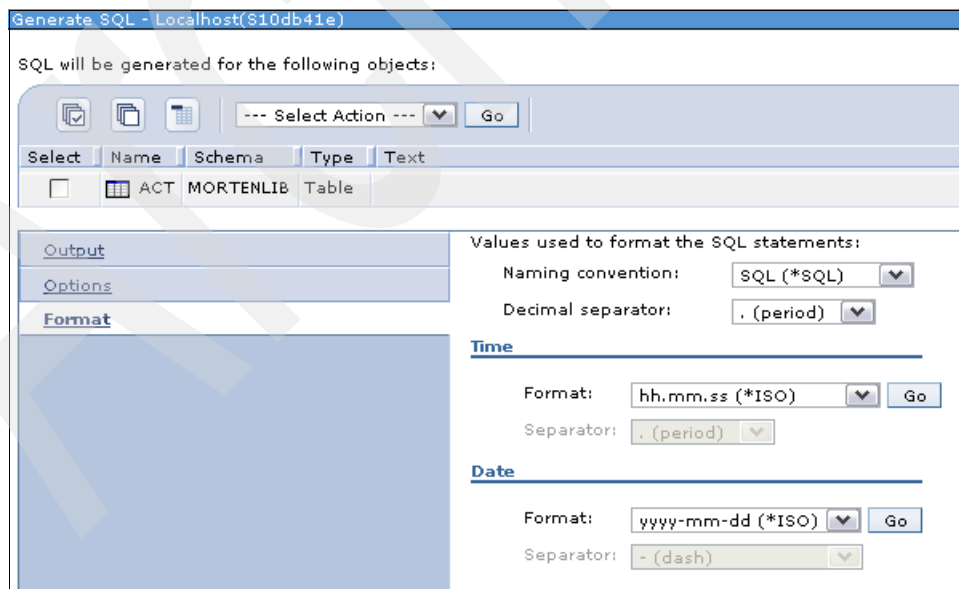
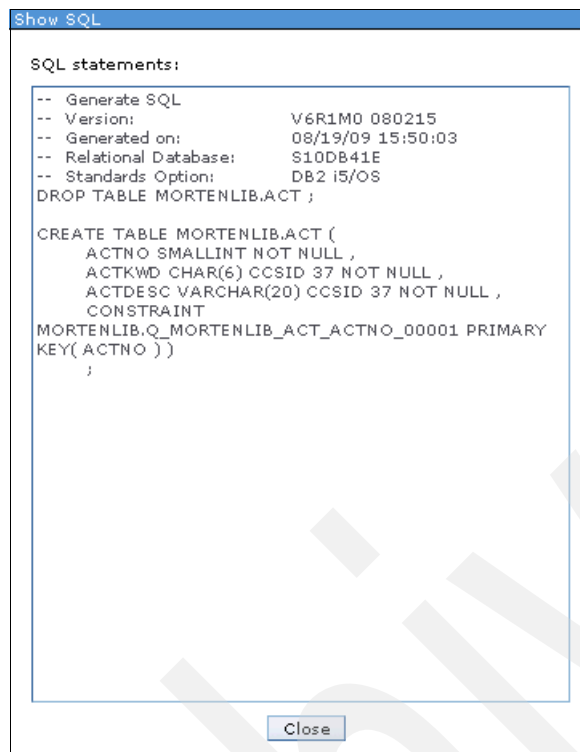


Figure 8-15 Generate SQL: format options

When you want to want to generate the SQL with the options that you have chosen, click **Generate** on the right side of the window. You will then have a text box in your browser with the generated SQL, as shown in Figure 8-16.



The screenshot shows a window titled "Show SQL" with a text area containing the following SQL statements:

```
SQL statements:
-- Generate SQL
-- Version:          V6R1M0 080215
-- Generated on:     08/19/09 15:50:03
-- Relational Database: S10DB41E
-- Standards Option: DB2 I5/OS
DROP TABLE MORTENLIB.ACT ;

CREATE TABLE MORTENLIB.ACT (
  ACTNO SMALLINT NOT NULL ,
  ACTKWD CHAR(6) CCSID 37 NOT NULL ,
  ACTDESC VARCHAR(20) CCSID 37 NOT NULL ,
  CONSTRAINT
MORTENLIB.Q_MORTENLIB_ACT_ACTNO_00001 PRIMARY
KEY( ACTNO ))
;
```

A "Close" button is located at the bottom center of the dialog box.

Figure 8-16 Example of generated SQL

This example only covers a simple table, but you can choose more tables and objects where you want to generate your SQL. When you generate it to the browser, you can copy and paste it from there to any editor or SQL interface depending of your purpose.

## 8.4 Change query attributes

The change query attributes (CHGQRYA) command specifies attributes for database queries and database file keyed access path builds, rebuilds, and maintenance that is run in a job.

You can change the query attributes for the jobs running on your system by selecting **Change Query Attributes** from the arrow button to the right of the local database name, as shown in Figure 8-17.

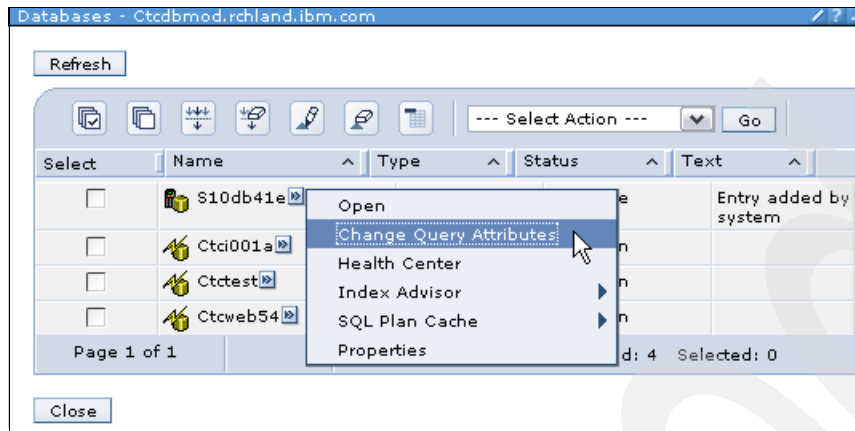


Figure 8-17 Change Query Attributes

This opens the window (Figure 8-18) where you can select which jobs to modify by clicking **Add** in the upper right corner.

You also have different options for the other query attributes, such as:

- ▶ Query processing time limit
- ▶ Query temporary storage limit
- ▶ Parallel processing degree Processing option

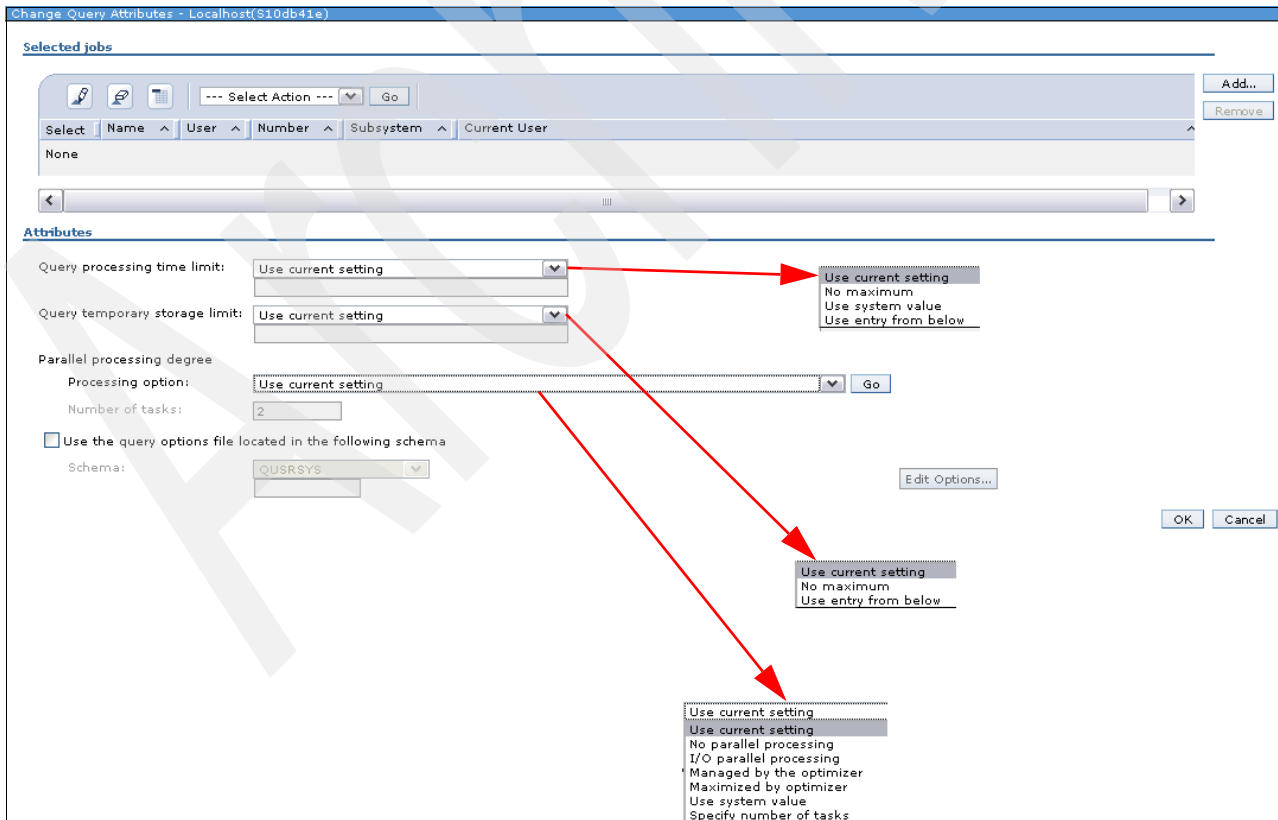


Figure 8-18 Change of query options



In Figure 8-18 on page 114 you can see the different query attributes that you can modify. You can also choose to make use of the query options file.

The query options file (QAQQINI) support provides the ability to dynamically modify or override the environment in which queries are executed. This can be modified through the System Director Navigator interface. The query options file QAQQINI is used to set some attributes used by the database manager.

For each query that is run the query option values are retrieved from the QAQQINI file in the schema specified on the QRYOPTLIB parameter of the CHGQRYA CL command and used to optimize or implement the query. This can also be modified through the System Director Navigator interface.

In the window shown in Figure 8-18 on page 114 you can check the box to activate **Use the query options file located in the following schema**. In this case you can either specify an existing query options file by typing in the library where it exists or you can specify a library where you want it created.

In the following example you can see how we create a new query options file in the schema specified. You will be asked whether you want to create the new query option file, and then you have a number of query options to choose from.

**Note:** Be aware of the required authority for the use of the query options file (Figure 8-19 on page 116). QAQQINI is shipped with a \*PUBLIC \*USE authority. This allows users to view the query options file, but not change it. Because changing the values of the QAQQINI file can affect all queries that are run on the system if the change is done to the file in the QUSRSYS, only the system or database administrator should have \*CHANGE authority to the QAQQINI query options file.

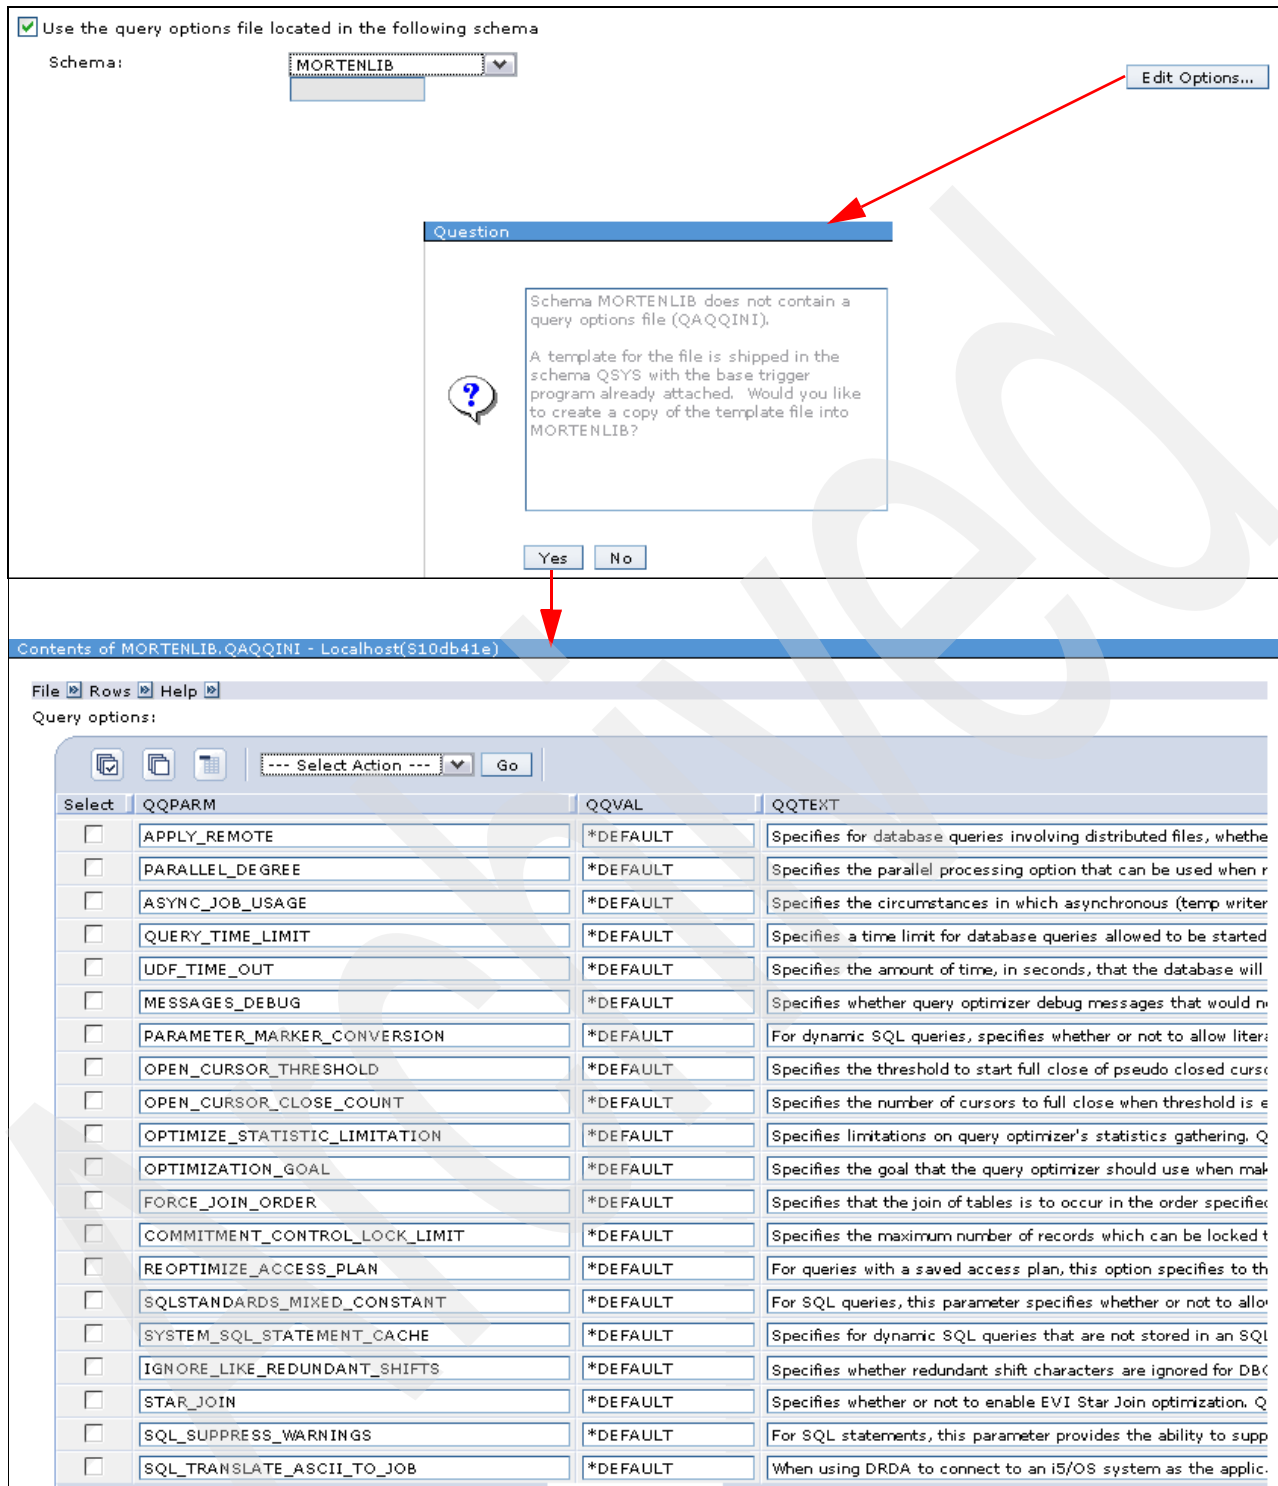


Figure 8-19 Creation of new query options file

The different query options are not covered in this book, but you can find much more information about the query options in the IBM i Information Center at:

<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/index.jsp>

## 8.5 SQL Performance Monitor

The SQL Performance Monitor is one of the best ways to work with and analyze the queries running on a System i. We look at how you can access the SQL performance monitor using the System Director Navigator.

### 8.5.1 Create a new SQL Performance Monitor

Figure 8-20 shows a quick overview of how to start a database monitor collection. The steps are:

1. In the navigation menu choose the **Databases** task, then choose **Show All Database Tasks** on the main databases page.

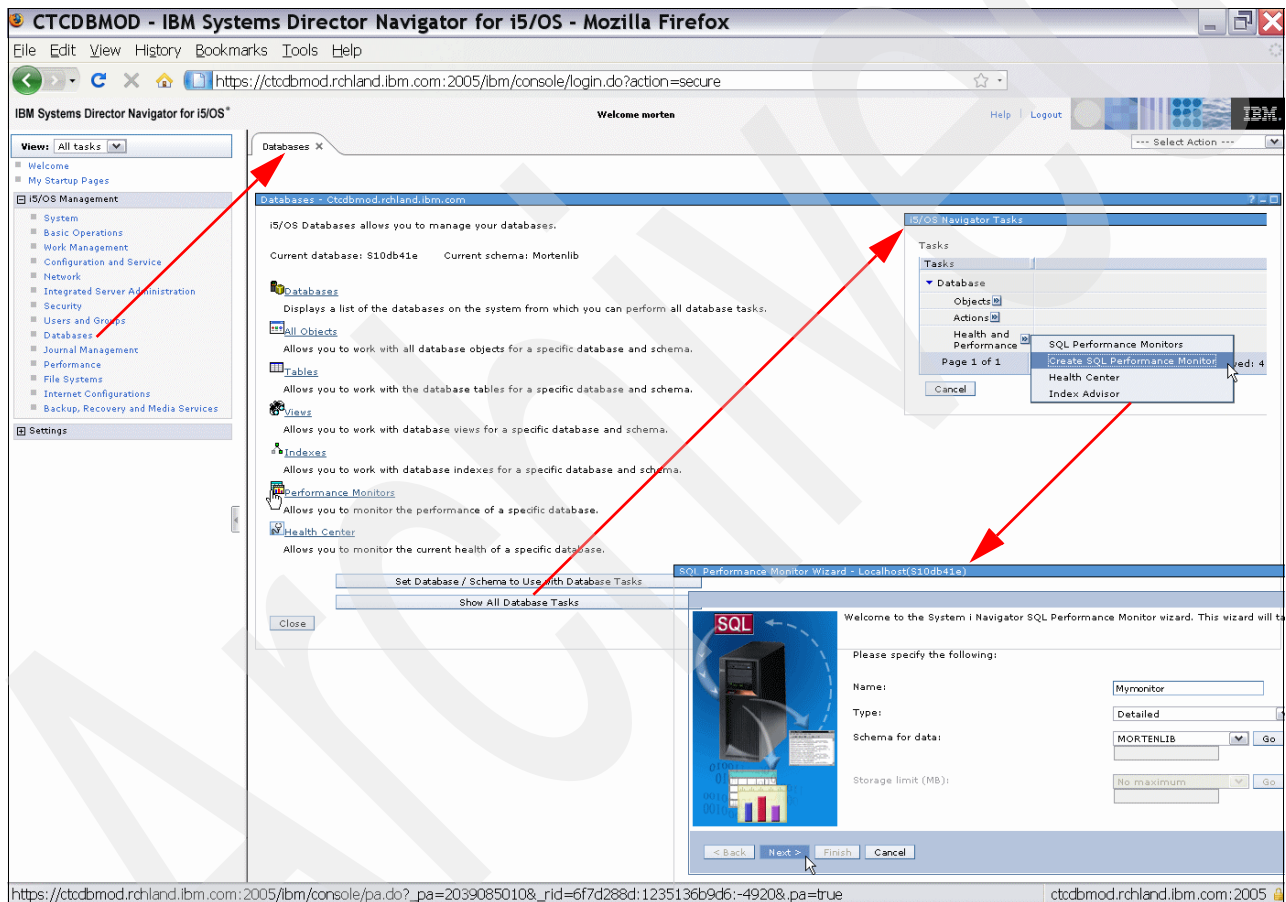


Figure 8-20 Start a database monitor collection

2. This brings you to the database tasks. Here you click the arrow button to the right of the Health and Performance menu item.
3. Select name, library, and description for your collection, then click **Next**.

- On the page shown, you can do your filtering to reduce your collection. You have many choices for the filtering, as shown in Figure 8-21.

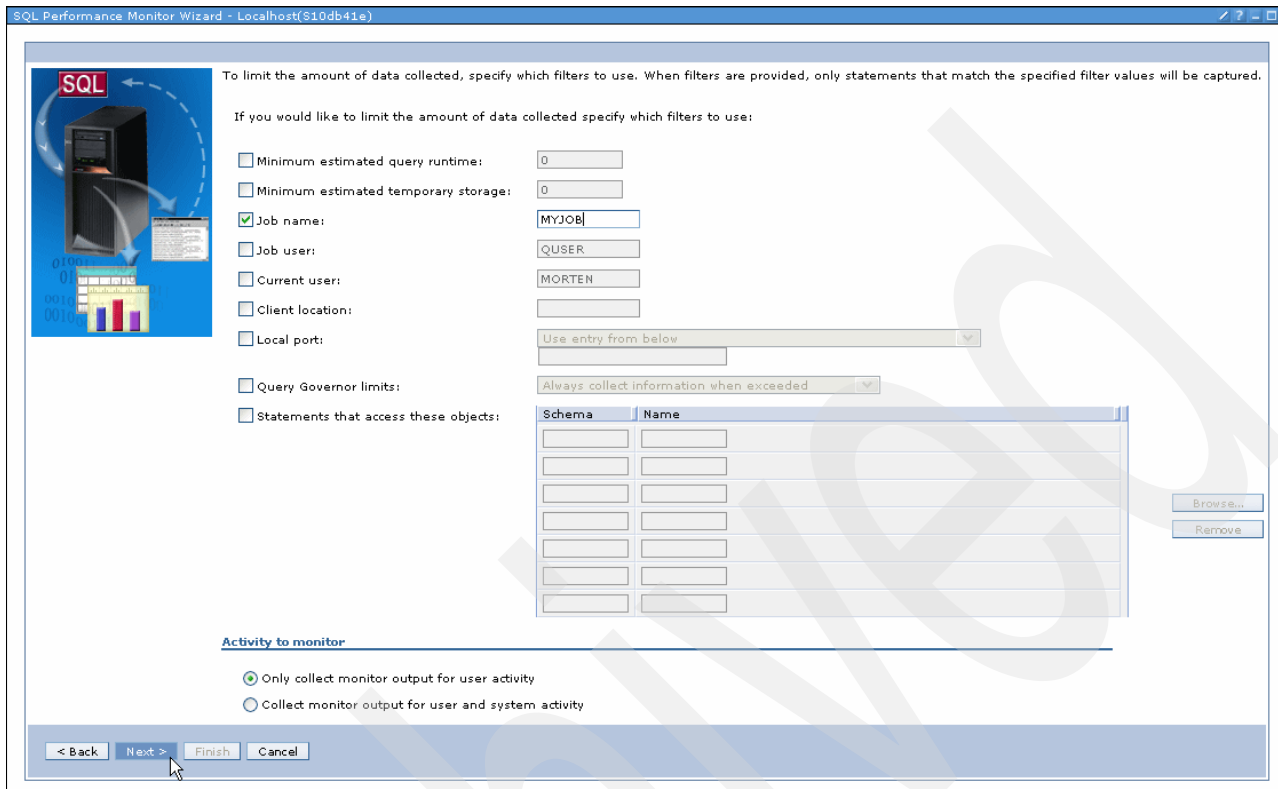


Figure 8-21 Filter the database monitor collection

- After filtering click **Next** and you will see a confirmation that the SQL performance monitor has started successfully (Figure 8-22).

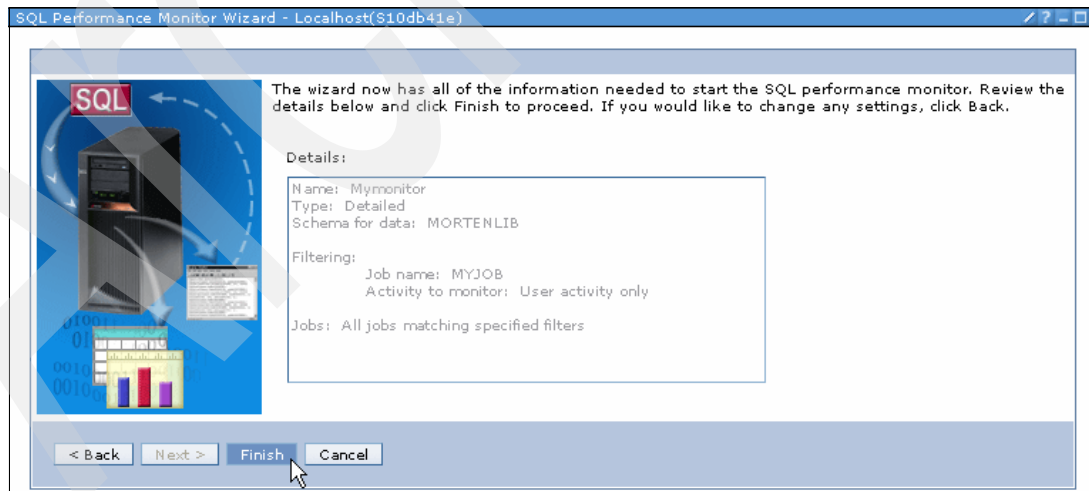


Figure 8-22 Confirm the database monitor collection

- Now you can let the monitor run and collect for the duration that you want. When you want to end the monitor, click the arrow to the right of your monitor from the SQL Performance Monitors menu, as shown in Figure 8-23.

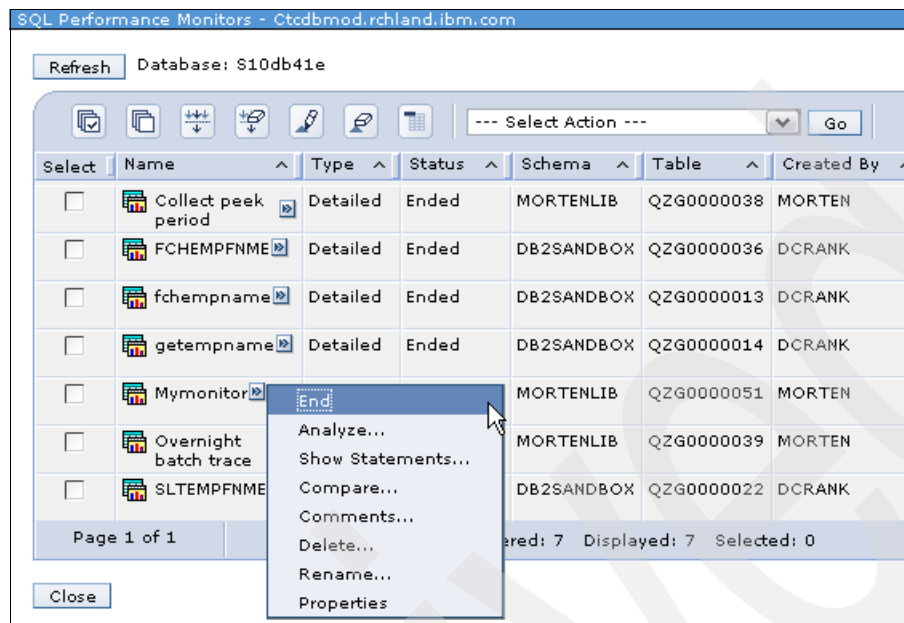


Figure 8-23 End database monitor collection

Now your database monitor collection is ready for analysis.

## 8.5.2 Work with SQL performance monitors

You can work with and analyze SQL performance monitors. The monitors are traces collected from a specific job or a number of jobs all depending on options that you selected when starting the SQL performance monitor.

## 8.5.3 Detailed database monitor

The database monitor collects information about a query in real time and stores this information in an output table. This information can help you determine whether your system and your queries are performing as they should or whether they need fine tuning. Database monitors can generate significant CPU and disk storage overhead when in use. The overhead can be up to 20–30% on additional query runtime, depending on available resources in the partition that you are running.

You can gather performance information for a specific query, for every query on the system, or for a group of queries on the system. When a job is monitored by multiple monitors, each monitor is logging rows to a different output table. You can identify rows in the output database table by each row's unique identification number.

When you start a database monitor, the monitor is automatically registered with the Systems Director Navigator and appears in the Systems Director Navigator SQL Performance Monitors list.

## 8.5.4 Example of database monitor analysis

In this section we look at a sample database monitor via the Systems Director Navigator interface. Let us assume that you have collected a database monitor for a night job, as shown in Figure 8-24. We click the arrow button to show the pop-up menu for the collection. Then we select **Analyze**.

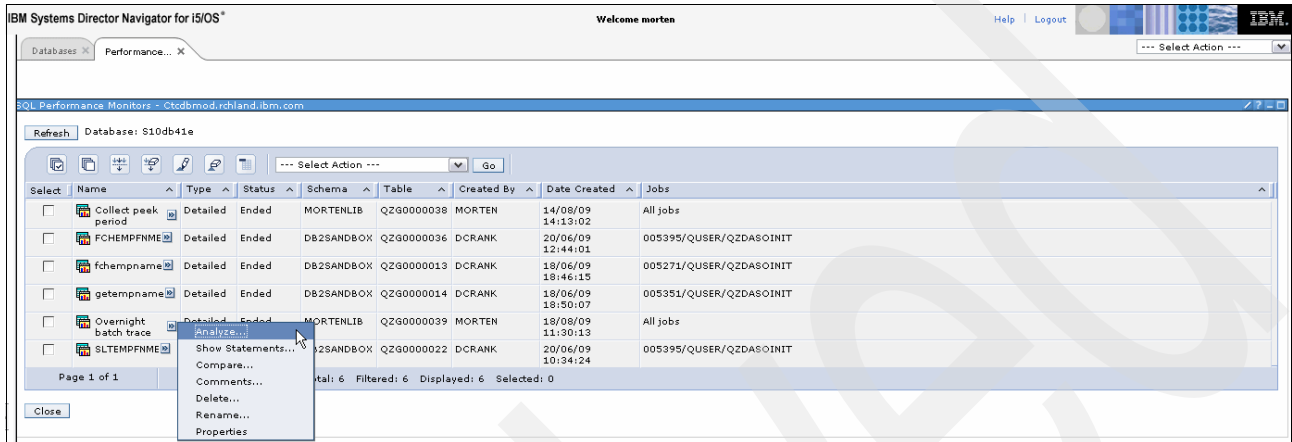


Figure 8-24 Start analyzing a database monitor

This provide us with a great overview for the database monitor collection, as shown in Figure 8-25.

Select		Value	Summary Available	Statements Available
<input type="radio"/>	18/08/09 11:30:13 to 20/08/09 13:21:11			
<input type="radio"/>	Overview			
<input type="radio"/>	SQL Statements	10097	✓	✓
<input type="radio"/>	Users	14	✓	
<input type="radio"/>	Jobs	87	✓	
<input type="radio"/>	Threads	87		
<input type="radio"/>	Average Table Rows	4195.296		
<input type="radio"/>	Average Rows Returned	30.812		
<input type="radio"/>	Average Runtime	0.008797		
<input type="radio"/>	Average Parallel Degree Used	1		
<input type="radio"/>	Maximum Parallel Degree	1		
<input type="radio"/>	SQE	2695	✓	✓
<input type="radio"/>	CQE	606	✓	✓
<input type="radio"/>	System Naming	5789	✓	✓
<input type="radio"/>	SQL Naming	2934	✓	✓
<input type="radio"/>	Unique Open Statements	305	✓	✓
<input type="radio"/>	Full Opens	728	✓	✓
<input type="radio"/>	Pseudo Opens	2256	✓	✓
<input type="radio"/>	Table Scans	783	✓	✓
<input type="radio"/>	Average MQTs Used	0		

Figure 8-25 Database monitor: analyze overview 1

We can now start analyzing the database monitor collection by raising questions about what we want to investigate further. From Figure 8-25 we can find the following information:

1. The number of SQL statements gives you a good overview of the collection.
2. The number of jobs and users having used SQL on the system during the collection is always good to know and to understand whether that was what you expected—maybe you see many more jobs or fewer jobs than expected, so you should investigate that.
3. The average runtime is only to give you an idea about what is going on, because some queries may take microseconds and others take seconds to run, so an average is only of interest if you are running exactly the same scenario every time that you collect data.

