

***Step-by-step Virtualization
configuration on an IBM eServer p5
server, part 1:
configure VIO Server***

**by:
Nam Keung
Terry Wang
Balaji Atyam**

IBM eServer Solutions Enablement

May 2005

Table of contents

Abstract	1
Introduction.....	1
Setup scenario	1
POWER5 machine type and hardware structure	2
Software	2
VIO Server and client partitions configuration.....	2
Steps to configure the p570 scenario	4
Create VIO Server partition	4
Conclusion.....	14
References	15
About the authors	16
Trademarks	17

Abstract

The IBM® Virtual I/O Server (VIO Server), announced in July 2004, facilitates the sharing of physical resources between partitions on the IBM eServer™ p5 family of servers. The innovative VIO Server technology is driven by the IBM POWER5™ Virtualization Engine™. This paper offers a step-by-step (screen-by-screen) explanation of the process for installing a VIO Server on an eServer p5 model 9117-570. This paper serves as a guide for installing a VIO Server on all models of the eServer p5 family of servers.

[NOTE: A step-by-step illustration of the process for configuring a client partition is explained in Part 2 of this set of three white papers. The third paper in this series demonstrates how to configure the Virtual devices.]

Introduction

Virtual I/O Server and Partition Load Manager are part of the Advanced POWER™ Virtualization Engine hardware feature. Together with Partition Load Manager, the VIO Server provides:

- Micro-Partitioning™ support for up to 10 logical partitions to share one processor
- Virtual SCSI (VSCSI) disks that allow partitions to share physical storage adapters and devices
- Automated CPU and memory reconfiguration
- Real-time partition configuration and load statistics
- Support for dedicated and shared processor logical partitioning (LPAR) groups
- Support for manual provisioning of resources
- Virtual networking

The VIO Server is a single function appliance that resides in a partition to provide Virtual SCSI target and shared Ethernet adapter Virtual I/O function to client partitions. This partition is not intended to run applications or for general user logins. The VIO Server partition is created using the Hardware Management Console (HMC). The VIO Server installation media ships with eServer p5 servers and is installed via CD. It supports the following operating environments as client partitions: IBM AIX 5L™ Version 5.3, SUSE Linux™ Enterprise Server 9 for POWER, and Red Hat Enterprise Linux AS for POWER, Version 3, as VIO clients.

Setup scenario

The sample installation process discussed in this paper involves setting up a VIO Server and three client partitions using the POWER5 Virtualization Engine features that are available on the latest set of IBM POWER-based eServer p5 hardware, running under the AIX 5L operating system. The VIO Server will own a physical SCSI and a physical Ethernet controller. Because the physical SCSI controller on the designated hardware system has three physical hard disks attached to it, this scenario will configure each client partition on a separate physical hard disk.

POWER5 machine type and hardware structure

The eServer p5 server is a model 9117-570. As mentioned, the system has two physical SCSI controllers, one located on slot T14 and the other located on slot T12. Each SCSI controller has three physical disks attached to it. Two physical Ethernet controllers are on slots T6 and C3.

For this scenario, the VIO Server will use the SCSI controller on slot T14 and the Ethernet controller on slot T6. (Figure 1)

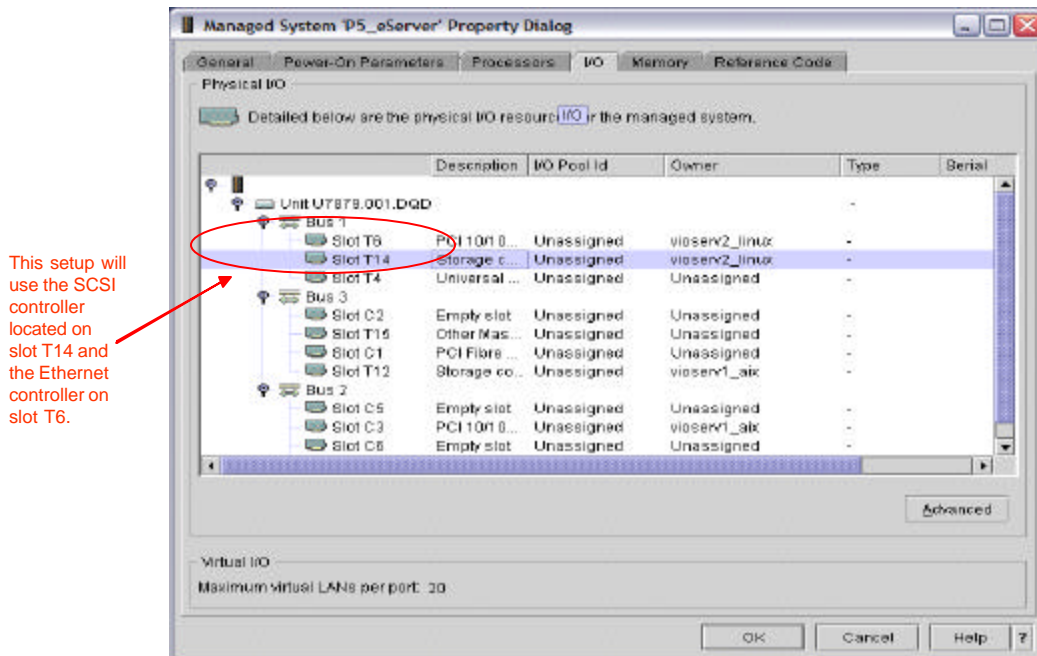


Figure 1: Slots T6 and T14 are owned by the VIO Server

Software

The following versions of software were installed on the eServer p5 server used in this scenario:

- VIO Server version 1.1 and latest fix pack 4, which includes IY58231:
<http://techsupport.services.ibm.com/server/virtualization/vios/download>
- HMC version 4 release 2 and PTF MH00138:
<http://techsupport.services.ibm.com/server/hmc/power5/fixes/v4r2.html>
- POWER5 microcode level:
<http://techsupport.services.ibm.com/server/mdownload/download.html>

VIO Server and client partitions configuration

For this scenario, the VIO Server, named `vioserv1_linux`, were installed on `hdisk2`. The three client partitions, named `svtlnx1`, `svtlnx2`, and `svtlnx3`, were installed on `hdisk2`, `hdisk1`, and `hdisk0`, respectively.

The `svtlnx1` client partition and the VIO Server shared the same `hdisk2` physical disk.

The VIO Server, `vioserv1_linux`, was configured with the following:

- One shared Ethernet adapter, using slot 5 and PVID 3
- Three Virtual SCSI adapters, using slot 10, 11, and 12
- Slot 10 on `vioserv1_linux` mapped to a logical volume on `hdisk2`
- Slot 11 mapped to `hdisk1`
- Slot 12 mapped to `hdisk0`

Each of the three client partitions were configured with a Virtual Ethernet adapter and a Virtual SCSI adapter. The Virtual Ethernet adapters also use slot 5 and PVID 3. The Virtual SCSI adapters use slot 10 locally, but the remote slot on `vioserv1_linux` was set as follows:

- `svtlnx1` --> slot 10 on `vioserv1_linux`
- `svtlnx2` --> slot 11 on `vioserv1_linux`
- `svtlnx3` --> slot 12 on `vioserv1_linux`

PV name	VG name	LV name	Virtual SCSI device	Server slot	Client partition ID	Client slot
hdisk0	rootvg_client1	svtlnx3_lv	vhost2	12	9	10
hdisk1	rootvg_client2	svtlnx2_lv	vhost1	11	8	10
hdisk2	rootvg	svtlnx1_lv	vhost0	10	7	10

You can review the figure below to see how the configuration was set. (Figure 2)

