SPECsfs2014 within Emerald

Don Capps

SNIA Emerald™ Training

SNIA Emerald™ Power Efficiency Measurement Specification

Version 3.0

February-March 2018

What is *not* in this presentation

- How to setup AC power
- How to setup the Power meter
- How to install and setup Windows or Unix systems.
- How to setup and properly configure DNS (Windows or Unix)
- How to setup the switches and NICs in your lab.
- How to install, and properly configure a Windows Active directory domain.
- How to properly setup a lab to handle high stress workloads.

Introduction to SPEC SFS2014

- Video: Application level benchmarking with SPECsfs2014 https://youtu.be/4wfeM1q0zHA
- Video: Using SPEC SFS with the SNIA Emerald Program for EPA Energy Star Data Center Storage Program https://youtu.be/7gDgcDYatvM

Setting up the environment (toc)

- Block diagram of the benchmark architecture
- What is sFlow®
- Setting up sFlow on switches
- Setting up sFlow collection
- Installing client operating systems
- Configuring client nodes
  - Installing Python (if needed)
  - Installing SPECsfs2014 software
  - Disable firewalls (Iptables and selinux)
  - Setting up ssh with keys, not passwords.
  - Client OS tunes
  - Two NICs is good idea. One for client access/control and the other for data movement.
  - How much RAM for the workloads.

Setting up the environment (toc)

- Ensure that DNS is properly configured
- Setting up storage server
  - Tuning storage server
  - Balancing load across spindles and NICs
  - Allocating enough space for the workloads.
- Configuring SPECsfs2014
  - CLIENT_MOUNTPOINTS=
  - BENCHMARK=
  - LOAD=
  - NUM_RUNS =
  - USER=
  - PASSWORD=
  - INCR_LOAD=
  - WARMUP_TIME=

Setting up the environment (toc)

- Running SPECsfs2014
  - Starting it running
  - Monitoring its progress
  - Examining the results
- Finding the optimal peak performance for each workload

Setting up the environment (toc)

- Merging the sFlow data with the power meter results
- Tag2014
- Filling out the TDR

If things go wrong (toc)

- Examining the SFS2014 logs
- Client process logs
- Summary logs
- Client performance logs
- Console output

Diagram of the benchmark architecture

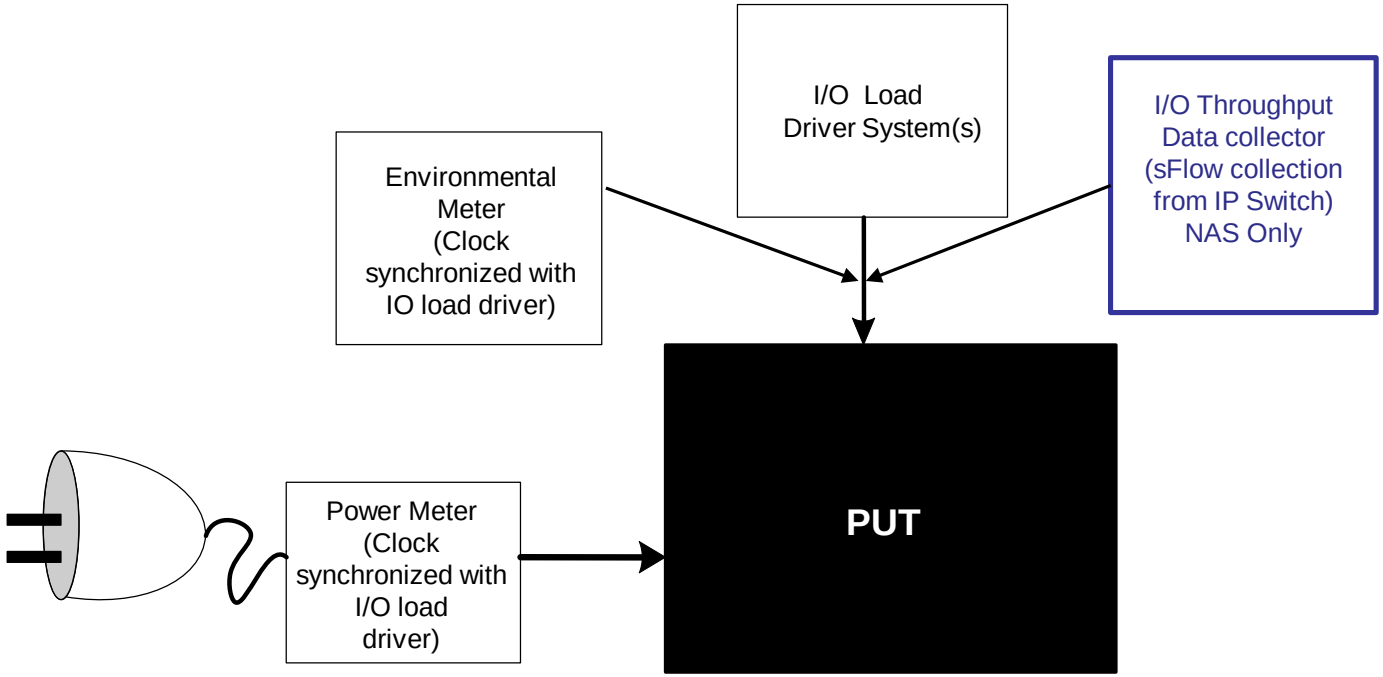

What is sFlow

- sFlow is the leading, multi-vendor, standard for monitoring high-speed switched and routed networks.
- sFlow technology is built into network equipment and gives complete visibility into network activity, enabling effective monitoring, management and control of network resources.
- sFlow is available from most leading network equipment vendors, including: Alcatel-Lucent, Allied Telesis, Arista, Brocade, Cisco, Dell, D-Link, Enterasys, Extreme, Fortinet, Hewlett-Packard, Hitachi, Huawei, IBM, Juniper Networks, NEC and ZTE. For a complete list of products supporting sFlow, see sFlow.org.

Setting up sFlow on the switch

Setting up sFlow on the switch

Quanta LB6M 10Gbe switch: Login to prompt

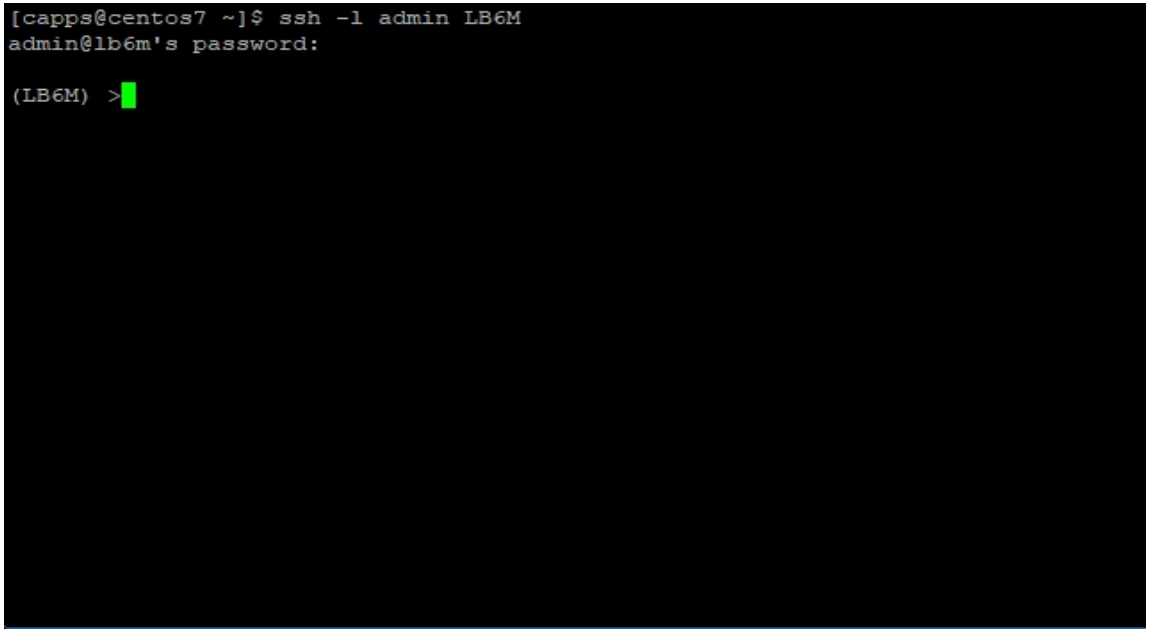

Setting up sFlow on the switch

Elevate privilege level

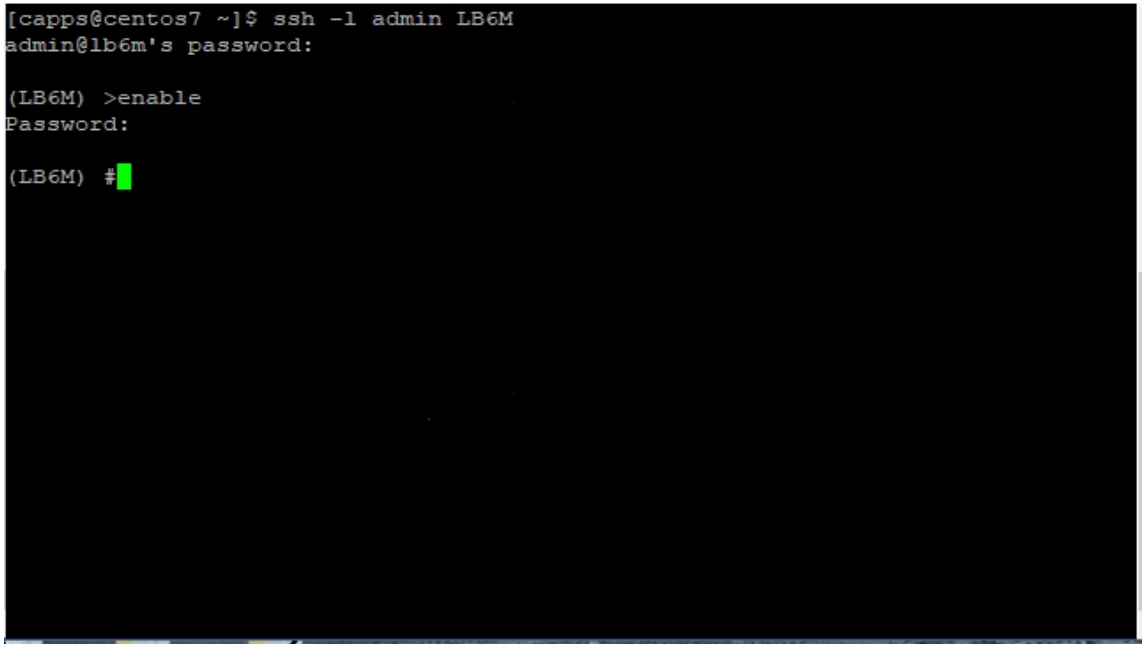

Setting up Sflow on the switch

Switch to configuration menu

```
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #configure

(LB6M) (Config)#
```

Setting up sFlow on the switch

Set the IP address for the collector

```
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #
(LB6M) #configure

(LB6M) (Config)#sflow receiver 1 ?

ip                          Configure IP Address of the Receiver.
maxdatagram                 Configure Maximum Datagram Size.
owner                       Configure Owner String of the Receiver.
port                        Configure Receiver Port.

(LB6M) (Config)#sflow receiver 1 ip ?

<ip>                        Enter Ipv4/ipv6 Address.

(LB6M) (Config)#sflow receiver 1 ip 10.0.0.231 ?

<cr>                        Press enter to execute the command.

(LB6M) (Config)#sflow receiver 1 ip 10.0.0.231
```

Setting up sFlow on the switch

Select ports to monitor on the switch

```
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#
(LB6M)  (Config)#interface 0/1-0/28

(LB6M)  (Interface 0/1-0/28)#
```

