

blog.scottlowe.org

The weblog of an IT pro specializing in virtualization, storage, and servers

 Subscribe to feed

[Home](#) [About](#) [Archives](#) [Tags](#) [Sponsors](#)

[NetApp Blog Aggregator?](#) • [Stolen Content](#)

SITE SPONSORS



RELATED POSTS

[VMware vSphere vDS, VMkernel Ports, and Jumbo Frames](#)

[Enabling Jumbo Frames on a Nexus 5000](#)

[New User's Guide to Configuring VMware ESX Networking via CLI iSCSI from NetApp to ESX Server Sanrad Configuration Basics](#)

ESX Server, IP Storage, and Jumbo Frames

Tuesday, April 22, 2008 in [Storage](#), [Virtualization](#) by [slowe](#) | [79 comments](#)

With the release of VMware Infrastructure version 3.5, VMware added support for [jumbo frames](#). Although the [documentation](#) states that jumbo frames “are not supported for NAS and iSCSI traffic”, jumbo frames for NFS and iSCSI does actually work. Here’s some information on getting it working.

HOW I TESTED

Keep in mind that this is not an “officially supported” configuration (see the section on the “Official” Support Statement below), so **use at your own risk**. I will not be held responsible if you blow up your production environment trying to make jumbo frames work.

Here’s how I tested the use of jumbo frames for software iSCSI and NFS datastores:

- For the physical switch infrastructure, I used a Cisco Catalyst 3560G running Cisco IOS version 12.2(25)SEB4.
- For the physical server hardware, I used a group of HP ProLiant DL385 G2 servers with dual-core AMD Opteron processors and a quad-port PCIe Intel Gigabit Ethernet NIC.
- For the storage system, I used a NetApp FAS940 running Data ONTAP 7.2.4.

The exact commands and/or procedures may be different for you depending upon the hardware and/or software versions that you’re running in your environment. Keep that in mind.

CONFIGURING THE PHYSICAL SWITCH

Fortunately for me, the Cisco Catalyst 3560G does indeed support jumbo frames. (Naturally, you’ll want to ensure that your switch supports jumbo frames.) Jumbo frames are not, however, enabled by default; they must be enabled using the following command in global configuration mode:

```
system mtu jumbo 9000
```

Note that 9000 bytes seems to be the generally accepted size for jumbo frames, so that’s what I used.

After running this command, *you must reboot the switch*. The change doesn’t take effect until a reload. Fortunately, IOS reminds you of this after you enter the command. Once the switch has rebooted, you can verify the MTU setting with this command:

```
show system mtu
```

This should report that the system jumbo MTU size is 9000 bytes, confirming that the switch is ready for jumbo frames. Now we’re prepared to configure the storage system.

CONFIGURING THE STORAGE SYSTEM

Using FilerView, increasing the MTU on the appropriate network interfaces to 9000 bytes is as simple as going to Network > Manage Interfaces and then clicking the Modify link for the interface to be changed. Set the “MTU size” to 9000 (from the default of 1500), click Apply, and you’re ready to roll.

You can verify the settings in FilerView using Network > Manage Interfaces > Show All Interface Details, or by using the “ifconfig -a” command from a Data ONTAP command prompt.

CONFIGURING ESX SERVER

There is no GUI in VirtualCenter for configuring jumbo frames; all of the configuration must be done from a command line on the ESX server itself. There are two basic steps:

1. Configure the MTU on the vSwitch.

2. Create a VMkernel interface with the correct MTU.

First, we need to set the MTU for the vSwitch. This is pretty easily accomplished using `esxcfg-vswitch`:

```
esxcfg-vswitch -m 9000 vSwitch1
```

A quick run of “`esxcfg-vswitch -l`” (that’s a lowercase L) will show the vSwitch’s MTU is now 9000; in addition, “`esxcfg-nics -l`” (again, a lowercase L) will show the MTU for the NICs linked to that vSwitch are now set to 9000 as well.

Second, we need to create a VMkernel interface. This step is a bit more complicated, because we need to have a port group in place already, and that port group needs to be on the vSwitch whose MTU we set previously.

```
esxcfg-vmknic -a -i 172.16.1.1 -n 255.255.0.0 -m 9000 IPStorage
```

This creates a port group called IPStorage on vSwitch1—the vSwitch whose MTU was previously set to 9000—and then creates a VMkernel port with an MTU of 9000 on that port group. Be sure to use an IP address that is appropriate for your network when creating the VMkernel interface.

