IBM

Integrated Virtual Ethernet Adapter

Technical Overview and Introduction

Unique and flexible up to 10 Gbps Ethernet industry standard virtualizable network solution

Hardware accelerated for improved throughput

Configurable connections between LPARs and the physical network without the Virtual I/O Server

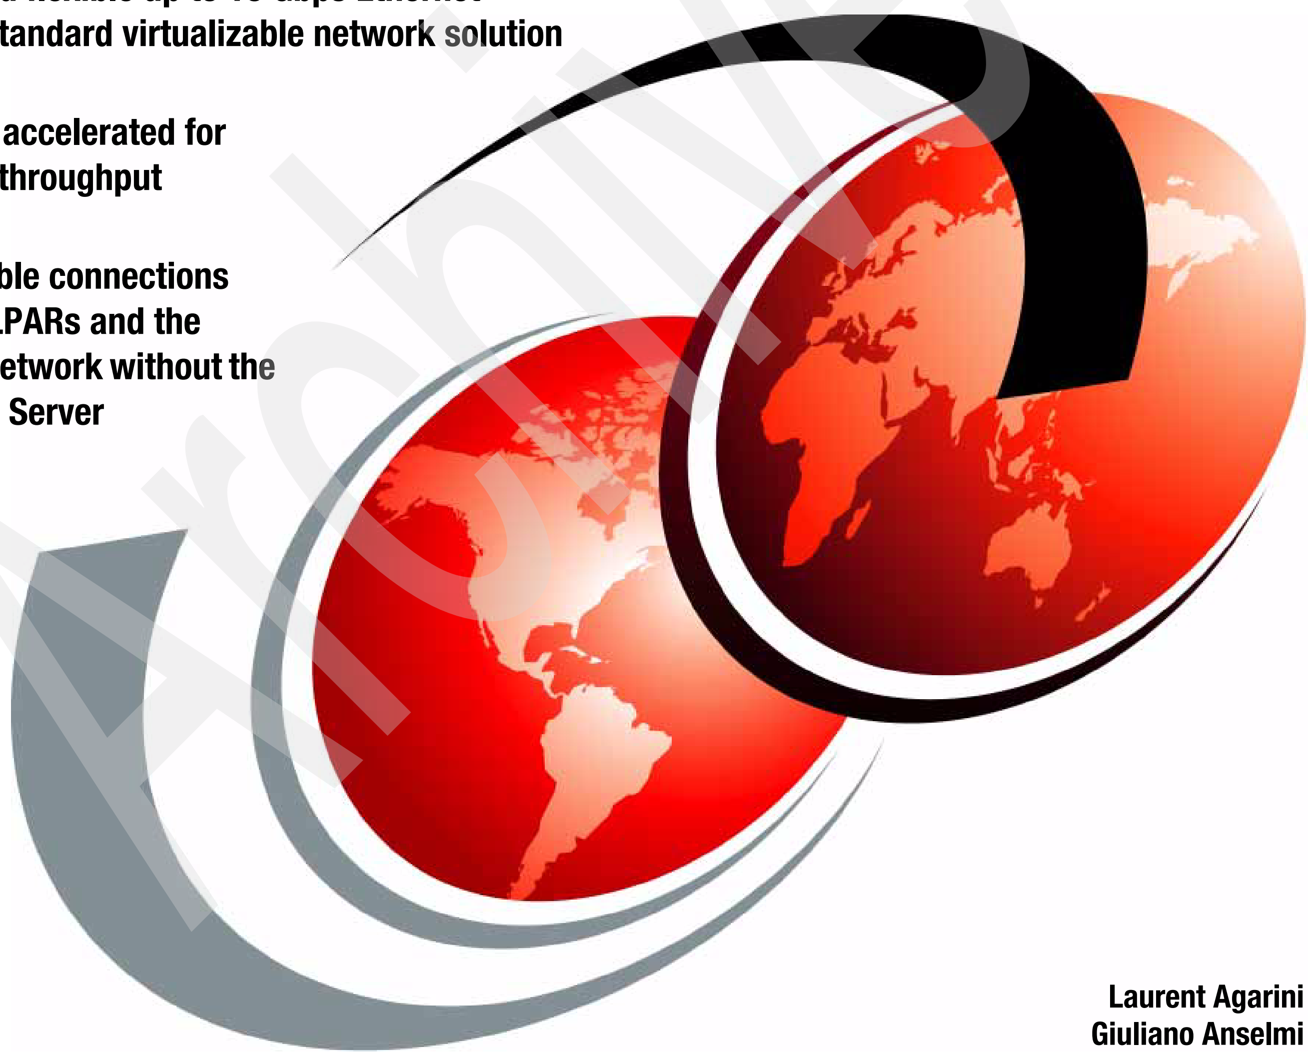

Laurent Agarini
Giuliano Anselmi

Redpaper

ibm.com/redbooks

International Technical Support Organization

Integrated Virtual Ethernet Adapter Technical Overview and Introduction

October 2007

REDP-4340-00

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

First Edition (October 2007)

This edition applies to IBM AIX Version 6, Release 1, the IBM System p 570, 560, 550, and 520, and the Hardware Management Console Version 7.

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

Redbooks (logo) ®	Micro-Partitioning™	Redbooks®
eServer™	POWER™	RS/6000®
pSeries®	POWER Hypervisor™	System i™
AIX 5L™	POWER5™	System p™
AIX®	POWER5+™	System p5™
IBM®	POWER6™	Tivoli®

The following terms are trademarks of other companies:

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

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

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

Preface

The introduction of the POWER6™ processor-based servers brought to market several advances in server design, performance, and function. One of these enhancements is the capability to virtualize Ethernet within a server. The feature that provides this function is named Integrated Virtual Ethernet adapter (IVE).

IVE is the integration of several technologies, including the Host Ethernet Adapter (HEA), advanced software, and updates to the hypervisor that provide integrated high-speed Ethernet adapter ports with hardware assisted virtualization capabilities.

IVE is a standard set of features offered on every IBM® System p™ POWER6 processor-based 570, 560, 550, and 520 server.

The IVE provides:

- Either two 10 Gbps Ethernet ports or four 1 Gbps ports or two 1 Gbps ports
- External network connectivity for LPARs using dedicated ports without the need of a Virtual I/O Server
- Industry standard hardware acceleration, loaded with flexible configuration possibilities
- The speed and performance of the GX+ bus, faster than PCI Express x16

This document is for technical professionals interested in understanding the many features of the IVE, and those who configure and support the servers that use it.

The team that wrote this paper

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

Laurent Agarini is a Solution Designer working for Global Storage Architecture (GSA) in la Gaude, France, and teaches IBM classes worldwide on AIX® and virtualization on System p. Laurent has been a team leader in the France command center for four years. He managed as an administrator the Tivoli® Standard Architecture (TSA) laboratory, implementing several platform operating systems, network and POWER™ technology for three years. Laurent has worked for IBM for six years.

Giuliano Anselmi is a Certified System p Technical Pre-Sale Support Specialist. He has worked in the Field Technical Sales Support group based in Rome, Italy, for four years. For seven years, he was an IBM eServer™ pSeries® Systems Product Engineer, supporting Web Server Sales Organization in EMEA, IBM Sales, IBM Business Partners, Technical Support Organizations, and IBM Dublin eServer Manufacturing. Giuliano has worked for IBM for 14 years, devoting himself to RS/6000® and pSeries systems with his in-depth knowledge of the related hardware and solutions.

The team that created this publication was managed by:

Scott Vetter, PMP
IBM US

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

Terry Brennan, Omar Cardona, Herman Dierks, Linda Flanders, Ron Gonzalez, Kiet H Lam, Shawn Lambeth, Thoi Nguyen, Veena Patwari, Jorge Rafa Nogueras, Amartey Pearson, Carolyn Scherrer, Jonathan Van Niewall, Colin Verrilli, Thomas Weaver
IBM US

Bruno Blanchard
IBM France

Marcus Eder
IBM Germany

Rita Ceresi, Carlo Costantini, Federico Vagnini
IBM Italy

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

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

1

Integrated Virtual Ethernet adapter overview

The POWER6 processor-based servers extend the virtualization technologies introduced in POWER5™ by offering the Integrated Virtual Ethernet adapter (IVE).

IVE, the name given to the collection of hardware (including the Host Ethernet Adapter (HEA)), software, and hypervisor provides integrated high-speed Ethernet adapter ports with hardware-assisted virtualization capabilities. It is a standard set of features offered on every IBM System p POWER6 processor-based 570, 560, 550, and 520 server.

The IVE developed to meet general market requirements for better performance and better virtualization for Ethernet. It offers:

- Either two 10 Gbps Ethernet ports or four 1 Gbps ports or two 1 Gbps ports
- External network connectivity for LPARs using dedicated ports without the need of a Virtual I/O Server
- Industry standard hardware acceleration, loaded with flexible configuration possibilities
- The speed and performance of the GX+ bus, faster than PCI Express x16

Note: Throughout this publication, when referring to a component of the IVE feature, such as the HMC dialogs, IVE is referred to as HEA.

1.1 Introduction

The Integrated Virtual Ethernet adapter is a standard feature of every POWER6 processor-based 570, 560, 550, and 520 server. You can select from different offerings according to the specific IBM System p server.

The IVE consists of a physical Ethernet adapter that is connected directly to the GX+ bus of a POWER6 processor-based server instead of being connected to a PCIe or PCI-X bus, either as an optional or integrated PCI adapter. This provides IVE with the high throughput and low latency of a bus imbedded in the I/O controller. IVE also includes special hardware features that provide logical Ethernet adapters. These adapters can communicate directly to logical partitions (LPARs), reducing the interaction with the POWER Hypervisor™ (PHYP).

In addition to 10 Gbps speeds, the IVE can provide familiar 1 Gbps Ethernet connectivity common on POWER5 and POWER5+™ processor-based servers.

Prior to IVE, virtual Ethernet provided a connection between LPARs. The use of an SEA and the Virtual I/O Server allowed connection to an external network. The IVE replaces the need for both the virtual Ethernet and the SEA. It provides most of the function of each.

Therefore, this eliminates the need to move packets (using virtual Ethernet) between partitions and then through a shared Ethernet adapter (SEA) to an Ethernet port. LPARs can share IVE ports with improved performance.

Figure 1-1 shows a Virtual I/O Server configuration on the left, and IVE on the right.

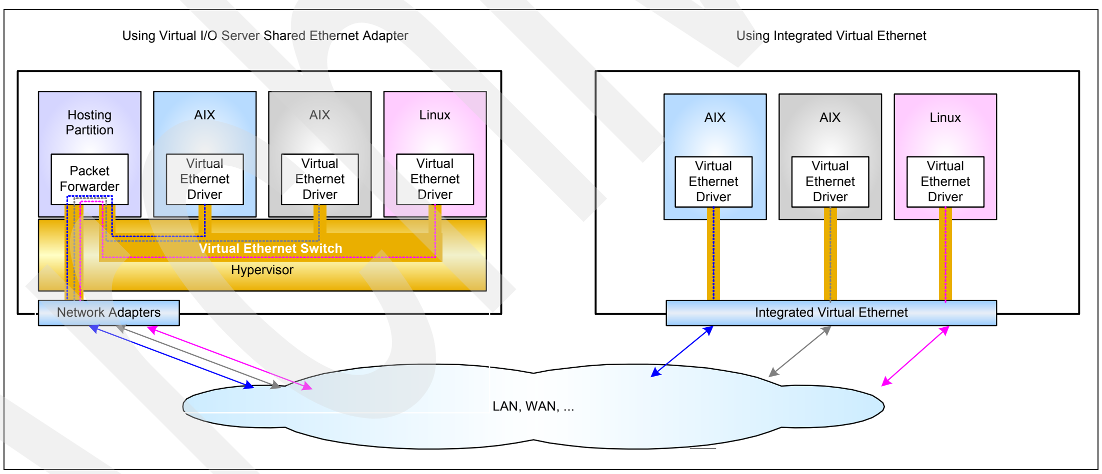

Figure 1-1 Virtual I/O Server Shared Ethernet Adapter compared to Integrated Virtual Ethernet

The IVE feature offers a choice of two or four Ethernet ports running at 1 Gbps and two ports running at 10 Gbps depending on the feature selected. In the case of an IBM System p 570 server, clients that are using 1 Gbps connection bandwidth in their IT infrastructure could move up to 10 Gbps infrastructure by adding a new System p 570 enclosure with the Integrated 2-port 10 Gbps virtual Ethernet.

A MES upgrade from one IVE feature code to another for single-CEC systems is not supported at the time of writing. In the case of a one or two-CEC 9117-MMA configuration, an MES of an additional CEC may contain any available IVE feature code as part of the updated configuration.

1.2 Advanced hardware features

IVE has several inherent features that provide configuration flexibility. This flexibility allows you to customize key network attributes.

Virtualized logical ports are described in 1.3.1, "Virtualized logical ports" on page 3, and Multiple Core Scaling (MCS) is described in "Multiple Core Scaling" on page 24. Together they provide scaling efficiencies.

IVE has traditional offload technologies such as large send (TCP segmentation offload), checksum offload, and Jumbo Frames.

Other IVE advanced hardware features include the following:

- TCP and UDP with IPv4 and IPv6
- Network/Protocol Management and RAS features
- Key advanced TCP features, such as large send offload, checksum offload, multiple receive queues, and low latency send and receive

iSCSI support is provided through the AIX and Linux® iSCSI software device driver.

1.3 Virtualization features

IVE extends the virtualization capabilities of a POWER6 processor-based system by providing:

- Internal network communication between logical partitions reducing hypervisor intervention using virtualized ports.
- Multiple Core Scaling that benefits highly parallel workloads.
- External network physical connection without the Virtual I/O server. Also, the hypervisor and SEA interaction is eliminated for data flow.

1.3.1 Virtualized logical ports

The Integrated Virtual Ethernet design comprises virtualization of a network port that can connect to external networks and share this connectivity to all virtualized logical ports without using the Virtual I/O Server, with less complexity, thus contributing to a better system optimization and better overall system performance.

According to the selected feature code, the IVE offers different configuration options, but from a logical point of view, an understanding of IVE can be simplified with a comparison to common hardware.

The logical port that is assigned is also named Logical Host Ethernet Adapter port, or LHEA port.

Figure 1-2 shows one of the available IVE feature codes offering up to 32 logical ports.

Logical ports have the following characteristics:

- Logical ports are grouped into a port group and any port group offers 16 logical ports, with one or two physical ports in the port group.
- An IVE has layer 2 switches, one per physical port.
- Each logical port can be assigned to any LPAR, and LPARs can have one logical port per physical port.
- Every logical port has a unique MAC address, and they are all stored in the VPD chip that is located on the IVE card (see Figure 1-3 on page 5 and Figure 1-5 on page 6).

> **Note:** If IVE is replaced, the new feature will have a new set of MAC addresses.

- A logical port belongs only to one physical port and one virtual switch, and there may be two logical ports in the same port group that do not share a virtual switch if located on different physical ports (as in the case with the 5639 feature code of the IBM System p 570).
- From a logical point of view, any logical port could be the Ethernet interface of a client device.
- LPARs that have logical ports associated with the same physical port (and therefore the same virtual layer 2 switch) and can communicate without the need of external hardware.

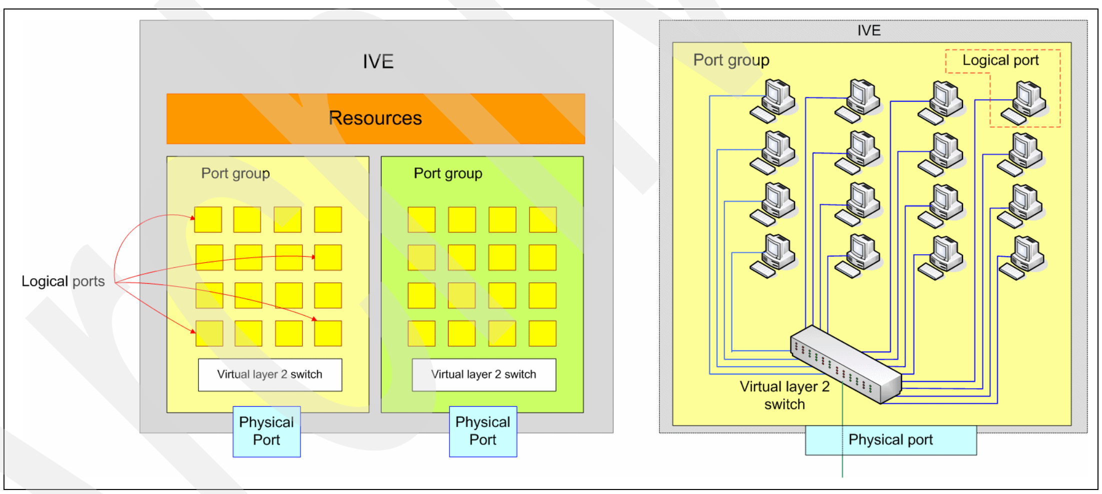

Figure 1-2 Example of an IVE port group with one physical port compared to physical network topology

1.3.2 Physical ports and system integration

The following sections discuss the physical ports and the features available at the time of writing on IBM System p product offerings.

System p 570

The System p 570 is a modular-built system made up from one to four enclosures. Each enclosure is four EIA[1] units tall, and each enclosure must have one IVE feature code selected as part of the base set of orderable options. Each enclosure can have a unique Integrated

[1] Electronic Industries Alliance (EIA). Accredited by American National Standards Institute (ANSI), EIA provides a forum for the industry to develop standards and publications throughout the electronics and high-tech industries. One Electronic Industries Association Unit (1U) is 44.45 mm (1.75 in.).

Virtual Ethernet adapter, so a fully configured System p 570 server can comprised several different IVE feature codes.

The following feature codes are available, at the time of writing, for each System p 570 enclosure:

- FC 5636 (standard), Integrated 2-port 1 Gbps

  16 MAC addresses, one port group
- FC 5639 (optional), Integrated 4-port 1 Gbps

  32 MAC addresses, two port groups
- FC 5637 (optional), Integrated 2-port 10 Gbps SR[2] (optical)

  32 MAC addresses, two port groups

Figure 1-3 shows the major components of the Integrated Virtual Ethernet adapter hardware and additional system ports, according to the different feature codes.

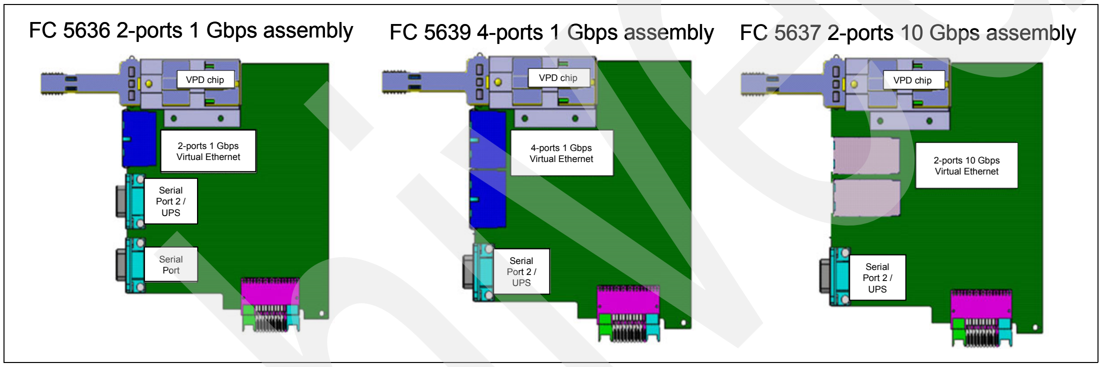

Figure 1-3 Integrated Virtual Ethernet feature codes and assemblies

Any IVE feature code located in the first enclosure of a System p 570 also includes the System VPD (Vital Product Data) Chip and system (serial) port (1 or 2 depending on the feature code).

The IVE feature code is installed by manufacturing. Similar to other integrated ports, the feature is neither hot-swappable nor hot-pluggable and must be serviced by a trained IBM System Service Representative.

[2] 10 Gbps SR (short range) is designed to support short distances over deployed multi-mode fiber cabling. It has a range of between 26 m and 82 m depending on cable type. It also supports 300 m operation over new, 50 µm 2000 MHz·km multi-mode fiber (using 850 nm).

Figure 1-4 shows the rear view of a basic System p 570 in a state of disassembly with some necessary components and covers removed to highlight the connection of the feature code assembly into the server enclosure I/O subsystem system board.

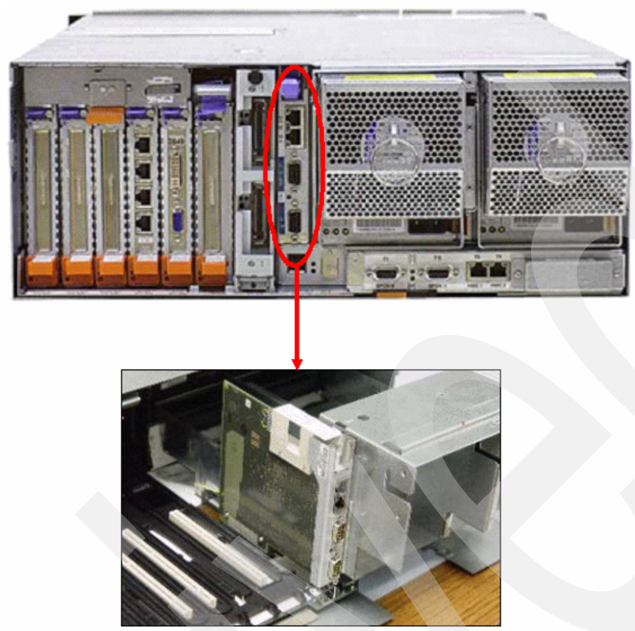

Figure 1-4 Integrated Virtual Ethernet adapter connection on System p 570 I/O system board

Feature code port and cable support

All the IVE feature codes have different connectivity options and different cable support (Figure 1-5).