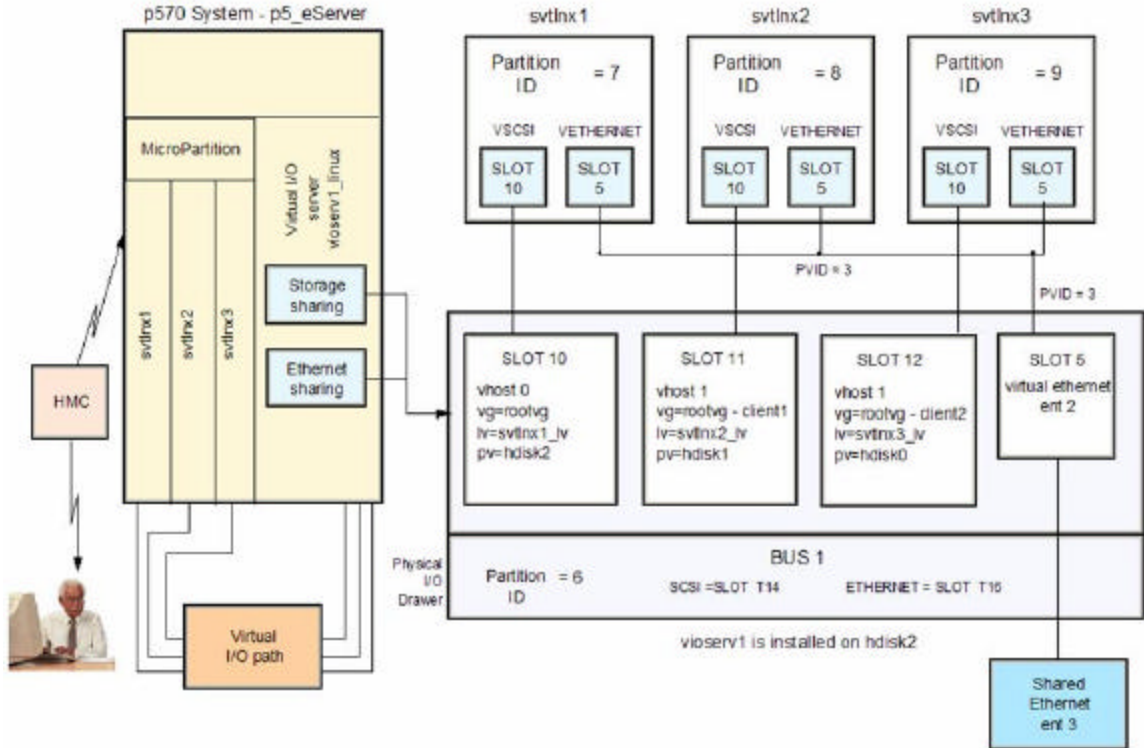


Figure 2: Configuration diagram of the VIO Server and client partitions

Steps to configure the p570 scenario

Now we will discuss the steps necessary to install VIO Server V1.1 and its fix pack, and configure the scenario. These steps are based on the hardware and configuration requirements mentioned above. We will go through each of these steps in detail:

- Part 1: Create VIO Server partition
- Part 2: Create client partitions
 - Modify the VIO Server profile
- Part 3: Configure the virtual devices in the VIO Server
 - Update fix pack
 - Install the operating system on the client partitions

Create VIO Server partition

The following processes create the VIO Server partition with a normal **run** and **install** profile:

1. From the Web-based System Manager (WebSM) navigation bar, expand **Management Environment**. Then, expand the managed system—in our case the managed system was named **svthmc1.austin.ibm.com**. Then, expand **Server and Partition**. Lastly, select **Server Management**.
2. Now, from the detail panel on the right (entitled **Server and Partition: Server Management**), expand the managed system. Then, right-click on **Partitions**.
3. From the pop-up menu, select **Create > Logical Partition**.

Note: Steps 1, 2, and 3 can be seen in the figure below. (Figure 3)

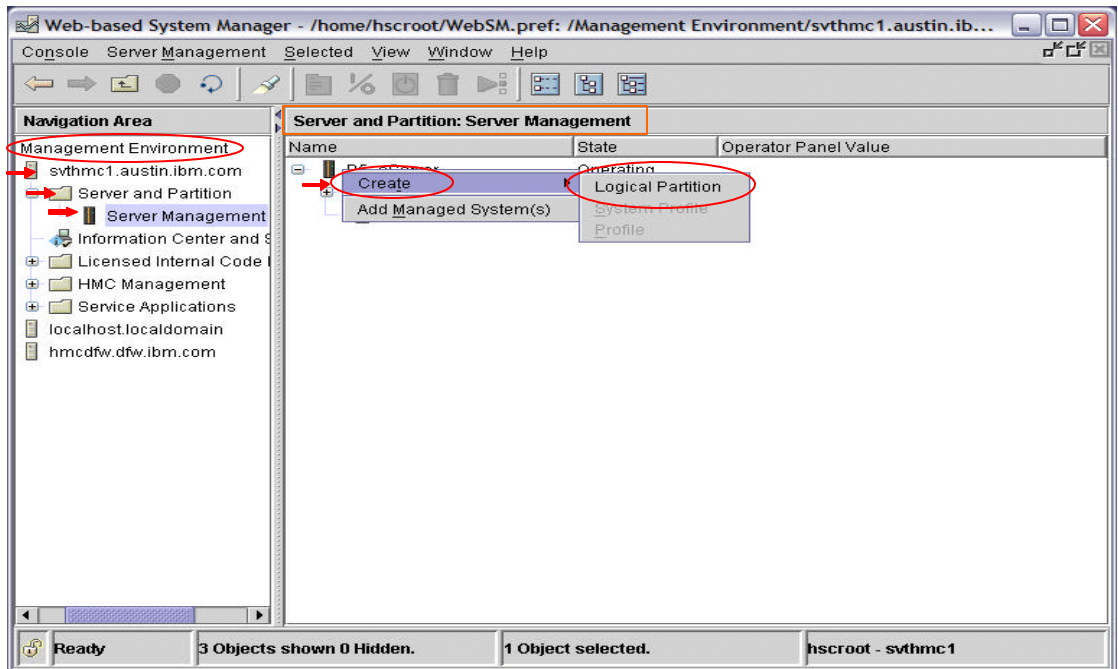


Figure 3: Configuration steps A1, A2, and A3 are illustrated on this screen capture

- From the **Create Logical Partition Wizard** panel, type `vioserv1_linux` as the **Partition name** and check **Virtual I/O Server**. Then click **Next**. (Figure 4)

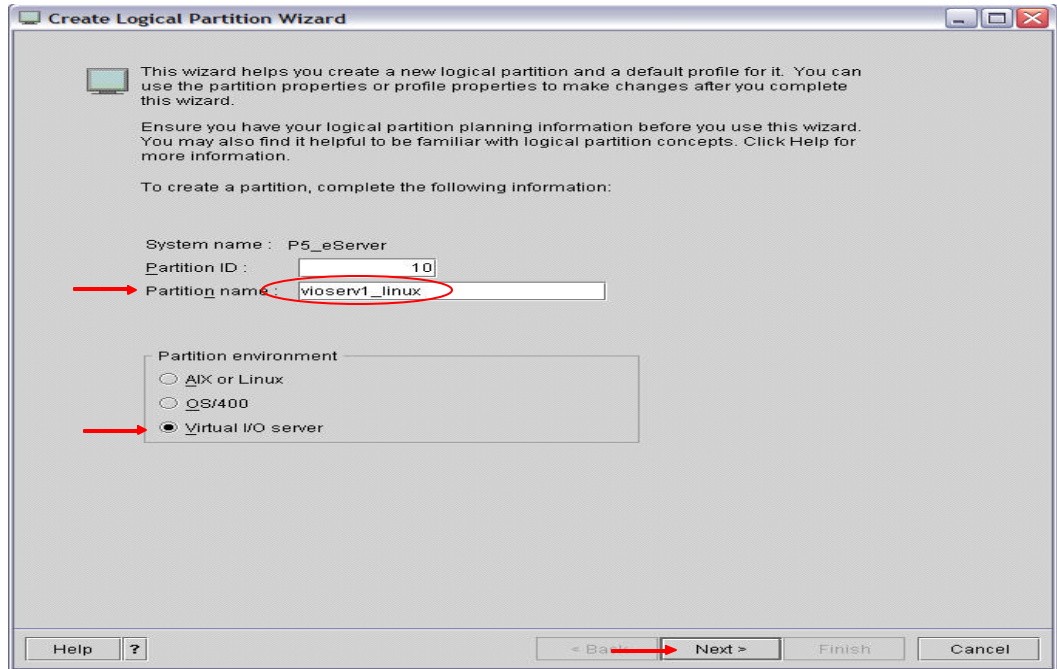


Figure 4: Configuration step A4 is illustrated on this screen capture

- From the **Workload Management Groups** panel, check **No**. Then click **Next**. (Figure 5)