To test that everything is working so far, use the `vmkping` command:

```
vmkping -s 9000 172.16.1.200
```

Clearly, you’ll want to substitute the IP address of your storage system in that command.

That’s it! From here you should be able to easily add an NFS datastore or connect to an iSCSI LUN using jumbo frames from the ESX server.

“OFFICIAL” SUPPORT STATEMENT

Officially, jumbo frames are only supported by VMware for use by virtual machines. Technically, VMware does not support the use of jumbo frames for the software iSCSI initiator or for use with NFS datastores. At least, that’s my understanding.

So, feel free to tinker around with jumbo frames for IP-based storage, and when VMware adds official support for it in the future—I can’t imagine why they wouldn’t—then you’ll be able to hit the ground running with the configuration steps necessary to make it work.

Tags: [ESX](#), [iSCSI](#), [NetApp](#), [NFS](#), [Storage](#), [Virtualization](#), [VMware](#)

79 comments

Comments feed for this article 

Trackback link: <http://blog.scottlowe.org/2008/04/22/esx-server-ip-storage-and-jumbo-frames/trackback/>

Duncan on Tuesday, April 22, 2008 at 8:24 am

Maybe, if you can find the time, you can do some testing with iometer with jumbo frames enabled and disabled to see what the difference in performance wise... would be very interesting to know I guess.



and for Cisco, does changing the MTU size system wide affect the other attached devices in anyway?

Ecio on Tuesday, April 22, 2008 at 9:07 am

Hi Scott,
have you done also some benchmarking in order to determine if/how much throughput improvement exists on jumbo frame configuration?



thanks

Andrew Miller on Tuesday, April 22, 2008 at 9:30 am

Fantastic — I’m really hoping for official support for this soon (would love to drop the # of interrupts my servers spend processing storage traffic...not enough to get an iSCSI HBA though ;-).



Dmitry Sorokin on Tuesday, April 22, 2008 at 11:30 am

Does anyone run any performance tests against iSCSI and NFS data stores with and without jumbo frames enabled?
We run tests against NFS data store (NetApp FAS3020) and didn’t see any significant performance improvements. We tested random read/write I/O workload using iозone with 8k blocks, but didn’t test any sequential loads.



Dmitry

Trackback from [Virtual Enthusiasm](#) on Tuesday, April 22, 2008 at 9:42 pm

slowe on Wednesday, April 23, 2008 at 7:08 am

Duncan,



As far as I am aware, changing the jumbo MTU on a Cisco switch won't affect non-jumbo traffic. It merely enables the switch to handle jumbo frames. Devices that are configured to use standard sized frames will still operate without any problems.



All,

I *DO* plan on running some performance testing with and without jumbo frames enabled. However, the results of this testing will be a few weeks out, so remain patient! Thanks to everyone for reading and commenting.

Dan on [Friday, April 25, 2008 at 12:40 am](#)

We are looking into using Jumbo frames on our 6070 and Sybase/SUN5200 over NFS. We have Cisco 6509's...



Do you have any general performance numbers on using Jumbo frames on Netapp?

Sture on [Friday, April 25, 2008 at 4:24 am](#)

Hi,



We've been running Jumbo/NFS in production for nearly two months now. No issues so far 😊

Bargi on [Tuesday, April 29, 2008 at 1:31 pm](#)

Hi,



Does anyone know if this also applies to HW iSCSI HBAs also?

On the topic of performance with jumbo frames I had 2 machines setup and tried with and without, while I saw no real improvement in performance the thing that I could see was that the amount of hit on the processor of the iSCSI target dropped noticable. (can't tell you specifics but was at least 10-15% easily.

I'll be setting up a full blown iSCSI/ESX setup over the next few. If anyone thinks it's worth checking again I'll put some stats up.

Ray

Jim Heuton on [Friday, May 2, 2008 at 4:04 pm](#)

I've got to add this one for the noobs (i.e. Me) to ESX3.5 and iSCSI w/ jumbo frames. You will need to actually add the empty port group to the new vSwitch using the command below BEFORE creating the VMkernel interface (otherwise you'll get weird errors about the port group not existing). This step is mentioned in the main post, but the command to actually do it isn't.



```
esxcfg-vswitch -AvSwitch1
```

And, thanks for the info! We have jumbo frames currently running on both the VMdata and iSCSI networks (EqualLogic arrays w/ 6509s) and haven't noticed any issues yet.

