IBM

LPAR Simplification Tools Handbook

- Review the new System Planning Tool for IBM System i and IBM System p
- Understand how to use the LPAR deployment wizard
- See how easy it is to design simple logically partitioned systems

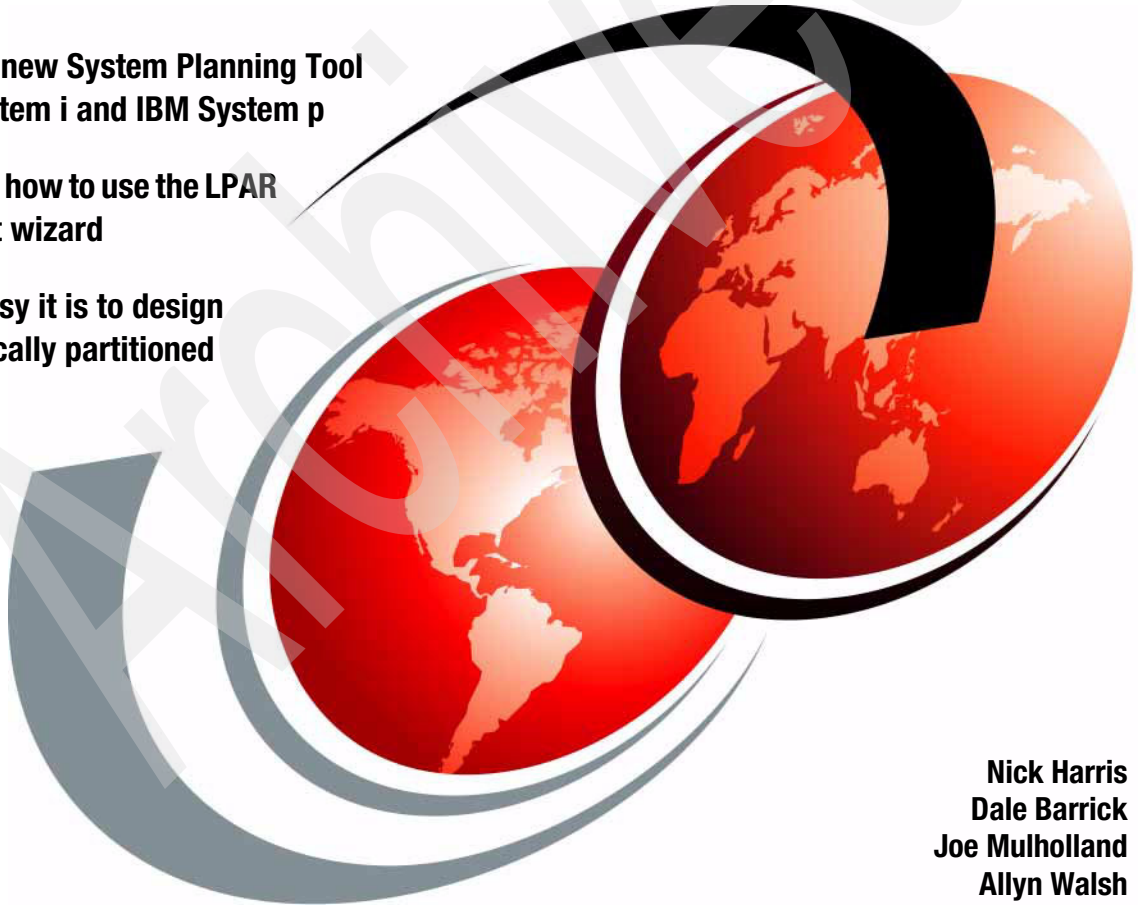

Nick Harris
Dale Barrick
Joe Mulholland
Allyn Walsh

Redbooks

ibm.com/redbooks

International Technical Support Organization

LPAR Simplification Tools Handbook

August 2006

SG24-7231-00

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

First Edition (August 2006)

This edition applies to IBM System i systems and the Hardware Management Console running Version 5 Release 2.

Contents

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates 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. 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 IBM's application programming interfaces.

Trademarks

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

Redbooks (logo)  ™	AS/400e™	POWER5+™
Eserver®	AS/400®	Redbooks™
eServer™	IBM®	System i™
iSeries™	OS/400®	System i5™
i5/OS®	PartnerWorld®	System p™
pSeries®	POWER™	System p5™
AIX®	POWER5™	WebSphere®

aThe following terms are trademarks of other companies:

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

Excel, Microsoft, Windows, 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.

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

Preface

In this IBM® Redbook we summarize the new functions designed to simplify the process of planning for and deploying logical partitions (LPARs) on the System i and System p™ models. The new LPAR Simplification functions include the new System Planning Tool (SPT) and the new Deploy System Plan wizard on Hardware Management Console (HMC). This book focuses on System i.

The System Planning Tool will be the replacement for the Logical Partition Validation Tool (LVT). It will provide the configuration mechanism for designing LPARs on System i™ and System p servers. This book covers the installation and use of SPT.

The new LPAR deployment wizard is a function of the HMC, and provides the capability to take the configuration produced in SPT and use it to create an LPAR configuration on the HMC managing the system. This significantly simplifies and reduces the work effort to create LPARs. The design work is done once in SPT, while the deployment or re-deployment can happen many times.

This book provides guidance information for IBM and IBM Business Partners who help customers plan, order, and deploy systems.

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

Nick Harris is a Consulting IT Specialist for the iSeries™ and has spent the last eight years in the International Technical Support Organization, Rochester Center. He specializes in eServer™ i5 and iSeries hardware, i5/OS® and OS/400® software, LPAR, high availability, external disk, Windows integration, and Linux®. He writes IBM Redbooks™ and teaches IBM classes at ITSO technical forums worldwide on all of these areas as they are related to system design and server consolidation. He spent 13 years in the United Kingdom (UK) AS/400® Business, where he worked with S/36, S/38, AS/400, and iSeries servers. You can contact him by sending an e-mail to niharris@us.ibm.com.

Dale Barrick is a Certified Consulting IT Specialist with IBM Techline in Atlanta, Georgia, and has 29 years of service. As a Systems Engineer, his background is S/36, AS/400, and iSeries. He is currently the iSeries 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 prior to its initial release.

Dale consults on LVT development and maintenance. He is also involved in the on-going Customer Specified Placement project for iSeries and pSeries® with IBM Rochester manufacturing. He has co-authored the IBM Redbook *LPAR Configuration and Management Working with IBM eServer iSeries Logical*, SG24-6251; the IBM Redbook, *IBM @server iSeries Migration: System Migration and Upgrades at V5R1 and V5R2*, SG24-6055; the IBM Redbook, *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000; the IBM Redpaper, HSL Architecture; two editions of the *IBM @server i5, iSeries, and AS/400e System Builder IBM i5/OS Version 5 Release 3 - October 2005*, SG24-2155, and *IBM @server i5 and iSeries System Handbook: IBM i5/OS Version 5 Release 3 October 2004*, GA19-5486; and the *Planning for IBM @server i5 Data Protection with Auxiliary Write Cache Solutions*, REDP-4003. For the past year he has been a consultant for the development team for the System Planning Tool (SPT), which is intended to replace the LVT. You can contact him by sending an e-mail to drbarri@us.ibm.com.

Joe Mulholland is a Support Center Specialist at the IBM Rochester Support center. He specializes in LPAR, CUOD, and i5/OS internals (VMC, Phyp, and main store analysis). He graduated from Winona State University with a B.S. degree in Computer Science in 2003.

You can contact him by sending an e-mail to jpmulhol@us.ibm.com.

Allyn Walsh has worked for over 30 years in the IT industry. He has been at IBM Rochester working with the AS/400/iSeries since 1991. His jobs have included working in Level 2 and the Support Center, PartnerWorld® for Developers, and he is currently assigned to the Americas Advanced Technical Support (ATS) organization.

His current focus areas include Hardware Management Console and System Management and High Availability. He has published a number of articles and IBM internal documents on HMC, system firmware, and Systems Management topics. He is a regular speaker at iSeries Technical Conference and User Group meetings. His prior experience includes extensive work in the areas of data communications, Systems Management, LPAR, and save/restore.

You can contact him by sending him an e-mail at Allyn_Walsh@us.ibm.com.

Thanks to the following people for their contributions to this project:
Dave Blaschke
Tim Covalt
Dave Gimpl

Rob Golley
Kyle Henderson
Beth Hettich
Kent Hofer
Tammy Van Hove
Edith Lueke
Kenton Lynne
Charles Marino
Jossie M McManus
Nathan A Rabe
Devaughn Rackham
Anand Raghavan
Sathyakumar S Ramalingam
Cale Rath
George Romano
Richard Sit
Christopher Smith
Tracy Smith
Tom String
Megan E Williams
Keith Zblewski
IBM Rochester Development Lab

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbook@us.ibm.com

- Mail your comments to:

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

1

LPAR simplification overview

In this chapter we address the question, What is LPAR Simplification? We provide a brief look at the history of logical partition tooling and the hardware order process.

We summarize the new functions designed to simplify the process of planning and deploying logical partitions (LPARs) on the System i and System p models. The new LPAR Simplification functions are included in the new IBM System Planning Tool (SPT) and the new Deploy System Plan wizard on Hardware Management Console (HMC).

The IBM System Planning Tool simplifies the preorder process of planning and designing your LPAR configuration. The Deploy System Plan wizard simplifies the LPAR creation process by importing and deploying the plan that you created with the SPT.

The purpose of the LPAR Simplification project is greater than just LPAR and its deployment. The full breadth of the project is from performance data gathering, capacity planning, hardware configuration, order processing, and deployment. The project creates bridges between these previously separate procedures, and thereby the process as a whole is streamlined and simplified.

Along with the information contained in this book you should also review the existing LPAR Validation Tool information (LVT) at:

http://www-03.ibm.com/servers/eserver/iseries/lpar/systemdesign.html

Who this book is for

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

A novice can use the information in this book, but it would be useful to view other sources for background and additional information. The eServer Hardware and iSeries Information centers are a good first source:

- eServer Hardware Info Center

  http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/index.htm

- iSeries Information Center

  http://publib.boulder.ibm.com/iseries/

Other excellent sources are the iSeries System Handbook and System Builder are:

- *IBM i5 and iSeries System Handbook*, GA19-5486

  http://www.redbooks.ibm.com/redpieces/abstracts/ga195486.html?Open

- *IBM i5, iSeries, and AS/400e™ System Builder IBM i5/OS Version 5* - SG24-2155

There are also a number of IBM Redbooks in the iSeries domain that would also prove very useful:

- *IBM iSeries Migration: A Guide to Upgrades and Migrations,* SG24-6055-01

  http://www.redbooks.ibm.com/redbooks/pdfs/sg246055.pdf

- *IBM iSeries Migration: A Guide to Upgrades and Migrations to POWER™ Technology,* SG24-7200-00

  http://www.redbooks.ibm.com/redbooks/pdfs/sg247200.pdf

- *Logical Partitions on System i5™: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000-01

- *High-speed Link Loop Architecture for the IBM @server iSeries Server: OS/400 Version 5 Release 2*, REDP-3652

1.1 Previous LPAR configuration tooling

Prior to this new LPAR Simplification tooling, there were four separate steps that needed to be completed to design and deploy logical partitions for IBM System i or System p systems.

1. Generate requirements.
2. Plan the system.
3. Order the system.
4. Create LPARs.

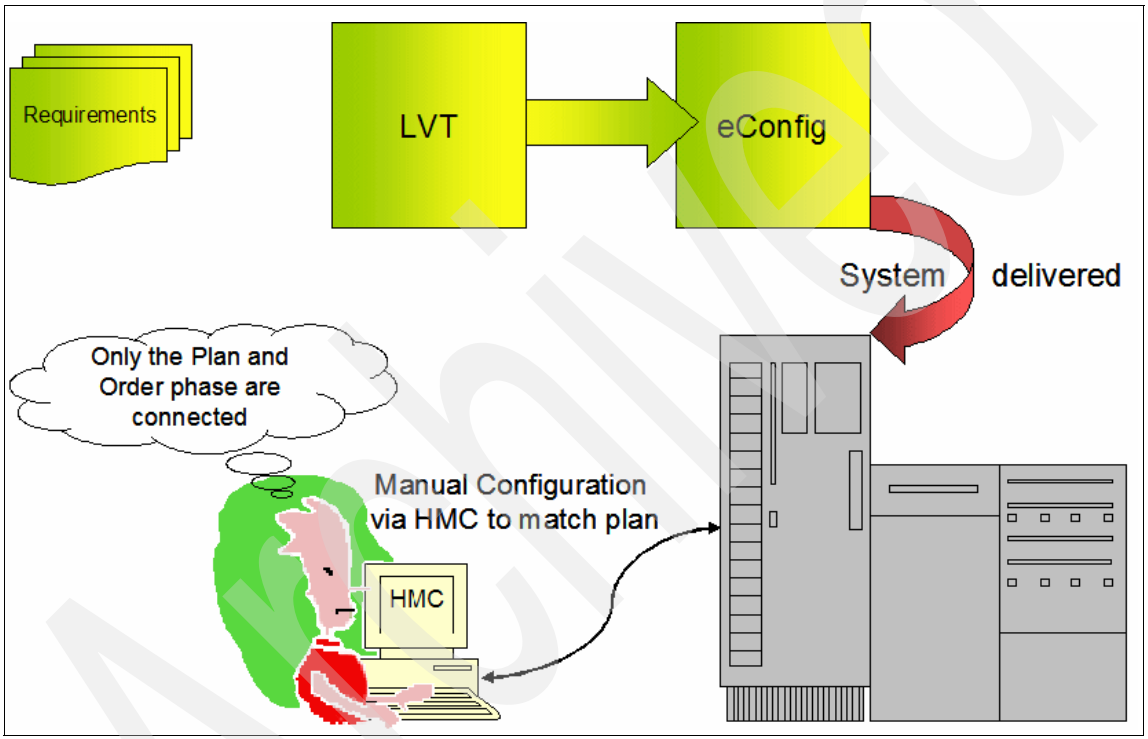

Figure 1-1 Without LPAR simplification effort

1.1.1 Generate requirements

This step involves using the IBM Systems Workload Estimator (WLE), performance tools, or some other means to get an idea of what is needed in a system. This step is vitally important to determine what the system capabilities must be. WLE has been enhanced with additional functions to enable the modelling of LPAR configurations.

1.1.2 Plan the physical layout of the system

The LPAR Validation Tool (LVT) is essential for logical partition hardware design and I/O placement. To use LVT the person designing the hardware configuration would take information from the requirements phase and manually enter it into LVT to design a system to meet the requirements. The LVT compares the input configuration to dependency and placement rules to verify that the system has a valid configuration.

1.1.3 Ordering the system

The IBM Configurator for e-business (eConfig) is the tool used to order a system. The ability to import a .cfr file generated by the LVT has recently been included in this tool. eConfig also provides rules to confirm that the configuration is valid from an IBM Sales view.

1.1.4 Create LPARs

Traditionally, a new system would arrive on site, and then someone from the customer, IBM, or an IBM Business Partner would perform the LPAR configuration. This involved using DST/SST on iSeries hardware, the serial HMC on pSeries hardware, or the LAN attached Hardware Management Console (HMC) beginning with POWER5™ hardware that included a service processor. This was a manual process, involving configuring the LPARs by hand to match the LVT and ensuring that the real hardware locations matched the LVT design.

There are also customer-orderable features that allow for Customer Specified Placement (CSP) for a system, which means that IBM can deliver a machine with its partitions and hardware in place. The customer can then simply load the operating systems.

1.2 LPAR simplification

The purpose of LPAR simplification is to bridge the gaps between the above steps. The LPAR simplification tooling provides a process in which a plan can be taken from the requirements stage through the creation of real LPARs with minimal effort (Figure 1-2). LPAR simplification has two basic parts:

- Plan.
- Deploy.

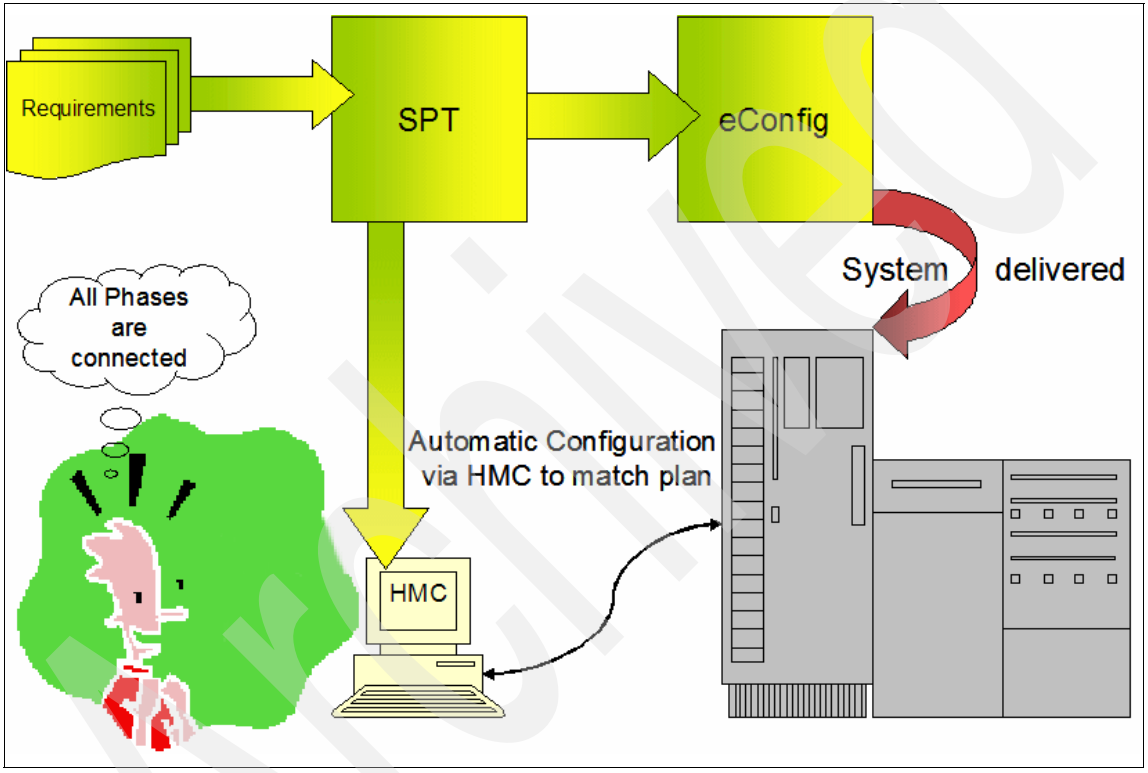

Figure 1-2 With the LPAR simplification effort

1.2.1 Plan

The plan portion of the LPAR Simplification project bridges the requirements and planning phases to the ordering phase. In order to facilitate the simplification, the LVT is being replaced by the System Planning Tool (SPT). SPT is discussed in greater detail in Chapter 2, “System Planning Tool (SPT)” on page 11.

Bridging the requirements to planning stages is accomplished by allowing output from the Workload Estimator (WLE), performance monitor, and performance

toolbox data to be used as input to SPT. SPT allows the creation of system plans from scratch. The SPT tool has sample configurations provided. These can be used as base systems that can be further modified to fit specific needs. With SPT the user is able to open and use existing LVT files. This is a great benefit to existing users, as previous configurations can be re-used with having to rekey the data.

Planning is bridged to ordering by allowing SPT output to be used as input into the eConfig tool. The LVT also provided this function.

The output of the SPT is a new file type called a system plan (sysplan) and has an extension of .sysplan. A brief description of the sysplan is given here.

The sysplan

There are two ways to create a sysplan. One is through SPT and save the output to a .sysplan file, and the second is through an HMC command run though the restricted shell command line. The HMC command used to create this sysplan is **mksysplan**. This is discussed in more detail in Chapter 4, “System plan management using restricted shell (CLI)” on page 73.

When created via the SPT, the sysplan contains all the information needed for use with eConfig, and is then used as the basis to place an order, and later to deploy the new system.

While the system plan is invaluable for ordering and deployment, it can also be used to produce a softcopy or hardcopy of the system design. Both the SPT and the HMC have the functionality to display the sysplan in a browser interface. This enables a user to easily see the hardware configuration and the LPAR information.

New interface screens have been added to the HMC in order to allow sysplan files to be manipulated. These include the ability to import, deploy, and manage the sysplan files. The manage system plans task includes the ability to import and deploy, as well as view, export, and remove sysplans from the HMC. The HMC can import a sysplan from media, including CD, disk, or USB device, as well as from a remote FTP site.

From the HMC the sysplan can be deployed and managed. The HMC also has the ability to export a sysplan to the same media it can import from. This is discussed in more detail in Chapter 4, “System plan management using restricted shell (CLI)” on page 73.

It is important to note that at this time a sysplan created on the SPT includes device information, but a sysplan generated from the HMC command line does not. The sysplan created by the **mksysplan** command only contains detail to the

IO Adapter (IOA) level. Details of devices below the IOA are not currently available on the HMC. Low-level information is only known to the operating system running within a particular partition. The mksysplan function cannot currently interrogate the running OS. Therefore one must use an operating system command or physically review the hardware layout manually to get the same level of detail that SPT provides.

A sysplan created via mksysplan can be used to deploy on other identical systems, or on the same system to regenerate or add to the current system plan.

It is also important to note that at this time, the sysplan created by the SPT cannot have overcommitted resources. Overcommitted resources occur when there are multiple partitions on the system, but there are insufficient system resources to allow all the partitions to be running at the same time. This could be because new resources have not been installed or by design where the powered-off partitions are shell partitions.

The HMC has the ability to overcommit resources as a partition can be powered off and have their resources remove and used by another running partition. Therefore, if you were to export a system plan with overcommitted resources, the sysplan would be invalid, as SPT validates a plan, and it assumes that all LPARs will be running at the same time.

1.2.2 Deploy

One of the most interesting features of the LPAR Simplification tooling is the ability to deploy a plan to an actual system. The HMC has been modified to take output from the SPT and automate the creation of the LPARs. This topic is discussed in greater detail from a GUI perspective in Chapter 3, “System plans and the hardware management console” on page 45, and from a command-line interface (CLI) perspective in Chapter 4, “System plan management using restricted shell (CLI)” on page 73.

Once a sysplan is moved to an HMC, it can be deployed if the physical system matches the plan. The HMC has the ability to deploy all or part of the sysplan. This allows for a subset of all LPARs to be deployed, as opposed to all LPARs.

The deploy process must do two very important validation steps before it will deploy a plan to a system: hardware validation and partition validation. A sysplan can be deployed to any system that successfully validates both steps. The validation process is discussed in greater detail in Chapter 5, “Deployment validation” on page 87. A brief over view is given here.

Hardware validation

During this step in the validation process, the HMC ensures that the physical system has the hardware to fulfill the plan. The HMC must validate that the main system unit and all expansion towers in the plan exist on the actual managed system. Validation is done to ensure that all system units exist with the correct IO hardware in the correct slots to fulfill the plan. This validation will be successful even if there is an equivalent number of or more hardware resources than are needed to fulfil the plan, but the validation will fail if any hardware that is in the plan is missing.

At this time, there is no way for the HMC to know about hardware below the IOA level. This is important to keep in mind, as the HMC cannot validate hardware it cannot see, such as disk drives. This becomes important if there is tower ambiguity. This is when there are multiple towers that may be identical to the IOA level and differ only at the device level, having a different number of disk drives, for instance. Tower ambiguity, and how to cope with it, is discussed in more detail in 5.1.3, “System or expansion unit ambiguity” on page 92.

Partition validation

Partition validation is done to ensure that any partitions that may already exist on the system also exist in the plan. If there are no preexisting LPARs, this validation will always succeed. If a partition exists on the system that does not exist in the plan, this validation step will fail and the plan will not deploy. In this case, the existing LPAR must either be deleted, or the sysplan put back into the SPT and modified to include it. If LPARs exist on the system that are in the plan, the deploy can continue. This is necessary to allow for the incremental deploy of a plan to the system.

Important: The default partition on a new system is considered a real partition by the deployment wizard during the partition validation step. This default partition is not required if you are deploying partitions created in SPT. Therefore it is important to delete the default partition before deploying a system plan on a new system.

1.3 Feature comparison between SPT and LVT

Table 1-1 shows a comparison between the new System Planning Tool and the current LPAR Validation Tool.

Table 1-1 SPT versus LVT

Feature	SPT	LVT
Available for free	X	X

Feature	SPT	LVT
IBM System i and IBM System p support	X	X
System specifications	X	X
Hardware I/O card and disk drive placement support	X	X
Partition specification and validation	X	X
Console specification and validation	X	X
Load source specification and validation	X	X
Mutli-system configuration support	X	X
Virtual memory and virtual resource support	X	X
Save and restore configurations	X	X
Export for order processing	X	X
Copy systems and expansion towers	X	Limited
Report generation	X	Limited
Output can be used for system documentation	X	Limited
HMC compatibility	X	Limited
Integrated system sizing and capacity planning	X	
System View capability	X	
Export configuration to MS Excel	X	
Interactive reporting	X	
Online help	X	
Report print capability	X	
System plan consolidation	X	
Save system plan for deployment	X	

2

System Planning Tool (SPT)

In this chapter we discuss the features, functions, and capabilities of the System Planning Tool.

2.1 Overview of the System Planning Tool

The IBM System Planning Tool (SPT) is a tool for designing logically partitioned iSeries and pSeries systems and is the replacement for the LPAR Validation Tool (LVT), but it can also be used for planning and documenting non-partitioned systems. The SPT is a browser-based tool that runs on your PC. For download and install information see Appendix A, “Installing SPT” on page 113. The GUI and order of operations are quite different from the LVT, 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, from planned workloads, from sample systems, 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 the LPAR Validation Tool (LVT), SPT is available for download from the http://www.ibm.com Web site. A subscriber list will be used to notify users when a new version is available.

The first time you download the SPT you must use the full version, which includes the required JVM™ code and other support files. For downloading subsequent versions of the SPT, you may download the update version of the SPT. 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, “Installing SPT” on page 113.

2.3 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 you start the SPT or when it is started at the end of the install process, you will be presented with be 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.

Click **Launch**, as shown in Figure 2-1.

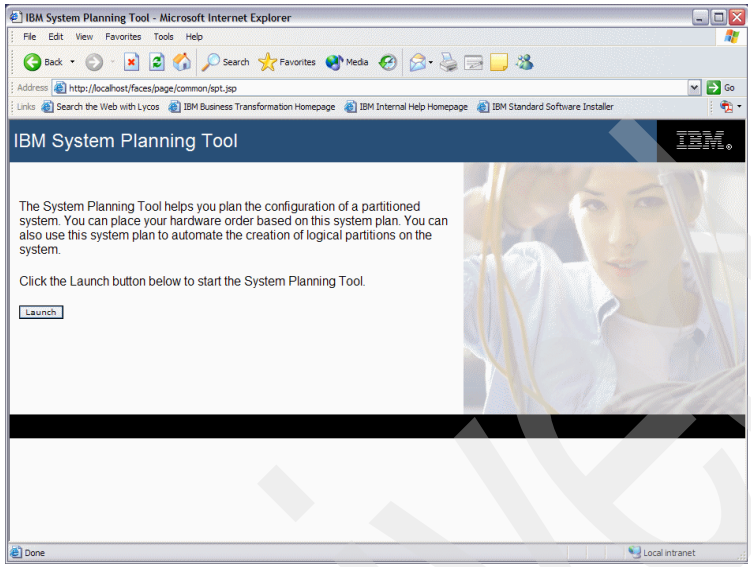

Figure 2-1 System Planning Tool Launch window

The Getting Started window initially has only two options. If you create a system plan and close the SPT window and then use Launch again, the Getting Started window will have a third option called Reconnect with open plan. Click **Create a new system plan** (Figure 2-2).

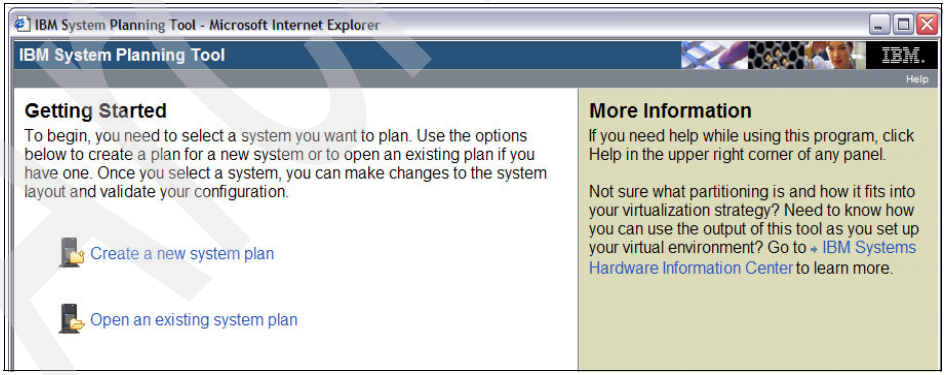

Figure 2-2 Getting started window

The Add System window opens and there are four options that allow the user to create a new system.

The first option, Create a new system (advanced), is very similar to that used in the current LVT. The user starts with a empty model and every detail of the system from type/model to processor and memory capacity and all the expansion units, drawers, adapters, and disk is required to be specified.

The second 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.

The last two options use IBM Workload Estimator (WLE). These options are to create a new system based on existing performance data or based on new workloads you want to run. One option allows you to create workloads in WLE from performance data saved on a local system, and the other option is used by invoking WLE to create workloads you can use.

2.3.1 The reconnect function

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 were being edited.

Note: If a user is in the middle of editing a system plan (and 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.

2.4 Creating a new system

To create a new system, select the first option, **Create a new system (advanced)**, and click **Next** (Figure 2-3).

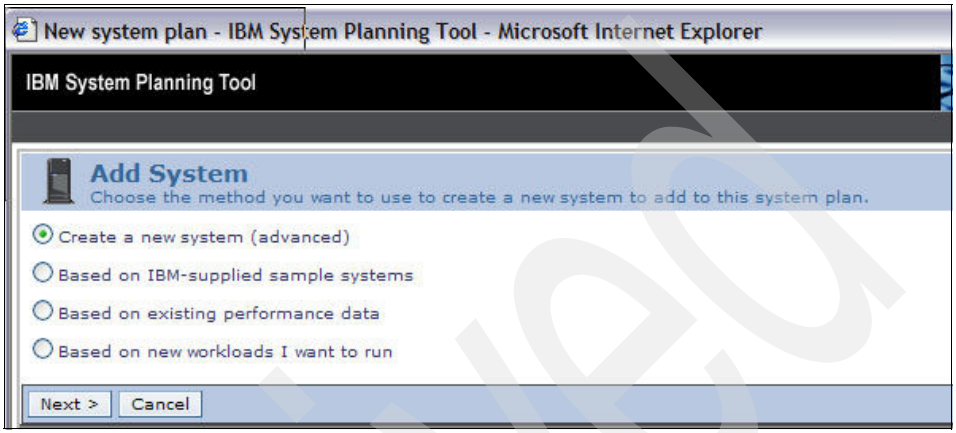

Figure 2-3 Add System selection window

2.4.1 Selecting the system platform and type/model

The System Type window will open as shown in Figure 2-4.

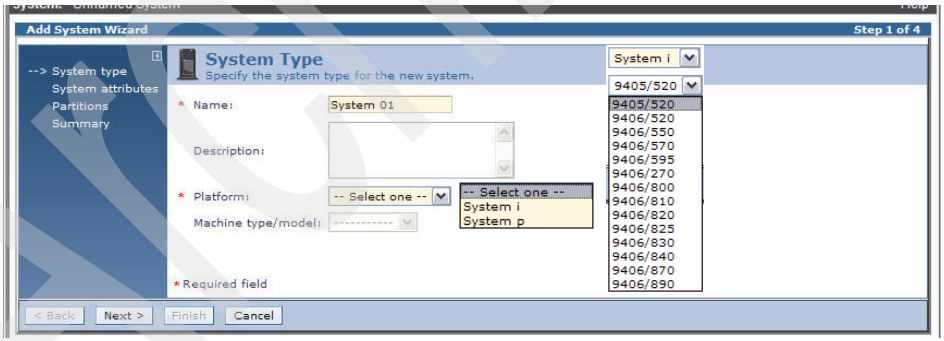

Figure 2-4 System Type window

On the left side of the window frame is a navigation bar. This is not a live area and only displays the progress through the steps. There is an arrow icon in the top right-hand corner of the navigation area that allows you to collapse this area to view more of the working area.

You can name the system as appropriate and enter a description. Note that if you plan to deploy this system plan, the name you give the system here will become the system name on the real system. You can also enter a description. This is useful if you are creating many system configurations and need to identify particular plans.

Next select the platform, System i or System p, as shown in Figure 2-4 on page 16. The machine type/model menu displays configurable systems appropriate for the system selected. Select your model from the menu and click **Next**. In our example we selected a Model 550.

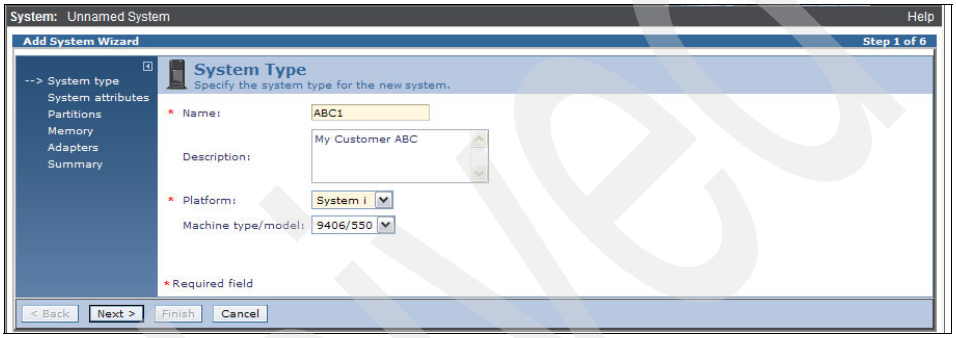

Figure 2-5 System type selected and named

2.4.2 Working with system attributes

The System Attributes window then appears, as shown in Figure 2-6.

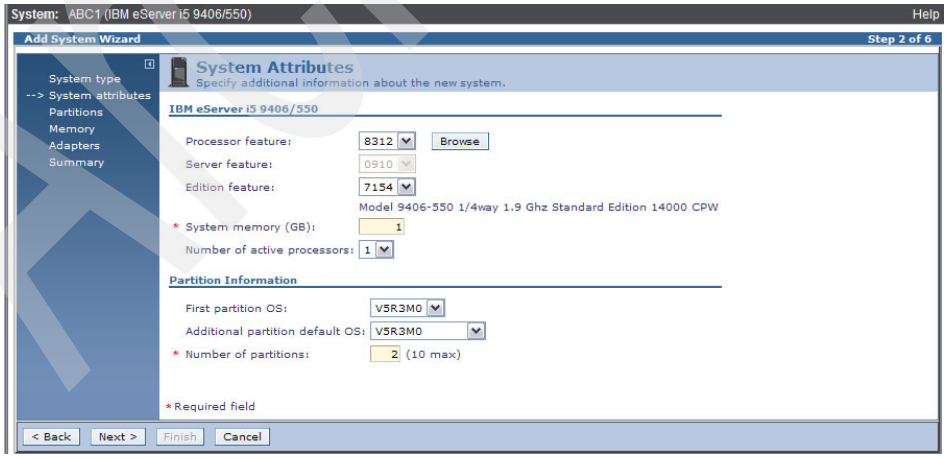

Figure 2-6 System Attributes window

Note that the system name has not changed to reflect the name we specified in the previous step, but the type/model have been selected as System i and Model 550.

1. Select the processor feature from the menu.

   If you are unsure of the capabilities the processor you are planning to select, click the **Browse** button and the Browse Processor Features window will open (see Figure 2-7).

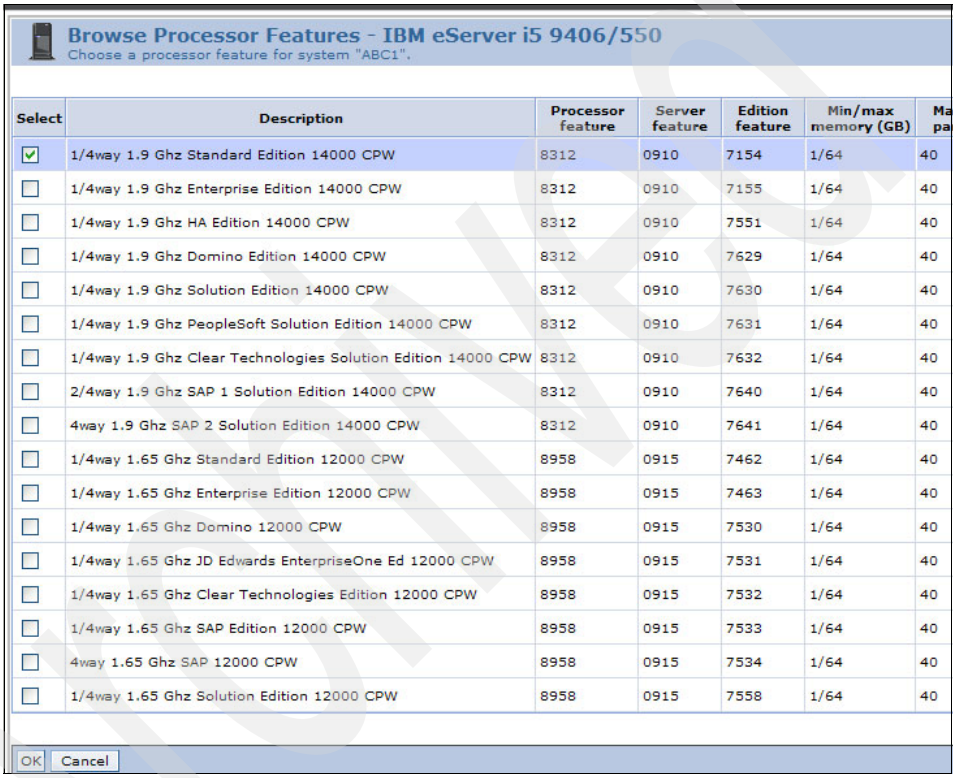

Figure 2-7 Browse processor features

2. This detail panel shows you the attributes of the various features for the system you selected. You can check the box adjacent to the processor feature to select it. Click **OK** to close the window.

   The server and edition feature combo boxes shown in Figure 2-6 on page 17 are hidden and displayed based on the selected machine type and processor feature. They will be hidden if the selected machine type or processor feature does not contain those features. Changing the processor feature causes the choices in the server feature to be updated. Changing the server feature causes the list of edition features to be updated. The same is true for

accelerator features and processor nodes. Also, the field will be visible but disabled (grayed out like in the picture) if there is only one choice.

3. Next select the OS for the first partition you would configure in the Partition Information section shown in Figure 2-8. This section of the window also allows you to select a second OS for other partitions. At this point the second OS is not important, as you have the option to select the OS on the Partitions window, which is the next SPT screen. The first partition and its OS is not significant for System i5 systems, unless this is an non-LPAR system or you are planning a virtual partition manager environment. On iSeries 8xx models it would indicate the primary partition OS and version, which is important, as on 8xx hardware there is still a requirement to have a primary or managing partition. On System i5 there is no requirement for a primary or managing partition.
4. Next select the number of partitions you are planning to configure on the system. If you decide later to add or remove partitions, this can be done on the Partitions window, which is the next screen in SPT, or from the edit partition link on the Work with Planned systems screen (see Figure 2-19 on page 29).
5. Make the OS and partition selections as needed from the drop-down boxes and click **Next**.

