

pVM53: Configuring Highly Available Virtual Networking Lab

(with Hints)

What This Exercise Is About

In this exercise, students set up a virtual networking environment which supports high availability. The exercise will use two VIO servers configured in either a POWER6 or POWER5 system to provide added resilience and serviceability. Two high availability configurations for virtual networking are described: Network Interface Backup (NIB) and Shared Ethernet Adapter (SEA) Failover.

What You Should Be Able to Do

At the end of the lab, you should be able to:

- Configure a dual virtual I/O server environment with Network Interface backup (NIB) on the Client partition (Scenario 1)
- Configure a dual virtual I/O server environment with SEA failover (Scenario 2)

General Description

This exercise describes two different high availability solutions related to the Virtual Ethernet environment on POWER6 and POWER5 systems. Students will choose to configure and test the failover mechanism for NIB or for SEA Failover environments.

Required Materials

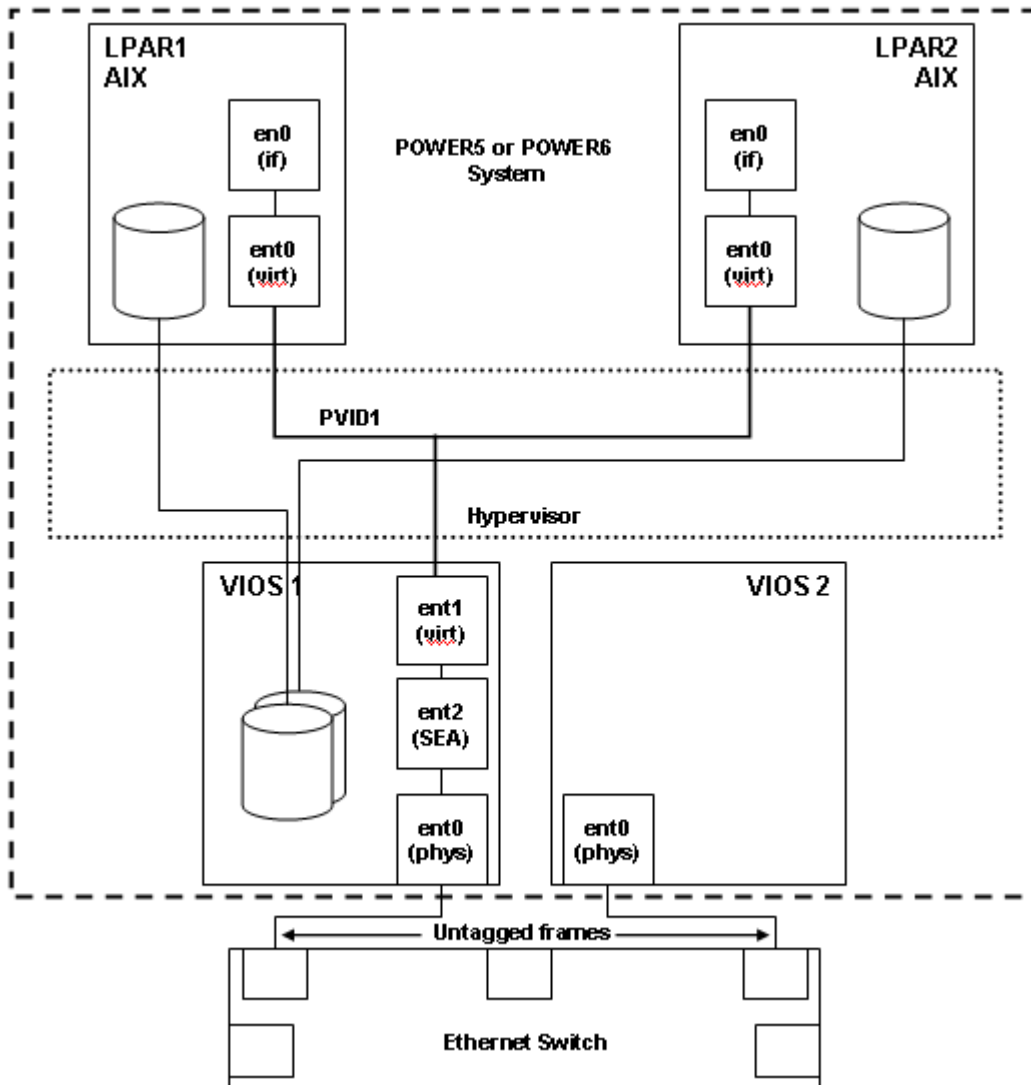
- This workbook
- A workstation with a browser, network connection, and access to an HMC configured to support an POWER5 or POWER6 system
- POWER5 or POWER6 system

System Information

Introduction: This lab was written assuming that each lab group has their own POWER5 or POWER6 managed system with two AIX partitions named **LPAR1** and **LPAR2**, and two Virtual I/O Server partitions named **VIOS1** and **VIOS2**. All four partitions should be *Running*.

VIOS1 has an SEA defined on **VLAN ID 1** and provides network connectivity for **LPAR1** and **LPAR2**. **VIOS2** does not have a virtual or Shared Ethernet adapter defined.

Here is a diagram of the network device configuration on the system at the start of this lab:



Students will need to be given the following information before the lab begins:

- The managed system name (as known to the HMC)
- IP address and/or hostname used to reach the following systems:
 - HMC
 - VIOS1
 - VIOS2
 - LPAR1
 - LPAR2
- The names of the partitions as known to the HMC for:
 - VIOS1
 - VIOS2
 - LPAR1
 - LPAR2
- Default profile names for the following partitions:
 - VIOS1
 - VIOS2
 - LPAR1
 - LPAR2
- The password for the following users and systems:
 - hscroot on HMC
 - padmin on VIOS1
 - padmin on VIOS2
 - root on LPAR1
 - root on LPAR2

Students can use the space above and label the diagram to record the information provided by the instructor.

Exercise Instructions with Hints

Preface

- The exercises depends on the availability of specific equipment in your classroom.
- All hints are marked with a » sign.

The diagrams help you in understanding the exercise flow. Tables help you define the ethernet adapter and interface names which should be used in your configuration. Adapter names shown on diagrams are examples.

IMPORTANT: Read all steps and configuration descriptions in this lab exercise thoroughly and execute all steps carefully. It is possible to disrupt the lab network by performing steps incorrectly.

