

**Main Menu**

- Home
- About Me
- Tech Blog
- Tech Presentations
- Pictures
- Carputer Project

**Sponsored Links**

Ads by Google

- [Network Diagrams](#)
- [Configure Router](#)
- [VM](#)
- [8 Port Router](#)

**Find Me On...**

**Sponsored Links**



Watch this video on [www.youtube.com](http://www.youtube.com)

**SQL server on VMware**  
 Use Veloxum to improve performance immediately, Free Trial  
[youtube.com/ks cruzen](http://youtube.com/ks cruzen)

Ads by G

## vSphere Host NIC Design - 6 NICs

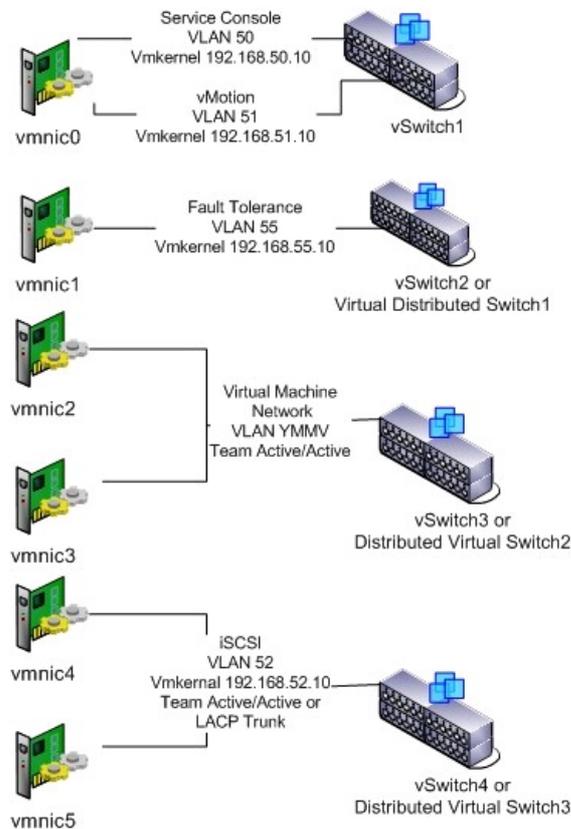
Blog - Tech Blog  
 Wednesday, 28 April 2010 15:27

5 Comments

Another free consulting gig from yours truly. I was asked for some help through [VMTN Forums](#) and [LinkedIn](#) on planning a host NIC design. The design incorporated 6 NICs and it was going to be a proof of concept based on all vSphere features. Well, I couldn't just do a design based on that; I figured I would diagram out a few different solutions for the masses.

So here we go. If you want to do just a proof of concept and don't care about doing it "right", you can always design it as pictured below. This design gives you the ability to incorporate all of vSphere's features and should have plenty of bandwidth to take care of everything. Of course, this is proof of concept so I didn't take NIC redundancy into account. This design is solely to give you the ability to use all of vSphere's features. One thing I constantly see mistaken on VMTN forums about NIC design is Fault Tolerance network. To my understanding, when you enable FT on a VM, there is going to be traffic flowing through that particular NIC and VLAN so you want to have it segregated from everything else. So if you are designing a vSphere environment to use Fault Tolerance, I would think about adding more NIC cards and checking out my blog post on [vSphere Host NIC Design - 10 NICs](#). **\*\*UPDATE 5/28/2010\*\*** Check the bottom of the page for a layout to use FT.  
**NOTE:** vmnic0, vmnic2, vmnic3 must all be configured as trunk ports on your physical switch, use tagging on your vSwitch Port Groups to allow traffic to flow.

