IBM

IBM System i and System p System Planning and Deployment: Simplifying Logical Partitioning

- Using the enhanced System Planning Tool
- Deploying with the HMC V7 interface
- Planning and deploying a Virtual I/O Server

Dale Barrick
Gene Meyer
Marc Rauzier
Chris Shomenta
Jim Cook

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM System i and System p System Planning and Deployment: Simplifying Logical Partitioning

August 2007

SG24-7487-00

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

First Edition (August 2007)

This edition applies to i5/OS V5R4, AIX 5L 5.3, System Planning Tool Version 2, Virtual I/O Server 1.4.0.0, and the June 2007 release of IBM Systems Workload Estimator.

Contents

Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information about 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 products and/or the programs 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

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

Redbooks (logo)  ®
eServer™
iSeries®
i5/OS®
pSeries®
AIX 5L™
AIX®
AS/400®
DS6000™
DS8000™
HACMP™
IBM®
Micro-Partitioning™
OpenPower™
OS/400®
POWER™
POWER Hypervisor™
POWER5™
POWER6™
Redbooks®
System i™
System i5™
System p™
System p5™
System x™
System Storage™
System/36™
System/38™
WebSphere®

The following terms are trademarks of other companies:

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

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

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

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

Logical Partitioning (LPAR) provides the significant capability to run multiple operating systems, each a partition on the same physical processor, memory, and I/O attachment configuration. LPAR is often discussed along with the concept of *server consolidation*. LPAR enables management across a single set of hardware and, when configured and managed correctly, can maximize efficient use of hardware resources all in a single place, often using resources in one partition when not needed by another partition.

By its nature, LPAR is powerful, but, as the number of and complexity of applications being run in each partition increases, can become complex to configure and achieve anticipated performance expectations.

This IBM® Redbooks® publication describes and provides examples of using the 2007 enhancements to the system planning and deployment tools and processes for planning, ordering, and deploying a partitioned environment on IBM System i™ and IBM System p™ configurations.

The objective is to help you order and IBM deliver a hardware configuration and get that configuration up and running your planned partition configurations with good performance in as short a time as possible.

This book and the tools and processes involved represent the next step in expediting this entire process, while still requiring sound knowledge of IBM System i and System p hardware processor and I/O capabilities for success.

Key new enhancements addressed in this book include:

- New System Planning Tool (SPT) Version 2, made available May 2007. This tools covers placing hardware in partitions on systems, including the latest IBM System i and System p support POWER5™ and POWER6™ processor technologies and I/O attachment capabilities
- New and improved interfaces between SPT Version 2 and HMC level V7.3 and deploying partitions of i5/OS®, AIX® 5.3, and supported Linux® operating systems.
- Improved graphical views of the hardware configuration
- Configuring and deploying a Virtual I/O Server partition on an IBM System p configuration.

Additionally, this book shows examples of using the IBM Systems Workload Estimator (WLE) sizing tool with SPT Version 2 to help define partitions with appropriate processor, memory, and disk configurations.

This book is intended for those familiar with IBM System i and IBM System p hardware capabilities and LPAR configuration considerations. However, for those new to the hardware and LPAR possibilities, use of this book and the tool examples within it can speed up gaining expertise in these areas.

The team that wrote this book

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

Jim Cook is a Senior Software Engineer Project Leader at the ITSO, Rochester Center. He leads teams that produce IBM System i announcement presentation sets that are maintained on the System i technical support Web sites and presents at ITSO iSeries® Forums internationally. Jim also produces IBM Redbooks about various System i and i5/OS-related topics.

Dale Barrick is a Certified Consulting IT Specialist for System i with Techline based in Atlanta, Georgia and has 30 years of experience. As a Systems Engineer his background is S/36, AS/400®, and System i. He is currently the System i Techline Focal Point for the Southeastern Area and was previously the National Lead for the Server Consolidation/LPAR Solutions Initiative team in Techline. He participated in the creation of the V5R2 LPAR Certification examination. He has also worked with the iTC performing LPAR plan validations and has assisted the iTC in the OS/400® V4R5 and V5R1 LPAR class rewrites and the original LPAR Validation Tool development.

Dale was a member of the development team for Customer Specified Placement (CSP), has co-authored several IBM Redbooks, and for the past two years he has been a consultant to the development team for the System Planning Tool (SPT). You can contact him by sending an e-mail to drbarri@us.ibm.com.

Gene Meyer is a Certified Senior IT Specialist with System i Techline in Dallas, Texas. He has 40 years of experience with data processing. He has 25 years with IBM working with VSE, System/34, System/36™, System/38™, AS/400, iSeries, and System i. He holds a degree in Computer Science from Texas A&M Commerce and an MBA from the University of Texas Pan American. His areas of expertise include iSeries connectivity, programming, System x™ connectivity, hardware configuration, and Linux installation on iSeries and Capacity Planning. You can reach him via e-mail at gmeyer1@us.ibm.com.

Marc Rauzier is a Certified Consulting IT Specialist with IT Delivery for Strategic Outsourcing business based in Lyon, France. As a Systems Engineer, his 20 years of experience is S/38, AS/400, iSeries, and System i. He is currently responsible for design, installation, and improvement of solutions for France Strategic Outsourcing customers, based on multipartitioned servers, external disks and tapes storage, and storage area network. His areas of expertise include System i, HMC, SAN, 35xx series tape libraries, and DS series external storage related to System i environment.

Chris Shomenta is an IT Specialist and Software Designer in Rochester, MN. He has 10 years of experience in the hardware and service field. He holds a degree in mathematics from The University of Texas at Austin. His areas of expertise include System i error analysis, System i and HMC interfaces, and System i High Availability.

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

International Technical Support Organization, Rochester Center

Tammy Van Hove
Greg Andrews
Dave Gimpl
Jossie McManus
Kent Hoffer
Allan Johnson
Sathyakumar (Andre) Ramalingam
George Romano
Nathan Rabe
IBM Rochester development

Nancy Milliner
Vasu Vallabhaneni
IBM Austin development

Steve Hochstetler
Narend Chand, resident from IBM New Zealand
Syamsul Hidayat, resident from IBM Indonesia
IBM Austin ITSO

A special thanks is expressed to the authors of the previous IBM Redbooks publication *LPAR Simplification Tools Handbook*, SG24-7231. The authors of this publication include:

Nick Harris, IBM Rochester Executive Briefing Center
Joe Mulholland, IBM Rochester Support center
Allyn Walsh, Americas Advanced Technical Support (ATS)

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

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

1

System planning and deploying logical partitions overview

This chapter provides an overview of the logical partitioning area of the system planning and deployment process. We overview 2007 enhancements in the various tools used to simplify the implementing of the planning, ordering, and deployment steps to get a system with logical partitions ordered and ready for production operation.

1.1 System planning and deployment for logical partitions overview

Logical Partitioning (LPAR) provides the significant capability to run multiple operating systems, each a partition on the same physical processor, memory, and I/O attachment configuration. LPAR is often discussed along with the concept of *server consolidation*. LPAR enables management across a single set of hardware and, when configured and managed correctly, can maximize efficient use of hardware resources all in a single place, often using resources in one partition when not needed by another partition.

Depending upon the particular POWER5 or POWER6 processor model, IBM System i supports i5/OS, AIX 5.3, and supported Linux distributions in a partition. Likewise, IBM System p POWER5 and POWER6 support AIX 5.3 and supported Linux distributions in a partition. IBM System p and System i also support a partition running a specialized IBM *operating environment* called the Virtual I/O Server (sometimes represented also as VIOS). Virtual I/O Server is much more commonly used on System p POWER5 and POWER6 models than on System i.

On IBM System p models you can select to use a specialized version of the Virtual I/O Server, known as the Integrated Virtualization Manager (IVM). The Integrated Virtualization Manager is a special partition that enables additional logical partition configuration and management without use of a Hardware Management Console (HMC) device.

Setting up partitions, depending upon your needs, can require thorough planning. In most cases LPAR requires significant expertise, not only in understanding specific hardware capabilities, but also understanding the type of application workloads to be run in each partition to help ensure satisfactory results in meeting performance expectations for each active partition.

As such, the process of sizing and ordering a system hardware configuration with the necessary hardware components and later configuring those resources into the appropriate operating system partitions can become very complex as the number of partitions you need increases, along with all the hardware needed by each partition.

The System Planning Tool can be used to simplify this entire process by minimizing the amount of hardware expertise needed when planning for partitions. The System Planning Tool also aids in partition configuration to help assure that you get your partitions up and running with satisfactory performance when configured with the help of the output of a capacity planning or sizing tool, such as the IBM Systems Workload Estimator (WLE).

System plan deployment leverages the output of the System Planning Tool to automate initial partition configuration upon system setup. Utilization of both system planning and deployment capabilities can improve productivity and help assure a workable configuration.

Utilization of the System Planning Tool can also be used as an ordering aid, as the system plan produced can be imported into IBM Sales Configurator.

The System Planning Tool is the key player in simplifying the system planning and deployment process for IBM System i and System p configurations, and especially important when logical partitions will be deployed. It helps minimize the level of hardware expertise needed when planning for partitions. The System Planning Tool also aids in partition configuration to help ensure that you get your partitions up and running with satisfactory performance. System plan deployment leverages the output of the System Planning Tool to automate initial partition configuration upon system setup. Utilization of both system planning

and deployment capabilities can improve productivity and help assure a workable configuration.

Note that the System Planning Tool can be used productively for planning and ordering a system with no partitions as well as a system with multiple partitions.

This book focuses on Version 2 of the System Planning Tool which takes advantage of other products involved in logical partitioning, such as the Hardware Management Console, Virtual I/O Server partitions, Integrated Virtualization Manager partitions, and the IBM Systems Workload Estimator tool.

From a System i and System p viewpoint, simplifying system planning and deployment of logical partitions made its first major step forward in 2006, providing better integration of the following tools:

- Use of the IBM Systems Workload Estimating (WLE) tool to size an IBM System i or System p configuration and using WLE's output as input into the System Planning Tool (SPT).

  Note that though there is no formal sizing output from the following non-IBM sizing tools into the IBM sales configurator tool, or the IBM System Planning Tool (SPT), you can also use the following products to help you configure your partitions to meet your performance expectations:

  - MPG Performance Navigator

    For more information about the Midrange Performance Group (MPG) Performance Navigator, refer to:

    `http://www.mpginc.com`

  - BMC PATROL for iSeries - Predict

    For more information about BMC Patrol for iSeries - Predict refer to:

    `http://www.bmc.com/`

    To assist you in using WLE, PATROL for iSeries - Predict, and Performance Navigator, consider using the IBM Redbooks publication *Sizing IBM i5/OS Work in IBM System i5™ Partitions,* SG24-6656, which contains examples of using these tools.
- System Planning Tool partition configuration tool output as input to the IBM hardware and software ordering tool — the IBM Sales Configurator (sometimes also referred to as e-Config) and process.
- Initial capability to import WLE output into the SPT for base processor and main memory configuration
- Deployment of partitions on the hardware configuration received from IBM manufacturing

These capabilities introduced during 2006 are documented in the *LPAR Simplification Tools Handbook,* SG24-7231.

In this new IBM Redbooks publication we document a second major step forward in simplifying the system planning and deployment of logical partitions. In this book we frequently use the term *Version 2* to collectively represent the 2007 set of functional and integration enhancements to all of the tools and processes involved in simplifying the system planning and deployment of logical partitions.

Version 2 includes:

- Normal enhancements to the IBM Workload Estimator and the IBM Sales Configurator through mid-2007.

- New SPT Version 2 capabilities (available as of May 2007) which include support of POWER5 and POWER6 processor technology on System p and i configurations. POWER6 System p is available now and support on System i will be available later in 2007.
- New HMC Version 7 release capabilities (available June 2007), which include support of POWER5 and POWER6 System p and i configurations.
- A new System Plan Viewer is available. This viewer is both part of the SPT tool and part of the Hardware Management Console Version 7. The viewer allows you to view and print complete summary or detailed views of the system's partitions and hardware. This topic is covered in multiple sections of the book, where viewer function is applicable.
- SPT V2 and HMC support of configuration and deployment of IBM System p Virtual I/O Server (VIOS) and Integrated Virtualization Manager (IVM) *partitions*.

This book covers:

- Use of the System Planning Tool (SPT) Version 2, made available May 2007.
- Interfaces between SPT Version 2 and HMC level V7.3 and deploying partitions of i5/OS, AIX 5.3, and supported Linux operating systems.
- Configuring and deploying a Virtual I/O Server partition on an IBM System p configuration.
- Using the IBM Systems Workload Estimator (WLE) with SPT Version 2 to help define partitions.
- Installing SPT Version 2 and other SPT operational considerations.

For more details on Virtual I/O Server, HMC V7.3 capabilities, and general POWER5 and POWER6 and i5/OS and AIX 5.3 capabilities we refer you to:

- eServer™ Hardware Info Center (POWER5, POWER6)

  `http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp`
- iSeries Information Center - System i Information Center

  `http://publib.boulder.ibm.com/iseries/`
- pSeries® Information Center - System p Information Center

  `http://publib16.boulder.ibm.com/pseries/index.htm`
- IBM Redbooks publications
  - *HACMP 5.3, Dynamic LPAR, and Virtualization*, REDP-4027
  - *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000-01
  - *Advanced POWERAdvanced POWER Virtualization on IBM System p5: Introduction and Configuration*, SG24-7940
  - *Advanced POWER Virtualization on IBM System p Virtual I/O Server Deployment Examples*, REDP-4224
  - *Integrated Virtualization Manager on IBM System p5™,* REDP-4061-01
  - *Hardware Management Console V7 Handbook*, SG24-7491
  - *IBM System i Overview: Models 515, 525, 595, and More*, REDP-5052
  - *PCI and PCI-X Placement Rules for IBM System i models with i5/OS V5R3 and V5R*4, REDP-4011-03.

We continue to have IBM Redbooks publication *LPAR Simplification Tools Handbook*, SG24-7231, available, as not all those doing LPAR configuration will upgrade to the newer

SPT and HMC versions of software immediately. We add this new Version 2 book to the existing work on LPAR management and best practices rather than replacing the existing book.

Who this book is for

This book is aimed at technical professionals who are involved in the planning, sizing, and configuration of System p and System i systems. There is an assumed level of hardware, operating system, and logical planning knowledge beyond the novice level.

This publication cannot make you an LPAR expert. Nor can this book make you an expert in the use of each of the system planning and deployment tools and the IBM sales configurator tool we describe in this book. To gain expertise requires repetitive use of each tool as well as perhaps attendance at specific IBM courses on each tool.

This book shows you examples of using the system planning and deployment tools and covers some problem determination issues. However, complete coverage of each tool's full range of capabilities and extensive treatment of problem determination techniques are beyond the scope of this book.

A novice can use the information in this book, but the novice or even an experienced LPAR architect would also want to have available all or most of the documents listed in the previous topic.

In this chapter, the following topics provide an overview level look at the LPAR process of defining, ordering, and installing (deploying) logical partitions on an IBM System i or IBM System p configuration.

Appendix A, "Managing System Planning Tool Version 2 installation" on page 189, contains information about ordering SPT Version 2.0 and other SPT management considerations.

1.2 The LPAR planning process

The process involved in planning for a system with logical partitions can get quite involved. If the system has not been ordered, there should be steps that collect requirements for the new system as well as steps to select hardware, software, and which partitions will be created with what resources. Then the plan should be ordered and, after the hardware arrives, the plan should be deployed. The whole process is shown in Figure 1-1.

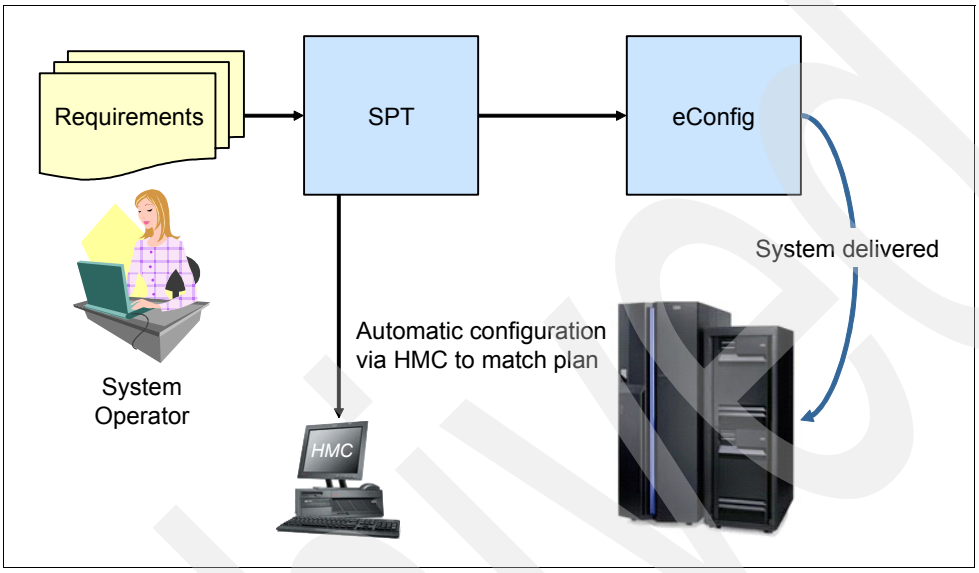

Figure 1-1 The logical partition planning process

If the system already exists in the data processing center, only a subset of these tasks need be done. The partitions can be reworked to handle new workloads and tasks. Then the new plan can be deployed on the existing system hardware. Here we examine each step of the process with emphasis on how this book covers the new function involved.

1.3 Gathering requirements

Performance and size requirements can be estimated from the tool called the IBM Systems Workload Estimator (WLE). With Workload Estimator i5/OS, AIX, and Linux partitions can be created and tasks, also known as workloads, can be assigned. Now Workload Estimator can be linked to from inside the System Planning Tool. When this process is selected much of the Workload Estimator's output recommendations, such as the type and model of system, processor, and memory allocation for the sized partition, can be carried back into the System Planning Tool.

This process is described in detail in Chapter 5, "System Planning Tool Version 2 and Workload Estimator" on page 149.

1.4 Planning partitions and allocating hardware

The System Planning Tool is the tool we use to define a system. The system's hardware can be defined in particular card positions within the system or expansion unit, and that hardware can be allocated to particular partitions. This book covers new function in version two of SPT.

New or improved functions in SPT version two include:

- Support for new system units and expansion units: As is standard practice when new hardware is announced, the SPT will have an update made available that supports that hardware. It happens that this time the initial announcement of new POWER6 6 hardware is covered by SPT with the release of the improved version two of the tool.
- Support for partitions using a Virtual I/O Server environment: The virtualization of both storage and communications resources for AIX and Linux partitions on System p systems is done through a special partition using a Virtual I/O Server as the operating system, also frequently referred to as a *VIOS server*. Chapter 4, “Virtual I/O Server” on page 103, covers the creation and deployment of these new partitions.

  One or more Virtual I/O Server partitions can be defined and deployed. Virtual I/O Server partition deployment also includes installing the Virtual I/O Server software itself. This topic is covered in Chapter 4, “Virtual I/O Server” on page 103.
- Support for Integrated Virtualization Manager: IVM is a specialized version of a Virtual I/O Server partition. IVM handles partition configuration on selected IBM System p, IBM eServer p5, and IBM OpenPower™ systems without the need for the Hardware Management Console (HMC). The latest version of VIOS, 1.4.0.0, adds a number of new functions, such as support for dynamic logical partitioning for memory and processors in managed systems, task manager monitor for long-running tasks, security additions such as viosecure and firewall, and other improvements.

  Although IVM with VIOS 1.3.0.0 is very functional, its capabilities remain a subset of the full range of LPAR capabilities offered with an HMC.

  An IVM partition is now deployable. See Chapter 4, “Virtual I/O Server” on page 103.
- New System Plan Viewer: The system plans created in SPT can now be viewed with a new System Plan Viewer. This viewer is both part of the SPT tool and part of the Hardware Management Console. The viewer allows you to view and print complete or detailed views of the system’s partitions and hardware. This topic is covered in multiple sections of the book, where viewer function is applicable.
- Improved usability design: The new SPT interface has new tabbed navigation, parts list tree expansion, and incremental validation. SPT can also add or remove multiple expansion units at once, and has improved ways to define virtual consoles, LANs, and SCSI connections. These improvements are covered in Chapter 2, “System Planning Tool (SPT) V2” on page 11.

1.5 Ordering the system

The SPT can create a configuration file (.cfr) that can be imported to the IBM sales configurator tool (sometimes referred to as e-Config) as input to the next step — the ordering process. Using SPT to create the configuration files also minimizes the need to make changes when deploying partition configurations on the delivered system.

Important: Starting with a system plan generated by SPT is the only way to enforce customer placement of hardware for your partitions as part of the IBM manufacturing system build process.

Recall that using the IBM Sales Configurator tool with an initial system order, you order the supported operating system software you want for your partitions. In the IBM Sales Configurator you can specify that either i5/OS or AIX software be installed. Examples of available order features include:

- #0140 Logical Partitioning Specify: #0140 is used to specify that this system is to be logically partitioned. The #0140 is only valid on n-way processors with OS/400 V4R5 or later. The IBM sales configurator adds a quantity of one #0140 to the order for each logical partition (LPAR).

  This is supported on System i models 520, 525, 550, 570, 595, 800, 810, 820, 825, 830, 840, 870, and 890. 0140 is a Customer Install Feature.

- #0145 AIX Partition Specify: #0145 is used to specify that this system is to be logically partitioned with an AIX partition. A quantity of one #0145 is required on the order/inventory records for each AIX partition required.

- #0454 - LPAR Partition Initialization: #0454 configures a partition and assigns the correct resources as specified by the customer. This feature is only available for Models 570 and 595. #0454 is not a Customer Install Feature.

- #0455 - LPAR operating system Preload: #0455 specifies to preload the operating system (i5/OS or AIX 5L™) specified by the client for a partition configured via #0454 LPAR. This is supported on Models 570 and 595. #0455 is not a Customer Install Feature.

- #0496 - Force i5/OS Preload: #0496 preloads i5/OS on a new server. The #0496 forces a preload of i5/OS on a single partition when Linux or AIX 5L partitions with virtual storage are on the order. i5/OS is preloaded on all of the disk drives in the configuration. Do not use this feature if the Linux or AIX 5L partition has dedicated disk controllers and drives in the on order configuration. #0496 is mutually exclusive with a #0006.

  Supported on Models System i 520, 525, 550, 570, and 595 models. #0496 is not a Customer Install Feature.

- #0006 LPAR Restrict Build Process: #0006 is added to an initial order where LPAR #0140 is requested. This #0006 instructs manufacturing to only load SLIC on the minimum number of disk drives.

Creating the configuration file (.cfr) for the configurator is covered in Chapter 2, “System Planning Tool (SPT) V2” on page 11.

1.6 System plans on the HMC

In addition to the configuration file mentioned in 1.5, “Ordering the system” on page 7, SPT can create an output file, which is called a *system plan*. When the system plan file, denoted by the .sysplan file extension, is on the Hardware Management Console, it can be used to deploy the partition configurations defined in the sysplan on the installed system. This file can be exported to media and used by an HMC to deploy the partition configurations stored within the sysplan file.

The association between the SPT configuration file used for a new system order and the sysplan file being used to deploy the partition configurations on the installed system offer significant advantages in getting your partitioned system up and running within a minimumized time period.

The HMC itself can create a system plan file for the system it is actively connected to when one or more partitions are already defined. If the partition is active running an operating system, the created system plan file will have more details than if the partition (operating system) is not active. For example, with an i5/OS partition active, the created system plan will include the disk device attachment information for its controlling IOA (controller).

The system plan can be viewed on the HMC or on a workstation where SPT is installed.

The HMC also has system plan management functions that allow system plans to be imported, exported, and deleted, as well as creating and deploying system plans.

You do not have to export the sysplan file to media. You can also send a system plan via FTP or import it via HTTP on running on another HMC. Chapter 3, “System plans and the hardware management console” on page 65, contains more information about management of system plan files.

Note: This publication uses the Version 7 Release Level 3.1.0 code level of the IBM Hardware Management Console, which provides a completely new browser-based interface to the HMC user from earlier V5 and V6 releases of HMC code. A remote workstation, using the appropriate URL to the V7R3 HMC can link to the HMC without using the IBM WebSM product on the remote workstation. An example of such a URL would be:

```
https://HMC hostname or
```

V7R3.1.0 or later is required to manage both POWER5 and POWER6 technology systems. Information about the new HMC V7R3 interface capabilities can be found at:

- IBM Systems Hardware Information Center

  http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp
- IBM Support Web site for the HMC at:

  https://www14.software.ibm.com/webapp/set2/sas/f/hmc/v7310notice.html
- IBM Redbooks publication *Hardware Management Console V7 Handbook*, SG24-7491

1.7 Virtual I/O Servers

An important part of AIX and Linux partitions sharing resources is the Virtual I/O Server, sometimes also referred to as a VIOS server. A Virtual I/O Server partition owns disk, Ethernet, and SCSI devices and can be configured to share those resources with multiple AIX and Linux partitions. The disk devices can be SAN-attached disk drives appropriately defined on an IBM System Storage™ Disk system, such as one of IBM DS6000™ or DS8000™ series configurations.

Multiple Virtual I/O Server partitions can be defined. Virtual I/O Server partitions can be defined, each with its own set of disk, Ethernet, and SCSI *devices* or as a backup Virtual I/O Server for another Virtual I/O Server partition.

For details on Virtual I/O Server creation, management, and deployment, see Chapter 4, “Virtual I/O Server” on page 103.

1.8 Virtual LAN and virtual SCSI definition enhancements

The System Planning Tool now has a simpler interface to create and allocate virtual LAN and virtual SCSI connections between partitions. The user now creates a table where each column is a shared virtual connection. The defined partitions make up the rows of the table. Boxes are checked or unchecked to include or exclude the partition in the virtual LAN or SCSI connection. For details on virtual connection creation and allocation, see 2.6.5, “Networking tab” on page 44.

2

System Planning Tool (SPT) V2

In this chapter we discuss the features, functions, capabilities, navigation, enhancements, and use of the System Planning Tool Version 2.

2.1 Overview of the System Planning Tool V2

The IBM System Planning Tool (SPT) V2 is an enhanced tool, replacing SPT V1, for designing new logically partitioned System i and System p systems. It can also be used for planning non-partitioned systems and documenting existing systems. The SPT is a browser-based tool that runs on your PC. The GUI and order of operations are quite different from SPT V1, but its purpose is the same. There is help text within the tool and a link to the eServer Hardware Information Center.

Users can design new systems from existing performance data or new workloads in the Workload Estimator (WLE) from sample systems provided in SPT and by using the advanced mode, which lets you design the system at the component level.

SPT creates a system plan that is saved as a .sysplan file. That system plan may be just one system or it may contain multiple systems, each with a unique system name.

The output of the SPT can be used to create a report or as input to the IBM configuration tool (eConfig) for order processing. The report function of SPT invokes the System Plan Viewer, which has a print option. You will also be able to use the .sysplan file to automatically create and deploy partitions on a Hardware Management Console (HMC).

2.2 Downloading SPT

Like SPT V1, SPT V2 is available for download from the following Web site:

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

A subscriber list is used to notify users when a new version is available. To be included on the subscriber's list send an email to rchspt@us.ibm.com.

The first time that you download the SPT you must use the full version, which includes the required JVM™ code and other support files. Once the full version is installed the current update file should then be downloaded and installed. For subsequent updates of the SPTit is only necessary to download and install the update. In either case when you run the .exe file, an install wizard is initiated to guide you through the installation. An icon for the SPT will be placed on your desktop when the installation is complete.

The full SPT installation information can be found in Appendix A, "Managing System Planning Tool Version 2 installation" on page 189.

2.3 What is new in SPT V2

SPT V2 contains extensive navigation changes and functional enhancements when compared to SPT V1. Below is a list of those to be covered in this book. Also, in the time that has passed since the V1 *LPAR Simplification Tools Handbook*, SG24-7231, was published there have been significant changes and enhancements to SPT V1 functions, which have been included in SPT V2. They are included in the section on new SPT V2 items.

2.3.1 Version 2 new items

The following items are new in SPT V2:

- Complex Virtual I/O Server configurations can be designed and deployed for System p systems including:
  - Multiple Virtual I/O Server partitions (for HMC-managed systems only).
  - Shared Ethernet Adapters (including failover) - EtherChannels.
  - SAN volumes and simple SAN volume multi-path I/O (MPIO).
  - Support of new Host Ethernet Adapters (HEA) on POWER6 models. On 9406 MMA models HEA may also be referred to as Integrated Virtual Ethernet (IVE).
  - Storage pools and simple storage pool mirroring.
  - Assignment/hosting of specific storage resources owned by Virtual I/O Server partitions to client partitions.
  - Virtual slots 2–10 not usable (reserved for future use) for VIrtual I/O Server partitions on POWER6 systems.
- Deployment Planning & Configuration of Integrated Virtualization Manager (IVM) managed partitions: Automated configuration of partitions from SPT system plan via deployment wizard on the IVM partition
- System Plan Viewer:
  - Ability to view and print report of a sysplan generated by mksysplan on HMC
  - Graphical representation of system hardware
  - Report of the configuration can be shown either in HTML or plain text when launched within SPT
  - Ability to print one or all systems in a system plan
- Enhanced ease of use:
  - Tabbed navigation between towers and system units
  - Tabbed navigation between different functional areas of a system: system, partitions, hardware, storage, networking, consoles, and summary
  - Ability to switch between and edit different systems in a plan in the edit view
  - Incremental validation to improve performance
  - Multiple add of partitions by OS type
  - Multiple add of expansion units
  - Multiple remove of expansion units
  - An Apply function (use frequently)
  - Subsequent adds of cards to which you must assign an LPAR (Remember that the last LPAR name used.)
  - Multiple add of virtual SCSI connections
  - Edit virtual slots for VLAN and virtual SCSI connections
  - VLANs related to partitions by use of check box grid
  - Automated creation of virtual Ethernet adapters to support VLANs
  - Virtual SCSI connections mapped between partitions as they are created
  - Partitions table in summary shows total partitions instead of only the first partition

Attention: SPT 2.0 will be the last release that will support .lvt (LPAR Validation Tool file) and .xml files. Customers should load their old .lvt and .xml plans and save them as .sysplan files.

2.3.2 Version 1 new items

The following lists SPT V1 functional enhancements available after *LPAR Simplification Tools Handbook*, SG24-7231, was published:

- Ability to copy expansion units, including the contents, within a system or to another system in the same system plan
- Hide validation messages
- Tabbed view for towers
- Expand and collapse disk drive bays
- Parts list tree expansion and selection remembered when switching between towers

2.4 System Planning Tool basics

In this section we review the basics of using the SPT to create a system design. The four design methods are discussed, but we show an example of using the advanced method for creating new a system.

When your SPT is started you are presented with the Launch SPT window. This may seem superfluous, but it is important. Separate SPT sessions are launched into new windows from this launch and these new windows do not have the address bar, links, and menu bar that the launch window has. Note also that there is an SPT icon placed in your system tray.

The launch window allows you to close the window you are working in and re-launch the SPT. If you also close the launch window the system tray icon can be used to restart the SPT without having to reload it. To entirely close the SPT you must exit from the tray icon.

Note: If you are working on a system plan and close the browser before you have completed the steps and finished, your partially completed work will not necessarily be saved. There is a reconnect feature built into SPT for exactly this purpose. So if you were working on a system plan and inadvertently closed the browser window, you could reconnect to the session and salvage some of your work. The reconnect function restores the system plan back to the state it was in when the Work with Systems panel was last displayed. Therefore, the only case when all work is lost is if the user was in the middle of creating the only system in a system plan and was still in the wizard when the browser was closed.

Figure 2-1 shows the first window displayed after starting SPT. Click **Launch** and the Getting Started window will open (Figure 2-2). You may launch multiple system plans from the Launch window by repeatedly coming back to this window and clicking **Launch**. You are responsible for ensuring that you are working on the correct plan.

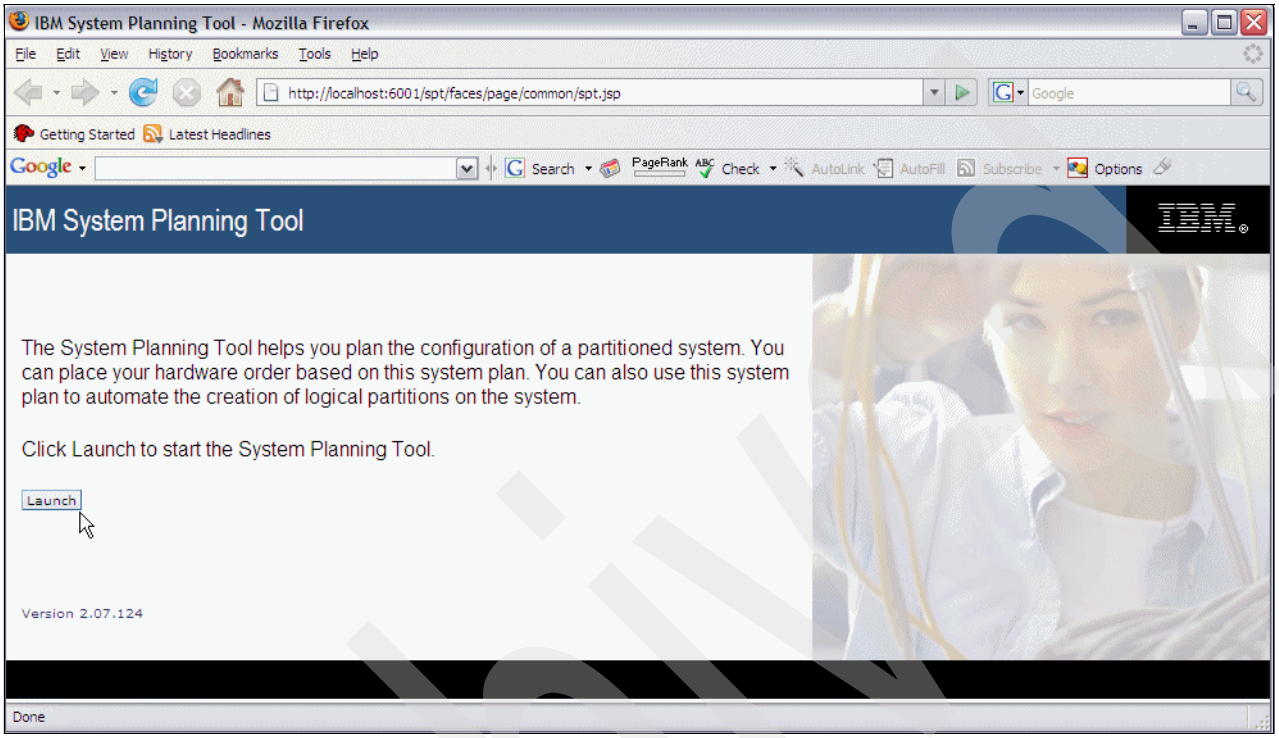

Figure 2-1 Launching the SPT

The Getting Started window initially has only two options. If you have created a system plan and closed the SPT window and then used Launch again, the Getting Started window will have a third option called Reconnect to open plan. In our example we click **Create a new system plan**, as shown in Figure 2-2.

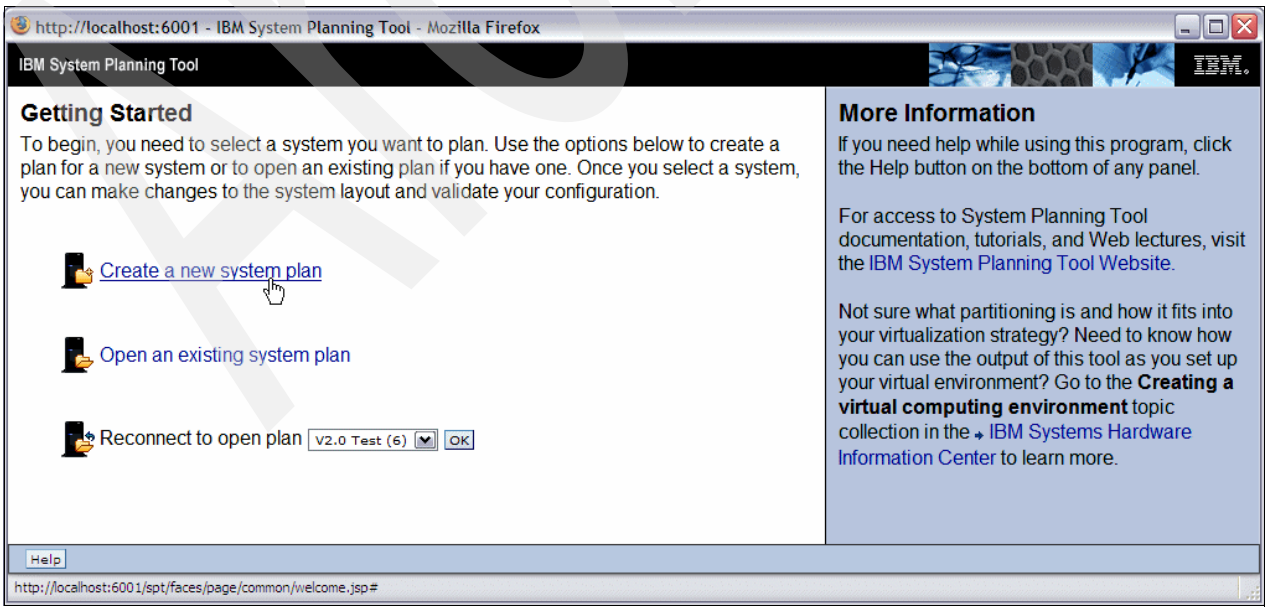

Figure 2-2 Getting Started

The Create a New System window opens and there are four options that allow the user to create a new system. This window is the same as it has been with SPT V1.

Before going through creating a new system plan we first discuss the reconnect option.

2.4.1 The reconnect option

As stated, the reconnect function is visible any time you close the browser window without properly closing the SPT session. You close the SPT session by choosing the **Close** action from the action bar on the Work with Systems panel.

The Reconnect selection contains a drop-down list showing all SPT sessions that have not been properly closed since the SPT application was started. The session names used are the names of the system plans that you were being editing.

Note: If a user is in the middle of editing a system plan (the plan is currently open in a browser window) and the user launches a new SPT session, the first plan you are editing will show up in the reconnect list, because SPT does not know whether you have closed the browser window. It just sees an active session.

The user should not reconnect to a plan that is open in another browser window, because the user will be updating the same plan in the same session from two different browser windows. If this happens, the system plan will get damaged.

2.5 Creating a new system plan

While your goal is to design a system, you are actually creating a new system plan and then selecting what systems you want to put in it. A system plan can contain multiple systems. You can and should name your system plan something unique, and you may add notes or comments in the Description field. There are four options for creating a new system:

- Option 1, Create a new system (advanced): This option provides for the user to select a system model and every detail of the system from type/model to processor and memory capacity and all the expansion units, drawers, adapters, and disk storage. Nothing is included automatically.
- Option 2, Based on IBM supplied sample systems: This option presents you with sample systems that have been included within the SPT. You can select one of these as a starting point for the system you intend to design. These systems have sample partition configurations and all the hardware included and I/O correctly positioned. These are intended to be used as starting points only and should be modified to meet your requirements.

  The last two options use IBM Workload Estimator (WLE). You must enter WLE from the SPT for these to work.
- Option 3, Based upon existing performance data: This option enables creation of a system based upon performance sizing output using the WLE. The sizing output is primarily based upon summarized performance data previously collected by the IBM service offering and product known as IBM PM Web Sales Tool. This tool, formerly known as PM for System i, PM for iSeries, or PM/400, transmits summarized performance data collected from an active i5/OS, AIX 5L, or supported POWER™-based Linux distribution partition to a centralized IBM repository that can be accessed by a browser interface to the IBM Web Sales Tool.

The performance data from the Web Sales Tool for each system or partition must have been previously exported to a file that can be accessed in your WLE session.

- Option 4, Based upon new workloads I want to run: This option enables creation of a system based upon performance sizing output using the WLE. The sizing output is based upon any new workloads that you define under WLE that do not include summarized performance data previously collected by the IBM PM Web Sales Tool.

In this publication we show examples of options 3 and 4 in the Chapter 5, “System Planning Tool Version 2 and Workload Estimator” on page 149.

Select **Create a new system (advanced)**, as shown in Figure 2-3, and click **Next**.

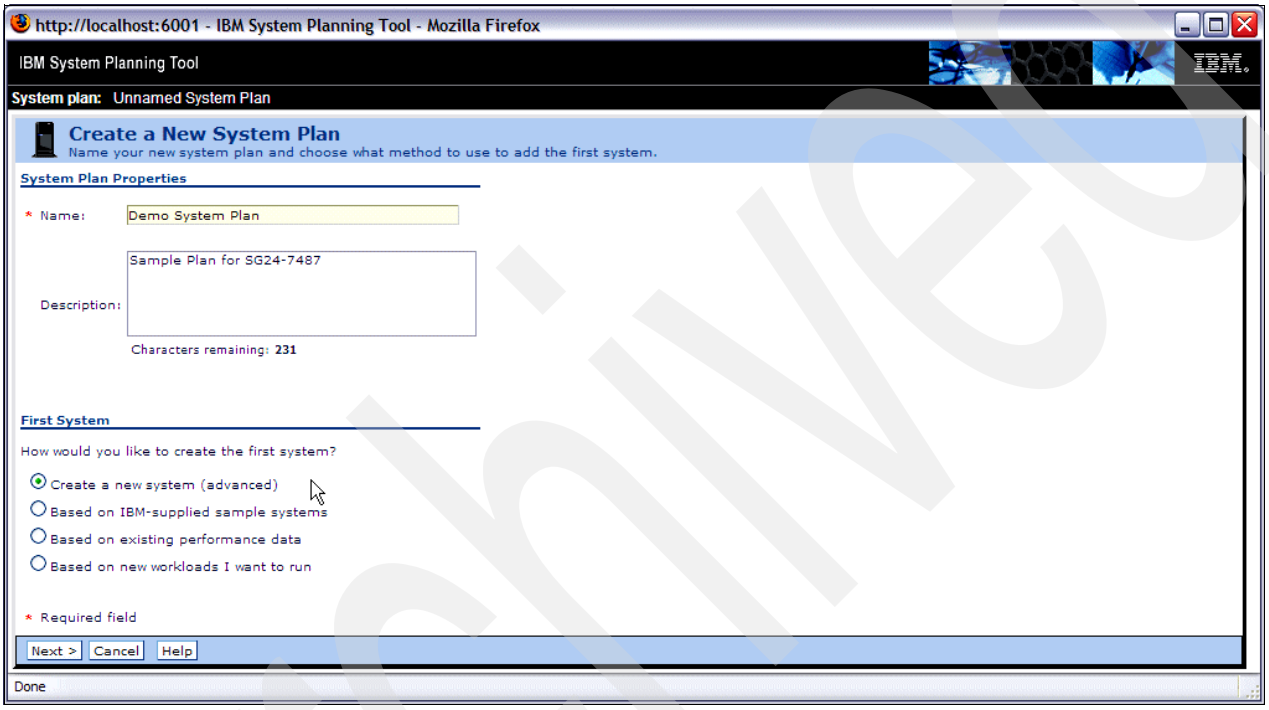

Figure 2-3 Create a new system plan

2.5.1 Selecting the system platform and type/model

The System Type window in Figure 2-4 allows you to name the system plan and select the platform, either a System i or System p. The sidebar at the left side keeps track of where you are in the process, but you cannot access anything from it. It can be minimized by clicking the twisty in the corner of the sidebar.

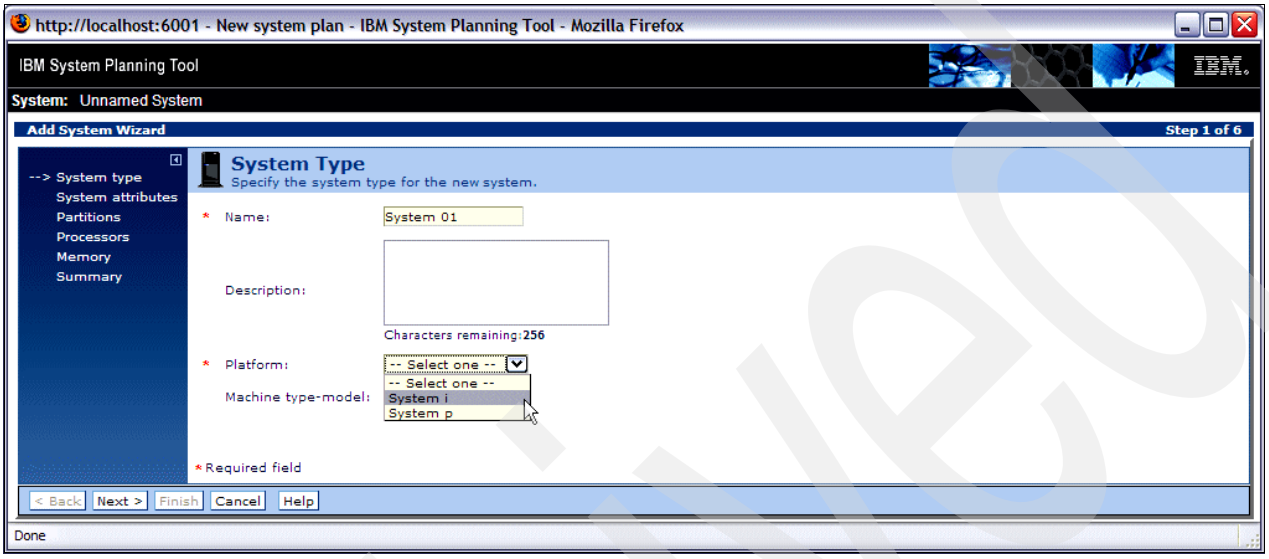

Figure 2-4 Selecting System i or System p

HMC is the default management interface for all systems in SPT. If you select System p, some of the low-end models also allow you to select the Management Interface type of either HMC or Integrated Virtualization Manager (IVM) (shown in Figure 2-5). See Chapter 4, “Virtual I/O Server” on page 103, for more information about IVM.

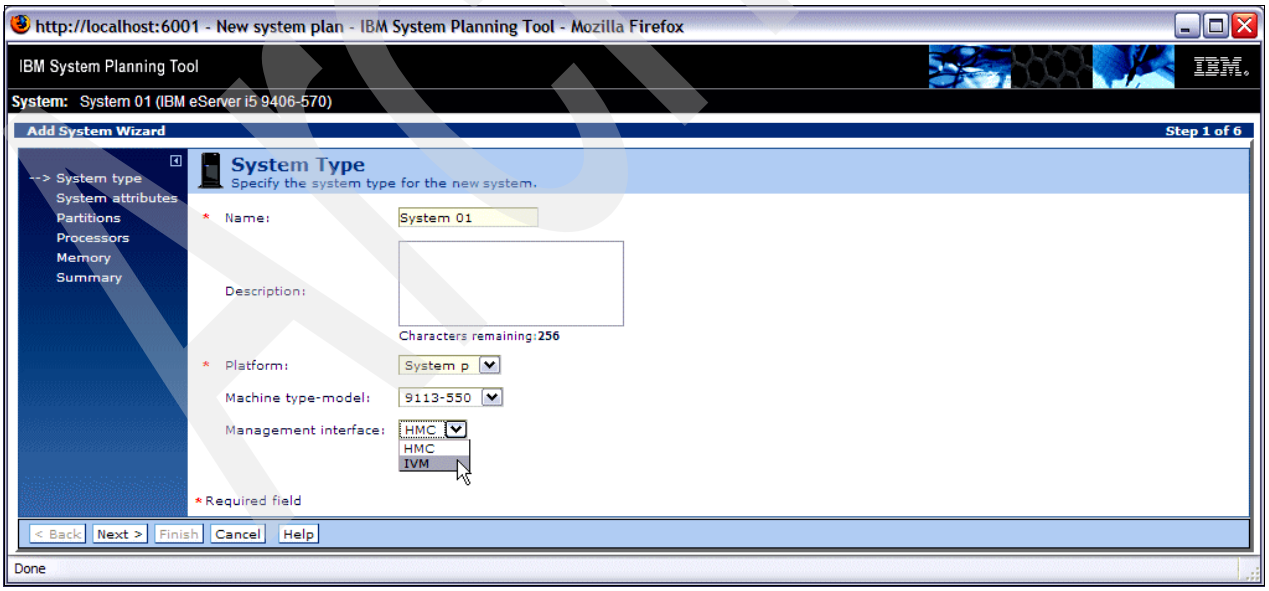

Figure 2-5 Selecting Management Interface on System p

After selecting the platform and the management interface, click the type/model drop-down box to select the particular machine that you are going to design. See Figure 2-6. Select the type/model and click **Next**.

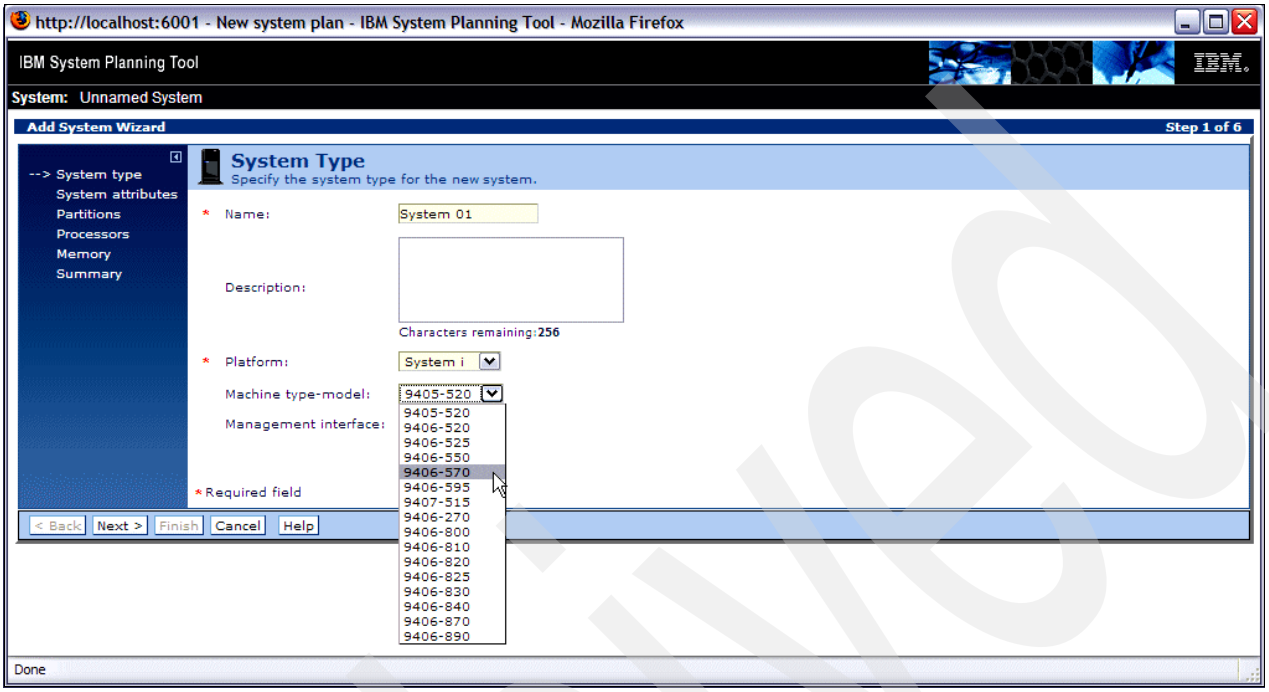

Figure 2-6 Selecting type/model

In our example we select a System i (9406) model 570. As you see in Figure 2-7 on page 20, this is a 4-way 570 with two processors initially activated.

2.5.2 Working with system attributes

Clicking **Next** brings up the System Attributes window. Here you select the processor, server, and edition features, as well as the total system memory and the number of active processors.

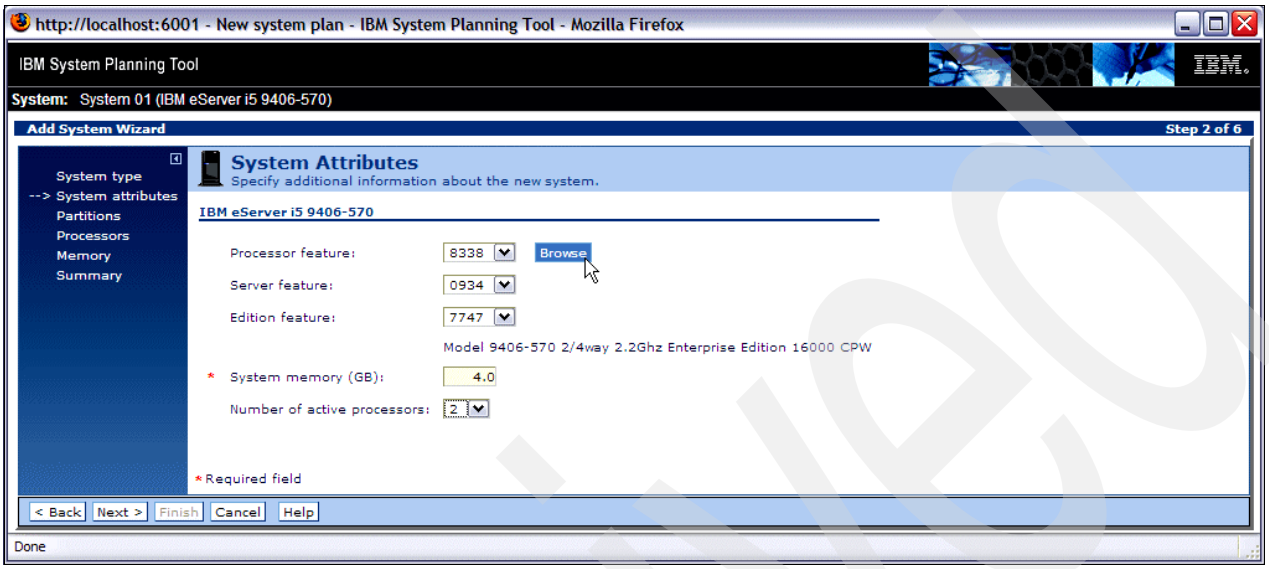

Figure 2-7 System Attributes

Click **Browse** for the system features (processor, server, and edition). If you use the drop-down boxes on the system features you will see the features that you can select, but they have no description. Use the browse function to see all of the features and their descriptions and make your selections from the table shown in Figure 2-8. Then scroll down and click **OK**.

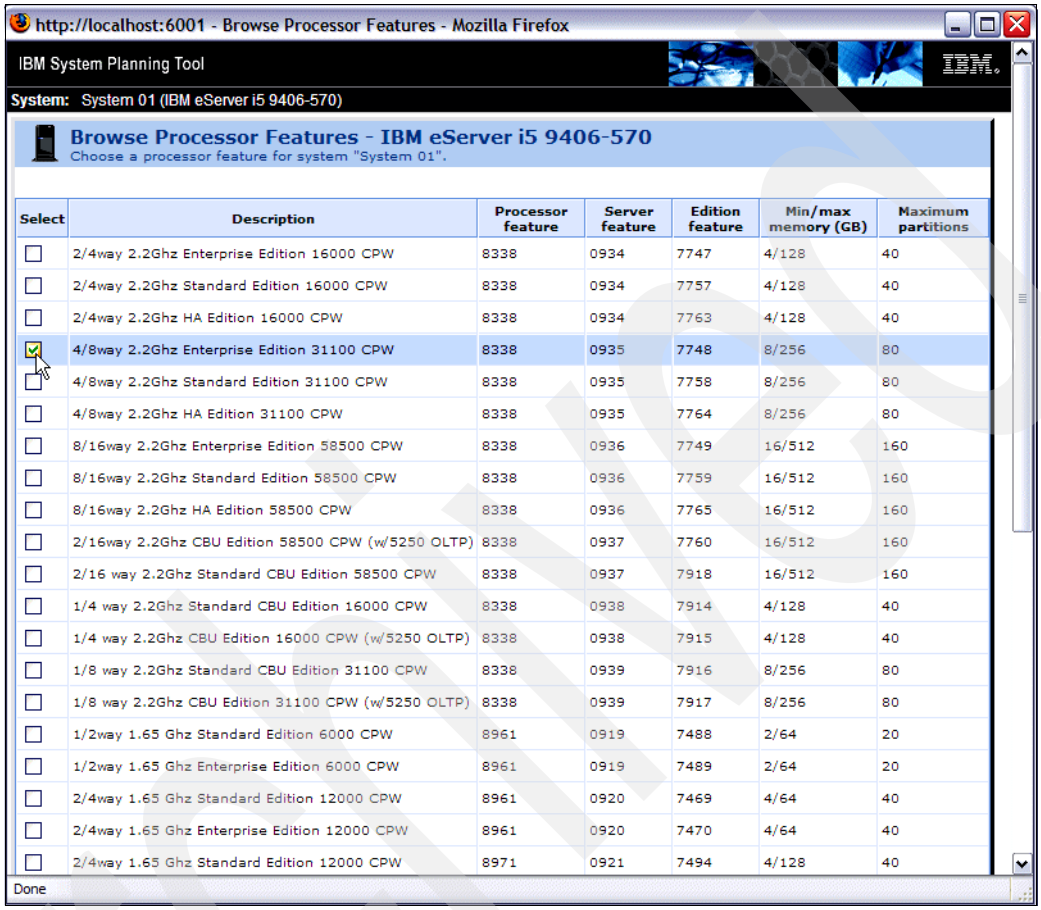

Select	Description	Processor feature	Server feature	Edition feature	Min/max memory (GB)	Maximum partitions
☐	2/4way 2.2Ghz Enterprise Edition 16000 CPW	8338	0934	7747	4/128	40
☐	2/4way 2.2Ghz Standard Edition 16000 CPW	8338	0934	7757	4/128	40
☐	2/4way 2.2Ghz HA Edition 16000 CPW	8338	0934	7763	4/128	40
☑	4/8way 2.2Ghz Enterprise Edition 31100 CPW	8338	0935	7748	8/256	80
☐	4/8way 2.2Ghz Standard Edition 31100 CPW	8338	0935	7758	8/256	80
☐	4/8way 2.2Ghz HA Edition 31100 CPW	8338	0935	7764	8/256	80
☐	8/16way 2.2Ghz Enterprise Edition 58500 CPW	8338	0936	7749	16/512	160
☐	8/16way 2.2Ghz Standard Edition 58500 CPW	8338	0936	7759	16/512	160
☐	8/16way 2.2Ghz HA Edition 58500 CPW	8338	0936	7765	16/512	160
☐	2/16way 2.2Ghz CBU Edition 58500 CPW (w/5250 OLTP)	8338	0937	7760	16/512	160
☐	2/16 way 2.2Ghz Standard CBU Edition 58500 CPW	8338	0937	7918	16/512	160
☐	1/4 way 2.2Ghz Standard CBU Edition 16000 CPW	8338	0938	7914	4/128	40
☐	1/4 way 2.2Ghz CBU Edition 16000 CPW (w/5250 OLTP)	8338	0938	7915	4/128	40
☐	1/8 way 2.2Ghz Standard CBU Edition 31100 CPW	8338	0939	7916	8/256	80
☐	1/8 way 2.2Ghz CBU Edition 31100 CPW (w/5250 OLTP)	8338	0939	7917	8/256	80
☐	1/2way 1.65 Ghz Standard Edition 6000 CPW	8961	0919	7488	2/64	20
☐	1/2way 1.65 Ghz Enterprise Edition 6000 CPW	8961	0919	7489	2/64	20
☐	2/4way 1.65 Ghz Standard Edition 12000 CPW	8961	0920	7469	4/64	40
☐	2/4way 1.65 Ghz Enterprise Edition 12000 CPW	8961	0920	7470	4/64	40
☐	2/4way 1.65 Ghz Standard Edition 12000 CPW	8971	0921	7494	4/128	40

Figure 2-8 Browse and select Processor Features

After selecting the processor features, you are taken back to the System Attributes window to select memory and active processors. Key in the memory amount in GB and use the drop-down box to specify the number of active processors, as shown in Figure 2-9. We select 8 GB of memory and five5 active processors.

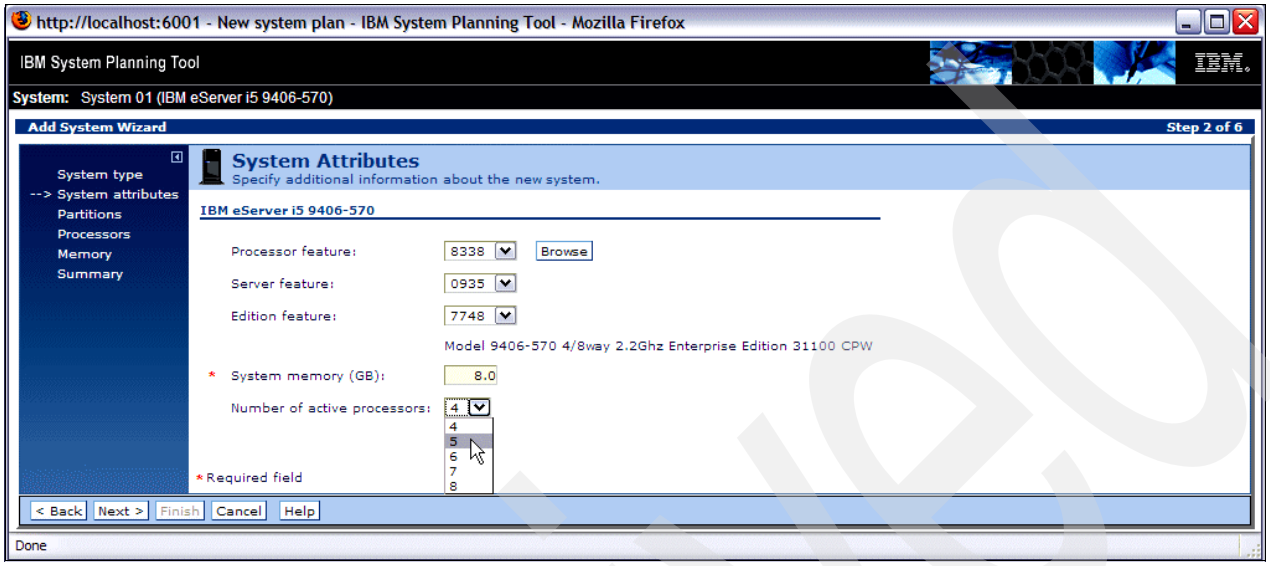

Figure 2-9 Selecting Memory and Processor quantities

Click **Next** and the System Partitions window appears.

2.5.3 Defining system partitions

In the System Partitions window (Figure 2-10 on page 23) you can add multiple LPARs by OS type. Note that it shows the number of partitions remaining available to add based on earlier system selections in Figure 2-8 on page 21. The partitions are given default names (LPAR1, LPAR2, and so on) as seen in Figure 2-11 on page 23.

Important: We recommend that you use the IBM Systems Workload Estimator (WLE) or other performance and capacity planning tools for obtaining the CPW information for a partition. Other sizing and capacity planning tools include BMC PATROL Predict for System i and MPG Performance Navigator. Experience shows that users are designing systems based on the CPW value shown by SPT, expecting that amount of CPW to be available to their operating system and workload. This is just an estimate and might work all right for bigger CPW allocations, but for small processor allocations like 0.2 processing units, there can be significant shared processing overhead incurred. WLE uses *tax tables* to account for that based on the system and the partitions and partial processor allocation that you are doing. SPT *does not* apply this tax in its estimations.

If you are creating a System p design the Virtual I/O Server (VIOS) would also appear as an OS type. See Chapter 4, “Virtual I/O Server” on page 103, for information about Virtual I/O Server.

On System i i5/OS can host (serve) AIX or Linux partitions as clients. On System p a Virtual I/O Server partition can host (serve) AIX or Linux partitions as clients. Notice that in SPT V2 there are no longer AIX-hosted or Linux-hosted OS typed options as there were in SPT V1. These options specified that some or all of these partition resources were going to be hosted (served) by an i5/OS partition. Therefore SPT V1 would not check for minimum requirements

for the hosted (client) partitions except for the tape requirement for an AIX-hosted (client) partition.

SPT V2 now always checks for normal minimum partition requirements unless you define a virtual SCSI adapter for a partition you intend to be a client. Then the client partitions get treated as *hosted*.

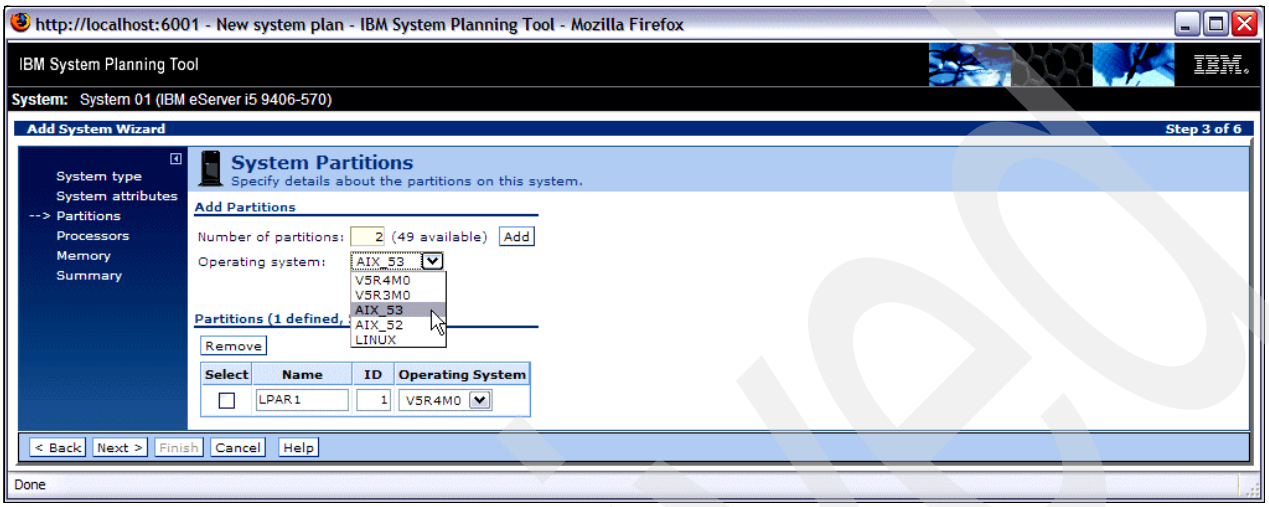

Figure 2-10 Adding multiple partitions by OS type

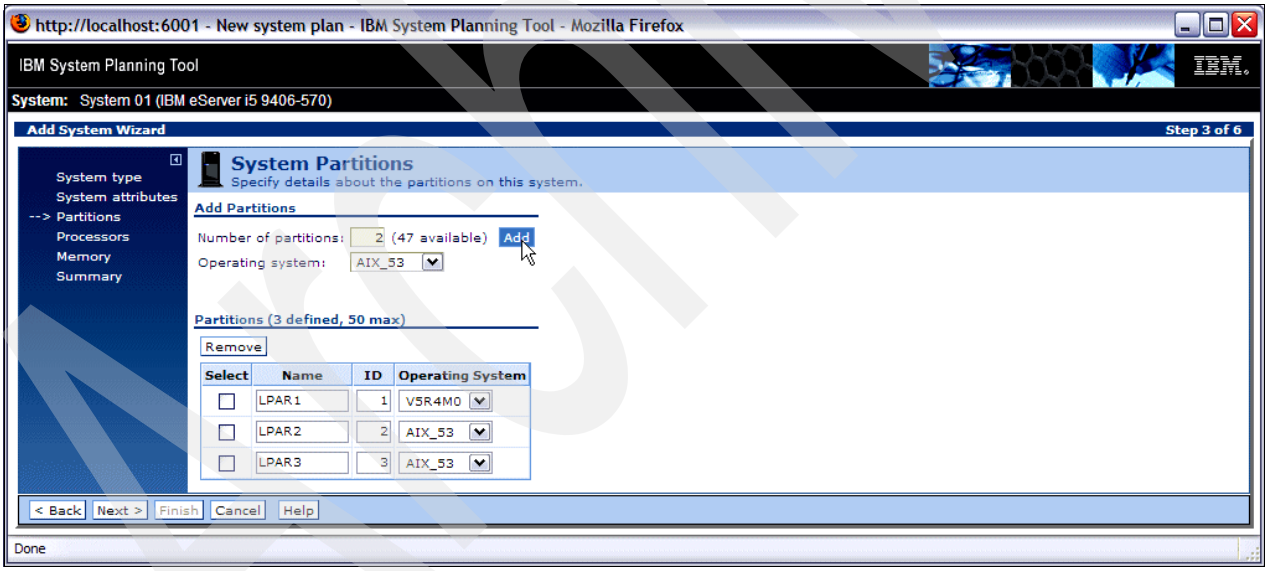

Figure 2-11 Default partition names

You can change the partition names, as our example in Figure 2-12 shows (LPAR1, First AIX, and Second AIX partition names), and at this point you can also change the OS in each partition. Once you use the Finish function (shown later) you will no longer be able to change the partition OS type. Partitions can also be removed by selecting one or more check boxes to the left of a partition name and clicking **Remove**. This removes the partitions and everything you may have placed in them. There is no undo function after doing a remove. You can always use the Back button to start over.

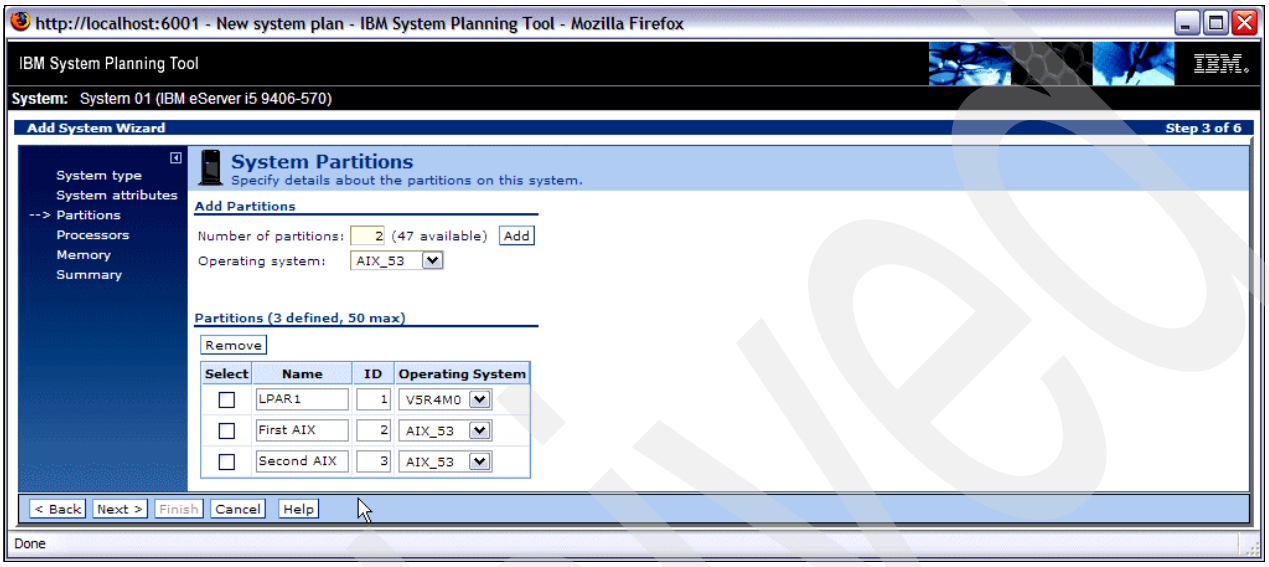

Figure 2-12 Changing partition names and OS type

Click **Next**. The Processors window (Figure 2-14 on page 26) is shown. Notice that the sidebar has been minimized. There is quite a lot of detailed information here. This is a good opportunity to use Help, which is *window sensitive*. Click **Help** to see definitions of the data elements, as shown in Figure 2-13. Scroll through the help information to see an explanation of all of the data and fields in the Processors window.

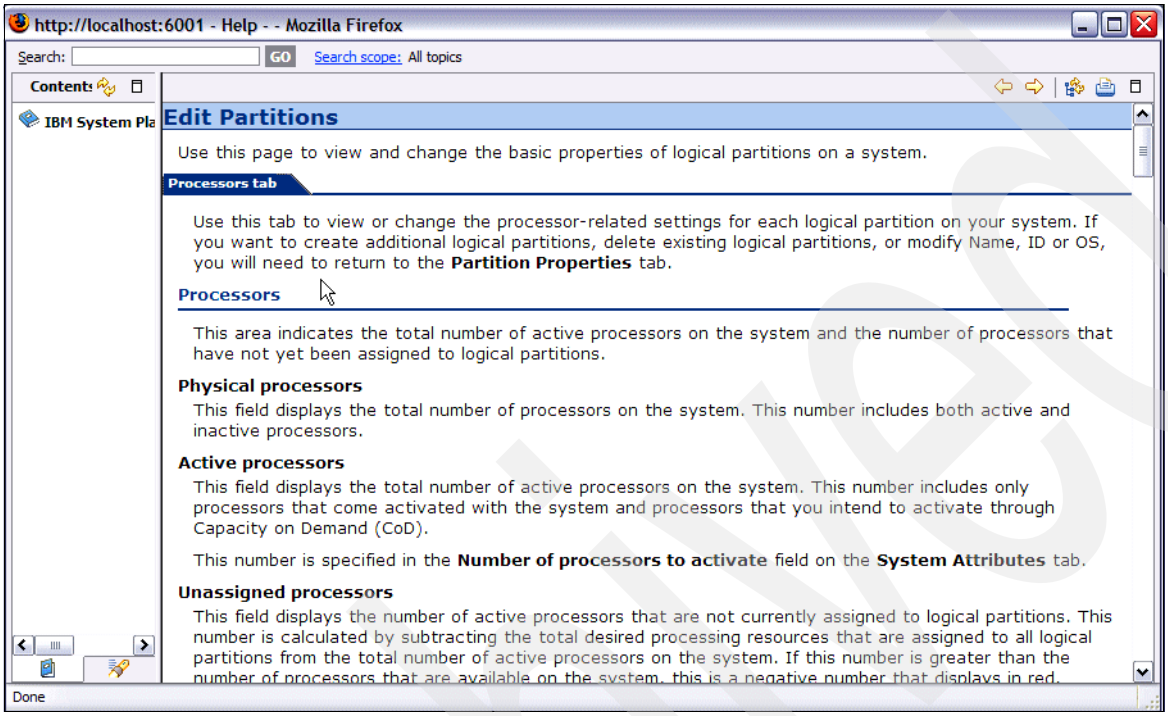

Figure 2-13 Help for Edit Partitions

In the Processors for Partitions portion of the window some default information is put in. You must change all of that data per your specifications. In Figure 2-14 the processors are over committed because we have specified six desired processors when there are only five active in this example. Therefore a negative number is in the Unassigned processors field and the shared processors number turns red showing the total over commitment. Also note that in the Licenses box there are asterisks (*) indicating that a recalculate should be done. To make things valid you would change the processor assignments to fit your requirements and eliminate the over commitment. SPT can only presume that all partitions will be running at the same time, so it has no choice other than to cite over commitment.

Attention: Best practices recommend that you *not* select a quantity of virtual processors greater than two times the actual assigned desired processing units. Be aware that with multi-threading that would equate to four threads per physical processor.

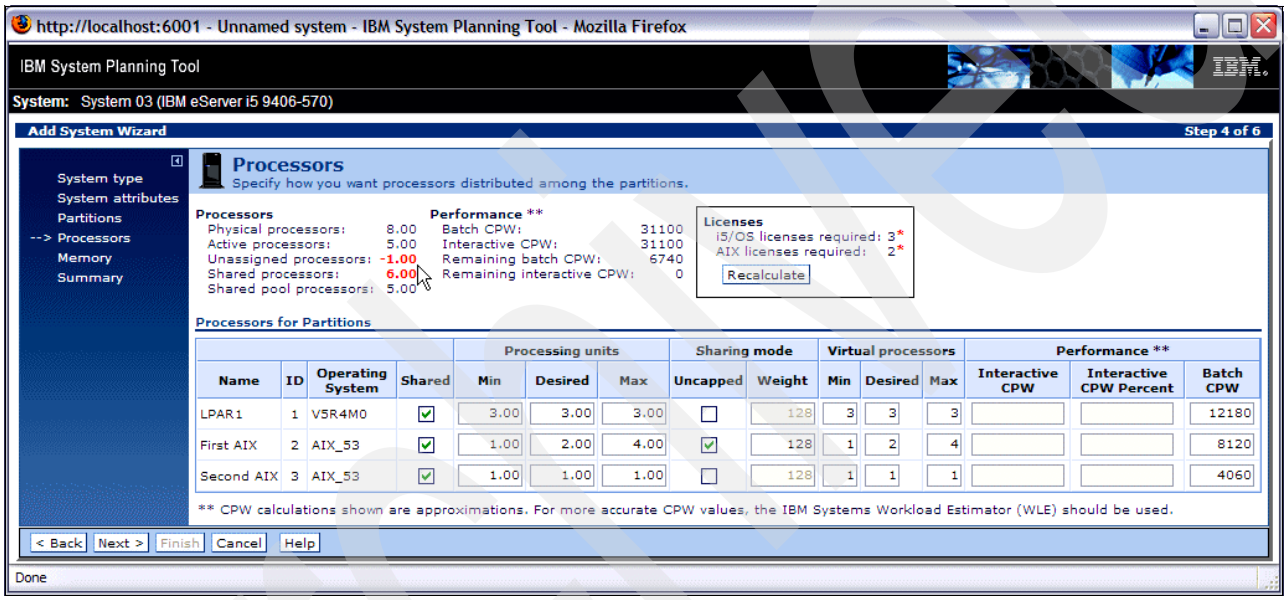

Figure 2-14 Partition processors and OS licensing

When you have manipulated everything needed in Figure 2-14 click **Next** to continue.

Figure 2-15 shows an example of errors listed when you try to go to the next step but have improper data specified. You are not permitted to advance to the next window until you resolve the errors.

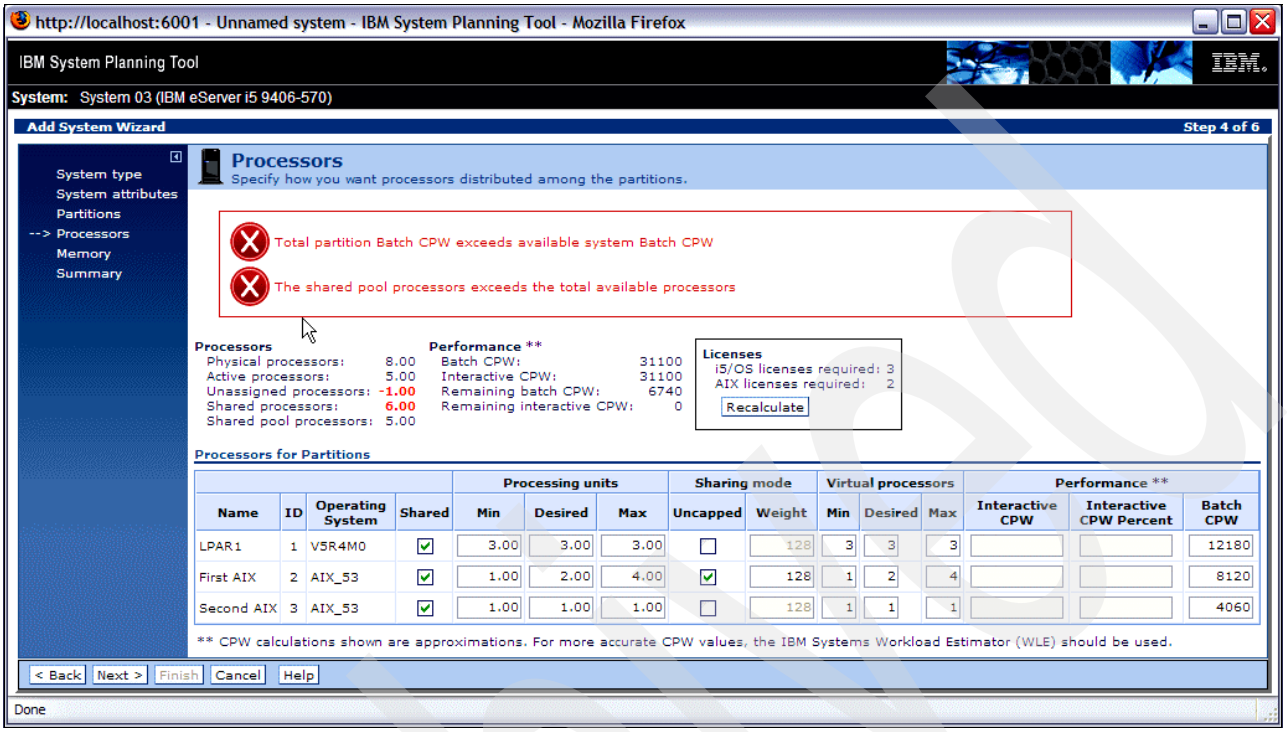

Figure 2-15 Errors regarding invalid data

Figure 2-16 shows errors resolved. Click **Next**.

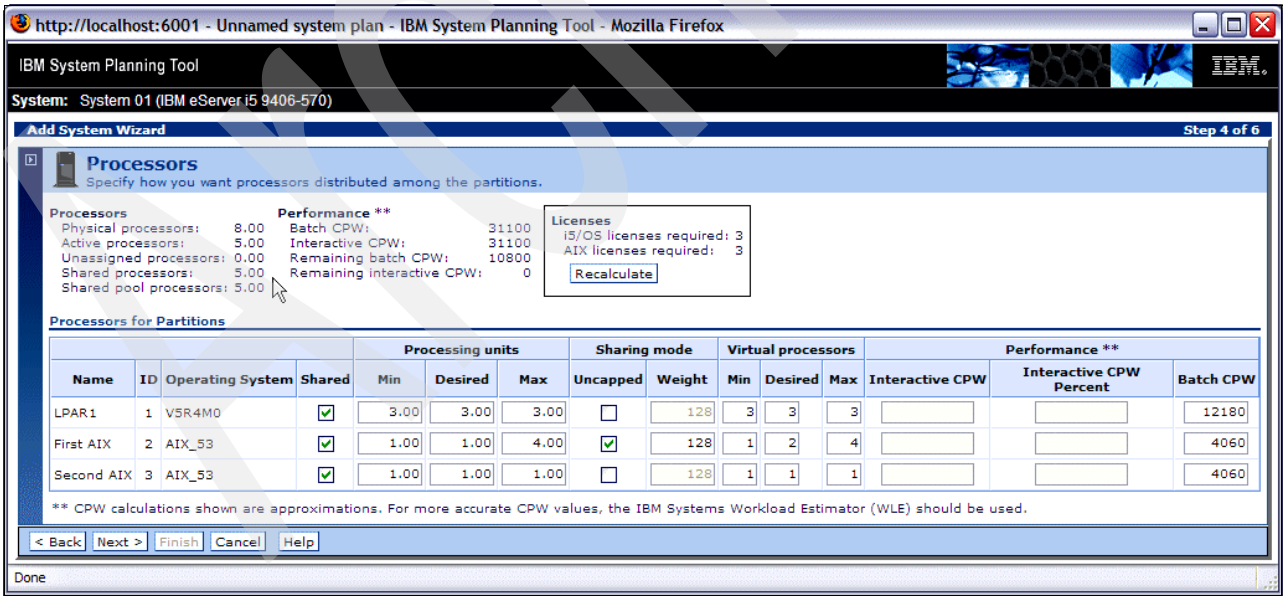

Figure 2-16 Resolved errors

2.5.4 Configuring memory

In the memory window the system, configured (assigned), hypervisor, and unassigned memory quantities are shown as well as the logical memory block size. The logical memory block size is defaulted to what the SPT perceived to be optimum based on the system model and amount of system memory. It can be changed by using the drop-down box. See *Logical Partitions on System i5*, SG24-8000, for an explanation of memory usage and block size.

The memory (MB) for the partitions is defaulted to the absolute minimum required by the selected operating system in each partition. You must change the amount of memory per your requirements. As you change the memory amounts the memory numbers above the logical memory block size will change accordingly, including the hypervisor memory. The amount of hypervisor memory needed is influenced by the system model, the amount of system memory installed, and the difference between the minimum and maximum memory specified. You should specify minimum and maximum as realistically as possible. The hypervisor memory number should only be used as an estimation. The hypervisor memory is used to manage the partitions.

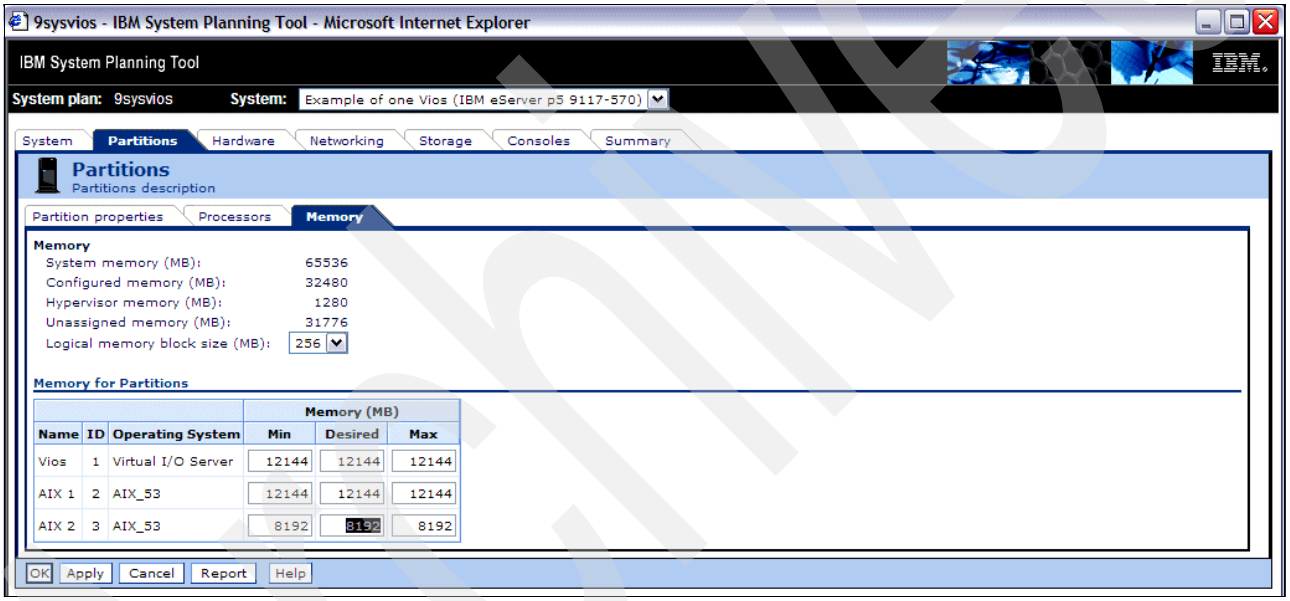

Figure 2-17 Assigning memory

Click **Next**. After all of the selections for the system and LPAR requirements are met, the next window (Figure 2-18) shows the System Summary window. Notice that so far there has been nothing about virtual slots, although the summary displays some that were automatically generated. This will be covered later.