See Figure 2-8 for selections used in our example.

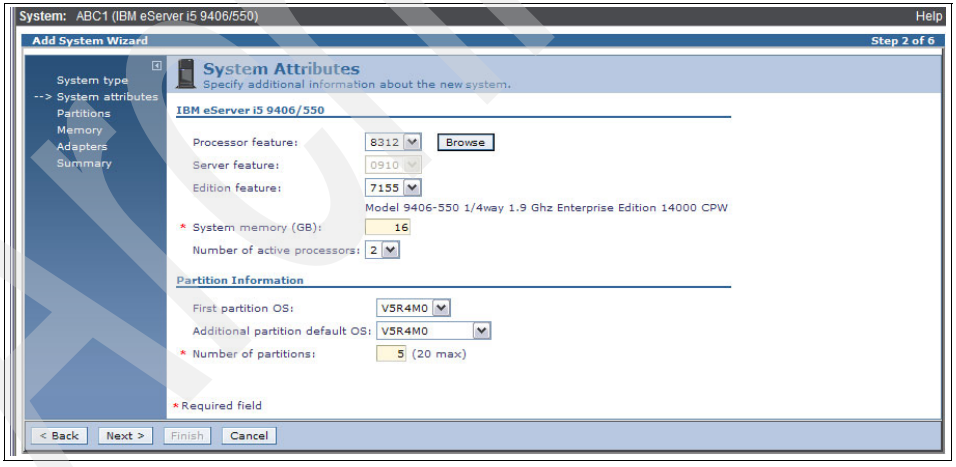

Figure 2-8 System Attributes completed

2.4.3 Defining system partitions

The System Partitions window is displayed next and shows default partition names (LPAR1, LPAR2, and so on). There are some default values placed in the Processing units fields and Virtual processors fields. You can change the partition names, as our example shows in Figure 2-9, and can edit all the other fields to meet your planned requirements. Entering a number in one of the interactive fields will calculate and place the corresponding value in the other interactive field.

Important: We recommend that you use WLE for obtaining the CPW information for a partition. Experience shows that users are designing systems based on the CPW value shown by LVT/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/LVT *does not* apply this tax in its estimations.

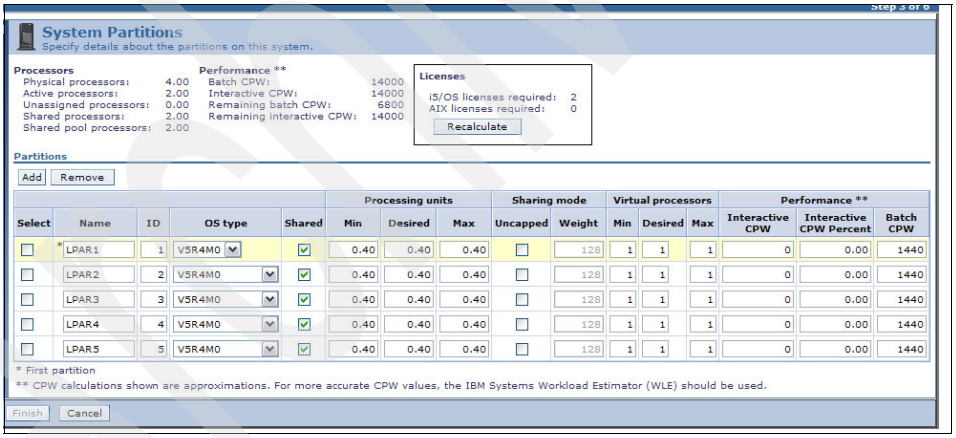

Figure 2-9 System Partition information

The Add button adds an partition row each time you click it. The remove button removes any partitions with the Select box checked and if any hardware has been assigned to the partitions being removed that will also be deleted from your design. Once you have removed a partition there is no undo function.

The top portion of the System Partitions display contains information about processors, performance, and licenses. In this example the processors are over committed and the resulting negative numbers are shown in red (Figure 2-10). Adjusting the processing units corrects that.

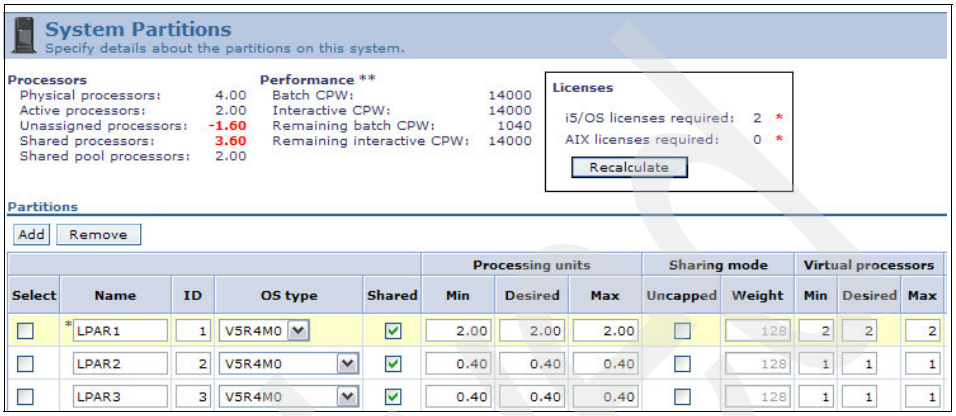

Figure 2-10 System Partitions window error

Once you have modified the processor assignments (physical and virtual), click **Recalculate** and the number of OS licenses required will be adjusted to reflect these assignments.

Note: Changes to the processing resources may also change the number of virtual processors. If you change the processors to remove a negative, you may receive and error message showing that the virtual processors are now incorrect. Change the virtual processors and click **Recalculate** to re-validate the configuration.

When all selections are complete, as show in Figure 2-11, click **Next**.

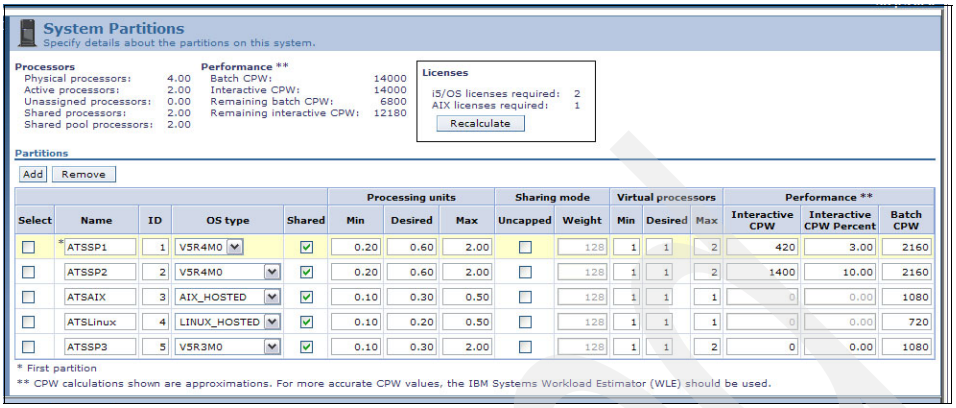

Figure 2-11 System Partitions information complete

2.4.4 Defining memory and virtual I/O

Continuing our design, the memory and virtual I/O need to be defined. The System Memory and Virtual I/O window is displayed (as shown in Figure 2-12) and the values in all of the fields are set to default. You must assign the memory per your requirements. Once you specify the memory requirements of the partitions, the memory information in the top portion of the window is calculated. The hypervisor memory is calculated by clicking **Recalculate**.

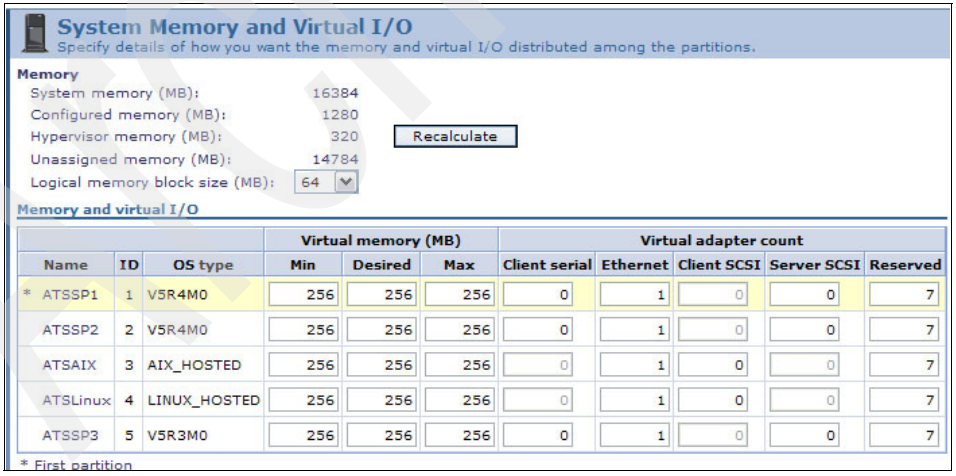

Figure 2-12 System Memory and Virtual I/O window

In our example we added the memory in large round numbers (5000, 4000, 1000). These values are converted to multiples of 64 MB as defined by the

Logical Memory Block (LMB) size value. For System i5 or System p5™ you can leave the LMB as default. SPT automatically calculates the values and shows the remaining unassigned memory, in out case 768 MB. While you can leave unassigned memory, with the dynamic resources movement capability it servers little purpose.

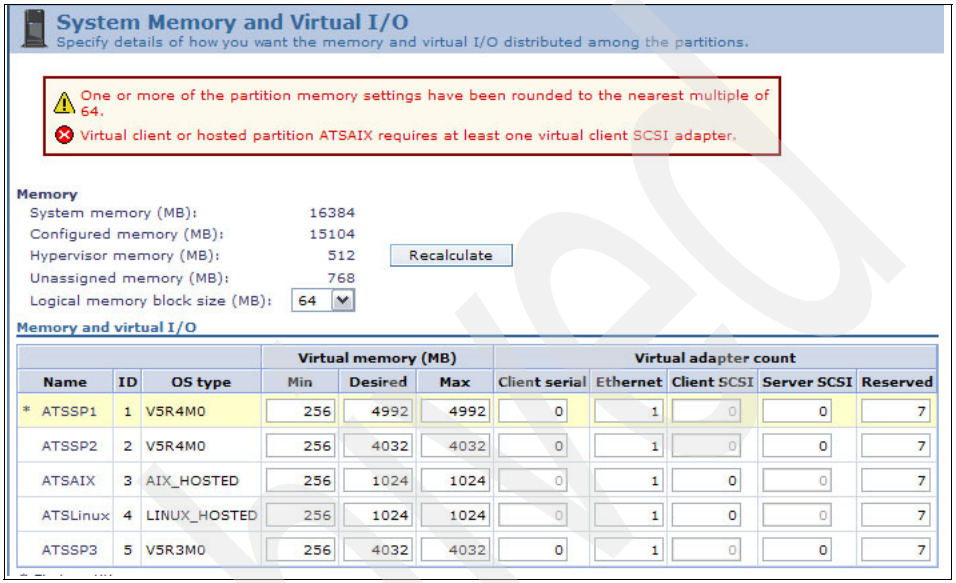

Figure 2-13 System Memory and Virtual I/O warning message

In Figure 2-13 on page 23 there is a second warning message indicating that the hosted AIX® partition does not have a client SCSI adapter defined. This is driven by our hosted partition information. By adding a client SCSI adapter the warning message will disappear and be replaced by another for the hosted Linux partition needing a Client SCSI adapter. Once the client side is fixed, further messages relating to the server SCSI adapters are displayed. In our case we responded to these messages by adding one to the Client SCSI fields for AIX and Linux partitions and two in the server SCSI for the first partition where we are planning to host out AIX and Linux systems.

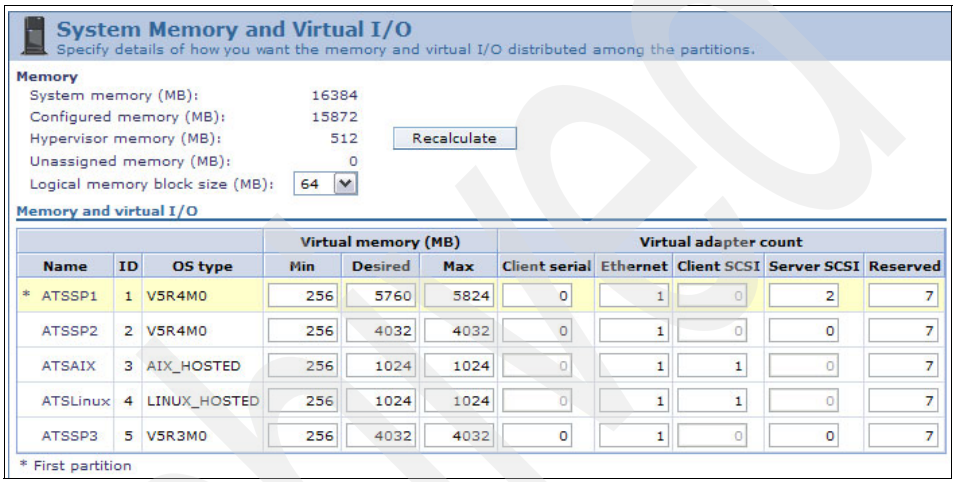

Figure 2-14 System Memory and Virtual I/O complete

The virtual I/O fields are populated based on the number of partitions defined and the OS type of those partitions. The use can update these fields if they have a different configuration planned. The next step allows the user to change the mapping that was automatically generated by SPT.

Click **Next** and the System Virtual Adapters window opens, as shown in Figure 2-15.

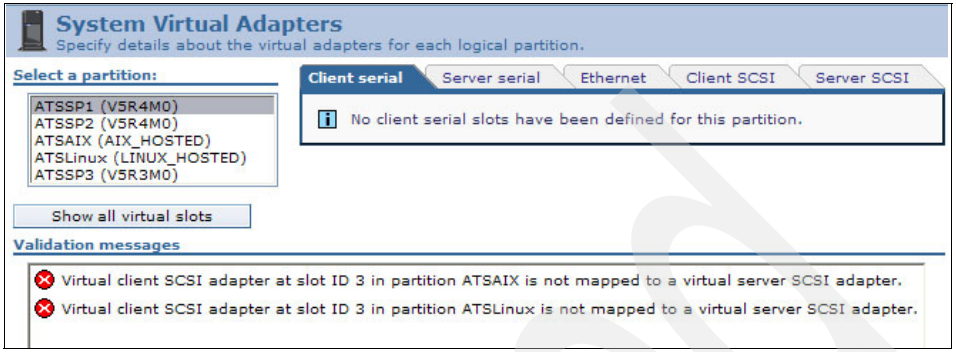

Figure 2-15 System Virtual Adapters

2.4.5 Specifying Virtual IO adapters

The System Virtual Adapters window is where you define the details of virtual I/O and the relationships between the serving and client partitions.

In our example we have two errors relating to the hosted partitions.

In the previous step we added client SCSI adapters and server SCSI adapters, two virtual server SCSI adapters in the ATSSP1 partition and a single virtual client SCSI adapter for each of the ATSAIX/ATSLinux partitions. The target partition/virtual slot ID defaults to <None>, and therefore we need to assign these adapters to the correct slot ID relating to the hosting partition.

Select the hosted partition as shown in Figure 2-16 and then use the pull-down menu to select the virtual slot you wish to use for that partition, as shown in the same figure. Repeat this for the ATSAIX partition using the other slot ID.

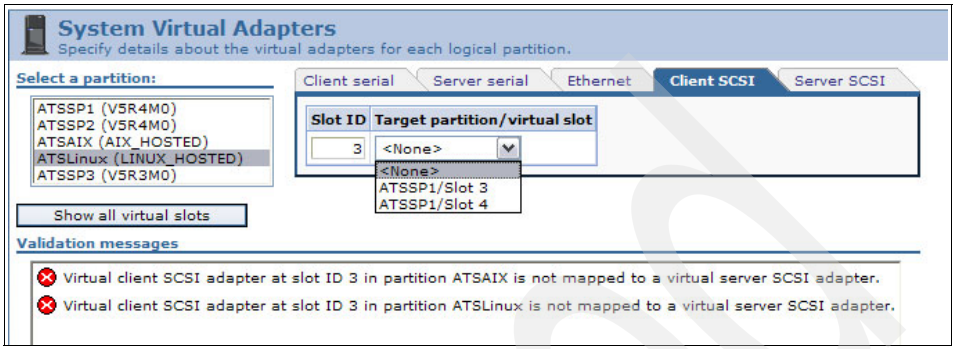

Figure 2-16 Assign Virtual Adapters to virtual slot window

In the System Virtual Adapters window there is an option to show all virtual slots. As the title suggests, this gives the user the ability to display all slots for a particular partition. The is very useful for systems that have many virtual adapters defined that are also different types of slot. Click **Select all virtual slots**. Select the partition you wish to display. In our example shown in Figure 2-17 we selected to display all of the virtual I/O slots for the ATSSP1 partition.

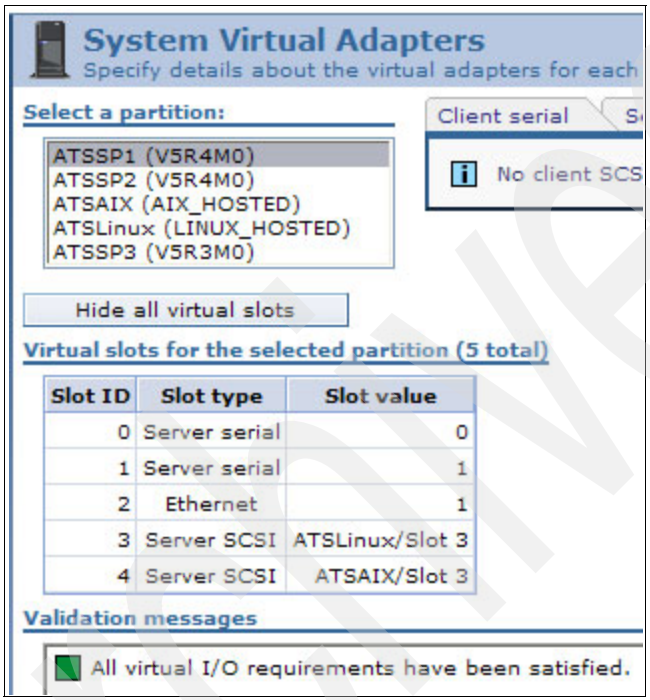

Figure 2-17 Show all virtual slots

2.4.6 System summary

After all the selections for the system and LPAR requirements are met, the next window presents you with a summary of the entire system as it is defined to this point. Click **Finish** and you will proceed to a window that allows you to complete the remaining details of your configuration.

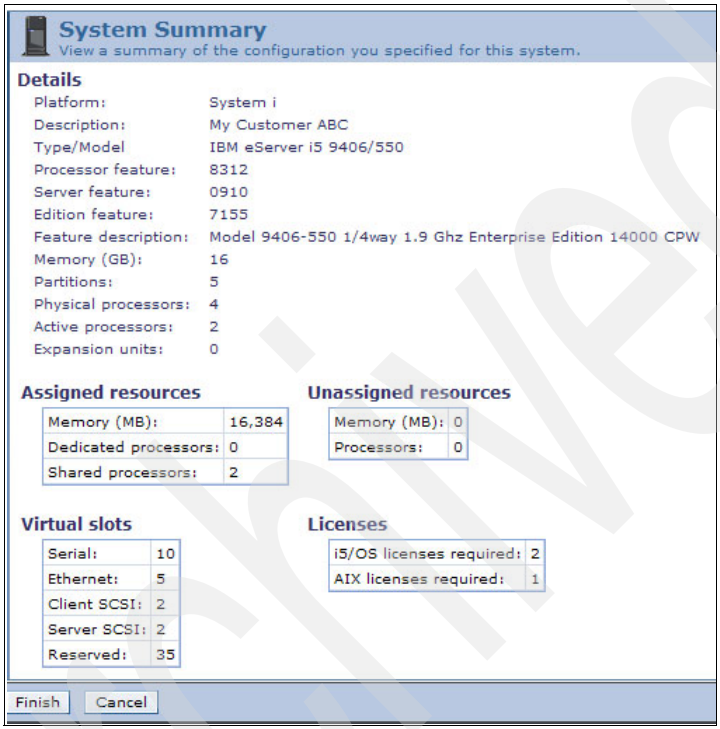

Figure 2-18 System Summary window

2.5 Work with planned system

The window shown in Figure 2-19 is Work with Planned Systems. From this window you can select to edit the partitions, place hardware (which is where you add I/O and expansion units to the system), define the load source (also known as the boot device), and define consoles.

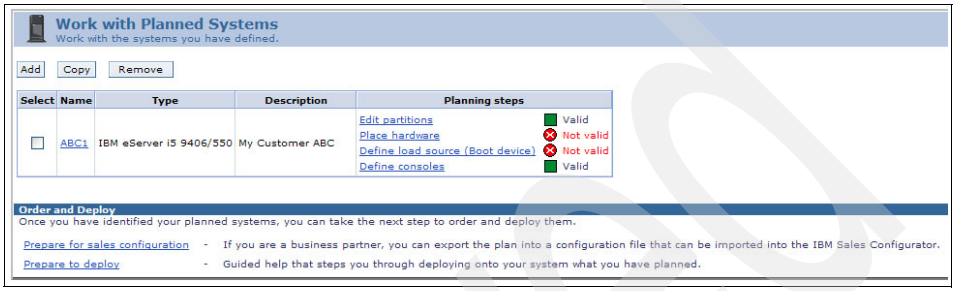

Figure 2-19 Work with Planned Systems window

You can click the system name and a pop-up window will display the system summary. This is particularly useful if you are configuring multiple systems and want to compare them. Each time you click a system name a separate pop-up opens with the summary details allowing for side-by-side comparison.

There are three buttons above the Work with System plans table. Add takes you back to the selection panel where you can create a new plan (sample, scratch, WLE). Copy creates a duplicate of the selected system plan named "Copy of <selected system plan>" and takes the user to the Edit system plan window. This is useful for cases where two identical or similar systems are needed. For similar systems, the user just needs to tweak the newly created system, rather than create a new one.

You will notice that Place Hardware and Define loud source are at an Invalid status. This is because no hardware has been placed and therefore no load source or console devices have been selected.

If you are a regular user of LVT the following screens should be very familiar to you.

Note: Along the header banner there are options that allow you to save, export, and produce a report of you system plan. The export function allows you to move the file to the sales configurator (eConfig) for order processing or for deployment of the system plan to the Hardware Management Console (HMC).

At this point you can go back and make changes to the partition information. Select the **Edit Partitions** link. You will be taken to a tab for of the partition wizard. You do not have to repeat the whole wizard process. You can jump in at any tab and make a change. See Figure 2-20 for a view of the Edit Partitions window.

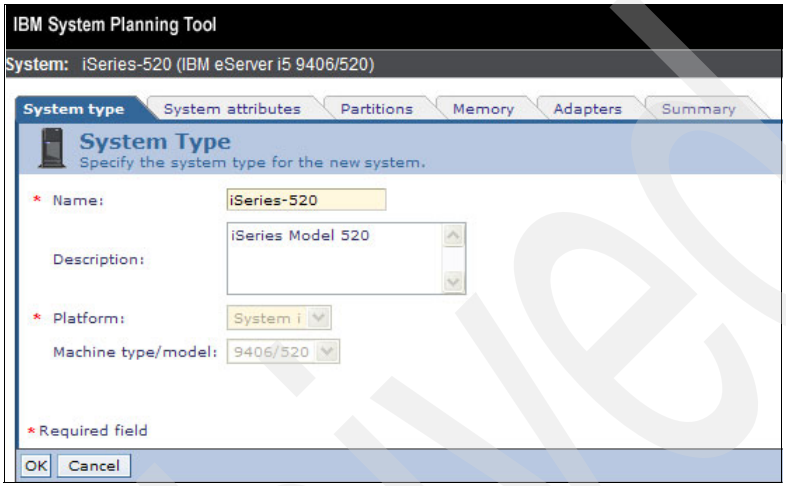

Figure 2-20 Edit partition information

In the next step you will place hardware devices into the system plan and allocate these devices to planned partitions.

2.5.1 Placing hardware with the System Planning Tool

The following figures show the addition of expansion units and placement of IOPs, IOAs, and disk.

The header panel shows the component of the system you are currently working with, in our case the Model 550 system unit, as no expansion units have been configured.

The action buttons in header frame are:

- The OS version will change the selectable devices for a specific OS and its release.
- The Show Expanded view toggles between collapsed and expanded. The collapsed view shows disk drives grouped into drive bays with a row for each drive bay, and the expanded view shows a row for each individual disk drive slot.
- Hide validation messages toggles the lower messages panel on or off.

- The Additional features displays features not specified in the left-hand navigation panel. The parts shown in the Additional Features list are components that are a necessary part of an order, but that will not be placed in a slot at the back of the machine (such as cables).

The left-hand navigation panel displays the hardware resources. Click the category name or the plus sign (+) to expand the selection.

The main panel displays the slot arrangement for the frame (either System Unit or Expansion Unit) you are working on.

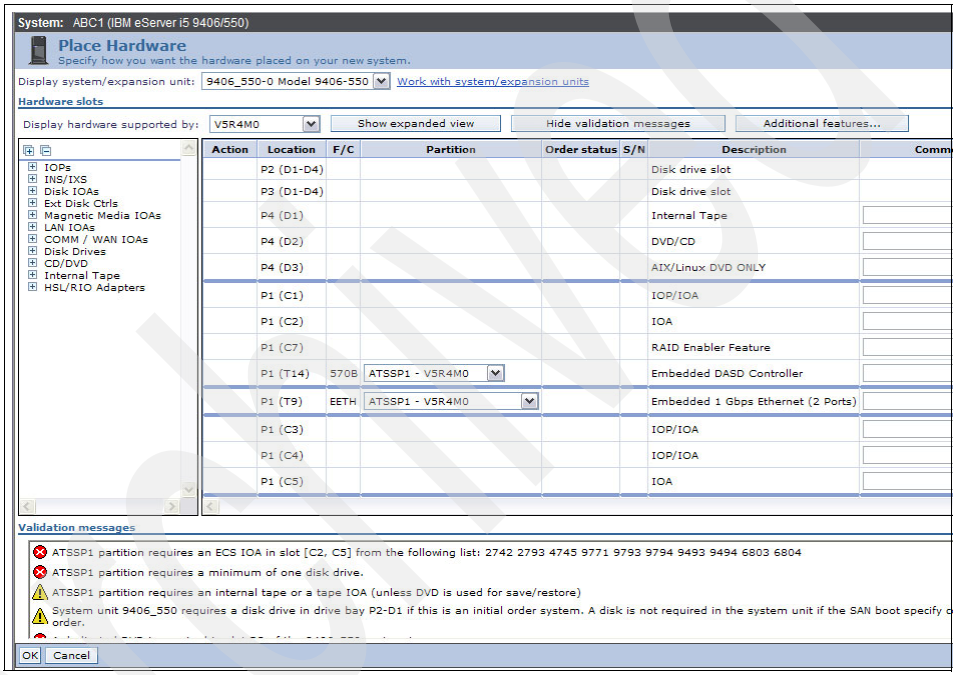

Figure 2-21 Place Hardware

Note: In the Place Hardware panel there are text boxes for serial numbers. Pre-ordering these is not important. 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 5.1.4, “Eliminating unit ambiguity” on page 93.

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 2-2 on page 14. These sample systems have hardware already placed. Reviewing these sample configurations would give a novice some insight into hardware placement.

2.5.2 Adding an expansion unit to the configuration

Select **Work with system/expansion units** in the header panel and you are taken to a window (see Figure 2-22) that provides the Add function for expansion units. Click the **Add** button box and click **OK** to add an expansion unit. The new unit will be added to the bottom of the list.

The selection check box in the row is only used for the copy and remove functions, so the user interface knows which units you want to perform the copy or remove operation on.

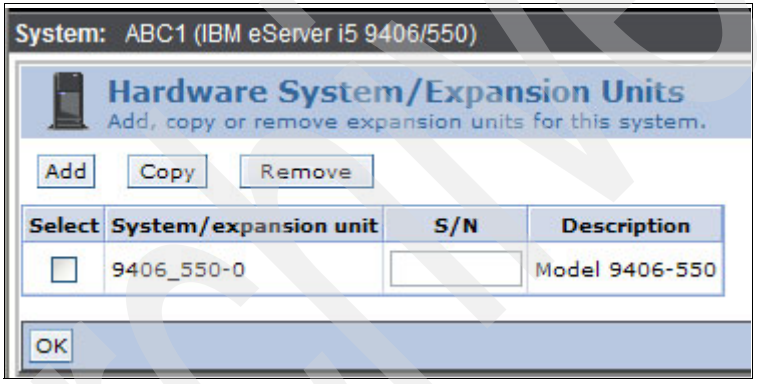

Figure 2-22 Hardware System/Expansion Units

The Add Expansion Unit(s) window opens. Select the pull-down to see the available expansion unit models (Figure 2-23). Select your model and click **OK**.

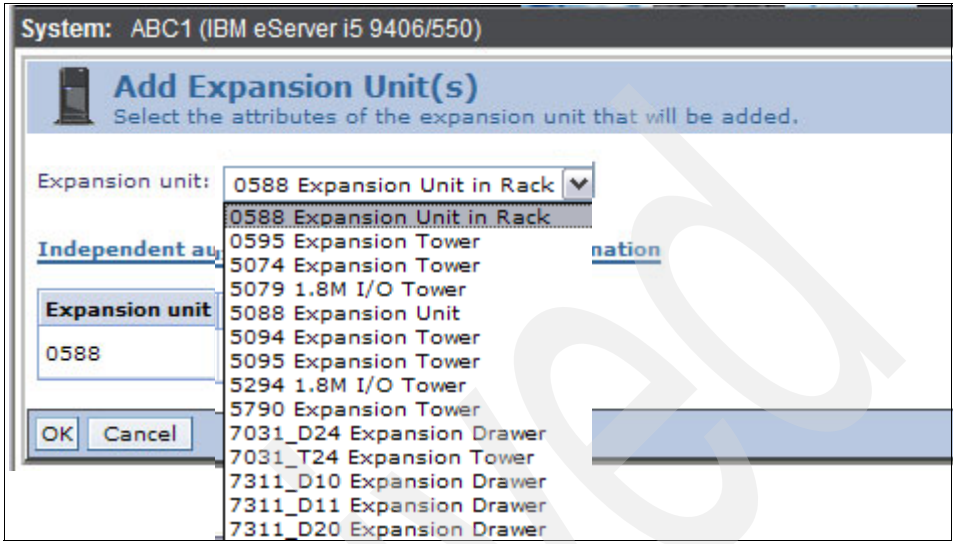

Figure 2-23 Adding an 0595 to the system

As you add expansion units you will see an additional entry in the table for expansion units. In this tabulated area you also have the ability to designate this expansion unit as part of an Independent ASP (IASP) configuration. Check the IASP box to add make this expansion unit an IASP entity. IASPs are used for System i5 availability. The availability design should have been decided before you started your system plan.

There is a Copy button. The purpose of this is to make a complete copy of an expansion unit, including all parts that have been placed in the unit, to another system unit (which could be the current unit if you want duplicate towers for one system). However, if you only need several of the same kind of expansion units for the current system unit, and the parts placed in those expansion units will be different, you really just want to use the Add button to add those units.

The remove function works in much the same way as copy, and when removing an expansion unit you will get a message telling you that any hardware already in the unit will also be removed, and you will have a Yes/No box to either proceed with or stop the remove.

Continue adding expansion units per your design requirements, then click **OK**. This takes you back to the main Place Hardware window (Figure 2-24) where you can select from a drop-down box the system unit or expansion unit you want to populate with I/O devices.

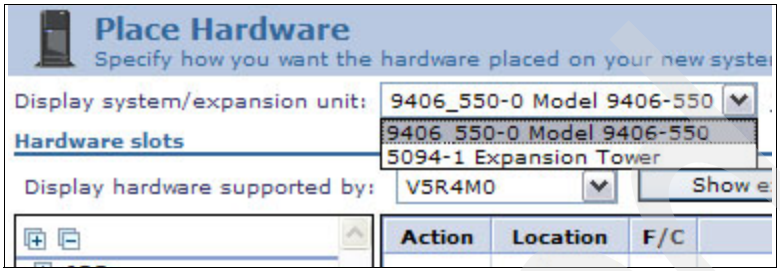

Figure 2-24 Selecting a unit to populate from the drop-down box

After selecting a unit from the drop-down box, 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. The placement rules can be found on the IBM Redbooks Web site at:

http://www.redbooks.ibm.com

The publication number is REDP-4011.

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.

Users of the LVT will be familiar with the actual placement process. From the sidebar select the category of hardware you want to add. Select the particular item from the list. For that item the allowable positions within the expansion unit will be highlighted. Then click the plus symbol (+) of the desired position.

After placing an item a remove symbol (X) appears, and if you select any item that can go in that slot, a replace symbol (<-<-) appears next to the item. The replace function actually performs a remove/add process.

As we are creating a working configuration 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.

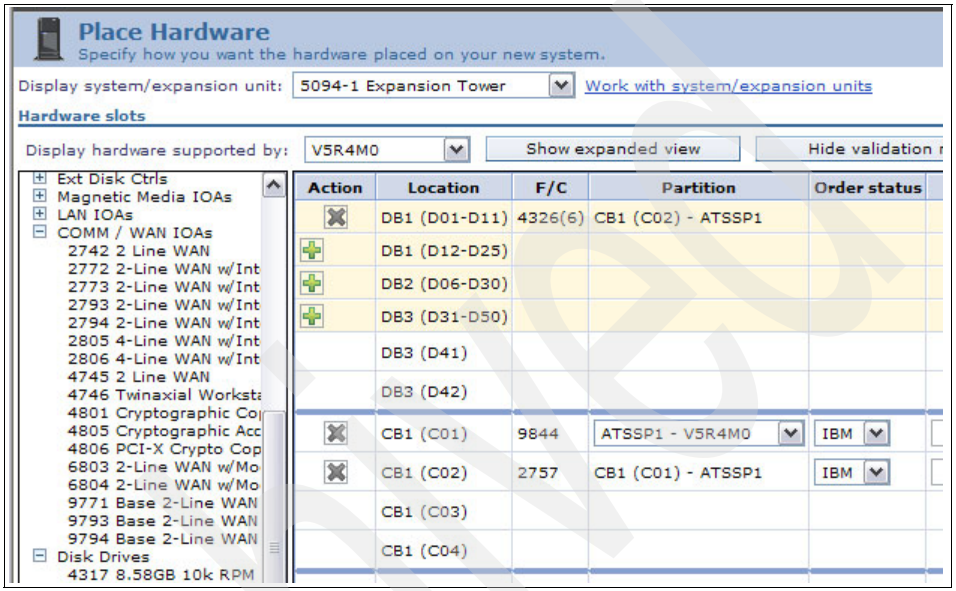

Figure 2-25 Place hardware disk components

2.5.3 Select load source unit

Now that we have placed an IOP and IOA we have the ability to define these devices as load source candidates for a partition. In Figure 2-26 we show the partition selection and load source selection options.

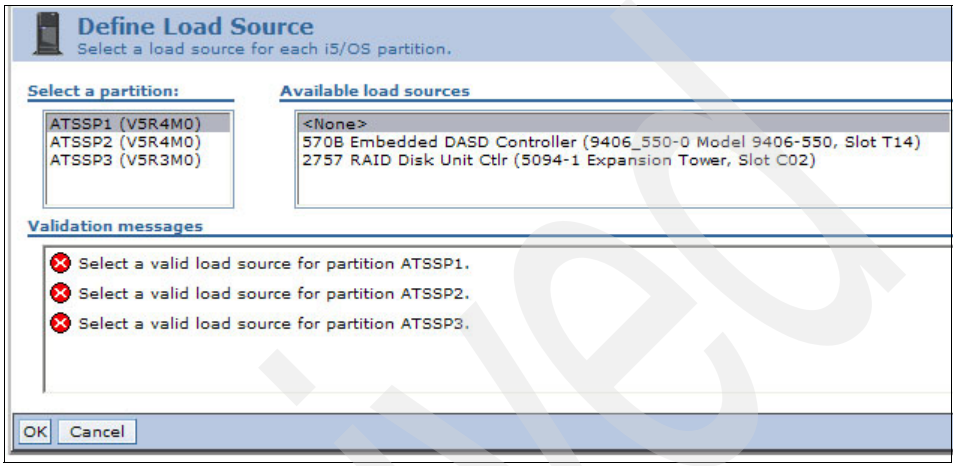

Figure 2-26 Select Load Source unit for a partition

The list of available load source capable devices can include an IOP or an IOA. You should check out the system and expansion unit capabilities, as some systems are now available with IOPless capability.

2.5.4 Select a console

Once you have defined the hardware placement for your partitions to complete your configuration you should now add a console for each partition.

In our example we see the HMC as the default in the pull-down menu. You can select from any supported console.

A console can be any supported console device. It could be a twinax IOA, LAN IOA, ECS IOA, or the HMC, as shown in Figure 2-27.

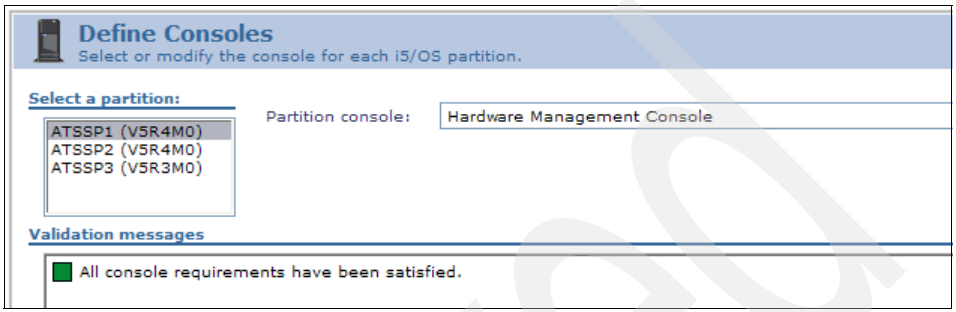

Figure 2-27 Define Consoles

2.6 System plan completion

Once all the components of the system plan are showing a valid status on the Work with Planned Systems display, the plan is deemed complete. The .cfr file should be created and saved to disk using the export function. Then it is necessary to import the system plan into eConfig to ensure that all the components in the plan are valid and the features are orderable. Once the system plan is validated by eConfig a priced proposal can be produced for the customer.

When the customer system arrives on site the system plan can be saved as a .sysplan file. This file can then be copied to media for deployment on the HMC during the system setup.

2.6.1 Exporting a system plan for eConfig

When all the planning steps are showing valid status, as shown in Figure 2-28, you can export the plan to eConfig. In the header there are six selections available. Click **Export** to begin the process to save the system plan as a .cfr file. You will use this file as input to eConfig.