4. If you have DB2 Symmetric Multiprocessing (SMP) you will be able to see the average and maximum parallel degrees used to extract the data.
5. The use of SQL Query Engine (SQE) and Classic Query Engine (CQE) is important for performance, as the SQE is newer technology and mostly performs better than the CQE. You should check whether the CQEs are other than SQL INSERT's and find out why the CQE is used, as you may be able to avoid using the CQE.
6. You cannot avoid full opens. The two first times that you are opening the same cursor in a job, it will be a full open, and any subsequent opens should be pseudo opens, if the way that the applications accessing the database are well programmed. If your number of full opens is high, you should investigate why the cursors are closed. Is it due to program settings or due to the connection settings? A rule of thumb is to look at the number of full opens compared with the number of unique opens. If the number of full opens is twice the number of unique opens then you are on right track and you can probably not make this much better in the application. If the number is high then it can be caused by many connects and disconnects to the database, which is not good, as it will result in many full opens when the cursors must be recreated.
7. The number of table scans can indicate areas where you may be able to tune the system. If table scans are due to the need of all or most entries in a table then it is most likely the fastest way to access the data, but when fewer rows are extracted, the use of indexes would be of interest.
8. If you have created materialized query tables (MQTs) and expect them to be used by the database engine, you can monitor whether any are used. If not, then check why the query optimizer chose not to use your MQTs. This requires a more in-depth analysis than this book covers.
9. Here you have the number of queries using system naming and SQL naming. Here you can check whether the queries using system naming are qualified with a library name, or else the system must make more lookups to find the name of the table referenced in the query.

Scrolling down on the Web page, you will see more information for the overview, as shown in Figure 8-26.

<input type="radio"/>	Average Indexes Used	1.015	←	✓	10
<input type="radio"/>	Full Indexes Created	113	←	✓	11
<input type="radio"/>	Sparse Indexes Created	0	←		12
<input type="radio"/>	Index From Index Created	6	←	✓	13
<input type="radio"/>	Index Creates Advised	286	←	✓	14
<input type="radio"/>	Advised Statistics	277	←	✓	15
<input type="radio"/>	Temporary Tables	258	←	✓	16
<input type="radio"/>	Sorts	226	←	✓	17
<input type="radio"/>	Access Plans Rebuilt	458	←	✓	18
<input type="radio"/>	Sort Sequence	0			
<input type="radio"/>	Call Statements	474	✓	✓	
<input type="radio"/>	Error	130	←	✓	19

Figure 8-26 Database Monitor analyze overview 2



10. In the average number of indexes used, you get an idea of whether indexes are used, but again this average does not tell you much as you do not know how many tables may have been joined.
11. The full indexes created is definitely something you should look deeper into. Often you will see those indexes in the index recommendations too.
12. When a sparse index is created, it is an index that covers a part of the rows in the table. Here in OS/400 Version 6.1 you can also create those sparse indexes as permanent indexes.
13. Here we have temporary indexes created from existing indexes. Here you can consider creating the indexes as permanent, but first of all it may be a good idea to understand why the optimizer choose to create a temporary index and not to use an existing one directly.
14. The indexes advised are very important to take a look at and understand whether the indexes are going to help a lot. In our example here we should expect that the table scans, the temporary indexes created, can be avoided by a number of those indexes recommended. You also must understand why the system suggests indexes. It may be due to lacking statistics about the metadata or it may be due to lacking indexes to use for extracting the data from the database.
15. The advised statistics can also provide a hint about missing indexes. The good thing about the missing statistic is the ability of the system to collect the data automatically.
16. The creation of the temporary tables may be avoided by having the right indexes in place, but here you must go deeper into the reason.
17. Sorts may be avoided by indexes, but it may be sorts of a few number of rows, so it all depends on the query.
18. The access plans rebuild is important. If it is the first time that a plan is created then the count will be one. The access plan rebuilt count is important. If this is the first time that a plan is created, then the count is one. The reason for an access plan rebuild is that the old access plan is not useful anymore. This can be caused by environment changes, like number of CPUs, memory size for the pool where the job is running, or due to other parameters, such as changes in number of records as the database grows. This can often be controlled by you, whereas other parameters, such as changes in number of records, cannot be controlled, as it is the way the database grows.
19. You should always investigate the errors and understand whether they are okay and handled by the applications.

You can also expand the information **How much work was requested**, as in Figure 8-27. Most of the information is in the overview.

<input type="radio"/>	How much work was requested?		
<input type="radio"/>	SQL Statements	10097	✓
<input type="radio"/>	Users	14	✓
<input type="radio"/>	Jobs	87	✓
<input type="radio"/>	Threads	87	
<input type="radio"/>	Average Table Rows	4195,296	
<input type="radio"/>	Average Rows Returned	30,812	

Figure 8-27 Expand How much work was requested?

Information about what the users and application provide to the optimizer about how to run the queries can be found by expanding **What options were provided to the optimizer**, as shown in Figure 8-28.

<input type="radio"/>	What options were provided to the optimizer?			
<input type="radio"/>	Average Parallel Degree	1		
<input type="radio"/>	Maximum Parallel Degree	1		
<input type="radio"/>	Query Options Specified	13		
<input type="radio"/>	Distinct Query Option Files	1	✓	
<input type="radio"/>	Governor Enabled	4	✓	✓
<input type="radio"/>	Allow Copy Data *YES	21	✓	✓
<input type="radio"/>	Allow Copy Data *OPTIMIZE	2539	✓	✓
<input type="radio"/>	Allow Copy Data *NO	0		
<input type="radio"/>	First IO	1783	✓	✓
<input type="radio"/>	All IO	777	✓	✓
<input type="radio"/>	Force Join Order	0		
<input type="radio"/>	Parameter Marker Conversion	2450	✓	✓
<input type="radio"/>	NLSS Specified	0		
<input type="radio"/>	Unicode Normalization	0		
<input type="radio"/>	Distributed	0		
<input type="radio"/>	Apply Query Attribute Remote	394	✓	✓
<input type="radio"/>	Blocking Enabled *READ	1	✓	✓
<input type="radio"/>	Blocking Enabled *ALLREAD	2600	✓	✓
<input type="radio"/>	Blocking Not Enabled *NONE	0		
<input type="radio"/>	Delay Prepare	5908	✓	✓
<input type="radio"/>	No Delay Prepare	2815	✓	✓
<input type="radio"/>	Close Cursor *ENDACT/*ENDJOB	74	✓	✓
<input type="radio"/>	Close Cursor *ENDMOD/*ENDPGM	1689	✓	✓
<input type="radio"/>	System Naming	5789	✓	✓
<input type="radio"/>	SQL Naming	2934	✓	✓

Figure 8-28 Expand the options that were provided to the optimizer

One of the interesting parameters to investigate is the query optimization goal. It can either be First IO or All IO. You should make sure that you understand why you are using one or the other. For interactive environments the first IO is often recommended, as it will return the first data to the application as fast as possible.

You should also investigate whether you really want to close your SQL cursors when you exit the programs and modules, and not when you are ending the jobs. Closing your SQL cursors too frequently can result in too much time spent rebuilding the necessary cursors.

The next line provides information about how the query optimizer chose to fulfill the SQL requests and what implementation was done. You can see the first part of this in Figure 8-29. Here we also have the indexes advised and information about temporary objects created.

<input type="radio"/>	What implementations did the optimizer use?			
<input type="radio"/>	Average Runtime	0.008797		
<input type="radio"/>	Runtime	88.833283		
<input type="radio"/>	Full Indexes Created	113	✓	✓
<input type="radio"/>	Sparse Indexes Created	0		
<input type="radio"/>	Index From Index Created	6	✓	✓
<input type="radio"/>	Index Creates Advised	286	✓	✓
<input type="radio"/>	Advised Statistics	277	✓	✓
<input type="radio"/>	Average MQTs Used	0		
<input type="radio"/>	Average Indexes Used	1.015	✓	✓
<input type="radio"/>	Temporary Tables	258	✓	✓
<input type="radio"/>	Sorts	226	✓	✓
<input type="radio"/>	Bitmap Creates	85	✓	✓
<input type="radio"/>	Bitmap Merges	9	✓	✓
<input type="radio"/>	Skip Sequential Scans	74	✓	✓
<input type="radio"/>	Table Scans	783	✓	✓
<input type="radio"/>	Nested Loop Join	121	✓	✓
<input type="radio"/>	Hash Join	19	✓	✓
<input type="radio"/>	Index Group By	254	✓	✓
<input type="radio"/>	Hash Group By	33	✓	✓
<input type="radio"/>	Index Order By	86	✓	✓
<input type="radio"/>	Sort Order By	225	✓	✓

Figure 8-29 First part of the information about what implementation the optimizer used

The second part of all the information about the used implementation can be found in Figure 8-30.

<input type="radio"/>	Average Parallel Degree Used	1		
<input type="radio"/>	Maximum Parallel Degree	1		
<input type="radio"/>	Parallel Table Scan	0		
<input type="radio"/>	Parallel Index Scan	0		
<input type="radio"/>	Parallel Hash Join	0		
<input type="radio"/>	Parallel Hash Group By	0		
<input type="radio"/>	Parallel Bitmap Create	0		
<input type="radio"/>	Parallel Bitmap Merge	0		
<input type="radio"/>	Parallel Index Create	0		
<input type="radio"/>	Parallel Prefetch	4110	✓	✓
<input type="radio"/>	Parallel Preload	65	✓	✓
<input type="radio"/>	Maximum Estimated Rows	131483		
<input type="radio"/>	Average Estimated Rows	1431.626		
<input type="radio"/>	Full Opens	728	✓	✓
<input type="radio"/>	Pseudo Opens	2256	✓	✓
<input type="radio"/>	Access Plans Rebuilt	458	✓	✓
<input type="radio"/>	Rebuilds Deferred	0		
<input type="radio"/>	Optimizer Time Outs	0		
<input type="radio"/>	Governor Limits Exceeded	4	✓	✓
<input type="radio"/>	Lock Escalations	0		
<input type="radio"/>	Reusable	2844	✓	✓
<input type="radio"/>	Nonreusable	145	✓	✓
<input type="radio"/>	SQE	2695	✓	✓
<input type="radio"/>	CQE	606	✓	✓

Figure 8-30 Second part of the information about what implementation the optimizer used

The type of SQL statements, grouping, and joining information can be found in Figure 8-31, Figure 8-32, and Figure 8-33 on page 128.

<input type="radio"/>	What types of SQL statements were requested?			
<input type="radio"/>	Row Selection	210	✓	✓
<input type="radio"/>	Subquery	0		
<input type="radio"/>	Group By	287	✓	✓
<input type="radio"/>	Order By	311	✓	✓
<input type="radio"/>	Inner Join	216	✓	✓
<input type="radio"/>	Outer Join	20	✓	✓
<input type="radio"/>	Exception Join	0		
<input type="radio"/>	Union	67	✓	✓
<input type="radio"/>	Except	0		
<input type="radio"/>	Intersect	0		
<input type="radio"/>	Distinct	62	✓	✓
<input type="radio"/>	Variables Specified	2119	✓	✓
<input type="radio"/>	Call Statements	474	✓	✓
<input type="radio"/>	Select Statements	2346	✓	✓
<input type="radio"/>	Update Statements	513	✓	✓
<input type="radio"/>	Insert Statements	174	✓	✓
<input type="radio"/>	Delete Statements	36	✓	✓

Figure 8-31 Part 1 of the types of SQL statement requested

<input type="radio"/>	Data Definition Statements	491	✓	✓
<input type="radio"/>	Other Statements	6063	✓	✓
<input type="radio"/>	Static SQL	4154	✓	✓
<input type="radio"/>	Extended Dynamic	0		
<input type="radio"/>	System-wide Cache Dynamic	900	✓	✓
<input type="radio"/>	Dynamic	2676	✓	✓
<input type="radio"/>	No Commit	7988	✓	✓
<input type="radio"/>	Uncommitted Read	735	✓	✓
<input type="radio"/>	Cursor Stability	0		
<input type="radio"/>	Cursor Stability KEEP LOCKS	0		
<input type="radio"/>	Read Stability	0		
<input type="radio"/>	Repeatable Read	0		

Figure 8-32 Part 2 of the types of SQL statement requested

<input type="radio"/>	Not Applicable	1374	✓	✓
<input type="radio"/>	Data Conversions	2441	✓	✓
<input type="radio"/>	Success	8364	✓	✓
<input type="radio"/>	Warning	173	✓	✓
<input type="radio"/>	No Data	1430	✓	✓
<input type="radio"/>	Error	130	✓	✓
<input type="radio"/>	Sort Sequence	0		
<input checked="" type="radio"/>	No Sort Sequence	10097	✓	✓

Figure 8-33 Part 3 of the types of SQL statement requested

Under Miscellaneous Information, shown in Figure 8-34, you can often find hints about single longer running requests, as you can see the longest running open and fetch. This may be of interest to investigate.

<input type="radio"/>	▼ Miscellaneous information		
<input type="radio"/>	Maximum Runtime	1.958408	✓
<input type="radio"/>	Maximum Open Time	1.854992	✓
<input type="radio"/>	Maximum Fetch Time	0.98996	✓
<input type="radio"/>	Maximum Close Time	0.034752	✓
<input type="radio"/>	Maximum Other Time	1.958408	✓
<input type="radio"/>	Maximum Table Rows	306859	✓
<input type="radio"/>	Maximum Rows Returned	2880	✓

Figure 8-34 DBMON miscellaneous information

In the I/O information area you can find a lot of information that may be of interest, but usually you need not look very closely into that.

<input type="radio"/>	▼ I/O information			
<input type="radio"/>	Average Row Retrieval CPU Time	0.001964	✓	✓
<input type="radio"/>	Average Row Retrieval Clock Time	0.006579		
<input type="radio"/>	Average Synchronous DB Reads	6.109352		
<input type="radio"/>	Average Synchronous DB Writes	0.273381		
<input type="radio"/>	Average Asynchronous DB Reads	5.250359		
<input type="radio"/>	Average Asynchronous DB Writes	0.230215		
<input type="radio"/>	Average Page Faults	6.109352		
<input type="radio"/>	Average Rows Returned To DB	144.535		
<input type="radio"/>	Average Calls To Retrieve Rows	8.825		

Figure 8-35 DBMON I/O information

Now its time to look more in-depth at those areas that you may have located. One that should always be investigated is the recommended indexes. Select **Index Creates Advised** and click **Summary**, as shown in Figure 8-36.

IBM Systems Director Navigator for i5/OS® Welcome morten

Databases x Performance... x

<input type="radio"/>	Average Runtime	0.008797		
<input type="radio"/>	Average Parallel Degree Used	1		
<input type="radio"/>	Maximum Parallel Degree	1		
<input type="radio"/>	SQE	2695	✓	✓
<input type="radio"/>	CQE	606	✓	✓
<input type="radio"/>	System Naming	5789	✓	✓
<input type="radio"/>	SQL Naming	2934	✓	✓
<input type="radio"/>	Unique Open Statements	305	✓	
<input type="radio"/>	Full Opens	728	✓	✓
<input type="radio"/>	Pseudo Opens	2256	✓	✓
<input type="radio"/>	Table Scans	783	✓	✓
<input type="radio"/>	Average MQTs Used	0		
<input type="radio"/>	Average Indexes Used	1,015	✓	✓
<input type="radio"/>	Full Indexes Created	113	✓	✓
<input type="radio"/>	Sparse Indexes Created	0		
<input type="radio"/>	Index From Index Created	6	✓	✓
<input checked="" type="radio"/>	Index Creates Advised	286	✓	✓
<input type="radio"/>	Advised Statistics	277	✓	✓
<input type="radio"/>	Temporary Tables	258	✓	✓

Figure 8-36 Select Summary information about the index creates advises

Then you will see overview information about the indexes recommended and the query that resulted in the index advice. You can also click the Statement radio button to see each individual SQL statement with the indexes recommended.

If you want to analyze the individual statement more closely, then you would now need to use Visual Explain in Navigator for i to get a graphical view of the queries.

When you have your SQL Performance monitor you can also compare it to other SQL Performance Monitors. Then you may be able to see trends for the SQL running during different periods.

## 8.6 SQL Plan Cache

The plan cache is a repository that contains the access plans for queries that were optimized by SQL Query Engine (SQE). Unlike the SQL Performance Monitor, the SQL Plan Cache is always available. There is nothing to start for collecting this information. This is a very good reason to use the SQL Plan Cache versus the SQL Performance Monitor. For example, if you had a query that ran slow an hour ago, you could still find the query in the SQL Plan Cache. The SQL Plan Cache will be a good starting approach for a specific query if you do not have an SQL Performance Monitor started already and you are concerned about starting one to try to capture a situation that may or may not be able to be recreated.



Access plans generated by Classic Query Engine (CQE) are not stored in the plan cache. They can be stored in SQL packages, the system-wide statement cache, and job cache. The purpose of the plan cache is to:

- ▶ Facilitate the reuse of a query access plan when the same query is re-executed.
- ▶ Store runtime information for subsequent use in future query optimizations.
- ▶ Provide performance information for analysis and tuning.

You can access the plan cache by clicking the arrow next to the local database, then clicking **SQL Plan Cache** in the pull-down window, and then **Show Statements**, as shown in Figure 8-37.

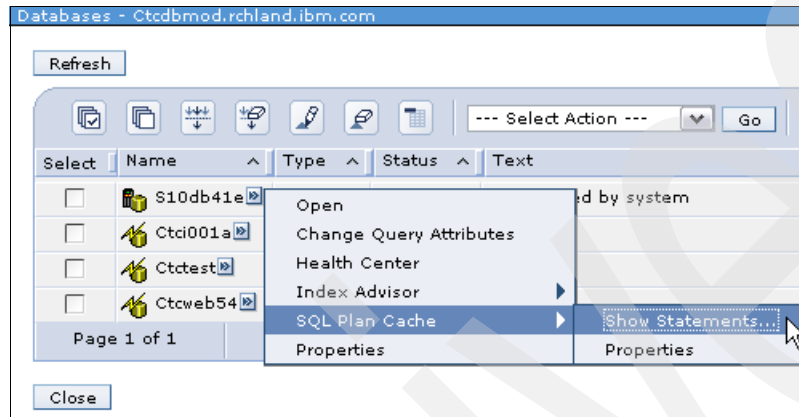


Figure 8-37 Access SQL Plan Cache

Then you get the options to filter your search in the plan cache, as shown in Figure 8-38.

SQL Plan Cache Statements - Localhost(S10db41e)

Filters to apply:

- Minimum runtime for the longest execution of the statement:  
0 0... Seconds
- Statements that ran on or after this date and time:  
25/08/09 18:07:07 Example: 12:30:00
- Top 'n' most frequently run statements:  
0 0...
- Top 'n' statements with the largest total accumulated runtime:  
0 0...
- Statements the following user has ever run:  
Morten
- Statements that are currently active
- Statements for which an index has been advised
- Statements for which statistics have been advised
- Include statements initiated by the operating system
- Statements that reference the following objects:

Schema	Name

- Statements that contain the following text:

Reset All Filters Apply

Figure 8-38 Filter the selection from the plan cache

When you click the CQE radio button (Figure 8-38 on page 132) your request will run, and you will be able to see the output in a format like that shown in Figure 8-39.

Last Time Run ^	Most Expensive Time (sec) ^	Total Processing Time (sec) ^	Total Times Run ^	Average Processing Time (sec) ^	Statement ^
20/08/09 12:42:44	0.9886	0.9886	1	0.9886	SELECT SYS_MNAME FROM QSYS2/SYSPSTAT WHERE SYS_DNAME = ? AND SYS_TNAME = ?
20/08/09 12:42:59	0.2773	0.2773	1	0.2773	SELECT ? AS MBRNAME, XITSO001 .* FROM QTEMP/PW9036_ITS0001 XITSO001 ...
25/08/09 17:50:50	0.0148	0.0307	4	0.0076	SELECT NAME_00001, STATU00001, JOBS_00001, DATA_00001, CREAT00001, SAVED000...
25/08/09 15:00:05	0.0036	0.0039	6	0.0006	SELECT ODLBNM, ODOBNM, ODOBTP, ODOBTX, ODOBOW, ODJRNM, ODJRLB, ODASP FROM Q...
25/08/09 13:26:23	0.0001	0.0008	9	0.0001	DECLARE X CURSOR FOR SELECT Partition_Name , REPLACE ( Leading_Column_Keys ...
25/08/09 13:10:32	0.0000	0.0001	2	0.0001	select * from mortenlib.act where actkwd = ?
25/08/09 13:05:53	0.0000	0.0001	3	0.0001	select * from mortenlib.act
19/08/09 13:28:26	0.0000	0.0001	2	0.0001	select * from mortenlib.act where actkwd = ? and actno > ?
20/08/09 12:43:18	0.0001	0.0001	1	0.0001	DECLARE Environ_y SCROLL CURSOR WITH RETURN FOR SELECT * FROM SESSION / Hea...
20/08/09 13:22:35	0.0001	0.0001	1	0.0001	DECLARE Environ_y SCROLL CURSOR WITH RETURN FOR SELECT * FROM SESSION / Hea...

Page 1 of 2 | 1 | Go | Total: 20 | Displayed: 10 | Selected: 0

Figure 8-39 Example of output from plan cache


You now have the ability to investigate the queries and find much information about the different runs.

If you want to do a more in-depth analysis of the queries, you will need Visual Explain in Navigator for i to be able to see a graphical view of the query plan and execution.

When you want to keep the information in the SQL Plan Cache, you have the ability to save that query to an SQL Plan Cache Snapshot™. A snapshot is a database monitor file generated from the plan cache at the time that a new snapshot is requested. It can be treated very much the same as the SQL Performance Monitors list. The same analysis capability exists for snapshots as exists for traditional SQL performance monitors. So you can use the analysis tools the same way that you did for the SQL Performance Monitor collection.

It is important to make the snapshots if you want to keep the information, because the SQL Plan Cache is cleared after an IPL and plans can be cleaned up and deleted by the system.

You can create the snapshot shown in Figure 8-40 with the following steps:

1. Click the arrow to the right of the local database name, then click **Open**.
2. Click the arrow to the right of the SQL Plan Cache, then click **Open**.
3. Click the arrow to the right of the SQL Plan Cache Snapshots, then select **New** .

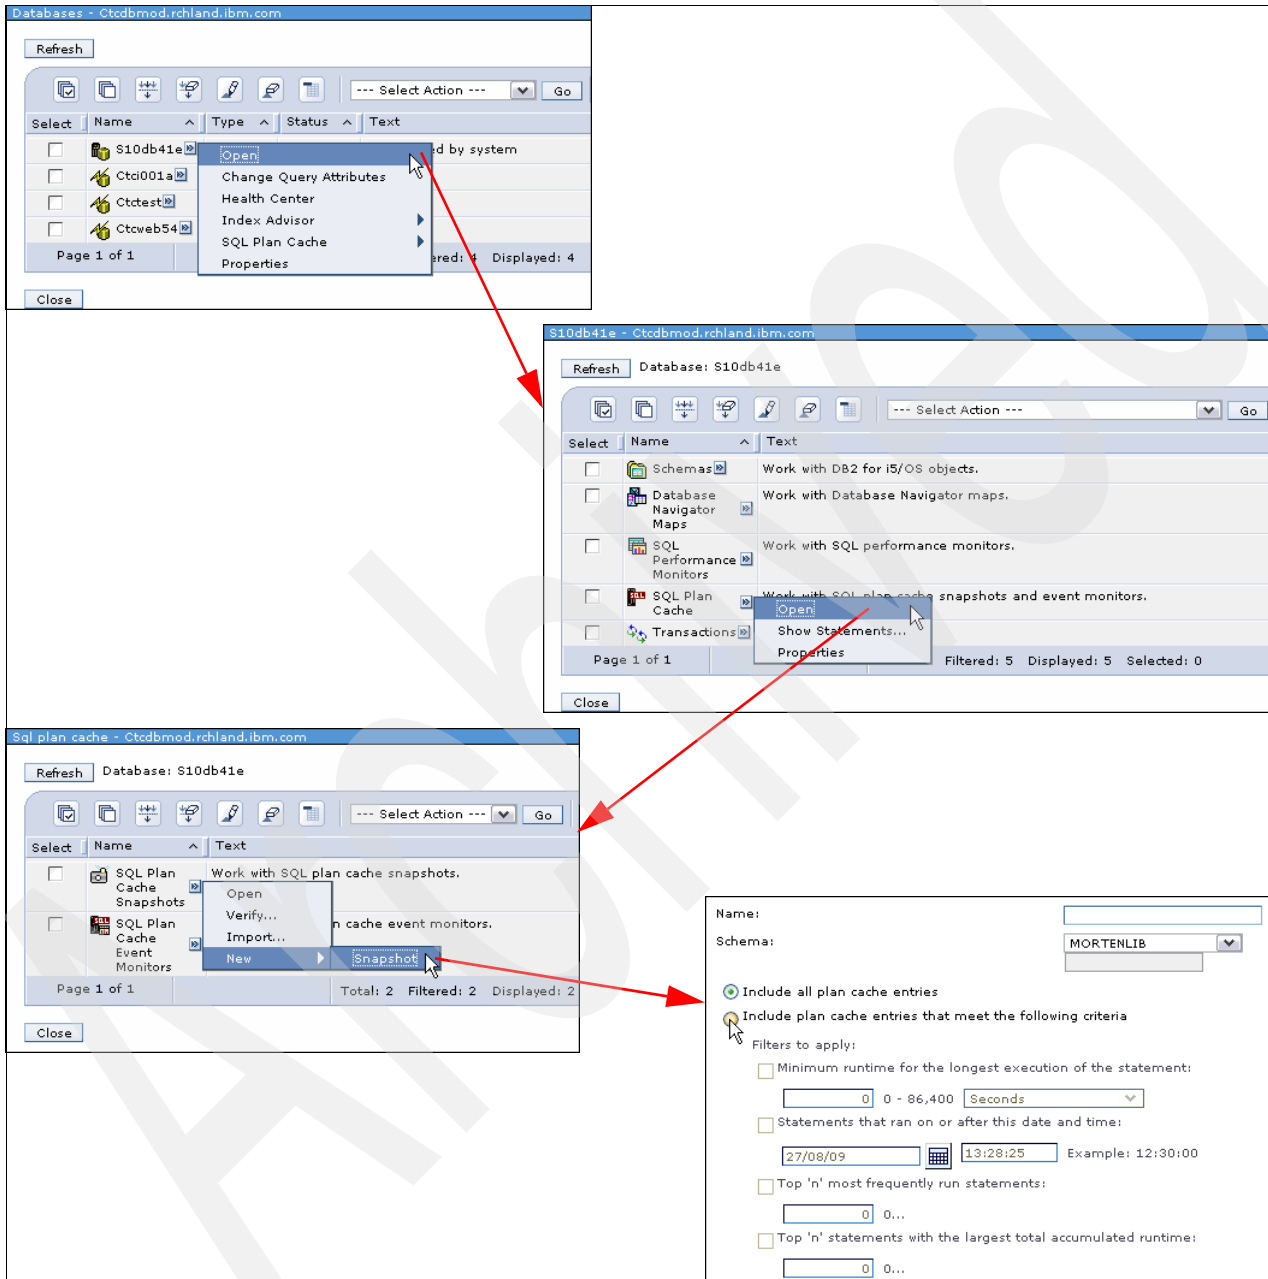


Figure 8-40 Make a SQL Plan Cache snapshot

4. The last step is to select how much you want to save from the SQL Plan Cache into the SQL Plan Cache Snapshot. As you can see in Figure 8-40, you can select all or you can filter your selection.

When you have your snapshot, you can also compare it to an other snapshot. Then you may be able to see trends for the SQL running during different periods. This could be weekly or

monthly trends. It could be a weekly or monthly capture of the top 10 longest running statements.

## 8.7 Show SQL details for job

On the database level, you can select **SQL Details for Jobs** from the pull-down menu, then click the **Go** radio button, as shown in Figure 8-41, which takes you to the view for the jobs.

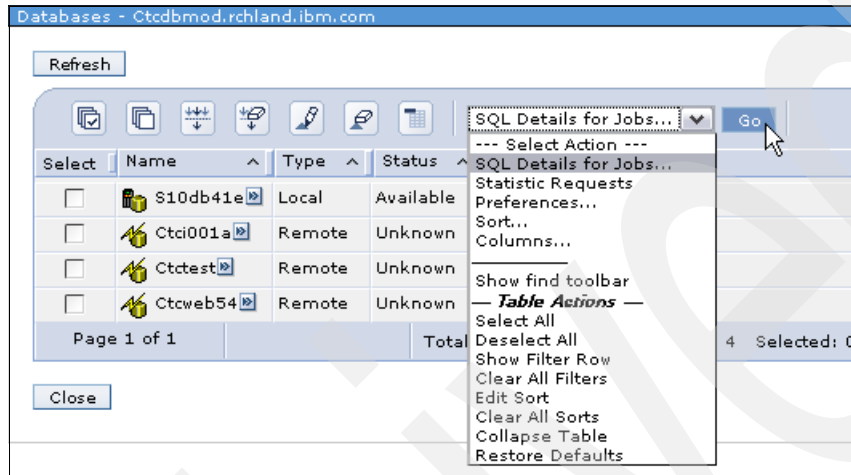


Figure 8-41 SQL Details for Jobs

Similar information can be found under the active jobs (Figure 8-42).

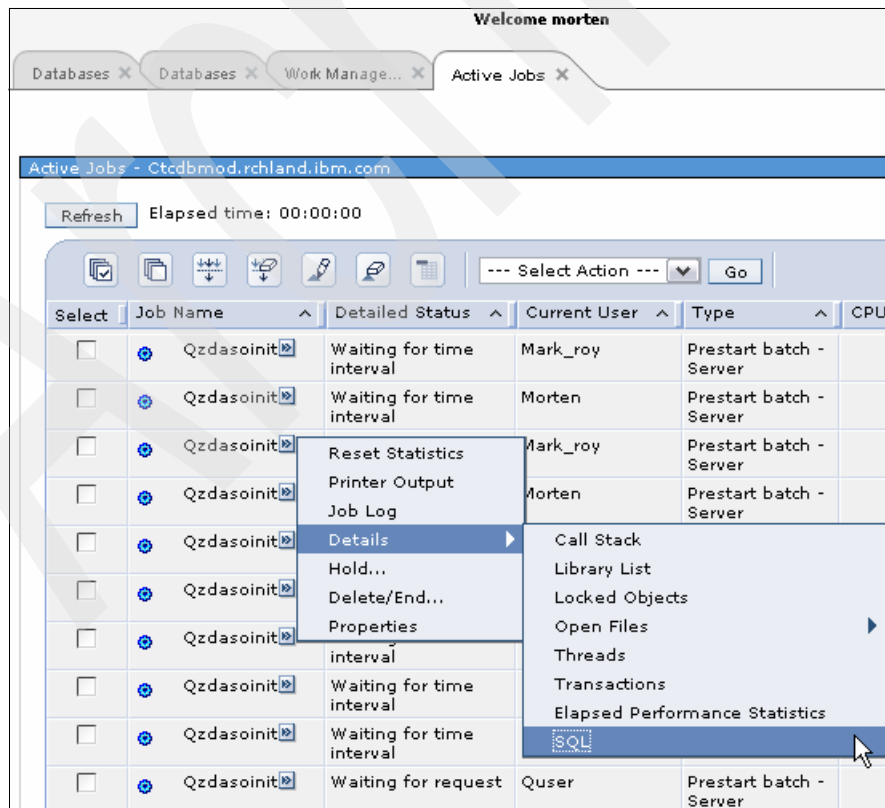


Figure 8-42 SQL details from active jobs

Both functions are good to get an idea of what a resource-demanding job is doing.

## 8.8 Index advisor

You can display index advisor information from the optimizer using Systems Director Navigator. Systems Director Navigator displays extracts from information found in the QSYS2/SYSIXADV system table. You can display index advice on a system level, a schema level, or a table level.

You can access all indexes advised on a system level by clicking the arrow icon to the right of the local database name, then the **Index Advisor** on the pop-up window and again **Index Advisor** on the second pop-up window, as shown in Figure 8-43.

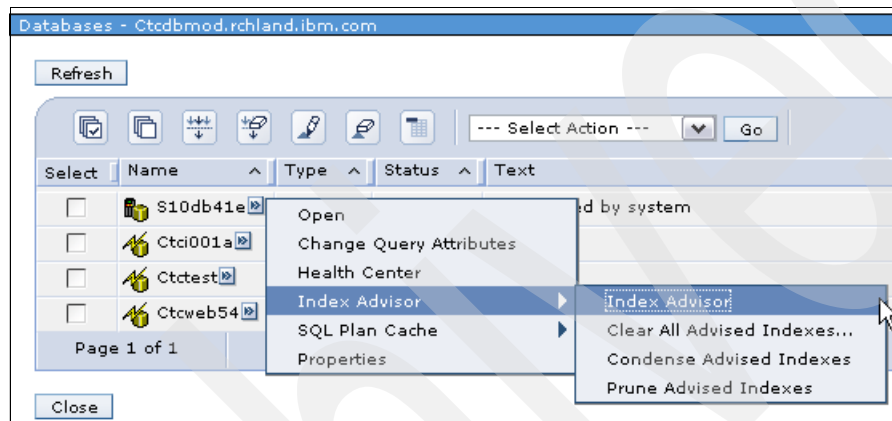


Figure 8-43 Index Advisor: system wide

Often you will work with one application or a single schema containing the tables, and then it would be a better choice to select only the indexes advised for all the tables in the schema. In that case you click the arrow to the right of the schema name, then the **Index Advisor** on the pop-up window, and again **Index Advisor** on the second pop-up window, as shown in Figure 8-44.