### All Features get plenty of bandwidth, but no redundancy 6 NICS



By: Kendrick Coleman  
[Http://www.kendrickcoleman.com](http://www.kendrickcoleman.com)  
 Twitter: @KendrickColeman

Next I decided to do a few different kinds of scenarios if you want an enterprise ready solution. Pl

**My Tweets**

about 3 hours ago  
 Hah RT @techmute: RT @ShitMyDevOpSays: you think doing it right is expensive? wait until you see what doing it wrong is going to cost...

about 23 hours ago  
 so much to get accomplished before VMworld. might have to take a day or two off work. maybe 4 days so i can study for VCAP

[vSphere Virtualization NIC Design](#)
[VMware Free Tools](#)
[Cloud Windows](#)
[Cisco Network](#)
[Carputer Security](#)
[SAN Backup](#)
[EMC NetApp](#)
[Virtual Machine](#)

NIC layout: Do you have 2 on board and 4 in 1 expansion card, or 2 2x expansion cards? And I will give you the option of doing something a bit different if you think you might hit a VM network bottle

These all have a few things in common:

1. The Layer 3 Switch.
  - If you are a SMB, you will most likely use some sort of stacked switch such as the Cisco 3750G or 3750E series. If you have a little bit larger enterprise, you will likely go with a 4500, 6500 or Nexus 7000 type of core switch. Do not think that 2 switches connected GB uplink port suffices as a "stack". A stack solution such as the 3750G has a 32Gb/s interconnect and the 3750E has a 64Gb/sec interconnect.
2. Trunk Ports
  - I found myself configuring every switch port as a trunk port and tagging the VLAN at the port group when designing for only 6 NICs. Teaming NICs on fewer vSwitches gives no redundancy, but keeping them all logically separated will cut down on broadcast traffic network noise.
3. Jumbo Frames
  - For jumbo frames to work, they must be configured from end-to-end. The vmnic, the physical switch, and the SAN must all be set to use 9000 mtu. If one piece is missing, network will fail. I particularly only deploy Jumbo Frames on the storage network because that you have the most to gain.
  - How to configure on vDS - <http://blog.scottlowe.org/2009/05/21/vmware-vsphere-vds-ports-and-jumbo-frames/>
  - How to configure on vSwitch - <http://blog.scottlowe.org/2008/04/22/esx-server-ip-storage-jumbo-frames/>
4. LACP
  - Link Aggregation will benefit you most on the iSCSI storage network.
  - To properly configure LACP, read KB Article: [1004048](http://www.kendrickcoleman.com/1004048)

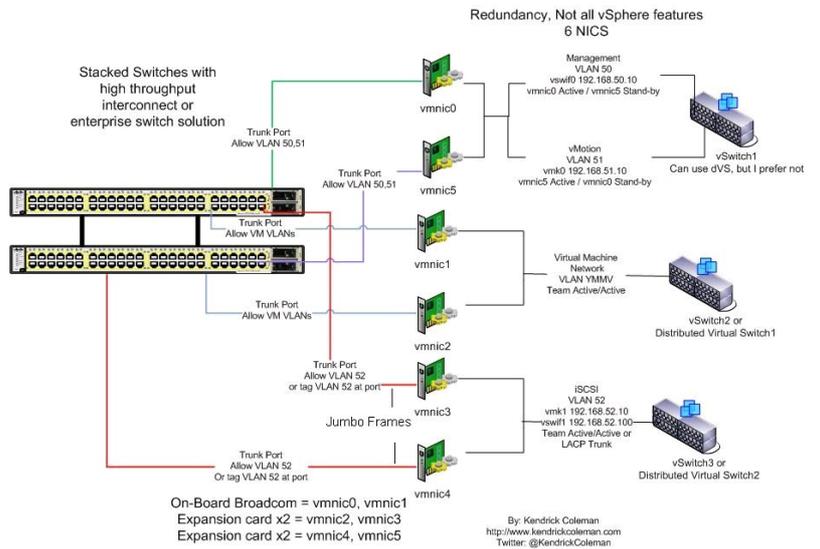
This is a traditional VMware 3.5 design and not based on vSphere alone. This is what I would call the "safe route", but doesn't take into account a Fault Tolerance network.