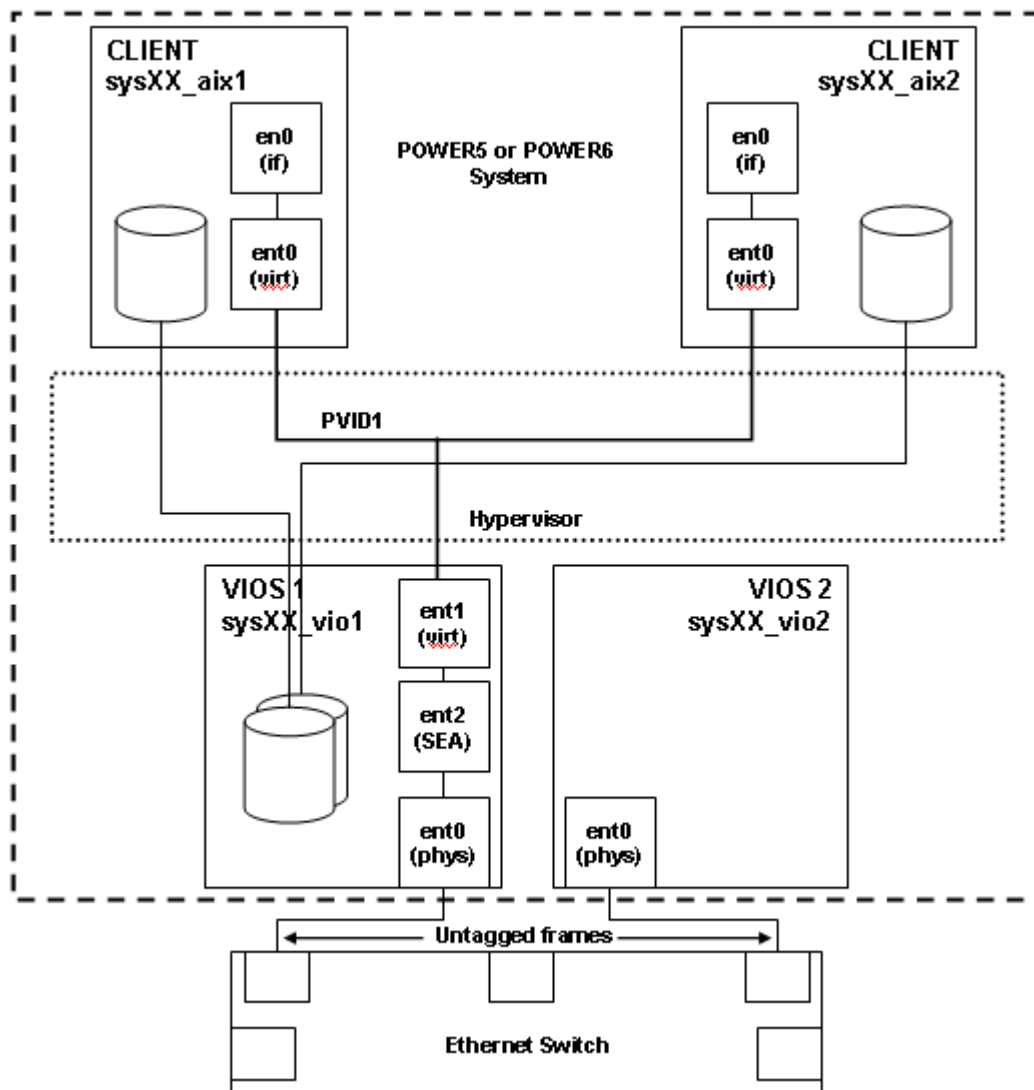
Lab Preparation: VLAN Configuration Verification

The diagram below represents the system configuration you should have at the beginning of this lab.

The names of Ethernet adapters that are shown in the diagrams are examples of what could be used. These could be different depending on your system configuration. (additional physical or virtual adapters would cause different adapter numbering as well).

Adapter names shown in the exercise steps are generic ones. For example ent<Phys> refers to the physical adapter in the VIOS partitions. You should fill in the tables to assist you with identifying the correct adapter numbers and names. One or both of your VIOS partitions may have more than one physical Ethernet adapter. You should identify ent<Phys> in your table as the physical adapter associated with the SEA on the VIOS partition, if an SEA exists.

In this exercise, you use the following partitions: VIOS1, VIOS2, LPAR1, and LPAR2. The partition names may actually be different and should be recorded at the beginning of this exercise. Ask your instructor if you do not have all the information you need to access and manage your system (the information listed under the **System Information** section).



As you identify your adapters in the following steps, complete this table with ethernet adapter names (ent0, ent1, . . .) and slots numbers (some values will be filled in much later in the exercise, or not at all depending on the configuration scenario you choose to try).

Partition	Physical Ethernet	Virtual Ethernet VLAN 1	Virtual Ethernet VLAN 2	Virtual Ethernet (SEA FO control) VLAN 12	SEA	NIB Link Agg Adapter
VIOS1	ent slot	ent slot	ent slot	ent slot	ent	N/A
VIOS2	ent slot	ent slot	ent slot	ent slot	ent	N/A
LPAR1	N/A	ent slot	ent slot	N/A	N/A	ent
LPAR2	N/A	ent slot	ent slot	N/A	N/A	ent

- ___ 1. Open a GUI session to your HMC using instructions provided by the instructor. As a first step, determine on which VLAN the virtual Ethernet adapter on VIOS1 is configured. You can look at the **Properties** dialog for VIOS1, and examine the properties of the virtual Ethernet adapter from the **Virtual Adapters** tab on the LPAR **Properties** dialog. Record the VLAN for the virtual Ethernet adapter here:
VIOS1 VLAN ID: _____
- » Select VIOS1 in the HMC GUI (how you do this will depend on the version of HMC you are using. Ask your instructor if you need help finding the LPAR in the HMC GUI)
 - » Show the **Properties** dialog for VIOS1
 - » Select the **Virtual Adapters** tab
 - » Select the virtual Ethernet adapter shown and view the properties of the adapter
 - » Record the **VLAN ID** in the space above. It should be **1**.
- ___ 2. On VIOS1, check the adapter name and adapter ID (also known as the slot number) for the virtual Ethernet adapter using following commands (record the applicable information in the table):
- » List virtual device names:
`$ lsdev -virtual`
 - » Output shows adapter device names and descriptions. Record the **ent** name of the virtual Ethernet adapter that you find in output in the table above. According to what you found in the last step, you would record this adapter information in the **Virtual Ethernet VLAN 1** column of the table for **VIOS1**. In this example output, the virtual Ethernet adapter device name is ent1:
`ent1 Available Virtual I/O Ethernet Adapter (1-lan)`
 - » Find the adapter ID of the virtual Ethernet adapter:
`$ lsdev -slots`
 - » Output shows the device names and location codes, which contain the adapter IDs. For example, in the following output, the virtual Ethernet adapter identified in `lsdev` output (**ent1**) has adapter ID **11**:
`U8203.E4A.10CD1F1-V1-C11 Virtual I/O Slot ent1`
- ___ 3. On VIOS1, determine the SEA device name and record it in the table. Also determine the physical adapter that is configured for the SEA and record it in the table under the “Physical” column for VIOS1.
- » Find **ent** device name for SEA adapter:
`$ lsdev -virtual | grep Shared`

- » Record the **ent** device name in the table under the “SEA” column for VIOS1. In the output below, this is **ent2**:
ent2 Available Shared Ethernet Adapter
- » Use **lsdev -attr** to determine the physical adapter configured for the SEA. If the SEA adapter has the name **ent2**, the command would look like this:
\$ lsdev -dev ent2 -attr
- » This is an example of the output from the last command on VIOS1 (note that you can also see the PVID for the SEA device in this output):

```

attribute      value      description user_settable
accounting     disabled  Enable per-client accounting of network statistics
True
ctl_chan              Control Channel adapter for SEA failover True
gvrp      no      Enable GARP VLAN Registration Protocol (GVRP) True
ha_mode  disabled  High Availability Mode True
jumbo_frames  no  Enable Gigabit Ethernet Jumbo Frames True
large_receive  no  Enable receive TCP segment aggregation True
largesend      0  Enable Hardware Transmit TCP Resegmentation True
netaddr        0  Address to ping True
pvid          1  PVID to use for the SEA device True
pvid_adapter  ent1  Default virtual adapter to use for non-VLAN-tagged
packets True
qos_mode      disabled  N/A True
real_adapter  ent0  Physical adapter associated with the SEA True
thread        1      Thread mode enabled (1) or disabled (0) True
virt_adapters ent1    List of virtual adapters associated with the SEA
(comma separated) True

```

- ___ 4. Open a telnet or terminal session to **VIOS2**. You should not find any virtual Ethernet adapters or Shared Ethernet adapters defined on VIOS2 at this time, but check to confirm. If you do find a virtual or Shared Ethernet adapter., remove them with **rmdev**.
- » Check for the existence of any virtual or Shared Ethernet adapters
\$ lsdev -virtual
 - » No virtual or Shared Ethernet adapters should be shown. If they are, remove them with **rmdev** (ask the instructor if you need assistance with this)
- ___ 5. On each VIOS, read IP routing configuration using **lstcpip -routinfo** or **netstat -routinfo** command and check that VIOS1 and VIOS2 are configured to communicate with the external network. Use **ping** to reach an external address like the HMC or default gateway.