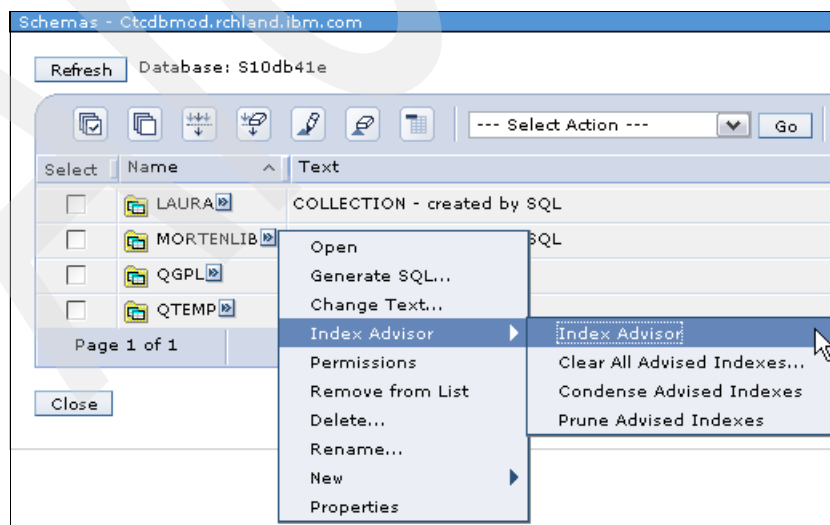


Figure 8-44 Index Adviser: library level

At other times you will have interest in checking indexes recommended for a single table. In that case you click the arrow to the right of the table name, then the **Index Advisor** on the pop-up window, and again **Index Advisor** on the second pop-p window, as shown in Figure 8-45. You can do this at any place where you have the table, and it does not matter how you access the table or from what menu.

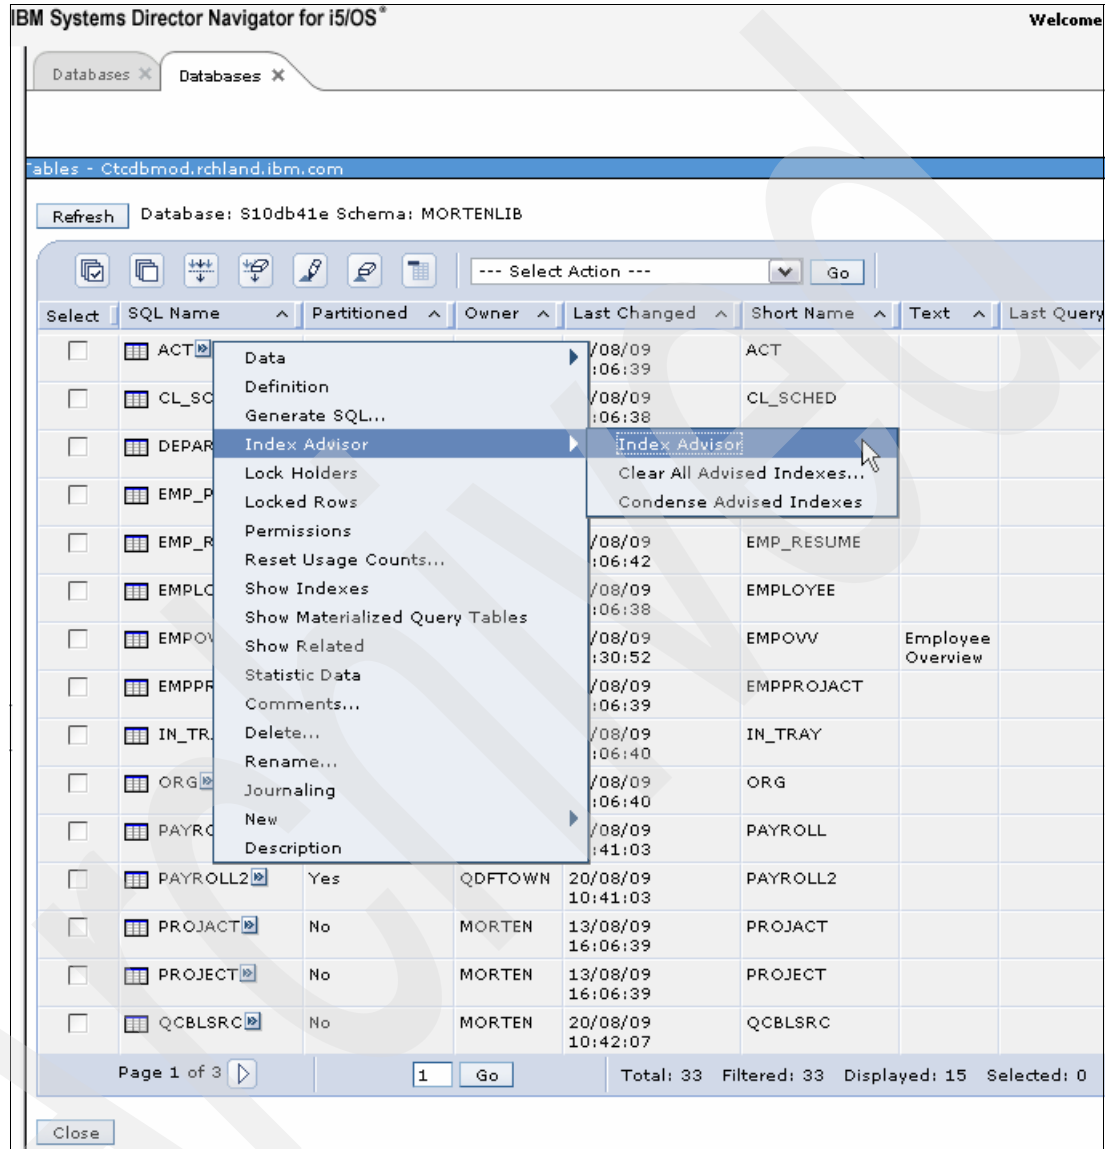


Figure 8-45 Index Advisor: table level

When you click the Index Advisor for the selected table you will see the indexes recommended for the table, as shown in Figure 8-46.

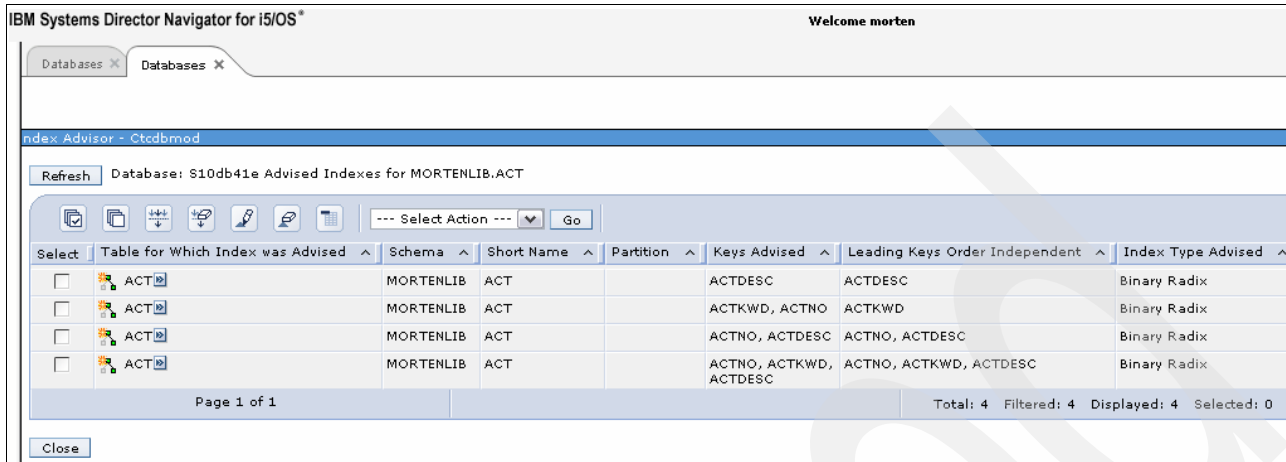


Figure 8-46 Index advisor: table-level example



The Index Advisor will often advise several different indexes for the same table. You can condense these advised indexes into the best matches for your queries. Again you can do that on system, schema, and table level. If we use an example of the table level, then you simply click the arrow radio button to the right of the table that you want to look at indexes for and select **Index Advised**, then **Condense Advised Index**, as shown in Figure 8-47.

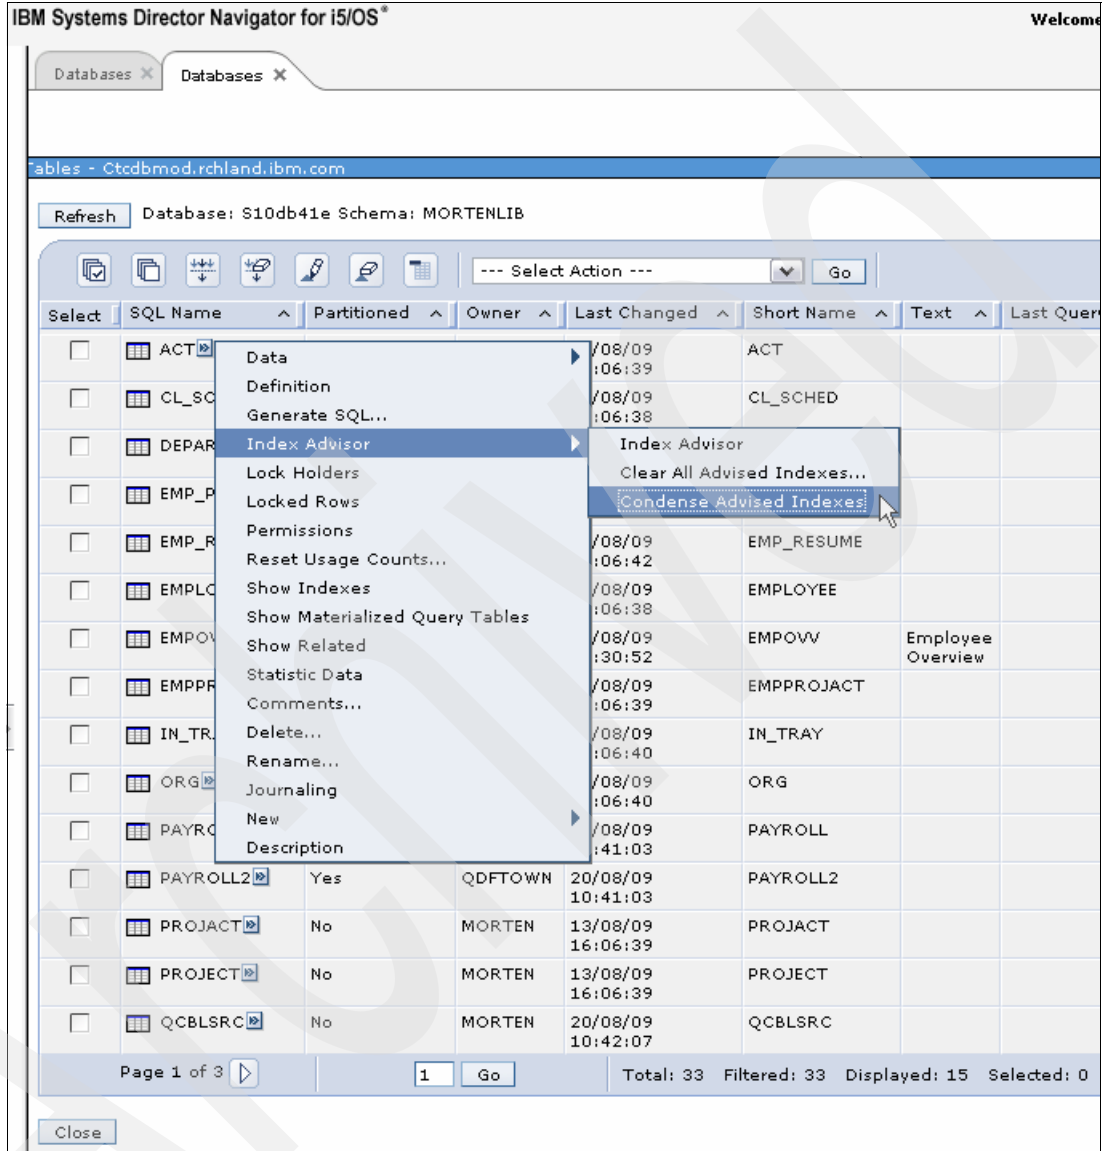


Figure 8-47 Select Condense Advised Indexes for a table

By clicking **Condense Advised Indexes** for the selected table you will have a result similar to the example shown in Figure 8-48.

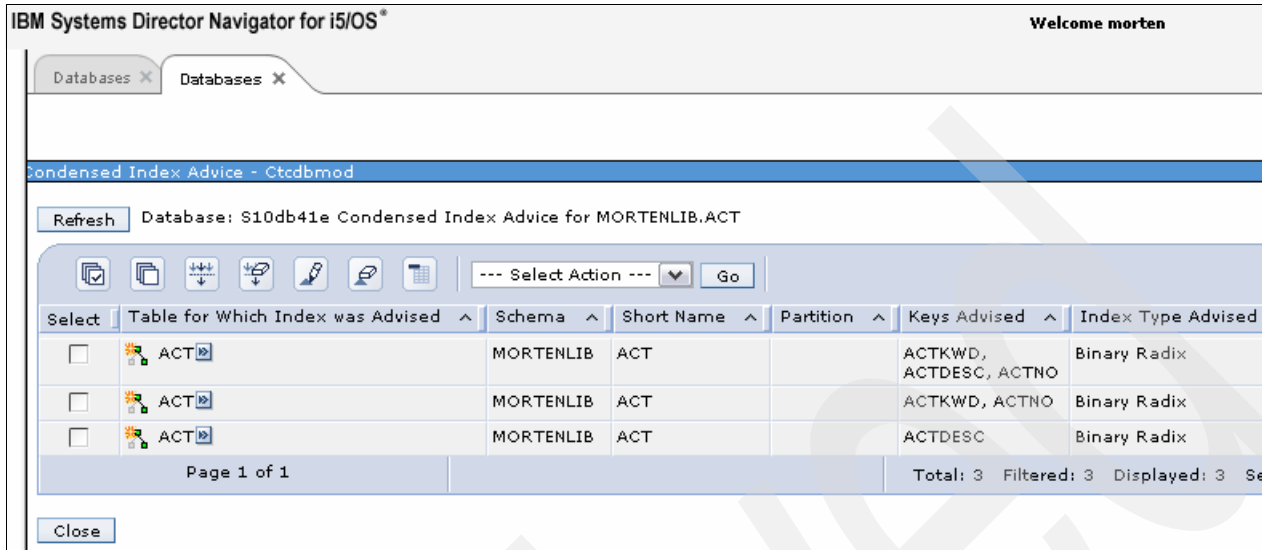


Figure 8-48 Example of Condensed Index Advice

The condensed index advice in Figure 8-48 shows that you have fewer indexes recommended than the indexes advised in Figure 8-46 on page 138. The condensed index advisor takes several factors into consideration and then eliminates indexes that can be fulfilled by other index recommendations.

The condensed index advised can be extracted on a system, schema, and table level, similar to the index advised.

## 8.9 Health center

The health center provides a great opportunity to have a good overview over your system from a database perspective. You can use the database health center to view different information about your database. You can view the total number of objects, the size limits of selected objects in your database, the design limits of selected objects, environmental limits, and activity level.

You can start the health center by clicking **Health Center** in the Database menu. That opens an overview, as shown in Figure 8-49. The default is an overview covering all libraries on the system.

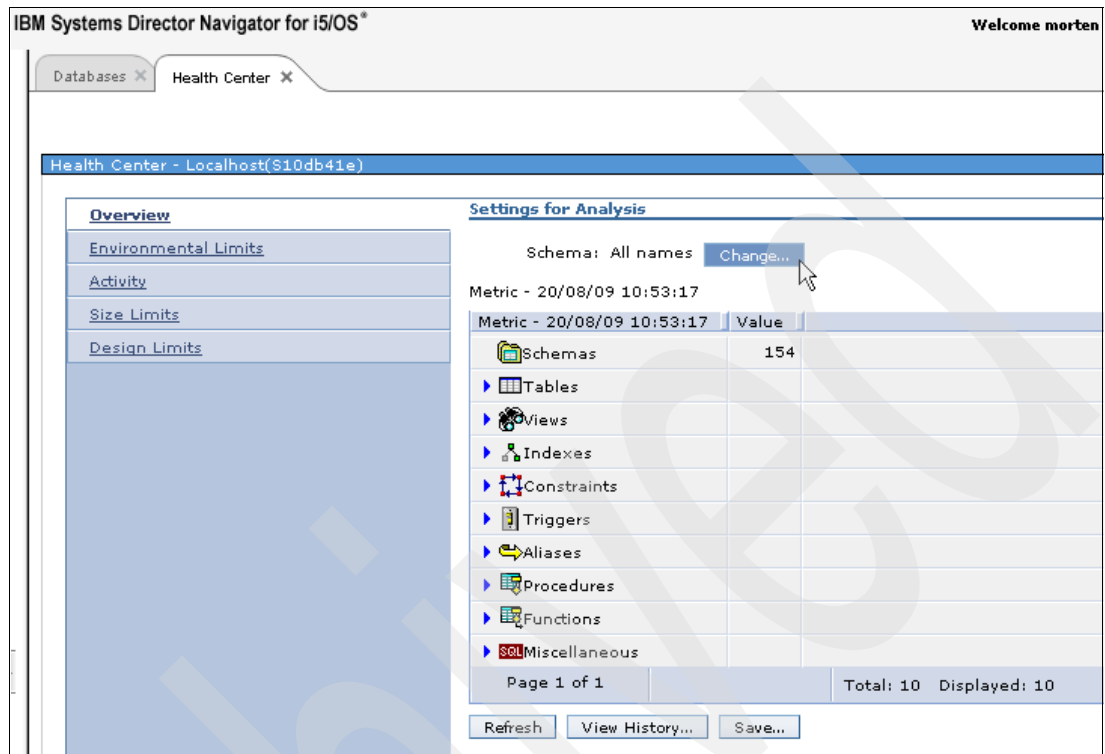


Figure 8-49 Health center overview

By clicking the arrows in front of the icons for the Tables, Views, Indexes, Constraints, Triggers, Aliases, Procedures, Functions, and Miscellaneous categories you will see addition information about the number of objects for the different types. You can change the schema that you are investigating by clicking the **Change** radio button, so you could select your own library, as shown in Figure 8-50.

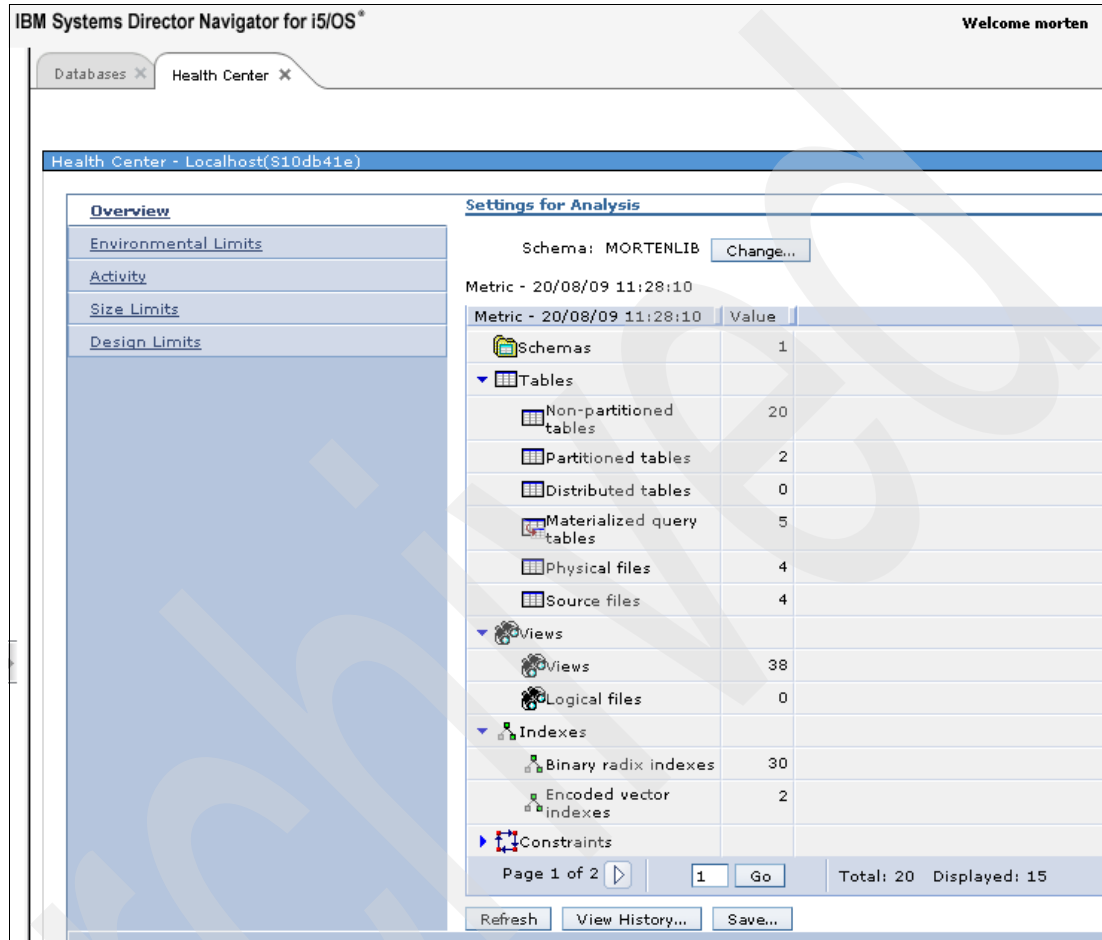


Figure 8-50 Health center overview for one library

You can change your preferences by clicking **Change** and entering filter information like schema and number of objects for activity.

You can now save your health center history to a table by clicking the **Save** button. This makes you compare your saved health center collections. This gives you an opportunity to save this information periodically, and then through the View History capability you can compare these collections to determine whether there are any trends to look into further.

If you click the Environmental Limits line, you will be able to acquire good information about the jobs running on the system right now. You will see different limits. If we take the maximum length of an SQL statement and expand that line, you will see the display shown in Figure 8-51.

Select	Environmental Limit - 28/08/09 10:12:36	Value	Percent of Limit	Status	When Value Was Recorded	Job Status
<input type="radio"/>	▶ Maximum number of activation groups to use SQL per job					
<input type="radio"/>	▶ Maximum number of active descriptors per job					
<input type="radio"/>	▶ Maximum number of CLI handles per job (160,000)					
<input type="radio"/>	▶ Maximum length of SQL statement per job (2 MB)					
<input type="radio"/>	● 017334/QUSER/QZDASOINIT (MORTEN)		2.14	Normal	28/08/09 10:07:31	Available
<input type="radio"/>	● 015717/QUSER/QZDASOINIT (MORTEN)		0.59	Normal	27/08/09 15:34:50	Not available
<input type="radio"/>	● 015716/QUSER/QZDASOINIT (MORTEN)		0.24	Normal	27/08/09 15:15:59	Not available
<input type="radio"/>	● 015179/QSYS/QDBSRV04 (QSYS)		0.08	Normal	27/08/09 14:58:21	Available
<input type="radio"/>	● 015180/QSYS/QDBSRV05 (QSYS)		0.08	Normal	27/08/09 14:58:23	Available
<input type="radio"/>	● 015286/QUSER/QSQSRVR (QSYS)		0.02	Normal	28/08/09 00:00:03	Available
<input type="radio"/>	015293/QSYS/CRTPFPRDTA (QSYS)		0.02	Normal	28/08/09 00:00:06	Not available
<input type="radio"/>	015287/QUSER/QSQSRVR (QDIRSRV)		0.02	Normal	27/08/09 14:58:34	Not available
<input type="radio"/>	● 016191/QSYS/CRTPFPRDTA (QSYS)		0.01	Normal	28/08/09 00:05:04	Available
<input type="radio"/>	015717/QUSER/QZDASOINIT (MORTEN)		0.01	Normal	27/08/09 15:34:11	Not available

Page 1 of 1      Total: 14    Displayed: 14    Selected: 0

Refresh    Clear...    View History...    Save...

Figure 8-51 Environmental limits under health center



The SQL details provide much interesting information (Figure 8-52 on page 144). You can find information about the SQL request, job information, and client information.

## 8.10 Journal management

You have all journal functionality for local journals within the IBM Systems Director Navigator for i. You can access the main journal menu from the main menu, as illustrated in Figure 8-53.

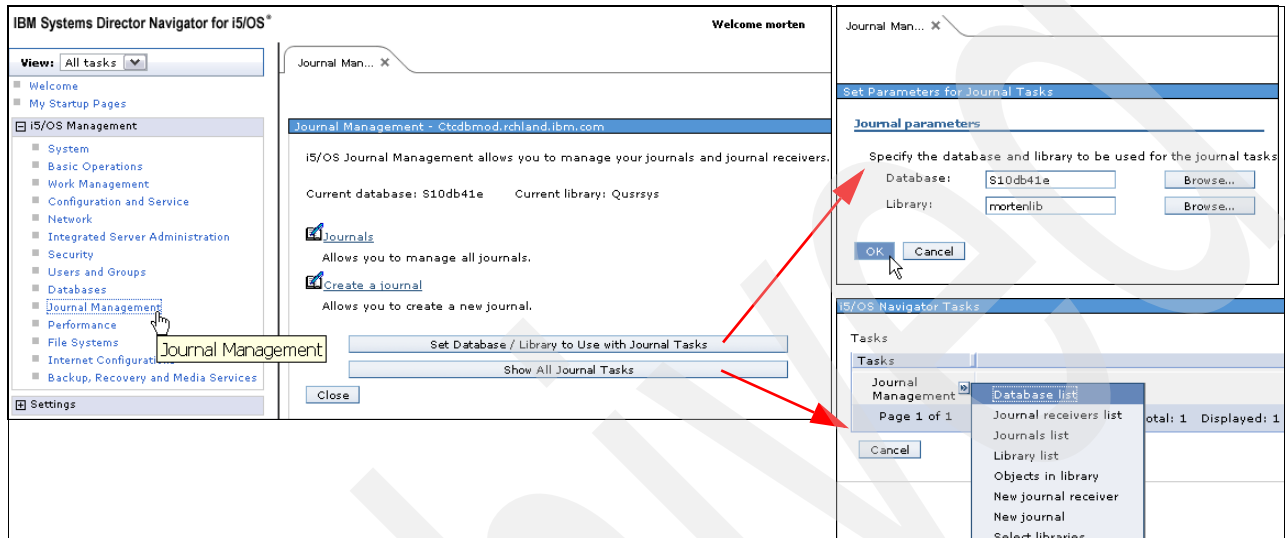


Figure 8-53 Access journal tasks from Director Navigator

Click **Journal Management** from i5/OS Management and you will see two buttons, one for setting the library for working with journals and one for showing all the journal tasks. You can also go directly to your journals in your chosen library by clicking the menu item or you can create a new journal by clicking the menu item for creating a journal.

When you choose to create a new journal, you will find an easy user interface and good help to get both journal and receiver created.

Figure 8-54 shows the easy steps to create a journal. You only must specify the journal name and the rest can be the default values. Of course you can change any of the suggested values on the three tabs covering the general settings, receiver options, and format options.

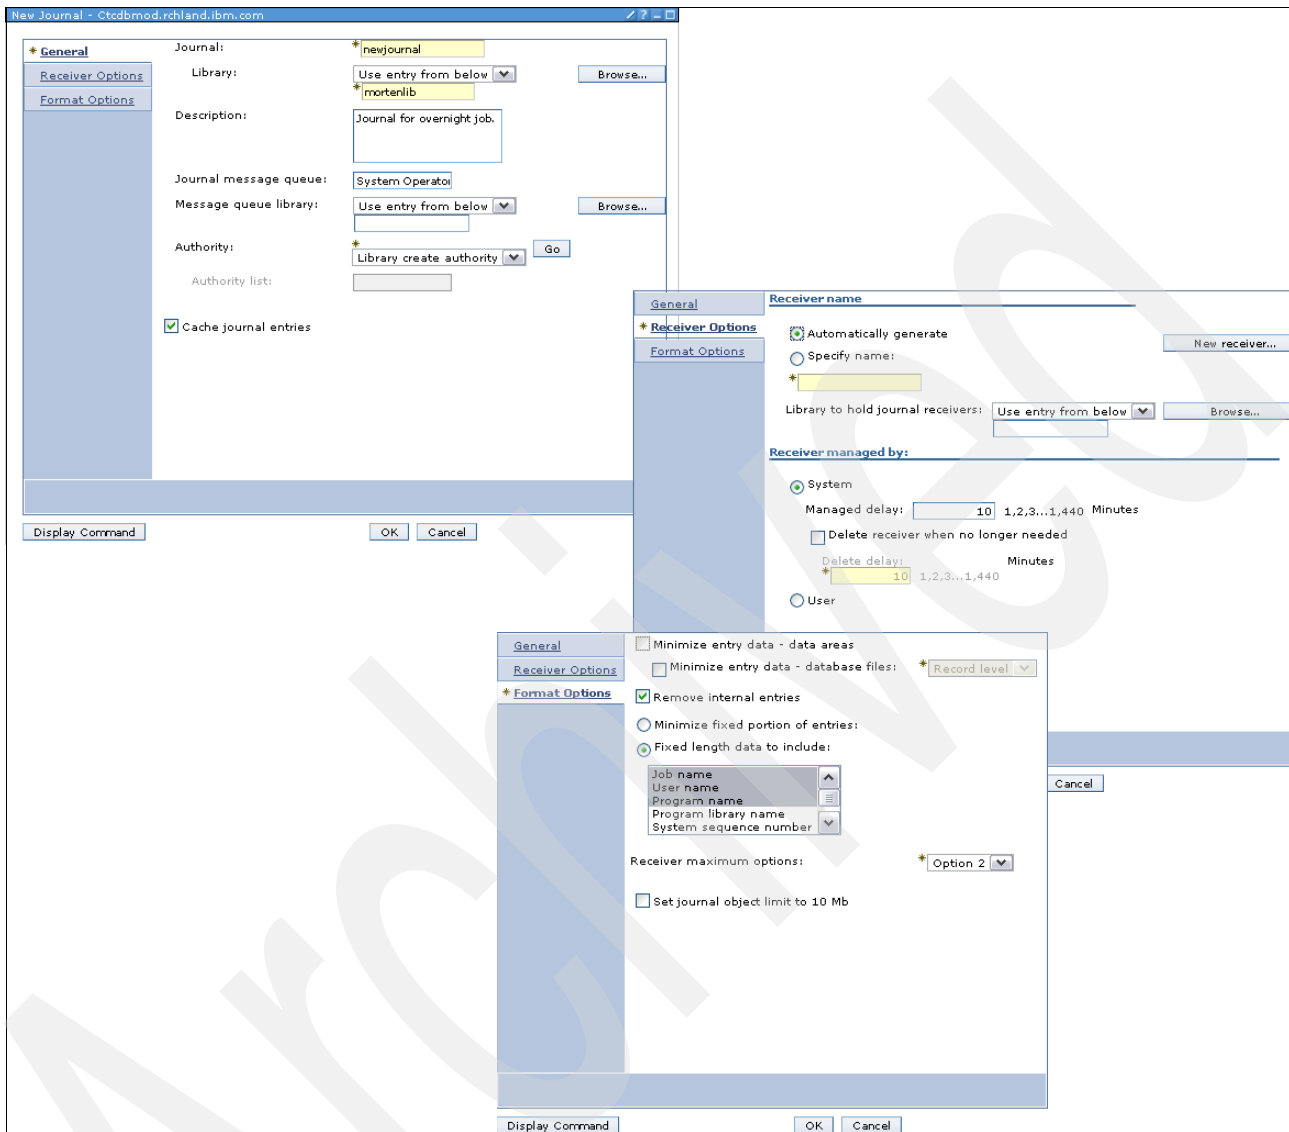


Figure 8-54 Creation of journal is easy using the default values





## Performance

This chapter contains the following:

- ▶ Overview of performance data views available through Performance Data Investigator in Systems Director Navigator
- ▶ Overview of collecting performance data with Systems Director Navigator
- ▶ Overview of viewing performance collections with Systems Director Navigator
- ▶ Overview of Health Indicators

## 9.1 Performance Data Investigator

Many different views of your performance data are provided through the Performance Data Investigator tool. They are naturally grouped according to task, either by collection type or for general health. Each of these sub-groups of views is provided through Performance Data Investigator to ensure a consistent interface, look-and-feel, and features set, as well as to allow you one place and one tool that can be used to analyze whatever performance data you choose. Performance Data Investigator also allows views to be added via extensions (as is the case for Disk Watcher, Job Watcher, and Performance Explorer). This tool can be launched by choosing **Investigate Data** (Figure 9-9 on page 160).

The performance data collectors are included in the base operating system. However, some of the analysis and monitoring views are packaged with a licensed program product. Table 9-1 provides the product name for each of the views.