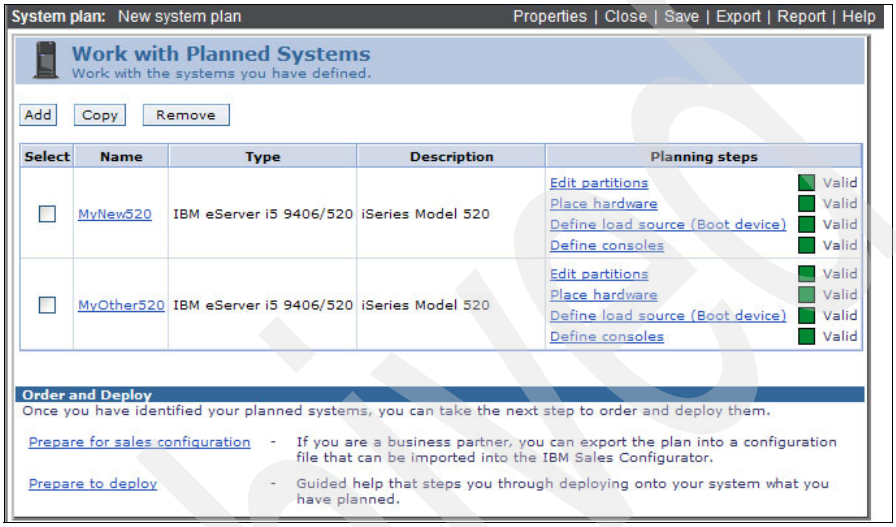

Figure 2-28 Valid status system plan

The Export panel opens and you have the ability to select the system you wish to export. By default all systems will be selected. If you have many systems you can select them all or check only the one you wish to export. In our example (see Figure 2-29) we have deselected MyOther520 and only the system plan for MyNew520 will be exported. Click **Export** to continue.

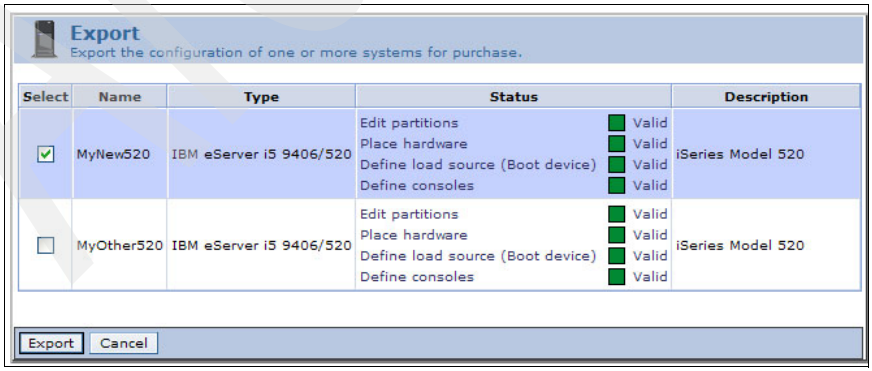

Figure 2-29 Select system to export

A standard file download panel appears with the default name of the .cfr file. Click **Save** to continue (Figure 2-30).

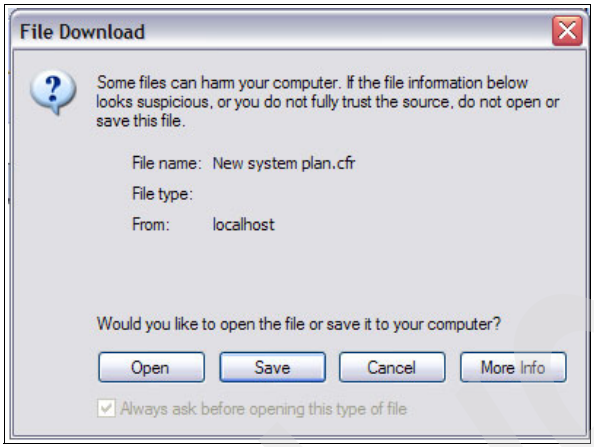

Figure 2-30 Download the exported system plan

When the Save As panel opens you have the ability to store the system in any folder. Select the default file name text **New system plan.cfr** and type over that with a name that is meaningful to your situation (maybe a combination of customer name and system model). In our example we use MyNew520 and then click **Save** (Figure 2-31).

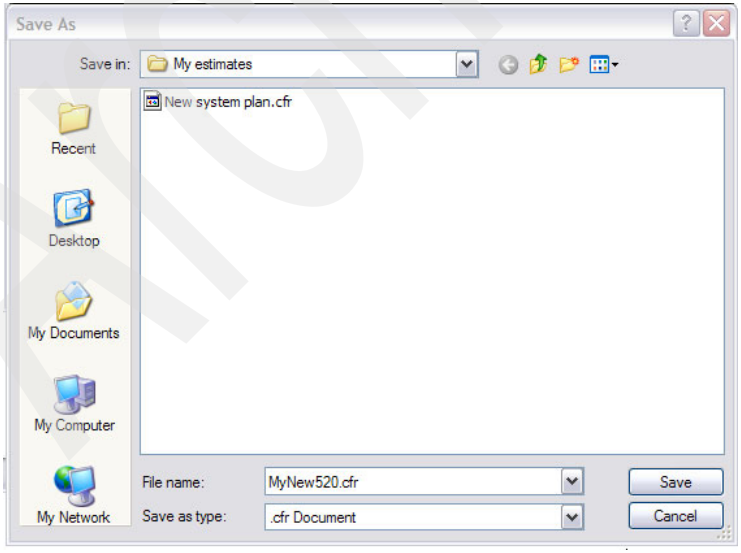

Figure 2-31 Saving the Exported system plan

The next step is to import this system plan into eConfig. Ensure that you have the latest eConfig file loaded on your PC. Open eConfig and browse to your saved system plan.

Once you have run eConfig you may find that there are some warning/errors not available in System Planning Tool. You may have to return to the system plan in SPT and make corrections to produce a valid eConfig. Once you have completed any changes required by eConfig, repeat the export steps to get the system into eConfig.

2.6.2 Saving a system plan for deployment

At some point the system is shipped and you will need to prepare to deploy the system plan. You should ensure that the HMC being used to manage the deployed is at the latest version of HMC code with all of the latest fixes installed. Visit IBM Fix Central to check for the latest information.

Open your existing system plan using the IBM System Planning Tool (SPT). It is very important that this system plan is the same plan that was used as input to order the system.

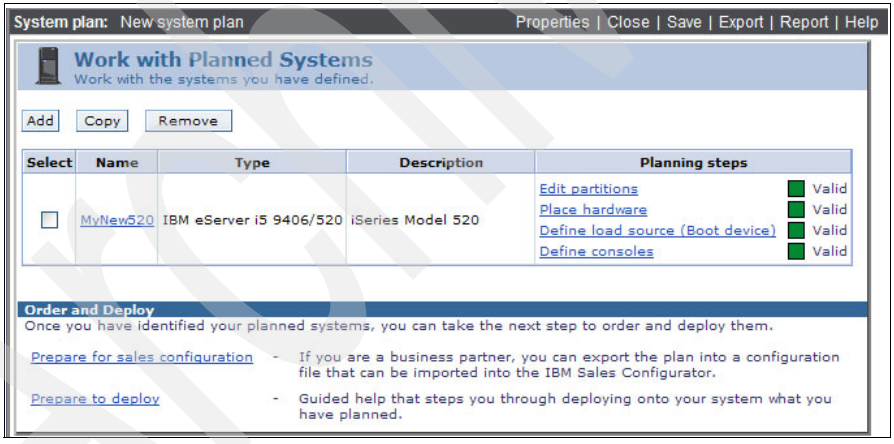

Figure 2-32 Existing system plan re-opened

From the header shown in Figure 2-32, click **Save**.

The File Download Panel will open as shown in Figure 2-33. Click **Save** to continue.

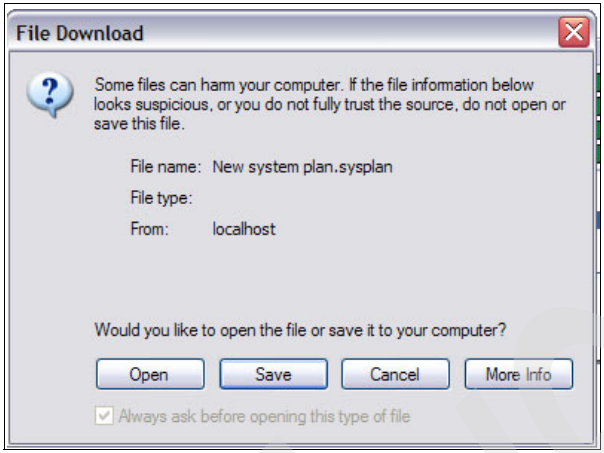

Figure 2-33 Save system plan

When the Save As panel opens you have the ability to store the system in any folder. Select the default file name text **New system plan.sysplan** and type over it with a name that is meaningful to your situation (maybe a combination of customer name and system model). In our example we use MyNew520.sysplan and then click **Save**.

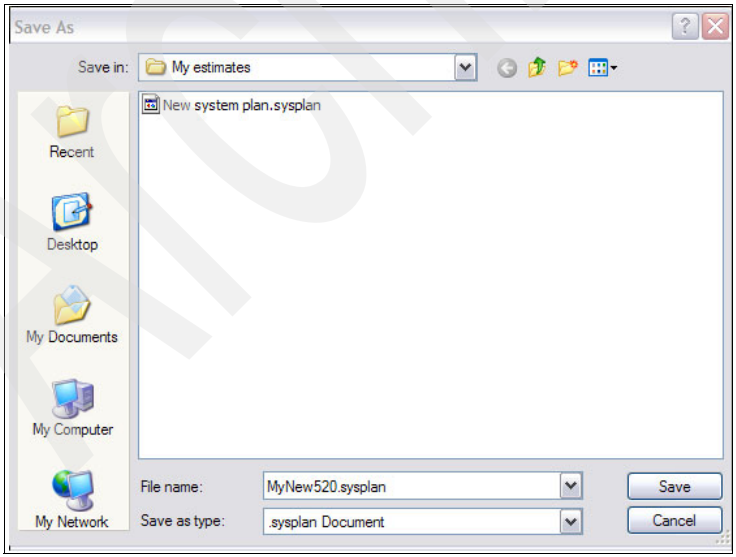

Figure 2-34 Editing the sysplan name during save

The .sysplan file is now available on disk and can be copied to media. You can then copy this file to the media required for the HMC, probably a USB memory key. Follow the instructions in 3.2.1, “Prepare to import a system plan” on page 49.

Depending on the system model and feature numbers ordered, the system may or may not come with all I/O in the correct positions. You should run the **mksysplan** command on the HMC before deployment to produce a view of the shipped hardware placement and configuration. The physical placement should also be checked, especially the disk locations. Devices like disk do not show in the system plan produced by the **mksysplan** command.

Important: If this is a new installation that includes a new HMC, the HMC guided setup tasks must be completed before a deployment can occur. If this is a new system being managed by an existing HMC, ensure that the HMC is at the correct code level to support the deployment process and the new system.

2.6.3 Viewing/printing a system plan for documentation

Once the system plan is complete (or has reached the *work with planned system* stage), you can view and print the information using the System Plan Report function.

From the Work with Planned Systems panel shown in Figure 2-35, click the **Report** link in the header.

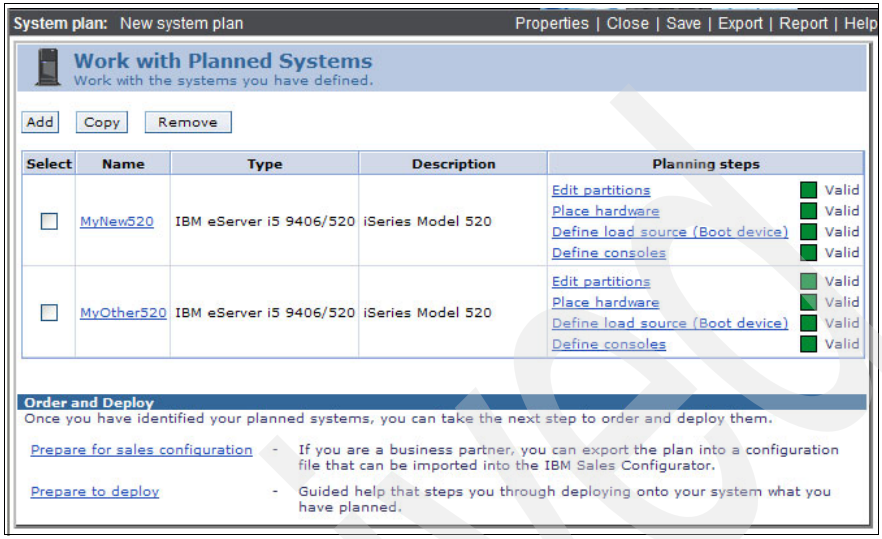

Figure 2-35 Producing a system plan report

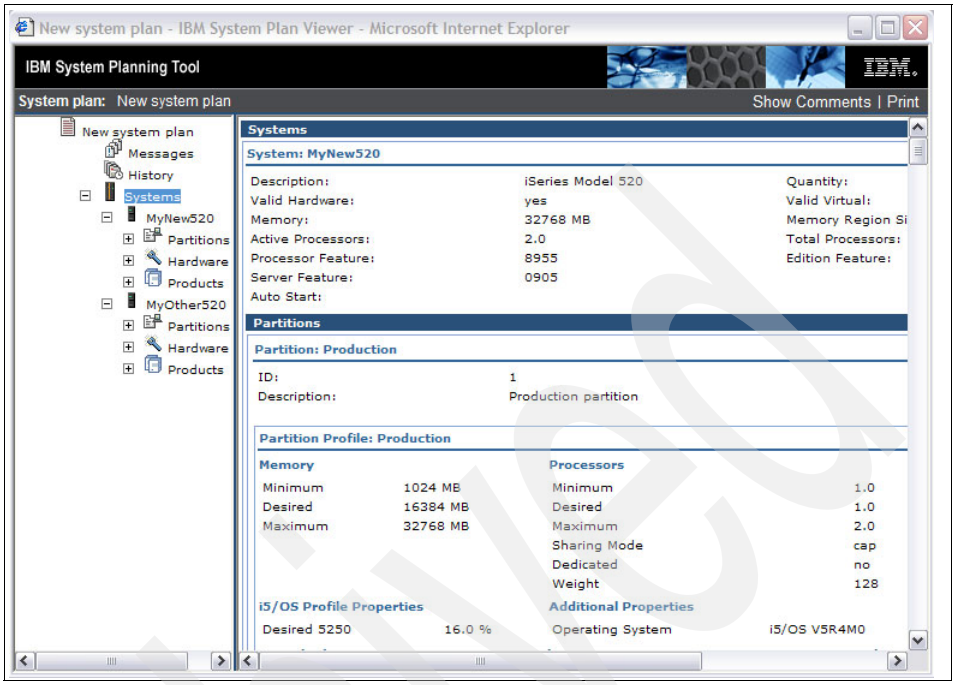

Figure 2-36 System plan report

The left-hand navigation pane is active and selecting a different line in the pane will display a view based on that selected in the right-hand work area. In our example we selected the system view. When you click the **Print** link in the header the view selected will be printed.

If you have Microsoft Excel installed on the PC running system plan you can export the System Plan Report view to an Excel spreadsheet.

3

System plans and the hardware management console

In this chapter we introduce another component of the LPAR Simplification Project that further utilizes the *Hardware Management Console (HMC)* as part of the LPAR design, deployment, and management processes. The HMC code level must be at V5R2 or later, and include the latest service packs.*System Plans* is a new feature that provides a graphical user interface to assist with the management and execution of the files that contain the LPAR configuration details. The utilities that comprise this new feature include import, export, view, remove, and deploy operations. An additional capability was added via the command-line interface that enables the HMC to create a system plan from an existing LPAR configuration.

The next few sections review the operations of each of the utilities for managing system plans, also referred to as *sysplans*. Most of the utilities are available through the HMC graphical interface. The view feature is only available using the graphical interface. Some functions are only available through the restricted shell, commonly referred to as the command-level interface (CLI).

Several new CLI commands are included in HMC V5R2. This chapter goes into detail about the commands that pertain directly to system plans and the

deployment of a system plan. Some of the new commands are **mksysplan**, **lssysplan**, **rmsysplan**, **deploysysplan**, and **cpsysplan**. The **mksysplan** command can be used to create a system plan from an existing POWER5 or POWER5+™ server. Once a sysplan exists on the HMC, the **lssysplan** command can list the plans, **rmsysplan** can remove or delete a sysplan, **cpsysplan** will create a duplicate copy either to or from removable media, and **deploysysplan** will validate hardware and LPAR configuration definitions and generate the series of commands that are used to create the LPAR configuration objects.

The **lsmediadev** command is also new and can be helpful to identify available media devices such as CD/DVD, diskette, or USB flash card or drive. By mounting one of these devices it can be used to export a sysplan file or to import a sysplan file that may have been created by SPT or created on another HMC. At this time there is no GUI equivalent to this command.

3.1 System plans

A system plan, also referred to as *sysplan,* is a representation or data model of what a system actually does consist of or what the LPAR design should consist of, depending on how the sysplan file is generated. If generated from an existing POWER5 or POWER5+ system using the **mksysplan** restricted shell command, the file will reflect the actual LPAR configuration of the server at that point in time. 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 I/O hardware required for each partition. Hardware allocations reflect whether it is defined as required for the partition to activate or whether it is hardware that is optional for the partition to activate and can be dynamically switched between two or more partitions.

The sysplan includes 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 Central Electronic Complex (CEC) and any I/O expansion towers or I/O drawers that will attach to the CEC using HSL or RIO loops. The card slots identify IOP or IOA feature codes, and this level of detail is used for 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 using the **mksysplan** command does not include any detail of what is attached to an IOA.

A system plan file always has a file extension type of.*sysplan*. 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 system plan graphical interface, the file is stored on the HMC in a predefined directory. The directory path is */opt/hsc/data/sysplan*.

3.2 System plans - HMC graphical user interface

Starting with the HMC V5R2 code level, there is a new task menu called *System Plans*, which is displayed on the HMC main task window. This feature can be accessed directly from the HMC or remotely using the *Web based System*

Management (WebSM) client connecting to the HMC. Either single-click the **System Plans** task in the navigation area or double-click the large icon to the right (Figure 3-1). This will display the management tasks window.

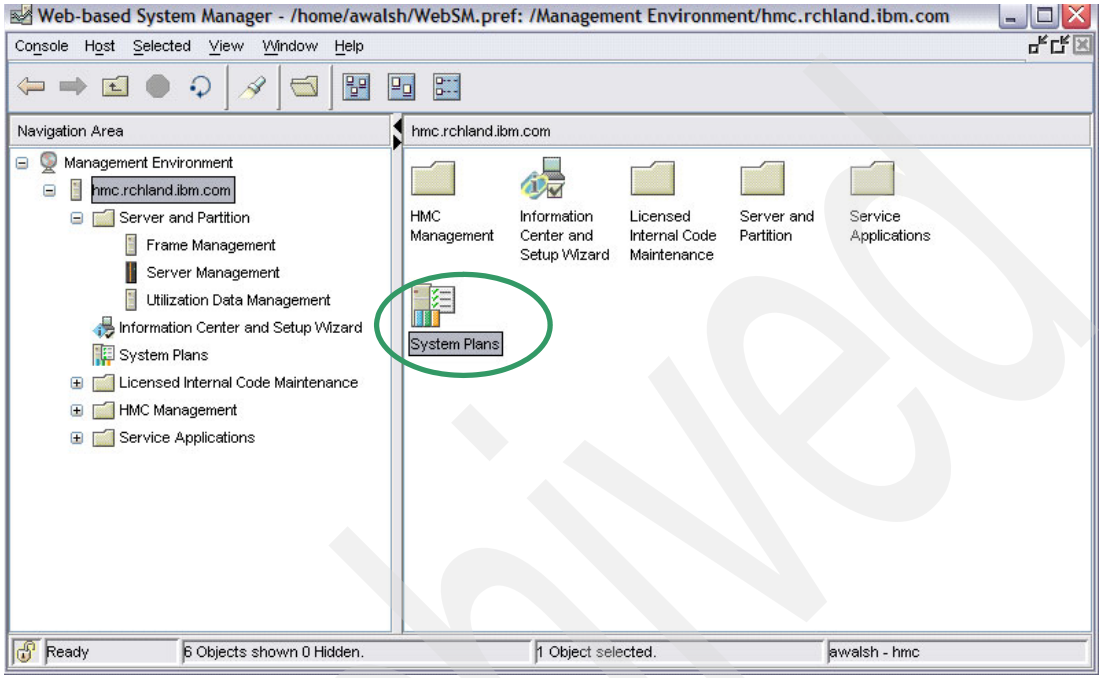

Figure 3-1 HMC main tasks panel

The main panel of the System Plans task menu provides options: Import System Plan, Deploy System Plan, and Manage System Plans (Figure 3-2). There is also an option *Learn more about system planning*, which is a link that takes you directly to the IBM Systems Information documentation on planning for logical partitions.

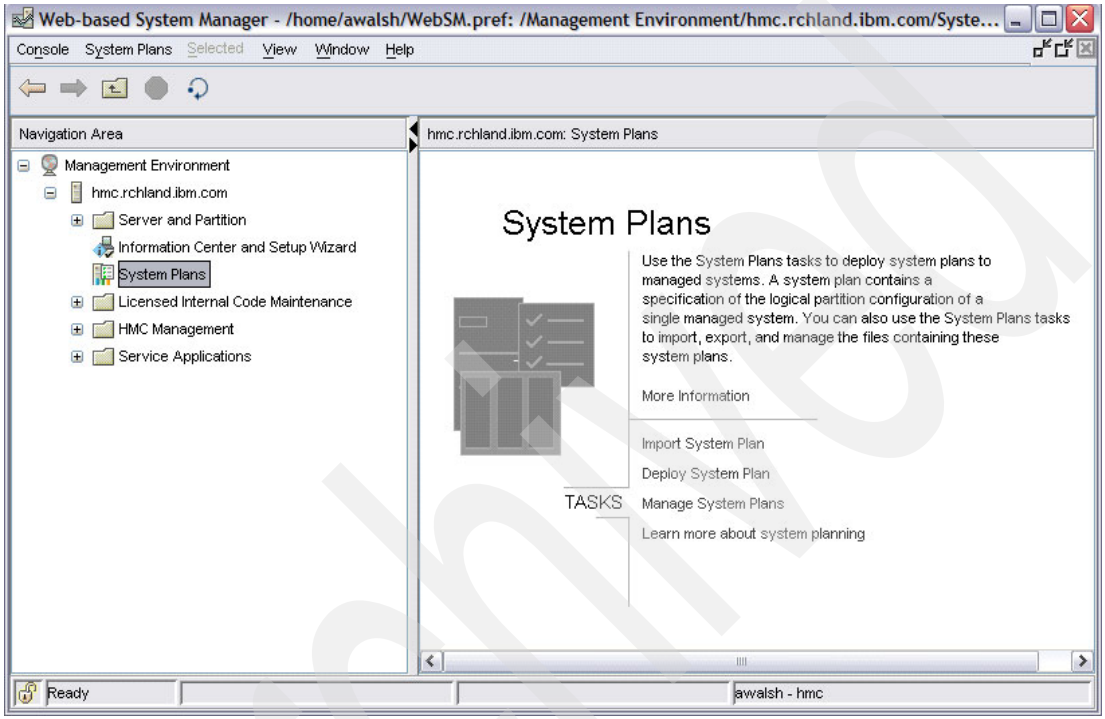

Figure 3-2 System Plans main menu panel

3.2.1 Prepare to import a system plan

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 then be transported on one of the supported media types, such as CD, DVD, diskette; a USB device such as a memory card; or a FTP server. This operation prompts for the location of the sysplan file to be loaded. If the sysplan file is on external media, the first step is to ensure that the media device is mounted. The easiest method of doing this is using the Format Removable Media task.

Go to **License Internal Code Maintenance → HMC Maintenance → Format Removable Media**. This feature scans for all available media devices and ensures that the device you select to format is mounted. See Figure 3-3.

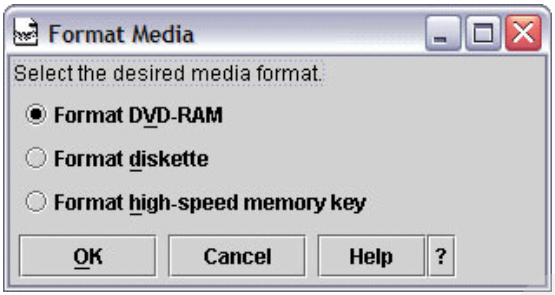

Figure 3-3 Formatting HMC media

If you have a USB memory key inserted in a USB slot on the HMC, select **Format high-speed memory key** and click **OK**. The memory format process starts and completes. The memory is now ready to take to the PC containing the system plan file. Insert the memory key in the PC and load the system plan file using the save function in SPT or by browsing to the file and copying to the memory stick.

If the memory stick is not found when you try to format it, you can check to see whether the HMC OS has found the device from an HMC command line. Simply position the cursor outside the HMC window and right-click. Then select **Terminals → rshterm**. To list available media devices the command is **`lsmediadev`**. In our example the memory stick device is not loaded and there is not type=3 device. Below is the list media command before the USB Flash Memory Device is inserted in the USB port:

```
lsmediadev
```

The results of running the command are:

```
device=/dev/cdrom,type=1,description=CD/DVD
device=/dev/fd0,type=2,description=internal diskette
drivedevice=/dev/hda,type=6,description=internal hard disk drive
```

To mount the USB Flash Memory Device, use the **`mount`** command. The first USB port will appear as SCSI device 1 (sda1), so the **`mount`** command would look like this:

```
mount /media/sda1
```

Now we again run the list media command:

```
lsmediadev
```

The results show that a type=3 device USB flash memory device is mounted:

```
device=/dev/cdrom,type=1,description=CD/DVD
device=/dev/fd0,type=2,description=internal diskette drive
device=/dev/sda1,type=3,description=USB flash memory device
device=/dev/hda,type=6,description=internal hard disk drive
```

3.2.2 Import a system plan

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 or from an FTP server. This example shows the import prompt window and that the system plan file is stored on a USB flash drive. The name of the file is jpmulhol.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/dev.

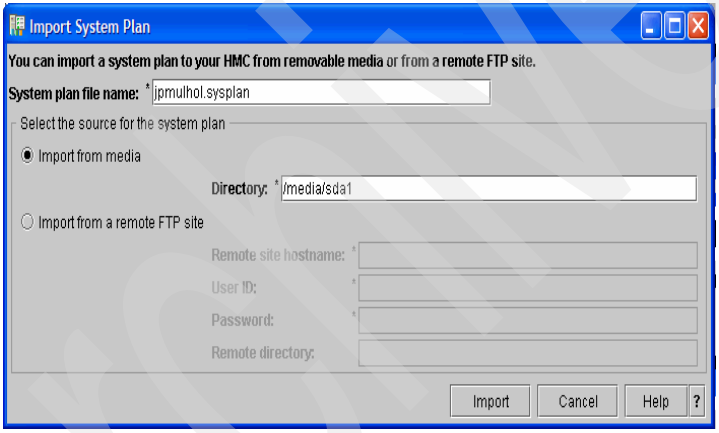

Figure 3-4 Import System Plan task panel

Following an attempted import operation, the following status/completion window is displayed (Figure 3-5).

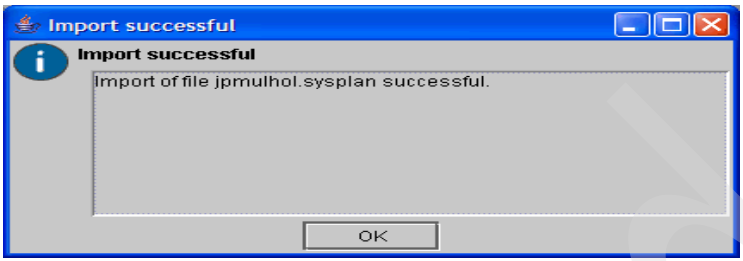

Figure 3-5 Import task completion panel

3.2.3 Manage system plans

On the main System Plans task menu are the additional operational choices of Deploy System Plan or Manage System Plans. If we select the Deploy option, the next window is the Deploy System Plan Wizard - Welcome panel. This option, along with the Import option, can be initiated from the Manage System Plans menu.

So let us take a look at the Manage System Plans panel next (Figure 3-6). Notice that if an existing plan is not highlighted, only the Import button is enabled. Once an existing plan is highlighted, the View, Deploy, Export, and Remove buttons become enabled.

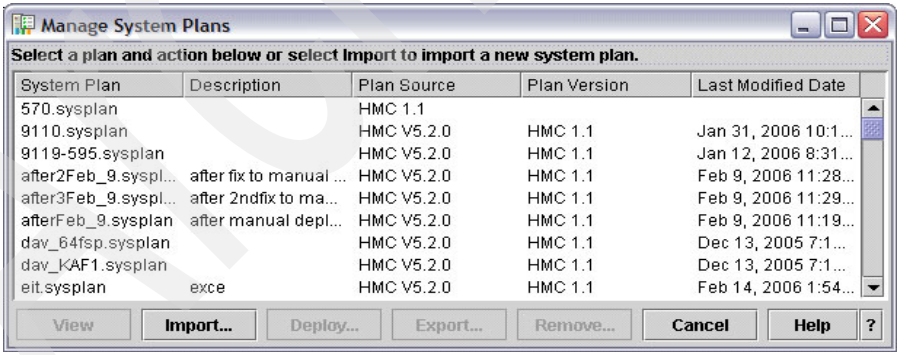

Figure 3-6 Manage system plan import button

3.2.4 View a system plan

From the Manage System Plans panel, highlight the sysplan that you wish to see, then select the **View** tab (Figure 3-7).

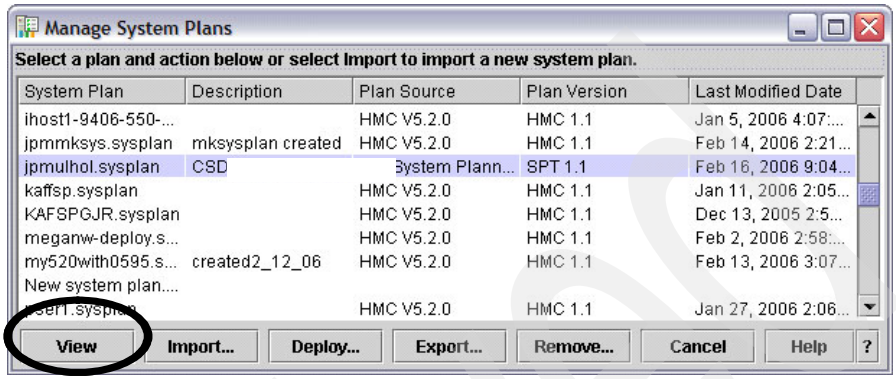

Figure 3-7 Manage system plan view button

This launches a browser session to display the details of the system plan (Figure 3-8). The first time you attempt to view a system plan using a WebSM client, you may be prompted for an authorized HMC user ID and password to access the System Plan Viewer. The HMC deploys a similar viewer to that used by the SPT, so the look and feel will be close, but less information is available.

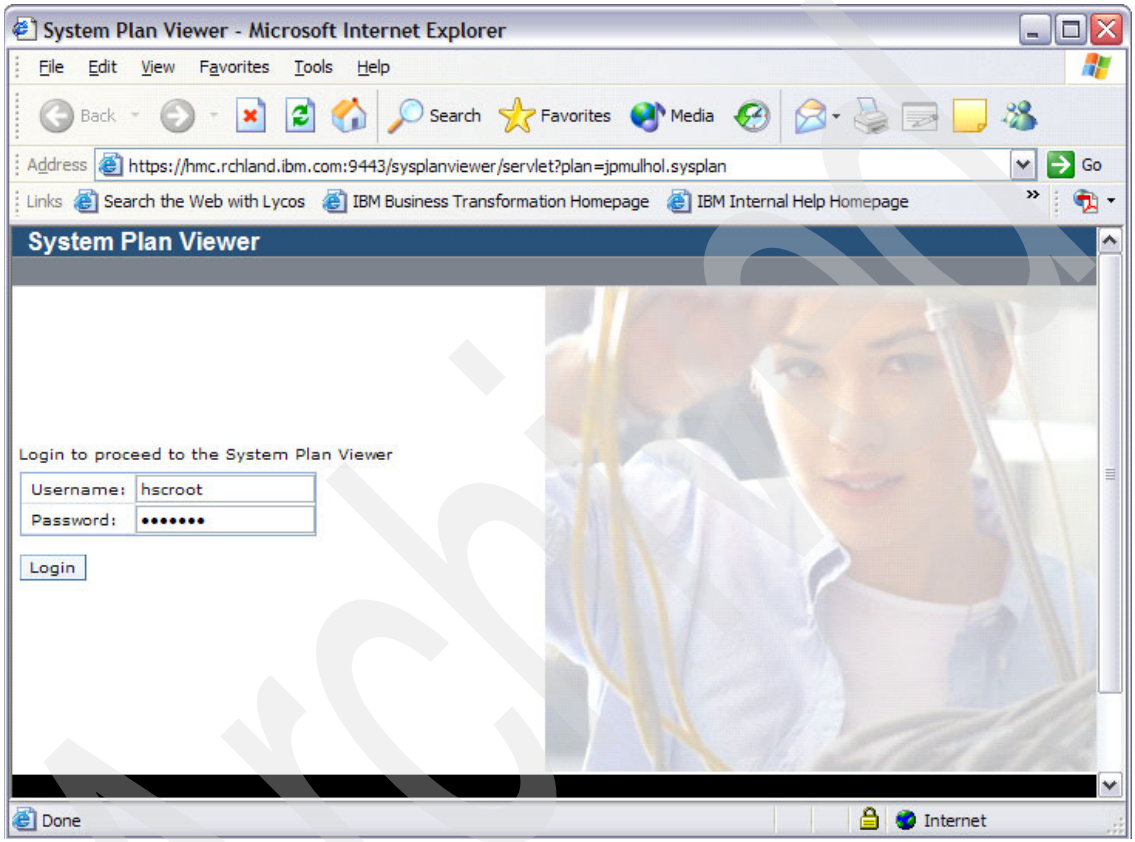

Figure 3-8 System Plan View signon

After satisfying the authentication process, the first page of the system plan is displayed. Notice the navigation bar on the left, which allows us to selectively view portions of the system plan. We can view a system summary, view details on a specific partition, or view the hardware within the CEC unit or by individual tower.

Figure 3-9 is the first page of a system plan that shows a system summary. This is an example of a plan that was created using SPT.

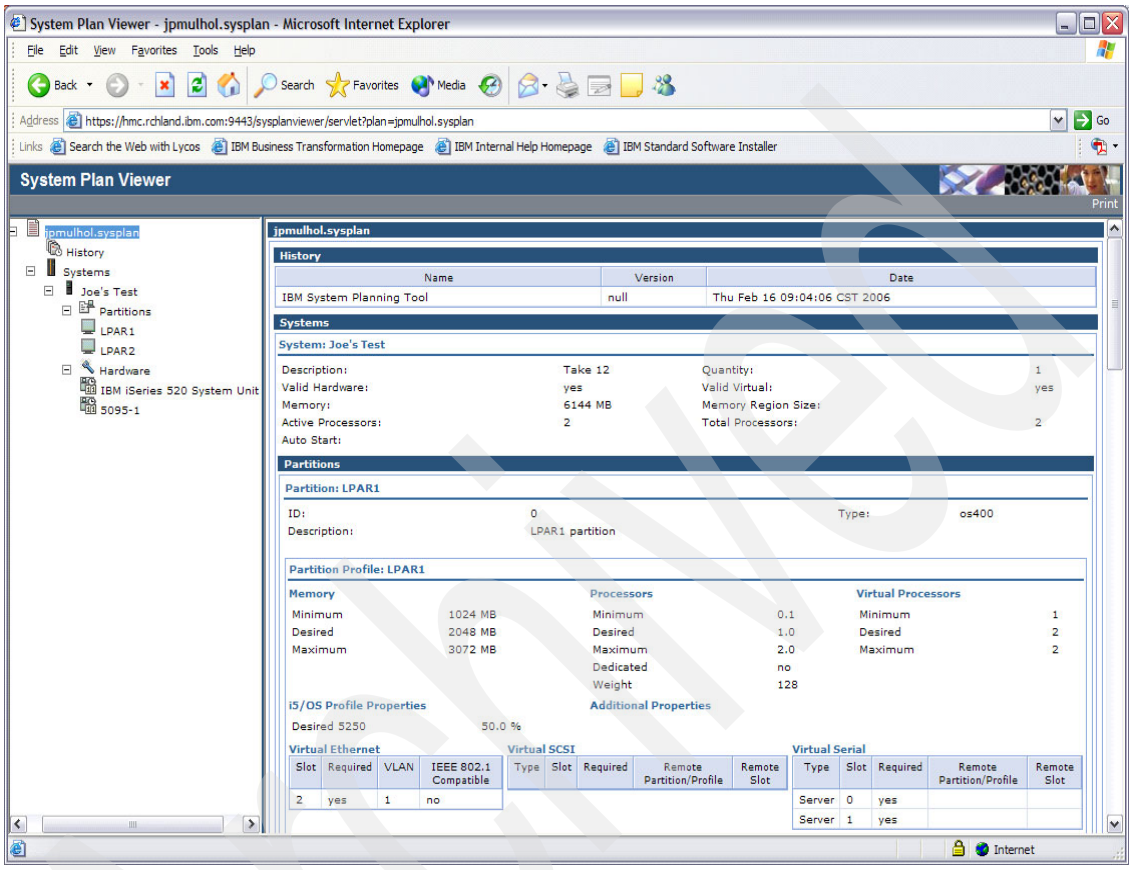

Figure 3-9 System view system summary

The main difference between a system plan created today using the CLI **mksysplan** is that the SPT generated sysplan file includes device-level information showing what is planned to be attached to the IOAs. The mksysplan sysplan file does not include that level of detail and so the viewer cannot show any detail below the IOA.

The advantage of viewing a sysplan (mksysplan) that represents the actual hardware installed on a system and be able to compare that to a sysplan plan that is about to be deployed or from a deploy attempt that failed, is that by comparing the two it may help reconcile any differences. The deploy process must meet system and hardware validation, and a misplaced card could cause a failure.

Figure 3-10 simply shows a view of the hardware installed on one of the 0595 towers. Notice that it does show a feature code for the IOP or IOA in each card slot. If this sysplan had been generated by mksysplan it is very likely that you will see more than one feature code listed for many of the IOAs. The reason for this is that the hardware will report in by the CCIN number, which may map to more than one feature code depending on which OS the card was ordered for. The column on the right will show whether the resource is assigned to a specific partition or whether it was configured as a switchable resource. More that one partition name may be displayed. This view also shows the feature code, characteristics, and number of disk units attached to the storage IOA. Again, this view is currently only available when the sysplan is generated by SPT.