Jim Heuton on [Friday, May 2, 2008 at 4:05 pm](#)

OK, the full command was chopped off. Let's try that again...



```
esxcfg-vswitch -Aport_group_name vSwitch1
```

Ernie Oporto on [Tuesday, May 13, 2008 at 12:09 pm](#)

So if jumbo frames are not supported for iSCSI or NFS, then a nice compromise would be to use normal frames for the Service Console port and its switch, and set jumbo frames for everything else. Thus, skipping the esxfg-vmknuc command?



slowe on [Tuesday, May 13, 2008 at 12:24 pm](#)

Ernie,



Yes, you could do that. That would allow you to use jumbo frames from within VMs—assuming the guest OS supports jumbo frames—and not from the Service Console or VMkernel ports.

Thanks for reading!

Felix on [Thursday, June 12, 2008 at 4:51 pm](#)

Hi



I though that the software iSCSI initiator running on the esx host did not support

jumbo frames?

If I just want jumbo frames from within my guest OS (Windows 2003 R2), I should only need to configure the MTU on the vSwitch right

(esxcli vswitch -m 9000 vSwitch1)

And no need to a vmkernel?

Thanks
Felix

slowe on Thursday, June 12, 2008 at 5:03 pm



Felix,

Jumbo frames are not supported by VMware for use by the VMkernel (NFS or iSCSI), but they do work just fine as outlined above.

Otherwise, you are correct—to configure jumbo frames for a VM you only need to set the MTU for the vSwitch (and configure the guest OS appropriately).

Thanks for reading!

Felix on Thursday, June 12, 2008 at 6:05 pm



Just need to be sure I got it right 😊

So jumbo frames actually do work from the iSCSI initiator on the ESX?

Thanks

slowe on Friday, June 13, 2008 at 12:52 am



Felix,

Jumbo frames *do* work. I'm using the ESX software iSCSI initiator with jumbo frames right now in my lab. I'm also using NFS with jumbo frames in the lab.

However, jumbo frames are *not* supported for use with iSCSI and NFS. If you run into problems with such a configuration, VMware Technical Support probably isn't going to help you much.

Does that help clarify things?

Felix on Monday, June 16, 2008 at 10:31 am



Hi

I can't get the jumbo/mtu 9000 to work with the microsoft iscsi initiator 😞 from my Windows 2003 ENT R2 SP2

It works with the ESXiSCSI initiator 😊

Thanks for a great post

Justin Grote on Sunday, June 22, 2008 at 2:22 pm



Hi Felix,

You need to use the enhanced vmxnet driver that comes with ESX 3.5 to get jumbo frame support for guest operating systems. You also get the new TSO TCP Offload feature which further reduces CPU utilization.

http://www.vmware.com/support/vi3/doc/whatsnew_esx35_vc25.html

Paul on Wednesday, August 27, 2008 at 3:14 am



Hi... I have our ESX server connected to our Dell PS5000 currently using standard frames, is there anyway to 'reconfigure' the existing vmkernel and vswitch, without having to create another ?

slowe on Wednesday, August 27, 2008 at 8:40 am



Paul,

The vSwitch can be reconfigured using esxcli vswitch, but you'll need to recreate the VMkernel NIC using esxcli vmknic.

Paul on Wednesday, August 27, 2008 at 8:11 pm



Thanks Slowe... the only problem I am now having is -

esxcli vmknic -a -i 172.16.1.1 -n 255.255.0.0 -m 9000 IPStorage tells me that -m isn't a valid switch. I'm doing this via the CLI on my pc and not the server

directly. Could this be the issue ?

slowe on Wednesday, August 27, 2008 at 9:53 pm



I take it by your comment that you are using the Remote CLI, then? I would ensure that I have the latest version of the Remote CLI or, if possible, run the command directly on the server. Obviously, this isn't possible with ESXi, but for "regular" ESX it shouldn't be an issue.

And keep in mind that this is for ESX 3.5 and higher ONLY.

Hope this helps!