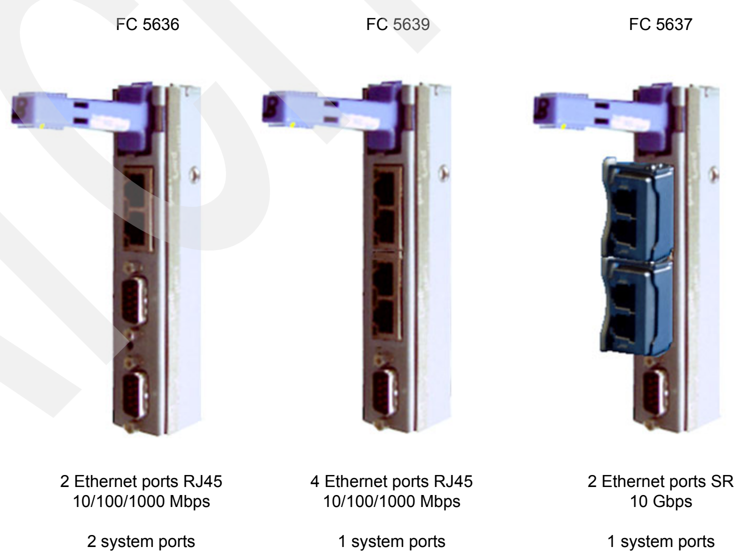

Figure 1-5 IVE physical port connectors according to IVE feature codes

FC 5636 and FC 5639 support:

- 1 Gbps connectivity
- 10 Mbps and 100 Mbps connectivity
- RJ-45 connector

Use the Ethernet cables that meet Cat 5e[3] cabling standards, or higher, for best performance.

FC 5637 supports:

- 10 Gbps SR connectivity
- 62.5 micron multi-mode fiber cable type
  - LC optical physical connector type
  - 33 meters maximum range

1.4 System p virtualization review

Starting with the POWER5 processor-based systems introduced in 2004, clients have realized the value of the System p Advanced POWER Virtualization. This proven technology is designed for stability and production workloads.

The Advanced POWER Virtualization feature enables administrators and companies to build their own flexible and reliable infrastructure. It is designed to simplify IT environments and operations.

One of the real values of Advanced POWER Virtualization is smooth systems operations and fast responses to client demands, leading to a substantial improvement in TCO[4] and simplified IT operations.

With the announcement of the IBM System p servers based on innovative IBM POWER6 processor technology, IBM is combining the fastest ever POWER Architecture technology with enhanced virtualization capabilities and reliability, availability, and serviceability (RAS) functions.

One of these enhanced virtualization capabilities is the virtualization of a leading-edge 10 Gbps speed physical Ethernet port and the ability to share it within a partitioned system.

Additional POWER6 processor-based server features include application mobility and partition mobility. For more information about these topics see:

- *Introduction to Workload Partition Management in IBM AIX Version 6*, SG24-7431
- *IBM System p Live Partition Mobility*, SG24-7460.

1.4.1 Virtualization features and options

The entire IBM System p product line offers full advanced virtualization features enabling customers to fit their needs, with greater performances and more capabilities in a smaller footprint.

Virtualization includes Logical partitioning, which includes dynamic LPAR and Micro-Partitioning™.

3 Category 5 cable, commonly known as Cat 5, is a twisted pair cable type designed for high signal integrity. Category 5 has been superseded by the Category 5e specification.
4 Total cost of ownership.

The Advanced POWER Virtualization feature includes a combination of hardware and software that support and manage the virtual environment on POWER5 and POWER6 processor-based systems, such as:

- Firmware enablement for Micro-Partitions
- Virtual I/O Server software to host and virtualize:
  - Storage
  - Ethernet adapter sharing

One of the new and important enhancements enabled with the POWER6 technology is the Integrated Virtual Ethernet adapter.

Configuration samples using virtualization features

Several publications provide Advanced POWER Virtualization features details and best practices. In this section we remind you of some of the Advanced POWER Virtualization essential features and capabilities.

We recommend the reference *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration*, SG24-7940, if you are new to virtualization on IBM System p servers.

Virtual I/O Server

As part of Advanced POWER Virtualization, the Virtual I/O Server allows sharing of physical I/O resources for the AIX and Linux operating system logical partitions. It enables better I/O resource optimization and utilization of disk and optical space.

Shared Ethernet Adapter

POWER5 processor-based server introduced the Shared Ethernet Adapter (SEA). This is a special Ethernet interface that can be created within a Virtual I/O Server that allows one or more Ethernet adapters (PCI, PCI-X, and PCIe) to be shared between several logical partitions (LPARs). Once the SEA is defined within a Virtual I/O Server, LPARs can access external local area networks that may or may not be using VLANs.

Shared Ethernet Adapter failover

Virtual I/O Server software Version 1.2 supports SEA failover function. There are two new SEA attributes that support this new function:

- Priority - Two values are possible — the numerically lower to act as the primary, and the numerical higher to act as the backup.
- Control channel, which must be on a separate and dedicated virtual Ethernet for SEA adapter control function only (SEA heartbeat).

When a loss of connectivity is detected, or in case the first Virtual I/O Server (primary interface of SEA failover) is down for maintenance, then the second Virtual I/O Server (backup interface of SEA failover) handles the traffic. This operation is transparent for all the client LPARs.

Shared Ethernet Adapter implementation

The performance of the Shared Ethernet adapter depends on the number of client LPARs, their network workloads, and the amount of assigned resources in the Virtual I/O Server.

In addition, if the Virtual I/O Server provides storage sharing and network connectivity to client LPARs, when CPU resources are constrained, the administrator has the following choices:

- Place an IP on a different interface from the SEA for management use and dynamic LPAR support.
- Verify that the thread attribute is set to 1 (shown with `lsdev -dev <SEA> -attr` command) that will equalize the priority between virtual disk and SEA network I/O. This throttles Ethernet traffic to prevent it from consuming a higher percentage of CPU resources versus the virtual SCSI activity. This is a concern only when CPU resources are constrained resources.

1.5 IVE and Virtual I/O Server positioning

The Integrated Virtual Ethernet adapter provides a combination of virtual Ethernet and external network communication through the VIOS. The purpose of this section is to assist you in differentiating the various virtualization resources.

Table 1-1 provides a comparison of functions and their application with respect to IVE and the Virtual I/O Server.

Table 1-1 IVE versus the Virtual I/O Server positioning

Virtualized function	IVE	VIOS
Basic functions		
Requires activation of the Advanced POWER Virtualization feature code		✓
Virtualizes Ethernet resources over a single server	✓	✓
Provides hardware acceleration	✓	
Managed by the HMC as a hardware device	✓	
Scalable and flexible for all virtualization requirements		✓
Operating system based - runs in an LPAR		✓
Enables advanced server functions, such as Live Partition Mobility		✓
Strategic solution for IBM Systems client requirements	✓	✓
Network functions		
Can be used as an SEA	✓	✓
Provides scalable threading for multiple-core systems	✓	
Provides 10 Gbps wire speed performance with very low latency	✓	
Provides secure network connections between LPARs	✓	✓
Supports industry standard networking attributes, such as VLAN, MTU, and so on	✓	✓
Does not requires a PCI-x or PCIe slot for external communication	✓	

Virtualized function	IVE	VIOS
Storage functions		
Virtualizes SCSI-based file system data		✓

2

Integrated Virtual Ethernet architecture

The Integrated Virtual Ethernet adapter (IVE) is available with 10 Gbps or 1 Gbps connectivity. The IVE provides virtualization and queuing per TCP connection and virtual addressing in descriptors. IVE also contains specific hardware for networking stack processing assists, and the system memory is used as the primary repository of control information.

2.1 Architecture introduction

Figure 2-1 shows a high level-logical diagram of the HEA.

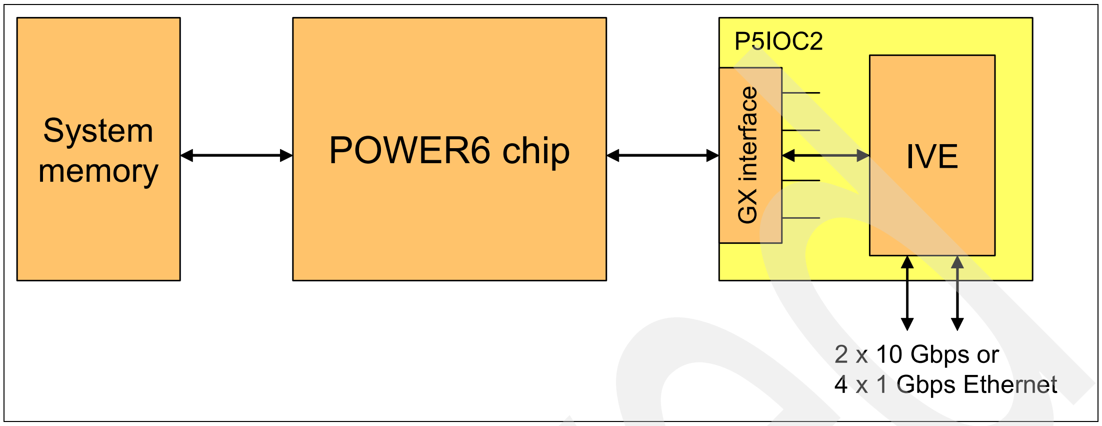

Figure 2-1 IVE system placement

Every POWER6 processor-based server I/O subsystem contains the P5IOC2 chip. It is a dedicated controller that acts as the primary bridge for all PCI buses and all internal I/O devices. IVE major hardware components reside inside the P5IOC2 chip.

The IVE design provides a great improvement of latency for short packets. Messaging applications such as distributed databases require low-latency communication for synchronization and short transactions. The methods used to achieve low latency include:

- GX+ bus attachment.
- Immediate data in descriptors (reduced memory access).
- Direct user space per-connection queuing (OS bypass).
- Designed for up to three-times the throughput improvement over current 10 Gbps solutions.
- Provides additional acceleration functions in order to reduce host code path length. These include header/data split to help with zero-copy stacks.
- Provides I/O virtualization support so that all partitions of the system can natively take advantage of the features.
- Allows one 10 Gbps port to replace up to 10 dedicated PCI 1 Gbps adapters in a partitioned system.

One of the key design goals of the IVE is the capability to integrate up to two 10 Gbps Ethernet ports or up to four 1 Gbps Ethernet ports into the P5IOC2 chip, with the effect of a low-cost Ethernet solution for entry and mid-range server platforms. The 10 Gbps, 1 Gbps, 100 Mbps, or 10 Mbps speeds share the same I/O pins and do not require additional hardware or feature on top of the IVE itself. Another key goal is the support of all the state-of-art NIC functionality provided by leading Ethernet NIC vendors.

IVE offers the following functions with respect to virtualization:

- Up to 32 logical ports identified by MAC address.
- Sixteen MAC addresses are assigned to each IVE port group.
- Each logical port can be owned by a separate LPAR.
- Direct data path to LPAR.

- Default send and receive queues per LPAR.
- Ethernet MIB and RMON counters per LPAR.
- VLAN filtering per logical port (4096 VLANs * 32 Logical Ports).
- Internal layer 2 switch for LPAR to LPAR data traffic.
- Multicast/Broadcast redirection to Multicast/Broadcast manager.

IVE relies exclusively on the system memory and CEC CPU cores to implement acceleration features. There is no requirement of dedicated memory, thus reducing the cost of this solution and providing maximum flexibility. IVE Ethernet MACs and acceleration features consume less than 8 mm^2 of logic in CMOS 9SF technology.

IVE locates its firmware in the Service Processor flash, which is then passed to the POWER Hypervisor (PHYP) control upon system startup. Therefore, flash code update is done by PHYP.

2.2 Logical and physical components

The IVE offering includes physical hardware and firmware components that provide transport functions to the operating systems.

- The Host Ethernet Adapter (HEA) is the major hardware component of the IVE, and it is included in the P5IOC2 chip. HEA also includes all the logical ports and the virtual layer 2 switches (see Figure 2-2 on page 14) and connects to the physical port.

  Once a logical port is assigned to LPAR, the LPAR operating system recognizes the HEA as the Logical Host Ethernet Adapter (LHEA).
- Physical ports are the hardware component to allow the external connectivity to LAN, VLAN, or switches.
- The daughter cards that contain the POWER Hypervisor and drivers for the IVE feature.

2.2.1 IVE logical component introduction

Depending on the IVE feature code installed, one or two physical ports are grouped into a port group. Any port group supports up to 16 LHEA ports.

Figure 2-2 shows the major logical components of an IVE.

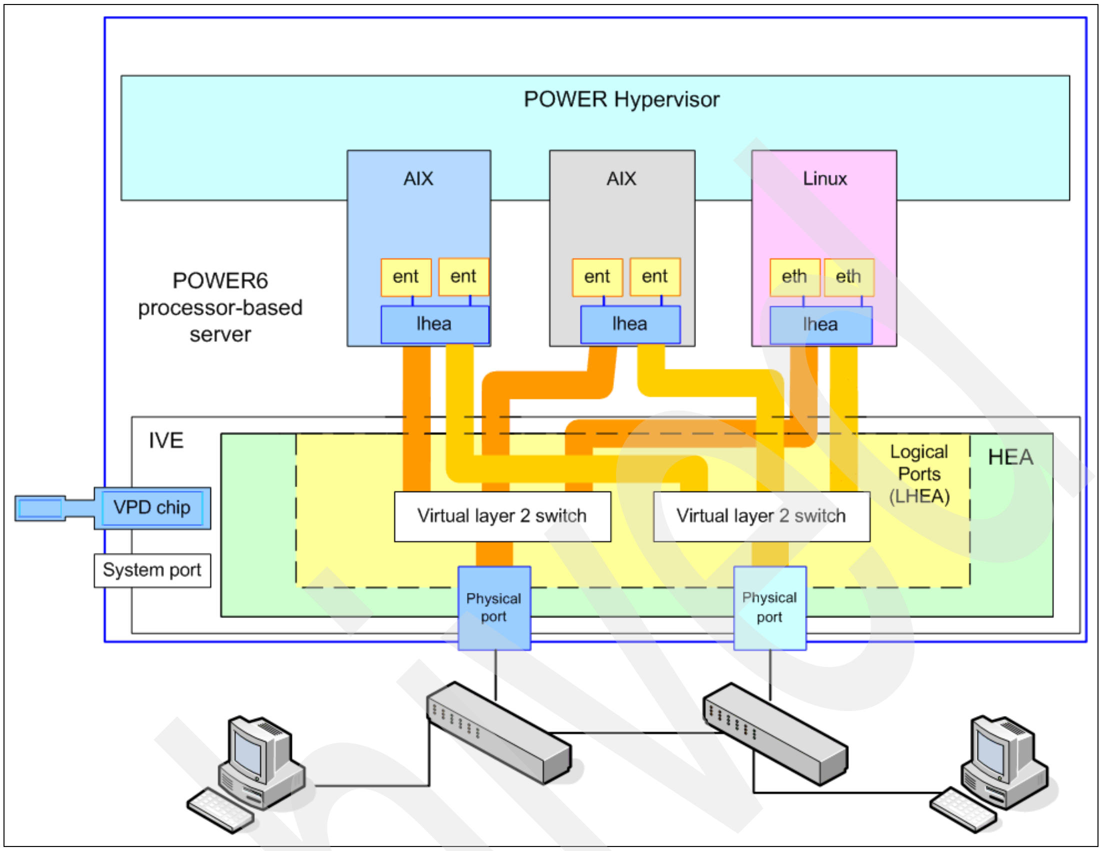

Figure 2-2 IVE logical components diagram

The major IVE concepts are defined as follows:

HEA Host Ethernet Adapter, the key imbedded function located on the P5IOC2 I/O controller chip.

LHEA Logical Host Ethernet adapter, a logical representation of a physical HEA adapter. This is the parent device of an LHEA port.

LHEA port Logical representation of a physical HEA port. The maximum number of ports you can assign to an LPAR depend on the IVE feature code.

Port group This is a group of logical ports that share one or two physical ports, depending from IVE feature code. Any IVE feature supports up to two port groups, and any port group supports up to 16 logical ports (LHEA port).

MCS Multiple Core Scaling (MCS) is a parameter that allows multiple receive and transmit queues (QP) to be supported in each LPAR. The MCS value must be defined at port group level.

2.2.2 IVE and physical, logical, and port groups

This section provides the possible configurations for each IVE feature codes (Table 2-1).

Table 2-1 Feature codes allowed configurations

Feature code	Physical ports availabl e	Port groups available	Port group MCS value and Queue Pair number (QPNs) (selectable)	Max logical ports per port group related to MCS value	Max logical ports available per LPAR per feature code
FC 5636 (1 Gbps)	2	1	1	16	2
			2	8	2
			4	4	2
			8	2	2
			16	1	1
FC 5639 (1 Gbps)	4	2	1	16	4
			2	8	4
			4	4	4
			8	2	4
			16	1	2
FC 5637 (10 Gbps)	2	2	1	16	2
			2	8	2
			4	4	2
			8	2	2
			16	1	2

The following are general considerations:

- One logical port per physical port per LPAR.
- The MCS value is set for a port group and all associated physical ports (you can also choose to use this enhancement on the OS level; see 4.2, "Configure LHEA on AIX" on page 53).
- All logical ports are capable of reaching the media speed. See also 4.2.3, "Performance considerations and recommendations" on page 57.

If you need to correlate the LHEA logical device with the IVE physical port, then the AIX physical location provides the physical port associated with the logical port assigned to the LPAR. Table 2-2 lists the AIX physical location for the LHEA beside the IVE feature code and physical port selected along with HMC and SMS port numbers.

Table 2-2 Identify IVE physical ports, logical port number, and AIX physical location

IVE feature code	Physical ports	Physical port number	Port group	HMC logical port number	SMS logical port number	AIX entstat -d	AIX physical location
5636	2	2	2	1 to 16	1 to 16	1 to 16	*enclosure*-P1-C10-T2
		1					*enclosure*-P1-C10-T1

IVE feature code	Physical ports	Physical port number	Port group	HMC logical port number	SMS logical port number	AIX entstat -d	AIX physical location
5637	2	2	1	1 to 16	1 to 16	1 to 16	*enclosure*-P1-C10-T2
		1	2	1 to 16	17 to 32	17 to 32	*enclosure*-P1-C10-T1
5639	4	4	1	1 to 16	1 to 16	1 to 16	*enclosure*-P1-C10-T4
		3					*enclosure*-P1-C10-T3
		2	2	1 to 16	17 to 32	17 to 32	*enclosure*-P1-C10-T2
		1					*enclosure*-P1-C10-T1

2.2.3 Manufacturer Default Configuration and IVE

Manufacturer Default Configuration (MDC) is the default profile stored in a System p server when it is shipped by manufacturing. The system is configured with only one LPAR that owns all the resources. If a System p570 is connected to the Hardware Management Console, but the system is non-partitioned, this configuration includes:

- One LHEA per each physical HEA
- One LHEA port (logical port) per physical port
- AIX/Linux device driver has access to HEA

The MDC physical port offload settings and driver attributes are the following:

- 1500 or 9000 MTU physical port setting (settable by the OS in MDC mode)
- Flow Control disabled (settable by the OS in MDC mode)
- Media speed (settable by the OS in MDC mode)
- Promiscuous mode
- MCS default value of 4

  This allows partitions to be configured with IVE ports without re-IPL of the initial partition.
- Duplex mode and speed set to auto
- All VLANs allowed

The HMC allows non-MDC configurations, where one LPAR can be configured with all system resources. This is also known as a *full system configuration*.

2.2.4 IVE device topology

Figure 2-3 on page 17 represents a view of a system installed with IVE features under SMS.

FC 5636 (Integrated 2-port 1 Gbps) has two physical ports configured into one port group and it is *Port Group 2*. The HMC, SMS, and AIX (in the case of the following command) list the logical ports from 1 to 16.

```
entstat -d entx | grep "Logical Port Number"
```

All the other IVE feature codes have two port groups, and for these other IVEs, the first logical port is located in the first port group (see also Table 2-2 on page 15).

FC 5639 (Integrated 4-port 1 Gbps) provides four physical ports. The first two physical ports (T1 and T2) are grouped in *Port Group 2*. The third (T3) and the forth (T4) physical ports are grouped in the *Port Group 1*.

Figure 2-3 shows that FC 5637 (Integrated 2-port 10 Gbps SR optical) is providing up to 32 logical ports (depending from MCS value). When you use the HMC to select a logical port from a port group, you are able to define which logical port you are associating to an LPAR. The HMC shows the number of the logical ports from 1 to 16 and the associated port group. SMS (also displayed with the **`entstat -d entx | grep "Logical Port Number"`** AIX command output, as shown in Example 4-13 on page 57) shows the logical port from 1 to 16 if it belongs to the first port group, but from 17 to 32 if it belongs to the second port group.

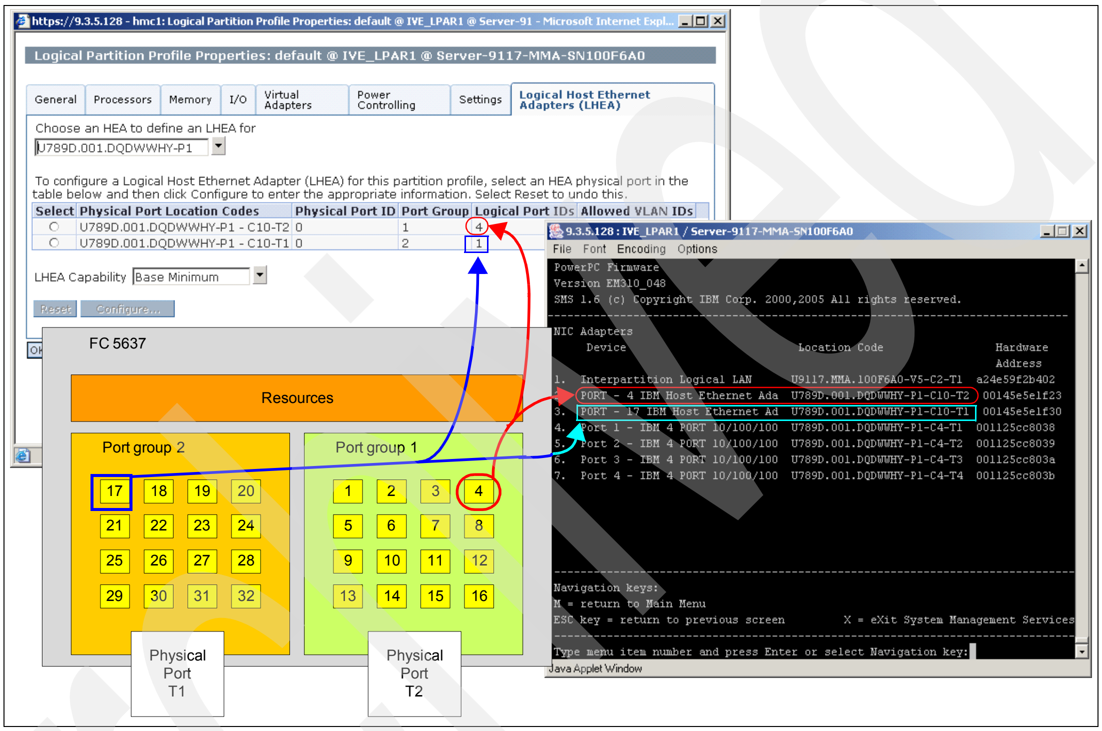

Figure 2-3 SMS hardware listing

2.2.5 AIX IVE device topology

As a logical child of the LHEA, the operating system recognizes the logical representation of a physical port as a standard Ethernet interface (see Example 2-1).