Setting up sFlow on the switch

Set the polling index to use

```
voice                         Configure Voice VLAN Parameters.

(LB6M)  (Interface 0/1-0/28)#port ?

lacpmode                      Enable/Disable the port's LACP mode.
lacptimeout                   Configure the port's LACP timeout.

(LB6M)  (Interface 0/1-0/28)#sflow ?

poller                        Configure poller options on interface.
sampler                       Configure sampler options on interface.

(LB6M)  (Interface 0/1-0/28)#sflow poller ?

<index>                       Enter Receiver Index <1-8>.
interval                      Configure poll interval.

(LB6M)  (Interface 0/1-0/28)#sflow poller 1 ?

<cr>                          Press enter to execute the command.

(LB6M)  (Interface 0/1-0/28)#sflow poller 1

(LB6M)  (Interface 0/1-0/28)#
```

Setting up sFlow on the switch

Set the polling rate to 5 seconds

```
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#
(LB6M) (Interface 0/1-0/28)#sflow poller interval 5 ?
```

Setting up sFlow collection tools

SNIA Emerald™ Training

SNIA Emerald™ Power Efficiency Measurement Specification

Version 3.0

February-March 2018

Setting up sFlow collection

- Download sflowtool from inmon.com
- Download sflowTrend from inmon.com

Download sflowtool

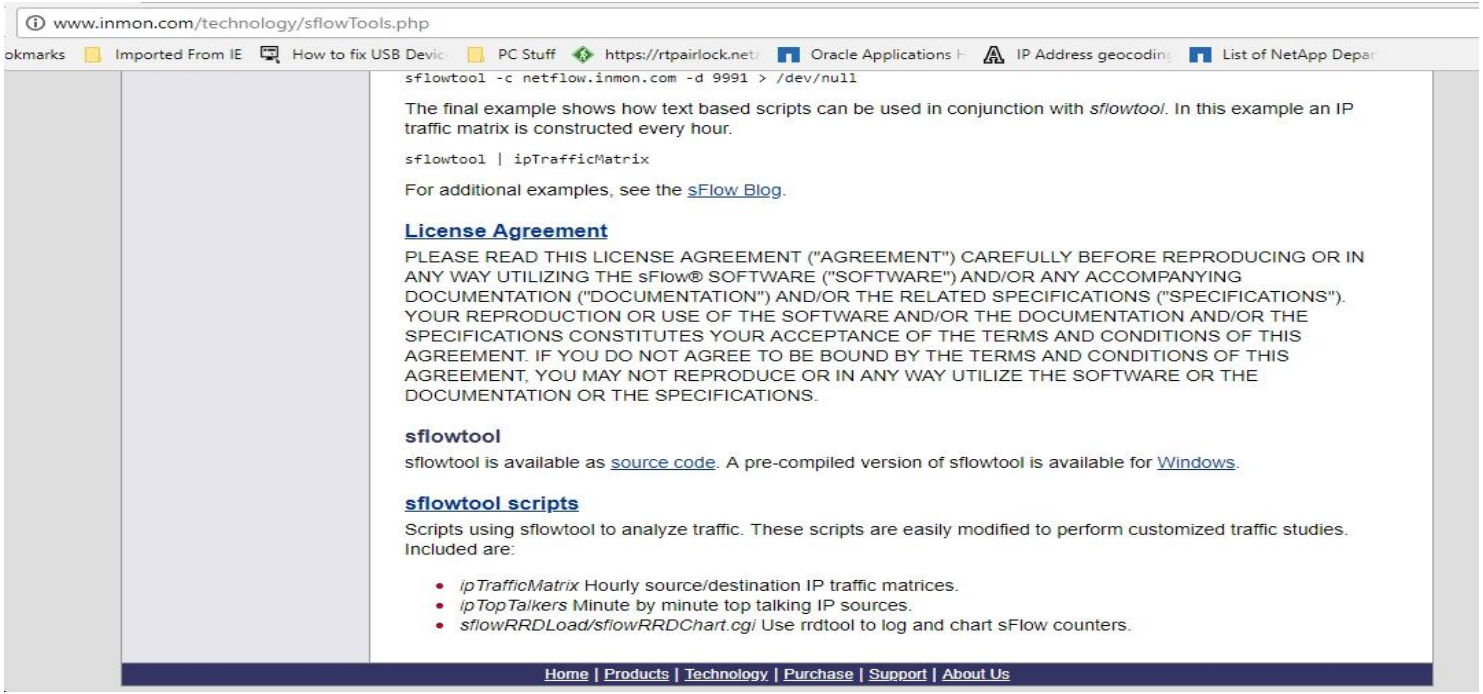

Running sflowtool

- sflowtool -4 –L > c:\tmp\sflowdata.txt

Running sflowtool

Data that is in the sflowtool file

SwitchIP	Port	InBytes	OutBytes
CNTR,10.0.0.248,2018-03-04 19:10:31,24,6,4294967295,1,3,153494289,570659,2838,1467,0,0,0,1122830633,919463,169,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:10:36,24,6,4294967295,1,3,164427808,650122,2860,1472,0,0,0,1296109924,1056224,170,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:10:41,24,6,4294967295,1,3,172979004,712033,2893,1478,0,0,0,1430569898,1162449,170,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:10:46,24,6,4294967295,1,3,184035482,792030,2898,1481,0,0,0,1604336944,1299946,170,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:10:51,24,6,4294967295,1,3,192577062,854056,2900,1484,0,0,0,1739303201,1406648,171,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:10:56,24,6,4294967295,1,3,203444482,933274,2906,1487,0,0,0,1912086457,1543110,171,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:01,24,6,4294967295,1,3,212235634,997451,2908,1489,0,0,0,2052267079,1653613,171,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:06,24,6,4294967295,1,3,220564553,1057904,2914,1490,0,0,0,2183741866,1757467,172,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:12,24,6,4294967295,1,3,231717183,1138865,2942,1493,0,0,0,2359889118,1896604,172,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:17,24,6,4294967295,1,3,240374506,1201242,2960,1495,0,0,0,2494988054,2003832,172,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:22,24,6,4294967295,1,3,251079469,1278233,2965,1498,0,0,0,2661588115,2136189,173,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:27,24,6,4294967295,1,3,259898164,1341837,2968,1499,0,0,0,2799269345,2245425,173,228,0,0,1			
CNTR,10.0.0.248,2018-03-04 19:11:32,24,6,4294967295,1,3,270594864,1419493,2970,1502,0,0,0,2968310559,2379005,173,228,0,0,1			

Download sflowtrend (optional)

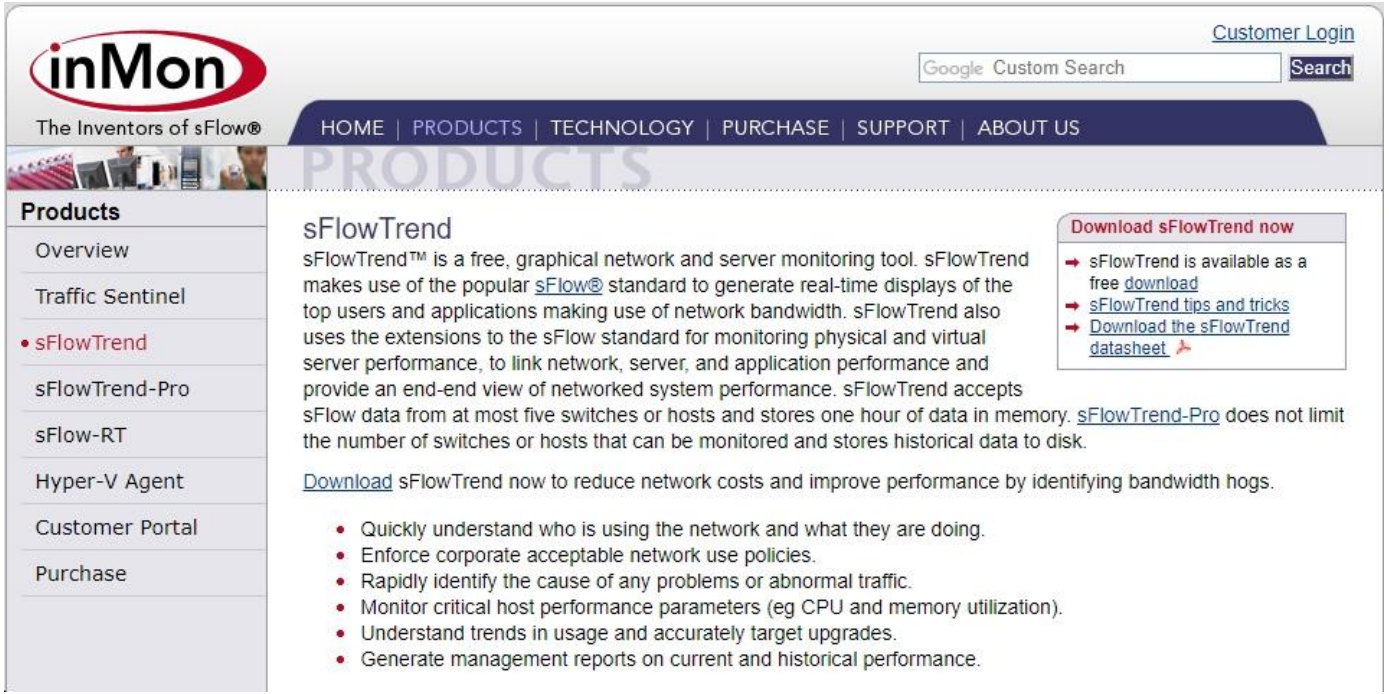

Running sflowtrend (It is nice to see the big picture)

The main dashboard view

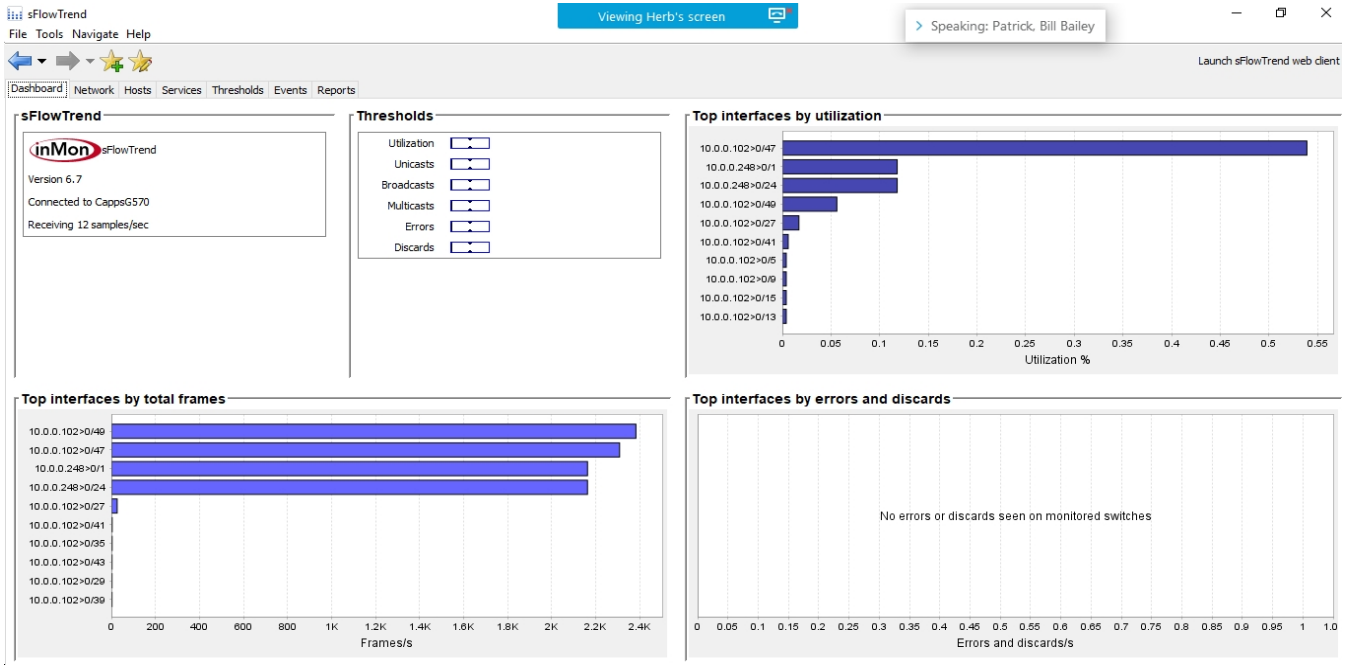

Running sflowtrend (handy reports)