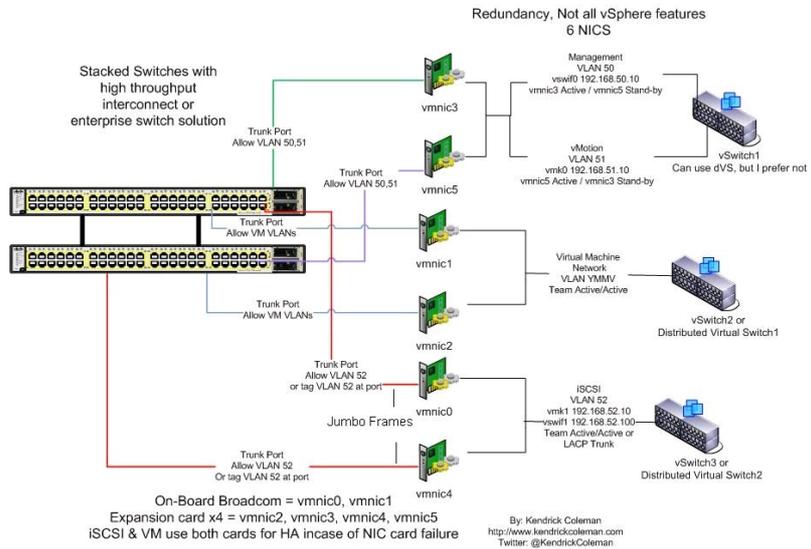
The difference in the diagrams takes into account the type of NIC expansion cards in your physics

6 NICs = 2x2x2 or 2x4

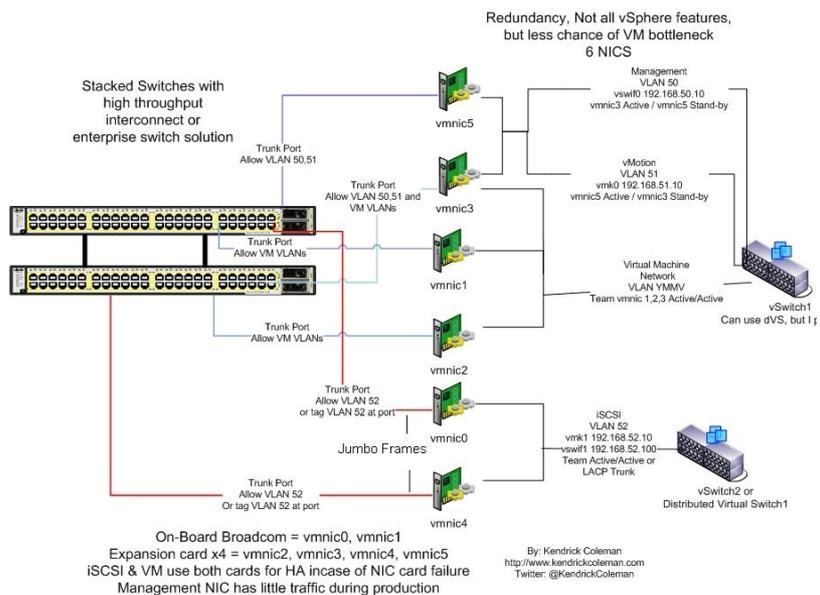
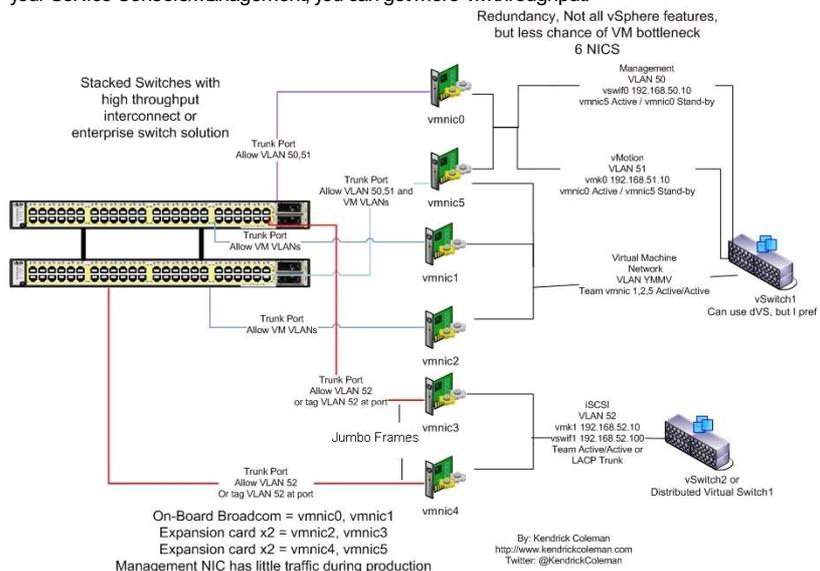
- If you have 6 NICs, you get there by having 2 onboard NICs, and either 2 x2 NIC expansion cards or 1 x4 NIC expansion card
- I particularly don't like to mix on-board with expansion in iSCSI because of fear of "flow control" only way I could get around a proper HA configuration on 2x4 is to have iSCSI and VM network port on both the on-board and expansion card.