Example 2-1 lsdev -Cc adapter | grep -i hea command output

```
ent0   Available  Logical Host Ethernet Port (lp-hea)
lhea0  Available  Logical Host Ethernet Adapter (l-hea)
```

The HEA that is providing the external connectivity is recognized from the operating system with a logical I/O slot location, and not with a virtual I/O slot location such as all the virtual devices provided from a Virtual I/O Server (see Example 2-2) or virtual Ethernet adapters. This logical location can be HEA1, HEA9, HEA17, or HEA25, according to which system drawer (in a multidrawer system) the IVE is located.

Example 2-2 lsslot -c slot command output

```
Slot                   Description      Device(s)
HEA 1                    Logical I/O Slot  lhea0 ent0
U789D.001.DQDWWHY-V2-C0  Virtual I/O Slot  vsa0
U789D.001.DQDWWHY-V2-C2  Virtual I/O Slot  vscsi0
```

Example 2-3 also shows that you may have two network interfaces, each from a different physical port, associated to a different port group of the same IVE. If both physical ports belong to the same IVE feature code so they are in the same server enclosure, the Ethernet interface parent is the same LHEA.

Example 2-3 Relation between parent devices

```
#lsslot -c slot | grep -i hea
Slot                   Description      Device(s)
HEA 1                    Logical I/O Slot  lhea0 ent0 ent1

#lsdev -l lhea0 -F physloc
U789D.001.DQDWWHY-P1

#lsdev -l ent0 -F physloc
U789D.001.DQDWWHY-P1-C10-T2

#lsdev -l ent1 -F physloc
U789D.001.DQDWWHY-P1-C10-T1
```

If you need to identify the IVE components in a multiple enclosure System p 570, correlate the LHEA logical device and the HEA logical I/O slot location (it is *lhea0* and *HEA 1* in Example 2-3 on page 18). Then you can identify where the IVE feature code assembly is physically located (Figure 2-4).

Figure 2-4 How to identify the IVE inside System p 570 enclosure

2.2.6 Linux IVE device topology

The Linux on POWER service aids for hardware diagnostics are available from IBM. The service aids allow system administrators to extract valuable information from the service processor for servicing. Many of the commands packaged inside the service aids are very similar to the AIX commands.

At the time of writing, the latest version of service aids package can be downloaded from the following Web site:

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

After installing the package, you find additional commands available to extract information from your System p server (or LPAR) using Linux. These commands are installed into the /usr/sbin/ibmras or /usr/sbin/ directory.

The `lscfg` command lists all the hardware information that is available in the system, and the following flags add additional details:

-v Displays Vital Product Data (VPD) when available

-p Displays system-specific device information

Example 2-4 shows how to query all of the defined Ethernet interface within a Linux on POWER LPAR or server.

Example 2-4 lscfg -vp|grep eth command output

```
eth0            U789D.001.DQDWWHY-P1-C10-T2
eth1            U789D.001.DQDWWHY-P1-C10-T1
eth2            U789D.001.DQDWWHY-V9-C4-T1
```

2.3 Physical port and logical port options

From a hardware perspective, the physical port is the connector where you can plug in the Ethernet cable (copper or optical, according to the selected IVE feature code). The logical port is the virtualization of the physical port that allows LPARs to share this physical port and its connectivity. In addition to that, any logical port owns its MAC address that it is stored in IVE hardware.

A physical port supports up to 16 logical ports. Dependant on the MCS value, the number of logical ports available through a physical port can be modified.

When the promiscuous mode option is set, the use of the physical port is available only to a single LPAR.

2.3.1 Physical port promiscuous mode

When a physical port is set to promiscuous mode through the Hardware Management Console (HMC), the port is reserved to an LPAR, and all the other LPARs cannot use this physical port.

Common uses for promiscuos mode are:

- Assigning a physical port with the promiscuous mode enabled to a Virtual I/O Server when you need to share the physical port between *more* than 16 LPARs using the virtualization of the Virtual I/O Server.
- When you want to use the physical port as a dedicated resource for a given LPAR.

Promiscuous mode adds processing that may affect performance (4.2.3, “Performance considerations and recommendations” on page 57).

When the promiscuous mode is enabled, the only logical port configurable for an LPAR is number 1 on the physical port. You will have a maximum of 16 logical ports (based on the MCS value) for the related port group that are usable if your port group has another physical port that is not set to promiscuous mode. If both physical ports are set to promiscuous mode, their port group will only offer one logical port per physical port.

If you assign a logical port to a Virtual I/O Server through a physical port without the promiscuous mode enabled, the Virtual I/O Server is not able to use the **`mkvdev`** command to define the SEA. SEA requires HEA as a dedicated adapter, capable of viewing all network packets in order to deliver them to all LPARs or to drop them when they do not belong to the defined network.

Promiscuous mode is not only for use as an SEA. It may be useful for other server functions. It is helpful for use by packet sniffing applications, such as **`tcpdump`** or **`iptrace`**.

In the case of Figure 2-5, the second physical port (the one on the right side) is set to promiscuous mode and assigned to a Virtual I/O Server. The Virtual I/O Server recognizes the first logical port associated to the physical port as a standard Ethernet interface that can be used to define a Shared Ethernet Adapter (SEA).

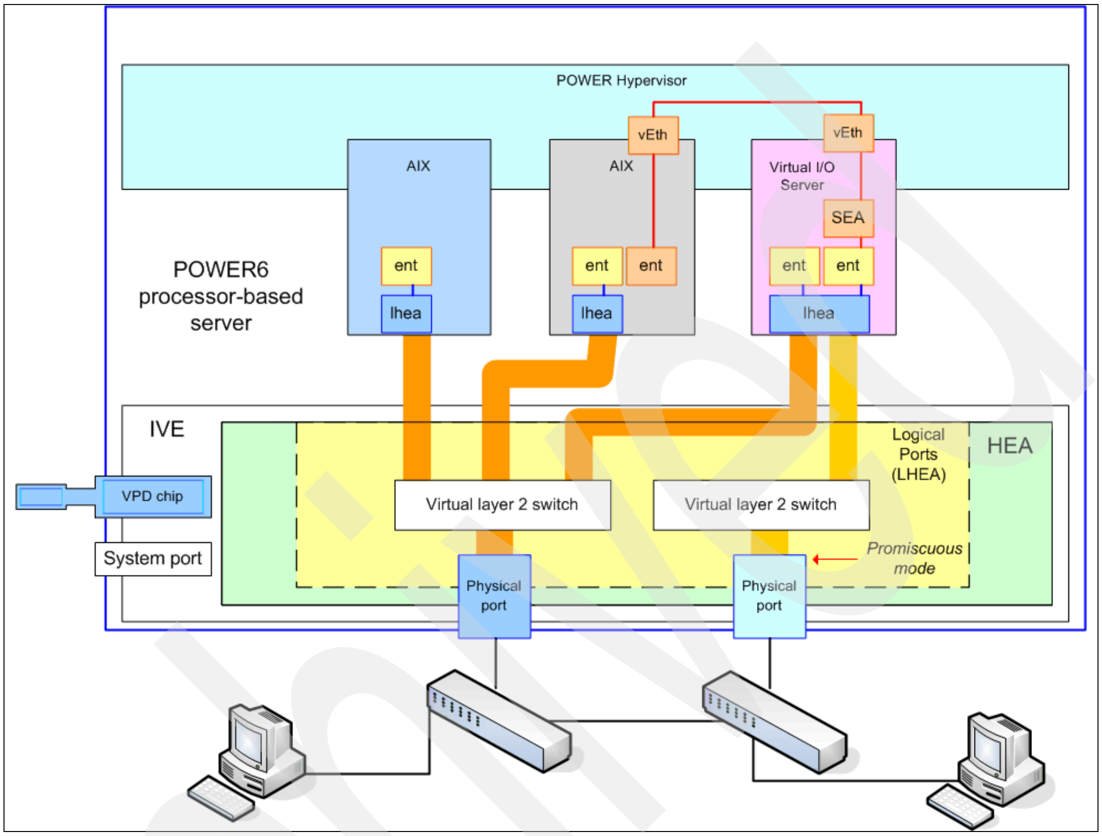

Figure 2-5 Example of physical port assigned to VIOS in promiscuous mode

2.3.2 SEA failover and EtherChannel with IVE

When you assign a physical port to a Virtual I/O Server by enabling the promiscuous mode, the physical port is then dedicated for SEA use.

The LHEA (Ethernet interface for operating systems) can be used to create an EtherChannel device, even as the backup adapter, but it is not supported to mix an LHEA with other kinds of Ethernet device other than LHEAs (such as dedicated adapter ports). All standard EtherChannel features are available, such as backup adapter and ping dead gateway detection, to name a few.

When using LHEA ports in aggregation, all other partitions using the LHEA ports from the same physical HEA ports must also be in aggregation.

2.3.3 Logical port options

The use of Integrated Virtual Ethernet adapters is based on a concept of logical ports. From an LPAR standpoint (OS device driver), a logical port appears as a basic Ethernet interface (it has a MAC address, MIB counters, VLAN filtering table, large send enabled or not, and so on).

A logical port is defined on a physical port and therefore competes for the bandwidth with all the other logical ports defined on the same physical port. This does not prevent a logical port from reaching the media speed if the other logical ports do not have intensive I/O workload.

A logical port should not be confused with a VLAN interface. Multiple VLANs (corresponding to a different IP subnet) are supported on a single logical port.

Logical ports are assigned to physical ports and to LPARs by one of the following mechanisms:

- Manufacturing Default Configuration (MDC) is programmed into the NVRAM during the System p manufacturing process (see 2.2.3, “Manufacturer Default Configuration and IVE” on page 16, for details).
- The Hardware Management Console (HMC) is used to perform the configuration (see “Configure an HEA with an HMC” on page 46).

VLANs

Ethernet supports up to 4096 virtual LANs (VLANs) on a network. HEA supports VLANs and can perform VLAN filtering on a logical port basis.

There are two aspects of VLANs that are addressed here:

- VLAN security
- VLAN filtering

VLAN security

VLAN security on the internal layer 2 switch is provided through HEA, HMC, and PHYP mechanisms. The POWER Hypervisor (PHYP) allows the HMC to identify up to 20 individual VLAN IDs that a logical port can use. PHYP uses this information in two ways:

- To modify the setting of VLAN filtering. PHYP does not allow the operating system to set the HEA VLAN filtering settings to allow VLANs other than those configured.
- To allow only those VLANs that are configured for a specific logical port. This action is optional since the non-allowed packets would be dropped by the inbound VLAN filters to the logical port.

The result is that the HMC, along with PHYP, can control the settings of the HEA hardware so that only permitted VLANs will be allowed.

PHYP does not reject any requests to allow VLANs that are not permitted. Rather, it silently modifies the requests. This behavior is to follow what occurs with a real Ethernet layer 2 switch.

If the set of allowed VLANs changes (after a dynamic LPAR operation with a new LHEA or after a profile modification and reactivation of the LPAR), then PHYP adjusts the VLAN filter settings for that logical port to either further restrict the set of VLANs or to add into the set (if the LPAR previously requested a VLAN that was not previously allowed). This behavior follows the behavior of a real Ethernet switch dynamically changing its port based on the VLAN configuration. A change in allowed VLANs can only occur when the logical port is first added to an LPAR (through activation or dynamic LPAR add) or for an existing logical port, when the profile is modified and the LPAR reactivated to use the modified profile.

Multicast and broadcast

Ethernet has a native broadcast (BC) and multicast (MC) capability. Packets in the network can be directed to all stations (BC) or to a group of stations (MC). Ethernet switches must replicate MC and BC packets to all appropriate ports.

Since the HEA contains one or more logical layer 2 switches, it is also responsible for replicating and forwarding MC/BC packets to all interested logical ports (on that logical switch). MC/BC packets received from the physical port must be forwarded to all interested logical ports. MC/BC packets transmitted from a logical port must be forwarded to other interested logical ports and to the physical port.

HEA implements MC/BC forwarding through a combined use of hardware (HEA) and firmware (PHYP). The HEA will direct all MC/BC packets that must be forwarded to a special queue named the default MC/BC queue (see Figure 2-6). There is one MC/BC queue for each physical port (and thus one for each logical switch). The MC/BC manager component of PHYP handles these queues by replicating the packets and forwarding them to the appropriate logical ports and physical port using the HEA hardware.

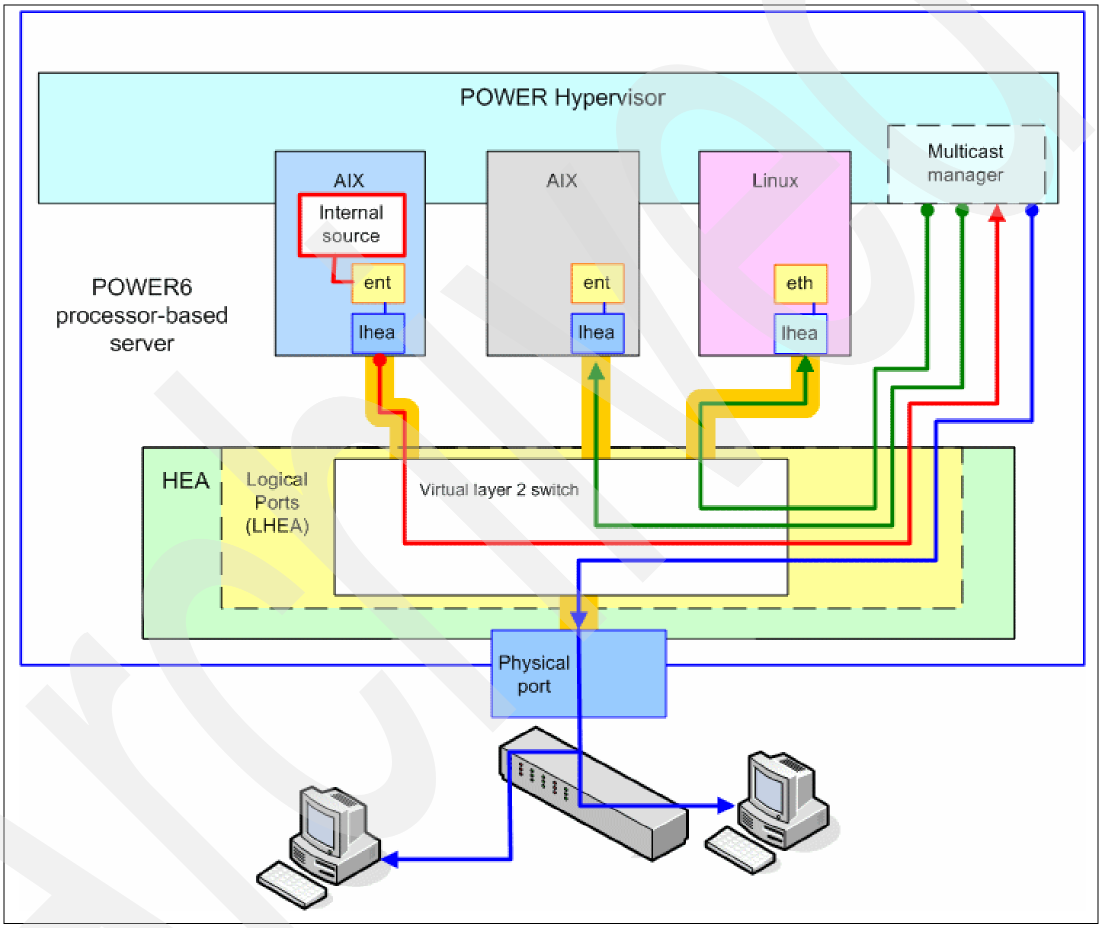

Figure 2-6 Logical port broadcast/multicast flow examples when the source is internal

The operating system device driver typically has an interface that allows upper layer protocols to register for particular multicast or broadcast addresses. The device driver then utilizes available hardware filters to permit those BC or MC packets. HEA implements MC/BC filtering by taking advantage of the fact that the switch is controlled by firmware and is part of PHYP.

2.3.4 Port group options and Multiple Core Scaling

Port group numbers and physical characteristics change according to IVE feature codes. The feature codes available for System p 570 present different port group topology. Figure 2-7 compares them. Note that the feature codes are listed in the order in which they appear in Figure 2-7:

- FC 5636 supports two physical ports, but both belong to the same port group.
- FC 5639 supports four physical ports and two port groups. Each port group contains two physical ports.
- FC 5637 supports two physical ports and any port belong to its port group, and each physical port belongs to one port group.

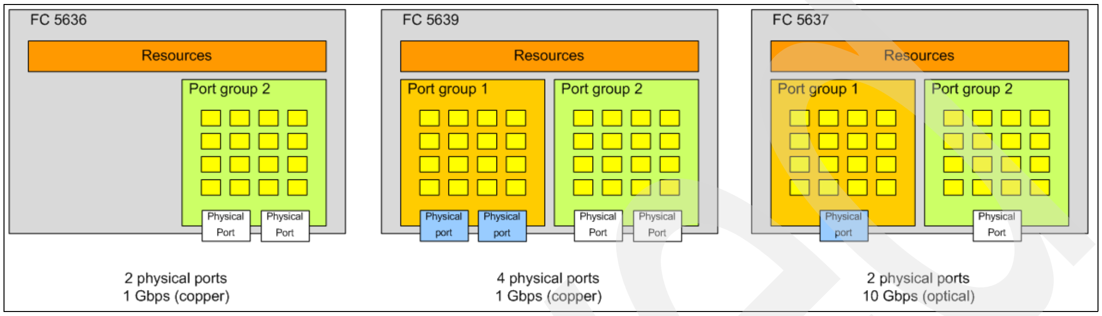

Figure 2-7 Port group option according to IVE feature code

Changing the Multi-Core Scaling (MCS) value affects a port group and changes the maximum number of logical ports available to LPARs.

Multiple Core Scaling

Multiple Core Scaling is part of the IVE design that provides a parameter to set the number of queue pairs (QPs) for use in each logical partition. A queue pair is a way to break network traffic into multiple streams that can be dispatched to multiple cores of the POWER6 microprocessor and take advantage of parallel processing.

An MCS value must be defined at the port group level. This enables the traffic from each physical port associated to this port group to be processed in parallel by interrupt handlers running on different processors on partitions configured to use them.

Figure 2-8 shows how the QP can be dispatched to the processor resources when MCS is set to 1 and enabled at the AIX level.

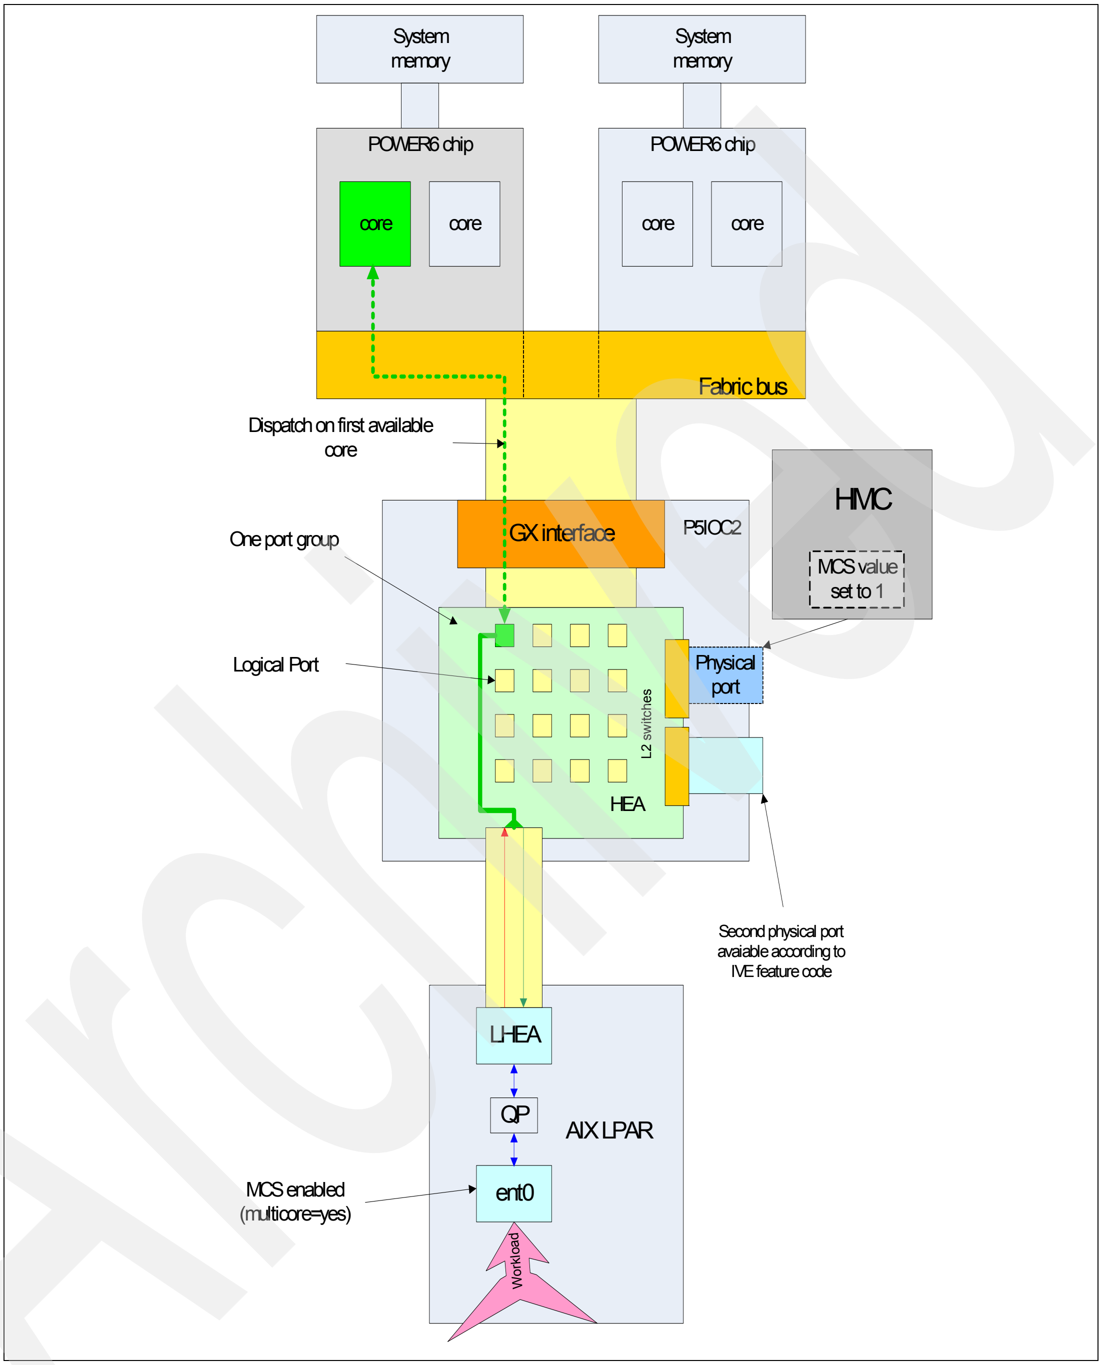

Figure 2-8 QP dispatch with MCS set to 1 and MCS enabled into AIX LPAR

According to the MCS value, the number of available logical ports changes (Figure 2-9). When the MCS is set to 1, this allows you to configure up to 16 logical ports, but when the MCS is set to 4, this allows you to configure up to 4 logical ports.