Table 9-1 Tool packaging

View	Product/option
IBM Systems Director Navigator for i: Health Indicators	Base operating system
IBM Systems Director Navigator for i: Collection Services	Base operating system
IBM Systems Director Navigator for i: Job Watcher	Performance Tools for i (5761PT1) - Option 3 - Job Watcher
IBM Systems Director Navigator for i: Disk Watcher	Performance Tools for i (5761PT1) - Option 1 - Manager feature
IBM Systems Director Navigator for i: Performance Explorer	Performance Tools for i (5761PT1) - Option 1 - Manager feature

**Note:** You can read more about the different performance tools available through IBM Systems Director Navigator for i in the IBM Redbooks publication *End to End Performance Management on IBM i*, SG24-7808.

## 9.2 Collection Services

Collection Services is an IBM i function that is used for collecting system management data. It can be used to gather performance data continuously and automatically with minimal system overhead. The Collection Services interval data can be application defined or user defined and can be managed and analyzed using multiple interfaces.

### 9.2.1 How Collection Services works

Collection Services collects system and job-level data at regular intervals, which can be set from 15 seconds to 1 hour. At every sampling interval, Collection Services probes and gathers metrics from various system resources including jobs, disk units, IOPs, buses, pools, and communication lines.

The raw performance data is stored in a single management collection object that takes advantage of teraspace support to make it a more efficient way to store large quantities of performance data. This also supports a release-independent design, which allows moving the performance data to a system at a different release level without requiring database file conversions.

To use the performance data for analysis, performance database files must be generated using the CRTPFRTA CL command. Starting with 6.1, database performance files are created by default in a low-priority CRTPFRTA job, as the data is being collected to facilitate real-time analysis. In releases prior to 6.1, you could have the performance data files created automatically by configuring Collection Services to create the database files during the collection.

Independent management of data collection and performance database file generation facilitate the most efficient way of collecting data from a large number of jobs and tasks at frequent intervals. The associated overhead is minimal enough to justify running Collection Services continuously in the background and capturing data needed to analyze and solve performance issues before they turn into a serious problem.

Collection Services does not provide object-level and program-level details such as call stacks, activation groups, locks, and SQL statements. Despite this Collection Services is useful for a big-picture look at how a system has been performing over longer periods of time.

Performance Data Investigator within the Web-based IBM Systems Director Navigator for i provides exclusive views with many charts and drill-downs that are available with Collection Services data. Collection Services Investigator within IBM iDoctor for IBM i GUI also offers similar views of wait statistics.

## 9.2.2 Configuring Collection Services

Collection Services comes with default configured options, but they can be changed to suit your environment and needs. Collection Services configuration allows for selecting the library where the data is stored, cycle time and frequency, collection profile, data retention, and so on. The configuration can be done using multiple interfaces:

- ▶ IBM Navigator for i (Windows client)
- ▶ IBM Systems Director Navigator for i (Web-based console)
- ▶ Configure Performance Collection (CFGPFRCOL) CL command
- ▶ Change Collection Services Attributes (QypsChgColSrvAttributes) API

### Configuration options

Collection Services configuration options can be divided into three main categories and options in each category are explained in the below sections.

#### ***Basic Collection Services attributes***

Below is a list of general Collection Services configuration options:

<b>Default Collection Interval</b>	Default collection sample interval.
<b>Collection Retention Period</b>	Specifies how long the management collection objects are retained.
<b>Collection Library</b>	Library where the management collection objects are stored. Performance data is stored in the QPFRDATA library by default. The data retention policy for management collection objects and performance database files applies only to the current configured library.

<b>Cycle Time</b>	Specifies the time when the current collection is ended and data is collected in a new management collection object.
<b>Cycle Interval</b>	The elapsed time between cycles after which the current collection is ended and data is collected in a new management collection object by the Collection Services server job (QYSPFRCOL).
<b>Create database files</b>	Specifies whether to submit the low-priority companion job (CRTPFRDATA) to create performance database files and process the data in the current management collection object as it is collected.
<b>Default Collection Profile</b>	Specifies the collector definition that determines what categories of data to collect and how often to collect them.

### ***Performance database file support attributes***

Below are additional Collection Services attributes that support performance database files:

<b>Standard Data Retention</b>	Specifies how long the standard data in performance database files within the configured library is retained.
<b>Summary Data Creation</b>	Specifies whether to create additional performance database summary data to facilitate quicker analysis.

### ***Management central monitor support attributes***

Below are the different options that the management central monitor supports:

<b>Graph Data Creation</b>	Option to generate detailed graph data automatically into a *MGTCOL object with the *PFRDTL attribute when the collection is cycled.
<b>Graph Data Retention</b>	Specifies how long the management central graph data is retained on the system.
<b>Historical Data Creation</b>	Option to generate historic graph data automatically into a *MGTCOL object with the *PFRHST attribute when the collection is cycled.
<b>Historical Data Retention</b>	Specifies how long the management central historical graph data is retained on the system.

These configuration parameters are only available on System i Navigator.

**Note:** The cycle and data retention attribute changes take effect immediately, while others take effect when the collection is cycled. Refer to the help text on the CFGPFRCOL command for more information.

## **Configure using IBM Systems Director Navigator for i**

Collection Services can be configured using the performance interface in IBM Systems Director Navigator for i using the below steps:

1. Go to the IBM Systems Director Navigator for i page by pointing your browser to the IBM i's IP address or DNS name and specify the default port 2001, as shown below:  

```
http://<systemname>:2001/
```
2. Enter your IBM i user ID and password to log in to IBM Systems Director Navigator for i.

- On the left pane, expand **i5/OS management** and click **Performance**, as shown in Figure 9-1. When the Performance tab is displayed for the first time, it will only show the Investigate Data and Collections options.

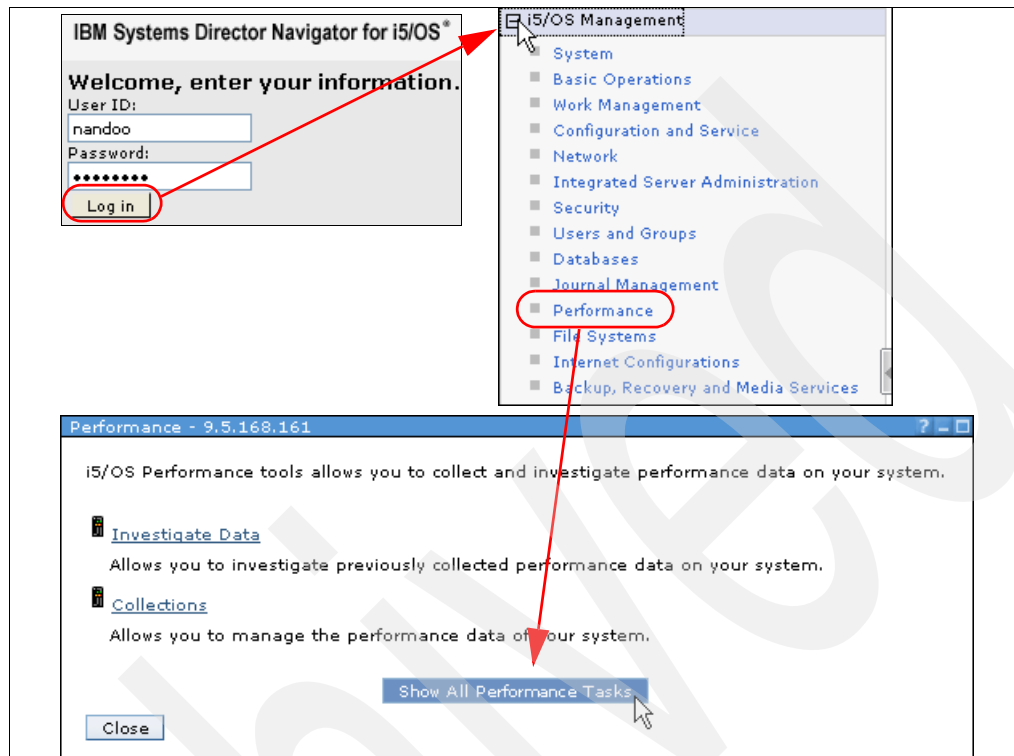


Figure 9-1 Navigate to Show All Performance Tasks

- On the right pane, click **Show all Performance Tasks** to display all available performance management tasks.
- As shown in Figure 9-2, expand **Collectors** and choose **Collection Services** **Configure Collection Services**.

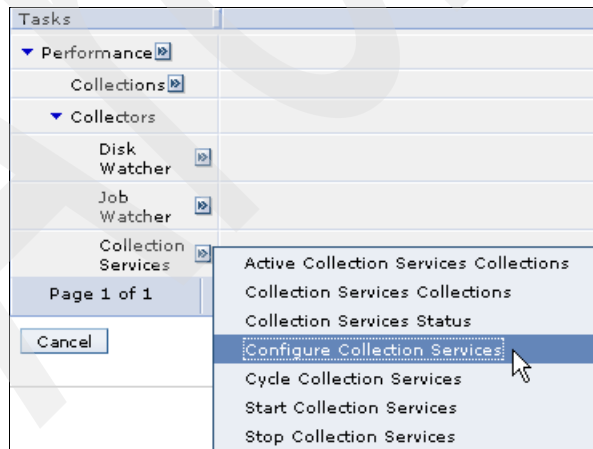


Figure 9-2 Configure Collection Services option

6. Make the changes as appropriate in the three tabs on the Configure Collection Services window (Figure 9-3) and click **OK**.

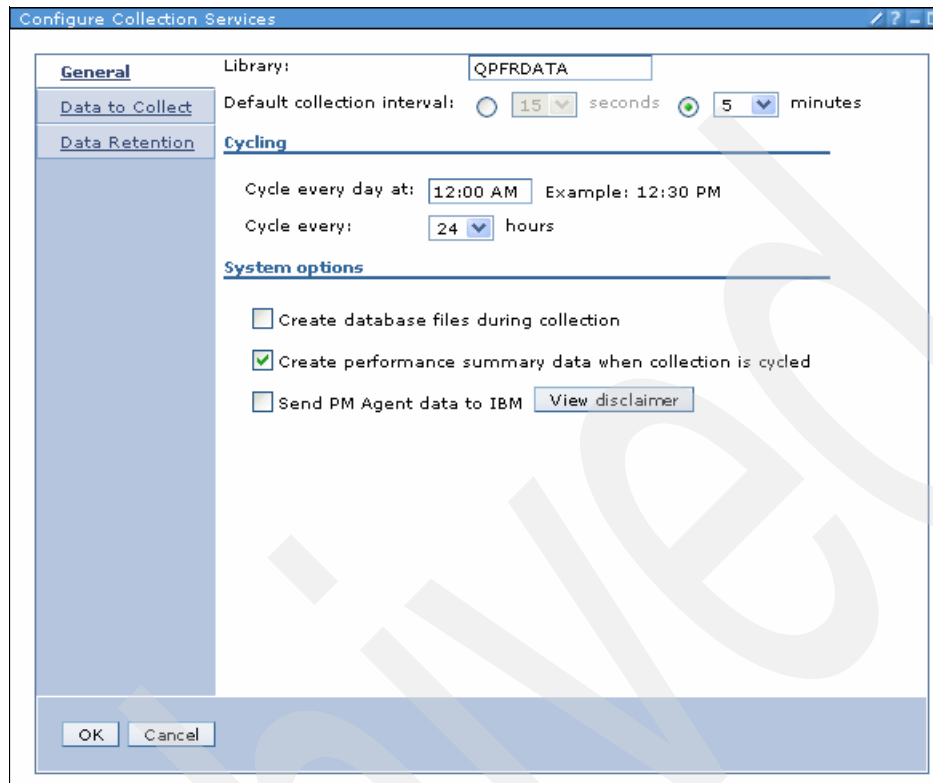


Figure 9-3 Configure Collection Services window

## 9.3 Job Watcher

Job Watcher provides you with the ability to collect many different types of information about a set of jobs (or all jobs and tasks) on a system *and* it does this while providing near real-time, detailed, and summarized views of job, thread, and task performance data that enable you to begin your analysis immediately. The data is collected by a server job, stored in database files, and displayed on a client via the iDoctor GUI presentation facilities or in the browser using the IBM Systems Director Navigator for i Performance interface.

### 9.3.1 How Job Watcher works

Job Watcher is similar in sampling function to the system commands WRKACTJOB and WRKSYSACT in that each refresh computes delta information for the ending snapshot interval. Refreshes can be set to occur automatically, as frequently as every 100 milliseconds. The data harvested from the jobs/threads/tasks being watched is done in a non-intrusive manner (similar to WRKSYSACT). The information harvested by Job Watcher includes:

- ▶ Standard WRKSYSACT-type information:
  - CPU
  - DASD I/O breakdown
  - DASD space consumption
  - Current user profile
  - And more

- ▶ Expanded details about types of waits and object lock/seize conditions
- ▶ Last-run SQL statements syntax
- ▶ Program/procedure call stack, 1000 levels deep
- ▶ Job Watcher also includes some data not available anywhere else in *real time*:
  - Contention time (This includes objects being locked at the operating system level and objects being seized at the microcode level in addition to the application-level object and record locks.): holder and waiter job and thread
  - A breakdown of all other waits that occurred:
    - Details about the current wait
    - Duration of the wait
    - Object being waited for
    - Conflicting job/thread/task information
    - Specific LIC block point ID
- ▶ SQL statements, host variables, prepared statement arrays, open cursors
- ▶ Communications data
- ▶ JVM Statistics
- ▶ Activation group statistics

Similar to the other collectors, Job Watcher also collects wait information, but the biggest advantage of using Job Watcher for performance analysis is its ability to use a combination of wait accounting, call stack information, and conflicting job/thread/task information during any contention. The buckets in Job Watcher consist of many waits that are generally considered good or bad, and seeing the bad ones on a graph like database record lock contention makes it easy to identify problem areas for further investigation.

### ***When to run Job Watcher***

Run Job Watcher when your system is running normally and you are not having performance problems. This gives you a health check of your system so that when it does have performance problems, you have a baseline to which to make comparisons.

You might want to run Job Watcher over your system for a period in the morning, afternoon, and possibly during the evening when you might have more batch-type jobs running.

You can also run Job Watcher when your system is not running normally to try to identify the problems.

**Tip:** You could also run Job Watcher over specific key jobs on your system when they are running well for baseline-comparison purposes. You will also want to run Job Watcher when you have an identified performance concern (for example, with Collection Services data) and need more detailed performance data for doing analysis and diagnostics.

## **9.3.2 Configuring Job Watcher**

Job Watcher can be configured and started using either IBM Systems Director Navigator for i, IBM iDoctor, or Job Watcher CL commands.

## Configure using IBM Systems Director Navigator for i

Job Watcher can be configured using the performance interface in IBM Systems Director Navigator for i using the below steps:

1. Go to the performance interface on the IBM Systems Director Navigator for i.
2. On the right pane, click **Show all Performance Tasks** to display all available performance management tasks.
3. Under the Performance Tasks, expand **Collectors** and choose **Job Watcher**  $\emptyset$  **Add a Job Watcher Definition**.

**Note:** The Job Watcher functions in the Systems Director Navigator Performance interface are only displayed if the Performance Tools Job Watcher option is installed.

4. Make the changes as appropriate using the Add Job Watcher Definition wizard and verify all the chosen options on the Summary page, as shown in Figure 9-4.
5. Click **Finish** to create the definition.



Figure 9-4 Add Job Watcher Definition

## Starting a Job Watcher collection

Job Watcher can be started using the performance interface in IBM Systems Director Navigator for i using the below steps:

1. Go to the performance interface on the IBM Systems Director Navigator for i.
2. On the right pane, click **Show all Performance Tasks** to display all available performance management tasks.
3. Under the Performance Tasks, expand **Collectors** and choose **Job Watcher**  $\emptyset$  **Start Job Watcher**.



4. Select a definition from IBM-supplied or user-defined definition type lists and click **Next**.
5. Specify new Collection interval or other Termination options to override values specified in the Job Watcher Definition and click **Next**.
6. Verify the settings again on the Summary page, as shown on Figure 9-5, and click **Finish** to start the Job Watcher.

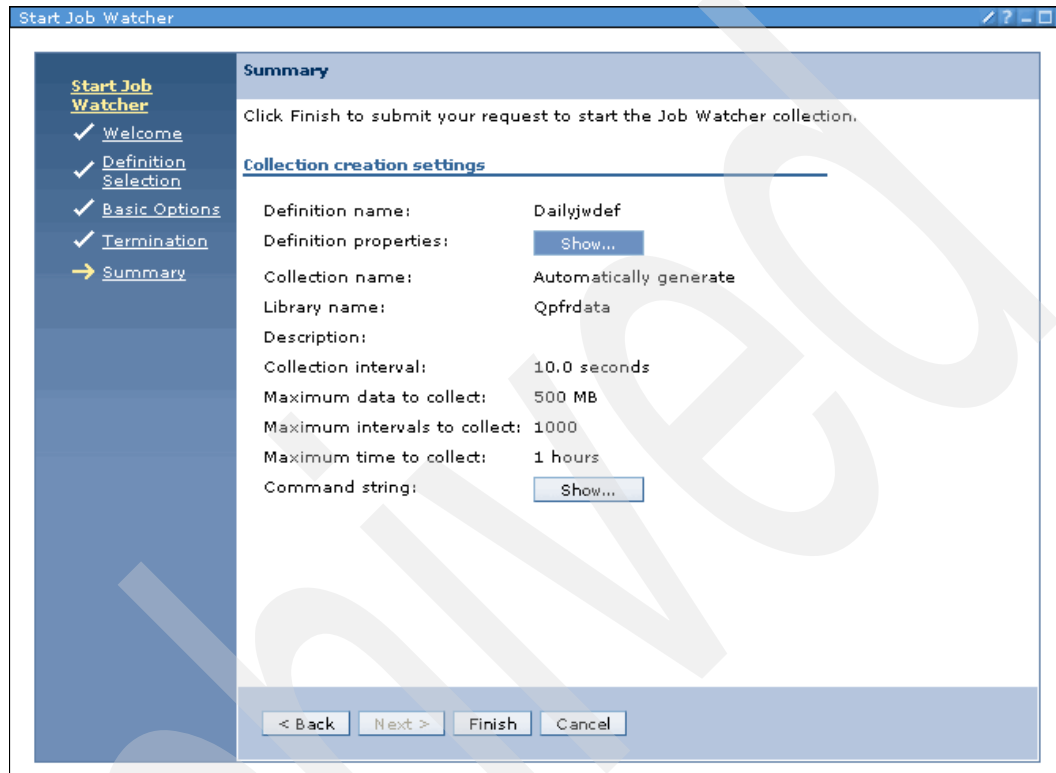


Figure 9-5 Start Job Watcher

For details about using the iDoctor Job Watcher, including topics on monitoring and analysis using Job Watcher, refer to the Redbooks publication *IBM iDoctor for iSeries Job Watcher: Advanced Performance Tool*, SG-24-6474, at:

<http://www.redbooks.ibm.com/abstracts/sg246474.html>

## 9.4 Disk Watcher

Disk Watcher provides you with the ability to collect either a statistical summary of disk performance data or a trace of all disk I/O events that occur on a system for near real-time analysis.

Disk Watcher goes beyond the WRKDSKSTS and WRKSYSSTS commands that provide information regarding disk I/O, paging rates, CPU usage, and temporary storage. It provides detailed information related to task dispatching elements (TDEs) and I/O queuing, objects, and programs involved with the I/O operation.

## 9.4.1 How Disk Watcher works

Disk Watcher collects detailed data on counts and durations of I/O operations to disk units. In addition, it provides useful runtime data regarding which objects are being read from disk or written to disk, and which jobs, threads, or system tasks are causing the operations. Figure 9-6 illustrates the high-level overview of Disk Watcher functions.

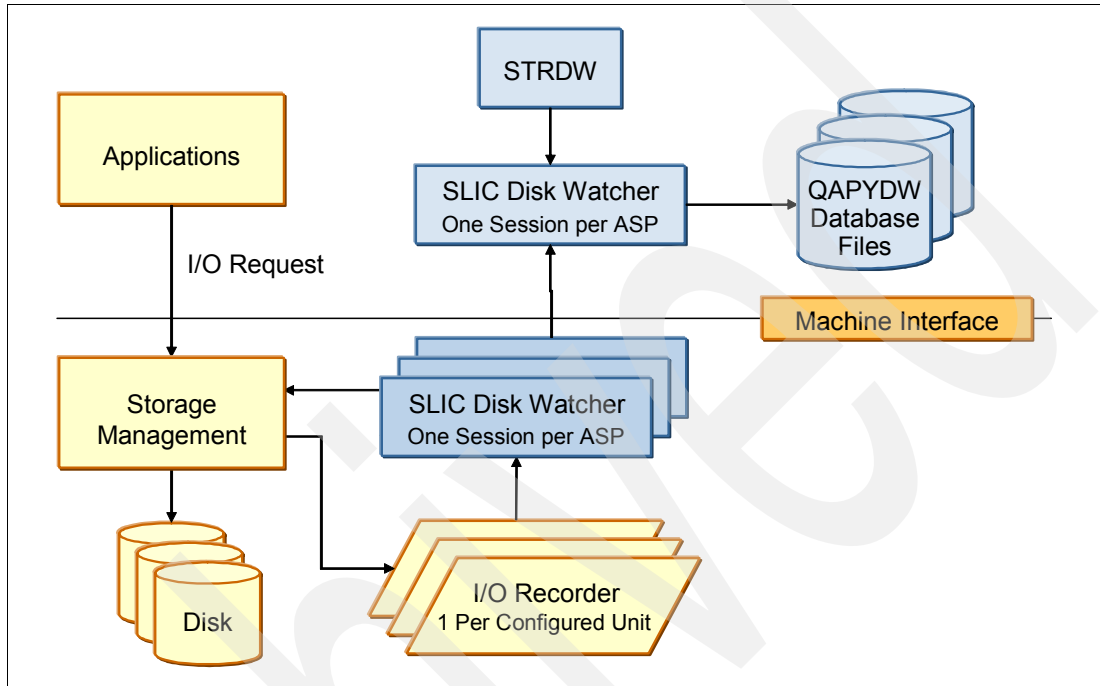


Figure 9-6 Disk Watcher high-level function overview

## 9.4.2 Configuring Disk Watcher

Disk Watcher can be configured and started using either IBM Systems Director Navigator for i, IBM iDoctor, or Disk Watcher CL commands.

### Configuration options

Disk Watcher configuration options can be divided into three main categories. The categories and their options are explained in the following sections.

#### **Basic Disk Watcher definition options**

The basic Disk Watcher definition options are:

<b>Collection Type</b>	Specifies the type of disk I/O collection to be performed, which determines the level detail of the data.
<b>Statistical</b>	I/O data is summarized into a set of counters based on the type of I/O operation for all I/O operations on the system.
<b>Trace</b>	I/O data is gathered for every I/O that which occurred on the specified disk or memory pool.
<b>Full</b>	Both statistical and trace data included.
<b>Collection Interval</b>	This is the interval between retrieval of disk I/O data. You can specify time or use dynamic intervals.

### ***Disk Watcher definition advanced options***

Below are additional Disk Watcher advanced options:

- ▶ **Data Availability**  
Specifies when the data is written to QAPYDW database files.
- ▶ **Collection file disk pool threshold**  
Specifies the percentage of disk pool that can be used before the collection will end. This threshold is for an auxiliary storage pool that contains the Disk Watcher database files. If the value specified is less than the current space available in the disk pool, the collection will end immediately.
- ▶ **System disk pool threshold**  
Specifies the percentage of the system disk pool that can be used before the collection is forced to end. Because Disk Watcher allocates temporary storage and the amount of data collected can be very large, this parameter allows you to limit how much of the system disk pool is consumed.

If the collection file exists in the system disk pool and values are specified in both the *collection file disk pool threshold* and *system disk pool threshold*, the value specified in the system disk pool threshold takes precedence.

**Note:** If the \*SYSTEM special value is specified for the disk pool thresholds, Disk Watcher uses the percentage that is configured with the change storage threshold function in the start system service tools (STRSST) command.

### ***Disk and memory pool selection***

Below are additional Disk Watcher selections to specify options for collecting disk I/O information:

- Disk pools** Specifies the disk pools for which disk I/O information is collected.
- Disk units** Specifies the disk units within the selected disk pools for which disk I/O information is collected.
- Memory pools** Specifies the main memory storage pools for which disk I/O information is collected (as data is read from disk into memory or written from memory back to disk).

### ***Object and program options***

The Disk Watcher option that controls whether object and program information is collected when running in *trace* mode is *object and program information*, which specifies that object and program information including their names is collected. It also specifies whether this data collection will be limited to objects and programs that are resident in main storage. If you want to view the trace information from within the Systems Director Navigator Web Console, you must specify \*RESIDENT or \*ALL for this parameter.

**Note:** Collecting information for objects that are not in main storage takes additional resources (CPU and I/O).

### Conditional collection options

Use the conditional collection options to specify additional conditional filters:

**Collection Condition** Specifies whether the data collected should be limited to I/O operations whose response time, wait time, or service time meet a specified condition.

**Collection Range** Specifies whether the data collected should be limited to I/O operations whose response time, wait time, or service time falls in a specified range.

### Add a definition using IBM Systems Director Navigator for i

Disk Watcher can be configured using the performance interface in IBM Systems Director Navigator for i using the below steps:

1. Go to performance interface on the IBM Systems Director Navigator for i.
2. On the right pane, click **Show all Performance Tasks** to display all available performance management tasks.
3. Under Performance Tasks, expand **Collectors** and choose **Disk Watcher**  $\otimes$  **Add a Disk Watcher Definition**.
4. Make the changes as appropriate using the Add Disk Watcher Definition wizard and verify all the chosen options on the Summary page, as shown on Figure 9-7.
5. Click **Finish** to create the definition.

**Note:** The Disk Watcher GUI options are only available if the Performance Tools License program manager feature is installed.

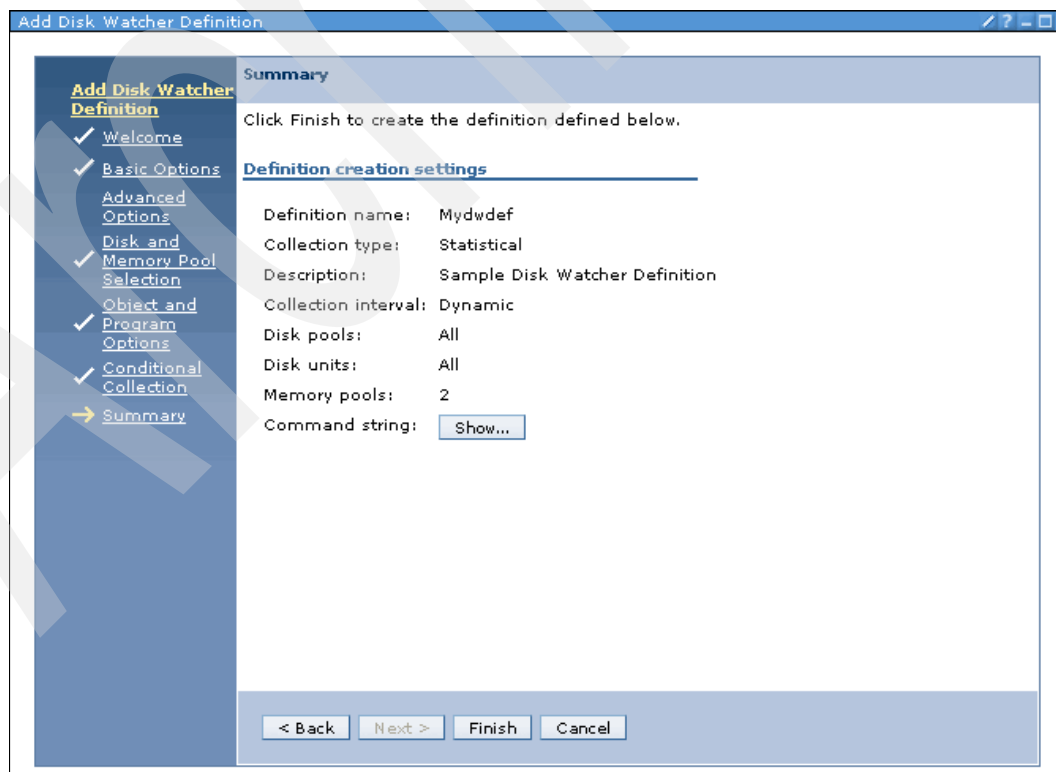


Figure 9-7 Add Disk Watcher Definition

## Start Disk Watcher using IBM Systems Director Navigator for i

Disk Watcher can be started using the performance interface in IBM Systems Director Navigator for i using the below steps:

1. Go to the performance interface on the IBM Systems Director Navigator for i.
2. On the right pane, click **Show all Performance Tasks** to display all available performance management tasks.
3. Under Performance Tasks, expand **Collectors** and choose **Disk Watcher**  $\otimes$  **Start Disk Watcher**.
4. Select a definition from the IBM-supplied or User-defined definition type lists and click **Next**.
5. Specify new a collection interval or other termination options to override values specified in the disk watcher definition and click **Next**.
6. Verify the settings again on the Summary page, as shown in Figure 9-8, and click **Finish** to start the Disk Watcher.

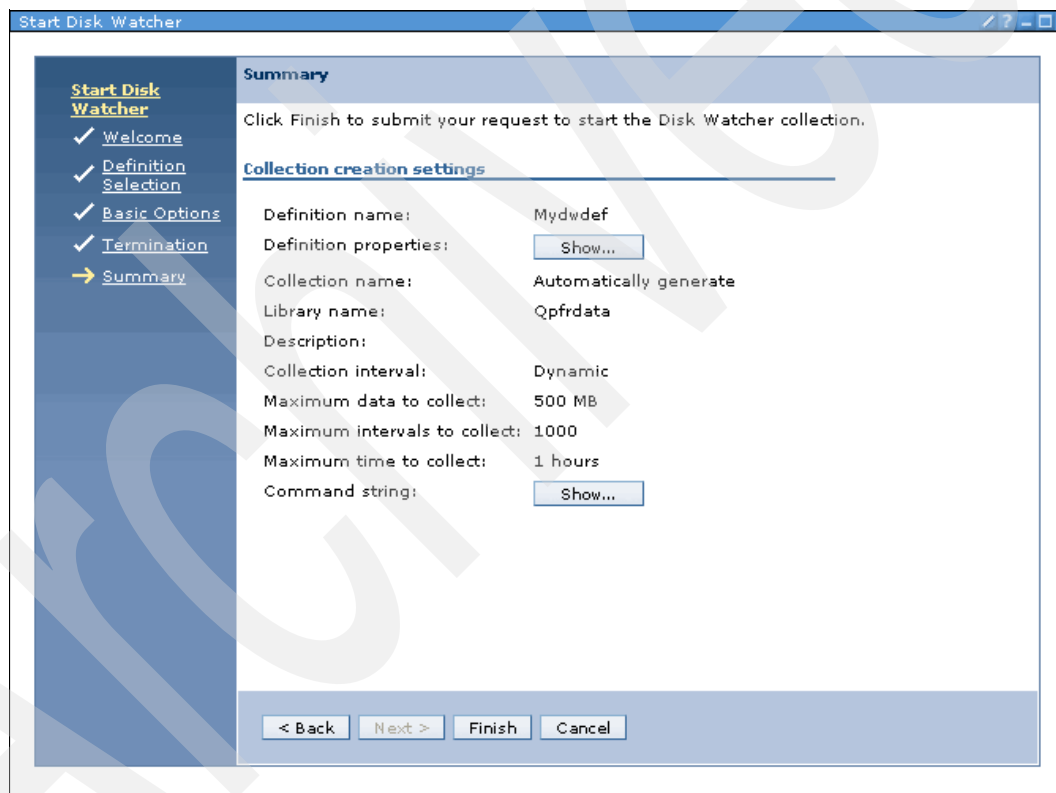


Figure 9-8 Start Disk Watcher

## 9.5 Performance Data Investigator

When you have collected data via Collection Services, Job Watcher and Disk Watcher, you can work with them to analyze the data. Systems Director Navigator provides a view of all of your collected performance data. Select **Investigate Data** from the main performance menu, as shown in Figure 9-9.