One of the reports that can be generated

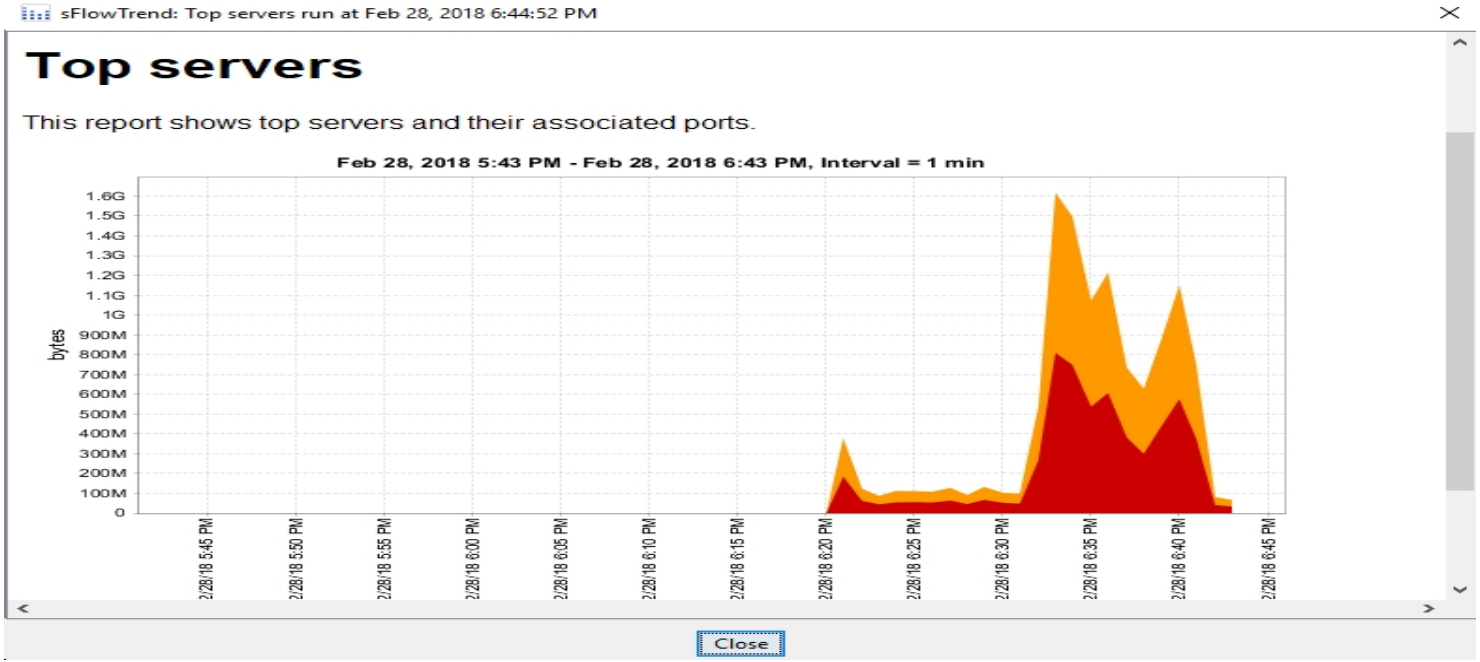

Running sflowtrend (handy report)

Port utilization view

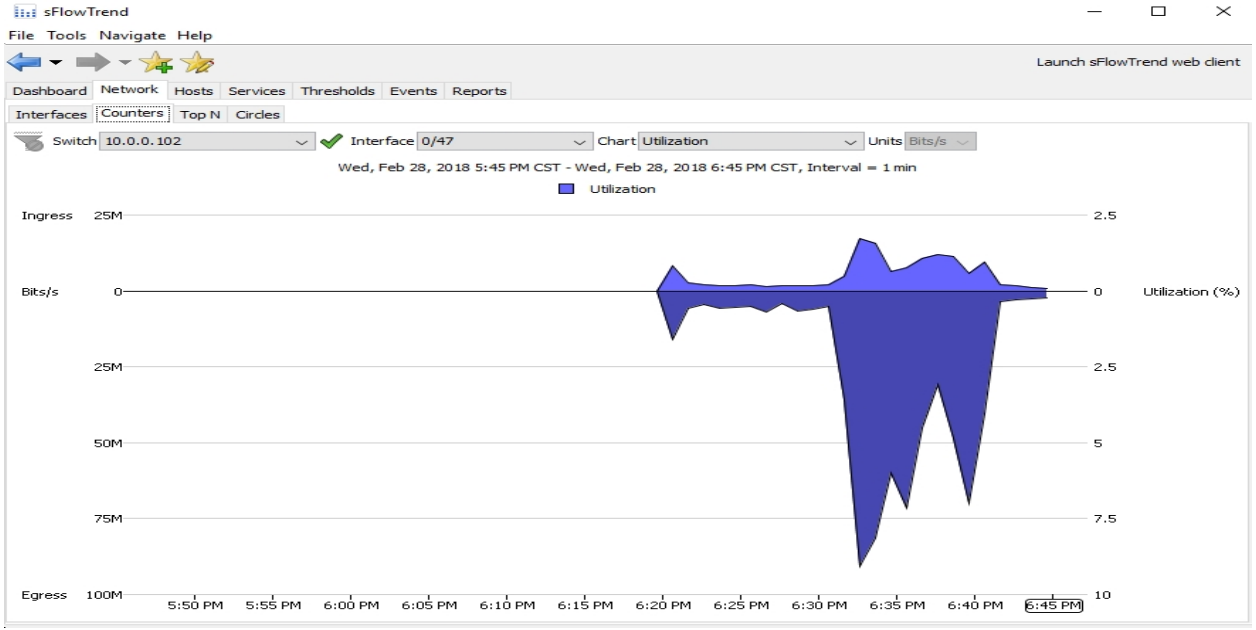

Installing the Client OS.