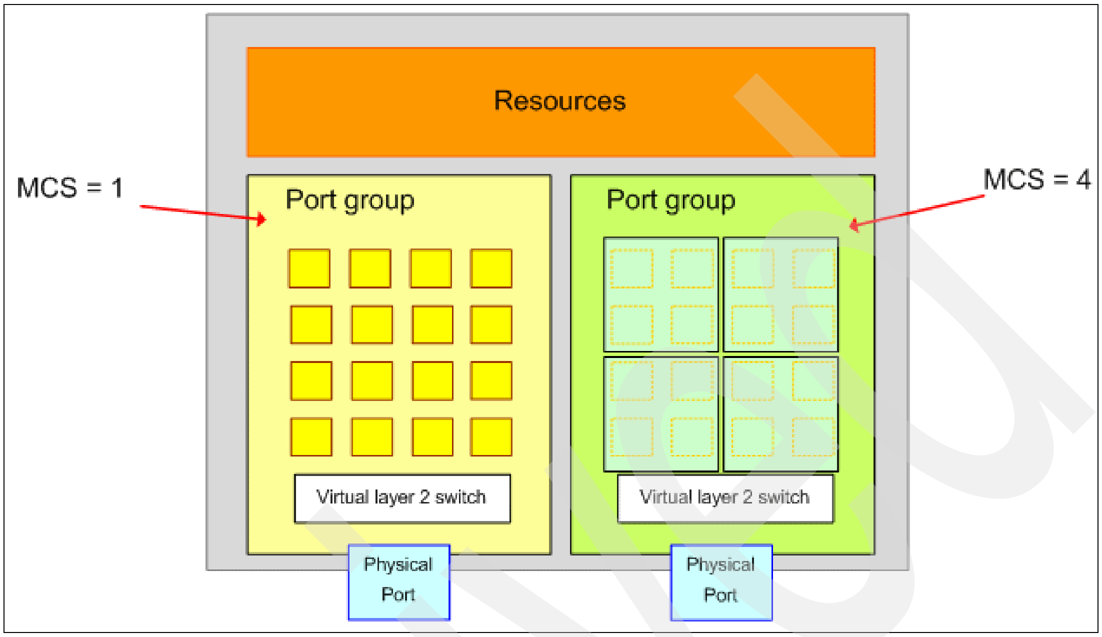

Figure 2-9 Logical ports available per port group according to MCS value

The MCS can be set to 4 or a higher value in the case of a high system workload where the number of real processor cores are available to a partition allowing the processing of network bandwidth to be distributed by using multiple queues. You should consider balancing the processing required to service a large number of queues versus the advantages it may provide on a busy network by allowing concurrent processing.

If the MCS value is equal to 4, the LPAR operating system uses four queues to receive and four queues to transmit. With this setting, the total number of logical ports that can be assigned to LPARs from the same port group is four. As shown in Figure 2-9, the second port group offers four logical ports because the MCS value is set to 4, and the first port group offers 16 logical ports because the MCS value is set to 1.

See Table 2-1 on page 15 for additional information about the relation between MCS values and the number of IVE logical ports available to LPARs.

Figure 2-10 shows how QPs are dispatched to the CPU resources when the MCS is set to 4 and the MCS is enabled at the AIX operating system level.

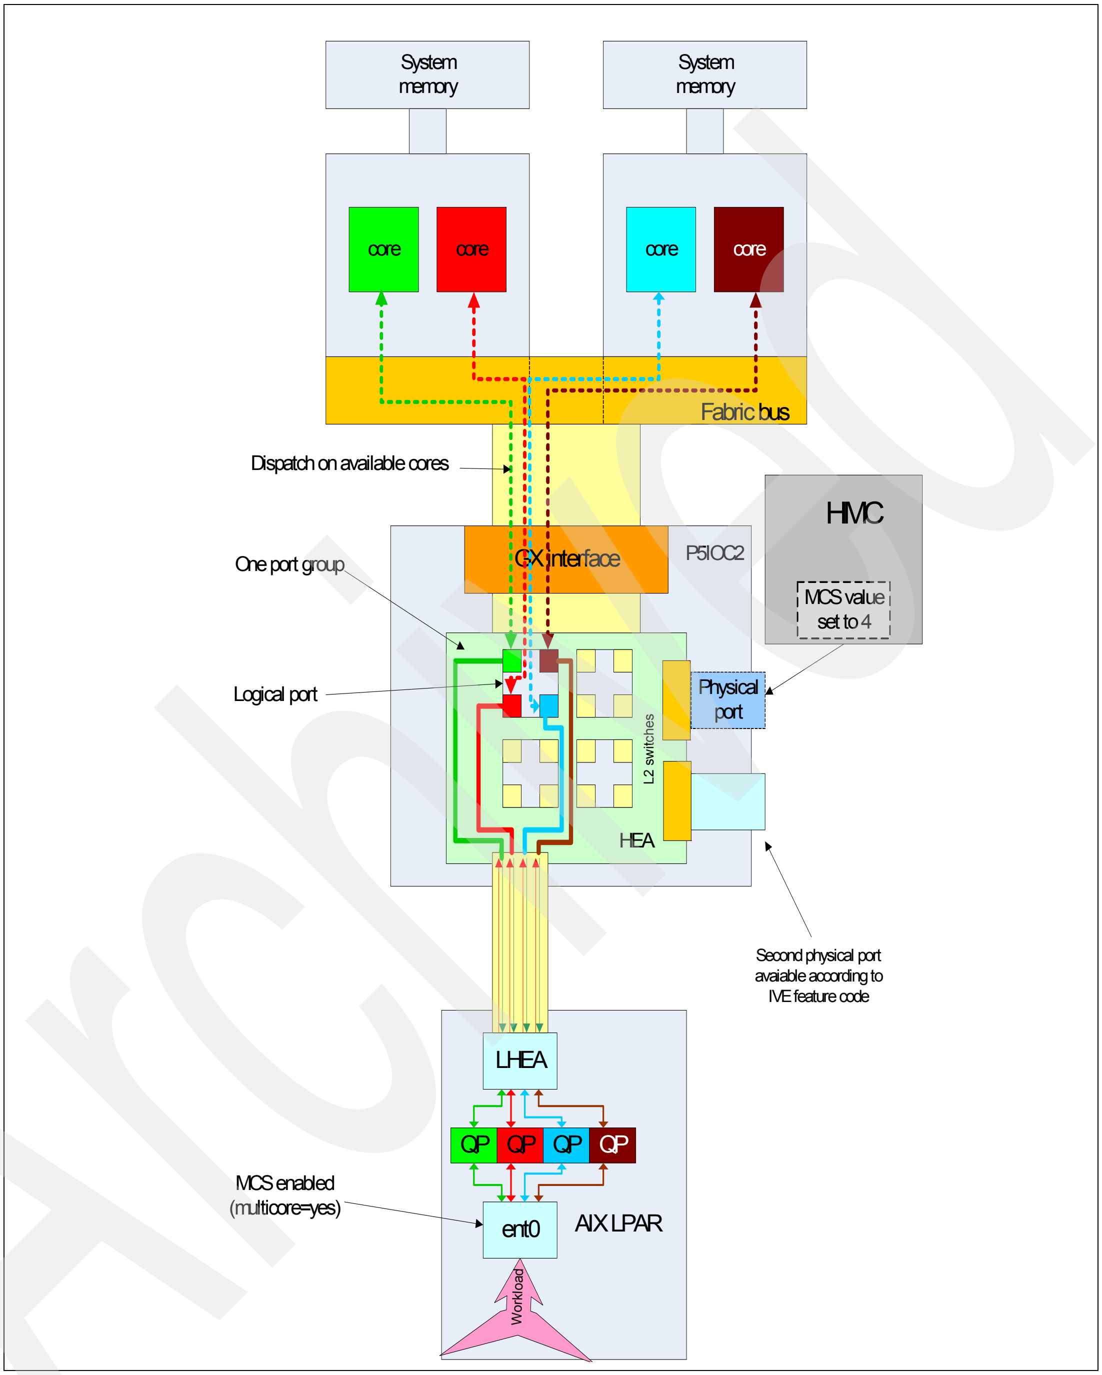

Figure 2-10 QPs dispatch with MCS set to 4 and MCS enabled into AIX LPAR

If you want AIX to disable MCS for a specific LPAR, you can change the multicore attribute value on the ent*n* interface from yes to no. In this case the LPAR uses only one QP and dispatches this QP as shown in Figure 2-11.

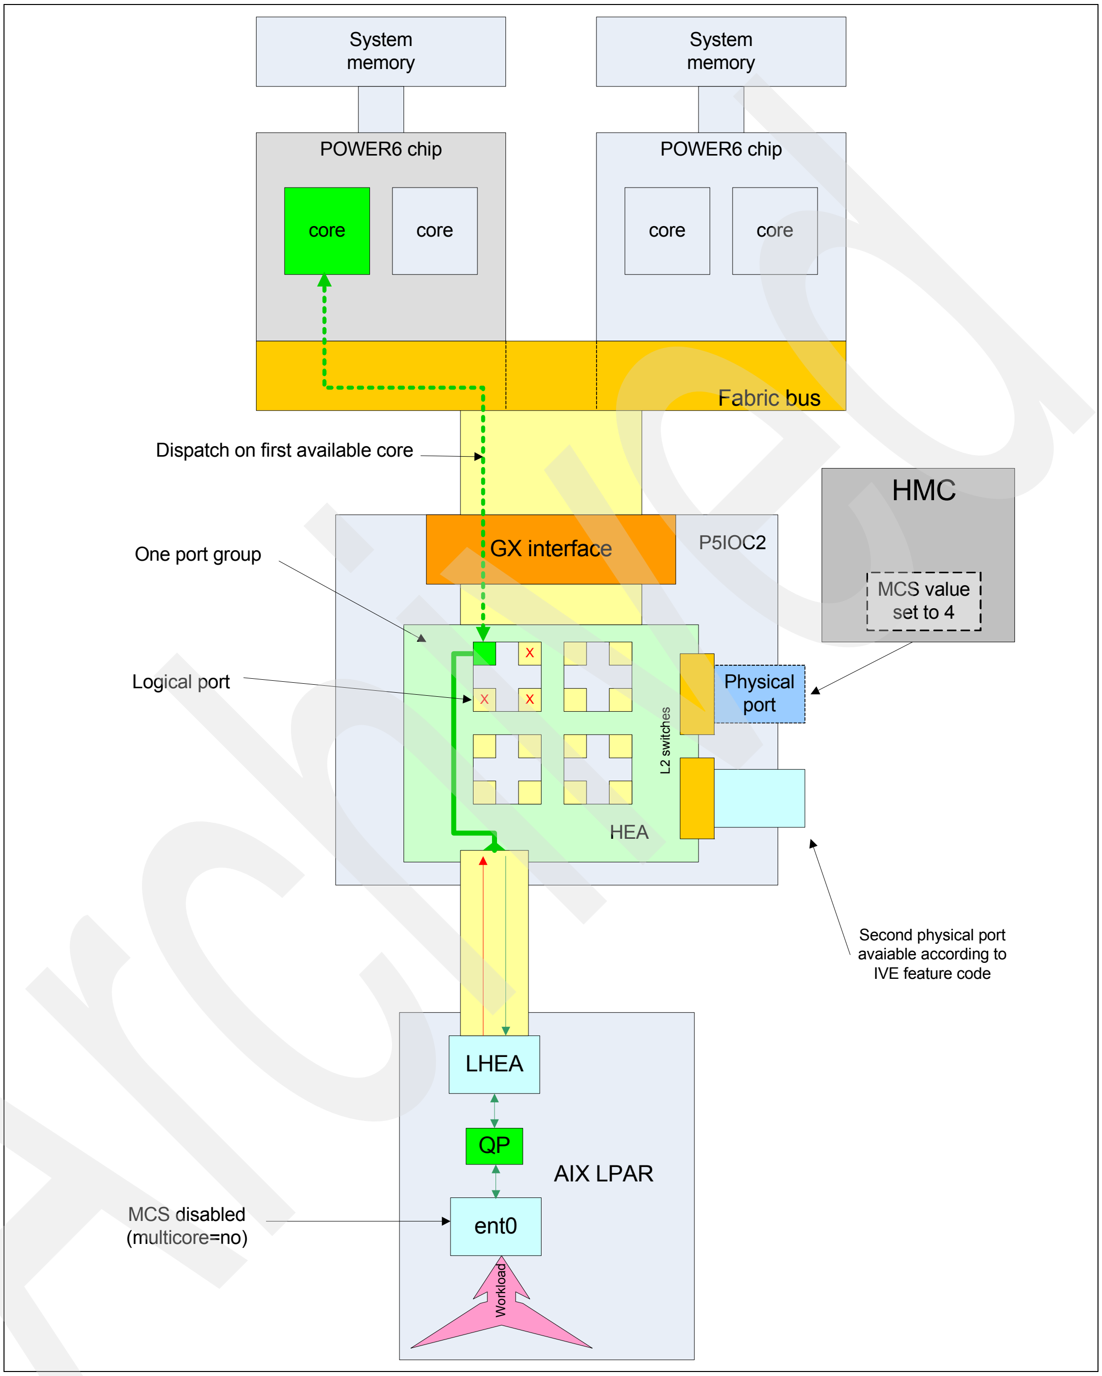

Figure 2-11 QP dispatch with MCS set to 4 but MCS disabled into AIX LPAR

The following are methods for chaging the MCS:

- On the HMC, changing the MCS value for a port group requires the managed system to be shut down and rebooted to put the new MCS values into effect on the firmware and hardware. Consider selecting the MCS as part of the initial system configuration.
- On AIX, changing the multicore attribute from yes to no (or no to yes) is dynamic and does not require the LPAR to be rebooted.

See also 4.2.3, “Performance considerations and recommendations” on page 57.

MCS allows loads to run at near wire speed, and considering the 10 Gbps IVE option, this is a significant advantage in throughput. One 10 Gbps IVE feature allows you to replace ten 1 Gbps Ethernet controllers, freeing up adapter slots and reducing complexity and IT infrastructure costs, and improving network utilization.

2.3.5 LPAR to LPAR internal communication

All IVE resources and functions can be allocated and enabled per partition. The same mechanisms are used to provide inter-partition protection and isolation.

IVE keeps the overall system solution simple (in particular, it provides transparency for software) and provides high performance. All IVE hardware acceleration and protection are available for partition-to-partition communication.

All the logical ports associated with the same physical port allow the partition-to-partition communication. LPARs that configure logical ports on the same physical ports can use internal LPAR-to-LPAR communications. This is of particular importance to the IVE features (FC 5636 and FC 5639) that have two physical ports associated with the same port group. With these features, LPARs that have an intercommunication requirement should use the same physical port.

See 3.1, “Hardware requirements and planning” on page 32, for additional information.

3

Integrated Virtual Ethernet planning and requirements

IVE presents many configuration options that are important for LPAR network connectivity and overall system performance, and the careful consideration of these options is important for the best infrastructure and solutions plans. This chapter explores the minimum requirements for IVE for the System p hardware and its system management, operating systems, and for the client browsers used for connecting to the HMC.

3.1 Hardware requirements and planning

IVE presents many configuration options to be considered as part of your planning efforts. This section describes some connectivity examples and the relationships between physical components and logical components to help you decide what is best for your implementation.

3.1.1 Feature code overview

Every IVE feature code offers a different logical layout of the logical ports, physical ports, and port groups. The examples in the previous section highlight the major differences between the IVE feature codes for use in system planning. This differences are listed below:

- FC 5636
  - 1 Gbps port speed
  - Supports one port group (up to 16 logical ports) with two physical ports
  - Provides two RJ-45 Ethernet connections
  - Provides two system ports, ports 1 and 2[1]
- FC 5639
  - 1 Gbps port speed
  - Supports two port groups (up to 32 logical ports, 16 per port group) with two physical ports per port group
  - Provides four RJ-45 Ethernet connections
  - Provides one system port, port 2
- FC 5637
  - 10 Gbps port speed
  - Supports two port groups (up to 32 logical ports, 16 per port group) with one physical port per port group
  - Provides two short range (SR) optical Ethernet Connections
  - Provides one system port, port 2[1]

3.1.2 Feature code 5636 2-port 1 Gbps IVE adapter

Figure 3-1 shows the FC 5636 (Integrated 2-port 1 Gbps virtual Ethernet) that offers two physical ports. The logical representation of the feature code (right side) shows that both ports belong to the same port group. FC 5636 supports up to 16 logical ports (if the MCS value is set to 1) per port group, and if 15 logical ports are assigned to LPARs through the first physical port, then you can assign only 1 logical port to the second physical port. To state this another way, you may configure a mix of the 16 logical ports across the two physical ports.

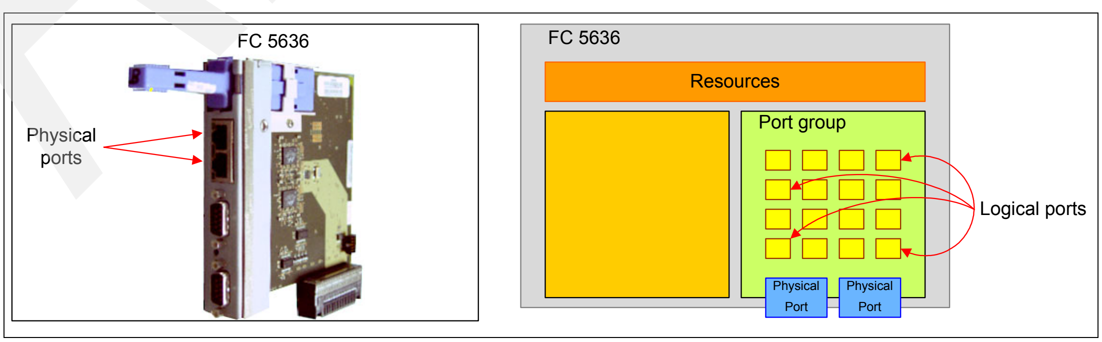

Figure 3-1 FC 5636 physical ports, port group, and logical ports representation

[1] Ports used for serial diagnostics, configuration, and UPS

3.1.3 Feature code 5639 4-port 1 Gbps IVE adapter

Figure 3-2 shows FC 5639 (Integrated 4-port 1 Gbps virtual Ethernet) that offers four physical ports. There are two port groups and any group includes two physical ports. Also with this feature code, every port group supports up to 16 logical ports. In the example in Figure 3-2, since the MCS value for the first port group is set to 8, the first port group offers two logical ports. If both logical ports are assigned to the same physical port, then the other physical port cannot have any logical port assigned. In our example, each logical port of the first port group of FC 5639 is assigned to a different physical port.

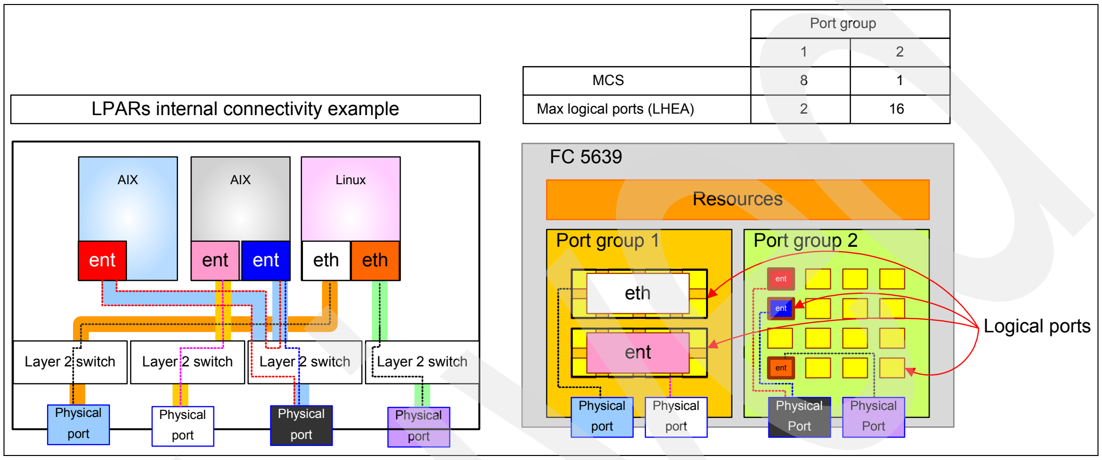

Figure 3-2 FC 5639 physical ports, port groups, logical ports representation and connectivity example

All the logical ports belonging to the same physical ports can communicate with each other through the IVE layer 2 switch associated to any physical port. The LPAR's internal connectivity example associated to FC 5639 (left side of Figure 3-2) shows that the Ethernet interface of the first AIX LPAR can communicate internally with the second Ethernet interface of the second AIX LPAR. This is possible because both Ethernet interfaces are defined through IVE logical ports associated to the same physical port.

In the same logical diagram example, instead, the Linux LPAR cannot communicate internally with the second AIX LPAR because, even if the physical ports used belong to the same port group, they are not sharing the same layer 2 switch. It is necessary to connect both physical ports to an external bridge to allow the communication between the Linux LPAR and the second AIX LPAR.