You can also look at the ethernet statistics using `entstat -a11` command, specifying the SEA adapter you recorded in the table from step 3 (for example, `entstat -a11 ent2`). Find the details for the Real Side Statistics and Virtual Side Statistics packets values. Incrementing values in these fields show that the virtual ethernet and SEA configuration is functioning.

- ___ 6. From the HMC, determine on which VLAN ID LPAR1 and LPAR2 are configured. You can follow the hints for step 1 if you need assistance.
- » Both LPAR1 and LPAR2 should be configured on VLAN ID 1
- ___ 7. Open a telnet or terminal window to **LPAR1**. Determine the adapter name and slot for the virtual Ethernet adapter and record them in the table in the row for **LPAR1**, under the **Virtual Ethernet VLAN 1** column.
- » Get the adapter device name:
`lsdev -c adapter -S a | grep ent`
 - » You should see output similar to the following. In this case the device name is **ent1** (it may be different on your system). Record the adapter name in the table:
`ent1 Available Virtual I/O Ethernet Adapter (1-lan)`
 - » Determine the adapter ID (or slot number) with `lscfg`. If the adapter is **ent1**, the command would be:
`lscfg -l ent1`
 - » The output shows the location code, which contains the adapter ID. In this case it is **4** (it may be different on your system). Record the slot number in the table
`ent1 U8203.E4A.10CD1F1-V2-C4-T1 Virtual I/O Ethernet Adapter (1-lan).`
- ___ 8. Open a telnet or terminal window to **LPAR2**. Determine the adapter name and slot for the virtual Ethernet adapter and record them in the table in the row for **LPAR2**, under the **Virtual Ethernet VLAN 1** column.
- » Get the adapter device name:
`lsdev -c adapter -S a | grep ent`
 - » You should see output similar to the following. In this case the device name is **ent1** (it may be different on your system). Record the adapter name in the table:
`ent1 Available Virtual I/O Ethernet Adapter (1-lan)`
 - » Determine the adapter ID (or slot number) with `lscfg`. If the adapter is **ent1**, the command would be:
`lscfg -l ent1`

- » The output shows the location code, which contains the adapter ID. In this case it is **4** (it may be different on your system). Record the slot number in the table
- ```
ent1 U8203.E4A.10CD1F1-V2-C4-T1 Virtual I/O Ethernet Adapter (1-lan).
```

- \_\_\_ 9. Verify that you can ping from LPAR1 to LPAR2 and back, and that you can ping from LPAR1 and LPAR2 to an external system (like your HMC, or VIOS1).

You have determined the current configuration of the virtual networking in the system and documented it in the table. Some of the cells in the table are still empty, and you will fill those in as you complete Scenario 1 or Scenario 2 below. You will only have time in this lab to complete one of the scenarios, so choose that one that interests you most.

### ***Description of Scenarios***

**Scenario 1** explores the Network Interface Backup configuration, in which each VIOS partition supports 2 different VLANS and 2 different sets of clients (in our case, VIOS1 supports LPAR1 on VLAN 1, and VIOS2 supports LPAR2 on VLAN 2). You will configure the clients to use a backup adapter in a link aggregation to provide the failover mechanism.

**Scenario 2** explores the Shared Ethernet Adapter Failover configuration, in which both VIOS partitions support the same set of clients on the same VLAN, and a control channel is used to determine failure on the primary server and trigger failover. In this scenario, configuration on the client is simplified, but bandwidth is lost because only one SEA adapter is ever used at a time.

Now you can choose your favorite scenario and get to work!

## Scenario 1: Configure Network Interface Backup

In the following steps, the Network interface Backup feature is configured using two different virtual LANs. This shows how an LPAR can use an AIX feature to provide improved network availability without sacrificing the bandwidth of available physical adapters.

Currently, VIOS1 provides access to the external switch for LPAR1 and LPAR2 through its Shared Ethernet adapter. All hosts are configured as VLAN unaware (without VLAN tagging) and they are defined using the same IP range as the machines on the external switch.

Remember that communication between LPAR1 and LPAR2 is done through the internal virtual Ethernet switch (the Hypervisor), and VIOS1 forwards any external traffic from LPAR1 and LPAR2 on VLAN 1.

### NIB Configuration startup

- \_\_\_ 1. Currently VIOS2 does not have a virtual or Shared Ethernet adapter. Add a virtual Ethernet adapter dynamically to **VIOS2** with the next available adapter ID, and a **VLAN ID 2**. Remember to check the **Access external network** checkbox since we will use this adapter in an SEA.

**NOTE:** Since we are configuring NIB, each VIOS partition will support a different VLAN. VIOS1 supports VLAN 1 and VIOS 2 supports VLAN 2.

Record the new virtual Ethernet adapter ID in the table, in the **VIOS2** row, under the column **Virtual Ethernet VLAN 2**.

- » From the HMC GUI, select the VIOS2 partition and add a virtual Ethernet adapter with the next available adapter ID, and a VLAN ID 2
  - » Check the **Access external network** checkbox
  - » Click OK twice to complete the operation
  - » Ask the instructor if you need assistance
- \_\_\_ 2. Add the virtual Ethernet adapter to the VIOS2 profile to make it persist across profile activations. Use the same adapter ID that you used to add the adapter dynamically (it should be recorded in the table).
- » Select **VIOS2** and edit the profile properties for the active profile
  - » On the **Virtual Adapters** tab add a virtual Ethernet adapter with the ID you recorded in the last step, and on **VLAN 2**
  - » Check the **Access external network** checkbox
  - » Click OK twice to complete the operation

- \_\_\_ 3. In the telnet or terminal session for **VIOS2**, run `cfgdev` and verify that the new virtual Ethernet adapter is available. Record the device name in the table at the beginning of the exercise.

» `cfgdev`

» `lsdev -virtual`

» You should see a virtual Ethernet adapter listed. Record its name in the table in the **VIOS2**, under the **Virtual Ethernet VLAN 2** column

- \_\_\_ 4. On **VIOS2**, identify the currently configured physical Ethernet adapter. You can use `lstcpip` to determine this, or another tool of your choice. Record the device name of the physical Ethernet adapter in use in the table at the beginning of the exercise.

» `lstcpip`

In this example output, the IP configuration is on the adapter **ent0**

| Name        | Mtu   | Network | Address           |
|-------------|-------|---------|-------------------|
| en0         | 1500  | link#4  | ba.ae.ce.cd.bd.c  |
| <b>ent0</b> | 1500  | 9.47.87 | <b>hurston105</b> |
| lo0         | 16896 | link#1  |                   |
| lo0         | 16896 | 127     | localhost         |
| lo0         | 16896 | :::1    |                   |

You can use the physical adapter you identified in the last step, and the virtual adapter you added in the previous step (step 12) to create the SEA on VIOS2. Before you can do this, you have to remove the IP configuration on the physical Ethernet interface.

- \_\_\_ 5. If you are currently logged into VIOS2 via telnet (or ssh), close this session and open a virtual terminal window from the HMC. From the HMC CLI, you can open a terminal window on VIOS2 with:

```
mkvtterm -m managed_sys_name -p VIOS2
```

» From the HMC GUI, select the **VIOS2** partition and choose to open a virtual terminal window from the selected menu

- \_\_\_ 6. From the virtual terminal session, remove the IP configuration from the physical adapter interface recorded in step 13 for VIOS2. Use `rmtcpip` to do this. For example, if the physical Ethernet adapter is **ent0**, the command to execute would be:

```
rmtcpip -interface en0
```

Answer 'y' to the prompt to continue

- \_\_\_ 7. Now you can create the SEA on **VIOS2**. Use the `mkvdev -sea` command. The default adapter will be the same as the virtual adapter, and the default ID specified should be **2** (because this is the VLAN on which the virtual Ethernet adapter is configured on VIOS2). For example, if the physical adapter is **ent0** and the virtual adapter is **ent1**, the command to execute would be:

```
mkvdev -sea ent0 -vadapter ent1 -default ent1 -defaultid 2
```

- \_\_\_ 8. Verify the new SEA exists on **VIOS2** with the `lsdev -virtual` command. Record the adapter name for the SEA in the table at the beginning of the lab, in the **VIOS2** row and under the **SEA** column.

```
» lsdev -virtual | grep Shared
```

```
ent2 Available Shared Ethernet Adapter
```

```
» In this case, the SEA device name is ent2
```

You completed configuration of the SEA on VIOS2, which is configured to forward traffic on VLAN ID 2. However, currently LPAR1 and LPAR2 both use VLAN 1 for communication. Next you will add virtual Ethernet adapters to LPAR1 and LPAR2 on VLAN 2 to support an NIB configuration.