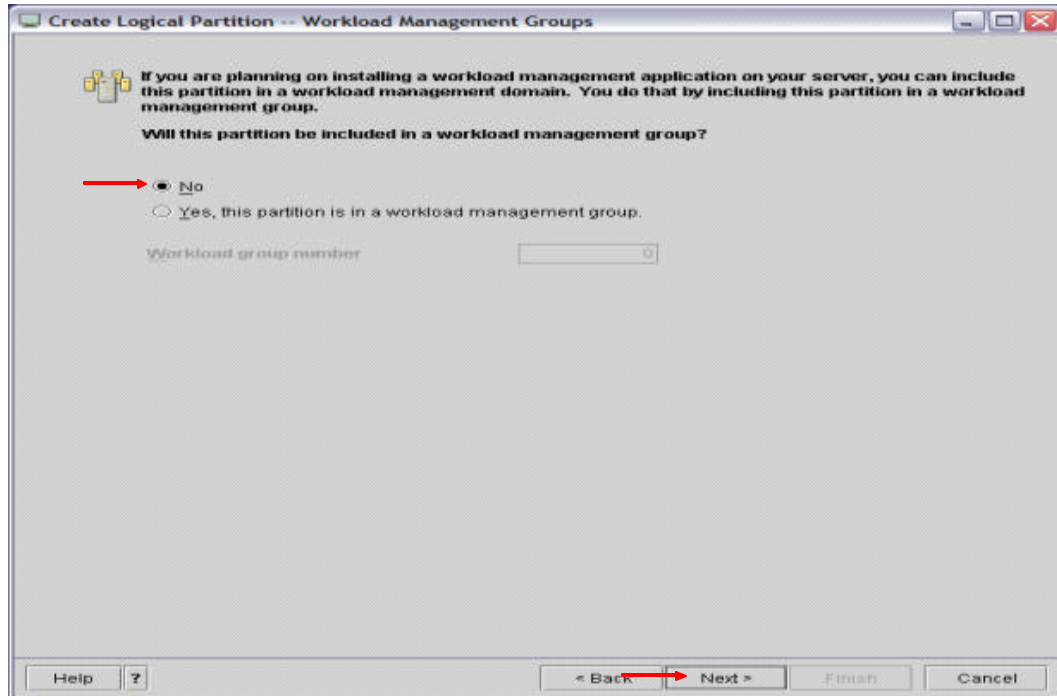


Figure 5: Configuration step A5 is illustrated on this screen capture

- From the **Create Logical Partition Profile** panel, type in the **Profile name** (for example: Run_VIO_Server_profile). Then click **Next**. (Figure 6)

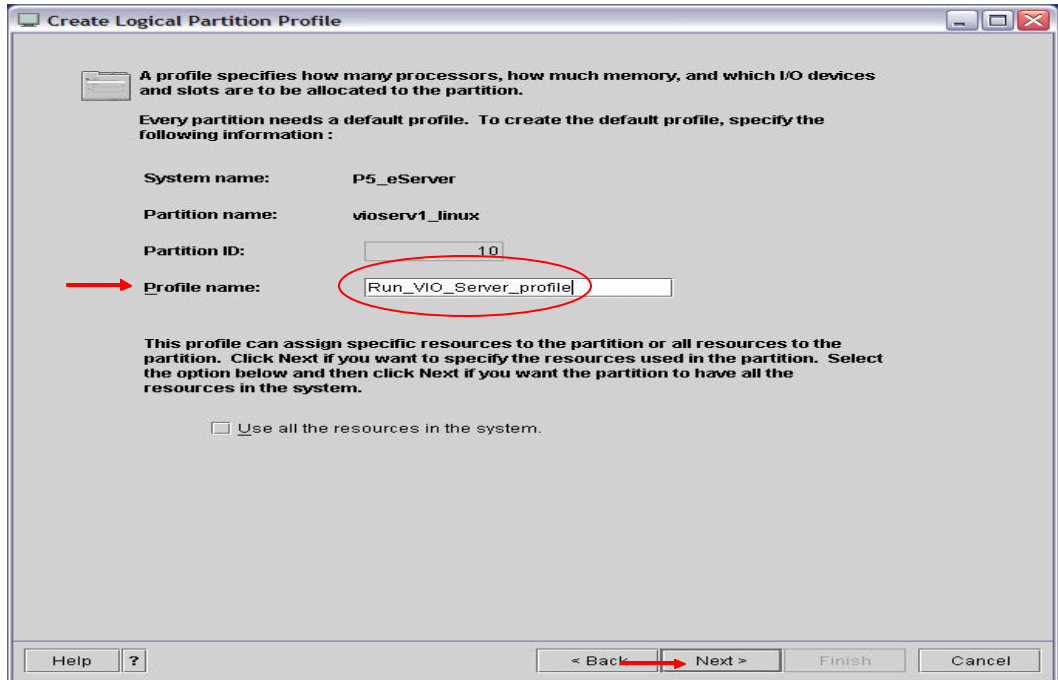


Figure 6: Configuration step A6 is illustrated on this screen capture

- From the **Memory** panel, set the minimum, desired, and maximum memory allocations (for example: 128 MB, 512 MB, and 1 GB, respectively). Then, click **Next**. (Figure 7)

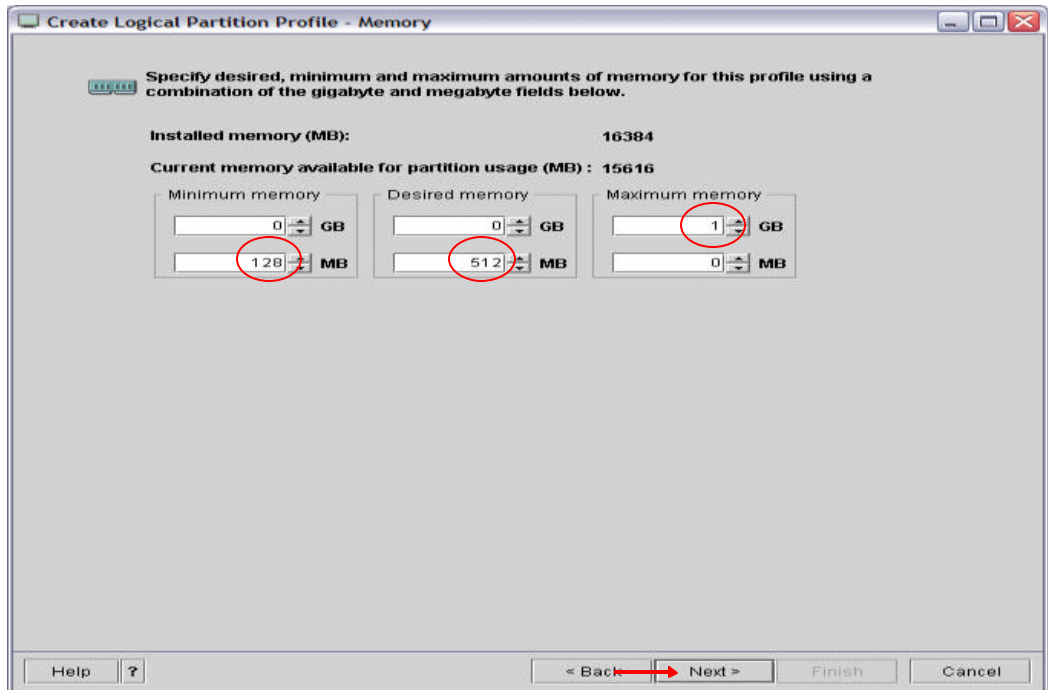


Figure 7: Configuration step A7 is illustrated on this screen capture

- From the **Processors** panel, check **Shared**. Then click **Next**. (Figure 8)