3.1.4 Feature code 5637 2-port 10 Gbps IVE adapter

FC 5637, the Integrated 2-port 10 Gbps (SR) virtual Ethernet shown in Figure 3-3 on page 34, offers two physical ports, but every physical port is associated to its port group. Therefore any port group supports up to 16 logical ports per any physical port, depending on the MCS value associated with the port group. In our example, the first port group MCS value is set to 1 and the feature offers 16 logical ports through the first physical port. The second port group MCS value is set to 4 and the feature offers four logical ports through the second physical port.

In the LPAR's internal connectivity example related to FC 5637 (left side of Figure 3-3 on page 34), the first AIX LPAR can communicate internally with the Linux LPAR because both LPARs have a logical port assigned that belongs to the same physical port. But the first AIX LPAR cannot communicate with the second AIX LPAR, unless there is an external switch configured to allow the connectivity between them, since they do not have their logical ports on the same physical port.

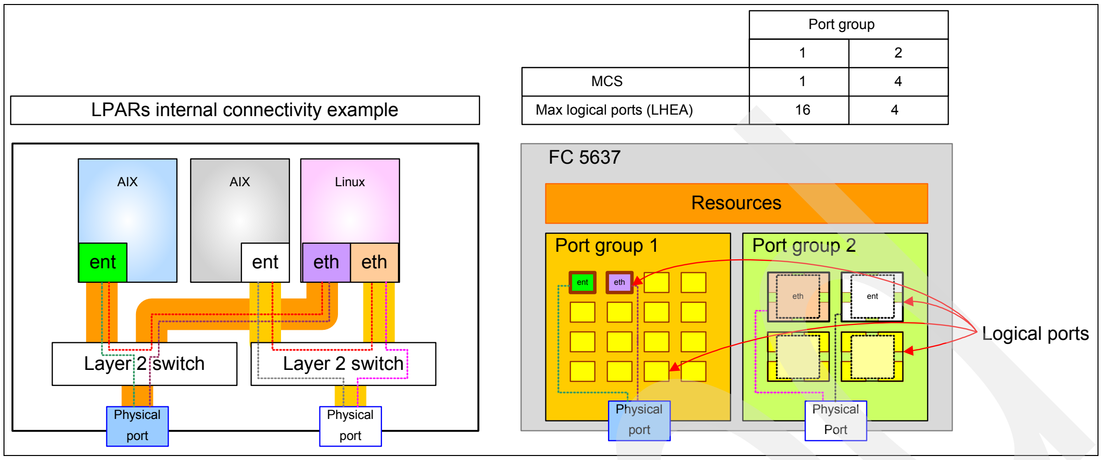

Figure 3-3 FC 5637 physical ports, port groups, logical ports representation and connectivity example

3.1.5 Enhanced availability

This section discusses methods for concurrent maintenance and additional hardware configurations to enhance serviceability.

Configuration to improve network availability

The System p product line offers many options to build a server infrastructure that is capable of reducing all the unexpected and unscheduled events faced by production systems. IVE feature codes offer different configuration possibilities that provide Ethernet interface redundancy to LPARs. In the case of a System p 570, a best practice is to build the logical partitions logical infrastructure based on two or more IVEs. Figure 3-4 shows just one of the possible configurations to improve the network availability using different operating systems with FC 5637 (Integrated 2-port 10 Gbps SR optical virtual Ethernet). This example improves network availability for the LPARs, and applies to all the other IVE feature codes as well.

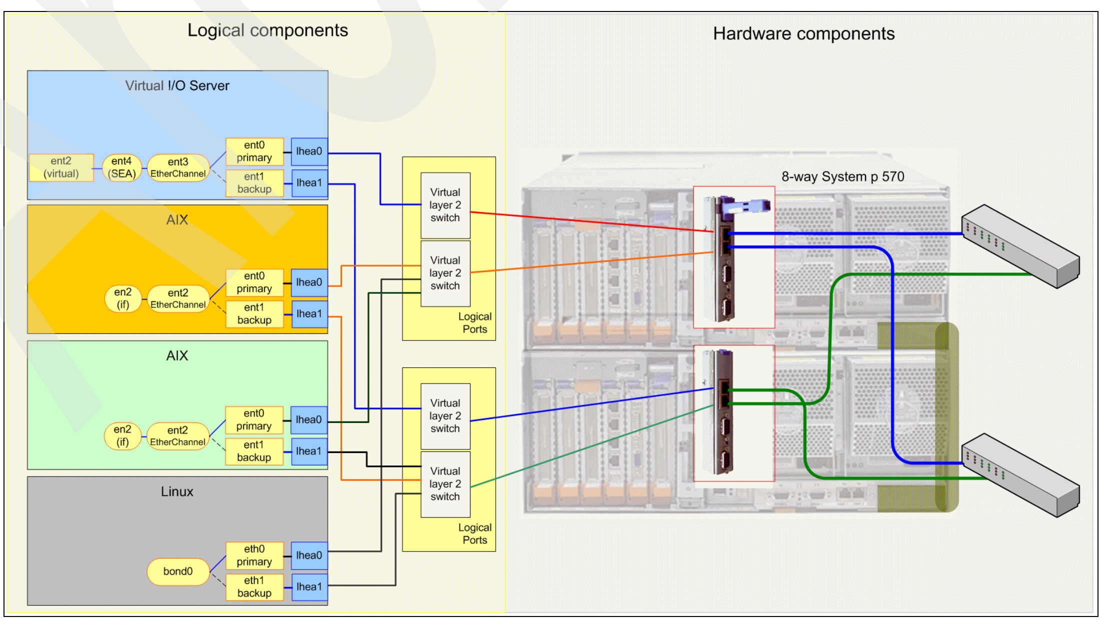

Figure 3-4 Example of logical connections to improve network availability

A Virtual I/O Server LPAR requires the IVE physical port to be set to promiscuous mode if that port will be used as the physical port in an SEA configuration. Therefore two physical ports set to promiscuous mode are required to improve the network connectivity of the SEA, and any physical port should be assigned from different IVEs located in different enclosures. In this case, the two LHEAs could also be used to make an EtherChannel associate with the Shared Ethernet Adapter.

An AIX LPAR does not require a physical port set to promiscuous mode (unless it is required to be used by packet sniffer applications, or for another reason), and it is possible to make an EtherChannel with two or more LHEAs, but with one LHEA per physical port. In the case where an LHEA is used to make an EtherChannel within an LPAR, then all the other LHEAs provided through the same physical port may also be used to create EtherChannel interfaces in other LPARs.

A Linux LPAR requiring Ethernet interface aggregation has the same requirements and the same behavior as an AIX LPAR.

3.1.6 System Planning Tool for IVE

The IBM System Planning Tool (SPT) is available to assist you with system planning, design, validation, and to provide a system validation report that reflects your system requirements while not exceeding system recommendations. SPT is a PC-based browser application designed to be run in a standalone environment. System plans generated by SPT can be deployed on the system by the Hardware Management Console (HMC) and Integrated Virtualization Manager (IVM). At the time of writing, the last version of SPT can be downloaded from the following Web site:

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

With regards to IVE, SPT assists in planning the IVE logical ports or IVE physical ports to assign to the planned LPARs.

Figure 3-5 shows an SPT panel while planning a System p 570 with four LPARs.

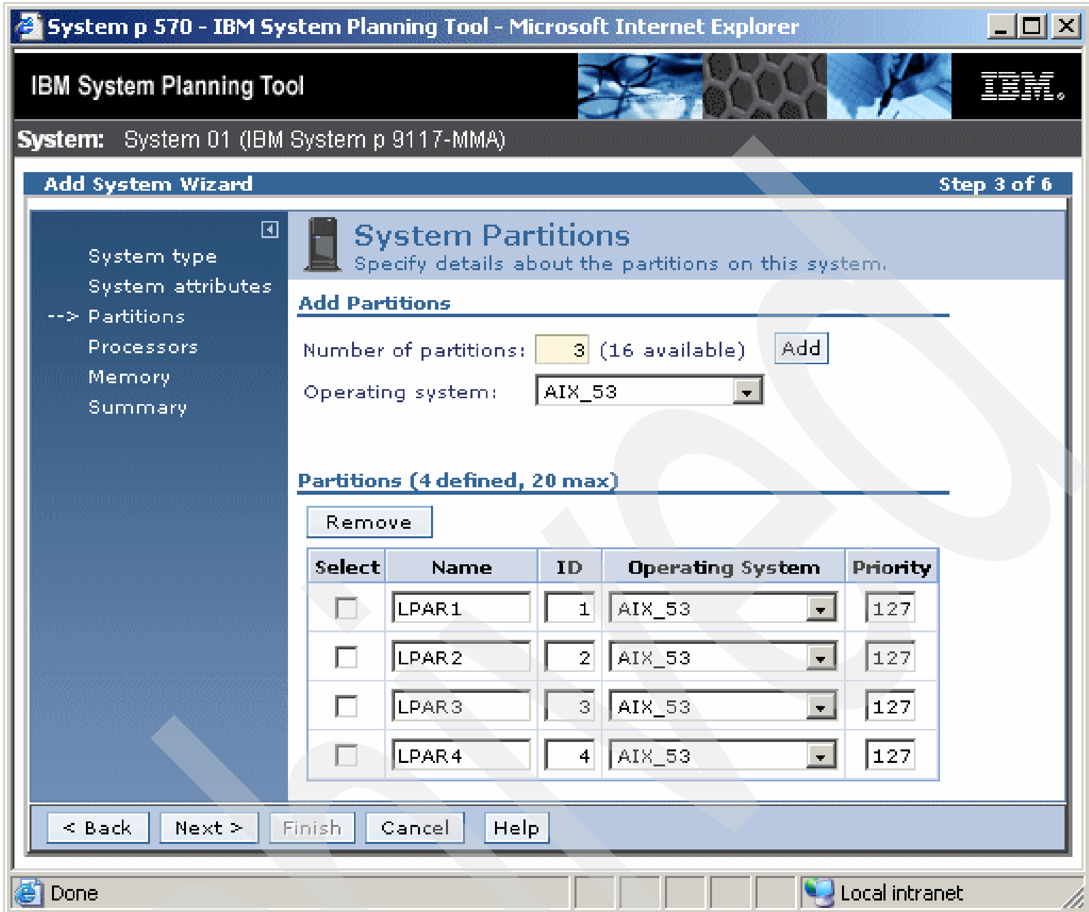

Figure 3-5 Planning a System p 570 with four LPARs

Once SPT has collected all the necessary information about the processor feature code and cores activation, or the amount of memory, and validated it against the LPAR requirements, then SPT asks to enter edit mode and set up the I/O resources. See *IBM System i and System p*, SG24-7487, for additional information about the System Planning Tool.

In this example, four LPARs require four dedicated Ethernet adapters to have network connectivity. Alternatively, you could plan to add an additional LPAR and use it as Virtual I/O Server (VIOS). Assigning the proper Ethernet device to the VIOS LPAR allows any LPAR to have network connectivity. IVE provides the capability to share its physical ports without a VIOS to offer network connectivity for any LPAR.

As shown in Figure 3-6, SPT set as a default one FC 5636 (Integrated 2-port 1 Gbps virtual Ethernet) for any System p 570 enclosure (the enclosure is one in our example). In order to change the IVE feature code if the default adapter is not the one planned, it has to be removed and replaced with the correct feature code. Click the button just beside the feature code to delete it.

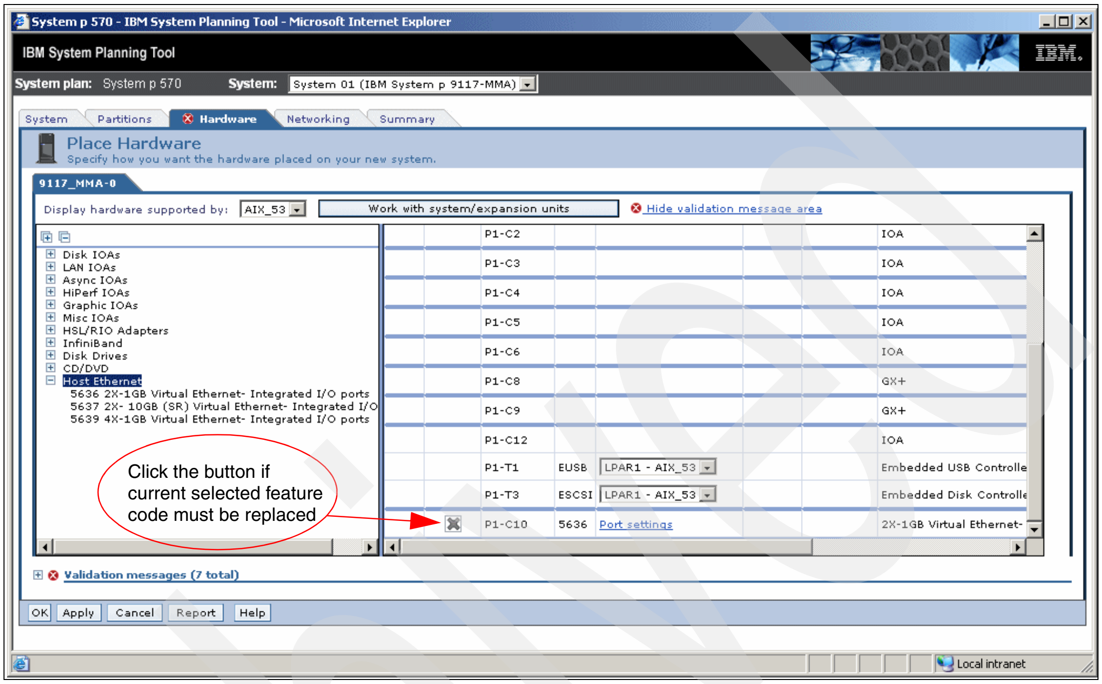

Figure 3-6 Example to delete IVE feature code

After the feature code is removed, the correct IVE feature code must be selected. Then click the button to add it in the list, as shown in Figure 3-7.

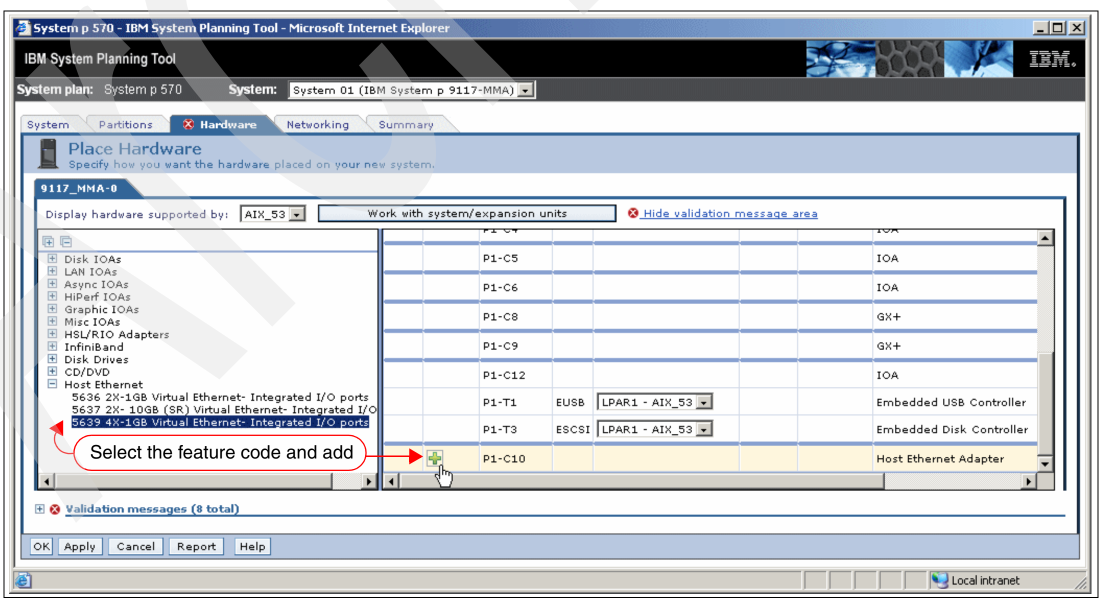

Figure 3-7 Example to add IVE feature code

When the planned IVE feature code is selected, the *port settings* option is available (see Figure 3-8) and must be selected.

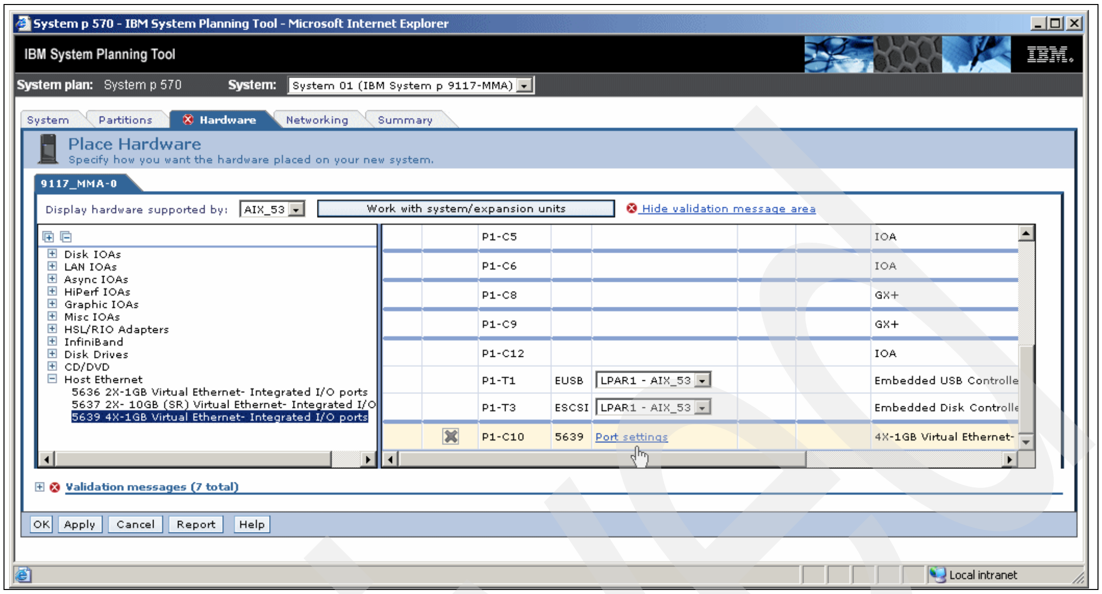

Figure 3-8 IVE feature code selection requires ports setting

According to the selected IVE feature, different numbers of physical ports and port groups appear in the port settings menu (see Figure 3-9).

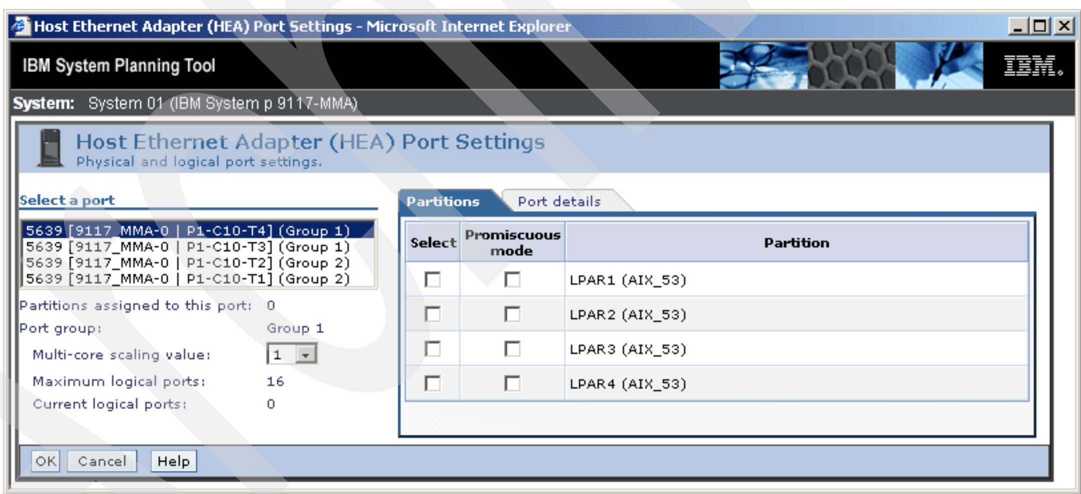

Figure 3-9 logical port setting menu

From the port settings menu, logical ports or a whole physical port can be assigned to the LPARs.

If a physical port is assigned to an LPAR in *promiscuous mode*, then the other LPARs cannot use it.

Note: A physical port must be configured in promiscuous mode to support the Virtual I/O Server's Shared Ethernet adapter. This would allow the Virtual I/O Servers virtualization capabilities to exceed the 16 LPARs per physical port maximum of the IVE.

Figure 3-10 shows the first physical port in the list assigned to the first LPAR in promiscuous mode. All other LPARs cannot use the same port, and the selection becomes greyed out.

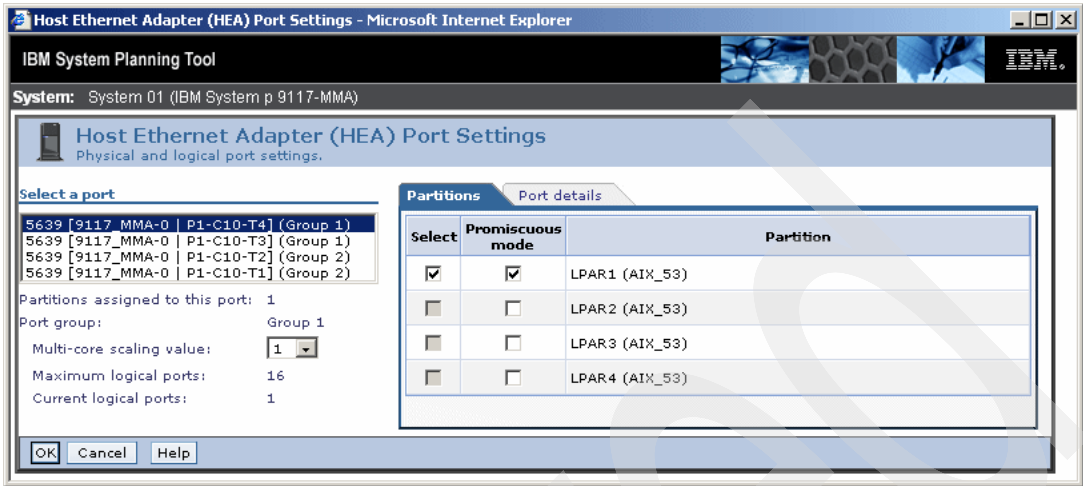

Figure 3-10 Assigning one physical port to LPAR in promiscuous mode

If a physical port is not assigned to an LPAR in promiscuous mode, then other LPARs may use this physical port. So, the physical port provides a logical port for any LPAR selected. Figure 3-11 shows all four LPARs that can be selected to use one logical port from the first physical port.