- \_\_\_ 9. Dynamically create an additional virtual Ethernet adapter on each client partition (**LPAR1** and **LPAR2**), using the VLAN ID for the VIOS2 VLAN. (**VLAN ID 2**). Record the adapter ID in the table at the beginning of the lab in the LPAR1 and LPAR2 rows, under the column labeled **Virtual Ethernet VLAN 2**.

- \_\_\_ 10. Save your new **LPAR1** and **LPAR2** configurations in new profiles. Name them **NIB**.

```
» Select LPAR1 and select Configuration -> Save Current Configuration. Type NIB in the dialogue box
```

```
» Repeat the action for LPAR2
```

- \_\_\_ 11. Run the `cfgmgr` command on **LPAR1** in order to configure the new virtual Ethernet adapter. Record the adapter name in the correct columns for LPAR1 and VLAN 2.

```
» First run cfgmgr to discover the new adapter
```

```
» You can list the new adapter with:
```

```
lsdev -c adapter -S a
```

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

```
ent1 Available Virtual I/O Ethernet Adapter (l-lan)
```

```
ent2 Available Virtual I/O Ethernet Adapter (l-lan)
```

```
» In the example above, the new adapter name is ent2. You can use this command to verify the adapter ID for ent2:
```

```
lscfg -l ent2
```

**ent2** U8203.E4A.10CD1F1-V2-C4-T1 Virtual I/O Ethernet Adapter  
(1-lan)

- » In this case the adapter ID is **4**. The value you see here on your system should match what you recorded in step 18.
- » Record the new adapter name in the table.

\_\_\_ 12. On LPAR1, you have to stop the network interface that is already configured in order to be able to define the same IP on the NIB device that is going to be configured. You should do this from a virtual terminal session, not a telnet or ssh session because you will lose network connectivity.

Before doing this, use **ifconfig**, **netstat** or **smit** to identify the current IP configuration (if don't already have it documented). Record the IP, netmask, and so forth. Once you have the information, you ifconfig to down and detach the interface. For example, if the currently configured interface is **en0**, the command would be:

```
ifconfig en0 down detach
```

\_\_\_ 13. On LPAR1, configure a Network Interface Backup using **smitty etherchannel**

- Select **Add An EtherChannel/Link Aggregation**.
- Select the virtual Ethernet adapter on VLAN1 (recorded in the table for LPAR1 in the **Virtual Ethernet VLAN1** column) in the **Available Network Interfaces** to select it as the Etherchannel adapter.
- On the **Add An EtherChannel/Link Aggregation** dialogue screen, set your new virtual Ethernet adapter on **VLAN2** as the **Backup Adapter**.
- Leave the default values in all of the other fields.

This makes VIOS1 the primary path to the external network, and VIOS2 the backup path. Record the name of the new NIB adapter in the table, in the row for **LPAR1**, under the **NIB** column.

\_\_\_ 14. On LPAR1, using **smitty tcpip**, configure the NIB interface with the same IP and netmask as previously used to communicate with this LPAR (the information you recorded in step 21). This reestablishes an external network connection. After this step, verify that you can ping the HMC from LPAR1.

\_\_\_ 15. On LPAR2, you have to stop the network interface that is already configured in order to be able to define the same IP on the NIB device that is going to be defined in the next step. Before doing this, use **ifconfig**, **netstat** or **smit** to identify the current IP configuration (if don't already have it documented). Record the IP, netmask, and so forth. Run the following commands from LPAR2's virtual terminal. For example, if

the currently configured interface is **en0**, the command would be:

```
ifconfig en0 down detach
```

\_\_\_ 16. On LPAR2, configure a Network Interface Backup using **smitty etherchannel1**.

- Select **Add An EtherChannel/Link Aggregation**
- Select the virtual Ethernet adapter on VLAN2 (recorded in the table in the **LPAR2** row under the **Virtual Ethernet VLAN2** column) in the **Available Network Interfaces** to select it as the Etherchannel adapter.
- On the **Add An EtherChannel/Link Aggregation** dialogue screen, set your virtual Ethernet adapter on **VLAN1** (get the adapter name from the table) as the **Backup Adapter**.
- Leave the default values in all of the other fields.

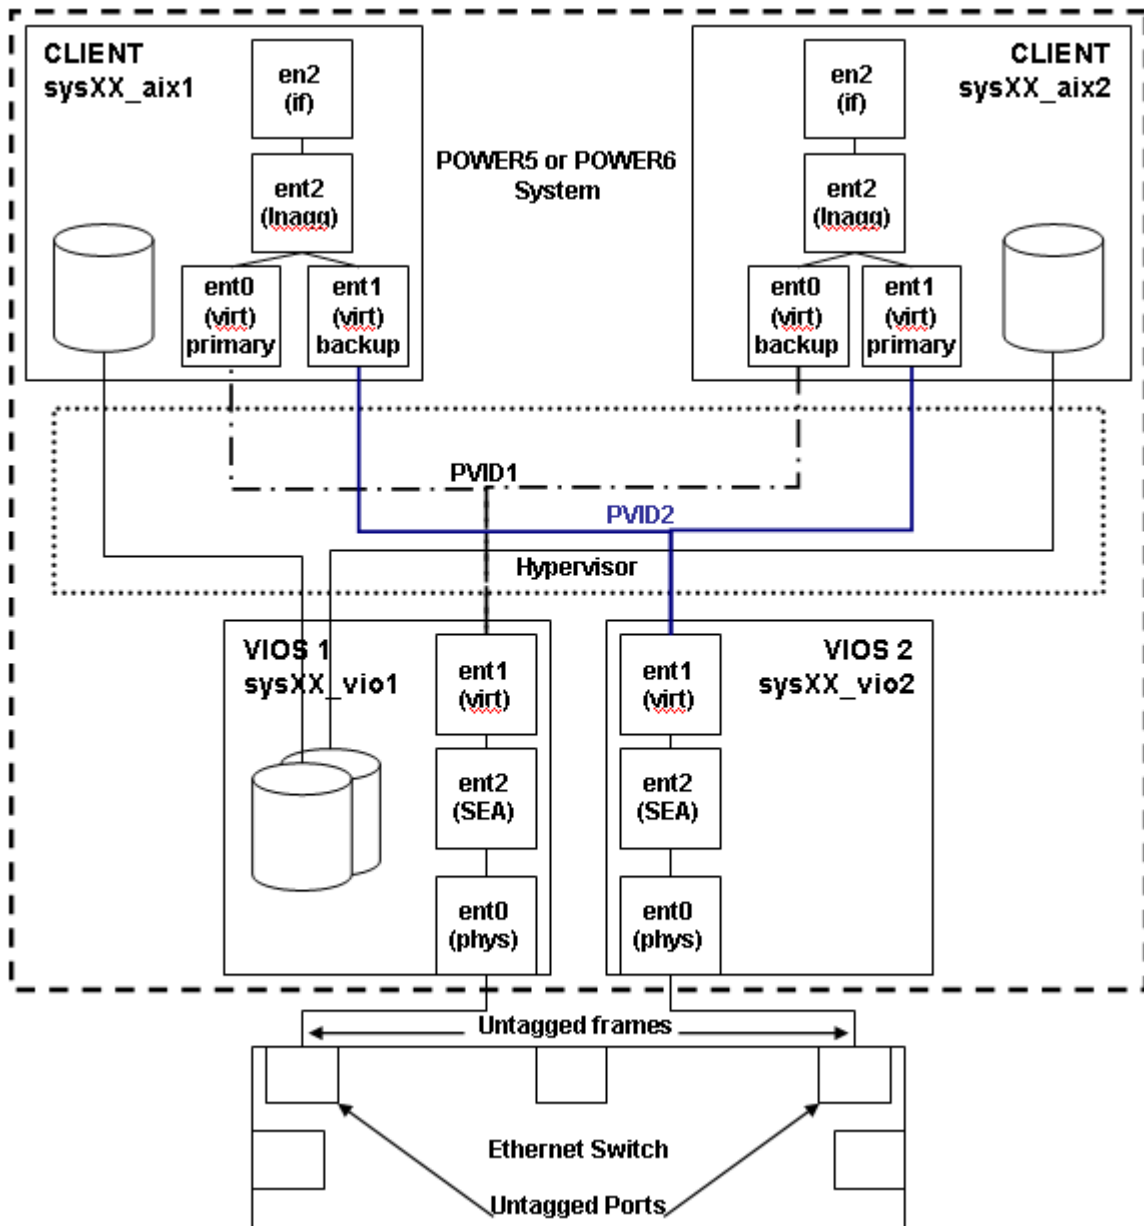
By configuring in this manner, LPAR1 and LPAR2 are using different Virtual I/O servers for their primary network access. This is a simple example of manually load-balancing the network traffic. Record the name of the new NIB adapter in the table, in the row for **LPAR2**, under the **NIB** column.