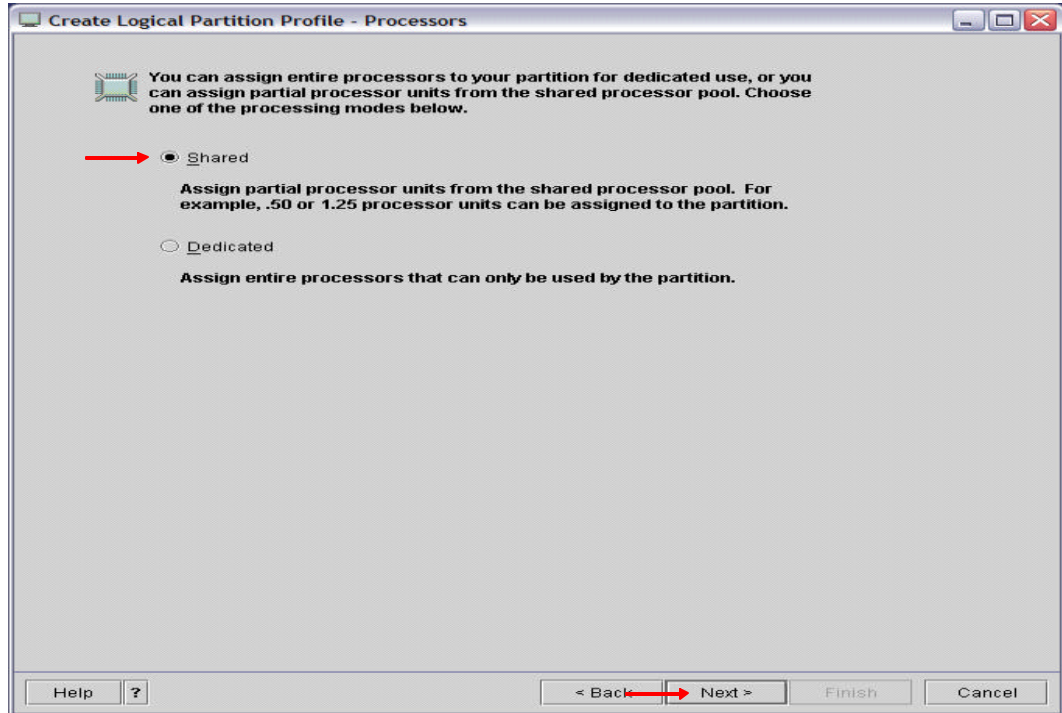


Figure 8: Configuration step A8 is illustrated on this screen capture

- From the **Processing Settings** panel, set the desired, minimum, and maximum processing units (for example: 0.5, 0.1, and 0.8, respectively). Then click **Next**. (Figure 9)

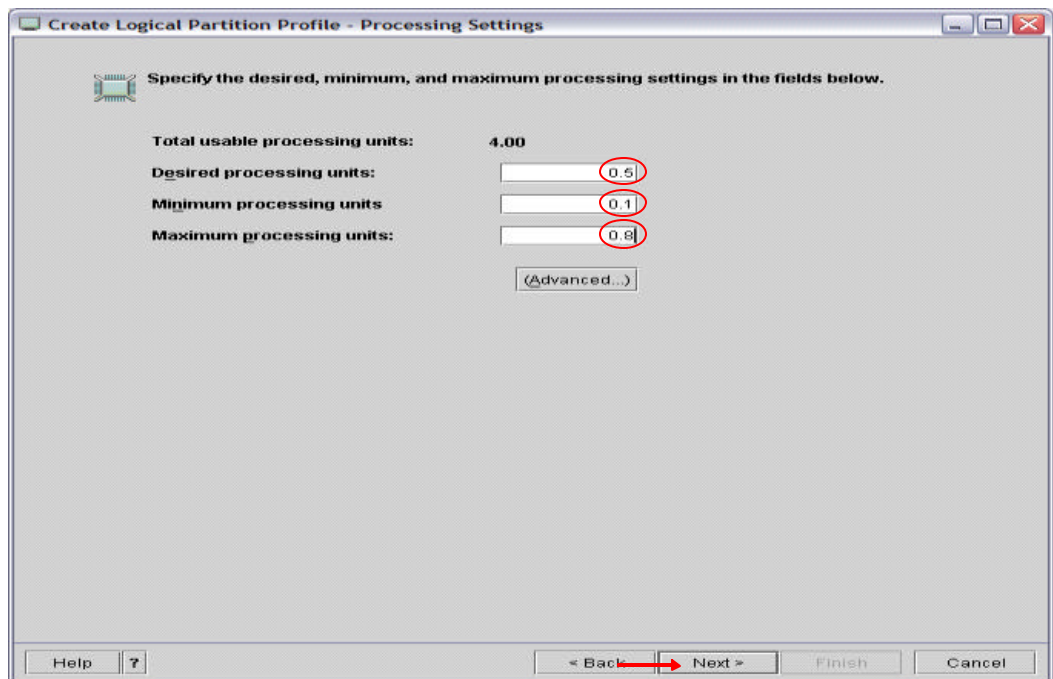


Figure 9: Configuration step A9 is illustrated on this screen capture

- From the **I/O** panel, select an Ethernet and a SCSI controller and add them as required. For example, add **Slot T6** as the Ethernet controller, and add **Slot T14** as the Storage controller. Then click **Next**. (Figure 10 and 11)