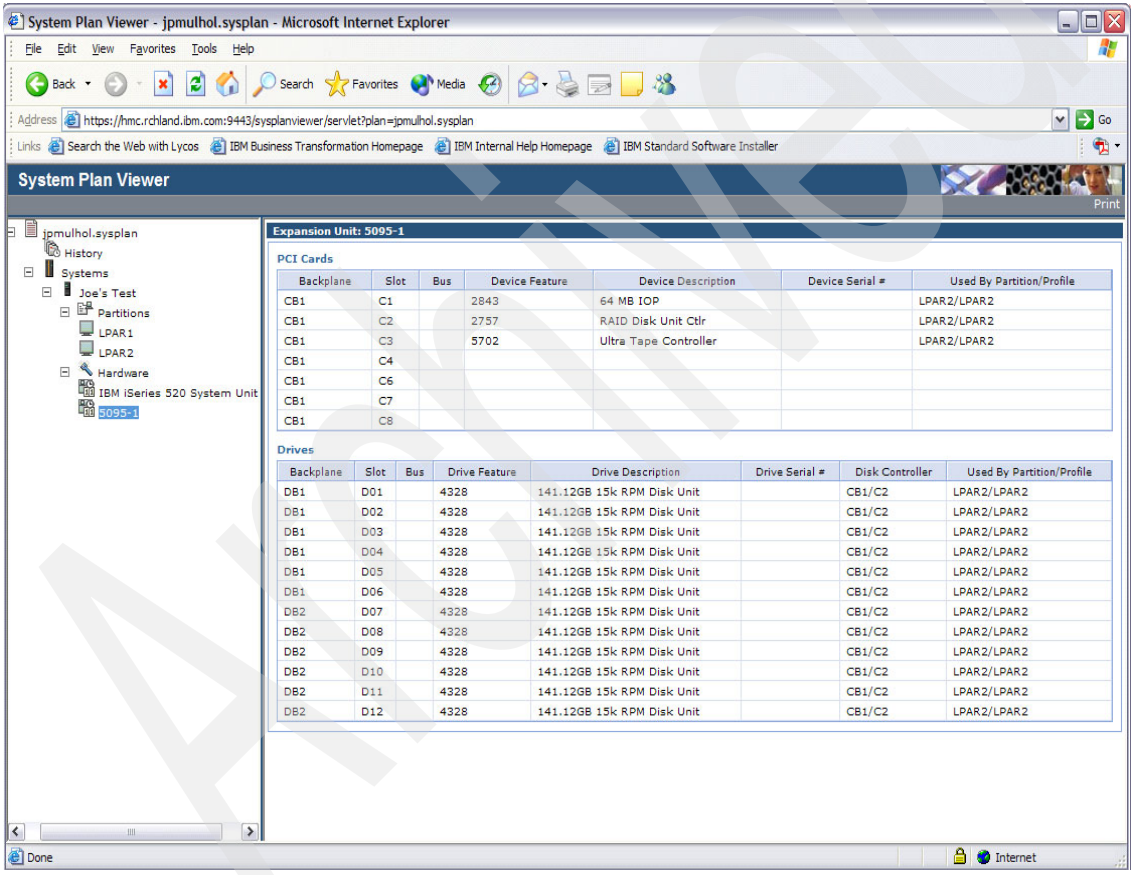

Figure 3-10 System view PCI Card view

3.2.5 Export a system plan

This operation is very nearly identical to the import operation, only in reverse. The export procedure is initiated from the Manage System Plan panel. The requirements are that the sysplan to be exported must exist on the HMC, and if the target for the export operation is media, the media must be on a mounted device.

In this example we export the sysplan that was previously imported from a flash drive. The export operation targets a FTP server where the plan could subsequently be imported by other MKS. Highlight the desired plan and select the **Export** button (Figure 3-11).

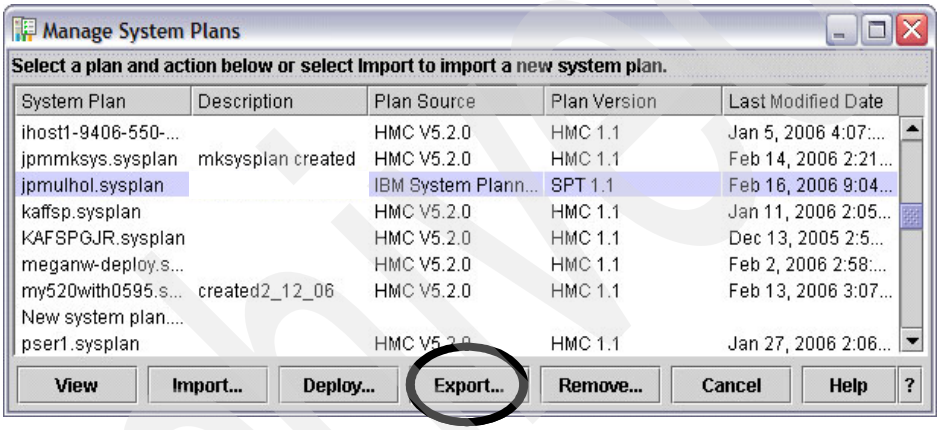

Figure 3-11 Manage system plans export button

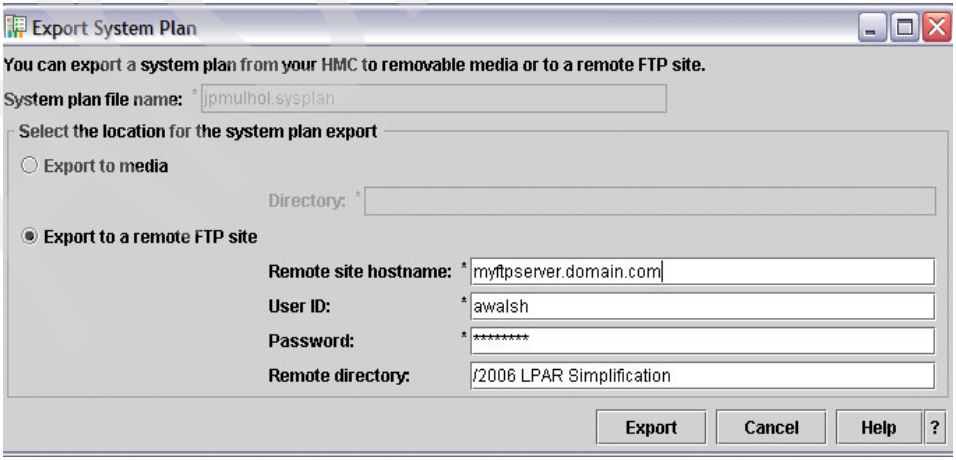

Figure 3-12 Export system plan to FTP server

The Export operation returns a message that the `Export of file xxx to (location) was successful` or an error message if it fails. For example, Figure 3-13 shows the error message when attempting to export to a flash drive that had not been mounted.

Figure 3-13 Flashdrive media error message

3.2.6 Remove a system plan

From the Manage System Plans panel, simply highlight the sysplan that you wish to remove and select the **Remove** button (Figure 3-14). This deletes the sysplan file from the HMC hard drive.

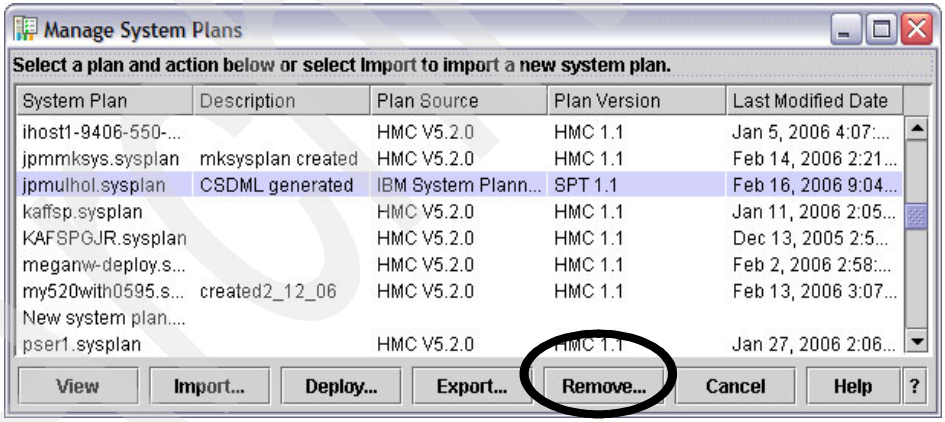

Figure 3-14 Manage system plan removal button

The confirmation panel will be presented showing the same of the sysplan that you highlighted. Select the button **Remove System Plan** and the file will be deleted (Figure 3-15).

Figure 3-15 System plan removal

3.2.7 Deploy a system plan

To successfully deploy a sysplan is the primary objective and may well be the most important component of the LPAR Simplification project. In this section we review the operational steps of deploying a sysplan. In a later section we delve deeper into the deployment validation process. Understanding the validation process can help ensure success and that the resulting LPAR configuration correctly matches the design. Deploy can be initiated from the main System Plans task panel or from the Manage System Plans panel (Figure 3-16).

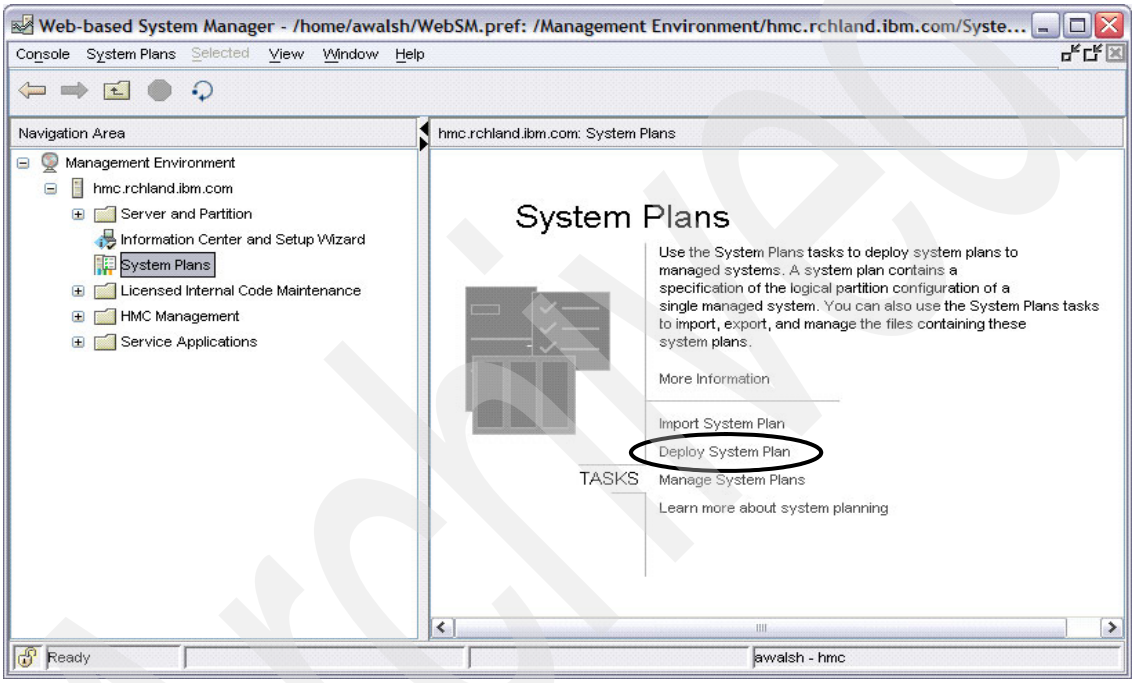

Figure 3-16 System Plans welcome

If you select the Deploy button from the Manage Systems Plans panel, the highlighted sysplan will be selected. If you select the Deploy System Plan from the main System Plan task window, the next panel to display will be the Deploy System Plan Wizard - Welcome screen. You can select the system plan to deploy or change the selection using the drop-down selection list (Figure 3-17).

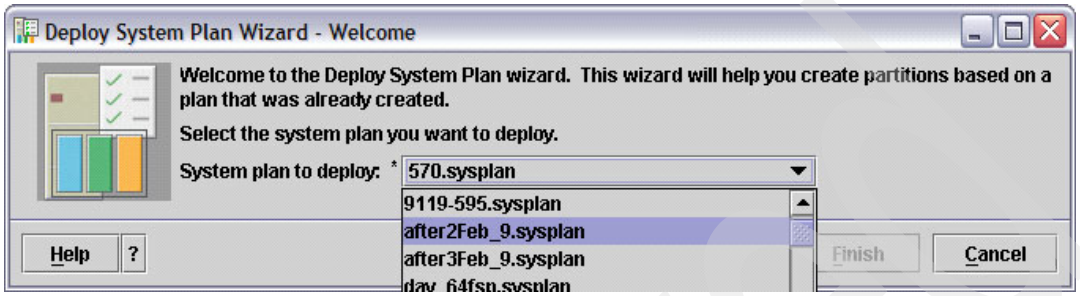

Figure 3-17 Deployment system plan selection

Click **Next** once the correct sysplan is verified.

The Deploy System Plan Wizard - Select System panel is displayed next. This screen asks you to select the managed server that will be the target of the verification and deploy. Use the drop-down menu to display all of the servers that are being managed by this HMC (Figure 3-18). Highlight your selection and then click **Next**.

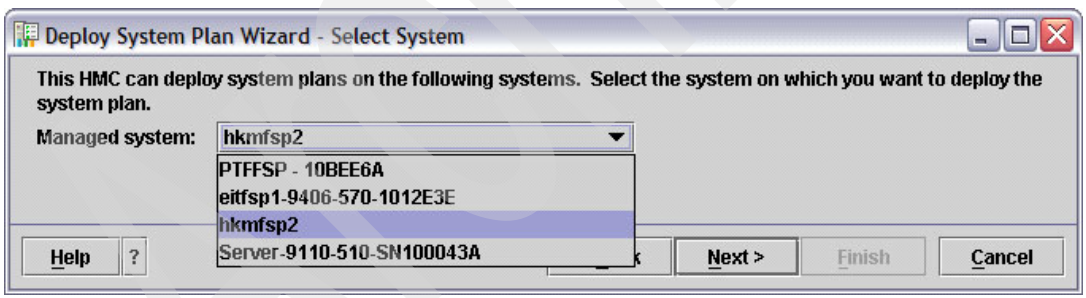

Figure 3-18 Deployment system selection

Now the Deploy System Plan Wizard - Validation screen appears. This shows the two major validation steps and the status for each. The first step is hardware validation and then partition validation. Initially each step has a status of *in progress*, which changes to either successful or failed.

If the hardware validation step should fail, the process is ended without going on to the partition validation. For example, Figure 3-19 is a failed hardware validation. The sysplan was designed for a model 520 with a single 0595 expansion tower. Two partitions would be created on this server if the deploy were successful. However, the wrong server was selected on the Deploy System Plan Wizard - Select System screen. Server-9110-510-SN100043A was picked, but the sysplan was designed for system hkmfsp2.

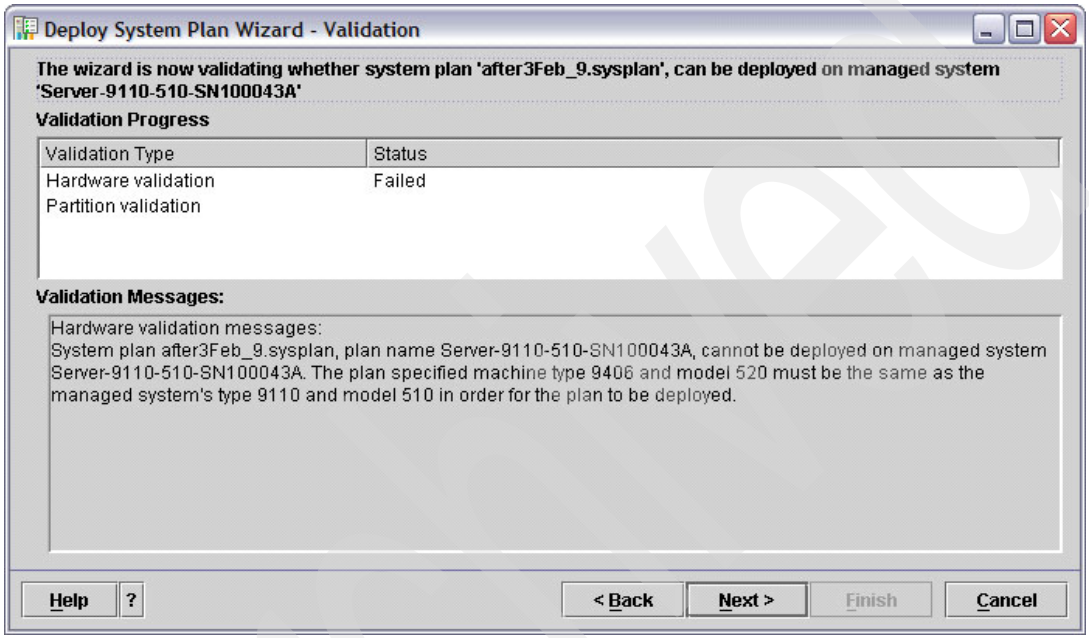

Figure 3-19 Deployment validation failure

In this example of a failed deploy, the validation messages show that the sysplan specified a machine type of 9406 and model 520. However the actual system that the plan was validated against, Server-9110-510-SN100043A, is a model 510. This explains why the system validation failed. Had this been a valid 9406, model 520, but a different server feature code, insufficient number of processors activated, or less than the specified total memory, the validation would have failed.

A successful deploy shows a status of successful for both validation steps and the validation messages show additional details. The sysplan in this example was created using **mksysplan** and then deployed back on to the same system. Before starting the deploy, both partitions that existed when **mksysplan** was run were deleted. This left the 520 in this test without any partitions configured. Notice that the detailed messages tell us that the sysplan included serial

numbers identified for the CEC unit and also with the tower, and these serial numbers matched the actual hardware found during the validation. This in an important point. There may be times when it is necessary to include the serial number for the physical towers attached to the system. This ensures that the partition resource allocations are assigned to the specific towers that the LPAR designer intended. More information about this topic is in Chapter 5, “Deployment validation” on page 87.

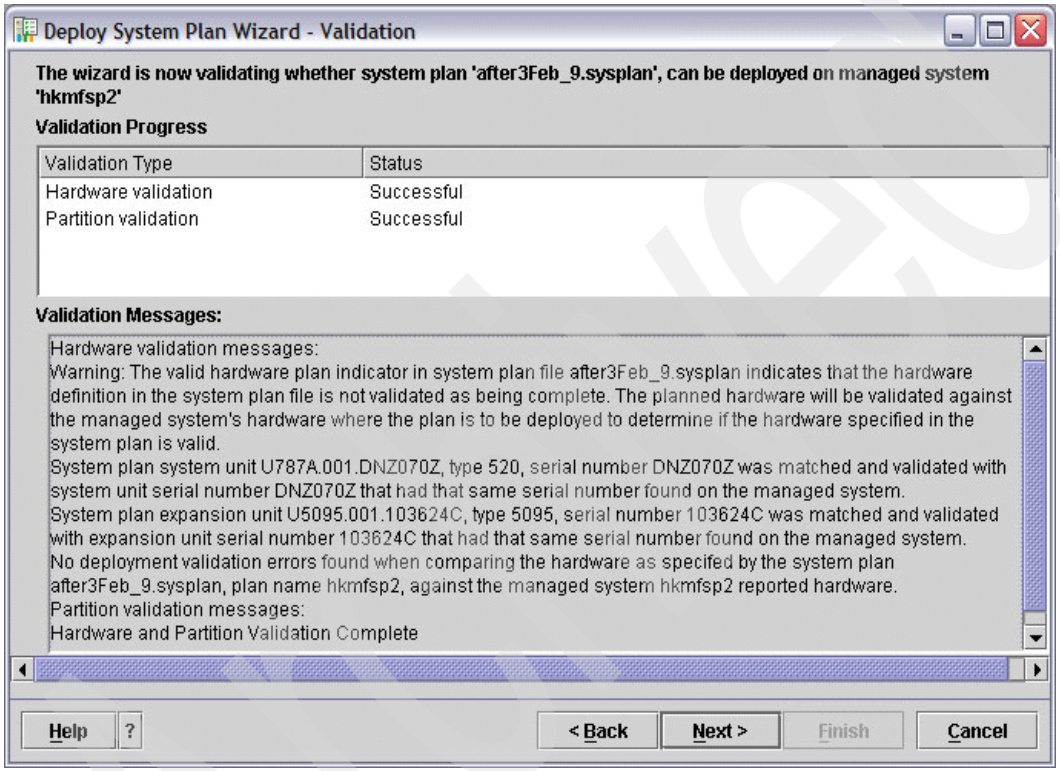

Figure 3-20 Hardware and Partition validation

Once the validation steps complete click **Next** and the Deploy System Plan Wizard - Select What to Deploy window appears. This shows a summary of the partitions and partition profiles that will be created. Notice the column marked Deploy that has everything check marked. From here you have the option to selectively deploy all or only part of the sysplan.

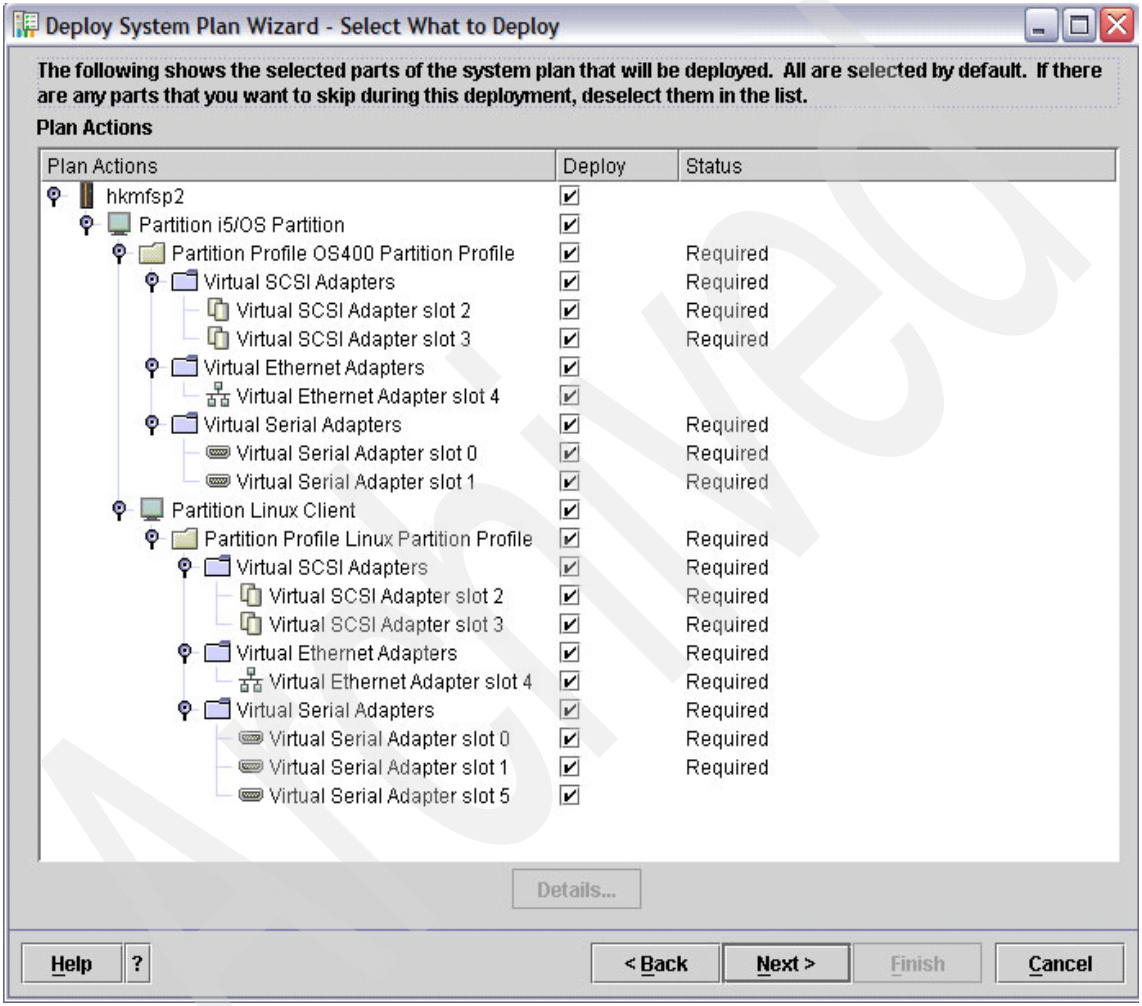

Figure 3-21 What to deploy summary

In the example above, if we choose to deploy only the i5/OS partition, we could simply uncheck all of the Linux partitions. However, the deploy process also validates dependencies. So in this example, if we tried to deploy only the Linux partition and not the i5/OS, the deploy would fail because the Linux partition is dependant upon the i5/OS as its hosting partition. To deselect the Linux partition

for deployment, uncheck the Linux Partition Client tag. This automatically unchecks all of the tags for the Linux partition. For this example we initiate the deploy only for the i5/OS partition. Select **Next** to activate the deploy (Figure 3-22).

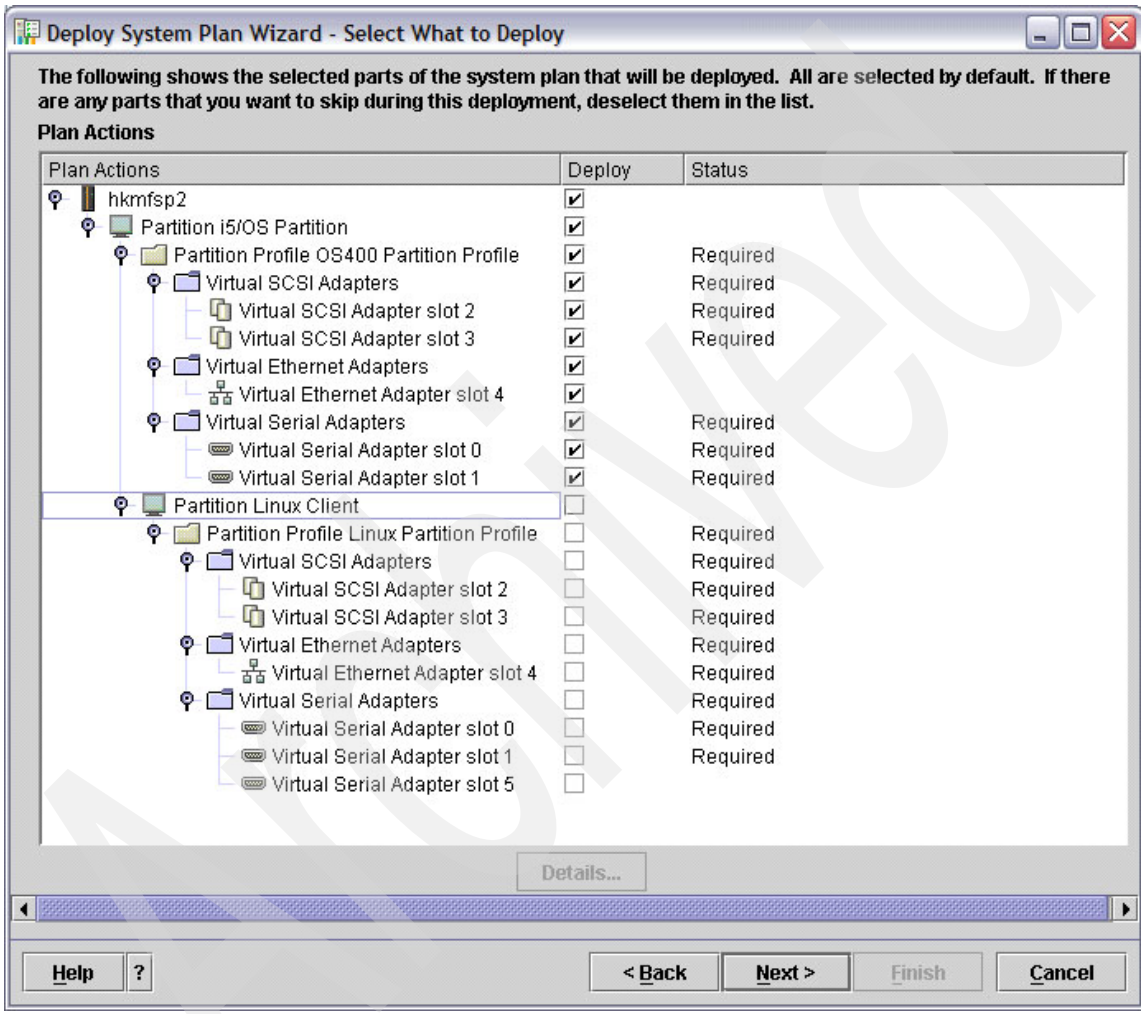

Figure 3-22 Partition deselected from deployment selection

Notice the Details button. From the Select What to Deploy panel you can highlight any component of the plan and click this button. This launches the System Plan Viewer and allows you to verify the specific details of the plan. For example, if you highlight the OS400 Partition Profile line the viewer takes you to the partition summary of the sysplan.

The Deploy System Plan Wizard - Summary panel is displayed and it shows each of the process steps that will be executed during the deploy (Figure 3-23). Selecting **Finish** on this panel activates the deploy process. Note the warning near the bottom of the panel that states `Please DO NOT perform any other action on this managed system while the deployment is running`. The deploy process time varies depending on the number and complexity of the partitions, but generally it ranges from 5 to 20 minutes.

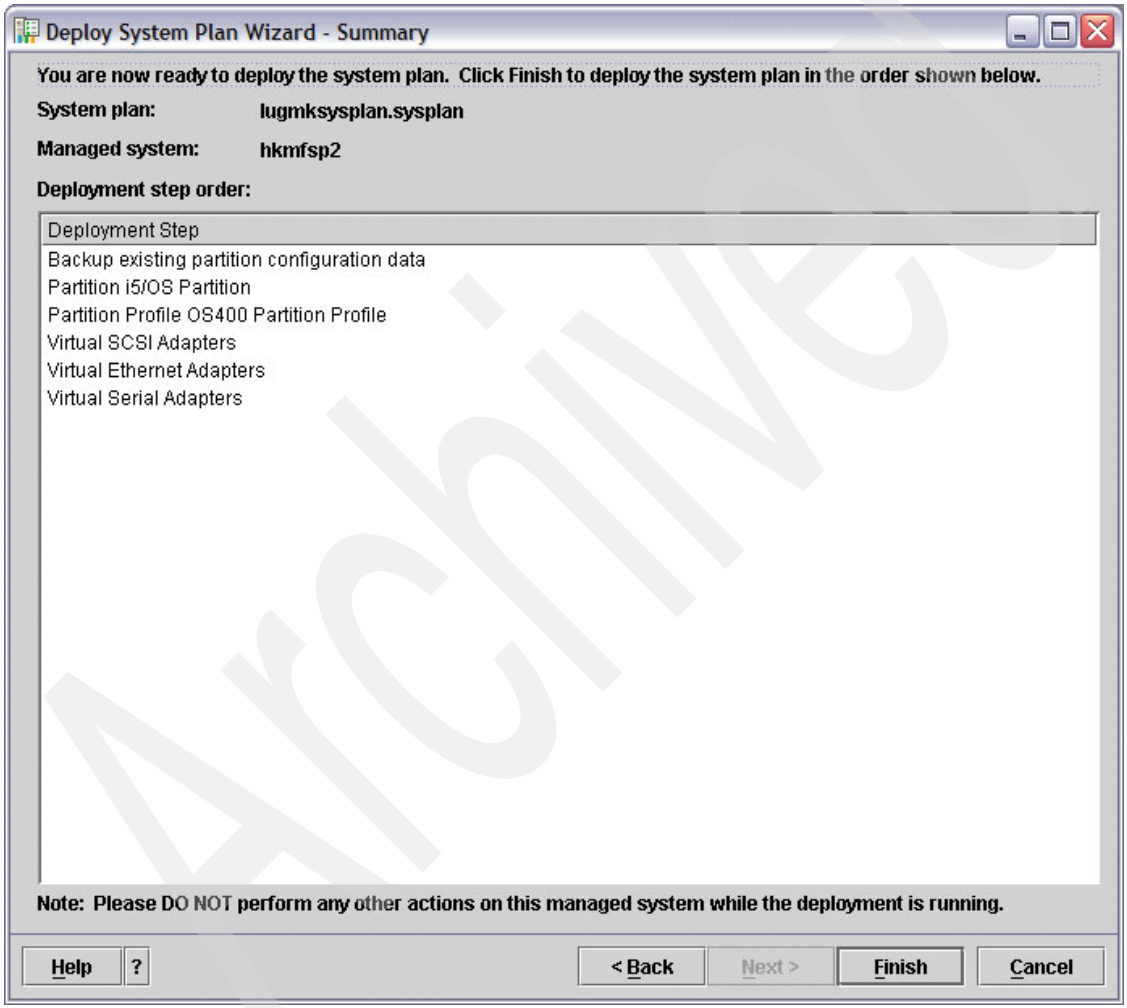

Figure 3-23 Deployment summary

The Deploy System Plan Wizard - Deployment Progress window shows the status for each step and any corresponding messages that provide the details on

each process (Figure 3-24). When the deploy completes the Deployment Complete panel is displayed (Figure 3-25 on page 68).

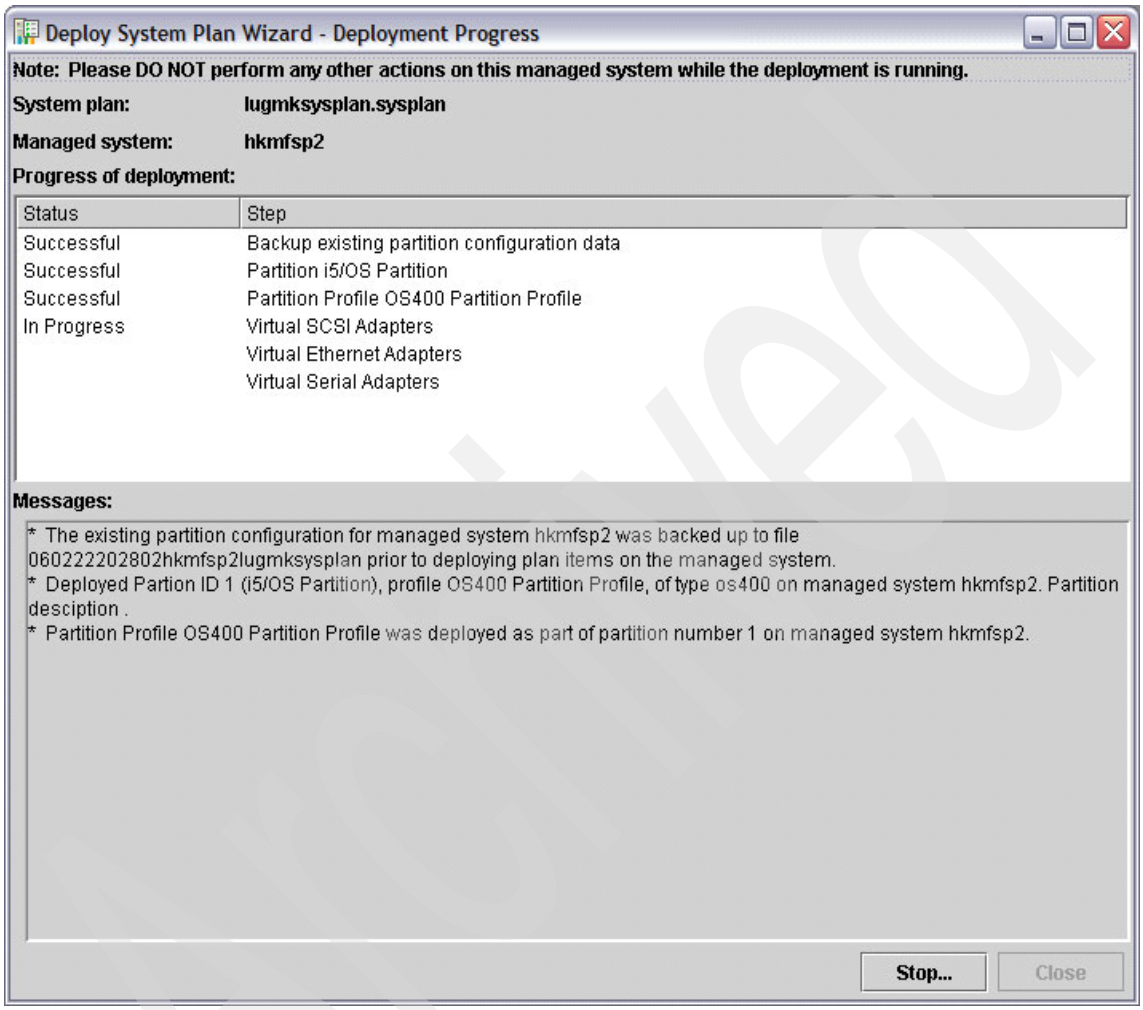

Figure 3-24 Deployment progress

Figure 3-25 Deployment complete

If we go to the HMC Server Management panel and look at the partition information for the managed server (hkmfsp2) that we just deployed this sysplan on, we will see a single i5/OS partition (Figure 3-26).

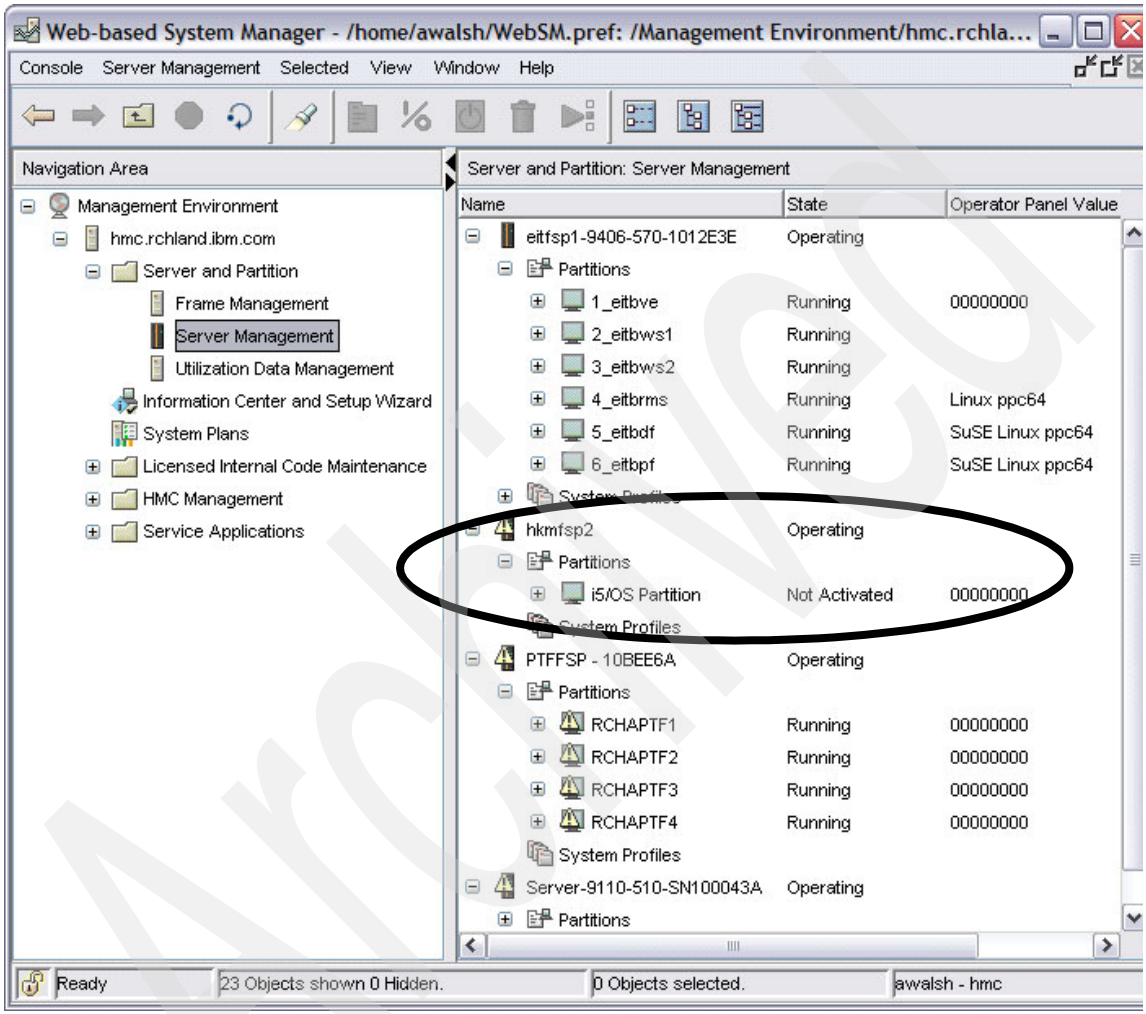

Figure 3-26 Deployed partition

Next we deploy the same sysplan file on the same server. This time the Deploy System Plan Wizard - Select What to Deploy window shows us that the i5/OS partition has already been deployed and the deploy column is grayed out (Figure 3-27). Select **Next** to bring up the Deploy System Plan Wizard - Summary panel. Select **Finish** to activate the deploy.