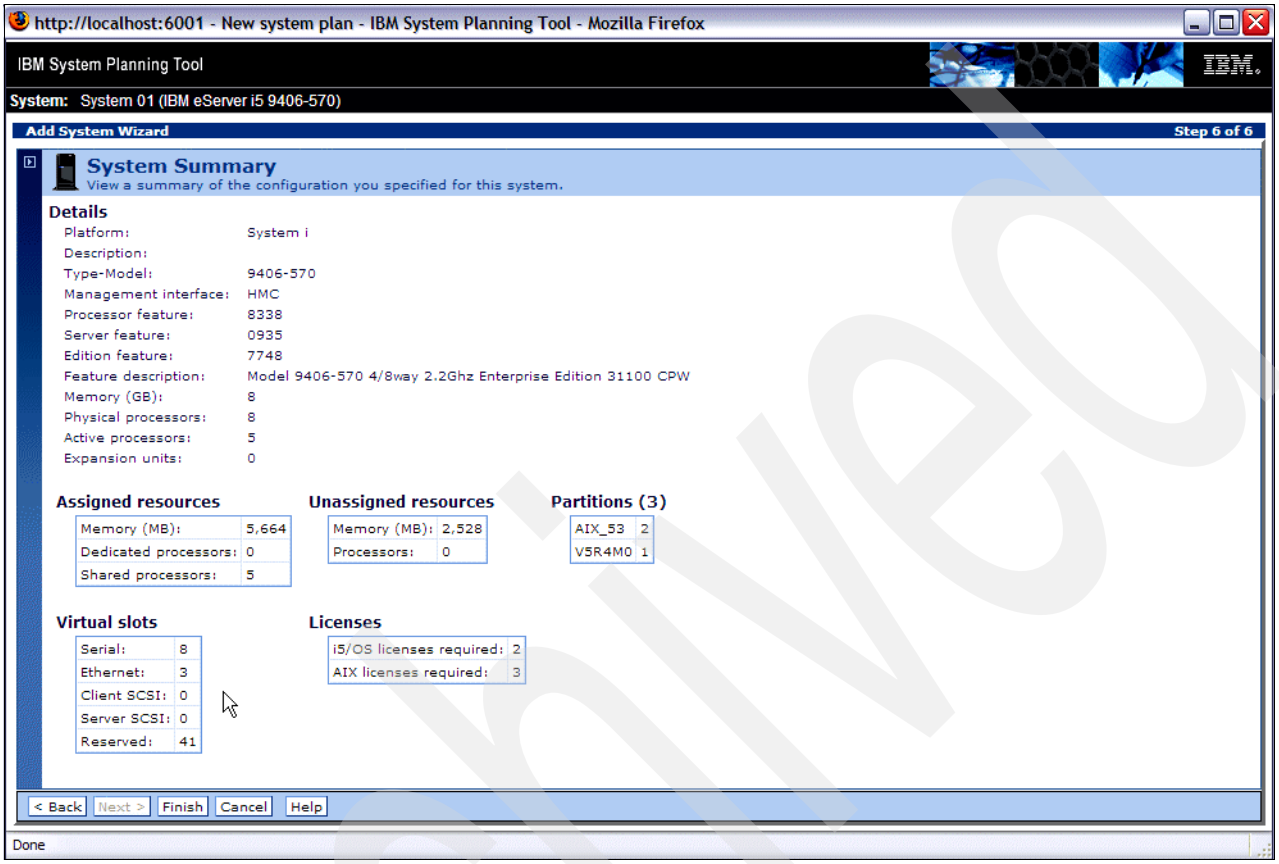

Figure 2-18 System Summary

Click **Finish** and you will proceed to the Work with Planned Systems window, which allows you to complete the remaining details of your configuration.

2.6 Work with planned systems

The window shown in Figure 2-19 is Work with Planned Systems. From this window you can select to edit the systems; add, copy, or remove systems; view properties (summary); view the report; and close, save, or export the file. The Work with Planned Systems window is where you begin filling out the details of a design and also where you end up when you are done so that you can save the .sysplan file, which can then be deployed from an HMC or an IVM partition, or export a .cfr file to be used for ordering purposes.

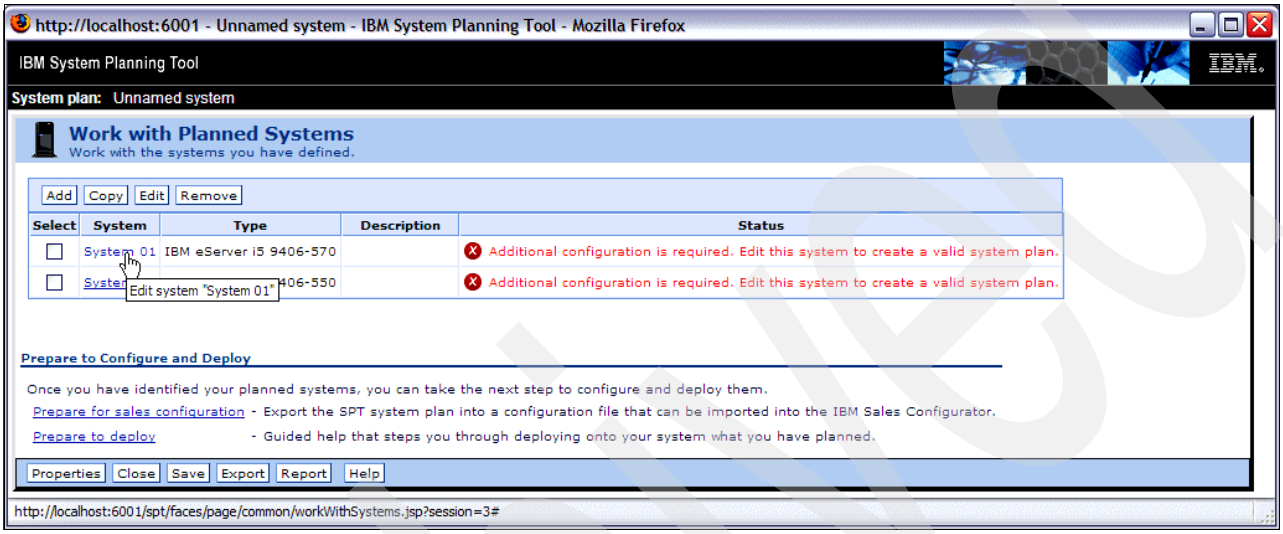

Figure 2-19 Work with Planned Systems

To edit a system you can either click the check box then click **Edit** or you can just click the system name as shown. Another system has been added to this plan for purposes of displaying some enhanced functions.

To add a system once you have created a system plan with one system, click **Add** and you will be taken through the normal *create a new system* process.

Copy allows you to create a duplicate of a system and puts you into edit mode of the duplicate system.

Remove allows you to remove one or multiple systems from the system plan.

The messages in the Status column are telling you there is more work to do. They will go away once you have created a valid design.

2.6.1 Placing hardware with the System Planning Tool

In our example we still need to add I/O hardware to complete the system configuration.

When you begin to edit a system you are taken to a new window, as shown in Figure 2-20 on page 31, and positioned on the Hardware folder tab. Notice the red x markers on some of the folder tabs. This indicates that the minimum requirements have not been met for these areas. The reason for being placed on the Hardware tab is that you have already created the information for the System and Partition tabs. However, you can select those tabs to make changes to those areas if needed. The validation messages at the bottom are tracking the minimum requirements for each partition and disappear as the requirements are met. At that point the red X is replaced by a green square.

The Hide validation message area toggles the lower messages panel off. This provides room on your display for all of the slots to initially be displayed for the larger I/O towers (for example, 5094).

The Additional Features tab displays features available at the system level rather than at the slot level. The parts shown in the Additional Features list are components or offerings that may be a necessary part of an order but that do not require placement by the SPT.

Placing hardware in SPT V2 is no different from how it was in SPT V1 except, for example, some of the navigation being slightly different. The process and the rules are the same. For your reference, placement rules are contained in the Redpaper *PCI and PCI-X Placement Rules for IBM System i models: i5/OS V5R3 and V5R4 (Fourth edition),* found at:

http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=redp4011

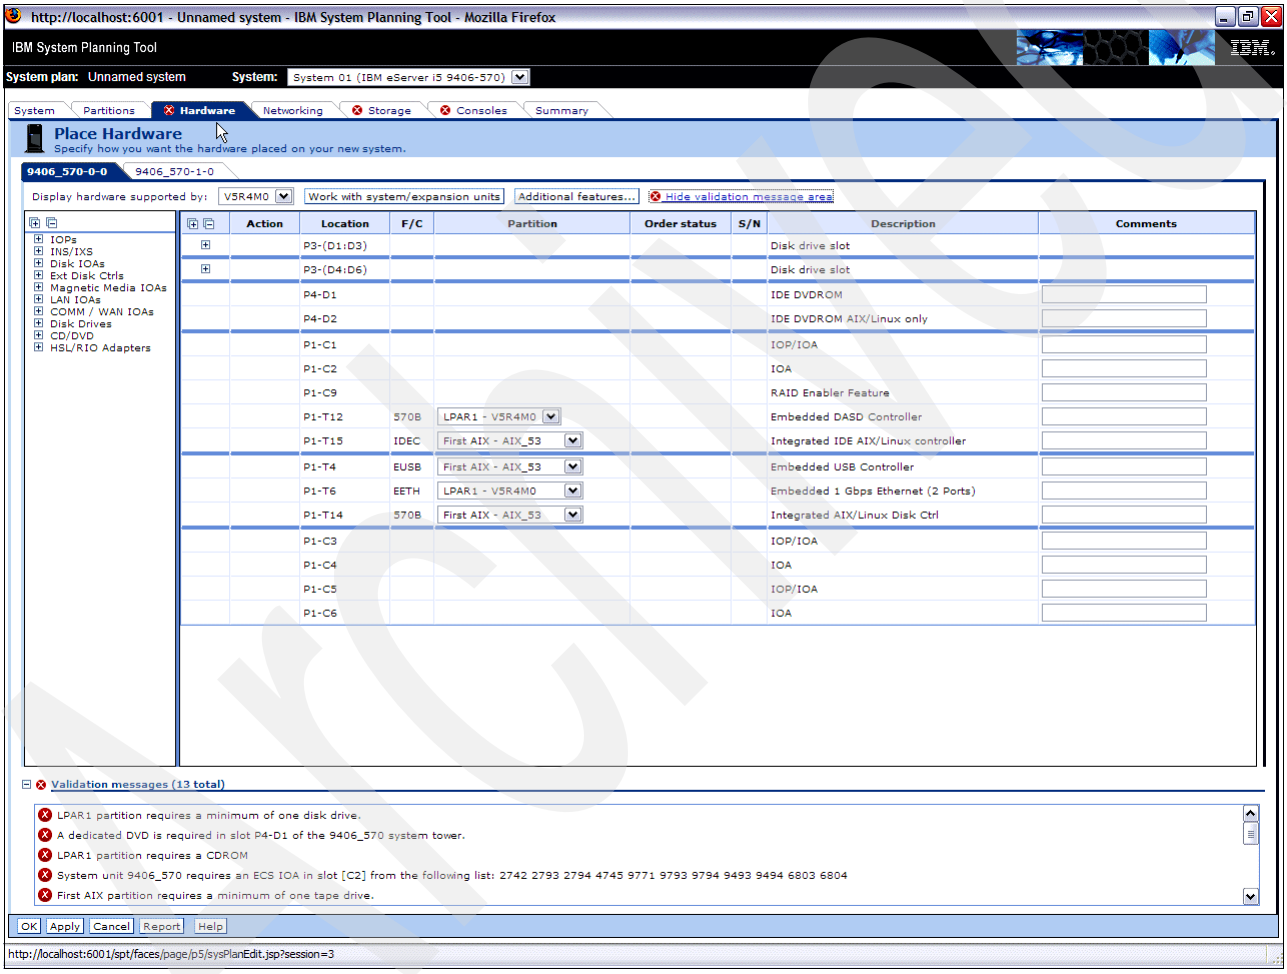

Figure 2-20 Place hardware

Note: In the Place Hardware panel there is a column for serial numbers (S/N). Once the system has been ordered, shipped, and installed, the serial numbers are available and can be input into the system plan.

Before we start to place hardware we would typically add any expansion units required for the configuration.

This book assumes that you have attended some logical partition education and are very familiar with hardware placement. You could select a sample system from Figure 2-3 on page 17. These sample systems have hardware already placed. Reviewing these sample configurations could give a novice user of SPT some insight into hardware placement.

2.6.2 Adding an expansion unit to the design

Select **Work with system/expansion units** in the header panel and you are taken to the Hardware System/Expansion Units window (see Figure 2-21). In this window you can add, copy, and remove expansion units.

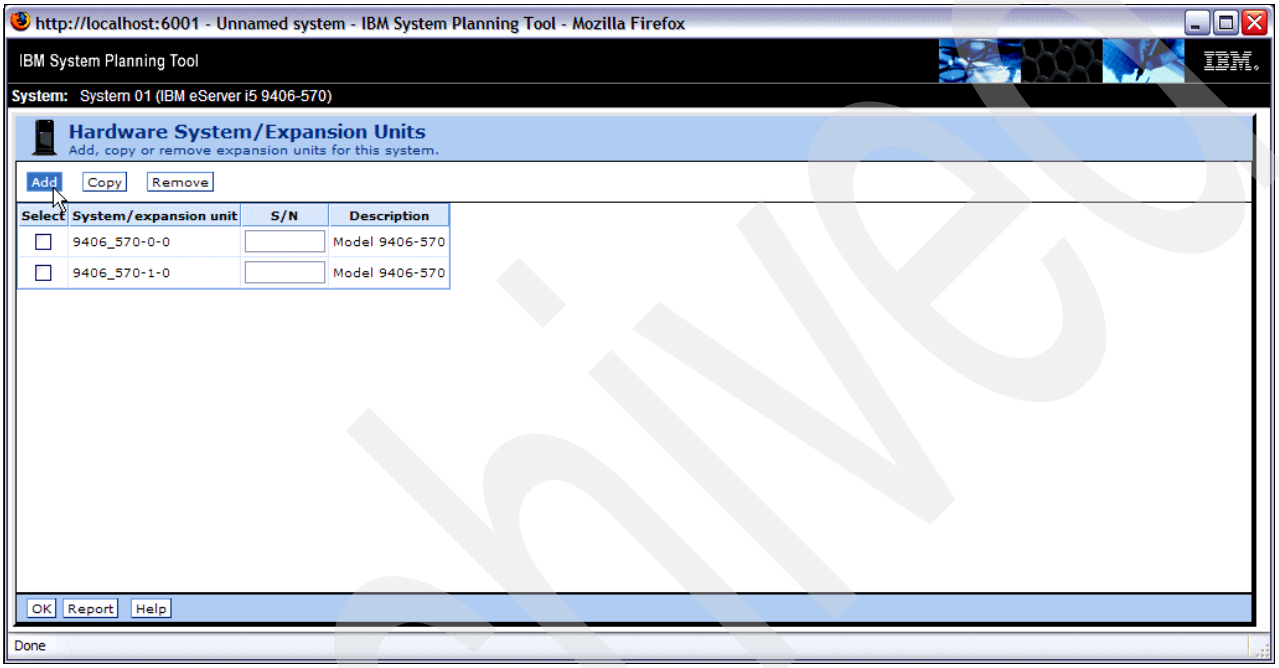

Figure 2-21 Add an expansion unit

In our example we click **Add** and you are presented with another window that provides a drop-down box from which to select the desired expansion unit.

Figure 2-22 shows the drop-down box selecting feature #0595. You can specify the quantity of the expansion unit, and if you want the expansion unit to be associated as an independent auxiliary storage pool you would click the IASP check box. When you have made your selection click **OK** and you will be returned to the previous window and the new expansion unit will then be added to the list. The Package feature drop-down box is hidden in this figure by the drop-down list. Certain expansion units are offered with some things automatically added to them, typically disk drives with a disk controller. The package feature box does not appear if the expansion unit you select does not have any package features available.

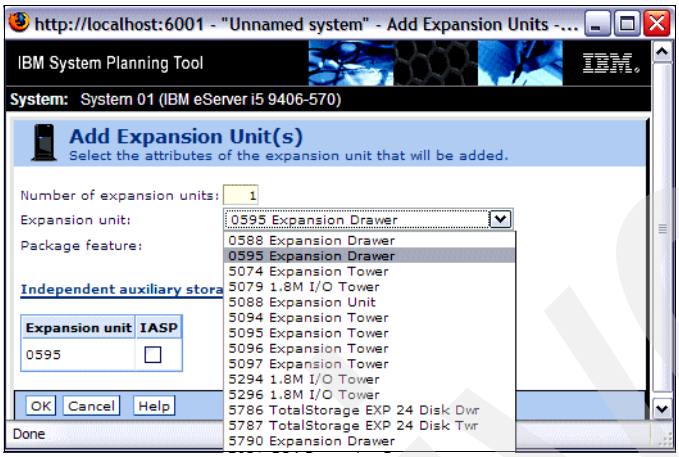

Figure 2-22 Selecting #0595 expansion

After selecting the desired expansion unit, click **OK** and your are taken back to the Hardware System/Expansion Units window.

Figure 2-23 shows the result of adding an expansion unit. Since the unit added has package features available, a new column has been added to the table and you still have the option of selecting a package feature.

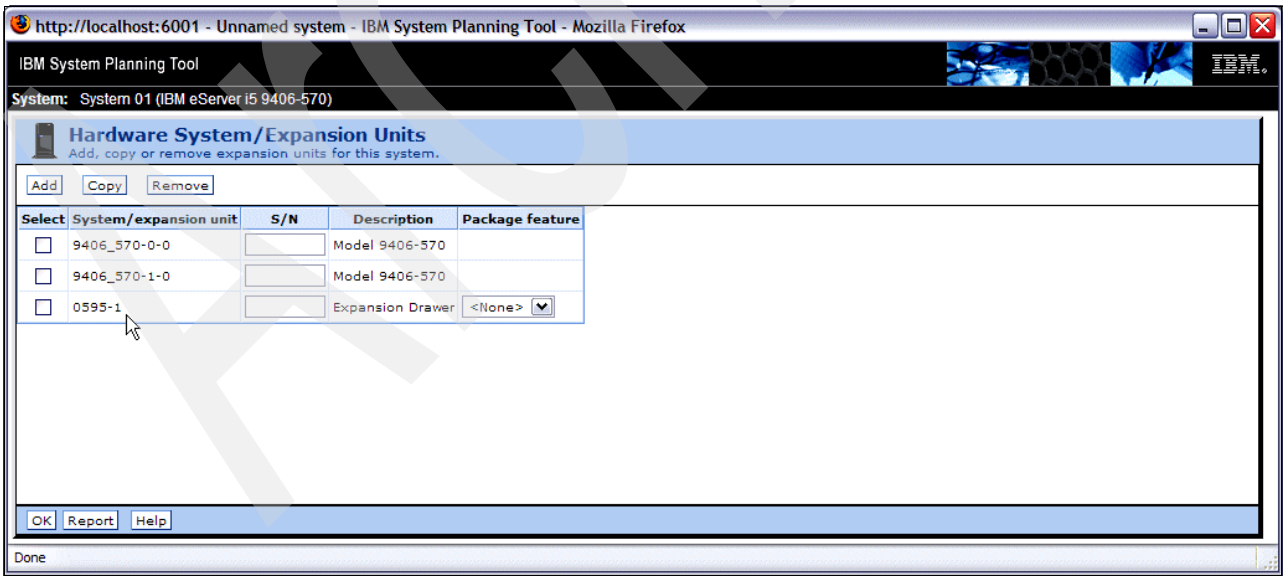

Figure 2-23 Results of adding an expansion unit

Now that we have added an expansion unit we can also copy it.

2.6.3 Copying expansion units

In our example shown in Figure 2-24 we have clicked the check box for the 0595 to copy it. When you copy an expansion unit the copy includes everything that has been placed in the unit.

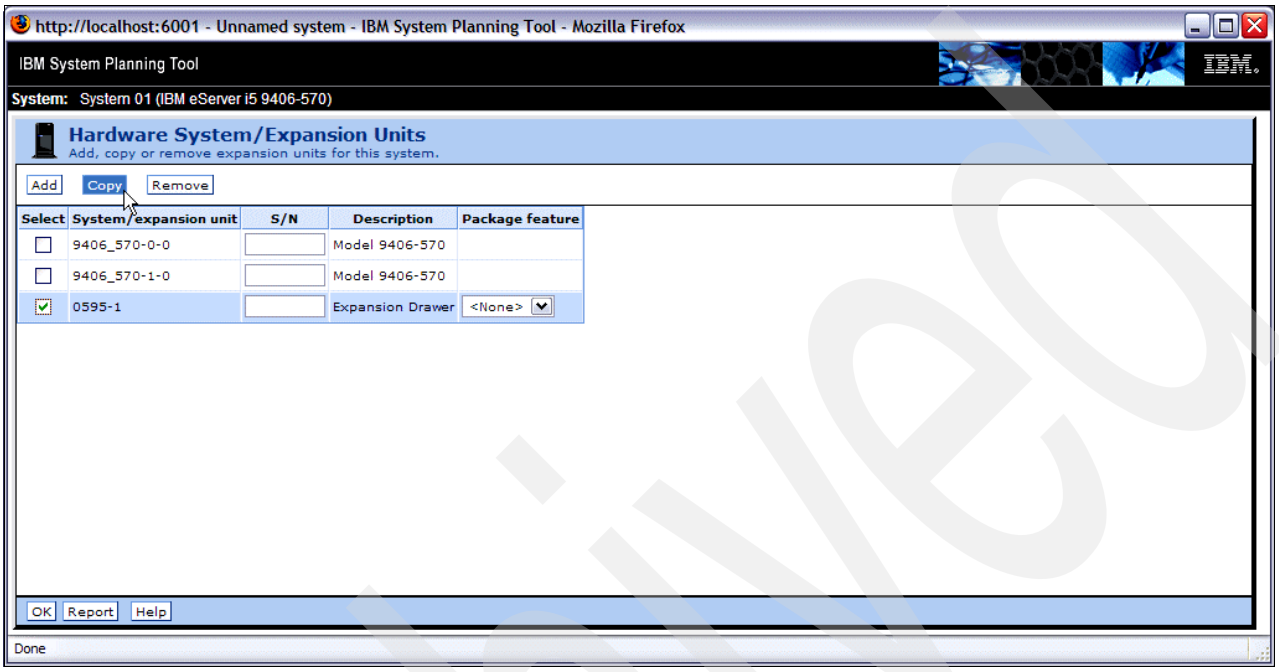

Figure 2-24 Copying an expansion unit

Note: In the Add Expansion Units panel there is a column for serial numbers. Once the system has been ordered and shipped the serial numbers are available and can be input into the system plan. This will help to prevent deployment failure through ambiguous system units and expansion towers. See 3.3.1, “Deployment validation process” on page 81, for further information about tower ambiguity.

Click **Copy** and the Copy Expansion Unit window appears. This window contains a list of all systems that exist in the system plan. This is because you can copy an expansion unit within a system or you can copy it to another system in the same system plan, as shown in Figure 2-25.

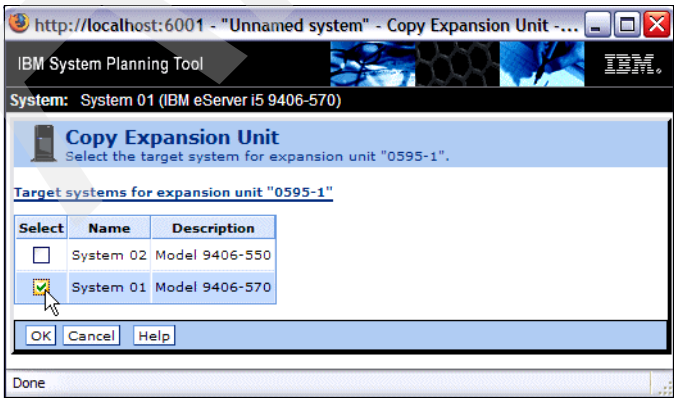

Figure 2-25 Copying an expansion unit

Click the check box of the system to receive the expansion unit copy and click **OK**. You will be returned to the previous window, which now includes the additional expansion unit, as shown in Figure 2-26.

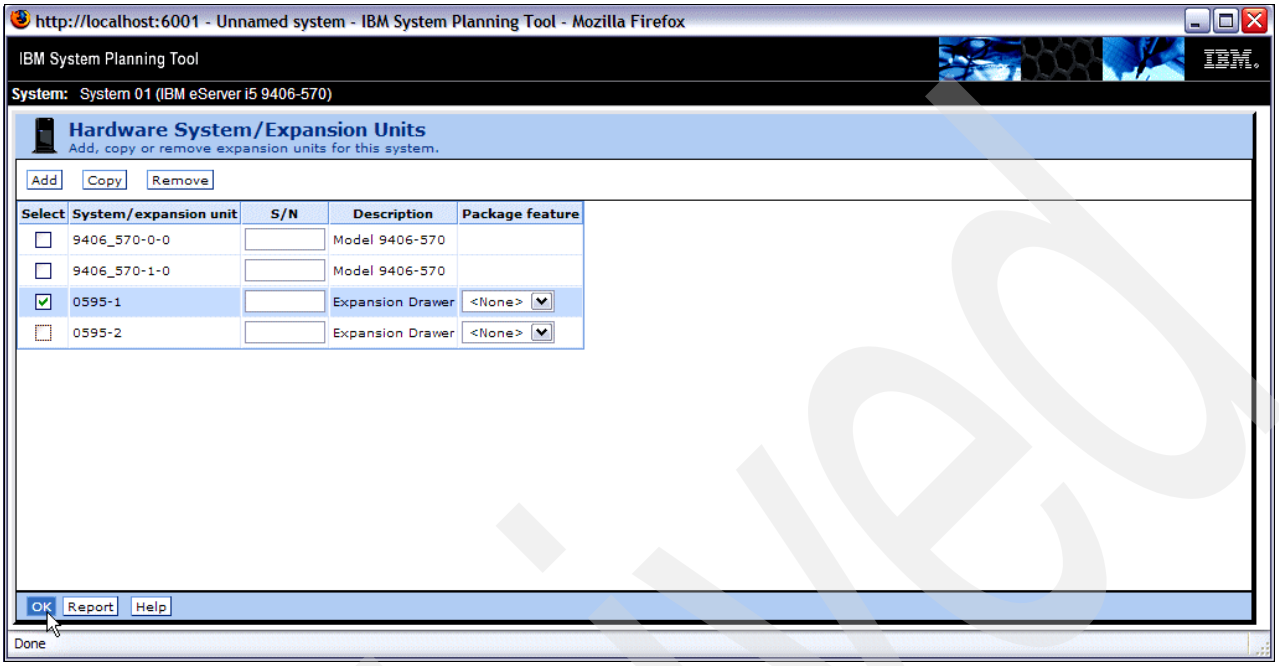

Figure 2-26 Results of copying an expansion unit

Click **OK** to go back to the Place Hardware window.

At this point our example assumes that you have added all the expansion units needed for your design. Now you need to populate them with I/O elements to meet your design requirements. If you are going to have several expansion units with contents that are identical or nearly so, you may want to populate one first and then use the copy function to create the others.

Note: It is not possible to copy 7031-D24, 7031-T24, 5786, or 5787 EXP 24 expansion units if you have already populated them because of the association that is built between a repeater and a disk IOA, which always resides in another expansion unit.

2.6.4 Hardware tab

Now we look at the process and functions used to place I/O in the system. In this example we add cards and disk drives to one of the #0595s.

After selecting a unit from the folder tabs, the physical layout of that unit appears and now you can place the I/O you need. The SPT follows all of the placement rules so that you cannot create a situation that will not work or, more importantly, will not be supported. It is not our intention to teach LPAR and hardware design in this publication, but you should remember that an expansion unit does not have to be dedicated to one partition. It can be shared between partitions at a bus level.

The left-hand navigation panel displays the hardware resource categories. Click the category name or the plus sign (+) to expand the selection. The list of features shown is controlled by the OS selected in the drop-down box labeled `Display hardware supported by` (to the left of

the “Work with system/expansion units” field) and by what system unit or expansion tower is currently displayed on the window. If a feature is not supported by the selected OS or is not allowed to be placed in the unit displayed, it will not be shown in the list.

From the sidebar select the category of hardware that you want to add. Select the particular item from the list.

In our example (Figure 2-27) we have selected our first 0595 (0595-1), and all the slots and disk positions in that unit are displayed. The IOP category has been clicked and a 2844 IOP has been selected. When you click the feature that you want to add, a large green plus symbol (+) appear in the card slots where it can be placed. You then click the plus symbol in the desired slot and that card will be placed there.

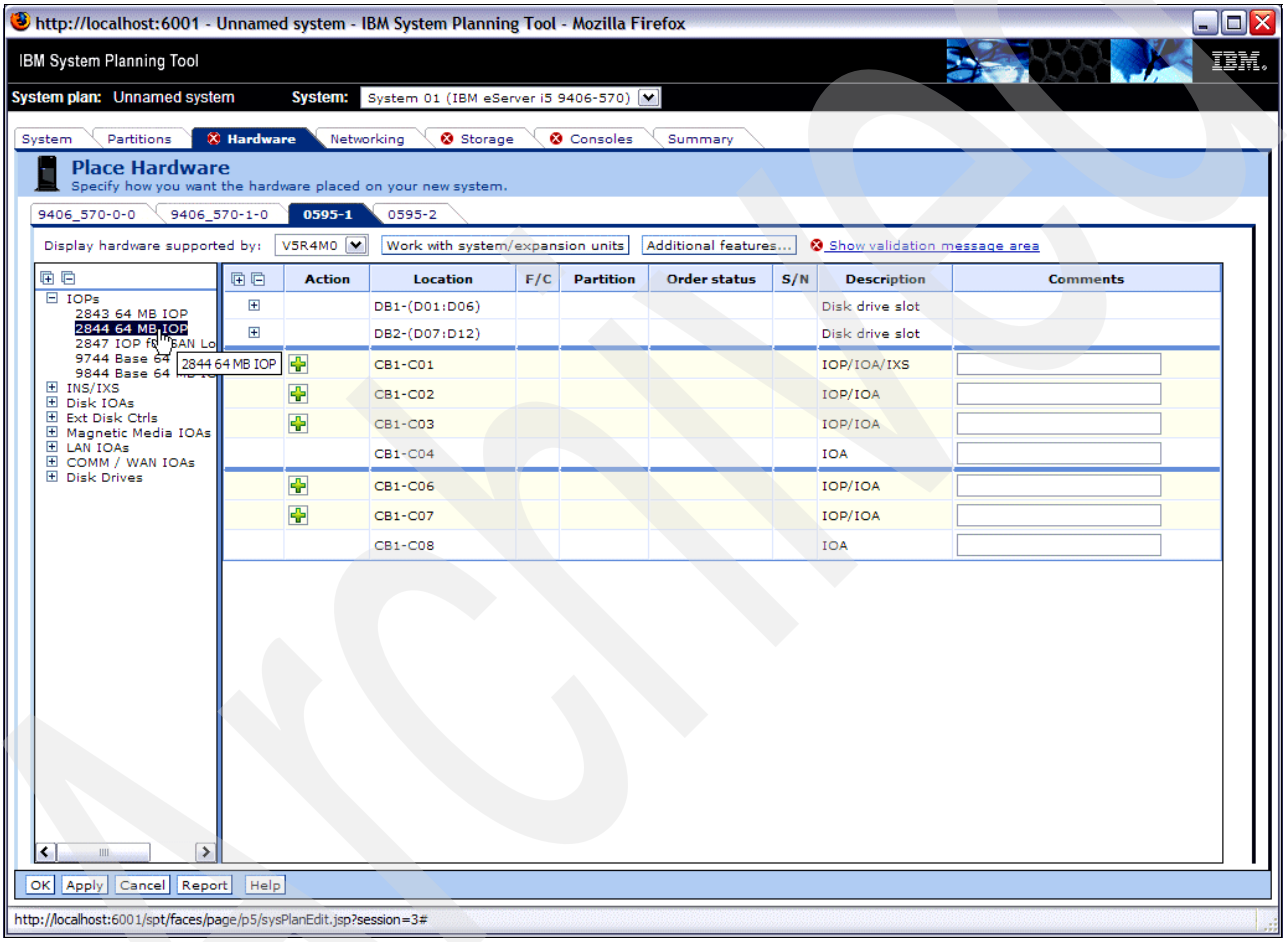

Figure 2-27 Selecting an IOP

When you place an IOP or an IOA that does not require an IOP (for example, a 5706 PCI-X 1Gbps Ethernet-TX IOA), SPT assigns it to the last partition used to add something previously. You can of course change the assignment by clicking the drop-down box and

selecting the appropriate partition, as shown in Figure 2-28. If you change the partition assignment of something that already controls other items, the partition assignment of those items changes as well.

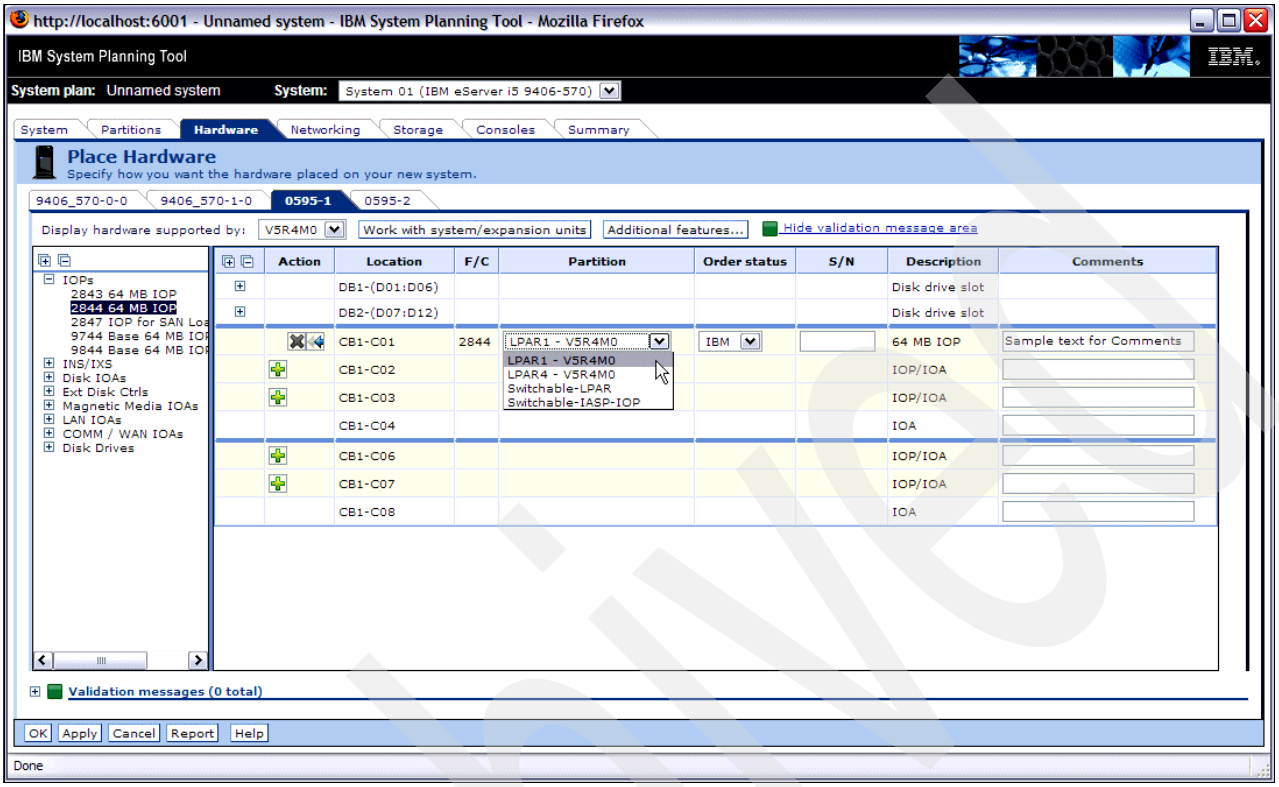

Figure 2-28 IOP/IOA partition assignment

The Order status column allows you to define one of three categories of any given feature, as shown in Figure 2-29. The categories are:

- IBM - This part is an IBM-manufactured part that must be ordered from IBM.
- OEM - This part is not an IBM-manufactured part.
- Own - This part is an IBM-manufactured part that is already owned by the customer.

The features designated as OEM or Own are included in all aspects of the SPT, but they are not included in any subsequent order based on any file created by and exported from the SPT.

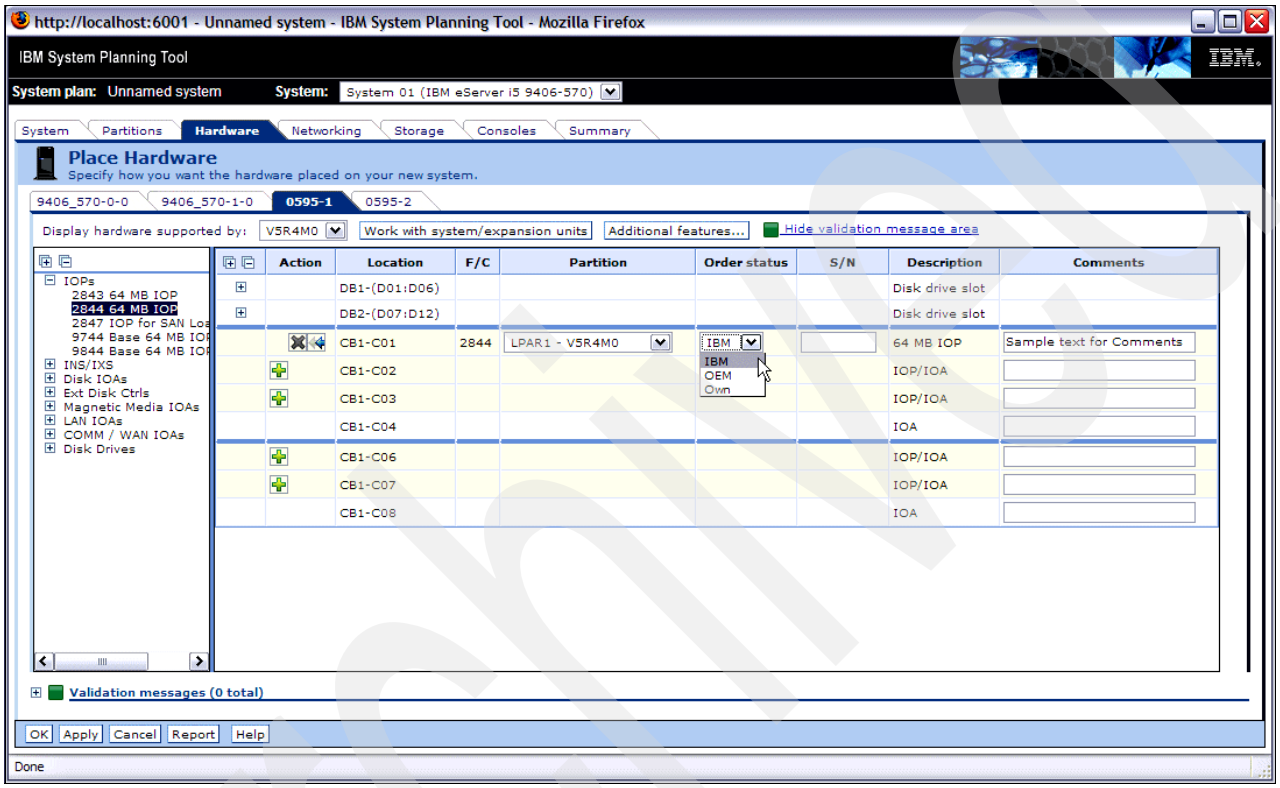

Figure 2-29 Order status

When you place a feature its description is added to the description column. You may also use the Comments fields to key in anything specific that you want to note, for example, “This controller is Load Source for LPAR1”. For our example we inserted sample text, as shown in Figure 2-30. Your comments can be longer than the field shown as your text will move over as you type. The comments can be shown or hidden in the System Plan Viewer. We cover the System Plan Viewer later in this publication.

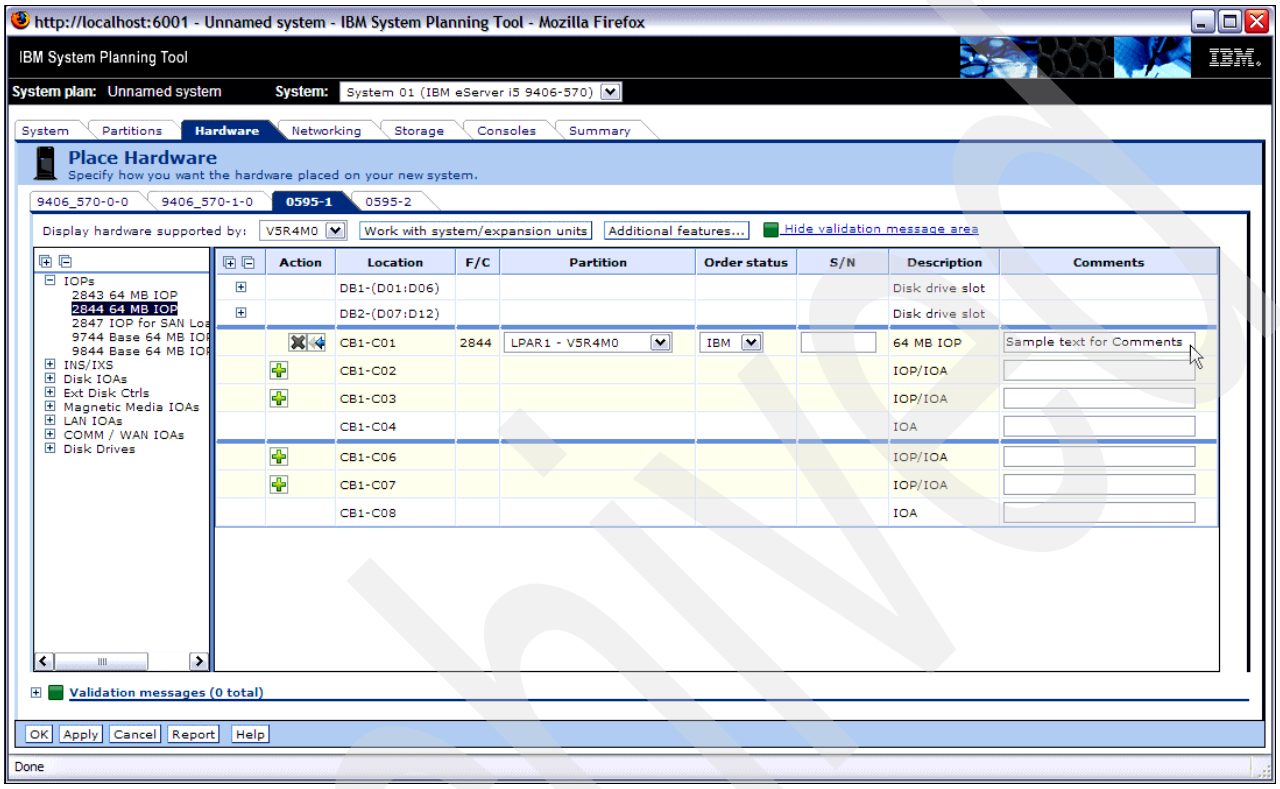

Figure 2-30 Descriptions and comments

After placing an item, a remove symbol (X) appears. If you select any item that is allowed to use that slot, a replace symbol (<-<-) appears next to the item, as in Figure 2-31. The replace function actually performs a remove/add process under the covers. As we create a working design you should remember that replacement of a controlling part will only be allowed if the new controller can control the parts that the current controller is controlling, and the controlled parts will not be affected. They do not need to be removed and re-added when the controller needs to be replaced. This works very well for changing a disk IOA when it is already controlling disk drives.

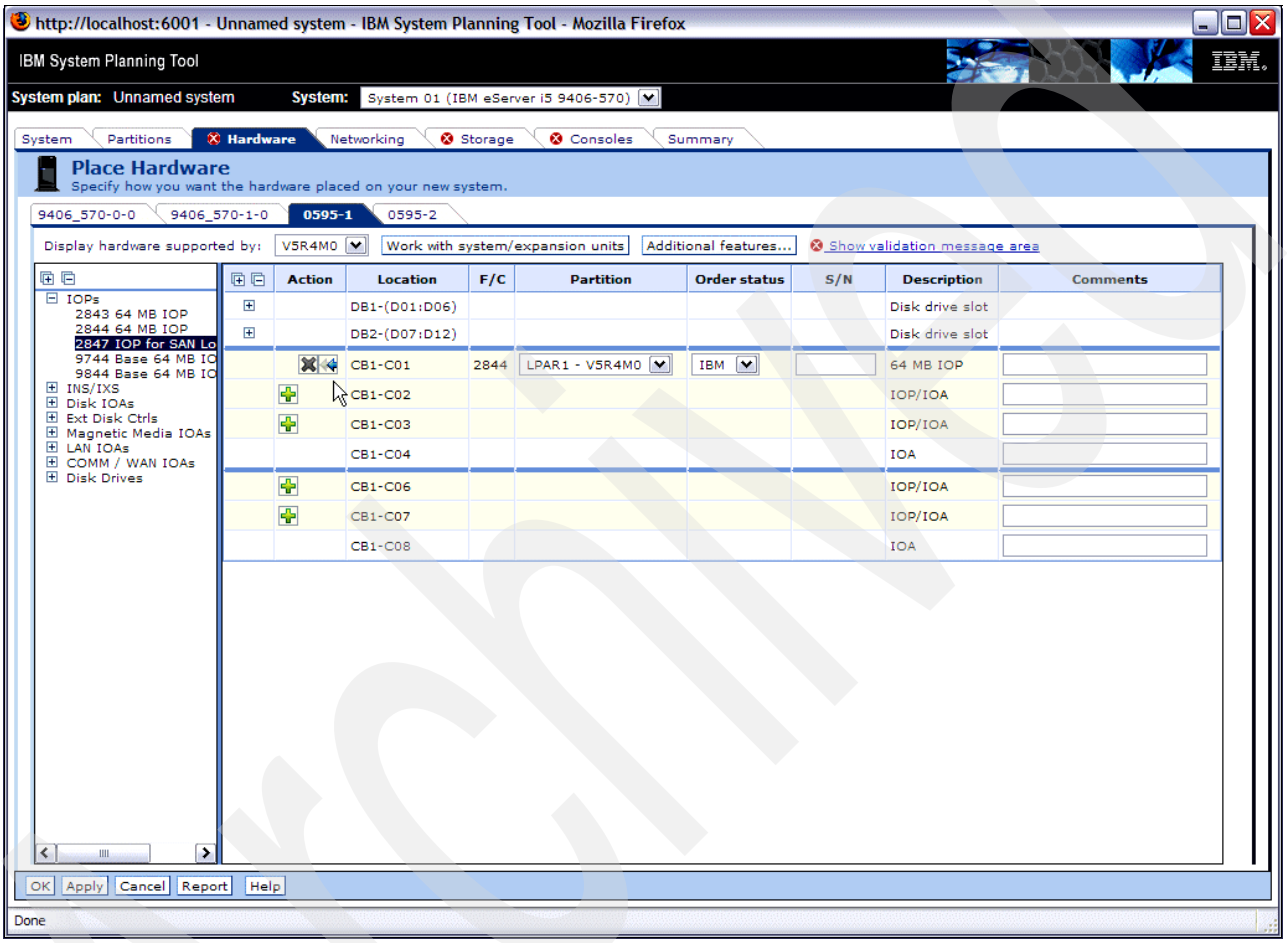

Figure 2-31 Remove and replace function

In the next frame (Figure 2-32) a disk IOA has been added to C02 and six disk drives have been added to disk bay 1 (DB1).

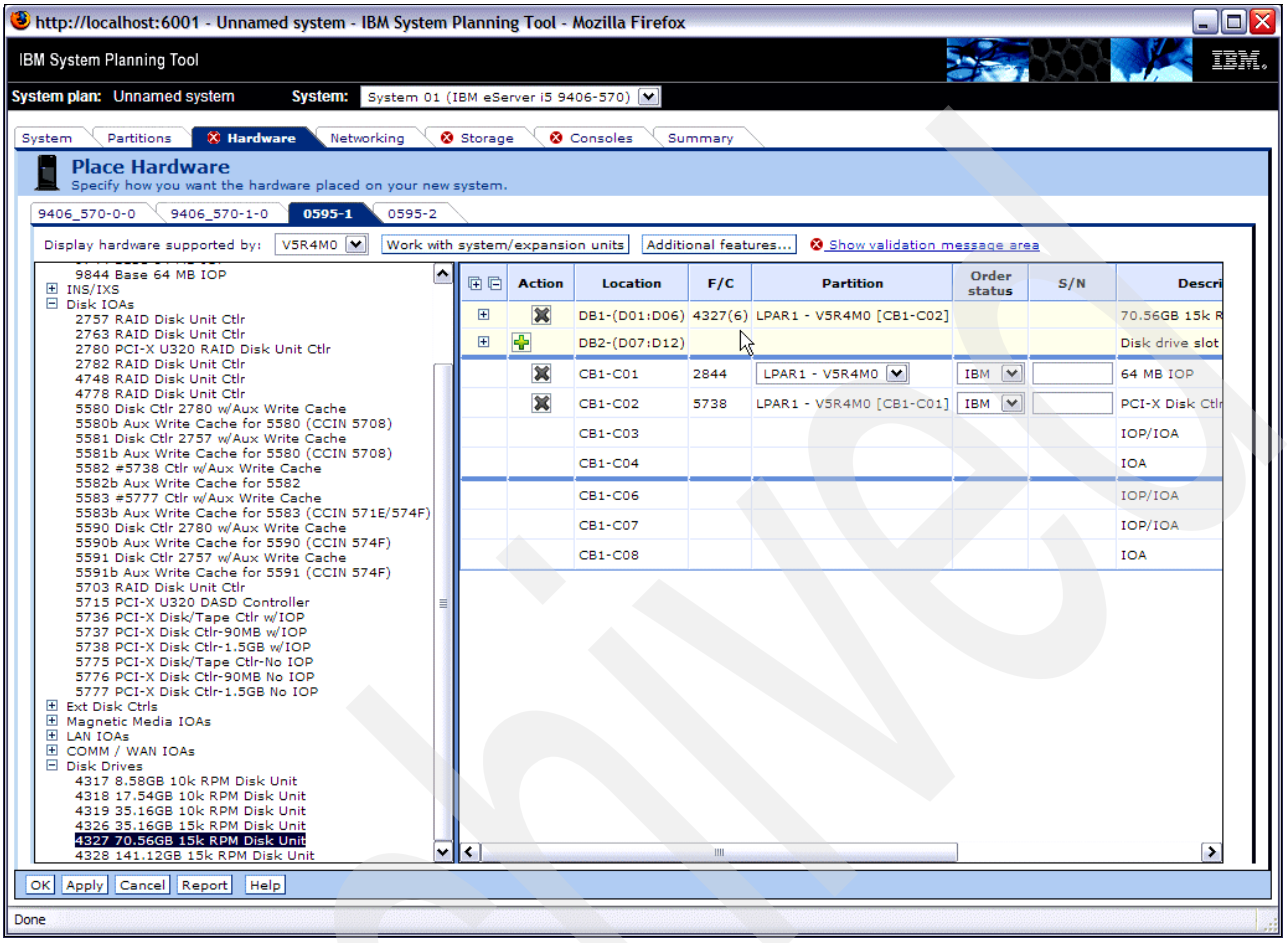

Figure 2-32 Added disk IOA and disk drives

For the disk positions in any given system or expansion unit you can expand the view to see the individual drives by clicking the small plus symbol (+) to the left of the disk drive bay locations, as shown in Figure 2-33. Disk drives can be added or removed individually in the expanded view. The expanded view also provides a comments field for each disk drive.

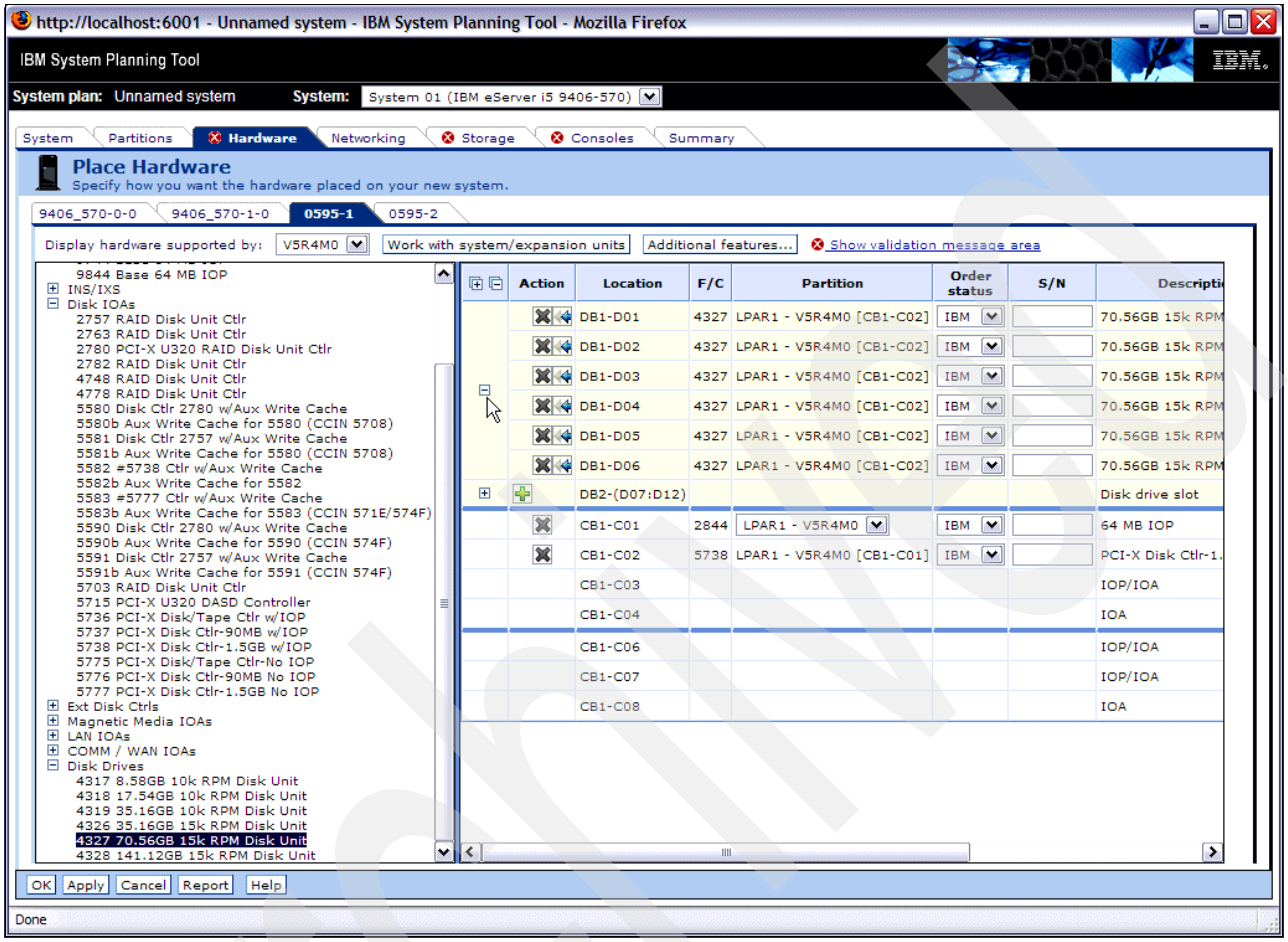

Figure 2-33 Expanded view of disk bay 1 (DB1)

You will need to continue this process of adding IOAs and disk to meet your requirements and then go to the next folder tab (Networking).

Notice a System drop-down box in the header line, shown in Figure 2-34. The drop-down box shows all the systems in the system plan. While in edit mode, in a multi-system system plan, you can switch over to a different system to work on. You can switch back the same way. This can be done in any window where this drop-down box appears.

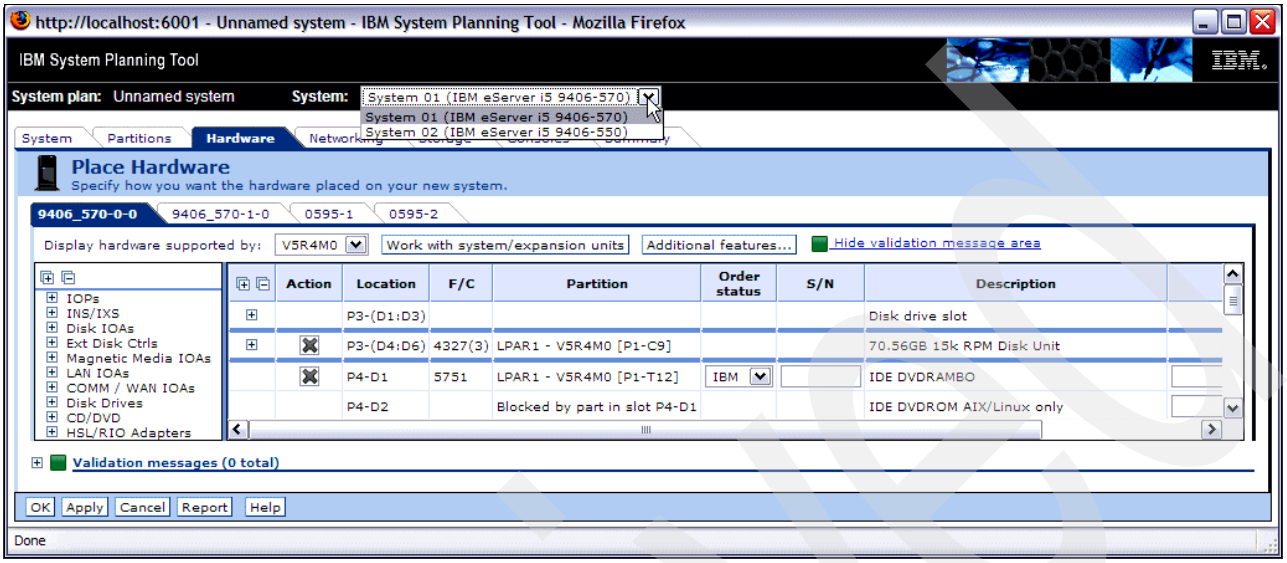

Figure 2-34 Selecting a different system for edit

2.6.5 Networking tab

The SPT automatically creates one virtual LAN (VLAN) connecting all of your partitions. See Figure 2-35. You can change this by unchecking partitions. You can also add VLANs of your own design. To do this click **Add** and put a check in the boxes of the partitions that you want to use that VLAN. If you click the partition name you will see the details of the VLANs with which the partition is associated and the physical Ethernet cards owned by that partition. If you have many partitions and many VLANS you will see a much larger grid to work with.

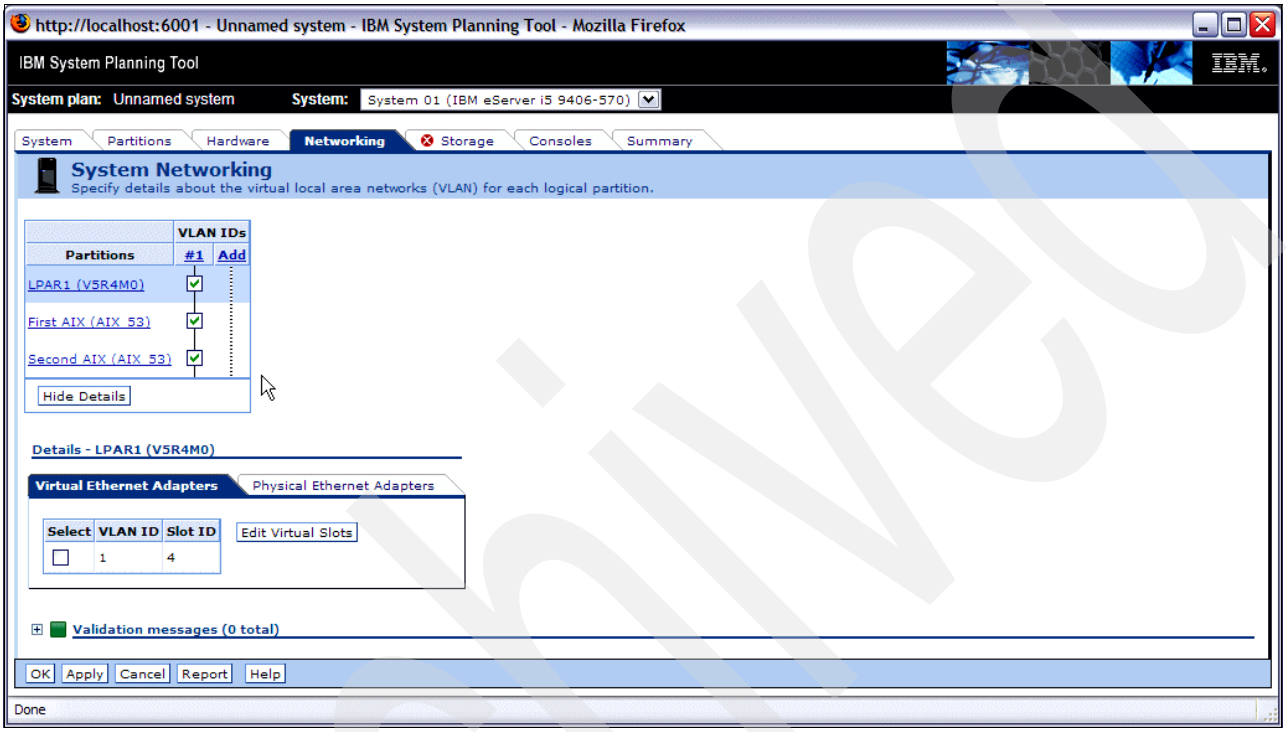

Figure 2-35 Networking

2.6.6 Storage tab

The Storage tab has some very important functions, especially if you have a Server Client relationship between some partitions. This is where you define the virtual SCSI connections between those partitions so that the client partition can use the server partition's resources. Also, this is where you will define the load source for whatever i5/OS partitions you have. In Figure 2-36 we begin the process by clicking **Add** and adding a virtual SCSI connection.

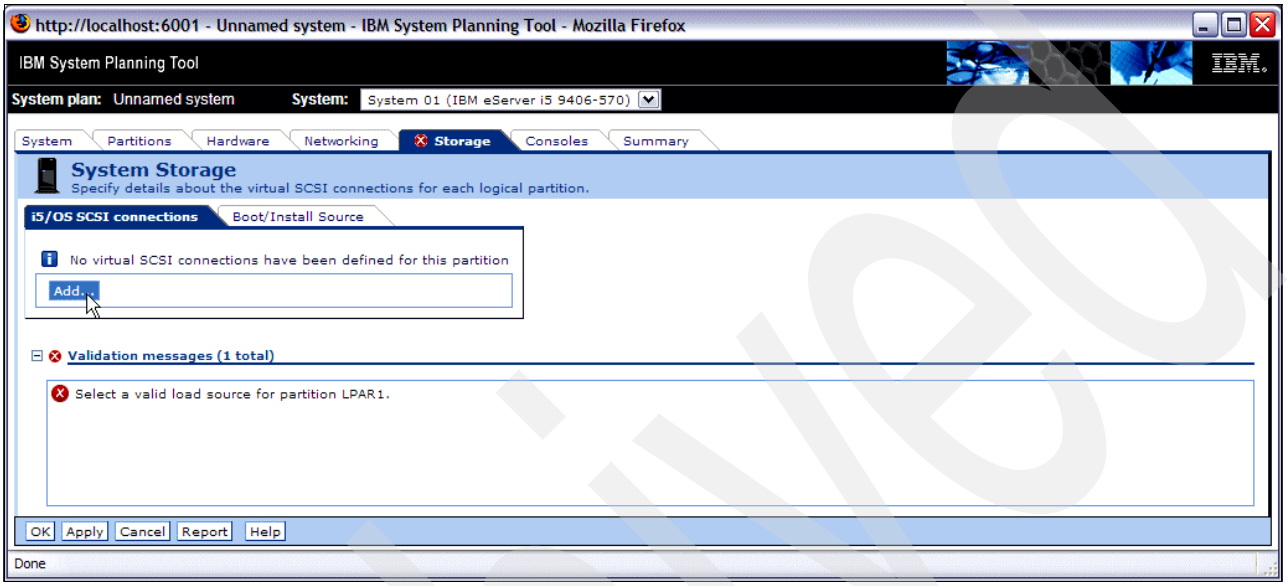

Figure 2-36 First page under Storage tab

In this case we only have one i5/OS partition, but if there were multiples of them or if you had any Virtual I/O Server partitions on a System p you would use the drop-down box to select the Server partition. Then select the number of connections, and then select the Client partition. In Figure 2-37 we add a connection between the LPAR1 and second AIX partitions.

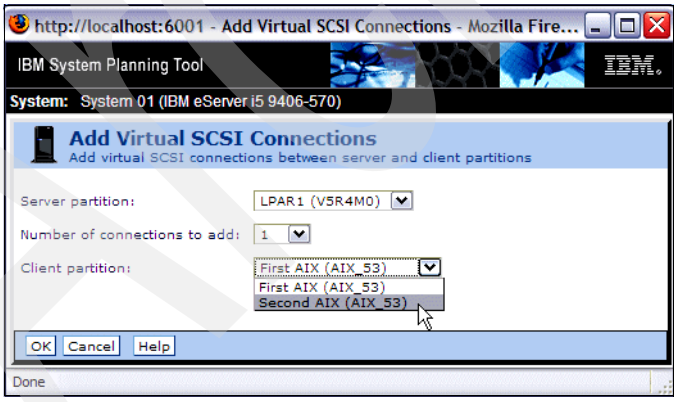

Figure 2-37 Adding virtual SCSI connections

When you click **OK** you will be taken back to the System Storage window, and the details of each connection can be viewed there. From that window you can also select **Edit Virtual Slots**, as shown in Figure 2-38.

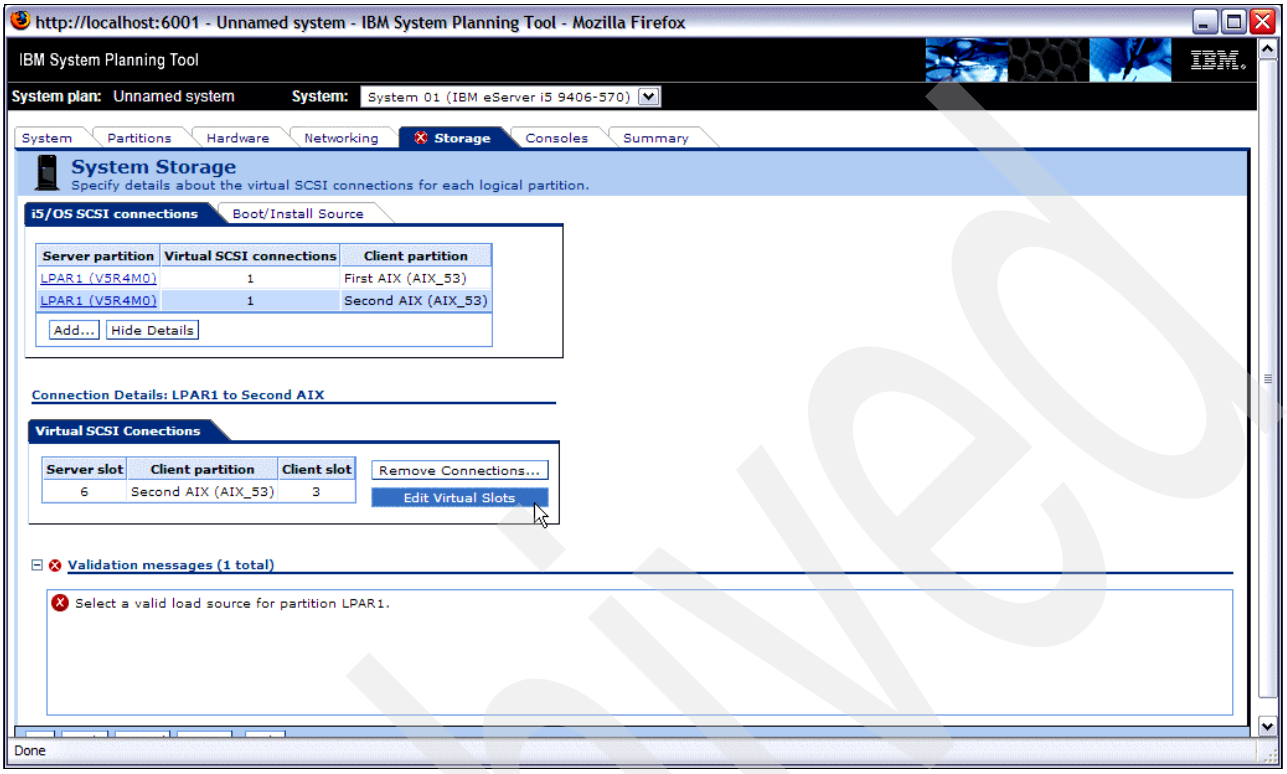

Figure 2-38 Details of virtual SCSI connections

Figure 2-39 is a sample of what you see next. All of the partition information tables can be collapsed or expanded and you can also move them from one side of the window to the other. This is very helpful in being able to compare and verify or edit the virtual connections between partitions.

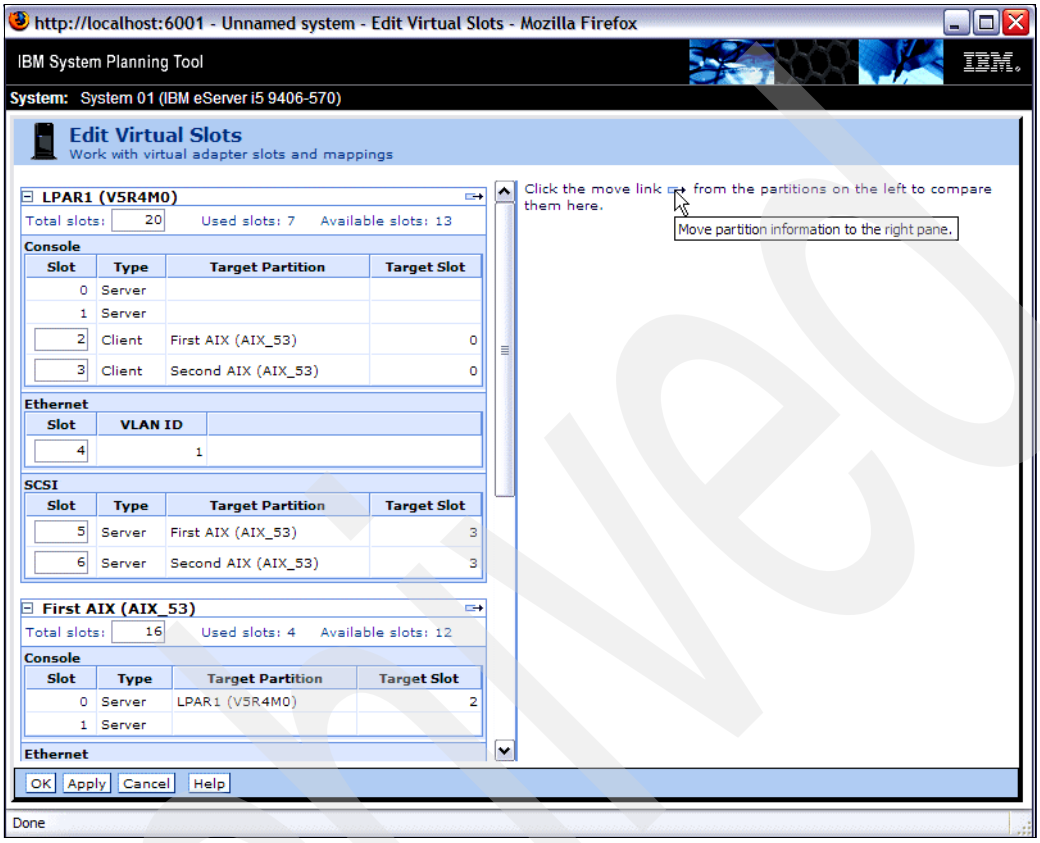

Figure 2-39 Edit virtual SCSI connections

In Figure 2-40 the table for the second AIX partition has been moved to the right-hand side. Note how the virtual slots of one partition match the target slots of its counterpart.

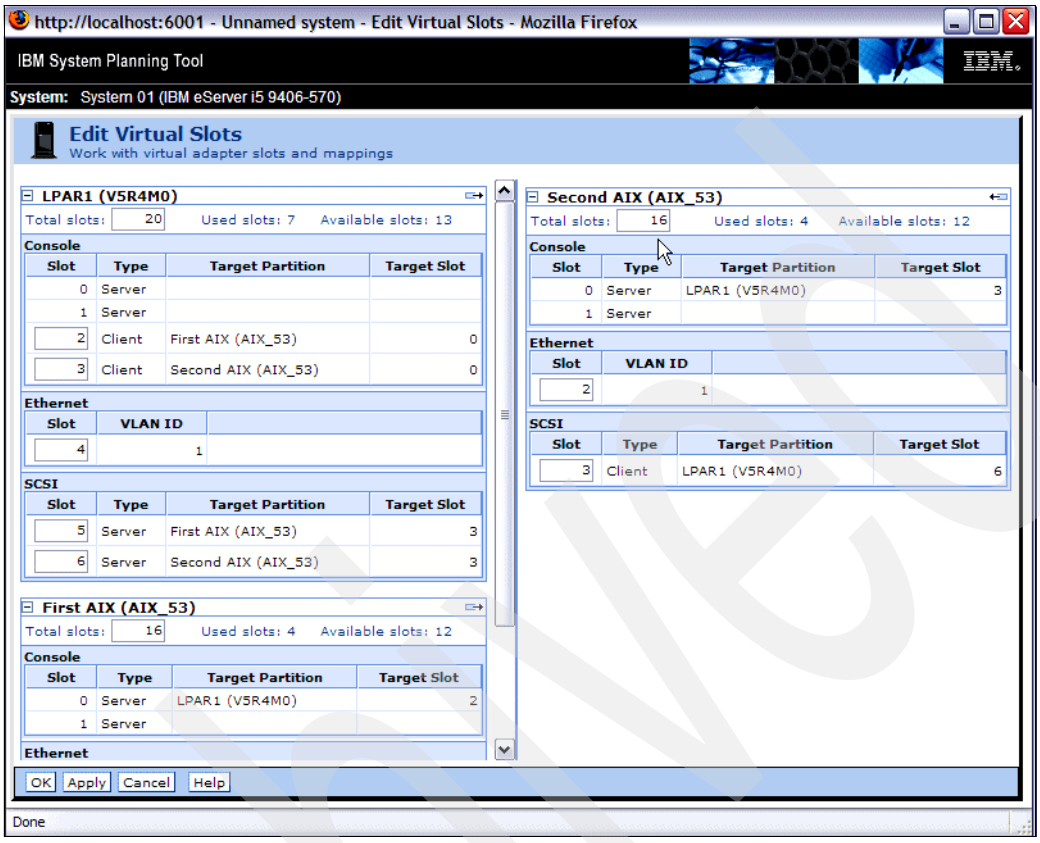

Figure 2-40 Editing and comparing virtual connections

Prior to SPT V2 you would have defined all of the virtual connections far earlier in the design process and you would have optionally specified how many virtual slots you wanted to reserve for future use. SPT V2 does not give you the opportunity to do any of this until you get to this window. From this window you can change the slot IDs of the virtual slots for the VLAN (Ethernet), virtual SCSI, and virtual console. As soon as you add a virtual SCSI to a client partition, SPT automatically creates a virtual console connection. The Total slots field for each partition can be changed as well, and this effectively perform the reserve function providing that you are specifying a quantity greater than the used slots number. The available slots number is the *reserved slots* number. If you change an ID or a quantity and click **Apply**, the corresponding information is updated, which also updates the target slot information.

2.6.7 Defining load source

You may have noticed a validation message on previous windows concerning a load source for LPAR1. In Figure 2-41 we have selected the Boot/Install Source tab and clicked the drop-down box. Remember that this only has to be done for i5/OS partitions in this particular example. However, if you are defining a System p with Virtual I/O Server partitions, that also needs to have a load source identified. See Chapter 4, “Virtual I/O Server” on page 103, for further details.

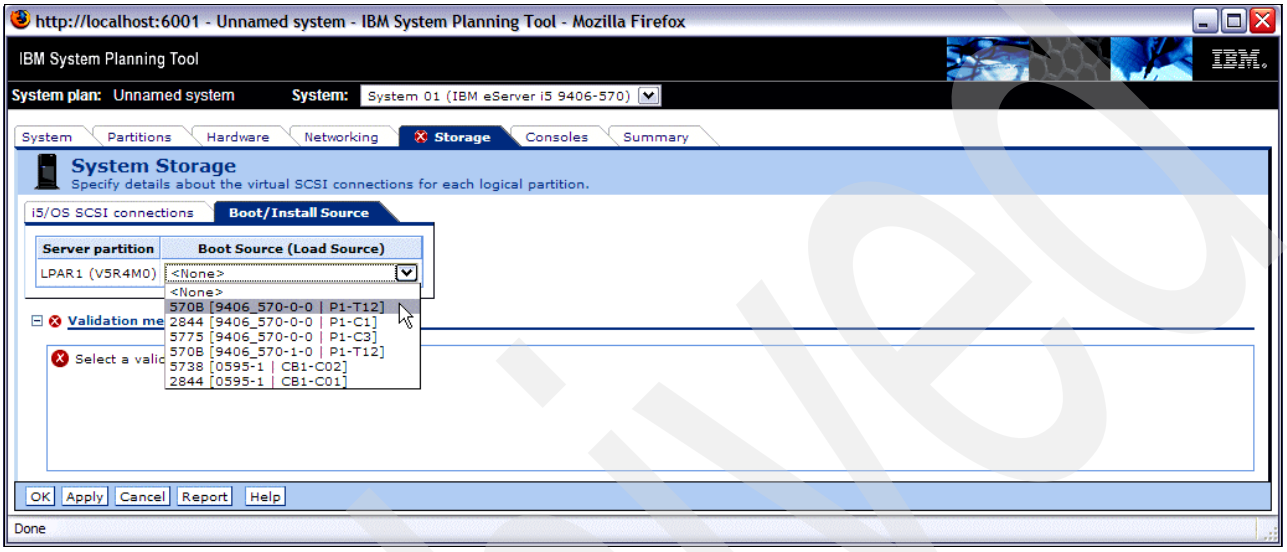

Figure 2-41 Selecting load source

While the i5/OS load source is actually one or more disk drives, what we select here is either the disk IOA that controls the desired disk or the IOP that controls the disk IOA. It is possible to have more than one disk IOA under an IOP, so to avoid confusion we recommend selecting the disk IOA rather than the IOP.

Note: On POWER6 MMA models a system unit can be ordered with no disk drives. On a System i order you can specify #0719 to indicate no disks within the system unit. If you want to do this, you must then specify a different specify number that indicates the enclosure that will contain a load source device:

- #0720 Load Source in #0595/5095
- #0721 Load Source in #5094/5294
- #0725 Load Source in #5786/5787

Keep this in mind when using SPT V2 to define your load source IOA.

In Figure 2-42 we see that all validation messages have been satisfied and there are zero remaining messages.

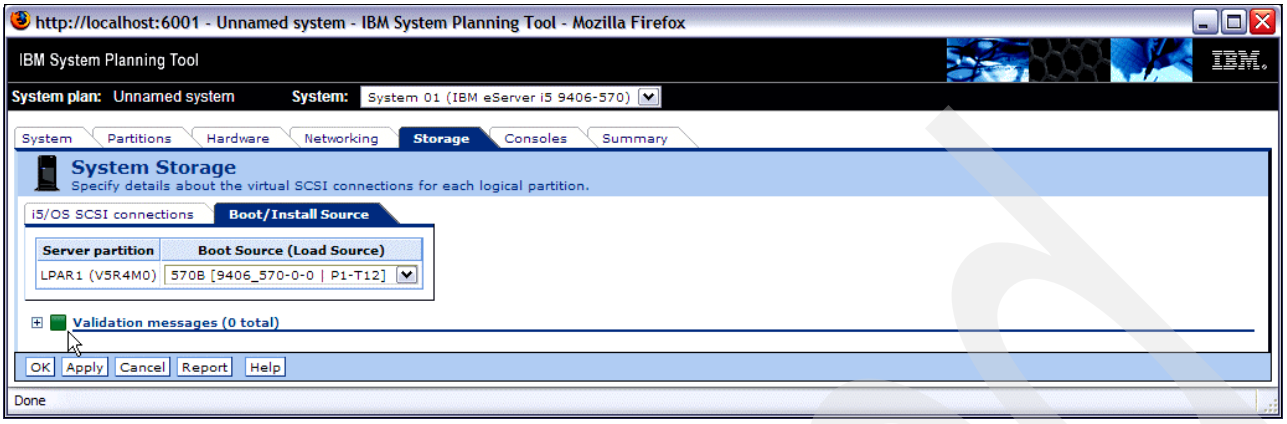

Figure 2-42 Load source selection completed

2.6.8 Consoles tab

On the following window there are two main functions to be completed.

Physical consoles

First is to assign Physical Consoles to partitions. In Figure 2-43 we have selected LPAR1 and clicked the Physical Console drop-down box and selected HMC. The drop-down box only contains whatever items in the selected partition are able to serve as the console.

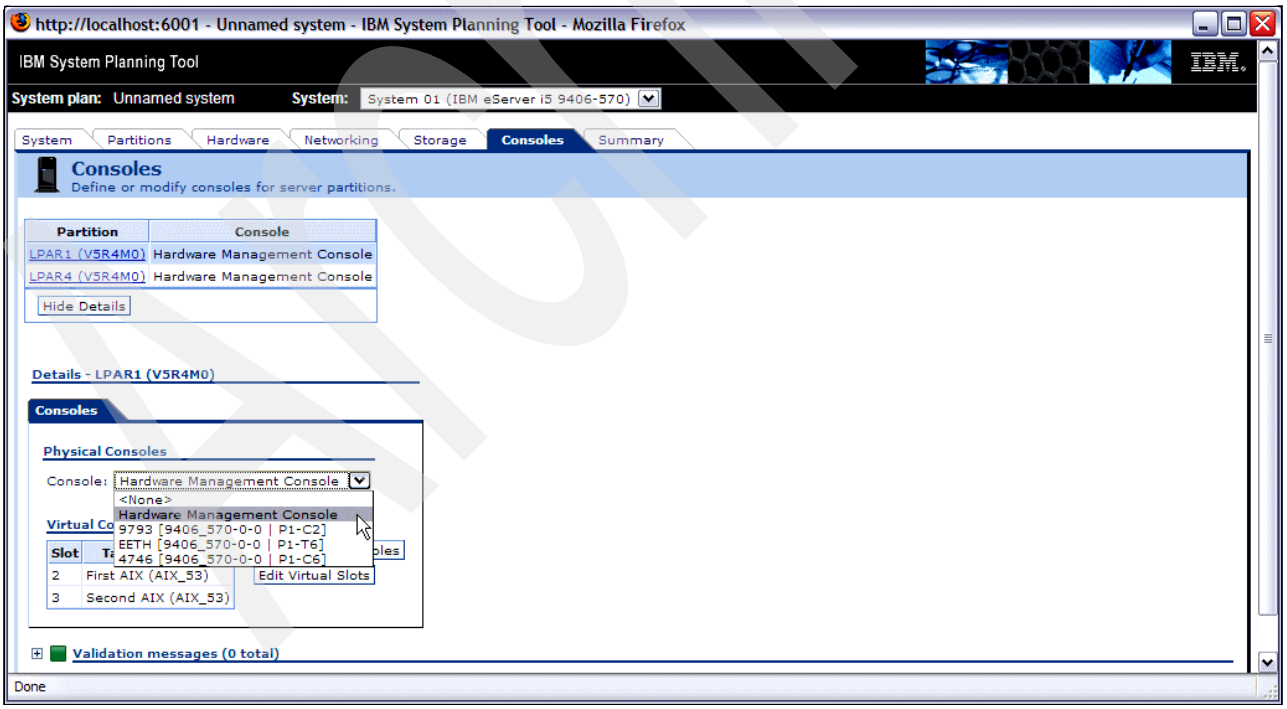

Figure 2-43 Assigning physical consoles

Virtual consoles

Second is to assign virtual consoles. In Figure 2-44 we have selected LPAR1 and clicked Edit Virtual Consoles. We do this to relate the virtual consoles of client partitions to a server partition.

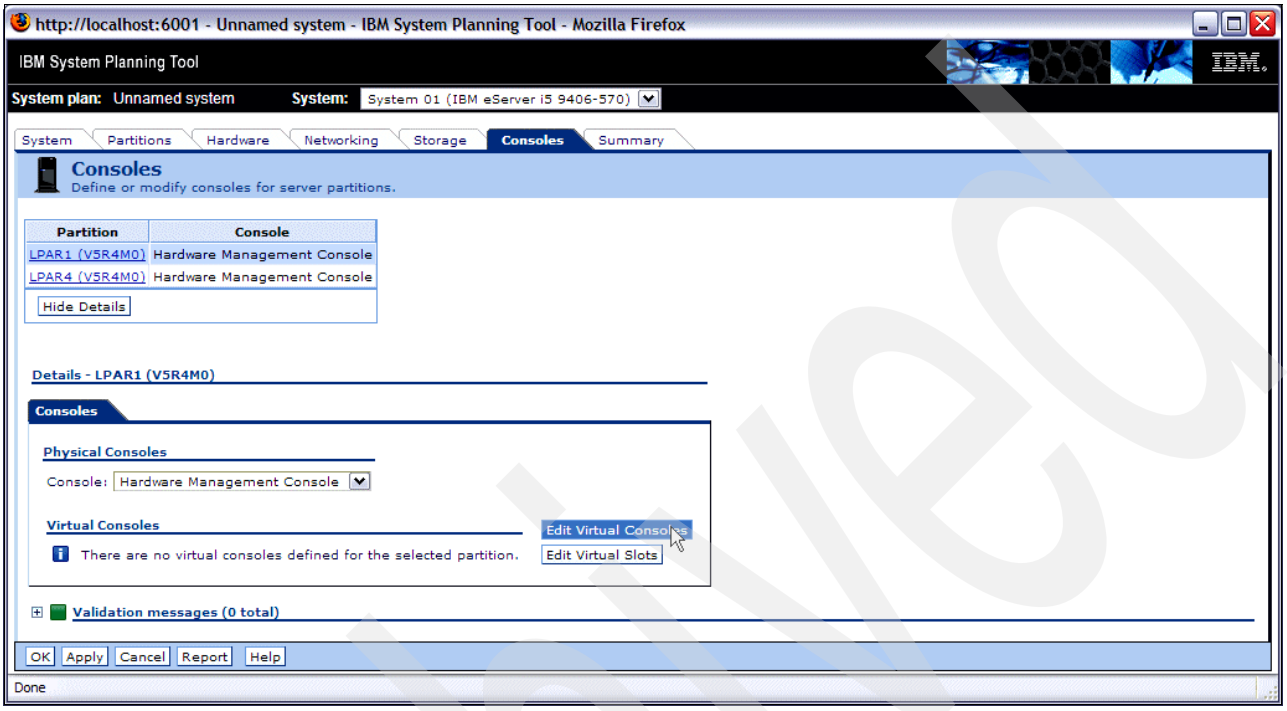

Figure 2-44 Edit Virtual console

Since we have more than one client partition being served by LPAR1 we are presented with the option to assign some, all, or none of their virtual consoles to the server partition. See Figure 2-45.

If you had already assigned them to LPAR1 but you really wanted these virtual consoles to be assigned to LPAR4, you would close this window, which puts you at the main Consoles window. Then select **LPAR4** and click **Edit Virtual Consoles**, check the check boxes, and click **OK**.

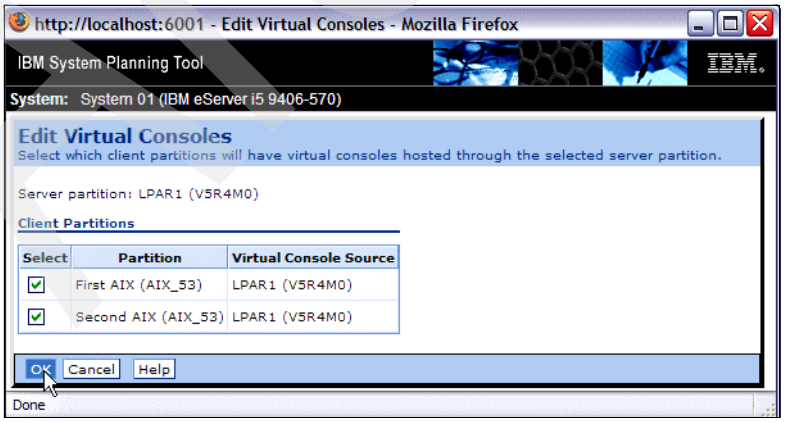

Figure 2-45 Assigning virtual consoles

When you do this the virtual console assignment is automatically removed from LPAR1 and moved to LPAR4, as shown in Figure 2-46.