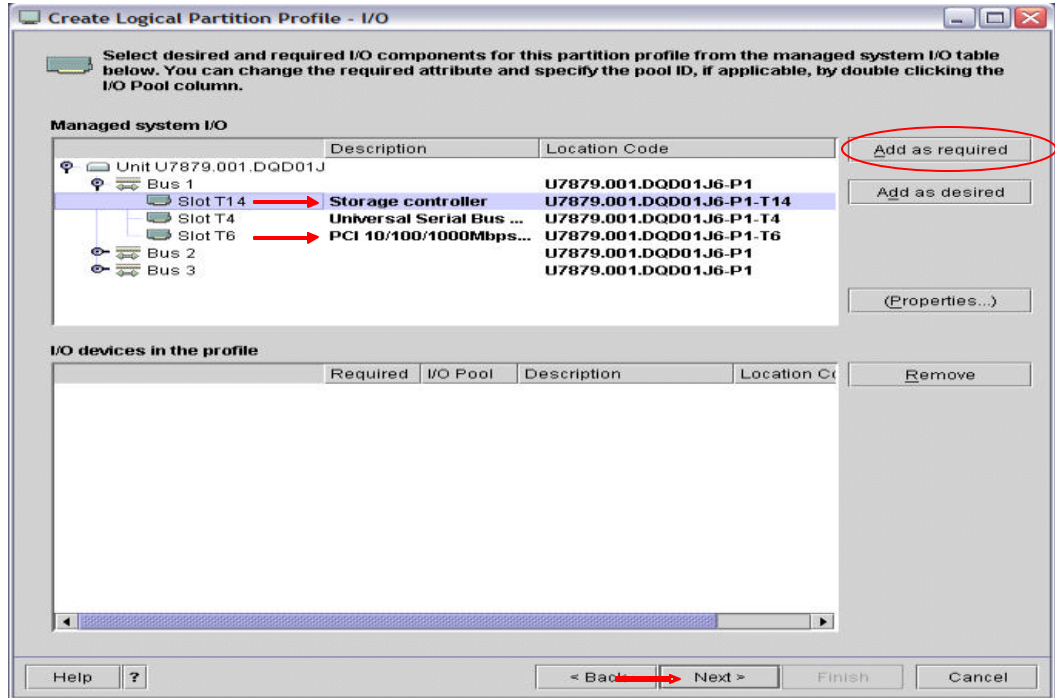


Figure 10: Configuration step A1 is illustrated on this screen capture

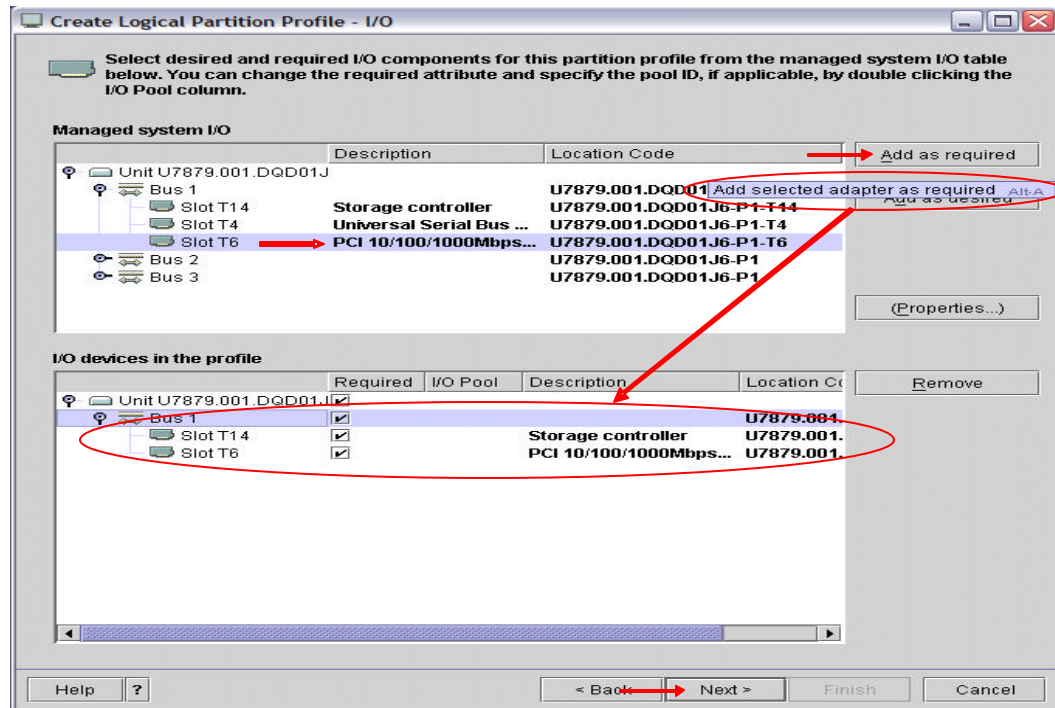


Figure 11: A continuation of configuration step A10 is illustrated on this screen capture