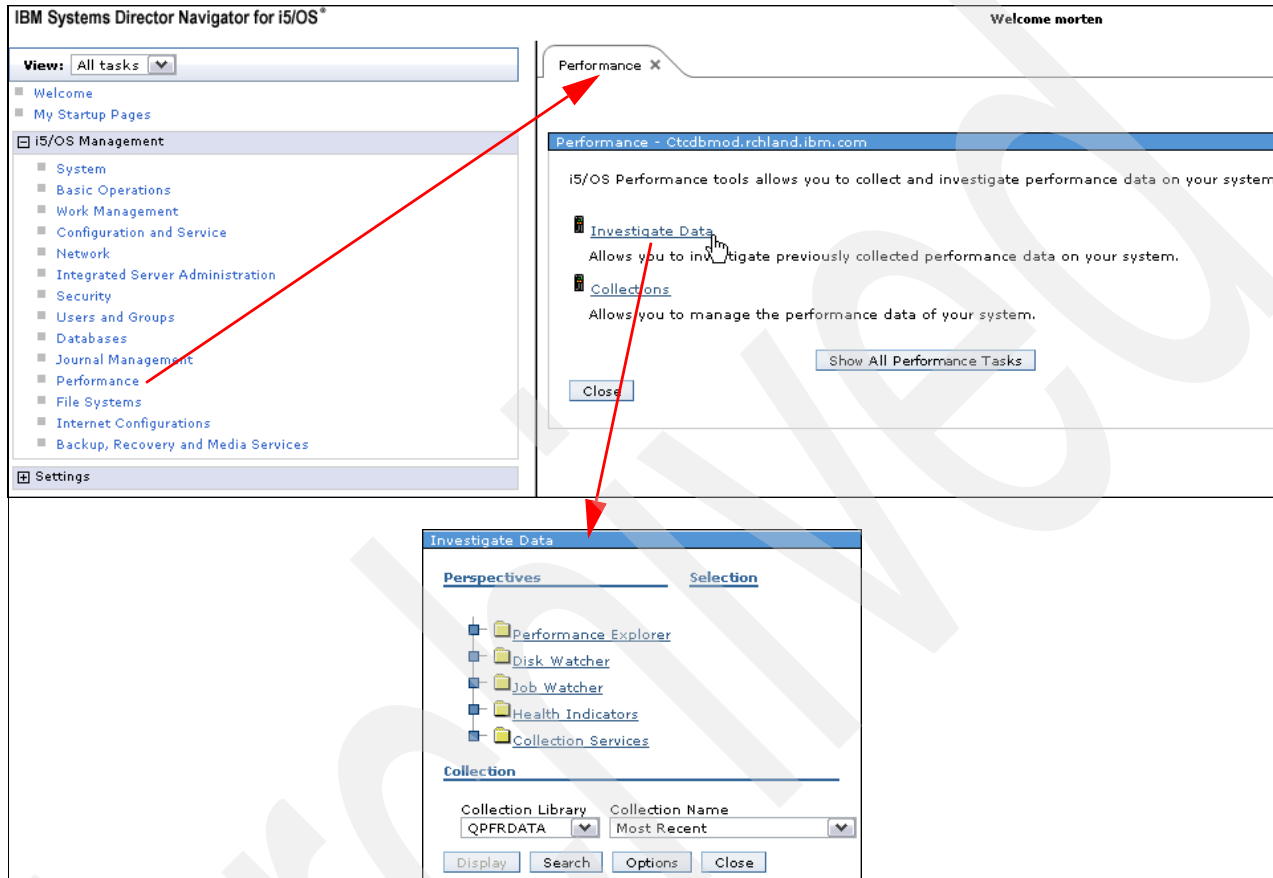


Figure 9-9 Investigate performance data

In the Investigate Data window, you will see a tree view of all collected performance data organized by perspective. Select a perspective by clicking one of the following links in the top level of the tree:

- ▶ Performance Explorer
- ▶ Disk Watcher
- ▶ Job Watcher
- ▶ Health Indicators
- ▶ Collection Services

Select which performance collection you want to use in the lower part of the window. The Display button will only be active if a collection exists for your perspective and collection. The perspective will expand to display several sub-categories under the top-level perspective, as shown in Figure 9-10. Figure 9-10 shows the Disk Watcher perspective as an example.

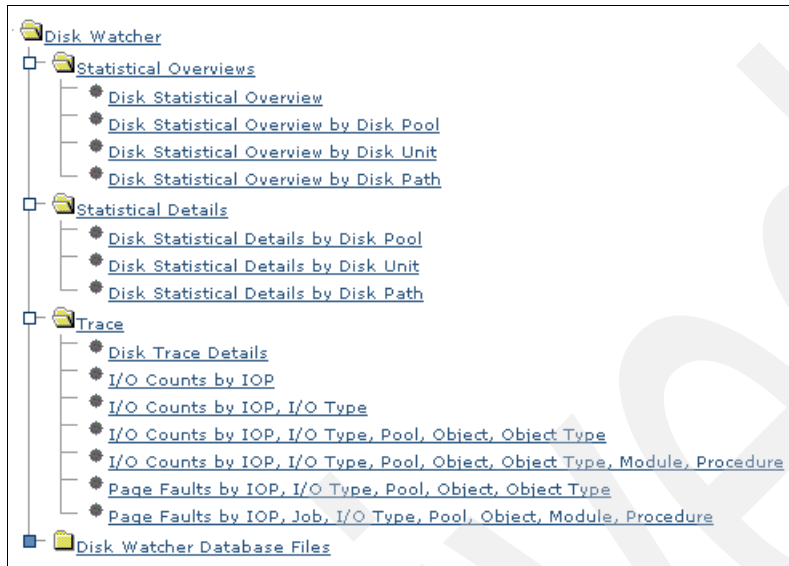


Figure 9-10 Disk Watcher overview

Similarly, with Job Watcher, you have many good graphs that you can view. Figure 9-11 shows an overview of the options.

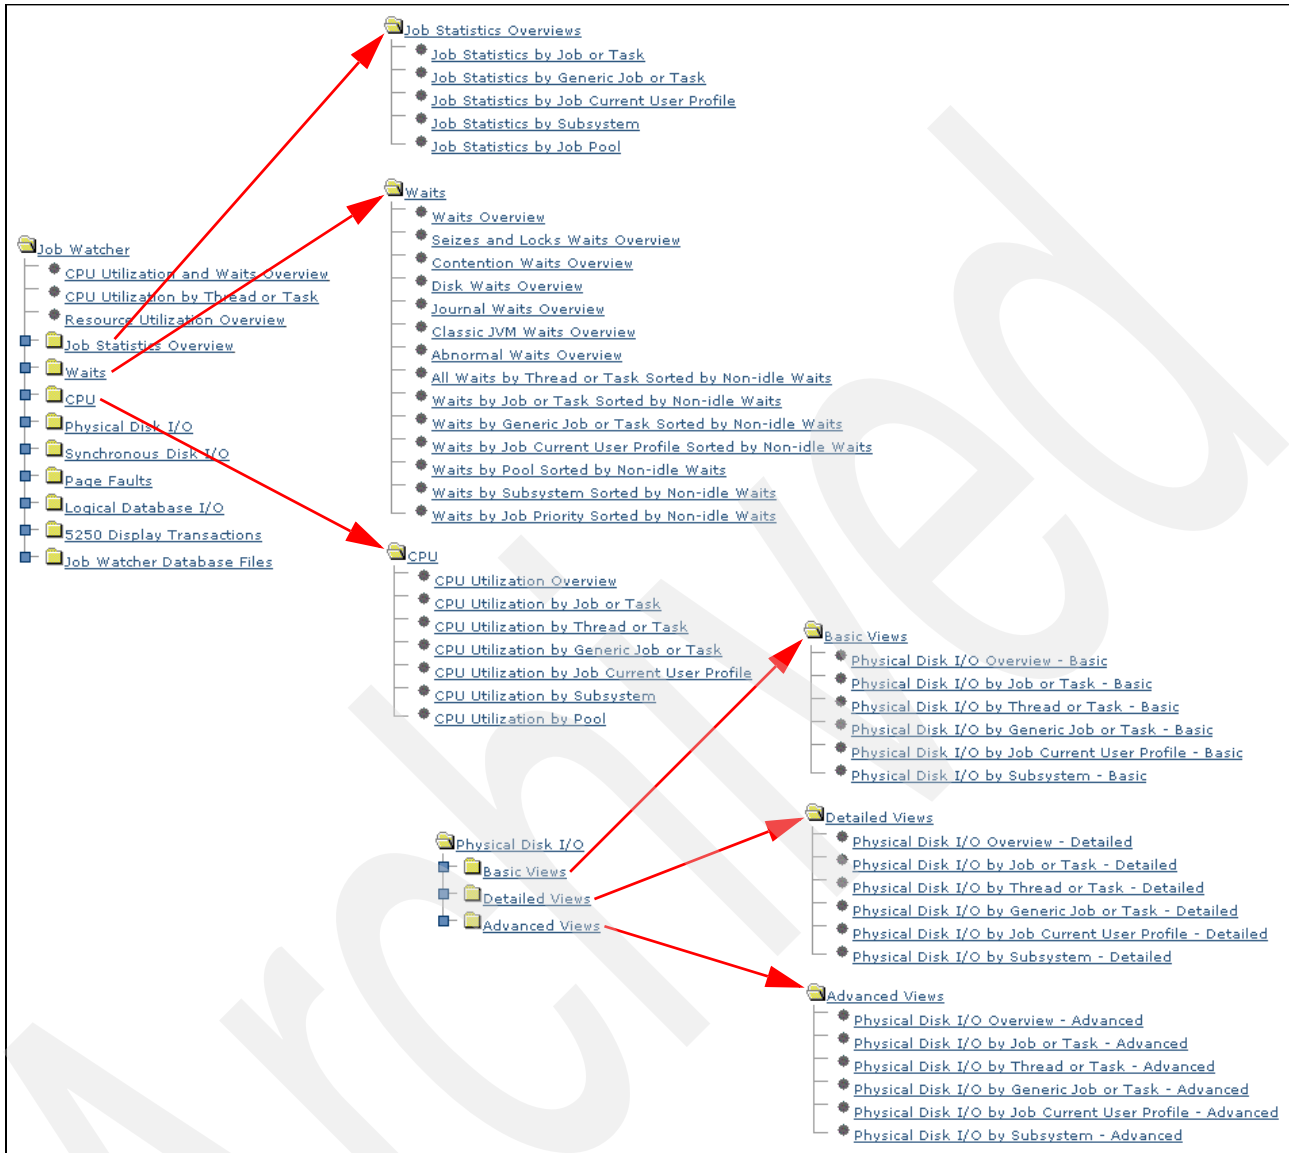


Figure 9-11 Job Watcher views (part 1)



The Job Watcher views continue in Figure 9-12.

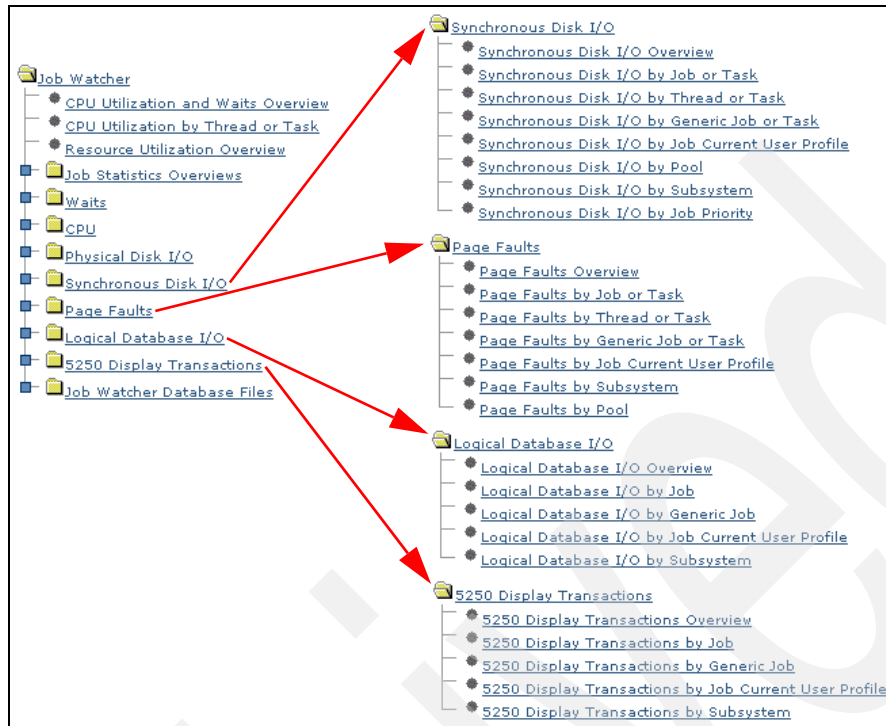


Figure 9-12 Job Watcher views (part 2)

Similar to Job Watcher, Collection Services also has a number of graphs that you can view. Figure 9-13 shows an overview.



Figure 9-13 Collection Services views

## 9.6 Health indicators

Health Indicators helps you to see an overview of a performance collection. Through this interface you can see whether your workload is within specified guidelines, as well as modify those guidelines for your system.

You have health indicators for the following areas:

- ▶ System resources
- ▶ CPU
- ▶ Disk
- ▶ Memory
- ▶ Response time

You can find health indicators through the Investigate Data link on the Performance main page, as shown in Figure 9-14.

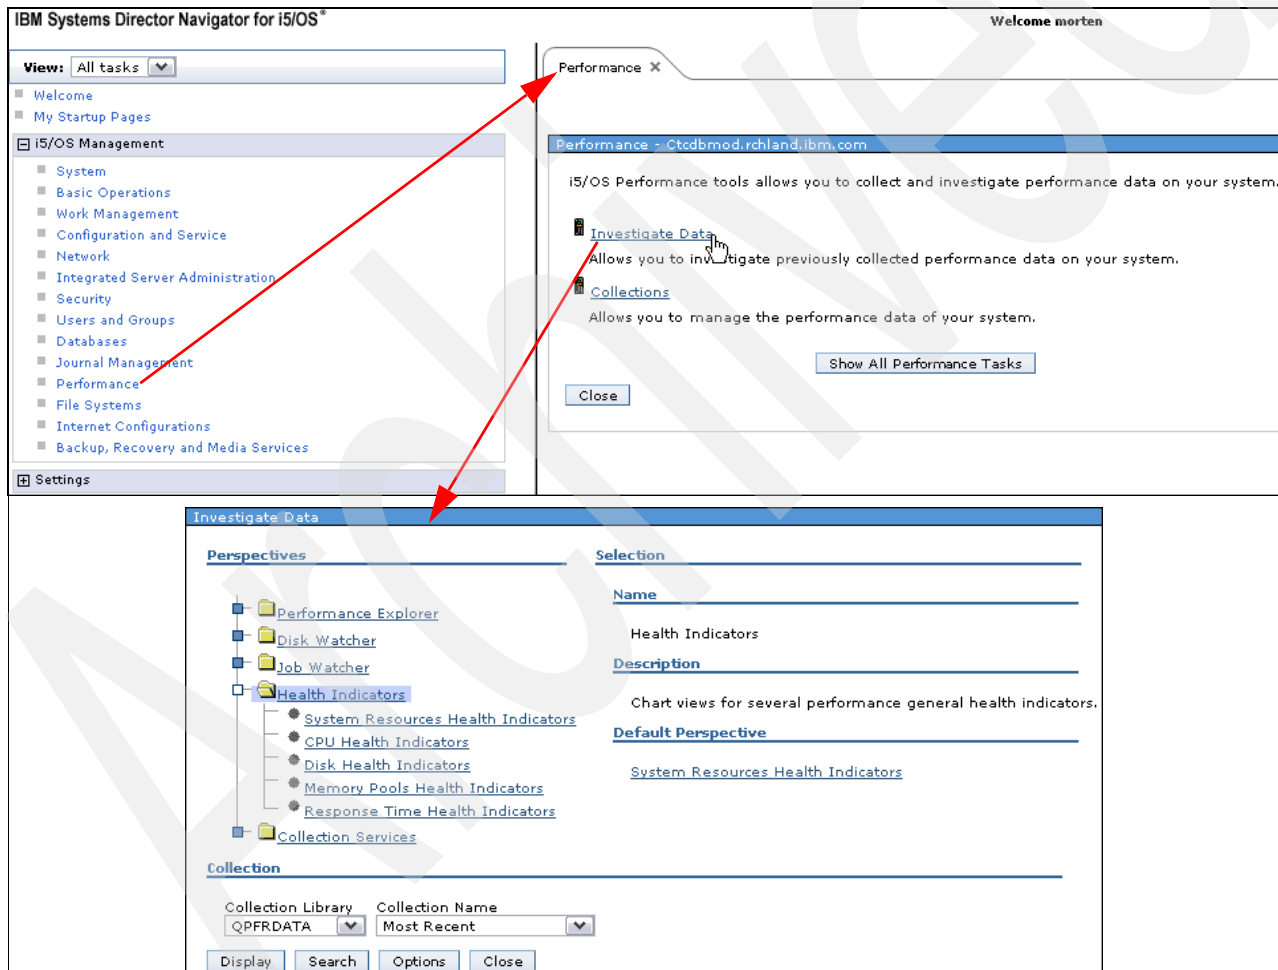


Figure 9-14 Health indicators

If you click the **Systems Resources Health Indicators** link, you will see indicators for the selected collection. In this case (Figure 9-14 on page 164), you see the most recent collection from the QPFRDATA. You will see what is in the guidelines and what is outside the guidelines for system resources, as shown in Figure 9-15.

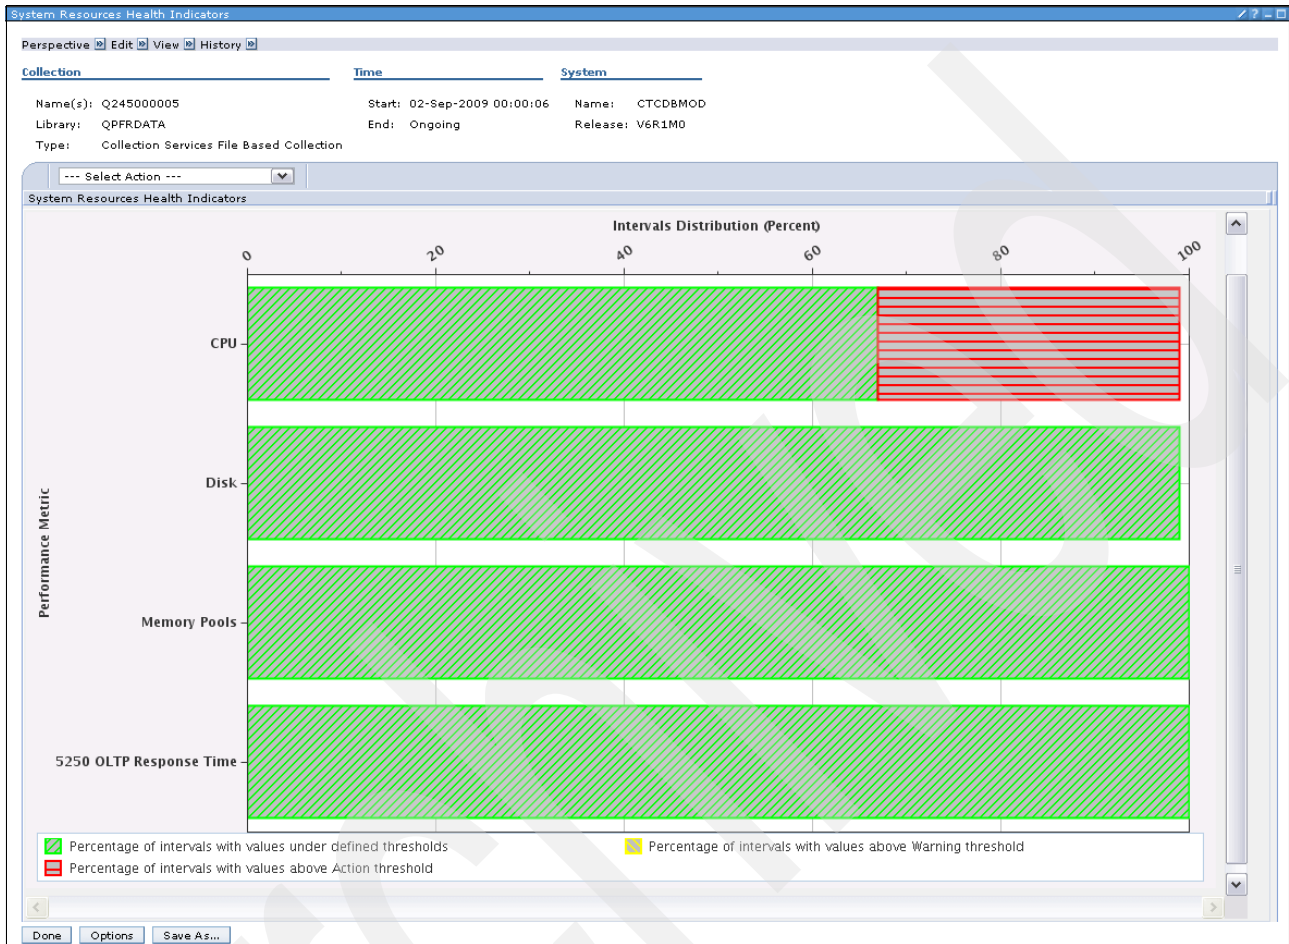


Figure 9-15 Health indicators for system resources

If you want to see or modify the guidelines, you can click the pull-down menu and select **Define Health Indicators**, as shown in Figure 9-16.

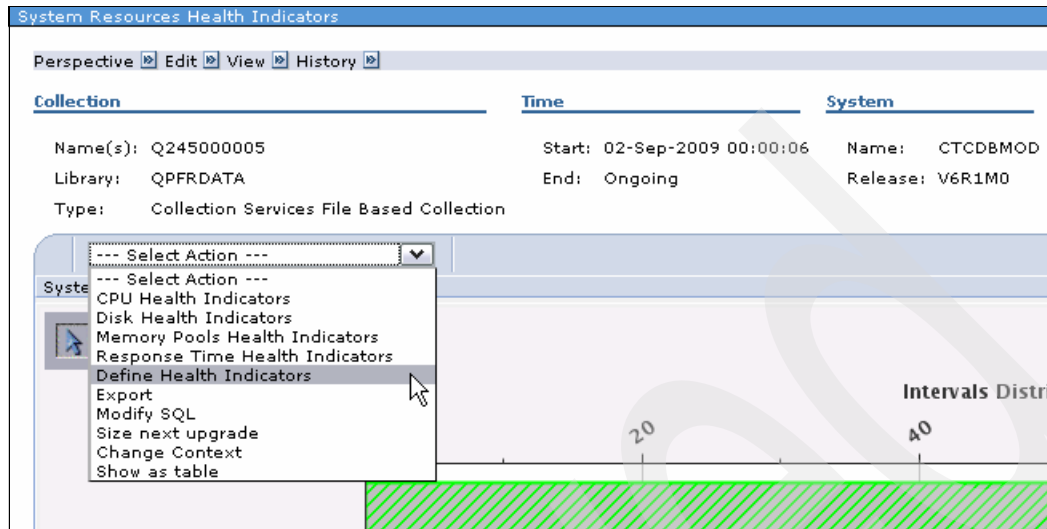


Figure 9-16 Define health indicators

You now have the option to define threshold values for a warning level and for an action level. Figure 9-17 shows an example of CPU health indicator.

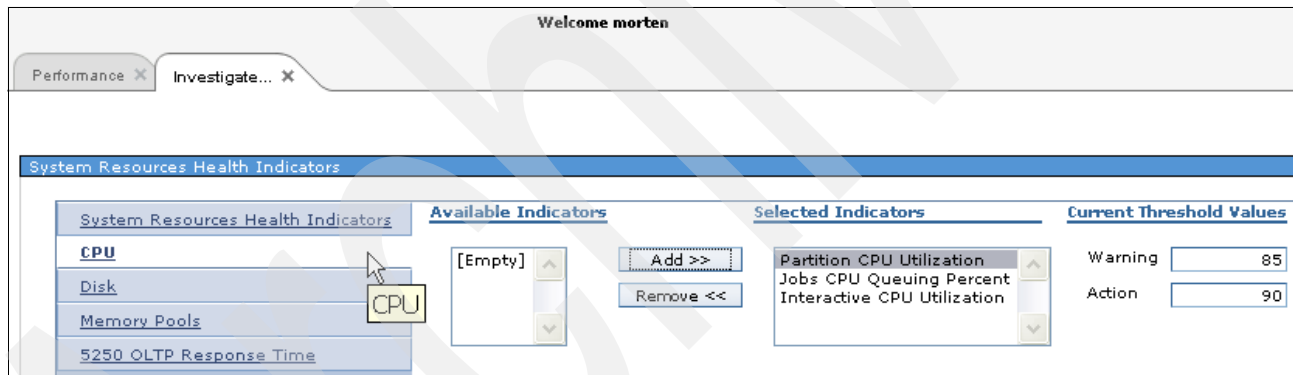


Figure 9-17 Health indicators: Set the CPU thresholds

You can also set thresholds for CPU, disk, memory, and 5250 response time.



## File systems

This chapter provides an overview of how Integrated File System (IFS) works and the new interfaces provided by IBM Systems Director Navigator for i.

## 10.1 Integrated File System

The Integrated File System is a part of the IBM i operating system that supports stream input/output and storage management similar to UNIX® operating systems and provides a similar hierarchical folder structure. Select **Integrated File System** from IBM Systems Director Navigator for i to display a list of all System i file systems to which you are authorized. You can access System i objects by specifying the path through the folder (Figure 10-1).

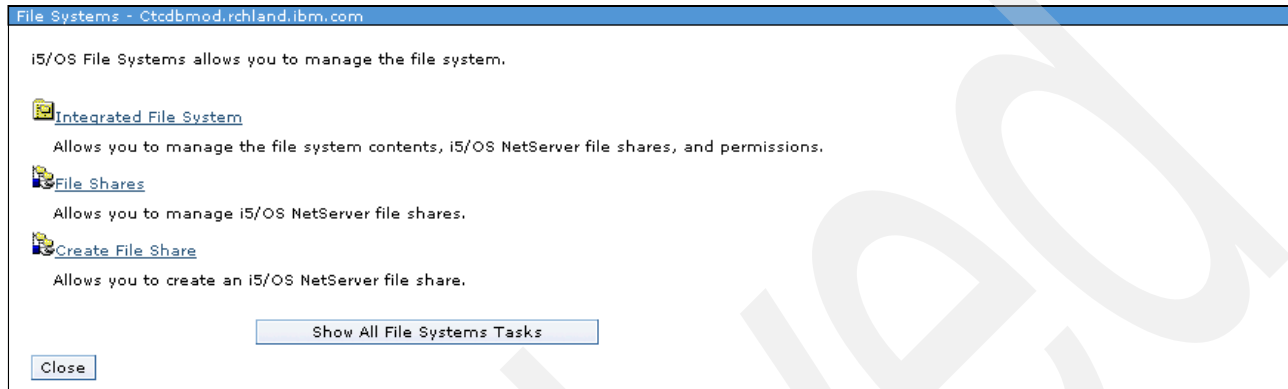


Figure 10-1 Integrate File System

### 10.1.1 What the Integrated file System is

The Integrated File System is a part of the IBM i operating system that supports stream input/output and storage management similar to UNIX operating systems, while providing an integrated structure over all information stored in your system.

The Integrated File System comprises 10 file systems, each with its own set of logical structures and rules for interacting with information in storage.

The key features of the Integrated File System are:

- ▶ Support for storing information in stream files that can contain long continuous strings of data. These strings of data might be, for example, the text of a document or the picture elements in a picture. The stream file support is designed for efficient use in the client/server applications.
- ▶ A hierarchical directory structure that allows objects to be organized like fruit on the branches of a tree. You can access an object by specifying the path through the directory to the object.
- ▶ A common interface that enables users and applications to access not only the stream files but also database files, documents, and other objects that are stored on your system.
- ▶ A common view of stream files that are stored locally on your system, on an Integrated xSeries Server (IXS), or on a remote Windows NT® server. Stream files can also be stored remotely on a local area network (LAN) server, on another remote System i product, or on a Network File System (NFS) server.

Figure 10-2 shows the Integrated File System menu.

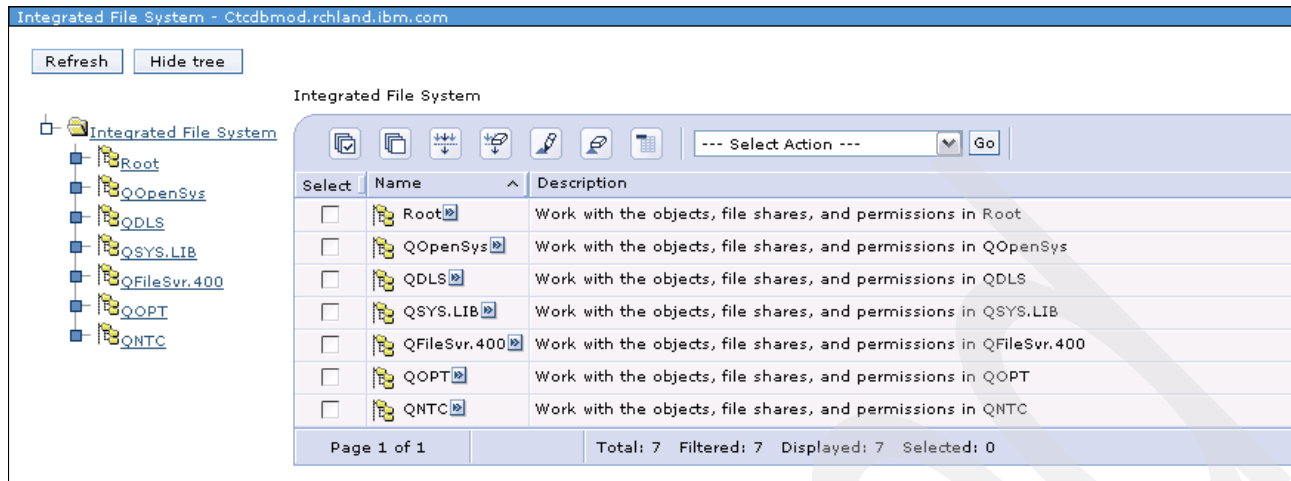


Figure 10-2 The Integrated File System menu

## 10.1.2 Root file system

The root (/) file system takes full advantage of the stream file support and hierarchical directory structure of the Integrated File System. It has the characteristics of the DOS and OS/2 file systems. In addition, it:

- ▶ Is optimized for stream file input and output
- ▶ Supports multiple hard links and symbolic links
- ▶ Supports local sockets
- ▶ Supports threadsafe APIs
- ▶ Supports \*FIFO objects
- ▶ Supports the /dev/null and /dev/zero \*CHRSF objects as well as other \*CHRSF objects
- ▶ Supports the journaling of object changes
- ▶ Supports the scanning of objects using the integrated file system scan-related exit points

### Accessing the root (/) file system

The root (/) file system can be accessed through the Integrated File System interface using either the IBM i file server or the integrated file system commands, user displays, and APIs.

When you click this section, you see the window shown in Figure 10-3.

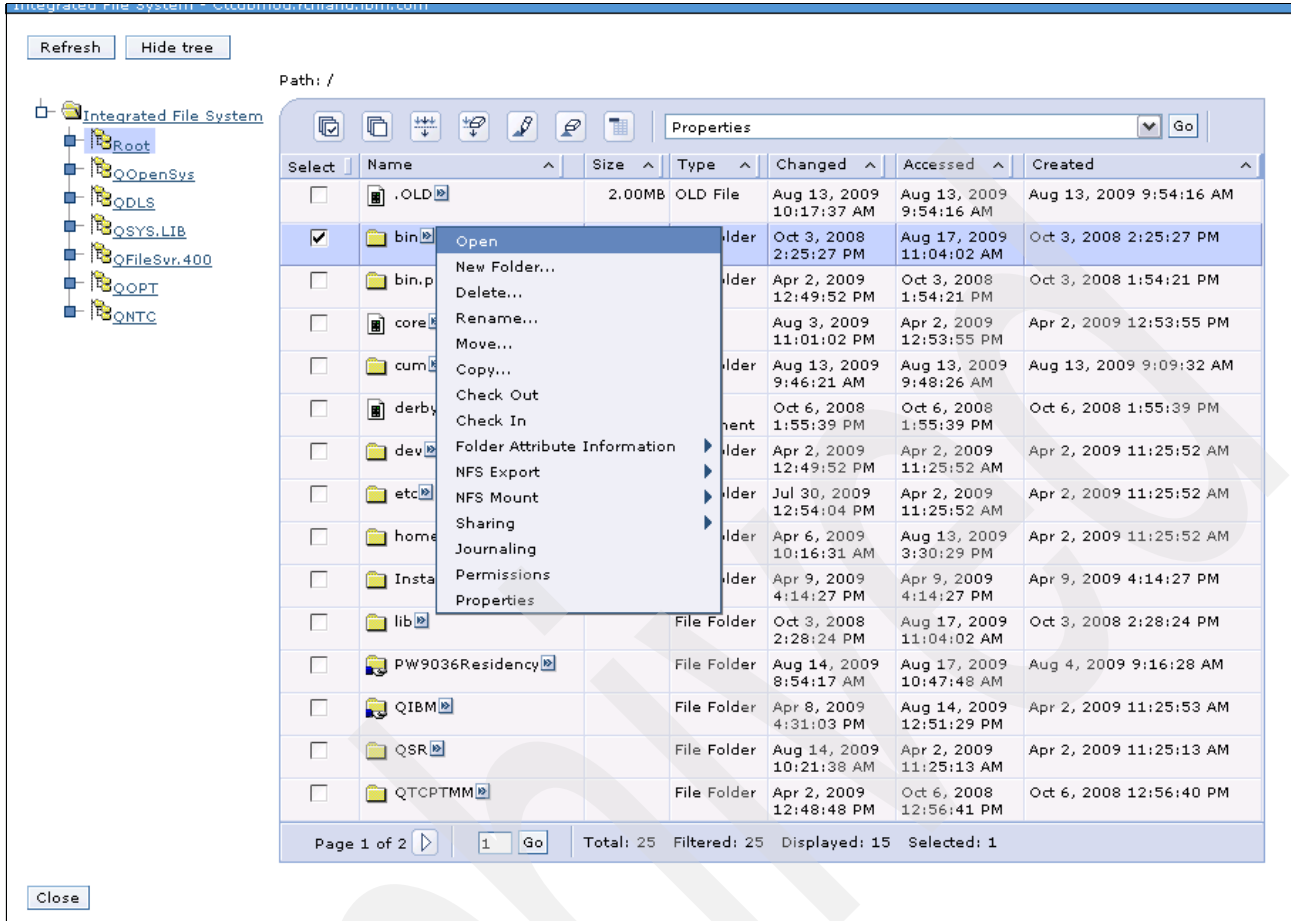


Figure 10-3 Root section selected



In Figure 10-3 on page 170 we selected a the **bin** folder and expanded its pop-up menu. You can select any of the specific functions. Also, you can use the drop-down list to select other functions, as shown in Figure 10-4.

