

QUESTION: How do you setup NIC Teaming

...



groundLoop 31 posts since

Jul 28, 2006 I just managed to implement NIC Teaming in conjunction with VLAN trunking on a pair of ESX 3.0 Servers. It works, but I'm not sure that I understand why.

The physical switch is a big Cisco 4507. What caught me off guard is that I did not need to enable Cisco's Gigabit EtherChannel functionality. I'm used to bonding together NICs in EtherChannels. So how is this working? Can anyone explain it?

Here's the basics of my config:

```
interface GigabitEthernet3/10
description VMware ESX - NIC 0 - Trunk A
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100,200
switchport mode trunk
switchport nonegotiate
speed 1000
spanning-tree portfast
end

interface GigabitEthernet3/11
description GRANT - VMware ESX - NIC 1 - Trunk B
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100,200
switchport mode trunk
switchport nonegotiate
speed 1000
spanning-tree portfast
end
```

Note that I'm not using the typical "channel-group XX mode on" command to bond both ports in to a Port-Channel. As a matter of fact, if I try to use EtherChannel, I start getting very strange behavior.

Would anyone be kind enough to explain the various Virtual Switch Load Balancing properties?

-groundLoop

Message was edited by: Updated topic to question?

groundLoop

QUESTION: How do you setup NIC Teaming ...



[Paul Lalonde](#) 767 posts since

Jan 16, 2006 1. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 7:05 AM

Hi,

If I'm not mistaken, most of the newer IOS-based Catalyst switches (including the Sup IV in the 4507) support auto PAGP and LACP, which means PAGP / LACP will detect directly connected switches / hosts with matching aggregation characteristics and build the EtherChannel automatically.

You can always check the log to see if EtherChannels have been built automatically... entries will exist if either PAGP or LACP have dynamically created the channel.

Regards,

Paul



[groundLoop](#) 31 posts since

Jul 28, 2006 2. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 7:34 AM

👤 in response to: [Paul Lalonde](#) Paul,

Thanks for the tip. Interestingly enough, I couldn't find any reference to automatically created PAGP, LACP, or EtherChannel groups on the switch.

I suspect this is working because of VMware's vSwitch default NIC teaming support. But just how it works, I'd like to know.

-gL



[Quotient](#) 394 posts since

Nov 30, 2005 3. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 9:03 AM

Ah, my pet favourite...

ESX doesn't support dynamic IEEE 802.3ad Link Aggregation Groups (pagp & lacp port groups).

It does however support static IEEE 802.3as LAGs (Gigabit / Fast EtherChannel).

Some "network" guys will usually tell you to avoid FEC/GEC, but from my experience this is misguided.

The again I am a "server" guy, so perhaps I'm a bit biased... 😊

You need to set up a channel group and observe the load balancing mechanism that is used, e.g. show etherchannel load-balance...

Hopefully, it is src-dst-ip. It doesn't matter if it's not.

This one is just the most flexible - especially if you're using NLB...

Consider changing it if there's little chance of impact...

The trick here is to setup your vSwitch load balancing policy to be compatible.

src-mac, dst-mac, src-dst-mac = MAC hash

QUESTION: How do you setup NIC Teaming ...

src-ip, dst-ip, src-dst-ip = IP hash


Use a pSwitch config like this:

```
!  
interface port-channel1  
description VMware ESX - Trunk A  
switchport trunk encapsulation dot1q  
switchport trunk allowed vlan 100,200  
switchport mode trunk  
switchport nonegotiate  
speed 1000  
spanning-tree portfast trunk  
!  
exit  
!  
interface GigabitEthernet3/10  
description VMware ESX - Trunk A - NIC 0  
switchport trunk encapsulation dot1q  
switchport trunk allowed vlan 100,200  
switchport mode trunk  
switchport nonegotiate  
speed 1000  
spanning-tree portfast trunk  
channel-group 1 mode on  
!  
exit  
!  
interface GigabitEthernet3/11  
description VMware ESX - Trunk A - NIC 1  
switchport trunk encapsulation dot1q  
switchport trunk allowed vlan 100,200  
switchport mode trunk  
switchport nonegotiate  
speed 1000  
spanning-tree portfast trunk  
channel-group 1 mode on  
!  
end
```

QUESTION: How do you setup NIC Teaming ...

That should do the trick!

Ben

EDIT: Corrected syntax error 



[Quotient](#) 394 posts since

Nov 30, 2005 4. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 8:07 AM

Just an update to say that without GEC / FEC trunks, you will end up with a broadcast storm under load.

You will also find that you are probably only transmitting / receiving data on one nic...


I believe it works in the default configuration without EC because the "originating virtual port ID" (whatever that is) uses a combination of MAC and IP load balancing algorithms or uses a src-dst-mac algorithm...

I wish these setting used industry standard terms or definitions. It would make life easier...



[groundLoop](#) 31 posts since

Jul 28, 2006 5. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 8:55 AM

 in response to: [Quotient](#) Quotient,

Thanks again. I've applied your recommended configuration to our "pSwitch".

At first, I struck out. We lost connectivity to the VMs. Then I attempted to identify the load balancing mechanism on the pSwitch. I'm not sure that I was able to establish the method being used. Here's the port-channel details:

```
myswitch#show etherchannel 25 detail
```

```
Group state = L2
```

```
Ports: 2 Maxports = 8
```

```
Port-channels: 1 Max Port-channels = 1
```

```
Protocol: -
```

```
Ports in the group:
```

```
-----
```

```
Port: Gi3/10
```

```
-----
```

```
Port state = Up Mstr In-Bndl
```

```
Channel group = 25 Mode = On/FEC Gchange = -
```

```
Port-channel = Po25 GC = - Pseudo port-channel = Po25
```

```
Port index = 0 Load = 0x00 Protocol = -
```

```
Age of the port in the current state: 00d:00h:26m:53s
```

```
Port: Gi3/11
```

```
-----
```

QUESTION: How do you setup NIC Teaming ...