\_\_\_ 17. On LPAR2, using **smitty tcpip**, configure the NIB interface with the same IP and netmask as previously used to communicate with this LPAR (the information recorded in step 24). After this step, you should be able to ping the HMC from LPAR2.

With production configurations like this, it is desirable to configure half of the clients with a primary path of their NIB through VIOS1 and the other half through VIOS2.

In this configuration, LPAR1 – LPAR2 inter-partition communication is now done through the external network until one NIB switches to its backup adapter.

At this step of the exercise you should have a configuration as shown in following figure:



\_\_ 18. To test the etherchannel failover: open a window for **LPAR2**, start a ping to the HMC,

```
ping <hmc_ip_address>
```

\_\_ 19. While the ping is running, log into the VIOS2 partition and use the HMC GUI or CLI to select **VIOS2** and perform a shutdown immediate. (VIOS2 is not needed by LPAR1 and LPAR2 as a virtual SCSI disk server).

» From the HMC CLI:

```
chsysstate -o shutdown -r lpar -m managed_sys_name -n VIOS2 --immed
```

» From the HMC GUI, select VIOS2 and choose **Shutdown** from the selected menu (on HMC V7, this is under **Operations**). Choose the immediate option.

\_\_\_ 20. Check that the ping fails to show a response, and the etherchannel doesn't switch to the backup adapter. Can you explain why?

» The etherchannel backup doesn't detect a network failure as it is currently built with only virtual adapters. You must specify a valid *internet address to ping*.

\_\_\_ 21. On LPAR2, change the Etherchannel attribute **internet address to ping** with the IP address of your HMC.

Use `smitty etherchannel` and select the following:

- **Change/Show Characteristics of an Etherchannel**
- **Internet Address to Ping** `<HMC_IP_Address>`

Click Enter to complete the operation.

\_\_\_ 22. Try pinging the HMC again. The ping works. Look at the error log using `errpt`. You can see that an error has been detected on the etherchannel and the backup adapter is now active.

```
errpt | more
```

```
IDENTIFIER TIMESTAMP T C RESOURCE_NAME DESCRIPTION
```

```
9F7B0FA6 0630173706 I H ent2 PING TO REMOTE HOST FAILED
```

```
8650BE3F 0630173706 I H ent2 ETHERCHANNEL RECOVERY
```

\_\_\_ 23. You can also check using `entstat` command, that the ethernet frames are routed through the backup adapter. Specify the NIB adapter to `entstat`.

```
entstat -d ent<NIB> | grep Active
```

```
Active channel: backup adapter
```

\_\_\_ 24. Activate VIOS2 partition from the HMC. From the HMC CLI you can run the command:

```
chsysstate -o on -r lpar -m managed_sys_name -n VIOS2 -f profile_name
```

» You can activate VIOS2 from the HMC GUI by selecting it and choosing the **Activate task**

- \_\_\_ 25. Notice that LPAR2 does not automatically switch back to its primary Ethernet adapter on VLAN2 when VIOS2 is available again. You must force it to switch back to the primary adapter. This is not required when you use an SEA failover configuration.

Do the following when VIOS2 is up and running:

On **LPAR2** force the etherchannel to switch back to the primary adapter using **smitty etherchannel** and **Force A Failover In An EtherChannel**.

You can also use `/usr/lib/methods/ethchan_config` command.

```
/usr/lib/methods/ethchan_config -f ent<NIB>

entstat -d ent<NIB> | grep Active
Active channel: primary channel
```

The AIX Error log contains the etherchannel switch notifications.

- \_\_\_ 26. To complete the NIB configuration at LPAR1, change the etherchannel attribute **internet address to ping** to identify your HMC IP address using **smitty etherchannel**.

You have successfully created and tested a Network Interface Backup Failover configuration - Congratulations! Let your instructor know that you are finished, and if you have any questions.

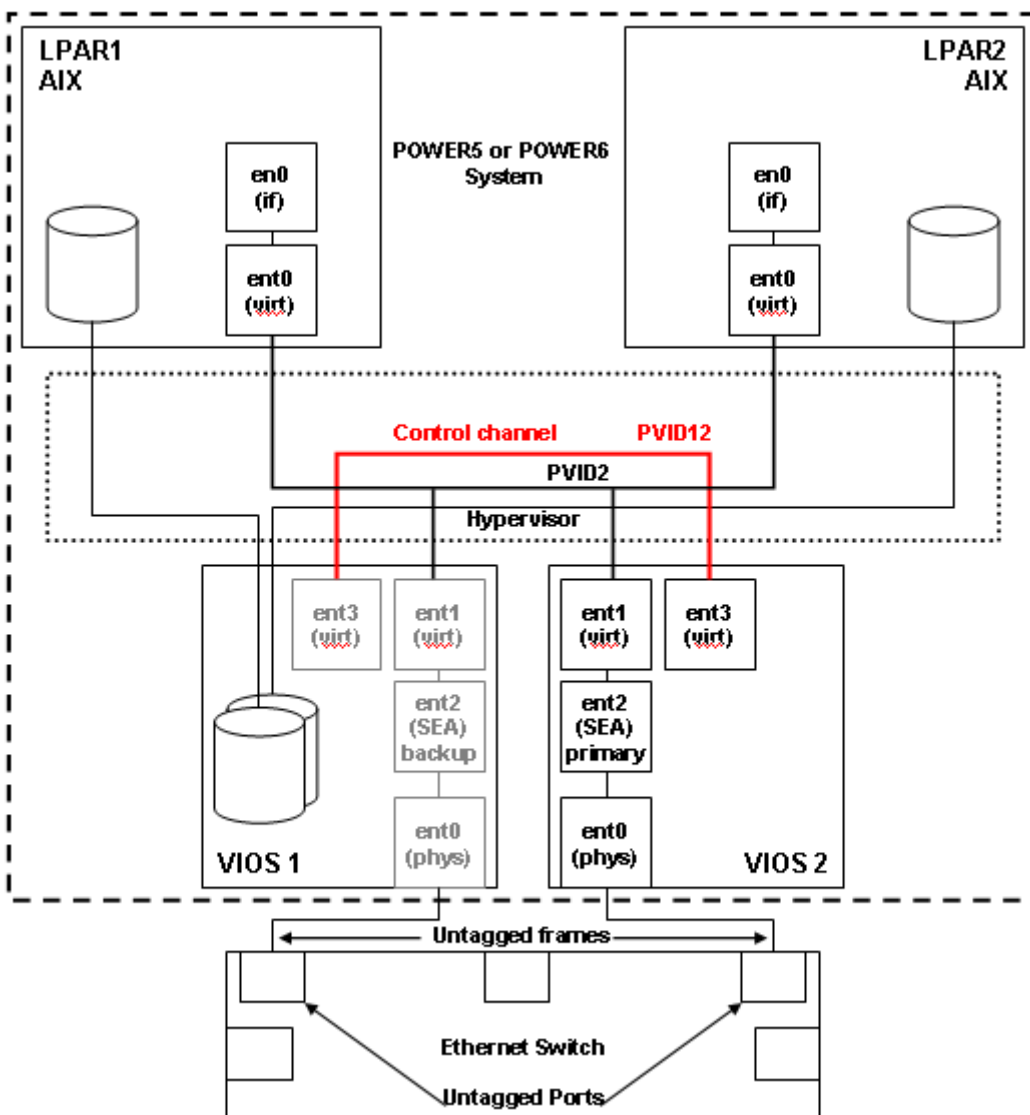
## Scenario 2: Dual VIOS SEA failover configuration