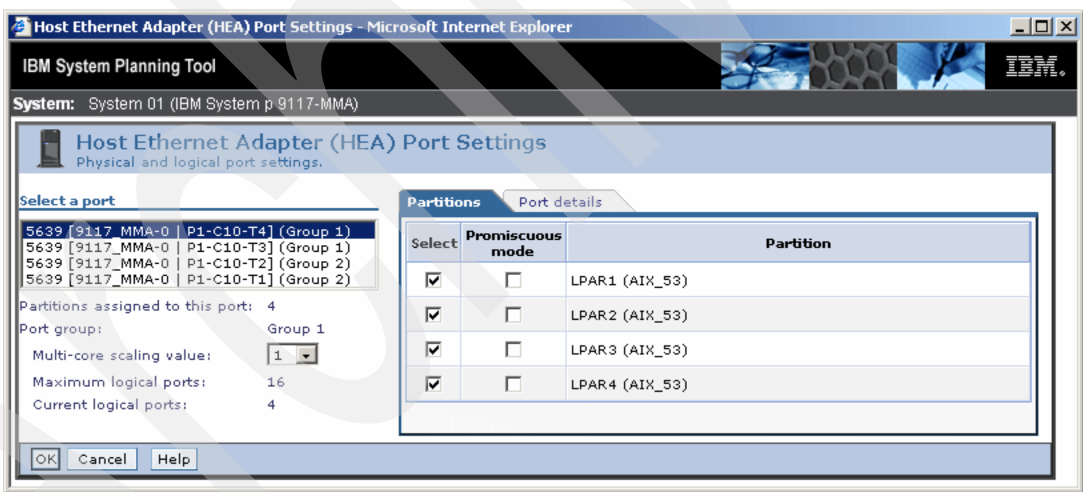

Figure 3-11 Assigning one physical port to four LPARs

Any physical port can provide up to 16 logical ports if the Multi-Core Scaling value (MCS) is set to 1 (see 2.2.2, “IVE and physical, logical, and port groups” on page 15).

In our example, FC 5639 (Integrated 4-port 1 Gbps virtual Ethernet) provides four physical ports (grouped in two port groups). Thus, the total number of logical ports available can be up to 32. However, any LPAR can have only one logical port assigned from any physical port available.

> **Note:** If an LPAR requires two logical ports to define two Ethernet interfaces at the operating system level, the logical ports must be from different physical ports, but they may be from the same port group.

SPT assists in the assignment of the correct number of physical ports or logical ports according to the system plan. Since a specific physical port belongs to a port group, it is possible to change the MCS value associated with a physical port, but it also affects the other physical port belonging to the same port group. SPT alerts you if you attempt to assign more logical ports than what the MCS value of the port group allows (see Table 2-1 on page 15). So, for example, if the MCS value is greater than 4, which allows fewer than 4 logical ports, then you cannot assign 4 or more logical ports to LPARs. In Figure 3-12, the MCS value is 8, 4 LPARs are selected, and SPT highlights the fact that the current logical ports are not enough to satisfy the four LPARs. The SPT ports setting menu highlights that the current logical ports are not enough to satisfy all four LPARs.

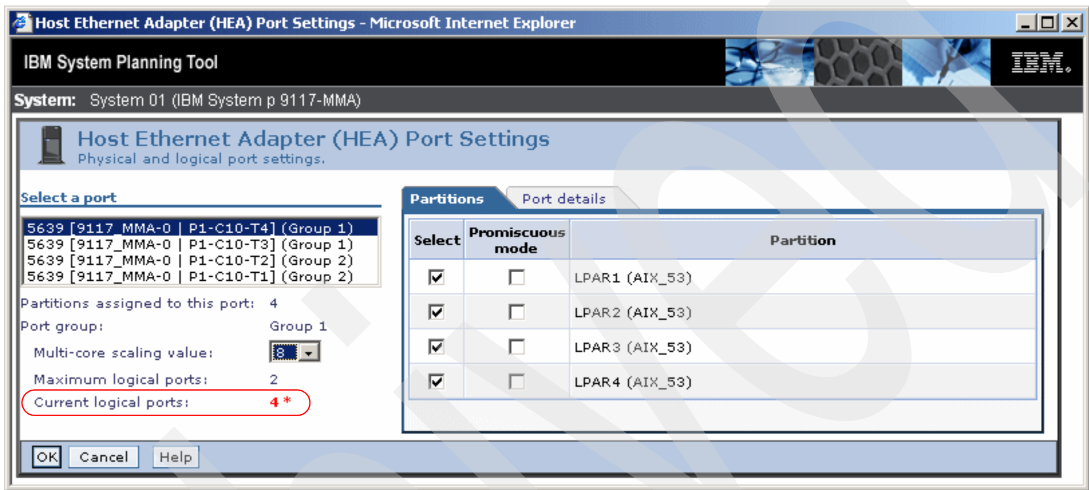

Figure 3-12 Warning message to highlight current logical ports are not sufficient

Still from this example, when the MCS value is set to 8, two LPARs can have the two available logical ports assigned. This is an acceptable combination, as shown in Figure 3-13.

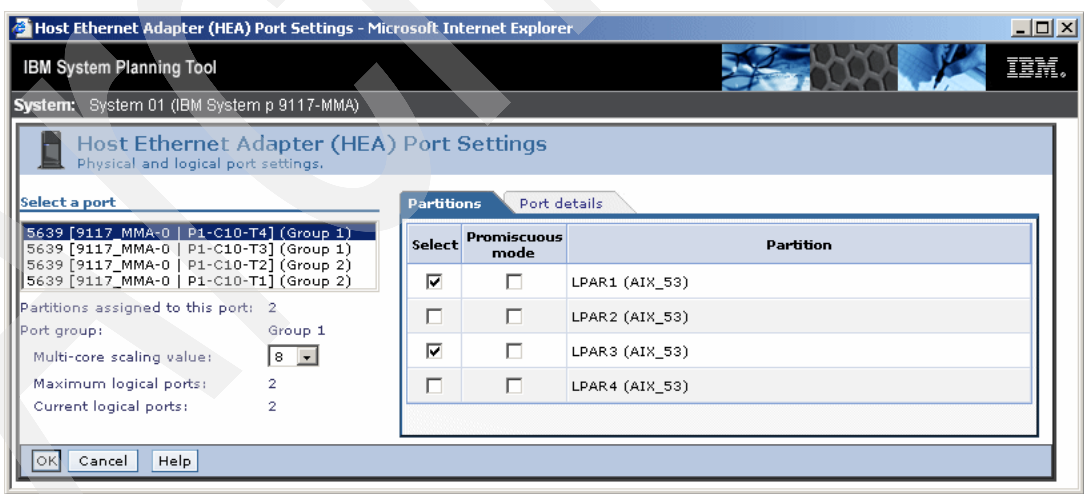

Figure 3-13 Port settings allowed combination

SPT quick reference

SPT helps you design a system to fit your needs. You can:

- Create a new system configuration.
- Create a system configuration based upon any performance data from an existing system, or performance estimates that anticipate future workloads.
- Determine what you must order to support your needs.
- Determine how you can partition a system that you already have.
- Save your system plan and import it into a Hardware Management Console, deploying the system plan to one or more systems.
- Make a report of the actual server configuration.

System plan deployment delivers the following benefits to you:

- Saves time and reduces the possibility of error
- Deploys multiple, identical systems almost as easily as a single system
- Archives the system plan as a permanent electronic record of the systems that you create in case you must recreate the configuration as part of a system recovery plan

In addition, with regards to IVE, SPT assists you in:

- Planning IVE physical ports
- Planning IVE logical ports
- Planning performance parameters

3.2 System management requirements

This section discusses POWER6 processor-based system management requirements related to IVE.

3.2.1 Hardware Management Console requirements

Hardware Management Console is mandatory for System p 570, and it is required to configure and use any IVE. At the time of writing, the HMC minimum level is code Version 7 Release 3 (HMC V7R3).

HMC V7R3 is the minimum HMC level to support any POWER6 processor-based system, but it also supports POWER5 and POWER5+ processor-based systems with firmware release at level SF240_299 or later.

In HMC V7R3, there are significant changes to handle the managed server LPARs. The configuration of the IVE (or HEA) for your POWER6 processor-based server has a separate configuration area within HMC.

To get to the Host Ethernet adapter area, select **Systems Management** → **Servers** → *server name*, then select **Hardware Information** → **Adapters** → **Host Ethernet**. See 4.1.1, “Configure an HEA with an HMC” on page 46.

For more detailed information about HMC, refer to *Hardware Management Console V7 Handbook*, SG24-7491, or ResourceLink:

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

https://www-01.ibm.com/servers/resourcelink/lib03030.nsf/pages/pHardwarePublicationsByCategory?OpenDocument&pathID=

3.2.2 Browser requirements for HMC remote connection

The HMC V7R3 uses a new framework. A major change is the Web browser-based user interface. This means that you do not have to install an application to access the HMC remotely, and you can connect to the HMC via your browser. Firefox and Internet Explorer® are supported.

3.2.3 IBM Director

IBM Director is an integrated suite of tools that provides you with a system management solution for heterogeneous environments, including IBM System p environment. IBM Director works with the HMC to provide a comprehensive system management solution.

IBM Director can be installed into a System p LPAR that uses an LHEA as the network interface, or IBM Director installed in a different IBM Systems platform can discover an LPAR using the LHEA as network interface.

At the time of writing, IBM Director 5.20 is available. For more information refer to:

- *IBM Director on System p5™*, REDP-4219-00
- IBM Director home page

  http://www-03.ibm.com/systems/management/director/
- *Implementing IBM Director 5.20*, SG24-6188

3.3 Operating system requirement

This section discusses operating system requirements related to IVE.

3.3.1 AIX

Since IVE is a set of standard features for any POWER6 processor-based server, the minimum AIX operating system levels are the minimum supported by the hardware platform:

- IBM AIX 5L™ Version 5.3
  - Service Pack 5300-06-02.
  - Packaging APAR for this package is IY99738.
  - IVE fileset devices.chrp.IBM.lhea.rte must be at level 5.3.0.2 or later.
- IBM AIX 5L Version 5.2
  - Service Pack 5200-10-02.
  - Packaging APAR for this package is IZ00822.
  - IVE fileset devices.chrp.IBM.lhea.rte must be at level 5.2.0.107 or later.

Access the following IBM public Web site for more information about the latest AIX fixes:

http://www14.software.ibm.com/webapp/set2/sas/f/genunix3/aixfixes.html

Note: Both AIX5L Version 5.3 and Version 5.2 support IVE. Consider the following additional information:

- AIX 5.2 LPAR requires the IVE physical port assigned with promiscuous mode enabled.
- AIX 5.3 is the target platform for performance requirements.

3.3.2 Linux on POWER

The following levels are required for support of IVE on IBM System p servers:

- Red Hat RHEL 4.5 or later
- SuSE SLES 10 SP1 or later

The following text applies to both distributions: Dynamic LPAR support for HEA adapters and logical adapters is not supported at the time of writing.

4

How to configure and work with Integrated Virtual Ethernet

This chapter describes the IVE functions and attributes, and provides directions on how to configure and use it.

You will find suggestions on how to work with IVE regarding:

- Live Partition Mobility
- The MCS value
- Flow control
- Hardware placement and interfaces numbering
- Performances considerations
- Major settings modification summary

4.1 Configure an HEA and an LHEA

The following sections describe the steps needed to configure an HEA and an LHEA. For the p570 you must use an HMC for configuration, as discussed in the following section.

4.1.1 Configure an HEA with an HMC

This section describes the steps for HEA configuration using both GUI and CLI on the HMC.

Take note: Plan carefully.

- The MCS value setting requires a restart of the managed system to make the change effective. It is a best practice to determine whether you need to change the default setting during your initial system configuration.
- Promiscuous mode reserves the physical port set with this attribute to the given LPAR. No other LPAR will be able to have a logical port on this physical port.

See also 3.1.6, “System Planning Tool for IVE” on page 35.

You can print these steps and use the check boxes to record your progress.

- To configure the HEA, your system must be LHEA capable:
  - ☐ Within the servers panel, select your managed system and use the **Tasks** button to find **Properties**, then choose the **Capabilities** tab, as shown in Figure 4-1.
    - If the following HMC command line command prints a 1, this indicates that the server is LHEA capable:

      ```
      lssyscfg -r sys -m <managed system name> -F lhea_capable
      ```

      If it is LHEA capable, the output will print a 1, 0.

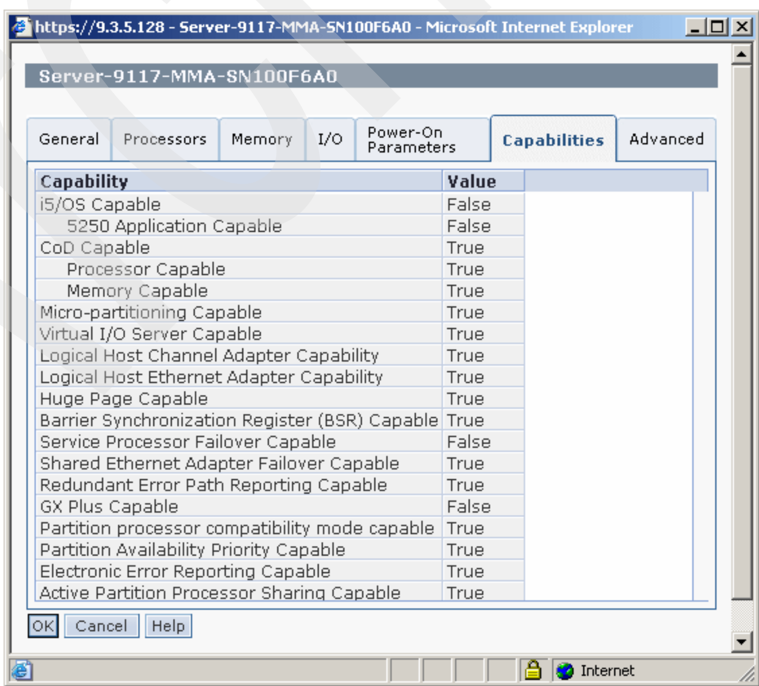

Figure 4-1 LHEA capability

- In the content of your managed server panel:
  - ☐ Select the check box of your managed system.
  - ☐ Select the **Tasks** button.
  - ☐ Go through **Hardware Information** → **Adapters** → **Host Ethernet**.
- On the new Host Ethernet Adapter window (as shown in Figure 4-2):
  - ☐ Select in the pull-down list the adapter to configure.
    - In in HMC command line use the following command to list the available adapters and ports:

      ```
      lshwres -r hea --rsubtype phys --level sys -m <managed system name>
      ```

      and choose in the display the adapter you want to configure.
  - ☐ Select the physical port to configure into the port group.
    - In HMC command line use:

      ```
      lshwres -r hea --rsubtype phys --level port_group -m <managed system name
      > --filter adapter_ids=<adapter id previously choosen>
      ```

      to display and then choose the port group to configure into the adapter:

      ```
      lshwres -r hea --rsubtype phys --level port -m <managed system name> --fi
      lter port_groups=<port group previously chosen>
      ```

      to choose the physical port to configure into the port group
  - ☐ Click the **Configure** button.

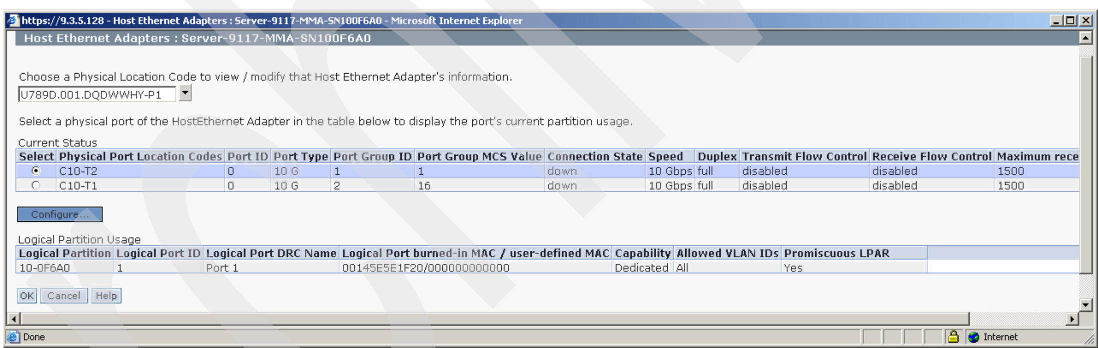

Figure 4-2 HEA configuration first step

- In the new HEA Physical Port Configuration window, as shown in Figure 4-3 on page 49:
  - ☐ Choose speed, duplex, and packet size according to your network infrastructure.

> **Note:** Considerations:
> - A 10 Gbps adapter cannot be set to 10 Mbps, 100 Mbps, or 1 Gbps.
> - You cannot use jumbo frames at 10 Mbps or 100 Mbps speed.
> - Half duplex is not supported by the IVE.

- Select the **Flow control enabled** check box.

> **Note:** IBM recommends turning on this control, which will enable the use of pause frame on the physical port. It is useful if network congestion occurs on the Ethernet link, to prevent packets from being dropped. For performance it is also recommended, especially for a 10 Gbps adapter, to let it reach the full media speed.

- Choose the Pending Port Group Multi-Core Scaling value (MCS) for this physical port (this change affects the port group and each physical port in the port group). We recommend the value 4 for a 10 Gbps adapter and 1 for a 1 Gbps adapter. See also 4.2.3, “Performance considerations and recommendations” on page 57.

> **Note:** MCS is the only function that affects all physical ports inside a port group. If the MCS value choosen is 16, you will have only one logical port allowed. Then you must assign logical port number 1 for an LPAR when configuring the LHEA. For the definition, see 2.2, “Logical and physical components” on page 13.

- Choose Promiscuous LPAR if needed (for details on this mode see 2.3.1, “Physical port promiscuous mode” on page 20).

> **Note:** Promiscuous mode can be used if you want to share this physical port between more than 16 LPARs, using an SEA through a VIOS. After you have set this mode on the HMC, the SEA automatically requests promiscuous mode to the device driver.
>
> With this mode on, you cannot set another LPAR to use an LHEA on this physical port.

☐ When all the settings reach your needs, click the **OK** button to complete the HEA configuration (Figure 4-3).

In HMC command line use:

```
chhwres -o s -m <managed system name> -r hea -l <adapter id previously choos
en> -g <port group previously choosen> --physport <physical port previously
choosen> -a “flow_control=1,pend_port_group_mcs_value=4,conn_speed=10000,dup
lex=auto,max_recv_packet_size=1500,promisc_lpar_name=none”
```

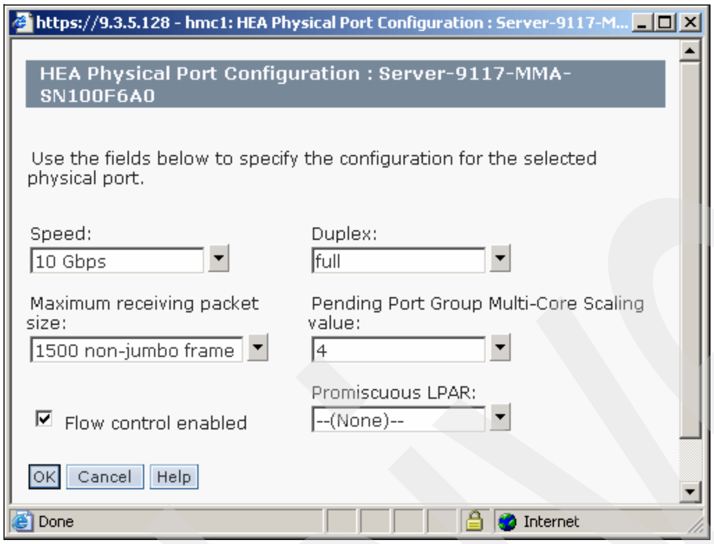

Figure 4-3 HEA configuration second step

4.1.2 Configure an LHEA port with an HMC

When you create a new LPAR using the Create Lpar Wizard, a panel allows you to configure the LHEA. For an existing profile, use the LHEA tab of the profile properties. These steps can also be performed using dynamic logical partitioning.

Use the steps below to configure it:

- In the Create Lpar Wizard window (as shown in Figure 4-4):
  - ☐ Select the HEA previously configured in the pull-down list.
  - ☐ Select the physical port you want to use.
  - ☐ Leave the LHEA Capability at the default setting of Base Minimum. At the time of writing, this function is not relevant.
  - ☐ Click the **Configure** button.

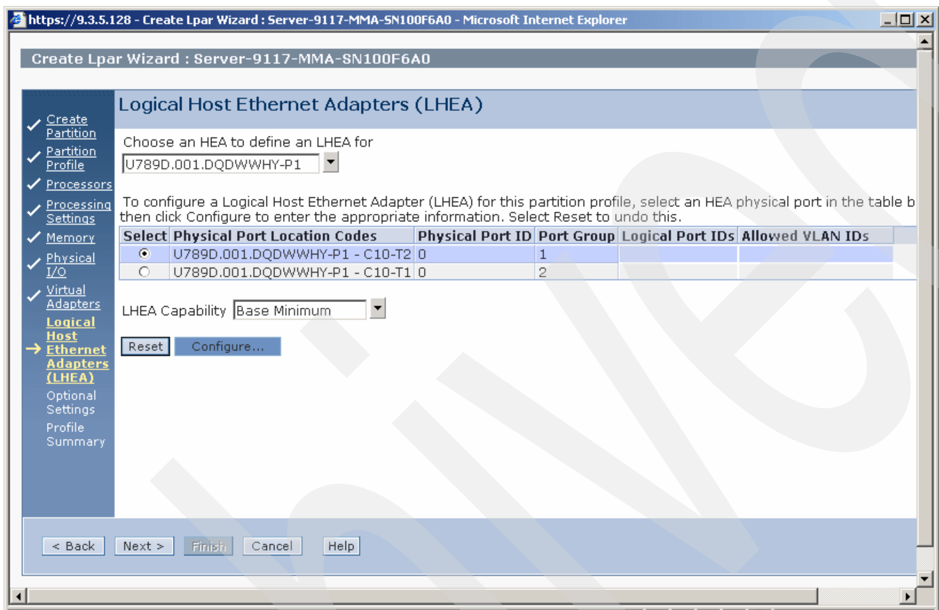

Figure 4-4 LHEA first step configuration

- In the Logical Host Ethernet Adapter Configuration window (shown in Figure 4-5 on page 52):
  - ☐ Select the logical port you want to use.