Installing client operating system

Boot install screen

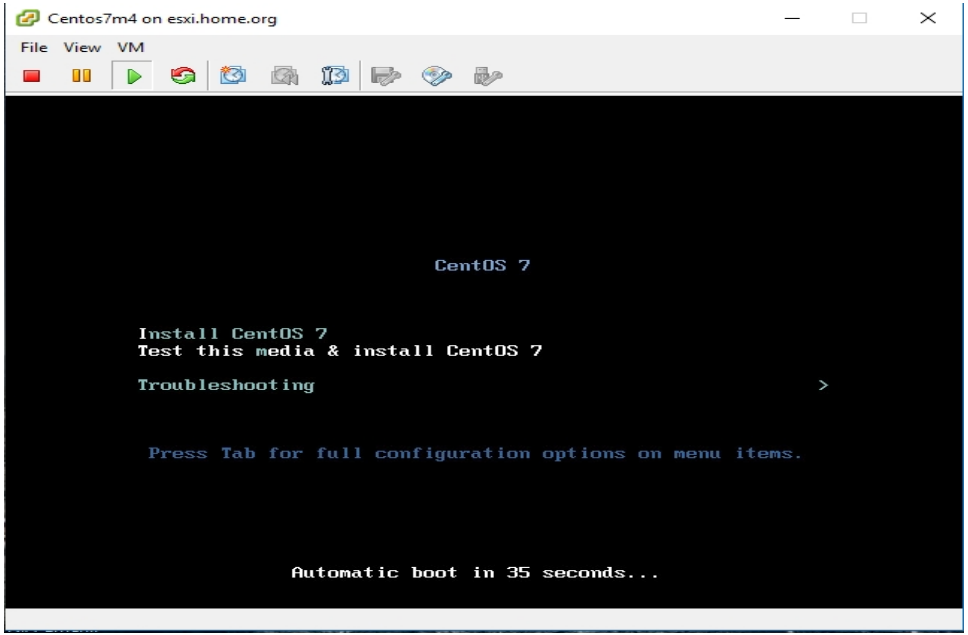

Installing client operating system

Select preferred language

Installing client operating system

Select software: NFS client, Development tools

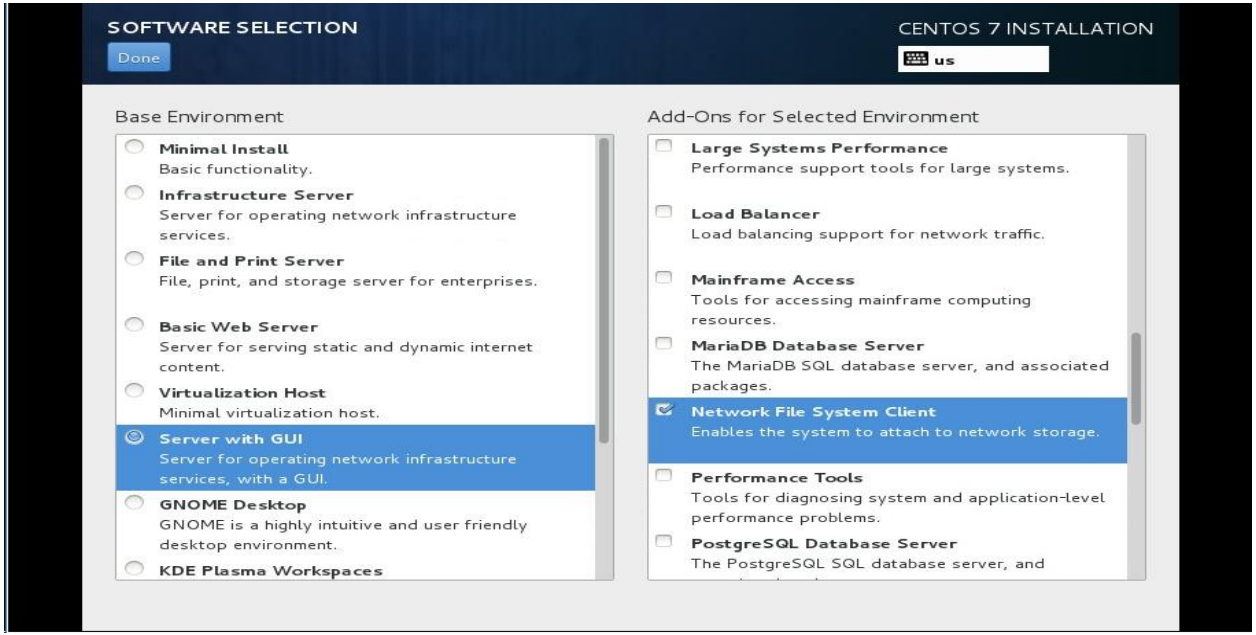

Installing client operating system

Setup the network interface

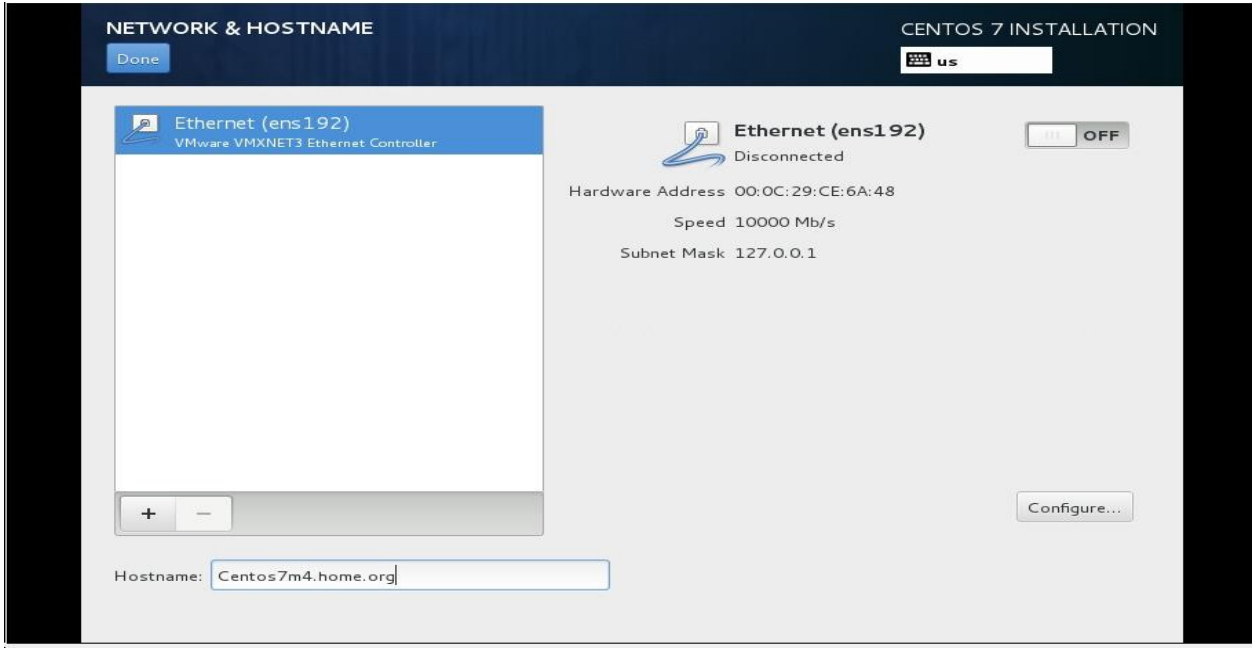

Installing client operating system

Select where to install the system

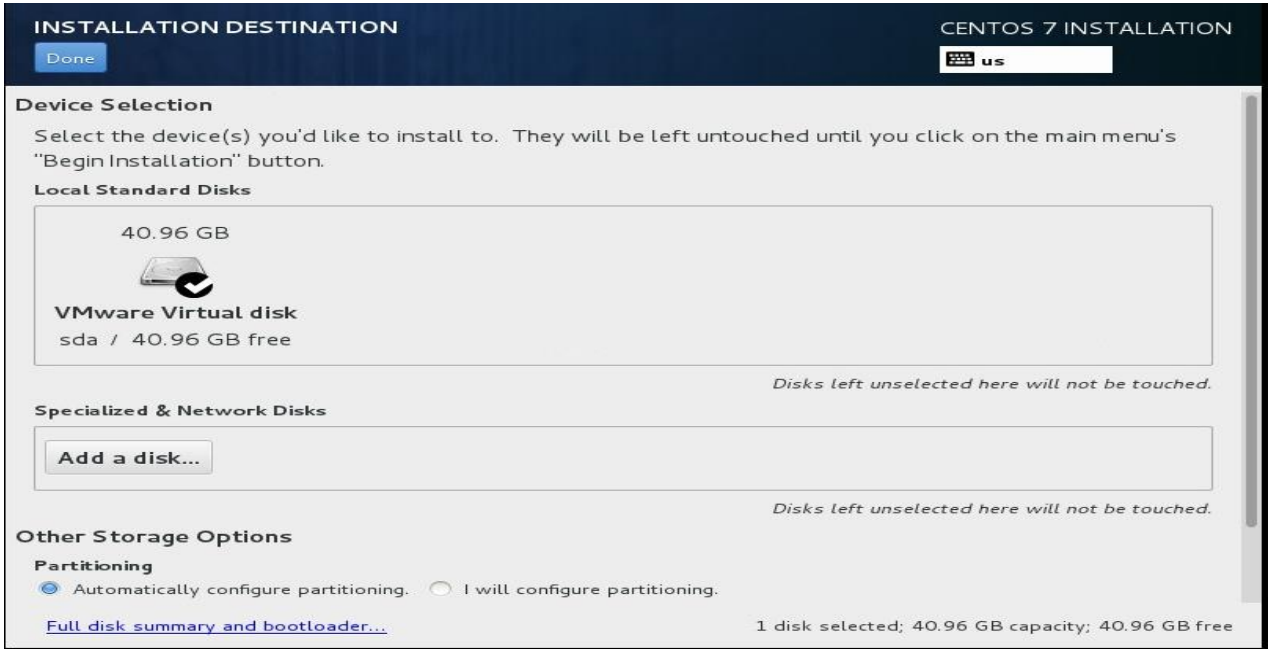

Installing client operating system

Setup root and user names and passwords

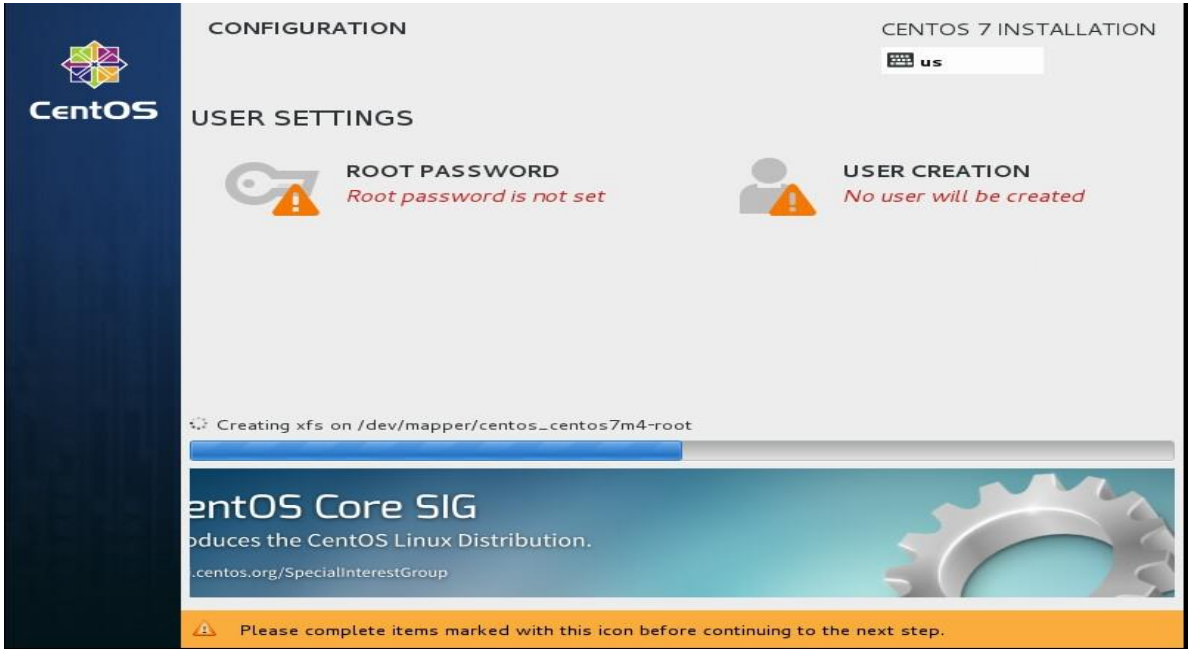

Installing client operating system

Login to start client configuration

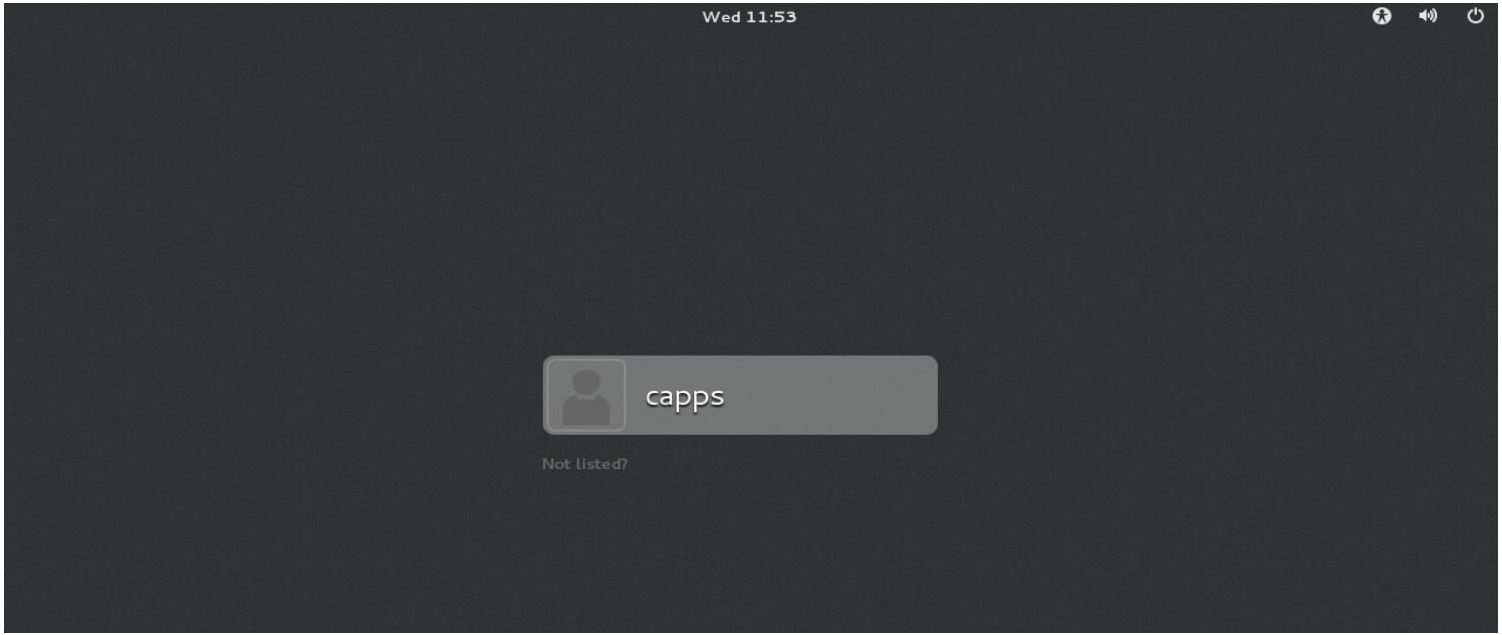

Installing client operating system (ssh)

Setup ssh keys

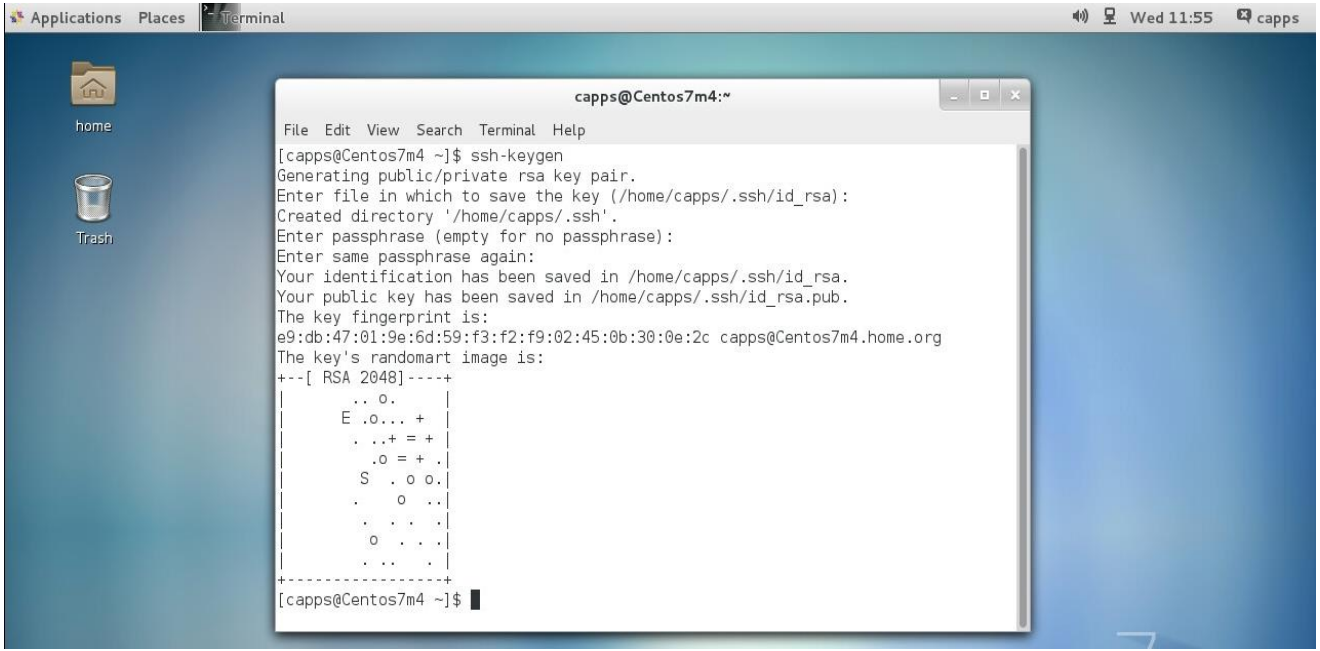

Installing client operating system (ssh)