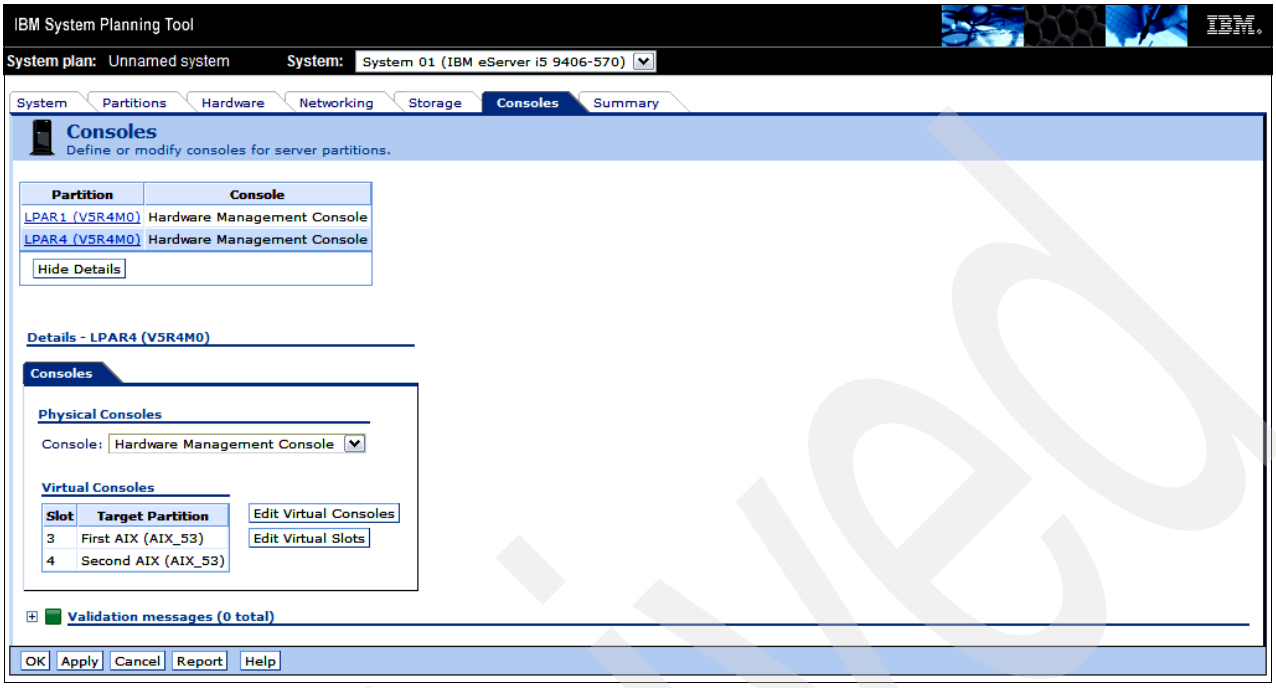

Figure 2-46 Virtual consoles moved to LPAR4

At the bottom of the window there are several action buttons.

If you click OK or Cancel you will be taken back to the Work with Planned Systems window.

The Apply button applies whatever changes you have made. You can use this after every change or a group of changes or not at all. However, if you do not use it and you accidently close the browser you are in or if SPT were to crash, your changes would be lost, even though you would still get a Reconnect option when you started again.

Attention: We recommend that you click **Apply** frequently.

Clicking Help presents you with help text that applies to the window you are on.

The Report button initiates the System Plan Viewer. It can be used at any time including, while you are creating a design. Many people use that capability to cross check what they have already placed and more importantly what they have not yet placed, by partition and as an overall total. Doing this can help you make sure you have all items to meet your requirements.

2.7 System Plan Viewer

As stated earlier, you can use the Report button to initiate the System Plan Viewer. This is available from many SPT windows.

Note: The System Plan Viewer is packaged as part of the System Planning Tool Version 2. This is the context in which we describe the System Plan Viewer in the next series of figures.

The System Plan View is also packaged as part of HMC Version 7.3 as part of the partition deployment based upon a saved system plan.

Click **Report**, as shown in Figure 2-47.

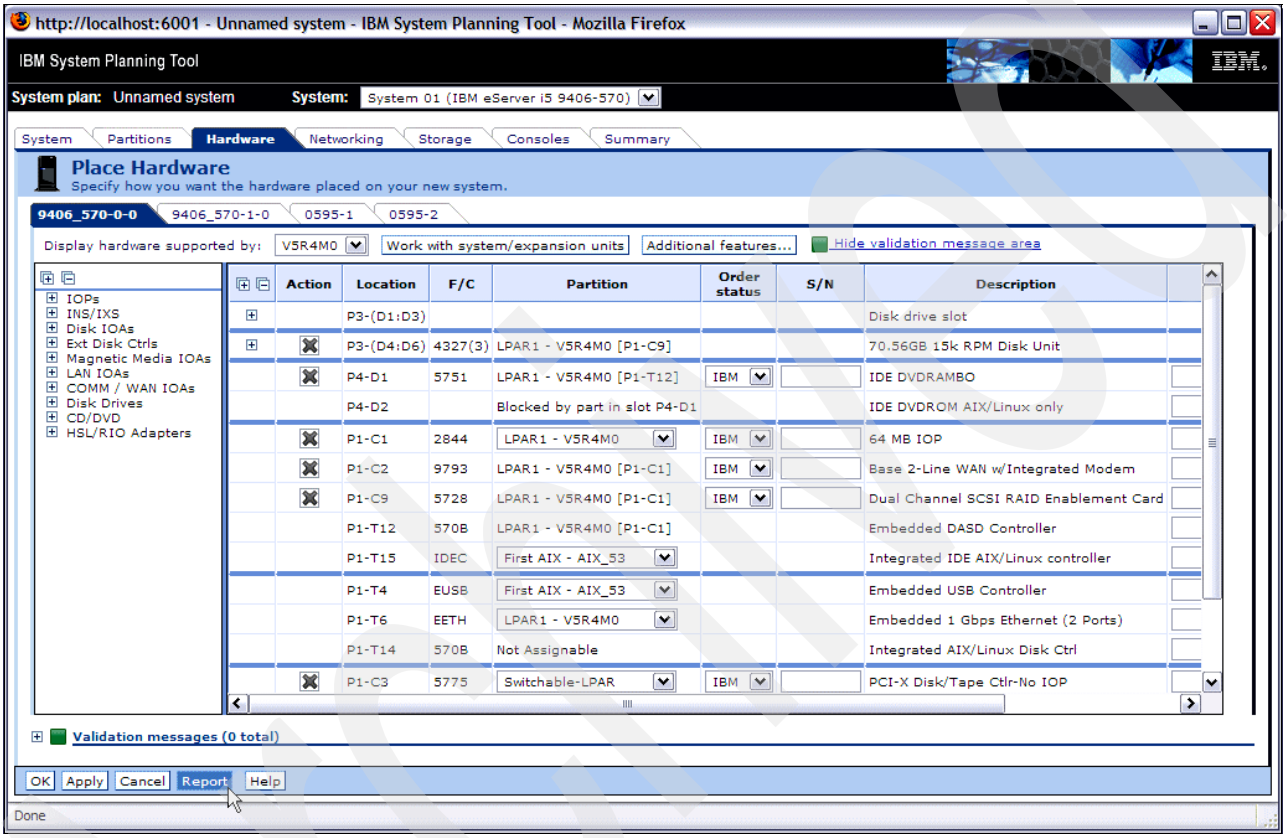

Figure 2-47 Selecting report

Figure 2-48 shows the first window of the report. In the left side bar there is a navigation tree for everything in the system plan. If you click an item in the tree you will be taken to the first page of that item.

Notes: When viewing the plan, the system/expansion units and the partition names may not be in the order you added them, as they are sorted alpha/numerically by the System Plan Viewer.

Once you have selected Report you can view a set of windows showing the hardware information in various graphical representations. This is considered the *Default Report* format, which is shown in Figure 2-48.

When launching the System Plan Viewer from within the System Planning Tool (not when launching the System Plan Viewer from the HMC), you have an additional *report option*. When the displayed format is the default report format, the “Report button” changes to the Text Report button (shown in the bottom left of Figure 2-48). Clicking the Text Report button will turn the HTML report into plain text. You can click the **Save** button, located at the footer of the window when the Text Report button appears, to save the report as a .txt file on your local workstation.

Figure 2-48 First System Plan Viewer page

In our example we clicked System 02 under Messages. This shows a table of the unresolved messages for System 02, as seen in Figure 2-49. There is no selection for System 01 under Messages because there are no unresolved messages for that system.

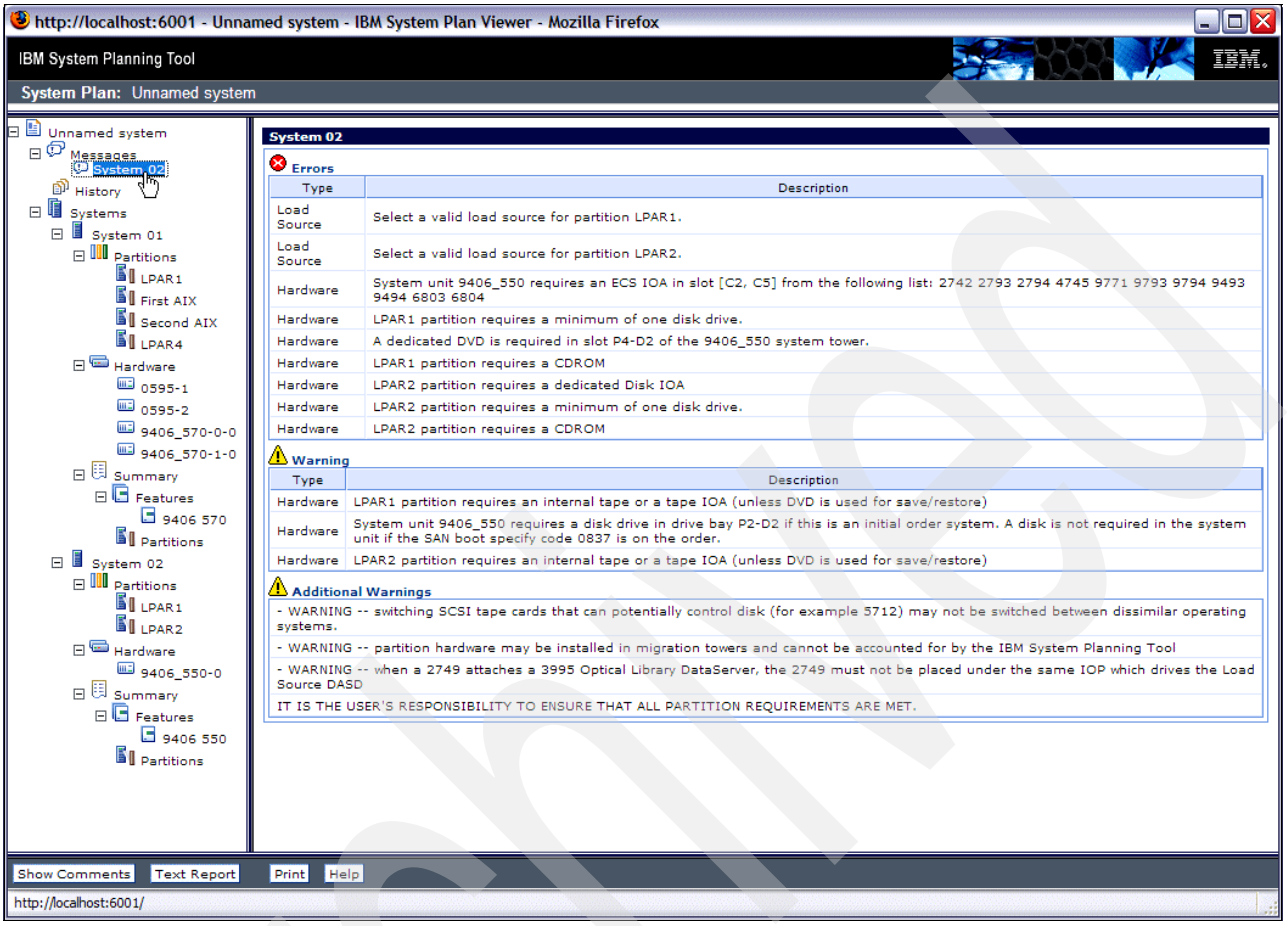

Figure 2-49 Unresolved messages regarding system 02

Next we click System 01 and see a listing of the basic system attributes, as shown in Figure 2-50. We also see the beginnings of LPAR1. This is because the report is essentially a long one page document with no page breaks.

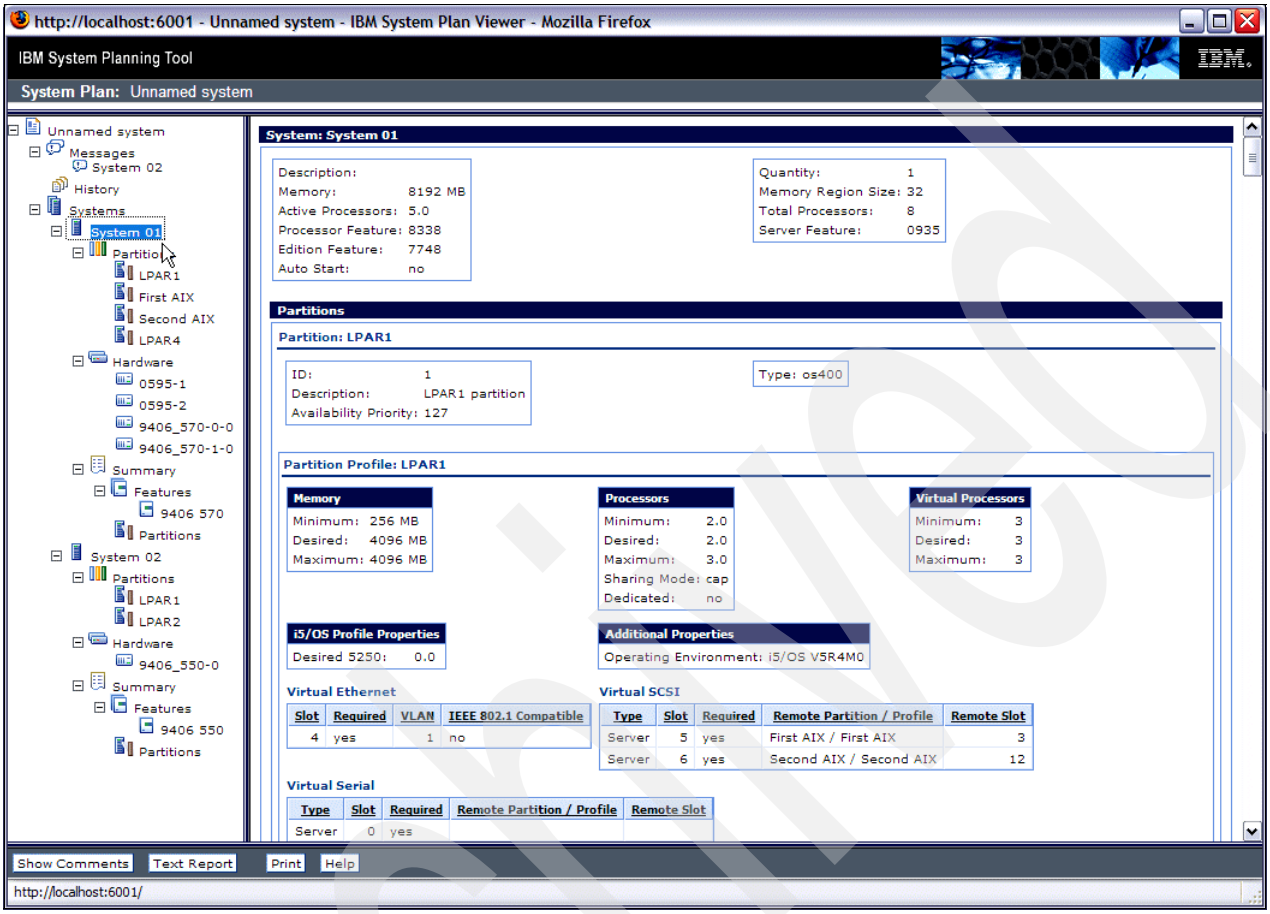

Figure 2-50 Initial viewer page for system 01

In Figure 2-51 we have clicked LPAR1 in the left navigation tree area and show the attributes of the partition and the physical contents (details) of that partition. Depending on the number of system and expansion units involved in a partition you may have to scroll down to see all of it.

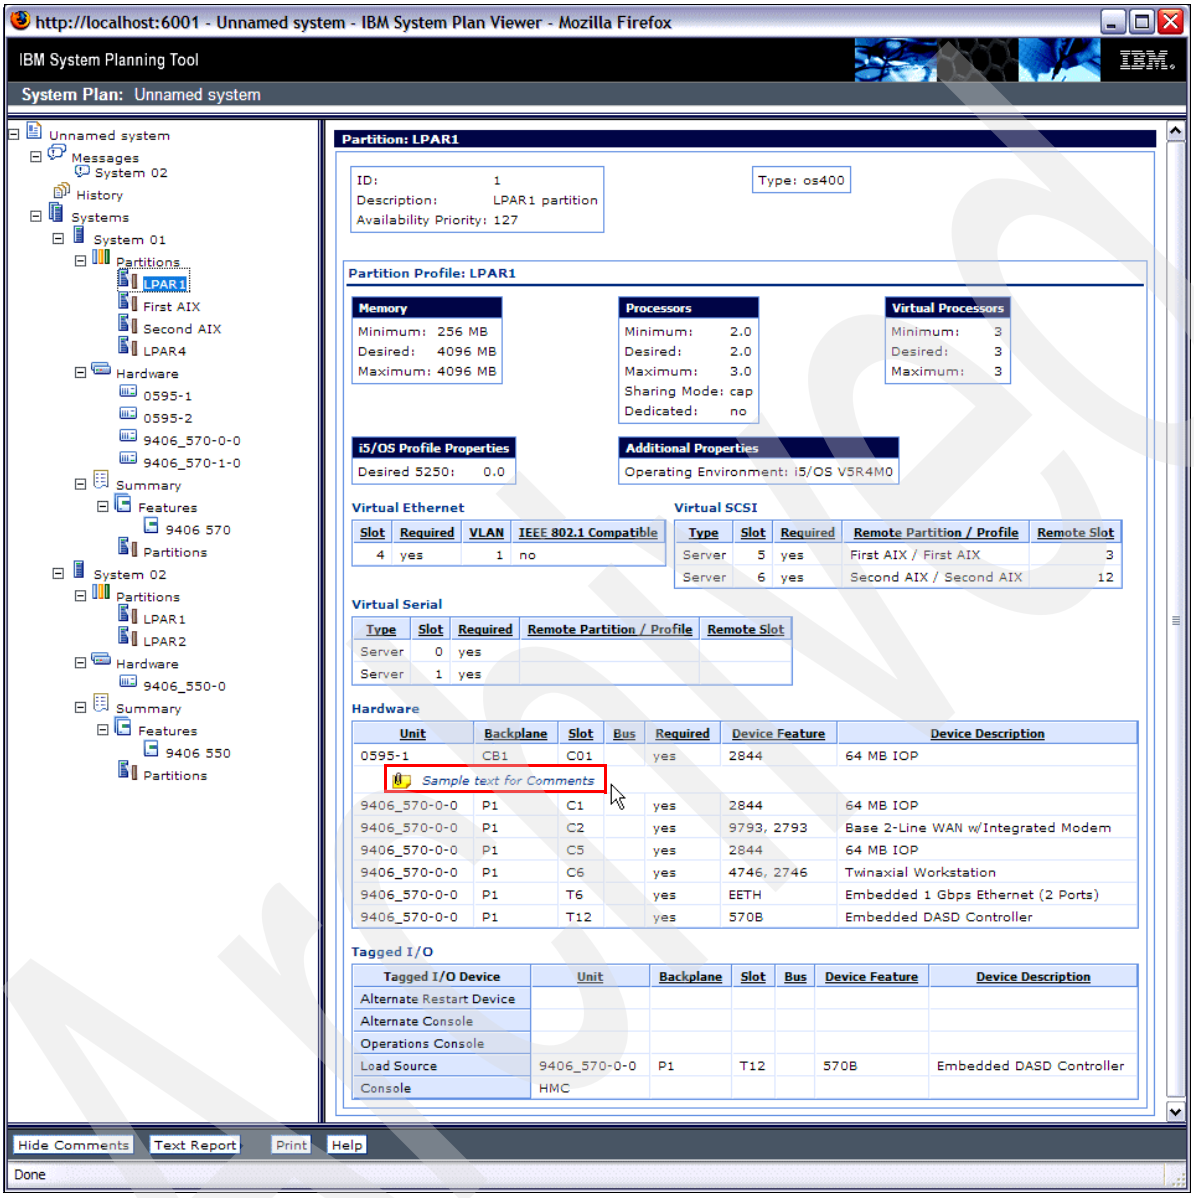

Figure 2-51 Attributes and detail page for LPAR1

Note that comments entered with the System Planning Tool can be seen by clicking the **Show Comments** button on the window footer, shown in the lower left of Figure 2-50 on page 56. For example, we entered `Sample text for Comments` in the Comments column for feature #2844 for I/O Expansion unit 0595-1 in Figure 2-30 on page 39.

You can see this text in Figure 2-51, just below the 0595-1 entry.

If you click a particular piece of hardware such as an 0595 IO enclosure, as shown in Figure 2-52, you are presented with the card/disk layout and contents in tables, as well as a diagram of the physical unit. The diagram can be collapsed to save space in the report.

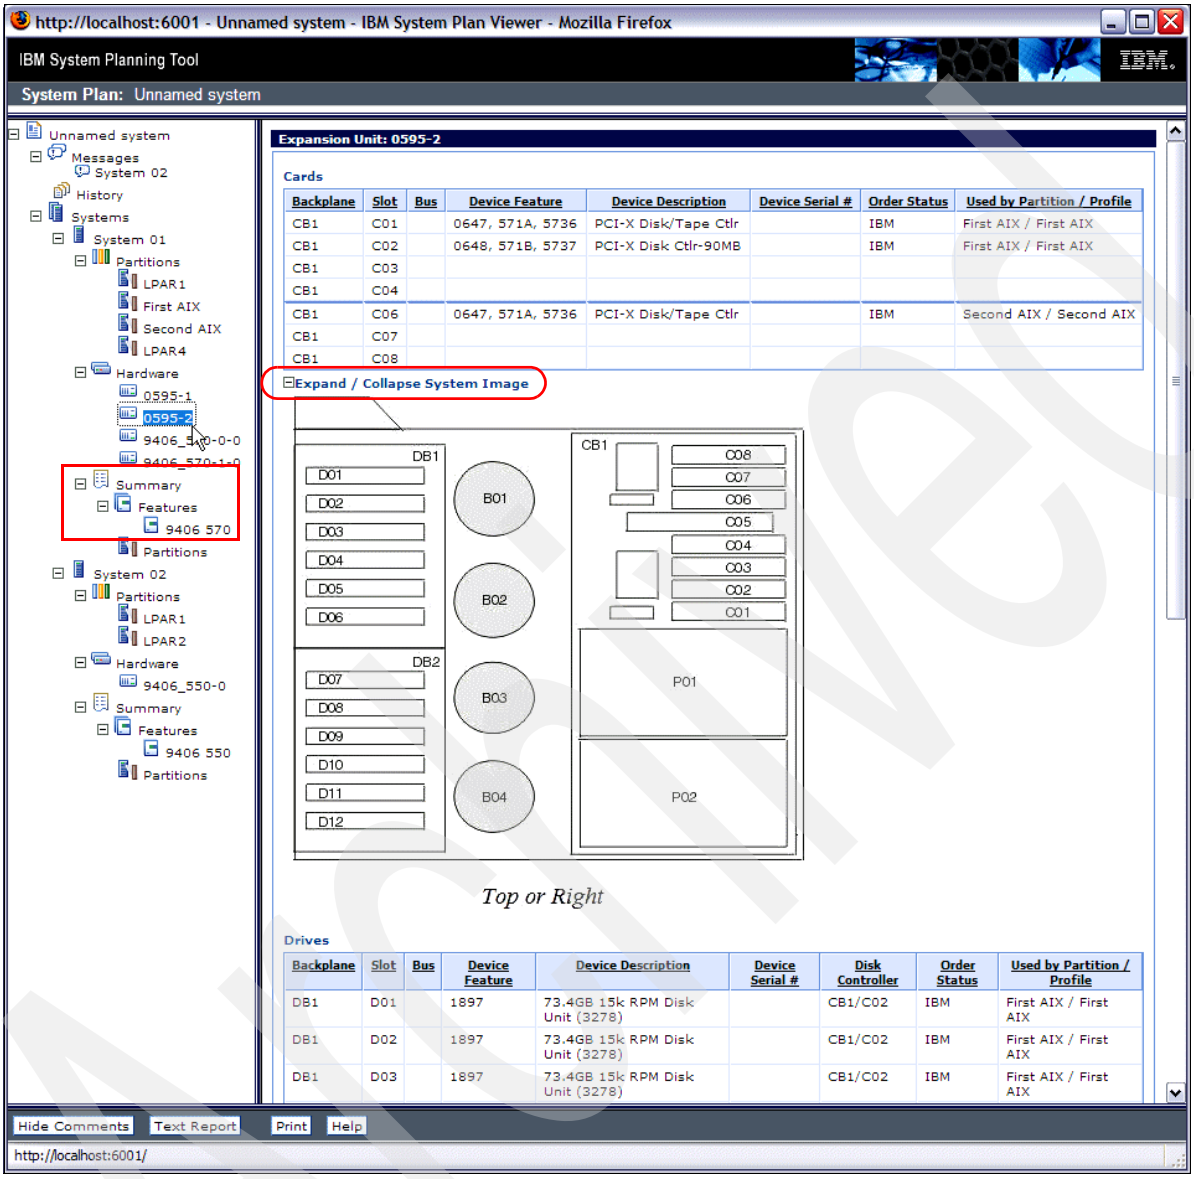

Figure 2-52 0595 example with expansion unit diagram

In our example, if you click Summary, Features, or 9406-570 (in the left pane), you are shown a table with the total count of features by partition and a table with the total feature count for the whole system. Figure 2-53 shows an example of Summary. If you are working on a system that has many partitions you will have to scroll down to see everything.

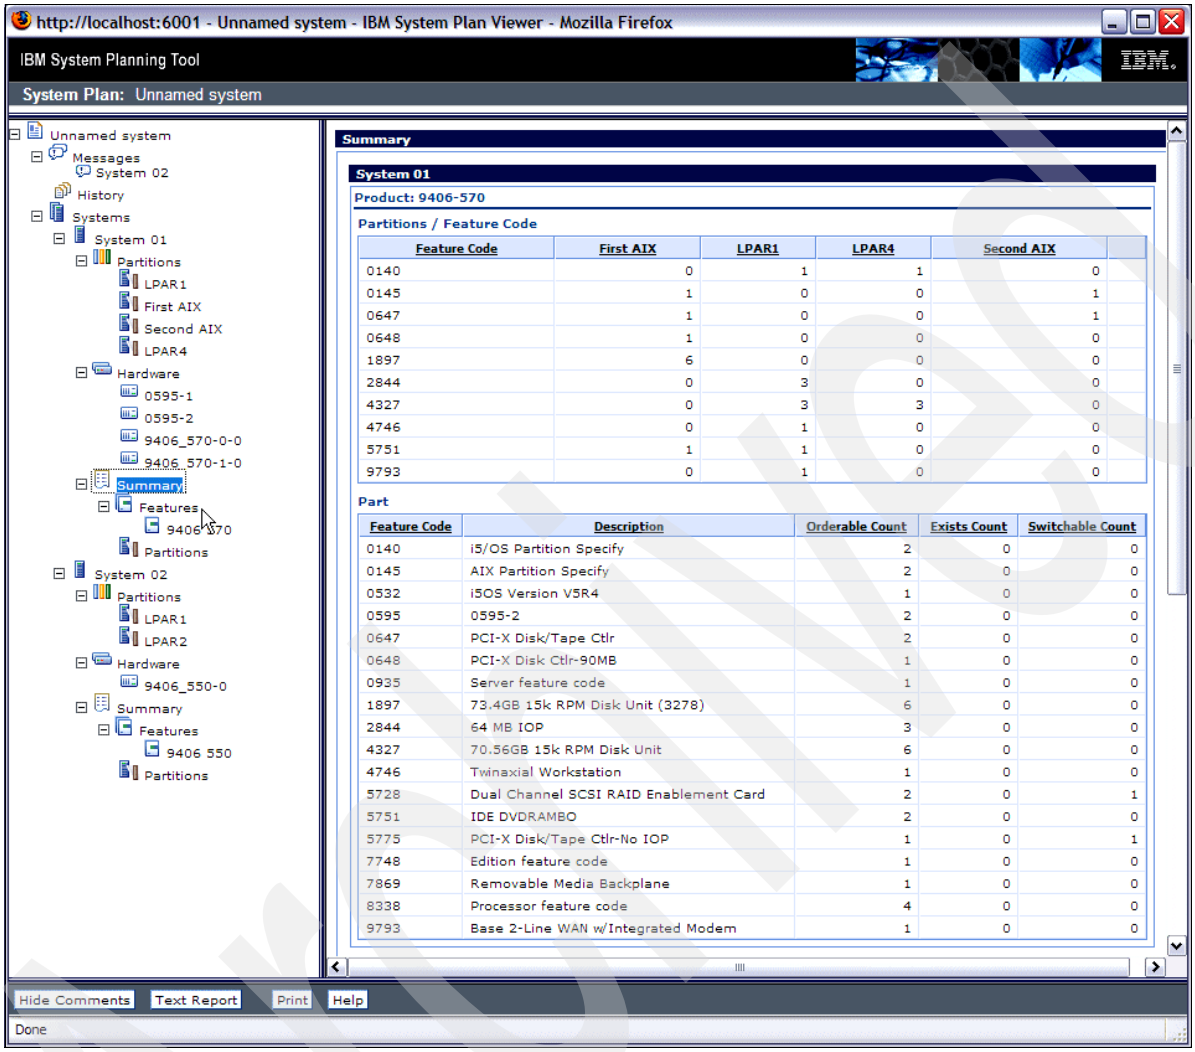

Figure 2-53 Totals by partition and by system

Clicking Partitions presents you with tables containing the attributes of all the partitions on the system. See Figure 2-54.

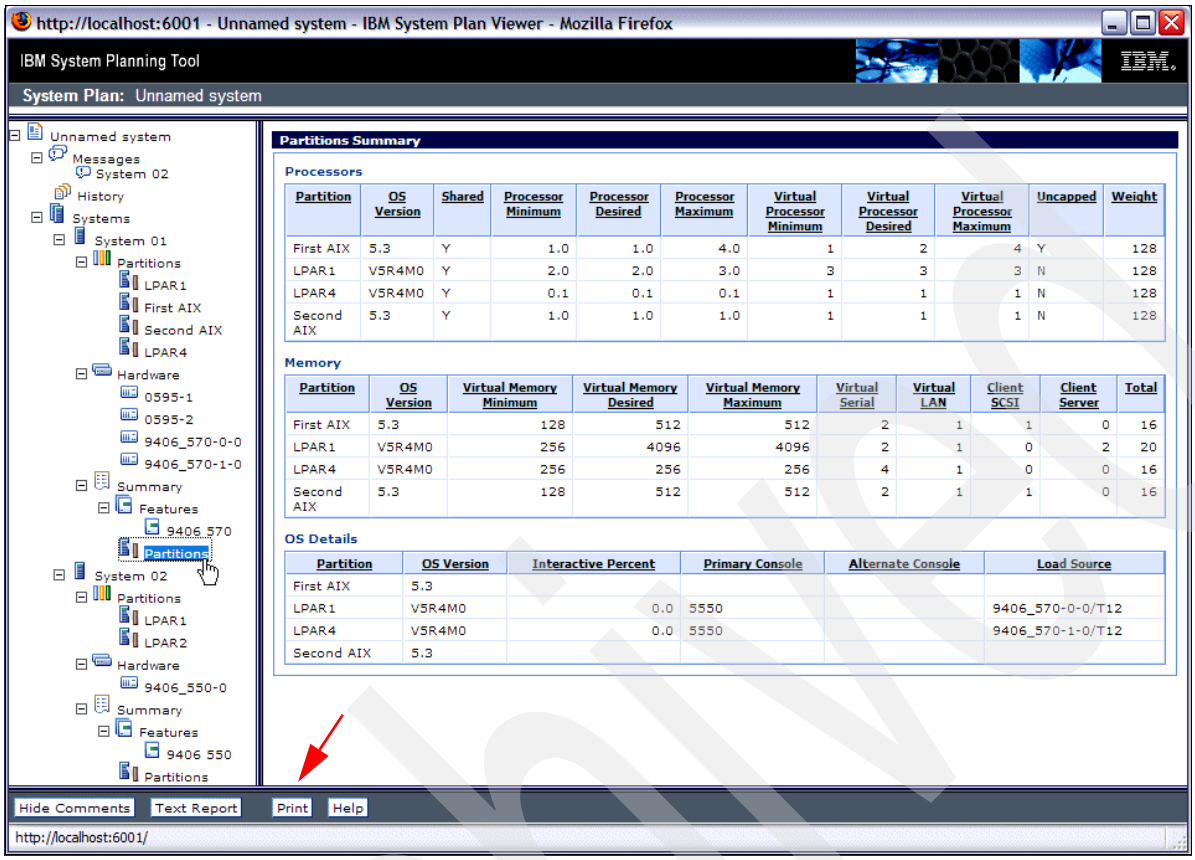

Figure 2-54 Partitions summary

These summary views of partition configuration should be reviewed before proceeding to ensure that your partitions are configured the way that you want them. This is especially important the more partitions that you are configuring. You will want to save the system plan and perhaps the System Plan Viewer output to assist you when you want to deploy the partition configuration on the actual hardware.

2.7.1 Printing the System Plan Viewer

Clicking the **Print** button in the window footer area will print the report to the printer of your choice.

You must understand that your output is very section specific. In the navigation tree, if you have selected Systems you will print the entire system plan including multiple systems, if you configured more than one.

If you have selected System 01 you will print only that system. Selecting LPAR1 in our example only prints the details of the LPAR1 partition. Selecting a specific piece of hardware (for example, 0595-2), the printout will only contain that expansion unit including the unit diagram, unless you have collapsed the diagram.

As stated earlier, the Text Report button will appear within the window footer when launching the System Plan Viewer from within the System Planning Tool. Click the **Text Report** button

to switch the report from HTML format to plain text format. Clicking the button again will switch back to HTML format, toggling between formats.

With the report shown in plain text format, you can click the **Save** button in the window footer to save the report as a .txt file on a disk drive on your local workstation.

2.8 System plan completion

When you are satisfied with your completed system plan configuration, you need to perform one or all of the actions described in the following sections:

- Save the system plan (.sysplan file).
- Export the system plan.

2.8.1 Save sysplan

It is very important to save your system plan (.sysplan file). This file is intended to be used later when deploying your system plan to the actual hardware configuration.

Once all of the components of the system are showing a valid status on the Work with Planned Systems display, the system is deemed complete. If you have not already done so, this is when you should use **Save**, as shown in Figure 2-55. There is no *save as* option, but you will get the opportunity to specify the file name before the actual save is accomplished. This file *must* be saved with a .sysplan file extension. You will not be able to reuse the file if you save it with any other extension.

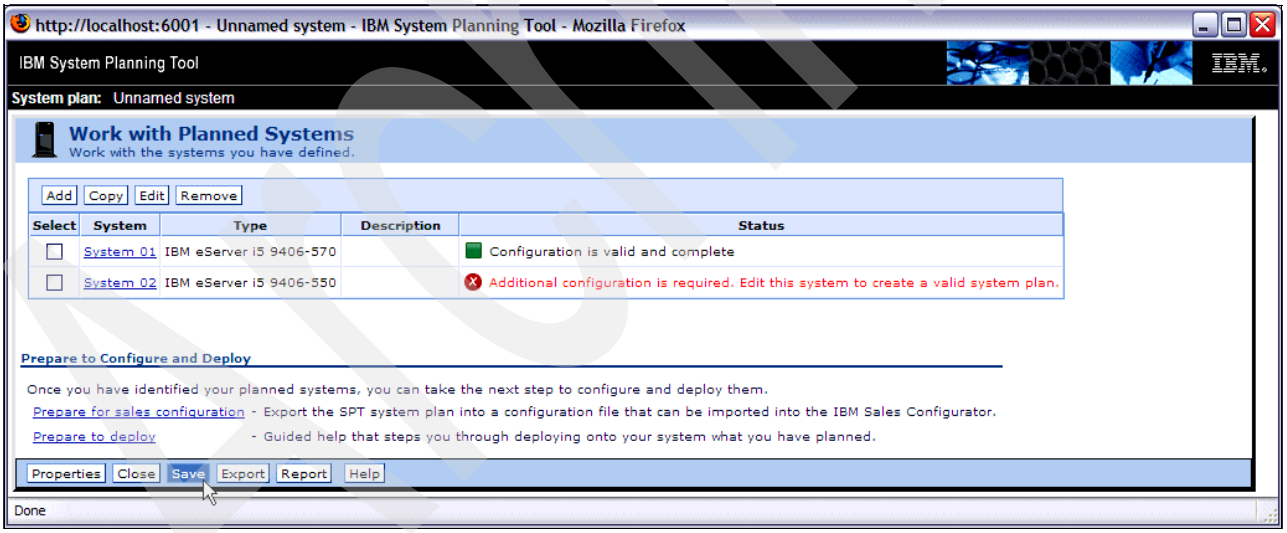

Figure 2-55 Save the sysplan

2.8.2 Export sysplan

One of the important uses of your system plan is to use it as input to the IBM sales configurator tool.

You do this by exporting your system plan as a .cfr file that can then be imported into the IBM configurator tool. When you have created and saved a .cfr file to disk, using the export function, you must import the system plan into the IBM configurator tool to ensure that all the

components in the plan are valid and the features are orderable. Once the system plan is validated by the IBM configurator tool as orderable, a priced proposal can be produced.

When using Export in a system plan that has multiple systems you can select the systems to be included in the export function, as shown in Figure 2-56.

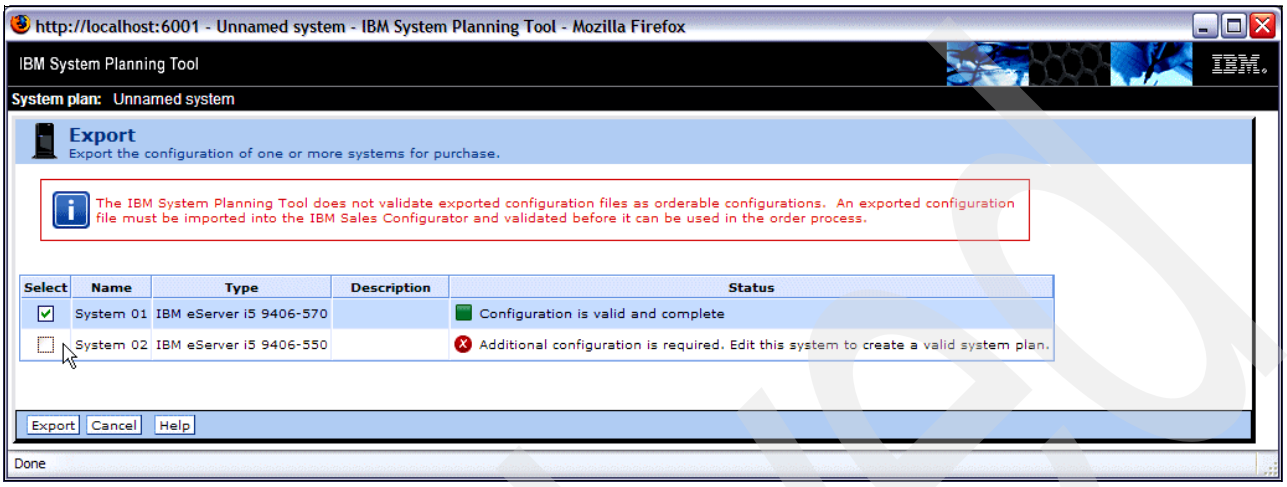

Figure 2-56 Export system selection

Click **Export** and a save window appears prepared to save a .cfr file, as shown in Figure 2-57. Save the file in a folder of your choice. This file contains all the information regarding the system you planned including all of the placement information. All of this information is passed on when it is imported into the IBM configuration tool.

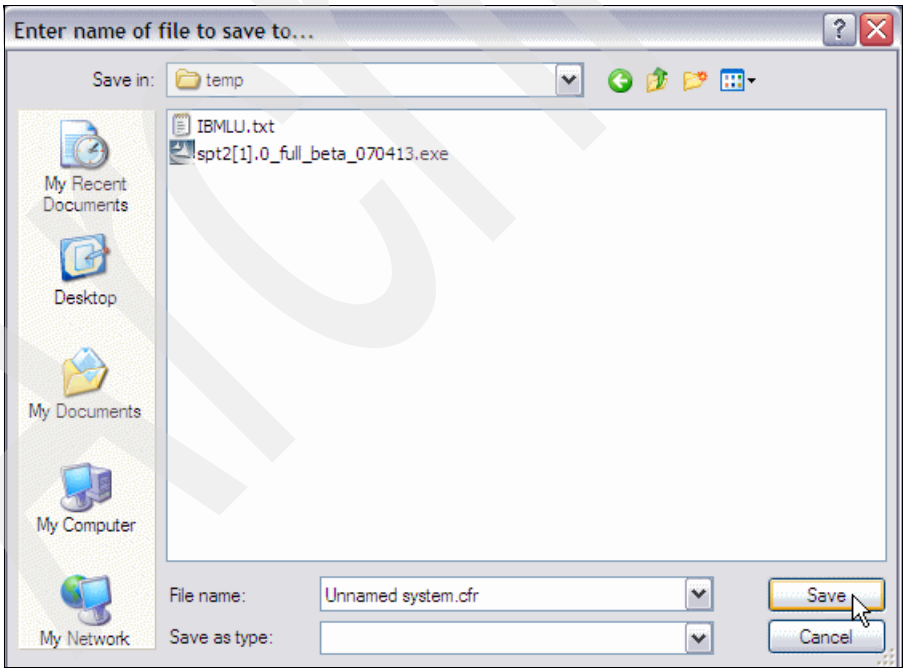

Figure 2-57 Saving file via Export

To exit the Export function click **Cancel** in the Export window (Figure 2-56 on page 62), and you will be returned to the Work with Planned Systems window (Figure 2-58).

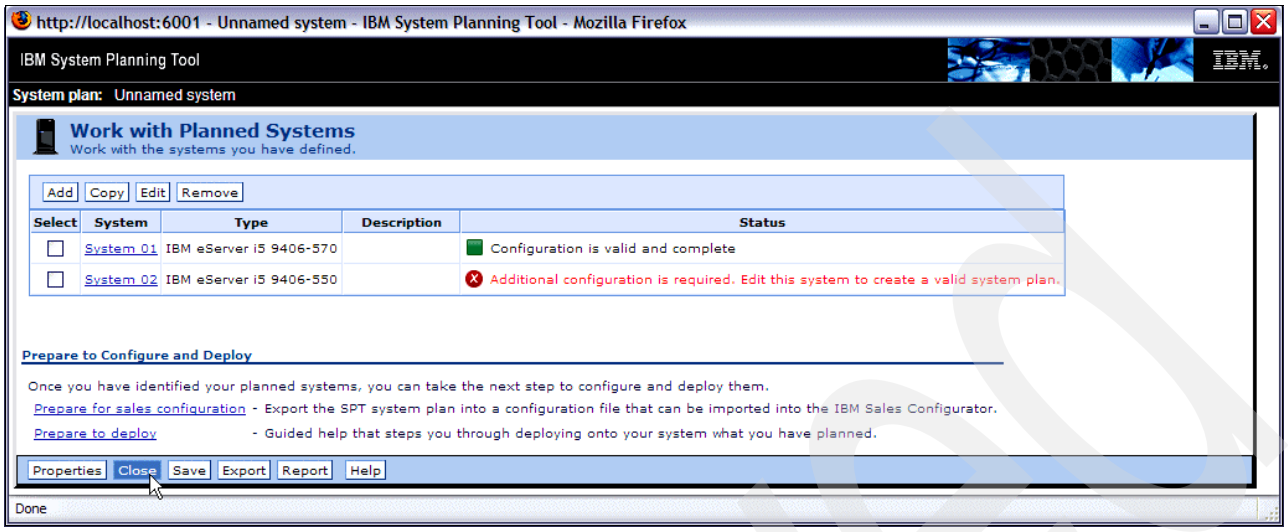

Figure 2-58 Work with planned systems

The Properties button shows only the system plan name and description box. These fields are both in edit mode, so you can make changes if needed.

To exit the SPT click **Close**. Respond with **Yes** to the *are you sure* message. Then close both browser sessions. To completely exit the SPT you must also close the SPT icon in the system tray.

When a system is ordered using the .cfg file exported to the IBM configurator tool, IBM manufacturing can build and deliver the system exactly as planned. This is called Customer Specified Placement (CSP). More information can be found on CSP at:

`http://www.ibm.com/servers/eserver/power/csp/`

2.8.3 Sysplan for deployment

When the customer system arrives on site, the system plan saved earlier as a .sysplan file in SPT can be imported to a controlling Hardware Management Console (HMC). You can then use the .sysplan file to deploy the system and its partitions through the HMC interface. See Chapter 3, “System plans and the hardware management console” on page 65, for further information about deployment.

Deployment of a Virtual I/O Server partition has special processing, as described in Chapter 4, “Virtual I/O Server” on page 103.

3

System plans and the hardware management console

In this chapter we introduce another important aspect of managing system plans. The Hardware Management Console (HMC) can be used to create import, export, view, create, remove, and deploy system plans. The HMC code level must be at v7r3.1 or later, and include the latest service packs. The HMC provides a set of graphical user interfaces (GUIs) for these LPAR management functions.

Note: If you are connecting to the HMC from a PC workstation, there is no longer a requirement to install WebSM on the PC. The HMC with Version V7R3.1.0 or later installed should be accessed through a Web browser, not WebSM. Refer to 1.6, “System plans on the HMC” on page 8, for references to additional information about HMC V7R3.

The chapter reviews the operations of each of these GUI interfaces for managing system plans, also referred to as *sysplans*. All of the utilities are available through the HMC graphical interface. Some functions are also available through the restricted shell, commonly referred to as the command-level interface (CLI). We do not document how to use those command-line functions that have not changed since they were documented in *LPAR Simplification Tools Handbook*, SG24-7231. A list of those command-line functions, and information about those command-line functions that have been added or changed, can be found in 3.3.3, “System plans management using restricted shell (CLI)” on page 98.

3.1 System plans

A system plan, also referred to as sysplan because of the .sysplan file extension, is a representation of the hardware and partition configuration currently on a system, or the plan for deployment of hardware and configuration of partitions on a system, depending on how the sysplan file is generated. If generated from an existing POWER5 or POWER6 system using the HMC, the file will reflect the actual LPAR configuration of the server at that point in time if all the partitions are active. There will be less detail if one or more partitions are not active. If the sysplan is generated by the System Planning Tool (SPT), the file will reflect the intended LPAR configuration for a target server. The sysplan includes details on partition allocations of memory, processors, and the hardware required for each partition. Hardware allocations may be defined as owned by the partition, and therefore required for the partition to activate. The only other choice is for the hardware to be defined as shared, in which case the hardware is optional for the partition, which can be activated without the hardware. Shared hardware can be dynamically switched between two or more partitions.

The sysplan also includes general information about the system, such as system type and model, total number of processors present, and the number that are activated, and the total installed and activated memory. It has detailed information about the card slots in the processor enclosure and any I/O expansion towers or I/O drawers that will attach to the processor enclosure. The card slots are shown as empty, or occupied by IOP or IOA feature codes, and this level of detail is used for the hardware validation during the LPAR deployment process. It should be noted that, at this time, the sysplan file created by SPT includes device-level detail (for example, what type and number of disk units are attached to a storage controller IOA). In contrast, a sysplan created by the HMC does not, by default, include any detail of what is attached to and controlled by an IOA. For more information about this, refer to “Enabling hardware inventory collection from active partitions” on page 73.

A sysplan file is a composite object, which means that it could possibly include many files. The file’s description is imbedded in the file, as is the file level and last modifying application information. When a sysplan file is created on the HMC or imported using the HMC GUI, the file is stored on the HMC in a predefined directory. The directory path is */opt/hsc/data/sysplan*.

3.2 Using the HMC graphical user interface

On the HMC running Version 7 Release 3.1 code level, there is a item in the left navigation frame called *System Plans*. This is where all of the graphical interfaces required to manage system plans on the servers can be accessed directly from the HMC or remotely using the *browser-based* client connecting to the HMC. Single-click the **System Plans** task in the navigation area. See Figure 3-1. This will display the system plans management tasks window.

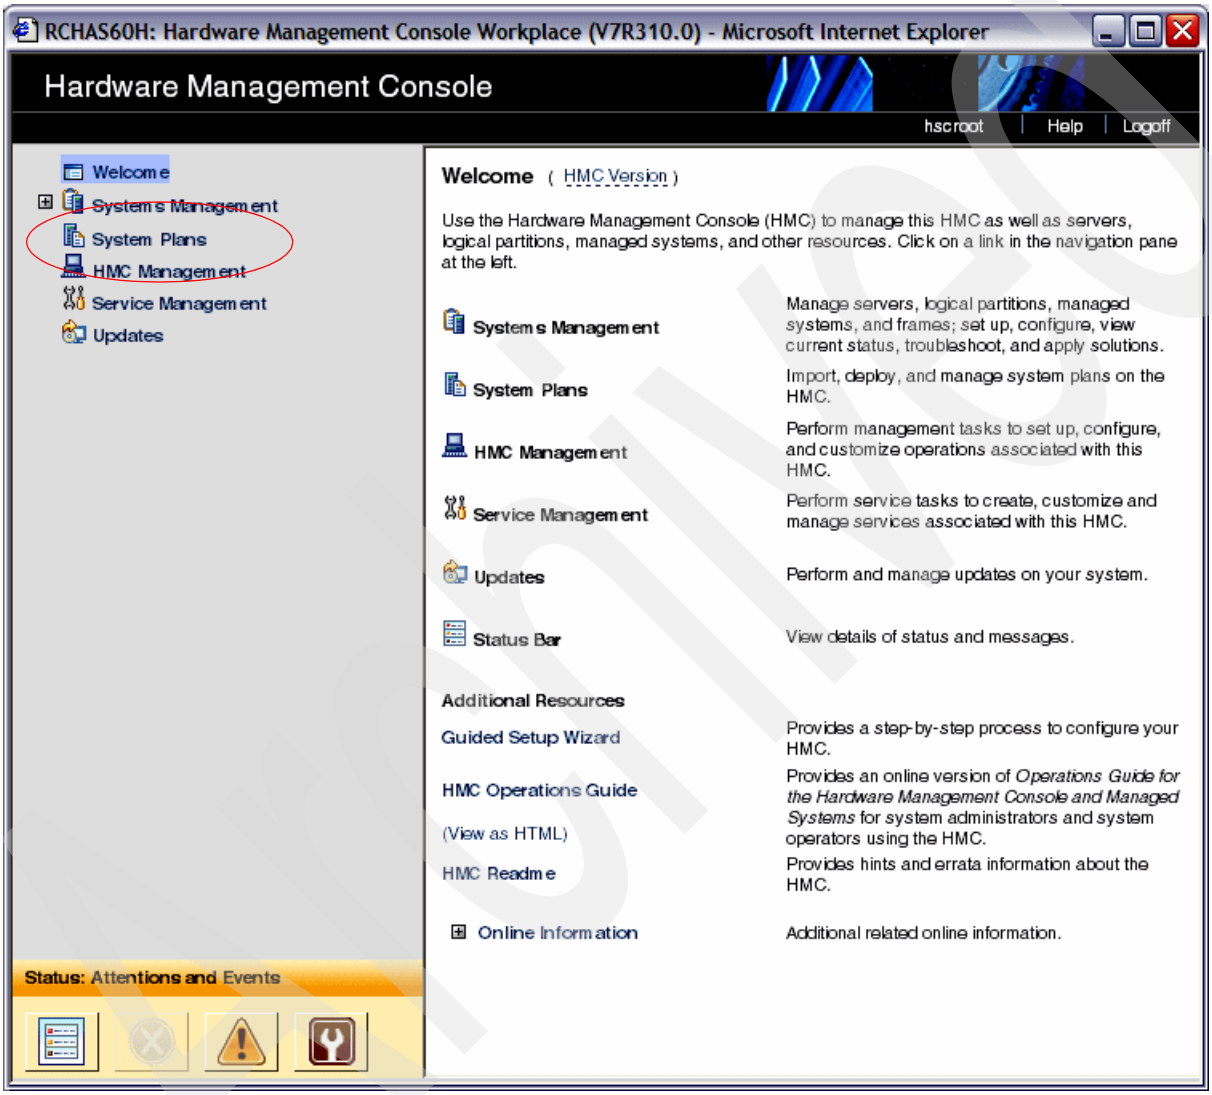

Figure 3-1 The HMC welcome page

The system plans management tasks window is shown in Figure 3-2. The upper section lists all of the system plans currently residing on the HMC. The buttons above the list lets the user select, deselect, sort, and filter, and manage the columns of the display table, and perform tasks on selected system plans. The task options are repeated in the lower Tasks section of the main system plans management window. Note that, with no system plan selected, the only options are to import a system plan or create a system plan.

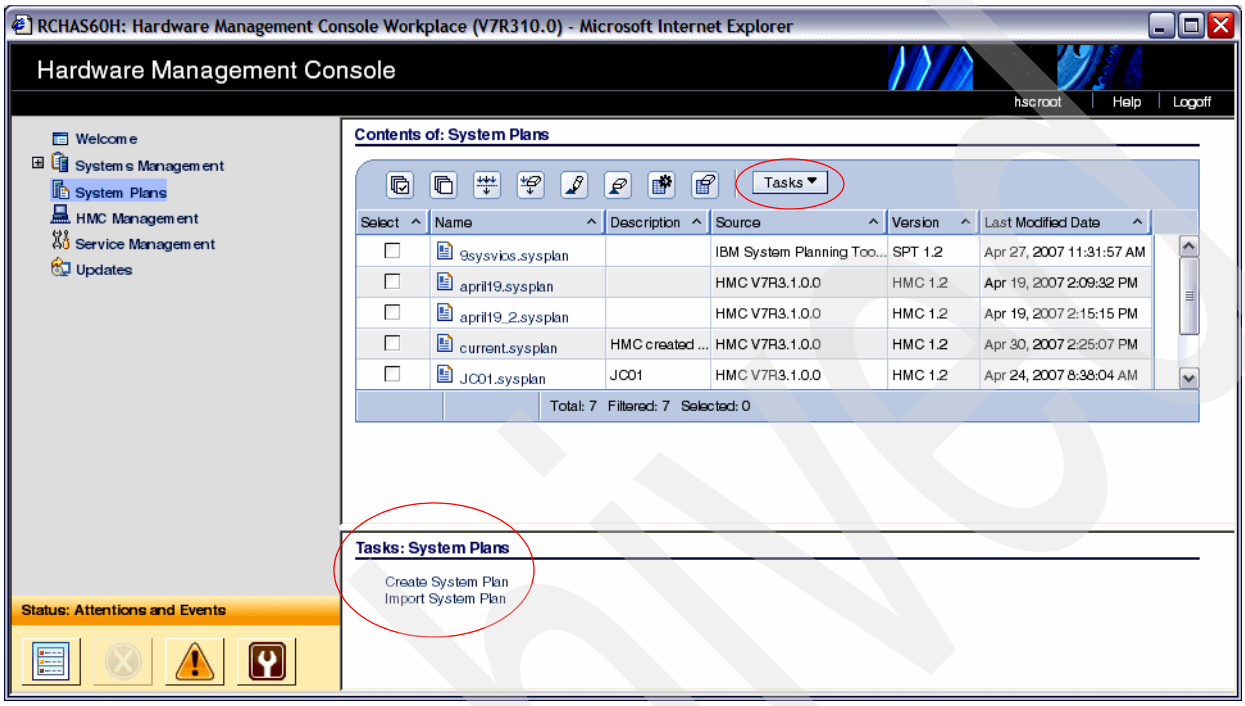

Figure 3-2 The main system plan management page

Using the HMC you can:

- Create a system plan.
- View a system plan.
- Deploy a system plan.
- Export a system plan.
- Import a system plan.
- Remove a system plan.

You can save a system plan created using the HMC interface as a record of the hardware and partition configuration of the managed system at a given time.

You can deploy an existing system plan to other systems that this HMC manages that have hardware that is identical to the hardware in the system plan.

You can export a system plan to another HMC (which imports the plan) and use it to deploy the system plan to other systems the target HMC manages that have hardware that is identical to the hardware in the system plan.

As listed previously, we have options (tasks) to view, create, deploy, export, import, or remove a system plan. Note that these tasks can be selected in either the Tasks drop-down menu or the Tasks links in the lower part of the right frame. The following sections provide more details for each option. Figure 3-3 will be a common starting point for each example.

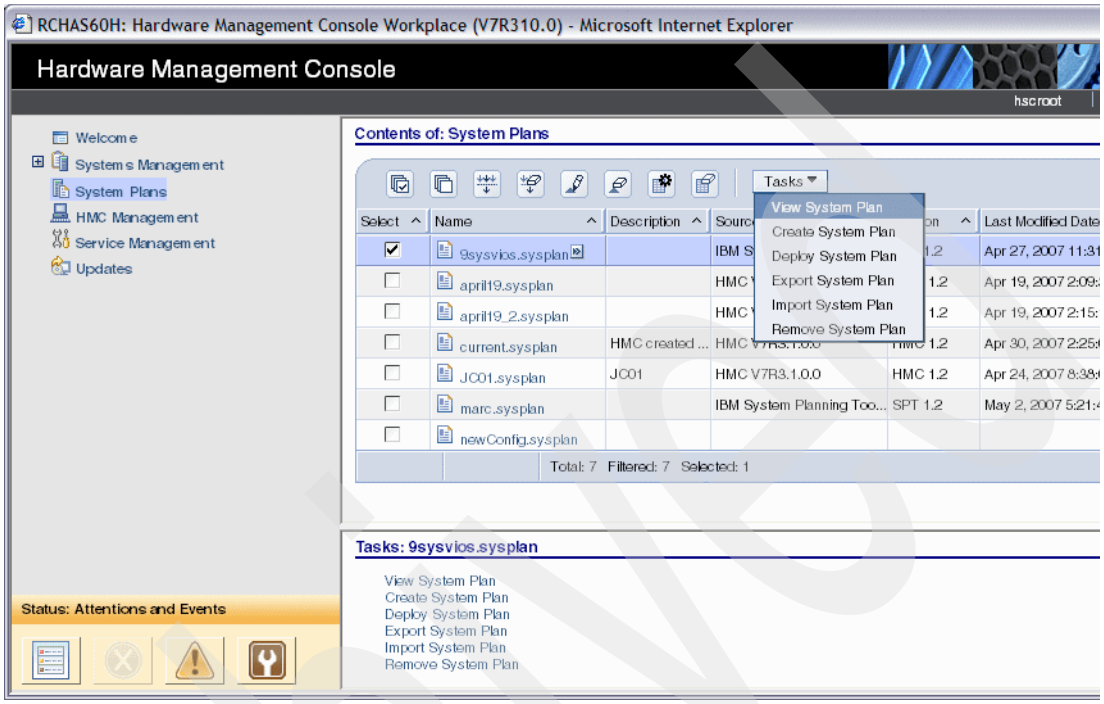

Figure 3-3 The system plan management page with a system plan selected

In our first example we have selected an existing system plan named 9sysvios.sysplan.

3.2.1 Importing a system plan to the HMC

The import operation provides the ability to load a system plan that was created using the SPT or created on another HMC. The system plan can be imported from one of the supported media types, such as CD, DVD, diskette; a USB device such as a memory card, a remote FTP site, or a PC connected to the HMC through a browser connection. First we prepare the media, if needed. Then we import the sysplan file.

From the System Plans task menu, select **Import System Plan**, which brings up the Import System Plan prompt window. Identify the system plan file name and whether it will be imported from media, an FTP server, or, if you are accessing the HMC through a PC-based Web browser, the sysplan file may be on that PC. In Figure 3-4 we show the import prompt window and that the system plan file is stored on a USB flash drive. The name of the file is newConfig.sysplan and the file was initially created using SPT and saved to the flash drive. The directory path to access the file on the flash drive is /media/sysdata.

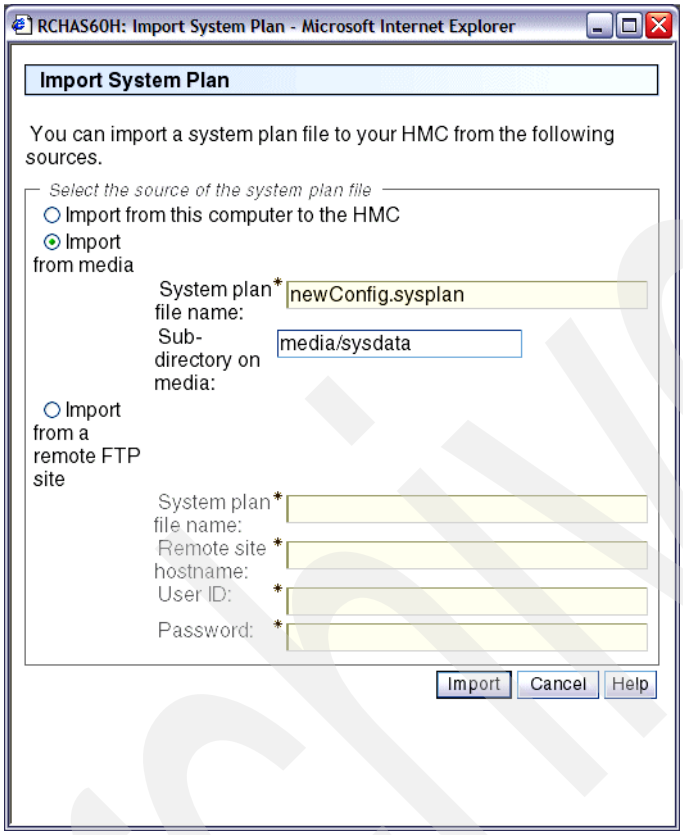

Figure 3-4 Import System Plan window

A successful import will result in this conformation window.

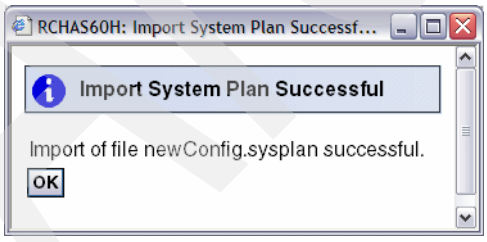

Figure 3-5 Importing a sysplan file success

3.2.2 Exporting a system plan from the HMC

System plans residing on the HMC may be exported to media, an FTP server, or, if you are using a PC to access the HMC through a browser, to a directory on the PC. The process is very much like the importing of a sysplan file. If you are exporting to media, that media will need to be formatted for use with the HMC.

Preparing the media

In order to export to external media, that media must be in a format available to the HMC. The easiest method of doing this is using the Format Removable Media task.

First select **HMC Management** from the left navigation frame. Then select **Format Media** in the right window. That brings up the media selection box. See Figure 3-6.

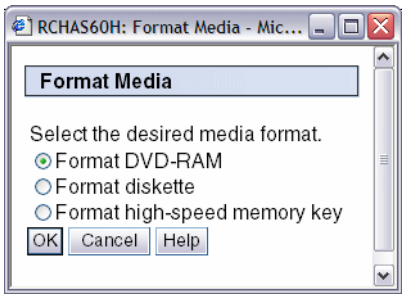

Figure 3-6 Format Media GIUI

If you have a USB memory key insert it into a USB slot on the HMC. If you have a diskette or CD that needs to be formatted, insert it into the disk or CD drive. Select the correct device to format and click **OK**. The memory format process starts and completes. Insert the media into the PC and load the system plan file using the save function in SPT or by browsing to the file and copying the sysplan file to the media.

Exporting the system plan

From our starting point in Figure 3-3 on page 69, we select a system plan to export. Click **Export System Plan** either from the Tasks drop-down menu or the content Tasks menu in the lower portion of the window. A dialog box will come up asking where the system plan is to be exported. See Figure 3-7 for an example. Here the name of the sysplan file is april19.sysplan and the target is media/USBstick directory.

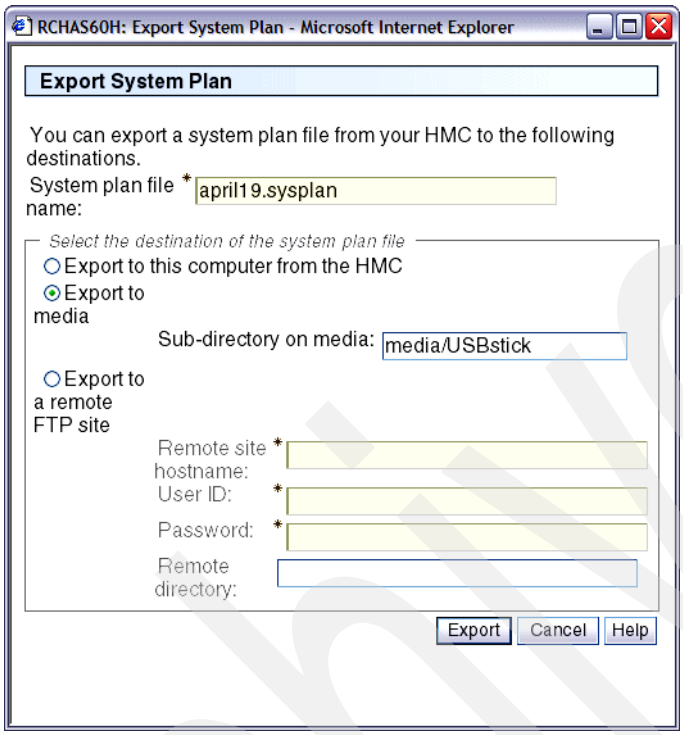

Figure 3-7 Export System Plan window

Click **Export** to initiate the export process. A results window with a success indication or an error message will indicates the result of the export.

3.2.3 Creating a system plan on the HMC

We can create a system plan for a system controlled by the HMC. The system plan will have information about the current partition definitions and hardware allocations. Processor, memory, and PCI cards will be identified in the system plan, even if they are not owned by a partition.

Notes: Hardware controlled through an IOA controller, such as disk units and external media devices, will not be represented in the system plan unless the owning partition is up and running. See “Enabling hardware inventory collection from active partitions” on page 73 for more information.

The sysplan file created by the HMC cannot be imported to the SPT for editing. The sysplan file can only be deployed and viewed. This can be done either on the HMC the file was created on, or an HMC to which the file has been moved.

From our starting point in Figure 3-3 on page 69, select **Create System Plan**. The Create System Plan window will prompt you for the system name, sysplan file name, a description, and a choice to view the system plan after creation. See Figure 3-8. After the requested information has been entered, click **Create**.

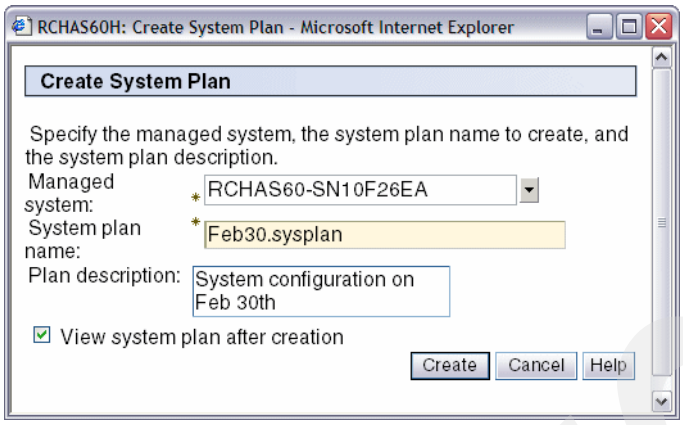

Figure 3-8 Create System Plan window

Following the successful creation of a system plan, the following window is shown (Figure 3-9), and the system plan is now in the list of plans on the HMC. Click **OK**.

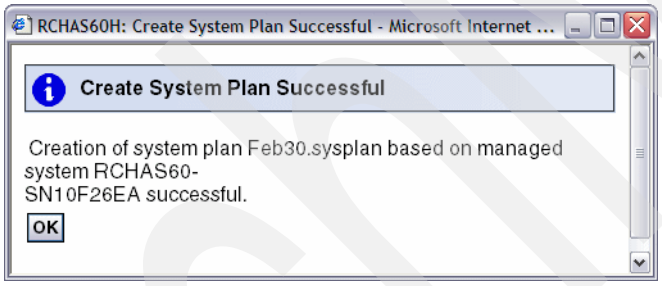

Figure 3-9 Success creating a system plan

Enabling hardware inventory collection from active partitions

When you use the HMC to create a system plan for a managed system, you can capture partition configuration information and a base set of associated hardware configuration information. If you have partitions already active, you can maximize the information the HMC can obtain about the hardware.

To maximize the information that the HMC can obtain from the managed system, power on the managed system and activate the logical partitions on the managed system, assuming that they already exist, before creating the new system plan.

Additionally, you should have previously set up Resource Monitoring and Control (RMC) on your HMC prior to creating a system plan to capture the most detailed information. Although it may cause the creation of the system plan to take several more minutes to finish processing, by using RMC you can capture disk drive and tape drive configuration information for a managed system in the system plan. You can view this more detailed hardware information using the View System Plan task.

The following steps are required to enable the HMC's internal inventory collection tool (invscout) to be able to perform its most detailed hardware inventory retrieval operations.

1. On the HMC, first open HMC Management to get the background right pane shown in Figure 3-10. Select **Change Network Settings**. On the Customize Network Settings window select the **LAN Adapters** tab.

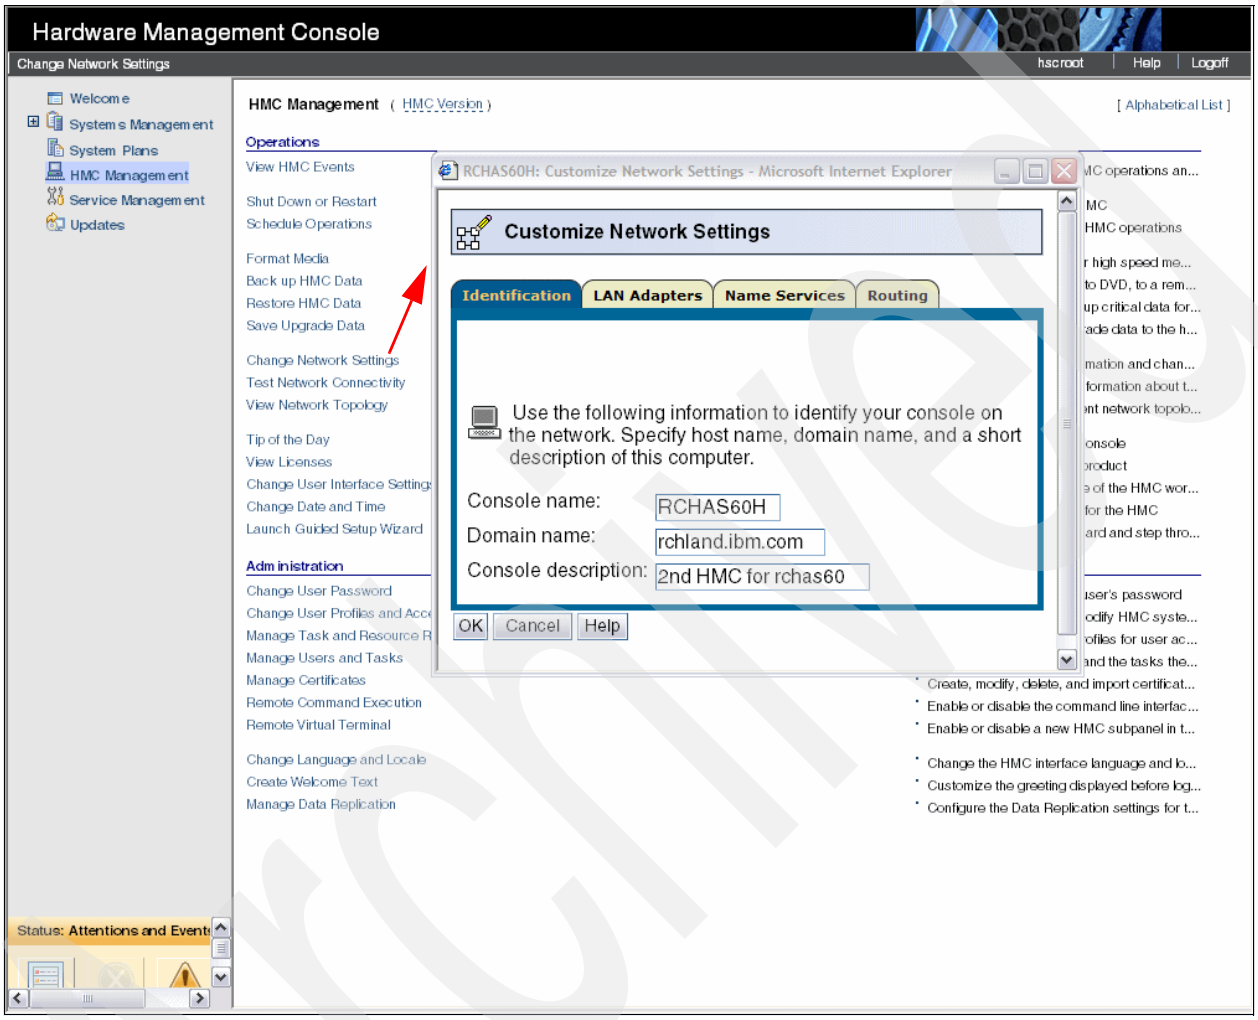

Figure 3-10 Customize Network Setting - LAN Adapters for enabling Remote Management and Control

2. On the LAN Adapters window (not shown) select the **eth0** LAN Adapter and click the **Details** button. The adapter details for eth0 are shown in Figure 3-11. First click the **Open** radio button within the Local Area Network information area. This enables the check box for Partition Communication. Check (activate) **Partition communication**.

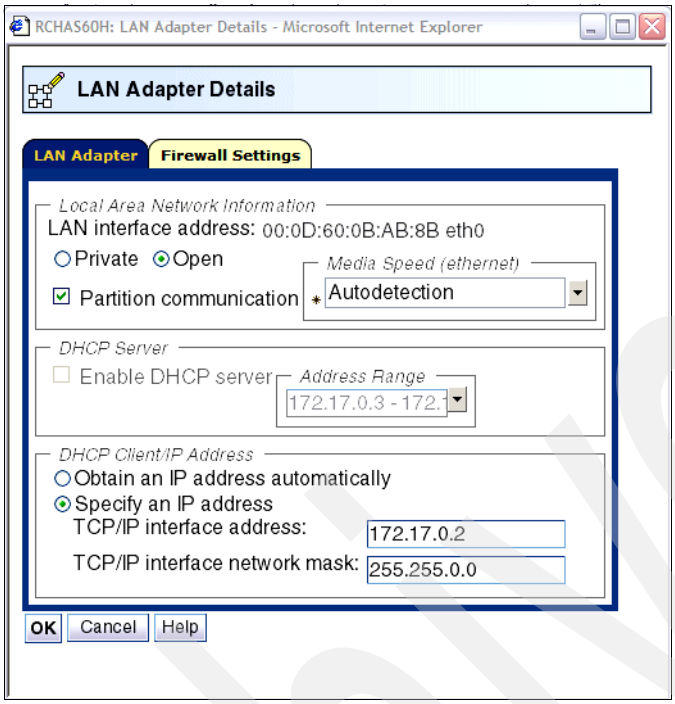

Figure 3-11 Customize Network Setting - LAN Adapters - partition communication

3. Click the **Firewall Settings** tab to bring up the window shown in Figure 3-12. Scroll down the Available Applications area to see if RMC is already specified as available. In our example we assume that RMC has not yet been made available. Therefore, select **RMC** in the lower Allowed Hosts pane and click **Allow Incoming**, as shown in Figure 3-12.

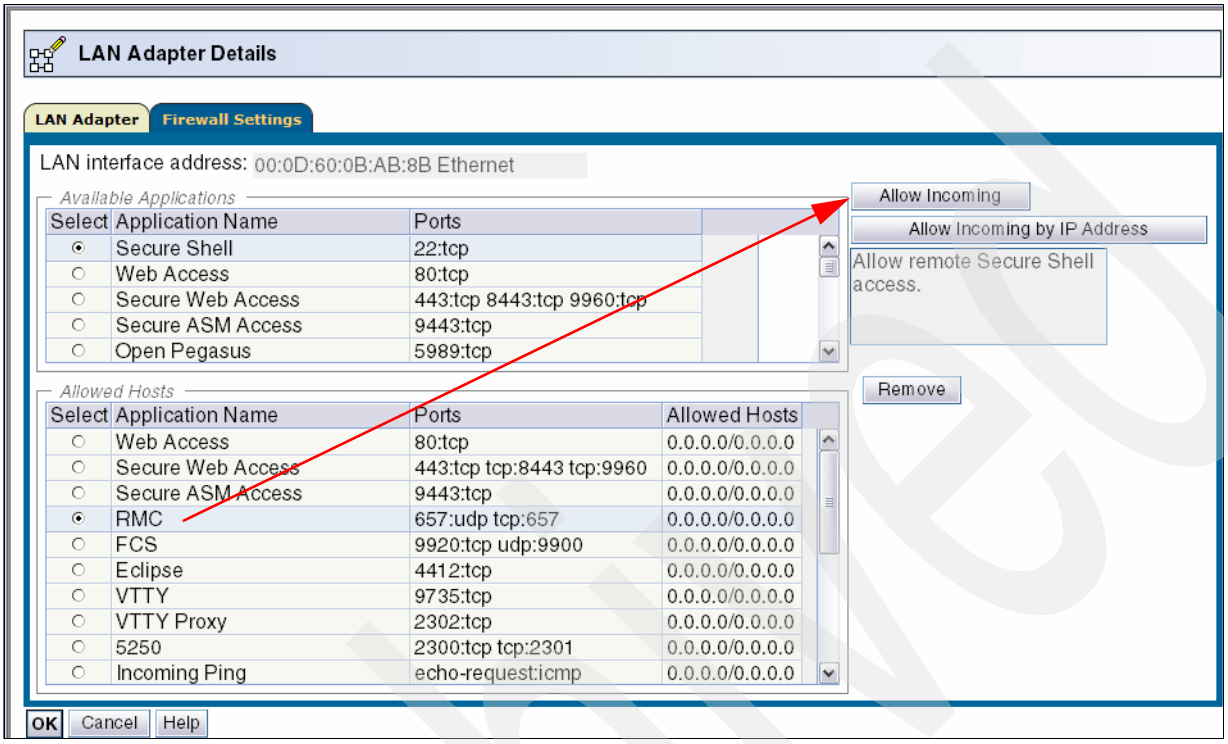

Figure 3-12 Customize Network Setting - LAN Adapters - enabling the RMC application

This moves RMC up into the Available Applications pane. Click **OK** twice. This brings up the window (not shown) stating that Network Settings Changes will be applied at the next HMC reboot.

4. Click **OK**. You are now back to the Hardware Management Console with just the HMC Management pane on the right.

You can verify that you have successfully enabled Resource Management and Control (RMC) by using the **lspartition** command on the HMC command-level interface (CLI). For more information about using the HMC CLI interface, refer to 3.3.3, “System plans management using restricted shell (CLI)” on page 98.

The list partition command is:

```
lspartition -c should show a list of all active partitions.
```

Example:

```
hmc:> lspartition -c 9117_MTM-10FZZD
```

In our example managed system, this results in:

```
<#0> Partition:<4, partn1.business.com, 1.2.3.444>
   Active:<0>, OS<, >
```

If this does not return any partitions, then the system may not be set up for RMC. Depending on whether the system is a System i or System p, the steps for RMC are different.

IBM Systems Hardware Information Center contains additional information about Resource Monitoring and Control. For background information about RMC, refer to the IBM Redbooks publication *A Practical Guide for Resource Monitoring and Control (RMC)*, SG24-6615. The content of this publication is based upon AIX 5L, 5.1.

If the create system plan from the GUI fails, and there is a need to create a system plan, use the underlying mksysplan CLI at the HMC command prompt, with the noprobe option. The noprobe option bypasses the default inventory collection of active partitions. So the resulting sysplan may not have IOA/IOP controlled disk units or media enclosures.

For example:

```
hmc:> mksysplan -m machineName -f filename.sysplan -v -o noprobe
```

Note that when creating a sysplan, if there is a failure due to a Virtual I/O Server error, the above noprobe option can be tried from the CLI.

3.2.4 Viewing a system plan on the HMC

The HMC has a system plan viewer similar to the viewer in the System Planning Tool. The viewer offers a non-editable presentation of the system's partitions and hardware. Using Figure 3-3 on page 69 as a starting point, select the desired plan in the main system plan management window. Click **View System Plan**.

When the HMC is being accessed remotely, you will be presented with a View System Plan sign-on window the first time the System Plan Viewer is launched. This *extra log-in* protects unauthorized users from looking at the system's configuration. It also prevents launching the Viewer from bookmarks without providing an appropriate user name and password.

You will get the login window shown in Figure 3-13. If you get a sign-on window, use a valid user name and password and click **Login**.

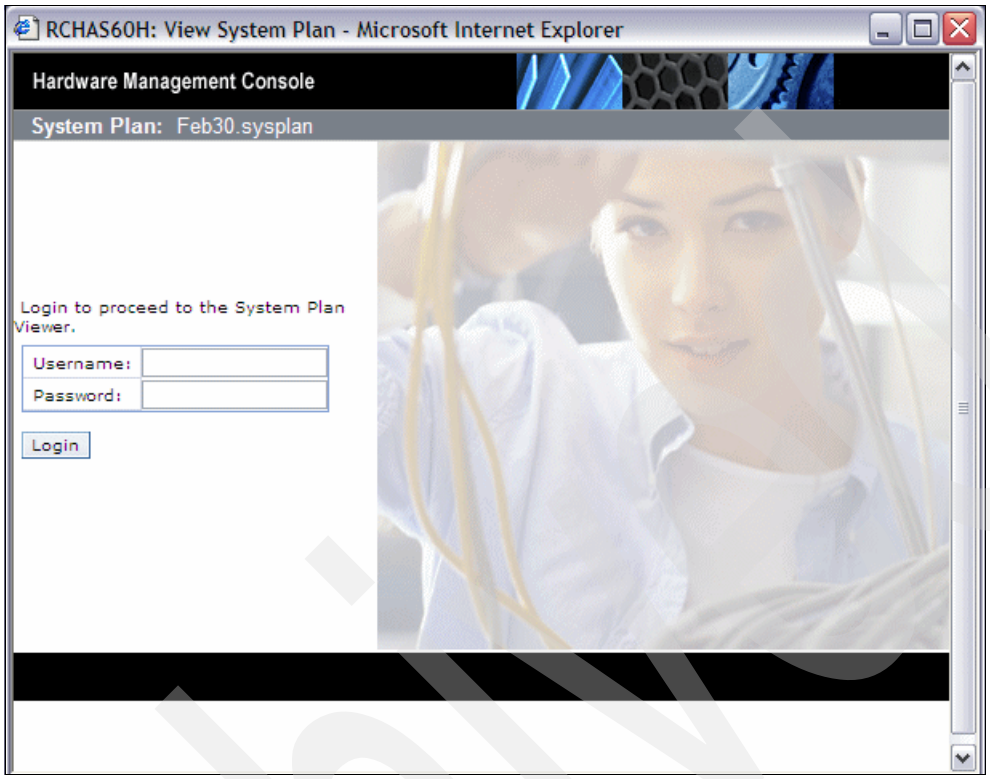

Figure 3-13 View system plan sign-on page

The system plan is shown in Figure 3-14. The left navigation frame allows for viewing of a single partition or the entire system. You can also choose just specific enclosures under the Hardware section. The file history is also viewable. The viewer also has a Print option and Show Comments/Hide Comments toggle, located at the bottom of the viewer window.

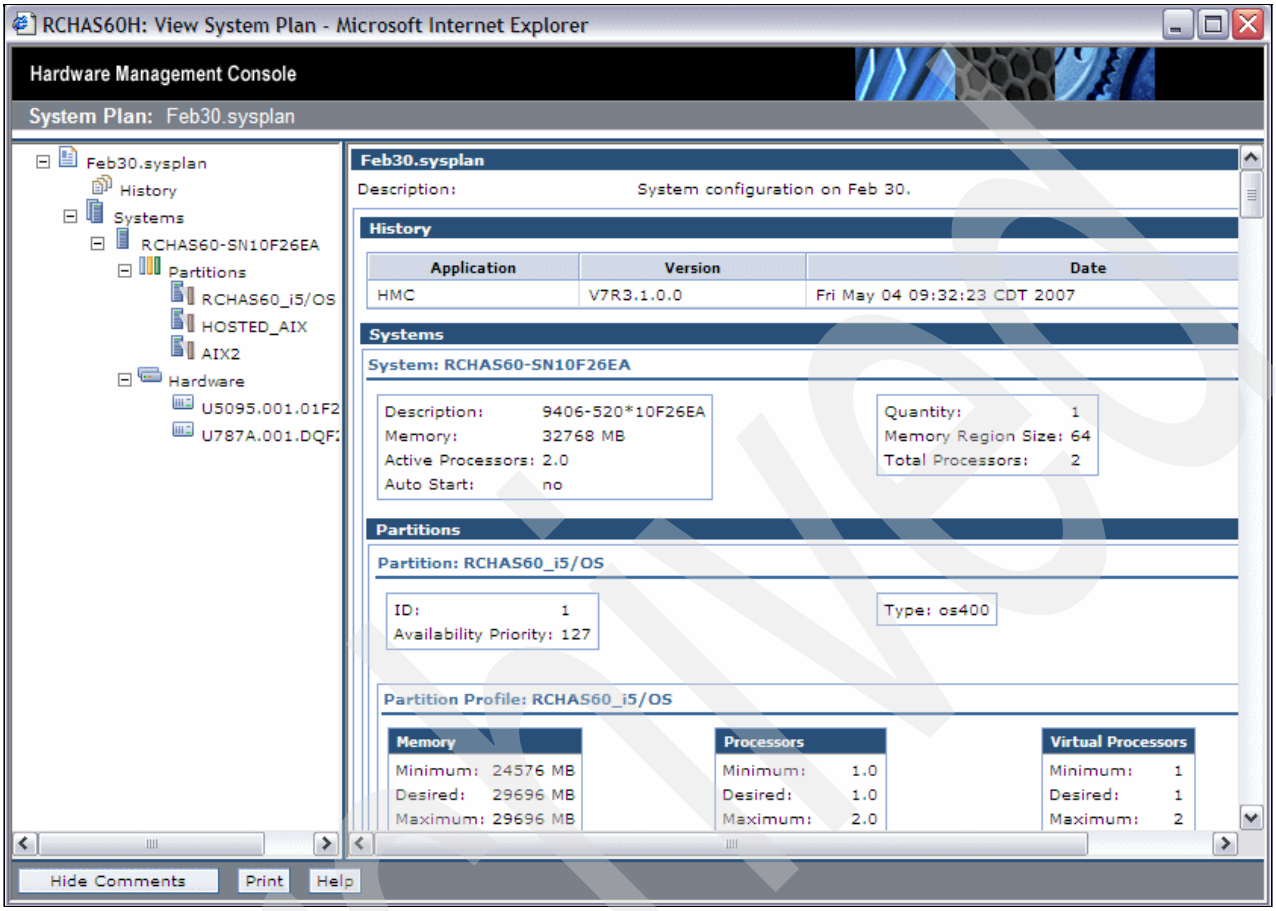

Figure 3-14 Viewing a system plan

If you are accessing the HMC from a PC browser, the print function will be through the PC's attached and network printers. If you are using the HMC terminal itself, the print function will be through printers connected to the HMC or network printers to which the HMC has access.

The system plan section shown in Figure 3-15 shows the system's disk units. Note that the controller for the disk unit is displayed in the table. This detail is only obtained if the i5/OS operating system controlling the disk units is up and running. Linux and AIX operating systems do not display disk controller information or location information.