> **Tip:** At the time of writing, to perform a NIM install you need to have untagged packets. The IVE layer 2 switch is not VLAN aware and will not strip the port VLAN ID (PVID). An external L2 switch, virtual Ethernet and SEA, or a LHEA that is not part of a VLAN are possible solutions. Therefore, you can create an LHEA without VLAN and use it to perform the NIM install. After the OS is installed, remove it with dynamic LPAR. Alternatively, you can also set up a virtual interface through a VIOS for use with NIM, or you could configure an external switch to strip the VLAN tags.

  If you plan to use VLAN's:

  - ☐Check the **Allow all VLAN IDs** box if you want all VLAN's (4096) to be allowed for this LHEA.
  - ☐If needed, enter the appropriate VLAN IDs in the VLAN to add field, clicking the **Add** button after each one. You can add a maximum of 20 unique VLAN IDs.
    - ☐Click the **Add** button if appropriate (adding manually VLANs).
  - ☐ Click the **OK** button.

  This completes the LHEA port configuration step when performed from the HMC.

– In the HMC command line for LPAR profiles use:

mksyscfg -r lpar For the first profile creation

mksyscfg -r prof For an additional profile creation

chsyscfg -r prof To modify an existing profile

All commands will support an attribute lhea_logical_ports, whose value is a list of strings of the format:

HEA-DRCIndex/port-group-ID/physical-port-ID/logical-port-ID/allowed-VLAN-IDs allowed-VLAN-IDs can be a list of 20 or less VLAN IDs.

In this case the value should be enclosed in quotation marks (“”). This attribute can have values none or all, or it may be omitted. lhea_logical_ports itself may be set to none.

These two commands also support another new, corresponding attribute, lhea_capabilities, whose format is a list of strings of the form:

```
HEA-DRCIndex/capability/interruptible-EQs/non-interruptible-EQs/QPs/CQs/MRs
```

The valid values for capability are:

- 0 - Base Minimum (default, also in case lhea_capability is not given)
- 1 - Low
- 2 - Medium
- 3 - High
- 4 - Dedicated
- 5 - Custom
- chsyscfg supports the += and -= syntax

For additional information see the HMC man pages, accessible from the HMC or on the following Web site:

http://www14.software.ibm.com/webapp/set2/sas/f/hmc/power6/related/hmc_man_7310.pdf

- Before creating a new LPAR with LHEA ports, you must determine the unused and available (defined by the MCS value) ports on the physical port. The output of the following command lists these for the physical port you specify:

```
lshwres -r hea --rsubtype logical --level port -m <managed system name>
--filter "adapter_ids=<adapter id previously chosen>,port_groups=<port
group previously chosen>" -F phys_port_id:logical_port_id:state:lpar_id
--header | grep -v none: | grep -E "none|id"
```

You will find into the output, unassigned and available logical ports for this HEA inside the port group choosen. You can then create your LPAR with:

```
mksyscfg -r lpar...-i "...,...,..."
```

Include the following code inside the double quotation marks regarding to the LHEA:

```
lhea_logical_ports=<adapter-ID/port-group/physical-port-ID/logical-port-ID> lhea_capabilities=<hea previously choosen>/0
```

Multiple LHEAs may be added this way, separated by commas.

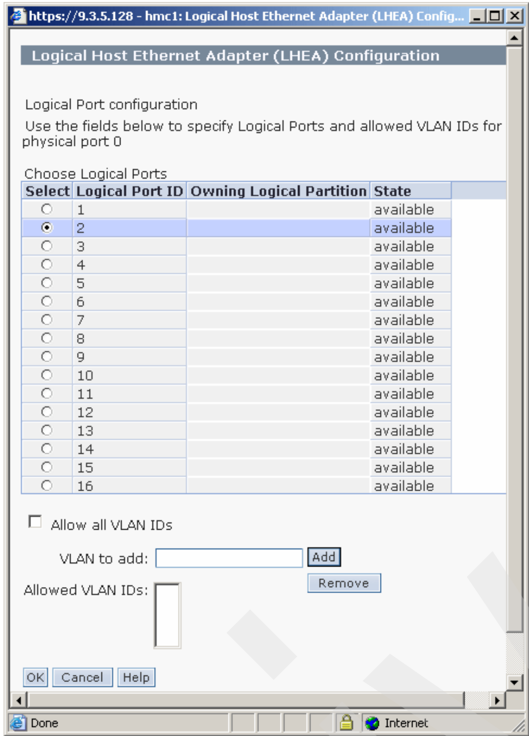

Figure 4-5 LHEA second step configuration

- Verification in the Create Lpar Wizard window, as shown in Figure 4-6.
  - ☐ Verify that you have configured the correct logical port IDs on the correct physical port.

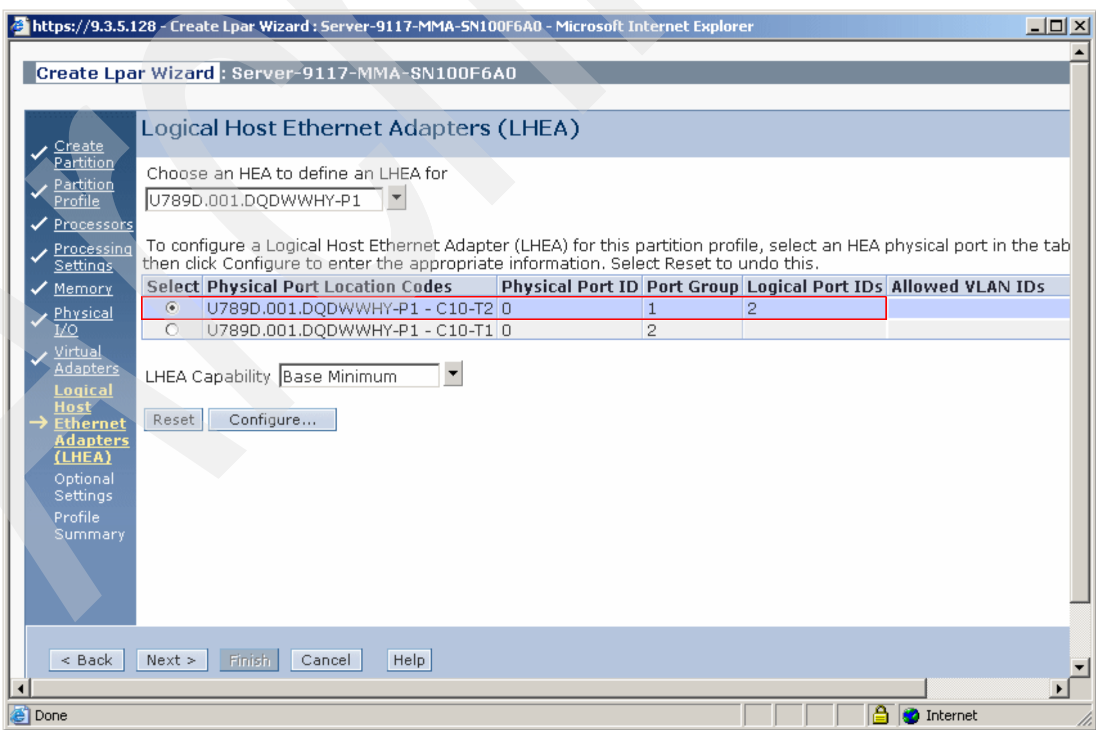

Figure 4-6 Show the new LHEA created on the HEA physical port

- In HMC command line use:

```
lssyscfg -m <managed system name> -r prof --filter "lpar_names=<lpar
name>,profile_names=<profile name>"
```

- On the HMC level, LHEA statistics can be view-only on CLI using:

```
lshwres -r hea -m <managed system name> --rsubtype phys --level port --stat
```

4.2 Configure LHEA on AIX

Review the following information prior to configuring TCP/IP on the LHEA interface:

- Partitions can own one logical port per physical port.
- From a partition, a logical port appears as a generic Ethernet interface, with its own resources and configuration.
- You cannot mix other kinds of Ethernet devices into an EtherChannel.
- You cannot have an LHEA in your LPAR if you want to use Live Partiton Mobility.

  As a work around, you can take the following actions:

  a. Add a virtual Ethernet adapter that will get its network connectivity through an SEA from a VIOS and then configure IP on it.

  b. Unconfigure the interfaces that belong to the LHEA, and remove them from your system with the **`rmdev -l lhe`*`ax`*` -R`** command.

  c. Use dynamic LPAR operation in order to remove the logical port definition from the hypervisor.

  d. If needed, you can add back an LHEA after the migration.

4.2.1 MCS enabled or disabled

According to your IT requirements, we recommend that you test and monitor the use of MCS to obtain better performance. See also 4.2.3, “Performance considerations and recommendations” on page 57.

- Attempt to use the queue pairs that match the virtual processors or dedicated processors that your LPARS can use. With AIX, the value you receive is either 1 or the number of queue pairs the MCS states.
- By default, AIX uses the number of queue pairs based on the MCS value set on the HMC.
- MCS value must be set first on the hardware level for the physical port by the HMC.
- In AIX, the default value for the multicore attribute is yes, which enables MCS. Example 4-2 on page 54 shows a multicore value of yes (enabled).

  To find which devices belong to an HEA, enter the command as shown in Example 4-1.

Example 4-1 The lsslot command for device relationship

```
# lsslot -c slot | grep -Ei “lhea|Desciption”
# Slot Description Device(s)
HEA 1 Logical I/O Slot lhea(0) ent(0) ent1
```

Example 4-2 The lsattr command shows if MCS enabled or disabled

```
# lsattr -El ent1 | grep multicore
multicore     yes             Enable Multi-Core Scaling
True
```

You can also see on the OS level the MCS value (4 in our case) set by the HMC and the attribute for each Queue Pair (QP), as shown in bold Example 4-3. The following values define those shown in the example:

QP# Index for QP structure within the driver space.

num: # QP number in decimal, which is provided/used by PHYP and driver.

bid: # Bus ID on which Device/QP resides. Used for external interrupt correlation in the error report (**errpt** command).

ISN/level: # Interrupt source number/interrupt level.

Example 4-3 The entstat command shows the MCS value used for a network interface

```
# entstat -d ent1 | grep QP
Number of Default QPNs for Port: 4
Default QPN Array:
    QP0 | num: 145 | bid: 90000300 | ISN/level: 303
    QP1 | num: 146 | bid: 90000300 | ISN/level: 304
    QP2 | num: 147 | bid: 90000300 | ISN/level: 305
    QP3 | num: 148 | bid: 90000300 | ISN/level: 306
```

- If you want to disable the MCS in AIX for a specific LPAR, change the multicore attribute from yes to no. Your interface must be unconfigured first. See Example 4-4.

Example 4-4 The chdev command to disable MCS for an individual LPAR

```
# ifconfig en1 down detach
# chdev -l ent1 -a multicore=no
ent1 changed
```

You can verify your change, as shown in Example 4-5.

Example 4-5 The lsattr command shows whether the MCS is used

```
# lsattr -El ent1 | grep multicore
multicore     no              Enable Multi-Core Scaling
True
```

A disabled MCS on AIX is equivalent to an HMC setting of MCS=1 (no matter what the currents HW value is), restricting the OS to use one queue pair number (QPNs), as shown in bold in Example 4-6.

Example 4-6 The entstat command shows the MCS value used for a network device

```
# entstat -d ent1 | grep QP
Number of Default QPNs for Port: 1
Default QPN Array:
    QP0 | num: 149 | bid: 90000300 | ISN/level: 303
```

Using flow control

At the time of writing, in MDC mode, the flow control for IVE ports is enabled by default. Once the system is partitioned, however, flow control is not enabled by default. If flow control is needed, enable it as shown in 4.1.1, “Configure an HEA with an HMC” on page 46.

Flow control must be set on the physical port by the HMC, or by the HMC and OS when in MDC mode. The setting can be viewed by the OS, as shown in bold in Example 4-7.

Example 4-7 The entstat command to view the HMC setting

```
# entstat -d ent1 | grep Pause
TX Pause Frame Negotiated: TRUE
RX Pause Frame Negotiated: TRUE
```

To disable the flow control, use the HMC (see 4.1.1, “Configure an HEA with an HMC” on page 46). That will take effect immediately, as shown in bold in Example 4-8.

Example 4-8 The entstat command to view the HMC setting

```
# entstat -d ent1 | grep Pause
TX Pause Frame Negotiated: FALSE
RX Pause Frame Negotiated: FALSE
```

4.2.2 Logical host Ethernet adapter information

Depending on the busses placement in the hardware memory map, the interface numbering can differ from one system to another. At the time of writing, the p570 uses this order:

1. LHEA (logical adapters)
2. PCI (physical adapters)
3. VIO (virtual adapters)

The HEA, LHEA, and interfaces relationship can be viewed as shown in Example 4-9 and Example 4-10 on page 56:

- For a p570 with one system drawer, as shown in Example 4-9.

Example 4-9 The lsslot command for slot and device relationship

```
# lsslot -c slot | grep -Ei "lhea|Description"
# Slot                  Description       Device(s)
HEA 1                   Logical I/O Slot  lhea0 ent0 ent1
```

The HEA 1 is the daughter card, the lhea0 is the logical adapter, and the ent0 and ent1 are the devices associated with the lhea0 and are placed on two different physical ports, as shown in Example 4-10.

Example 4-10 The lsdev command for interfaces physical placement

```
# for i in 0 1
> do
> echo "### ent$i ###"
> lsdev -l ent$i -F physloc
> done
### ent0 ###
U789D.001.DQDWWHY-P1-C10-T2
### ent1 ###
U789D.001.DQDWWHY-P1-C10-T1
```

- For a p570 with four system drawers, as shown in Example 4-11.

Example 4-11 The lsslot command for slot and device relationship

```
# lsslot -c slot | grep -Ei "lhea|Description"
# Slot                    Description       Device(s)
HEA 1                     Logical I/O Slot  lhea0 ent0 ent1
HEA 9                     Logical I/O Slot  lhea1 ent2 ent3
HEA 17                    Logical I/O Slot  lhea2 ent4 ent5
HEA 25                    Logical I/O Slot  lhea3 ent6 ent7
```

  – The HEA*x* lines in the output above represent the daughter cards, where *x* numbers are system level, unique, non-zero, and given by the firmware.

    The lhea*x* are the logical devices that owned the ent*x* interfaces.

    ent*x* hardware placement is shown in Example 4-12.

Example 4-12 The lsdev command for physical interfaces placement

```
# for i in 0 1 2 3 4 5 6 7
> do
> echo "### ent$i ###"
> lsdev -l ent$i -F physloc
> done
### ent0 ###
U789D.001.DQDVXGL-P1-C10-T2
### ent1 ###
U789D.001.DQDVXGL-P1-C10-T1
### ent2 ###
U789D.001.DQDVXGA-P1-C10-T2
### ent3 ###
U789D.001.DQDVXGA-P1-C10-T1
### ent4 ###
U789D.001.DQDVXGB-P1-C10-T2
### ent5 ###
U789D.001.DQDVXGB-P1-C10-T1
### ent6 ###
U789D.001.DQDVXHB-P1-C10-T2
### ent7 ###
U789D.001.DQDVXHB-P1-C10-T1
```

– As shown in Example 4-12 on page 56, the ent*x* are placed on:
  - Two different physical ports (T1 and T2)
  - On HEA card slot C10 on the system drawer back
  - Planar P1 (only one on p570)
  - Different system drawers (DQDVX*xx*)
  - A single System p (U789D.001 in our case)
- State and connectivity can be view as shown in bold in Example 4-13.

Example 4-13 The entstat command for connectivity and state

```
# entstat -d ent1 | grep Port
Logical Host Ethernet Port (l-port) Driver Properties:
Logical Port Link State: Up
Physical Port Link State: Up
Logical Host Ethernet Port (l-port) Specific Properties:
Logical Port Number: 2
Port Operational State: Up
External-Network Switch-Port Operational State: Down
```

Note that the logical port number is related to the port group:

– Ports 1 to 16 are related to port group 1.

– Ports 17 to 32 are related to port group 2.

– External-Network-Switch-Port Operational State is the state of the physical link.

– Port Operational State is the state of the logical port on the L2 switch (used for partition-to-partition network traffic).

- Configure the IP on the interface (such as en0) as you normally would for a network interface.

4.2.3 Performance considerations and recommendations

This section contains considerations and recommendations with respect to performance.

- MCS is enabled by default on AIX and uses the QPs value set on the hardware by the HMC.

  – The use of MCS is workload related and affects all physical ports belonging to the port group.

  – Use the smallest MCS value that meets the performance goal and keep in mind that the value affects all physical ports inside a port group, therefore, all LPARs that use an LHEA inside that port group.

  – For a logical port, a higher MCS value yields a higher raw throughput until it reaches media speed, but this also results in more CPU cycles per Gbps throughput due to less effective interrupt coalescing. The physical port is able to run at media speed when it has three receive queues (Rx) actively receiving data, regardless of whether the QPs belong to one, two, or three logical ports.

  – For a 10 Gbps physical port, an MCS value of 4 is the minimum value we recommend to reach media speed for the logical ports. In contrast, if there are four LPARs sharing the physical port and all of them are active, an MCS value of 1 allows those LPARs to reach the media speed through the ports.

– For a 1 Gbps physical port, an MCS value of 1 is the minimum value we recommend to reach the media speed for the logical ports.
– The small message transaction rate is higher with a higher MCS value.
– As a general recommendation, use a number of QP that match the virtual processors your LPAR can use, which does not exceed the physical processors number inside your server.

- The maximum transmission unit (MTU) rate must be set on the HW by the HMC and then inside the OS, according to your network infrastructure.
- We recommend flow control usage and allowing the interface to reach the full media speed (especially important for 10 Gbps adapter), and also to handle network traffic congestion.
- Set the promiscuous mode on a physical port with the HMC for a VIOS LPAR if you want to share this physical port with more than 16 LPARs, using an SEA.
  – This mode is not used for performance.
  – You can see the setting inside the OS using the command:

```
entstat -d entx | grep Promiscuous
```

  – LPARs using this SEA for external network connectivity will be able to sniff the network.
- You can set promiscuous mode temporarily on a physical port if you want to use sniff tools such as tcpdump or iptrace from an LPAR, but this will dedicate the whole physical port to the given LPAR.

4.3 Dynamic LPAR operations

This section provides information about dynamic LPAR operations related to the IVE. It describes the steps required to add or remove an LHEA logical port or completely remove an LHEA on your LPAR, using both GUI or CLI.

The examples provided show you a configuration for unassigned logical ports, but you can also select from assigned logical ports when they are not in use, deconfigured, and removed from one LPAR or when the LPAR is not activated.

Note: In order to perform all dynamic LPAR operations, your LPAR must have an interface configured and can reach the HMC. In addition, you must verify that you have at least one HMC resource as the management server, as shown.

For AIX systems prior to Version 6.1::

```
# lsrsrc IBM.ManagementServer
Resource Persistent Attributes for IBM.ManagementServer
resource 1:
Name = "<your HMC IP>"
Hostname = "<your HMC hostname or IP>"
ManagerType = "HMC"
LocalHostname = "<your LPAR hostname or IP>"
ClusterTM = "9078-160"
ClusterSNum = ""
ActivePeerDomain = ""
NodeNameList = {"<your LPAR>"}
```

For AIX systems 6.1 or later:

```
# lsrsrc IBM.MCP
Resource Persistent Attributes for IBM.MCP
resource 1:
        MNName            = "172.16.21.123"
        NodeID            = 15025524784985301456
        KeyToken          = "hmc8"
        IPAddresses       = {"172.16.20.114"}
        ConnectivityNames = {"172.16.21.123"}
        HMCName           = "7042CR6*107627C"
        HMCIPAddr         = "172.16.20.114"
        HMCAddIPs         = "192.168.128.1"
        HMCAddIPv6s       = ""
        ActivePeerDomain  = ""
        NodeNameList      = {"p740_lpar03"}
```

4.3.1 Add a logical port on LHEA

To perform a dynamic LPAR operation to add a new LHEA take the following steps:

Note: You cannot have more than one logical port on each physical port.

1. In your managed server panel, select the **LPAR** check box and use the **Task** button, as shown in Figure 4-7.

   In the HMC command line use the following command to choose in the list the managed system:

   ```
   lssyscfg -r sys -F name
   ```

   It is useful to define a variable to name it as in the following (use double quotation marks if there is space in the name):

   ```
   name=”your managed system name”
   ```

   List LPARs to choose the one you want (using the previously defined variable):

   ```
   lssyscfg -r lpar -m $name -F name
   ```

It is useful to define a variable to name your LPAR:

```
lpar="your lpar name"
```

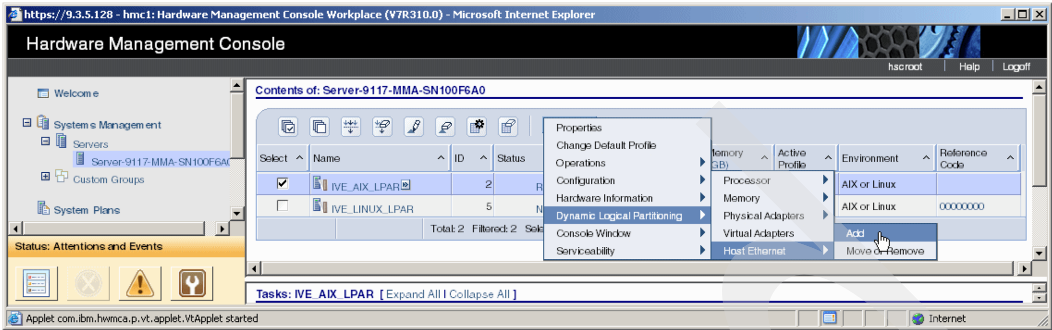

Figure 4-7 First step to add a logical port

2. In the new window Add Logical HEA Resources, as shown in Figure 4-8:
   - Choose in the pull-down list an HEA to select Logical Ports from.

     In the HMC command line use:

     ```
     lshwres -r hea --rsubtype phys --level sys -m $name
     ```

     Choose in the display the HEA to add the Logical Port from.

     It is useful to define a variable for the HEA you choose:

     ```
     hea=adapter_id
     ```

   - Select an HEA Physical port, then click **Configure** to choose the logical port to add to your LPAR.

     In the HMC command line use:

     ```
     lshwres -r hea --rsubtype phys --level port_group -m $name --filter adapter_
     ids=$hea
     ```

     Choose in the display the port group to use (we use port group 2 in our example).