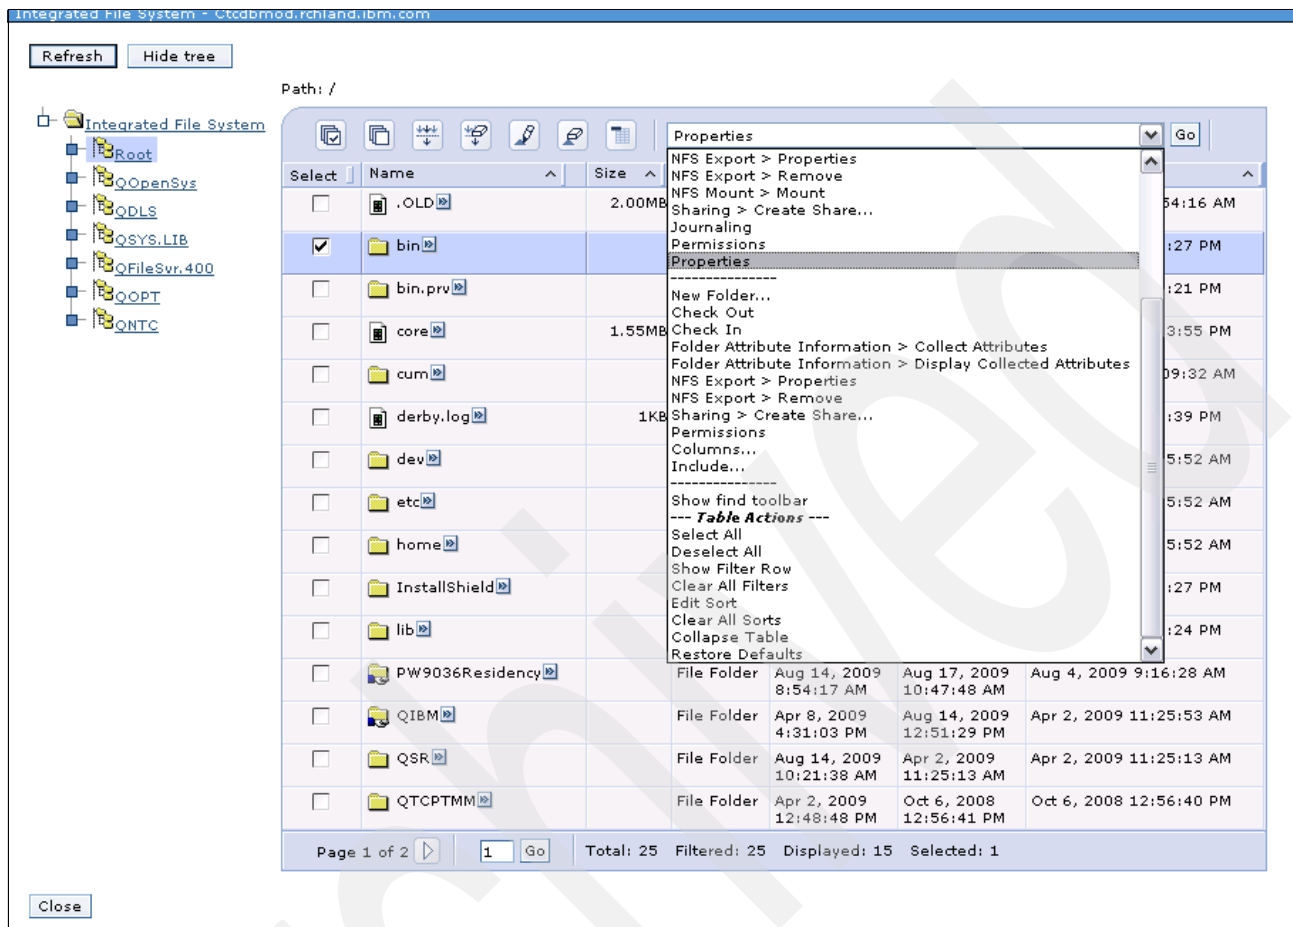


Figure 10-4 Drop-down list selected

**Note:** For further information see the following Web site:  
<http://publib.boulder.ibm.com/infocenter/iseres/v6r1m0/index.jsp>

### 10.1.3 QopenSys file system

The QOpenSys file system is compatible with open system standards based on UNIX, such as POSIX and X/Open Portability Guide (XPG). Like the root (/) file system, this file system takes advantage of the stream file and directory support that is provided by the Integrated File System. In addition, it:

- ▶ Is accessed through a hierarchical directory structure similar to UNIX systems
- ▶ Is optimized for stream file input and output
- ▶ Supports multiple hard links and symbolic links
- ▶ Supports case-sensitive names
- ▶ Supports local sockets
- ▶ Supports threadsafe APIs
- ▶ Supports FIFO objects

- ▶ Supports the journaling of object changes
- ▶ Supports the scanning of objects using the integrated file system scan-related exit points

The QOpenSys file system has the same characteristics as the root (/) file system, except that it is case sensitive to enable support for UNIX-based open systems standards.

### Accessing QOpenSys

QOpenSys can be accessed through the Integrated File System interface using either the IBM i file server or the Integrated File System commands, user displays, and APIs. When you click this section, you see the window shown in Figure 10-5.

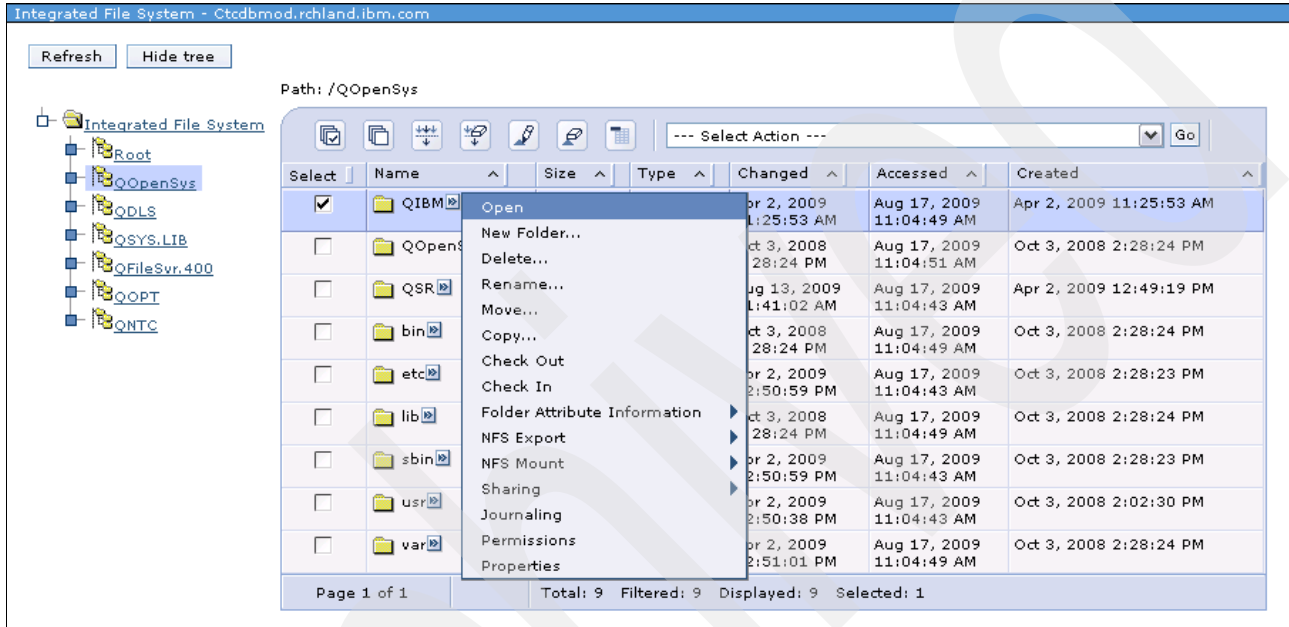


Figure 10-5 QIBM folder selected

In this example we selected a folder called QIBM and expanded its pop-up menu. You can select any of the specific functions. Also, you can use the drop-down list to select other functions, as shown in Figure 10-6.

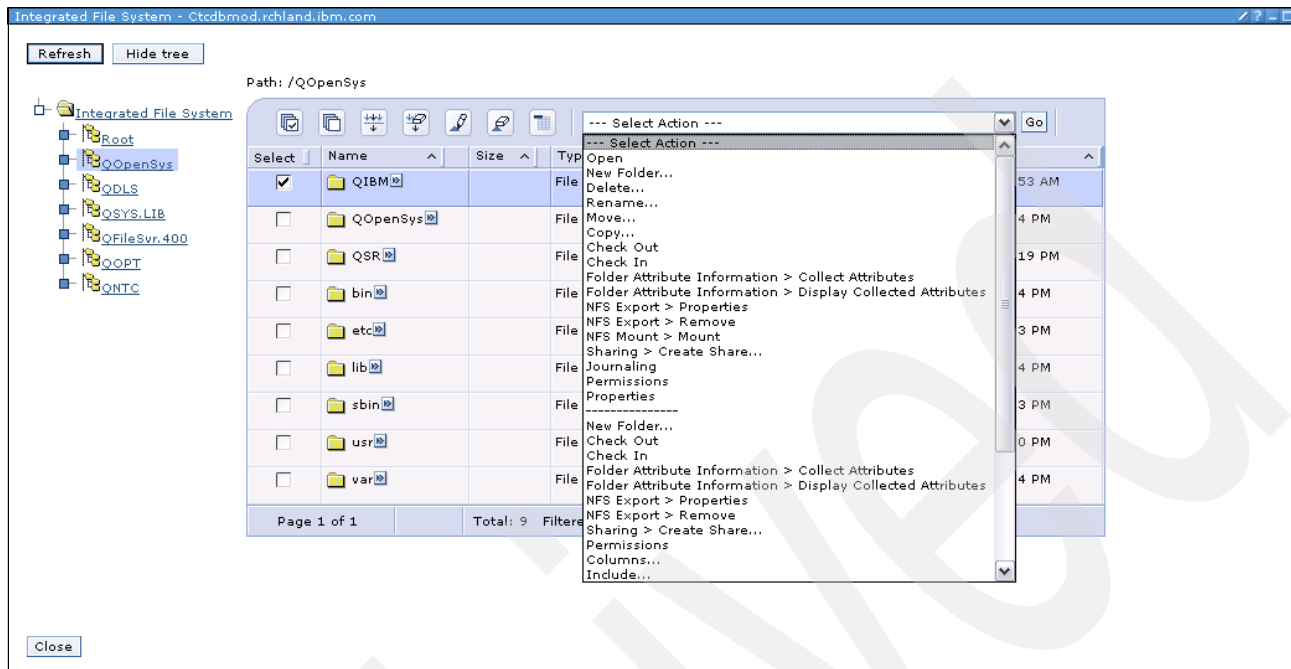


Figure 10-6 Drop-down list selected

**Note:** For further information see the following Web site:

<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>

### 10.1.4 QDLS file system

The QDLS file system supports the folders structure. It provides access to documents and folders. In addition, it:

- ▶ Supports IBM i folders and document library objects (DLOs)
- ▶ Supports data stored in stream files

#### Accessing QDLS through the Integrated File System interface

The QDLS file system can be accessed through the Integrated File System interface using either the IBM i file server or the Integrated File System commands, user displays, and APIs.

When you click this section you see the window shown in Figure 10-7.

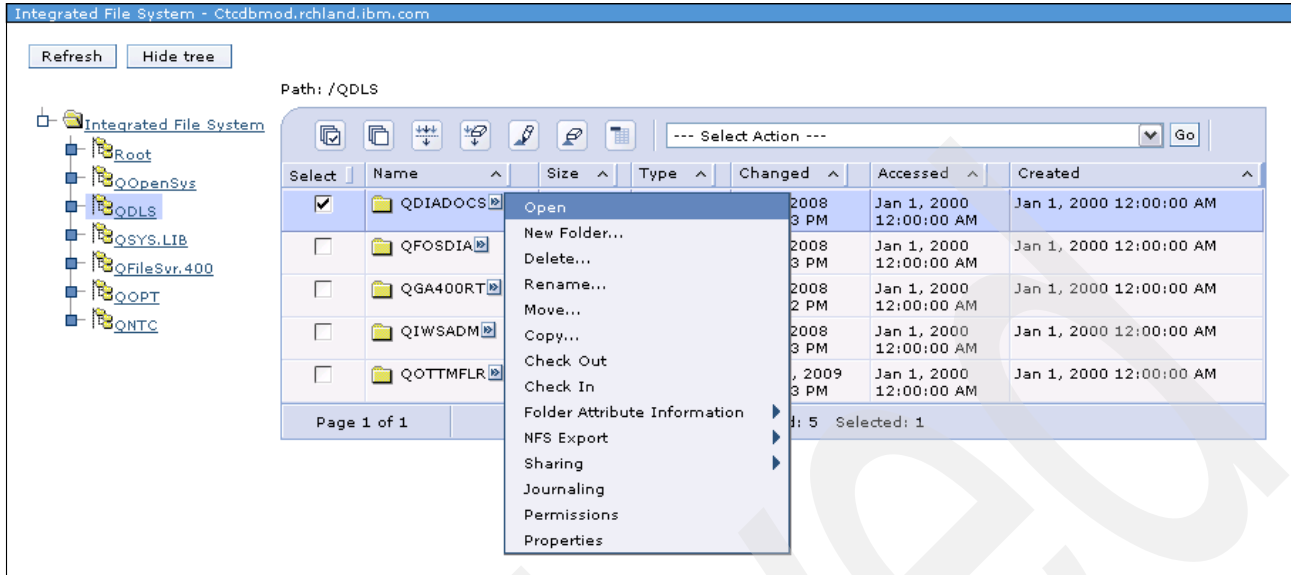


Figure 10-7 QDLS section selected

In this example we selected a folder QDIADOC and expanded its pop-up menu. You can select any of the specific functions. Also, you can use the drop-down list to select other functions, as shown in Figure 10-8.

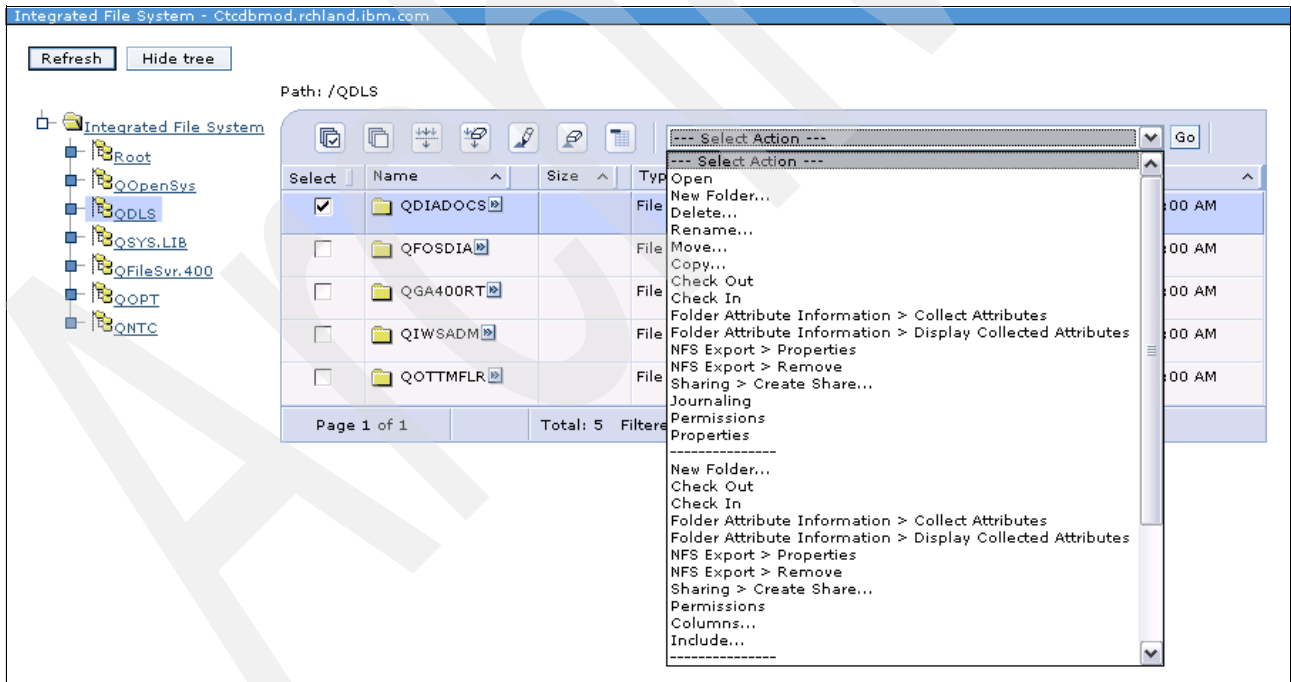


Figure 10-8 Drop-down list selected

**Note:** For further information see the following Web site:

<http://publib.boulder.ibm.com/infocenter/series/v6r1m0/index.jsp>

## 10.1.5 QSYS.LIB file system

The QSYS.LIB file system supports the IBM i library structure. This file system provides you with access to database files and all of the other IBM i object types that the library support manages on the system and in the basic user auxiliary storage pools (ASPs). In addition, it:

- ▶ Supports all user interfaces and programming interfaces that operate on IBM i libraries and objects in those libraries
- ▶ Supports all programming languages and facilities that operate on database files
- ▶ Provides extensive administrative support for managing IBM i objects
- ▶ Supports stream I/O operations on physical file members, user spaces, and save files

Before the Integrated File System was introduced in Version 3 of OS/400, the QSYS.LIB file system was the only file system. Programmers who used languages, such as RPG or COBOL, and facilities, such as DDS, to develop applications were using the QSYS.LIB file system. System operators who used commands, menus, and displays to manipulate output queues were using the QSYS.LIB file system, as were system administrators who were creating and changing user profiles.

All of these facilities and the applications based on these facilities work as they did before the introduction of the Integrated File System. These facilities cannot, however, access QSYS.LIB through the Integrated File System interface.

### **Accessing QSYS.LIB through the Integrated File System interface**

The QSYS.LIB file system can be accessed through the Integrated File System interface using either the IBM i file server or the Integrated File System commands, user displays, and APIs.

When you click this section, you see the window shown in Figure 10-9.

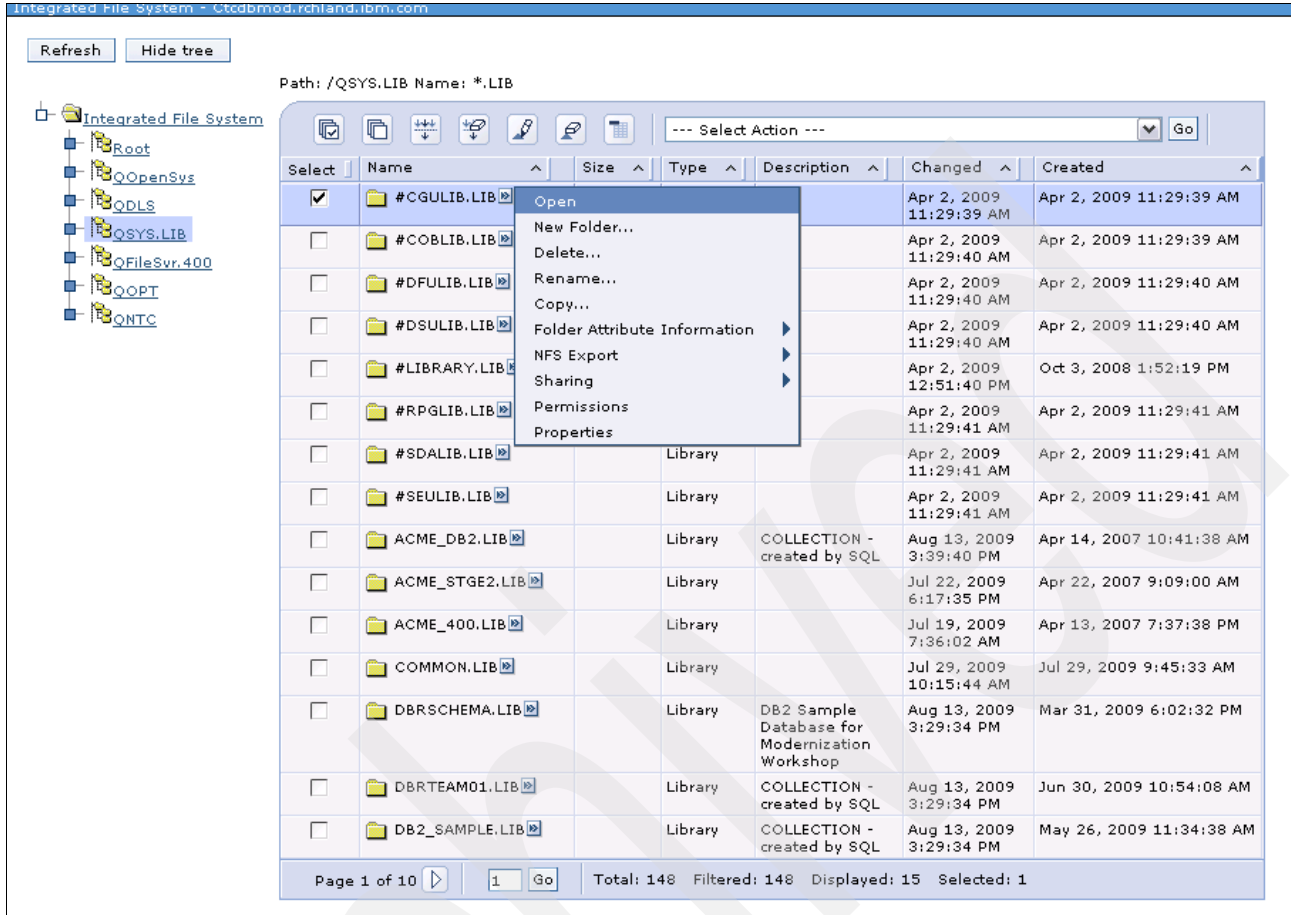


Figure 10-9 QSYS.LIB section selected

In this example we selected a folder #CGULIB.LIB and expanded its pop-up menu. You can select any of the specific functions. Also, you can use the drop-down list to select other functions, as shown in Figure 10-10.

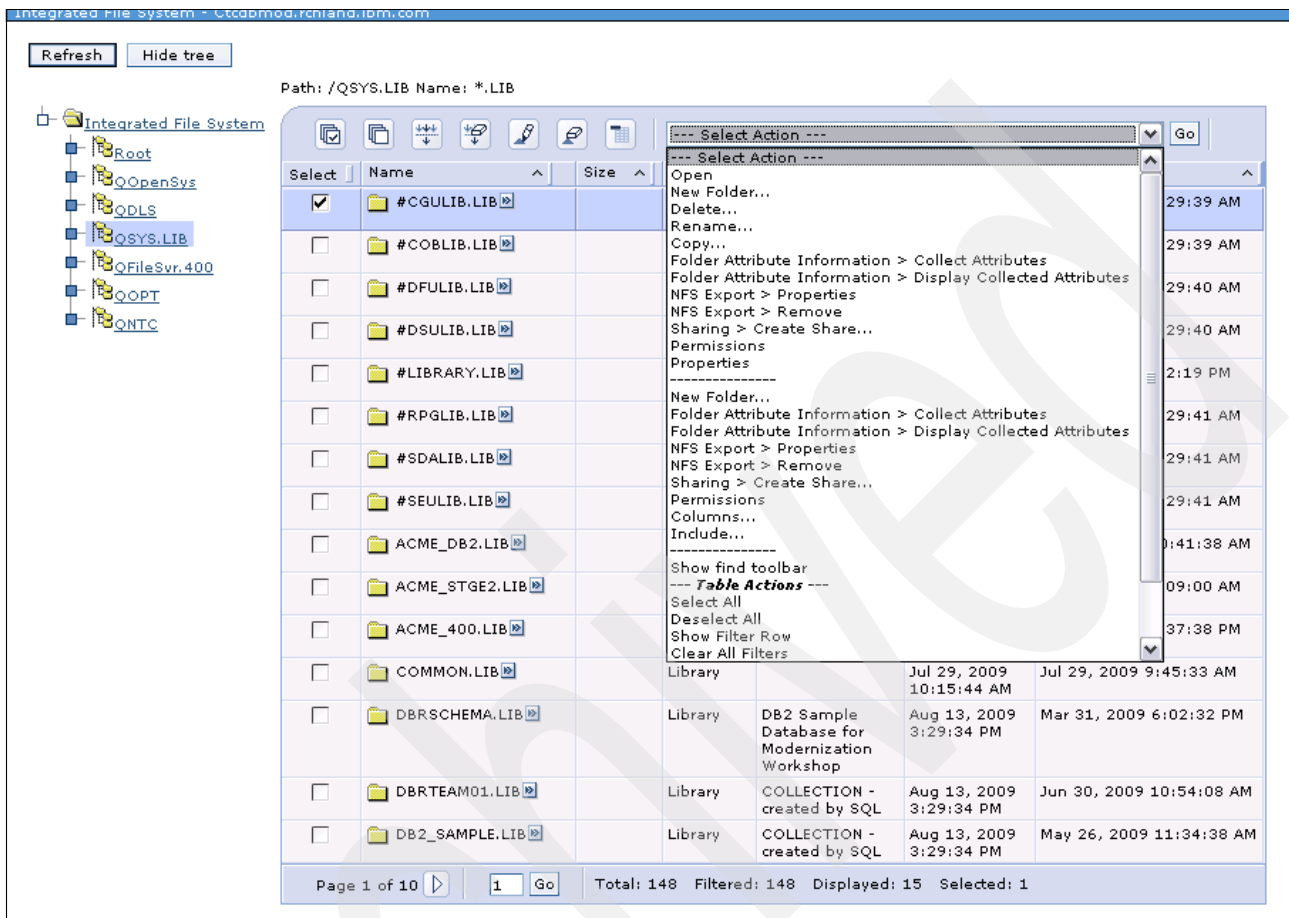


Figure 10-10 Drop-down list selected

**Note:** For further information see the following Web site:  
<http://publib.boulder.ibm.com/infocenter/iseres/v6r1m0>

### 10.1.6 QFileSvr.400 file system

The QFileSvr.400 file system provides transparent access to other file systems that reside on remote System i platforms. It is accessed through a hierarchical directory structure.

The QFileSvr.400 file system can be thought of as a client that acts on behalf of users to perform file requests. QFileSvr.400 interacts with the IBM i file server on the target system to perform the actual file operation.

#### Accessing QFileSvr.400 through the Integrated File System interface

The QFileSvr.400 file system can be accessed through the Integrated File System interface using either the IBM i file server or the Integrated File System commands, user displays, and APIs.

## 10.1.7 QOPT file system

The QOPT file system provides access to stream data that is stored on optical media. It:

- ▶ Provides a hierarchical directory structure similar to PC operating systems such as DOS and OS/2
- ▶ Is optimized for stream file input and output
- ▶ Supports data stored in stream files

### Accessing QOPT through the Integrated File System

The QOPT file system can be accessed through the Integrated File System using either the IBM i file server or the Integrated File System commands, user displays, and APIs. When you click this section, you see the window shown in Figure 10-11.

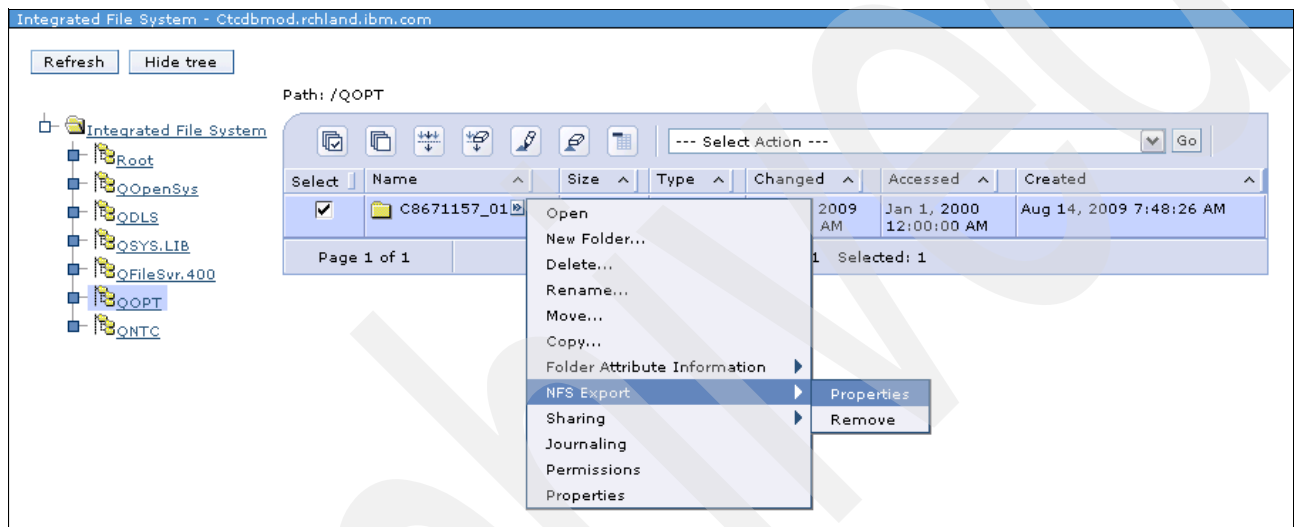


Figure 10-11 QOPT section selected



In this example we selected a folder C8671157\_01 .LIB and expanded its pop-up menu. You can select any of the specific functions. Also, you can use the drop-down list to select other functions, as shown in Figure 10-12.

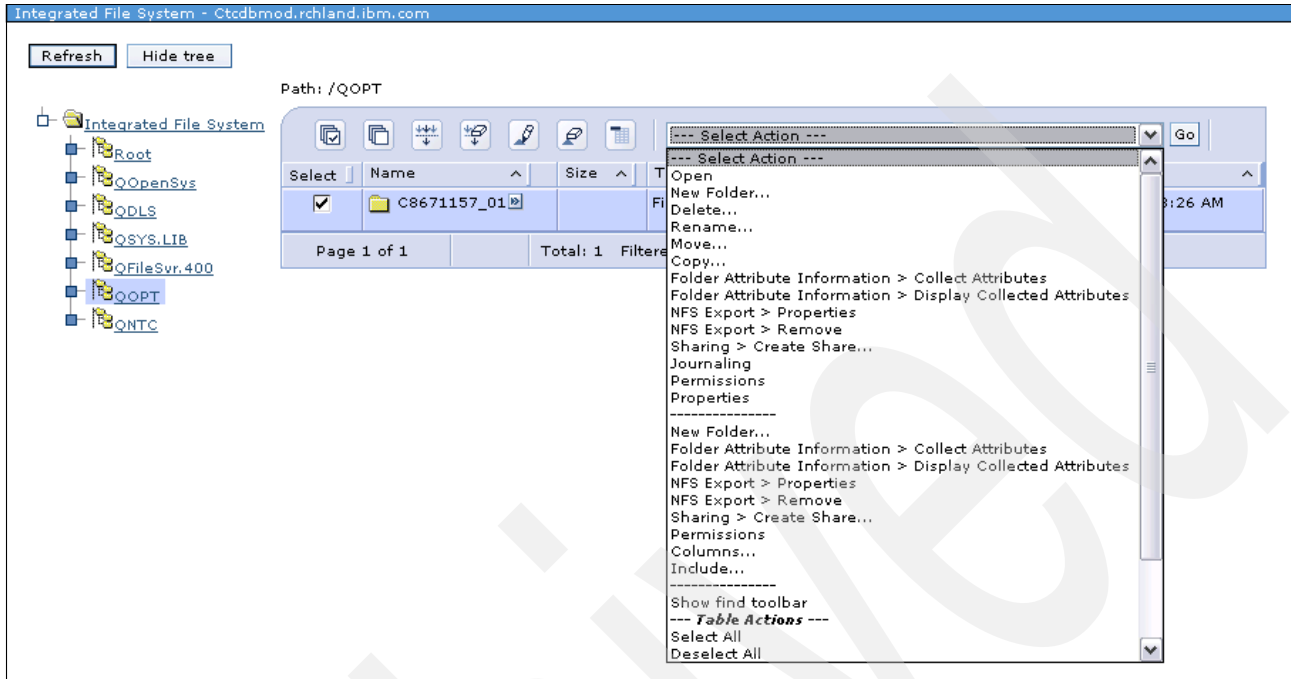


Figure 10-12 Drop-down list selected

**Note:** For further information see this the following Web site:

<http://publib.boulder.ibm.com/infocenter/series/v6r1m0/index.jsp>

### 10.1.8 QNTC file system

The QNTC file system provides access to data and objects that are stored on an Integrated xSeries Server (IXS) running Windows NT 4.0 or later or the Linux operating system.

The QNTC file system also provides access to data and objects that are stored on remote servers running Windows NT 4.0 or later, Linux Samba 3.0 or later, or supported versions of i5/OS NetServer.

The QNTC file system is part of the base IBM i operating system. It is not necessary to have the Integrated Server Support, option 29 of the operating system, installed to access /QNTC.