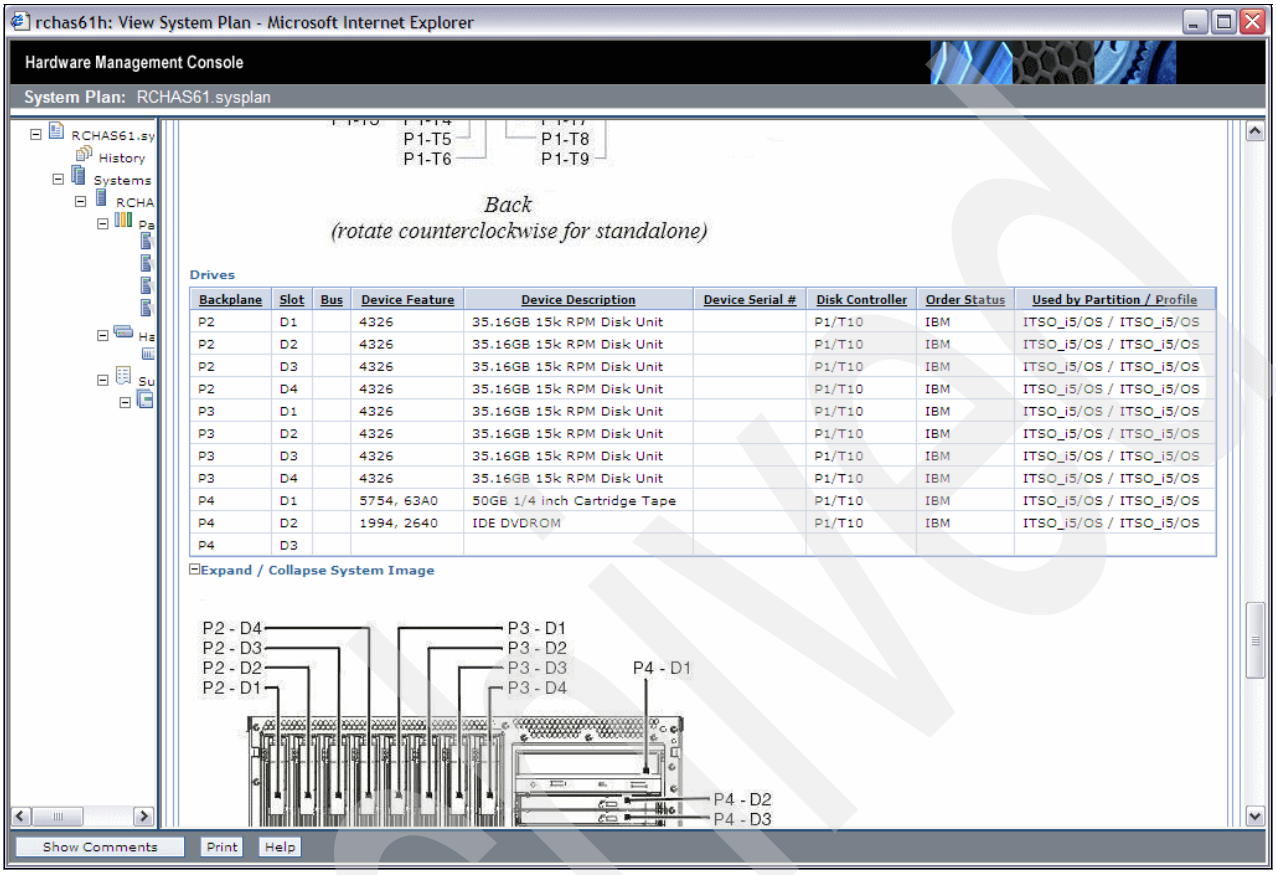

Backplane	Slot	Bus	Device Feature	Device Description	Device Serial #	Disk Controller	Order Status	Used by Partition / Profile
P2	D1		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P2	D2		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P2	D3		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P2	D4		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D1		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D2		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D3		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D4		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P4	D1		5754, 63A0	50GB 1/4 inch Cartridge Tape		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P4	D2		1994, 2640	IDE DVDROM		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P4	D3							

Figure 3-15 Viewing a system plan

3.2.5 Removing system plan on the HMC

When a system plan is no longer needed, the sysplan file can be easily removed from the HMC. Using Figure 3-3 on page 69 as a starting point, select the desired plan in the main system plan management window. Click **Remove System Plan** either from the Tasks drop-down menu or the content Tasks menu in the lower portion of the window. You will see a conformation window asking whether you are sure that you want to delete the file. See Figure 3-16.

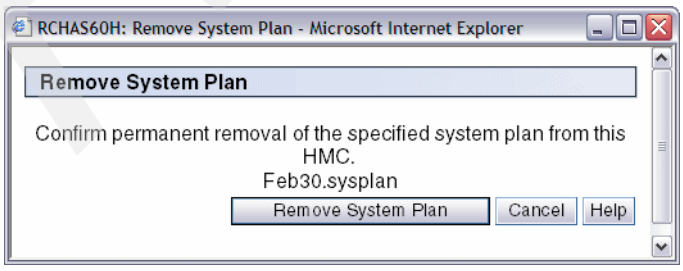

Figure 3-16 Confirm removal of system plan window

Clicking **Remove System Plan** will remove the selected sysplan file from the HMC.

3.3 System plans deployment

Since the *LPAR Simplification Tools Handbook*, SG24-7231, from a general point of view, the deployment process did not change a lot. Of course, due to the updates of System Planning Tool Version 2 and HMC software, the details are not the same. On the other hand, the major improvements of the process are related to Virtual I/O Server implementation.

In this section, we provide a summary of the deployment validation process. We cover the new deployment wizard, by using examples. And last, the updates to the restricted shell command-line interface are detailed.

You will find specific deployment information about Virtual I/O Server in 4.4, “Deployment of Virtual I/O Server on HMC” on page 129.

3.3.1 Deployment validation process

Before deploying any system plan, it must be validated. There are two steps in this validation process. Hardware validation is done first, then, if it is successful, partition validation is done.

This process is very detailed in the *LPAR Simplification Tools Handbook*, SG24-7231, and you can refer to it at any time. But, because it is fundamental to fully understand how it works, we summarize, in this section, its main concepts.

Hardware validation

When running hardware validation, the HMC checks that any planned hardware exists on the managed server and that all of the IO processors and adapters are physically located in the planned slots. It does not necessary mean that we need an exact match between the planned and the existing hardware. For example, we may plan using less processors or memory than physically installed. We may plan not using all of the physically installed IO units.

The validation includes all the following items:

- Server type, model and processor feature: An exact match is required.
- Number of processors: At least the planned number must exist.
- Memory: At least the planned amount must exist.
- Expansion units: All the expansion units in the plan must exist.
- Slots: All the IO processors and adapters in the plan must exist in a correct expansion or in the Central Electronic Complex (CEC) and must be at the same location.
- Any serial number: An exact match is required.

Important: The HMC is not aware of the devices that are connected to the IOA, and therefore there is no validation at a lower level than the IOA. When using the System Plan Tool, you *must* specify devices like disk drives, CD/DVD drives, and tape drives. The validation process *cannot* perform any validation about these devices.

At this point it is important to take actions to *avoid any ambiguity* about the expansion units or the processor enclosures (CECs). You could have multiple CECs, for example, on a 16-Way model 570. In that case you would have four CECs.

This ambiguity takes place when two or more installed expansion units or CECs have the same type and contain exactly the same IO processors and adapters in the same slot. You may plan a partition to use specific expansion units due, for example, to their physical location

in the racks or on the floor, or to specific disks drives that the HMC cannot see. The validation process allows such a system plan but there is no guarantee for the deployment to allocate the right expansion to the partition.

The best way to eliminate expansion units or CECs ambiguity is to specify, in the system plan, their serial number.

You must eliminate any hardware validation error for the partition validation to start.

Partition validation

When running partition validation, the HMC checks that any *existing* partition on the server exactly match with one of the planned partition.

The validation includes all the following items:

- Partition name
- Partition ID
- Name of the default profile
- Processing resources in the system plan
- Memory resources in the system plan
- Physical hardware in the system plan
- Virtual adapters, including slot IDs and maximum adapters, in the system plan

If any of the above items fail, the partition validation is unsuccessful and the deploy will fail. Some of the corrections to allow the deployment must be applied on the server. This is the case for the name of the default profile, which cannot be changed in the System Planning Tool and is the same as the partition name. This is also the case for some hardware features like the USB controller or the IDE CD controller that the HMC allows you to assign to an i5/OS partition (while it cannot use them), but the SPT does not.

3.3.2 Deploy a system plan using the graphical wizard

In this section, to show the new deployment wizard related to V7R3 HMC software, we run three deployment examples:

- The first fails due to hardware errors.
- The second fails due to partitions errors.
- The third is successful.

Deployment can be initiated when you are using any right pane of the HMC by clicking **System plans** on the left pane, as shown in Figure 3-17.

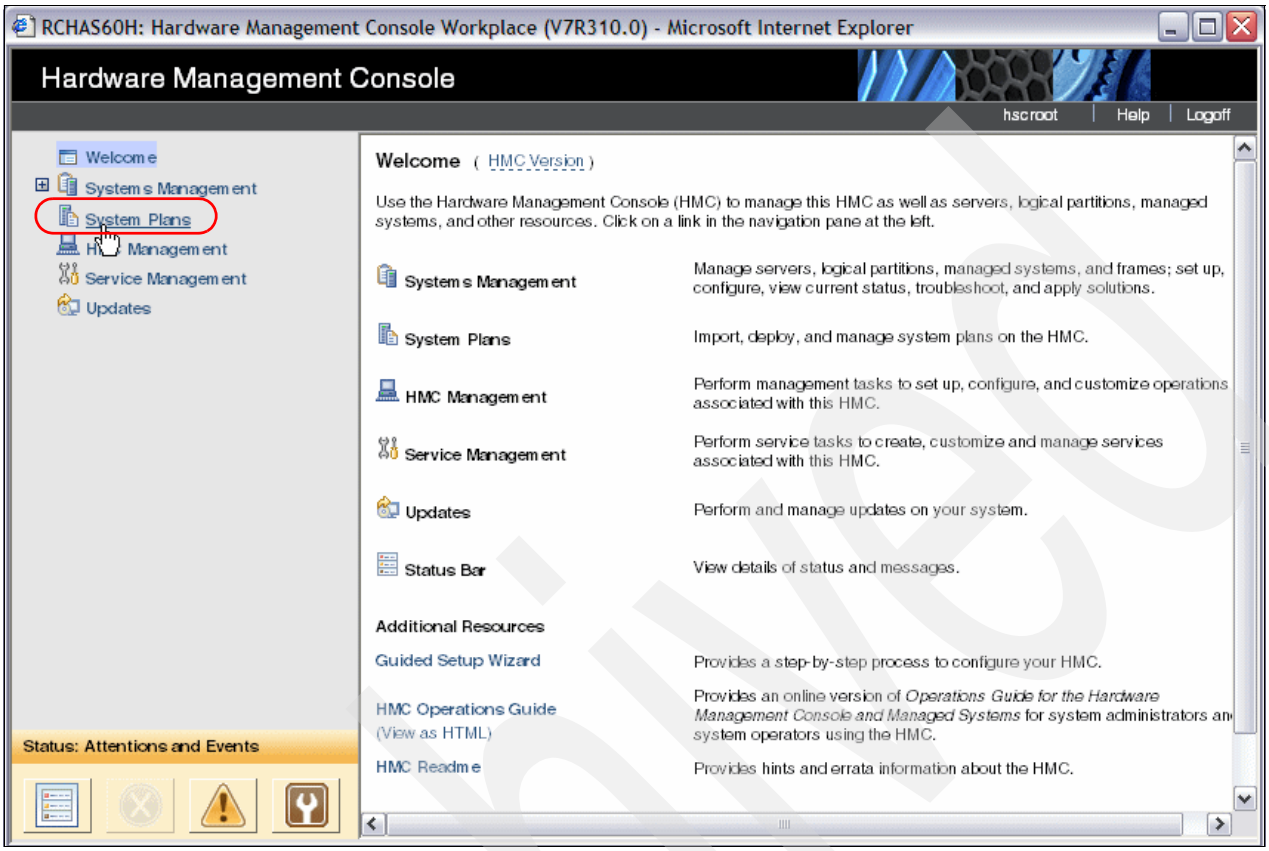

Figure 3-17 Launch deployment

The steps are:

1. On the list of the system plans, you first select the one that you want to deploy by clicking the check box to the left of the system plan, as shown in Figure 3-18. In our example, we select the marc.sysplan file.

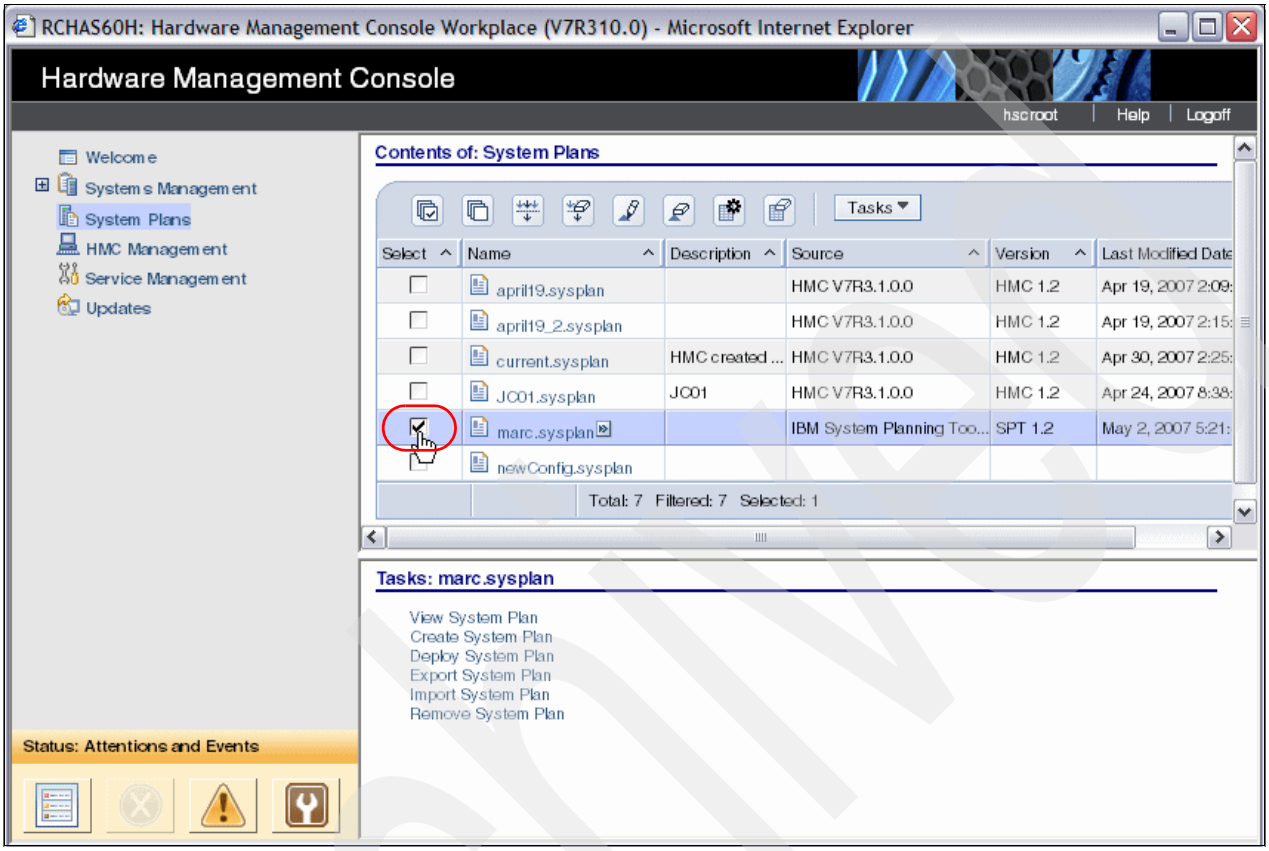

Figure 3-18 Select the system plan to deploy

There are three different ways to start the deployment of the selected system plan, as shown in Figure 3-19.

a. Click the contextual pop-up menu immediately to the right of the system plan name and select **Deploy System Plan**.

b. Click **Deploy System Plan** in the bottom Tasks panel.

c. Click **Tasks** at the top of the System Plans list panel and select **Deploy System Plan**.

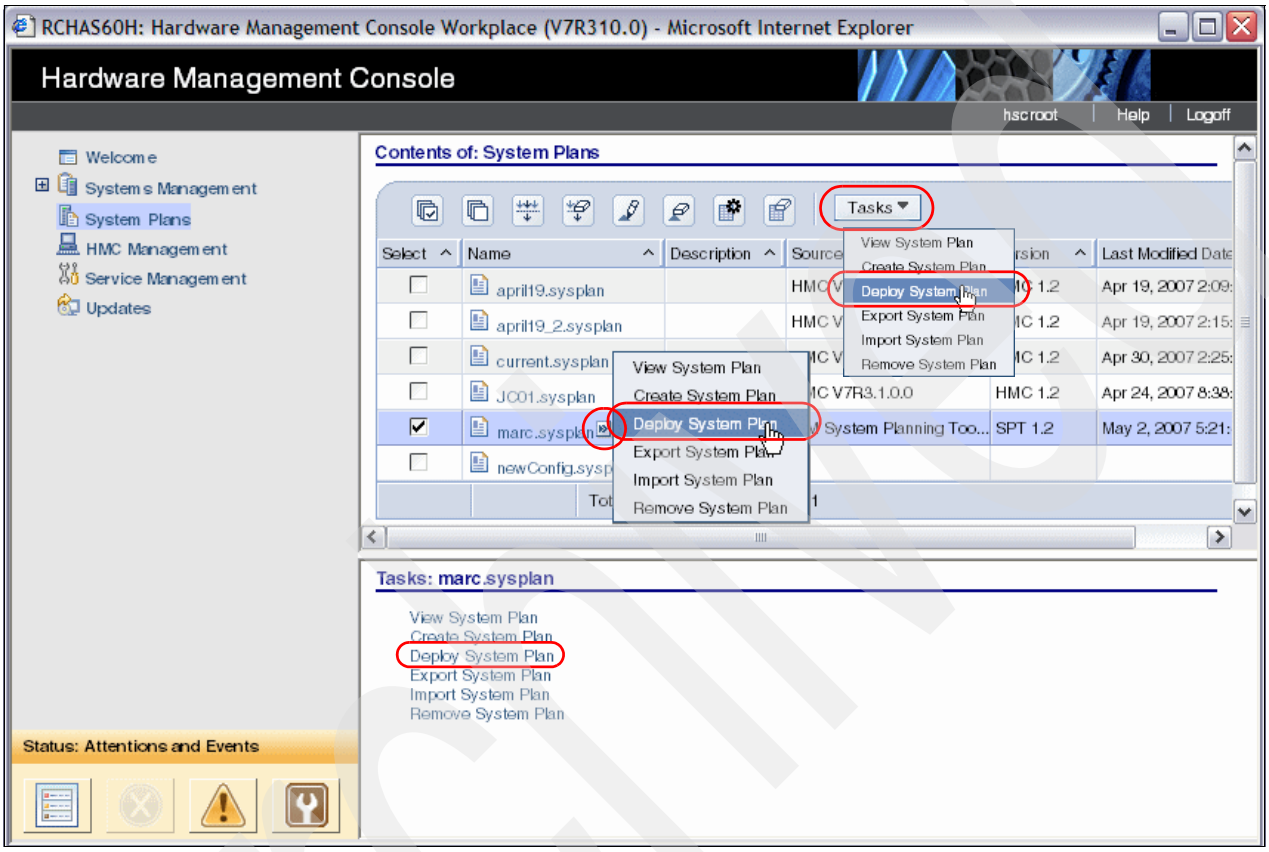

Figure 3-19 Launch the deployment

2. Once the wizard is started, its welcome page, as shown in Figure 3-20, requests you to confirm the system plan to deploy and choose the managed server to be the target of the procedure. When your choices are done (in our example, we want to deploy the marc.sysplan file to the RCHAS60-SN10F26EA server), click **Next** to continue.

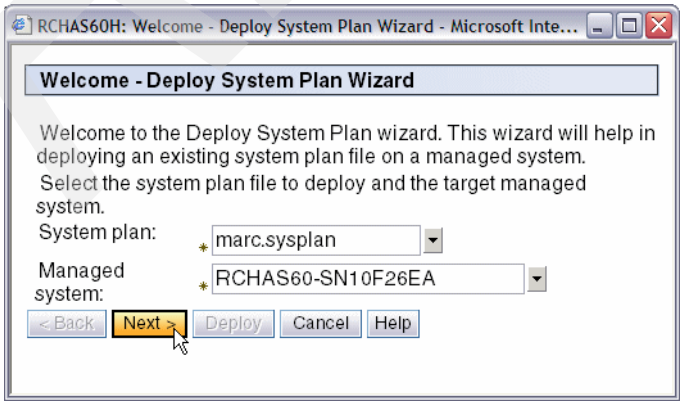

Figure 3-20 Confirm the deployment startup

3. The next display (Figure 3-21) shows you the validation progress. When this has completed, you can examine all the related messages — those that are successful validation ones as well as those that are unsuccessful.

   The validation types are detailed in 3.3.1, “Deployment validation process” on page 81.

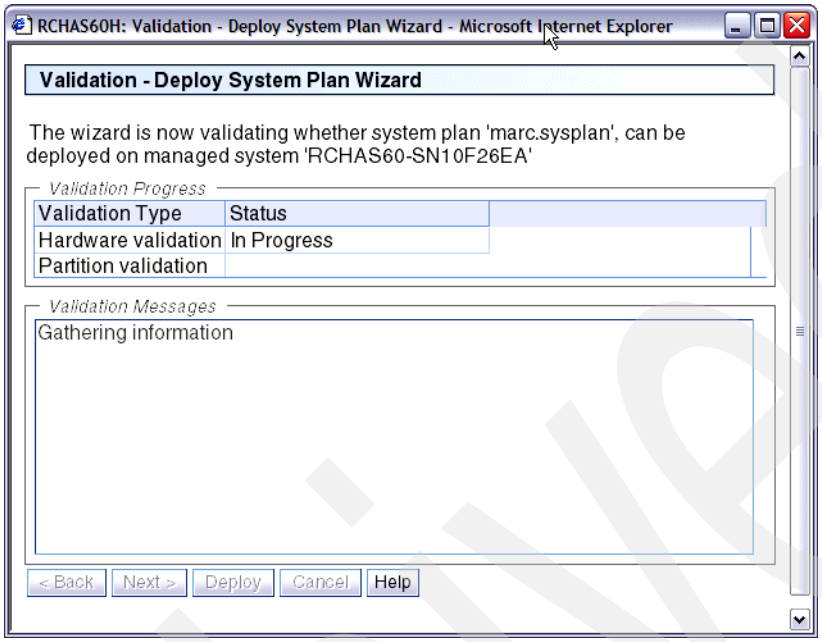

Figure 3-21 Deployment validation in progress

We show now three examples of validation results:

– The first example is related to a system plan that fails to deploy due to hardware errors. See our example in Figure 3-22, which shows unsuccessful hardware validation.

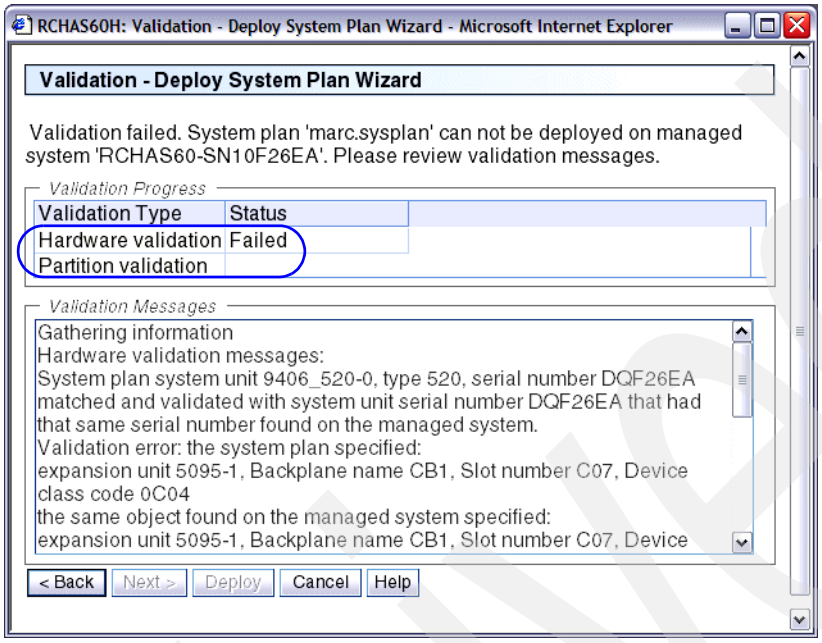

Figure 3-22 Deployment validation summary (hardware failed)

In the Validation Messages panel, you see the result of each particular validation step. You may need to scroll down/up the list to find the root cause of the failure and, in most cases, you can easily correct the system plan to get a successful validation. Here (see Figure 3-23 on page 88), first hardware validation error message says that, in the expansion unit 5095, the slot C7 is occupied by a device that is not the same that the one specified in the system plan. To correct this specific error, the best way is to update the system plan with SPT to match the hardware. Another option could be to physically install the correct device according to the system plan, and adjust all the hardware to the system plan. This may lead you to order such a device and move I/O cards from a slot to another.

After reading all the necessary messages, click **Cancel** to exit the window.

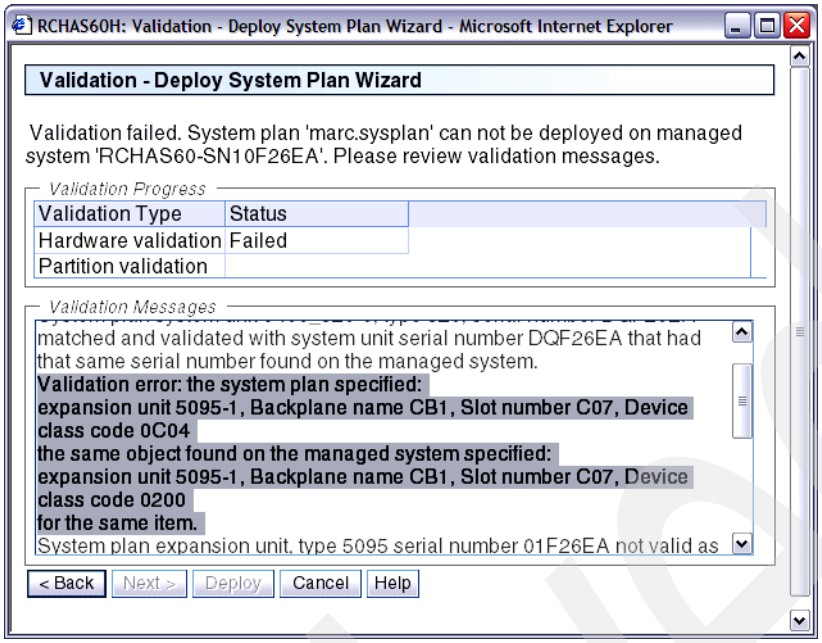

Figure 3-23 Example of hardware validation error

Note: The partition validation does not begin while there are errors on the hardware validation type.

– The second example is related to a system plan that fails to deploy due to partition errors. See our example in Figure 3-24, which shows successful validation for the hardware part, and failed validation for the partition validation type.

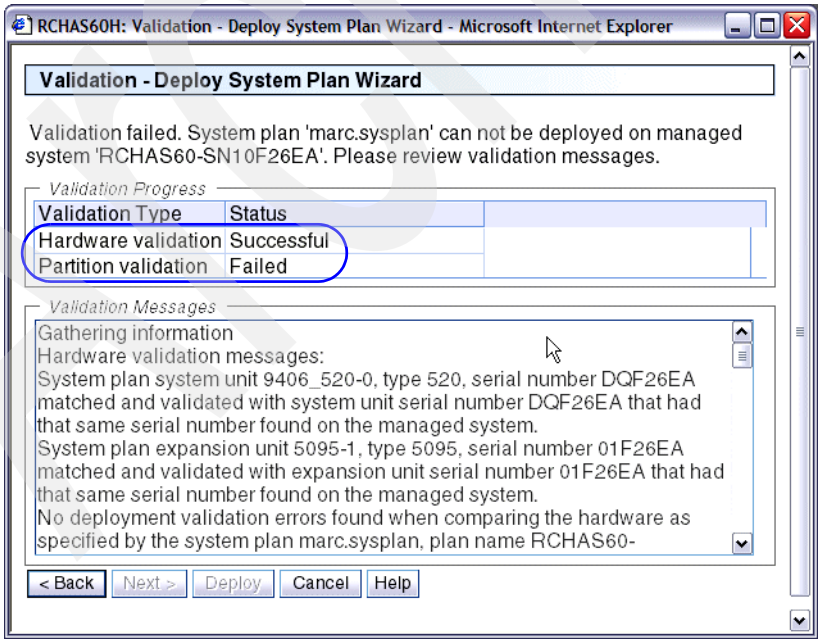

Figure 3-24 Deployment validation summary (partition failed)

In the Validation Messages panel, you see the result of each particular validation step. You may need to scroll down/up the list to find the root cause of the failure and, in most

cases, you can easily correct the system plan to get a successful validation. Here (see Figure 3-25) the first partition validation error message says that the partition 3 name of the system plan does not match with existing partition 3 name. To fix this problem, you can either rename the partition in the system plan with SPT, or rename the partition with the HMC.

After reading all the necessary messages, click **Cancel** to exit the window.

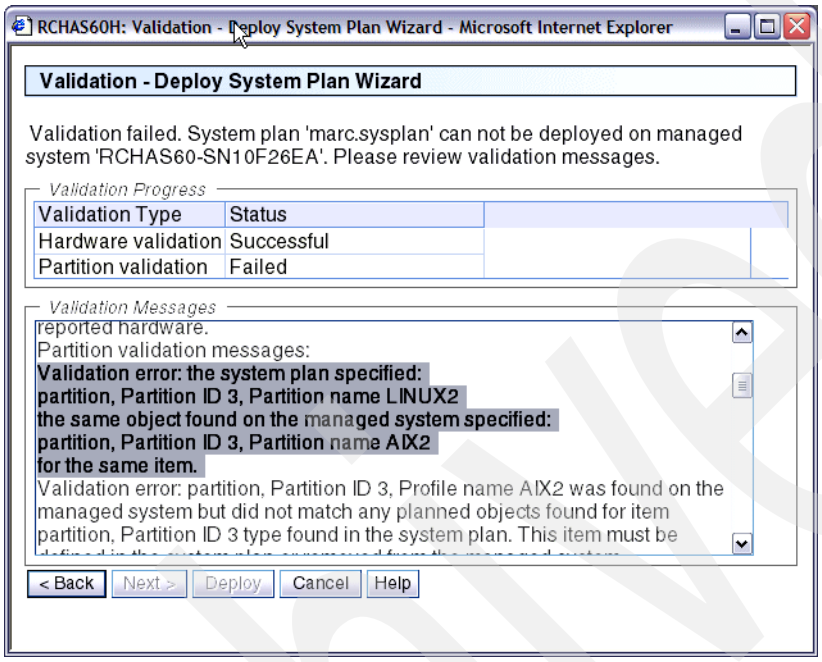

Figure 3-25 Example of partition validation error

d. The third example is related to a system plan that gets successful validation.

See our example in Figure 3-26, which shows successful validation for both hardware and partition steps. You can review all the messages in the Validation Messages panel, then click **Next** to continue.

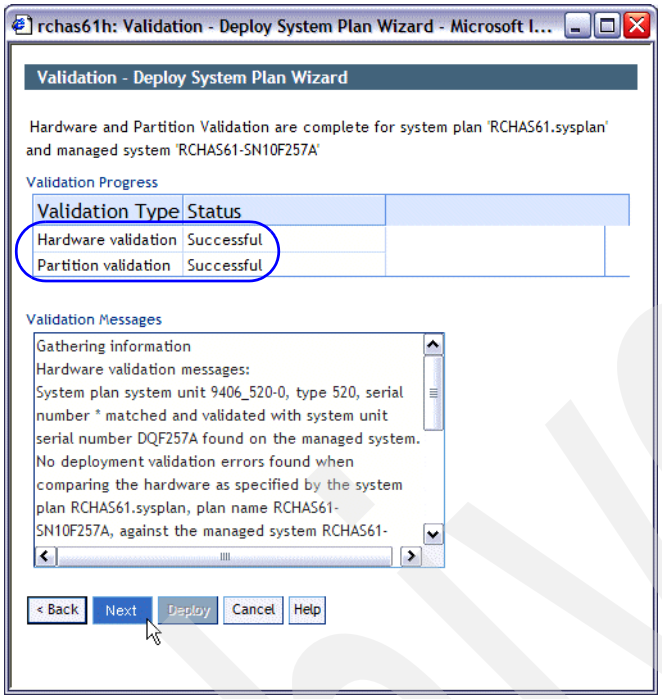

Figure 3-26 Example of successful validation

4. After the validation, the next panel is the starting point of the deployment. There are two portions in this panel.
   a. The first contains the list of all the actions that are planned (see Figure 3-27).

      In our example, the system plan is built with one i5/OS partition, which hosts three Linux partitions.

      Notice the partially deployed status of the i5/OS partition. It means that some items of the partition exist on the server and do not need to be deployed. Specifically, here, the partition and the profile are created (and the LPAR is running at the time of the window capture). These items will not be deployed again.

      Notice also the Deploy column. Each action of the plan may be deselected if you prefer to run it at a later deployment. Notice that, even if there are deselected items, the dependency is checked before running the deployment. So, you cannot, for example, deploy an hosted partition while the hosting one does not exist.

      If you need to review the details of a specific action, you can select this one in the radio boxes of the Select column and click **Details**.

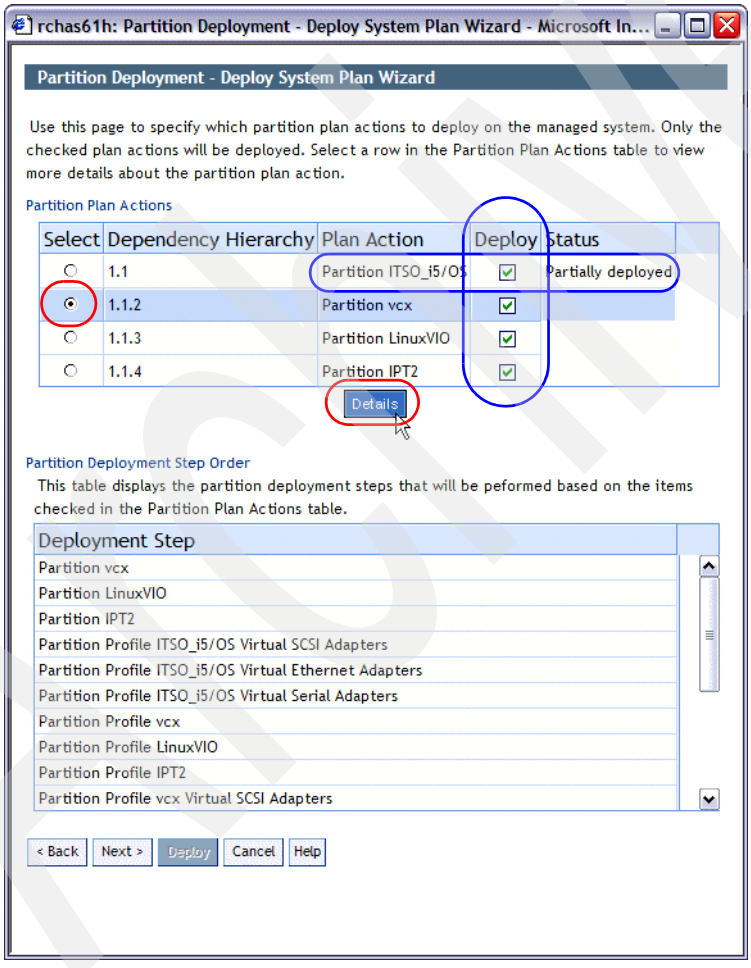

Figure 3-27 Request the details of a specific action

Click the **Details** button. The HMC links to the System Plan Viewer, restricted to the view associated with the selected action (1.1.2, in our example). The System Plan Viewer *report* for the action is displayed in a separate window, as shown for our example, in Figure 3-28. In our example, we show the planned profile for the vcx partition. Once the review is finished, just close this window.

Note: You may need to authenticate again to the HMC when accessing this option.

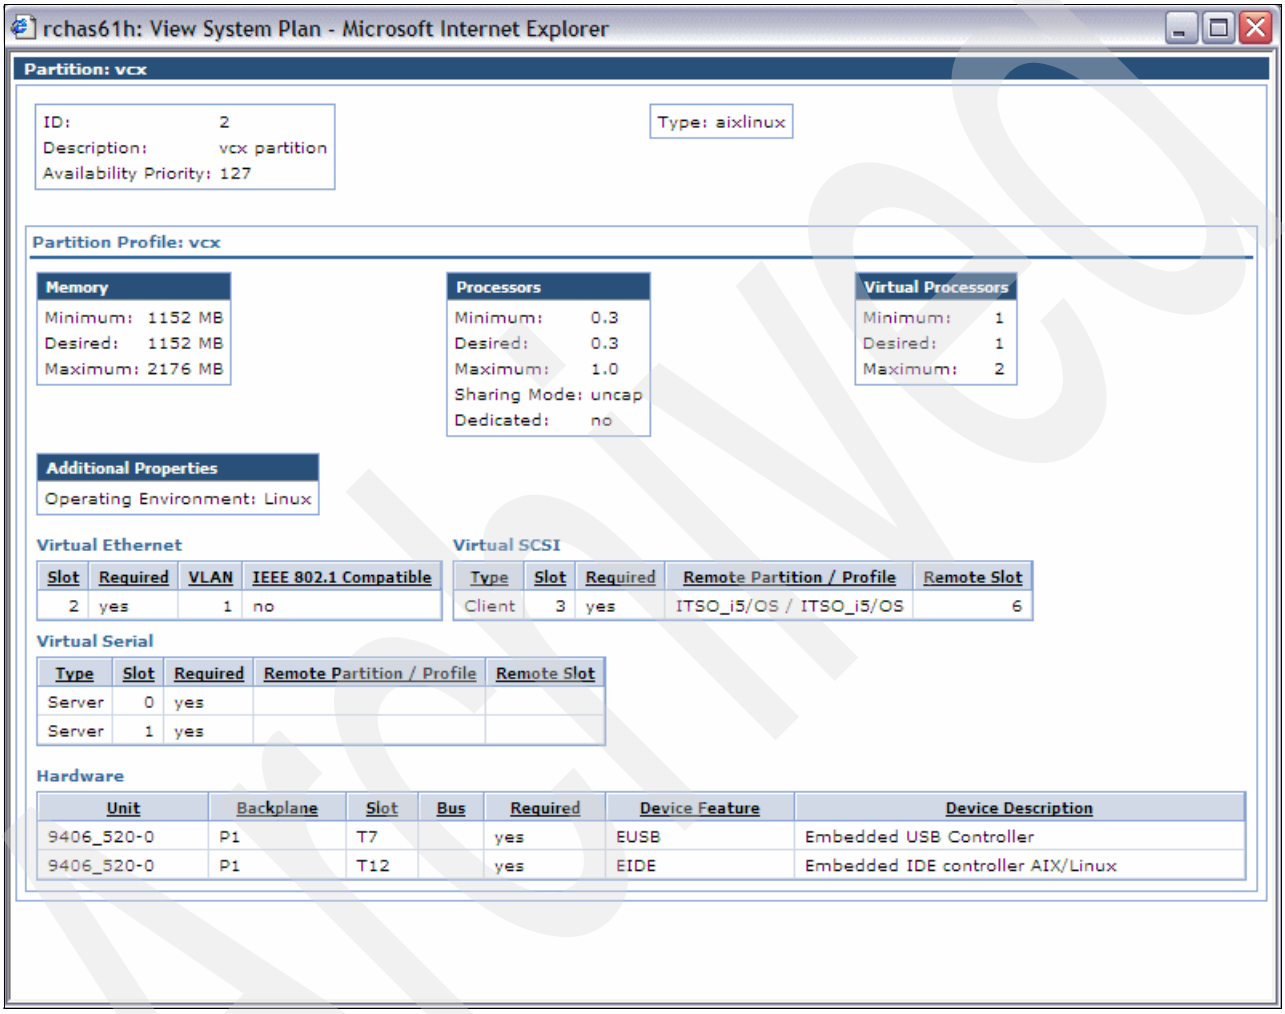

Figure 3-28 Details of an action, shown by a restricted view System Plan Viewer

b. The second portion contains the detailed list of all the steps the HMC will perform to deploy the plan, as shown in Figure 3-29. You can scroll down and up to review all the steps.

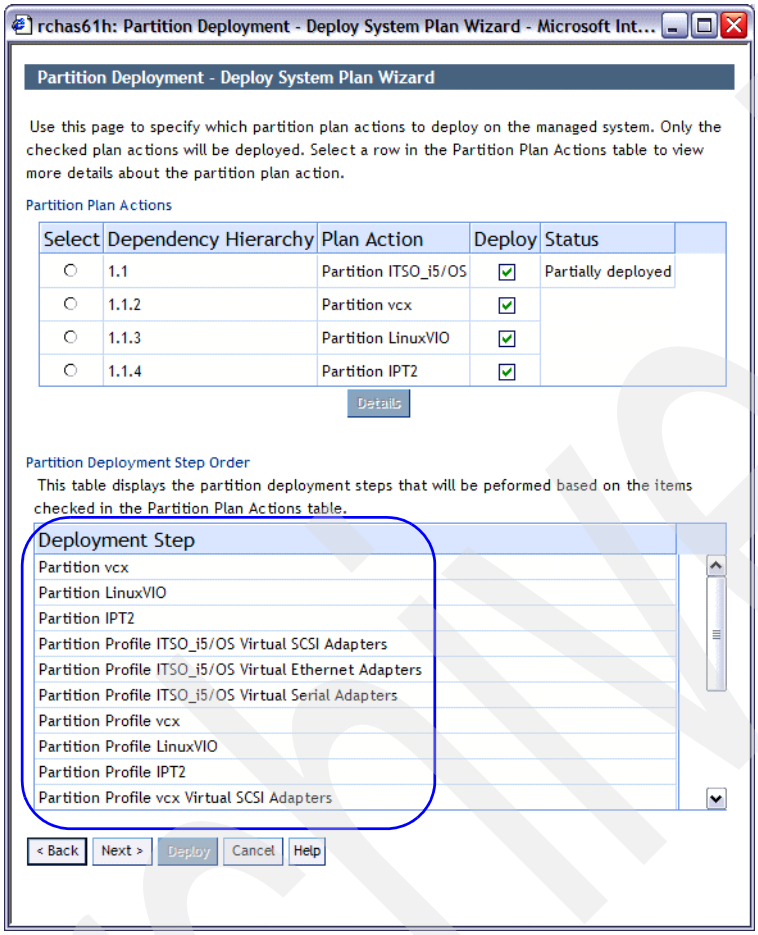

Figure 3-29 List of the deployment steps

Once you are ready to deploy, after reviewing details of the actions, eventually deleting some of the actions, and reviewing the deployment steps, as shown in the Figure 3-30, click **Deploy** to start the deployment.

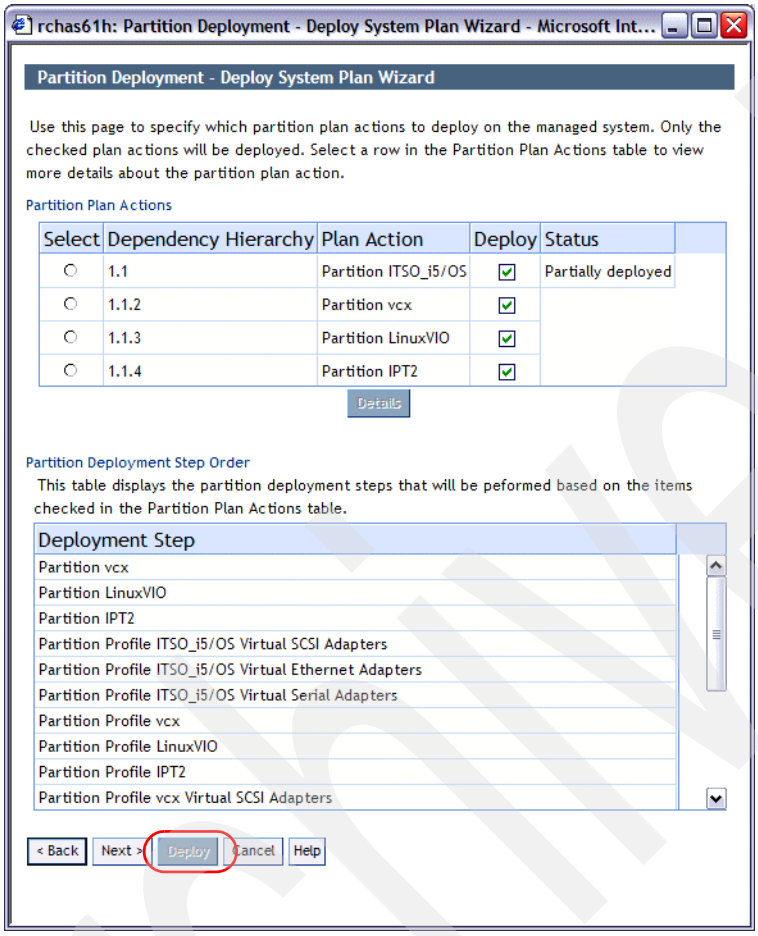

Figure 3-30 Partition Deployment - Deploy System Plan Wizard

5. The next panel, Figure 3-31, is the last panel before running the deployment. There is a summary of all the steps that the HMC will perform. Click **Deploy** to start the operations.

   Notice the warning just above the buttons at the bottom of the window.

   The deploy process time, depending on the complexity and the number of partitions, generally ranges from 5 to 20 minutes, and may be longer for specific deployments (when using Virtual I/O Server partitions, for example).

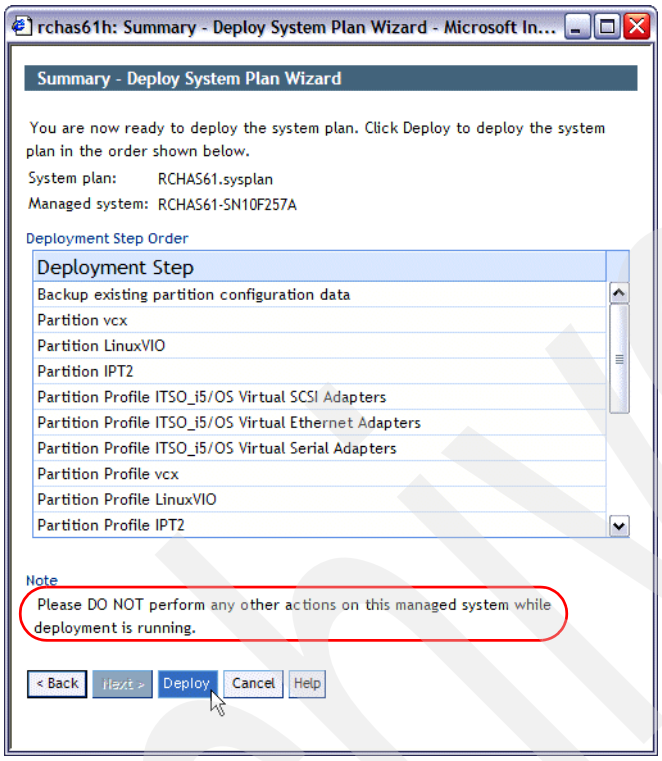

Figure 3-31 Ready to deploy

While the deployment is running, as shown on the Figure 3-32 and Figure 3-33 on page 97, you can follow its progress.

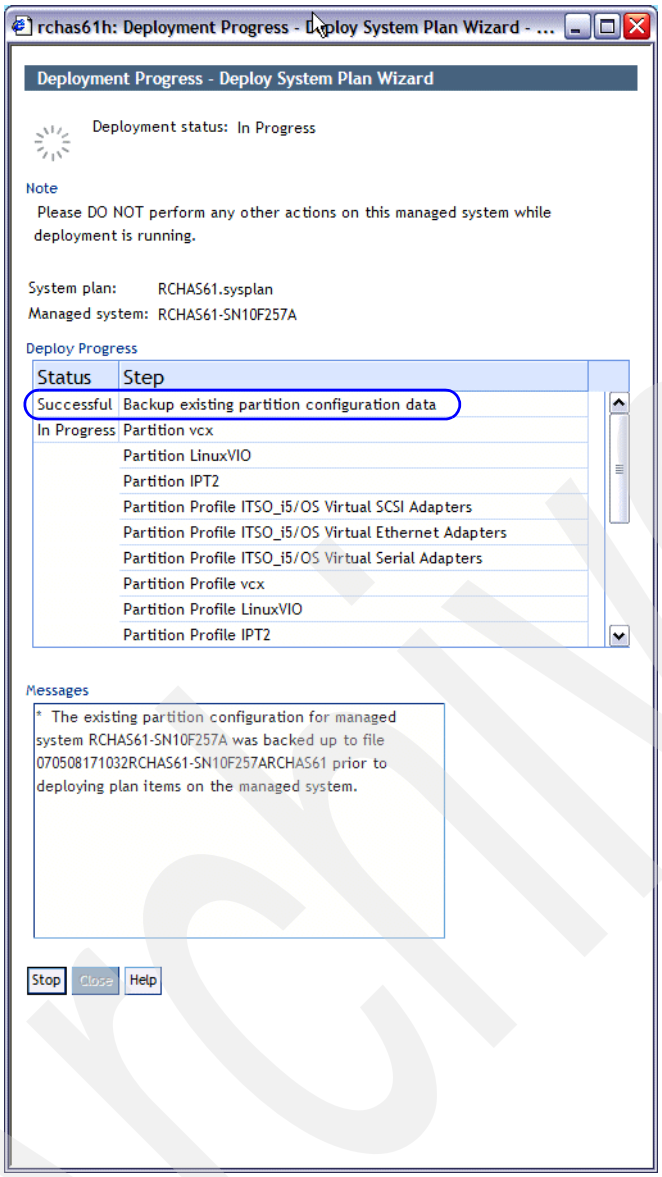

Figure 3-32 Deployment starting

Each step is `In progress` then `Successful`, from the top to the bottom of the Deploy progress portion of the panel. Notice that you have to move the cursor of this list by yourself to see those steps that do not fit in the initial window.

Note: The first step, which you cannot disable, is always to perform a backup of the actual partitions configuration.

In the Messages portion of the panel, we see the detailed results of each step. In our example, when running the step “Partition Profile ITSO_i5/OS Virtual Serial Adapters“, you can see that there is no result yet. But, when running the previous step, we can see that one of the results was “`Virtual Ethernet Adapter deployed for slot 5 on profile ITSO_i5/OS of partition 1 on managed system RCHAS61-SN10F257A.`“. Notice that the

display sort, for this list, is the opposite of the steps. You see the latest event first and the first one last, and therefore, there is no need to scroll down/up this list during the deployment.

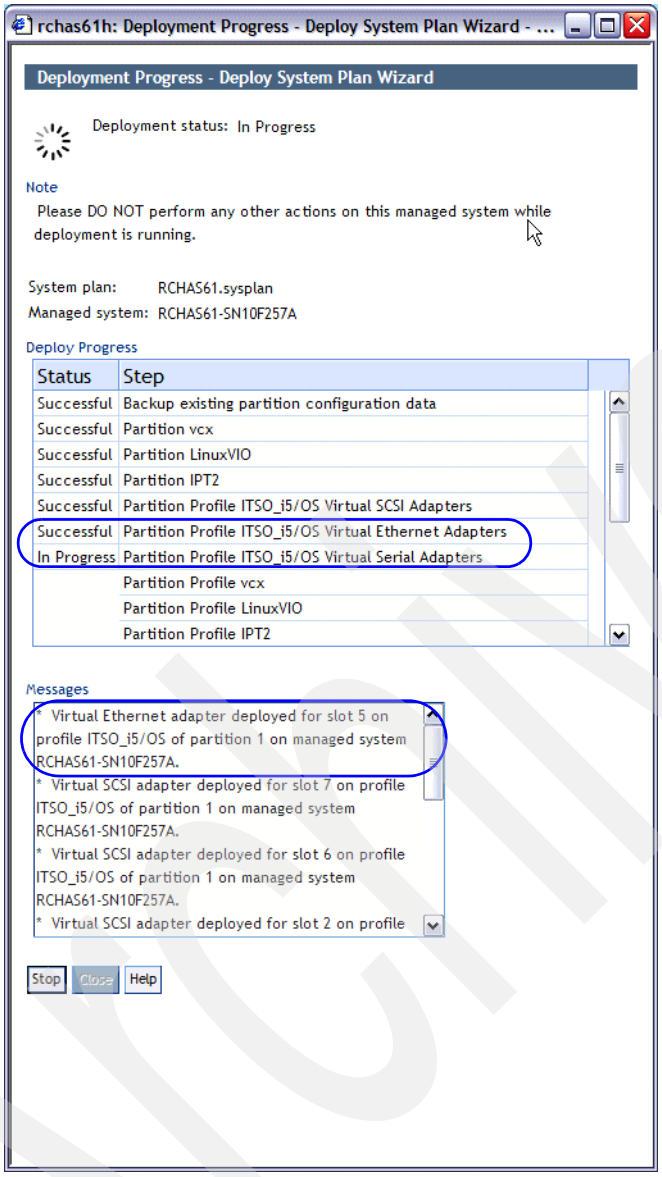

Figure 3-33 Deployment in progress

Once the deployment has completed, as shown in Figure 3-34, you may review all the steps and messages to make sure that everything is okay. Just click **Close** to finish the session.

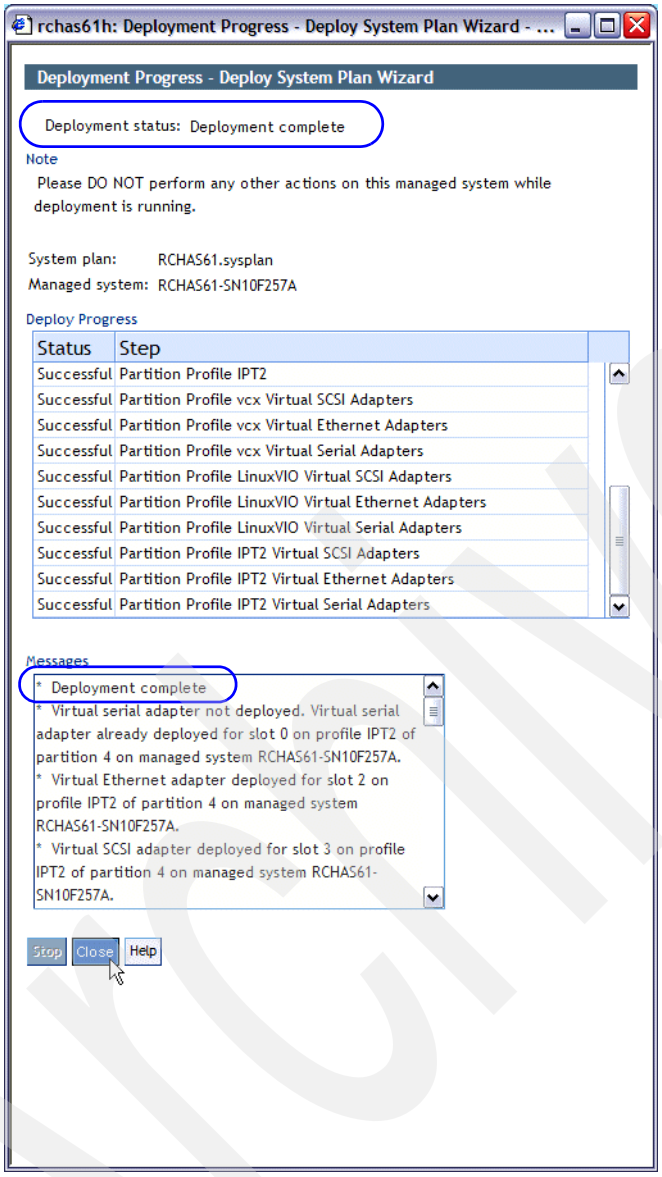

Figure 3-34 Deployment complete

3.3.3 System plans management using restricted shell (CLI)

This section focuses on the HMC command-level interface (CLI) commands designed for managing tasks relating to system plans.

Four of them were introduced in HMC V5R2 code level and remain unchanged. Since System Planning Tool Version 1, their function and syntax did not change. We do not document existing parameters in HMC V5R2 code level.

One of them was introduced in HMC V5R2 and is updated in the HMC V7R3 code level. It has new parameters. We document only these new parameters.

To get more information about the unchanged commands and parameters, or about the way to use the restricted shell, you may refer to *LPAR Simplification Tools Handbook*, SG24-7231.

Unchanged commands

The unchanged commands are the basic ones. They are:

- **lssysplan** is used to display a list of the system plans.
- **deploysysplan** is used to deploy a system plan on a managed server.
- **rmsysplan** is used to delete a particular system plan.
- **cpsysplan** is used to export (copy from the HMC) or import (copy into the HMC) a system plan; because this command is unchanged, unlike the GUI, there is no way to specify exporting/importing to/from the current PC.

Updated commands

The **mksysplan** command has two new optional parameters. This command is used to create a system plan that contains the actual LPAR configuration of a specific managed server.

1. -o is used to specify an option for inventory collection.

   By using this parameter with the noprobe value (the only allowed one), you can request not to scan the devices that are attached to IOA cards. This way, you do not retrieve in your system plan any disk drives attached to the disks controllers, any tape drives that are unknown by the HMC, but are only known by the operating system of the partitions. Using this parameter allows the command to run much more faster, if you do not need the devices' level details.

   If you do not specify the -o parameter, the HMC requests the operating system of each active partition to return the devices information for each IOP.

   Important: To be able to receive device information from the operating system, the corresponding partition *must be running*.

   With or without using the option, the display output is the same as you can see in Example 3-1. The first command does not specify the -o parameter, and therefore gather all the devices information. The second one specifies the -o parameter, and therefore does not request the operating system of the partitions to provide devices information.

Example 3-1 mksysplan command with or without -o parameter

```
hscroot@RCHAS60H:~> mksysplan -f marc.sysplan -m 9406-520*10F26EA
Started inventory gather process ...
System plan marc.sysplan created successfully for the system 9406-520*10F26EA.

hscroot@RCHAS60H:~> mksysplan -f marc.sysplan -m 9406-520*10F26EA -o noprobe
Started inventory gather process ...
System plan marc.sysplan created successfully for the system 9406-520*10F26EA.
hscroot@RCHAS60H:~>
```

   Both commands create the same system plan whose name is marc.sysplan, but their contents are different.

Without specifying -o noprobe, like the first one, you get the device information from the running partitions, as you can see in Figure 3-35. You can see the disk drives details.

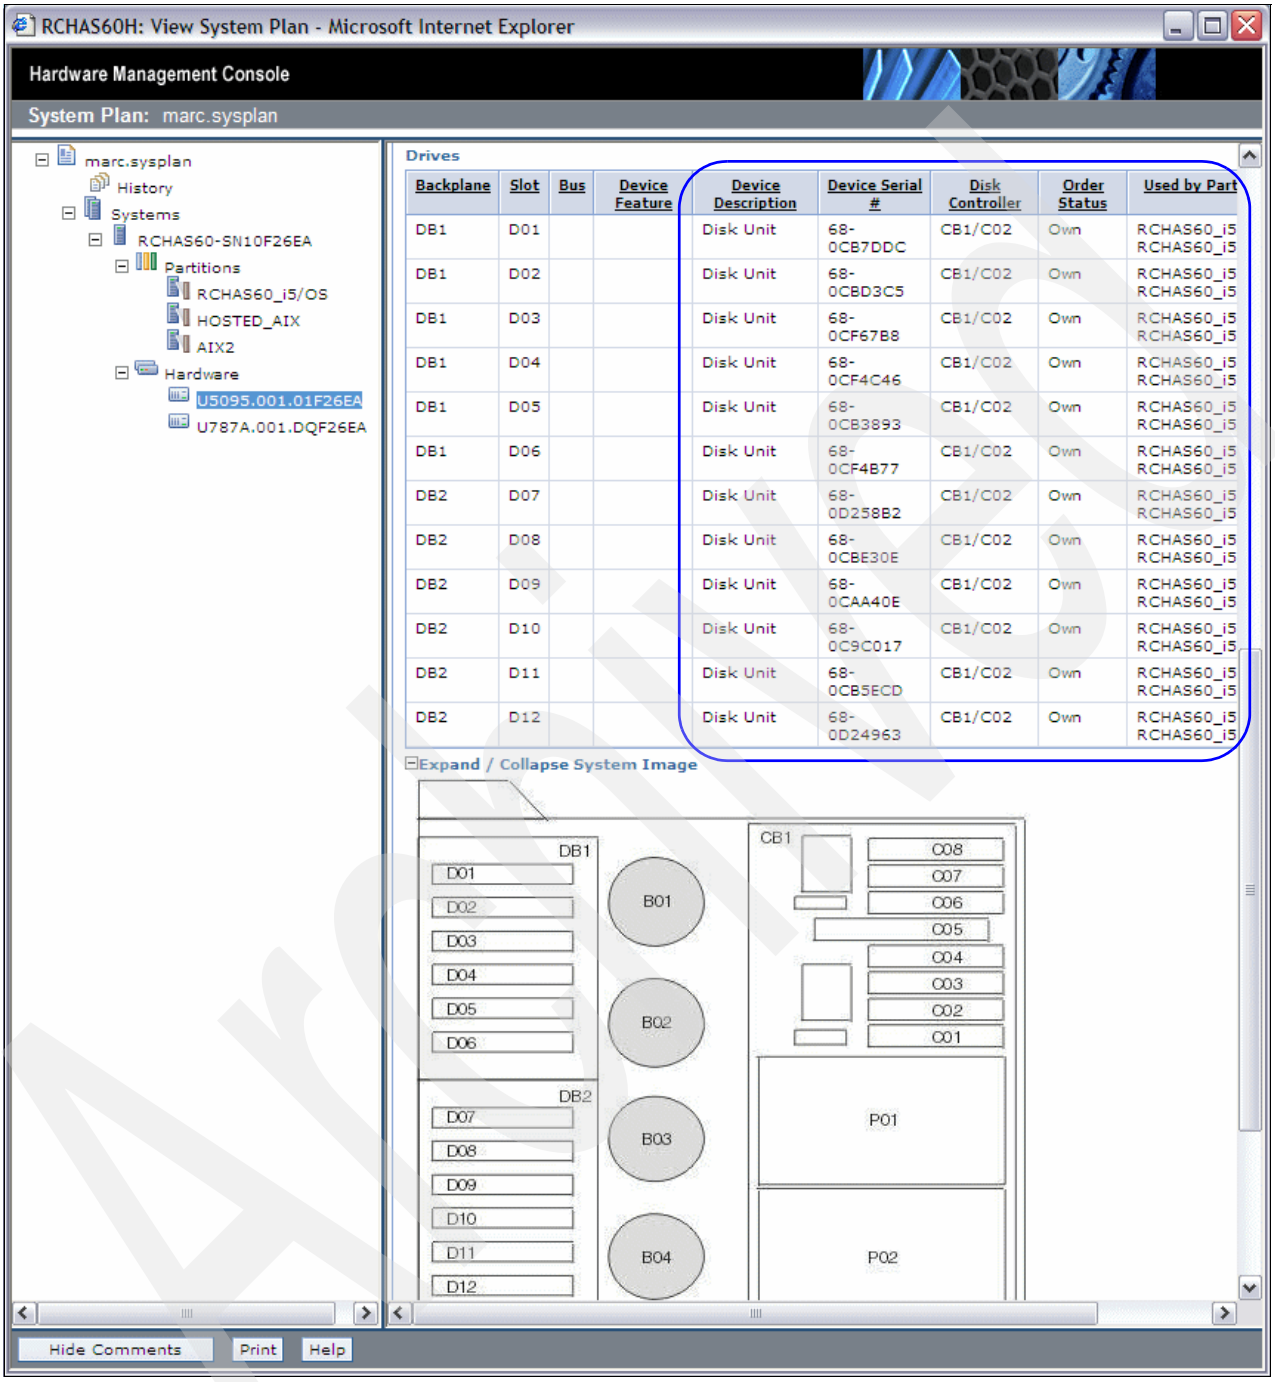

Figure 3-35 Disk information is gathered when not using -o noprobe on mksysplan

When specifying -o noprobe, like the second one, the same query does not retrieve those devices' information, as shown in Figure 3-36.

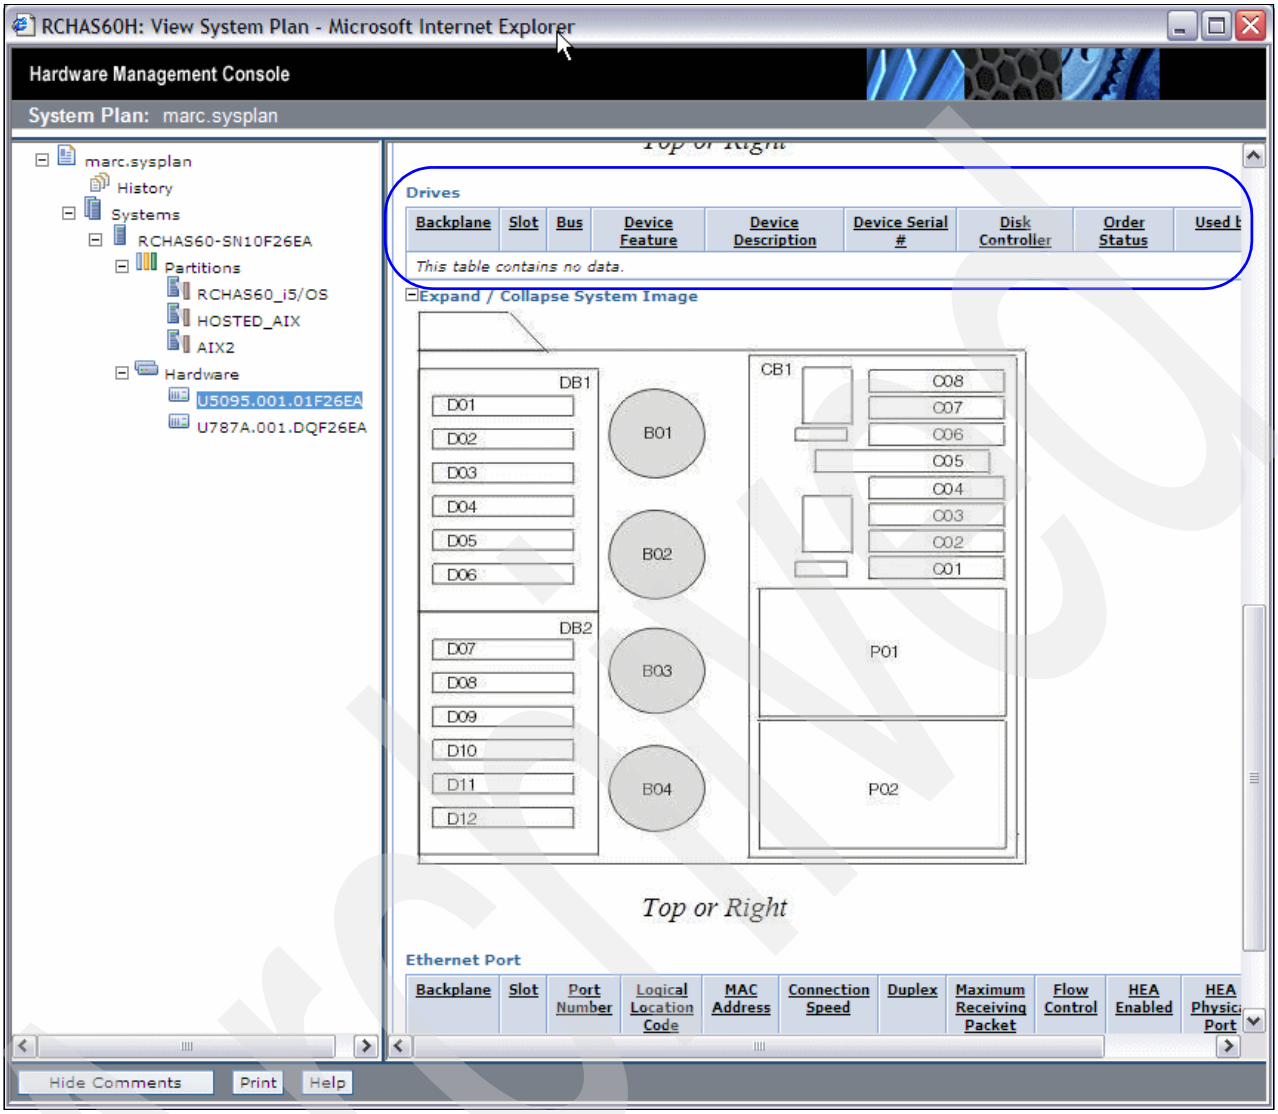

Figure 3-36 Disk information is not gathered when using -o noprobe on mksysplan

2. -v is used to specify verbose display output.

   This parameter does not have any value. You specify it or not. When specified, the display output contains more information about the steps that the HMC performs. As you can see in the Example 3-2, depending on the -o parameter usage, the display output is not exactly the same. Both commands specify the -v parameter. The first gathers the devices information, while the second does not.

Example 3-2 mksysplan with -v parameter and with or without -o parameter

```
hscroot@RCHAS60H:~> mksysplan -f marc.sysplan -m 9406-520*10F26EA -v
Started inventory gather process ...
Gathering slot level hardware and logical definitions ...
Creating new VPD files for the system ...
Adding inventory sensed from each active partition....If you dont need a system
plan to include realtime inventory, retry using the -o noprobe option
```

```
Gathering VIOS identified hardware and logical definitions ...
Completed inventory gathering without errors ...
Writing system plan ...
System plan marc.sysplan created successfully for the system 9406-520*10F26EA.

hscroot@RCHAS60H:~> mksysplan -f marc.sysplan -m 9406-520*10F26EA -v -o noprobe
Started inventory gather process ...
Gathering slot level hardware and logical definitions ...
Gathering VIOS identified hardware and logical definitions ...
Completed inventory gathering without errors ...
Writing system plan ...
System plan marc.sysplan created successfully for the system 9406-520*10F26EA.
hscroot@RCHAS60H:~>
```

Note: The equivalent GUI action that creates system plans makes use of the default of the **mksysplan** command. Therefore, it will always gather all the device information without any verbose output.

4

Virtual I/O Server

This chapter discusses the functions of the Virtual I/O Server on the System p and System i relative to defining and deploying a Virtual I/O Server partition using the System Planning Tool Version 2 and HMC Version 7.3 or later. Through July 2007, Virtual I/O Server deployment is supported only on a System p POWER5 or POWER6 model system. Virtual I/O Server Version 1.4 is required for the functions described in this chapter.

This chapter cannot make you an expert on Virtual I/O Server (sometimes also referred to as VIOS) capabilities, which are included with the Advanced Power Virtualization (APV) feature available for either System i or System p. Virtual I/O Server is supported only for client partitions running supported POWER-based AIX 5.3 and Linux distributions.

In this chapter we provide a short summary of APV capabilities. More details on APV and Virtual I/O Server can be found at the IBM Systems Information Center, with many links to additional information:

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

We also recommend reviewing the IBM Redbooks publications *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration*, SG24-7940; *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194; and *Advanced POWER Virtualization on IBM System p Virtual I/O Server Deployment Examples*, REDP-4224.

This chapter shows examples of using System Planning Tool Version 2 to define a Virtual I/O Server for deployment and then deploying the Virtual I/O Server partition on a System p using the HMC. Example Virtual I/O Server configurations used were selected from the IBM Redbooks listed above.

Topics we cover include:

- Virtual I/O Server requirements
- Defining a Virtual I/O Server partition using System Planning Tool Version 2 for System p
- Defining Virtual I/O Server on System i
- Deploy using Hardware Management Console (HMC)
- Overview of system planning for Integrated Virtualization Manager (IVM)

4.1 Advance Power Virtualization capabilities summary

The Advanced POWER Virtualization feature is a combination of hardware enablement and software that includes the following components that are available together as a single priced feature:

- Firmware enablement for Micro-Partitioning™. This is required for partitioning on all System p POWER5 and POWER6 models and the IBM System i Model 515.
- Installation image for the Virtual I/O Server software, which supports:
  - Virtual LAN support, including Shared Ethernet Adapter and EtherChannel capabilities
  - Virtual SCSI server
  - Integrated Virtualization Manager (IVM) for supported systems
- Partition Load Manager (only supported for HMC managed systems and not part of POWER Hypervisor™ and Virtual I/O Server FC 1965 on IBM OpenPower systems).

APV is not required on all other System i POWER5 and POWER6 models for logical partitioning. These System i models include Micro-partitioning, Virtual LAN and Virtual SCSI capabilities. Virtual LAN is also available on System p models without this feature for a system managed by an HMC or IVM without an HMC.

When installed on a System i, neither APV's Partition Load Manager (PLM) nor IVM support an i5/OS partition.

Table 4-1 shows the current set of APV feature numbers of System p and System i.

Table 4-1 Advanced Power Virtualization Feature number - system model cross reference

System model	APV feature number	Included in base configuration?	IVM supported?
System p POWER5 and POWER6 Models			
9115-505	7432	No	Yes
9110-510	7432	No	Yes
9123-710	1965	No	Yes
9111-520	7940	No	Yes
9131-52A	7940	No	Yes
9124-720	1965	No	Yes
9113-550	7941	No	Yes
9133-55A	7941	No	Yes
9117-570	7942	No	Yes
9118-575	7944	No	Yes
9119-590	7992	Yes	No
9119-595	7992	Yes	No
9117-MMA (POWER6)	7942	No	Yes
System i POWER5 and POWER6 Models			

System model	APV feature number	Included in base configuration?	IVM supported?
System i 9407-515	#7966	No	No
All other System i POWER5 models	#7940	No	No
POWER6 models	7942	No	No

Note that APV on POWER6 models has enhanced functions over APV on POWER5 models. See IBM System p announcement letter 107-288, dated May 22, 2007, for more POWER6 information.

For more detailed APV and Virtual I/O Server information see the references listed on the first page of this chapter. One introductory PDF on Virtual I/O Server capabilities available at the IBM Systems Hardware Information Center is *System i and System p: Using the Virtual I/O Server.*

4.2 Virtual I/O Server requirements

The Virtual I/O Server is a partition that can virtualize LAN and device I/O for an AIX or Linux partition. It allows for the management of the DVD/CD, Ethernet, and disks in a virtual environment. At the writing of this book, to utilize the SPT V2 to set up the Virtual I/O Server for deployment requires Virtual I/O Server Version 1.4.0.0 or later.

New capabilities with the Virtual I/O Server V1.4 also include:

- Virtual I/O Server support on POWER6 processors
- Lightweight Directory Access Protocol (LDAP): centralized user management
- Simple Network Management Protocol (SNMP): standard interface used for monitoring and management
- Support of new storage subsystems: nSeries and NetApp® subsystems and Fibre Channel attachment
- New I/O interconnects for:
  - Serial Attached SCSI (SAS) support for disk attachment
  - Serial Advanced Technology Attachment (SATA) optical controller support
- Support of new Host Ethernet Adapters (HEA): Shared Ethernet Adapter (SEA) support under Virtual I/O Server provides bridging Virtual Ethernet to HEA

To check for the latest Virtual I/O Server release level, fix pack, and installation instruction level information, visit:

`http://www14.software.ibm.com/webapp/set2/sas/f/Virtual I/O Server/home.html`

The supported Virtual I/O clients include:

- AIX 5L Version 5.3 or later
- SUSE LINUX Enterprise Server 9 for POWER
- SUSE LINUX Enterprise Server 10 for POWER
- Red Hat Enterprise Linux AS 3 for POWER (update 2 or later)
- Red Hat Enterprise Linux AS 4 or POWER or later

4.3 Define Virtual I/O Server in System Planning Tool Version 2

As an example for SPT V2 to define the Virtual I/O partition (VIOS), we use the Virtual I/O configuration shown in 4.3.1, “Major steps to define a Virtual I/O Server” on page 107, which uses a single Virtual I/O Server server partition for two client partitions running AIX or Linux. Later we also show Virtual I/O Server configuration examples for Logical Volume Mirroring (LVM) and dual Virtual I/O Server partitions. The dual Virtual I/O Server partition configuration used in this publication has been taken as a best practice from *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194.

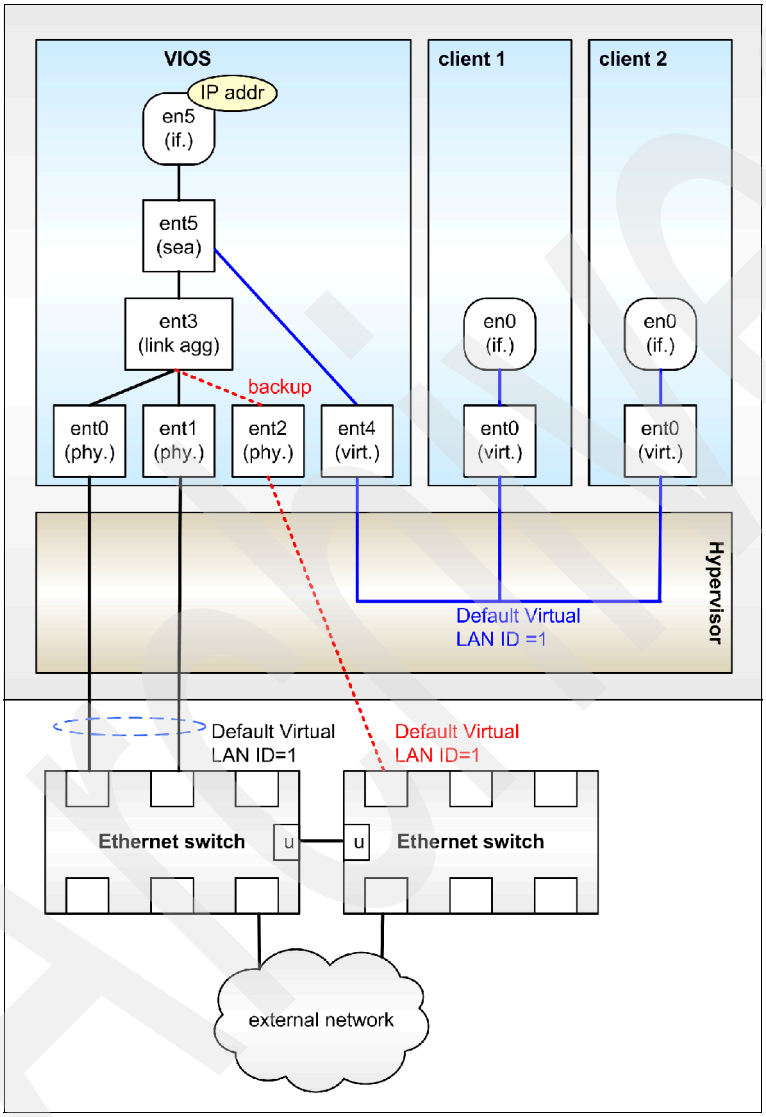

Figure 4-1 Example of a single Virtual I/O Server

In our Virtual I/O Server partition examples several terms and acronyms are used in our figures. We explain them here.

Each client partition’s virtual LAN adapter is identified as an Ethernet adapter (such as device *ent0 virt*), which is associated with one network interface (such as *en0 if*). An IP address can be defined on a network interface.

We show a Virtual I/O Server partition owning several physical Ethernet adapters and one virtual adapter. One of the physical adapters can be used as a backup adapter, as indicated by a dashed line that is implemented in association with the link aggregation device.

en0 if Ethernet interface providing programmed access to the virtual Ethernet adapter.

ent0 virt Virtual Ethernet adapter that connects to the POWER Hypervisor VLAN-aware Ethernet switch.

ent4 virt Virtual Ethernet adapter that connects the Virtual I/O Server with the relevant VLAN used in the POWER Hypervisor switch.

en5 if Ethernet interface providing programmed access to a Shared Ethernet Adapter.

ent5 sea Shared Ethernet Adapter (SEA) device. This device can connect directly to a physical Ethernet adapter and functions as a layer-2 bridge transferring packets from the virtual Ethernet adapter to the physical Ethernet adapter. It can also connect and transfer traffic to a link aggregation device.

ent3 link agg This is a link aggregation device. This is a component that implements IEEE 802.3ad Link Aggregation or Cisco EtherChannel. This link aggregation device typically has two (maximum of eight) physical Ethernet adapters connected and manages them as an single Ethernet connection.

ent0 phy. ent1 phy,

ent2 phy Dedicated physical Ethernet adapters that connect to the physical network.

Backup The backup adapter (ent2 phy) of the link aggregation device is a connection between the interface and a second Ethernet adapter that provides a secondary path to the network if the primary path has a network failure.

Consult with a networking professional for more information about networking with shared Ethernet, link aggregate, and Ethernet switch product capabilities.

4.3.1 Major steps to define a Virtual I/O Server

To define the Virtual I/O Server several steps are required. These steps include:

1. Define the system attributes.
2. Set up all the partitions.
3. Set up processors.
4. Set up memory.
5. Define the networking.
6. Define the storage.

Define the system attributes

The first step in defining a Virtual I/O Server is to chose the system type, model, memory, processor capacity (number of processors), and active processors. The first set of figures for this topic does not include explanations for how all values shown were selected. Use the extensive information in Chapter 2, “System Planning Tool (SPT) V2” on page 11, to understand all the information shown in these figures, except those that are unique to a Virtual I/O Server, which are addressed in this chapter.

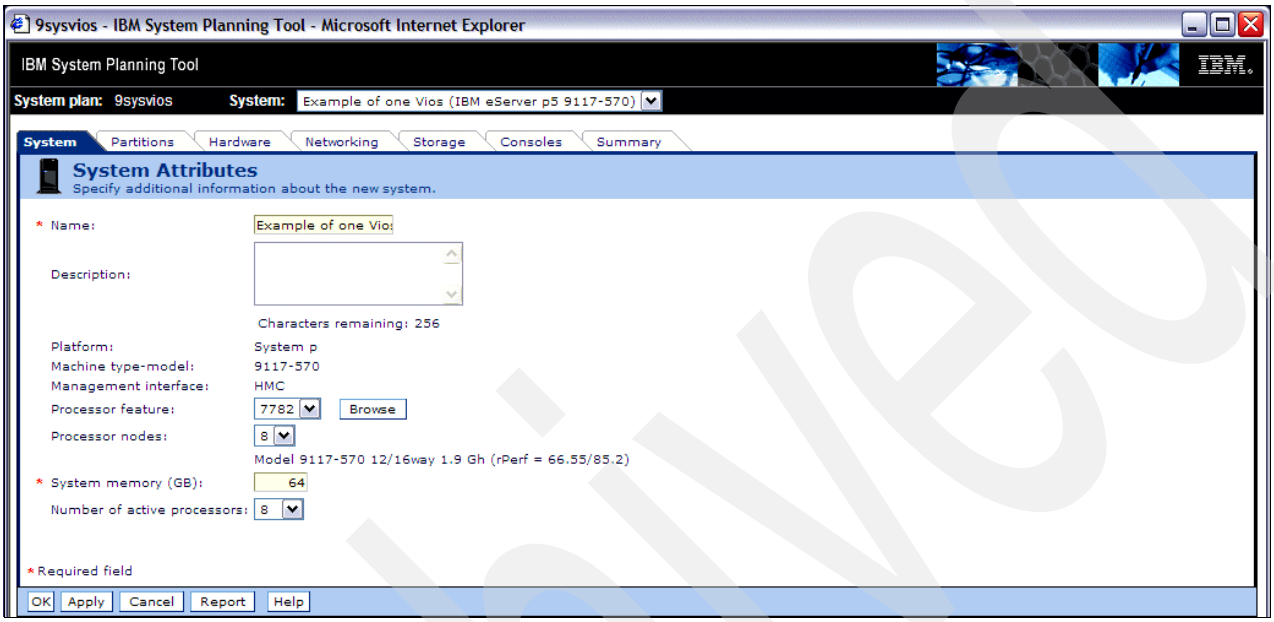

Figure 4-2 System attributes

4.3.2 Set up partitions

After choosing the model, number of processors, memory, and active processors, click the **Partitions** tab at the top of the window. Now define the partitions for names and the operating system running that partition. You may add one partition or up to the maximum supported partitions for the chosen system. Select the operating system for each partition and you can change the name from the generic name of LPARn to a name more meaningful to you.

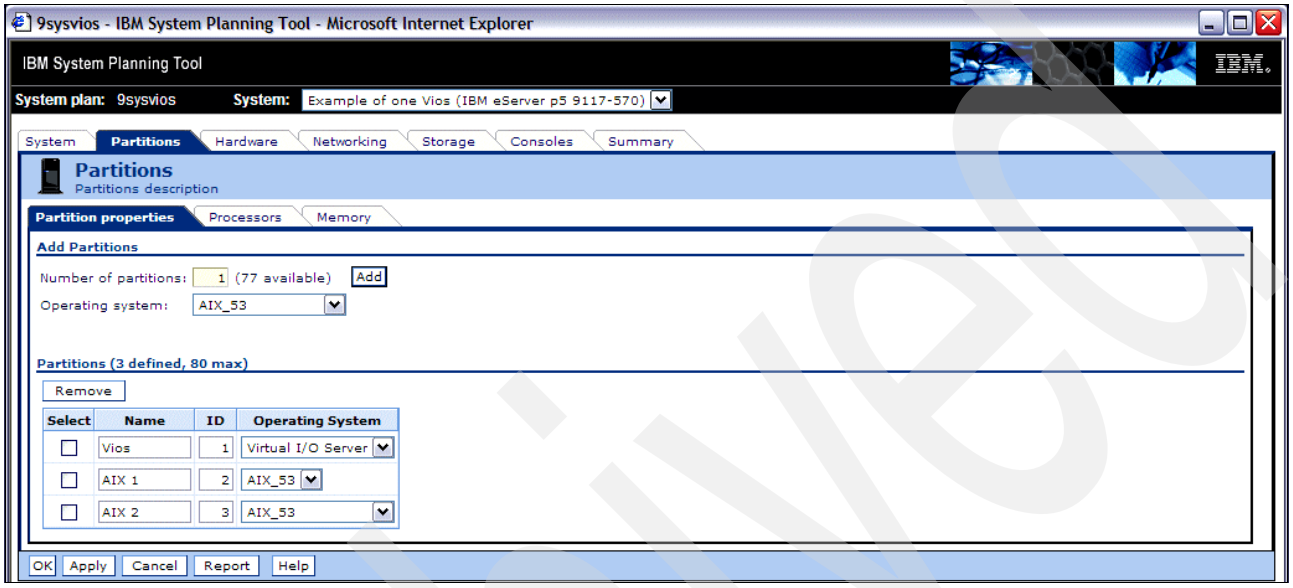

Figure 4-3 Partition definition

4.3.3 Set up processors

After the completion of the Partition tab click the **Processor** tab near the top of the window. Assign processing units, sharing mode, and virtual processor values that you have determined are appropriate for your run-time environment. For detail information about the input into each of the fields and their options shown in Figure 4-4 see 2.5, “Creating a new system plan” on page 16.

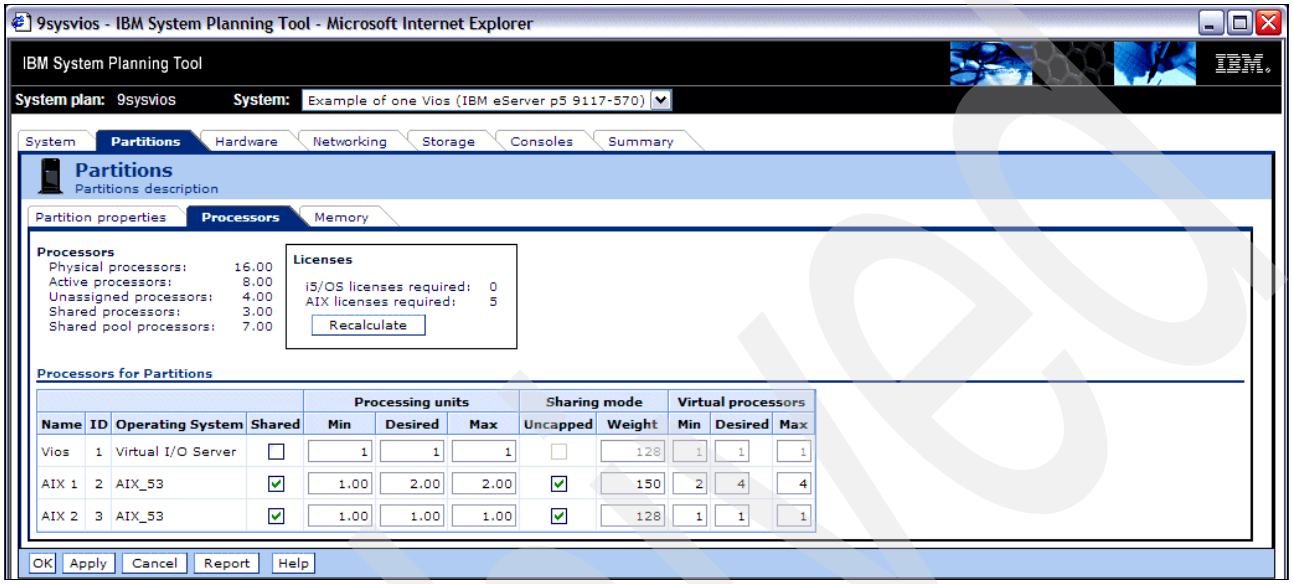

Figure 4-4 Processor definition

Important: As for any partition, sizing the processor capacity needed for good performance of applications running in that partition needs to be done. For sizing processor capacity required for a Virtual I/O Server you should consider:

- Using the IBM Systems Hardware Information Center planning assistance information for Virtual I/O Servers at:

  `http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphb1/iphb1_vios_planning.htm`

- Using the IBM Systems Workload Estimator tool available at the Web site:

  `http://www.ibm.com/systems/support/tools/estimator`

- Reviewing the Performance and planning chapter in *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194

4.3.4 Set up memory

Once the processors are defined, Click the **Memory** tab to specify the memory allocation to each partition. For more information about the options shown in Figure 4-5, see 2.5, “Creating a new system plan” on page 16.