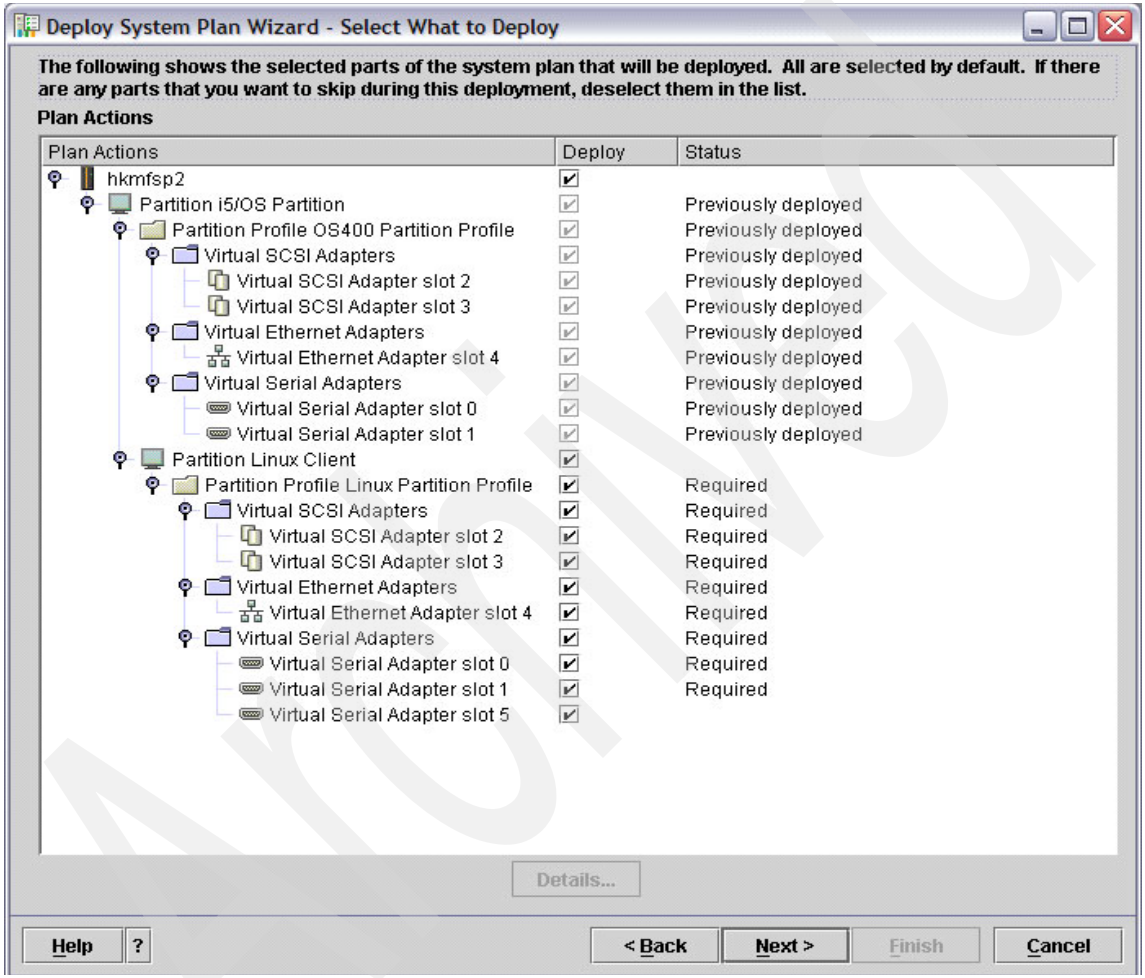

Figure 3-27 Select what to deploy

The Deploy System Plan Wizard - Deployment Progress window shows the status for each step of the deploy process for the Linux partition (Figure 3-28 on page 71). When the deploy completes the Deployment Complete panel is displayed.

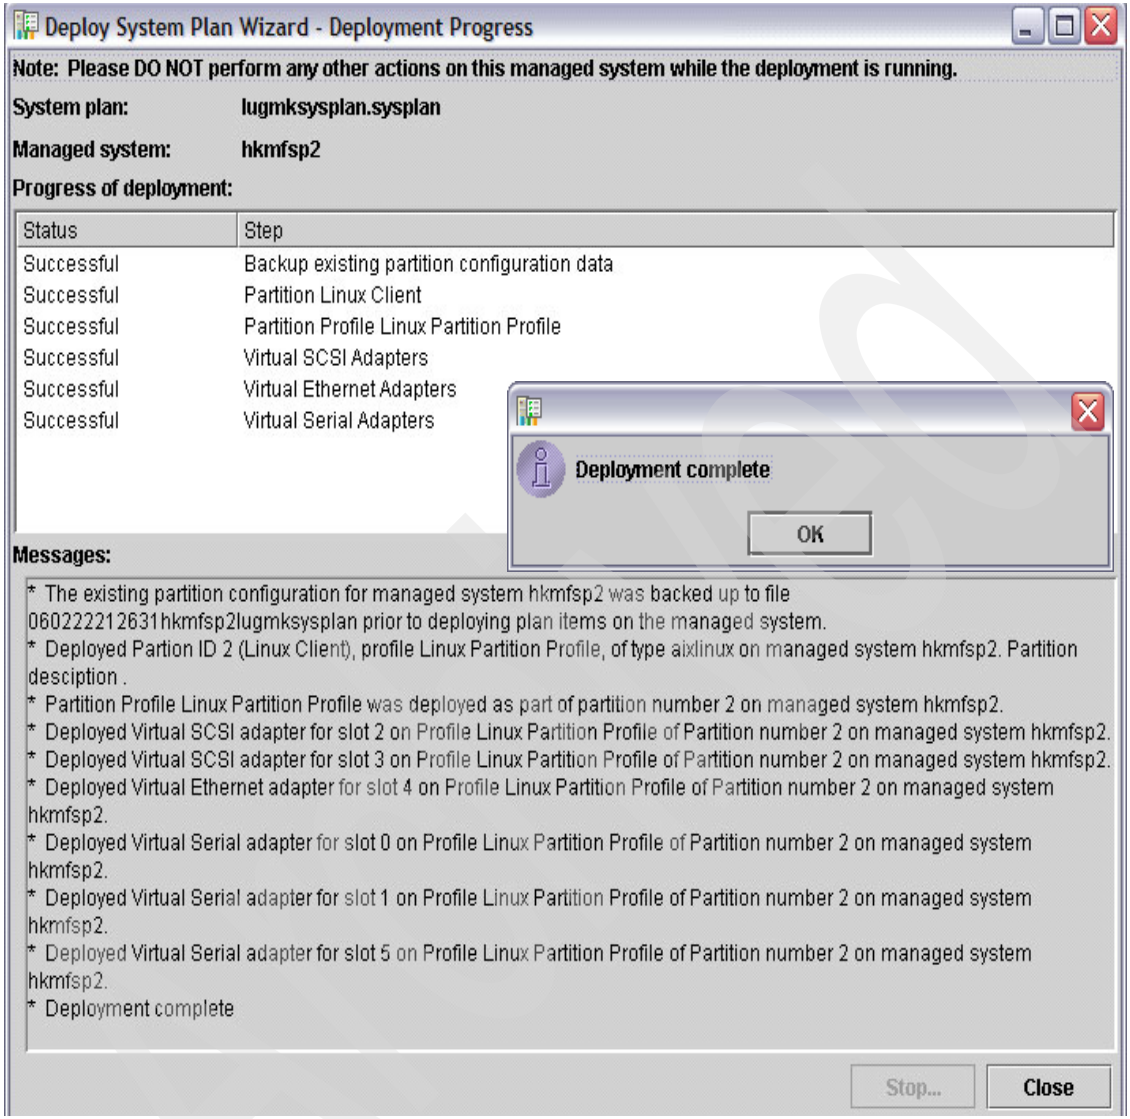

Figure 3-28 Deployment completion wizard

If we return to our HMC Server Management panel we now see both the i5/OS partition and the Linux partition for the hkmfsp2 server (Figure 3-29 on page 72). The state shows a status of Not Activated. The partition profile name is the same

at the partition name and at this time the deployment process only supports creating a single partition profile. If you need additional profiles, these must be created manually using the HMC LPAR wizard after the deploy has completed.

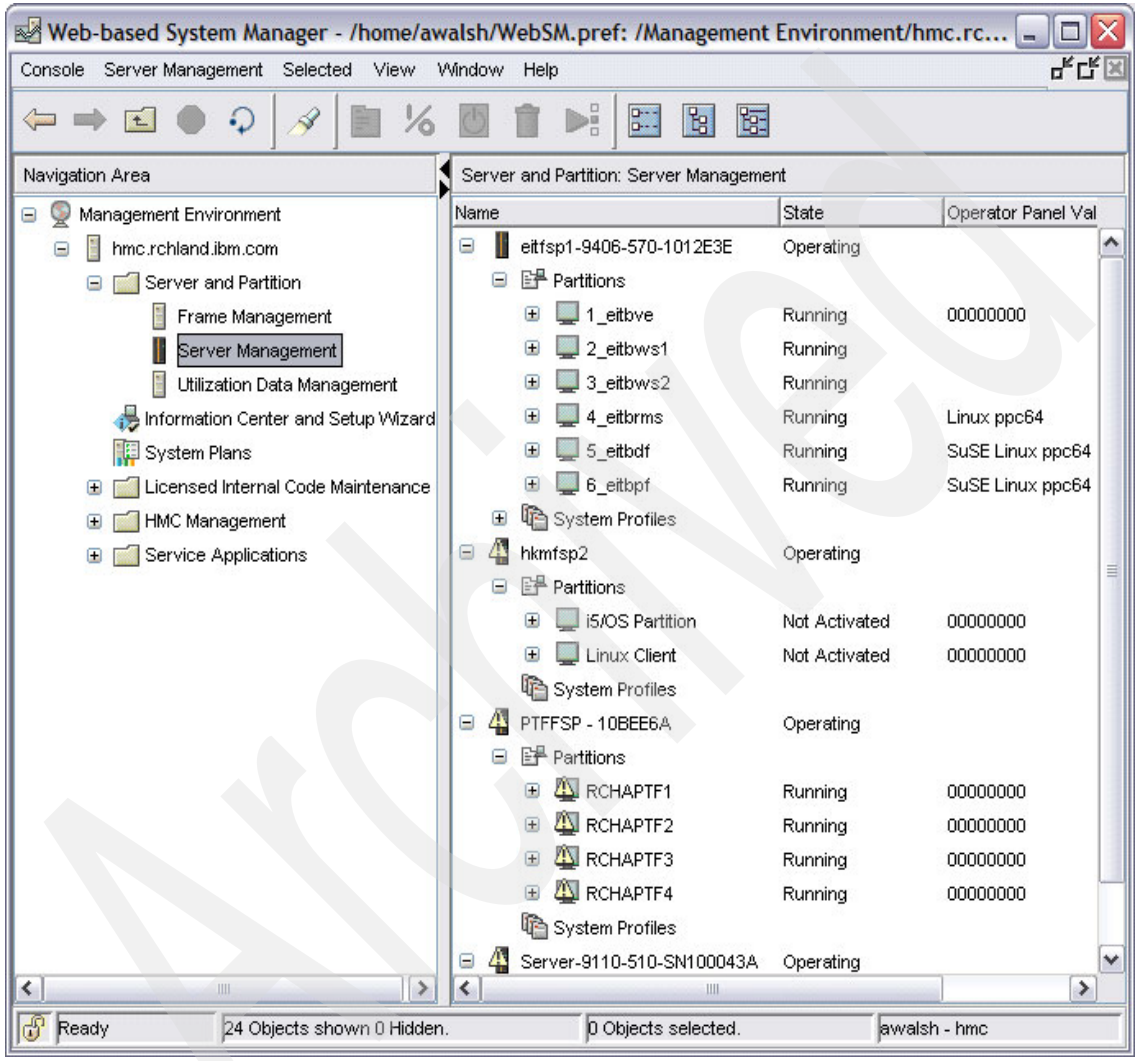

Figure 3-29 View completed deployment

4

System plan management using restricted shell (CLI)

In previous chapters we introduced system plans — what they are and how to manage them using the HMC graphical user interface. Most of the same functions included in the GUI can be invoked using the restricted shell command line interface (CLI). CLI can be accessed directly from the HMC or remotely by using a secure shell (SSH) client. The HMC does not provide a GUI equivalent to all of the command-level capabilities. In contrast, there is no equivalent in CLI to the view system plan feature available in the HMC GUI or in the System Planning Tool (SPT).

Previously, we identified several of the new commands added in the V5R2 HMC code level that are specifically in place to help manage system plans. These are **mksysplan**, **lssysplan**, **rmsysplan**, **deploysysplan**, and **cpsysplan**. In this chapter we step through an example in which we copy a sysplan file from removable media, write a sysplan file to removable media, list the sysplan files that exist on the HMC, remove a sysplan file, and then deploy a system plan.

4.1 Command line interface (CLI)

This chapter focuses on the CLI commands designed for managing tasks relating to system plans. Most of these are new commands that were introduced with HMC V5R2 code level. We also look at additional commands that were already available and are used as part of the preparation or within the deploy process. Finally, we deploy a sysplan using the CLI and step through the validation steps and show what commands are built to create the LPAR configuration objects.

To access the restricted shell command line interface on the HMC, simply position the cursor outside of the HMC window, right-click, and then select **Terminals** → **rshterm**. To access CLI on the HMC remotely you must connect to the HMC using a secure shell (SSH) client. In the examples to follow, we use a popular SSH client called *PuTTY*. This is a freeware product and generally satisfies most of the requirements for CLI on the HMC. PuTTY is available for Windows-based clients and Linux-based clients. The client software is available for free from a wide variety of Web sites. A simple Internet search locates a number of download selections. The SSH client used in the following examples is a Windows version of PuTTY and was downloaded from:

`http://www.putty.nl/download.html`

The CLI commands that are supported on the HMC are documented in the online manual. Simply type `man` before any command to display the command definition, parameter formats and examples. There is also a PDF version of the manual that is updated with each new HMC release. At the time of this writing, the latest version of this manual is for HMC V5R1 and is available for download at:

`http://techsupport.services.ibm.com/server/hmc/power5/tips/hmc_man_GA6.pdf`

Since most CLI commands (certainly most of the commands discussed in this chapter) require the managed server name as a parameter, changing the default name to a shorter name is convenient. The default name generated for each managed server, although very useful in its design, is very lengthy and therefore cumbersome when used in CLI. For example, the default server name is ‘Server-Machine type-model-SN’ and an i520 default name would be ‘Server-9406-520-SN10xxxxx’. Use the HMC Server Management panel to rename the server. Simply highlight the server, right-click, select **Properties** → **General**, and change the Name field.

4.2 Prepare to copy system plan file

The **`cpsysplan`** command is used to copy sysplan files either to or from the system plan directory on the HMC and another storage location. This is

equivalent to both the import and export tasks provided by the GUI interface. The system plans directory on the HMC is /opt/hsc/data/sysplan.

The external storage location can be an FTP server or one of the supported media types, which are CD/DVD, diskette, or USB flash memory device. The cpsysplan operation requires that the user know which external storage location will be used and to ensure that if it is a media device that it is mounted. The **lsmediadev** command is new with V5R2 of the HMC code and can be useful to identify available media devices:

```
lsmediadev

device=/dev/cdrom,type=1,description=CD/DVD
device=/dev/fd0,type=2,description=internal diskette
drivedevice=/dev/hda,type=6,description=internal hard disk drive
```

If you need to mount the necessary media device (for example, a USB Flash drive) use the **mount** command. The first USB port appears as SCSI device 1 (sda1):

```
mount /media/sda1

lsmediadev

device=/dev/cdrom,type=1,description=CD/DVD
device=/dev/fd0,type=2,description=internal diskette drive
device=/dev/sda1,type=3,description=USB flash memory device
device=/dev/hda,type=6,description=internal hard disk drive
```

4.3 Copy a system plan

LPAR Simplification is delivered in several phases. The initial phase is primarily targeted for deployment of a system plan that is created using the System Planning Tool (SPT) on a new POWER5 server that has not been previously configured. The SPT uses an export facility to copy the sysplan file to an external storage location to deliver the file to the HMC that will manage the new system. The first phase also introduces the ability to build a sysplan file directly on the HMC based on the existing hardware attached to any of the servers managed by the HMC. Ultimately, the SPT is enabled to import an HMC generated sysplan, and this is extremely useful for planning system upgrades. V5R2 HMC System Plans ennoblement provides the ability to copy in a sysplan created by SPT or a sysplan created by another HMC. It also supports the ability to copy out a sysplan from the HMC to removable media. The next section demonstrates the **cpsysplan** command format and features.

Like most for commands, there are a variety of parameters available. Some are required and many are optional. Let us take a look at the format and parameters of the copy system plan command.

4.3.1 Format and example of cpsysplan command

The **cpsysplan** is designed to further simplify movement of system plans to or from either removable media or an FTP server. The command parameters are shown here:

```
cpsysplan
```

Usage:

```
cpsysplan -r ftp | media
                -o import | export
                -f <file name>
                [-h <host name>]
                [-u <user ID>]
                [--passwd <password>]
                [-p <port>]
                [-d <directory>]
                [--check]
                [--help]
```

Copies system plan files to and from the system plan directory on the HMC and another storage location.

-r The location the system plan files are copied to or from:

- **ftp** Remote FTP site
- **media** Removable media

-o The direction of the copy:

- **Import** - from specified location to HMC
- **Export** - from HMC to specified location

-f <file name> The system plan file name to be copied

-h <host name> The host name or IP address of the remote FTP site

-u <user ID> The user ID for remote FTP site login

--passwd <password> The password to use for remote FTP site login

-p <port> The port number for the remote FTP site

-d <directory> The directory to use on the remote FTP site

--check Checks the system plan directory to see if the system plan file already exists when importing a file

The first example is to copy a sysplan from a USB flash memory card that has been inserted in the first USB port and is mounted at directory path /media/sda1. First we list the contents of the system plan directory that currently exist on the HMC using a standard **ls** command and the directory name:

```
ls /opt/hsc/data/sysplan
ATS111sysplan.sysplan  atsi5_111.sysplan  Myhome.sysplan  New system
plan.sysplan  sysplan2_15.sysplan
```

Now that we know what sysplans currently exist on the HMC hard drive, we copy a new sysplan (import) file from the flash card. The location of the plan (-r) specifies media. The direction of the copy (-o) is import. In our example the sysplan file name is BIGexample.sysplan, where -d is an option that allows you to specify the directory path where the source file is located. --check says to check the existing system plan directory for duplicate file names. If the command completes successfully, there is no confirmation message. However, if the copy fails there will be an error message:

```
cpsysplan -r media -o import -f BIGexample.sysplan -d /media/sda1 --check
```

Below is an example message of the command failing (in this case the sysplan file already exists in the System Plan directory):

```
File BIGexample.sysplan already exists in the system plan directory. Remove the
file or specify another file name.
```

The same command is used to copy a sysplan file (export) from the HMC system plans directory to removable media. For this example we copy a sysplan file called atsi5_520.sysplan to a FTP server. The location of the plan (-r) specifies FTP. The direction of the copy (-o) is export. The sysplan file name is atsi5_520.sysplan. The FTP sever host name (-h) is atsi5p1. The user name for accessing the FTP sever is (-u) is ibmtest. The password (--passwd) is java2day. (-d) /hmcsysplan is the directory path on the FTP server where the source file is located.

The first attempt is another example message of a failed command because file names are case sensitive.

```
cpsysplan -r ftp -o export -f ATSi5_111.sysplan -h atsi511 -u Mytest --passwd
pwd2day -d /hmcsysplan
The file atsSi5_520.sysplan was not found or was empty.  Please specify a valid
file name and try again.
```

This is a successful example. Again, notice that three is no confirmation message generated:

```
cpsysplan -r ftp -o export -f atsi5_111.sysplan -h atsi511 -u Mytest --passwd
pwd2day -d /hmcsysplan
```

4.3.2 List system plans

Another new command to help identify system plans that already exist on the HMC is the list system plans command **lssysplan**. This command generates the same results as the **ls** command. However, it is not necessary to know the system plans directory path. At this time of this writing, the online manual does not have any details on this command.

```
man sysplan

No manual entry for sysplan
```

An alternative method to view details of most command is to use the --help option:

```
lssysplan --help
```

Usage:

```
ssysplan [-F [<attribute names>]]
                 [--header]
                 [--help]
```

This lists system plan files in the system plan directory on the HMC.

-F [<attribute names>] Delimiter-separated list of the names of the attributes to be listed for each system plan file. If no attribute names are specified, then all attributes are listed.

--header Prints a header of attribute names when -F is also specified.

--help Prints this help.

Execute the command:

```
lssysplan

name=BIGexample.sysplan,description=feb_20,source=HMC V5.2.0,version=HMC
1.1,"date=Feb 21, 2006 10:40:06 PM"
name=sysplan2_15.sysplan,description=feb15,source=HMC V5.2.0,version=HMC
1.1,"date=Feb 15, 2006 3:11:19 PM"
name=atsi5_111.sysplan,description=Feb16_2006,source=HMC V5.2.0,version=HMC
1.1,"date=Feb 16, 2006 2:36:45 PM"
name=New system plan.sysplan,description=,source=,version=,date=
name=Myhome.sysplan,description=CSDML generated,source=IBM System Planning Tool
null,version=SPT 1.1,"date=Feb 16, 2006 9:04:06 AM"
name=ATS111sysplan.sysplan,description=Feb20_2006,source=HMC V5.2.0,version=HMC
1.1,"date=Feb 20, 2006 10:39:49 PM"
```

Notice that the sysplan 'BIGexample.sysplan' that we copied from the USB flash memory previously now shows up on the MHC in the system plans directory.

4.3.3 Remove system plan

Another new command, **rmsysplan**, is available to make the process of deleting a sysplan file very simple. As with all of the xxxsysplan commands, we do not need to specify the directory path to the system plans directory. Simply specify the sysplan file name.

To view the online manual for this command, simply precede the command with the man qualifier:

```
man rmsysplan

NAME
       rmsysplan - remove system plan

SYNOPSIS
       rmsysplan -f file-name [--help]

DESCRIPTION
      rmsysplan removes a system plan file from the system plan file directory
on the Hardware Management Console (HMC).

OPTIONS
       -f     Specifies the name of the file that will be removed.  Only files
in the system plan file directory on the HMC will be removed.

      --help Display the help text for this command and exit.

EXAMPLES
      Remove the system plan file sysplan.sysplan:

       rmsysplan -f sysplan.sysplan

ENVIRONMENT
       None

BUGS
       None

AUTHOR
       IBM Austin

SEE ALSO
       deploysysplan, cpsysplan, lssysplan, mksysplan
```

To delete the sysplan the we copied (imported) from flash memory, the command format is simply:

```
rmsysplan -f BIGexample.sysplan
```

4.3.4 Create a system plan file

The **mksysplan** command is a very useful tool when combined with the viewing/printing capability of the System Plans View function. At the time of writing, the command is able to create a sysplan file that shows the system information, LPAR definitions (including memory and processor allocation), as well as IO resource assignments. To execute this command the HMC must be at V5R2, but the target POWER5 or POWER5+ server can be on a back level of firmware.

The sysplan file can then be 'VIEWED' on the HMC that it was created on, or it can be exported and viewed on another HMC that has imported the sysplan file. If you are using WebSM, you have the ability to view, export, and print System Plans. Select **VIEW** and then select the system plan you require from the Manage System Plans panel. The sysplan can be printed from this Web browser session on your workstation.

One current limitation is that the sysplan file can only identify the I/O resources down to the IOA level.

Another very useful function that this tool provides today is simply the ability to capture hardware configuration data and report this in a visual and printable format. So even on a brand new server that has not yet been loaded with any OS or any LPAR definitions, **mksysplan** could be run to assist in verifying that the hardware that was ordered is present. Also, there may be times when the successful deployment of a SPT-generated sysplan will need to know details such as the serial number assigned to IO towers or drawers. SPT has the ability to specify serial number assignments, and this can help ensure that the correct towers will be allocated for the I/O requirements of the planed partitions. The topic of tower ambiguity (a situation in which there could be multiple towers with identical or nearly identical card layouts) is discussed further in Chapter 5, "Deployment validation" on page 87.

To review the online manual use the man call or the --help parameter:

```
mksysplan --help
```

Usage:

```
mksysplan -f <file name>
                  -m <managed system>
                  [--check]
                  [-d "<description>"]
```

```
[--help]
```

This creates a system plan file that represents known information about a managed system's hardware, partitions, profiles, and partition provisioning information.

-f <file name>	The system plan output file name. This output file will be created in the system plan directory on the HMC. Existing files will be overwritten unless --check is specified.
-m <managed system>	The managed system's name.
--check	Checks the system plan directory to see whether the system plan output file already exists.
-d "<description>"	Specifies a description that will be added to the system plan file.
--help	Prints this help.

For example, we create a system plan on a model 520 that currently has five partitions. The server name is ATSi5_520. It has two 0595 IO towers attached. The (-m) identifies the server parameter, (-f) will be the name of the sysplan file created in the system plans directory (note that the file name must have a file extension of .sysplan), and (-d) is a file description.

```
mksysplan -m ATSi5_111 -f ats111Feb24.sysplan -d i520 with 2 0595 and 5 LPARs
```

The above example will be successful as long as the system name is found under the list of managed servers for the HMC that issues the command. Keep in mind that the command parser is case sensitive. If any of the parameters have a space, the parameter must be enclosed in double quotes (“ “). For example, if the system name was ATS 111, the command parameter would be -m “ATS 111”.

4.3.5 Deploy a system plan

The command **deploysysplan** performs the validation of the sysplan on the actual target server of the sysplan. This command generates error messages for mismatches in the configuration. Just as with the HMC GUI, the deploy process stops at validation. The GUI shows the details of the sysplan that successfully completed validation on the Deploy System Plan Wizard - Select What to Deploy window. To begin the process of creating the LPAR configuration objects the administrator would simply select **Next**.

The **deploysysplan** command can be instructed to perform validation only. This allows the administrator to review all messages generated, either reporting whatever error may have prevented the validation from completing or confirming that the validation was successful. If the validation is successful the

deploysysplan displays each of the individual commands that it generated to create the sysplan configuration. This allows the administrator to manually create the configuration, step-by-step, by simply cutting and pasting the individual commands. It also allows the administrator to selectively create just a portion of the sysplan. Caution: Doing a partial deploy in this manner requires that the administrator understand any partition dependencies, such and Client-SCSI and Server-SCSI dependencies.

Another option is that the **deploysysplan** can skip the validation step and proceed directly with creating the configuration objects. This method can be used following a successful execution of this command for validation only. It would *not* be wise to attempt to execute this command without a successful validation first.

A third option is to have the **deploysysplan** perform validation and automatically complete the process by executing the commands to create the LPAR configuration objects. This would of course require a the first phase, a successful validation of the sysplan, before the configuration commands are executed.

The online manual is available to view the command format or the format can be reviewed using the --help parameter, as follows:

```
deploysysplan --help
```

Usage:

```
deploysysplan -f <file name>
                    -o dv | v | d
                    [-m <managed system>]
                    [--force]
                    [--plan <plan number>]
                    [-v]
                    [--help]
```

This deploys a managed system's partition plan from a system plan file.

-f <file name> The system plan file name that is to be deployed.

-o Specifies the deployment option to use when deploying the system plan on the managed system.

dv Validates the system plan for the managed system. If it is valid, then deploy the system plan.

v Validates the system plan for the managed system but does not deploy it.

d Skips validation of the system plan and just deploys the specified partitions and profiles.

-m <managed system> The managed system's name.

--force Deploys the system plan even if there are active partitions.

--plan <plan number> Specifies which system plan in the system plan file to deploy if the file contains more than one system plan.

-v When the -o v option is specified, lists the system plan step identifiers that are part of the system plan but not yet implemented on the managed system. When the -o dv or d option is specified, lists the deployment steps that are performed.

--help Prints this help.

4.4 Deploy system plan example

The following example is a successful execution of the **deploysysplan**. We specify the deployment option (-o v)), which means validation only. This allows us to review the commands that the **deploysysplan** creates. Now we can review these commands, and in some cases the administrator may choose to manually execute these commands. After we run just the validation and review the commands, we issue the **deploysysplan** command a second time and specify a deploy option (-o d), which is deploy only, without validation.

The following command identifies a sysplan file called hkmfsp2.sysplan. The target server name is hkmfsp2 and the deploy option is -o v, which is validation only. The -v means to list the system plan step identifiers and -c mean to list the individual commands.

```
deploysysplan -f hkmfsp2-SPT.sysplan -m hkmfsp2 -o v -v -c

System plan system unit IBM iSeries 520 System Unit
IBM iSeries 520 System Unit, type 520, serial number * was matched and
validated with system unit serial number DNZ070Z found on the managed system.
System plan expansion unit 5095-1, type 5095, serial number * was matched and
validated with expansion unit serial number 103624C found on the managed
system.
No deployment validation errors found when comparing the hardware as specified
by the system plan hkmfsp2-SPT.sysplan, plan name hkmfsp2, against the managed
system hkmfsp2 reported hardware.


The following commands will be used when deploying plan file
hkmfsp2-SPT.sysplan, plan name hkmfsp2, on managed system hkmfsp2:

Plan item Partition LPAR1, will be deployed with command
```

```
mksyscfg -r lpar -m hkmfsp2 -i
"lpar_id=1,name=LPAR1,profile_name=LPAR1,lpar_env=os400,min_mem=512,desired_mem
=1024,max_mem=1024,proc_mode=ded,sharing_mode=share_idle_procs,min_procs=1,desi
red_procs=1,max_procs=1,max_virtual_slots=10,io_slots=\"21020002//1,2103000B//1
,21020003//1,2102000B//1,21040003//1,21050003//1,21050002//1,21040002//1,210300
03//1,21030002//1,2104000B//1,2101000B//1\",auto_start=0,all_resources=0,conn_m
onitoring=0,load_source_slot=21020003,console_slot=hmc,desired_5250_cpw_percent
=0.0,min_5250_cpw_percent=0.0,max_5250_cpw_percent=0.0".


Plan item Virtual SCSI Virtual SCSI, will be deployed with command
chsyscfg -r prof -m hkmfsp2 -i
"lpar_id=1,name=LPAR1,virtual_scsi_adapters+=\"3/server/2//3/1,4/server/2//4/1\
"".

Plan item Virtual Ethernet Virtual Ethernet, will be deployed with command
chsyscfg -r prof -m hkmfsp2 -i
"lpar_id=1,name=LPAR1,virtual_eth_adapters+=\"\"2/0/1//0/1\"\"\"".


Plan item Partition LPAR2, will be deployed with command
mksyscfg -r lpar -m hkmfsp2 -i
"lpar_id=2,name=LPAR2,profile_name=LPAR2,lpar_env=aixlinux,min_mem=256,desired_
mem=256,max_mem=512,proc_mode=shared,sharing_mode=uncap,uncap_weight=128,min_pr
ocs=1,desired_procs=1,max_procs=1,min_proc_units=1.0,desired_proc_units=1.0,max
_proc_units=1.0,max_virtual_slots=12,auto_start=0,all_resources=0,conn_monitori
ng=0,boot_mode=norm,power_ctrl_lpar_ids=\"1\"".



Plan item Virtual SCSI Virtual SCSI, will be deployed with command
chsyscfg -r prof -m hkmfsp2 -i
"lpar_id=2,name=LPAR2,virtual_scsi_adapters+=\"3/client/1//3/1,4/client/1//4/1\
"".

Plan item Virtual Ethernet Virtual Ethernet, will be deployed with command
chsyscfg -r prof -m hkmfsp2 -i
"lpar_id=2,name=LPAR2,virtual_eth_adapters+=\"\"2/0/1//0/1\"\"\"".

Plan item Virtual Serial Virtual Serial, will be deployed with command
chsyscfg -r prof -m hkmfsp2 -i
"lpar_id=2,name=LPAR2,virtual_serial_adapters+=\"5/client/0/1//0/1\"".
```

The validation was successful and each command that would be required to create the configuration is listed above. At this point, you could cut and paste these commands in the same sequence as listed and create the LPAR configuration.

Figure 4-1 is the HMC view of the Server management panel after deploying the system plan we just validated.

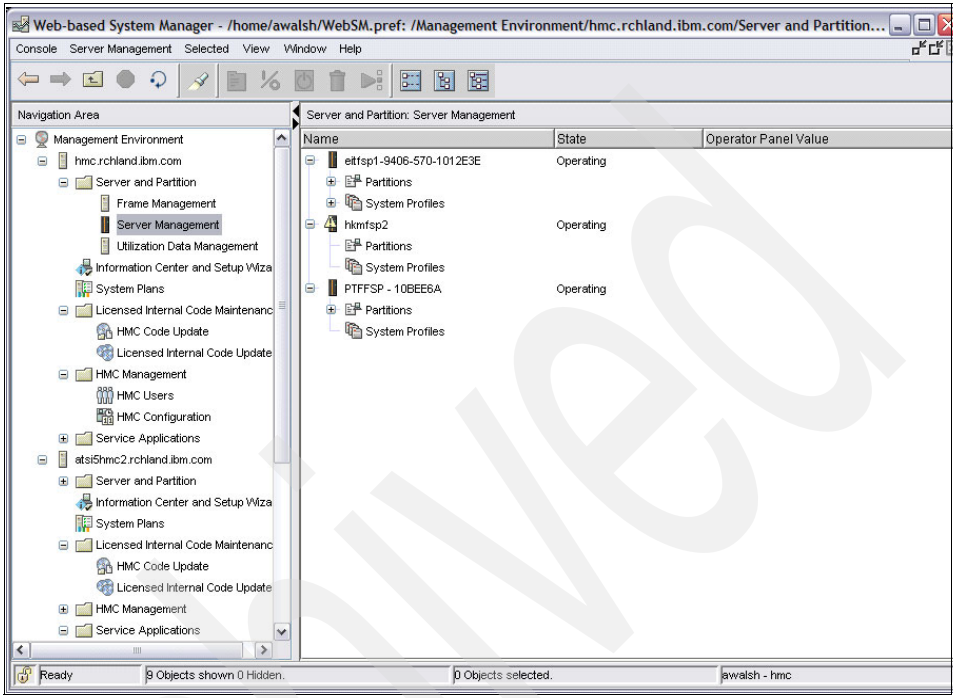

Figure 4-1 Server Management partition status before sysplan deployment

To finish this example, since we successfully ran the validation and do not have any need to do a partial deploy, we issue the **deploysysplan** again and specify the -o d option to do a deploy only, without validation:

```
deploysysplan -f hkmfsp2-SPT.sysplan -m hkmfsp2 -o d
```

Figure 4-2 is the HMC view of the Server management panel after deploying the system plan we just completed.

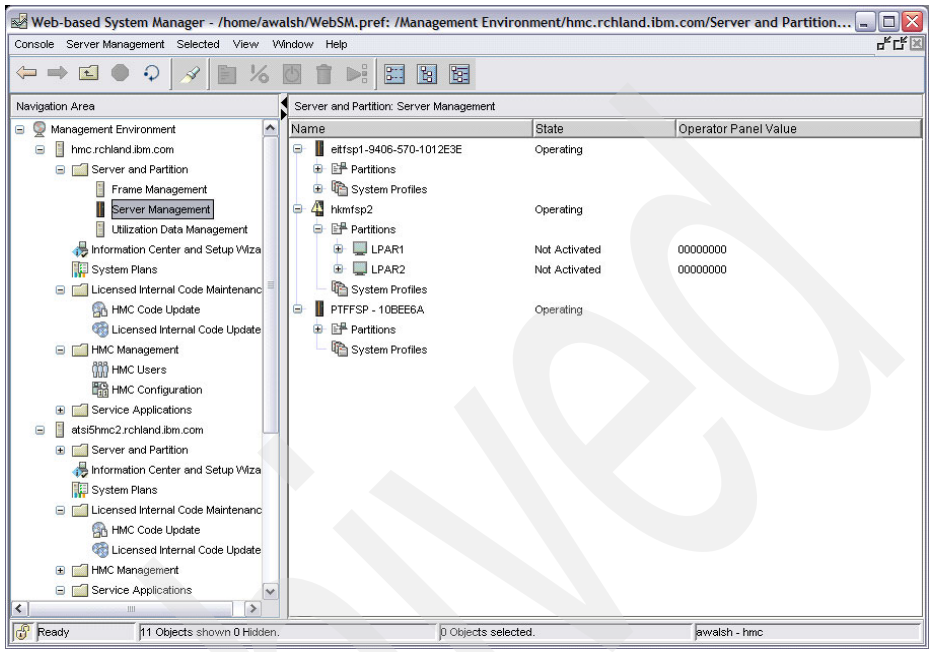

Figure 4-2 Server Management partition status after sysplan deployment

5

Deployment validation

When an attempt is made to deploy a sysplan on a system, the plan must be validated. The validation consists of two steps. Hardware validation is done first. If a plan passes hardware validation, it then goes through partition validation. This chapter discusses the validation process in more detail. 5.1, “Hardware validation” on page 88, and 5.2, “Partition validation” on page 93, describe the two phases of validation in detail, while 5.3, “Validation process” on page 94, gives a simplified overview of the entire process.

When validation errors are encountered, deployment fails and detailed messages are generated. This allows modifications to be made to allow the next validation to succeed. Once the hardware and partition validation have succeeded, the plan can be deployed.

5.1 Hardware validation

Hardware validation is the first validation a sysplan must go through in the deploy process. This validation is not doing rules checking or placement checking, as this is done via SPT. It only verifies that the hardware on the physical system satisfies the hardware in the plan. A failure at any step of the validation process generates a descriptive message to help debug why the deploy failed. At that time, the necessary changes can be made to the plan or the system to allow the deploy to succeed. It is important to note that extra physical hardware does not cause the deploy to fail as long as the existing hardware *satisfies* the plan.

Note: The terms *satisfy* and *satisfies* in the context of validation indicate that there exists at least enough physical hardware to fulfill the plan requirements. It does not imply that all physical hardware exists in the plan. Therefore a physical system can contain extra hardware, above and beyond what the plan calls for, and still *satisfy* the plan.

There are two cases in which the hardware validation must be taken into consideration, depending on the information in the sysplan from the SPT:

- Validation with accurate serial numbers
- Validation without accurate serial numbers

5.1.1 Validation with accurate serial numbers

Serial numbers for system and expansion units can be put into the plan via the SPT. A sysplan with valid serial numbers for all system and expansion units is the ideal case. It eliminates the possibility for unit ambiguity by ensuring that the unit is uniquely identified in the plan, even if multiple units are otherwise identical, down to the IOA level. The only drawback is that the serial numbers are not known until the system is delivered, so the plan must be updated through the SPT after delivery, prior to deploy.