In this scenario, you will configure dual SEA failover on the system. Currently both client partitions are using VIOS1 to get access to an external network using VLAN ID 1.

First you will create the control channel adapters on each VIOS partition. These control channel adapters are used to determine the health of the SEAs and are required to avoid a broadcast storm (which can result when two trunking virtual adapters are available on the same VLAN).

Next you will recreate the SEA on VIOS1 to support a failover configuration for VLAN 2, then create the SEA adapter on VIOS2 to be the primary SEA on VLAN ID 2.

At the end of this part, your system is configured as shown in the following diagram:



## Creating control channel adapters on VIOS1 and VIOS2

In this section, you will create the control channel virtual Ethernet adapters on VIOS1 and VIOS2, which will communicate on VLAN ID 12. It is very important to create this adapter on both VIOS partitions before creating SEA adapters to support failover for the same VLAN. Failing to have proper control channel configuration can result in causing a broadcast storm when both SEA adapters are activated on the same VLAN (VLAN ID 2 in this case).

- \_\_\_ 1. Using the HMC, dynamically add a virtual ethernet adapter to **VIOS1** . Set the VLAN ID to **12** for this adapter and configure as shown in the visual below. Notice that although this control channel adapter will be part of an SEA, you **do NOT check the Access external checkbox**, because this adapter will not be used to forward VLAN traffic to an external network. Record the control channel adapter ID (slot number) in the table at the beginning of the exercise under the **Virtual Ethernet (SEA FO control) VLAN 12** column for **VIOS1**..:

Virtual ethernet adapter  
 Adapter ID : \* 14  
 VLAN ID : \* 12 System VLAN Info...

**Optional Settings**

IEEE 802.1q compatible adapter  
 Maximum number of VLANs :  
 New VLAN ID : 1 Add  
 Additional VLANs : Remove

Access external network  
 Trunk priority : 0 ▲ ▼

OK Cancel Help

- \_\_\_ 2. From the HMC, add the control channel adapter to VIOS2. Follow the instructions in the previous step to add a virtual Ethernet adapter to **VIOS2** on **VLAN ID 12**. Record the control channel adapter ID (slot number) in the table at the beginning of the exercise under the **Virtual Ethernet (SEA FO control) VLAN 12** column for **VIOS2**.
- \_\_\_ 3. On VIOS1 , configure the new adapter using `cfgdev` command. Check for the new virtual Ethernet adapter with `lsdev`, and record the adapter device name in the table at the beginning of the exercise under the **Virtual Ethernet (SEA FO control) VLAN 12** column for **VIOS1**.
  - » `cfgdev`
  - » `lsdev -virtual | grep ent`

» Record the device name of the new virtual Ethernet adapter (you can verify it is the right adapter by matching the adapter ID shown in the location code (with `lsdev -vpd`) with the value you recorded in the table).

4. On VIOS2 , configure the new adapter using `cfgdev` command. Check for the new virtual Ethernet adapter with `lsdev`, and record the adapter device name in the table at the beginning of the exercise under the **Virtual Ethernet (SEA FO control) VLAN 12** column for **VIOS2**.

» `cfgdev`

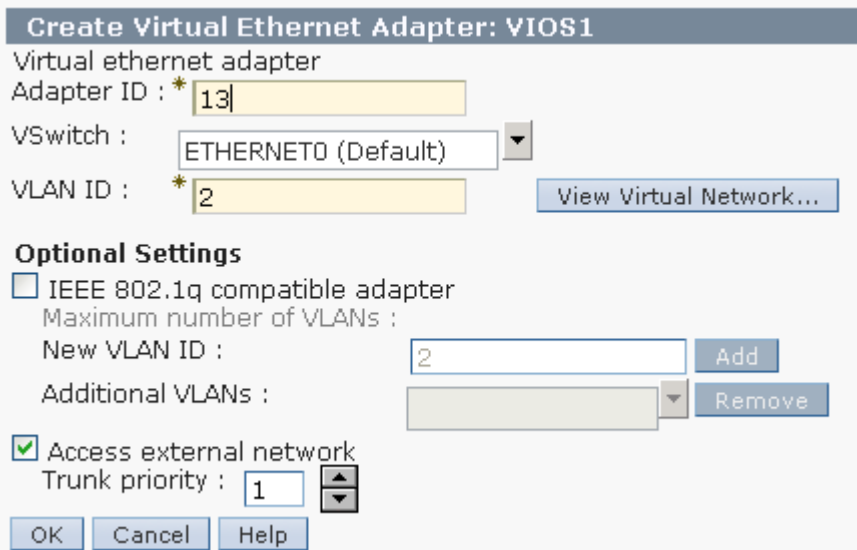
» `lsdev -virtual | grep ent`

» Record the device name of the new virtual Ethernet adapter (you can verify it is the right adapter by matching the adapter ID shown in the location code (with `lsdev -vpd`) with the value you recorded in the table).

### Implementing the Shared Ethernet Adapter failover on VIOS1

We will make VIOS2 the primary SEA provider in this configuration so that we can test failover without disrupting the AIX partitions' access to storage (which is provided by VIOS1). VIOS1 will serve as the backup SEA, but currently VIOS1 is configured as the only SEA supporting VLAN1. In this section you will replace the VIOS1 SEA with one that supports a failover configuration as the secondary adapter on VLAN 2.

5. From the HMC, dynamically add another virtual Ethernet adapter to **VIOS1** on **VLAN2**. Use the next available adapter ID, and record this ID in the table at the beginning of the exercise in the **VIOS1** row, under the **Virtual Ethernet VLAN 2** column. Check the **Access external network** checkbox, and set the **trunk priority** to **2** (VIOS1 will provide the backup SEA in this configuration). The dialog should resemble this one:



- \_\_\_ 6. On VIOS1, run `cfgdev` to discover the new adapter and write down its adapter name in the table at the beginning of the exercise in the **VIOS1** row under the **Virtual Ethernet VLAN 2** column.

» `cfgdev`

» `lsdev -virtual | grep ent`

» Write down the name of the new virtual Ethernet adapter in the table

Now VIOS1 has a virtual Ethernet adapter to support an SEA backup on VLAN 2. Next you will remove the current SEA adapter on VIOS1 and replace it with one that supports VLAN 2 and uses the control channel adapter to support failover mode.

- \_\_\_ 7. If the VIOS1 network access is currently configured on the SEA interface, you will lose external connection when you remove the SEA device. Get the name of the SEA adapter from the table at the beginning of the exercise. Determine on which adapter the IP information is configured for VIOS1 (use `lstcpip`). If it is the SEA adapter, record the configured IP information before you remove the interface in the next step (if you do not have it recorded at the beginning of the lab already).