11. From the **I/O Pools** panel, click **Next**. (Figure 12)

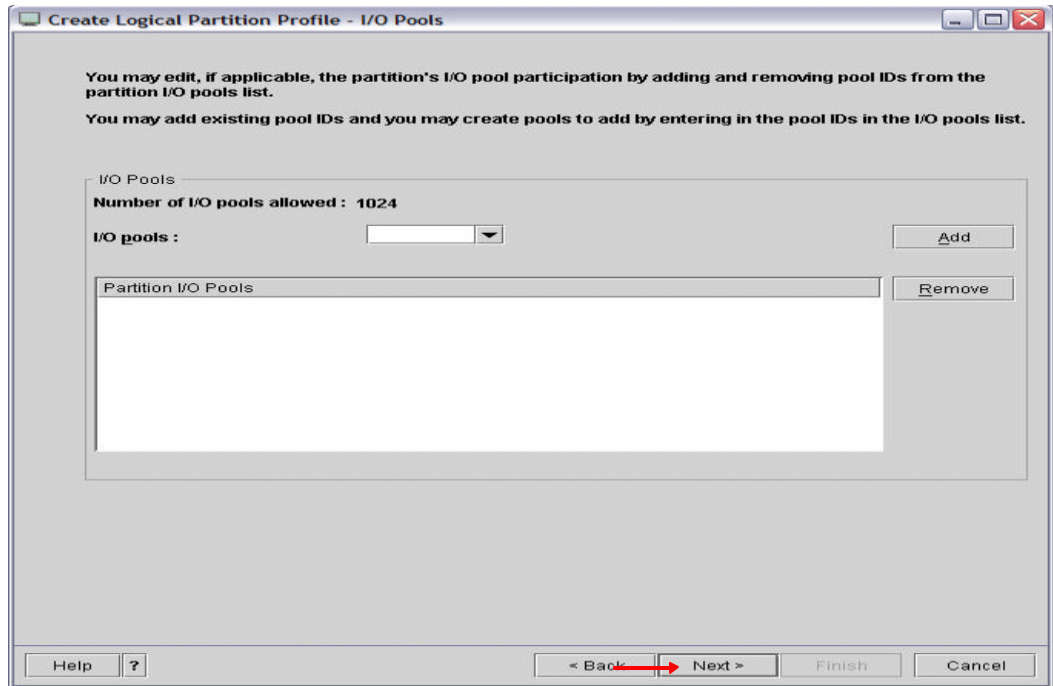


Figure 12: Configuration step A11 is illustrated on this screen capture

12. From the **Virtual I/O Adapters** panel, check **No**. Then click **Next**. (Figure 13)

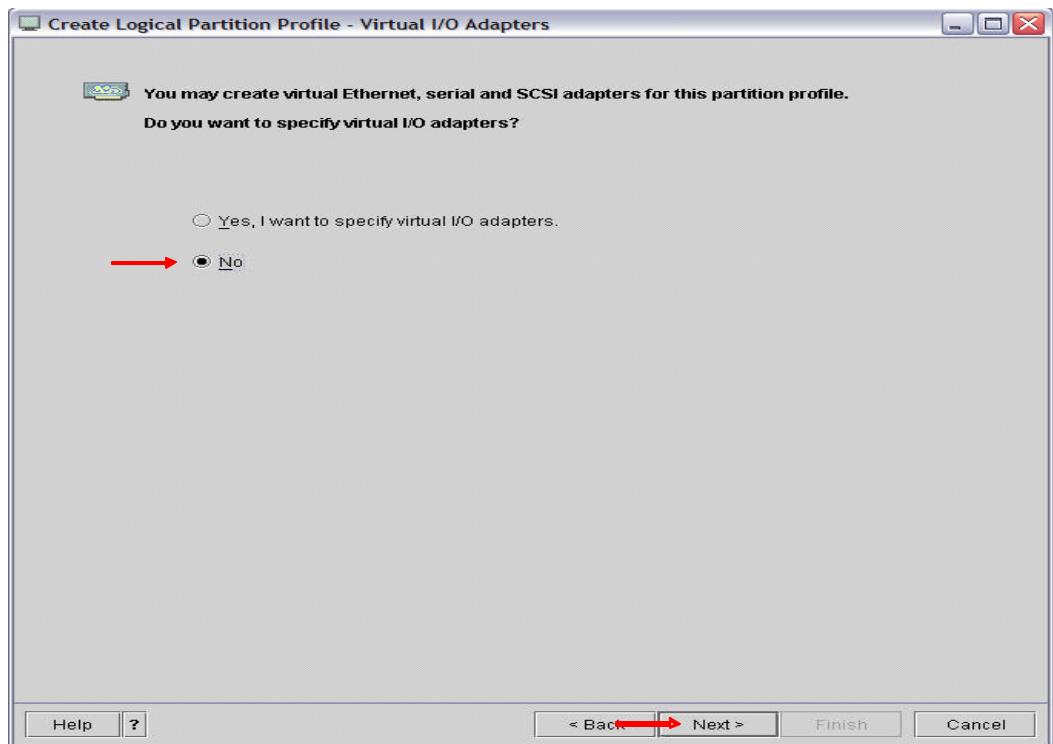


Figure 13: Configuration step A12 is illustrated on this screen capture

13. From the **Power Controlling Partitions** click **Next**. (Figure 14)

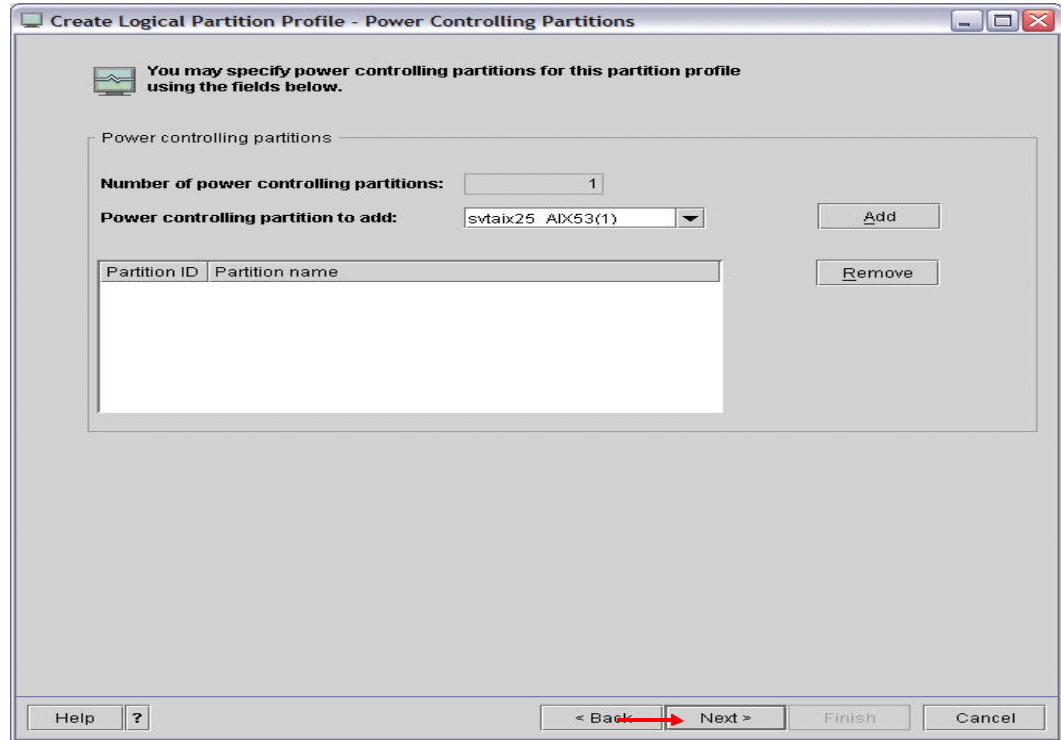


Figure 14: Configuration step A13 is illustrated on this screen capture

14. From the **Optional Settings** panel, in **Boot modes** check **Normal**. Then, click **Next**. (Figure 15)

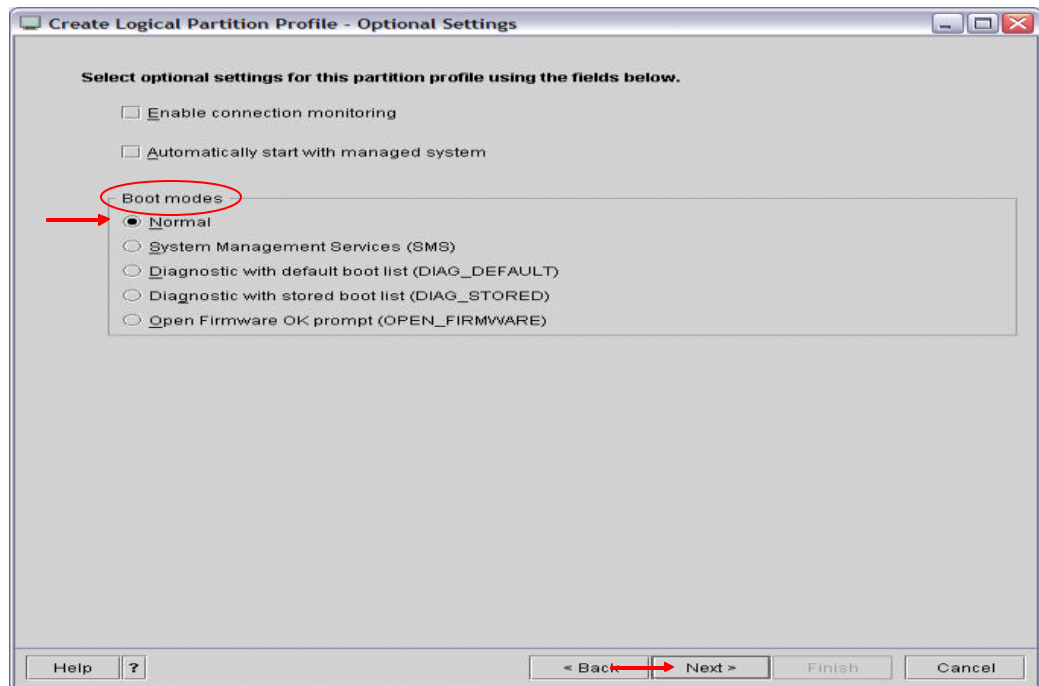


Figure 15: Configuration step A14 is illustrated on this screen capture

- From the **Profile Summary** panel, confirm the settings. Then, click **Finish**. If the partition is created successfully, it will add `vioserv1_linux` into the **Partition** with profile name `Run_VIO_Server_profile`. (Figure 16)

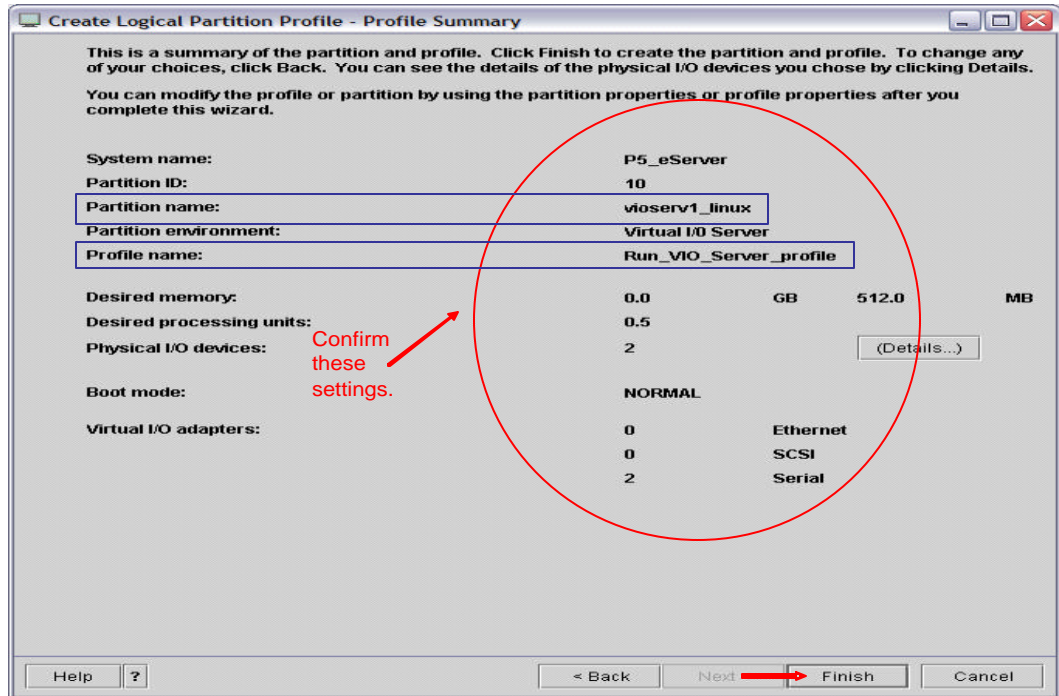


Figure 16: Configuration step A15 is illustrated on this screen capture

- Right-click `Run_VIO_Server_profile` under `vioserv1_linux`. Then select **Copy** from the pop-up menu. (Figure 17)