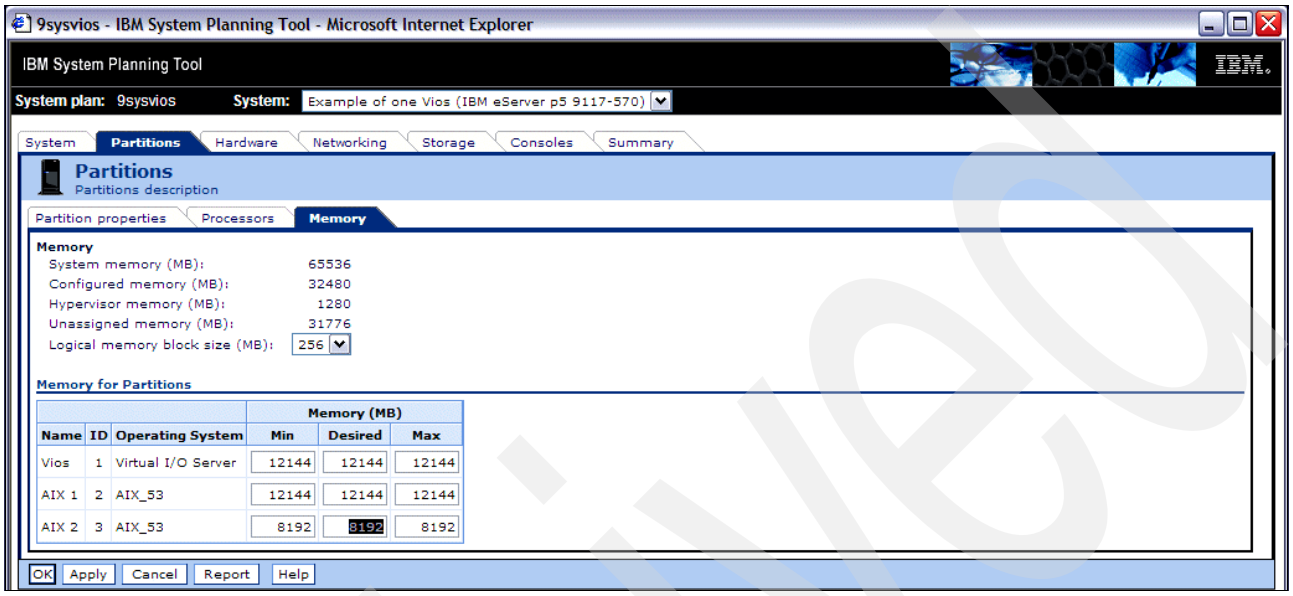

Figure 4-5 Memory definition

Important: The minimum partition memory size for a Virtual I/O Server is 512 MB. As for any partition, sizing the main memory needed for good performance of applications running in that partition needs to be done. For sizing memory required for a Virtual I/O Server, you should consider:

- Using the IBM Systems Hardware Information Center planning assistance information for Virtual I/O Servers at:

  `http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/topic/iphb1/iphb1_vios_planning.htm`

- Using the IBM Systems Workload Estimator tool available at the Web site:

  `http://www.ibm.com/systems/support/tools/estimator`

- Reviewing the “Performance and planning” chapter in *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194

4.3.5 Set up hardware

By clicking the **Hardware** tab, you define the hardware required for the Virtual I/O Server/ You also define all of the physical hardware that belongs to each partition in this part of the System Planning Tool Version 2.

For the Virtual I/O Server make sure that the requirements for physical disk and Ethernet are assigned to that partition. If a SAN disk configuration is to be used, make sure that the fiber adapter cards are placed in this section. Also specify any physical Ethernet cards that are to be used for LAN aggregation.

Add the local disks for the Virtual I/O Server to be installed on.

Figure 4-6 shows an example that includes disk, LAN adapter, and Fibre Channel adapter features.

As you can see in the Validation messages area, at the time we captured this window, an AIX partition needs additional definition.

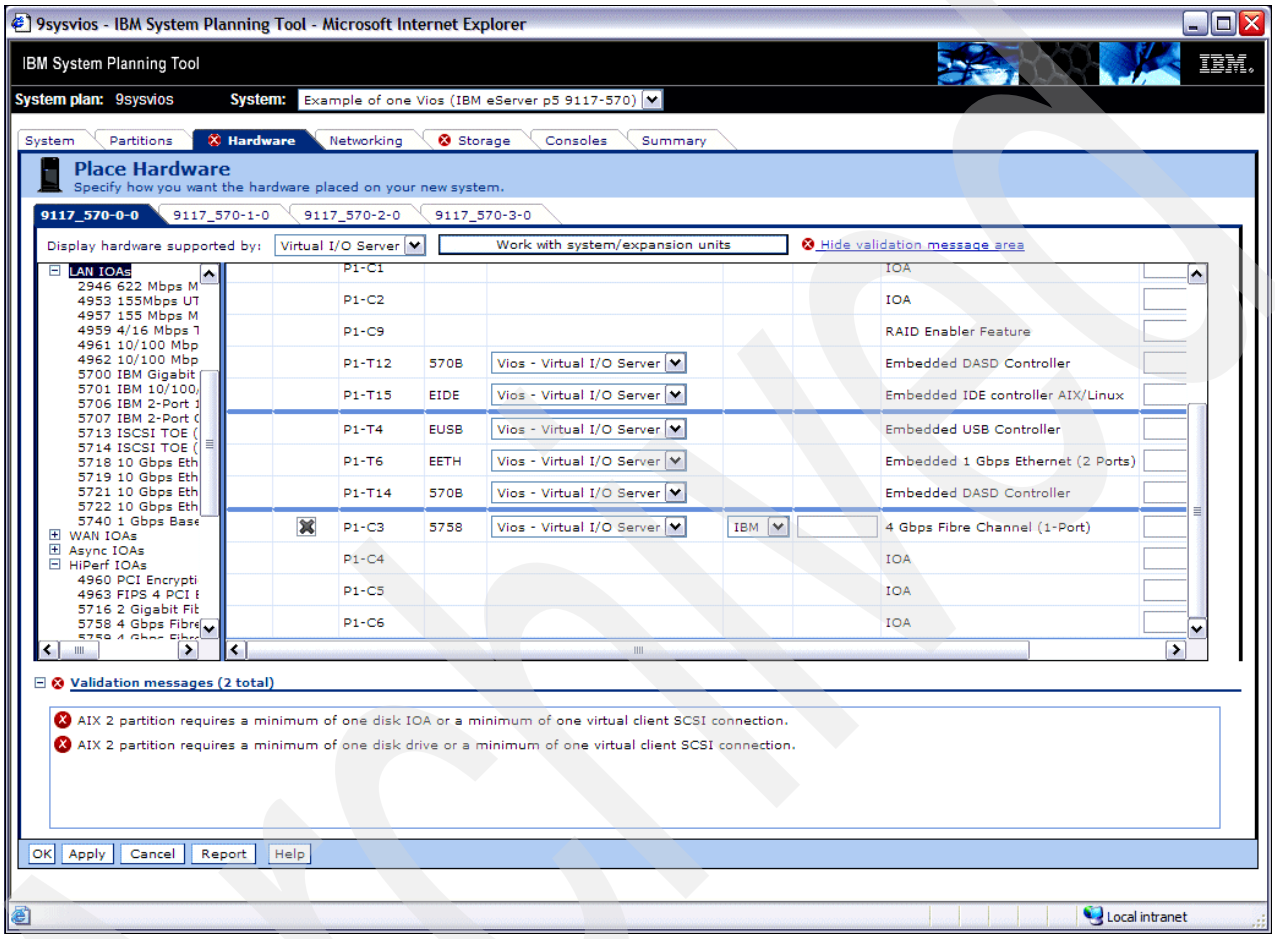

Figure 4-6 Hardware definition example for a Virtual I/O Server partition

4.3.6 Networking

This is the first major area where the power of the Virtual I/O Server partition capabilities can be put to greatest advantage. Requirements for this partition are typically different from other partitions running i5/OS, AIX 5.3, or a supported Linux distribution.

By clicking the **Networking** tab you start setting up the virtual network used by each partition serviced by this Virtual I/O Server partition. The Networking tab allows you to define virtual LANs, and set up EtherChannels and Shared Ethernet Adapters.

The combining of physical Ethernet cards can be specified with two options — EtherChannel (Ethernet Channel) and 802.3adLInk. The use of EtherChannel or 802.3ad Link must be supported by the switch and configured on the switch before these capabilities under the Virtual I/O Server. If you use a Shared Ethernet Adapter (SEA), the aggregated LAN needs to be defined before defining the SEA.

Within the Networking area you can specify to use and set up:

- Shared Ethernet Adapter Failover.
- EtherChannel.
- Network Interface Backup.
- Host Ethernet Adapter (available only on POWER6 models). Details on configuration and efficient usage of HEA virtualization capabilities are beyond the scope of this publication.

For further explanation and a best practices approach about configuring these features of a Virtual I/O Server, refer to *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration*, SG24-7940, and IBM Redpaper *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194.

The “Networking” chapter of *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194, compares SEA Failover and Network Interface Backup.

Start of network definition

The first step of the network definition is to look at the Virtual LAN (VLAN) connection. Each VLAN can be added by clicking **Add** (circled in red), as shown in Figure 4-7. We show VLAN #1 already defined. By putting a check in the box to the left of each partition under the VLAN number column you specify which partitions can connect to each other on that VLAN. SPTV2 provides an easy-to-specify interface, which builds the virtual slot connections.

Each VLAN ID - Slot ID can be edited by putting a check in the Select box to the left of the VLAN ID and clicking **Edit Virtual Slots**. See the Details - VIOS (Virtual I/O Servers) area of the window shown in Figure 4-7.

Attention: Be careful if you do editing. You must be very knowledgeable about the relationship of all the slots. This is especially important if you change slot locations after the system plan has been completed.

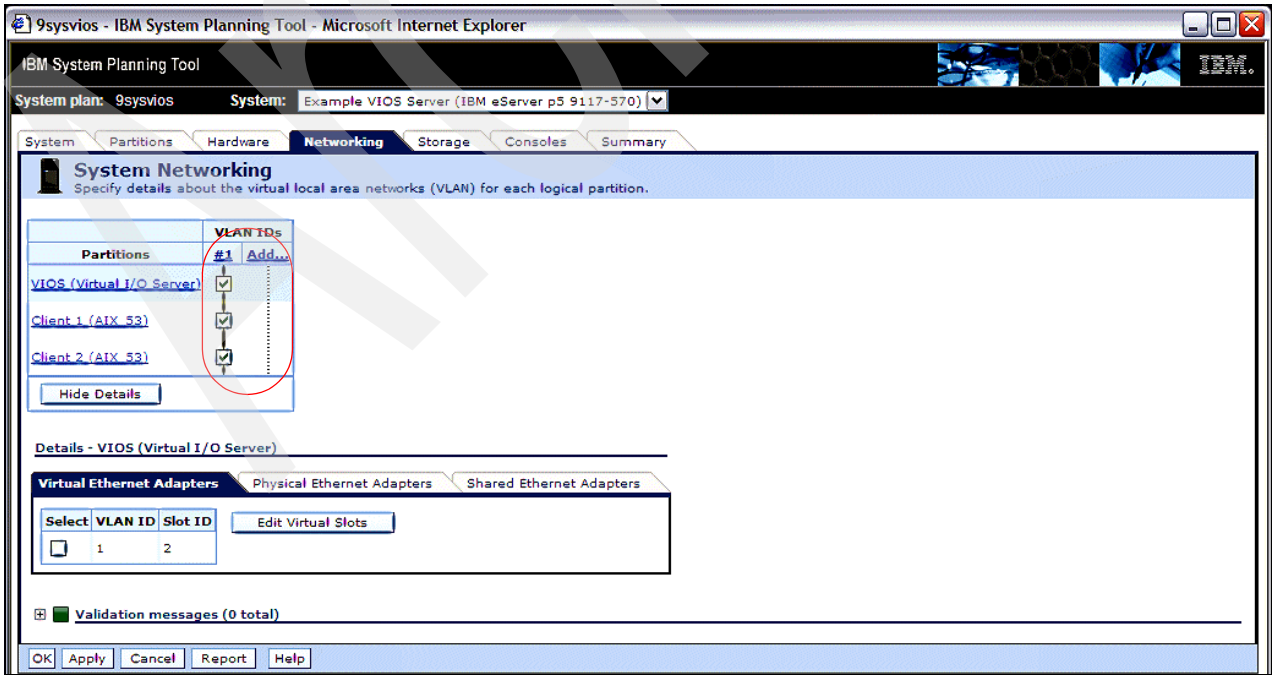

Figure 4-7 First page of defining Virtual LAN network attributes

Aggregation of the physical adapters

The Virtual I/O Server supports aggregation of physical LAN adapters. When setting up aggregation you must ensure that the switch that the physical adapters attach to supports link aggregation. The following windows would apply to the definition of an EtherChannel.

Note: EtherChannel technology was invented by Kalpana in the early 1990s. Cisco Systems acquired the company during 1994. During 2000 the IEEE passed 802.3ad, which is an open standard version of EtherChannel.

EtherChannel is a port trunking technology primarily used on Cisco switches. It allows grouping of several physical Ethernet links to create one logical Ethernet link for the purpose of providing fault-tolerance and high-speed links between switches, routers, and servers.

Referred to *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194, for background information about networking capabilities and parameters under Virtual I/O Server support.

Click the **Physical Ethernet Adapter** tab shown in Figure 4-7 on page 113 to get the window shown in Figure 4-8.

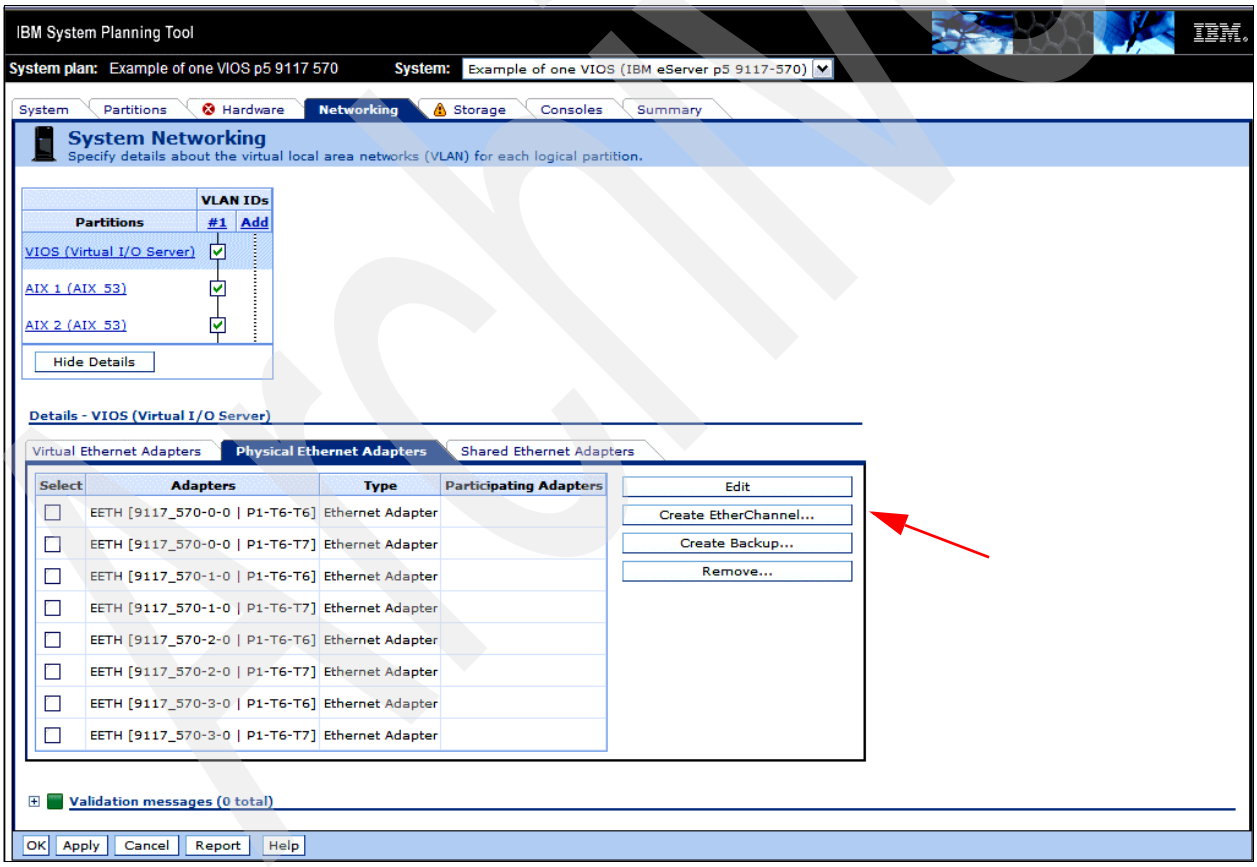

Figure 4-8 Starting to assign physical adapters to be aggregated as an EtherChannel

Note: The two-port physical LAN adapters shown in our example (Figure 4-8) are those that come with each Model 570 processor enclosure (CEC). You can tell this by looking closely under the Adapters heading and noting 9117 570-0-0, 9117 570-1-0, 9117 570-2-0, and 9117 570-3-0. On a system you could certainly have additional adapters to in an I/O expansion tower or expansion drawer to select to be managed by your Virtual I/O Server partition.

Click **Create EtherChannel** to get the Create Etherchannel window shown in Figure 4-9. This window shows the base EtherChannel settings.

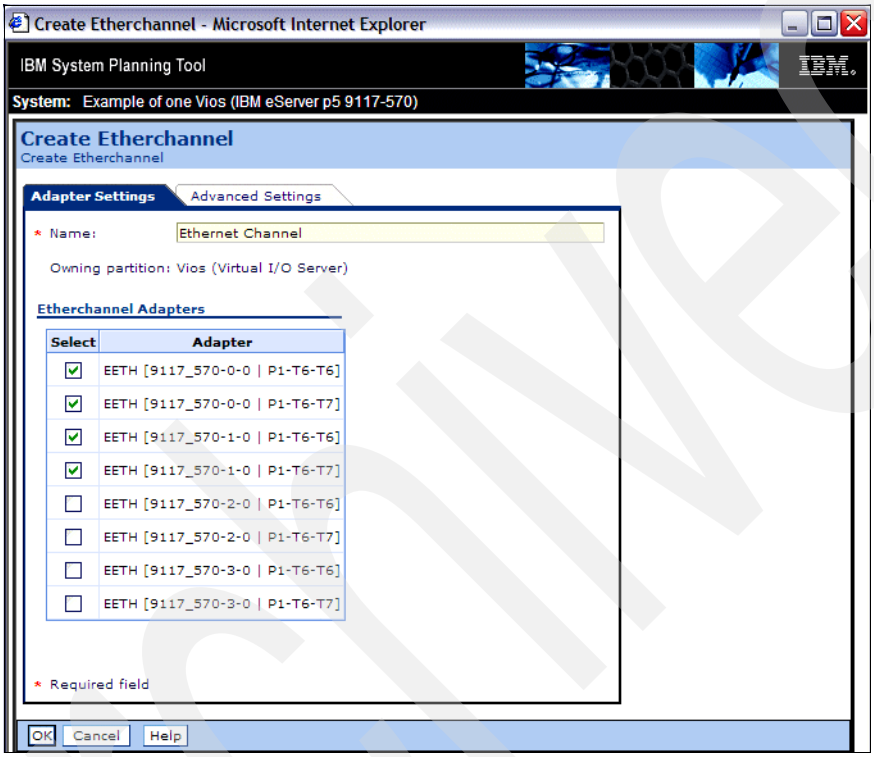

Figure 4-9 EtherChannel definition, basic adapter settings

Give the aggregated LAN a name and put a check beside each physical adapter that is to be included in the EtherChannel.

The name you specify for EtherChannel, Shared Ethernet (SEA), or Interface Backup is used only when defining and viewing the sysplan you create. The devices created when you have deployed the sysplan do not show this name.

Note: At the time the windows shown in Figure 4-8 through Figure 4-11 were captured, the adapter locations were shown as Pn-Tm-Tm (P1-T6-T6 and P1-T6-T7 in our example). This is incorrect. An update to System Planning Tool Version 2 is planned that will properly display a location as Pn-Tm, or, using our example, as P1-T6 and P1-T7.

Go to the SPT Web site at:

`http://www.ibm.com/systems/support/tools/systemplanningtool`

See 2.2, “Downloading SPT” on page 12, for additional subscription to updates information.

By clicking the **Advanced Setting** tab, you see the advanced settings for our EtherChannel defined in Figure 4-10. On this Advanced Settings window you define settings, which include:

- Alternate address setting
- Enablement of gigabit jumbo frame
- Backup adapter for the primary aggregated adapter

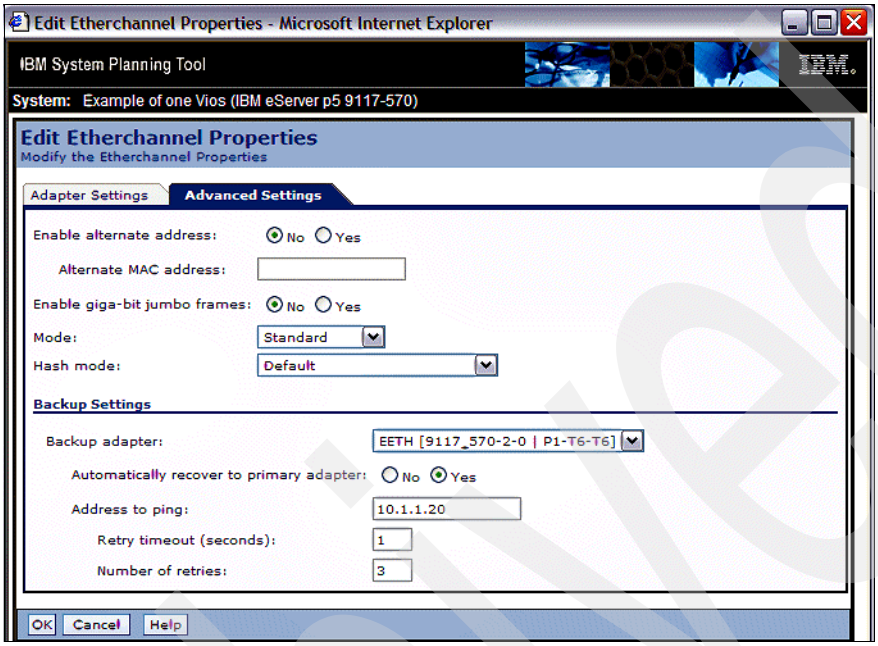

Figure 4-10 Advance settings for EtherChannel

Once you have finished reviewing all of your settings, click **OK**. You are returned to the initial Networking window, which look similar to the one in Figure 4-11.

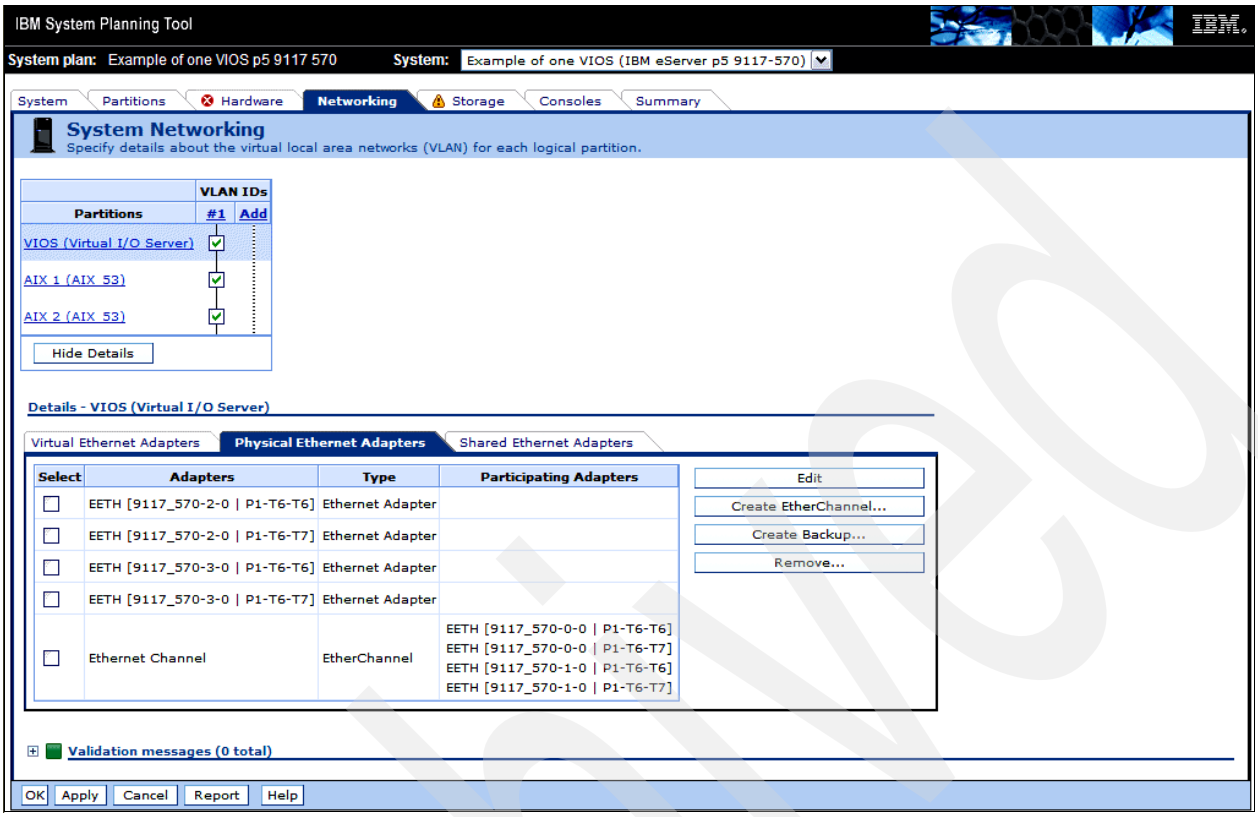

Figure 4-11 VIrtual I/O Server with an EtherChannel defined

Shared Ethernet Adapter

On the initial Networking window (for example, Figure 4-7 on page 113, or Figure 4-11 in our EtherChannel example) you can separately set up under a Virtual I/O Server an adapter that is shared between partitions.

In the window shown in Figure 4-12 we have already clicked the Shared Ethernet Adapters tab, which shows that we have not yet defined a Shared Ethernet Adapter. Click **Create Adapter**.

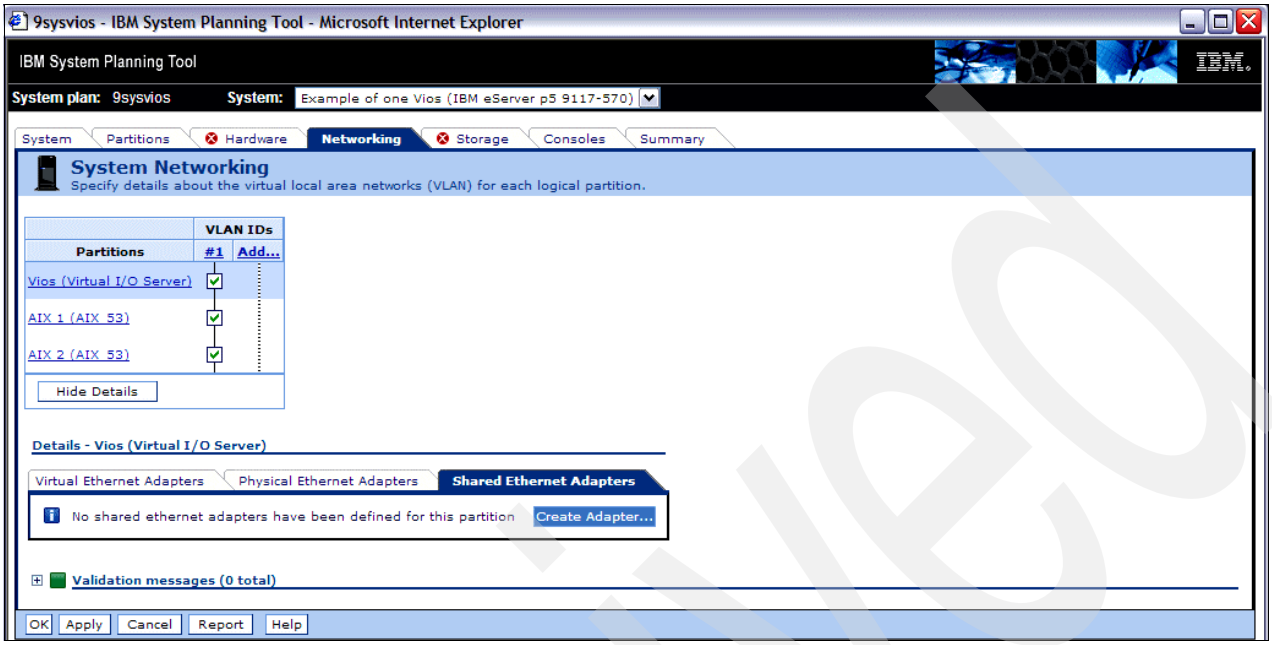

Figure 4-12 Create Shared Ethernet Adapter

Clicking **Create Adapters** brings up the Create Shared Ethernet Adapter window shown in Figure 4-13. Enter a meaningful name for your SEA in the required name field.

Clicking the down arrow on the Physical adapter field shows a list of available physical adapters, as shown in our example. Select the adapter to be shared. The selected adapter may be a single physical Ethernet adapter or it may be one of the named aggregated adapters (EtherChannel).

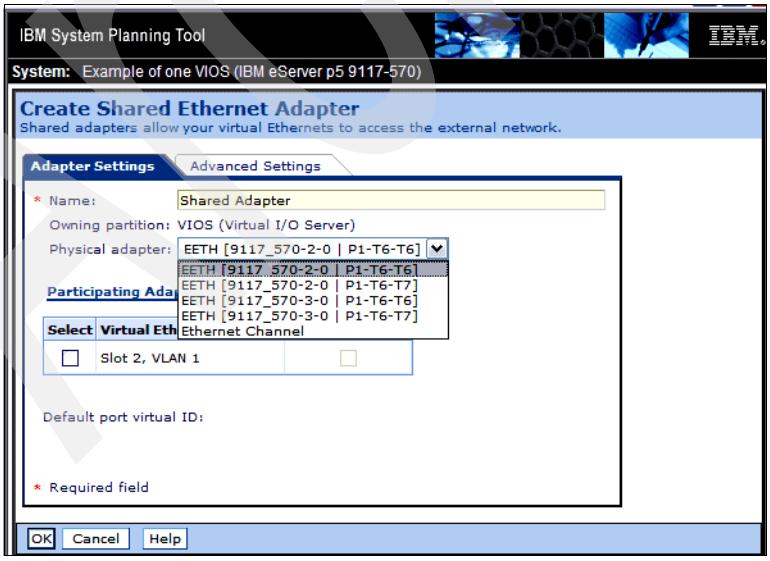

Figure 4-13 Create the SEA, base settings

By clicking the **Advanced Settings** tab you may define:

- Threading mode
- Trunk priority
- Failover mode
- Control channel
- Ping address for a live connection

As seen in Figure 4-14, we show the default values. Make any changes your network administrator has determined. When complete, click **OK** to return to the initial Networking window.

Important: Setting up SEA failover correctly requires a knowledgeable network administrator. Additional information is beyond the scope of this book. The IBM Redpaper *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194, contains an example of using this capability and other advanced settings parameters.

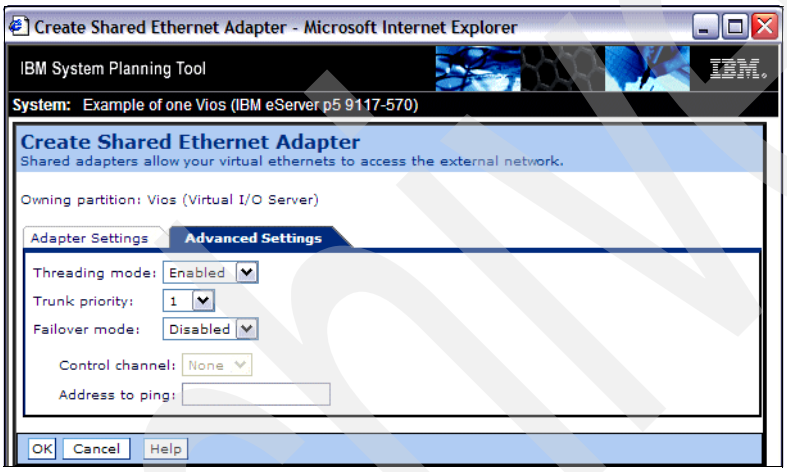

Figure 4-14 Advance settings for Shared Ethernet Adapter

When finished reviewing all of your SEA settings, click **OK**. You are returned to the Networking window showing your defined SEA, as shown in Figure 4-15.

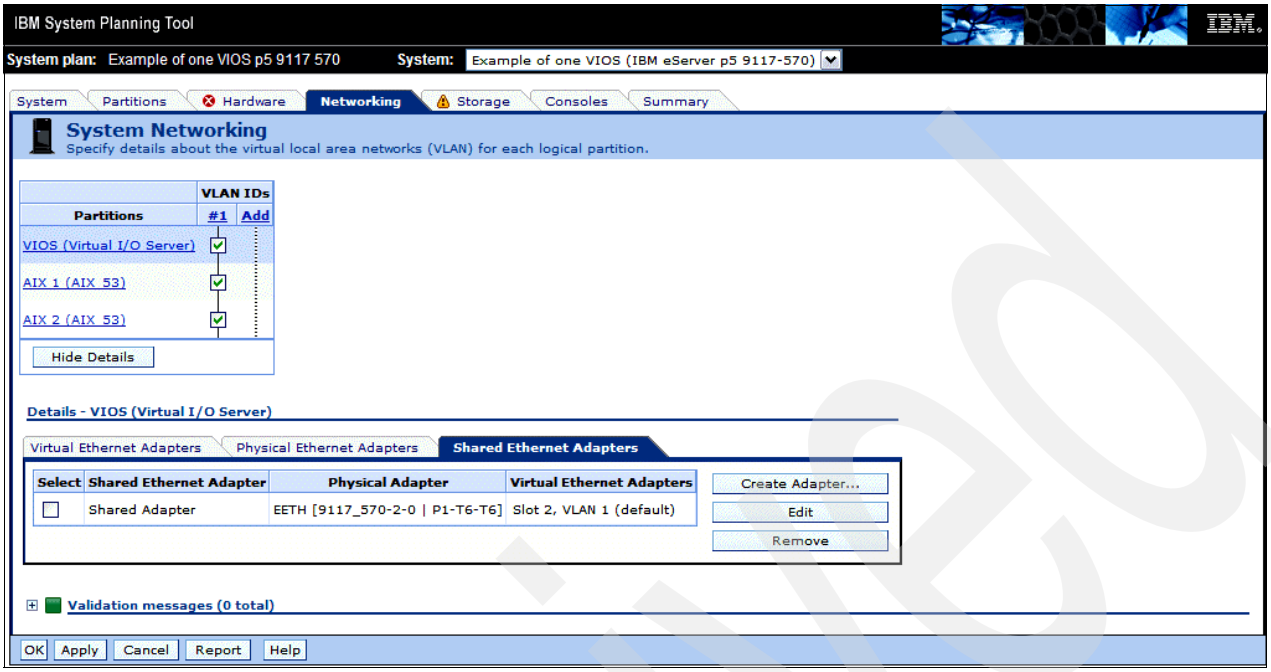

Figure 4-15 Defined Shared Ethernet Adapter

This ends our coverage of setting up the Network configuration for a Virtual I/O Server partition.

4.3.7 Storage

The second part of the Virtual I/O Server definition is the storage section. In this section you can:

- Define Virtual I/O Server install source and target disk.
- Define the disk.
- Create disk volumes.
- Assign host/server storage (physical disk drives, SAN volumes, or storage from a storage pool) from Virtual I/O Server partitions to client partitions.
- Set up a SAN volume for Multiple Path Input Output (MPIO).
- Set up storage pool mirroring.

The steps for defining are:

1. Using the Storage tab, as shown in Figure 4-16, define the source Ethernet adapter for the Virtual I/O Server install and the target disk volume for the Virtual I/O Server.

   a. Specify an Ethernet adapter in the Install source field. The Ethernet adapter you chose must be one of the physical Ethernet adapters you defined as being managed by this Virtual I/O Server partition.

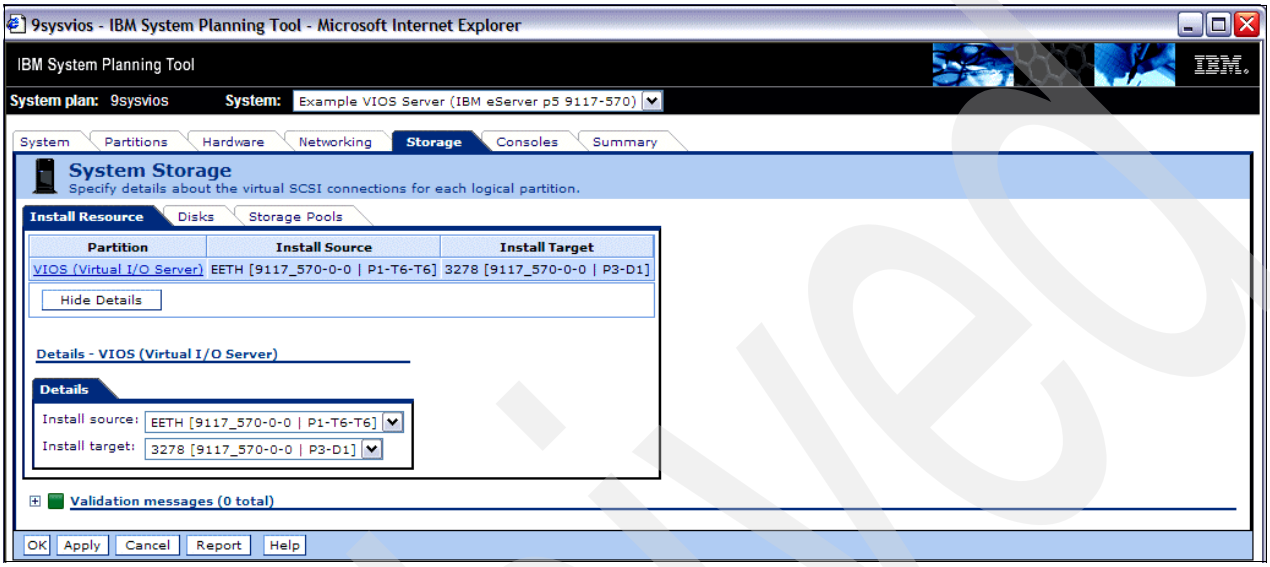

Figure 4-16 Install location for the source and the target

You have to specify that this adapter will be used by the deployment code to configure for network install and this adapter can be used for creation of SEA while doing the provisioning step of deployment. The IP address configured on the interface will be migrated to the adapter that is created on top of this adapter.

b. Specify a disk drive in the Install target field. The target disk can be a physical SCSI or SAS (for example, within a 9117 MMA POWER6 processor enclosure) disk or a SAN disk.

   If the target disk is a SAN disk then the system plan deployment process will not install the VIOS on this partition. The VIOS partition has to manually install an active IP address configured on the interface before redeploying the system plan.

Provisioning under a Virtual I/O Server means creation of:

- Shared Ethernet Adapter
- Etherchannel
- Storage pools
- Backing devices (the actual physical storage devices, for example, physical disks or logical volumes when using Storage Area Network (SAN) storage servers)

2. Click the **Disks** tab to get a window showing any disks assigned to this partition. In our example we include creating a SAN volume. Therefore, on the initial Disks window (Figure 4-17) click **Create SAN Volume**.

   You can create one or more SAN volumes owned by this Virtual I/O Server partition. Note that to create a SAN volume you must have previously assigned a Fibre Channel Adapter to this partition. SAN volumes are used to create a storage volume for each of the clients. Figure 4-17 shows the window to create the SAN volume.

   The SAN volume that you create is one that you plan for your client (hosted) partitions to use.

   Your volume must match the corresponding volume that you defined on the actual external storage *device* hardware configuration.

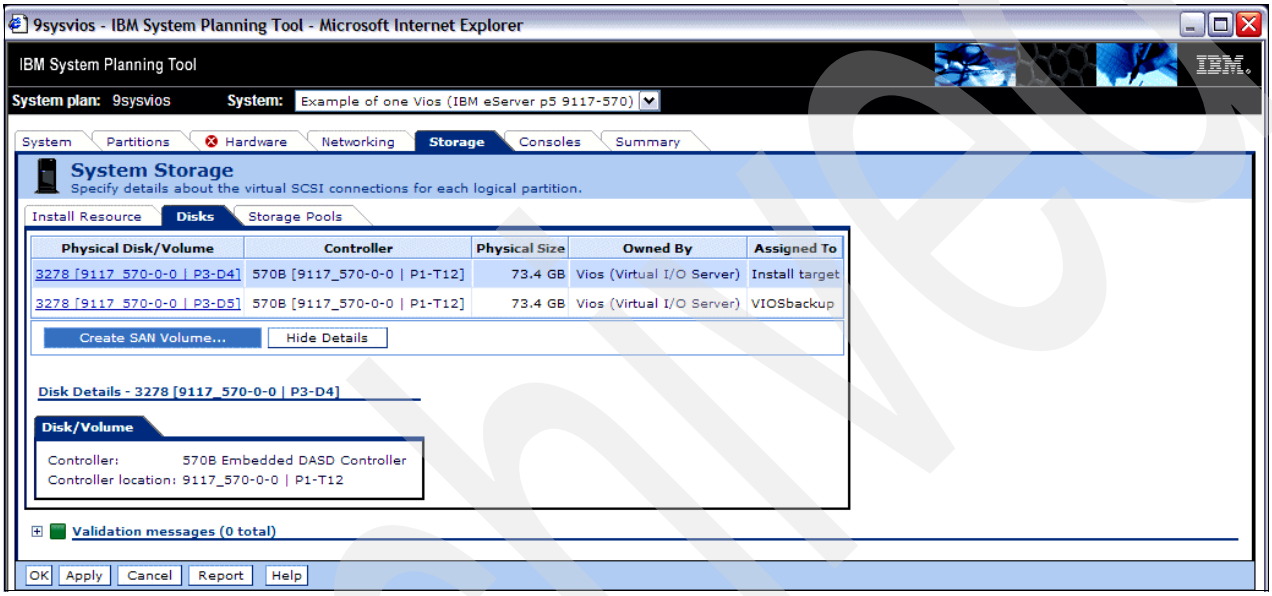

Figure 4-17 Create SAN volume

3. On the window shown in Figure 4-18 you describe the volume characteristic.

   a. Give it a name.
   b. Set the size of the volume.
   c. Set the number of volumes.
   d. Chose the fibre change adapter connected to the SAN.

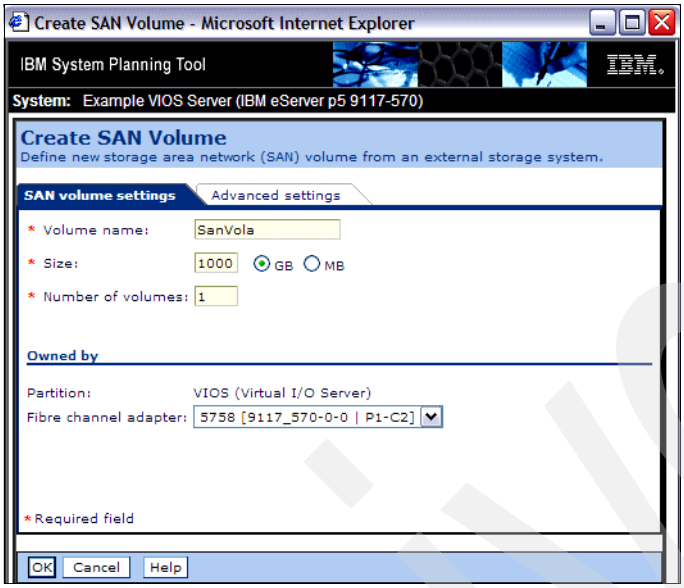

Figure 4-18 Definition of a SAN volume

4. On the Advance Setting tab (Figure 4-21 on page 125) you can enter information in the World-wide name and Logical unit number fields. They are defined as:

   – World-Wide name
     - The unique name for the connection to the SAN and if used must match.
     - The field is limited to 16 alphanumeric characters.
     - If this field is blank then the SAN volume assignment during deployment is made based on size.
   – Logical unit number
     - The logical unit number of the SAN volume.
     - This field is limited to 18 alphanumeric characters and must be properly formatted starting with the two characters 0x and followed by a 16-digit hexadecimal address.

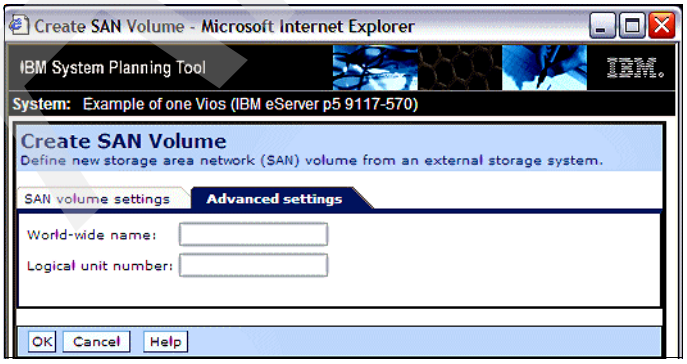

Figure 4-19 Advance setting of the SAN volume

Click **OK** to return to the Storage - Disks window.

In our example (Figure 4-20) the line showing SANVola has been added to the Disks window. We already double-clicked SANVola to get its details, shown in the lower area of the Disks window.

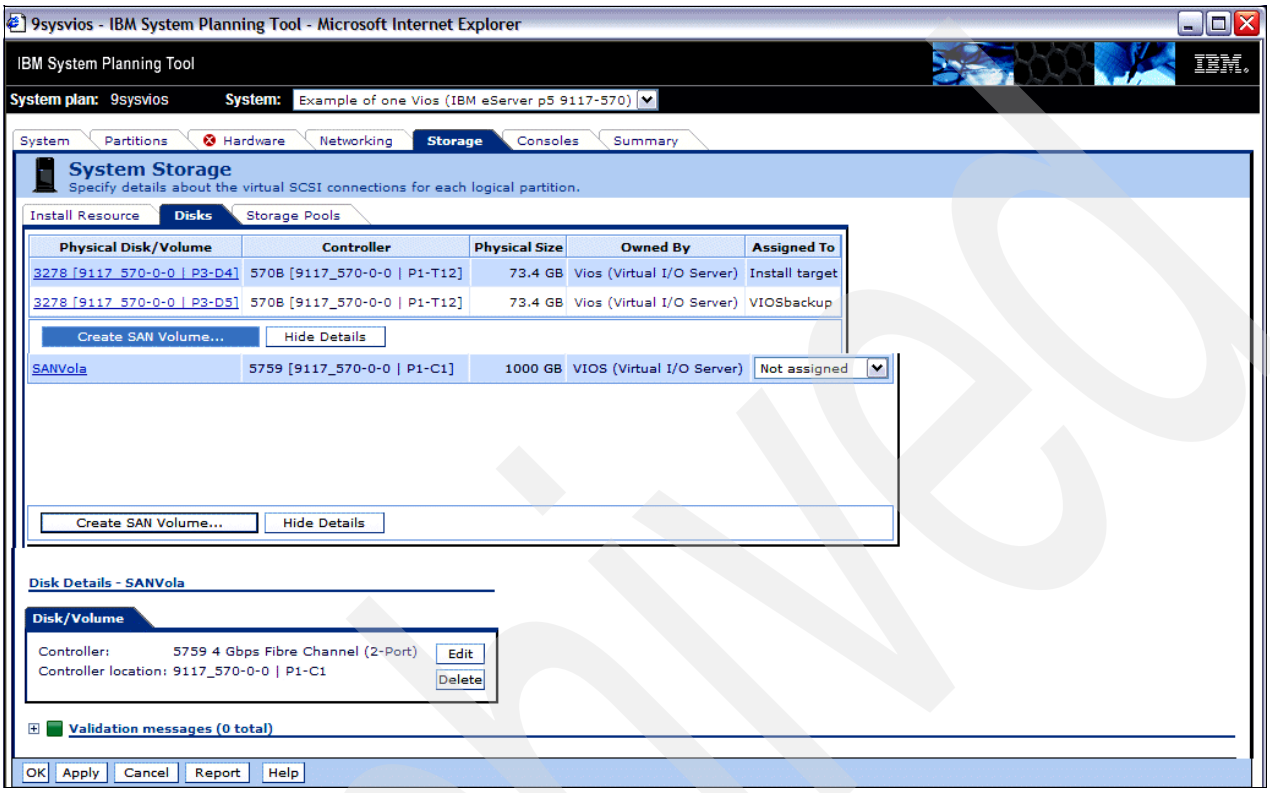

Figure 4-20 Virtual I/O Server disks with new SAN volume shown

You can now define storage pools from which storage can be assigned or hosted or served (all terms used to mean the same thing) to client partitions. Click the **Storage Pool** tab.

You Virtual I/O Server can define storage pools for any real disks and SAN volumes already defined for use by this partition. Each client partition must have storage volumes assigned. Figure 4-21 shows the naming of a storage pool, which contains the storage for the disks assigned to each partition.

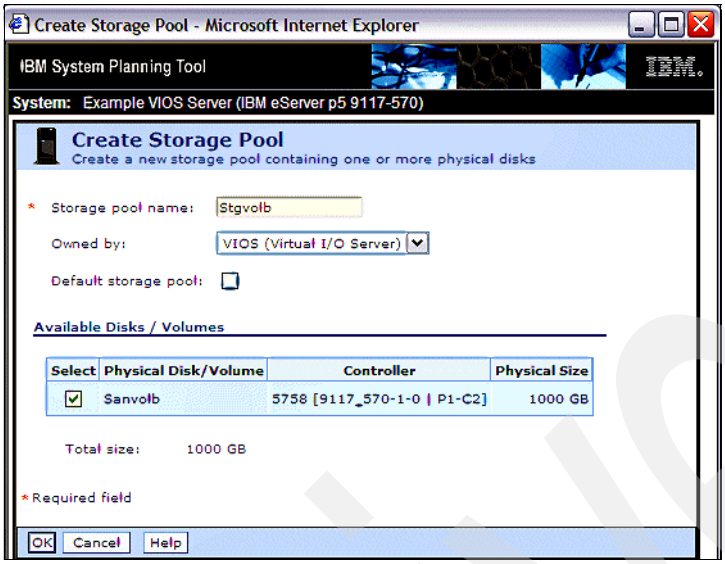

Figure 4-21 Create the storage pool

The Default storage pool box can be checked. Select whether you want to make this new storage pool the default storage pool for the partition that owns it. Click **OK** and you are returned to a window that looks similar to the one in Figure 4-22.

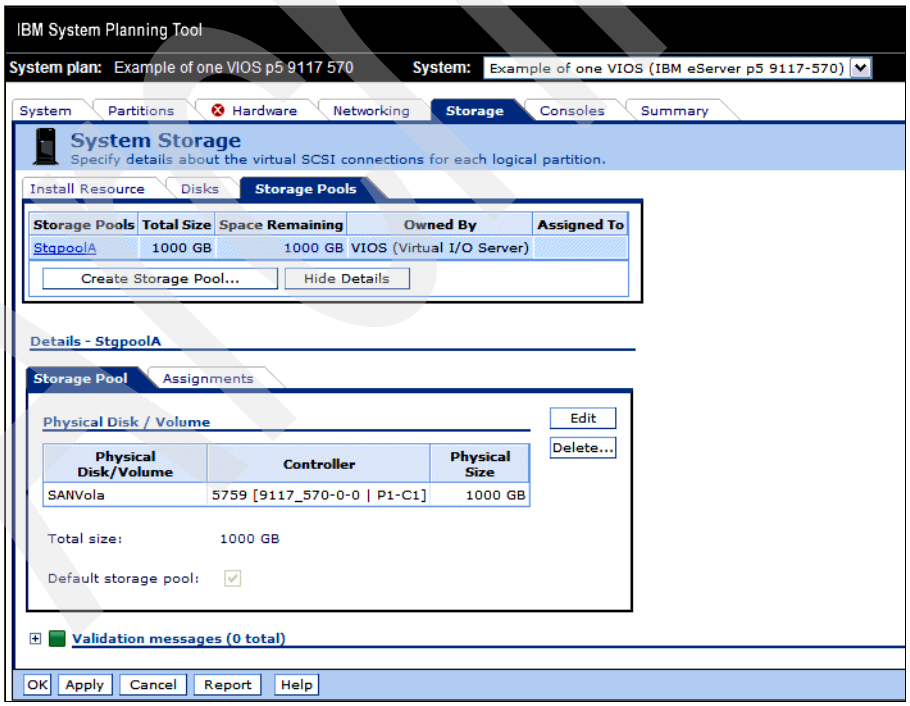

Figure 4-22 Storage pool example for SAN volume

In the Details area for the storage pool (StgpoolA) in our example, click the **Assignment** tab to define the storage for each of the partitions associated with this Virtual I/O Server. You first

get a window not shown, on which you must click **Add** to get the example window shown in Figure 4-23.

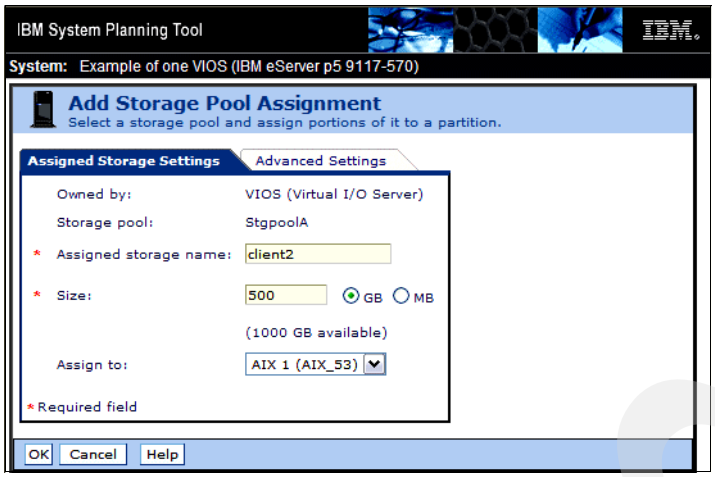

Figure 4-23 Assigning storage to a partition

In the assignment window you name the assigned storage, specify its storage amount, and specify the partition name the pool is assigned to.

The storage pool name is used to create a volume group under the Virtual I/O Server with that name.

The assigned storage name is used to create the mapping name for the assignment and also the logical volume name.

On the Advanced Setting tab you may set up mirroring. We do not show the advanced settings window.

Click **OK** to return to a window that looks similar to the one shown in Figure 4-24. You can see the total available storage and the amount of assigned storage for each partition.

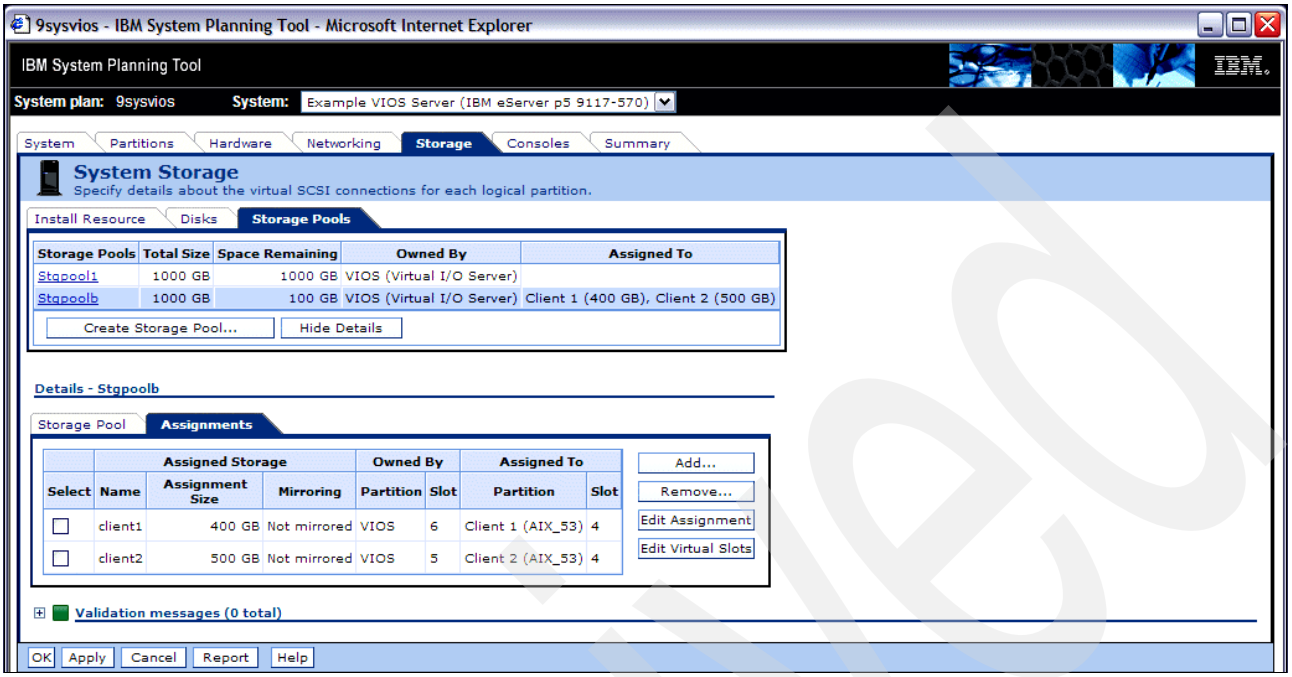

Figure 4-24 All assignments

In our example we are ready to complete our system plan, but first have to assign a console. Click the **Console** tab near the top of the window shown in Figure 4-24 and follow the steps documented in 2.6.8, “Consoles tab” on page 50.

You can refer back to Chapter 2, “System Planning Tool (SPT) V2” on page 11, which follows the Consoles Tab information, to fully complete your system plan and save for possible use in ordering a system and later deploying your system plan containing a Virtual I/O Server partition.

We show two additional network configurations with Virtual I/O Server partitions that are relatively complex. We show them because they are examples included in *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194. Setting up the complete configurations for these example Virtual I/O Server, networking and client partitions is beyond the scope of this book. However, they are shown to help illustrate the capabilities of using Virtual I/O Server partitions.

Figure 4-25 shows a possible configuration supporting network interface failover using dual Virtual I/O Servers.

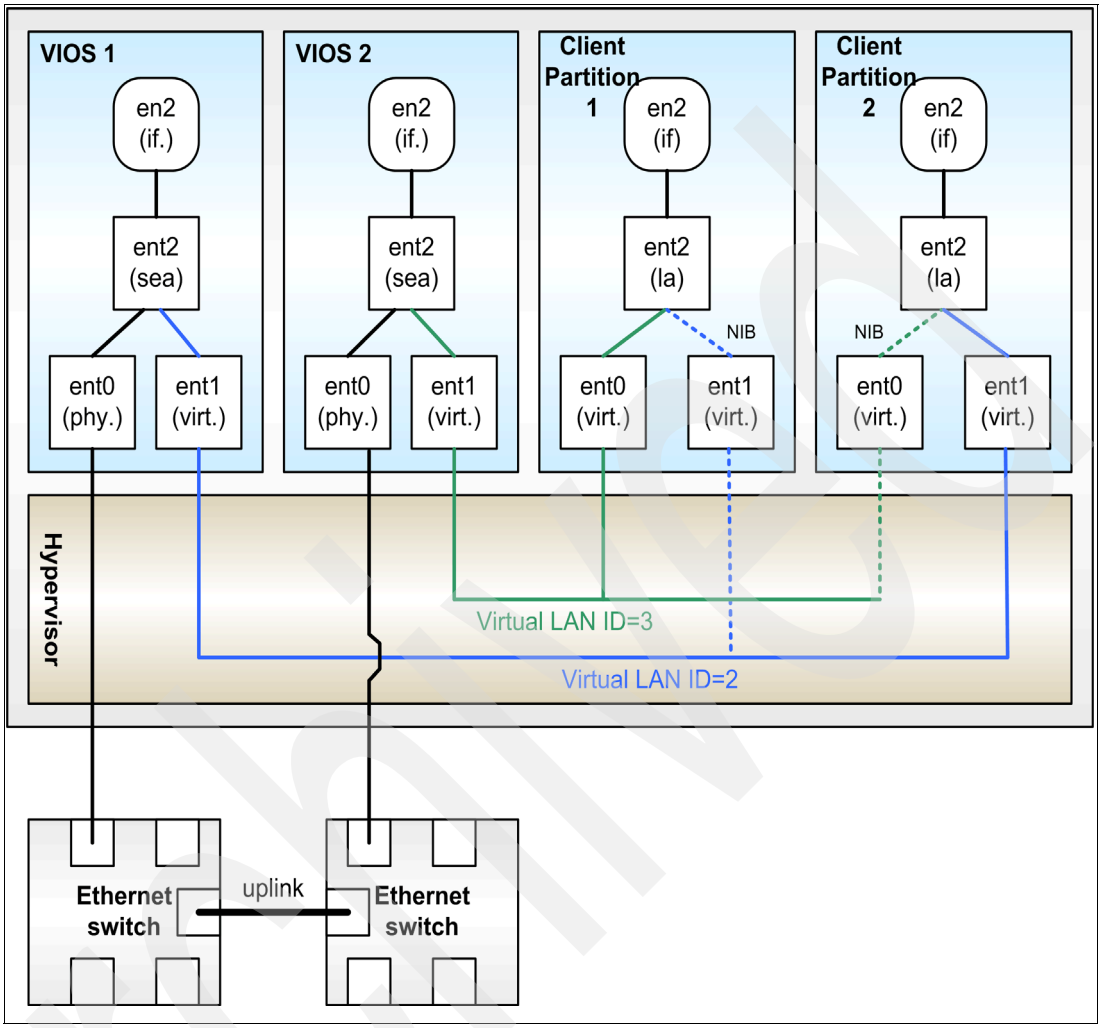

Figure 4-25 Network interface failover using dual Virtual I/O Servers

Figure 4-26 shows a configuration with dual Virtual I/O Servers connected to SAN storage using multipath I/O (MPIO). This configuration is not supported in System Planning Tool Version 2.

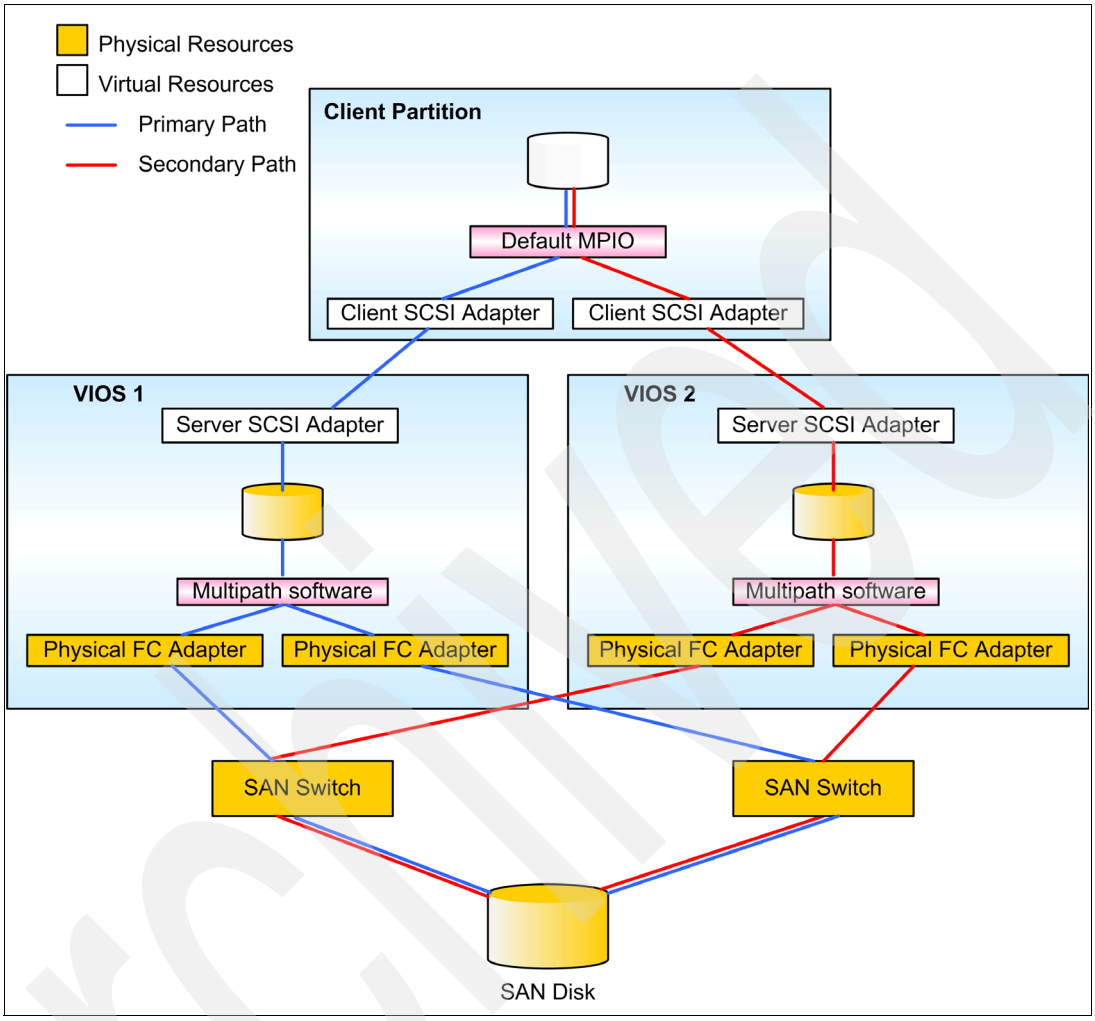

Figure 4-26 Dual Virtual I/O Servers connected to SAN storage using MPIO

Other examples using Virtual I/O Server can be found in *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194, providing data protection resiliency to each client and to a Virtual I/O Server itself.

This ends the sections describing setting up a system plan with a Virtual I/O Server partition. The following topics cover deploying a system plan with a VIrtual I/O Server on the actual system hardware configuration.

4.4 Deployment of Virtual I/O Server on HMC

The deployment of a system plan with a Virtual I/O Server takes several steps for the full deployment. This deployment assumes a *genesis system*, that is, a System p as received from IBM manufacturing. If partitions have been configured and operating system software has already been installed, deploying a partition using our example may fail.

Note: The VIOS Install DVD is needed for the process being described here. You cannot use the update CDs. This DVD is not available as a download and must be ordered.

To keep up to date on your software entitlements for your machine, consider registering and using the IBM Web site:

https://www-05.ibm.com/servers/eserver/ess/OpenServlet.wss

As part of the Advance Virtualization Manager product, you receive a DVD that contains the load image for the Virtual I/O Server. During the setting up for a Virtual I/O Server partition deployment, and actually deploying the system plan, there are some deployment steps you have do additional planning for, before starting the deployment. We describe them here:

1. The loading of the Virtual I/O Server software image to the HMC and the name you assign to that image on the HMC's disk drive. Place the Virtual I/O Server software DVD in the DVD drive of the HMC. From the HMC chose the **HMC Management** link on the left side of the window, then chose **Open Restricted Shell Terminal** from the right side. This takes you to a command-line entry. Enter the command:

   ```
   OS_install -o define_resource -a type=AIX -a version=1.4.1.0 -a
   location=/export/resources/vios -a source=/dev/cdrom vios1_res
   ```

   Make sure that the names remain vios and vios1_res on the command. This loads the load image of the Virtual I/O Server on the hard drive of the HMC. This image is used later in the deployment of the Virtual I/O Server partition. You will see viso1_res later on in this deployment topic. You must use the upper and lower case characters as shown. The VIOS version (1.4.1.0 in our example) must be entered exactly.

2. Understand the IP addressing and subnet mask for your network to ensure that the install of the Virtual I/O Server over an IP address you specify does not conflict with other IP addresses and subnet masks you use in your network.

Figure 4-27 shows what a typical network might look like. During use of the HMC GUI install you will be presented with a window to enter information about the IP address, mask, and gateway of the physical adapter where the Virtual I/O Server is to be installed. The HMC must have two Ethernet ports, as shown in Figure 4-27, for the following definition to work.

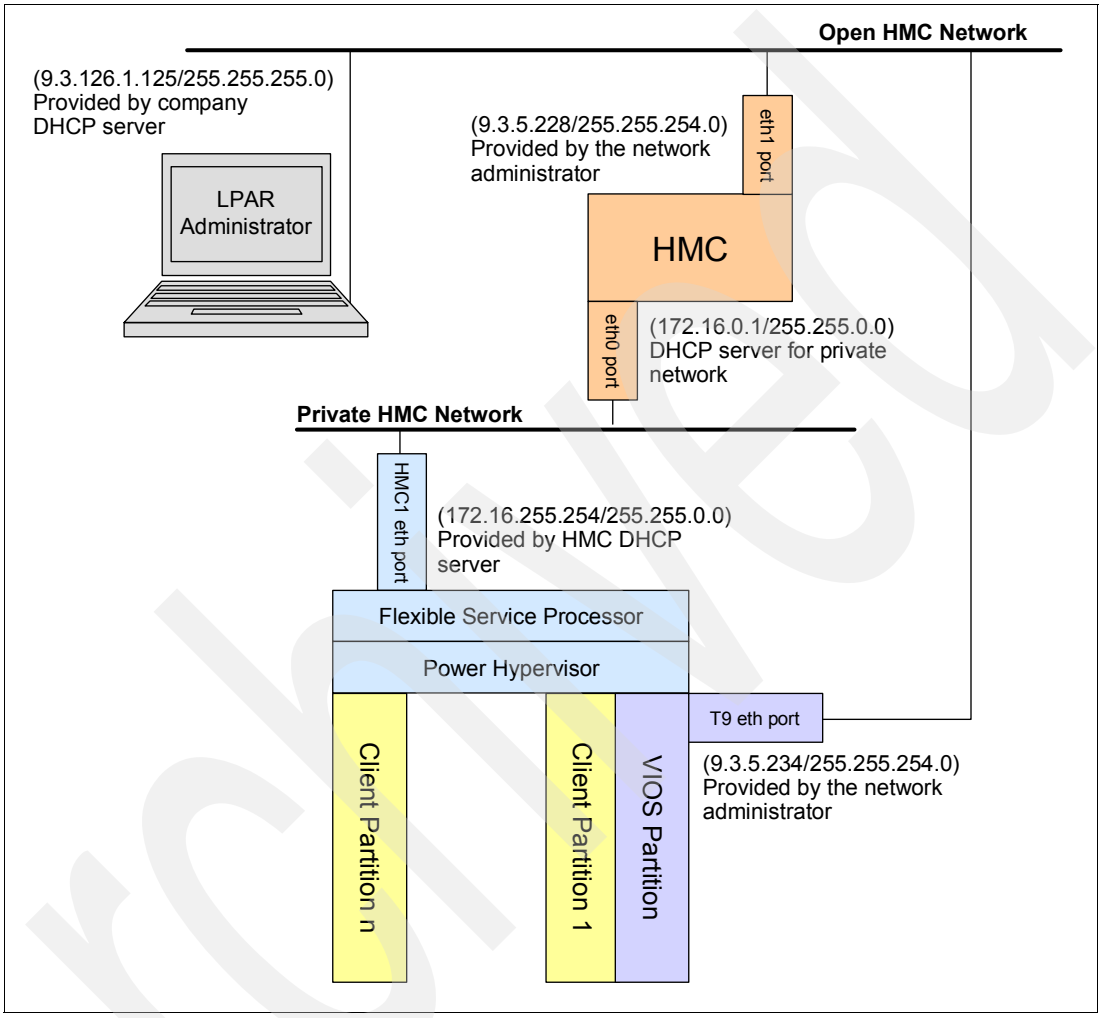

Figure 4-27 Example LAN network used to deploy our example Virtual I/O Server partition configuration

Your HMC must be configured with two LAN ports. For deskside HMCs this may require specifically ordering the second port. One port, eth0, will be connected to a private network. The second port will be connected an open network.

You must configure HMC eth0 physical port as a DHCP server. You specify a range of IP addresses for this DHCP server. Figure 4-28 shows the range selected in our example and that this adapter is on a private network. When you turn on the System i or System p and the service processor (Flexible Service Processor (FSP)) this port's cable is plugged into this private network, the System i or System p receives an a IP address. This IP address is the one used to connect to the FSP. You would not want to use this network segment to install the Virtual I/O Server. DHCP is not supported on the install. You would need to specify an IP address within this segment.

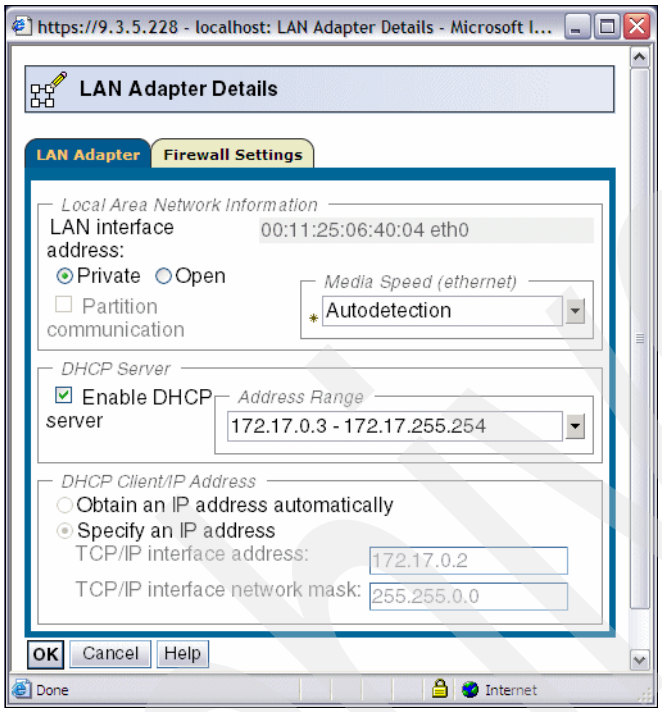

Figure 4-28 Ethernet 0 on the HMC connected to the Flexible Service Processor

The second HMC port - eth 1 is connected to a public network. It is assigned a fixed address that is typically provided by the network administrator. Figure 4-29 shows the address that matches the one shown in Figure 4-27 on page 131. The HMC can now be used by someone on the public network to manage the systems the HMC controls.

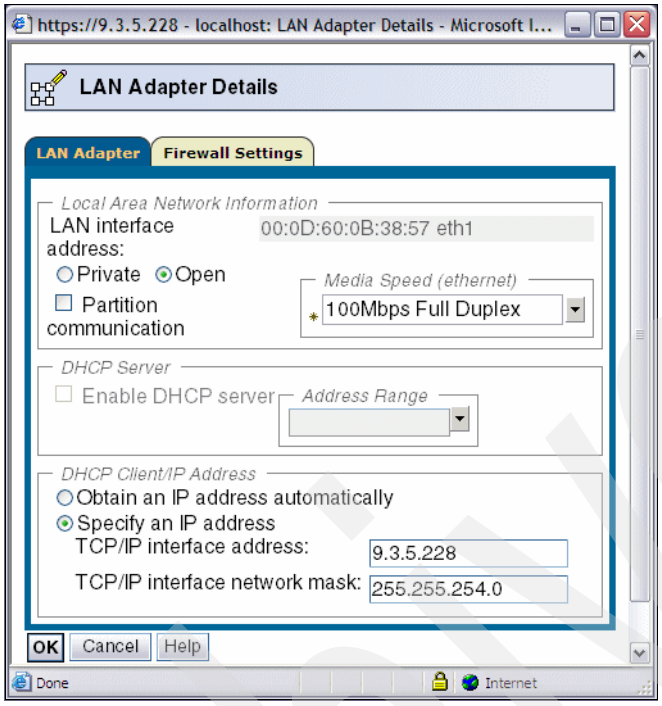

Figure 4-29 Ethernet 1 second port of HMC for Public Network

Another important IP item needs to be set up on the HMC for the public, and that is the gateway to get to other network segments. The gateway is the port in a router that allows you to connect to the other segments. Figure 4-30 shows the setup of the gateway for our example network. Once the gateway and the Ethernet adapters are configured, then you can install the Virtual I/O Server.

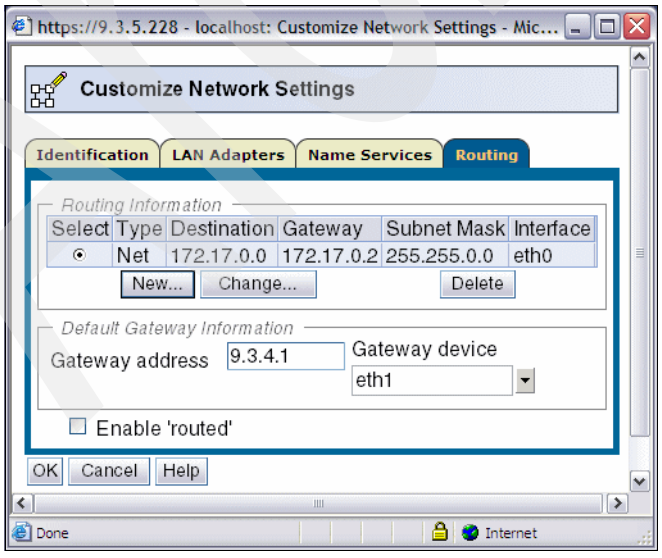

Figure 4-30 Define Ethernet gateway

The network administrator also must provide a fixed IP address for an Ethernet port owned by the Virtual I/O Server partition. This port and IP address will be the one used to install the Virtual I/O Server image software.

3. Another important step is to make sure of the cabling of the disk drives in the system assigned to the Virtual I/O Server partition. For a Virtual I/O Server to find the DASD that you allocate to storage pools, to client partitions, and that are used for the install, you must have the internal drive bay cables attached to the correct ports on the disk IOA.

   For example, if using disks in a 7311_D20 expansion unit, the System Planning Tool Version 2 picks a port number on the disk IOA. You must have the cable attached to that port on the disk IOA or the deployment can have unpredictable results.

   The System View report has cabling instructions that show you exactly how to cable it.

Once the information is gathered you are ready to start the deployment on the HMC.

4.4.1 Starting deployment of system plan containing a Virtual I/O Server partition

Once you have loaded the Virtual I/O Server image to the HMC and the IP address of the Ethernet LAN port over which the Virtual I/O Server software will be installed, you are ready to start the deployment. Figure 4-31 shows an example of the make directory, mount, and OS_install instructions to load the Virtual I/O Server image onto an HMC disk.

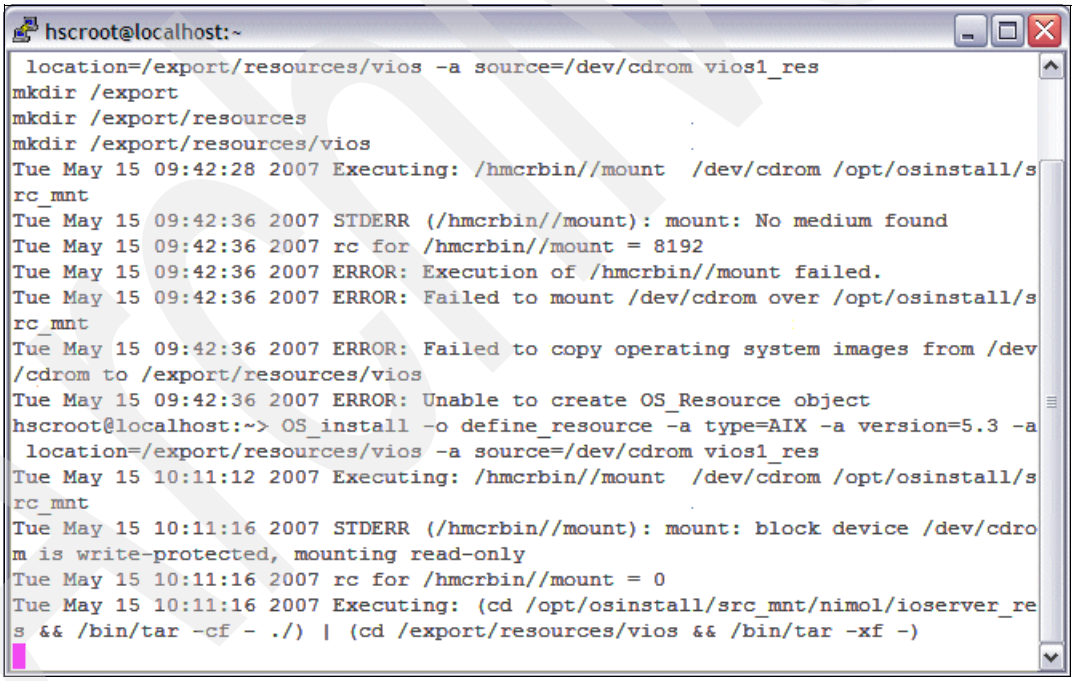

Figure 4-31 OS-install command result

Figure 4-32 is where we start the deployment process, first specifying important parameter values and then clicking the **Deploy** button to start the automated deployment process.

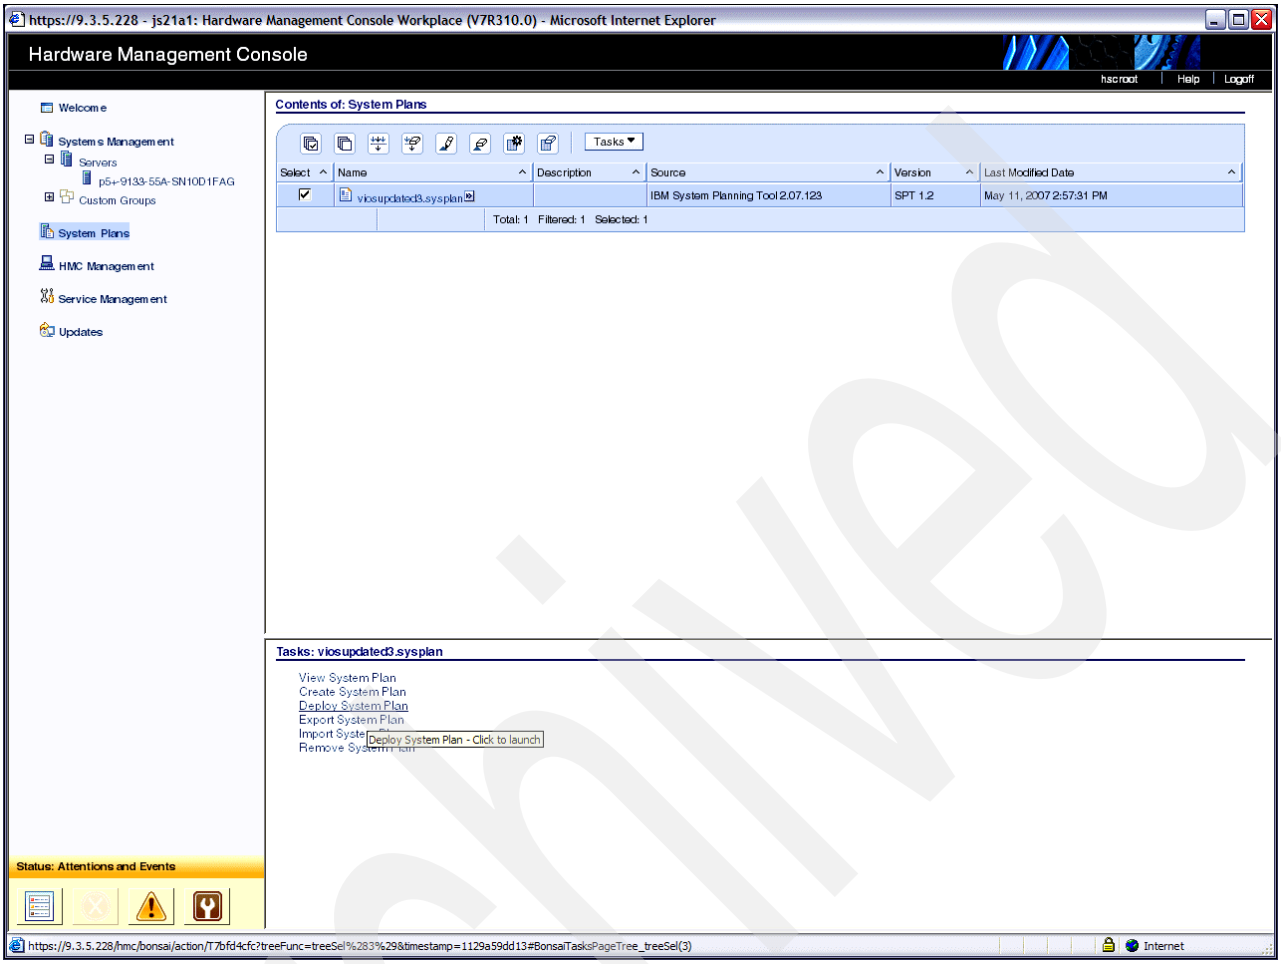

Figure 4-32 Chose the Virtual I/O Server to install

Most of the Virtual I/O Server deploy steps are the same as deploying any system plans. However, the are some important differences.

Figure 4-32 shows the a selected system plan. Once a plan has been selected, you can select **Deploy System Plan - Click to launch**, as shown in our figure.

Launching to the deployment process brings up the Welcome - Deploy System Plan Wizard window shown in Figure 4-33. Specify the system plan name and managed system. These values may already be filled in for you or they may be blank.

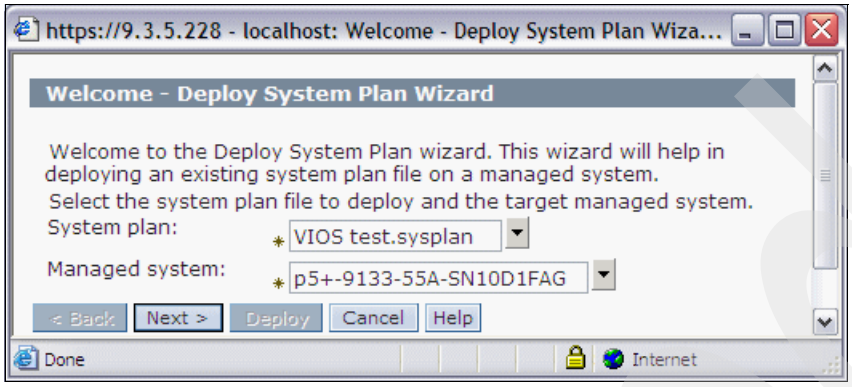

Figure 4-33 Select the system to install Virtual I/O Server on

After specifying your system plan and managed system, click **Next**. The system plan is validated for hardware and partition correspondence. Validation progress and messages are shown. Figure 4-34 shows a successful validation.

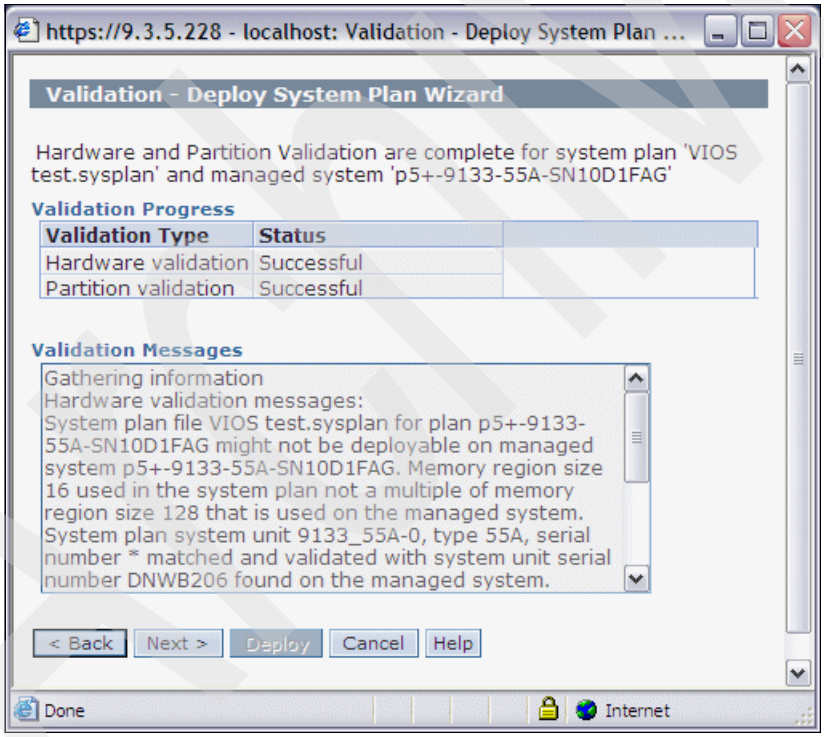

Figure 4-34 Sysplan validation

Click **Next** to continue.