Paul on Wednesday, August 27, 2008 at 10:14 pm



Yeah Remote CLI and i d/l the latest version yesterday and the same issue.. I am using 3.5.. What directory do I need to be in on the ESX server if i try and run these commands.. Logging on as Root and running them in the root directory doesnt work, cannot find command or the like... I really appreciate the help you are giving me here !

Paul on Wednesday, August 27, 2008 at 11:07 pm



Never mind i found it... YAY for the linux nuffies... thanks again for the great help...

Jason on Friday, August 29, 2008 at 9:31 am



Got it working.

What is the command to add this VMkernel port group to the new vSwitch1 with mtu 9000 and VLAN 200.

I use 6 nics to each of my esx servers. 2 are port-channel trunked to vSwitch0 nic0, nic1. 4 nics are port-channel trunked to vSwitch1 which is used for my storage network.

I got everything working by deleting my current vmkernel port group. then using the command to add a empty port group. then using scotts command to add the new IPStorage port channel with mtu 9000. then i just used the GUI in VC to make the port group vian 200.

Thanks

Jason

Steve on Monday, September 1, 2008 at 8:13 pm



Per the instructions above, I configured jumbo frames on my celeros san and esx 3.5 update 2 using a pair of 10g chelsio nic, one in my san and one in esx.

WOW..

Without Jumbos enabled, my throughput was about 200 MB/s. I was questioning if I should continue using 10G. After enabling Jumbo frames, my throughput went to over 400 MB/s writes and 500 MB/s reads..

I can't believe the difference. I wish VMWare would officially support it.

Thanks

Steve,

sjrieks@sy.edu

Thoma on Wednesday, September 10, 2008 at 1:31 pm



Hi,

I try to run this Switch command on my ESX 3.5 update2 on HP BL25p servers.

```
esxcfg-vswitch -m 9000
```

gives the error:

Unable to set MTU to 9000 the following uplinks refused the MTU settings: vmnic3

Any suggestion of what can be wrong??

The other commands works fine!

//Thomas

Thoma on Wednesday, September 10, 2008 at 1:33 pm



Sorry correct command I run is

```
esxcfg-vswitch -m 9000 vSwitch3
```

still the same error:

Unable to set MTU to 9000 the following uplinks refused the MTU settings: vmnic3

slowe on Wednesday, September 10, 2008 at 3:43 pm



That error indicates that vmnic3 doesn't support jumbo frames. You'll need to be sure that vmnic3 is running at 1Gbps. If it is and still refuses to accept the MTU setting, there isn't a whole lot you can do about it.



Good luck!

Thomas on Friday, September 12, 2008 at 4:15 am



Thanks slowe!

Seems that it's not much to do now 😞

I am using the NetXtreme BCM5703 Gigabit Ethernet in full 1Gb mode. I can see on the Internet more than me have the same problem. The card is supporting Jumbo Frames in Windows but not with ESX. I guess the limitations is in the network driver in ESX.

//Thomas

Mike on Tuesday, September 16, 2008 at 5:43 pm



Hi Scott,

This info is valuable!
I am running into this:

When running vmkping -s 9000, the amount sent is actually 9008, and doesn't pass through.

I was able to push through a command of vmkping -s 8954 (sending 8962 bytes) and that was successful.

And idea why that may be?

I am using a Cisco switch (mtu 9000) and a Dell (also set to mtu 9000), and my 3par machine is definitely set to mtu 9000.

Pierre on Tuesday, October 7, 2008 at 1:08 pm



Thanks for this blog!

After enabling the JumboFrame support like you said, the command vmkping -s 9000 10.0.10.1 do not return value. But I can't ping at 8991 because ping seem to add a 8 bytes overhead. It is normal?

```
PING 10.0.10.1 (10.0.10.1): 8991 data bytes
8999 bytes from 10.0.10.1: icmp_seq=0 ttl=64 time=1.640 ms
8999 bytes from 10.0.10.1: icmp_seq=1 ttl=64 time=1.363 ms
8999 bytes from 10.0.10.1: icmp_seq=2 ttl=64 time=1.362 ms
```

Thanks

orb on Saturday, October 18, 2008 at 1:36 pm



@Mike

I do have also the same problem with PowerConnect 6248 switches. vSwitch and vmnic return well a MTU of 9000. The Equallogic storage device has its mtu set to 9000 and log returns a successful initiator connection. I get all volume presented to my ESX machines but I am unable to create a storage. It times out, ESX complains it cannot find a valid path.