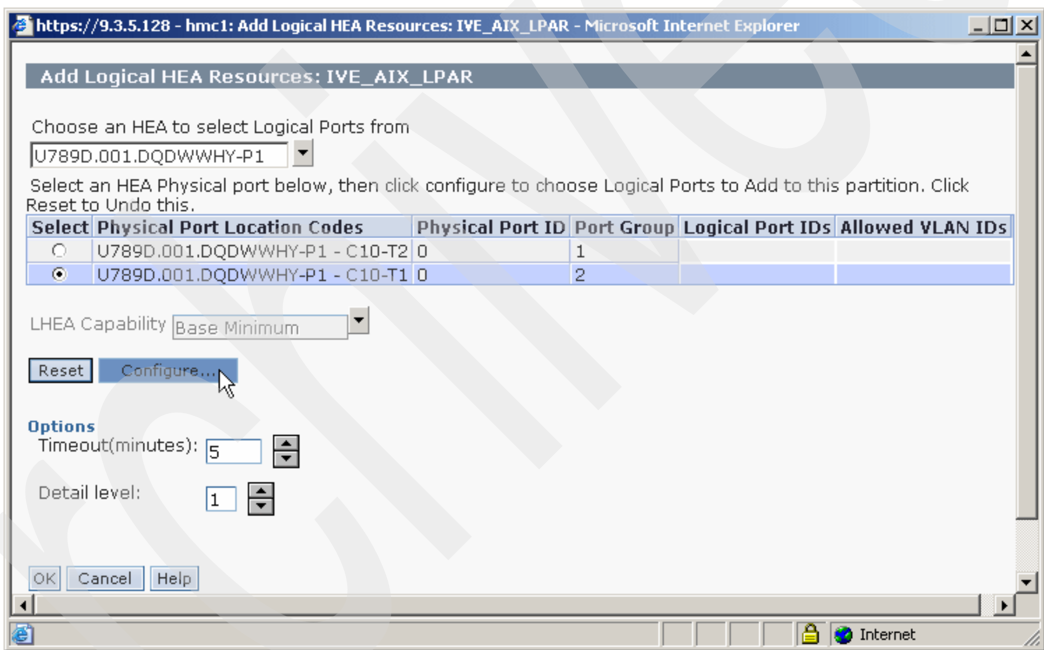

Figure 4-8 Second step to add a logical port

3. In the new window Logical Host Ethernet Adapter (LHEA) Configuration:
   – Specify Logical Port.
   – Specify allowed VLAN IDs if needed.

   As shown in Figure 4-9, in the HMC command line, use:

```
lshwres -r hea --rsubtype logical --level port -m $name --filter
"adapter_ids=$hea,port_groups=2" -F phys_port_id:logical_port_id:state:lpar_id
--header | grep -v none: | grep -E "none|id"
```

   This command shows you all logical ports that are unassigned (lpar_id=none) and available (state=1, depending on MCS value) for each physical port in a port group (port group 2 in our example).

   Choose a physical port and one logical port from the previous command output.

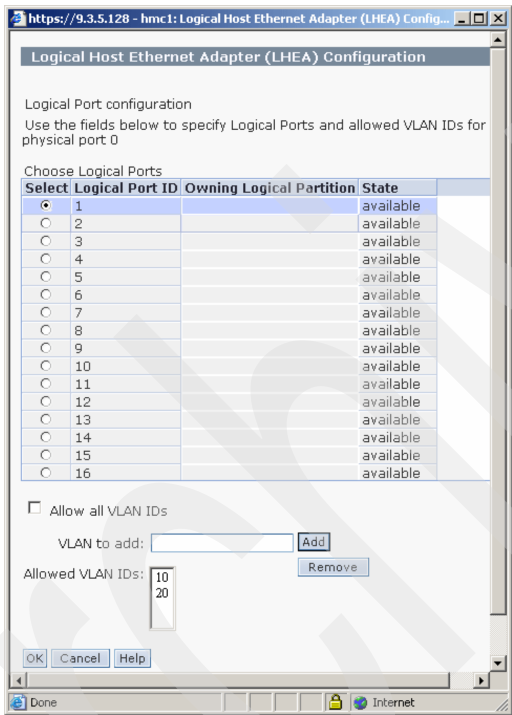

Figure 4-9 Third step to add a logical port

4. On the new window Add Logical HEA Resources, verify the logical port configuration. Click the **OK** button to complete your dynamic LPAR add operation, as shown in Figure 4-10.

   In the HMC command line use:

```
chhwres -r hea -m $name -o a -p $lpar -l $hea -g 2 --physport 0 --logport 1 -a
"vlan_id_list=10,20"
```

   This command completes the dynamic LPAR add step, adding in our example logical port 1 belonging to physical port 0 inside the port group 2 from the HEA that you have choosen in the given managed system for your LPAR.

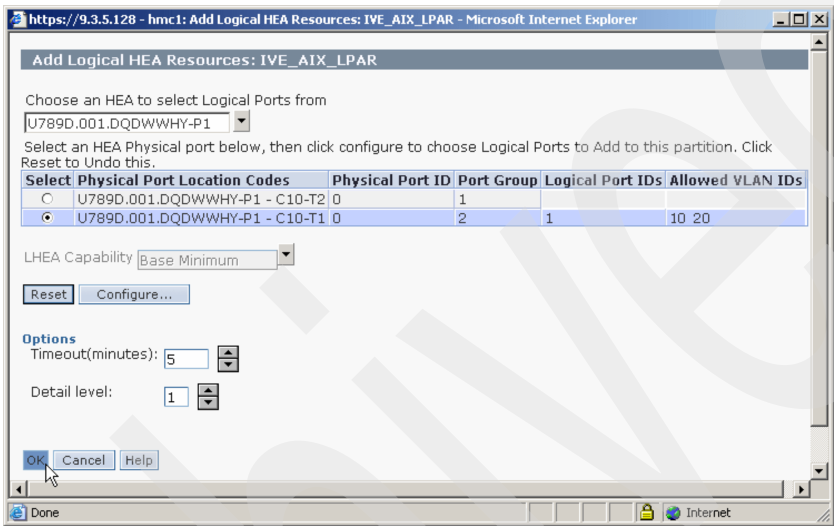

Figure 4-10 Fourth step to add a logical port

5. Log into AIX and run the **cfgmgr** command to discover your new interface. The following HMC command line lists and verifies the logical port configuration:

```
lshwres -r hea --rsubtype logical --level port -m $name --filter
"adapter_ids=$hea,port_groups=2,lpar_names =$lpar" -F lpar_name port_group
phys_port_id logical_port_id state mac_addr vlan_id_list --header
```

6. Configure your new interface.

4.3.2 Remove a logical port on LHEA

To remove a logical port, perform the following steps:

> **Note:** On Linux-based systems, steps 1 and 2 are not required.

1. Log onto AIX and note the physical location (T1 is port group 2, T2 is port group 1) and the logical port number (logical port 17 viewed by AIX is the logical port 1 of port group 2 viewed by the HMC) of the interface you want to remove, as shown in Example 4-14 on page 64.

2. Unconfigure then remove the interface from your system (you do not have to remove the LHEA, just the ent*x*, en*x,* and et*x* interfaces), as shown in Example 4-14.

Example 4-14 entstat, ifconfig, and rmdev commands to document and remove interfaces ent1 en1 et1

```
# lsdev -l ent1 -F physloc
U789D.001.DQDWWHY-P1-C10-T1
# entstat -d ent1 | grep "Logical Port Number"
Logical Port Number: 17
# ifconfig en1 down detach
# for i in et1 en1 ent1
do
rmdev -l $i -d
done
et1 deleted
en1 deleted
ent1 deleted#
```

3. In the server panel, check your LPAR check box and use the **Task** button, as shown in Figure 4-11.

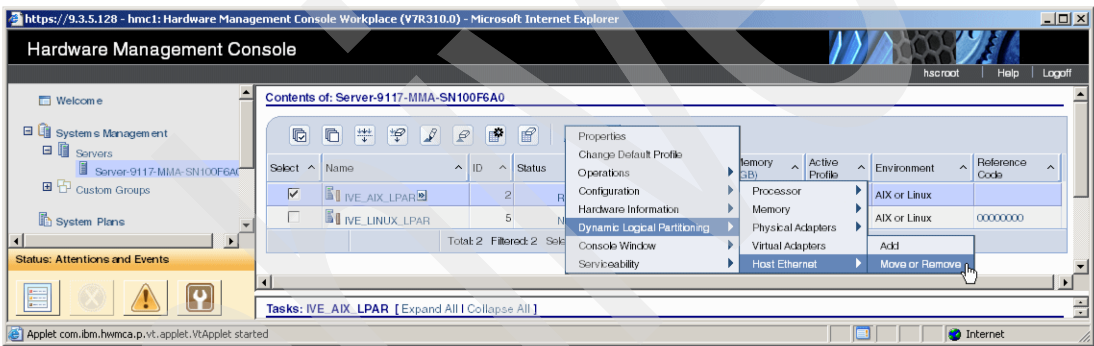

Figure 4-11 Second step to remove a logical port

4. In the next window Remove/Move Logical HEA Resources:

   a. Choose the appropriate HEA from the pull-down list.

> **Tip:** During our testing using a dynamic LPAR operation, the logical port that is assigned to an LPARs LHEA is not varied on. When this happens, the LHEA will have no logical ports. The logical port assignment can be deleted by using the command:
>
> `rsthwres -r hea -m $name -p $lpar -l $hea -g 2 --logport 1`
>
> - This command also applies to *orphaned* LHEAs, or LHEAs with no logical ports.
> - When specifying a logical port, -l, -g, and --logport are required together.
> - If the user specifies no LPAR, this command applies to all LPARs in the managed system.
>
> You can configure an LPAR with logical port, in which case the operation applies to all logical ports assigned to a LPAR.

i. In the HMC command line use **lssyscfg -r sys -F name** to choose in the list the managed system. It is useful to define a variable to name it:

```
name="your managed system name"
```

Use double quotation marks if there is space in the name.

ii. List LPARs to choose the one you want (using the previously defined variable):

```
lssyscfg -r lpar -m $name -F name
```

It is useful to define a variable to name your LPAR, as below:

```
lpar="your lpar name"
lshwres -r hea --rsubtype phys --level sys -m $name
```

iii. Choose in the display the HEA to remove Logical Port from, using the physical location noted for your interface in step 1.

It is useful to define a variable for the HEA you choose:

```
hea=adapter_id
```

According to the information note in step 1, you can now remove the logical port. Logical port number 17 indicates that is logical port 1 of port group 2 (T1 indicates also that the port group is 2).

iv. Run the dynamic LPAR removal command:

```
chhwres -r hea -m $name -o r -p $lpar -l $hea -g 2 --logport 1
```

This step completes the logical port removal with the CLI.

v. You can also move the logical port, adding the target LPAR in the command using the format:

```
-t target-partition-name or --tid target-partition-ID
```

The user can also just use -r hea to indicate all HEAs.

Another useful command is:

```
lshwres -r hea  --rsubtype logical --level port -m $name  -R
```

The R option lists information for partitions with hardware resources that can be restored due to a dynamic logical partitioning operation failure or activation failures because of configuration errors.

b. Choose the appropriate physical port.

c. If moving a port to another LPAR, select the destination LPAR in the Move to pull-down list.

d. Optionally change the time-out value or detail level.

e. Click the **OK** button to complete the logical port removal, as shown in Figure 4-12.

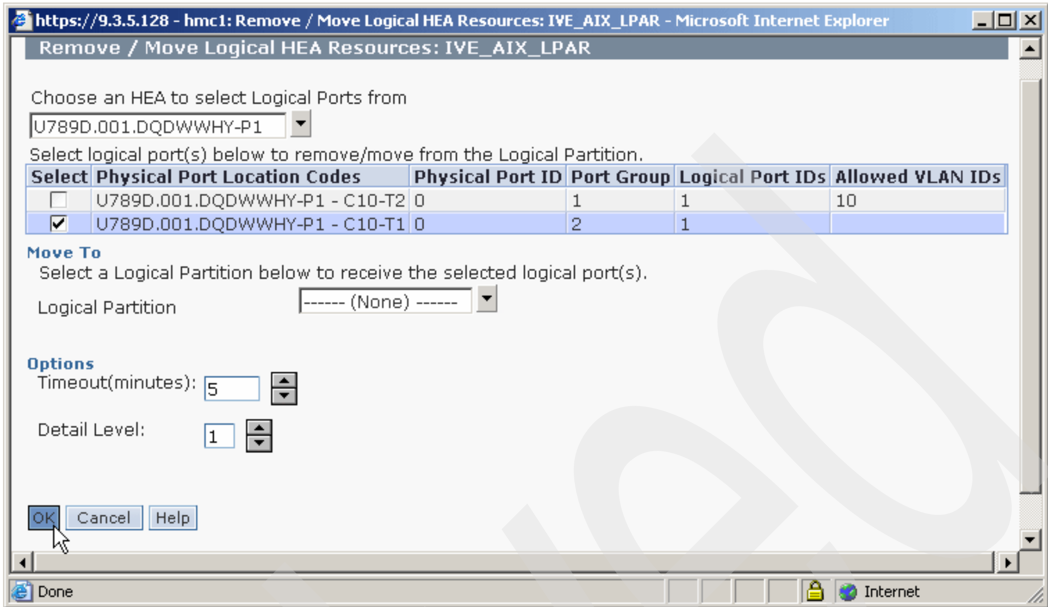

Figure 4-12 Third step to remove a logical port

4.3.3 Remove an LHEA

You may need to remove an LHEA, the logical representation of the physical adapter, if you need to alter the network configuration.

> **Note:** On Linux-based systems, step 1 is not required.

1. Log into AIX and perform the following command to find which LHEA you want to remove with all associated interfaces and logical ports:

   ```
   lsslot -c slot | grep lhea
   ```

   a. Note the physical location and the logical port number of each interface belong to the LHEA you want to remove, as shown in Example 4-14 on page 64.

   b. Unconfigure interfaces that belong to the LHEA:

   ```
   ifconfig enx down detach
   ```

   c. Remove the LHEA chosen and all associated interfaces:

   ```
   rmdev -l lheax -R
   ```

2. With GUI or CLI, remove all logical ports that belong to the LHEA, referring to 4.3.2, “Remove a logical port on LHEA” on page 63, in order to complete the LHEA removal.

4.4 How to change major settings and components summary

Table 4-1 provides a quick summary of the settings that are dynamic and the steps to follow to modify them, or to take into account the change (action order).

Table 4-1 Major settings and actions summary

Operation change	Where	HMC	AIX
Change MCS	Action order	1	N/A
	Dynamic	N	N/A
	Comments	Power cycle needed	Enable (default) or Disable by changing multicore attribute from yes to no on the ent*x*
Change promiscuous mode	Action order	1	N/A
	Dynamic	Y	N/A
	Comments	Typically used for an LHEA linked to an SEA in the VIOS	Command to see the setting: **`entstat -d entx	grep Promisc`**
Change flow control	Action order	1	1 in MDC mode
	Dynamic	Y	Y As result of HMC action (from AIX in MDC mode)
	Comments		Command to see the setting: **`entstat -d entx	grep Pause`**
Change Speed	Action order	1	N/A
	Dynamic	Y	Y As result of HMC action
	Comments		Command to see the setting: **`entstat -d entx	grep Speed`**
Change MTU	Action order	1	2
	Dynamic	Y	Y, but action required
	Comments		Command to change the interface attribute to take effect by: **`chdev -l entx jumbo_frames=yes`** Command to see the setting in use: **`entstat -d entx	grep Jumbo`**
Add LHEA or logical ports	Action order	1	2
	Dynamic	Y	Y but action required
	Comments	Run a DLPAR operation	Command to run: **`cfgmgr`**

Operation change / Where		HMC	AIX
Remove LHEA	Action order	2	1
	Dynamic	Y, but action required	Y, but action required
	Comments	► If you have another interface configured run a DLPAR operation, and modify your profile for future activation. ► If no other interface is available, you cannot DLPAR to remove an LHEA or a logical port definition from the hypervisor. In this case you must shut down your LPAR, modify the profile, and reactivate it.	Command to find the LHEA and associated interfaces to remove: **`lsslot -c slot | grep lhea`** Command to run on interfaces that belong to the LHEA: **`ifconfig enx down detach`** Command to remove the LHEA chosen and all associated interfaces: **`rmdev -l lheax -R`** Command to remove only one logical port: **`entstat -d entx | grep "Logical Port Number"`** To retreive the logical port number: **`rmdev -l entx -d`** **`rmdev -l enx -d`** **`rmdev -l etx -d`** to remove the device and the interfaces

<table>
<tr><th colspan="2">Where
Operation change</th><th>HMC</th><th>AIX</th></tr>
<tr><td rowspan="4">Add/remove VLANs on LHEA port
Two methods offered:
▶ Dynamic LPAR</td><td colspan="3">This scenario will preserve the MAC address.
Dynamic LPAR used</td></tr>
<tr><td>Action order</td><td>1/5/6/7</td><td>2/3/4/8/9/10</td></tr>
<tr><td>Dynamic</td><td>N</td><td>N</td></tr>
<tr><td>Comments</td><td>1- Note the actual LHEA port used and associated VLANs to modify.
5 - Remove the LHEA port from the LPAR using DLPAR.
6 - Add the new LHEA port (using the LHEA port noted in step 1) and add all VLANs needed using DLPAR.
7- Save the profile to keep the modifications using: configuration > save current configuration</td><td>2 - Commands to find your VLANs interfaces:
<code>for i in `lsdev -l en* | grep\</code><br><code>-i vlan | awk '{print $1}'`</code>
<code>do</code>
<code>print "### $i ###"</code>
<code>lsattr -El $i</code>
<code>done</code>
3 - Command to run on each VLANs, devices, and interfaces associated with the LHEA port to modify
<code>ifconfig enx down detach</code>
<code>rmdev -l enx -d</code>
<code>rmdev -l etx -d</code>
<code>rmdev -l entx -d</code>
4 - Unconfigure and remove the LHEA port device and interface (on which one the VLANs, devices, and interfaces were built).
<code>ifconfig enx down detach</code>
<code>rmdev -l enx -d</code>
<code>rmdev -l etx -d</code>
<code>rmdev -l entx -d</code>
8 - Run <code>cfgmgr</code> command to discover the LHEA port.
9 - Use the command <code>smit addvlan</code> (fast path) to add your VLANs and devices on top of the LHEA port device.
10 - Set yours IPs on the needed interface.</td></tr>
</table>

<table>
<tr><th colspan="2">Where
Operation change</th><th>HMC</th><th>AIX</th></tr>
<tr><td rowspan="4">► Profile modification and reactivation</td><td colspan="3">This scenario will preserve the MAC address.
Profile modification and reactivation used</td></tr>
<tr><td>Action order</td><td>1/6/8</td><td>2/3/4/5/7/9/10</td></tr>
<tr><td>Dynamic</td><td>N</td><td>N</td></tr>
<tr><td>Comments</td><td>1- Note the actual LHEA port used and associated VLANs to modify.
6 - Remove the LHEA port from the LPAR profile and reactivate the LPAR.
8 - Add the new LHEA port (using the LHEA port noted in step 1) and all VLANs needed in the LPAR profile and reactivate it.</td><td>2 - Commands to find your VLANs interfaces:
for i in `lsdev -l en* | grep\</b><br><b>-i vlan | awk '{print $1}'`
do
print "### $i ###"
lsattr -El $i
done
3 - Command to run on each VLAN, device, and interface associated with the LHEA port to modify:
ifconfig enx down detach
rmdev -l enx -d
rmdev -l etx -d
rmdev -l entx -d
4 - Unconfigure and remove the LHEA port device and interface (on which one the VLANs, devices, and interfaces were built:
ifconfig enx down detach
rmdev -l enx -d
rmdev -l etx -d
rmdev -l entx -d
5 - Shut down your LPAR:
shutdown -F
7 - Shut down your LPAR:
shutdown -F
9 - Use the command smit addvlan (fast path) to add your VLANs and devices on top of the LHEA port device.
10 - Set yours IPs on the needed interfaces.</td></tr>
</table>

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 72. Note that some of the documents referenced here may be available in softcopy only.

- *Introduction to Workload Partition Management in IBM AIX Version 6*, SG24-7431
- *IBM System p Live Partition Mobility*, SG24-7460
- *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration*, SG24-7940
- *Hardware Management Console V7 Handbook*, SG24-7491
- *IBM Director on System p5*, REDP-4219
- *Implementing IBM Director 5.20*, SG24-6188

Other publications

These publications are also relevant as further information sources:

- *System p Logical Partitioning Guide*, SA76-0098
- *System p Advanced POWER Virtualization Operations Guide,* SA76-0100
- *System p Overview*, SA76-0087
- *Operations Guide for the Hardware Management Console and Managed Systems*, SA76-0085

Online resources

These Web sites are also relevant as further information sources:

- IBM quick links for AIX fixes

  http://www14.software.ibm.com/webapp/set2/sas/f/genunix3/aixfixes.html
- IBM Linux on POWER service aids for hardware diagnostics

  http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html
- IBM System Planning Tool

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

  http://www-03.ibm.com/systems/management/director/

How to get Redbooks

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

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Integrated Virtual Ethernet Adapter

Technical Overview and Introduction

Redpaper™

Unique and flexible up to 10 Gbps Ethernet industry standard virtualizable network solution

Hardware accelerated for improved throughput

Configurable connections between LPARs and the physical network without the Virtual I/O Server

The introduction of the POWER6 processor-based servers brought to market several advances in server design, performance, and function. One of these enhancements is the capability to virtualize Ethernet within a server. The feature that provides this function is named Integrated Virtual Ethernet adapter (IVE).

IVE is the integration of several technologies, including the Host Ethernet Adapter (HEA), advanced software, and updates to the hypervisor, that provide integrated high-speed Ethernet adapter ports with hardware-assisted virtualization capabilities.

IVE is a standard set of features offered on every IBM System p POWER6 processor-based server. The IVE provides:

- Two 10 Gbps Ethernet ports or four 1 Gbps ports or two 1 Gbps ports
- External network connectivity for LPARs using dedicated ports without the need of a Virtual I/O Server
- Industry standard hardware acceleration, loaded with flexible configuration possibilities
- The speed and performance of the GX+ bus, faster than PCI-e x16

This IBM Redpaper is for technical professionals interested in understanding the many features of the IVE, and those who configure and support the servers that use it.

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

REDP-4340-00