On the next window, select the operating system to install. Select the Virtual I/O Server by clicking the radio button, as shown on Figure 4-35. Click **Next**.

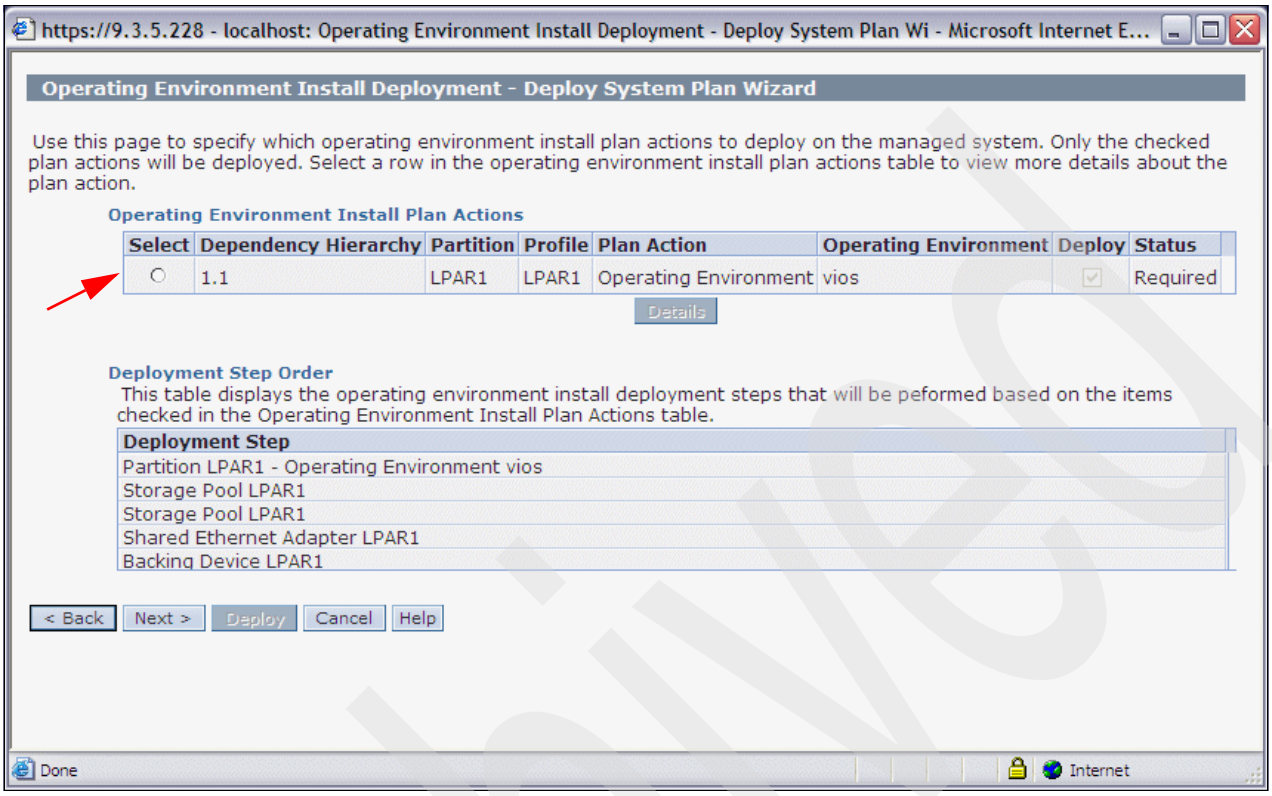

Figure 4-35 Start of the deployment

You are presented with the partitions that are to be built. Figure 4-36 shows our example. You can select each partition and then click **Details** to see the details of your partition definition.

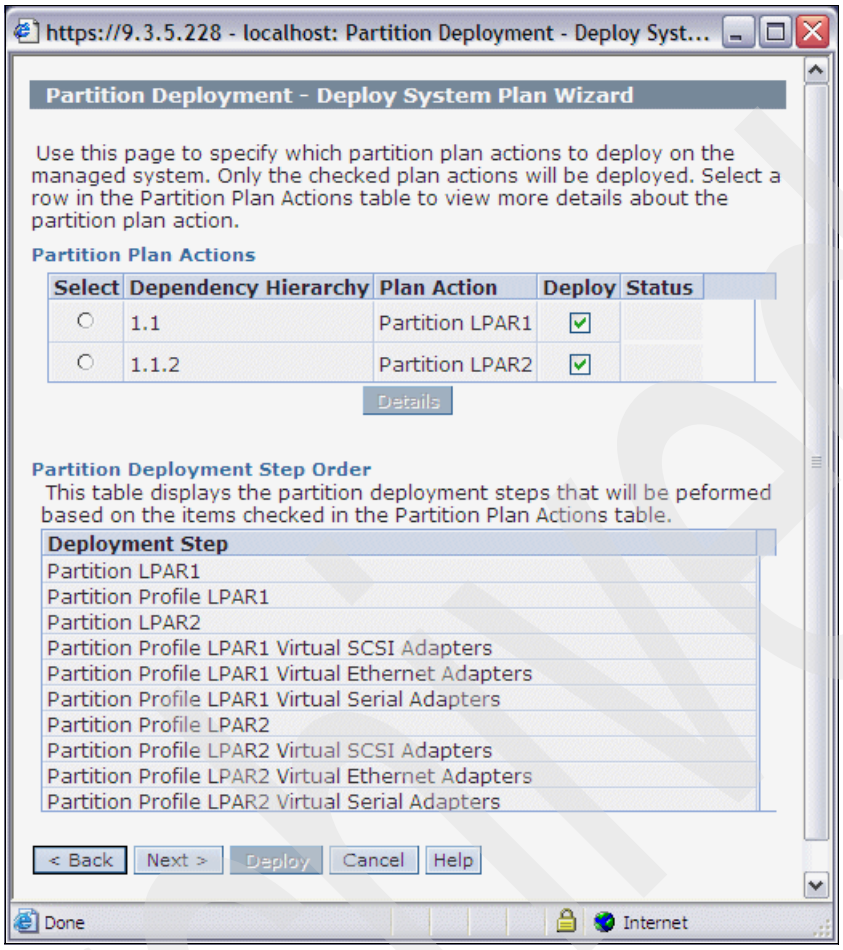

Figure 4-36 Partition build

Once you are satisfied with your partition details, examine the Deployment step area to see the deployment steps that will be performed, based upon you currently shown partition plan actions shown above.

Click **Next** to continue.

The next window (Figure 4-37) is one that is different from deployment of non-Virtual I/O Server partitions. This window is used to gather information for a Virtual I/O Server partition install. Figure 4-37 shows two white X on red background icons next to text that indicates that additional definitions have to be specified.

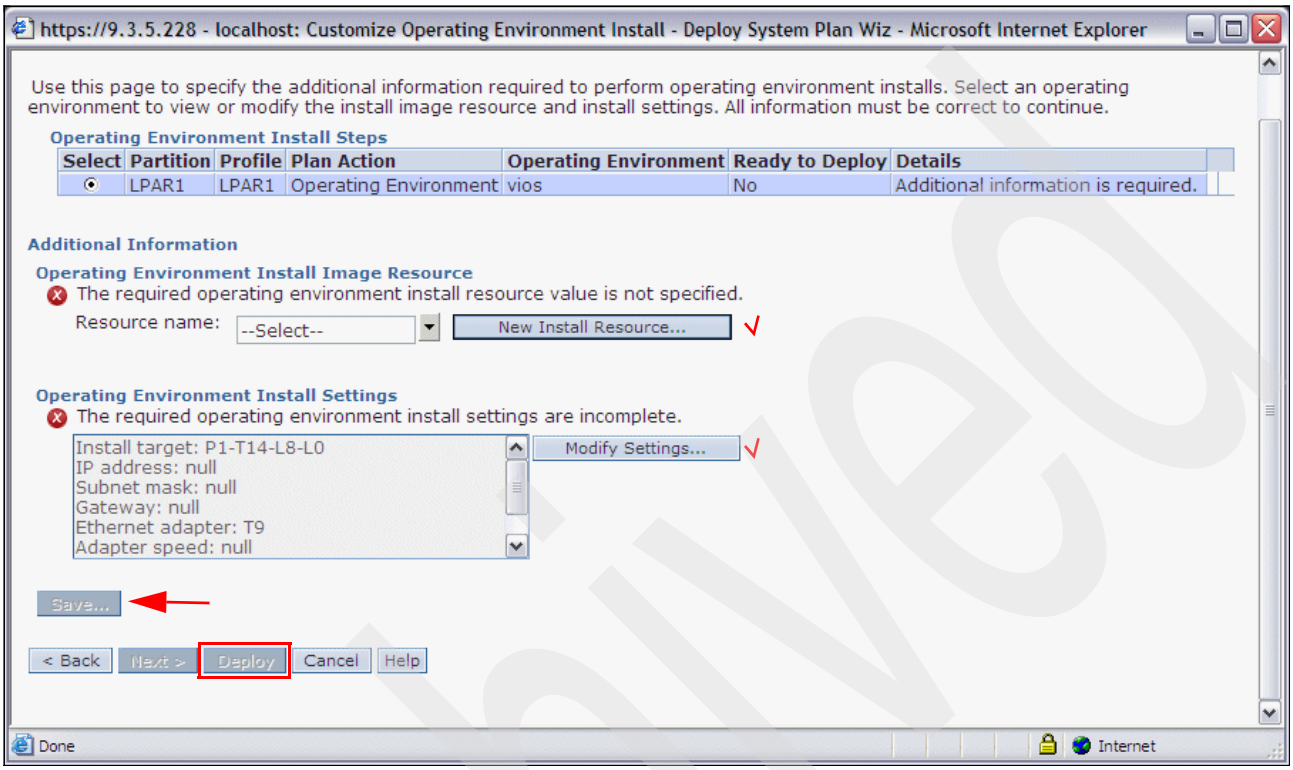

Figure 4-37 Unique Virtual I/O Server deployment information

On the first attempt to deploy a Virtual I/O Server partition you must select a resource name by first clicking **New Install Resource**.

Important: Figure 4-37 is referenced several times in succeeding topics of this chapter. However, the error indications shown here, identified by white x on red background icons, will disappear as the errors shown in Figure 4-37 are resolved in these succeeding topics.

Figure 4-38 shows the window presented by clicking **New Install Resource**.

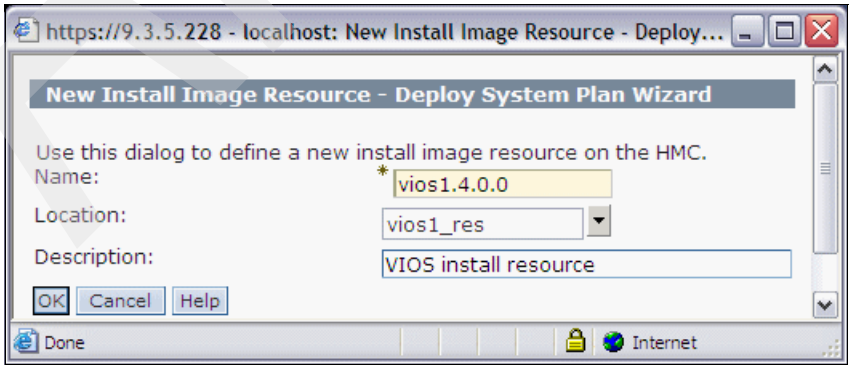

Figure 4-38 Set resource name

At the time this book was published you must enter `vios1.4.0.0` for the name field.

You also must enter the HMC disk directory name (containing the Virtual I/O Server image) in the Location field, which you specified back in 1 on page 130, under 4.4, “Deployment of Virtual I/O Server on HMC” on page 129.

If you have not completed the first step of doing the OS_install command, the Location field would be blank. You would have to go back to the first step and start over. When the name is entered correctly, click **OK** to return back to Figure 4-37 on page 139.

Now you must complete the network information needed to install of the Virtual I/O Server code, as indicated by the white X on red background icon under Operating Environment Install Settings.

Click **Modify Settings**, which leads to a window such as the one shown in Figure 4-39, which initially contains some information you specified previously for your Virtual I/O Server partition, but is also missing some parameter values.

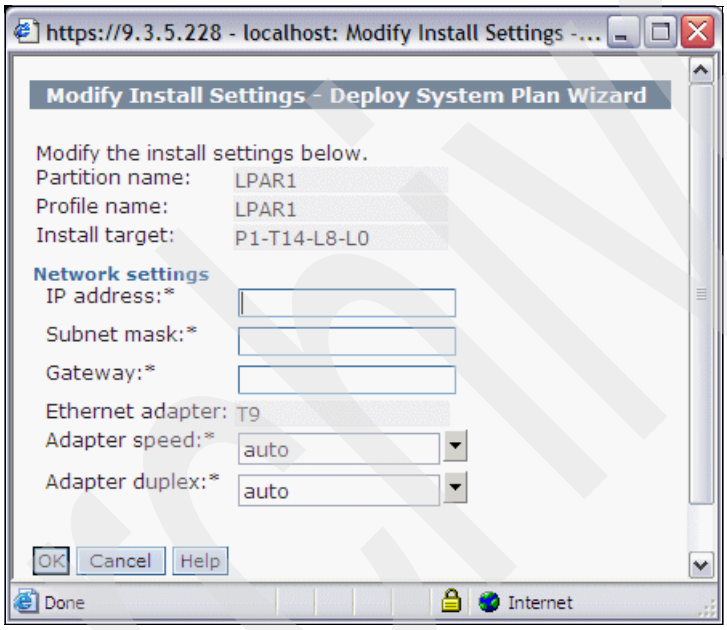

Figure 4-39 Enter network information

In Figure 4-40, we have entered the IP address, subnet mask, and gateway parameter values, as previously discussed in Figure 4-29 on page 133 and Figure 4-30 on page 133.

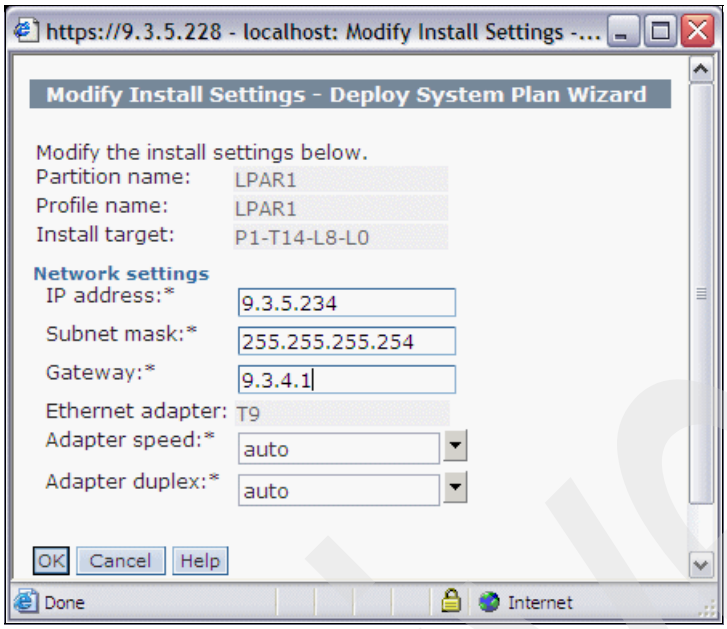

Figure 4-40 Filled in information

Click **OK**, which returns us to a window (not shown in this book) similar to the one shown in Figure 4-37 on page 139, but with the missing information now completed and the error messages no longer showing.

On that window click **Save**. Figure 4-41 shows the window to save the updated install information in our system plan file.

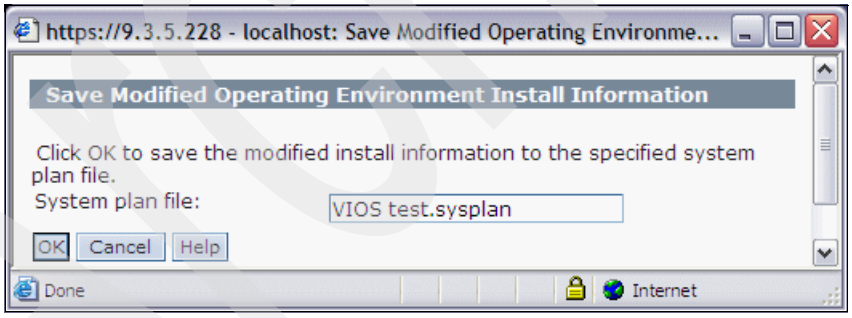

Figure 4-41 Save changes

Enter or verify your system plan name to be saved. Click **OK**.

This returns again to the window (not shown) similar to the one in Figure 4-37 on page 139, but with no error messages.

Once the information is entered and saved, you are ready to start the deploy of the Virtual I/O Server install.

You can start the deployment from the window show in Figure 4-37 on page 139.

Alternatively, you may have previously attempted a Virtual I/O Server deployment and failed for some reason. In that case, you have probably already validly entered the install image and

IP addressing information, so you could start the deployment process using the window shown in Figure 4-35 on page 137, which is repeated here in Figure 4-42.

In any case, click the **Deploy** button. The automated deployment process begins.

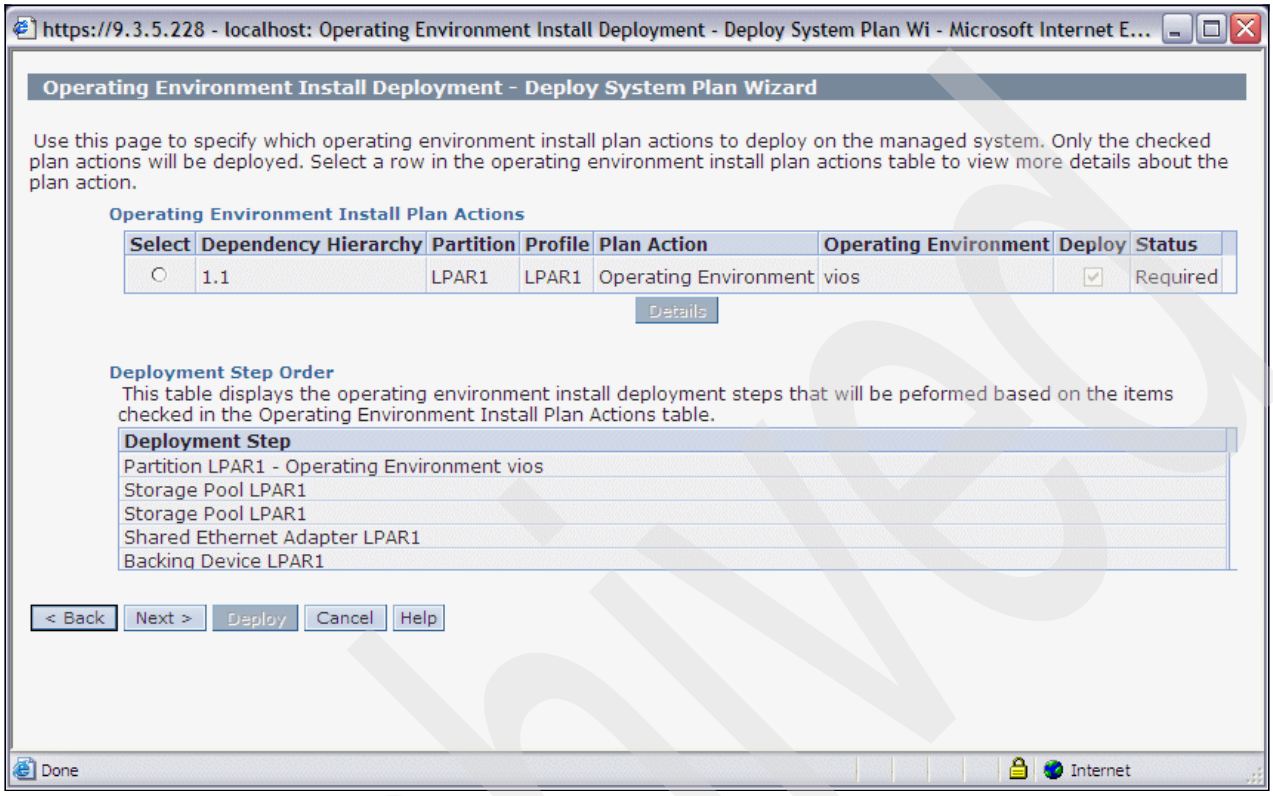

Figure 4-42 Select the install resource

Figure 4-43 shows the initial stages of the system plan. Remember that there could be partitions in addition to a Virtual I/O Server partition within the same system plan.

You see two areas of information:

- Deploy Progress area
- Messages area

The window refreshes approximately every 10 seconds with the newest step status and messages appearing at the top of each area. That is, newest at the top with older status and messages in the lower portion of each area.

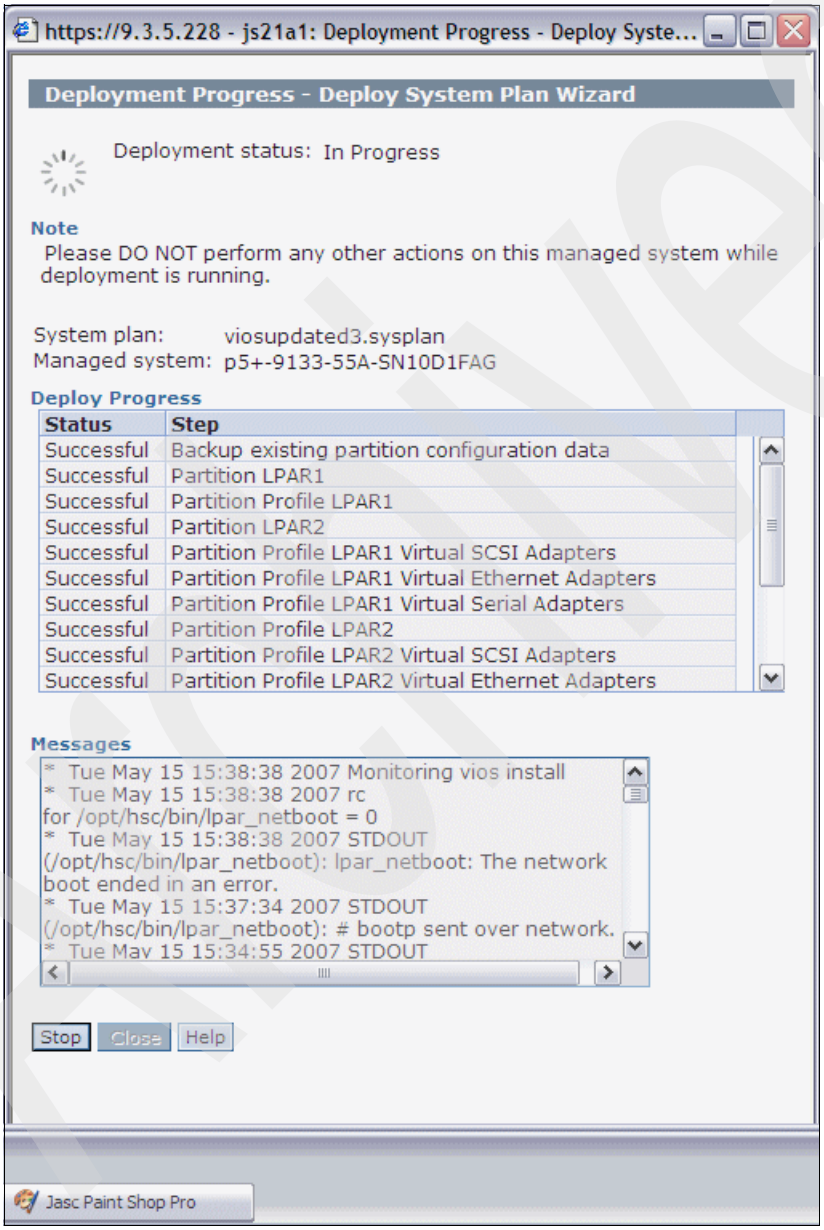

Figure 4-43 Starting the deploy process in log page

The next steps of the automated deployment are the same for all partitions until the step when installation of the Virtual I/O Server software is being performed. Figure 4-44 shows some messages indicating that the net booting of the Virtual I/O Server partition is being done.

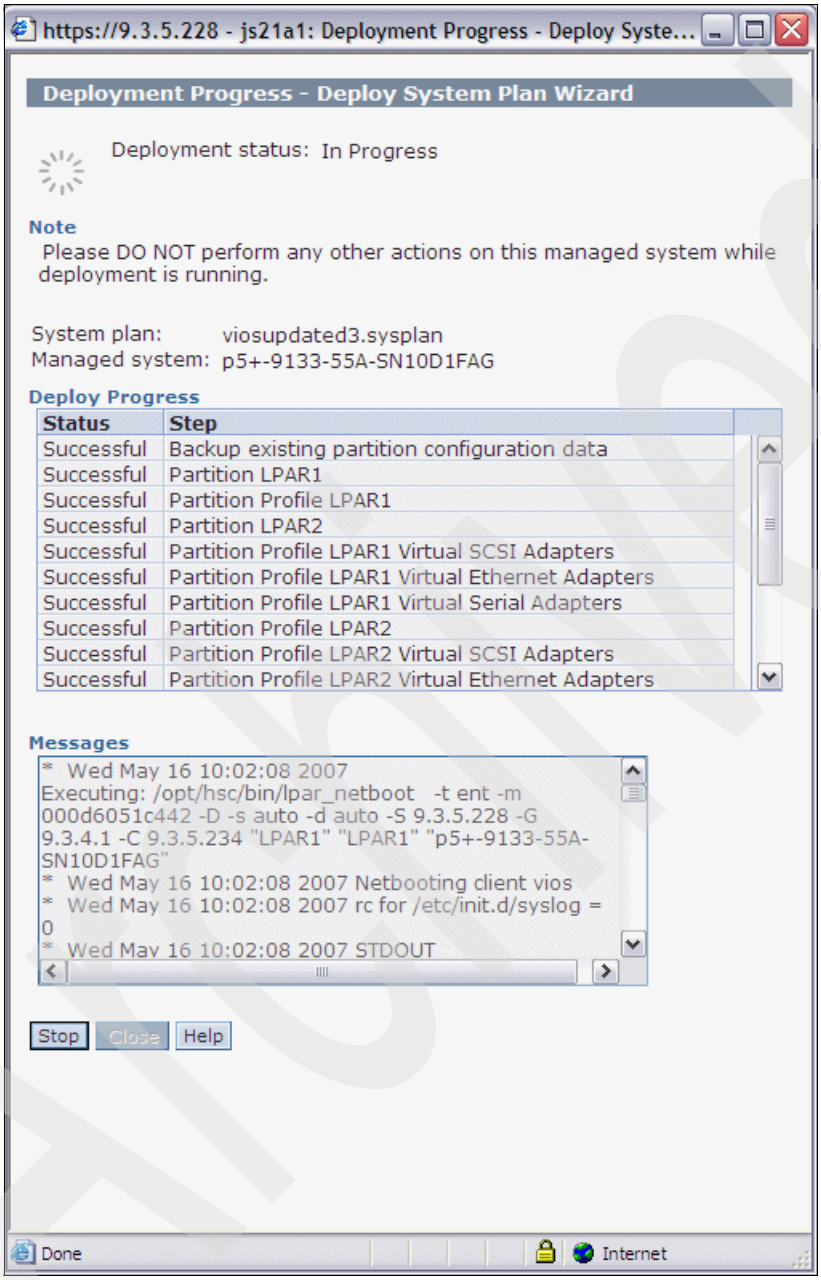

Figure 4-44 Log page of the progress of deploy

At this time there is a wait for the Virtual I/O Server system to install the root portion of the operating system and then it does a restart before the system is installed. This installation is from the Virtual I/O Server image that was loaded as a first step. A Network Installation Management (NIM) of the Virtual I/O Server takes place.

The source matches the name from the HMC image and the target is the assigned IP address given previously. Figure 4-45 shows an updated progress view of installing a Virtual I/O Server.

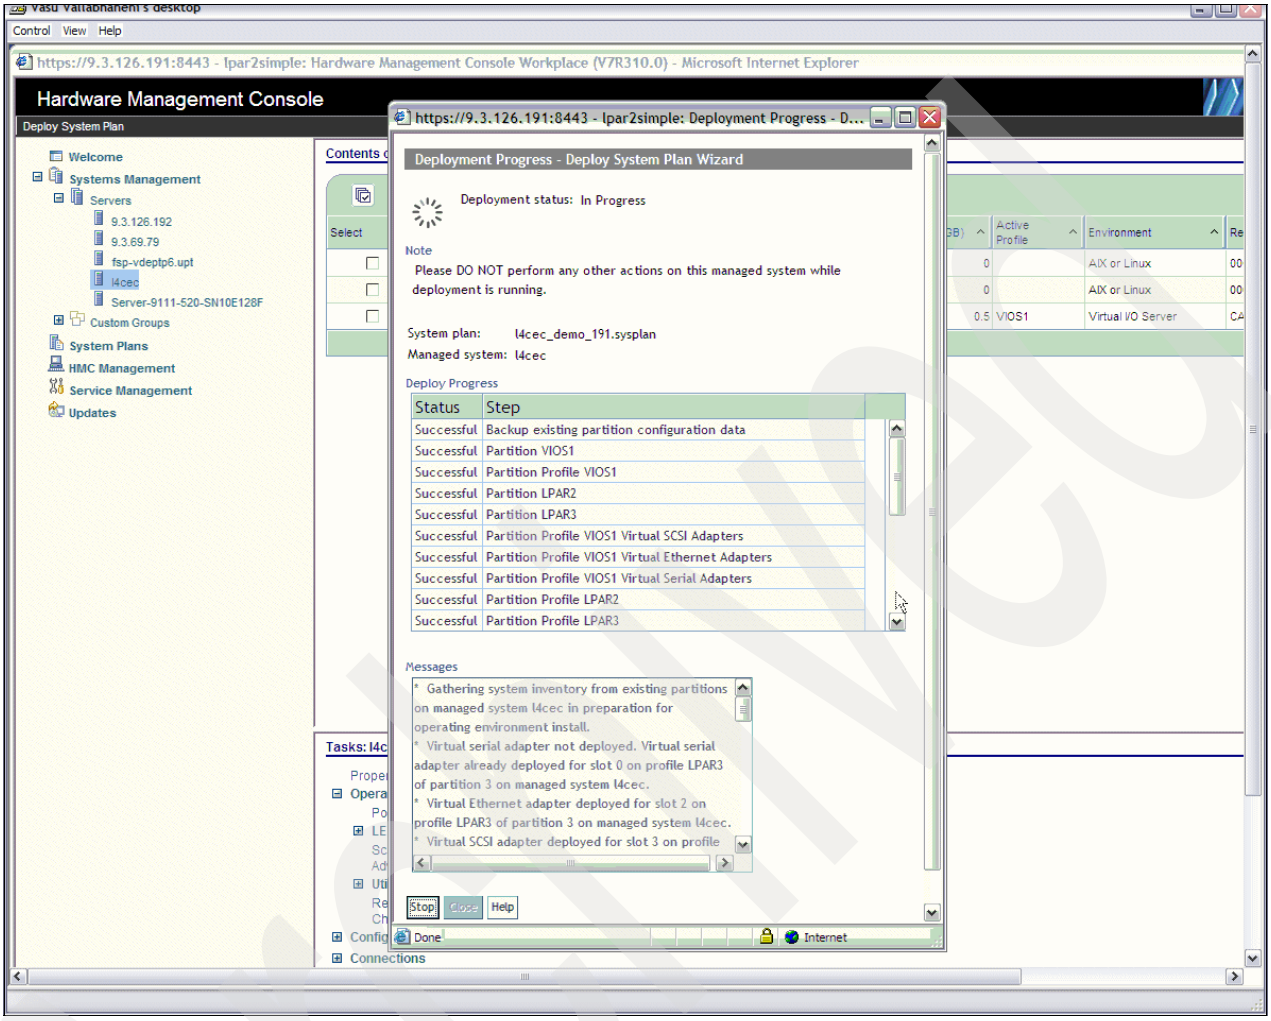

Figure 4-45 Start of Virtual I/O Server install

Once the Virtual I/O Server software has been installed, the provisioning of its resources takes place. This includes creating the following according to your Virtual I/O Server partition's configuration settings for:

- Shared Ethernet Adapter
- EtherChannel
- Storage pools
- Backing devices (for example, CD/DVD device, physical disks, SAN-attached logical volumes)

Once all of the provisioning has completed for that specific Virtual I/O Server partition, a message appears indicating this. If more than one Virtual I/O Server partition is to be created, the deployment of the next Virtual I/O Server partition is started. The Virtual I/O Server partitions are done in sequence.

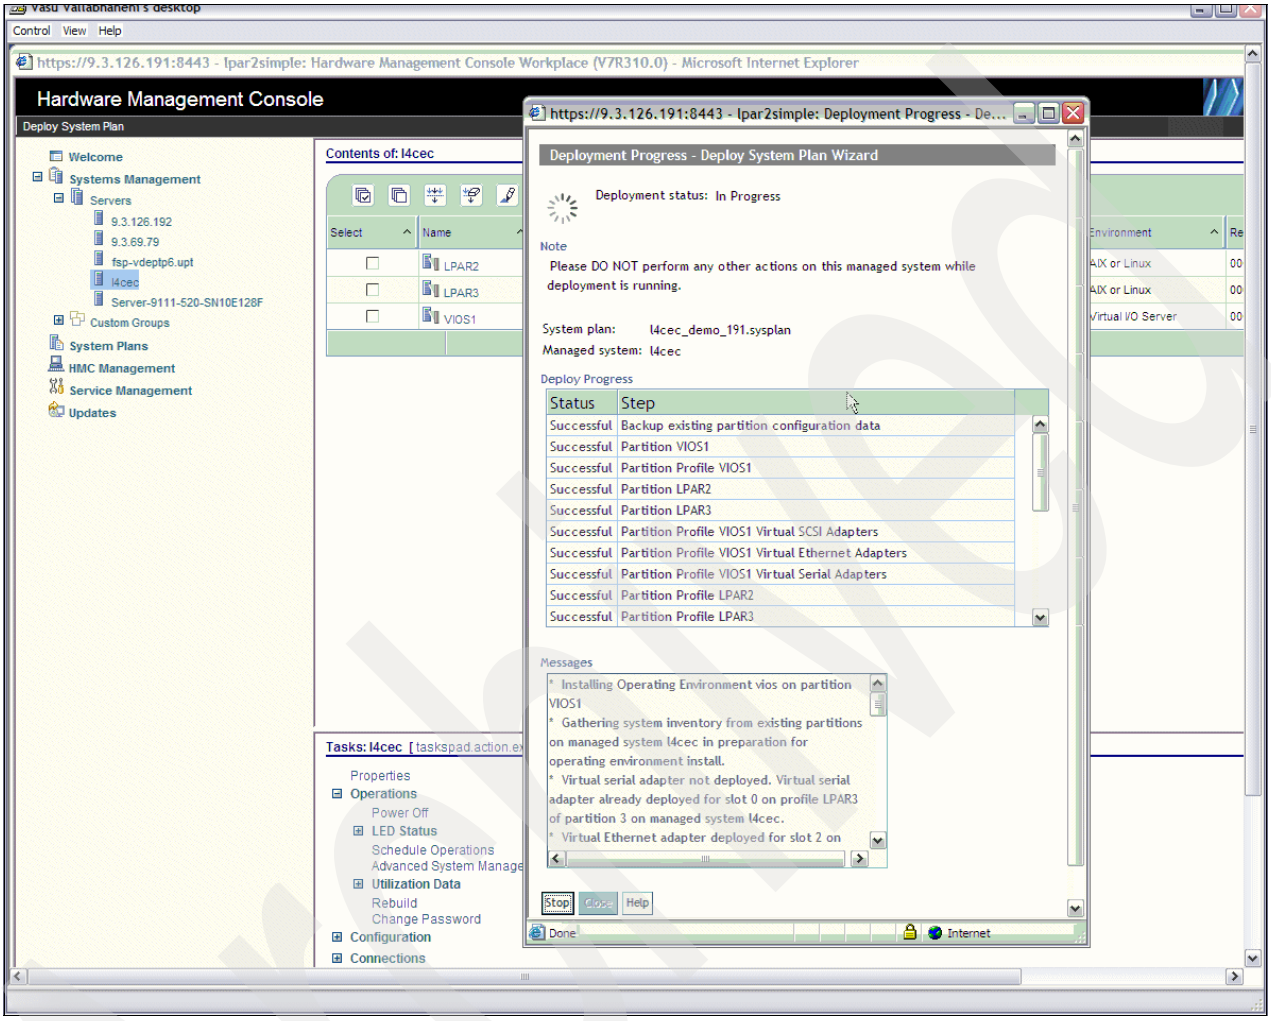

Figure 4-46 Completion of the Virtual I/O Server deploy

This ends our coverage of defining and deploying a system plan with a Virtual I/O Server partition defined.

4.5 Deployment on Integrated Virtualization Manager (IVM)

IVM is a specialized Virtual I/O Server partition that, by its definition, owns all system resources that may be partitioned. Essentially, IVM on a system is mutually exclusive with managing the system with an HMC. Both cannot be in effect at the same time.

IVM provides a very robust set of partitioning functions for AIX and Linux partitions, but this function set is a subset of HMC capabilities. Providing more details on IVM capabilities and comparing partitioning capabilities between IVM and an HMC is beyond the scope of this publication.

However, you should know that the System Planning Tool Version 2 supports planning and configuring and ordering a system managed by IVM.

IVM installs Virtual I/O Server as the first partition instead of using the HMC. This partition controls all of the hardware. It provides Virtual SCSI and virtual Ethernet to the other partitions. IVM has a Web interface to its management capabilities. IBM Systems Hardware Information Center and *Integrated Virtualization Manager on IBM System p5*, REDP-4061, provide additional information.

5

System Planning Tool Version 2 and Workload Estimator

This chapter provides instructions to the interface of SPT V2 and Workload Estimator (WLE).

With the interface from SPT V2 you can link between SPT and WLE to:

- Add a new workload to an existing partition.
- Add (create) a new partition with one or more new workloads.
- Consolidate one or more systems or partitions onto a different system with a new partition configuration.
- Importing PM for System i (new browser interface used in this book now termed IBM Web Sales Tool) allows you to bring in current performance information to create partition requirements.

5.1 Sample use of SPT V2 and WLE from performance data

The following examples go through a series of windows and discuss options that include bringing summarized performance data from the IBM Web Sales Tool (software and services, formerly known by several other names, such as PM/400, PM eServer for iSeries, and recently as IBM Performance Management (PM) for the IBM i5 and IBM p5) into an SPT V2 system plan.

As summarized in our example, each i5/OS partition had the *PM software* installed and active.

Each i5/OS partition has the software installed to collect performance data and send that data to IBM on a specified time frame. By using IBM Web Sales Tool you can select the system or partitions to bring down to WLE. Save the data in a directory on your PC. This data can be from a partitioned system or from multiple systems to be consolidated. The interface is started from SPT V2 installed on a workstation.

If you start in WLE first, the option to bring the result from WLE into SPT V2 is not available.

Figure 5-1 shows the first SPT V2 window with “Based on existing performance data” selected.

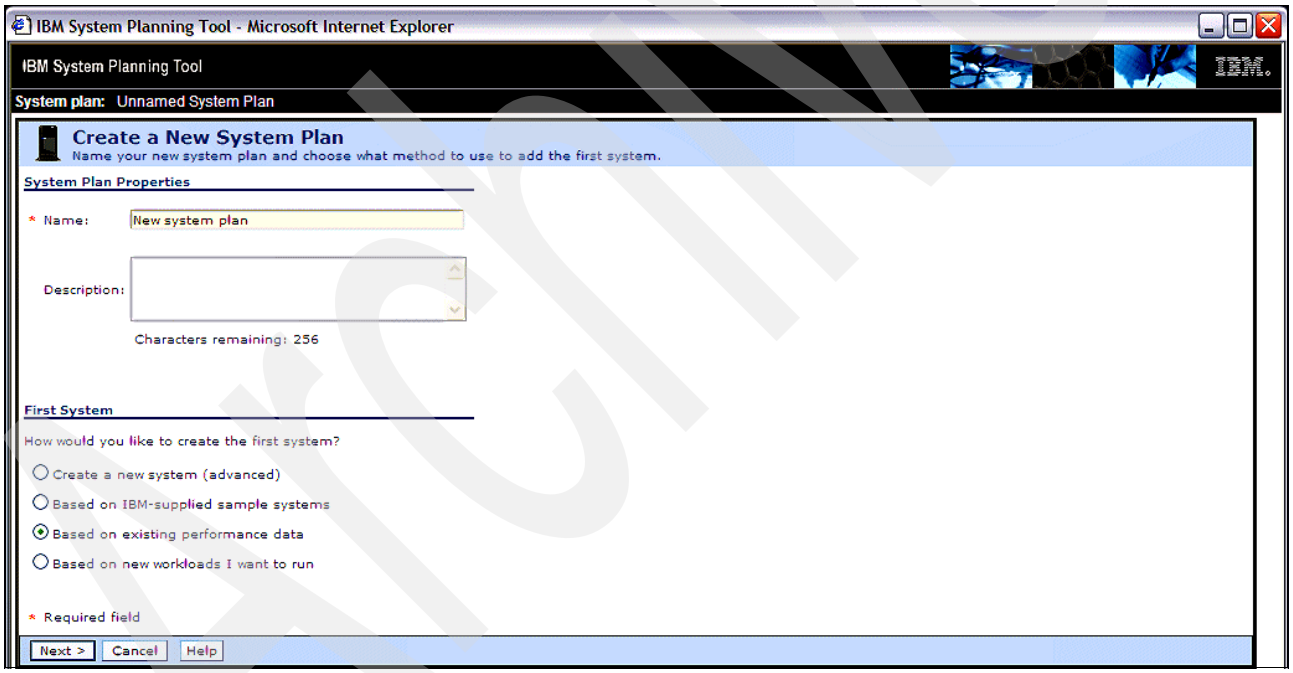

Figure 5-1 Chose to start WLE for collected performance data

The steps are:

1. You need to load the PMdata into WLE. The window to start loading the data looks the same if you start in WLE except there is no mention of the SPT V2. Figure 5-2 shows the interface window.

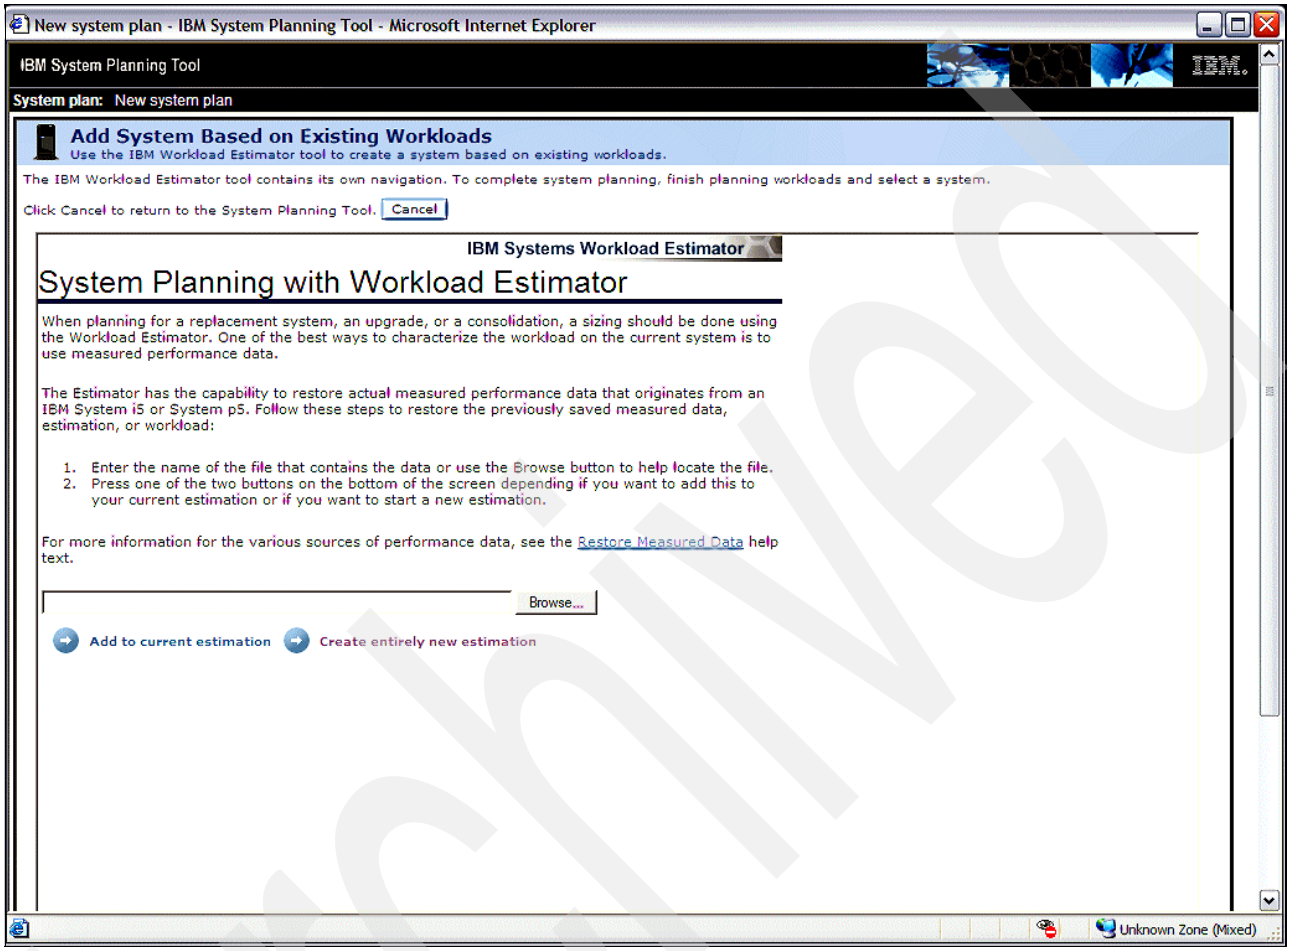

Figure 5-2 Import the performance data from the PM tool

2. Click **Browse** and a selection of files that are on your PC are presented to you. Figure 5-3 shows the collected files of performance data we use in our examples. You can chose a system to consolidate the partition data. The partitions data should have the extension of a two-digit number appended to the end (00–99).

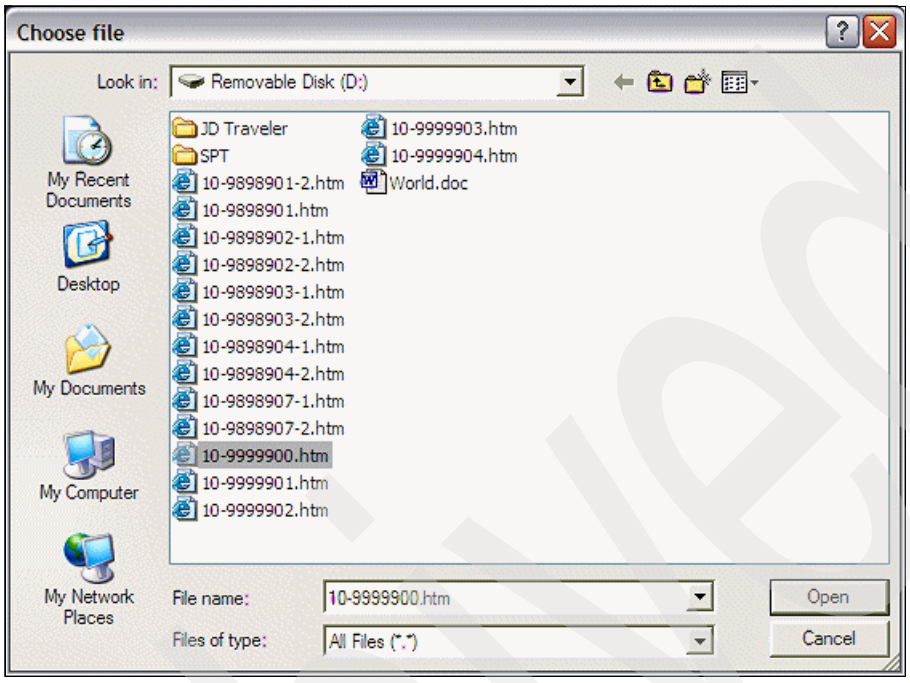

Figure 5-3 Browse for the desired performance data

If only one system is desired, you can continue in WLE. However, if more systems or partitions are desired, select the option to **Restore Saved Estimation** until you have completed adding all systems or partitions. Figure 5-4 show the window to select the option to add more systems.

- By default, if you import different systems or a partition of a system they will show as one file per system with a partition and a workload. To present a single system with multiple partitions with their workload requires moving the partitions to the first system.
- With the new IBM Web Sales Tool the systems and partitions go directly to WLE. The PMiSeries tool would place the information on your PC then go to WLE. With the new system you would have to save your workload in WLE to be able to use SPT to use the option *work with collected data*.

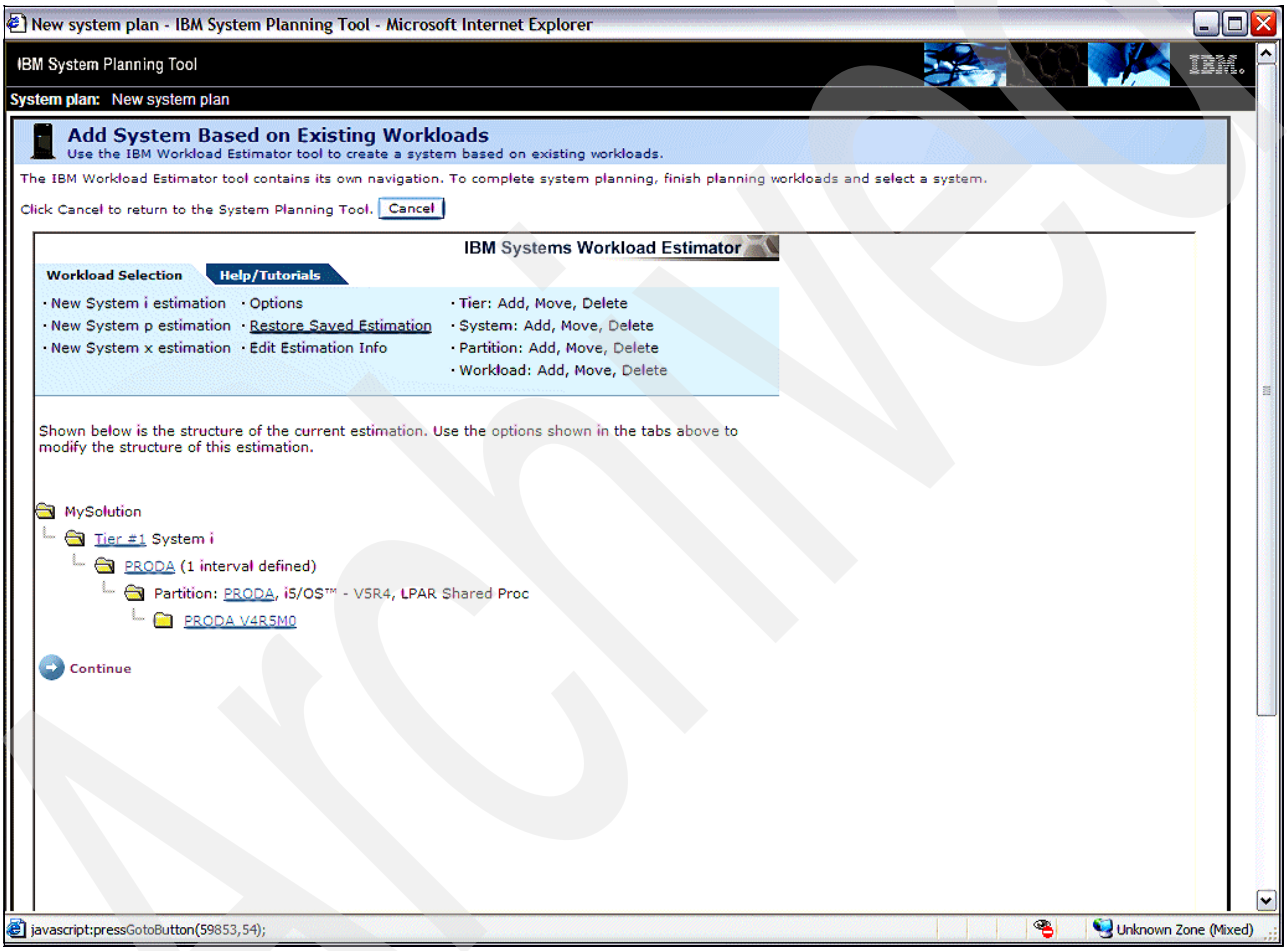

Figure 5-4 Restore saved estimates to add more performance data

3. After all desired systems or partitions are added to WLE you can click **Continue**. However, if you want the system to be in a single system at the end, you have to move each partition up to the first system. Select the option to move partitions. After moving the partitions make sure to click **Return**. On all the options to add, move, and delete remember to click **Return**.

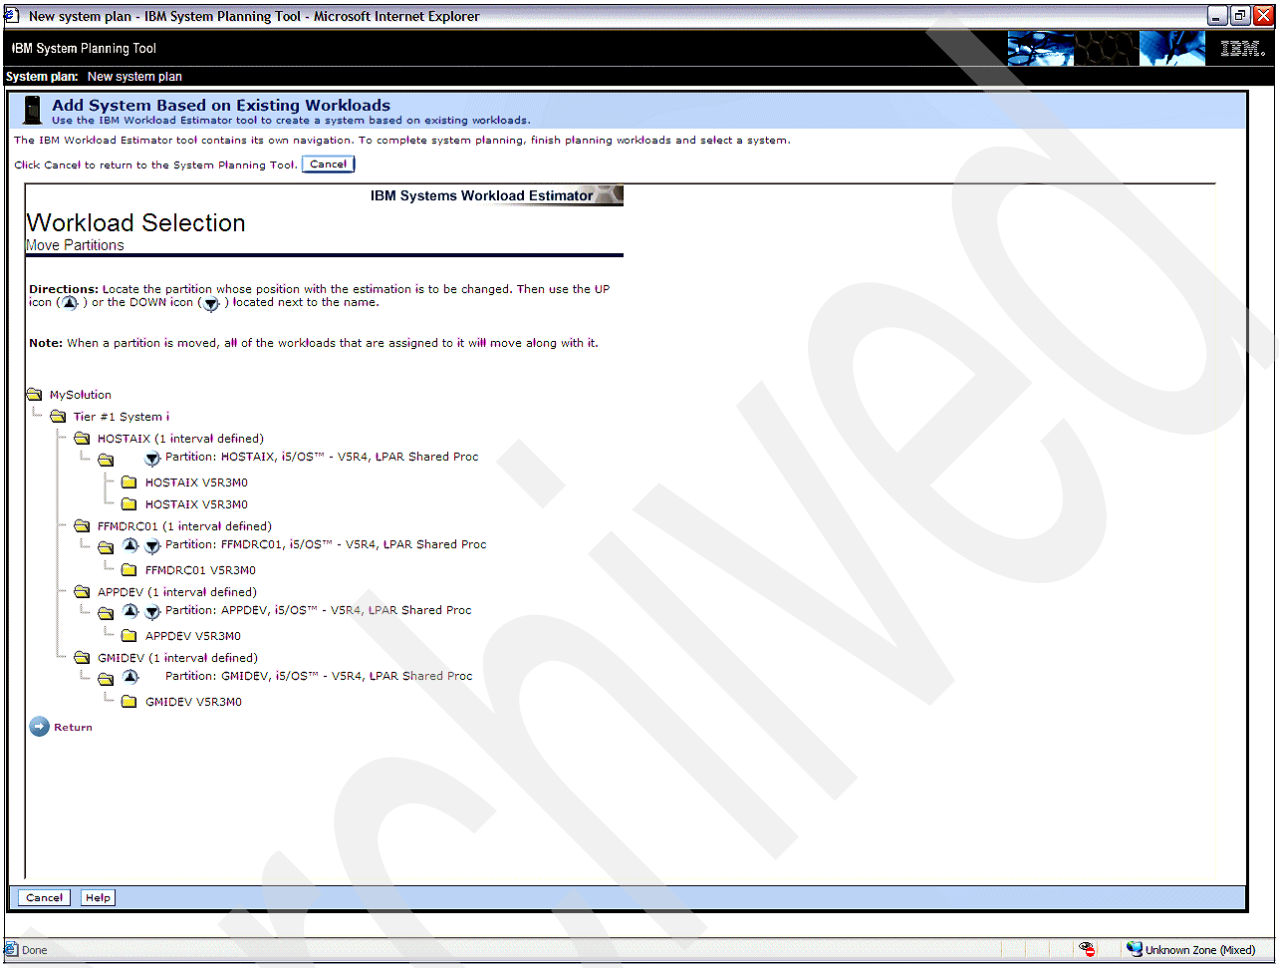

Figure 5-5 Partition information move to top system

4. In our example we have set up multiple partitions in a single system. Once the partitions have been moved, some of the original systems do not have any workload. These need to be deleted. Select the option to delete a system. Make sure that all systems without workload are deleted. Figure 5-6 shows systems that need to be deleted.

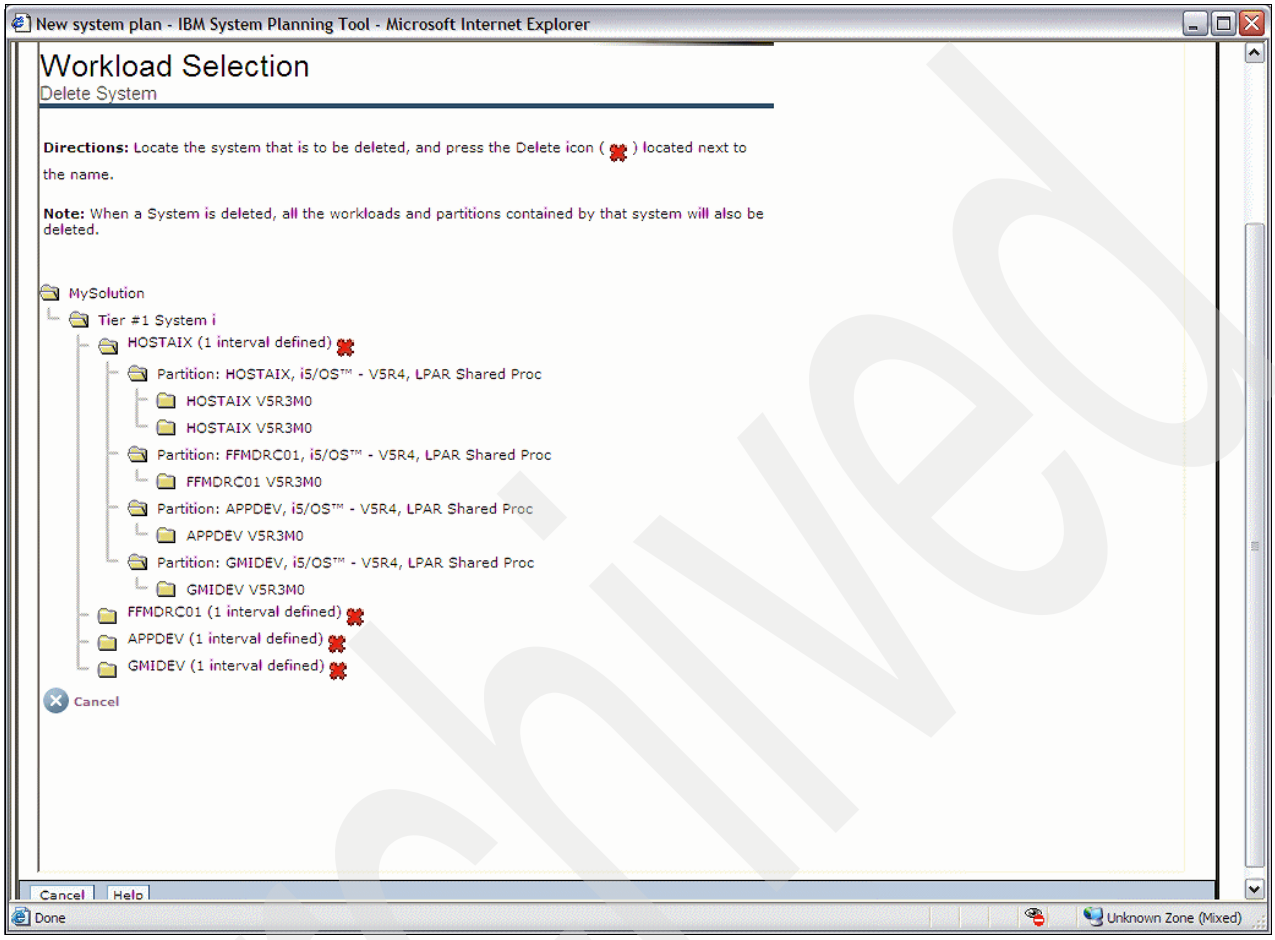

Figure 5-6 Delete excess systems

5. Next you accept or change the default parameter values that WLE uses to determine the upgrade for a system. Some options are:
   - OS version
   - DBCS support
   - Disk protection
   - Disk arm busy
   - Target processor utilization
   - Target 5250 utilization
   - Target LPAR CPU utilization
   - Disk full percent
   - Disk controller
   - Drive speed
   - Model family for upgrades

Figure 5-7 shows most of these options. You can click at the bottom of the window to change the options permanently or just use the changes during this analysis. Once selected, you are presented with the details of the collected data.

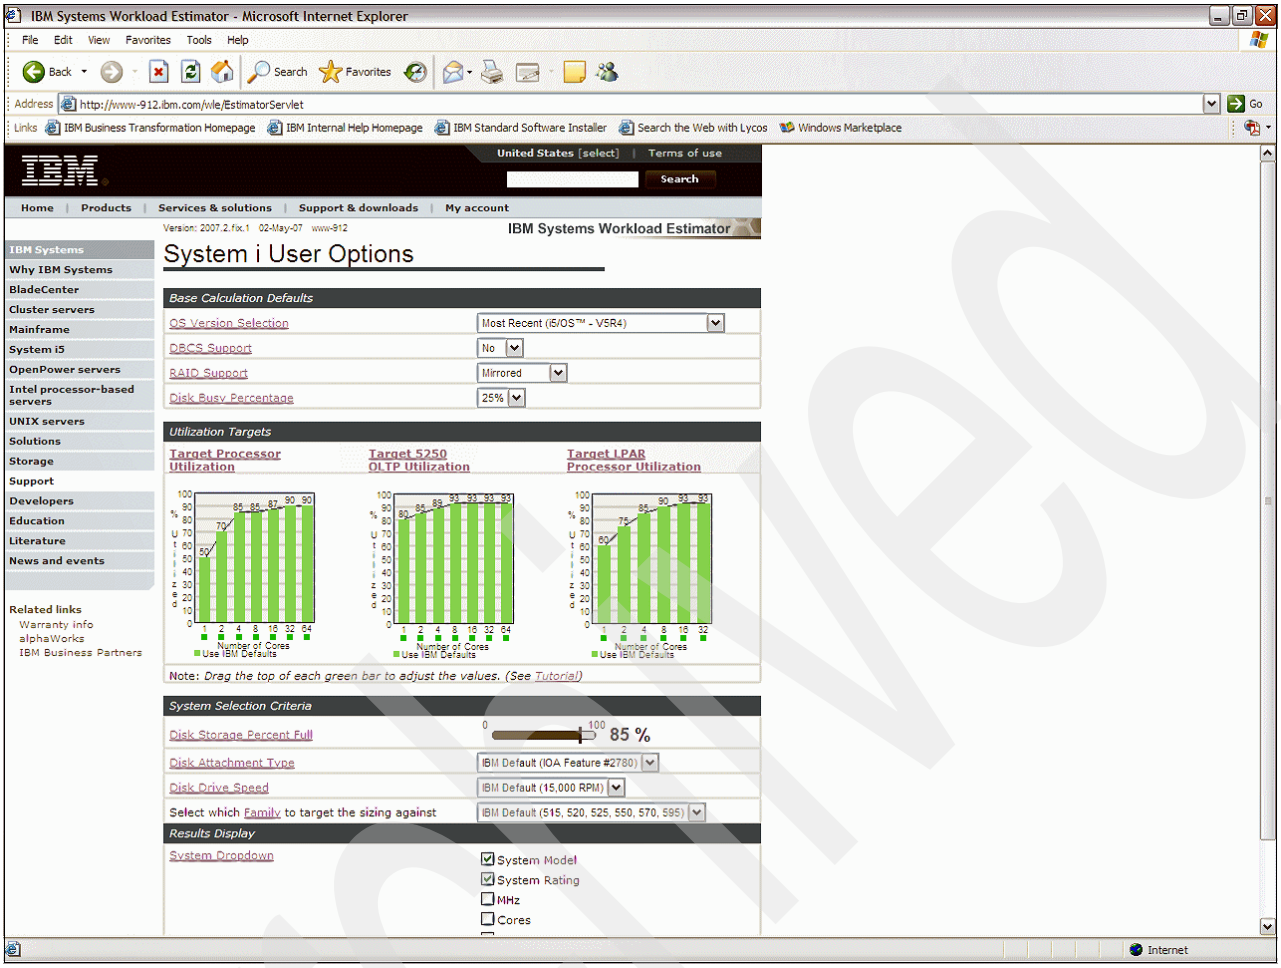

Figure 5-7 User options

6. The detail data of each partition is shown. This data can be modified to reflect a heavier workload on CPU, interactive, memory, and disk. Figure 5-8 shows the fields that can be changed. The data point reflects the current trend data on a graph. If the trend is negative, the field shows a negative value. This can be a problem in the growth system that is projected. You need to understand whether workload or disk are being removed.

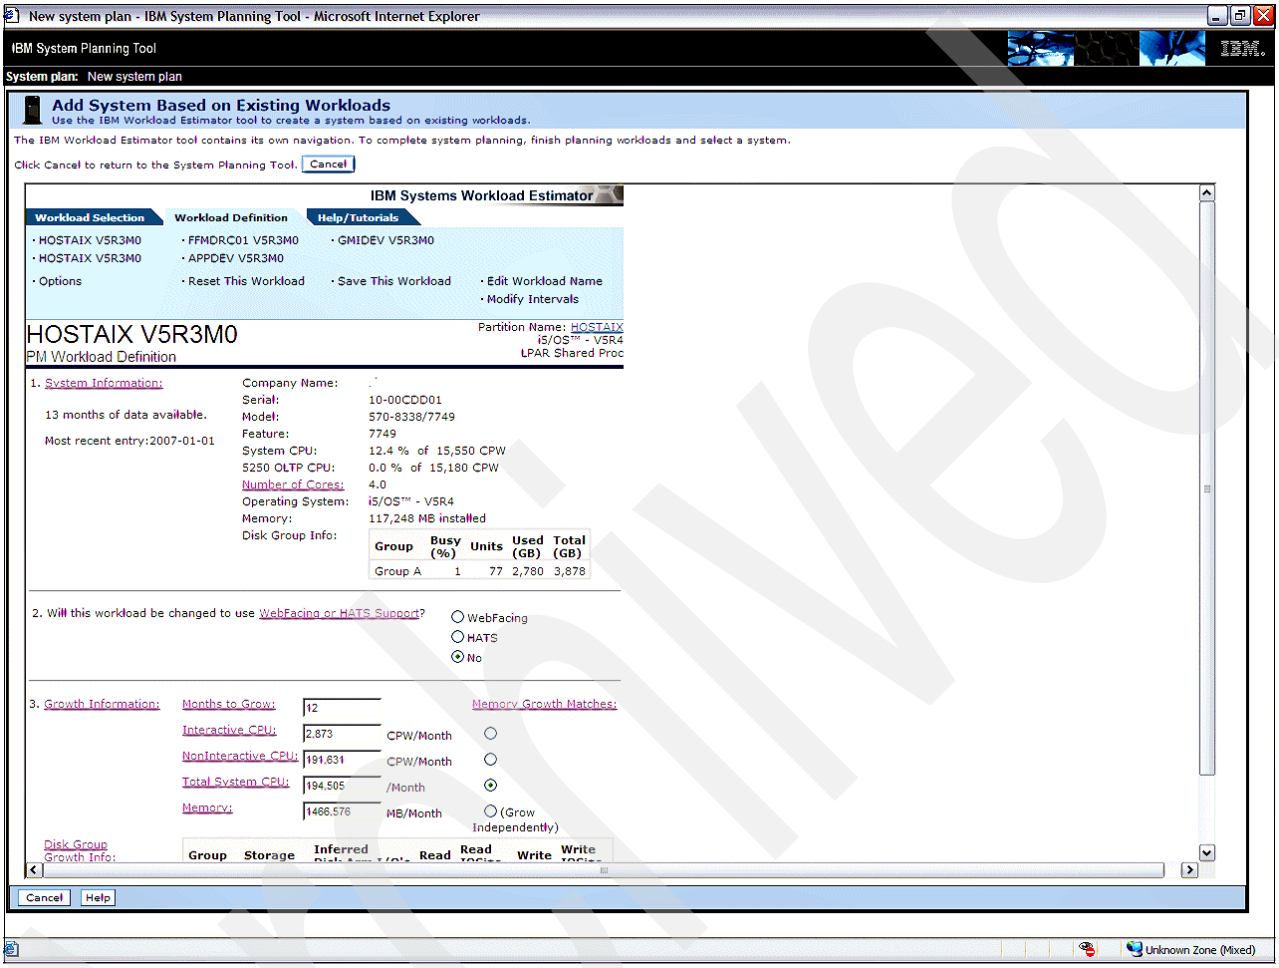

Figure 5-8 Partition data

In our example, several PM for System i workloads have been imported and we have processed WLE screens up through the last workload (HOSTAIX V5R3M0) window. At the bottom of this WLE window there are four buttons available (not shown in Figure 5-8):

- Back
- Advanced Growth Options
- Detailed Data
- Continue

You can click **Continue** to proceed to the estimator results window showing an immediate solution and a growth solution.

Before doing that, however, we show you an example of WLE's Advanced Growth Options window (Figure 5-9) and Detailed Data window (Figure 5-10 on page 159).

The data point reflects the current trend data on a graph. If the trend is negative the fields show a negative value. This can be a problem in the growth system that is projected. You need to understand whether workload or disk are being removed. You see the relationship of the data in Figure 5-8 on page 157 and the data in Figure 5-9, which can help to make a determination if the data needs to be changed. To return click **Continue**.

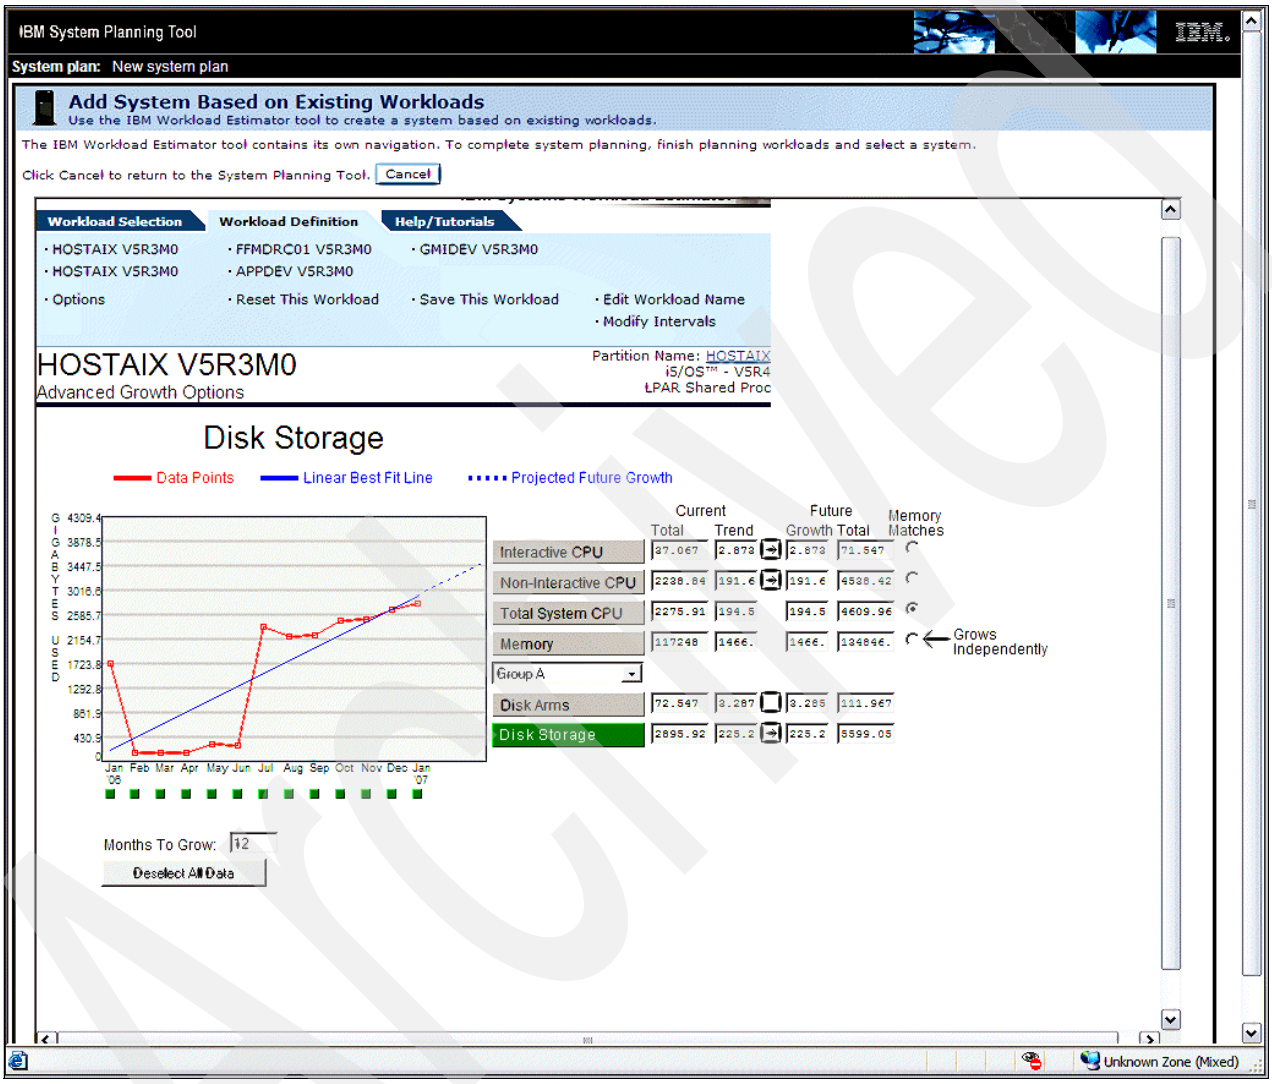

Figure 5-9 Data graph

7. You can also look at the detailed data to see the collected data days. You can change the disk controller, disk protection, and disk speed on the collected data. If you change the information you have to select to enable the change. Always use Return or Continue to return to the previous window or to continue to the next window. Figure 5-10 shows the window for the detailed collection periods.

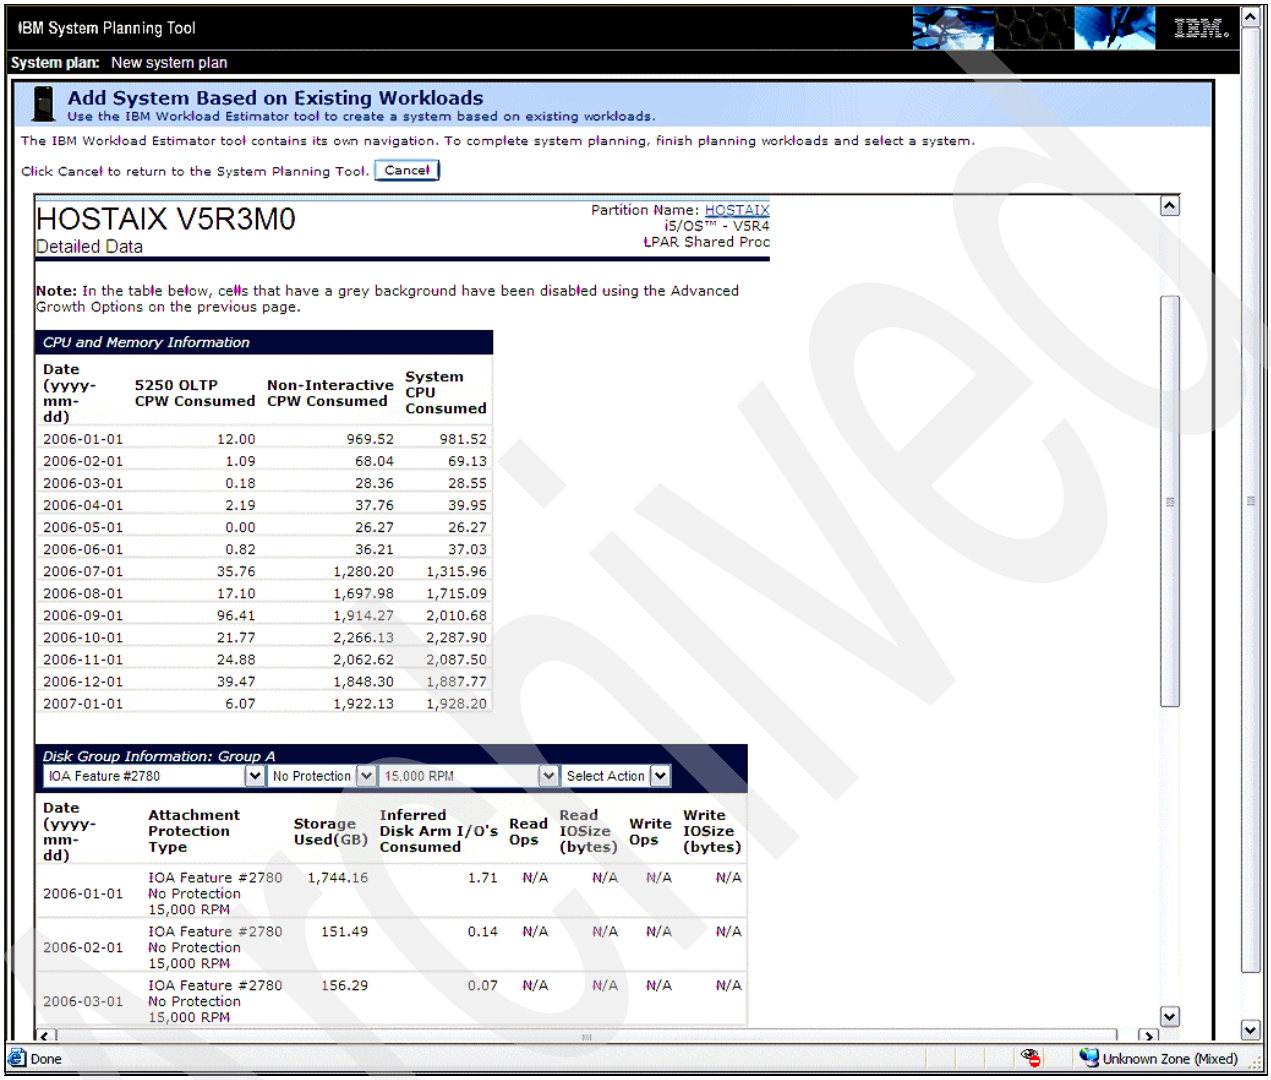

Figure 5-10 Detail partition data

8. You can select the upgrade path for WLE to find the best fit for the upgrade. Figure 5-11 shows the selection options.

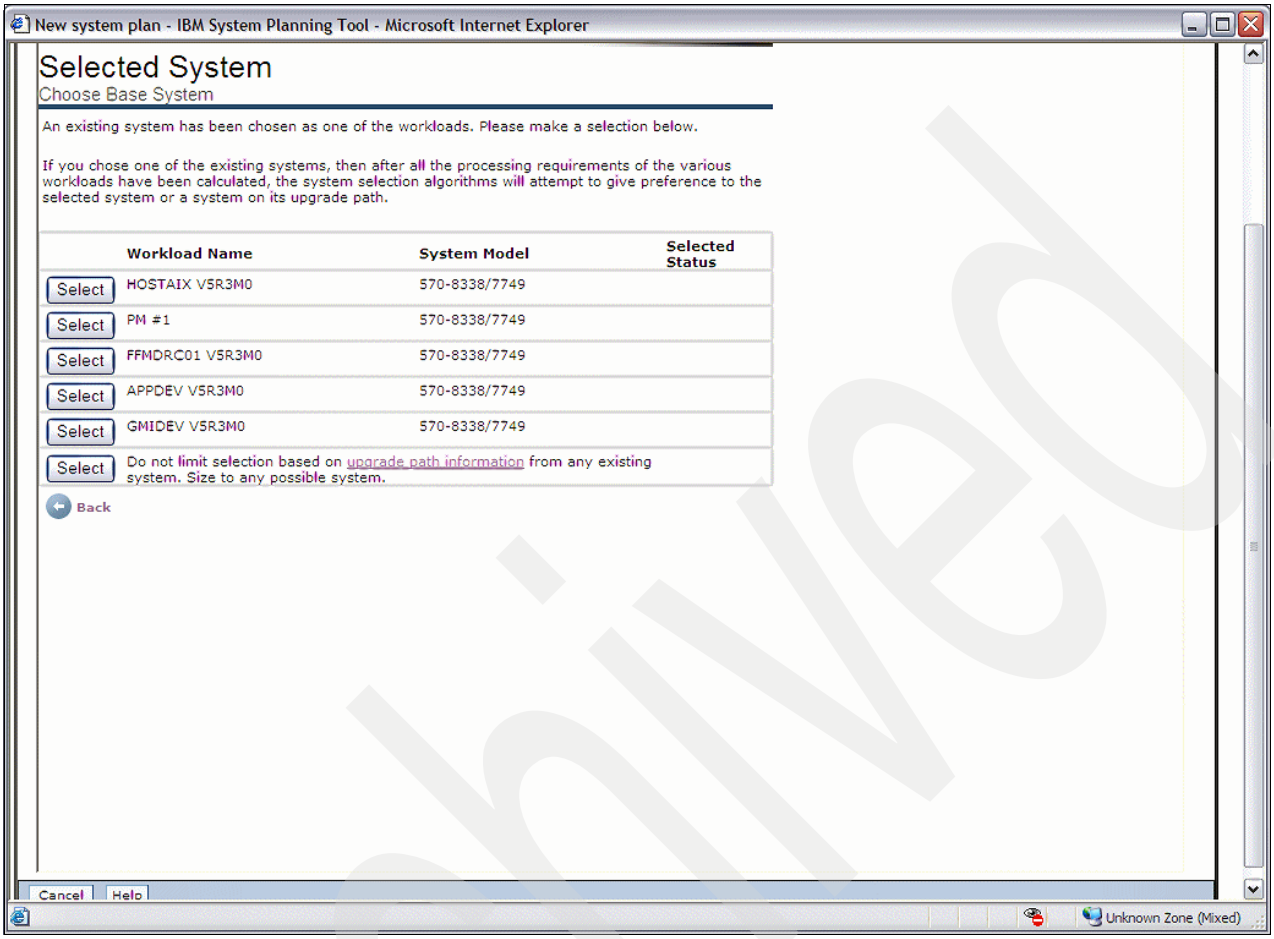

Figure 5-11 Select system for upgrade

9. If you click Back on this window you will go back to the detail data window for additional changes. When you click Select, WLE selects a system for the immediate upgrade and one for a growth upgrade. See Figure 5-12 for the selection. On this window you can create a PDF report of the selected system, save the workloads, go back to change a workload, or chose a different model for the upgrade.

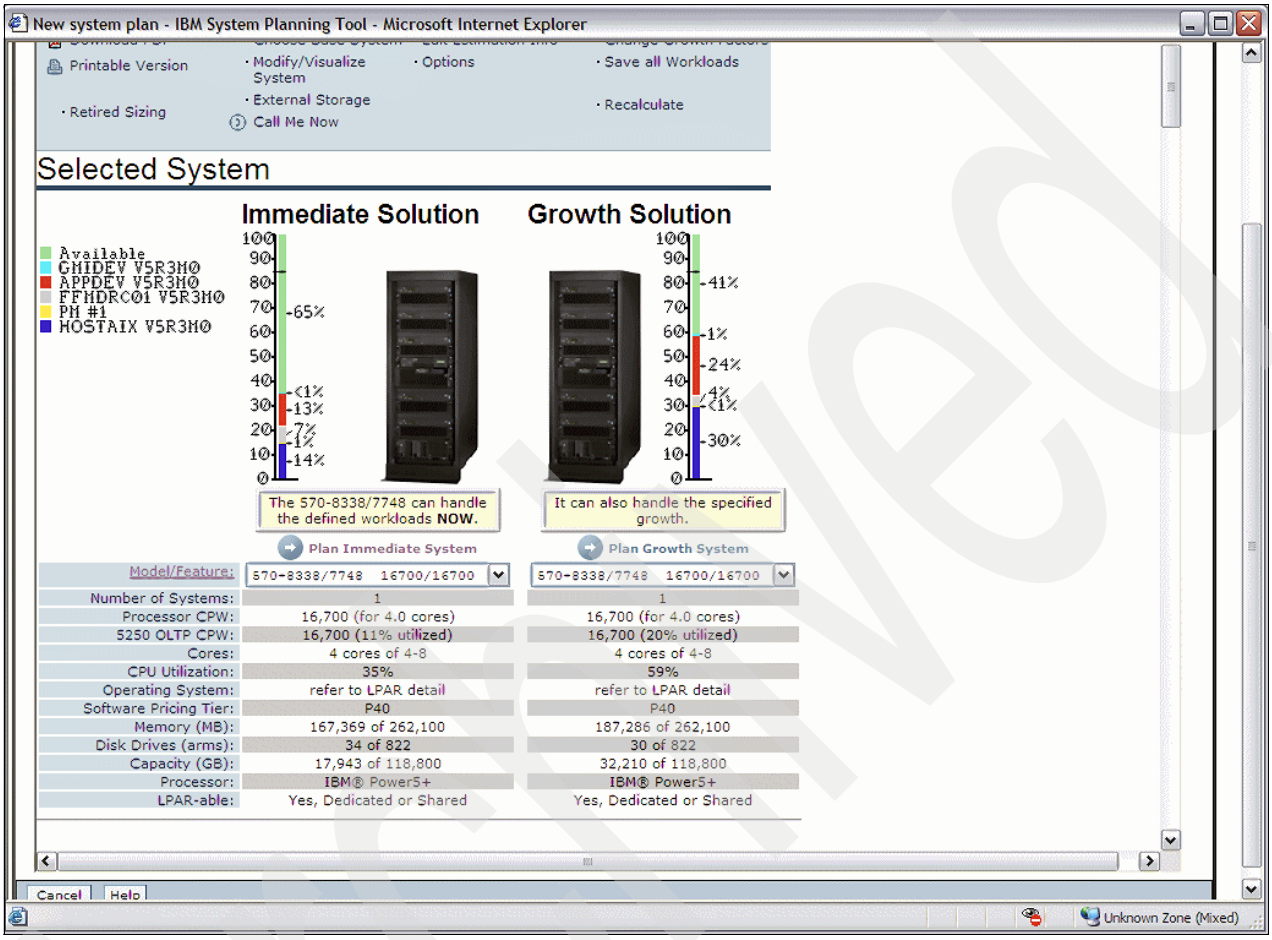

Figure 5-12 WLE selected system

10. If you are satisfied with the result then chose either the immediate or growth model to return to SPT V2. Figure 5-13 points out the selection options.