Port state = Up Mstr In-Bndl
Channel group = 25 Mode = On/FEC Gcchange = -
Port-channel = Po25 GC = - Pseudo port-channel = Po25
Port index = 1 Load = 0x00 Protocol = -
Age of the port in the current state: 00d:00h:26m:53s
Port-channels in the group:

Port-channel: Po25

Age of the Port-channel = 06d:19h:09m:23s
Logical slot/port = 11/25 Number of ports = 2
GC = 0x00000000

Port state = Port-channel Ag-Inuse

Protocol = -

Ports in the Port-channel:

Index Load Port EC state No of bits

-----+-----+-----+-----+-----

0 00 Gi3/10 On/FEC 0

1 00 Gi3/11 On/FEC 0

Time since last port bundled: 00d:00h:26m:53s Gi3/11

Time since last port Un-bundled: 01d:02h:24m:17s Gi3/11

Figuring trial and error couldn't hurt, I tried adjusting the Load Balancing options on the vSwitch.

Route Based on Originating Virtual Port ID fails.

Route Based on ip hash works

Route Based on source MAC hash fails.

If I compare and contrast these results with my prior efforts, Route Based on Originating Virtual Port ID appeared to work when I did not explicitly define the GEC/FEC port-channel.

I wish I understood this stuff better. And I agree completely with your assessment, it would be great to have some agreed-upon nomenclature.

So now I have a working Port-Channel, and I am using the "Route Based on IP hash" algorithm. Am I headed in the right direction?

Thanks again for the advice.

-gL

QUESTION: How do you setup NIC Teaming ...



[Quotient](#) 394 posts since

Nov 30, 2005 6. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 9:07 AM

↑ in response to: [groundLoop](#) Definitely headed down the right path...

In fact you've arrived...!



FYI:

To display the load balancing algorithm on the pSwitch issue:

show etherchannel load-balance

For algorithm options, enter global configuration mode (conf t) and issue:

port-channel load-balance ?

Ben



[groundLoop](#) 31 posts since

Jul 28, 2006 7. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 10:07 AM

↑ in response to: [Quotient](#) Quotient,

You've got this stuff down cold. Thanks.

We are currently running at:

#show ether load-balance
Source XOR Destination IP address

So, it appears that the IOS option for "Source XOR Destination IP address" load balancing is equivalent the VMware's "IP hash".

I take it that this is default behavior on the Sup IVs?

I just caught up on your earlier solution too, <http://www.vmware.com/community/thread.jspa?threadID=49110>.

This is starting to make sense.

I just learned another interesting fact from VMware support. ESX supports **outbound** load balancing by default. EtherChannel is required for inbound load balancing.

Is there any advantage to the different load-balancing algorithms?

-gL [edit: updated with info from Support Request]



[Quotient](#) 394 posts since

Nov 30, 2005 8. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 10:17 AM

QUESTION: How do you setup NIC Teaming ...

 in response to: [groundLoop](#) Ah, the cream - here it is:

Use vSwitch ip hash with:

src-ip#Load distribution on the source IP address

dst-ip#Load distribution on the destination IP address

src-dst-ip#Load distribution on the source XOR destination IP address

Use vSwitch source MAC hash with:

src-mac#Load distribution on the source MAC address

dst-mac#Load distribution on the destination MAC address

src-dst-mac#Load distribution on the source XOR destination MAC address

Use vSwitch originating virtual port ID with:

src-port#Load distribution on the source port

dst-port#Load distribution on the destination port

src-dst-port#Load distribution on the source XOR destination port

From experience src-dst-ip has better interoperability with unicast and multicast load balancing solutions, e.g. NLB, radware, F5 Big-IP, etc.

In addition, this seems to be the default for most new switches.

I just wish it was VMware's because then the switch ports could be fully commissioned (with GEC) rather than having to change the vSwitch property first.


It would also mean that ESX would support receive load balancing out of the box in most new environments.

note: <http://www.vmware.com/community/thread.jspa?threadID=49308>



[groundLoop](#) 31 posts since

Jul 28, 2006 9. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 10:35 AM

 in response to: [Quotient](#) Quotient,

Excellent input. You've not only given me the **correct answer**, but provided more than a few very helpful answers. I can't seem to figure out the points/award system, sorry, but I'd hook you up with the full measure (22 points?) if I could.

Thanks again.

-gL



[Quotient](#) 394 posts since

Nov 30, 2005 10. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 10:44 AM

 in response to: [groundLoop](#) no problem, gL...

You need to post the topic as a question to use the points system...

... but that's okay, because all good things come to those who wait 😊

QUESTION: How do you setup NIC Teaming ...

Glad I could help.



[groundLoop](#) 31 posts since

Jul 28, 2006 11. **Re: NIC Teaming and VLAN Trunking with Cisco Switches** Jul 28, 2006 10:57 AM

👤 in response to: [Quotient](#) Looks like its too late now to convert this thread into a question. I tried editing the original post, but the option to flag the thread as a question is not available. Maybe a benevolent moderator can reverse my mistake and award you the points you deserve.

Thanks again!

-gL



[Quotient](#) 394 posts since

Nov 30, 2005 12. **Re: QUESTION: How do you setup NIC Teaming and VLAN Trunking w/Cisco gear?** Jul 28, 2006 12:12 PM

Just on a side note, you should also consider using a dummy VLAN for the native VLAN.

Create the VLAN and then issue the following command for the port-channel group and all switchports that are members:

switchport trunk native vlan <dummy VLAN ID>