Maikel on Sunday, November 9, 2008 at 8:16 am



Thanks for the great info, I made the adjustments to my vSwitch and the ping works.

Is there any other action needed on the Virtual Machines, to actually use these settings, or are they forced by the vSwitch?

slowe on Sunday, November 9, 2008 at 8:46 pm



Maikel,

Understand that this article is about getting VMware ESX to use jumbo frames for the VM storage, not about using jumbo frames from within the VMs. That's handled separately.

To get VMs to use jumbo frames, you configure the vSwitch(es) and then configure the guest OS appropriately.

I hope that makes sense to you. Thanks!

Daniel on Tuesday, November 11, 2008 at 1:58 pm



Is this supposed to work on ESXi too? I was able to change the MTU on the

vSwitch, but calling `esxcfg-nics -l` shows that the two NICs assigned to that vSwitch are still using a value of 1500. In addition to that, running `esxcfg-vmknick` with the `-m` flag returns the error "Unknown option: m".



Has anybody tried this on ESXi? We have a normal ESX 3.5 license, but I don't want to install and activate the full-blown until we've tested it carefully, so we were testing ESXi on a test server (same hardware as the production one, though).

Thanks,

Daniel

JShire on [Tuesday, November 11, 2008 at 5:37 pm](#)



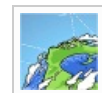
Hey Scott,

Very useful post!

Unfortunately I'm running into the same problem as Pierre and Mike. When MTU is set to 9000 on the vmknick, you can only transmit 8991.

Any suggestions would be greatly appreciated.

slowe on [Wednesday, November 12, 2008 at 12:26 am](#)



Daniel,

It may be the case that the NICs do not support jumbo frames.

Jshire,

I haven't seen that behavior. What symptoms are you experiencing?

Daniel on [Wednesday, November 12, 2008 at 6:42 am](#)



Scott,

Thanks for the reply. According to Intel, the NICs I'm using should support jumbo frames; the machine has an Intel quad-port PCIe GbE controller (and two ports integrated in the motherboard, but I wasn't using those). For some strange reason, VMware is using the e1000 driver for the NICs on the PCIe card (82571EB) and the igb driver for the ones on the motherboard (82575EB).

At any rate, even if I did get ESXi to change the MTU on the NICs, I would still have the problem of not being able to set the MTU on the VMkernel port.

I guess I'll try using the integrated ports for the iSCSI connection to see if I can change the MTU on those, and if I still can't change it for the VMkernel port I'll try using ESX instead of ESXi.

Daniel on [Wednesday, November 12, 2008 at 7:01 am](#)



OK, this gets weirder by the minute: the NICs I'm using for iSCSI connectivity are on vSwitch1, which I created specifically for that purpose according to the ESXi docs. Trying to change the MTU for that vSwitch doesn't affect the NICs, but if I change it for vSwitch0 (currently being used for the management network and the RCL VM), the change is applied to all of the NICs on the server, including the ones which are not assigned to any vSwitch!

Wlad on [Wednesday, November 26, 2008 at 11:39 am](#)



I have the same Problem with 8991 ,

if i ping with 8992 (+8) it have to work but no answer , only with 8991

Frank on [Thursday, December 4, 2008 at 12:57 pm](#)



Hi Scott,

Have you had the chance to run any performance tests on jumbo vs. non jumbo frames? I am specifically interested in NFS improvements.

Thanks

umer on [Wednesday, December 10, 2008 at 1:42 pm](#)



Hi Gentlemen,

This is really driving me insane..

My Setup: MD3000i, 2x PowerConnect 6xxx series
Jumbo Frames with 9000 MTU enabled end to end on all devices

I am testing for Jumbo frames, and here are the steps i take

1. Add a vSwitch :
esxcfg-vswitch -a vSwitch3

2. Add PortGroup
esxcfg-vswitch -A "IPStorage" vSwitch3

3. Add VMKernel Int.
esxcfg-vmknick -a -i 192.168.2.24 -n 255.255.255.0 -m 9000 IPStorage

4. Attach vmnics
esxcfg-vswitch -L vmnic7 vSwitch3
esxcfg-vswitch -L vmnic8 vSwitch3

NOW!!