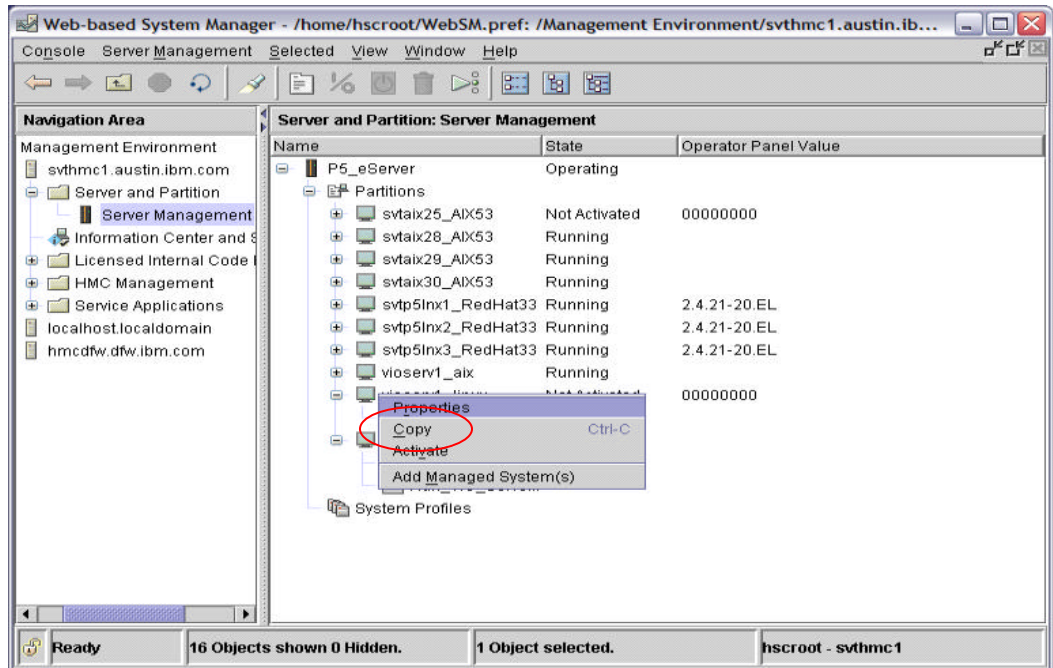


Figure 17: Configuration step A16 is illustrated on this screen capture

17. From the **Copy Profile** panel, type the install profile name (for example, `Install_VIO_Server_profile`) in the **New profile name** field. If successful, **Install_VIO_server_profile** will be refreshed under **vioserver1_linux**. (Figure 18)

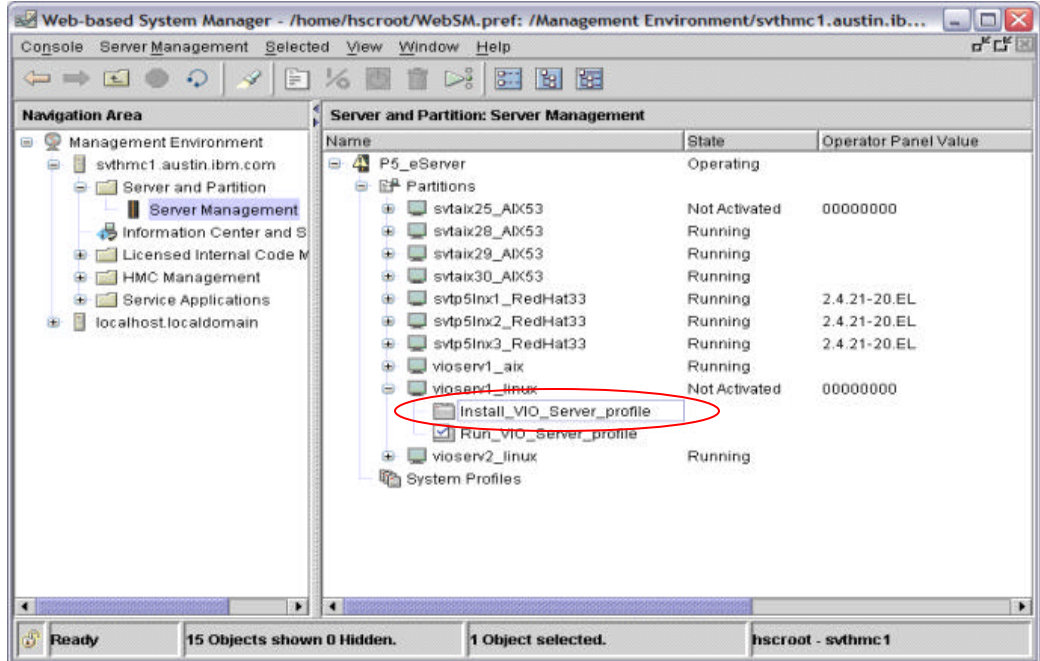


Figure 18: Configuration step A17 is illustrated on this screen capture

18. Right-click **Install_VIO_Server_profile** and select **Properties** to modify the install profile. After the **Properties** panel comes up, select the **Physical I/O** tab and add the CD/DVD controller to the desired list. For example: add **Slot T15, Other Mass Storage Controller**, as desired. (Figure 19 and 20)