#### Accessing QNTC through the Integrated File System interface

By using the IBM i NetServer, System i Navigator, Integrated File System commands, user displays, or APIs, you can access the QNTC file system through the Integrated File System interface.

## 10.2 File shares

An IBM i NetServer file share is a directory path that IBM i NetServer shares with clients on the network. A file share can consist of any Integrated File System directory on the IBM i operating system. You can create, display, configure, and end IBM i NetServer file shares. In general, all Integrated File System limitations and considerations apply when you access shared directories with IBM i NetServer. Starting in IBM i 6.1, IBM i NetServer supports IPv6 access in the file sharing function. Figure 10-13 show the file shares functions.

Select	Name	Path	Description	Current Users	Access	Max Users
<input type="checkbox"/>	Qdirsrv	/QIBM/ProdData/OS400/DirSrv	OS/400 -- Directory Services	0	Read/Write	2147483647
<input type="checkbox"/>	Qibm	/QIBM	IBM Product Directories	1	Read/Write	No maximum
<input type="checkbox"/>	Pw9036	/PW9036Residency	PW9036 Redbook Residency File Share	3	Read/Write	No maximum

Page 1 of 1      Total: 3   Filtered: 3   Displayed: 3   Selected: 0

The information in this list is 7 minutes old. Press the Refresh button to update.

Figure 10-13 File Shares page

We can select a file and expand its pop-up menu, as shown in Figure 10-14.

Select	Name	Path	Description	Current Users	Access	Max Users
<input checked="" type="checkbox"/>	Qdirsrv	/QIBM/ProdData/OS400/DirSrv	OS/400 -- Directory Services	0	Read/Write	2147483647
<input type="checkbox"/>	Qibm	/QIBM	IBM Product Directories	1	Read/Write	No maximum
<input type="checkbox"/>	Pw9036	/PW9036Residency	PW9036 Redbook Residency File Share	3	Read/Write	No maximum

Page 1 of 1      Total: 3   Filtered: 3   Displayed: 3   Selected: 1

Figure 10-14 Function options

## Overview of functions for managing file shares

As shown Figure 10-14 on page 180, you have four options to choose from:

- ▶ **New:** This option allows you to create a new file share based on an existing share. Figure 10-15 shows the new file share window.

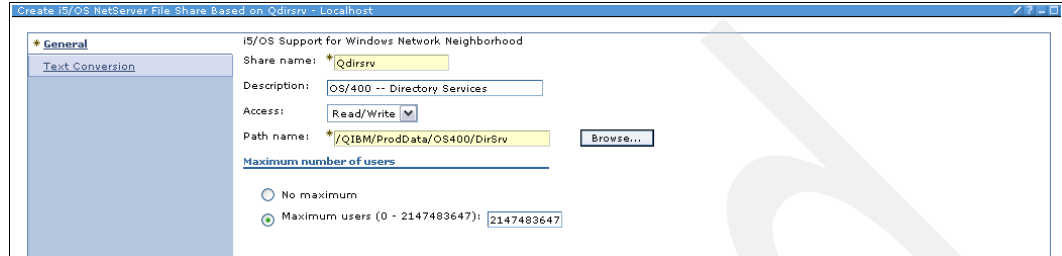


Figure 10-15 Create file share

As Figure 10-15 shows, you can set up the share name, the access (read/write or read only), the path name, and the number of users who can use it.

- ▶ **Properties:** Use the General Properties page to configure the new file share with a name, description, access, maximum number of users, and directory path name. See Figure 10-16.

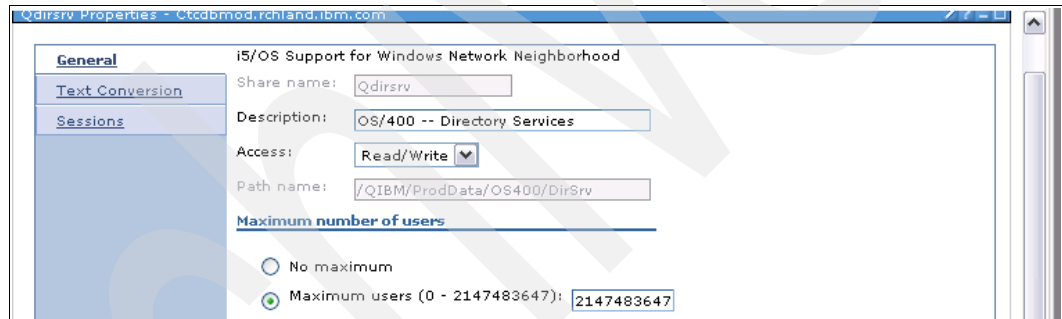


Figure 10-16 General Properties

Use the Text Conversion page to identify which file types must have their contents converted from the IBM i file coded character set ID (CCSID) to the CCSID that you specify for the share. See Figure 10-17.

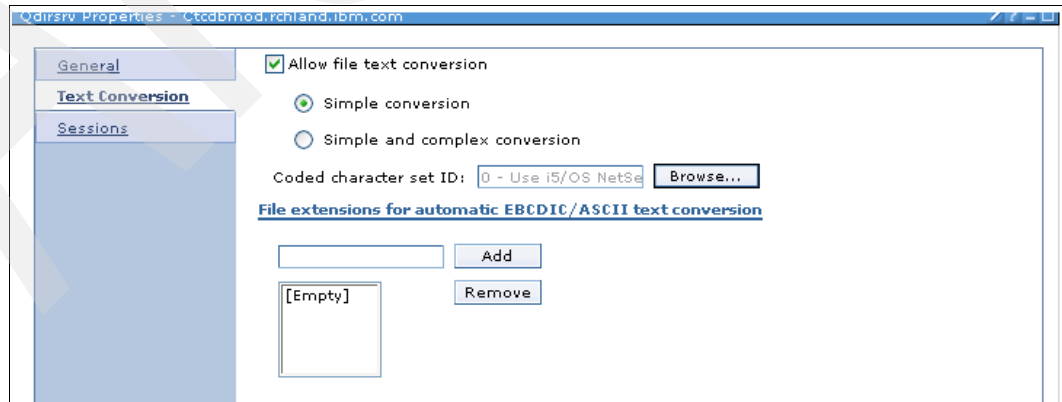


Figure 10-17 Text conversion

- ▶ **Session:** Use this option to see the data for the jobs using the file shares. See Figure 10-18.




Figure 10-18 Session

### 10.3 Create file share

To create a new file share you must select the **Create File Share** task, as shown in Figure 10-19.

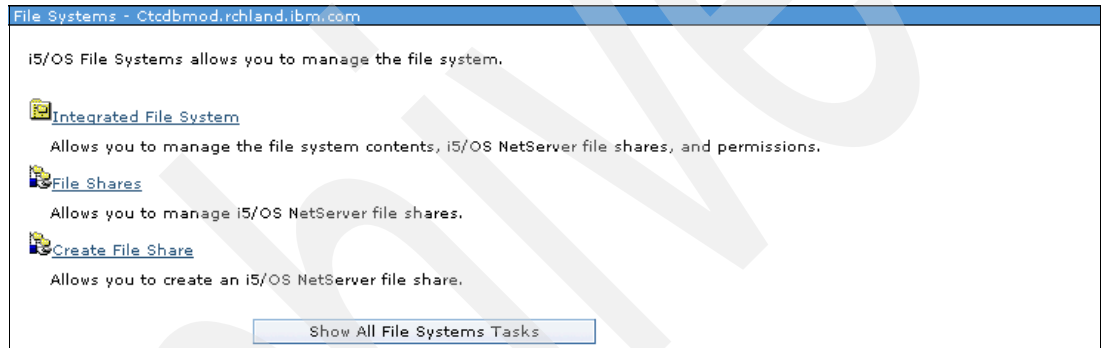


Figure 10-19 File Systems task page

When you select this option you see the window shown in Figure 10-20.

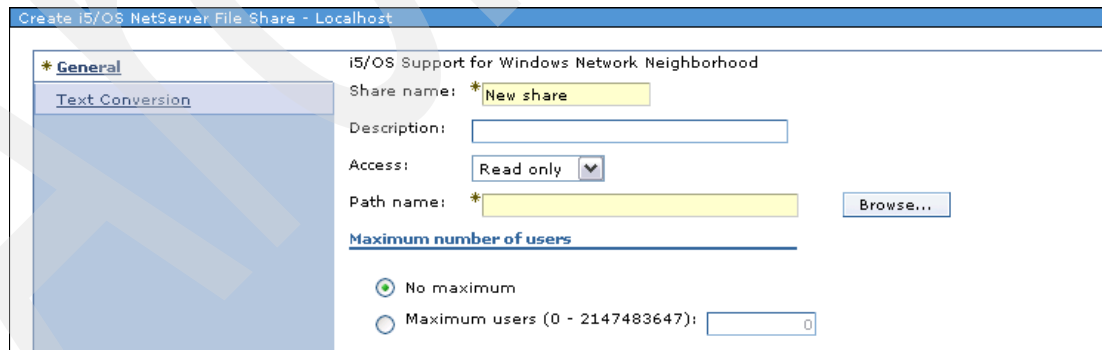


Figure 10-20 Create new file share

Use the General properties page to configure the new file share with a name, description, access, maximum number of users, and directory path name. In this case you cannot base the new share on an existing share.

When you select the Text Conversion option you see the window shown in Figure 10-21.

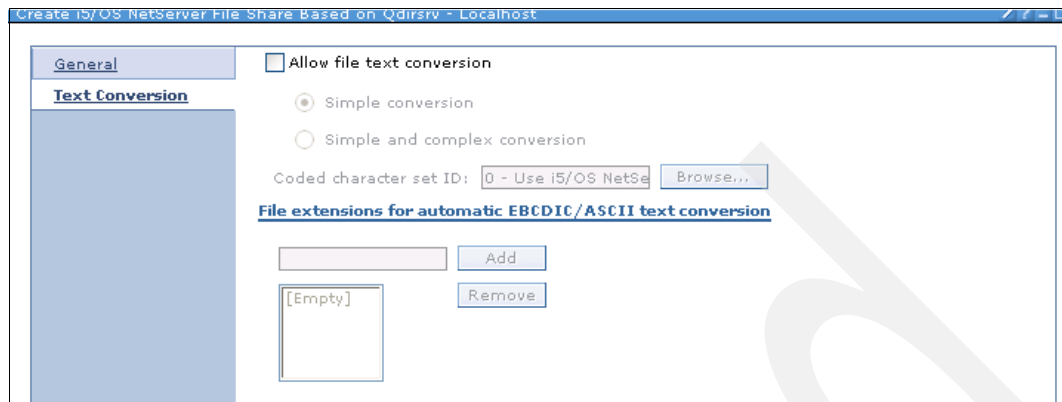


Figure 10-21 Text Conversion option

**Note:** Before selecting this option you must complete the required fields that appear on the General properties page.

Archived



## Advanced Job Scheduler for IBM i

From within Systems Director Navigator for i5/OS, the Advanced Job Scheduler for IBM i (AJS) can be used to schedule complex batch job flow, distribute spooled files via e-mail attachment, and send job status notifications. Jobs can run based on a schedule or an event. Jobs can run on local or remote IBM i systems and commands can be processed on non-IBM i platforms.

This chapter describes the following topics:

- ▶ Accessing AJS from within Systems Director Navigator for i5/OS
- ▶ Scheduling a job
- ▶ Job groups
- ▶ Job dependencies
- ▶ Notification and report distribution
- ▶ Remote command jobs
- ▶ Multiple scheduling environments

# 11.1 Accessing AJS from Systems Director Navigator

To access AJS from Systems Director Navigator, do the following:

1. Expand **IBM i Management** on the left from Systems Director Navigator for i5/OS.
1. Select **Work Management**.
1. At the bottom of the Work Management page, click **Show all the Work Management tasks**.
1. Each of the AJS folders contains tasks that can be used to list or perform functions in the scheduler. Figure 11-1 shows the AJS folders within all the work management tasks. Each folder has options in its drop-down menu to display a list or perform a task.

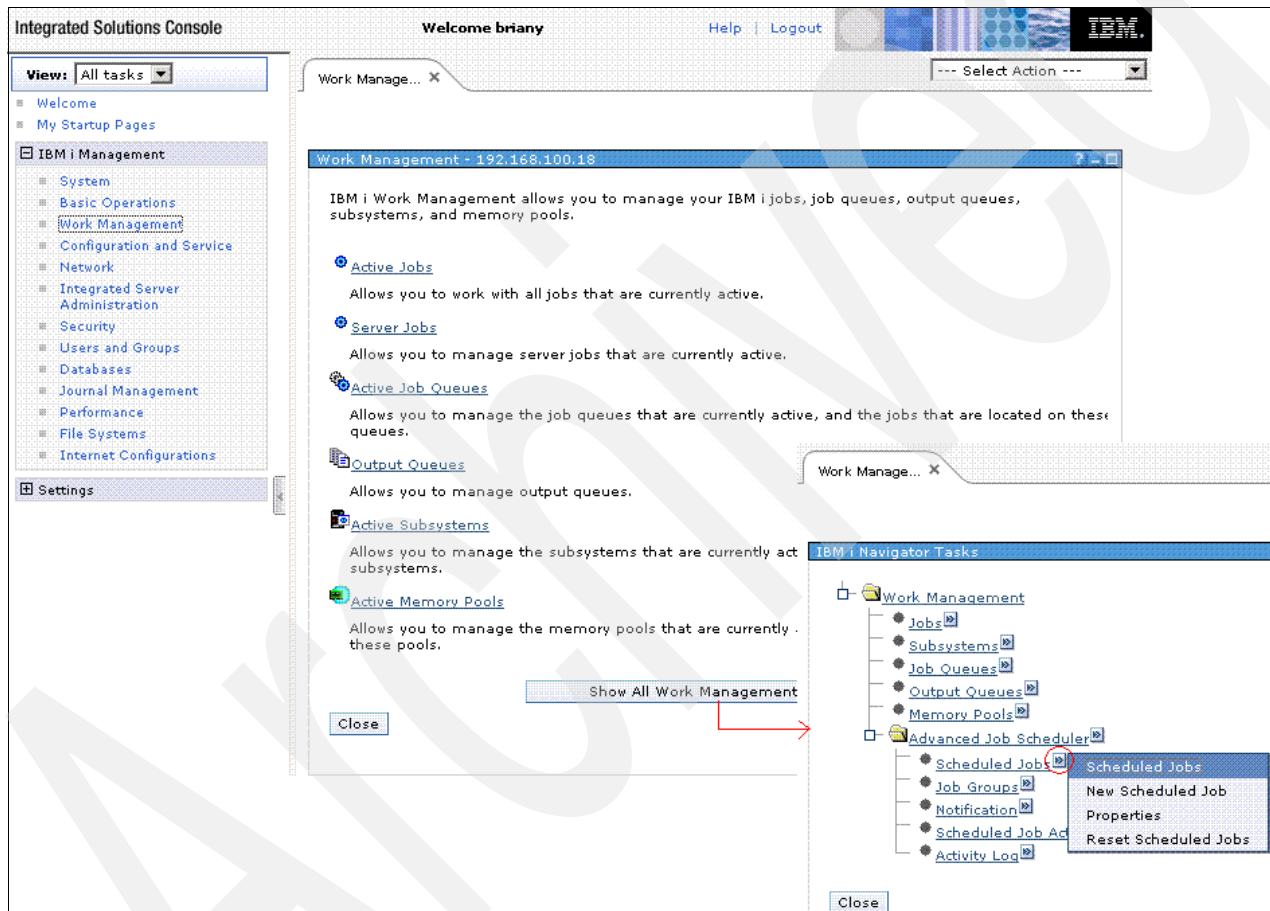


Figure 11-1 Accessing Advanced Job Scheduler for IBM i

## 11.2 Scheduling a job

Using the Web interface for AJS from IBM Systems Director Navigator for i, you can schedule jobs with a simple schedule such as every week Monday through Friday or more complex schedules such as the seventh working day every month. If the standard scheduling conventions within AJS are not enough, you can create a scheduling calendar that contains all the dates that a job is to run. Holiday calendars can be used for holidays or exceptions to the normal scheduling. These are the dates that you do not want a job to run. These calendars can be used in multiple jobs.



In this section we discuss three main areas in creating a scheduled job:

- ▶ The work
- ▶ The schedule
- ▶ The notification

## 11.2.1 The work

Each scheduled job can run up to 999 commands. See Figure 11-2. If any of the commands end in error, the job will end be marked as abnormally completed. Each command can contain a list of messages for monitoring. During the run of a job, if a command ends in error but has the error message in the list, the job will continue with the next command. See Figure 11-3 on page 188. Dynamic variables can be used in the command string. These variables will be resolved at run time. There are system-defined variables and user-defined variables. In the command string, the variables begin with an ampersand (&) and should normally have single quotation marks before and after.

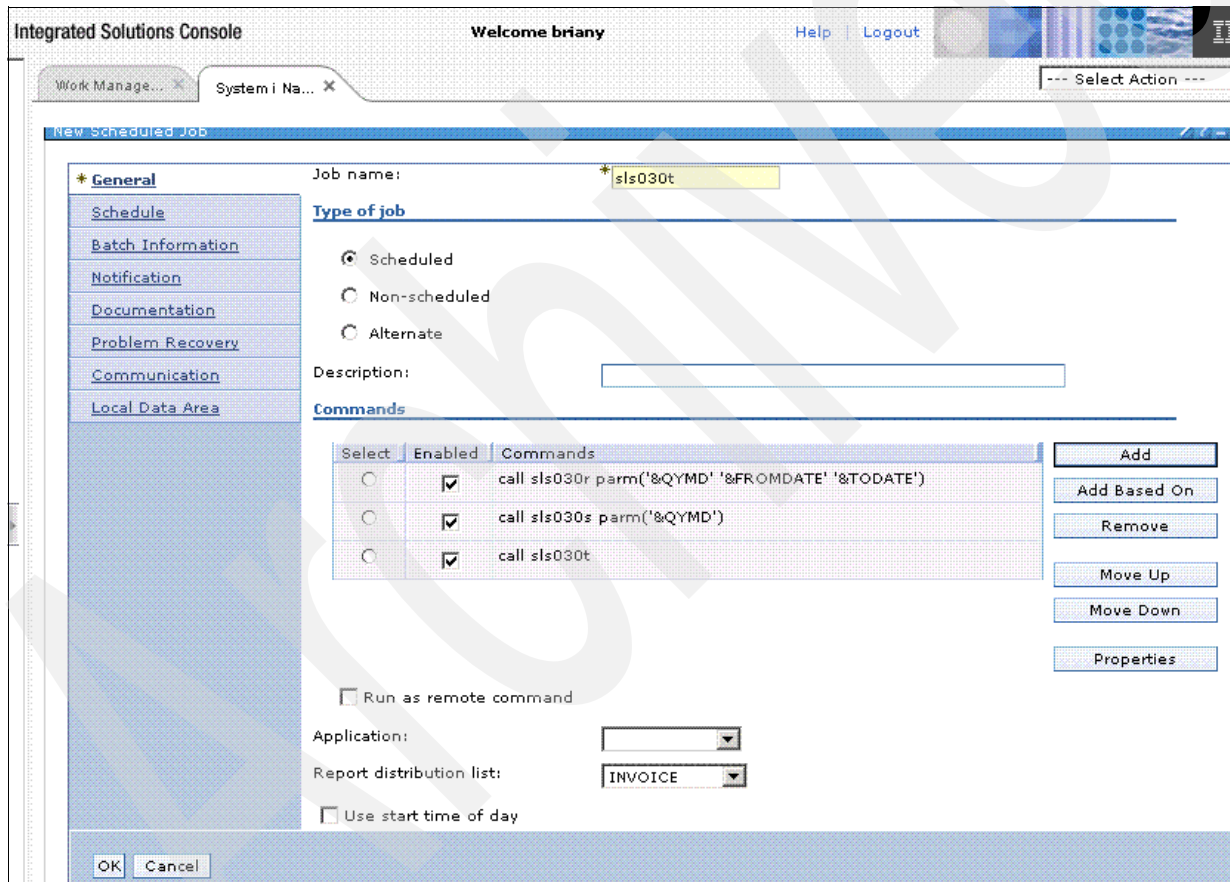


Figure 11-2 Command list in a scheduled job

Figure 11-2 shows the General tab of a job. This job has three commands to process. The first command is using a system-defined variable (QYMD) and two user-defined variables (FROMDATE and TODATE). The system-defined variable will resolve to the system date in year, month, day format. There are many system-defined variables in various formats.

Figure 11-3 shows a list of messages to monitor. If the command to process ends with one of these messages the job will continue.

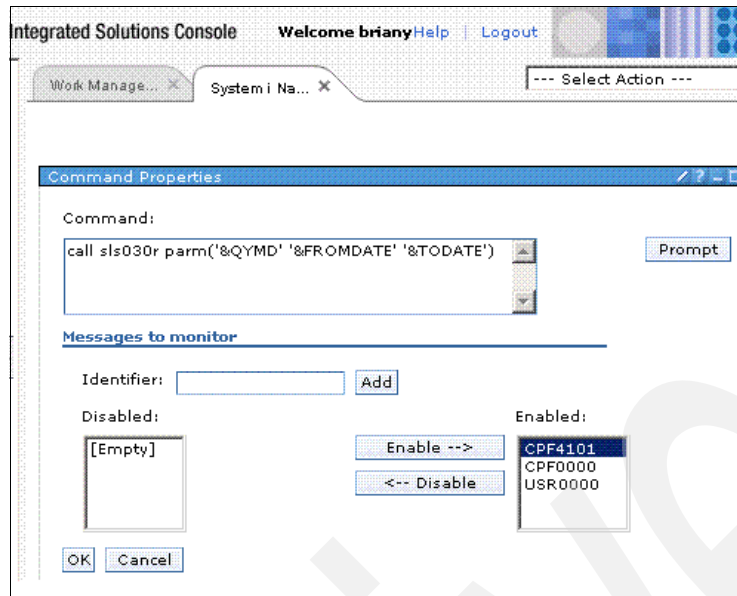


Figure 11-3 Messages to monitor

## 11.2.2 The schedule

Each job can contain a schedule in which it will run. The interactive calendar on the Schedule tab, along with the Frequency and Details box, is used to define the schedule. See Figure 11-4. A scheduled job can run up to 10 times per scheduled day or it can be defined to run every 30 minutes between the hours of 6:00 p.m. and 7:00 a.m.

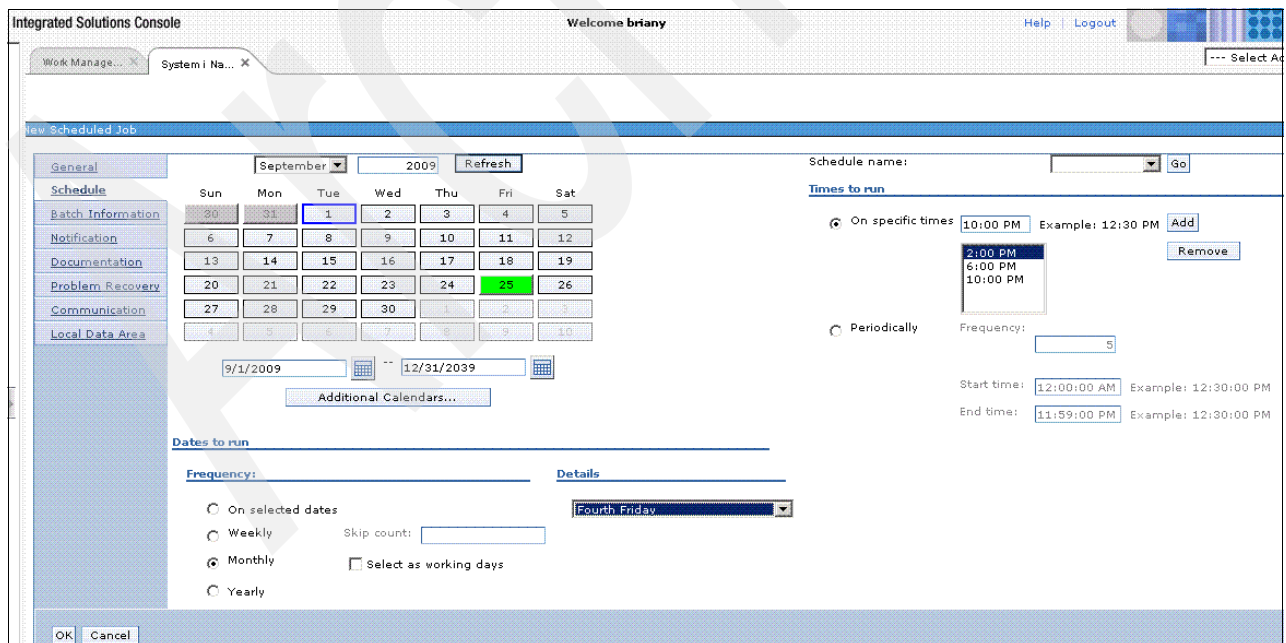


Figure 11-4 Schedule tab

Figure 11-4 on page 188 shows the Schedule tab of a job. In this example the job is scheduled to run at 2:00 p.m., 6:00 p.m., and 10:00 p.m. the fourth Friday of every month.

### 11.2.3 The notification

Notification comes in two types in respect to a scheduled job: the status of the job while it is running and after it has completed is one type. Distributing spooled files produced by the job is another type. On the Notification tab you can specify who will receive status notification via e-mail. Recipients are shortcut names for e-mail addresses. Recipients can be notified if the job completes normally, fails to complete normally, does not start within a specified number of minutes after it has been submitted, and if a job is waiting on an error message. See Figure 11-5. A report distribution list is specified on the General tab. This predefined list defines the spooled files to distribute along with the recipients. See 11.5, “Notification and report distribution” on page 192, for more information.

Figure 11-5 shows the Notification tab of a scheduled job. You can see the different recipients that will receive the various notifications.

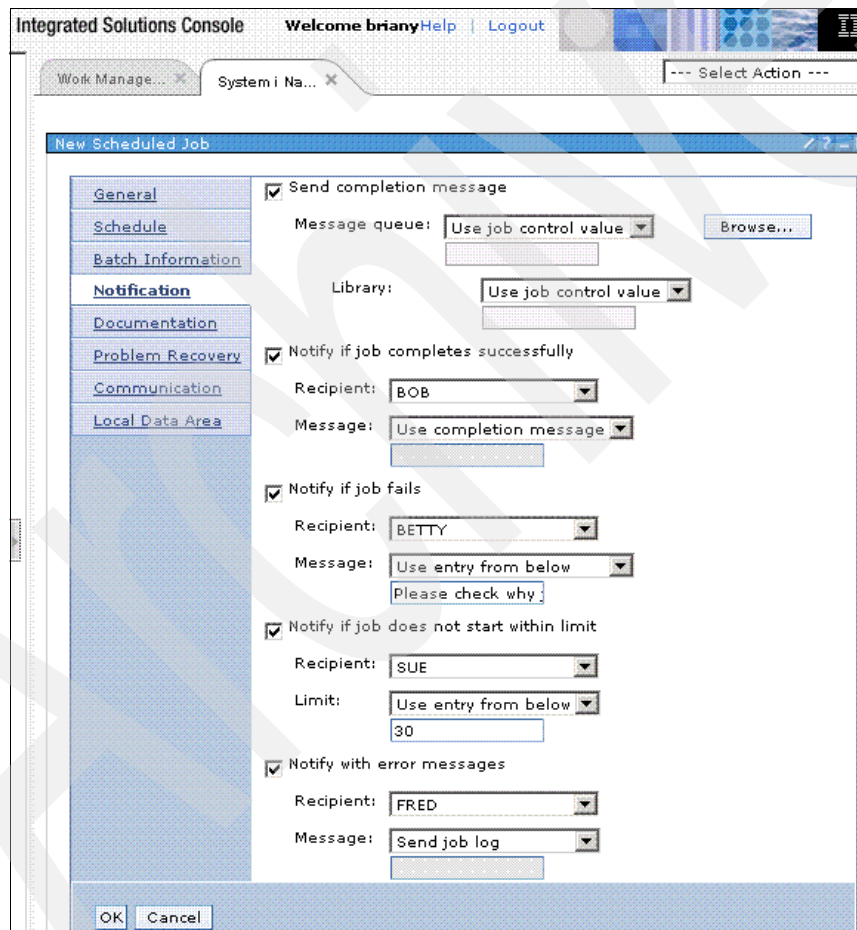


Figure 11-5 Notification tab

## 11.3 Job groups

Job groups are made up of AJS jobs grouped together to run consecutively in the order specified in the group sequence field. This is one form of dependency in AJS. A normal

completion is required before the next job in the group will be submitted. The first job within a group must have a sequence of 1 and will reflect the scheduling information. Changes to the schedule should be done at the group level. Any subordinate job (sequence greater than 1) will be bypassed during processing if the job's status is held.

There are several items to consider when creating a job group:

- ▶ A group definition must be created first. This will contain the scheduling information for the group.
- ▶ A job can be added to a group by listing the job groups and selecting the **Add New Job to Group** option on the group.
- ▶ An existing job can be added to a group by listing the schedule jobs and selecting the **Add Selected Job to Group** option on the job.
- ▶ When jobs are added to a group, they are added to the end.
- ▶ Subordinate group jobs can have their own schedule. This means that a subordinate group job may not run when the group is started.
- ▶ Jobs in a group can be resequenced by selecting the **Resequencing Jobs in Group** option.
- ▶ Holding the first job in a group will hold the schedule of a group.

## 11.4 Job dependencies

AJS allows you to set up jobs or job groups that are dependent on one another. For instance, you could set up a payroll job that will not run at a centralized site until all remote sites have transmitted their payroll data to the central site. Another job that prints the payroll checks could be dependent on the central site processing job being completed before the check printing job can start. See Figure 11-6. These job dependencies are called predecessor and successor relationships. A predecessor job is a job that must run before another job (called a successor job) can be started. A successor job is a job that is run after a predecessor job finishes. These relationships can be extended indefinitely to define the interrelationships of the jobs that you want to process.

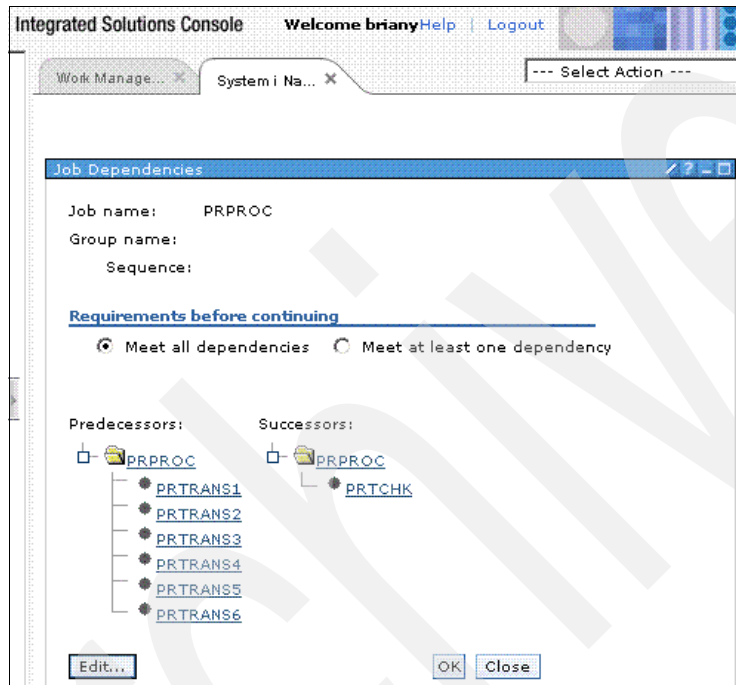


Figure 11-6 Job dependencies

There are several items to consider with job dependencies:

- ▶ You can have any number of predecessor jobs for a successor job.
- ▶ You can have any number of successor jobs for a predecessor.
- ▶ Unlike jobs in a group, you can define whether a normal completion is required before running the next job.
- ▶ You can define job dependencies so that at least one predecessor job must complete before a successor job will run or all predecessor jobs must complete before a successor job will run.
- ▶ A scheduled job that is a successor to another scheduled job should not have the same time. The scheduled time for the successor job should be at least one minute after the scheduled time for the predecessor job.
- ▶ A subordinate job in a group should not be a successor job.
- ▶ You can specify that the last job in a group is the predecessor. This means that no matter what job runs last in a group it is the predecessor to a successor job.

Figure 11-6 on page 191 shows all the dependencies for job PRPROC. You can see that PRTRANS1 through 6 must run before PRPROC will run, and PRPROC must run before job PRTCHK.

## 11.5 Notification and report distribution

The notification feature in AJS allows you to notify people of the status of jobs via e-mail. Important messages can be sent with an escalation list attached that specifies the length of time that will pass before the next person on the list is notified. Escalation will continue until someone stops escalation. Notification coupled with report distribution allows you to automatically distribute spooled files to one or more e-mail addresses or to one or more output queues.

Items to consider for notification and report distribution are:

- ▶ In order to be able to send e-mail, an outgoing mail server and default reply address must be specified. Access the main Properties page from the notification folder to enter the values.
- ▶ Specify a report distribution list on the General tab of a job.
- ▶ Recipients can contain one or more e-mail addresses.
- ▶ Through report distribution, you can split up a spooled file and send it to different e-mail addresses.
- ▶ You can specify in report distribution to remove the original spooled file after it has been distributed.
- ▶ Spooled files can be distributed in TXT format or PDF.
- ▶ Only basic text spooled files can be distributed.
- ▶ Spooled files that contain graphics cannot be distributed.
- ▶ Joblog spooled files must be explicitly defined in a report distribution list in order to be distributed.
- ▶ AJS does not save spooled file data for future distribution.
- ▶ It is possible to distribute reports for a scheduled job that has already run by locating the correct scheduled job activity entry and selecting the Distribute Reports menu option.