I can vmkping " 192.168.2.1" the IP of my switch stack
PING 192.168.2.1 (192.168.2.1): 9000 data bytes
9008 bytes from 192.168.2.1: icmp_seq=0 ttl=64 time=3.918 ms
9008 bytes from 192.168.2.1: icmp_seq=1 ttl=64 time=4.036 ms
9008 bytes from 192.168.2.1: icmp_seq=2 ttl=64 time=3.868 ms

esxcfg-vswitch output
Switch Name Num Ports Used Ports Configured Ports MTU Uplinks
vSwitch3 64 4 64 9000 vmnic7
PortGroup Name VLAN ID Used Ports Uplinks
IPStorage 0 1 vmnic7

HOWEVER —

I just CANNOT get a ping response from my storage IP (MD3000i)
I get a 100% packet loss.

Jumbo frames are INDEED enabled on the MD3000i.

Please help me as I am losing my mind over this.

Any help is MUCH appreciated in advance.
Thanks.

Jim Nitterauer on Thursday, December 11, 2008 at 11:29 am



1. First try pinging the IP of a switch stack on another host to be sure your switch fabric is allowing the jumbo frames across the storage network. If not, then the issue is with the connected switch.
 2. Next try the vmkping -s 9000 command
 3. If that fails, lower the packet size to say 8000 (-s 8000) and see if that works. If it does, you know that jumbo frames are getting passed over the storage lan to the storage IP. The issue is now a setting for the maximum size of the jumbo frames allowed. You must account for the VLAN tagging.
 4. Go back to the storage device and increase the MTU to 9014
 5. Now try your pings using vmkping -s 9000
 6. You can sequentially lower the -s value until you get success. Once you do, add the spread to the MTU on the storage server's interface.
- Ex. if you get success with -s 8982 but not with -s 8983, this means that your MTU or max packet size on the receiving end is too low by 18. Increase it and ping again.

Good luck!

umer on Thursday, December 11, 2008 at 4:06 pm



Thanks Jim.

The largest ping size i can get a response from my storage nic is
8954

— 192.168.2.10 ping statistics —
3 packets transmitted, 0 packets received, 100% packet loss
[root@SRVXXX-root]# vmkping -s 8954 192.168.2.10
PING 192.168.2.10 (192.168.2.10): 8954 data bytes
8962 bytes from 192.168.2.10: icmp_seq=0 ttl=64 time=1.421 ms
8962 bytes from 192.168.2.10: icmp_seq=1 ttl=64 time=1.398 ms
8962 bytes from 192.168.2.10: icmp_seq=2 ttl=64 time=1.385 ms

The storage software doesnt allow any value above 9000.

any further pointers?

Jim Nitterauer on Thursday, December 11, 2008 at 10:23 pm



Do you have more than 1 ESX host on your storage network?

Jim Nitterauer on Thursday, December 11, 2008 at 10:31 pm



Also, have you done this?

```
esxcfg-vswitch -m 9000 vSwitch3
```

This sets the vSwitch to handle jumbo frames. I did not see this in your steps.

umer on Friday, December 12, 2008 at 11:48 am



Hello Jim,

Yes, I have Three ESX that are sharing this SAN at this time.

The MTU is set to 9000, but i forgot to mention that step

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch3	64	6	64	9000	vmnic8,vmnic7

Thanks again!! Please let me know if you can think of anything else.

David F. on Friday, December 12, 2008 at 11:59 am



Since I'm in the middle of deploying a set of Equallogic SANS for our vmware environment, I'm extremely interested in this topic.

I don't understand VMWare's logic on this, specifically given that the command to enable jumbo frames for a vmkernel doesn't seem to have any way to distinguish between iSCSI traffic and non-iSCSI traffic. So, if you turn it on for the vmkernel, it's on for iSCSI/NFS traffic also if I understand this correctly. I don't see how it can be on for some parts of vmkernel traffic, and not for others.

For now though, until jumbo frames are officially support, I'm going to skip enabling them on our network.

Dave

1 · 2 · [Next >](#)

Reply

Name (required)

Email (required)

Website

Your comment

[SITE ADMIN](#)

[Log in](#)

[Entries RSS](#)

[Comments RSS](#)

[WordPress.org](#)



Scott's Weblog by [Scott Lowe](#) is licensed under a [Creative Commons Attribution-ShareAlike 3.0 United States License](#).

 [Subscribe to feed](#)

Powered by [WordPress](#) and [Tarski](#)