The hardware validation occurs any time a deploy operation is done. Even if a partial deploy has been done, all hardware validation is redone when another deploy is performed. If the hardware has been changed, the hardware validation can fail, even if the change only affected previously deployed LPARs. If this is the case, the sysplan should be altered to reflect the hardware through the SPT.

The steps in the validation process where accurate serial numbers are given are as follows:

1. System validation

   This step validates four things: type, model, number of processors, and memory pool size. If any of these three fail to validate, the validation process will end, and the deploy will fail.

   – Type and model

     The type and model of the physical system are checked simultaneously. The system being deployed to must satisfy the type and model in the sysplan. You cannot deploy a plan created for an i5 520 onto an i5 550, for instance. If this step fails, the validation immediately ends with a message indicating that the physical system is of the wrong type and model.

   – Number of processors

     The validation process ensures that there are at least the number of physical processors installed on the system as the plan indicates there should be. It also verifies that the number of activated processors on the physical system is at least the number specified by the plan.

   – Memory pool

     The validation process ensures that the memory installed on the physical system is at least the amount specified by the plan.

2. Unit validation

   Unit validation verifies that each unit with the given serial number in the plan exists in the physical hardware. The validation process also verifies that each slot in those serial number identified units contains the planned hardware. It is important to note that at this time the hardware validation only validates to the IOA level. The HMC cannot validate hardware that it cannot see, such as disk drives.

3. Slot dynamic reconfiguration connector (DRC) index assignment

   Each slot in a physical system unit has a slot-DRC-index (DRC) assigned to it. This number is required by the deploy process to do the actual creation of the LPARs. This identifier is used by the HMC and the system to uniquely identify every slot on the system, and is not assigned to a slot until the unit is powered up. Therefore, there is no way to know these values at plan time.

   The validation process must obtain the DRCs of the hardware system and expansion units for the resources it is going to assign to the LPARs as part of validation of these units. The DRCs are used in the command that creates the LPARs as part of their deployment. This information is added to the in memory model of the system plan during unit validation. The actual system plan file being deployed from the HMC is not altered.

5.1.2 Validation without accurate serial numbers

Obviously it must be possible to have a system plan without serial numbers for system and expansion units defined. This is necessary because we cannot know the serial number of all units before a system is manufactured. It is therefore possible to attempt to deploy a system plan where all system and expansion unit serial numbers are not accurately defined in the system plan being deployed. This fact is also what allows the same system plan to be deployed against multiple target systems, as long as the plan is satisfied.

This flexibility introduces the potential for a major problem with ambiguous units. This is discussed in more detail in 5.1.3, “System or expansion unit ambiguity” on page 92. The hardware validation process must be able to validate the hardware regardless of the possibility of ambiguity.

There are four possibilities for the serial number definition in any given sysplan. They are as follows:

- Accurate

  The serial number and underlying hardware in the plan is satisfied by the physical system.

- Missing

  The serial number was not entered into the sysplan with the SPT.

- Inaccurate

  The serial number is given in the plan but one of the following is true:

  - The serial number in the plan does not match any serial number on the physical system.
  - The serial number in the plan exists on the system but the hardware defined in the plan for that serial number is not satisfied by the physical hardware in the unit with that serial number on the actual system.
  - The serial number in the plan exists for a unit on the system but the serial number for the unit type in the plan is not the same unit type on that system.

- Duplicate

  The same serial number is used more than once to describe different units, in either the plan or on the physical system.

These situations are all dealt with during unit validation. The following steps show the validation process when the serial numbers for units in the system plan do not match the serial numbers for the actual system hardware or when the system plan has no serial numbers for units at all. In fact, this is the general case

for hardware validation, where 5.1.1, “Validation with accurate serial numbers” on page 88, is actually a special case.

1. System validation

   Refer to step 1 on page 89.

2. Unit validation

   When all serial numbers are not accurately reflected in the plan prior to deploy, the unit validation step must still match units in the plan to the units that physically exist on the system. The unit validation step attempts to do this in the following way:

   a. Compare any serial numbers that *are* included in the plan to those that are on the physical system.

      - If there are duplicate serial numbers given in the plan, or on the physical hardware, deploy fails. If the same serial number is given twice in the plan, it is likely an oversight, and the plan can be corrected with the SPT. The physical hardware should never have duplicate serial numbers. If it does, Hardware Support should be contacted to resolve the issue.
      - If the hardware associated with a serial number in the plan is satisfied by a physical unit with the same serial number as in the plan it is considered valid.

        It is possible that the unit in the system plan has specific requirements and these are not met by the unit on the physical system that has that same matching serial number. An example of this would be PCI cards that are part of the system plan that appear in specific slots for that planned unit. It the planned requirements for that unit are not met that plan will fail validation. It is important to note that the physical unit is not eligible to validate any other unit in the plan. If the serial number of a unit in the plan matches the serial number of a unit in the physical hardware, those units are compared only to each other.
      - In the case of inaccurate or missing serial numbers the unit is validated in step b for system units or c for expansion units.

   b. Validate system unit

      This step ensures that the system units are valid according to the plan beyond the validation described in step 1. It verifies that the specific requirements of the system units that appear in the plan are met by system units that are part of the physical system hardware.

   c. Validate expansion units

      The expansion units that have not already been validated are now analyzed. They are validated in groups by feature or model/type. It is

pertinent to again mention that at this time there is no validation done below the IOA level. For example, the actual system disk drives or how any drives are related to the IOAs on the actual system hardware is not known to the HMC. That means that this information cannot be used to determine whether this part of the hardware plan is met by the drives that might be part of the system where the plan is being deployed.

When a unit in the physical hardware is found that satisfies only one unit in the plan, it is considered validated. If there are only ambiguous units that satisfy the part of the plan being validated, the validation process simply chooses one without regard to what the user expects to be chosen. This step is where the presence of ambiguity can be problematic. This is discussed in detail, with an explanation of how to eliminate it as a problem, in 5.1.3, “System or expansion unit ambiguity” on page 92.

3. Slot-DRC-index assignment

   Refer to step 3 on page 89.

5.1.3 System or expansion unit ambiguity

There have been many references to unit ambiguity in this book. This section discusses it in more detail and explains how to eliminate it.

System or expansion unit ambiguity refers to having two or more units in which the hardware is identical to the IOA level. This means that they are the same feature (0595, for example) with the same IOPs and IOAs in the same slot locations. This is ambiguous in that these two 0595 units could have entirely different disk drive configurations in the system plan, but that difference on the actual system hardware cannot be seen by the HMC. At this time there is no way for the validation process to know that the drive configuration is different. Therefore there is no way for the validation process to ensure that the planned units are always correctly matched to the expected unit that is part of the physical system's hardware.

It is possible that the validation process will assign slot-DRC-indices from the “wrong” unit to the memory model of the plan. The deploy process would then create the partition, and the profile would reference the unexpected hardware. While the profile would be valid from the standpoint that it would activate and could be used, it would not have the planned hardware. In this case the disk drive configuration would be wrong.

Unit ambiguity would also be a problem if it is desired that a specific unit in a specific physical location on the floor was expected to be assigned to a given LPAR. It is common for a system to be physically arranged so that the proximity of the expansion units to the system unit, for example, should indicate what

LPAR the expansion unit belongs to. If there are ambiguous units, this result cannot be guaranteed without taking action to eliminate ambiguity.

When LPAR2 is deployed, there is no guarantee that the validation process will not choose the unit already in use by LPAR1, and create LPAR2 to point to those same resources. Consider the case where the plan for LPAR1 and LPAR2 define ambiguous units that satisfy both partitions. LPAR1 has been deployed, but the deploy of LPAR2 was postponed.

When LPAR2 is deployed, there is no guarantee that the validation process will not choose the unit already in use by LPAR1 and create LPAR2 to point to those resources. If LPAR1 happened to be inactive when LPAR2 was activated to install and an initialize was done, LPAR1 data would be lost. Even if the error was caught before an install, the situation must be remedied. The configuration for LPAR2 would need to be manually fixed or redeployed once ambiguity was resolved.

Important: If partial deploys are done to a system with ambiguous units, the plan must be updated to include serial numbers for the ambiguous units to avoid the above described scenario.

5.1.4 Eliminating unit ambiguity

To eliminate the possibility of unit ambiguity simply modify the system plan through SPT and specify the serial numbers found on the actual managed system's units before importing it on the HMC to deploy. Technically, only enough serial numbers need to be entered so that each unit can be uniquely identified. In other words, n–1 ambiguous units need serial numbers to eliminate ambiguity between n units.

5.2 Partition validation

Where hardware validation is concerned with ensuring that what is in the plan is on the physical system, partition validation is concerned with ensuring that what is configured on the physical system exists in the plan. The partition validation phase checks the partition profiles that may exist on the system, possibly due to a previous partial deploy, to ensure that they exist on the plan. This validation step verifies that any partition that is configured on the physical system matches the following parameters on a partition in the plan:

- Partition name

  Any existing partition on the actual system must match a partition with the same name in the plan.

- Partition ID

  Any existing partitions must match the partition ID in the plan that it is supposed to satisfy.

- Default profile

  Partition validation is done on the default profile. If there are multiple profiles in existence for a partition that exists at deploy time, only the one marked as the default is validated against the sysplan.

- Processing resources

  All processor settings, except the weight of uncapped processors, are validated to ensure that the plan is satisfied by the existing partition.

- Memory resources

  All planned memory settings are checked to verify that they are satisfied by the existing partition.

- Virtual I/O adapters

  The validation ensures that all virtual I/O adapters in the plan are satisfied by the existing partition, including slot number assignments.

- Physical hardware

  The validation ensures that the existing partition has the hardware to satisfy the plan.

If any of the above fail to satisfy the plan, the deploy will fail. An error message is generated detailing what caused the validation to fail. With this information, either the existing partition profile or the plan via the SPT can be modified to allow the next validation to succeed.

Note: Restarting or completely deploying a system plan that was partially deployed should never result in a validation error. Any existing partitions on the managed system should have been created by the same system plan and therefore should always be valid when validated against that same system plan. Also note that if there is a hardware validation problem that results in validation failure then the partition validation steps are not performed at all.

5.3 Validation process

While 5.1, “Hardware validation” on page 88, and 5.2, “Partition validation” on page 93, accurately describe the validation process, it can be difficult to understand. This section presents the information graphically in an attempt to simplify the process. For more in-depth information about a given step, reference the appropriate section in the aforementioned sections.

The validation process starts out with a .sysplan file on the HMC, as represented by Figure 5-1. The sysplan has been discussed at length previously in this book. The arrows in the diagram represent the relationship that the planned partitions have to the planned hardware. Prior to a deploy operation, the .sysplan is just a flat file on the HMC.

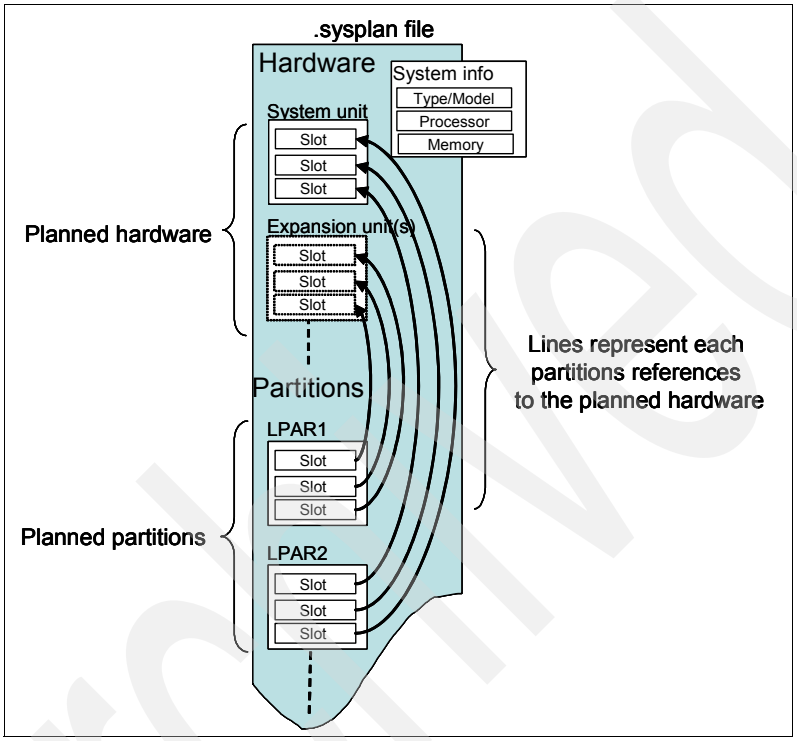

Figure 5-1 The .sysplan file

When a deploy operation is initiated, the HMC accesses the .sysplan file to create a model of the plan in memory. This *in memory model* of the sysplan is used for validation. This is represented in Figure 5-2. As the diagram shows, the partition relationships to the planned hardware are maintained. At this point, the hardware described by the in-memory model has not been associated with actual hardware.

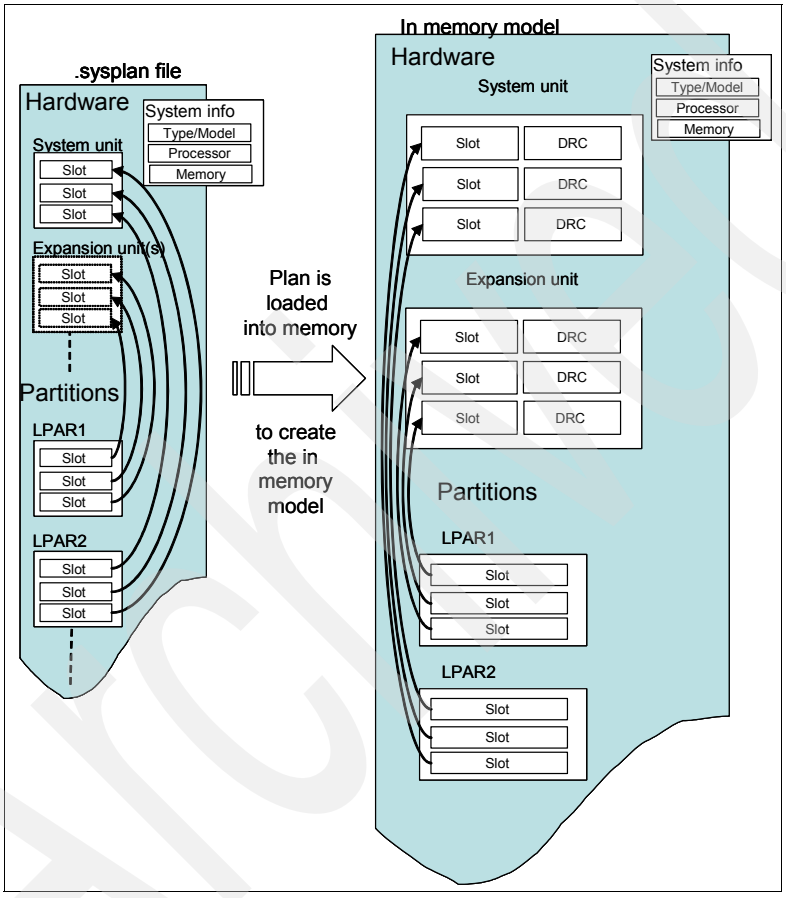

Figure 5-2 System plan loaded into memory to be used for validation

The in-memory is then compared to the actual physical system. Hardware validation is attempted. First the HMC retrieves information from the actual system to do system validation. As indicated by Figure 5-3, the HMC accesses information from the managed system regarding its type and model as well as processor and memory configurations. System validation is described in further detail in step 1 on page 89.

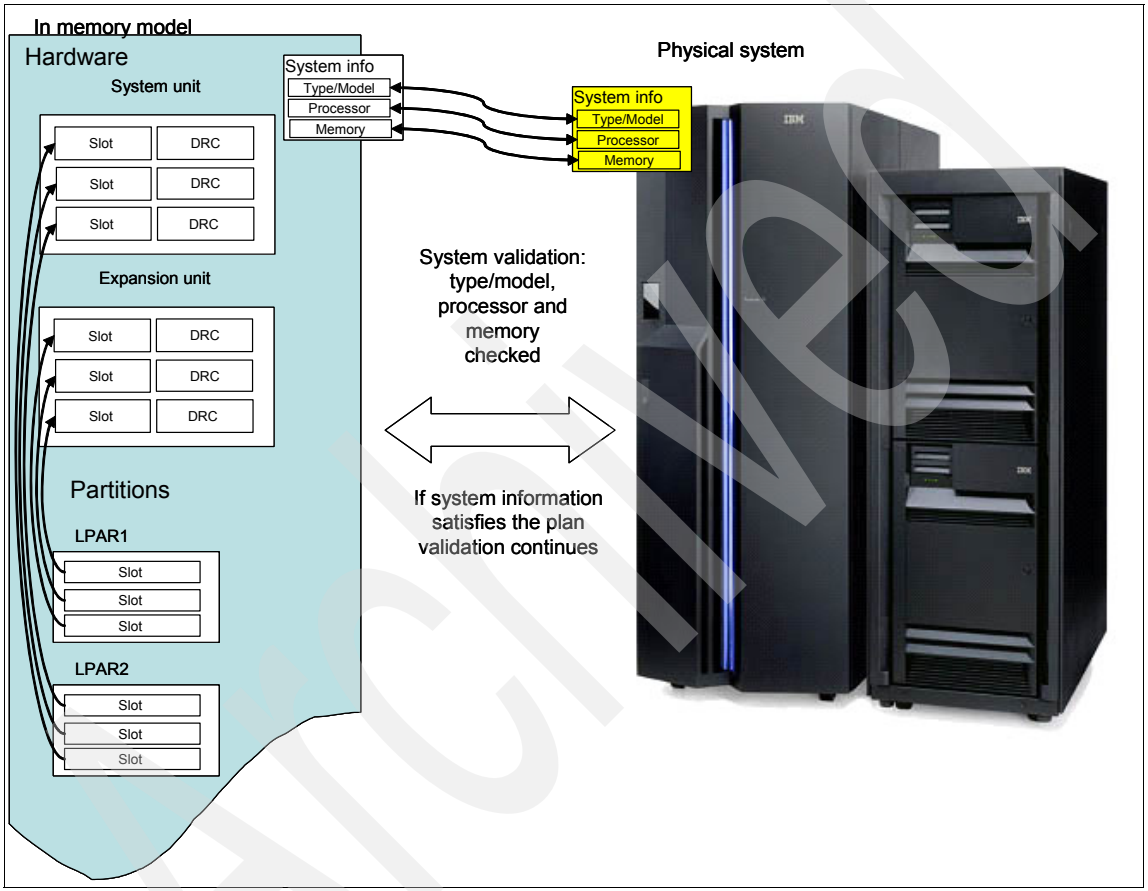

Figure 5-3 System validation step

Once the system validation is completed, the system moves onto unit validation. The HMC gathers data about the physical units and ensures that they satisfy the plan, as indicated by Figure 5-4. A detained description of what is done for unit validation can be found in step 2 on page 91.

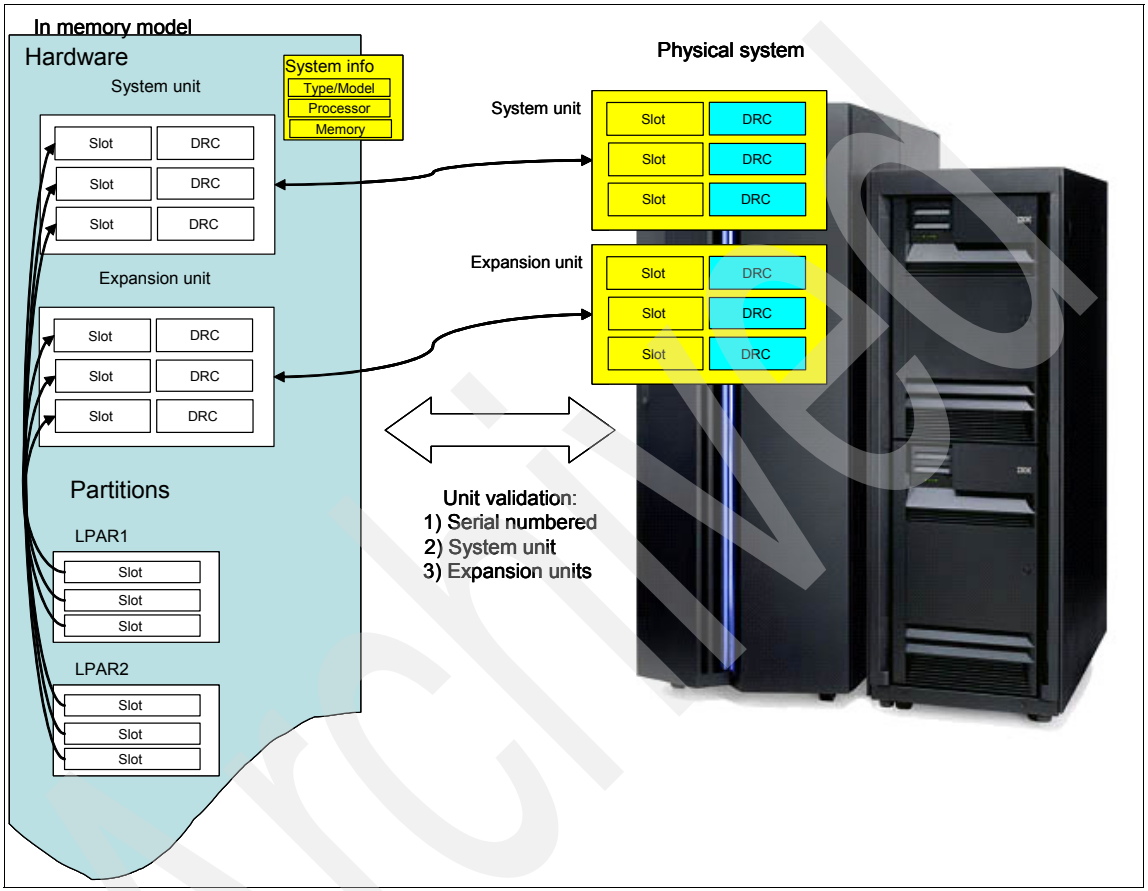

Figure 5-4 Unit validation step

During the unit validation process, slot-DRC-indices from the physical system are assigned to the in-memory model. See Figure 5-4 on page 98. A discussion on slot-DRC-indices and their importance to the deploy process can be found in step 3 on page 89.

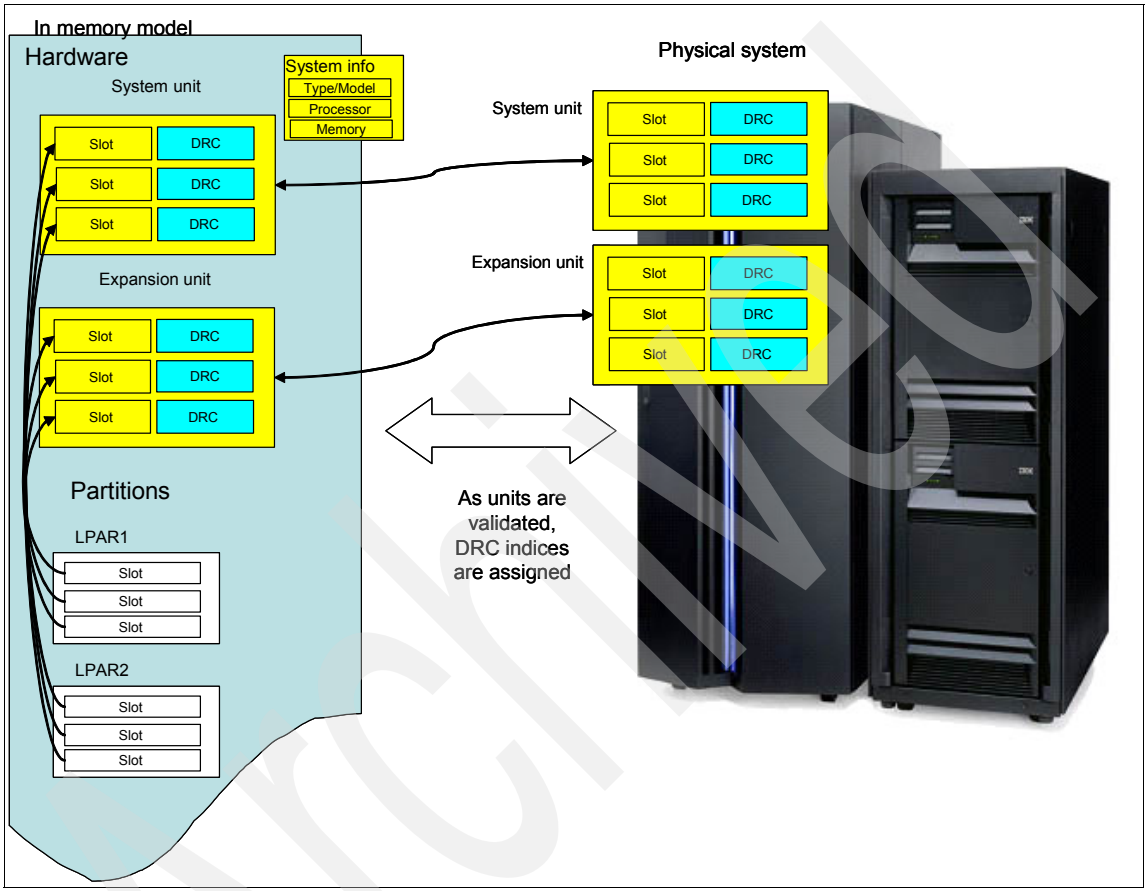

Figure 5-5 Slot-DRC-indices assigned to the in memory model

This is a critical step, as it is how a system plan that represents partitions that are mapped to the planned resource are now mapped to those actual resources as they were found on the managed system. When a planned partition is created, the assigned DRC index is how that partition *owns* or allocates the resource identified for that index and assigns it to the new partition.

After completion of unit validation, the HMC gathers data about any partitions that already exist on the managed system at the time the deploy was issued. Partition validation is attempted. If any partitions exist, they are validated as described in 5.2, “Partition validation” on page 93. See Figure 5-6.

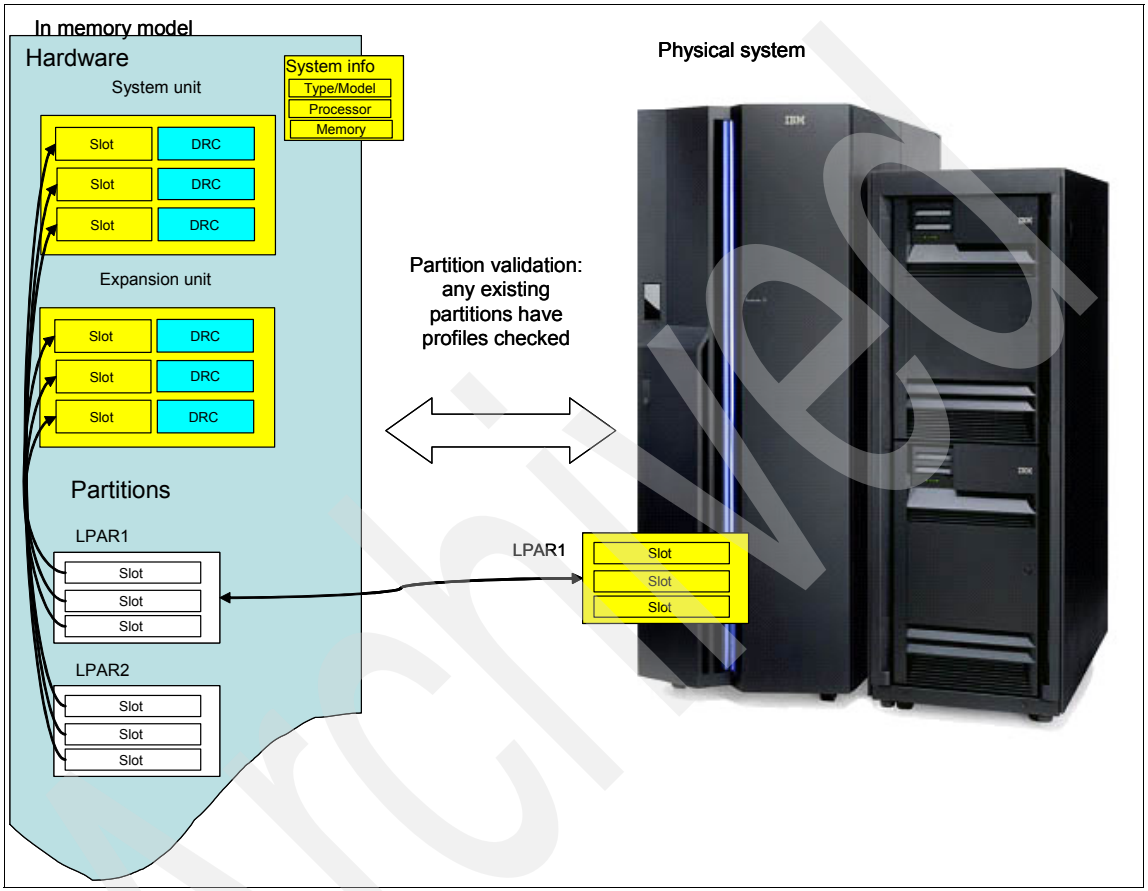

Figure 5-6 Existing partitions are validated

Partition validation is the final step in the validation process. Once it is complete, the deploy process can continue. See Figure 5-7.

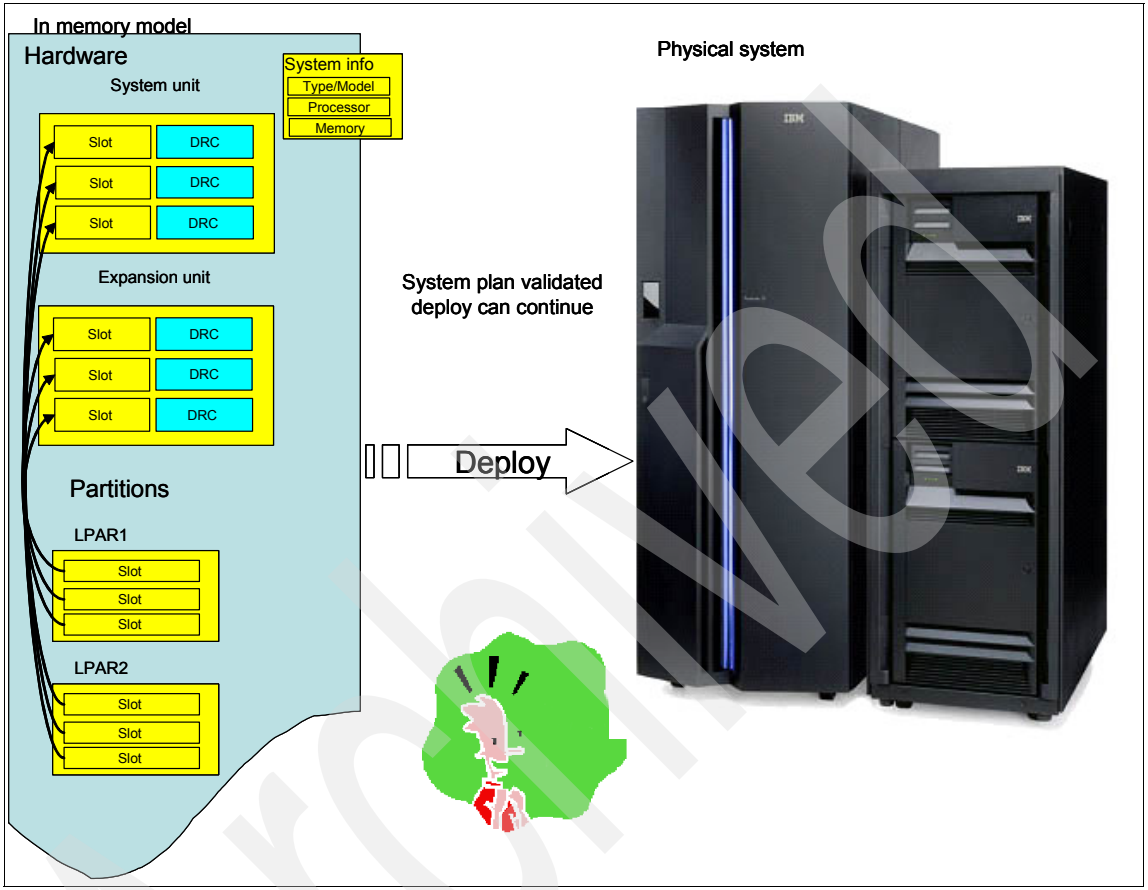

Figure 5-7 Validation complete, deploy can continue

5.4 Validation message examples

The following are examples of various validation messages one could expect to see when attempting to deploy a plan. This is by no means an exhaustive list. but attempts to give a general idea of the types of error messages to expect. It is possible to eliminate any validation errors by one of two methods: alter the plan via the SPT or modify the existing hardware or partitions on the actual system.

5.4.1 Successful validation without serial numbers

Figure 5-8 is an example of the messages you would expect to see for a successful validation. In this example the plan did not contain serial numbers. No partitions existed on the system prior to the deploy.

1. This message is posted when the system validation stage is completed. The validation step completed at this point is described in step 1 on page 89.

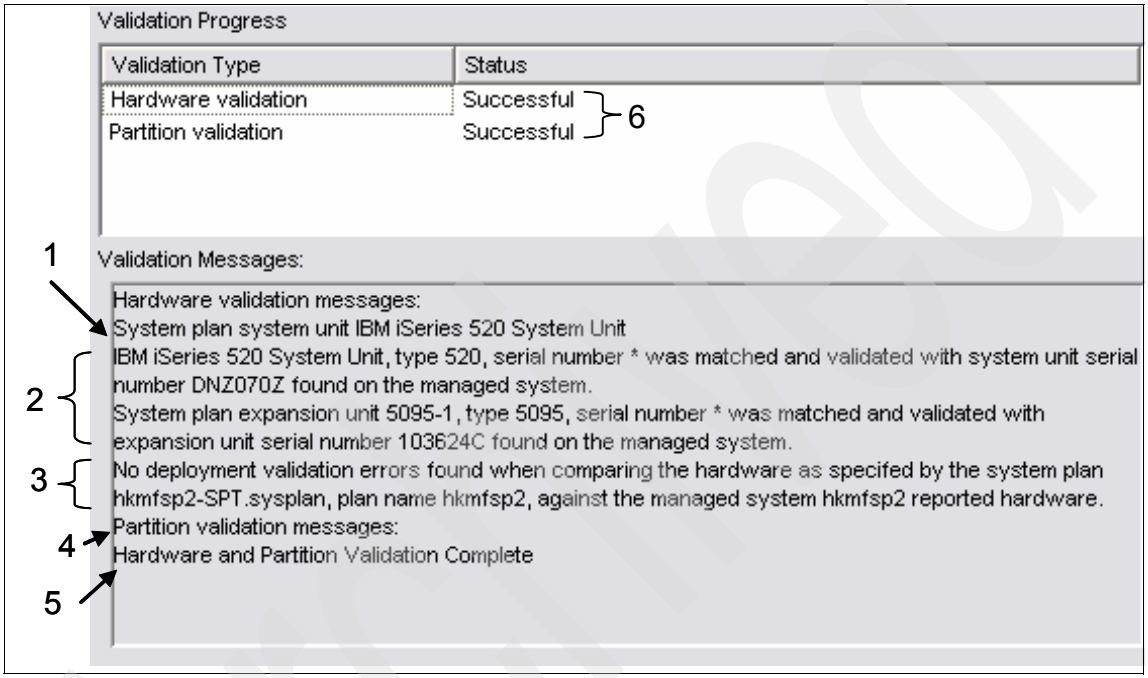

Figure 5-8 Example of successful validation

2. These messages are posted as the plan is satisfied by the physical units. Note that the messages report the serial number of the physical units that where matched. If the plan contained serial numbers, they would be shown instead of an asterisk (*). In this example the system plan deployed successfully but the plan that was deployed did not specify the system or expansion unit serial numbers.

   This is important. If inaccurate serial numbers or no serial numbers appeared in the system plan, these messages indicate which planned units where mapped to which physical units. As each planned unit is satisfied, the slot-DRC-indices are assigned.

3. This message is posted when the hardware validation has completed successfully.

4. This message indicates that partition validation is has started.
5. This message indicates that the validation process has completed.

These messages indicate that the validation completed successfully.

5.4.2 Unsuccessful hardware validation due to system verification failure

Figure 5-9 shows a validation error due to insufficient memory to satisfy the plan. This example is a failure of the system-level validation that is described in step 1 in 5.1.1, “Validation with accurate serial numbers” on page 88:

1. This message indicates that the type and model of the system unit validated successfully.
2. This message indicates that the physical system does not have enough memory to satisfy the plan.
3. This message explains that the validation failed, and to see the previous messages to determine why.

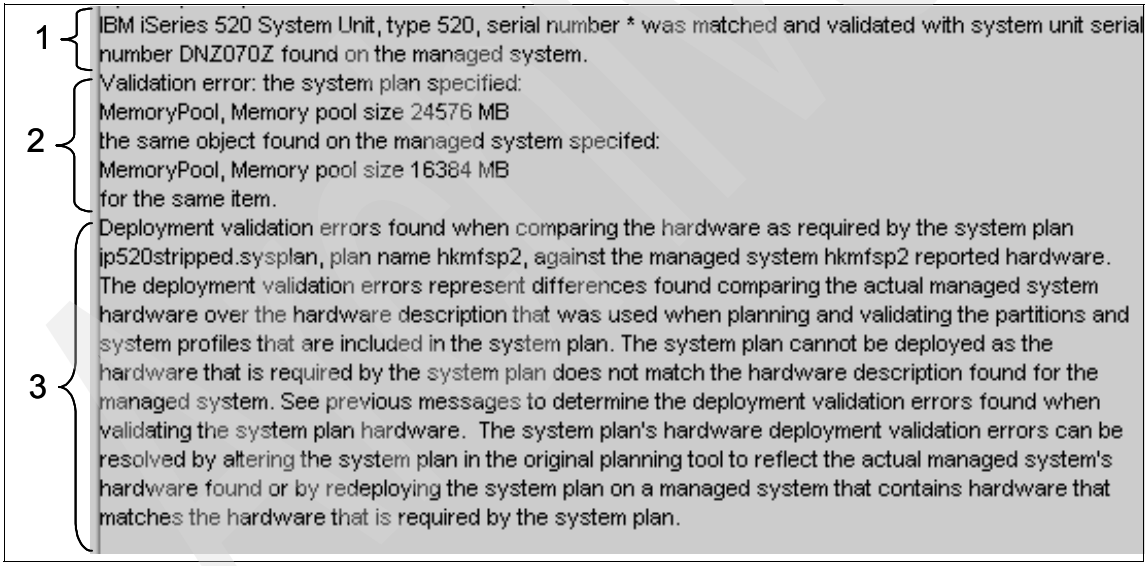

Figure 5-9 Validation failure due to insufficient memory

5.4.3 Unsuccessful hardware validation due to hardware mismatch

Figure 5-10 shows a validation error due to planned hardware not being satisfied by the physical system. In this example, the plan calls for a 5761 in slot C07 in the 5095 expansion unit. The physical hardware has a 626 in that slot.

1. These messages indicate that the system unit validated successfully.
2. This message calls out that a planned card does not match the physical card in the system. We only see this message because there is only one 5095 to check. If there where multiple 5095 units to check, we would only get a message stating that the planned unit was not satisfied by any physical unit.
3. This message explains that the 5095 in the plan is not satisfied by the only possible physical unit available.
4. This message explains that the validation failed, and to see the previous messages to determine why.

