

Communities

All Communities ▾

Browse ▾

LOGIN / REGISTER

▲ Up to Discussions in VI: VMware ESX™ 3.5

This Question is **Answered** ★ 2 "helpful" answers available (6 pts)

17 Replies Last post: Jul 22, 2008 8:55 AM by eric.heilig

1 2 Previous Next

Nic Teaming load balancing policy... posted: Aug 23, 2006 2:40 PM

**ErMaC**300 posts since
Mar 29, 2006

So I read the ESX Server Configuration manual regarding the load balancing policy. I decided that load balancing by IP hash, as I understood it, would give the best load balancing. From what I understand:

Port ID based: all traffic from one VM will go out from one uplink.
IP hash based: traffic from one source IP to a destination IP goes out the same uplink
Mac hash based: traffic from one MAC address goes through the same uplink.

This sounded to me like IP-based would be the best. But that doesn't appear to be the case. I was using port-based before, and once I switch to IP-based, the traffic on my secondary nics dropped to 0. What gives? Has anyone done testing between all three types to determine under what loads which balancing method is best?

Search VMware Communities

VI: VMware ESX™ 3.5

Search

Advanced Search

Actions

View as PDF

Only VMware ...
Delivers Nexus
1000V

Ensure consistent, policy-based network capabilities to virtual machines across your data center.

Learn More

Re: Nic Teaming load balancing policy...

1. Aug 23, 2006 2:43 PM in response to: ErMaC

**soleblazer**234 posts since
May 25, 2006

Take a look at this thread:

<http://www.vmware.com/community/thread.jspa?messageID=445938񬷲>

Are you looking at inbound? Without etherchannel only outbound traffic will be load balanced, but that thread should answer your questions.

Re: Nic Teaming load balancing policy...

2. Aug 23, 2006 4:20 PM in response to: soleblazer

**ErMaC**300 posts since
Mar 29, 2006

From talking with a VMware technical rep they said we don't need to configure anything on the switches and we would get load balancing.

We have a pair of Cata 2960G's and each ESX Host is connected to both for failover. When I use port-based balancing, I'm seeing load balance across both nics for **both inbound and outbound**, but when I select the other two I see no balancing on either.

Note we didn't configure any sort of etherchannel or anything on those switches, because as far as I know you cannot etherchannel between two switches, correct?

Re: Nic Teaming load balancing policy...

3. Aug 23, 2006 5:16 PM in response to: ErMaC

**soleblazer**234 posts since
May 25, 2006

Yep, no etherchannel between switches.

The only way (from what I have gathered) to load balance inbound is with etherchannel, you can tweak outbound load balancing by adjusting the options you have used. I personally have left to the default and have seen very good performance.

Re: Nic Teaming load balancing policy...

4. Aug 25, 2006 8:20 AM in response to: ErMaC



Quotient
394 posts since
Nov 30, 2005

Just so we're all clear on the facts... 😊

The load-balancing policy (frame distribution) can be any of the three options (MAC, IP or Port) to effectively achieve the same outcome, provided EtherChannel is used.

MAC based solutions balance according to the L2 address.
IP based solutions balance according to the L3 address.
Port based solutions balance according to the L4 address.

When using port-based (L4) on the ESX host, on even one pSwitch and without EtherChannel, it could appear that you are achieving bi-directional load-balancing because the outbound traffic is sent to the same port that it originated on (note the wording of the option).

HOWEVER, during a failover the solution will DIE without EtherChannel, with recovery taking several seconds to complete. Disabling STP or using a portfast policy will minimise this but a significant delay will still be observed. This might be acceptable in some situations.

In all configurations, including those with and without EtherChannel, the maximum bandwidth available between any **two** PHYSICAL hosts is the maximum bandwidth of **one** pNIC.

When attempting to achieve switch redundancy you should use beacon probing and your choice of originating virtual port ID or explicit failover order. Just expect long delays during failover as the pSwitch or pSwitches relearn their frame distribution topology.