## 11.6 Remote command jobs

Commands can run on non-IBM i platforms such as PCs. They can be processed on a single system or a group of systems. The run remote command (RUNRMTCMD) is used to issue the commands on the remote systems. RUNRMTCMD allows server users to run a command on a remote system that is running the target portion of this function. The target portion of this function can be a remote executing daemon (REXECD). We recommend using the incoming remote command (IRC) service of the IBM i Access for Windows on the remote systems to process the commands.

Items to consider when running remote command jobs:

- ▶ Commands processed on non-IBM i platforms may report that the command ended normally even though it did not. The main reason is that the REXECD being used does not return the appropriate error-code information to the RUNRMTCMD. The process will produce a spooled file with information about the process that may help to determine whether the command completed normally. AJS cannot use the spooled file to determine job completion because of translation issues and inconsistencies.
- ▶ You can add systems and system groups to AJS by selecting **Properties** on the Advanced Job Scheduler folder. Select the **Systems and System Groups** tab on the left.
- ▶ You must check the **Run as remote command** check box on the General tab of a scheduled job to specify that the job is running on a non-IBM i system.
- ▶ The Batch Information tab for a job contains the system or system group to run the job.
- ▶ When processing commands on a group of systems, each command must complete normally on each system within the group before the next command is processed in the scheduled job.

## 11.7 Multiple scheduling environments

AJS has the ability to manage multiple scheduling environments. A scheduling environment is basically all the objects in the QUSRIJS data library duplicated into another library. QUSRIJS is the main data library created during the install of the AJS product. It contains the physical files that store all the information about the scheduled jobs. Each scheduling environment is its own entity and can be accessed and managed using AJS interfaces. When defining a scheduling environment within AJS, a monitor switch is available to designate whether the scheduling environment can be active. A scheduling environment must be active to automatically submit jobs. You probably would not want a replicated environment to become active if the source system was also active. This may cause the same job to be submitted from each active environment (source and replicated environments). However, there are many times when you would want multiple environments active at the same time. For example, you could have a production and test environment active at the same time. You can have any number of scheduling environments.

Items to consider when using multiple scheduling environments:

- ▶ A user can only access one scheduling environment at a time. A list of users and their assigned data library can be found on the Users tab on the main properties page for AJS. By default, the main data library (QUSRIJS) will be used by all users who have not been added to this list. The data library for a user can also be changed by using the *set data library using js* (SETDLJS) command.
- ▶ A job monitor job will exist for each data library that allows the data library to be monitored (active). A meaningful name can be specified for these jobs. The default is QIJSSCD. These jobs run in the QSYSWRK subsystem.
- ▶ The change data library using JS (CHGDLJS), set data library using JS (SETDLJS), and start job scheduler (STRJS) commands can be used in failover scripts on a backup system to change the attributes of a data library to allow monitoring, after which the job monitor for the data library can be started and the users set to use the replicated data library.
- ▶ Jobs can be copied from one data library to another using the copy job using job scheduler (CPYJOBJS) command. Be sure to specify \*DATALIB in the to job entry (TOJOB) parameter and the name of the data library.

- ▶ A data library must exist before it can be added to the list of data libraries in AJS. Here are some ways to create a data library for a new scheduling environment:
  - To create a clean data library you can check the Create data library box when adding a new data library to the list. This will first create a data library with all the necessary files and then add the data library to the list. This is done from the main properties page for AJS. See Figure 11-7.
  - You can use the copy library (CPYLIB) command to copy the QUSRIJS library. You would probably use this method when creating a test environment. The new data library could be named QUSRIJSTST.
  - You can use the save library (SAVLIB) and restore library (RSTLIB) commands to save the QUSRIJS library from one system and restore to another system. You may need to use the restore to library (RSTLIB) parameter in the RSTLIB command to restore the library as a different name.
  - You can use data replication software to replicate the QUSRIJS data library from the source to a different name on the target system. The data library on the target could be QUSRIJSPRD.

Figure 11-7 shows adding data library QUSRIJSABC. In this example the scheduling environment can be monitored (active) and the job monitor job will be named ABCSCD and will automatically be started after an IPL.

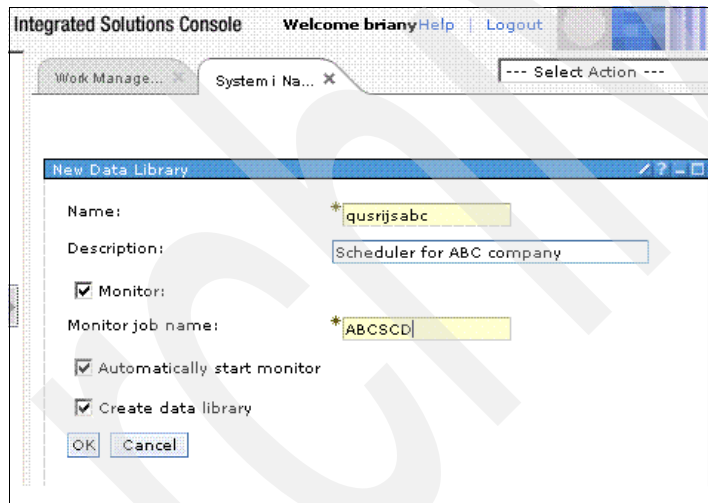


Figure 11-7 New data library





## Plug-ins and add-ons

This chapter contains information about Backup, Recovery, and media services, a high-availability solution, a cluster solution, and the plug-ins necessary to use on IBM Systems Director Navigator for i.

## 12.1 Backup, Recovery, and Media Services

Backup, Recovery, and Media Services helps you manage saves and media. It also enables you to restore objects that you have previously saved. BRMS enables you to manage your most critical and complex saves simply and easily, ensuring that you can recover your system fully in the event of a disaster or failure.

BRMS assists you in establishing a disciplined approach to designing and managing your save operations and provides you with an orderly way to retrieve lost or damaged data. It also enables you to track all of your media from creation to expiration and keep your system running smoothly by performing daily maintenance activities.

Within BRMS you can create control groups to set criteria for certain types of save operations. You can create either a backup control group or an archive control group. Backup control groups perform save operations on active data that is necessary for your day-to-day operations. Backup control groups set criteria on when saves occur, how these saves are performed, and where information is saved. You can create multiple backup control groups to perform full system saves to changes-only saves. By running backup control groups regularly, you can use a disaster recovery report that steps you through your recovery procedures.

Unlike backup control groups, archive control groups save objects that are used infrequently, but may need to be accessed for later use if needed. The saved objects are removed from the system. For example, you may have legacy customer information stored in a database and you want to free disk space by archiving this data to media. Another difference between backup and archive is the difference between restore and retrieve. In general, objects saved as part of a backup are recovered from the save media. Archived objects are retrieved from the media and then after they are used can be re-archived back to the media.

### ***LPPS for using BRMS functions within IBM Systems Director Navigator for i***

The LPPS for using BRMS functions within IBM Systems Director Navigator for i are:

- ▶ 5761BR1 \*BASE Backup Recovery and Media Services
- ▶ 5761BR1 1 BRMS - Network Feature
- ▶ 5761BR1 2 BRMS - Advanced Feature

### ***Important information for character-based interface users of BRMS***

You can manage backups, recovery operations, and your media with BRMS by using either the BRMS plug-in to System i Navigator (BRMS plug-in) or through the character-based or *green screen* interface. If you have used the character-based interface in the past and want to use the BRMS plug-in, you should be aware of the differences between these two interfaces.

### ***Terminology differences***

The following list describes the differences between terminology used in the character-based interface and the BRMS plug-in.

- ▶ The BRMS plug-in term *backup control groups* refers to the combined attributes of a backup control group and a media policy in the character-based interface.
- ▶ The BRMS plug-in term *changes-only backup* is the same as a cumulative backup in the character-based interface.
- ▶ The BRMS plug-in term *media pool* is the same as a media class in the character-based interface.
- ▶ The BRMS plug-in term *disk pool* is the same as an auxiliary storage pool in the character-based interface.

- ▶ The BRMS plug-in term *disk pool group* is the same as auxiliary storage pool class in the character-based interface
- ▶ The BRMS plug-in term *container pool* is the same as a container class in the character-based interface.
- ▶ The BRMS plug-in term *global policy properties* is the same as the system policy in the character-based interface.

### Compatibility differences

If your current character-based interface control groups do not reference the system policy (\*SYSPCY), archive policy (\*ARPCPY), or backup policy (\*BKUPCY) and do not share media policies among control groups, then using the BRMS plug-in to edit backup and archive control groups should have no functional impacts to your character-based control groups. If this is not true, read the following information.

The backup and archive control groups created by the BRMS plug-in can be viewed and run from the character-based interface. Also, control groups created from the character-based interface can be viewed and edited via the BRMS plug-in backup and archive control group properties. A message will be presented to the user when a character-based created or edited control group is about to be changed by the BRMS plug-in. The user will have the opportunity to cancel the operation or continue. However, we do not recommend changing control groups created on the character-based interface using the BRMS plug-in unless you understand the following:

- ▶ The BRMS plug-in backup and archive control groups are designed to be independent of each other. Thus, if you make changes to one control group, it will not affect the others. To ensure this independence, control groups updated with the BRMS plug-in will have all references to the character-based interface system policy (\*SYSPCY), archive policy (\*ARPCPY), and backup policy (\*BKUPCY) removed. These references will be resolved and replaced with the actual values when the backup control group is saved by the BRMS plug-in. The archive policy, backup policy, or system policy do not affect backup or archive control groups created or edited by the BRMS plug-in. The exceptions to this are the network feature, sign off exceptions, and subsystem to check controls in the system policy, which are used by the BRMS plug-in.
- ▶ Changes made to other, non-IBM supplied policies are reflected in the BRMS plug-in.

**Note:** For further information see the following Web site:

<http://www-03.ibm.com/systems/i/support/brms/pluginfaq.html>

When you select in the IBM Systems Director Navigator for i main menu Backup Recovery and Media System option you will see the window shown in Figure 12-1.

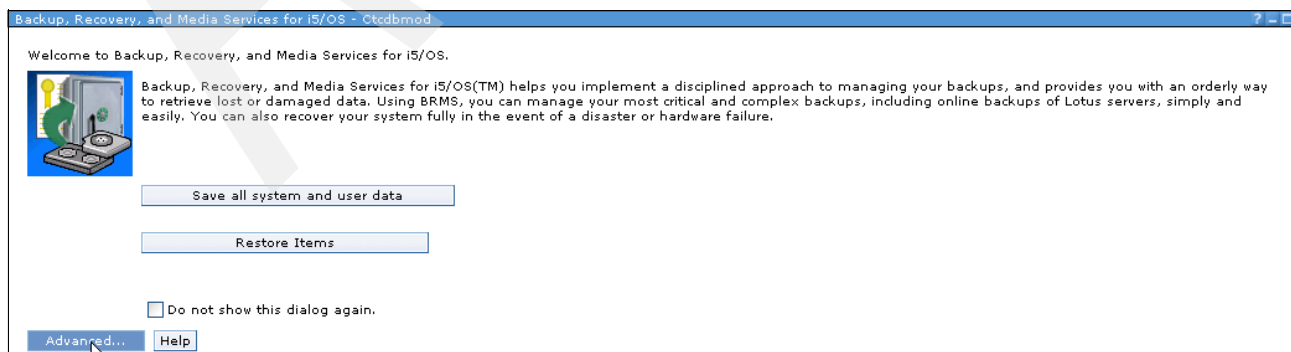


Figure 12-1 Main Backup Recovery and Media System menu selected option advance

As shown in Figure 12-1 on page 197, if you click the **Advanced** button, you see the window shown in Figure 12-2.

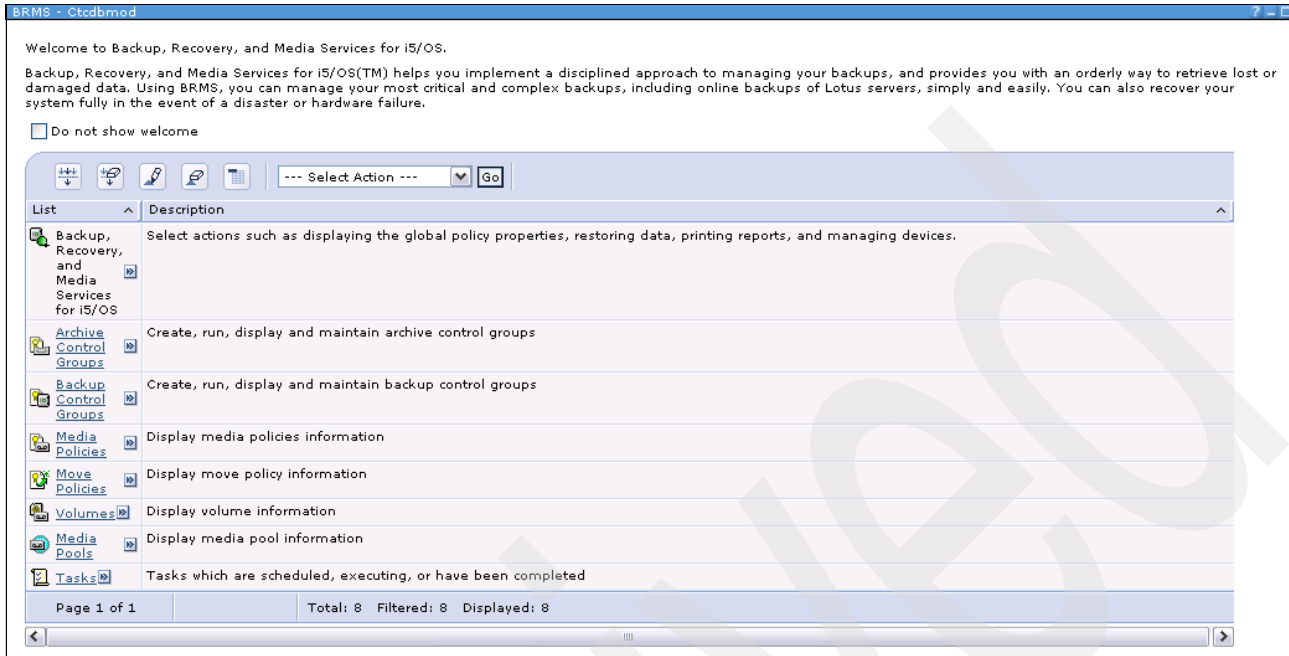


Figure 12-2 BRMS interface menu

**Note:** For more details about the options for the BRMS interface menu, see the IBM Redbooks publication *IBM i 6.1 Technical Overview*, SG24-7713.

## 12.2 High-availability solution

In this section we discuss a high-availability solution.

## High Availability Solution Manager

If you have the HASM LPP installed (5761HAS \*BASE iHASM) you can use the new High Availability Solution Manager GUI to select, configure, and manage a complete high-availability solution. You can choose from a list of six high availability solutions. This new GUI is intuitive and does not require extensive knowledge of the underlying system technology. When you select the High Availability Solution Manager option you see the window shown in Figure 12-3.

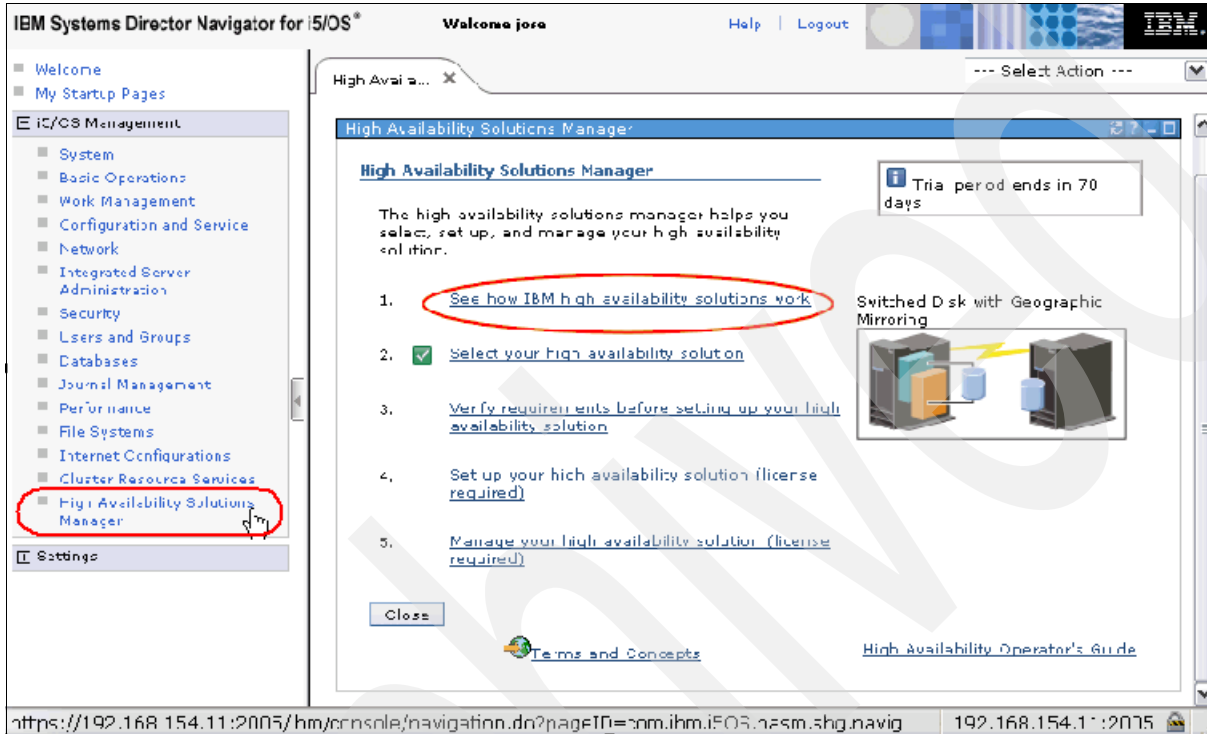


Figure 12-3 High Availability Solution Manager menu

**Note:** For more information about the High Availability Solution Manager functions see *Implementing PowerHA™ for IBM i*, SG24-7405.

## Cluster Resource Services

If you have the High Availability Solution Manager (HASM) LP installed, you can use the new Cluster Resource Services GUI to perform various clustering functions, such as working with clusters, nodes, cluster resource groups, switchable devices, administrative domains, and monitored resource entries. Clusters are a key component of setting up a high-availability solution to protect your system from system outages and disasters. If you do not have the HASM LP installed, you would use the Clusters GUI on System i Navigator instead to perform the clustering tasks.

You can interoperate between the Cluster Resource Services GUI and the character-based interface.

You can use the Cluster Resource Services GUI to create a cluster with up to 128 nodes. Previously, the largest cluster that you could create using the GUI was two nodes.

Previously, only disk pools were switchable. Now you can switch other types of devices, such as tape devices, optical devices, communications lines, controllers, and network servers, to

the backup node when an outage occurs. The cluster ensures that the resource names and hardware resource information is identical on all recovery domain nodes. You can use the New Device Cluster Resource Group wizard to add switchable devices to a new device cluster resource group, or use the Add an Existing Switchable Device panel to add the devices to an existing device cluster resource group.

You can define a single failover message queue in Cluster Properties or Node Properties for sending failover messages for an entire cluster or backup node.

You can specify new types of monitored resources, such as user profiles, system values, and device descriptions, for an administrative domain to ensure that the operational environment for your application is consistent across the cluster nodes. Use the Add Monitored Resource Entry panel to add new resources to monitor.

You also can specify the Synchronization Option in Administrative Domain Properties to control how resource synchronization will occur, for example, when a node rejoins a cluster. See Figure 12-4.

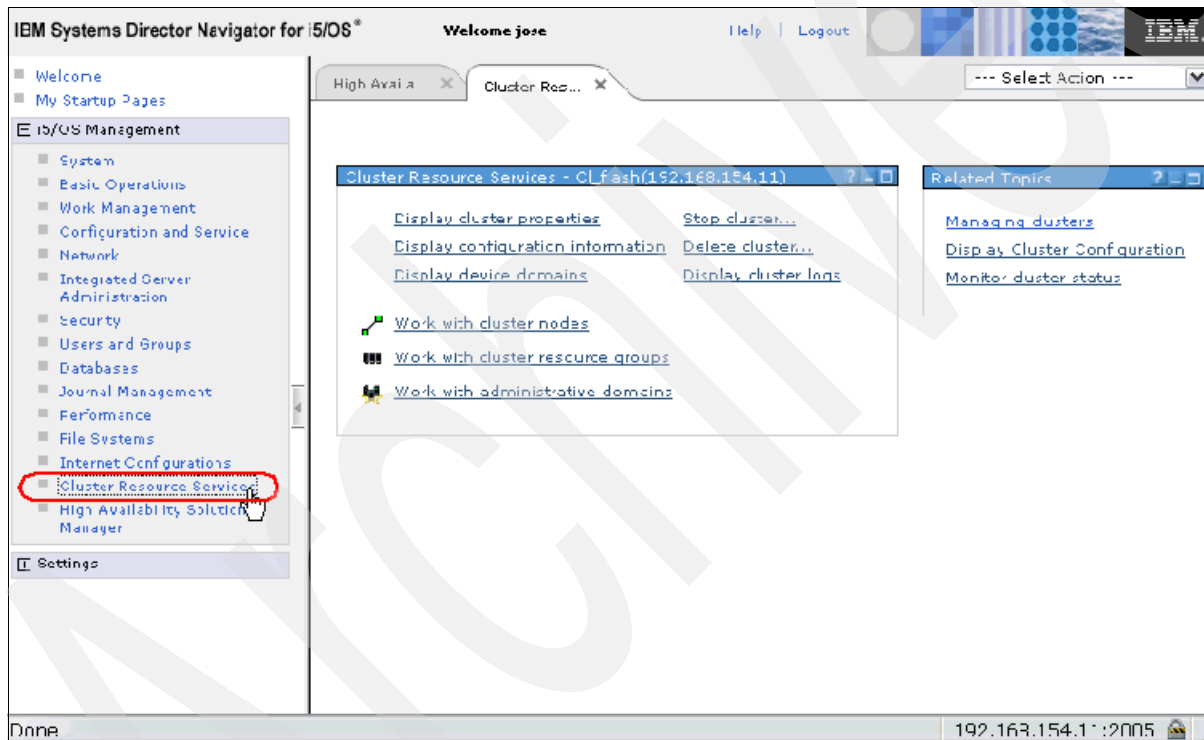


Figure 12-4 Cluster Resources Services

**Note:** For more information about the Cluster Resources Services functions see *Implementing PowerHA for IBM i*, SG24-7405.

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

## IBM Redbooks publications

For information about ordering these publications, see “How to get Redbooks publications” on page 201. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *End to End Performance Management on IBM i*, SG24-7808
- ▶ *IBM iDoctor iSeries Job Watcher: Advanced Performance Tool*, SG24-6474
- ▶ *IBM i 6.1 Technical Overview*, SG24-7713
- ▶ *Implementing PowerHA for IBM i*, SG24-7405

## Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM Systems Director Navigator for i Information Center  
<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/topic/rzatg/rzatgdirectoroverview.htm>
- ▶ IBM i Information Center  
<http://publib.boulder.ibm.com/infocenter/iseriess/v6r1m0/index.jsp>
- ▶ BRMS Graphical User Interface FAQ  
<http://www-03.ibm.com/systems/i/support/brms/pluginfaq.html>

## How to get Redbooks publications

You can search for, view, or download Redbooks publications, Redpapers publications, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](https://ibm.com/support)

IBM Global Services

[ibm.com/services](https://ibm.com/services)



# Index

## A

- Access for Windows 192
- Access plans 131
- access plans 123
- active connections 45
- add 24
- Add Job Watcher Definition 154
- administrative domains 199
- Advanced Job Scheduler 185
- AES 85
- AJS 185
- archive control groups 196
- ASP 175
- ASP master key 85
- audit journal 88
- audit records 88
- authorization list 82
- authorization lists 33
- auxiliary storage pool 196
- auxiliary storage pool class 197
- auxiliary storage pools 175

## B

- backup control group 196
- Backup control groups 196
- Basic Operations 30
- batch 30
- BRMS 6, 196

## C

- CA 10
- CCA Cryptographic Service Provider 87
- CCA CSP 87
- CCSID 181
- Central server 59
- certificate 10
- Certificate Authority 10
- Change Query Attributes 113
- changes-only backup 196
- CHGQRYA 113
- Classic Query Engine 122, 131
- Client side requirements 9
- Clone master keys 88
- Cluster Properties 200
- cluster resource groups 199
- Cluster Resource Services 199
- clusters 199
- coded character set ID 181
- Collection Library 149
- Collection Retention Period 149
- Collection Services 160
- comparison 2
- Configuration 7

- Connection type 50
- Console Logging and Tracing 25
- constraints 112
- container class 197
- Container Pool 197
- CPYJOBJS 193
- CPYLIB 194
- CQE 122, 131, 133
- Create DNS Server 53
- Create User-Defined Server 53
- Create User-Defined Server wizard 57
- Credential Store 25
- CRTPFRDTA 149
- Cryptographic Coprocessor 87
- cryptographic functions 33
- cumulative backup 196
- current library 105
- Cycle Interval 150
- Cycle Time 150

## D

- Data queue server 59
- Database Monitor 119
- database navigator maps 106
- DB2 Symmetric Multiprocessing 122
- DDM 59
- DDS 110
- default library 104
- default table size 23
- denial of service 88
- dependencies 191
- dependency 189
- dependent 191
- digital certificate 38
- digital certificates 38
- Disk Pool 196
- Disk Pool Group 197
- disk status 29
- disk units 31
- Disk Watcher 160
- DiskWatcher 6
- Distributed Data Management 59
- DLO 173
- DMZ 66
- DNS 32, 53, 57
- DNS dynamic update 58
- DNS Dynamic Update Keys 53
- document library objects 173
- domain 42
- Domain Name System 32, 53, 57
- Domino Servers 53
- Domino@ Servers 53
- dynamic IP address 67

## E

Edit mode 23  
EIM 77  
Electronic Customer Support 50  
Electronic Service Agent™ 50  
electronic support 50  
enrollment 77  
Environmental Limits 143  
escalation list 192  
extrusions 88

## F

file share 180  
file system 37  
Forbidden 14  
Full Masquerading 67

## G

generate SQL 110  
global policy properties 197  
Graph Data Creation 150  
Graph Data Retention 150  
group sequence field 189  
groups 34, 77  
guidelines 164

## H

HASM 199  
Health and Performance 117  
Health Center 35, 140  
Help 18  
help server 38  
hide 25  
hierarchical key system 84  
High Availability Solution Manager 199  
Historical Data Creation 150  
Historical Data Retention 150  
Holding 190  
host name 42  
HTTP 403 14  
HTTP 404 14  
HTTP 500 14  
HTTP Admin Server 8  
HTTP proxy server 67  
HTTP server 38

## I

i5/OS Tasks Page 18  
IBM i required products 9  
IBM i Task Page 4  
IDS 34  
iHASM 6  
Incoming Remote Command 192  
Independent Auxiliary Storage Pool 85  
index advisor 136  
Installation 7  
integrated and virtual servers 71

Integrated File System 168  
integrated file system 37  
Integrated Server Administration 32  
Integrated Server Support 179  
Integrated Solutions Console 2  
Integrated xSeries Adapter 73  
Integrated xSeries Server 168, 179  
Integrated xSeries® 73  
interactive 30  
interfaces 31, 41  
Internet Setup Wizard 66  
Internet setup wizard 38  
intrusion detection 88  
intrusion detection system 34  
IP Packet Filtering 67  
IP packet filtering 62, 67  
IP policies 31  
IPP 38  
IPv4 31, 41  
IPv6 31, 41  
IRC 192  
ISC 2  
iSCSI 77  
iSCSI Integrated Server 73  
IXA 73  
IXS 73, 168, 179

## J

job cache 131  
Job groups 189  
Job Watcher 160  
JobWatcher 6  
Journal Management 145  
Journals 36  
Journals Recivers 36

## K

Key Constraints 109  
key verification value 86  
keystore 86  
keystores 33, 84

## L

L2TP 50  
Leased line 50  
left navigation area 19  
library list 106  
License Program Products Plugins 6  
line description 47  
lines 41  
listening ports 45  
Logout 18

## M

Manage Dynamic Update Keys 58  
Management Central 150  
Management Module 79  
master keys 33, 84

- Materialized query table 122
- media class 196
- media policy 196
- Media Pool 196
- messages 29–30
- monitored resource entries 199
- Mounting 37
- MQT 122
- multiple tabs 21

## N

- Net Print server 59
- NetServer 37, 180
- NetServer Disabled User IDs 68
- NetServer Sessions 68
- Network Address Translation 67
- Network File System 37, 168
- Network Server Host Adapters 77
- networking servers 31
- New Controls 28
- NFS 37, 168
- NFS Mount 55
- Node Properties 200
- nodes 199
- Notification 189
- NWSH 77

## O

- object permissions 33
- On Demand 35
- Operating mode 50

## P

- packet priority 65
- Packet Rules Editor 62
- passwords 29
- PC5250 2
- performance 37
- Performance Data Investigator 149
- permissions 84
- ping 31, 43
- Plan Cache 106
- Point-to-Point 49
- policies 88
- portlet 2
- Positioning 2
- POSIX 171
- PPP 49
- PPP profile 67
- PPPoE 50
- predecessor 191
- printer output 30
- printers 30
- private network 67
- pseudo opens 122
- PTF Levels 9
- Public authority 84

## Q

- QAQQINI 115
- QDLS file system 173
- QFileSvr.400 file system 177
- QHTTPSVR subsystem 8
- QNTC file system 179
- QOpenSys file system 171
- QOPT file system 178
- QoS 65
- QSYS.LIB file system 175
- quality of service 65
- query options file 115
- QUSRIJS 193

## R

- RC2 85
- Redbooks Web site 201
  - Contact us xi
- Remote Access Services 49
- remote access services 31
- Remote Command server 59
- remote databases 106
- remote executing daemon 192
- remote Windows commands 73
- remove 24
- Required jobs 8
- resequenced 190
- REXECD 192
- Root File System 169
- routes 31, 41
- RSA 85
- RSTLIB 194
- Run Remote Command 192
- Run SQL Statements 35
- RUNRMTCMD 192

## S

- SAPI 87
- Save System 85
- save/restore master key 85
- SAVLIB 194
- SAVSYS 85
- schedule 188
- scheduling environments 193
- schema 105
- Secure Sockets Layer 10
- security 33
- security programming interface 87
- security warning 10
- Servers Properties 53
- Service Processor 79
- SETDLJS 193
- SHA1-HMAC 85
- Sign on server 59
- single sign-on 2
- SLIP 49
- slow performance 15
- sort 24
- sparse indexes 123

- spoofed packet 88
- spooled files 189, 192
- SQE 122, 130
- SQL packages 131
- SQL Performance monitor 117
- SQL Plan Cache 130
- SQL Plan Cache Snapshot™ 133
- SQL Query Engine 122, 130
- SSL 10
- statement cache 131
- status 189
- stream input/output 168
- subordinate job 190
- successor 191
- Support Tips 7
- switchable devices 199
- Switched line 50
- Synchronization Option 200
- system defined variable 187
- system defined variables 187
- System i Access 53–54
- System i Navigator 2
- system load 65
- system policy 197
- system status 29
- System tasks 29
- system values 31

## T

- table 108
- tabs 21
- task categories 28
- TCP/IP Configuration 41
- TCP/IP configuration 31
- TCP/IP Servers 53
- TCP/IP servers 55
- teraspace 149
- Text Conversion 181
- time management 31
- trace route 31
- traceroute 43
- transactions 106
- triggers 112
- troubleshoot 43

## U

- Universal Connection wizard 50
- user defined variable 187
- user defined variables 187
- user-defined servers 56
- users 34, 77

## V

- Virtual IP 67
- virtual private network 64
- virtual servers 71
- Visual Explain 133
- VPN 64

## W

- warnings 10
- Web Serving 67
- welcome page 18, 20
- Windows domains 77
- wizards 31
- Work Management 30
- workload 164

## X

- X/Open Portability Guide 171
- XPG 171

Archived







# IBM Systems Director Navigator for i



**Redbooks®**

**Learn about the new Web console for IBM i**

**Explore the functionality of the new Web console**

**Leverage this new powerful tool to manage IBM i**

In this IBM Redbooks publication we discuss IBM Systems Director Navigator for i, which is a Web console interface for System i administration where you can work with the Web-enabled tasks of System i Navigator. IBM Systems Director Navigator for i includes a number of welcome pages that allow you to quickly find the task that you want to perform.

The IBM Systems Director Navigator for i interface is not just a set of URL addressable tasks, but is a robust Web console from which you can manage your System i model. However, the System i Navigator Tasks on the Web, which are a set of URL-addressable tasks, can be accessed by using the URL or from within the IBM Systems Director Navigator for i interface.

The information in this IBM Redbooks publication is intended to help you start using the Web-based console, IBM Systems Director Navigator for i, by providing you with a look at the new interface as well as some tips for working with various parts of the new console.

## **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](http://ibm.com/redbooks)

SG24-7789-00

ISBN 0738433802