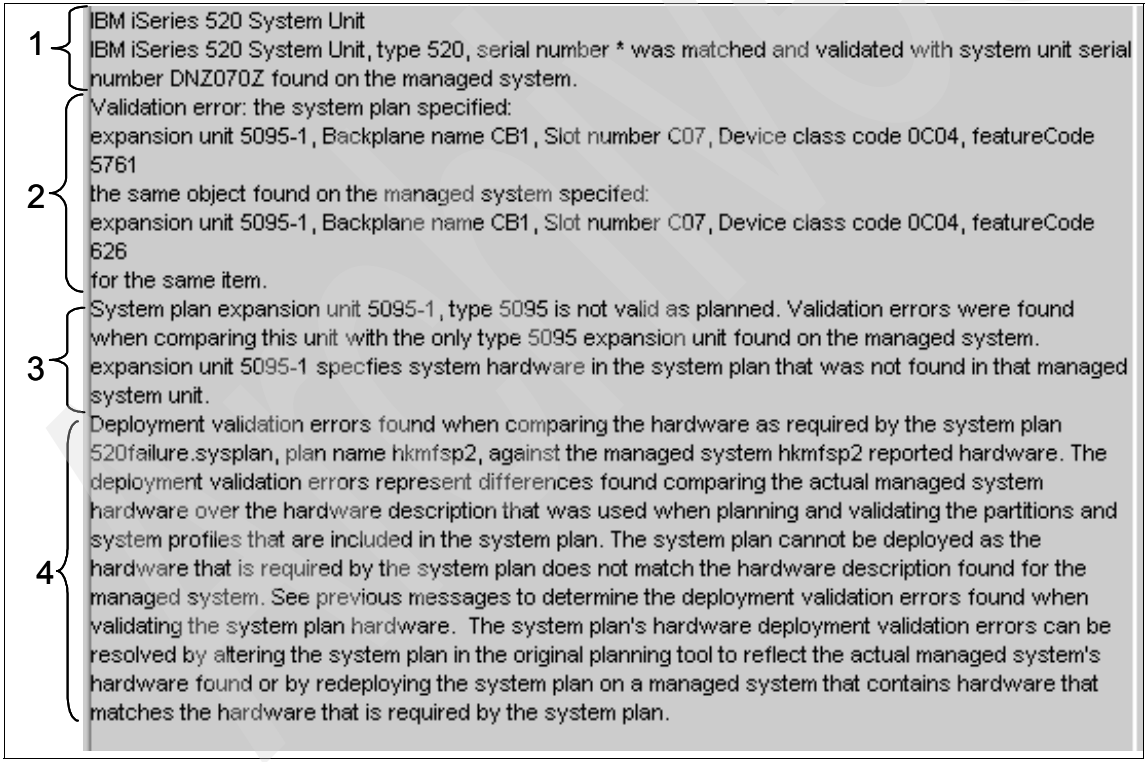

Figure 5-10 Validation failed do to hardware mismatch

5.4.4 Unsuccessful hardware validation due to missing unit

Figure 5-11 shows messages associated with a failure due to the physical hardware not having a unit that is in the plan. Figure 5-11 shows the messages that are generated when the validation fails.

1. These are typical messages for a satisfied system unit and an expansion unit.
2. This message indicates that the plan has a 5094 unit specified, but the physical hardware does not have any 5094.
3. This message explains that the validation failed, and to see the previous messages to determine why.

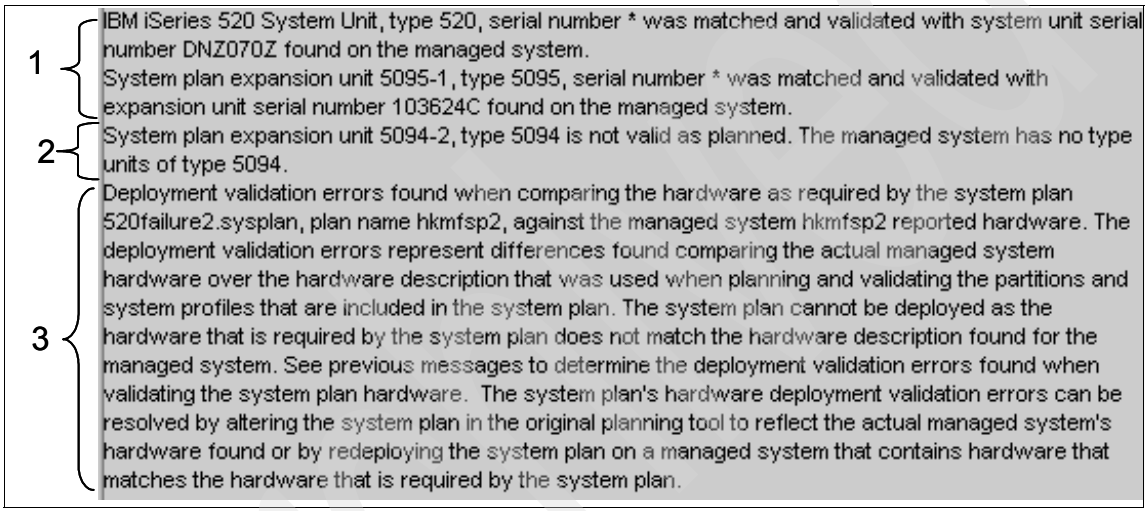

Figure 5-11 Validation failure due to missing unit

5.4.5 Successful validation with preexisting partitions

Figure 5-12 shows what messages to expect in a successful validation with preexisting partitions. This would happen on subsequent deploys after an initial partial deploy.

1. These are typical successful hardware validation messages.
2. This message explains that existing partitions where found and that they satisfied the plan.

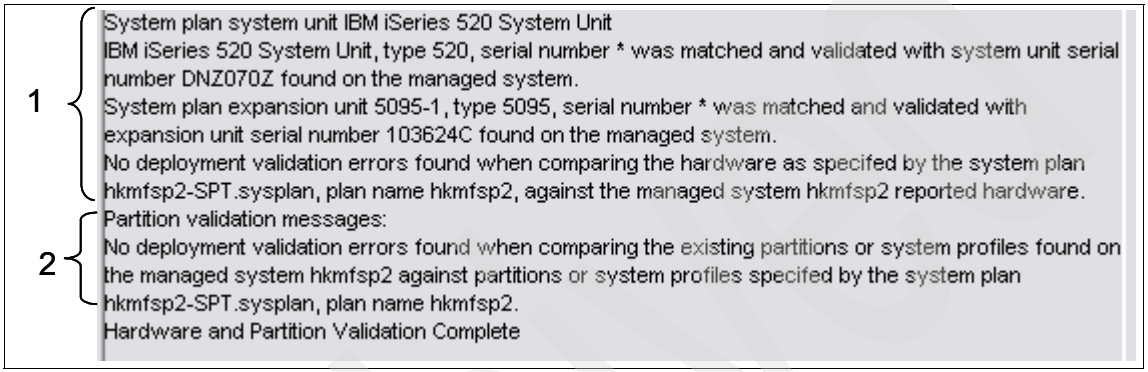

Figure 5-12 Successful validation with preexisting partitions

5.4.6 Unsuccessful partition validation due to existing partition not satisfying plan

Figure 5-13 on page 108 shows a failed partition validation due to the existing partition not satisfying the plan. In this instance, the existing partition has lower values for minimum and desired processor resources than the plan allows for.

1. These messages show typical successful hardware validation errors.
2. These messages explain that the minimum and desired processor values do not satisfy the plan.
3. This message is displayed whenever a partition validation fails. It directs the user to see the previously posted messages to determine why the validation failed. In this example a planned partition LPAR2 already exists on the managed system but the attributes for the actual existing LPAR2's minimum and desired processors did not match those planned values found in the system plan for LPAR2. It also explains that either the plan or the existing partition must be changed to allow the validation to succeed.

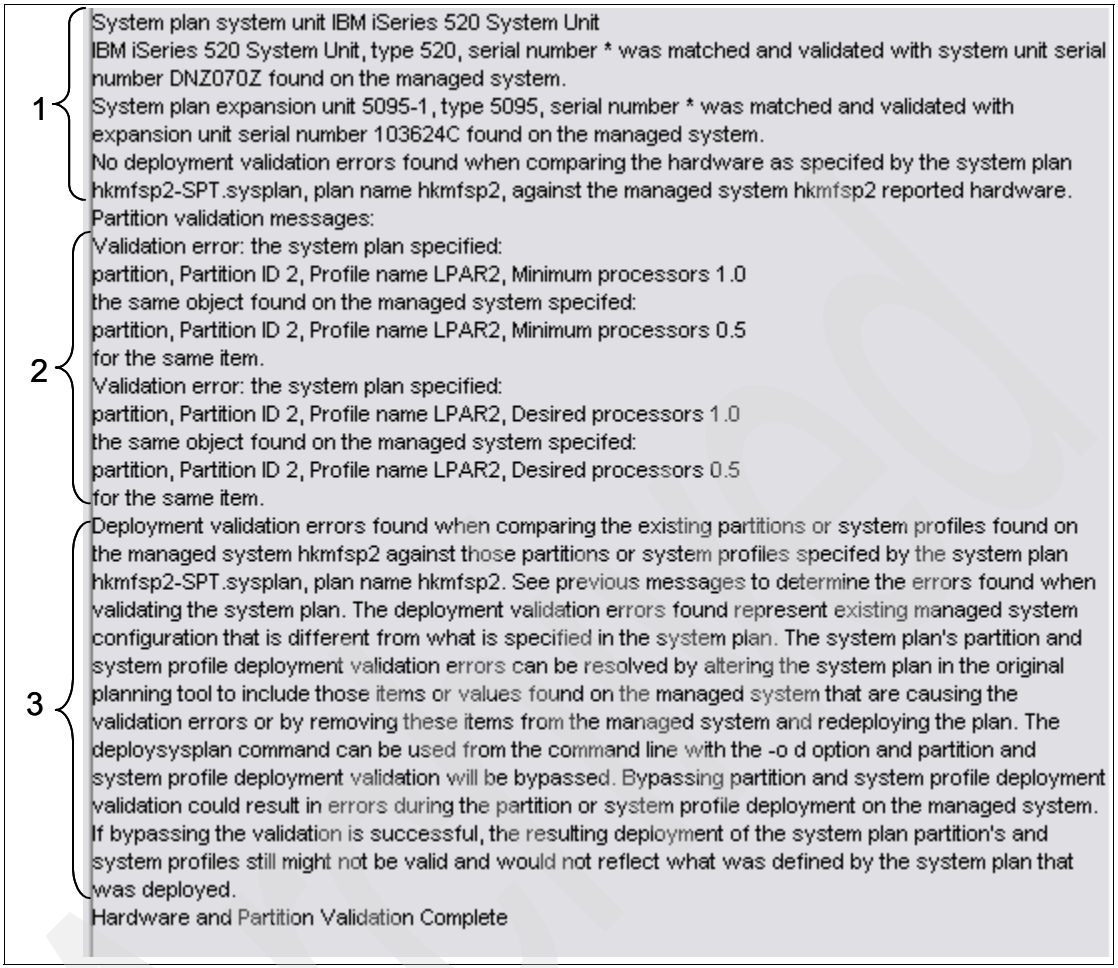

Figure 5-13 Validation failure due to existing partition failing to satisfy planned partition

5.5 Recovery from Validation Errors

The recommended recovery action for any hardware or partition validation error is to do a **mksysplan** against the managed system that was being used to do the validation. This can be done from the command line on the HMC where the system plan was being deployed. The **mksysplan** command would look like:

```
mksysplan -f <myplanname.sysplan> -m <myvalidationsystem>
```

Where <*myvalidationsystem*> would be the name of the managed system that was validated against the system plan.

Once this new system plan is generated the Manage System Plans menu item from the WebSM interface can be opened and that new system plan can be viewed. Using the viewer on this system plan and on the system plan that failed to validate is a way to compare the system plans side by side to look for those attributes that were called out as validation errors.

Figure 5-14 is an example of a partition that failed. In this case partition 1 found on the managed system was called *LPAR1*, and that partition's profile was also called LPAR1. The system plan specified that the first partition and its profile should be called *test*. Since what was found on the managed system for the first partition was not part of the system plan being deployed, that system plan failed partition validation.

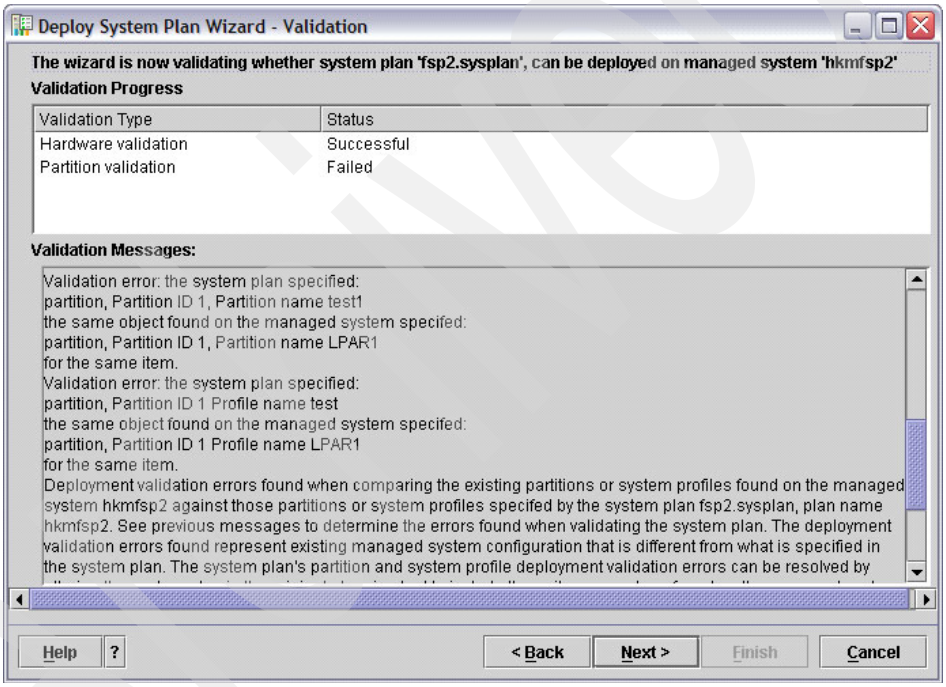

Figure 5-14 Example validation message

To try to diagnose this problem the command **mksysplan** was run on the **hmc** command. In this example the managed system that was being validated was called hkmfsp2. The **mksysplan** command to create a system plan for this system looks like this:

```
mksysplan -m hkmfsp2 -f hkmfsp2currentLPARs.sysplan
```

The system plan that is created is called hkmfsp2currentLPARs.sysplan. This system plan should reflect the hardware and the partitions that already exist on the system called hkmfsp2.

The next step to diagnosing the problem is to open both the system plan that failed to deploy with the view option from the Manage system plans menu. In this example the system plan that failed is called fsp2.sysplan, and the system plan that was just created, hkmfsp2currentLPARs.sysplan, is opened with the view option as well. See Figure 5-15.

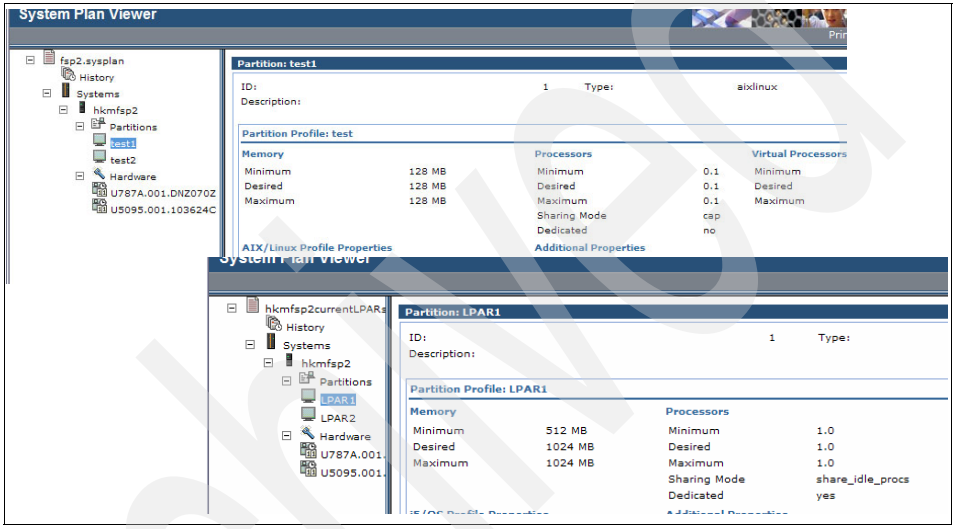

Figure 5-15 System view

Opening both plans with the viewer is the best possible way to try to figure out what to do about the validation error found. In this case it is obvious that partition ID 1 is not contained in the managed system, as shown in the system plan hkmfsp2currentLPARs.sysplan. This also shows that it is not the same configuration as the partition ID 1 deployed in system plan fsp2.sysplan. Viewing these two system plans (the failing plan and a plan that was made with **mksysplan** from the managed system) side by side allows not only this difference but other differences between what is in the system plan and what is actually on the managed system to be seen. This view not only shows that LPAR1 is different from the system plan, but also that there is another partition LPAR2 that is also not part of the fsp2.sysplan.

In this case a possible resolution is to bring the fsp2.sysplan system plan back to an SPT planning tool and re-configure the first partition in that plan to match the partitions LPAR1 and LPAR2 that are currently configured on the managed system. Another possible resolution might be if LPAR1 and LPAR2 are no longer

needed on the managed system hkmfsp2, remove partitions LPAR1 and LPAR2 from the hkmfsp2 managed system and then redeploy the system plan that failed.

This example showed an example of using the **mksysplan** and the HMC system plan "view" to diagnose a partition validation error, but the same technique can be used to try to diagnose hardware validation problems.

A

Installing SPT

In this appendix we describe the steps necessary to download and install the IBM System Planning Tool product.

Downloading the SPT file

The new IBM System Planning Tool is available for download at:

http://www-03.ibm.com/servers/eserver/support/tools/systemplanningtool/

Figure A-1 IBM System Planning Tool Web site

On the SPT home page browse down to the SPT download link and click it to start the process.

You will be presented with a standard disclaimer. If you choose to accept the legal notices, page down to the bottom of the notice and click **I agree**. You will be presented with the file download panel. Click **Save** and store the file on your PC.

Running the SPT install .exe

Now that you have downloaded SPT, browse to the folder where you saved the .exe file. To start the install in our example we browse to the LPAR Simplification lab folder and then find the v1r20-061506.exe, as shown in Figure A-2.

Note that the value v1r2_061606 is specific to our example and may differ from the real downloadable file name and its versioning.

1. Double-click the icon to start the install.

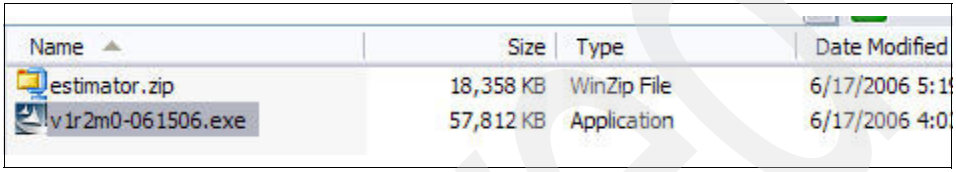

Figure A-2 Browse to install exe

InstallSheild will start as shown in Figure A-3.

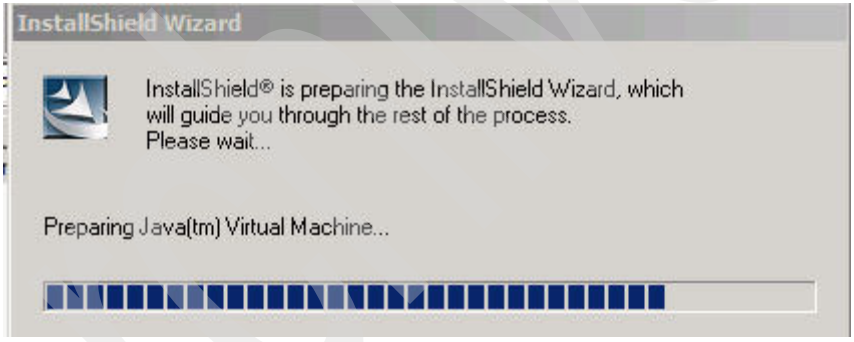

Figure A-3 InstallShield start

2. As the InstallShield product starts you may get a warning message from your PC firewall. Check the **Remember this setting** box and then click **Allow**.

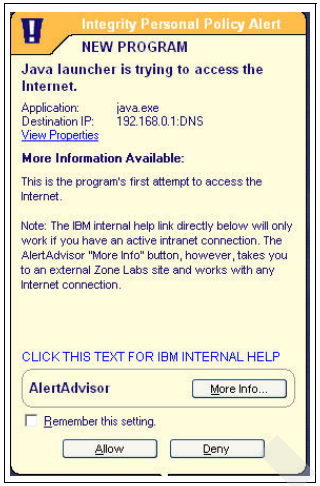

Figure A-4 Firewall warning message for Java™.exe

The SPT splash screen will appear, as shown in Figure A-5.

Figure A-5 System Planning Tool splash screen

3. At this point you may see a second firewall warning message asking about SPT.exe. Check the **Remember this setting** box and then click **Allow** (Figure A-6).

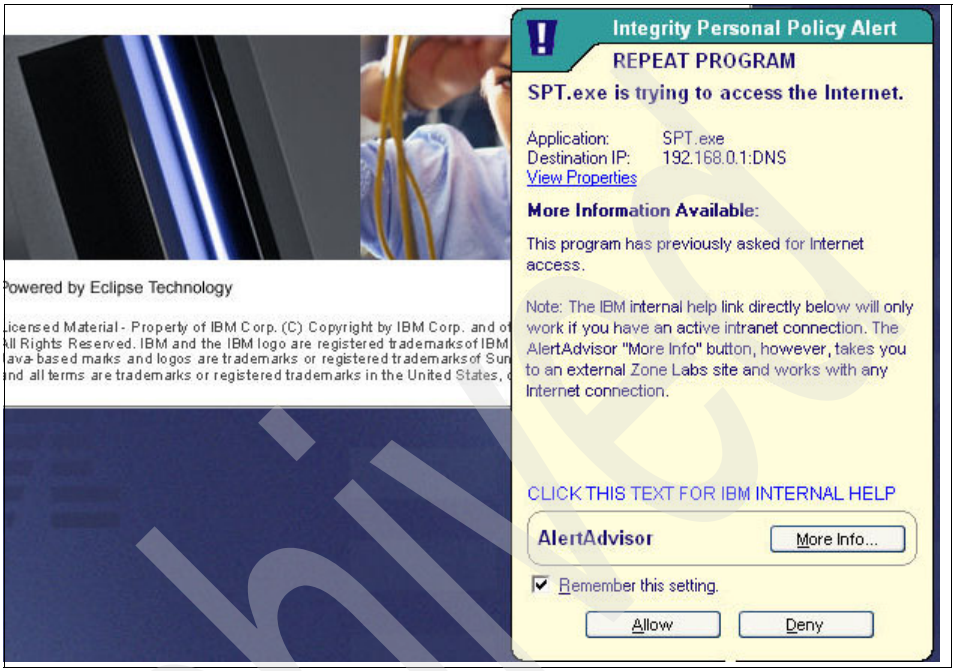

Figure A-6 Firewall warning message for SPT.exe

The SPT product install will initialize and a message will display, as shown in Figure A-7.

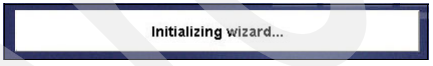

Figure A-7 Install initialization message

The SPT installation wizard will start, as shown in Figure A-8.

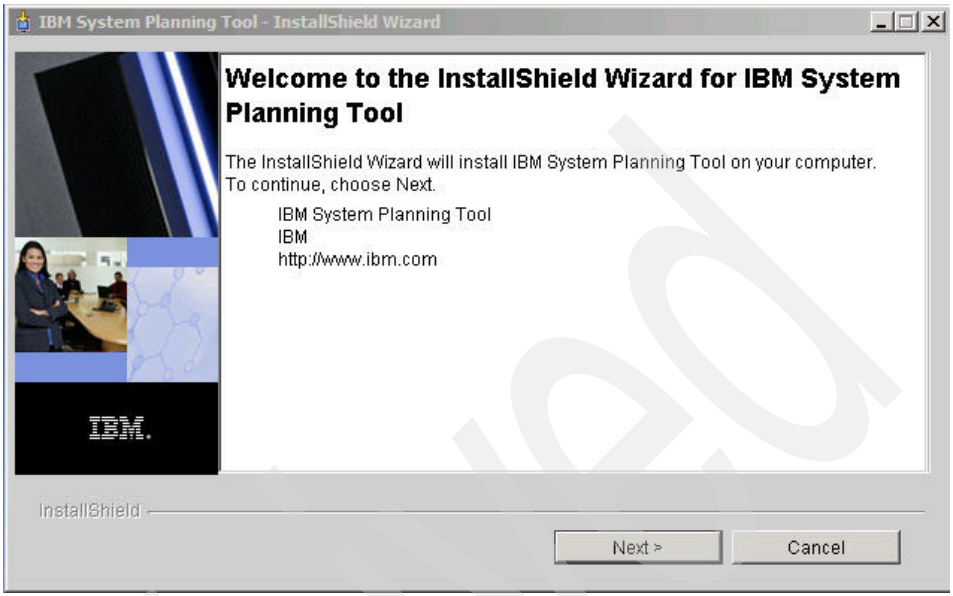

Figure A-8 SPT install wizard 1

4. Next comes the standard software license acceptance panel (Figure A-9). Click the **I accept** radio button and click **Next** to continue. For non-English speakers the license agreement is translated. Access translations from the pull-down menu.

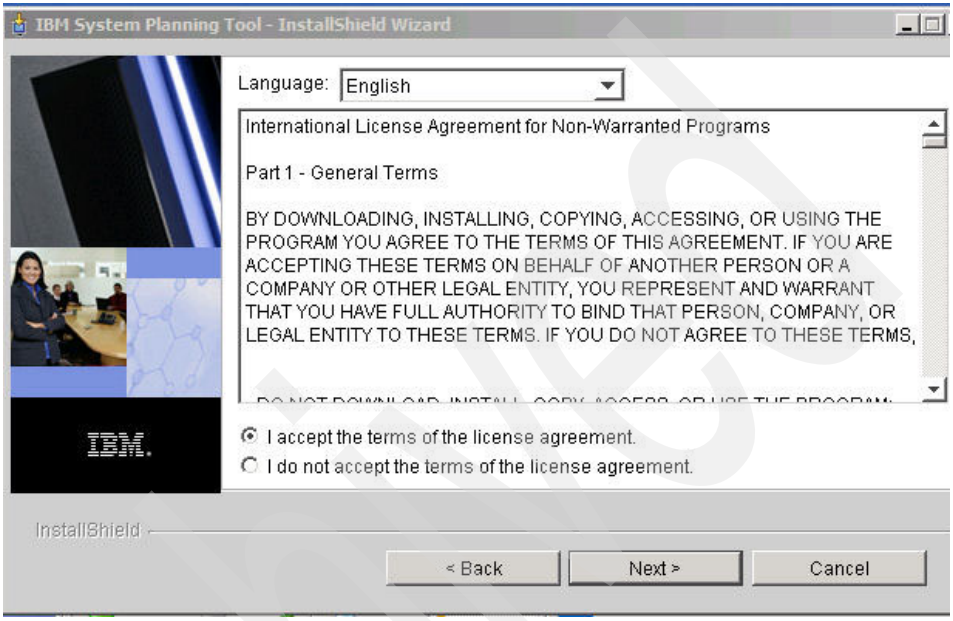

Figure A-9 SPT install accept license

5. The SPT install process will now present the default install path c:\Program Files\IBM\IBM System Planning Tool. You could browse to another directory for install, but we suggest that you click **Next** to accept the default path and continue.

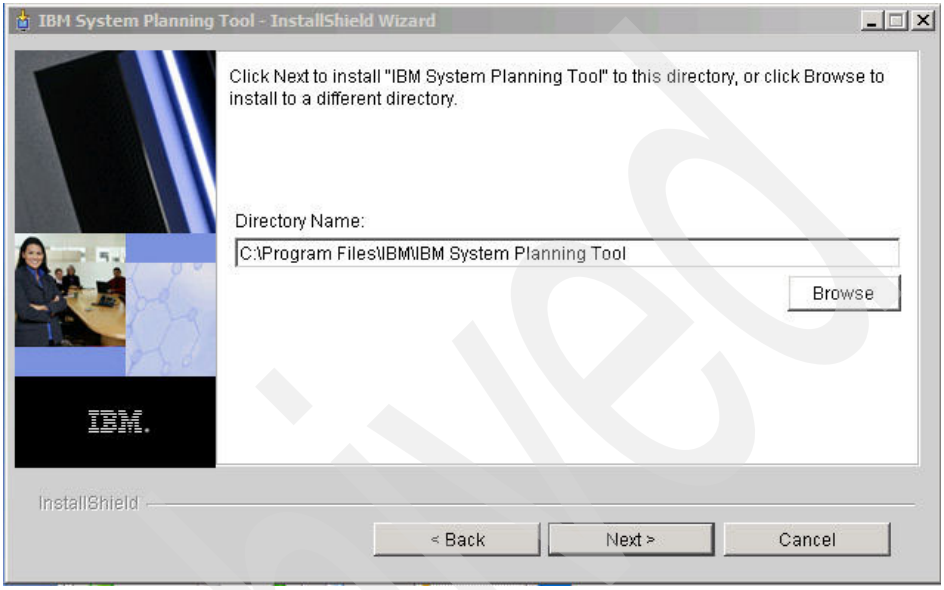

Figure A-10 SPT install folder

6. You will now be presented with a port selection screen (Figure A-11). This allows you to change the default port that System Playing Tool will use depending on your particular firewall scheme. Click **Next** to continue with the default or change to a different working port and click **Next**.

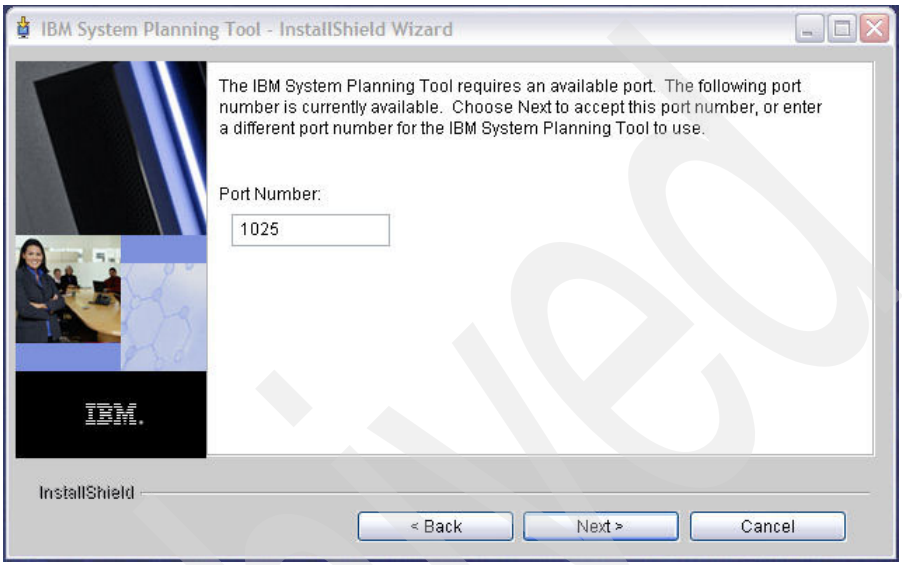

Figure A-11 SPT port selection

The install process displays the summary of install selections that are deployed (Figure A-12). SPT will take approximately 77.3 MB of disk space.

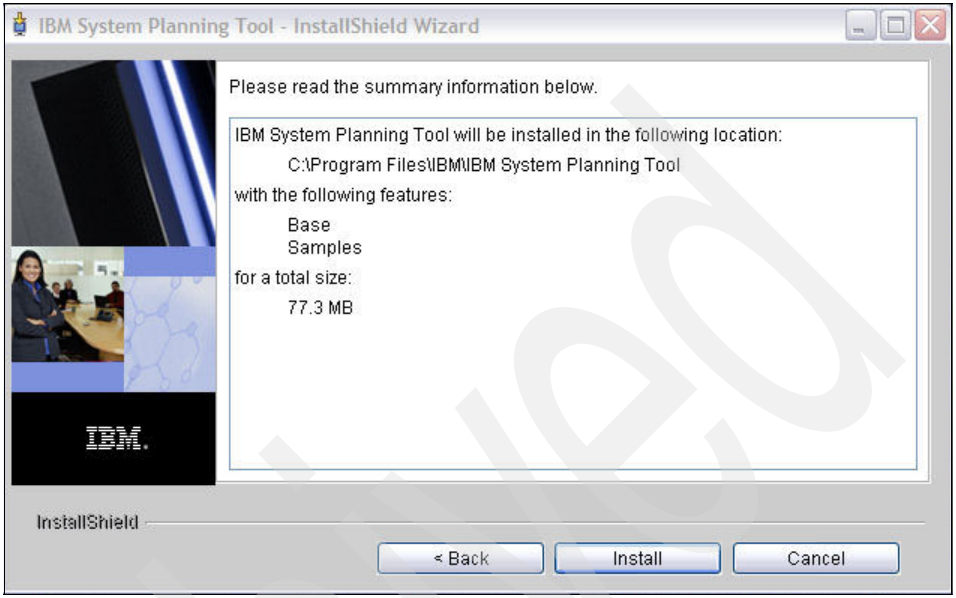

Figure A-12 SPT install summary

7. Click **Next** to start the install process.
8. The Installing IBM System Panning tool install progress panel is displayed (Figure A-13). When this completes a next button is turned on. Click the **Next** button and continue.

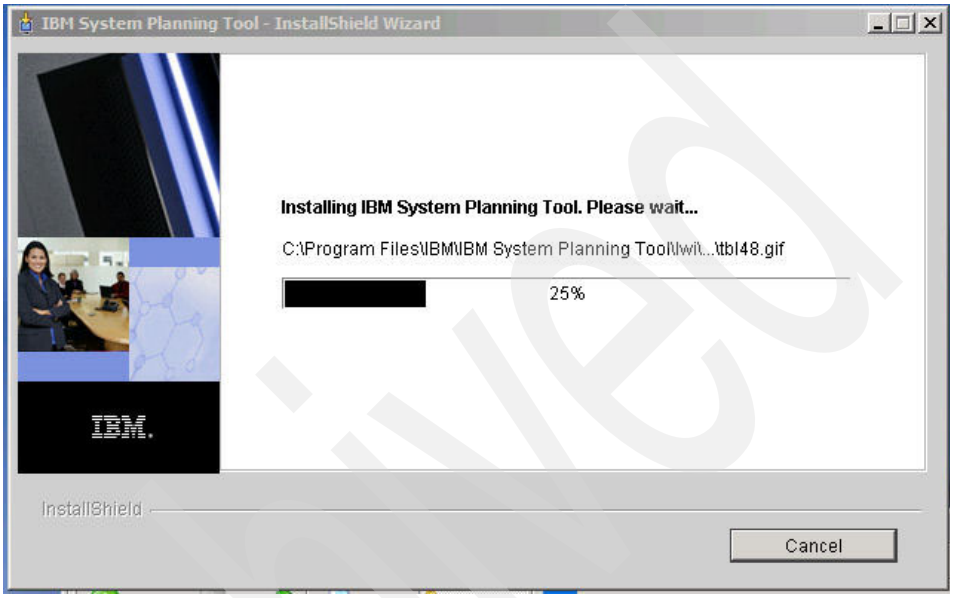

Figure A-13 SPT install progress

9. You are now presented with the final screen (Figure A-14). Click **Finish** to complete the install of the System Planning Tool.

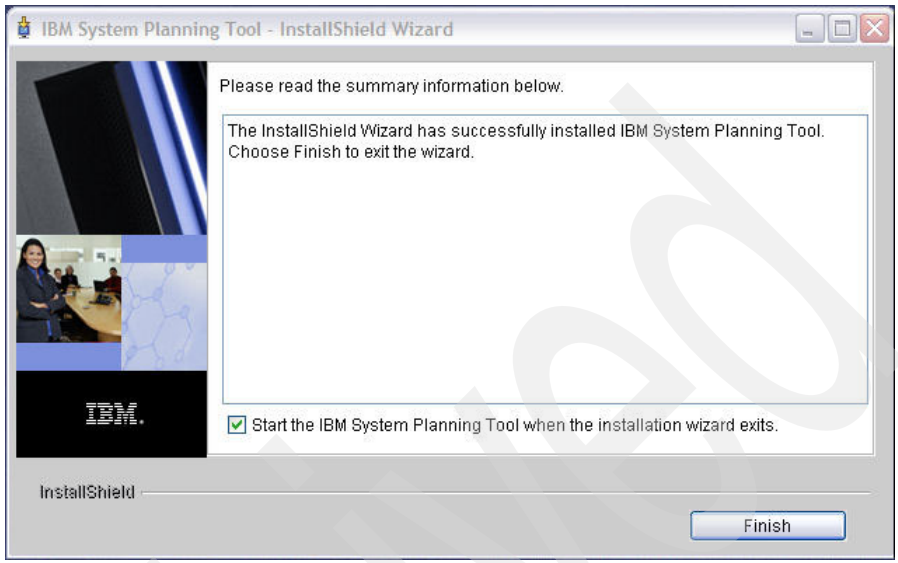

Figure A-14 SPT install complete

10. Note that if you leave the Start System Planning Tool check box a browser will open as soon as you click Finish. This will start the SPT product. You do not want to start SPT if you are planning to install WLE. In this case you should uncheck the start box and click **Finish**.

See Appendix B, “Introduction to using WLE with SPT” on page 129, for installation instructions.

Starting the System Planning Tool

From **Start** → **Programs** → **IBM System Planning Tool**, click **System Planning Tool** to start (Figure A-15).

Figure A-15 Start System Planing Tool

The tool takes a few seconds to initiate.

Note: If SPT fails to open from the programs selection it may already be running. There is a small System Planning Tool icon in the lower right-hand quick launch area. Click this icon to start the tool.

The Welcome splash screen (shown in Figure 2-1 on page 14) displays.

If you have just installed SPT and had the start SPT box checked as you finished the install, SPT will be running. Even if you closed the browser window this does not close SPT.

You can check whether SPT is running. If there is an SPT icon in the system tray, SPT is already started. See Figure A-16.

Figure A-16 SPT running icon in system tray

Ending SPT

To end SPT right-click the SPT icon in the system tray. You will be presented with a menu list, as shown in Figure A-17. Select **Exit** to end SPT.

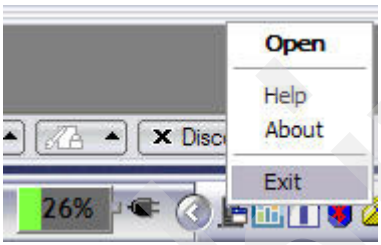

Figure A-17 Ending SPT

You should end SPT before installing WLE.

B

Introduction to using WLE with SPT

This appendix discusses the installation of Workload Estimator and its integration with IBM System Planning Tool.

This is a completely optional installation. You can configure system plans from scratch via the Create new option or with one of the sample system plans.