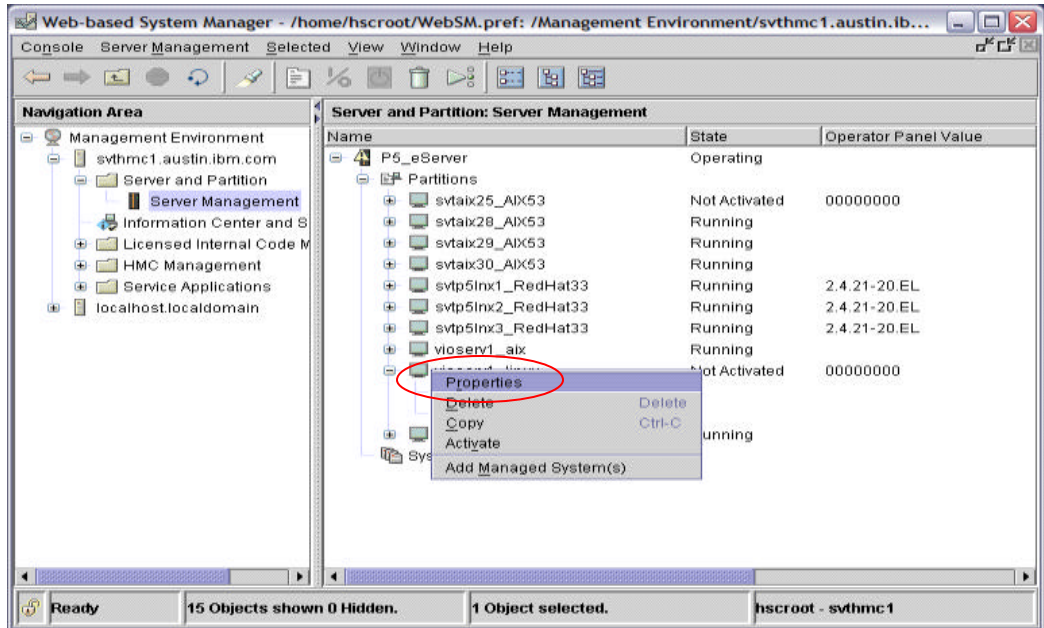


Figure 19: Configuration step A18 is illustrated on this screen capture

Part 1: Configure VIO Server

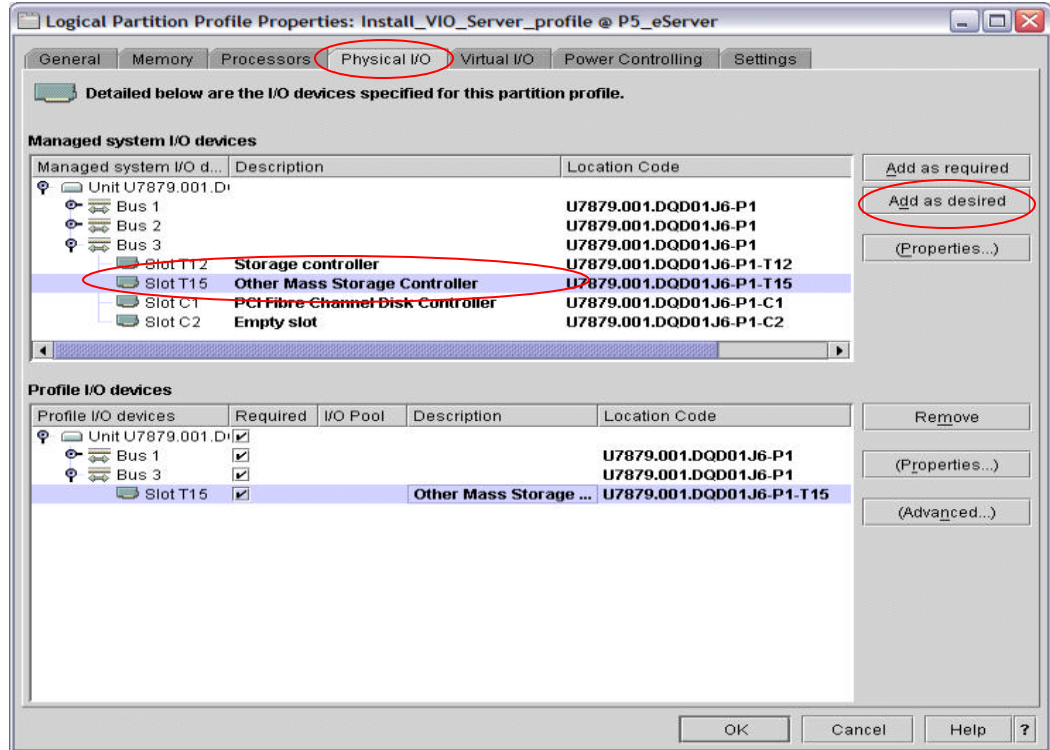


Figure 20: Continuation of configuration step A18 is illustrated on this screen capture

19. Select the **Settings** tab, check **System Management Services (SMS)** for the **Boot Modes**. Then click **OK**. (Figure 21)

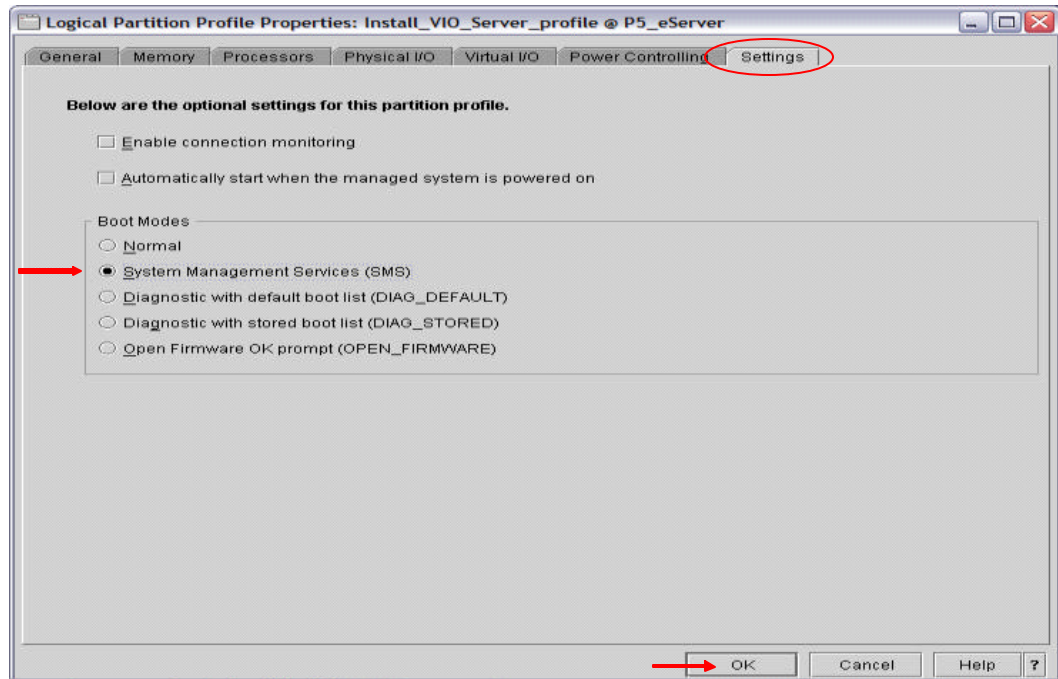


Figure 21: Configuration step A19 is illustrated on this screen capture

Now, the VIO Server partition with two profiles, **run** and **install**, are created successfully.

Conclusion

The IBM Virtual I/O Server facilitates the sharing of physical resources between partitions on an IBM eServer p5 server. This capability is available as the result of leading technology delivered with the POWER5 processor. The VIO Server offers many attractive features:

- Allows for the creation of partitions without requiring additional physical I/O resources.
- Enables the creation of more partitions than there are actual I/O slots or physical devices. Partitions may enjoy dedicated I/O, Virtual I/O, or both.
- Provides Virtual SCSI and shared Ethernet adapter function to client partitions.
- Helps maximize utilization of physical resources.

For these and many other reasons, more and more enterprises which run their aggressive IT missions on the eServer p5 family of servers are moving to virtual partitioning.

Through a screen-by-screen approach, this paper has provided the reader with a sense of the relative simplicity and speed with which a VIO Server can be implemented, along with a demonstration of setting up a number of client partitions to interact with it.

With the sophistication of the mainframe-like partitioning technology that is now available with POWER5, there is no reason to postpone embarking on the “virtual” path—allowing your enterprise to enjoy higher and more efficient use of its many server resources.

References

- White paper: *Implementing Virtualization on an IBM eServer p5 server (part 2): Configure client partitions*
ibm.com/servers/enable/site/peducation/abstracts/abs_4372
- White paper: *Implementing Virtualization on an IBM eServer p5 server (part 3): Configure Virtual devices*
ibm.com/servers/enable/site/peducation/abstracts/abs_4376
- IBM Redbook: *Introduction to Advanced POWER Virtualization on IBM p5 Servers, Introduction and basic configuration (SG24-7940)*
ibm.com/redbooks/abstracts/sg247940.html?Open
- IBM Redbook: *Server Consolidation on IBM pSeries Systems (SG24-6966)*
ibm.com/redbooks/abstracts/sg246966.html?Open
- White paper: *IBM eServer POWER5 Processors Virtual SCSI Throughput Analysis*, by Elizabeth Stahl, January 2005
ibm.com/eserver/pseries/hardware/whitepapers/virtual_scsi.pdf
- IBM Redbooks Technote: *Server Consolidation: A Comparison of Workload Management and Partitioning (TIPS0426)*:
ibm.com/redbooks/abstracts/tips0426.html?Open
- Advanced POWER Virtualization on eServer p5 Web site
ibm.com/servers/eserver/pseries/ondemand/ve/resources.html
- IBM Virtual I/O Server support center
<http://techsupport.services.ibm.com/server/virtualization/vios/download>

About the authors

Nam Keung is a senior technical consultant for IBM in Austin, Texas. He has worked in the area of AIX ISDN communications, AIX SOM/DSOM development, AIX multimedia development, Microsoft® Windows® NT clustering technology, and Java™ performance. His current assignment involves helping IBM Business Partners and solution providers in their efforts to port and deploy applications to the pSeries platform. He also consults in performance tuning and other educational needs for the pSeries platform.

Terry Wang has a master's degree in computer science from the University of Wisconsin, Madison. Terry joined IBM in 1988 working on AIX development. He is currently with the IBM System and Technology Group, solution enablement team. His most recent interest is in the areas of AIX virtualization technology and the application of this technology to solution provider applications.

Balaji V. Atyam is a senior technical consultant in the Systems and Technology Group. His responsibilities are porting, benchmarking, performance tuning, parallel programming, and technical consulting services to key solution providers in the area of high performance computing on IBM eServer platforms.

Trademarks

© IBM Corporation 1994-2005. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: IBM, eServer, pSeries, AIX, AIX 5L, POWER, POWER5, Micro-Partitioning, and Virtualization Engine.

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

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

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

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

Information is provided "AS IS" without warranty of any kind.