Test ssh login to self (and all other client nodes after distributing the ssh keys)

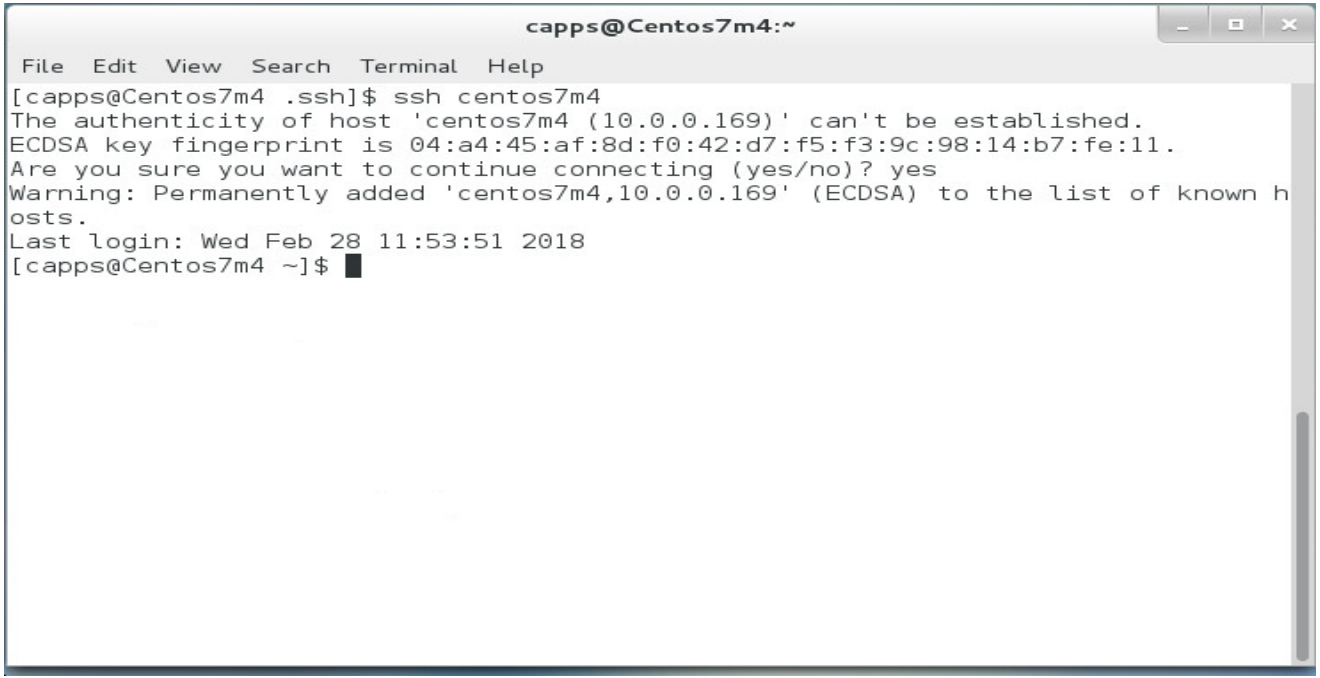

Installing required software

Installing Python 2.6 not 3.x (If needed)