If the IP information is configured on an additional virtual adapter on VLAN1, you will also lose your external network connection.

- \_\_\_ 8. From a virtual terminal session for VIOS1 remove the SEA adapter and interfaces. Note that LPAR1 and LPAR2 will lose external network connectivity when you do this.

Remove the existing SEA device for VLAN 1. For example, if the **SEA** is **ent2**, the commands to run would be

```
rmdev -dl en2
```

```
rmdev -dl et2
```

```
rmdev -dl ent2
```

- \_\_\_ 9. From the terminal session on VIOS1, recreate the SEA with the new virtual Ethernet adapter on VLAN 2 (use the device name you discovered in step 6 and recorded in the table for VIOS1). When you run the `mkvdev -sea` command, it is very important that you specify the `ha_mode` and `ctl_chan` attributes. If you fail to do this, creation of the primary adapter on VIOS2 could result in a network broadcast storm.

To run this command, in addition to the virtual adapter on VLAN 2, you will need the name of the physical Ethernet adapter on VIOS1 and the control channel adapter you created in the first part of this scenario. These should be recorded in the table

at the beginning of the exercise. For example, if the physical adapter is **ent0**, the control channel adapter on **VLAN 12** is **ent4**, and the virtual adapter on **VLAN 2** is **ent5**, the command to run is:

```
mkvdev -sea ent0 -vadapter ent5 -default ent5 -defaultid 2 -attr
ha_mode=auto ctl_chan=ent4
```

\_\_\_ 10. Write down the name of the new SEA adapter in the table if it is different from what you had recorded before. It should be the same.

\_\_\_ 11. Examine the SEA attributes. Verify that the correct control channel adapter is specified and that the `ha_mode` is **auto**. For example, if the SEA adapter you recorded in step 9 is **ent2**, and the control channel adapter is **ent4**, you would expect this output from `lsdev`:

```
lsdev -dev ent5 -attr
```

```
. . .
ctl_chan ent4 Control Channel adapter for SEA failover
. . .
ha_mode auto High Availability Mode
. . .
```

**IMPORTANT:** If you do not see an `ha_mode` of **auto**, or do not see the correct adapter listed as the control channel adapter, ask your instructor for assistance before you move on.

\_\_\_ 12. If the IP information was previously configured on the SEA interface, use the `mktcpip` command to configure the IP information on the new SEA interface. If the IP information is currently configured on an additional virtual Ethernet adapter on VLAN 1, use `chdev` to unconfigure that adapter before configuring the SEA interface. For example, if the virtual Ethernet adapter interface is **ent3**, the command to run would be:

```
chdev -dev en3 -attr netaddr='' netmask='' state=detach
```

Use the IP information you recorded for VIOS1 and the `mktcpip` command to configure the SEA interface. Run `help mktcpip` from the VIOS1 CLI if you need to see the usage statement.

\_\_\_ 13. From VIOS1, check the connectivity to the external network with a ping to your HMC (using the hostname may not work, but you should be able to ping the IP address).

```
>> $ ping <hmc_ip_address>
```

- \_\_\_ 14. Now that you have an external network connection from VIOS1 again, dynamically remove the virtual Ethernet adapter configured on VLAN1, since it is no longer needed. If there is a second virtual Ethernet adapter defined on VLAN1, remove that one too. Remember that you must remove the ent devices from the VIOS1 CLI before you can remove them dynamically from the HMC.
- » On VIOS1, if the virtual Ethernet adapter on VLAN1 is ent1, you can remove the adapter and interfaces with these commands:
    - » `rmdev -dev en1`
    - » `rmdev -dev et1`
    - » `rmdev -dev ent1`
  - » From the HMC, dynamically remove the adapter from VIOS1. Select VIOS1, and choose **Dynamical Logical Partitioning -> Virtual Adapters** from the selected menu.
  - » Select the virtual Ethernet adapter on VLAN 1 (the adapter ID should be recorded in the table at the beginning of the exercise, and delete it
  - » If there are any other virtual Ethernet adapters defined on VLAN 1, repeat the same steps to remove them.
- \_\_\_ 15. From the HMC session, save your new VIOS1 configuration in a new profile. Name it **SEAFO** (select VIOS1, choose **Configuration -> Save Current Configuration** task).

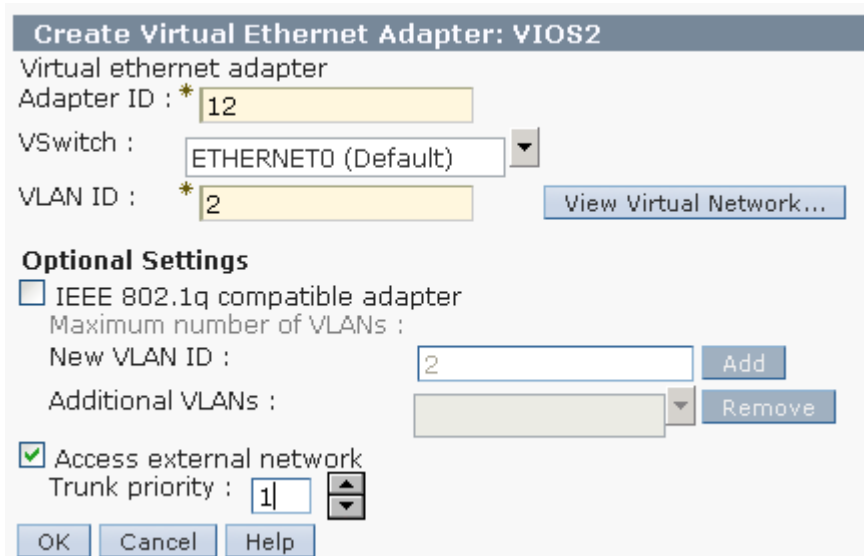
Notice that you are still unable to communicate on the external network from both AIX partitions. This is because VIOS1 is providing an SEA for VLAN 2, and the AIX partitions are using VLAN1. When the SEA Failover configuration is complete, you will configure the AIX partitions to use the highly available VLAN 2.

### **Implementing the primary Shared Ethernet Adapter on VIOS2**

In the following steps you add a virtual Ethernet adapter on **VLAN 2** with **trunk priority 1** to **VIOS2**, to implement the primary adapter in an SEA failover configuration. You will create the virtual Ethernet adapter on VLAN 2, then create the SEA device, specifying the `ha_mode` and control channel adapter. (the virtual adapter on VLAN 12).

- \_\_\_ 16. From the HMC, dynamically create a virtual Ethernet adapter on **VIOS2** on **VLAN ID 2**. Check the **Access external network** checkbox and set the **trunk priority** to **1**. This adapter will be the primary adapter in the SEA failover configuration. You can use the next available adapter ID. Record this adapter ID (or slot number) in the

table at the beginning of the exercise in the **VIOS2** row under the **Virtual Ethernet VLAN 2** column. The dialog should resemble the one shown here:



- \_\_\_ 17. Configure the new adapter on **VIOS2** with `cfgdev`. Use `lsdev -virtual` to list the new virtual Ethernet adapter, and record the device in the in the **VIOS2** row of the table at the beginning of the exercise, under the **Virtual Ethernet VLAN 2** column.
- \_\_\_ 18. It is likely that the VIOS2 partition is currently configured to use its physical Ethernet adapter for the external network connection. Use the `lstcpip` and `lsdev -type adapter` commands to determine if the IP information is configured on the physical adapter (you may have the physical adapter name recorded in the table at the beginning of the exercise).

» `lstcpip`

» Example output:

```
Name Mtu Network Address
en0 1500 link#2 0.14.5e.51.74.23
en0 1500 9.47.87 hurston106.beaver
lo0 16896 link#1
lo0 16896 127 localhost
lo0 16896 ::1
```

» Note the device name for the configured interface. In the example output it is **en0**.