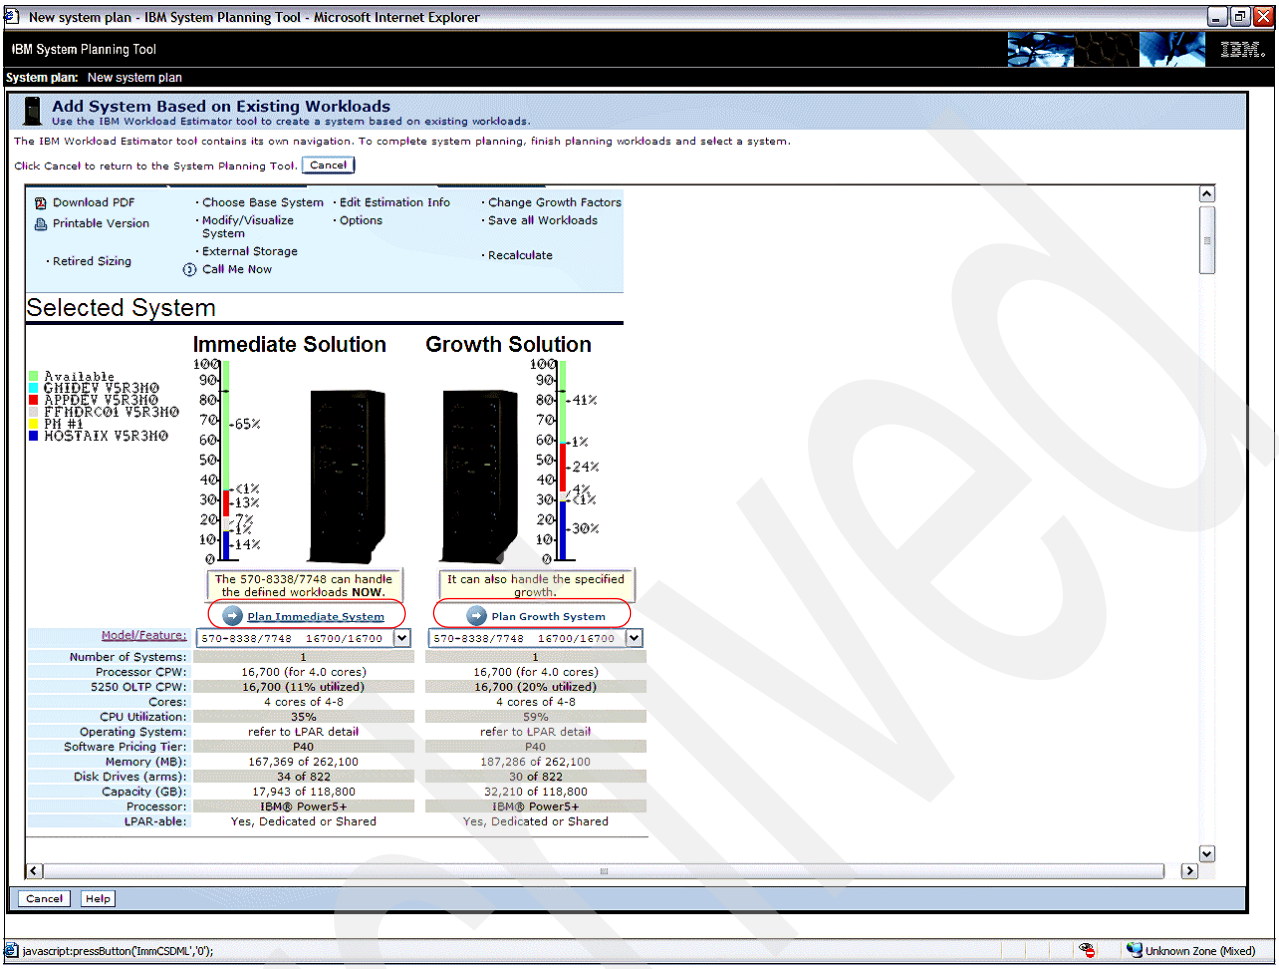

Figure 5-13 Chose immediate or growth system

11. Once you have made your selection, you are returned to System Planning Tool Version 2. WLE has produced a system that needs to be completed. The selected system is generated with the system partitions allocated, memory allocated, and the processor allocations. The partition allocation and naming is shown in Figure 5-14.

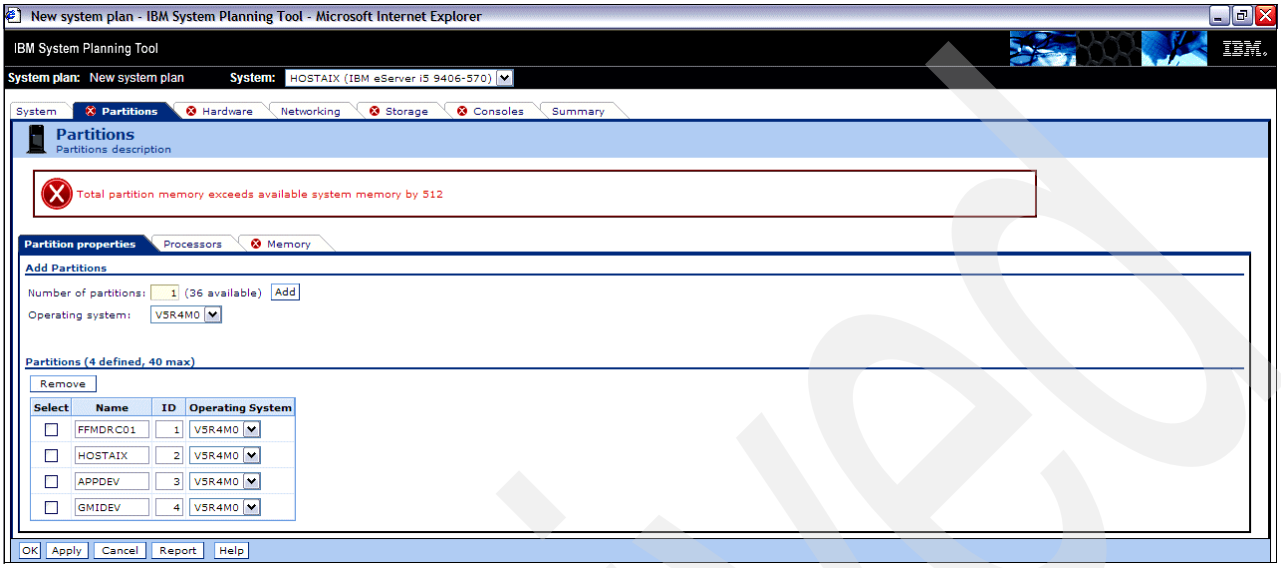

Figure 5-14 WLE generated partition information

See Figure 5-15 for how the processors are allocated based on WLE analysis of the workload of each partition.

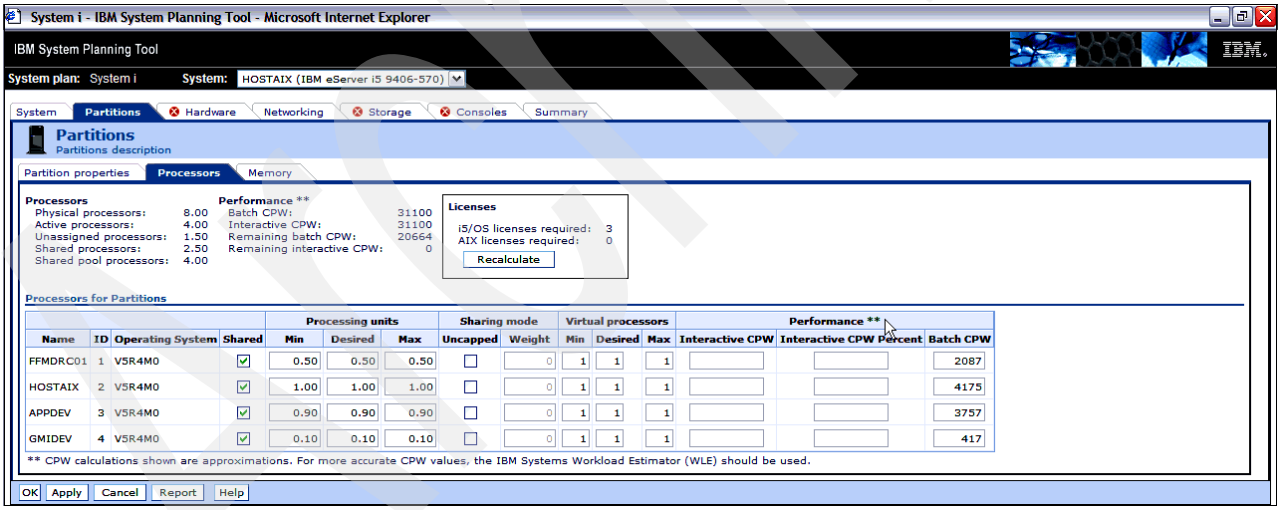

Figure 5-15 WLE processor information

12. Click the **Memory** tab for all of the partitions. Always verify that the memory allocations are what you want. In Figure 5-16 you see the memory allocations by WLE in our example.

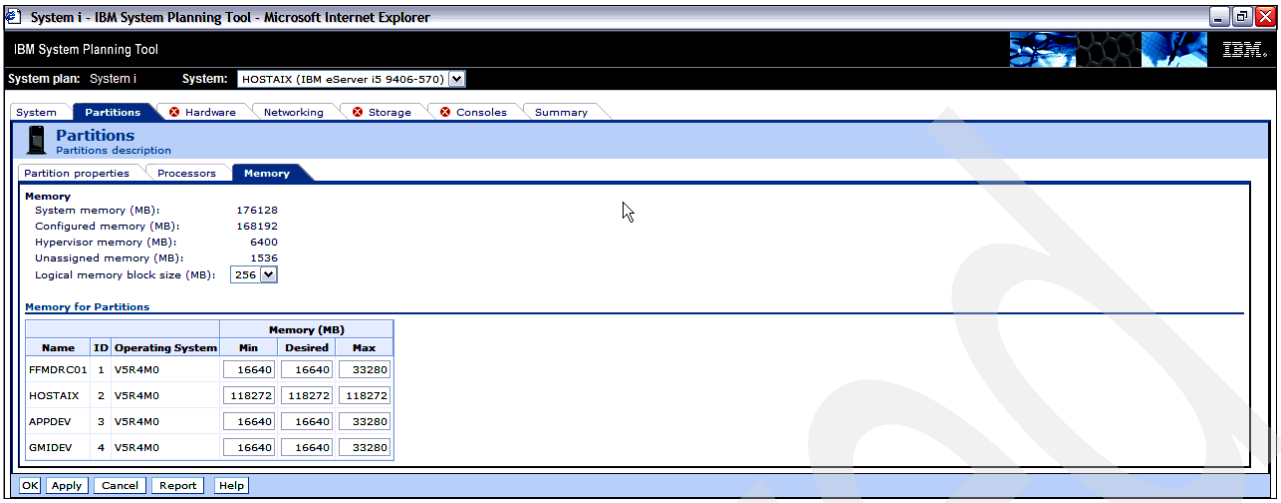

Figure 5-16 WLE generated memory allocation

You can always make adjustments in SPT for processor capacity as well as memory capacity. The WLE output can be viewed as a good starting point for your full configuration under SPT.

You need to make further configuration specifications in SPT before the system configuration can be ordered and later deployed when the system arrives from IBM manufacturing. This includes completing configuration of items such as specifying the number of actual disk controllers, disk drives attached to each controller, Ethernet hardware features, and card locations for these features. You also need to select an i5/OS load source for any i5/OS partition, a console type, any desired tape hardware, and so forth.

In this book we provide information about the range of SPT capabilities in Chapter 2, “System Planning Tool (SPT) V2” on page 11.

When completing your configuration in SPT always save your own copy of the configuration file used as input to the IBM Sales Configurator tool to verify the hardware configuration received from IBM manufacturing. Also save the sysplan file itself to have it available to deploy your LPAR configuration on that system.

5.1.1 Example of SPT V2 and WLE for new workload

In the second example we show the option to create new workloads in WLE. This time we illustrate using a System p example.

Once in WLE we take the option to create a new SPT V2 plan for a System p. Remember that the only way to get to the WLE interface and to return to SPT V2 is to start in SPT V2. Figure 5-17 shows the selection of the radio button for a new workload you would like to run. The following windows show the steps to create a new System p with several new workloads.

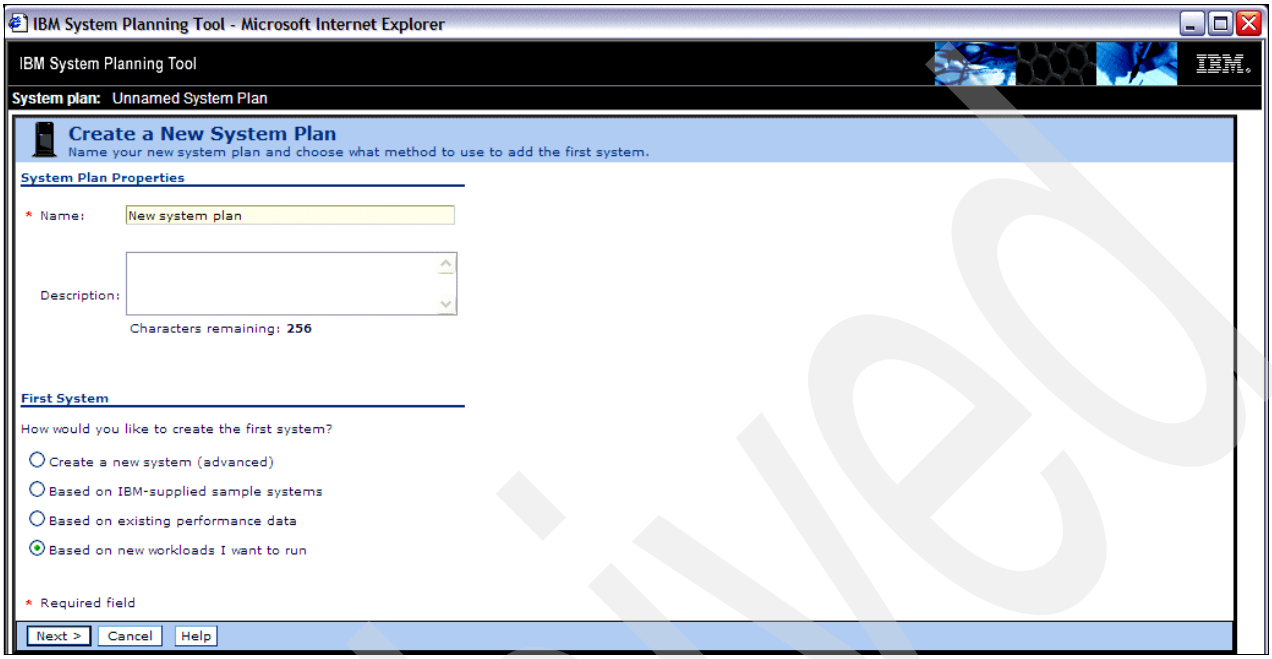

Figure 5-17 SPT interface to WLE

The steps are:

1. Once you leave System Planning Tool Version 2 and start a session in WLE, you are presented with a System i option. To start a System p you need to select the option to create a new System p system, as seen on Figure 5-18.

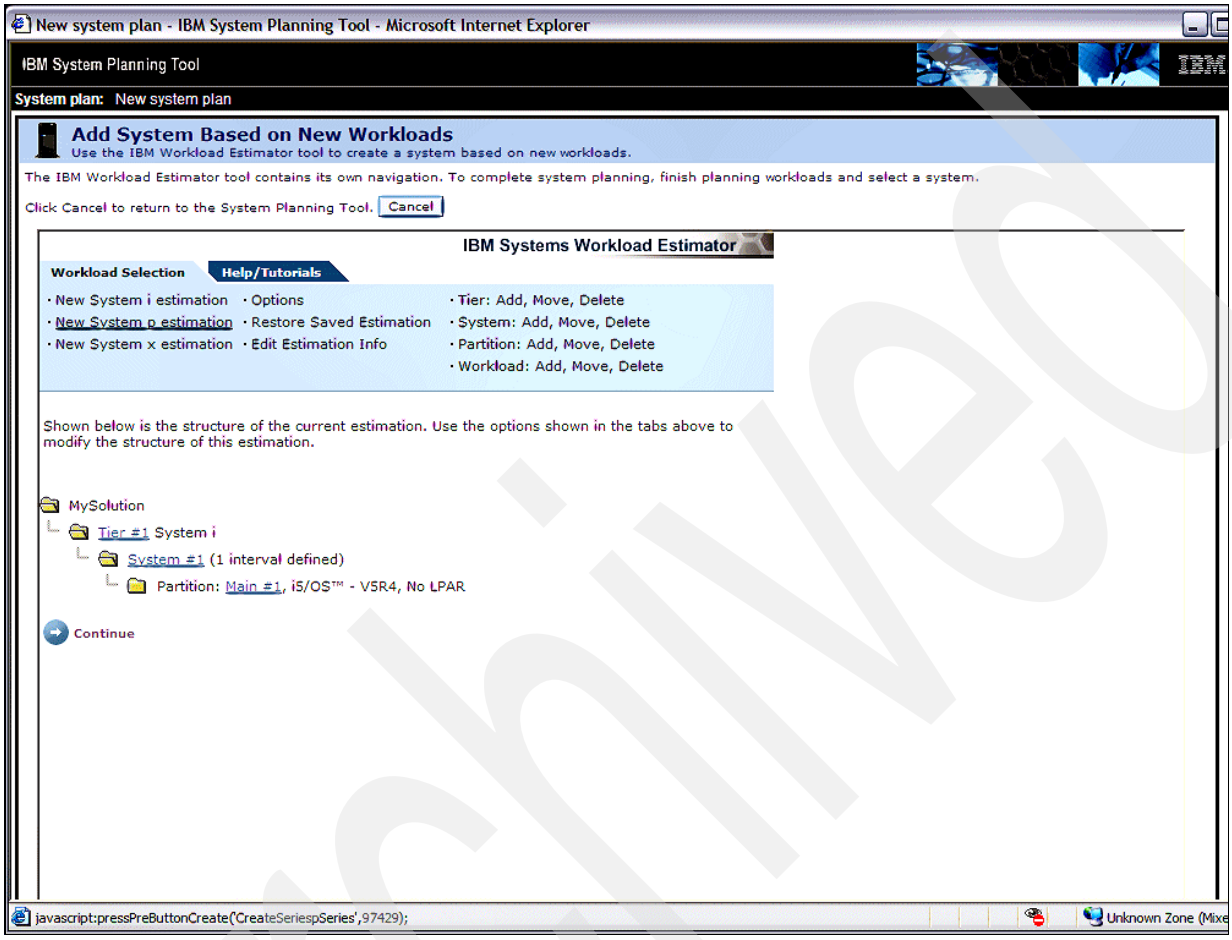

Figure 5-18 Change system type

2. The first window has one system and one partition defined. If you want each of the new workloads to be in a separate partition, then you need to add partitions, as seen in Figure 5-19.

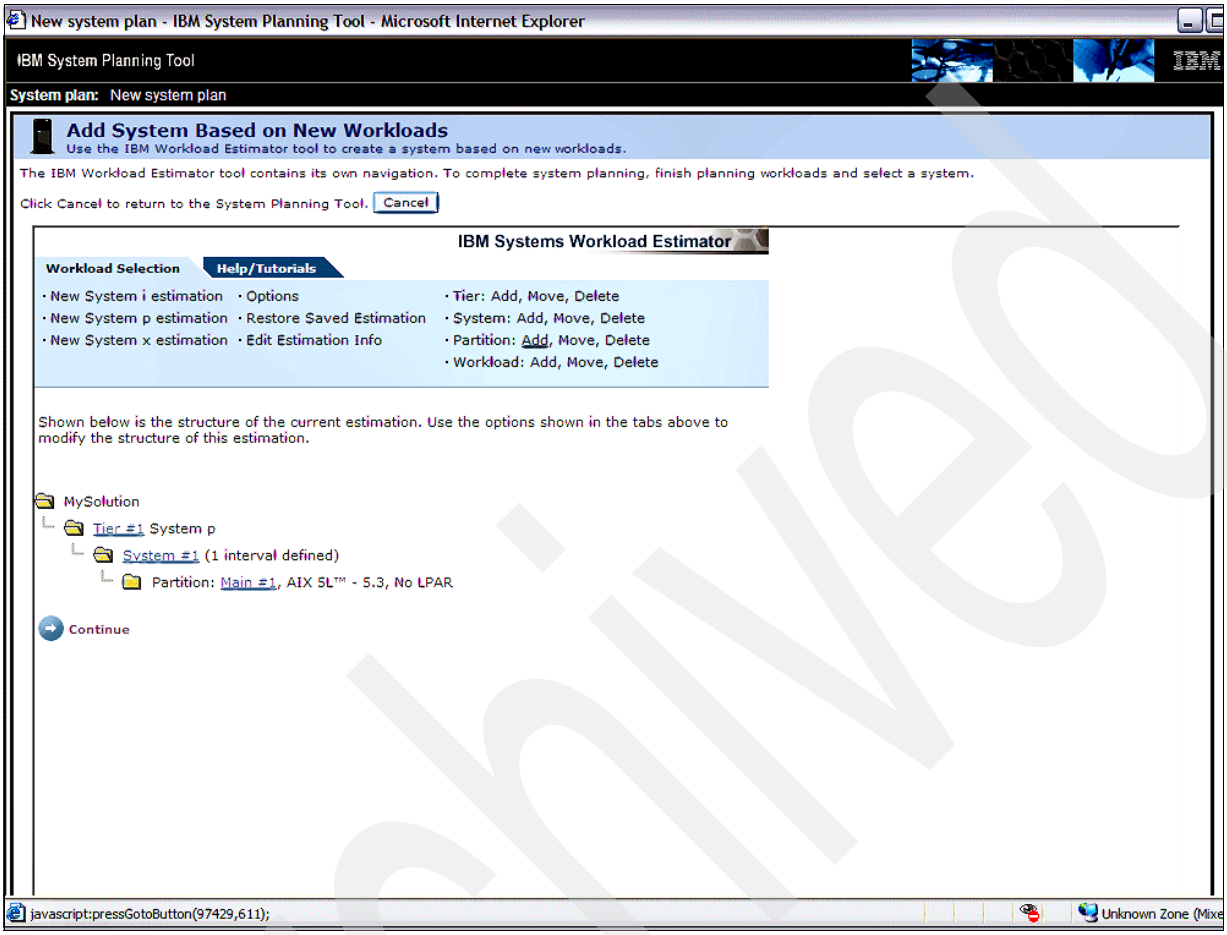

Figure 5-19 Add more partitions

3. You may chose the type of each of the partitions that are added. This can be seen in Figure 5-20. Make sure to click **Go** for each partition. After all the partitions are added click **Return**.

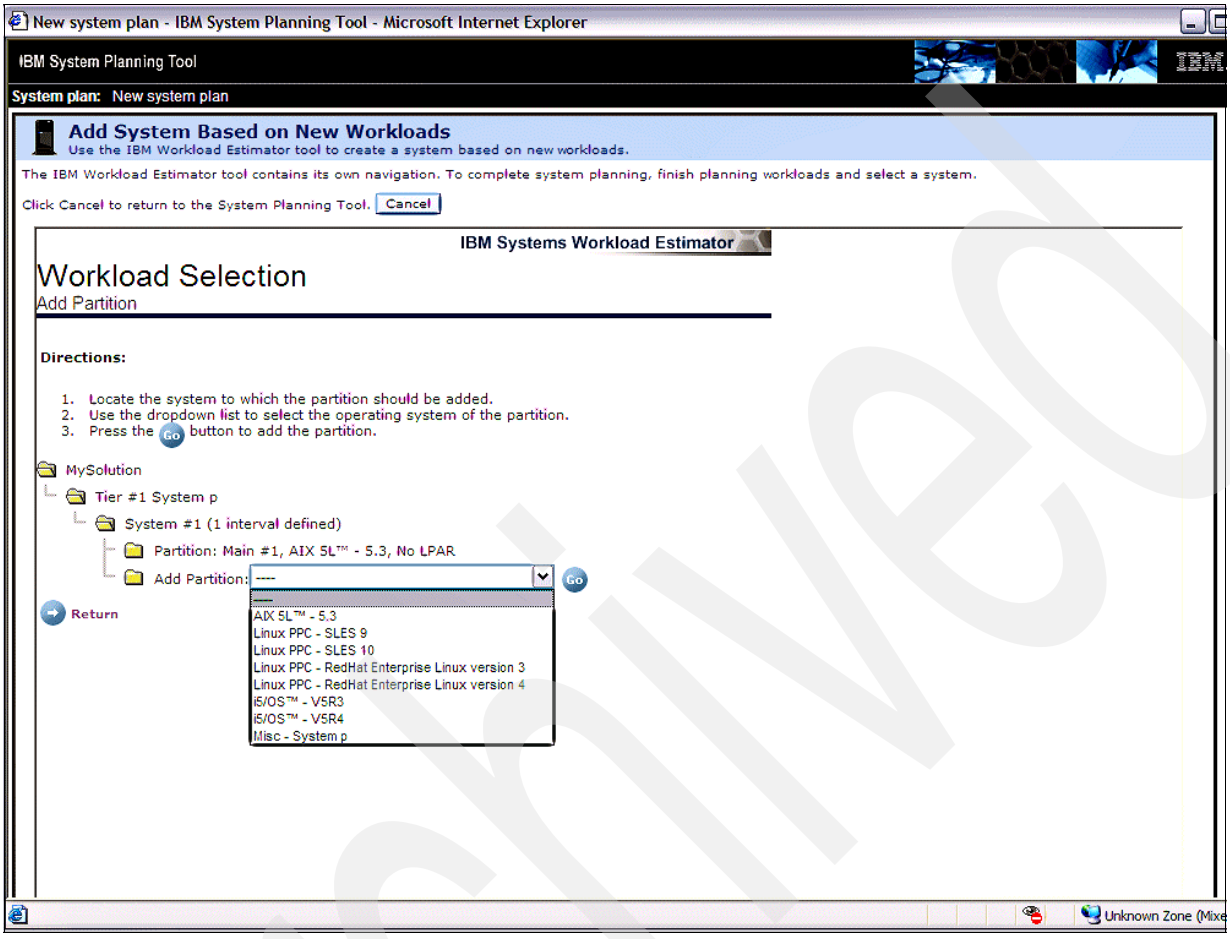

Figure 5-20 Select the type of partition

4. Define the partitions that need to be on the system. Figure 5-21 shows the partition selection. Before you click **Continue**, workloads for each partition need to be defined.

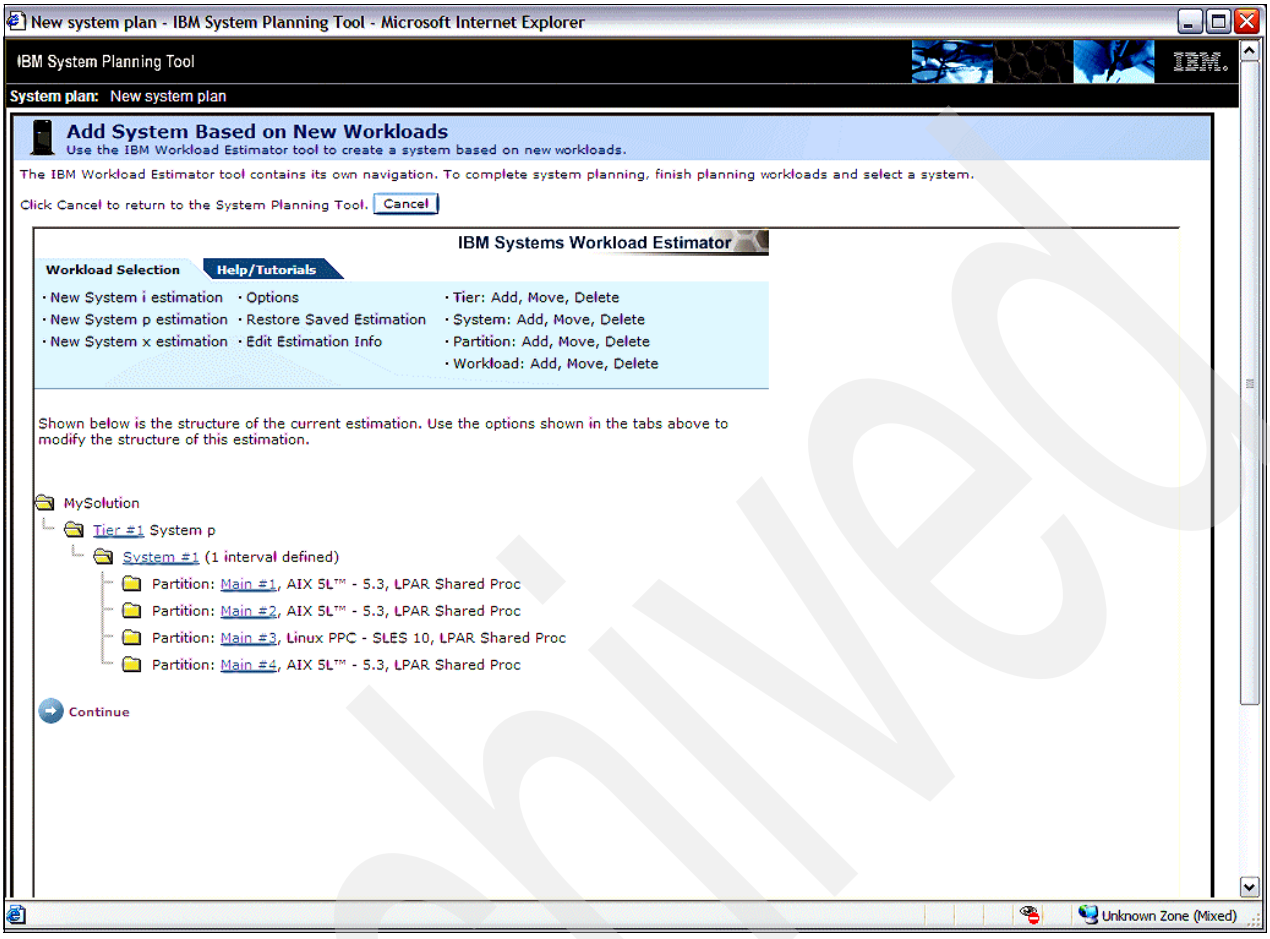

Figure 5-21 Add partitions

5. The next thing that has to be selected for each partition is a workload. The selection options are seen on Figure 5-22. You must click **Go** for the selected workload to be accepted. When all workloads are added to each partition, click **Return**.

Figure 5-22 Select workload type

6. After reviewing the required partitions and workloads, you are ready to start defining the workloads parameters. Figure 5-23 shows the assigned partitions and the addition of your workloads. Click **Continue** for the next window.

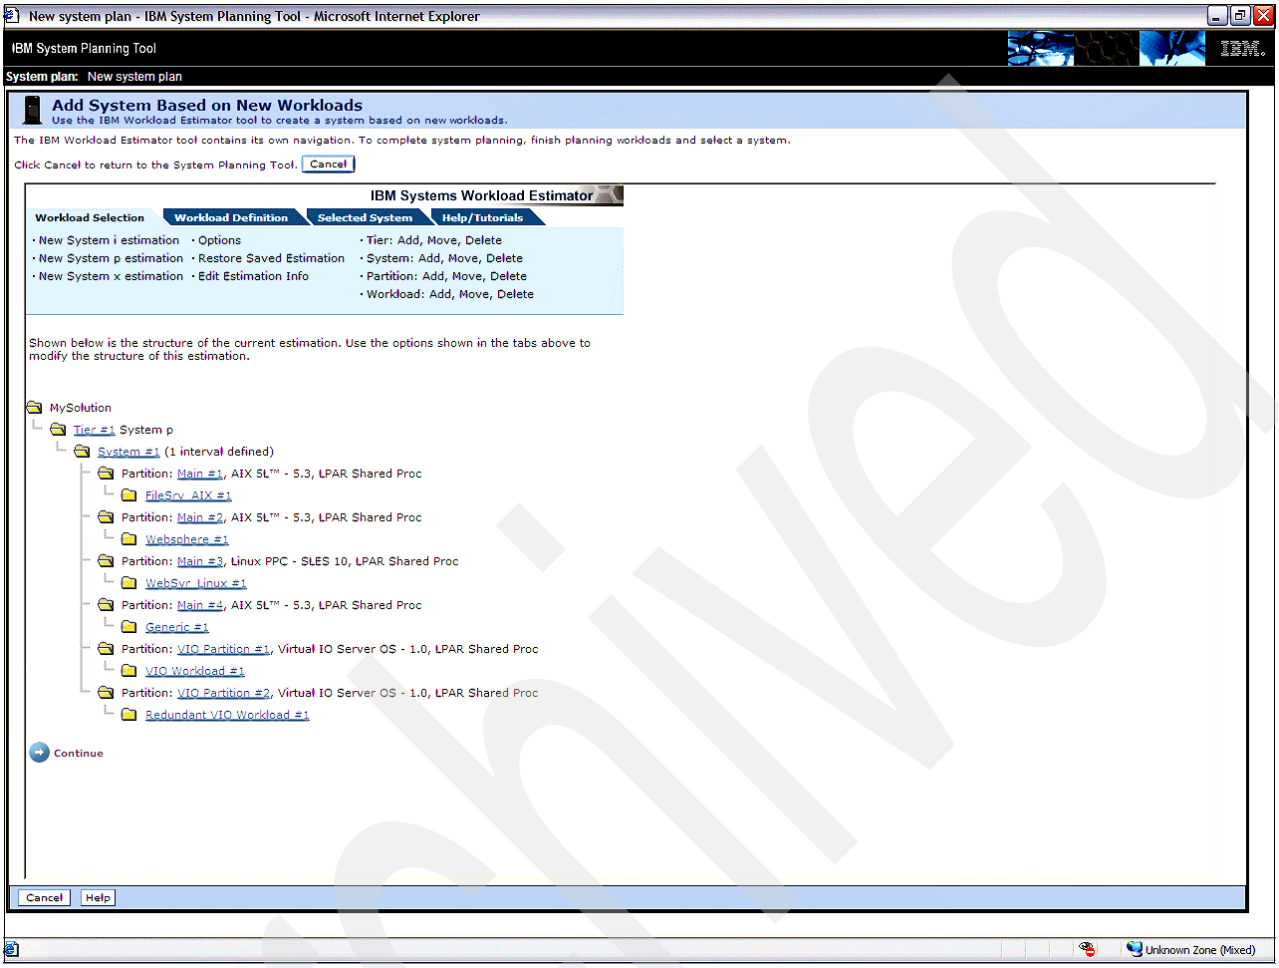

Figure 5-23 Add workload

7. When you chose an AIX file server, the option to create the Virtual I/O Server is an available selection. Figure 5-24 shows the selection options. Each of the following workload definitions has options that you answer. These options are used to determine the CPU, memory, disk, and system selection that the workload requires. In the following examples the values entered do not reflect any specific workload. The value entered for a sizing should reflect a through analysis of the application. Virtual I/O Server allows the virtualization of the disk and networks for AIX and Linux. By making this selection WLE generates a Virtual I/O Server partition. The parameters for these partitions are defined in System Planning Tool Version 2.

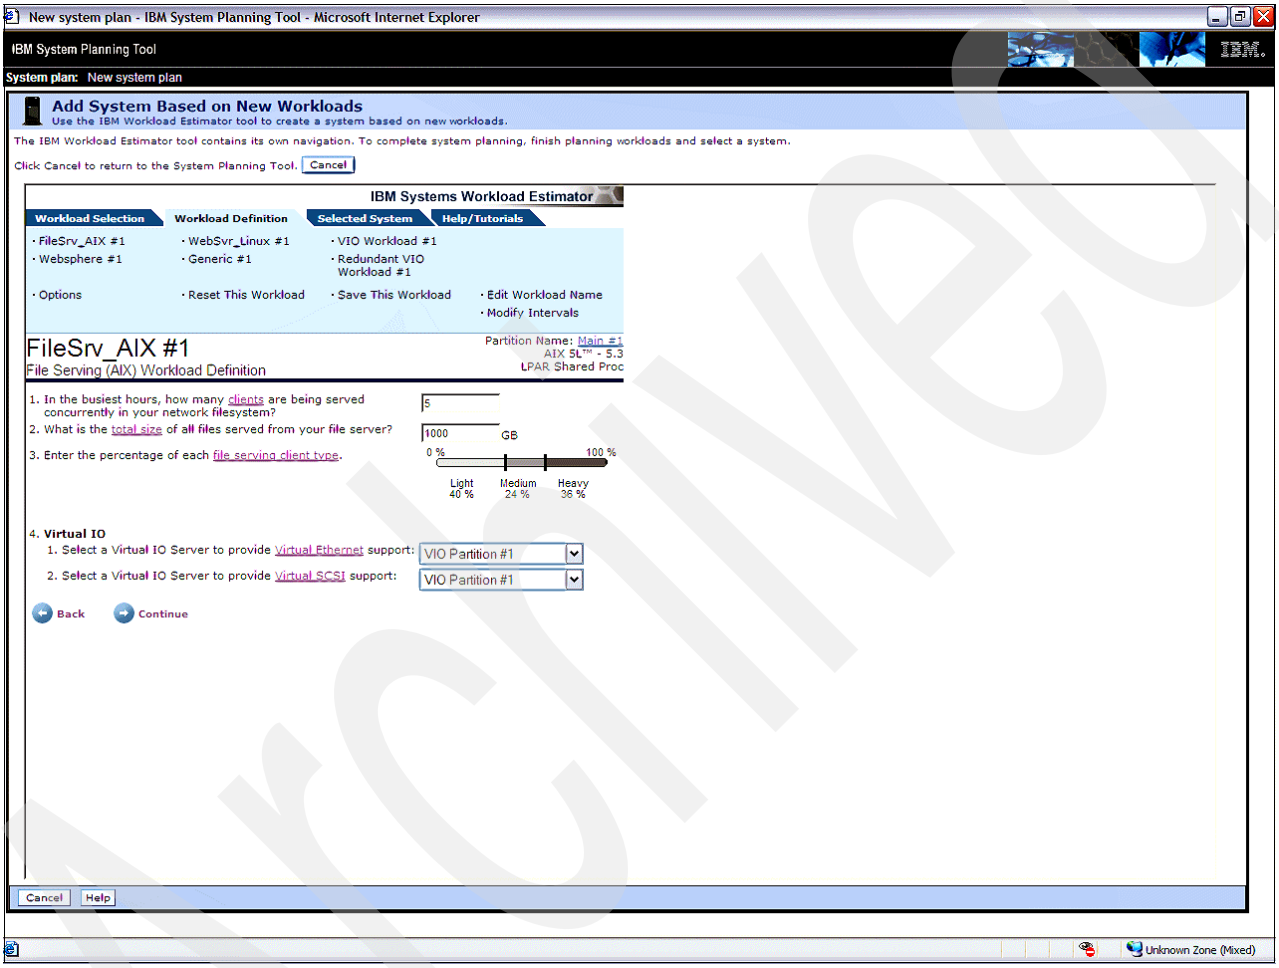

Figure 5-24 AIX file serving and the building of a Virtual I/O Server partition

8. The next four windows show the questions for sizing an AIX WebSphere® application. The option to add this partition to a Virtual I/O Server is one of the selectable option. The examples in this selection are intended to show the options available, not show how to size a WebSphere application. Figure 5-25, Figure 5-26 on page 174, Figure 5-27 on page 175, and Figure 5-28 on page 176 show the question used for the sizing of a WebSphere application server.

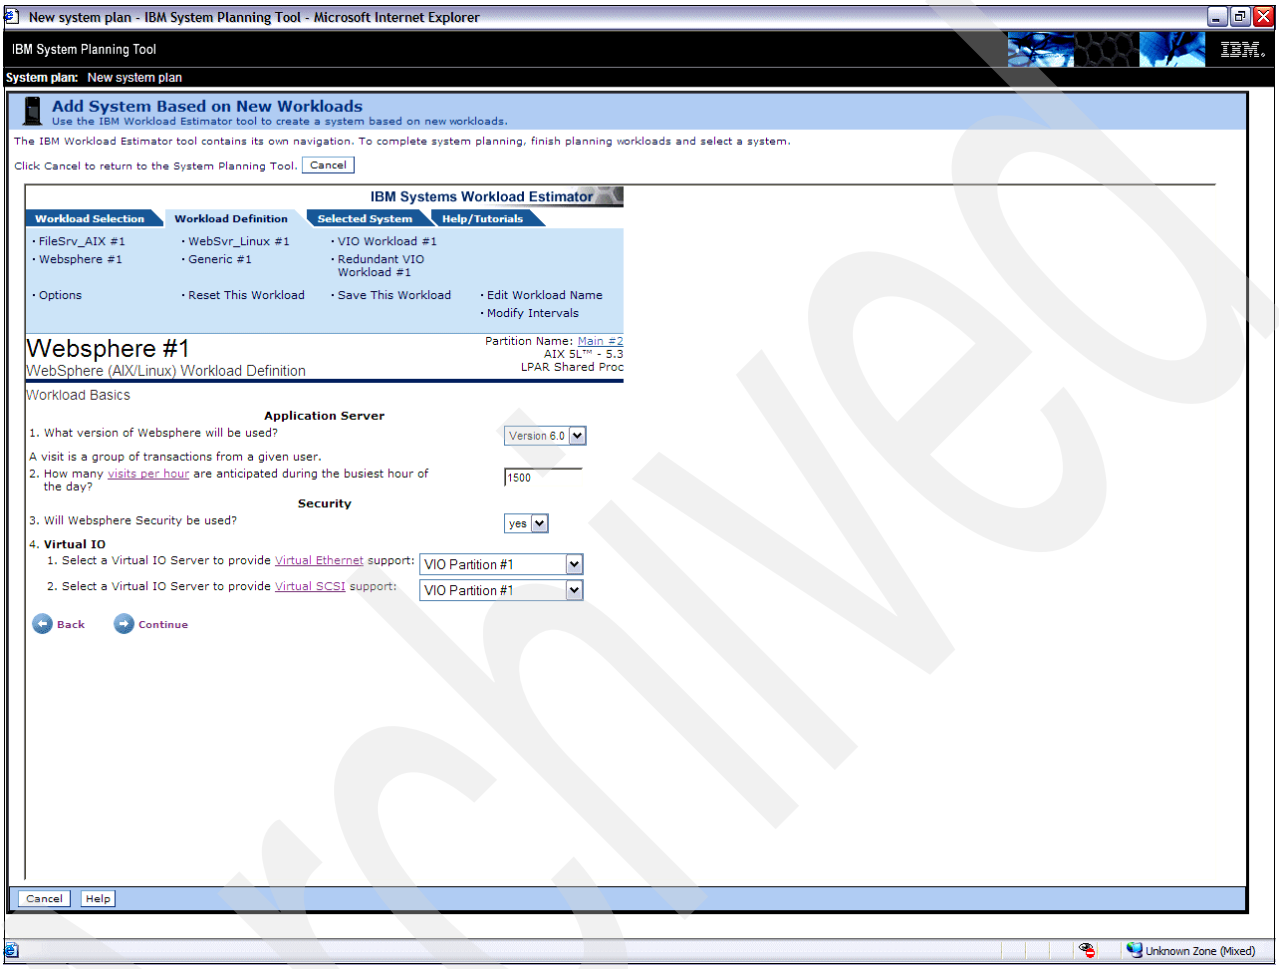

Figure 5-25 Defining a AIX WebSphere partitions - window 1 of 4

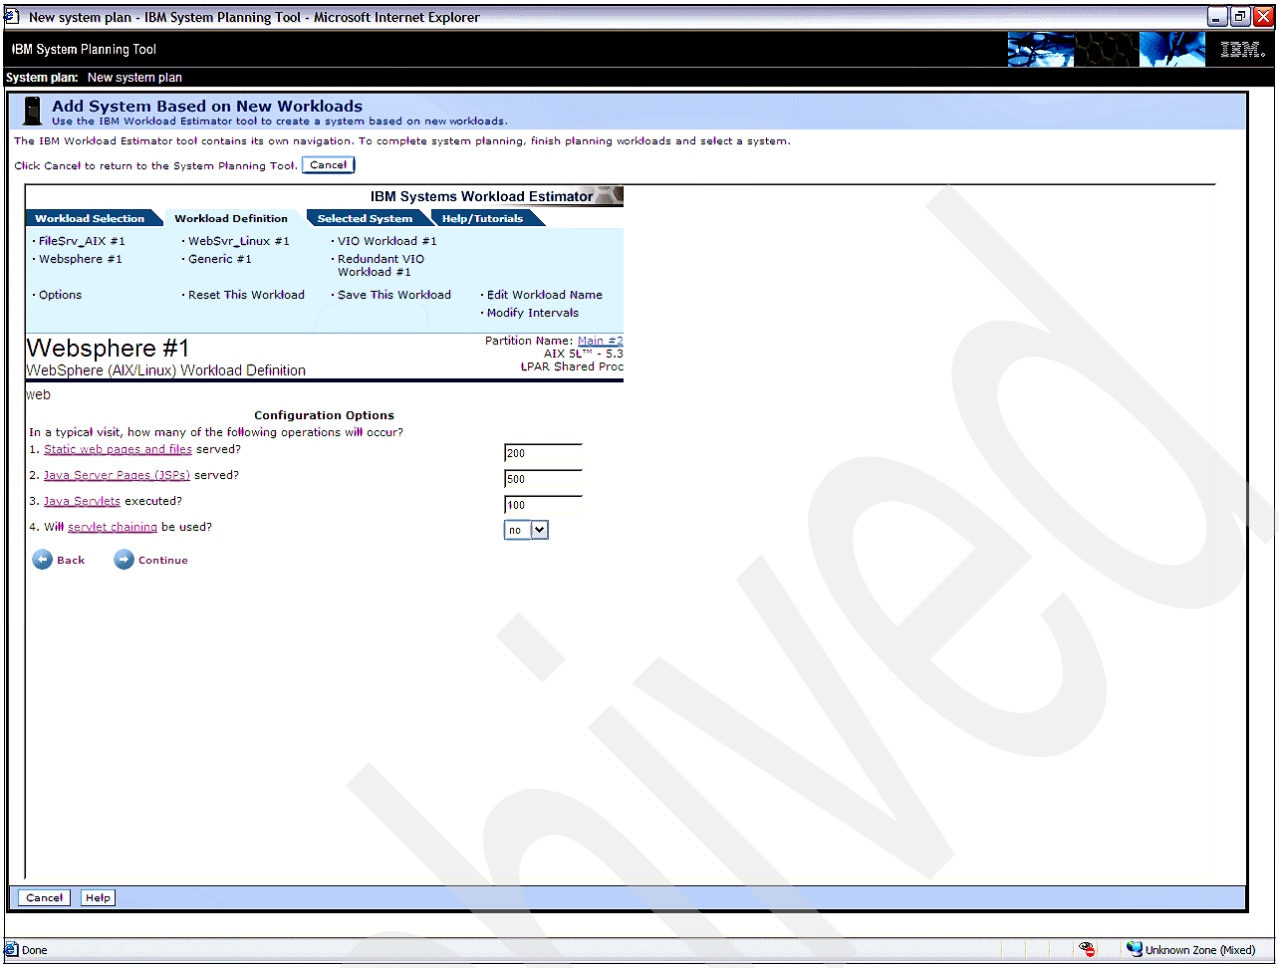

Figure 5-26 Defining a AIX WebSphere partitions - window 2 of 4

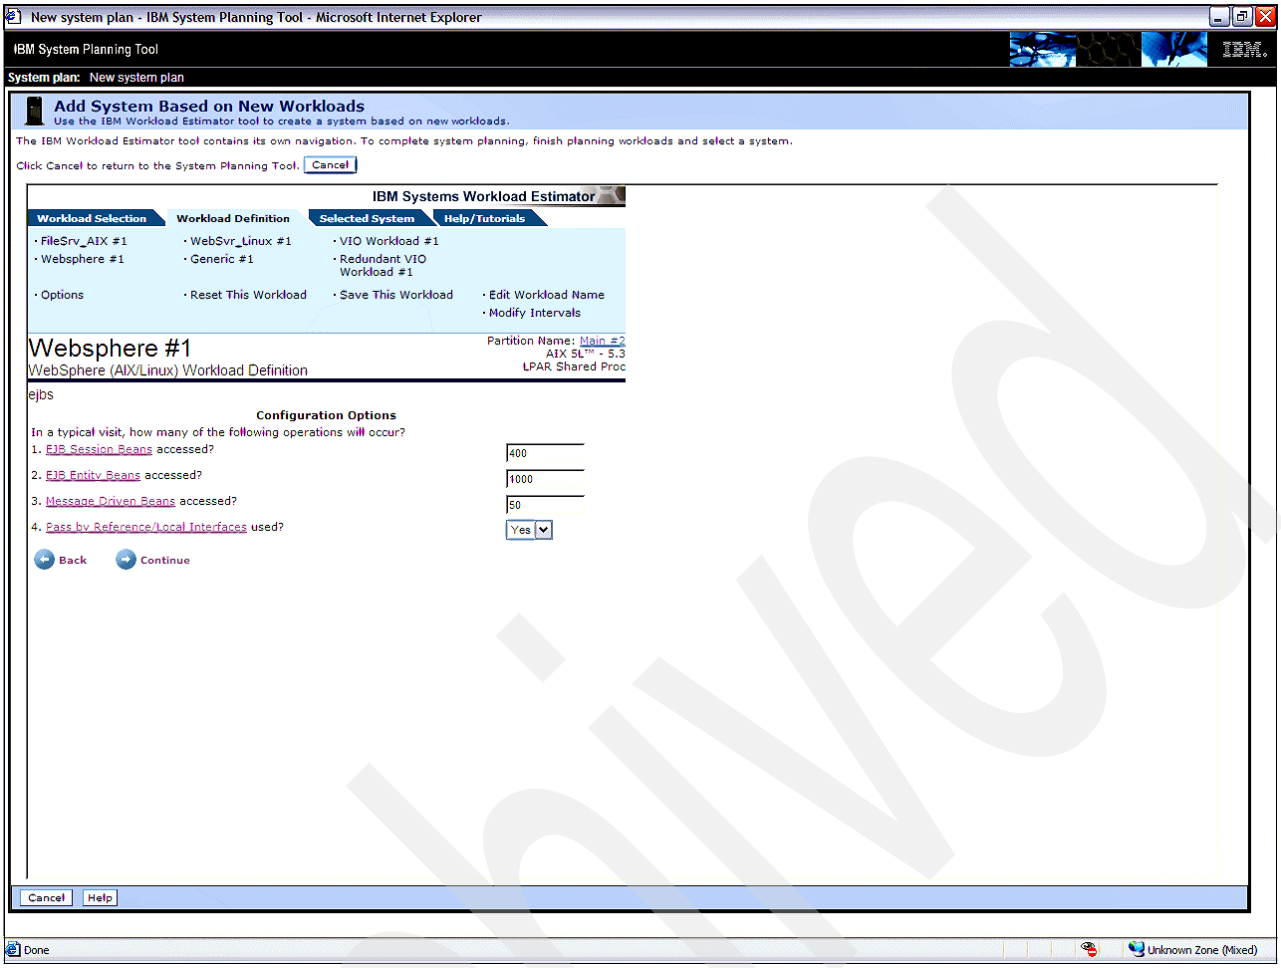

Figure 5-27 Defining a AIX WebSphere partitions - window 3 of 4

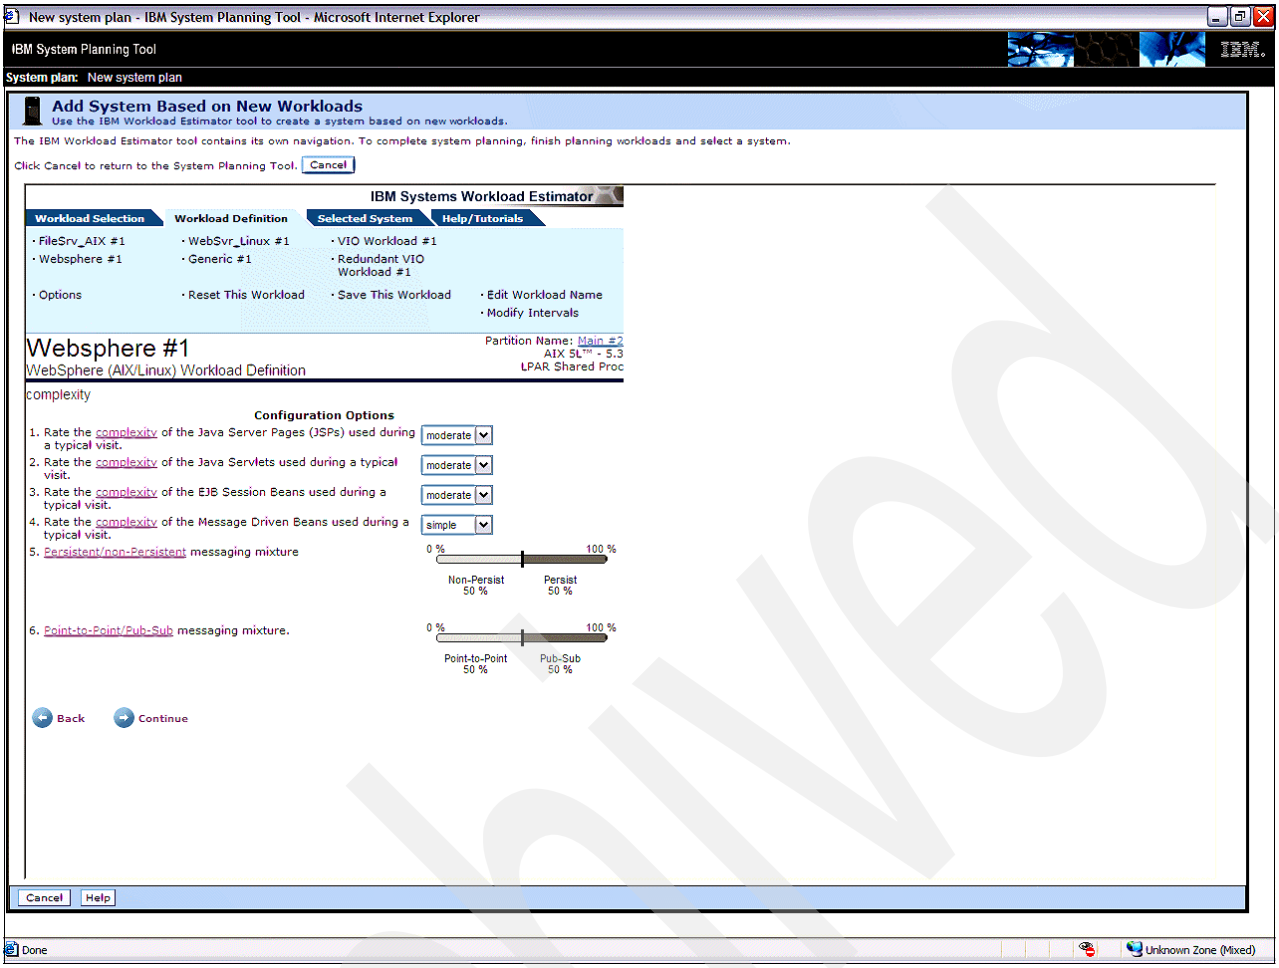

Figure 5-28 Defining a AIX WebSphere partitions - window 4 of 4

9. The next partition is a Linux Web server, as seen when you defined the partitions and workload in the window shown in Figure 5-23 on page 171. Again the option to add this partition to the Virtual I/O Server is a selectable item. Figure 5-29 shows the options.

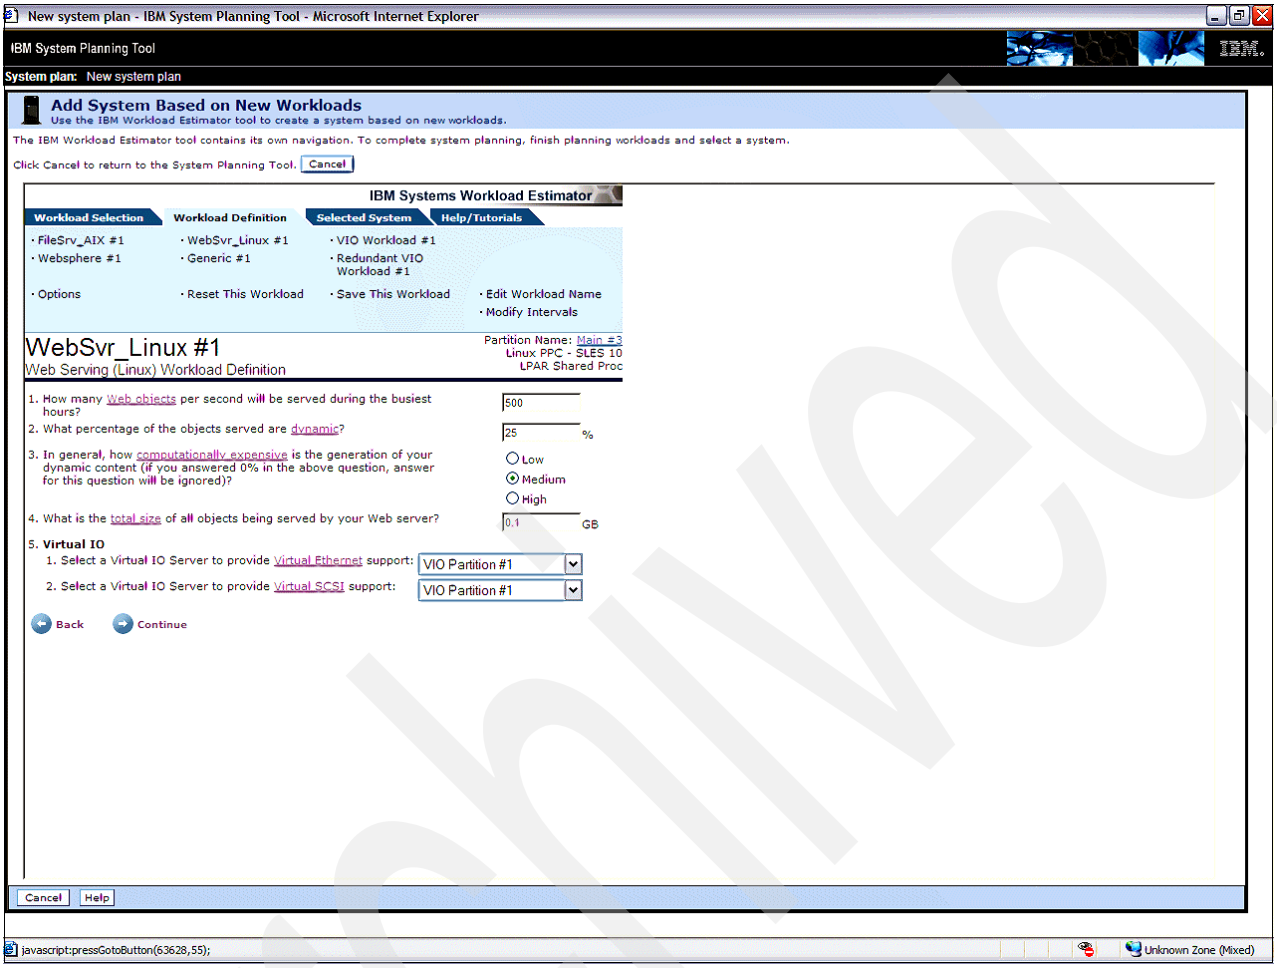

Figure 5-29 Linux Web server

10. For a generic AIX partition information is required to determine how much CPU is required, memory, and disk usage. The CPU requirements are expressed in rPerf. Memory is the amount for the partition and disk can be defined with capacity and number of arms or by the amount of disk reads and writes. Figure 5-30 shows the questions.

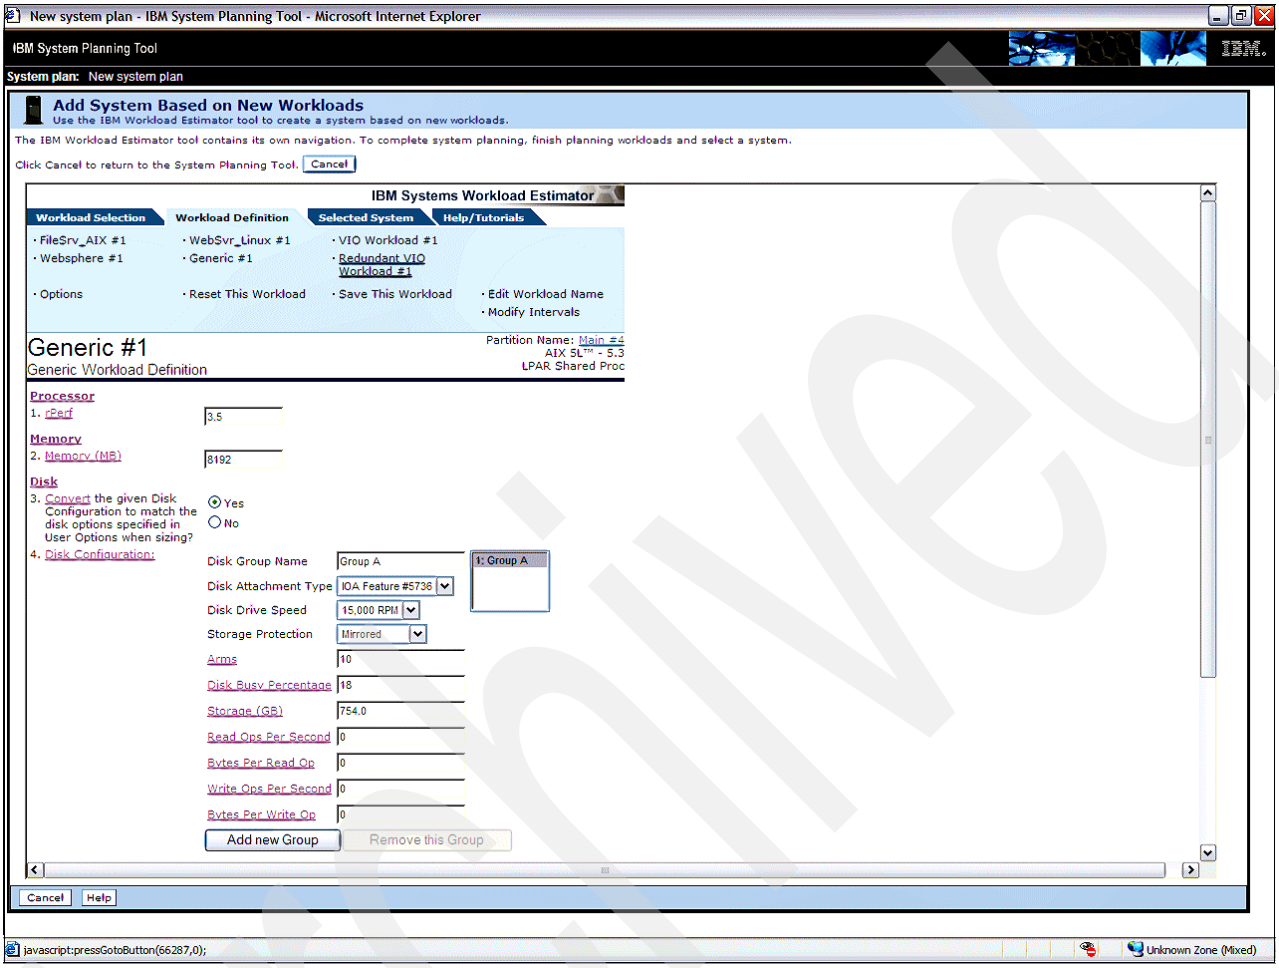

Figure 5-30 Generic AIX workload

11. The question is asked if a redundant Virtual I/O Server is needed. In the Redpaper *IBM System p Advanced POWER Virtualization Best Practices*, REDP-4194, we recommend a redundant Virtual I/O Server for mirrored volumes and to allow a more robust use of Multiple Path I/O connections. Figure 5-31 shows the selection options. After making your selection click **Continue**.

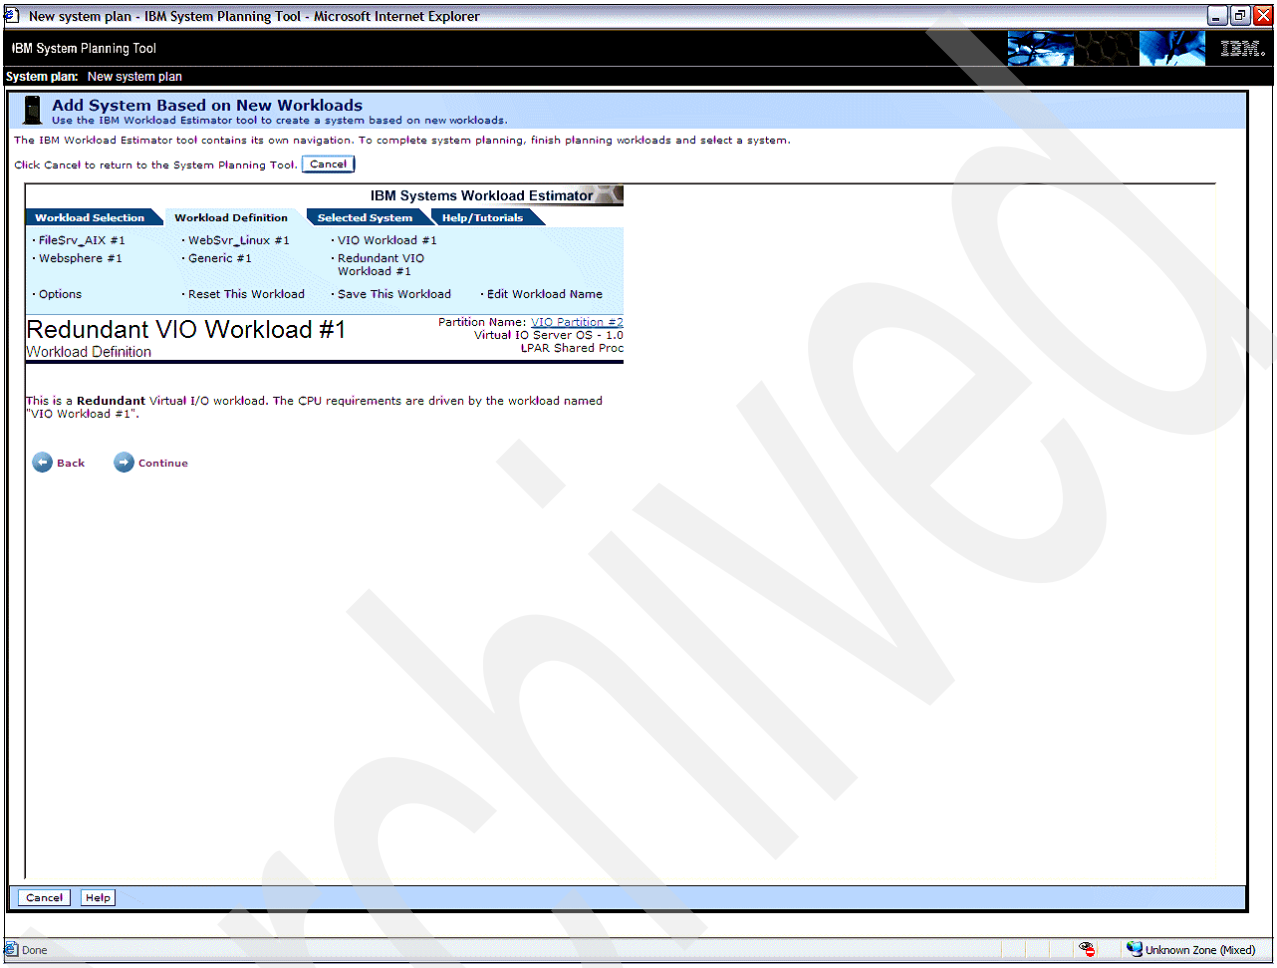

Figure 5-31 Question for redundant Virtual I/O Server partitions

12. You are then presented with a window showing the partitions that use the Virtual I/O Server partitions. Click **Continue**.

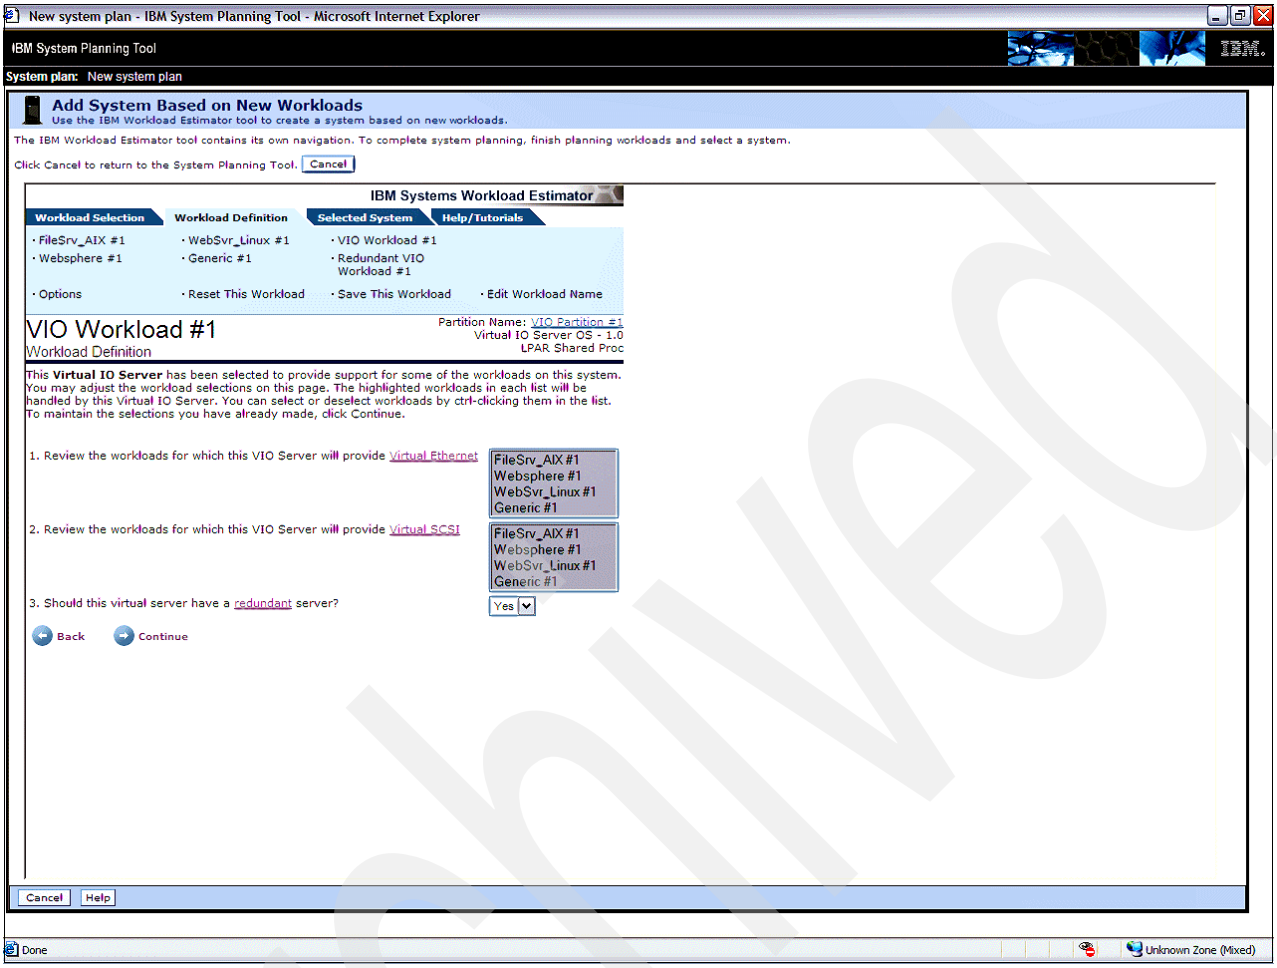

Figure 5-32 Virtual I/O Server supported partition from selected options

13. WLE presents the selected system based on the workload analysis. Figure 5-33 shows the selected system. From this window you create a PDF document, return back to change a workload, chose to export the information for external storage, and save the results of the analysis. If you select to see the external disk requirements you are presented with the window shown in Figure 5-34 on page 182 to see the requirements. You can select the immediate or growth result, which are circled in red, to return to System Planning Tool Version 2.

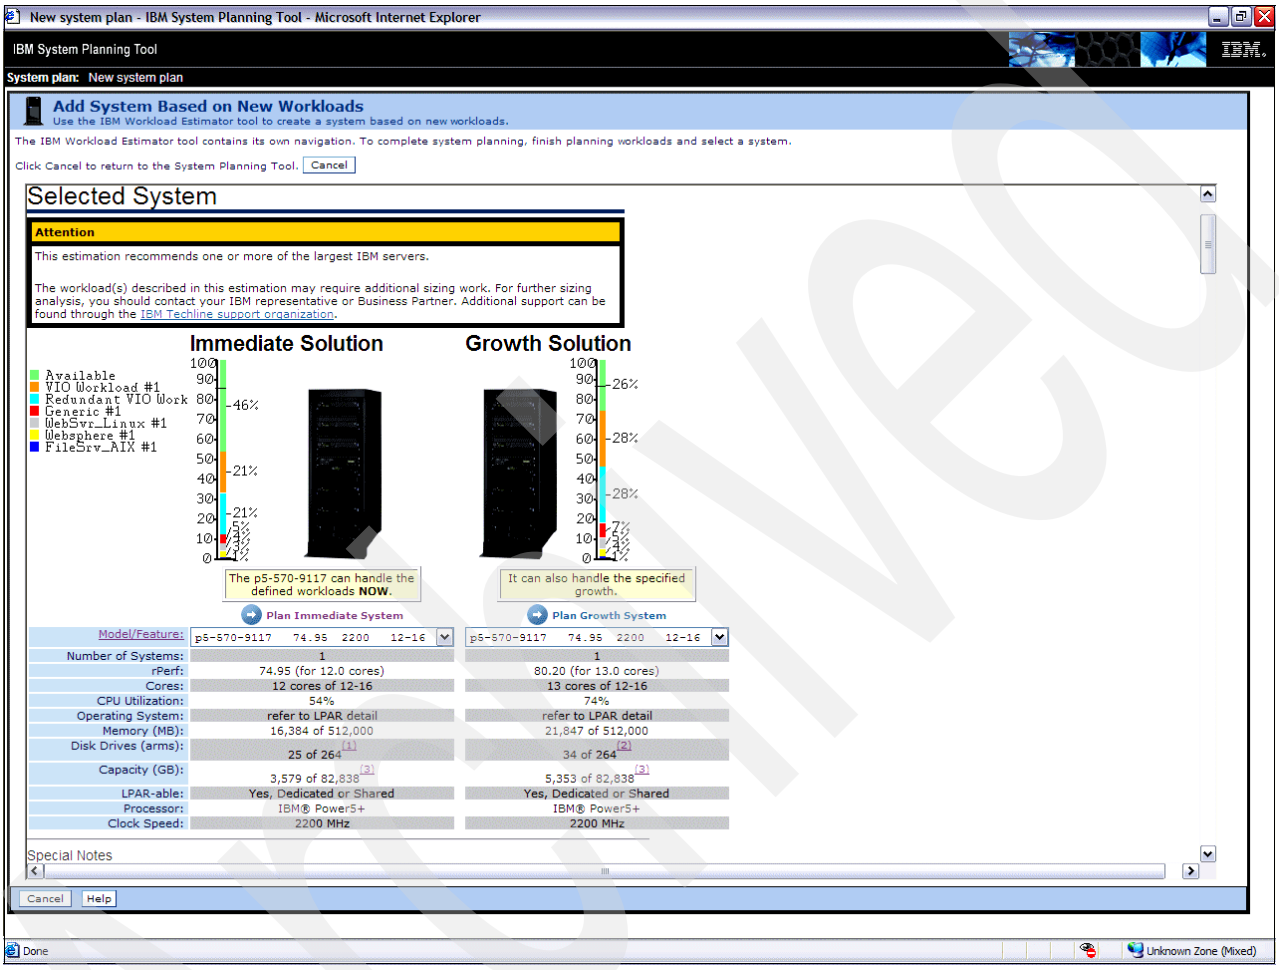

Figure 5-33 WLE selected system

14. Since the Virtual I/O Server can use external storage, WLE calculates the amount of disk needed. From this window you can download the information to go into DiskMagic. Figure 5-34 shows the option. In this example we show no further processing under DiskMagic, which is beyond the scope of this publication. Click **Return** to get to the selection of the system to return to System Planning Tool Version 2.

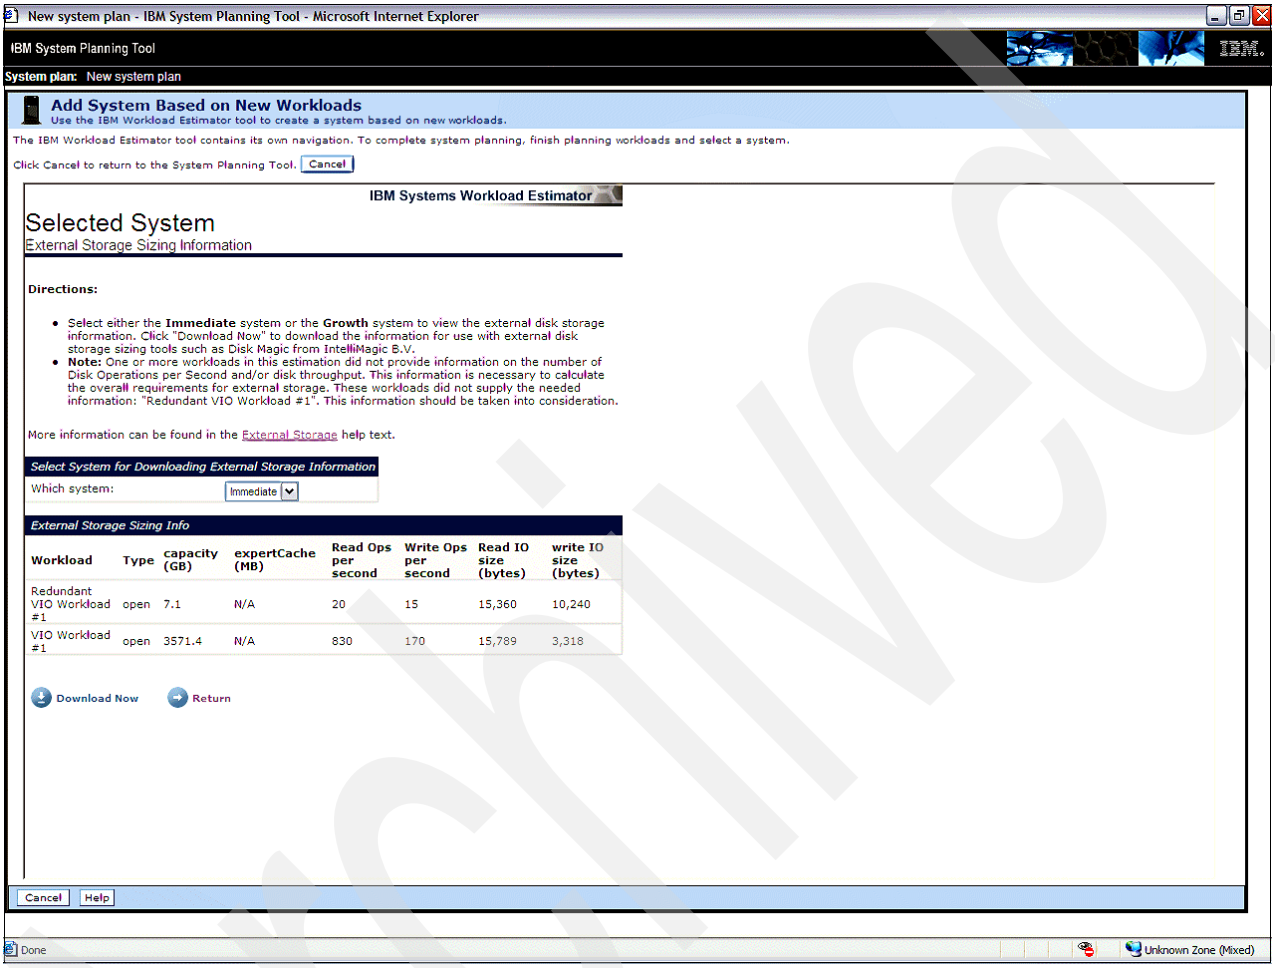

Figure 5-34 WLE show calculated information for a SAN

15. After returning from the storage window and the selection is made for the immediate or growth system, the following window shows the result in SPT what WLE generated for the desired workloads. Figure 5-35 shows the required system attributes for the desired workloads.

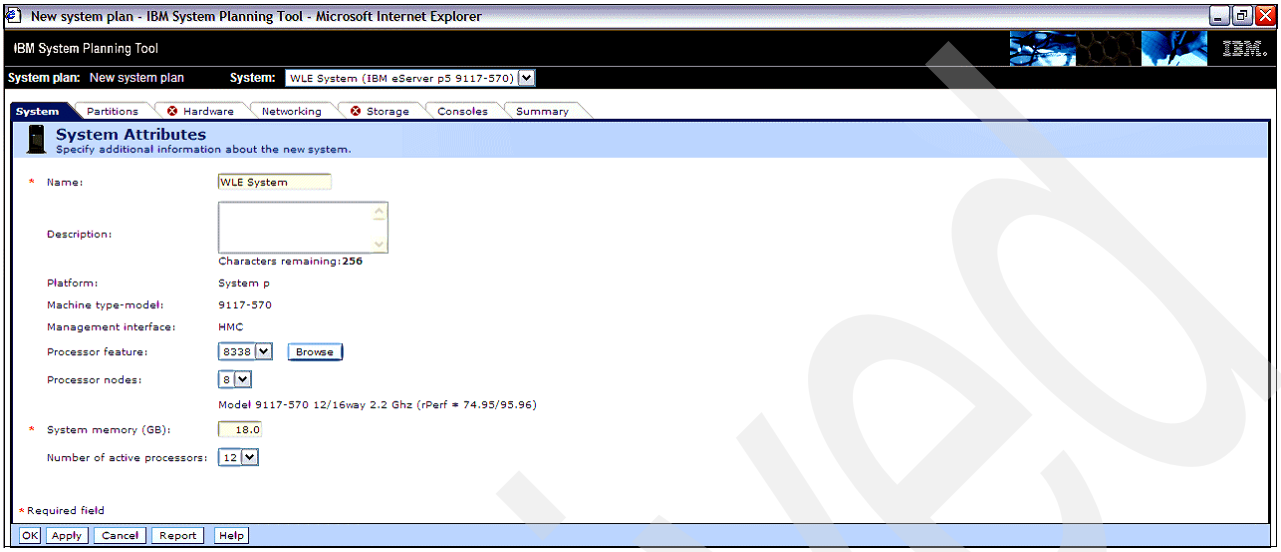

Figure 5-35 SPT V2 shows returned model selection

16. Click the **Partition** tab to see the selection of the assigned names that is generated from WLE. You may change the name. The names could also be changed on WLE when you added partitions. See Figure 5-36.

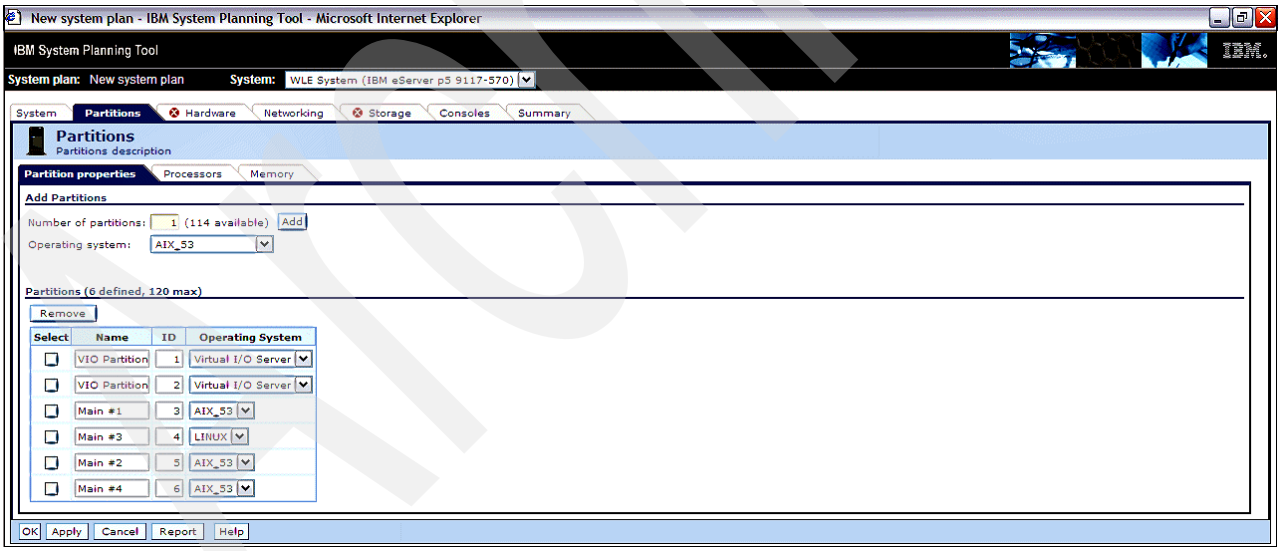

Figure 5-36 Partition definition

17. Click the **Processor** tab to see the definition of each partition processor allocation. These are the allocations from WLE and are shown on the window (Figure 5-33 on page 181) by using the slide bar on the right of the window to see the partition information based on the desired workload. Figure 5-37 allows you to change the allocation if more or less processors are desired for a partition. The virtual processors are rounded up to the next whole number, which is the recommended value. However, they may be increased to a larger number after careful evaluation.

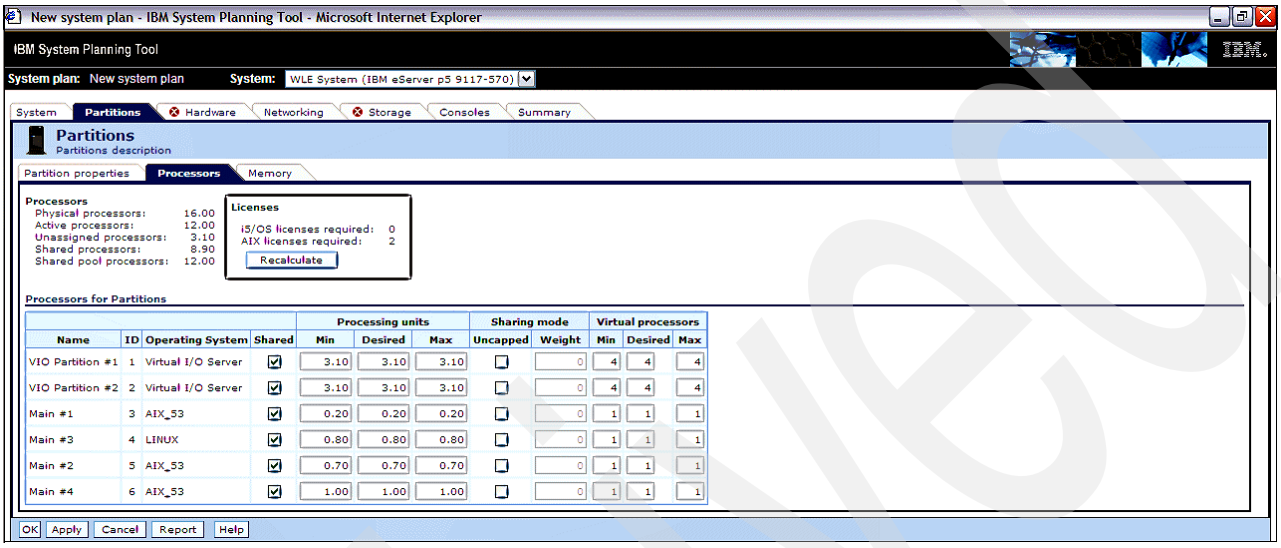

Figure 5-37 Processor definition

WLE takes the default action to set the partition's sharing mode to capped (by not checking the box under the Uncapped column. There are pros an cons to selecting a partition to be capped or uncapped. Specifying capped means that the partition will not automatically try to use more processor capacity (if available) during run time. If you do not understand your system's total operation for all active partitions, you should select capped.

However, if you know your operating environment is such that work in one active partition decreases CPU utilization for significant periods while the workload in other active partitions significantly increases, you should take advantage of the system's capability to assign processor power to a partition - based upon the sharing mode Weight relative value among active partitions.

Further discussion of how performance is managed with uncapped partitions is beyond the scope of this publication.

18. Click the **Memory** tab to see the memory allocation set by WLE. Figure 5-38 shows the allocation by partition. These allocations can be changed if you determine that allocations should be changed or you want more memory added to the system.