When you have highly available switch complexes, including modular and stackable switch clusters, that provide similar or better redundancy as two separate physical switches, you can take advantage of the benefits of EtherChannel / IEEE 802.3ad STATIC Link Aggregation Groups (LAGs).

VMware's implementation of IEEE 802.3ad allows for the distribution of packets according to the following algorithms: src-port (L4), src-dst-mac (L2) & src-dst-ip (L3).

EtherChannel can only be used to achieve equal distribution of packets across all ports if the portgroup contains 2, 4 or 8 ports. All other combinations, i.e. 3, 4, 5, 6 & 7, will result in uneven distribution of packets.

Last point is just a matter of semantics... 😊

You can EtherChannel between two switches. This is what it was primarily designed for. You cannot however, connect more than two devices together in any IEEE 802.3ad compliant solution. Therefore, a solution where two pSwitches are configured to use EtherChannel when connected to a single ESX vSwitch is not possible.

Hope this helps,
Ben

Re: Nic Teaming load balancing policy...

5. Aug 25, 2006 8:17 AM in response to: [Quotient](#)



ErMaC
300 posts since
Mar 29, 2006

Well I guess we're sticking to the port-based one since we can't enable Etherchannel in our current setup. That's fine, so far the traffic's been normal enough and we're really only doing it for redundancy purposes. A few seconds of downtime (I counted 1 or 2 pings during our testing) is acceptable for us.

Re: Nic Teaming load balancing policy...

6. Sep 22, 2006 11:26 PM in response to: [Quotient](#)



vmwareuser11
2 posts since
Sep 19, 2006

This thread is informative. Can switch redundancy be obtained with IP load balancing?

Re: Nic Teaming load balancing policy...

7. Sep 24, 2006 1:04 AM in response to: [Quotient](#)



howie
107 posts since
Feb 23, 2004

Port based solutions balance according to the L4 address.

This is NOT true. The "port" in Port based LB refers to the virtual switch port, not the L4 port.

I also wanted to point out that a user should only use IP based LB if EtherChannel or equivalent is configured on the physical switch.

Also, VMware does not really recommend EtherChannel in ESX3.x in general, as it does not really buy anything in most configurations.

-howie

Re: Nic Teaming load balancing policy...

8. Mar 5, 2007 12:53 PM in response to: [howie](#)



vmproteau
259 posts since
Aug 12, 2006

I somewhat understand the Load Balancing Policies for a vSwitch and for us L4 port based load balancing will be what we generally use however; is there anyway to increase the bandwidth to more than one NICs worth?

Here is the physical layout:
vSwitch in a load balanced configuration (2 NICs)
Physical NIC1 goes to trunk port on Cisco SwitchA
Physical NIC2 goes to trunk port on Cisco SwitchB

We have a Cisco team so I only have a rudimentary knowledge of the physical switch setup but as I understand it, even if you Ether Channel the 2 physical switch ports, you will still have a spanning tree problem because of the loop that is created between the 3 switches.

Am I wrong....and (in this HAswitch configuration) is there anyway to combine the bandwidth of 2 physical NICs so that a virtual Machine is utilizing both cards?

Re: Nic Teaming load balancing policy...

9. Apr 3, 2007 10:52 AM in response to: [vmproteau](#)



MikeAvery
19 posts since
Nov 17, 2006

If those two pNics are connected to a single pSwitch, yes.

If those pNics are each connected to a separate pSwitch, you will be able to the bandwidth of both pNics to the vSwitchPort, with a caveat: you cannot exceed the bandwidth of one pNic between the two MACs (1:1). You can use the combined bandwidth for 1:N

-Three machines (A,B,C) try to push a large file to one VM that is configured as you describe above.

Using Ip-hash for load balancing, M<-->A will never exceed bandwidth of 1 pNic

Using Ip-hash for load balancing, M<-->A,B,C combined can achieve the aggregate bandwidth of both pNics

If you use Etherchannel on the pSwitch establish a match between pSwitch and vSwitch load balancing policy, you can use all bandwidth for M<-->A, so long as both pNics are on the same switch.