Go get Python, if you need it. Not needed in this example

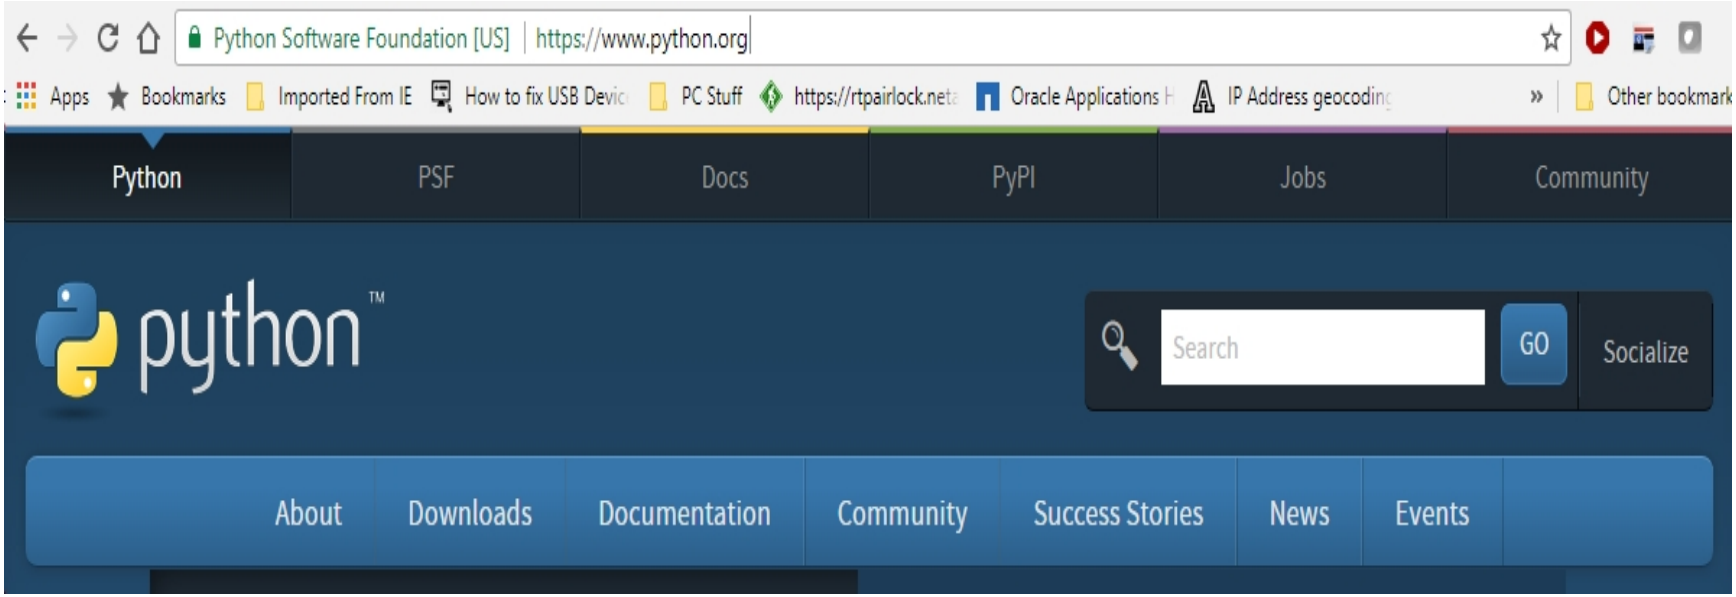

Installing SPECsfs2014
Do this on all client nodes

Installing SPECsfs2014

Mount the DVD, or filesystem that contains the SFS2014 kit

Installing SPECsfs2014

Copy to location where one wishes to run

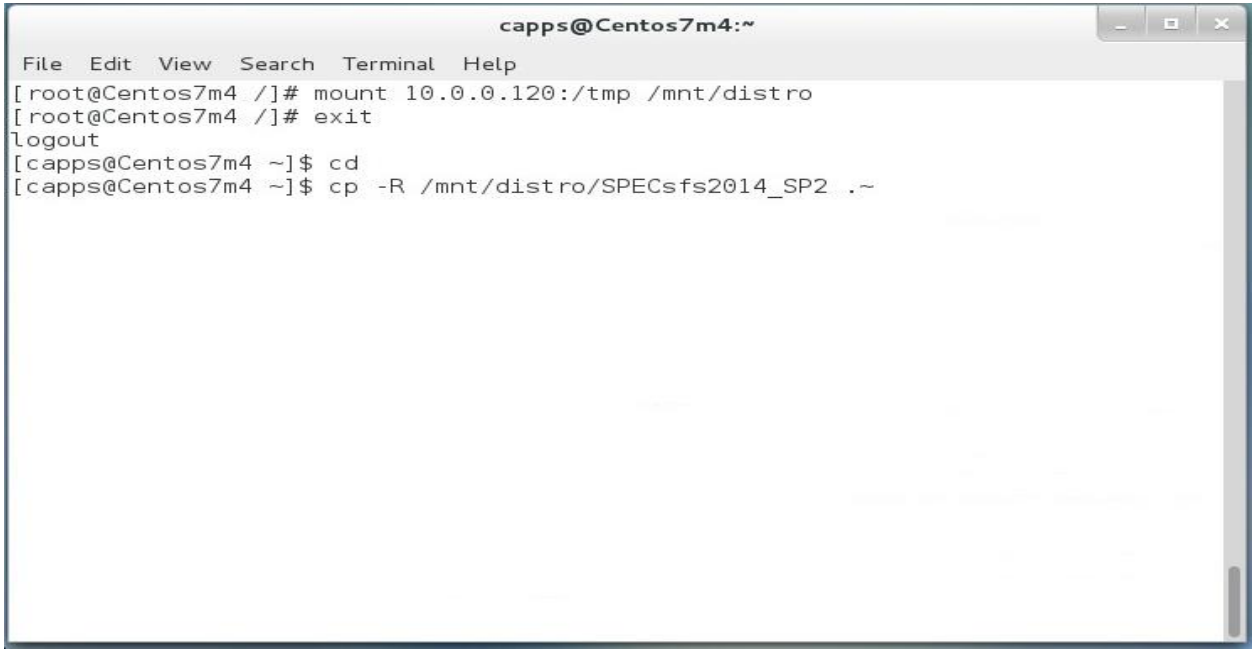

Installing SPECsfs2014

Validate the SPECsfs2014_SP2 directory exists

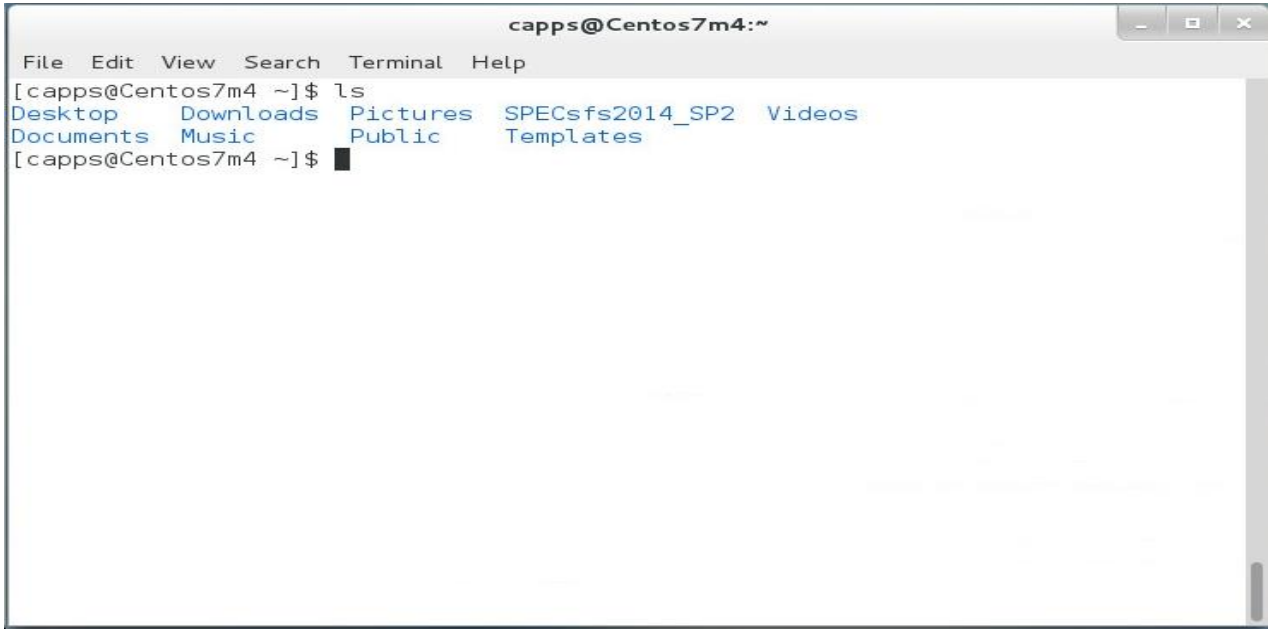

Installing SPECsfs2014

Validate the contents of the kit look like this

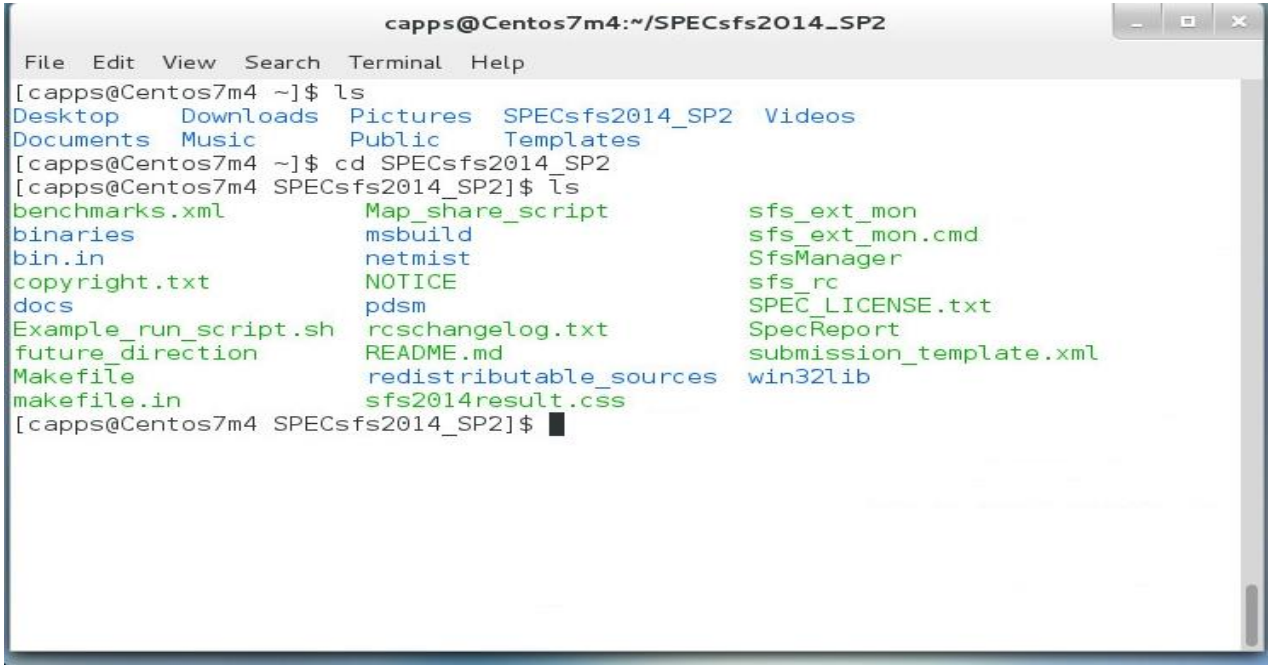

Configuring the client

Client configuration (ssh keys)

The script from the SFS2014 User's guide that sets up the ssh keys

See SFS2014 User's guide appendix "B"

Appendix B – Setting up password-less SSH

Here is a sample script that can be used to set up password-less SSH on Linux clients.

```
# Define the hosts to be involved in the trust here
# DO NOT include the host you are running, it is added by default

hosts="s2 s3 s4 s5 s6"

echo ""
echo ""
echo "This script will generate SSH keys for the specified machines,"
echo " and set up password-less authentication between them."
echo " You will be prompted for passwords several times during this process."
echo ""
```

Client configuration (ssh keys)

Cont appendix “B”

```
# Get current user
user=`who -m | awk {'print $1'}`
echo "Trust will be configured for user $user"
echo ""
echo "If this is not correct, stop and login as the appropriate user"
echo -n "(RETURN to continue, CTRL-C to exit) "

read continue

# Configure keys on current host
cd $HOME
ssh-keygen -t rsa
cat .ssh/id_rsa.pub >> .ssh/authorized_keys
chmod 700 .ssh
chmod 600 .ssh/*
```

Client configuration (ssh keys)

Cont appendix “B”

```
for host in $hosts
do
  ssh $user@$host 'ssh-keygen -t rsa'
  ssh $user@$host 'cat .ssh/id_rsa.pub' | cat - >> ~/.ssh/authorized_keys
done

for host in $hosts
do
  scp .ssh/authorized_keys $host:.ssh
  ssh $user@$host 'chmod 700 .ssh ; chmod 600 .ssh/*'
done

exit
```

Configuring client operating system

SNIA GSI | GREEN STORAGE

Apply the recommended client tunes from the SFS2014 User's guide

- Apply the tunes:

```
if  your NIC type == 10 GigE
  echo "300000" > /proc/sys/net/core/netdev_max_backlog
fi
echo "131071" > /proc/sys/net/core/rmem_default
echo "131071" > /proc/sys/net/core/rmem_max
echo "131071" > /proc/sys/net/core/wmem_default
echo "131071" > /proc/sys/net/core/wmem_max
echo "4096 87380 8388608" > /proc/sys/net/ipv4/tcp_rmem
echo "4096 87380 8388608" > /proc/sys/net/ipv4/tcp_wmem
echo "128"    > /proc/sys/sunrpc/tcp_slot_table_entries
echo "65536"  > /proc/sys/net/core/somaxconn
echo "5"      > /proc/sys/net/ipv4/tcp_fin_timeout
```

Listing firewalls

Check to see if you have any firewalls running

Disabling firewalls

Disable any/all firewalls on every client

```
root@Centos7m4:~
File  Edit  View  Search  Terminal  Help
[root@Centos7m4 ~]# systemctl disable firewalld.service
rm '/etc/systemd/system/dbus-org.fedoraproject.FirewallD1.service'
rm '/etc/systemd/system/basic.target.wants/firewalld.service'
[root@Centos7m4 ~]#
```

Checking the amount of RAM

Check to make sure the clients have enough RAM for the workloads

- Client memory requirements per business metric:

DATABASE = 55 Mbytes per LOAD increment
SWBUILD = 400 Mbytes per LOAD increment
VDA = 10 Mbytes per LOAD increment
VDI = 8 Mbytes per LOAD increment

Checking DNS works

Ensure DNS works !!!!!

This is the most common problem that folks have. You MUST have a working DNS

- Tests:
  - Ssh from the prime to every client node.
    Ensure the password-less ssh works.
  - nslookup every client and ensure that its name gave you an IP that will nslookup back to that name.

Setting up the Storage server

SNIA Emerald™ Training

SNIA Emerald™ Power Efficiency Measurement Specification

Version 3.0

February-March 2018

Setting up the storage server

Ensure that the storage server has sufficient space, and that the accesses are balanced across all of the filesystems and NICs

Setting up storage server

- Tuning storage server (Best practices from the vendor)
- Balancing load across spindles and NICs (Use all data paths) By using the CLIENT_MOUNTPOINTS and the mounts done by each client.
- Configure enough space for the workloads.(Reminder, it's a bummer to run out)

Configuring SPECsfs2014

SNIA Emerald™ Training

SNIA Emerald™ Power Efficiency Measurement Specification

Version 3.0

February-March 2018

Configuring SPECsfs2014

- Make 4 copies of the sfs_rc/configuration file, from here:
  https://www.snia.org/sites/default/files/emerald/download/Spec_v3_0/SNIA_Emerald_FileAccessIOLoadDriverConfigurationFile_Version_2017_08_30.txt
  and name them:
  - sfs_rc_database
  - sfs_rc_vdi
  - sfs_rc_vda
  - sfs_rc_swbuild

Configuring SPECsfs2014

Sfs_rc variables that you need to set in each sfs_rc file

Configuring SPECsfs2014

- CLIENT_MOUNTPOINTS= clientname:/mountpoint clientname:/mountpoint or clientname:\\servername\sharename for Windows.
- USER= The valid user name to use for login and running the tests.
  Unix: Username
  Windows: Domain\Username
- BENCHMARK= SWBUILD | VDA | VDI | DATABASE **(have this match the sfs_rc_name)**
- PASSWORD= (only needed for Windows)
- LOAD= Load value. from1 to big. Where big does not produce INVALID results.
- INCR_LOAD= (set this the same as you did for LOAD)
- NUM_RUNS = 10 (Is the minimum set of load points)
- WARMUP_TIME= (Optional. It will default to 300 seconds)

Getting ready to run SFS2014

Ensure that all of the clients have the filesystems mounted before starting

- Mount the test area from the storage server on each client node.
- You must mount the storage on every client node before you start the test. I highly recommend you put the mount in /etc/fstab so you won't have to type this over if you reboot.

Getting ready to run SPECsfs2014 (mount)

Example of mounting the filesystem

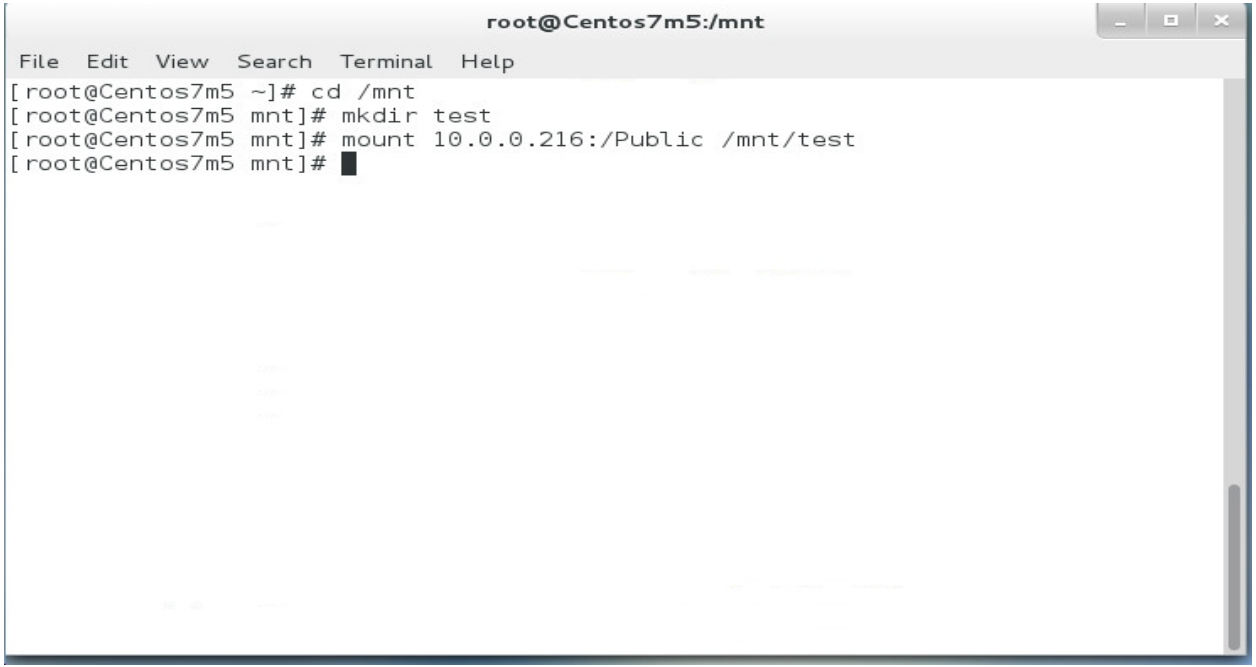

Free space on the storage server

You need this much space * LOAD value in the sfs_rc file

- You need this much storage space:

DATABASE	=	24 Gigabytes per DATABASE
SWBUILD	=	5 Gigabytes per BUILD
VDI	=	12 Gigabytes per DESKTOP
VDA	=	24 Gigabytes per STREAM

Ready to run SPECsfs2014 (check free space)

Example of checking free space before starting test

Start the sFlow, power, env, collection and GO.