WLE download site

To download the IBM Workload Estimator:

1. Browse to the IBM Workload Estimator home page:

   http://www-912.ibm.com/supporthome.nsf/document/16533356

2. The WLE home page is shown in Figure B-1. Use the link shown in the browser window to start the WLE download. Follow the on-screen instructions.

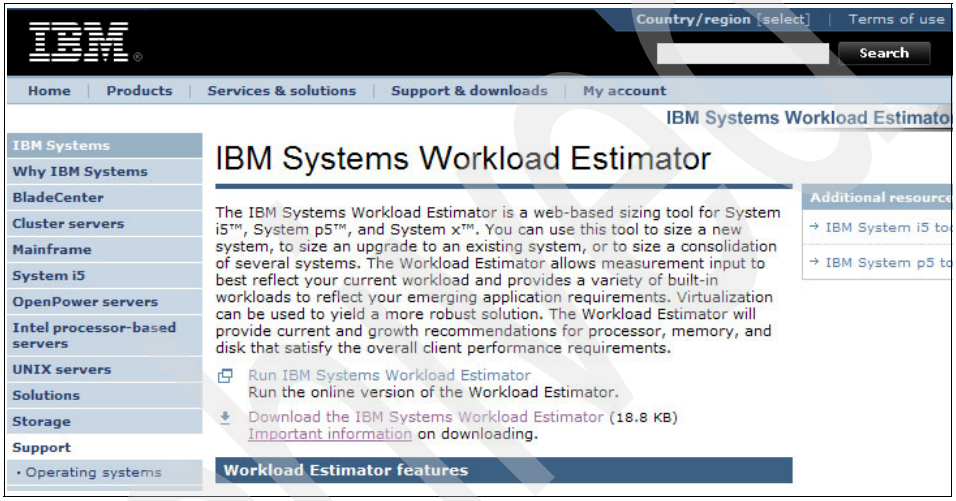

Figure B-1 IBM Workload Estimator home page

This will download the estimate.zip file, similar to that shown in Figure B-2.

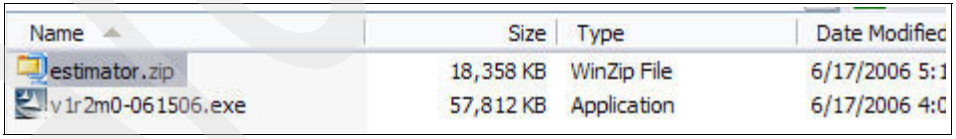

Figure B-2 Estimator aip file saved to folder

3. Use an extract program to decompress the zip, and the installation files will be extracted to the default directory estimator, as shown in Figure B-3.

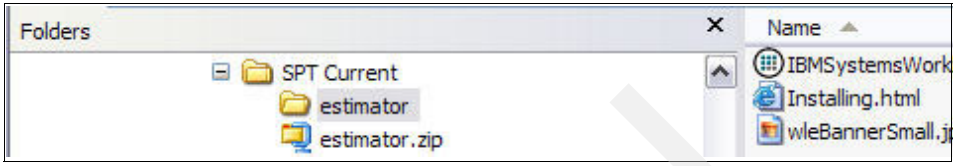

Figure B-3 Estimator zip file extracted

4. Before you start the WLE install process you should check that the IBM System Planning Tool is not running. Look for the System Planning Tool icon in the system tray on the lower right-hand side of the PC desktop. The icon is shown in Figure B-4.

Figure B-4 SPT running icon

5. To stop the System Planning Tool, right-click the icon and a small menu appears, as shown in Figure B-5. Click **Exit** to end the System Planning Tool.

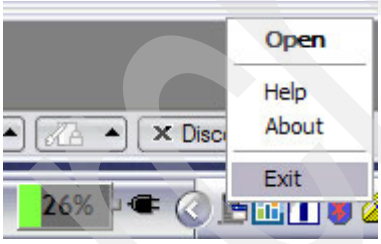

Figure B-5 Exit SPT

6. Now we are ready to start the WLE product install. Double-click the **IBMSystemWorkloadEstimator.exe** file, as shown in Figure B-6, in the estimator folder.

Figure B-6 Select IBM Systems Workload Estimator exe file

This will start and InstallShield process similar to the SPT install and a splash screen will appear, as shown in Figure B-7.

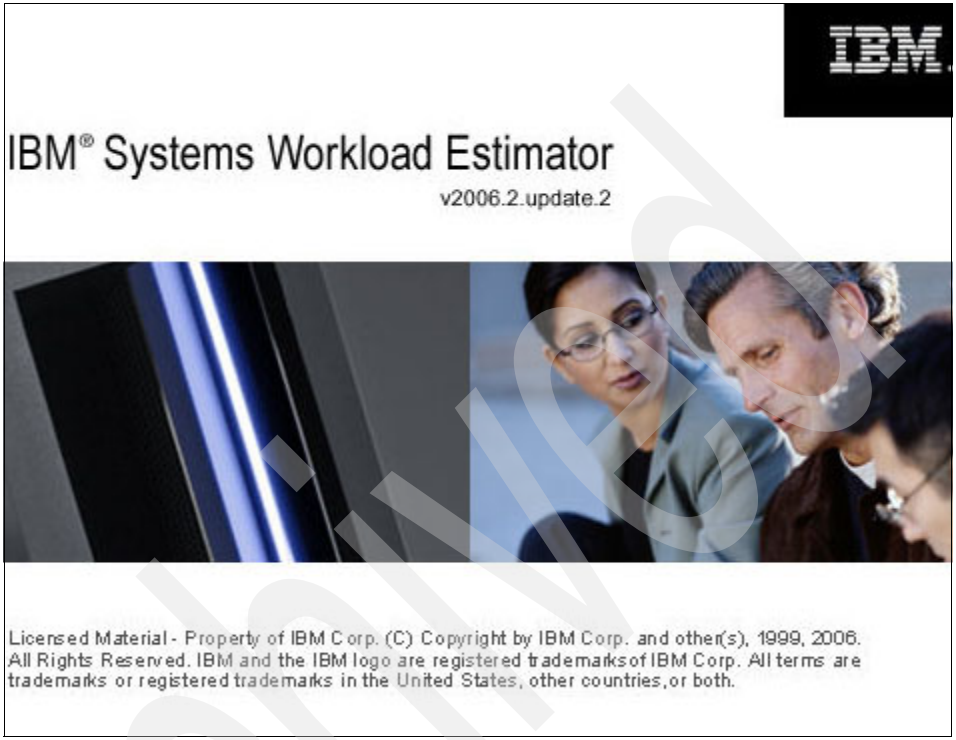

Figure B-7 IBM Systems Workload Estimator splash screen

7. As this process begins you may see firewall warning messages similar to the one shown in Figure B-8. You must respond to this message to continue with the installation based on your particular firewall scheme.

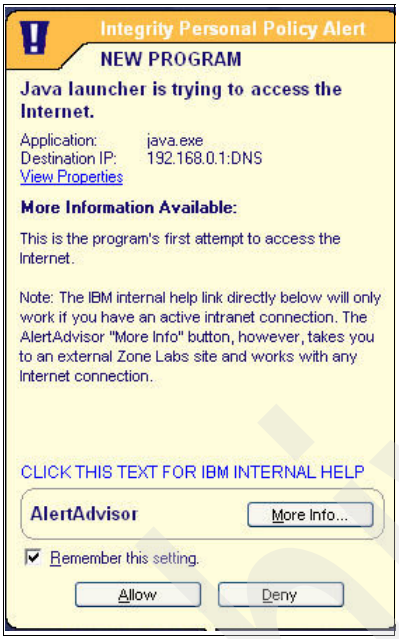

Figure B-8 Firewall detection message

In our case we checked the remember box and then allowed the program to access the Internet. A wizard initialization panel is momentarily displayed. This shows that the process is continuing.

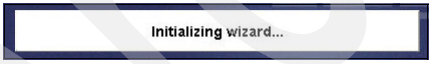

Figure B-9 Wizard initialization

The next screen is a additional wizard progress message (Figure B-10).

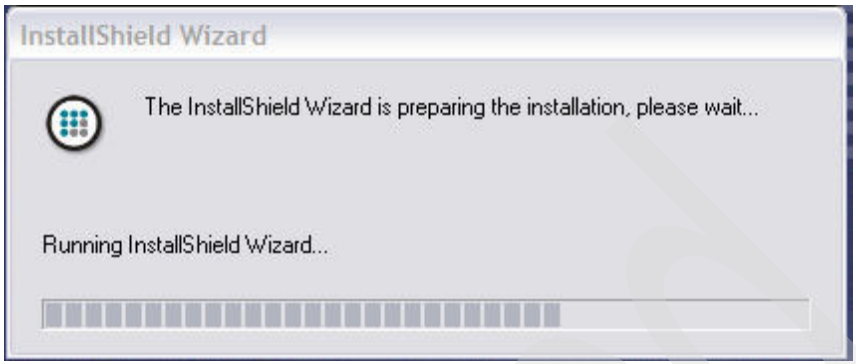

Figure B-10 InstallShield running panel

8. Once initialization is complete, the WLE Welcome screen displays, as shown in Figure B-11. Click **Next** to continue.

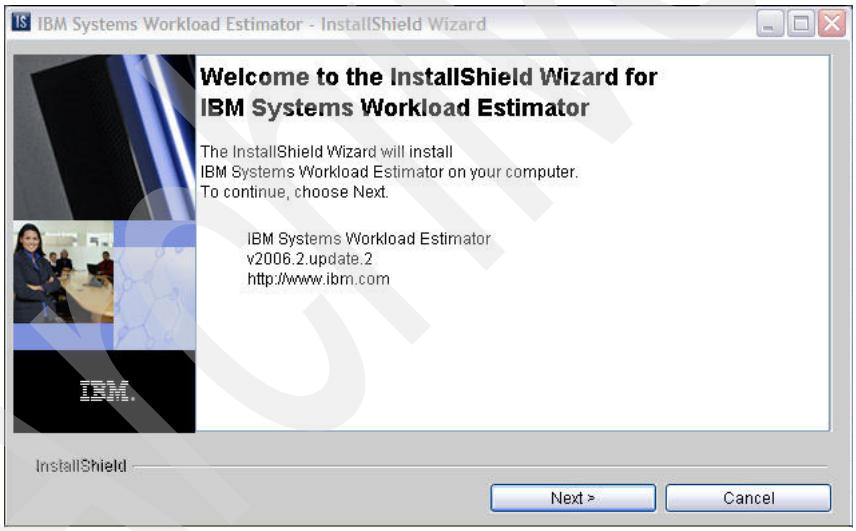

Figure B-11 Workload Estimator install welcome screen

9. The next step in the process asks you to accept the WLE license agreement. Click the **I accept** radio button and **Next** to continue (Figure B-12).

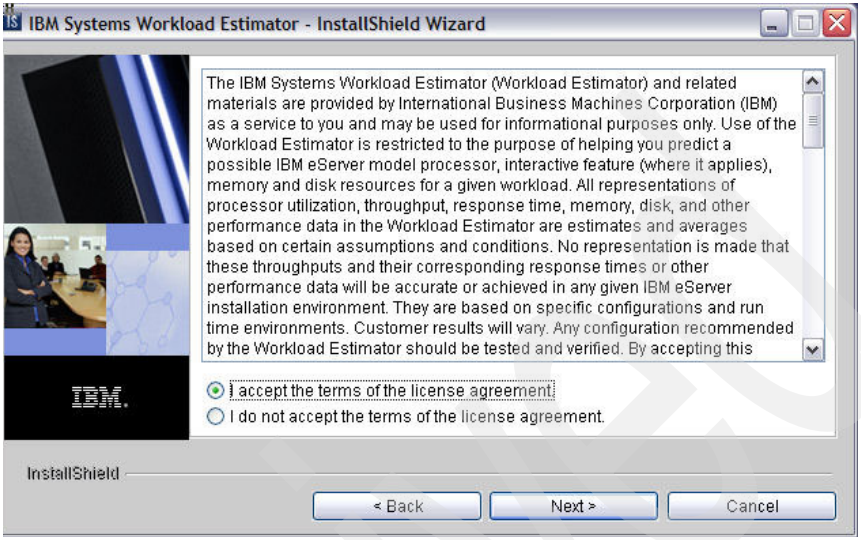

Figure B-12 WLE licence acceptance

10. The next step is to decide whether you want to accept the default installation directory (Figure B-13). Our suggestion is to accept this directory and click **Next**.

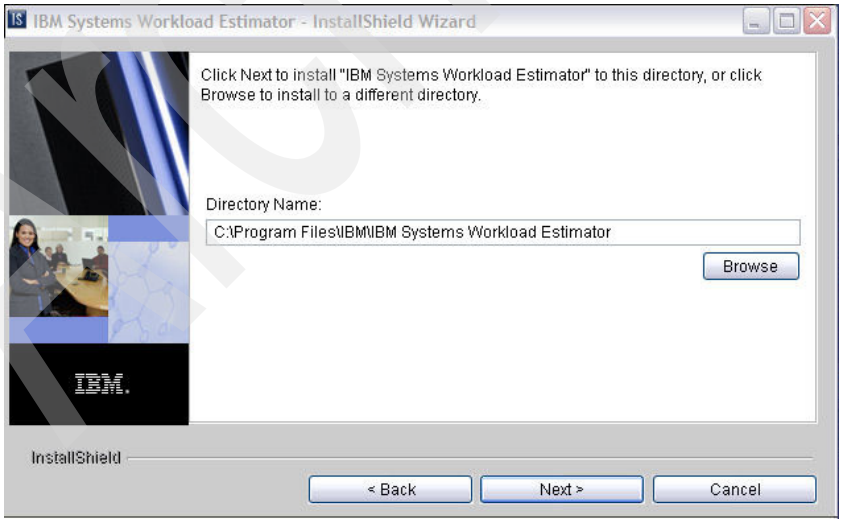

Figure B-13 WLE default install directory

11. The install process detects that the IBM Systems Planning Tool is already installed in the PC and therefore it will only install the WLE support programs and not the full WLE install. This means that WLE can only be started from inside SPT. There is also a warning message. This reminds you to make sure that SPT is not running. We assume that you have stopped SPT, so click **Next** to continue (Figure B-14).

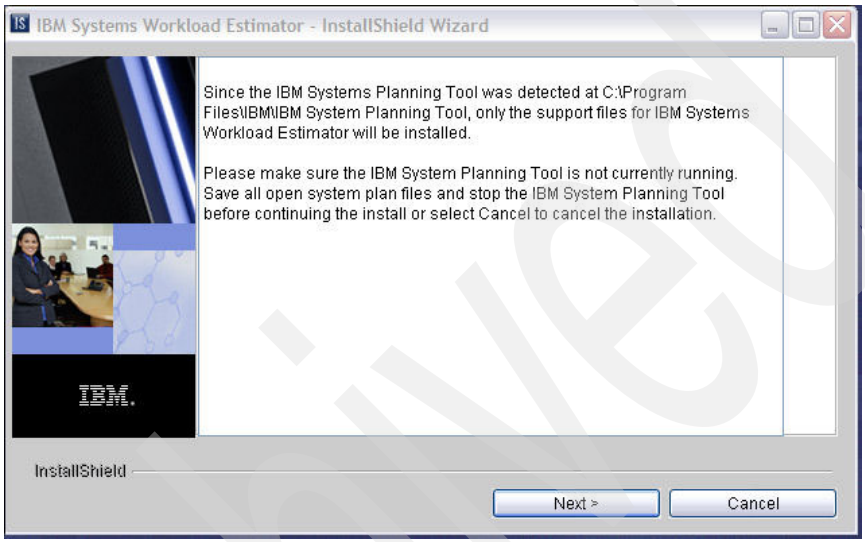

Figure B-14 SPT detection and warning message

12. The next step provides summary information including the install size of WLE. Click **Next** to continue and start installing the files (Figure B-15).

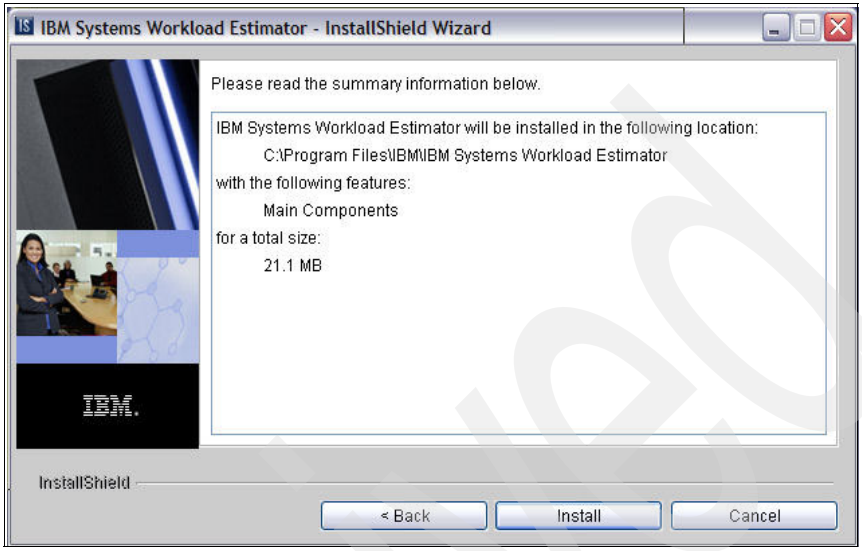

Figure B-15 WLE install summary

An install progress indicator displays (Figure B-16). This takes about 2 minutes, depending on the PC on which you are installing WLE.

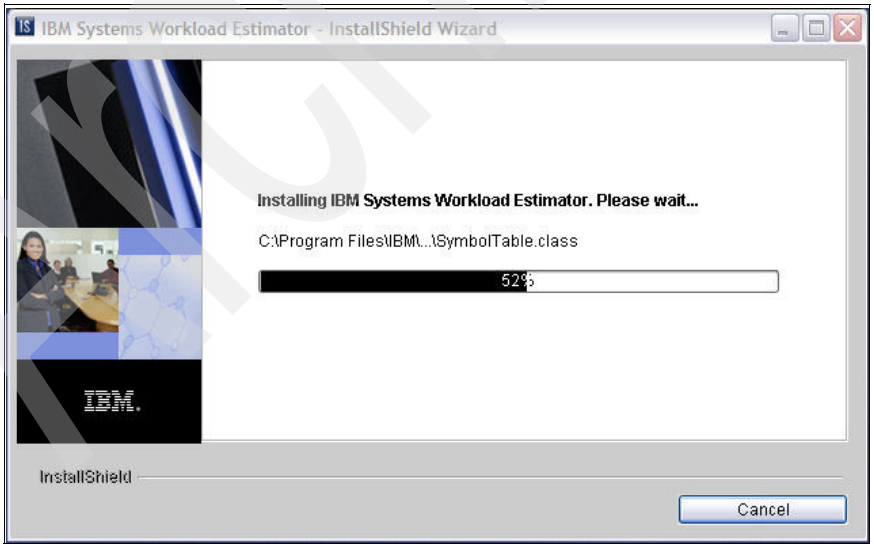

Figure B-16 WLE install running

13. This process completes and the finalization screen displays (Figure B-17). Click **Finish** to complete the install.

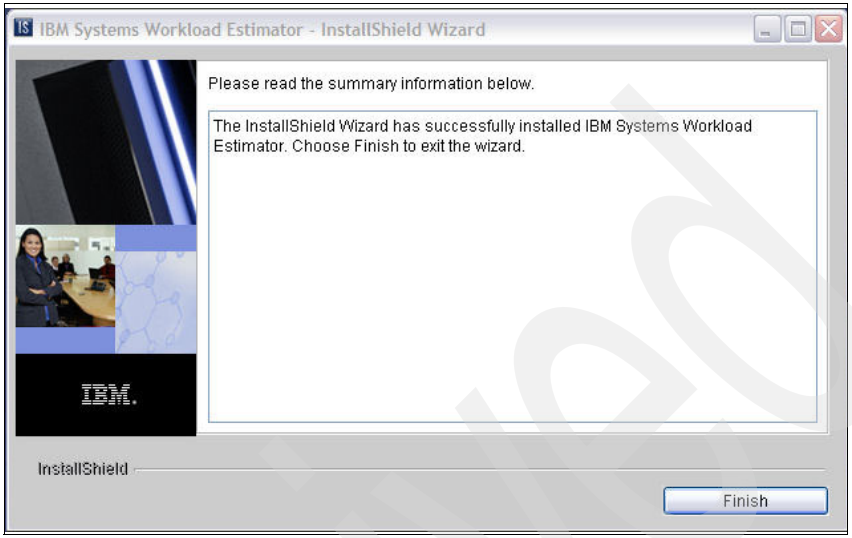

Figure B-17 IBM Workload Estimator install complete

At this point we have successfully installed WLE, and in the next section we show how to start the product.

Starting WLE in SPT

To start IBM Workload Estimator when it is installed in conjunction with IBM System Planning Tool:

1. Start SPT. Open the windows **Start** → **Programs** menu and then **IBM System Planning Tool** and **System Planning Tool**, as shown in Figure B-18.

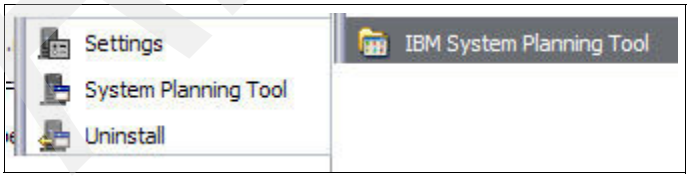

Figure B-18 Starting SPT

2. The SPT launch window opens and you should click **Launch** to display the Getting Started window, as shown in Figure B-19.

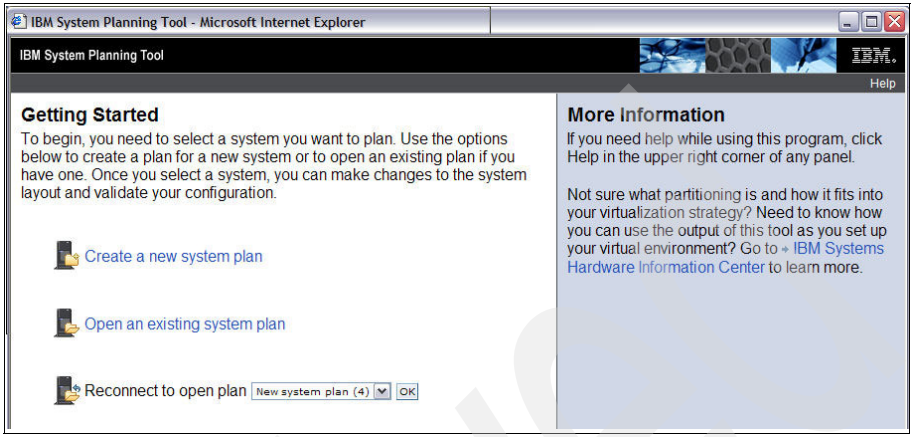

Figure B-19 SPT Getting started window

3. Click **Create a new system plan** to open the Add System panel, as shown in Figure B-20.

Add system based on new workloads

Assuming that you are planning to add a new workload or application, click **Based on new workloads I want to run** and then **Next** to open the new workload function of WLE (Figure B-20).

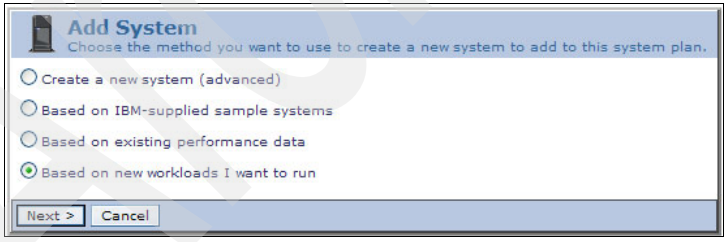

Figure B-20 SPT Add System panel

In the next step the WLE Add systems based on new workloads navigation panel will open to the Workload selection tab, as show in Figure B-21. On this first screen there is one additional tab, Help/Tutorial. If you are unfamiliar with WLE we recommend the you read the tutorial and examples before continuing with the configuration.

In the Workload Selection header there are a number of links. You can select the family that you will model in WLE (System i, p, or x). You can customize the WLE environment with the Options link. If you have already created a WLE model you can restore it for further work. You can also edit estimator information. This allows you to personalize this WLE model. The Workload Selection screen (shown in Figure B-21) has defaulted to My Solution with a single tier that has a System i that has one i5/OS V5R4 partition and not LPAR. You can change these defaults by using the Add, Remove, and Delete links on the right of the Workload Selection header section.

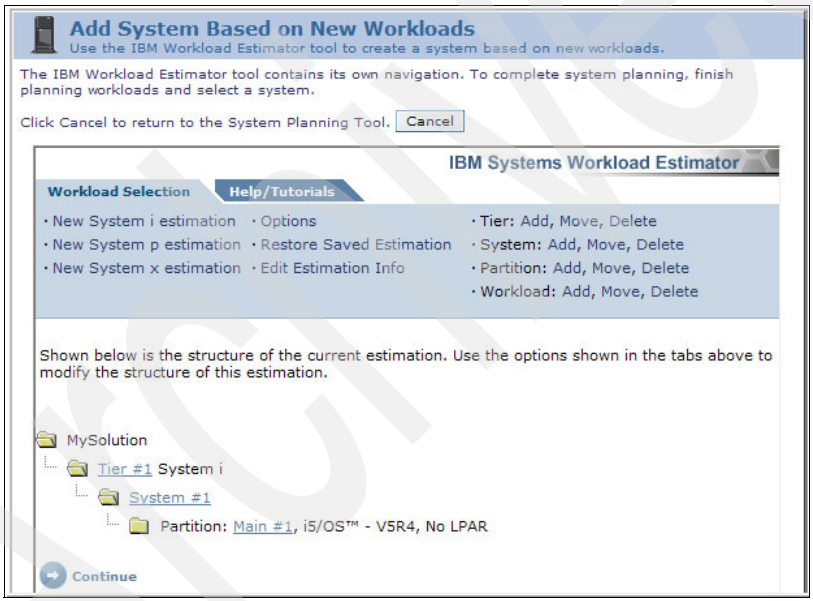

Figure B-21 Select workloads to add

At this point the default selections are:

- My solution - This is the overall container for the WLE estimate. You can add more systems with the Tier - Add link.
- Tier 1 - The system family, System i, p, or x. You can add multiple system as additional tiers with System - Add.
- System #1 - This relates to the System i with a default single default i5/OS partition that will model a V5R4 workload with no LPAR, as there are no

additional partitions. You can add partitions within System #1 with Partition - Add.

- Main #1 - This relates to the partition 1. There are no workloads running in this default partition. Select **Workload - Add** to put model workloads in this partition.

To show how the Workload Selection screen changes we have added two partitions (Main #2 and Linux #1). In the three partitions in our example we changed their properties to use shared processing resources, plus we have added workloads for Traditional, WebSphere®, and Web Serving, as shown in Figure B-22.

Figure B-22 WLE model with workloads added

We will skip to the point in the WLE model where we re-enter the System Planning Tool. You must save your WLE model before you continue to the System Planning Tool.

When you reach the Selected System panel in WLE, click **Continue** to move the model into SPT, as shown in Figure B-23. There is no back button once you enter SPT. Therefore, if you wanted to re-access your WLE model, you would have to close SPT and reopen the WLE application.

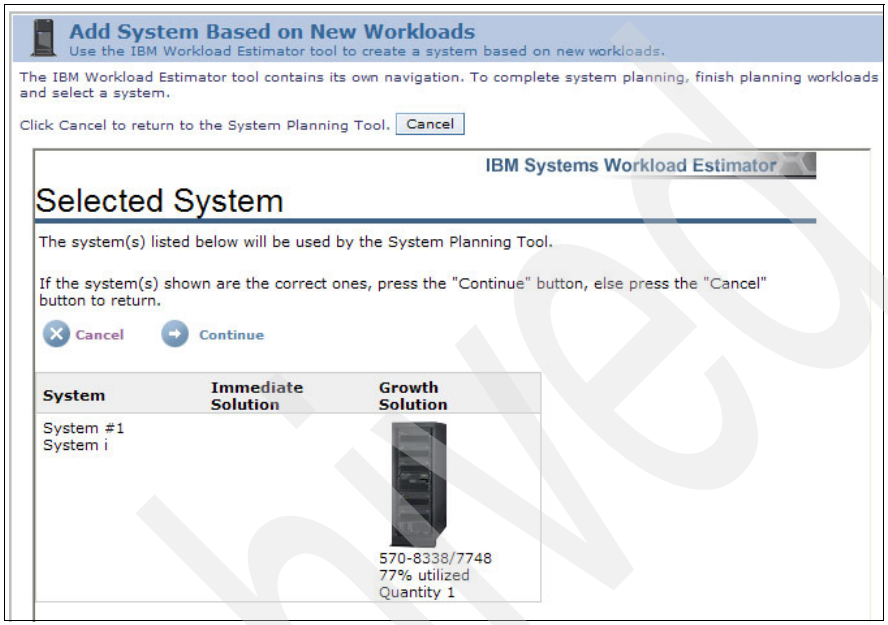

Figure B-23 WLE selected system

The SPT - Work with Planned System panel opens. From this point edit partitions and follow the SPT planning process described in 2.4, “Creating a new system” on page 16. This allows you to edit the partition information and complete the detailed processing, memory, and resources information.

Once the partition information is complete, you must place hardware and select the load source (for i5/OS partitions) and the console devices for each partition.

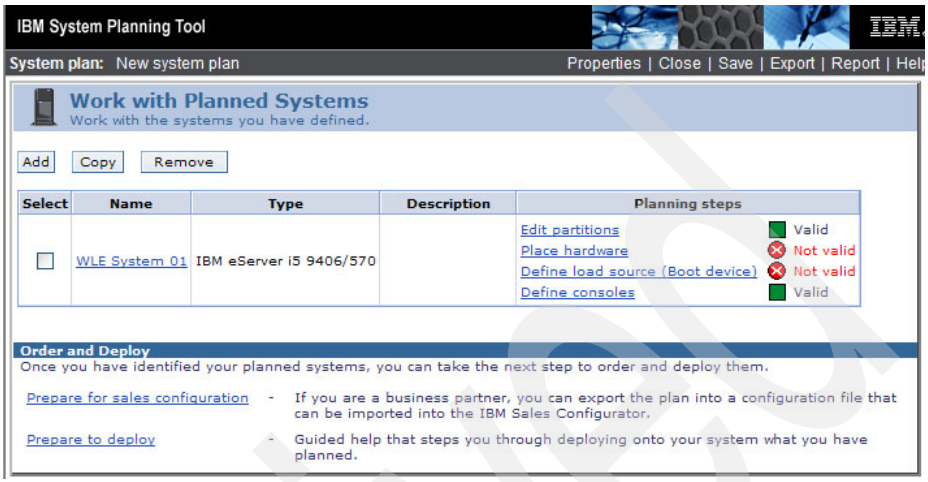

Figure B-24 WLE model in SPT

This completes the movement of a WLE model into SPT.

Add system based on performance data

In this section we describe the movement of existing performance data into an SPT plan. Use the option to work on existing performance data from PM/400, PM/AIX, or PM/Linux. This is where you also select a WLE model based on new workloads that were previously created.

1. Click the **Based on existing performance data** radio button and click **Next** (Figure B-25).

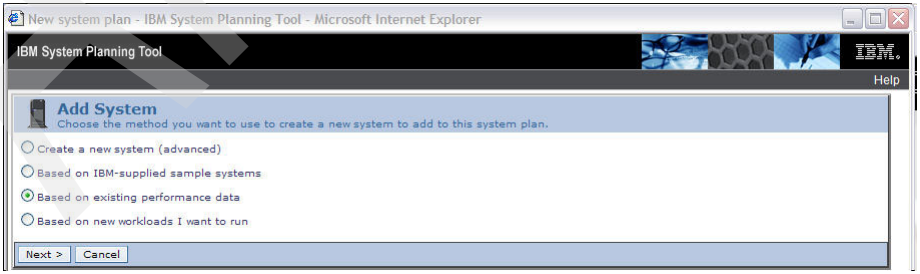

Figure B-25 Select WLE - Based on existing performance data

2. The Add System Based on Existing Workloads window opens and offers you an option to browse to your saved workload, as shown in Figure B-26. Click the **Browse** button and select your saved WLE file.

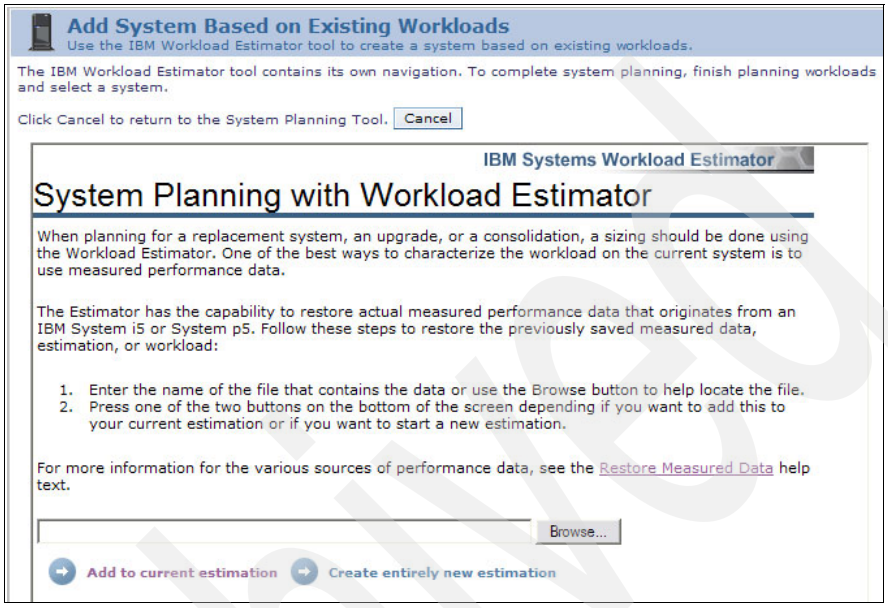

Figure B-26 Selecting WLE model in SPT

3. In our example shown in Figure B-27 we see the wle .xml file. We selected this file and clicked **Open**.

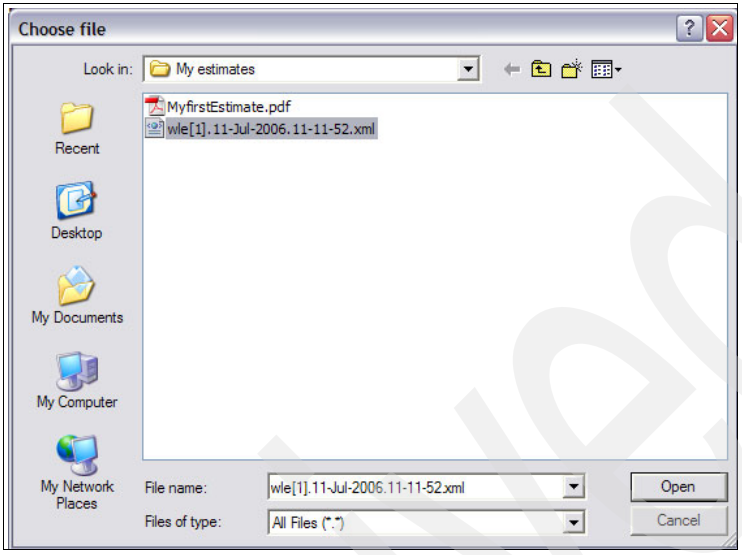

Figure B-27 Browse to existing WLE file

4. The file you selected is now loaded into the Add System based on Existing Workloads panel. Click Add to current estimate or Create an entirely new estimate (Figure B-28).

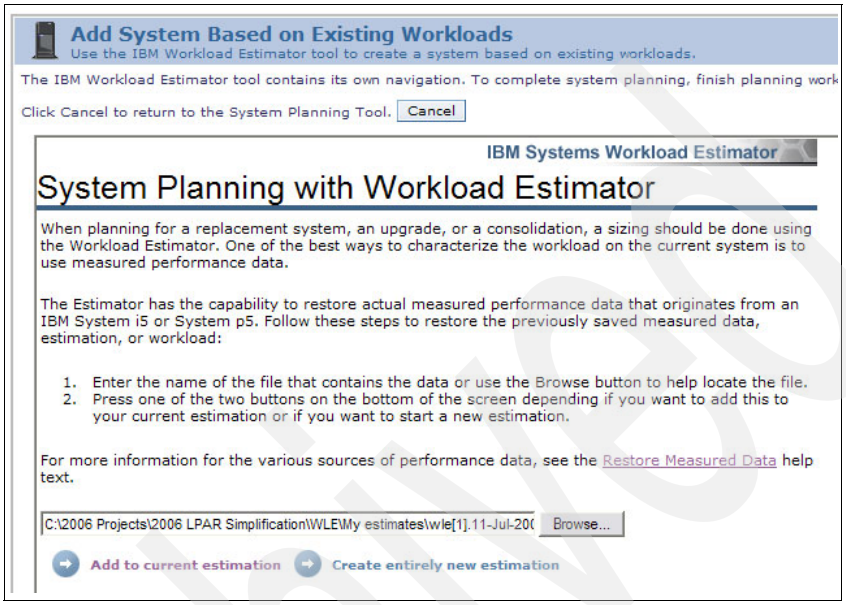

Figure B-28 Add WLE to to current estimate

The next step takes you through the same process that starts in “Add system based on new workloads” on page 139. Follow the steps there and then progress to system planning.

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 148. Note that some of the documents referenced here may be available in softcopy only.

- *IBM @server i5 and iSeries System Handbook: IBM i5/OS Version 5 Release 3 October 2004*, GA19-5486

  http://www.redbooks.ibm.com/redpieces/abstracts/ga195486.html?Open

- *IBM @server i5, iSeries, and AS/400e System Builder IBM i5/OS Version 5 Release 3 - October 2005*, SG24-2155

- *IBM iSeries Migration: A Guide to Upgrades and Migrations,* SG24-6055-01

  http://www.redbooks.ibm.com/redbooks/pdfs/sg246055.pdf

- *IBM @server iSeries Migration: A Guide to Upgrades and Migrations to POWER Technology*, SG24-7200

  http://www.redbooks.ibm.com/redbooks/pdfs/sg247200.pdf

- *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000

- *High-speed Link Loop Architecture for the IBM @server iSeries Server: OS/400 Version 5 Release 2*, REDP-3652

Online resources

These Web sites and URLs are also relevant as further information sources:

- eServer Hardware Info Center

  http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/index.htm

- iSeries Information Center

  http://publib.boulder.ibm.com/iseries/

How to get IBM Redbooks

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

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

LPAR Simplification Tools Handbook

Redbooks

Review the new System Planning Tool for IBM System i and IBM System p

Understand how to use the LPAR deployment wizard

See how easy it is to design simple logically partitioned systems

In this IBM Redbook we summarize the new functions designed to simplify the process of planning for and deploying logical partitions (LPARs) on the System i and System p models. The new LPAR Simplification functions include the new System Planning Tool (SPT) and the new Deploy System Plan wizard on the Hardware Management Console (HMC). This book focuses on System i.

The System Planning Tool is the replacement for the Logical Partition Validation Tool (LVT). It provides the configuration mechanism for designing LPARs on System i and System p systems. This document covers the installation and use of SPT.

The new LPAR deployment wizard is a function of the HMC and provides the capability to take the configuration produced in SPT and use it to create an LPAR configuration on the HMC managing the system. This significantly simplifies and reduces the work effort to create LPARs. The design work is done once in SPT, whereas the deployment or redeployment can happen many times.

This book provides guidance information for IBM and IBM Business Partners who help customers plan, order, and deploy systems.

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

SG24-7231-00 ISBN 0738494763