Hope this helps

Message was edited by:
MikeAvery

Re: Nic Teaming load balancing policy...

10. Apr 3, 2007 11:01 AM in response to: [MikeAvery](#)



Paul Lalonde
767 posts since
Jan 16, 2006

Actually Mike, you're mostly right, but not entirely.

Etherchannel (or, 802.3ad static) provides load balancing for multiple flows. A flow is defined as a traffic stream with the same SRC and DST. With Etherchannel, a single flow will always take a single path, regardless of the # of outbound network interfaces. There is no load balancing or load aggregation of a single flow.

The benefit of etherchannel is realized with multiple flows. Each unique pairing of SRC / DST streams can be handed off to their own outbound interface for transmission onto the network.

(Warning: oversimplifying the following for illustration, but 802.3ad hash mechanism isn't this cut-and-dry):

For example, take an ESX server with 3 outbound NICs for VMs connected to a Cisco switch with Etherchannel enabled. Our virtual machine is called VM. If VM talks to HostA on the network, it will use physical NIC #1. If VM talks to HostB, it will use physical NIC #2. If it talks to HostC, it will use physical NIC #3. This way, a single VM can balance its outbound traffic across different physical NICs when talking to different hosts, but it will only ever use one path when talking to an individual host.

As more VMs are added to ESX and there are more destinations for those VMs to talk to, the more the flows get "shared" amongst the outbound physical NICs.

Etherchannel was never designed to "aggregate" traffic for a single flow because this would have resulted in out-of-order packet delivery, ultimately grinding such upper layer protocols like TCP/IP down to a fraction of available bandwidth.
Paul

Re: Nic Teaming load balancing policy...

11. Apr 3, 2007 11:22 AM in response to: [Paul Lalonde](#)



Mike Avery
19 posts since
Nov 17, 2006

I think we agree Paul! Maybe you can answer a question for me if you're around?

I have 4 pNics, with two connecting to separate Cisco switches.

I currently have a half-baked implementation:

All 4 pNics in one vSwitch using port groups, including console, vmkernel and vmnet. The pSwitches are configured as dot1q trunk ports.

I'm looking to use Etherchannel and configure a compatible load balancing policy. Must I use a different vSwitch for each channelled pair? I can see having to manually "handle" spreading out the port group assignment among virtual machines. Perhaps I am better to put on channelled pair in standby mode within the one vSwitch?

Any suggestions would be great.

Mke

Message was edited by:
MikeAvery

Re: Nic Teaming load balancing policy...

12. Apr 3, 2007 12:13 PM in response to: [Mike Avery](#)




Paul Lalonde
767 posts since
Jan 16, 2006

Honestly, Mike, I think leaving Etherchannel *off* would be best in this scenario. Just keep the one vSwitch and rely on the default port-based load balancing teaming method across the 4 pNics. This way, you should see the best distribution of traffic across all interfaces while retaining the obvious benefit of failover.

Paul

Re: Nic Teaming load balancing policy...

13. Apr 3, 2007 12:57 PM in response to: [Paul Lalonde](#)




Mike Avery
19 posts since
Nov 17, 2006

Thanks for your thoughts.

Mke

Re: Nic Teaming load balancing policy...

14. Apr 3, 2007 1:15 PM in response to: [Mike Avery](#)




Paul Lalonde
767 posts since
Jan 16, 2006

You're welcome, Mike. As 'howie' says (above), Etherchannel doesn't always buy better performance or reliability. In your split pSwitch config, it doesn't make much sense.

Paul

[Go to original post](#)

1 2 [Previous](#) [Next](#)

[Go to original post](#)

[Up to Discussions in VI: VMware ESX™ 3.5](#)

Community Resources

Developer Center
Documentation
Technical Papers
Compatibility Guides

VMworld

VMworld.com
VMworld Session & Labs

Get Support

Downloads & Patches
Product Support Centers
Manage Licenses
Search the KnowledgeBase

Other Links

Contact Us
About VMware
Jobs at VMware
Customer Case Studies
Events