SNIA Emerald™ Training

SNIA Emerald™ Power Efficiency Measurement Specification

Version 3.0

February-March 2018

Start collection tools

- Start the sflowtool collection now.
  Save its output file in a safe location.
  Example: sflowtool –L -4 > c:\temp\sflowdata.txt
- Start the Power meter
  Save its output in a safe place.
- Start the environmental monitor and
  save its output in a safe location.

Running SPECsfs2014 (Go for it)

Example script to run the 4 mandatory workloads. You need to edit this to tailor it for your environment.

```
########################################################################
# Example script to run the 4 SPECsfs2014 workloads required by Emerald.
########################################################################
#
#
########################################################################
# Run the SWBUILD workload
########################################################################
#
python SfsManager -r sfs_rc_swbuild -s swbuild
#
# Edit this to tailor it for your configuration
# Cleanup space for use by the next workload.
########################################################################
```

Running SPECsfs2014 (cont)

```
#
# Example:
# cd /mnt/workdir
# for i in *
# do
#   rm -rf $i &
# done
# wait
#
#########################################################################
# Run the VDA workload
#########################################################################
python SfsManager -r sfs_rc_vda -s vda
#
# Edit this to tailor it for your configuration
# Cleanup space for use by the next workload.
#
```

Running SPECsfs2014 (cont)

```
#
# Example:
# cd /mnt/workdir
# for i in *
# do
#   rm -rf $i &
# done
# wait
#
#########################################################################
# Run the VDI workload
#########################################################################
python SfsManager -r sfs_rc_vdi -s vdi
#
# Edit this to tailor it for your configuration
# Cleanup space for use by the next workload.
#
```

Running SPECsfs2014 (cont)

```
# Example:
# cd /mnt/workdir
# for i in *
# do
#   rm -rf $i &
# done
# wait
#########################################################################
# Run the DATABASE workload
#########################################################################
python SfsManager -r sfs_rc_database -s database
#
# Cleanup space for use in later runs
#
# Example:
# cd /mnt/workdir
```

Running SPECsfs2014 (cont)

```
# for i in *
# do
#   rm -rf $i &
# done
# wait
##########################################################################
```

Running SPECsfs2014 (console output)

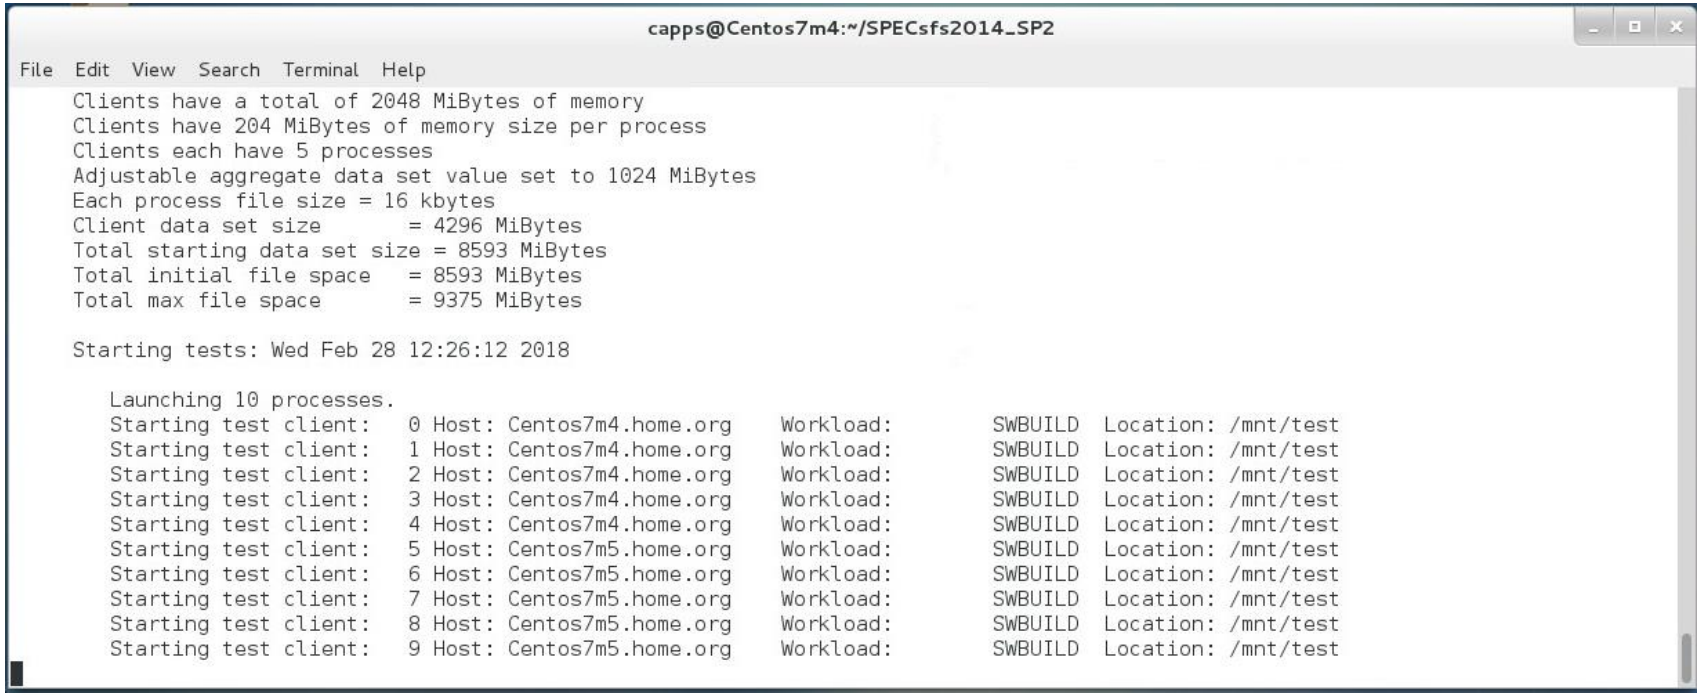

Running SPECsfs2014 (console output)

SNIA GSI | GREEN STORAGE

```
capps@Centos7m4:~/SPECsfs2014_SP2
File  Edit  View  Search  Terminal  Help
    Each process file size = 16 kbytes
    Client data set size       = 4296 MiBytes
    Total starting data set size = 8593 MiBytes
    Total initial file space   = 8593 MiBytes
    Total max file space       = 9375 MiBytes

    Starting tests: Wed Feb 28 12:53:19 2018

       Launching 10 processes.
       Starting test client:   0 Host: Centos7m4.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   1 Host: Centos7m4.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   2 Host: Centos7m4.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   3 Host: Centos7m4.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   4 Host: Centos7m4.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   5 Host: Centos7m5.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   6 Host: Centos7m5.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   7 Host: Centos7m5.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   8 Host: Centos7m5.home.org    Workload:        SWBUILD  Location: /mnt/test
       Starting test client:   9 Host: Centos7m5.home.org    Workload:        SWBUILD  Location: /mnt/test
       Wed Feb 28 12:53:28 2018 Prime's GO 'start comm' average message latency 4 ms
       Waiting to finish initialization. Wed Feb 28 12:53:28 2018
       Wed Feb 28 12:53:30 2018 Starting INIT phase
       Wed Feb 28 12:53:30 2018 Init Heartbeat __/\_/\__ from client 5
```

Running SPECsfs2014 (console output)

The heartbeat messages tell you that it's running. These arrive every minute

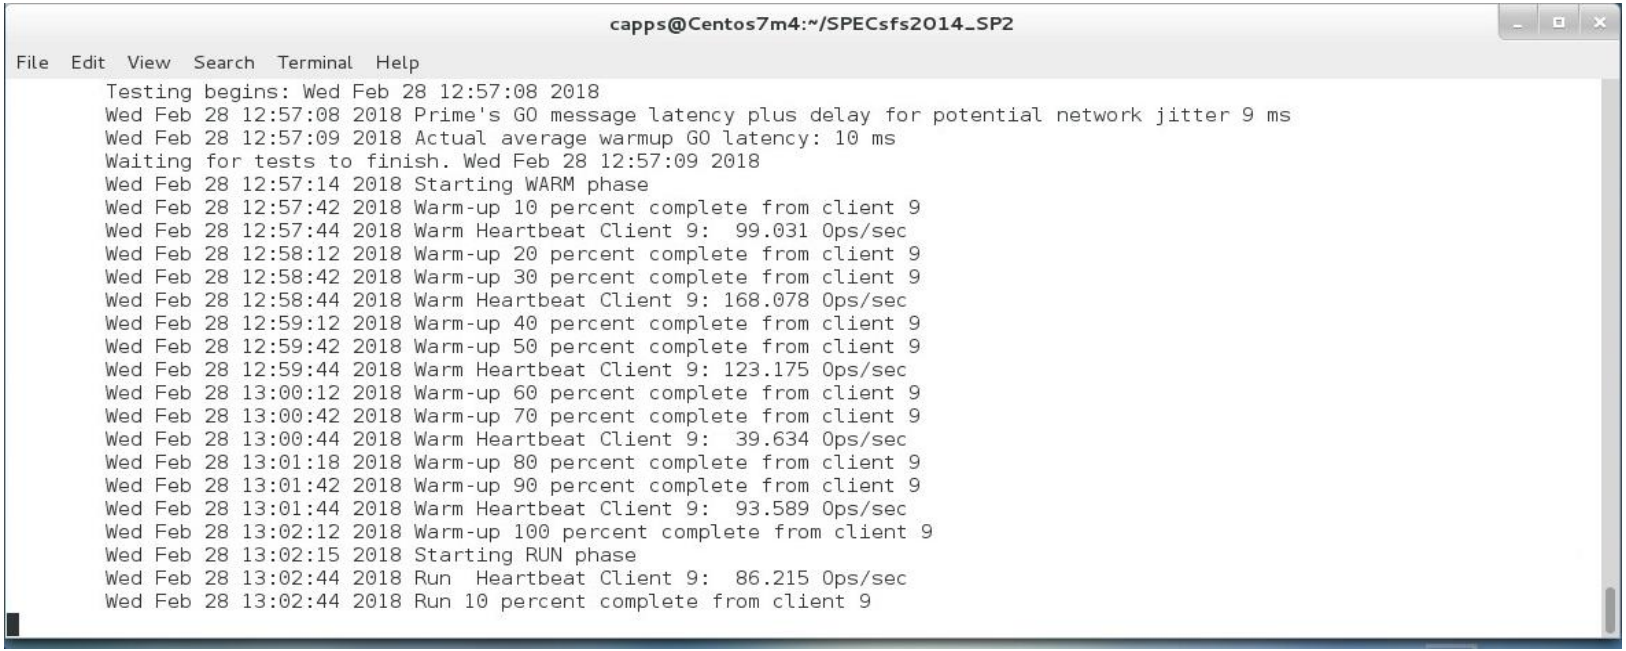

Running SPECsfs2014 (console output)

```
Wed Feb 28 13:03:44 2018 Run  Heartbeat Client 9:  99.905 Ops/sec
Wed Feb 28 13:04:15 2018 Run 40 percent complete from client 9
Wed Feb 28 13:04:44 2018 Run  Heartbeat Client 9: 121.495 Ops/sec
Wed Feb 28 13:04:45 2018 Run 50 percent complete from client 9
Wed Feb 28 13:05:14 2018 Run 60 percent complete from client 9
Wed Feb 28 13:05:44 2018 Run  Heartbeat Client 9: 164.049 Ops/sec
Wed Feb 28 13:05:45 2018 Run 70 percent complete from client 9
Wed Feb 28 13:06:15 2018 Run 80 percent complete from client 9
Wed Feb 28 13:06:44 2018 Run  Heartbeat Client 9: 156.090 Ops/sec
Wed Feb 28 13:06:45 2018 Run 90 percent complete from client 9
Wed Feb 28 13:07:14 2018 Run 100 percent complete from client 9
Wed Feb 28 13:07:14 2018 Prime receiving results from child 9
Tests finished: Wed Feb 28 13:07:14 2018
Shutting down clients, and communications layer...
----------------------------------------------------------------
Overall average latency              2.021 Milli-seconds
Overall SPEC SFS2014_SP2          1000.032 Ops/sec
Overall Read_throughput        ~   7119.807 Kbytes/sec
Overall Write_throughput       ~   1829.166 Kbytes/sec
Overall throughput             ~   8948.973 Kbytes/sec
Public Finger Print             24267
----------------------------------------------------------------
Band 1:  20us:56349    40us:19799    60us:1765     80us:431     100us:62
Band 2: 200us:10      400us:88684   600us:43643   800us:25863     1ms:12711
Band 3:   2ms:17970     4ms:19643     6ms:2595      8ms:1160     10ms:712
Band 4:  12ms:749      14ms:313      16ms:199      18ms:161      20ms:141
Band 5:  40ms:838      60ms:410      80ms:338     100ms:0
Band 6: 200ms:318     400ms:164     600ms:129     800ms:78        1s:39
```