**NOTE:** On the next 4 pictures, your SAN must be directly connected into this layer 3 switch. There are not routing your iSCSI traffic, but only switching it, vastly increasing performance.





I thought I would try out this next scenario if you feel that only having 2 NICs for VM Network Traffic not be enough. The elimination of one vSwitch puts Service Console/Management, vMotion, and Network all on 1 vSwitch or 1 dVS if you prefer. By actively using vmnic5 as both VM Network Traffic your Service Console/Management, you can get more VM throughput.



If you belong to a medium or larger enterprise, I would physically separate the storage network or layer 2 stacked switch solution. Granted, the examples above will not have any routing going on if SAN is also directly connected into the stacked layer 3 switch, but if you have the means to physically separate the storage network, you have more options to play with.  
**NOTE: Below is an example of physically separating the storage network. ANY of the previous examples will work (I just didn't want to reiterate the same diagrams) and should be configured the way you have the means.**



I see that you have the service console & vmotion on vswitches rather than on distributed and the shared. I know that the service console has very little traffic and is probably safer on the vswitch : system is accessible on its own but I'm surprised to see vmotion there. Wouldn't that need its o and be suited to be on the dv? Or is it not considered to have a good deal of traffic normally as w Also, I see you have created multiple dvswitches instead of creating 1 with multiple portgroups. to better segregate the physical NICs?

**From Kenny:**

vmotion will ONLY have traffic on it during a vMotion. It shouldn't be happening all the time and DF it for you automatically (depending on licensing of course).

You have the option of putting EVERYTHING on a dVS, but I prefer to have my service console and sitting on 1 vSwitch with 2 NICs and use each others NICs as active/stand-by. This way you can c using only 2 NICs for both, where as if you want to have redundancy, you have to have 2 NICs for taking up 4 NICs total. vCenter is the control plane for vDS while the ESX host is the data plane. I goes down, your dVS does not go down. Traffic will still flow, but since vCenter isn't available, you make any changes.

I would always recommend using a dVS for VMNetwork traffic

It's your choice whether or not you want to create just 1 dvSwitch or multiple ones. I do it to have a look inside of the VI Client and it makes you separate the traffic a bit more on that level. Remember virtual switch is kind of like a physical switch. Whatever is plugged into that switch can talk before and going back to the physical network (not like a hub, but like a switch). So having multiple vSwitch and dVS is just further segregating traffic.

Loading comments...

blog comments powered by DISQUS

[< Prev](#)

[Next >](#)

[b](#)