» `lsdev -type adapter | grep ent`

» Example output:

```
ent0 Available Logical Host Ethernet Port (lp-hea)
ent1 Available Virtual I/O Ethernet Adapter (l-lan)
```

» Does the interface name you recorded belong to a physical adapter? In the example output above it is a logical port on an HEA adapter (the output came from a POWER6 system). On your POWER5 system, you will see an Ethernet adapter in a PCI slot for the physical adapter.

- \_\_\_ 19. From a virtual terminal session on VIOS2, remove the IP configuration from the physical adapter interface with `rmtcpip -interface`. For example, if the physical adapter interface is `en0`, you would run the command:

```
rmtcpip -interface en0
```

Answer 'y' to the prompt to continue.

- \_\_\_ 20. On VIOS2, confirm the IP configuration is removed by running `lstcpip`. You should only see a listing for `lo0`.

- \_\_\_ 21. In this step you will create the primary SEA adapter for VLAN 2 on VIOS2.

When you run the `mkvdev -sea` command, it is very important that you specify the `ha_mode` and `ctl_chan` attributes. If you fail to do this, creation of the primary adapter on VIOS2 could result in a network broadcast storm.

To run this command, in addition to the name of the virtual adapter on VLAN 2, you will need the name of the physical Ethernet adapter on VIOS2 and the control channel adapter you created in the first part of this scenario on VLAN 12. These should be recorded in the table at the beginning of the exercise. For example, if the physical adapter is `ent0`, the control channel adapter on **VLAN 12** is `ent1`, and the virtual adapter on **VLAN 2** is `ent2`, the command to run is:

```
mkvdev -sea ent0 -vadapter ent2 -default ent2 -defaultid 2 -attr
ha_mode=auto ctl_chan=ent1
```

A new SEA adapter is created. Write its name in the table at the beginning of the exercise under the **SEA** column for **VIOS2**.

- \_\_\_ 22. Verify that the SEA is properly configured on VIOS2 by checking the output of `lsdev -dev -attr` for the SEA device. For example, if the SEA name you recorded in the table for VIOS2 in step 21 is `ent3`, the command to see the attributes would be:
- ```
lsdev -dev ent3 -attr
```

Verify that the `ha_mode` attribute is set to **auto**, and the `ctl_chan` attribute is set to the control channel adapter on **VLAN 12** that you have recorded in the table for

VIOS2. For example, if the control channel adapter is **ent1**, the output should look like this:

```
. . .
ctl_chan      ent1      Control Channel adapter for SEA failover
. . .
ha_mode       auto       High Availability Mode
. . .
```

STOP!!! Before you continue to the next step, ask a lab instructor to determine that you have the correct adapter configuration. Failure to properly configure an SEA failover scenario can result in a broadcast storm that can affect the entire lab network.

___ 23. Configure TCPIP for VIOS2 on the SEA interface. Use `mktcpip` to configure the interface. Consult `help mktcpip` if you need assistance with the command usage.

___ 24. From a VIOS2 terminal, check connectivity to the external network with:

```
$ ping <hmc_ip_address>
```

___ 25. Using the HMC, save your new VIOS2 configuration in a new profile. Name it **SEAFO** (select **VIOS2**, choose **Configuration** -> **Save Current Configuration** task).

LPAR1 and LPAR2 are still not able to communicate on the external network. Next you will create virtual Ethernet adapters on the AIX LPARs that can communicate on VLAN 2.

Configure AIX LPARs for communication on VLAN2

In this part of the exercise, you will create a new profile for LPAR1 and LPAR2 that has a virtual Ethernet adapter on VLAN 2. You will shutdown and reactivate the LPARs with the new profiles, then test the SEA failover configuration in the next part of the exercise.

___ 26. Shut down LPAR1 and LPAR2, either from virtual terminal sessions with the `shutdown -F` command (*this is the preferred method*), or from the HMC with:

```
chsysstate -o shutdown -r lpar -m managed_sys_name -n LPAR1
chsysstate -o shutdown -r lpar -m managed_sys_name -n LPAR2
```

» You can also shut them down from the HMC GUI

You can continue to the next step while the AIX LPARs are shutting down.

- ___ 27. From the HMC, copy the active profile for both LPAR1 and LPAR2 to a new profile named **SEAFO**:

- ___ 28. Edit the **SEAFO** profile for both LPAR1 and LPAR2. Change the VLAN ID for the virtual Ethernet adapter to VLAN 2 and save the changes. The virtual Ethernet adapter dialog should resemble this one for both LPARs:

- » On the **Virtual Adapters** tab of the profile, select the Ethernet adapter shown and choose Edit
 - » Change the **VLAN ID** to **2**
 - » Make this change in the **SEAFO** profiles for both **LPAR1** and **LPAR2**
- ___ 29. When the AIX LPARs are in the *Not Activated* state, activate them with the modified SEAFO profiles:
- ```
chsysstate -o on -r lpar -m managed_sys_name -n LPAR1 -f SEAFO
chsysstate -o on -r lpar -m managed_sys_name -n LPAR2 -f SEAFO
```

» You can also activate them from the HMC GUI

- \_\_\_ 30. When the AIX LPARs are active, verify that you can reach them via telnet or ssh. Verify that you can ping from LPAR1 to LPAR2, and from each AIX LPAR to the HMC, VIOS1 and VIOS2.

Note that when you send traffic between any two LPARs in the system, it goes over the VLAN in the Hypervisor only (which is really a memory to memory copy from the send buffer of one partition to the receive buffer of another). However, when you send traffic from an LPAR to the HMC or gateway IP address, you are using the active SEA. Right now the active SEA is on VIOS2. In the next part of the exercise, you will test failover to VIOS1.

### **Check network failover and recovery**

In the last part of this scenario, you will test the failover mechanism of the configuration. You will shut down VIOS2, which is currently providing all LPARs with access to the external network. When you do this, the Hypervisor will make the backup SEA active on VIOS1 by detecting the outage on VIOS2 through the control channel adapter on VLAN 12. The AIX LPARs should not experience any disruption of their network connection.

- \_\_\_ 31. Use **entstat** on both of the VIO Servers to verify the Primary SEA is bridging the ethernet frames/packets to the external network.

» `$ entstat -all ent<SEA> | grep Active`

You should see something like the following on VIOS2:

```
Priority: 1 Active: True
```

And something like this on VIOS1:

```
Priority: 2 Active: False
```

This implies the primary trunk adapter on VIOS2 is active.

- \_\_\_ 32. Start a ping session from LPAR1 and LPAR2 to your HMC and verify service is not disrupted.

»# `ping <hmc_ip_address>`

- \_\_\_ 33. Open a virtual terminal window to shut down VIOS2 using the **shutdown** command. This will cause the primary SEA to fail.

- \_\_\_ 34. Check that the SEA Failover feature has switched to the Backup SEA and that VIOS1 is bridging ethernet frames to the external network. It may take up to a minute to see the status of the backup adapter to change to Active. Use the following command on VIOS1

```
$ entstat -all ent<SEA> | grep Active
```

You should see something like the following on VIOS2:

```
Priority: 2 Active: True
```

This implies the secondary trunk adapter is active.

- \_\_\_ 35. Stop the ping commands from LPAR1 and LPAR2 by typing `Ctrl^c` in both windows, and bring VIOS2 back online with the SEAFO profile:

```
chsysstate -o on -r lpar -m managed_sys_name -n VIOS2 -f SEAFO
```

Does the active SEA change? Check with `entstat` on both VIO servers.

- » The primary adapter on VIOS2 becomes active when VIOS2 comes online. The Hypervisor automatically switches back to the primary adapter, with complete transparency to the network clients.

You have completed the SEA Failover scenario - Congratulations!. Let the instructor know that you are finished, and if you have any questions.

## **END OF LAB**