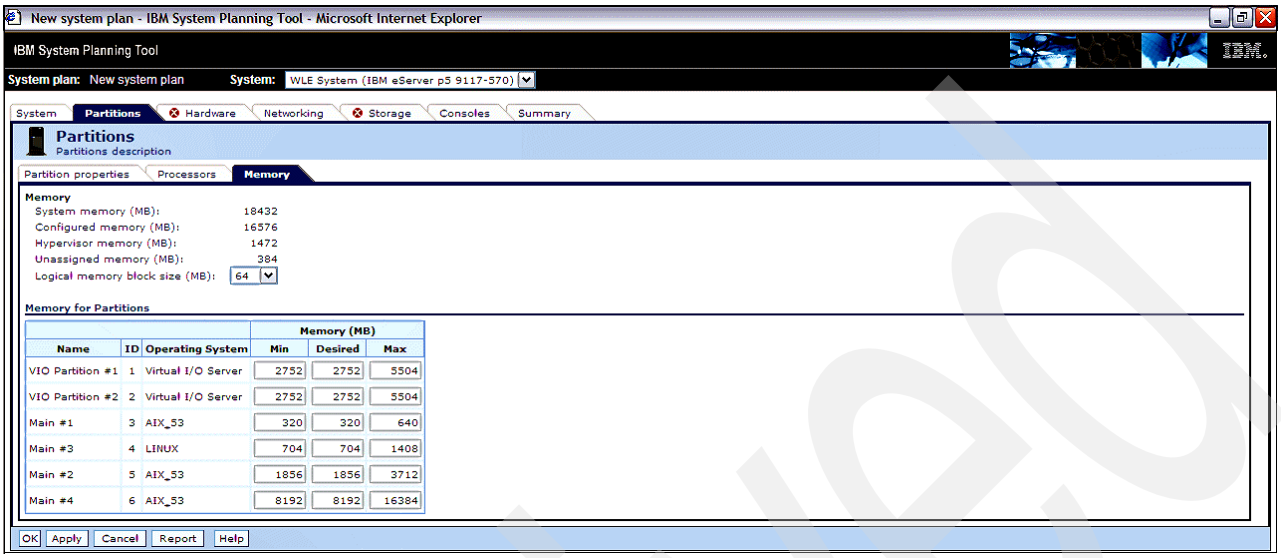

Figure 5-38 Memory definition

19. Based up on the system recommended by WLE, a default layout of the system hardware is placed under the Hardware tab. If additional adapters are required add them to the hardware. Figure 5-39 shows the default hardware of the WLE recommended system.

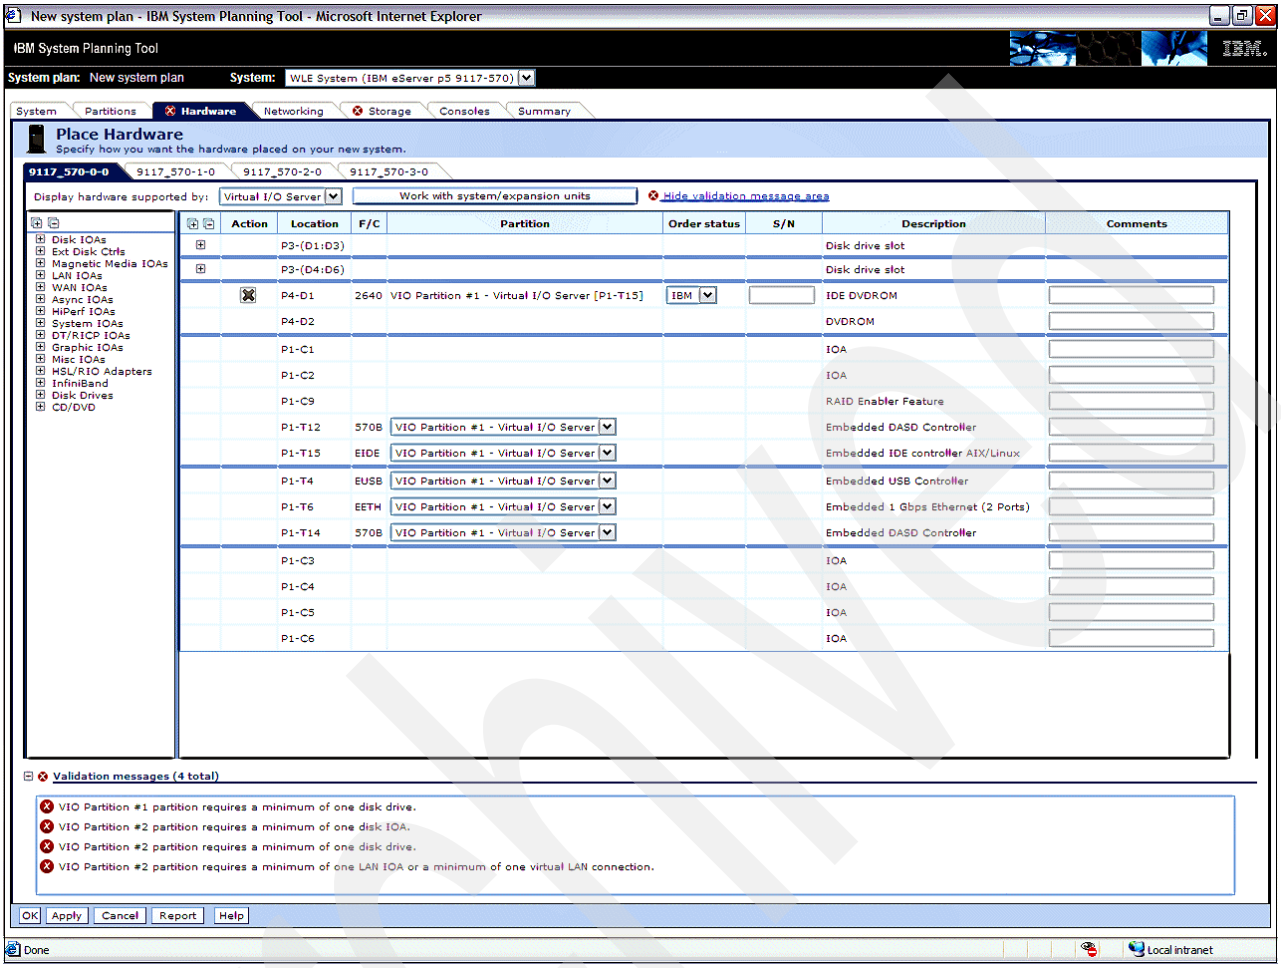

Figure 5-39 Hardware definition

20. WLE has defined the VLAN for you. Figure 5-40 shows the VLAN. For further explanation see 4.3, “Define Virtual I/O Server in System Planning Tool Version 2” on page 106.

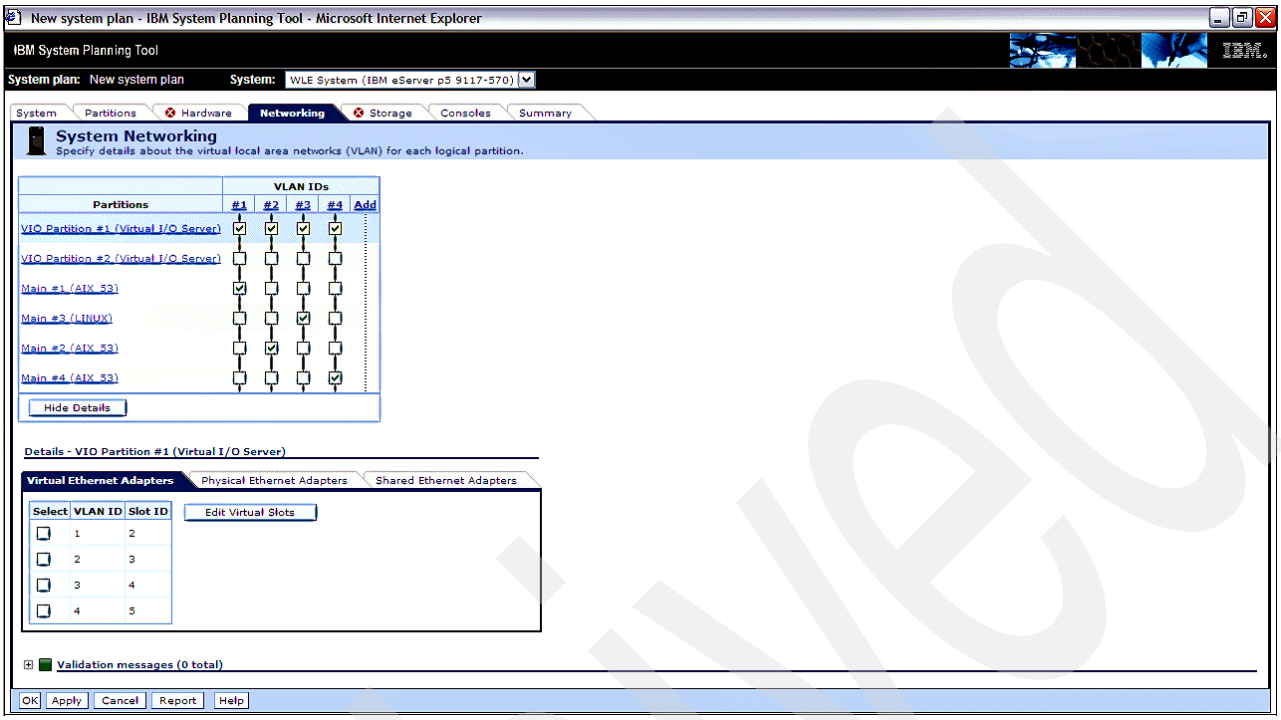

Figure 5-40 Virtual LAN definition

21. WLE provided partial definition for the storage of the Virtual I/O Server partition. If you want to use a storage pool, you need only to add physical volumes or SAN volumes to provide the backing devices for the assigned virtual storage. If you do not wish to use the storage pools, simply delete the storage pool and configure the allocation of the physical volumes or SAN volumes directly to the client partition as desired. Complete the information for the SAN and volume definition. See 4.3, “Define Virtual I/O Server in System Planning Tool Version 2” on page 106, for further explanation of the values.

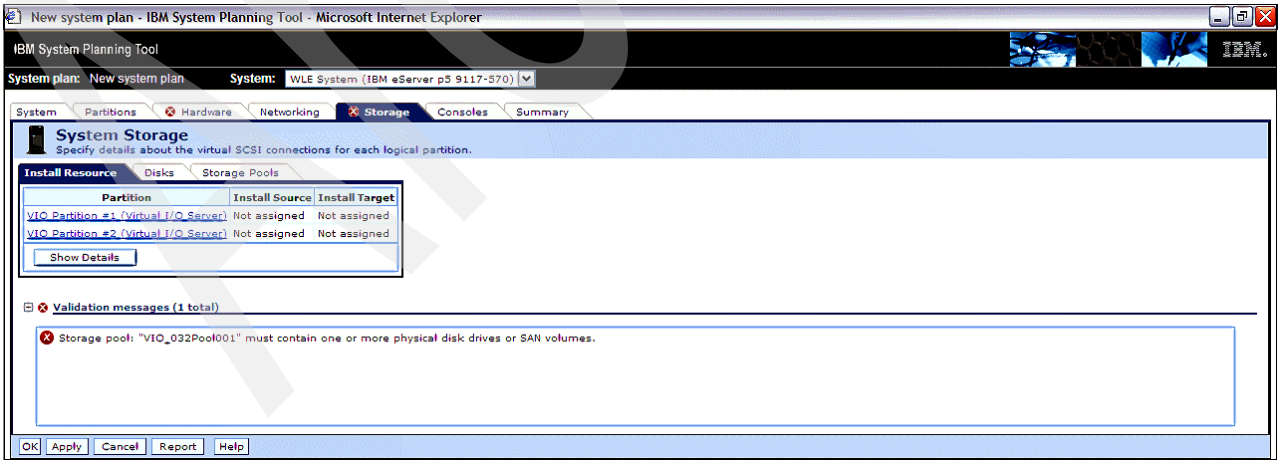

Figure 5-41 Storage definition

A

Managing System Planning Tool Version 2 installation

This appendix describes the steps necessary to download, install, start, end, and uninstall the IBM System Planning Tool Version 2.

We also provide a specific note about setup that may be useful.

Downloading the SPT installation files

Attention: File names, sizes, and patch levels are examples, and real ones will probably be different than those found in this appendix.

The IBM System Planning Tool and documentation are available for download at:

http://www-304.ibm.com/jct01004c/systems/support/tools/systemplanningtool

Installing the SPT is a two-step process:

1. The first file to install contains the base software, including the necessary Java™ code. Its name contains the word *setup*.
2. The second file to install contains all of the cumulated updates to the base code. Its name contains the word *patch*.

Both files must be downloaded. On the SPT Web site (see Figure A-1 and Figure A-2 on page 191) scroll down to the Getting started section, and click each of the links to the files.

Figure A-1 The SPT Web site

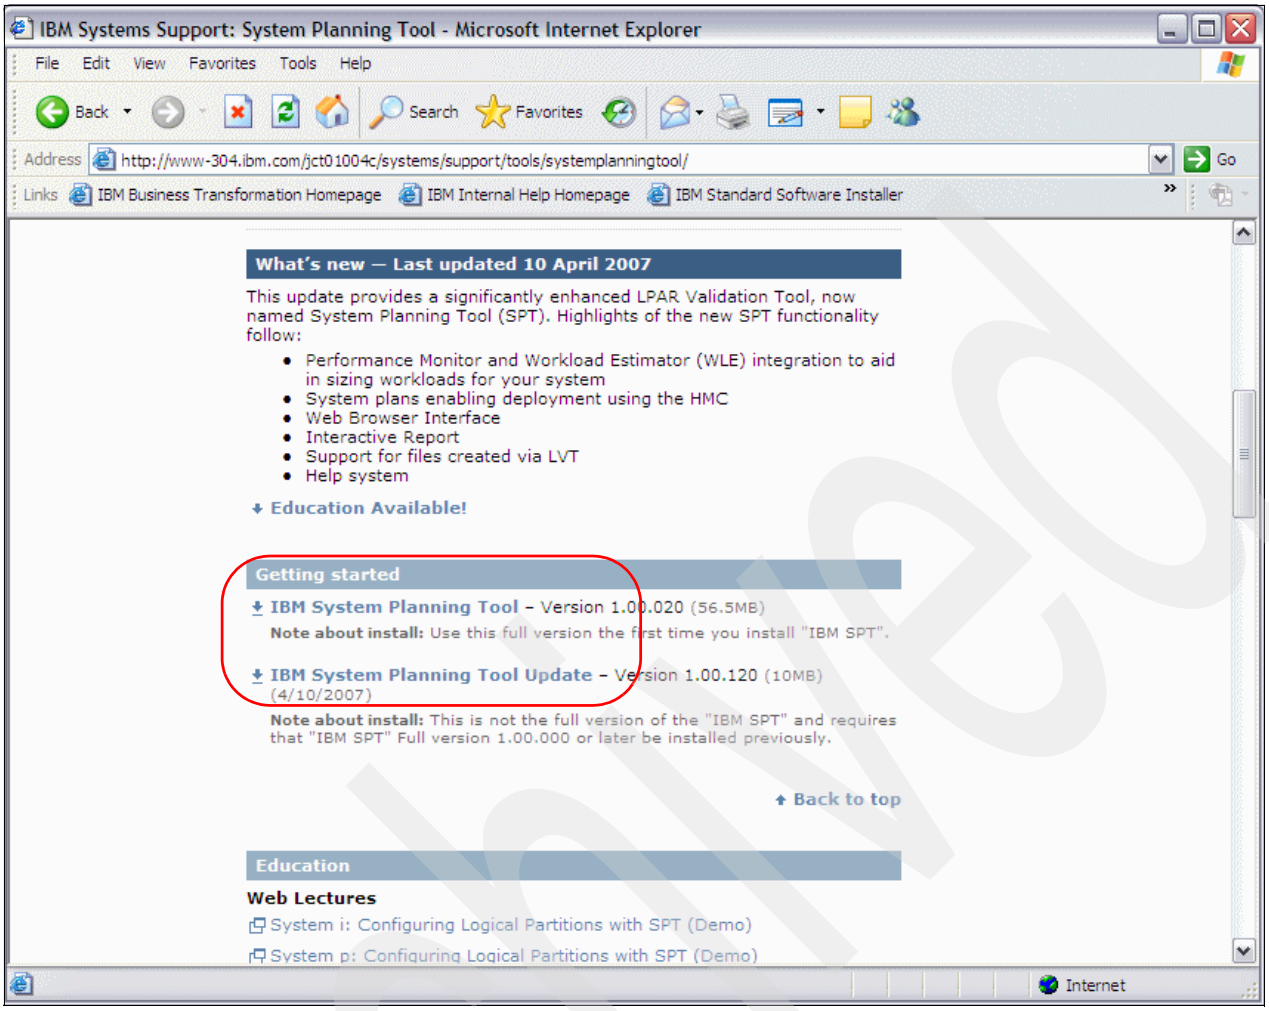

Figure A-2 Getting started section

Each one will request you to accept the license in order to start the download, as shown in Figure A-3. Scroll down the window and click **I agree** to download the file onto your PC.

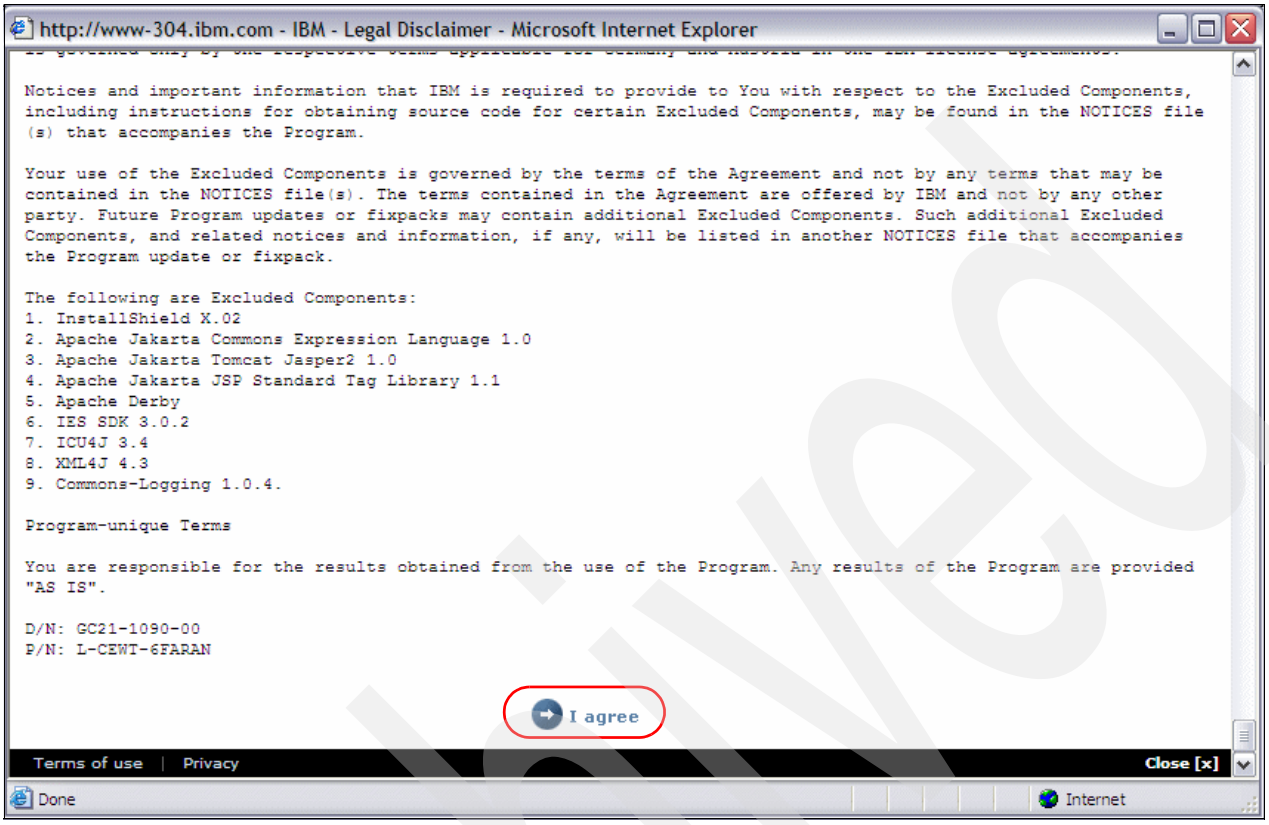

Figure A-3 Download license agreement

Save the downloaded files in a convenient folder on your computer, as shown in Figure A-4 and Figure A-5 on page 193.

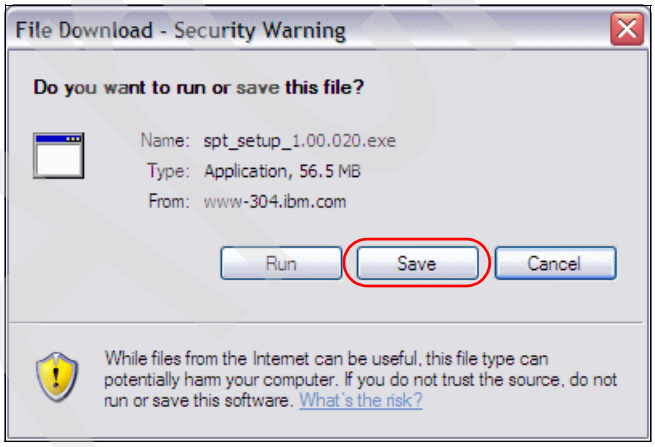

Figure A-4 Save the downloaded files

In our example, we use the temp folder.

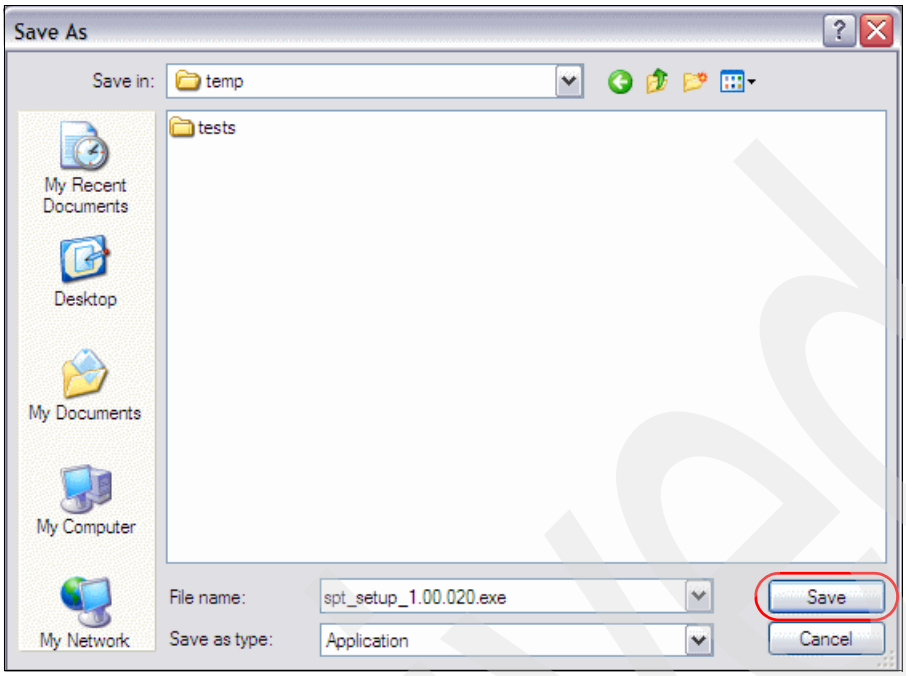

Figure A-5 Select a convenient folder

Make sure that the file sizes are correct after downloading. This confirms that the download completed successfully.

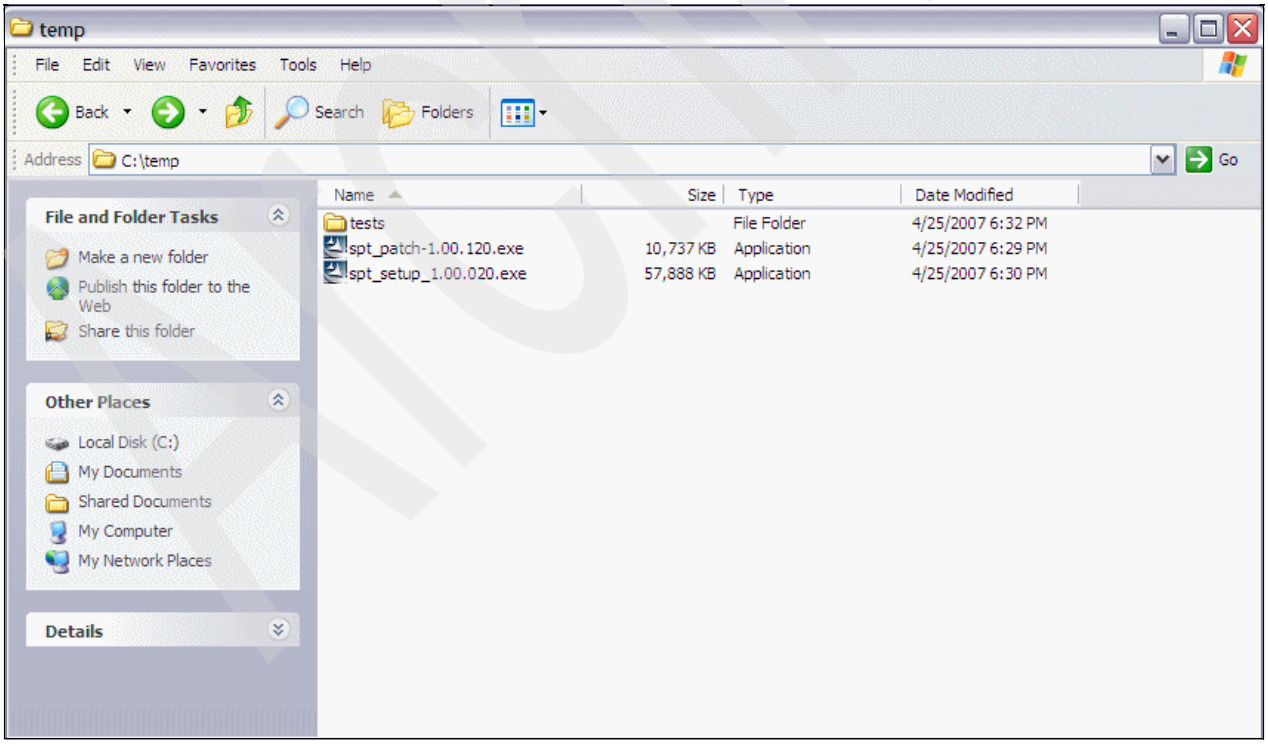

Figure A-6 The downloaded files

Tip: If you want to receive a notice each time a new patch is available, scroll down to the bottom of the Web page (see Figure A-1) to find the contact information section. You see a link in the Receive notification paragraph to send a predefined e-mail to a dedicated mailbox. This link is rchspt@us.ibm.com. Clicking this link from the Web site brings you directly in your mail software and you just have to send the e-mail, as it is automatically built with all necessary information.

Installing the System Planning Tool Version 2

Note: Depending on your PC general settings, at various steps during the installation process, you may receive alerts from your firewall software. Use the related procedures to allow any connection required by System Planning Tool Version 2.

Normal installation procedure

It is possible to install V2 over V1 or from scratch. In both cases the first file to run is spt_setup_xxxx.exe.

1. Double-click this file.

   Three panels are successively displayed, for which no action is needed.

2. The Preparing JVM display appears (Figure A-7).

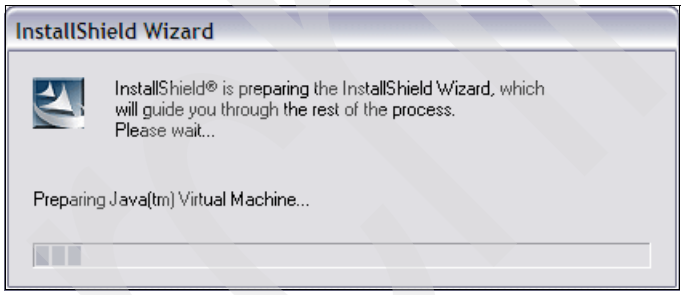

Figure A-7 Preparing JVM

3. The SPT splash window appears (Figure A-8).

Figure A-8 The SPT splash window

4. The Initializing wizard message appears (Figure A-9).

Figure A-9 The initializing wizard message

5. The wizard requests your first action, as shown in Figure A-10. Click **Next**.

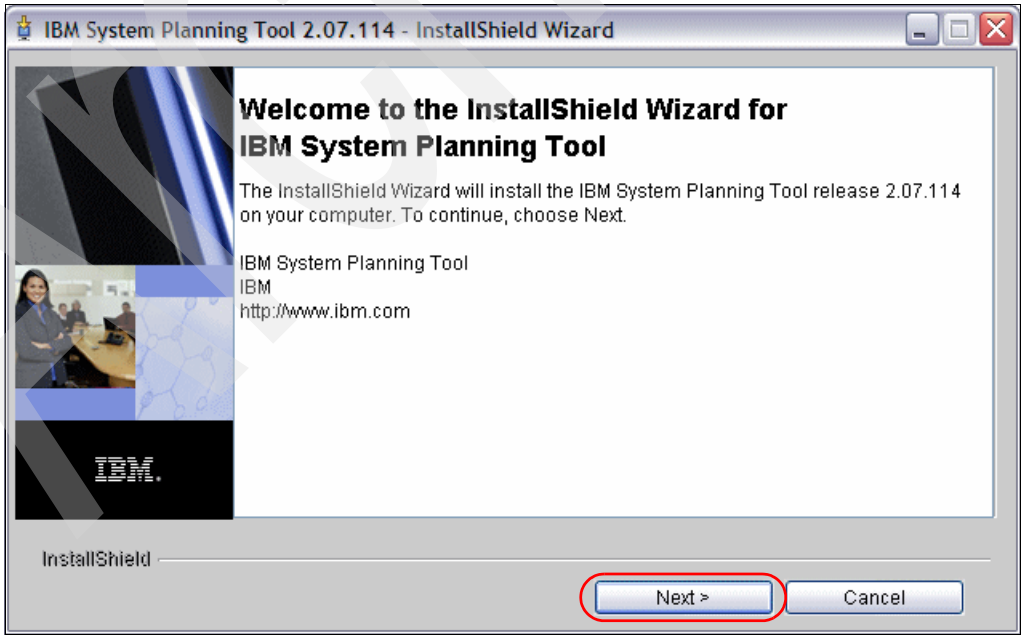

Figure A-10 Start the installation process

6. As shown in Figure A-11, you have to read the license agreement and accept the terms to continue. Select **I accept** and click **Next**.

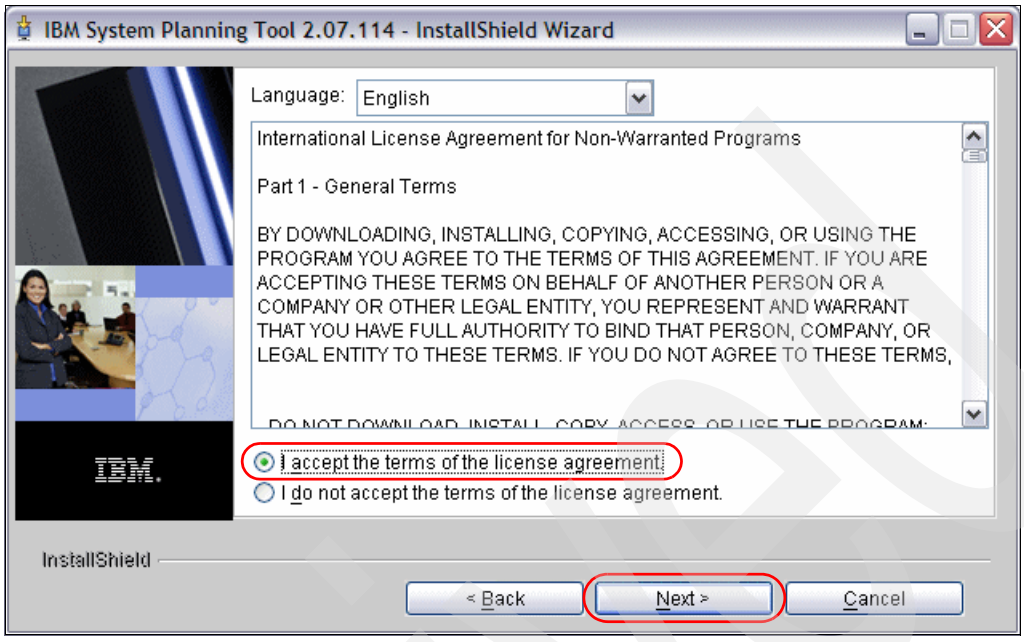

Figure A-11 Accept the license agreement

If System Planning Tool Version 1 exists on your PC, the wizard tells you that it needs to uninstall it before running the System Planning Tool Version 2 installation process. In our example we have System Planning Tool Version 1 installed. Click **Next** to accept the uninstall process (see Figure A-13 on page 197), as shown in Figure A-12.

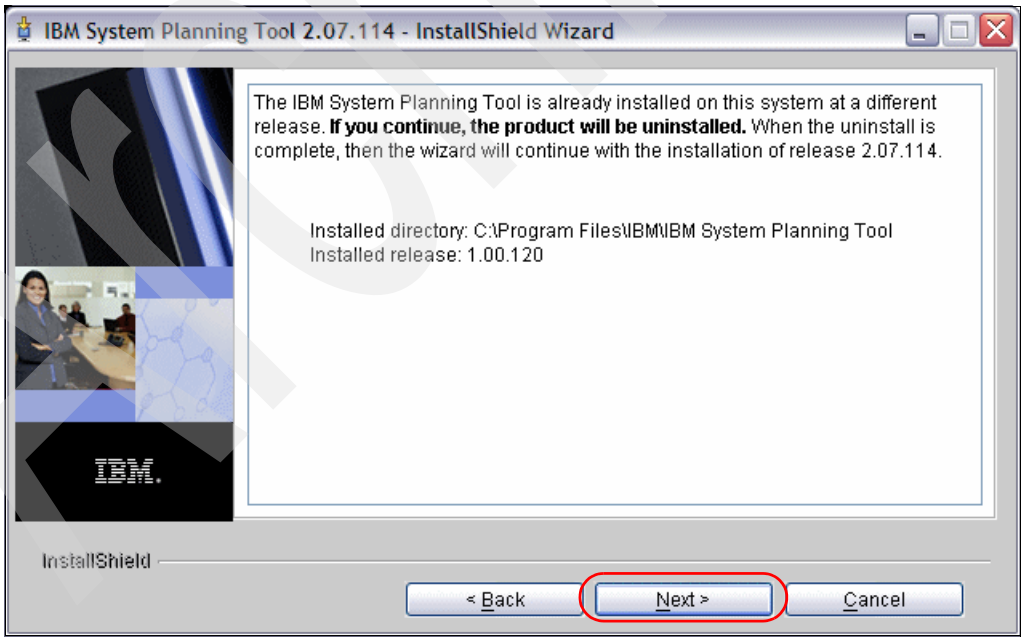

Figure A-12 Confirm System Planning Tool Version 1 uninstallation

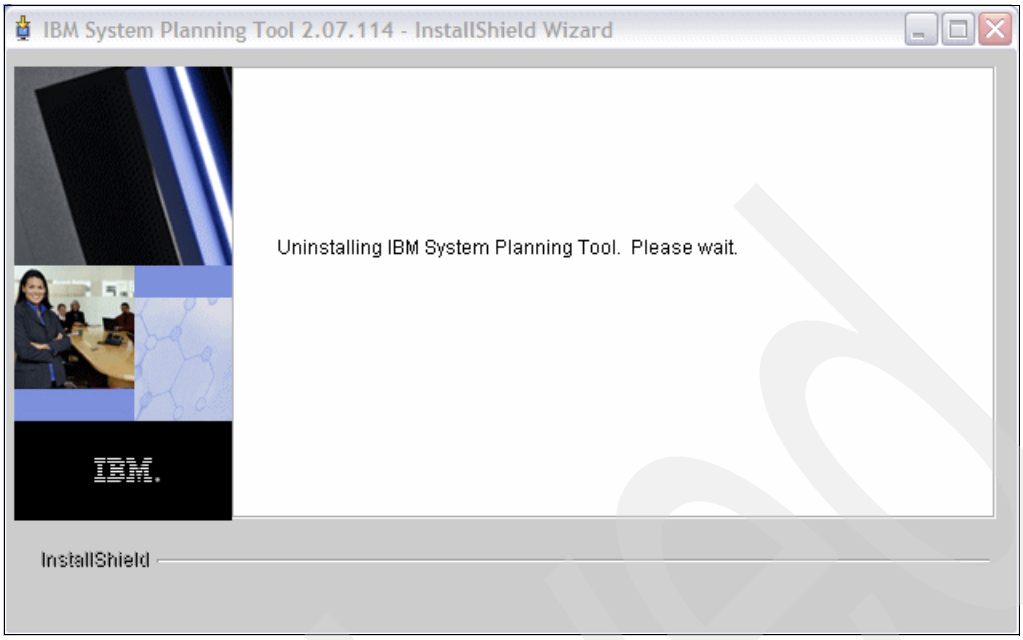

Figure A-13 System Planning Tool Version 1 uninstallation in progress

7. With or without the uninstall process, the next window presents the location in which the new files are to be installed. We suggest that you accept the default folder, and then click **Next** to continue, as shown in Figure A-14.

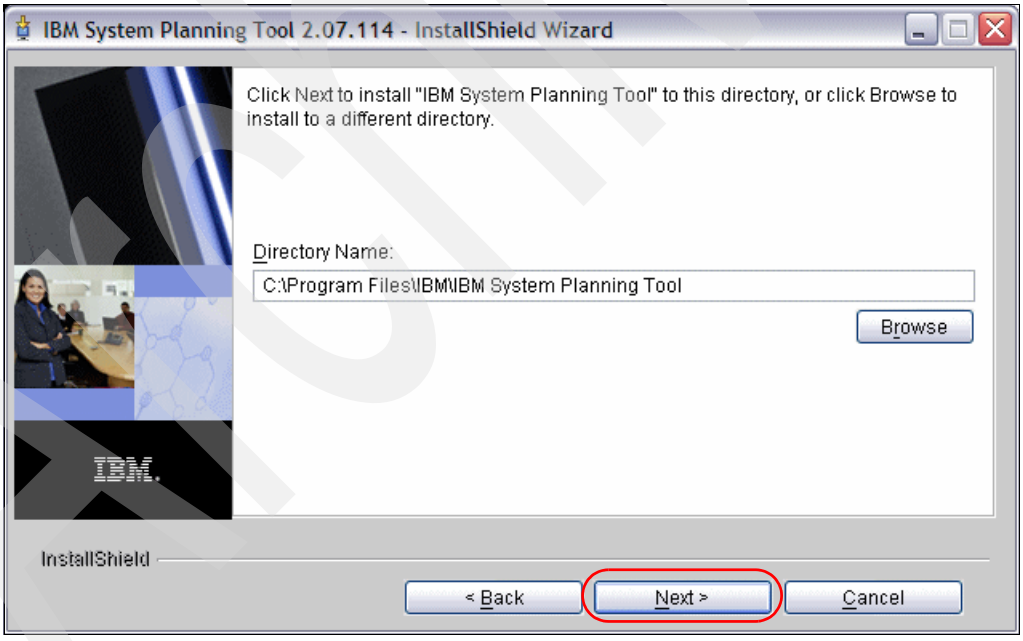

Figure A-14 Installation folder

8. The next step is to select an available IP port number on your PC for the System Planning Tool Version 2 daemon to listen on (see “Setting up the System Planning Tool Version 2” on page 208 for more details). At this time, you can accept the proposed value and click **Next** to continue, as shown in Figure A-15.

Figure A-15 IP port setting wizard

9. The last window requesting an action summarizes the folder installation and the installed features and their size. Click **Install**, as shown in Figure A-16, to confirm the installation procedure.

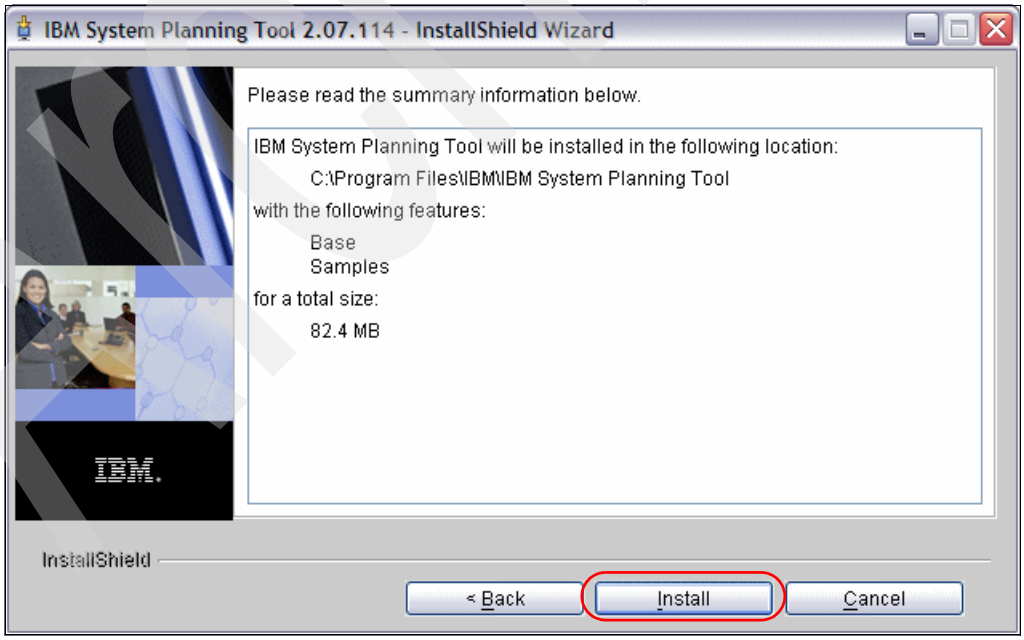

Figure A-16 Installation confirmation

During the installation process, you will see the various files being extracted and copied into the installation folder, as shown in Figure A-17.

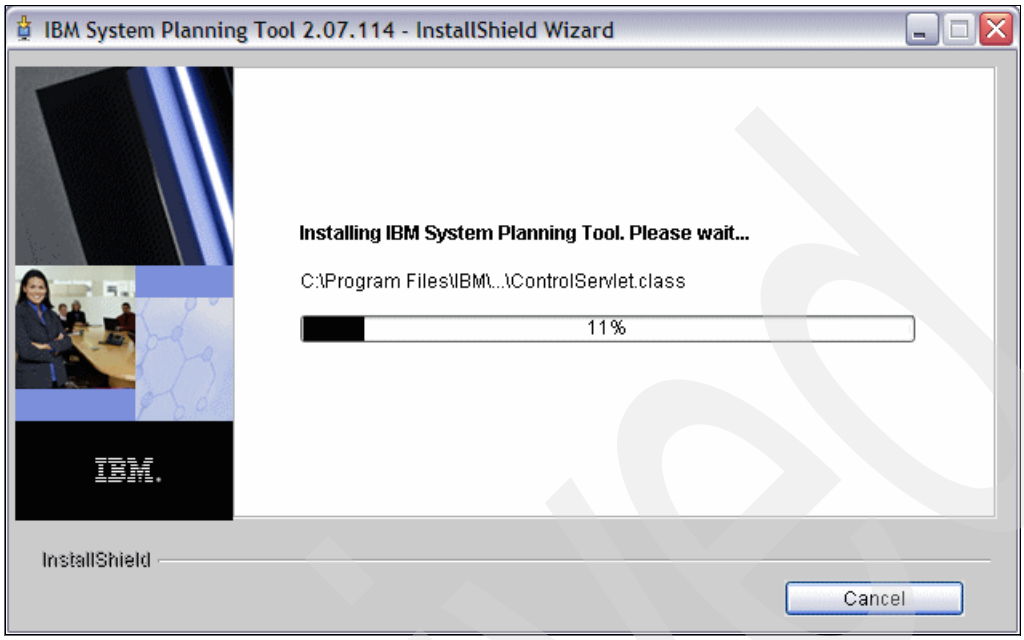

Figure A-17 Extracting and copying files

10. When the System Planning Tool Version 2 base code has been installed (see Figure A-18), you need to apply the patch. Therefore, uncheck the Start the IBM System box and click **Finish** to end this first step, the install of System Planning Tool Version 2 base code.

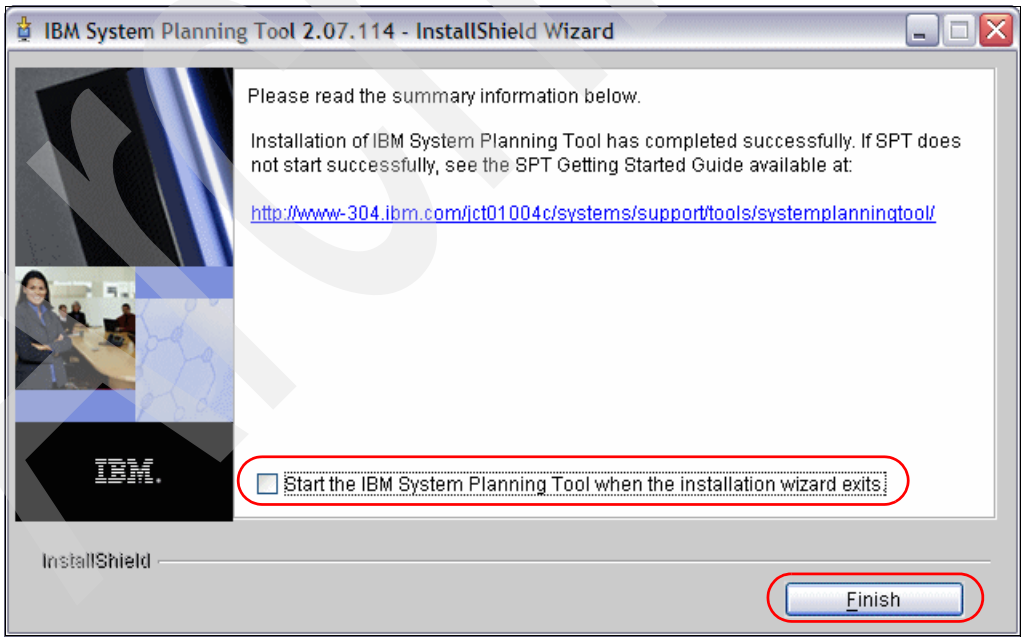

Figure A-18 Installation is finished

Now you have to install the patch file.

Important: This is an operation that you may have to perform several times, as patches are produced on a regular basis. If you do not subscribe to the information letter, you may need to manually check the download Web page several times a month.

The steps are:

1. Go back to your download directory (see Figure A-6 on page 193), temp in our example, and double-click the patch file, **spt_patch_xxxx.exe**.

   The process is generally the same as that described for the setup file installation. You will see the same first three panels that do not require any input (see Figure A-7 on page 194, Figure A-8 on page 195, and Figure A-9 on page 195). The next display requires you to activate the real installation procedure, as shown in Figure A-19. You can verify the patch level in the text. Click **Next** to start.

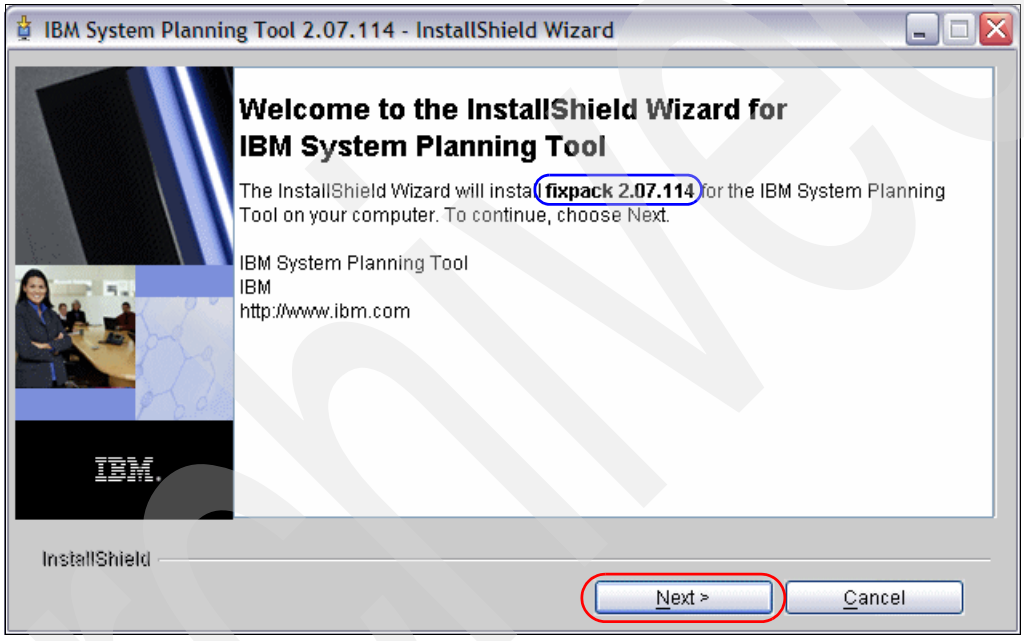

Figure A-19 Start the patch installation

2. The last window summarizes the installed features and their size and the installation folder, as shown in Figure A-20. Click **Install**.

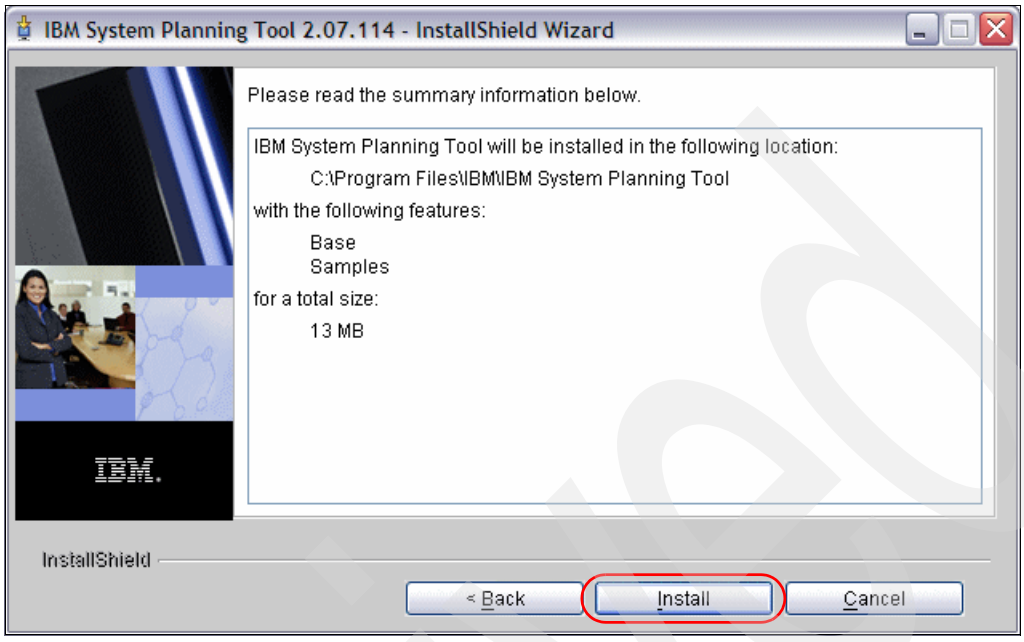

Figure A-20 Confirm the patch installation

3. The same progress bar as for the setup file is displayed when extracting and copying the files (see Figure A-17 on page 199). The same final display leaves the choice to immediately start System Planning Tool Version 2 when clicking **Finish**, as you can see in the Figure A-21.

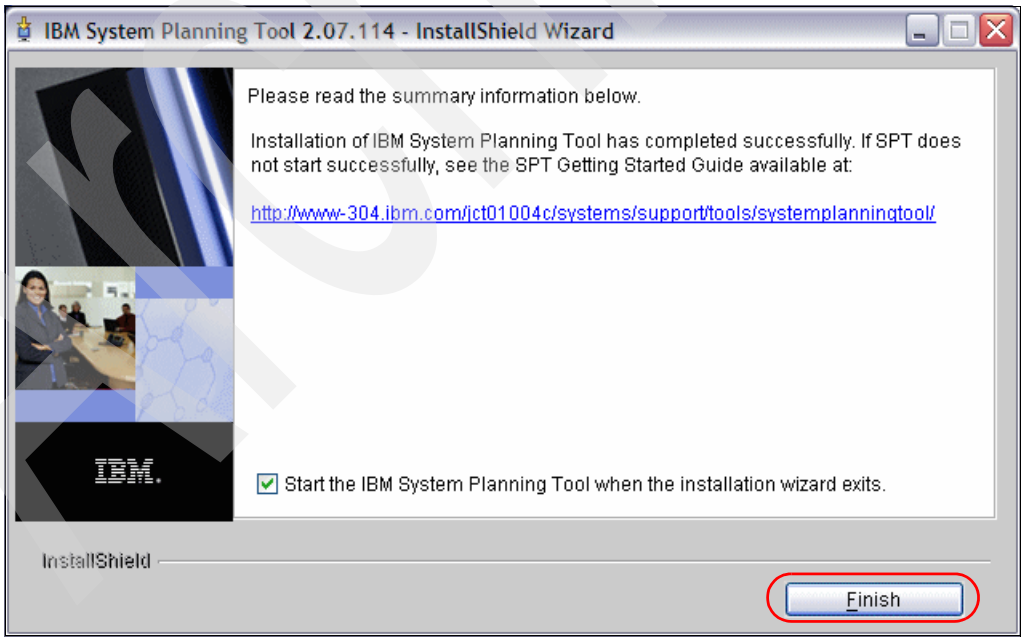

Figure A-21 Patch installation is finished

Wrong installation procedure

If you run the patch first (file name spt_patch_xxx.exe), depending on whether the previous release is installed, you will get the following warnings, and the better solution is to cancel the installation to start it again with the correct file (the setup file):

- *If V1 is already installed*, the SPT splash window and initialization wizard message are displayed, and then the installation process halts, as shown in Figure A-22.

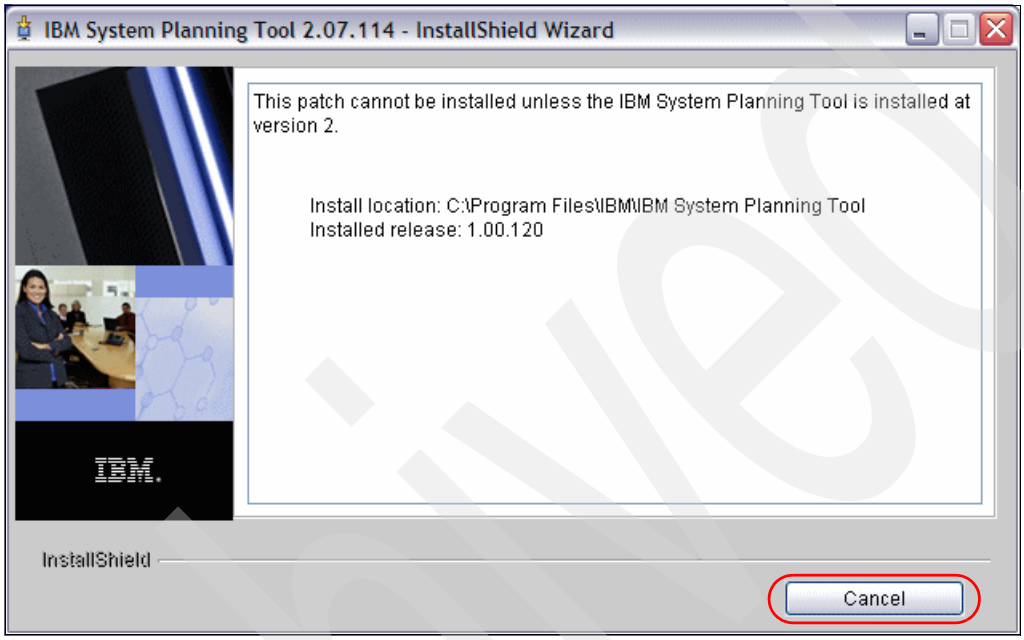

Figure A-22 Unable to install patch

- *If V1 is not installed*, the wizard warns you that suitable JVM files are missing (see Figure A-23, where you have no choice but to click OK) and requests you to specify a location where it may find what it needs (see Figure A-24 on page 203). The best way is to cancel here.

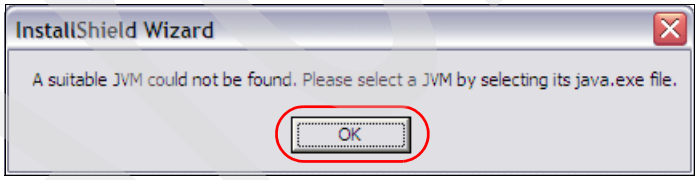

Figure A-23 JVM cannot be found

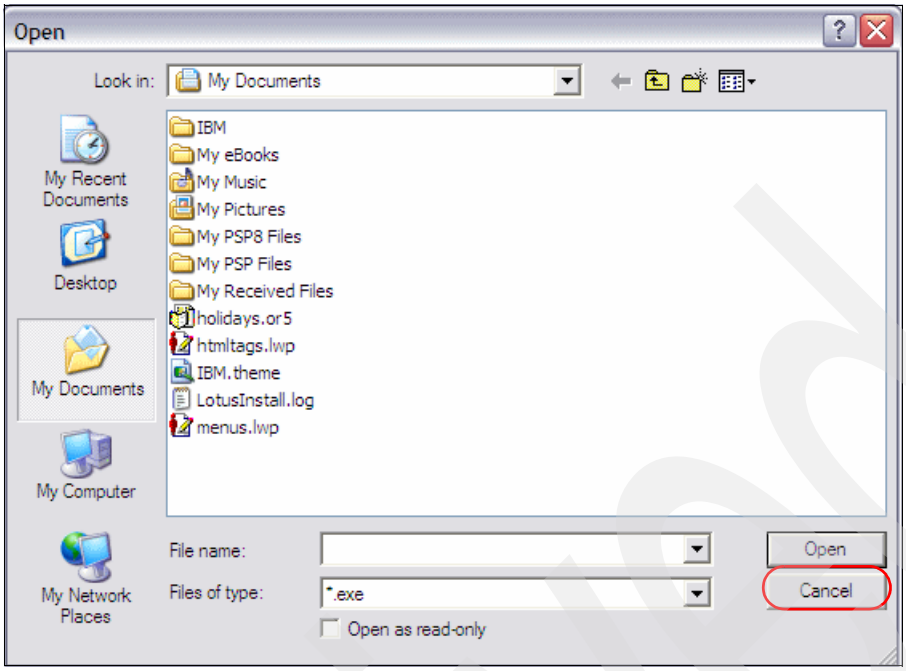

Figure A-24 Location to look for the JVM files

Browser cache special considerations

As you have seen when using SPT in Chapter 2, “System Planning Tool (SPT) V2” on page 11, in reality it is a server-like application running on your PC, to which you connect with a browser. Any browser has caching capabilities. It means that some of the HTML pages or JavaScript™ files used during previous SPT sessions may be stored in the cache directory of the browser and used in place of new installed ones.

If it happens, and it is rather unpredictable, you may receive errors or warnings, or experience incoherent behavior of SPT.

Note: The best way to avoid these problems is to clear the cache data of the browser after installing a new version or new patch of SPT.

Here are two examples of ways to clear cache data for the most common browsers.

- For Internet Explorer® 6 (IE)

  a. When using IE, to clear the cache data, select **Tools** → **Internet options** and make sure that the General tab is active. Then, as shown in Figure A-25, click **Delete files**.

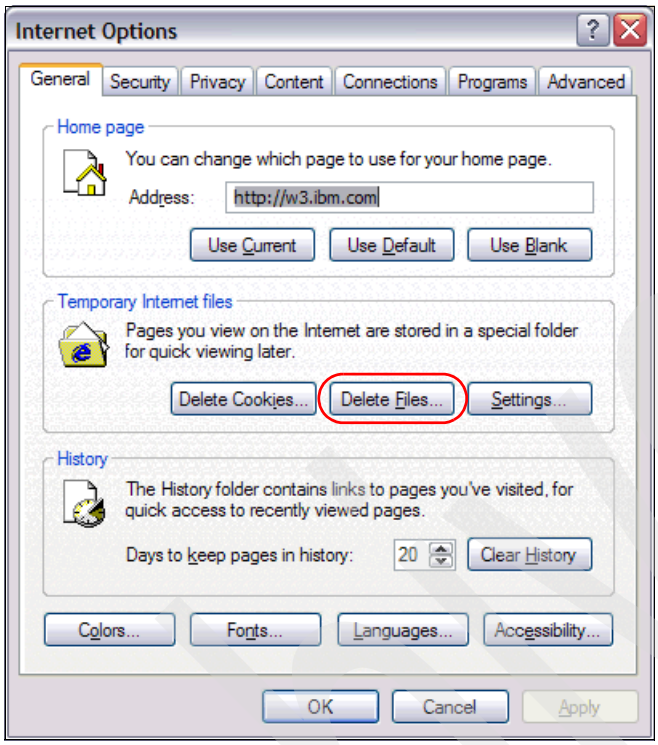

Figure A-25 Access to IE cache options

  b. To confirm, as shown in Figure A-26, select **Delete all offline content** and click **OK**.

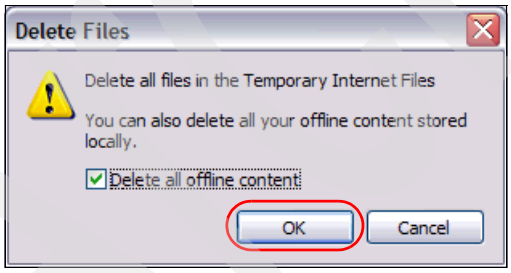

Figure A-26 Confirm delete cache data with IE

- For Mozilla Firefox (FF) (at least up to 2.0.0.3 release)

  When using FF, to clear cache data, select **Tools** → **Clear Private Data**. Then, as shown in Figure A-27, uncheck all but the Cache boxes and click **Clear Private Data Now**.

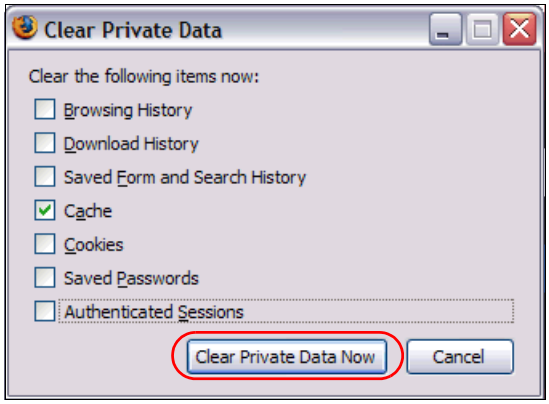

Figure A-27 Clear cache data with FF

Starting the System Planning Tool Version 2

You have three ways to start the System Planning Tool Version 2:

- As shown in Figure A-28, double-click the **IBM System Planning Tool** icon that was created during the installation procedure.

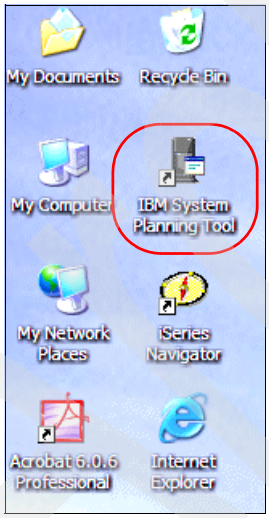

Figure A-28 Double-click the IBM System Planning Tool icon

- As shown in Figure A-29, look for the command to start in the programs menu. Click **Start → All programs → IBM System Planning Tool → System Planning Tool**.

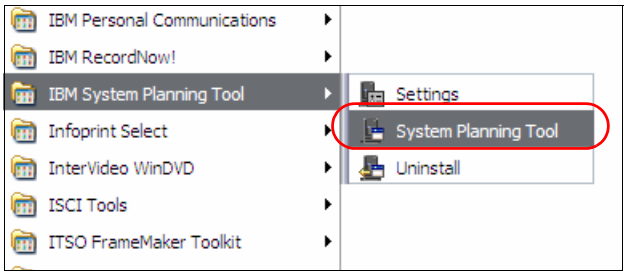

Figure A-29 System Planning Tool Version 2 start command

It takes a few seconds for the initialization to complete until you get the browser window displayed in Figure A-30.

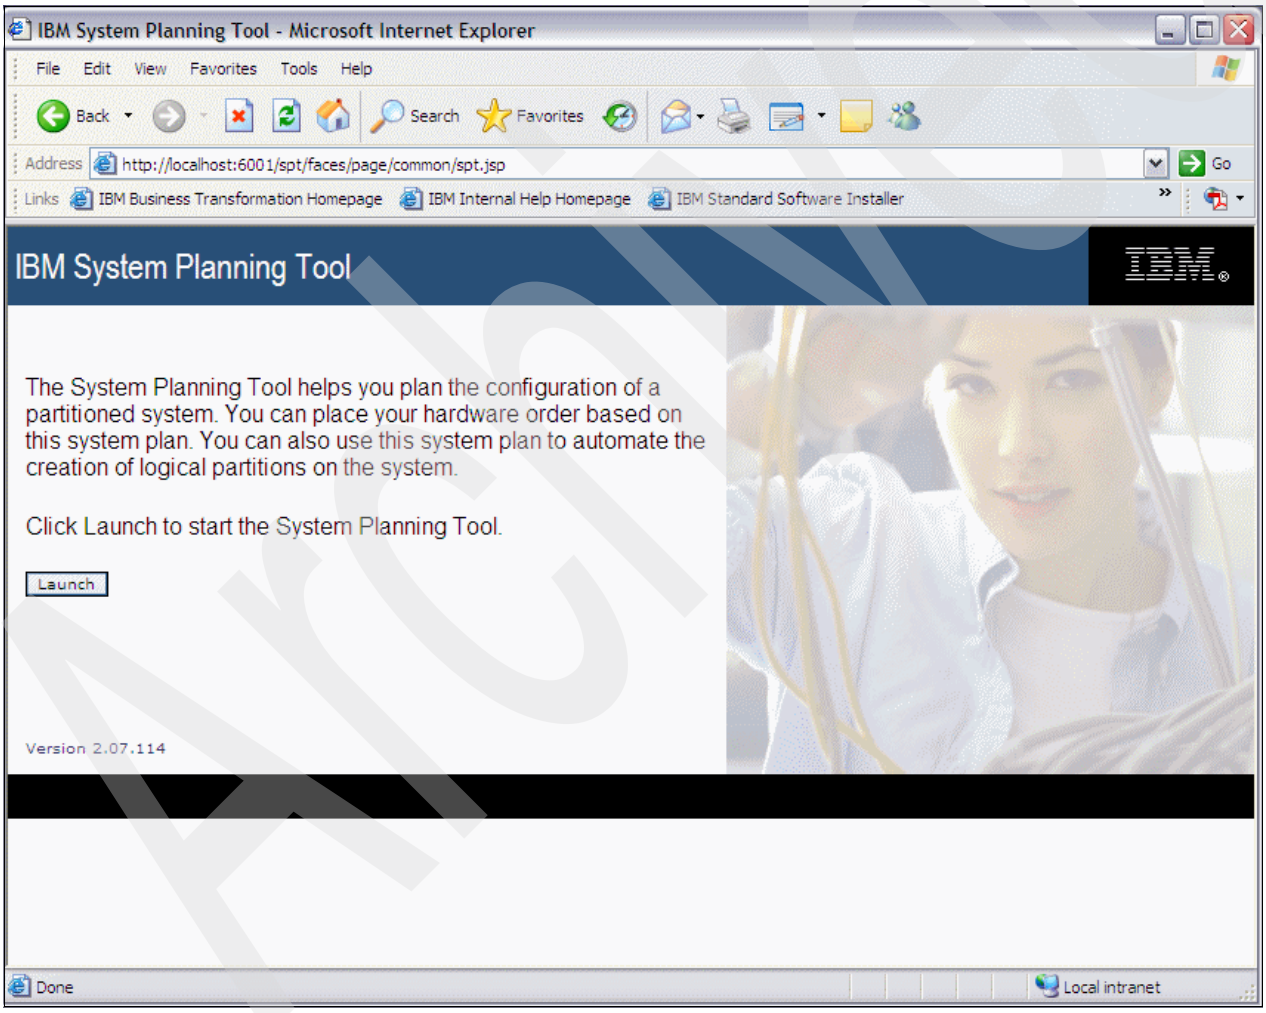

Figure A-30 The System Planning Tool Version 2 starting window

Also, notice the icon in the system tray, as shown Figure A-31. As long as this icon exists, it means that the System Planning Tool Version 2 is active, even if any previously connected browser window is closed.

Figure A-31 The activity icon

- Right-click the system tray icon and select **Open**, as shown in Figure A-32.

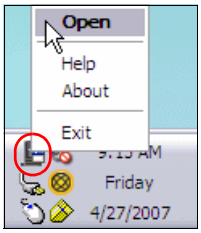

Figure A-32 Open SPT using the system tray icon

The daemon of System Planning Tool Version 2 fails to start if the IP port number it should use is already open by another software. In this case, you receive an alert, as shown in Figure A-33. Click **Cancel**, set up another IP port number (see “Setting up the System Planning Tool Version 2” on page 208), and retry the procedure.

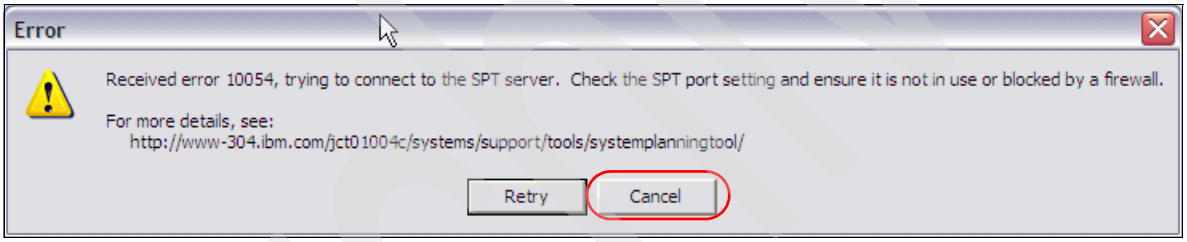

Figure A-33 Error when starting SPT

Ending the System Planning Tool Version 2

Right-click the system tray icon and select **Exit**, as shown in Figure A-34.

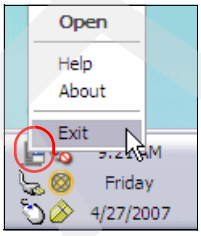

Figure A-34 Ending SPT

Setting up the System Planning Tool Version 2

There is only one parameter that needs to be setup to run System Planning Tool Version 2. It is the IP port number used by the daemon to listen on. Any browser window related to the tool will connect to the daemon (using localhost IP name) to that particular port.

During installation time, an available port above the well-known ones (up to 1023) is selected by the wizard.

This port may be suitable all the time. But, you may have other software on your PC that also uses this particular port. If System Planning Tool Version 2 is active, this other software may fail to start. On the other hand, if this other software is running when you try to start SPT, this one will fail.

To set up this port, run **Start** → **All programs** → **IBM System Planning Tool** → **Settings**, as shown in Figure A-35.

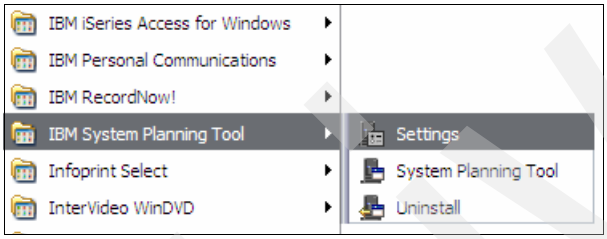

Figure A-35 Start SPT setup

As shown in Figure A-36, you can type the port you want and click **OK**. Note that the availability of the IP port number is not checked at this point.

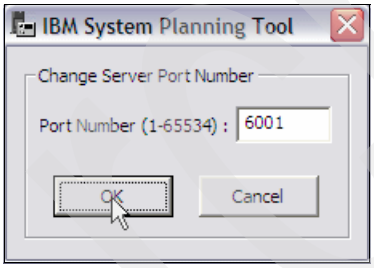

Figure A-36 SPT IP port number setting

The next information window (see Figure A-37) just clarifies that the new IP port number takes affect the next time System Planning Tool Version 2 starts. Click **OK**.

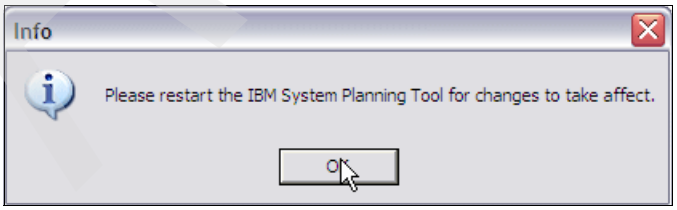

Figure A-37 Information about the IP port number

Tip: High IP port numbers are rarely used, so you may want to set up numbers like 65534, 65533, and so on.

Uninstalling the System Planning Tool Version 2

Prior to uninstall the System Planning Tool Version 2, make sure that it does not run (check “Ending the System Planning Tool Version 2” on page 207 if necessary). Then click **Start** → **All Programs** → **IBM System Planning Tool** → **Uninstall**, as shown in Figure A-38.

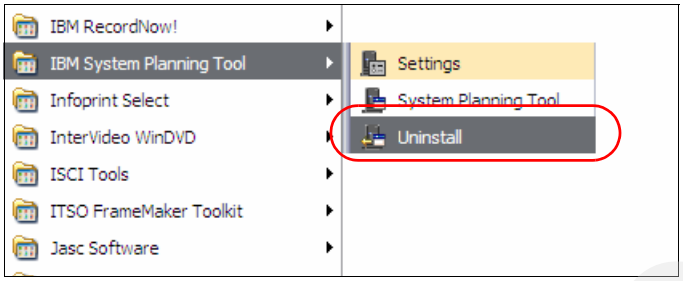

Figure A-38 Uninstall command

The SPT splash window, then the Initializing wizard message appear, as shown in Figure A-39 and Figure A-40.

Figure A-39 The SPT splash window

Figure A-40 The initializing wizard message

At this time, on the first uninstall wizard, you can still cancel the uninstall by clicking **Cancel**, as you can see in Figure A-41. Click **Next** to continue the procedure.

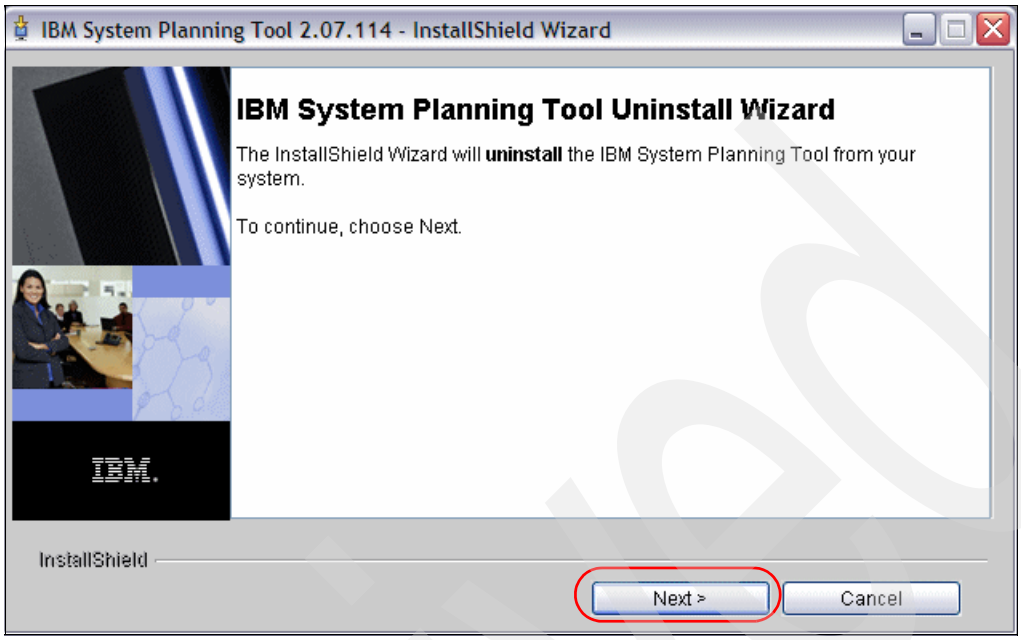

Figure A-41 First uninstall wizard panel

The Confirm uninstall panel shows you which features will be uninstalled and which folder is involved in the procedure. See Figure A-42. Click **Uninstall** to continue the procedure.

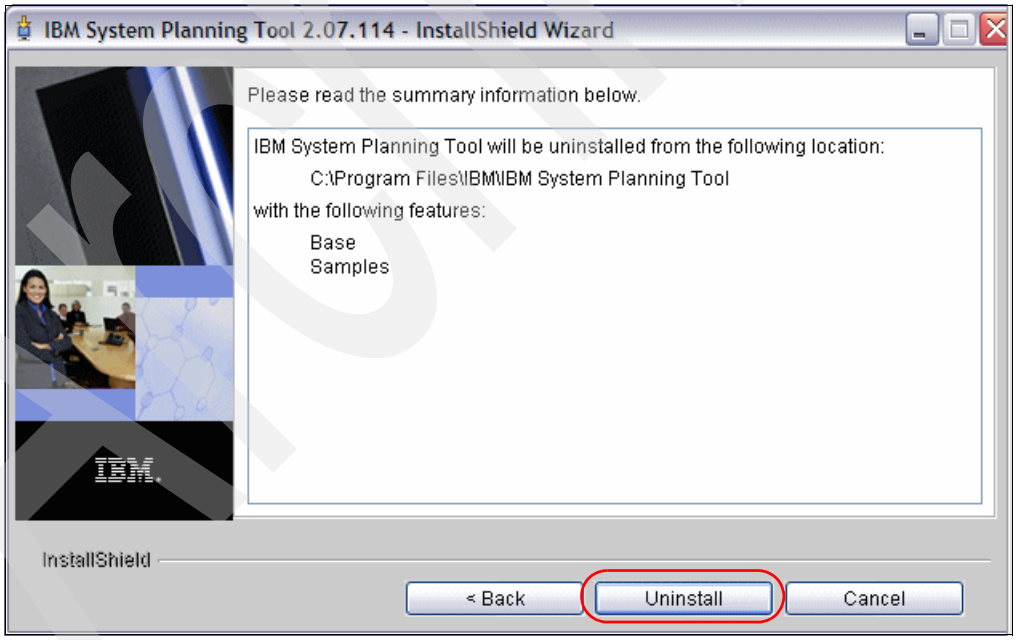

Figure A-42 Confirm uninstall panel

Until the end of the procedure, the work in progress panel is displayed, as shown in Figure A-43.

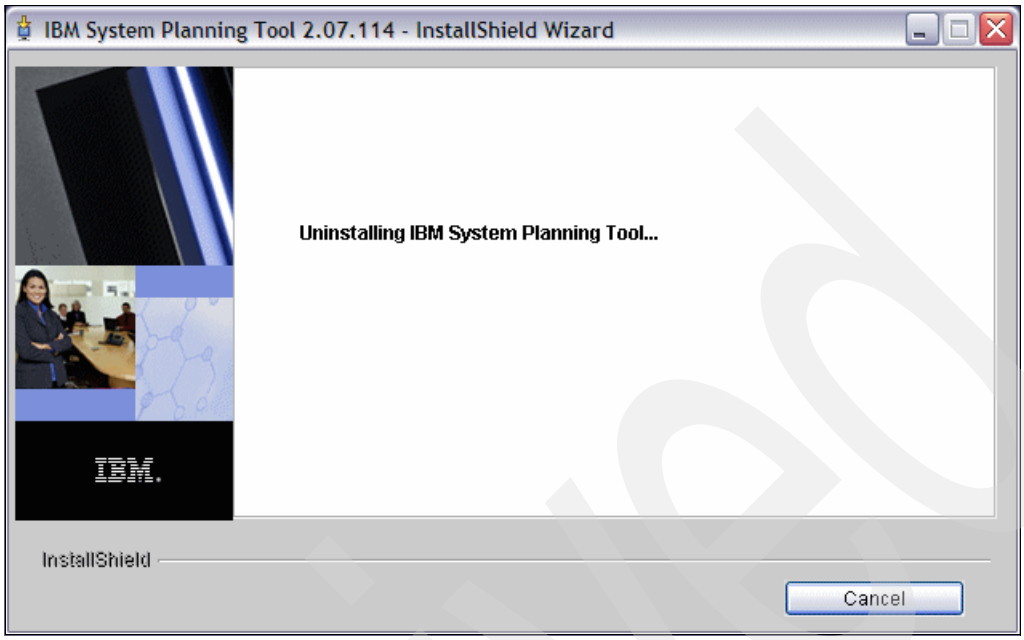

Figure A-43 Uninstall in progress

You are now presented with the final panel (Figure A-44) to confirm the removal of the System Planning Tool Version 2. Click **Finish** to exit the wizard.

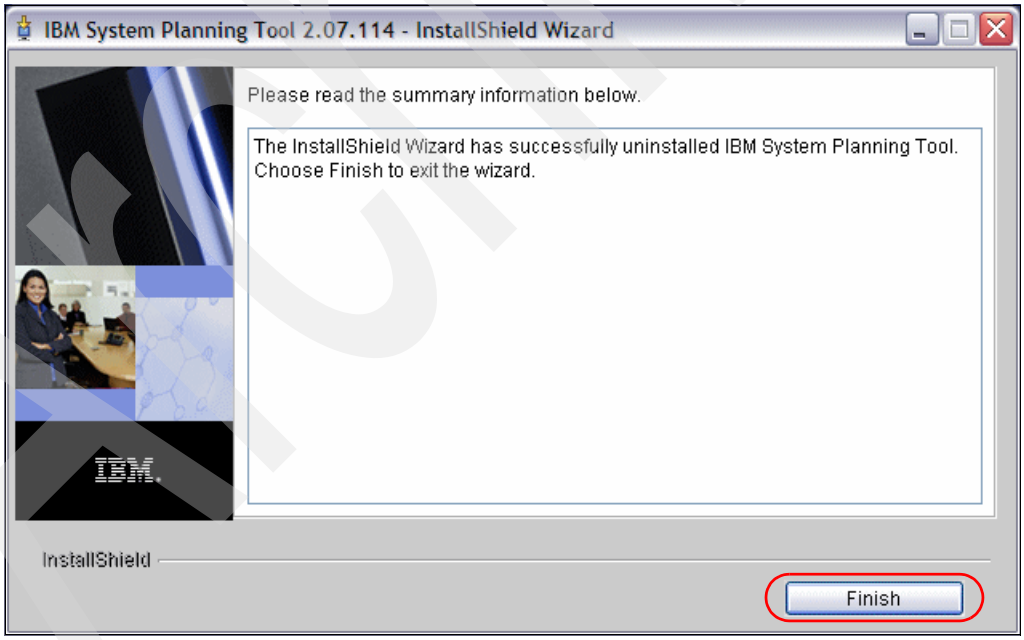

Figure A-44 Uninstall complete

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 215. Note that some of the documents referenced here may be available in softcopy only.

- *Hardware Management Console V7 Handbook*, SG24-7491
- *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration*, SG24-7940
- *Advanced POWER Virtualization on IBM System p Virtual I/O Server Deployment Examples*, REDP-4224
- *Integrated Virtualization Manager on IBM System p5*, REDP-4061
- *IBM System i Overview: Models 515, 525, 570, 595 and More*, REDP-5052
- *PCI and PCI-X Placement Rules for IBM System i5, eServer i5, and iSeries servers with i5/OS V5R4 and V5R3, REDP-4011*
- *IBM System p Advanced POWER Virtualization Best Practices,* REDP-4194x
- *HACMP 5.3, Dynamic LPAR, and Virtualization*, REDP-4027
- *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000
- *LPAR Simplification Tools Handbook*, SG24-7231
- *WebSphere Message Broker V6, Best Practices Guide: Bullet Proofing Message Flows*, REDP-4043

Other publications

These publications are also relevant as further information sources:

- *System p and System i Site and Hardware Planning Guide* SA76-0091
- *IBM System p Overview,* SA76-0087
- *System p Logical Partitioning Guide*

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System i and System p Partitioning for AIX with an HMC*

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System i and System p Partitioning for i5/OS with an HMC*

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System i and System p Partitioning for Linux with an HMC*

  This PDF focuses on managing POWER5 and POWER6 technology systems.

- *System p Partitioning with the Integrated Virtualization Manager*

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System i and System p Using the Virtual I/O Server*

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System p Installation and Configuration Guide for the Hardware Management Console Version 7 Release 3.1.0 Maintenance Level 0,* SA76-0084

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System p Operations Guide for the Hardware Management Console and Managed Systems Version 7 Release 3.1.0,* SA76-0085

  This PDF focuses on managing POWER5 and POWER6 technology systems.
- *System i and System p Managing the Hardware Management Console (HMC)*

  This PDF focuses on managing POWER5 technology systems.

Online resources

These Web sites are also relevant as further information sources:

- IBM Systems Hardware Information Center

  This Web site is focused primarily on POWER5 technology hardware on System i and System p. Pre-POWER5 technology hardware can be found at the respective pSeries and iSeries informations centers listed later in this list.

  POWER6 technology hardware information can be found as described later in this list.

  `http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp`
- iSeries Information Center - System i Information Center

  This Web site is focused primarily on pre-POWER5 technology hardware on System i and all i5/OS-based software documentation (V5R2, V5R3, V5R4). This includes i5/OS-based software that runs on POWER5 and POWER6 technology configurations.

  `http://publib.boulder.ibm.com/iseries/`
- pSeries Information Center - System p Information Center

  This Web site is focused primarily on pre-POWER5 technology hardware on System p and all AIX-based software documentation. This includes AIX-based software that runs on POWER5 and POWER6 technology configurations.

  `http://publib16.boulder.ibm.com/pseries/index.htm`
- IBM Systems Support Web site for System p

  This support Web site offers a wide range of tools and resources, including hardware and software maintenance (for example, Fix Central capabilities) and technical documentation. For accessing POWER6 hardware information, this site provides the most effective links to available documents such the System p Overview (POWER6), and HMC documentation specific to Version 7 Release 3.1.0.

  We recommend going to this Web site first to find the PDF documents listed in this publication under “Online resources” on page 214.

  This site also has links to the other information centers listed here.

  `http://www.ibm.com/systems/support/p`

Click the **Product documentation** link to access documentation grouped under several categories (libraries), such as:

- Systems Support Site: POWER6 specific documentation
- Systems Information Center:
  - Hardware - POWER6 and POWER5 and Hardware Management Console
  - Software - AIX 5L V5.3, i5/OS V5R4, Linux on System p
- pSeries and AIX Information Center
  - Hardware - POWER4, RS6000, Hardware Management Console
  - Software - AIX 5L V5.2, V5.1, AIX 4.3

- IBM Systems Support Web site for System i

  This support Web site offers a wide range of tools and resources, including hardware and software maintenance (for example, Fix Central capabilities) and technical documentation. For accessing POWER6 hardware information, this site provides the most effective links to available documents such the System i Overview (POWER6), and HMC documentation specific to Version 7 Release 3.1.0.

  We recommend going to this Web site first to find the PDF documents listed in this publication under “Online resources” on page 214.

  This site also has links to the other information centers listed here.

  `http://www.ibm.com/systems/support/i`

How to get Redbooks

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

ibm.com/redbooks

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services

IBM Prerequisites Web site. This site provides compatibility information for hardware features. This tool helps you to plan a successful system upgrade by providing you with the prerequisite information for features you currently have or plan to add to your system:

`http://www-912.ibm.com/e_dir/eServerPrereq.nsf/`

Index

L

M

N

O

P

R

S

T

U

V

W

IBM System i and System p System Planning and Deployment: Simplifying Logical Partitioning

Using the enhanced System Planning Tool

Deploying with the HMC V7 interface

Planning and deploying a Virtual I/O Server

Logical Partitioning (LPAR) provides the significant capability to run multiple operating systems, each a partition on the same physical processor, memory, and I/O attachment configuration. LPAR is often discussed along with the concept of server consolidation. LPAR enables management across a single set of hardware and, when configured and managed correctly, can maximize efficient use of hardware resources all in a single place, often using resources in one partition when not needed by another partition.

By its nature, LPAR is powerful, but, as the number and complexity of applications being run in each partition increases, can become complex to configure and to achieve anticipated performance expectations.

This IBM Redbooks publication describes and provides examples of using the 2007 enhancements to the system planning and deployment tools and processes for planning, ordering, and deploying a partitioned environment on IBM System i and IBM System p configurations.

The objective is to help you order and IBM deliver a hardware configuration and get that configuration up and running your planned partition configurations with good performance in as short a time as possible.

This book and the tools and processes involved represent the next step in expediting this entire process, while still requiring sound knowledge of IBM System i and System p hardware processor and I/O capabilities for success.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7487-00 ISBN 0738488771