Running SPECsfs2014 (console output)

Each load point will create one of these. You need 10 load points per workload type

```
capps@Centos7m4:~/SPECsfs2014_SP2
File  Edit  View  Search  Terminal  Help
        Wed Feb 28 13:07:14 2018 Prime receiving results from child 9
        Tests finished: Wed Feb 28 13:07:14 2018
        Shutting down clients, and communications layer...
        ----------------------------------------------------------------
        Overall average latency              2.021 Milli-seconds
        Overall SPEC SFS2014_SP2          1000.032 Ops/sec
        Overall Read_throughput        ~   7119.807 Kbytes/sec
        Overall Write_throughput       ~   1829.166 Kbytes/sec
        Overall throughput             ~   8948.973 Kbytes/sec
        Public Finger Print             24267
        ----------------------------------------------------------------
        Band 1:  20us:56349    40us:19799    60us:1765     80us:431     100us:62
        Band 2: 200us:10      400us:88684   600us:43643   800us:25863     1ms:12711
        Band 3:   2ms:17970     4ms:19643     6ms:2595      8ms:1160     10ms:712
        Band 4:  12ms:749      14ms:313      16ms:199      18ms:161     20ms:141
        Band 5:  40ms:838      60ms:410      80ms:338     100ms:0
        Band 6: 200ms:318     400ms:164     600ms:129     800ms:78       1s:39
        Band 7:    2s:29        4s:4          6s:0          8s:0         10s:0
        Band 8:   20s:0         40s:0         60s:0         80s:0        120s:0
        Band 9: 120+s:0

        ----------------------------------------------------------------
netmist completed successfully, summarizing.
[capps@Centos7m4 SPECsfs2014_SP2]$
```

Steps in the test procedure

- Per-Workload Calibration
  - Find maximum valid load for each workload on PUT (SPEC SFS 2014 metrics)
- Measurement
  - Run each workload as ten evenly-spaced load points up to the maximum valid load point
  - Collect environmental, power, and sFlow data
- Data Reduction
  - Derive efficiency metrics for each workload using data reduction methods
- SNIA Emerald Metrics
  - Find “sweet-spot” for all four workloads

Collecting the results

- The SFS2014 results are in $TOP/results on the Prime.
  sfssum.[swbuild,vda,vdi,database].txt
  sfslog.[swbuild,vda,vdi,database].txt

  Save these in some safe place.
- Save the sflowtool output file in this same location.
- Save the power monitor results in this same location.

Example results

Example results (sFlow and Power)

- http://www.demartek.com/private/SNIA_Emerald/Test4_files.zip
- sFlow raw data:
  CNTR,10.0.0.248, 2018-03-12 13:14:26
  IP addr Timestamp
  7,6,4294967295,1,3,66153,389,247,1,0,0,0, 598533882572,1274548005,30285,11610,44,0,1
  Bytes_in Bytes_out
- Power meter raw data

Store No.	Date	Time	Millisecond	WT1:U-1	WT1:I-1	WT1:P-1	WT1:S-1	WT1:Q-1	WT1:PF-1	WT1:Phi-1	WT1:Freq U-1	WT1:Freq I-1	WT1:U-1	WT1:Uthd-1	WT1:Ithd-1
1	2/27/2018	16:18:33	353	230.29	2.4751	5.35E+02	5.70E+02	-1.97E+02	0.9387	-20.17	60.002	60.003	230.29	0.439	13.76
2	2/27/2018	16:18:38	353	230.29	2.4622	5.32E+02	5.67E+02	-1.96E+02	0.9383	-20.23	60.002	60.002	230.29	0.422	12.256
3	2/27/2018	16:18:43	353	230.29	2.4735	5.35E+02	5.70E+02	-1.97E+02	0.9385	-20.19	60.002	60.002	230.29	0.431	13.894
4	2/27/2018	16:18:48	354	230.3	2.4497	5.29E+02	5.64E+02	-1.96E+02	0.9375	-20.36	60.003	60.002	230.3	0.442	12.421

Example results (SPEC sfs2014 results)

SNIA GSI | GREEN STORAGE

Example of the sfssum.vda.txt contents
Used for calibration. Not included in Emerald results

Business Metric	Requested Op Rate	Achieved Op Rate	Avg Lat (ms)	Total KBps	Read KBps	Write KBps	Run Sec	# Cl	Cl Proc	Avg File Size KB	Cl Data Set MiB	Start Data Set MiB	Init File Set MiB	Max File Space MiB	Workload Name	Valid Run
9	90.00	90.071	0.798	41599.242	3647.205	37952.037	300	1	18	1048576	202752	202752	202752	221184	VDA	
18	180.00	180.116	2.875	83026.161	6817.453	76208.708	300	1	36	1048576	405504	405504	405504	442368	VDA	
27	270.00	270.191	4.763	124814.636	10843.302	113971.334	300	1	54	1048576	608256	608256	608256	663552	VDA	
36	360.00	360.239	4.554	164445.594	13859.482	150586.112	300	1	72	1048576	811008	811008	811008	884736	VDA	
45	450.00	450.304	7.286	208721.185	17872.621	190848.564	300	1	90	1048576	1013760	1013760	1013760	1105920	VDA	
54	540.00	540.359	8.675	249397.387	21335.173	228062.214	300	1	108	1048576	1216512	1216512	1216512	1327104	VDA	
63	630.00	630.426	10.091	291766.929	24809.548	266957.381	300	1	126	1048576	1419264	1419264	1419264	1548288	VDA	
72	720.00	720.483	12.211	333282.071	29162.061	304120.009	300	1	144	1048576	1622016	1622016	1622016	1769472	VDA	
81	810.00	810.540	11.029	373242.740	31460.608	341782.132	300	1	162	1048576	1824768	1824768	1824768	1990656	VDA	
90	900.00	900.606	13.486	416206.961	35670.615	380536.346	300	1	180	1048576	2027520	2027520	2027520	2211840	VDA	

Data reduction

- Use sflowtool -4 –L to capture the sFlow data from the switch.
- Use tag2014 to merge the power and sFlow data into a time correlated data set of the RUN/measurement phase.
  https://github.com/powernap/tag2014/releases/tag/v1.5
  Author Nick Principe at iXsystems
- Use Excel to produce the MiBytes/sec/Watt efficiency values for the peak load point RUN phase of each workload.
- Patrick Stanko presents spreadsheet walkthrough.

Data reduction

```
C:\tmp>
C:\tmp>c:\"program files (x86)"\python36-32\python tag2014.py
Must specify an input file
USAGE: tag2014.py {-a|-c} [-f ts_col ... ] [-m] [-r] [-n] -i in_file -l sfslog -o out_file [-t time_shift]
     -a : Analyzer data (CSV data produced by Unisphere Analyzer)
     -c : CSV data

     -f ts_col : field(s) that contains timestamp information

     -m : restrict output, include WARMUP data
     -r : restrict output, include RUN data
     -n : restrict output, include RUN_TAIL data

     -i in_file  : input data file
     -l sfslog   : sfslog file
     -o out_file : output file, omit for STDOUT

     -t time_shift : shift the time in the data file by time_shift seconds
```

Data reduction. Example SWBUILD

Run #	sFlow MiB/sec	Avg Power	MiB/sec/Watt
1	3000	1500	2.00
2	5000	1502	3.32
3	7000	1504	4.65
4	9000	1506	5.97
5	11000	1508	7.29
6	13000	1510	8.61
7	15000	1512	9.92
8	17000	1514	11.22
9	19000	1516	12.53
10	21000	1518	**13.83**

Data reduction. Example VDA

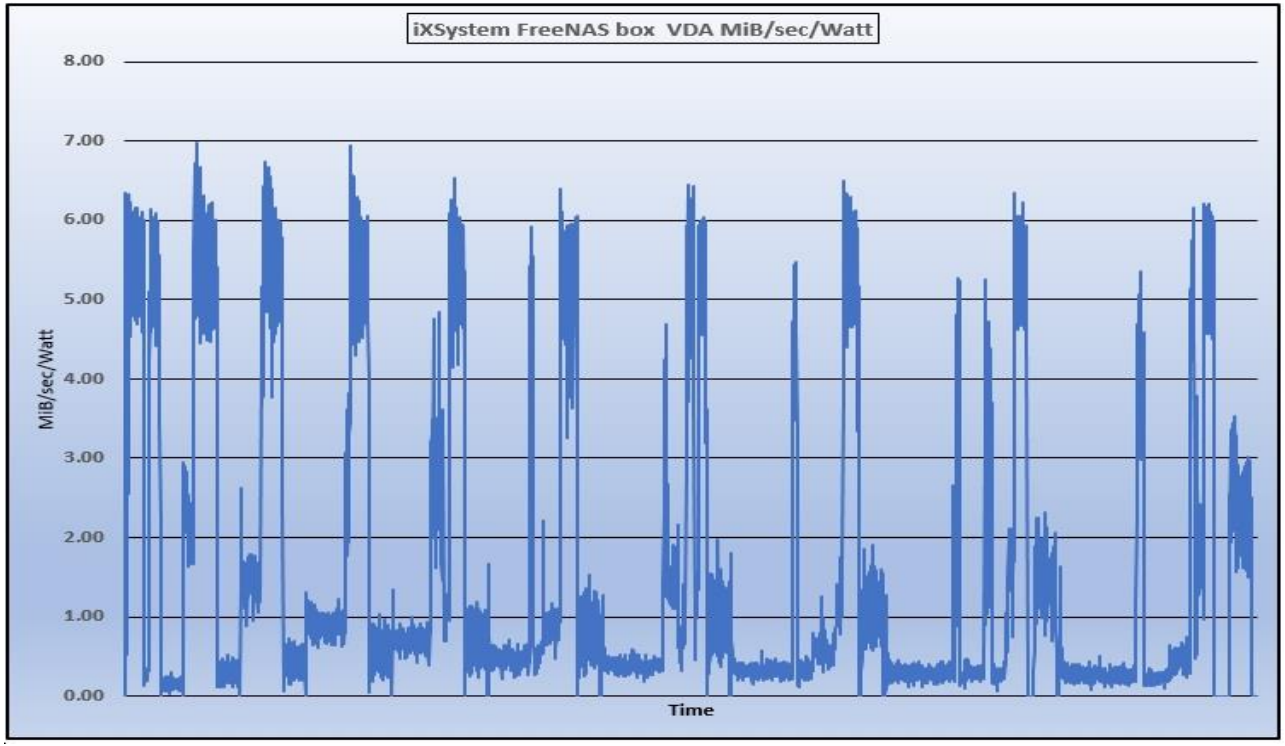

Data reduction

Workloads	SWBUILD	VDI	VDA	DATABASE
Maximum Efficiency	13.83 MiB/sec/Watt	18.75 MiB/sec/Watt	20.20 MiB/sec/Watt	17.33 MiB/sec/Watt

BFF and range submitted

- After one has found the peak MiB/sec/Watt from each of the SFS2014 workloads. Find the best foot forward by increasing and decreasing drive count and run suite again. This may take several iterations.
- Required by the EPA's Energy Star program.
  For the low and high range points use either of the below
  - Efficiency results with -40% drive count and + 15 % drive count.
  - Storage device count points where Perf/W value is 15% < the BFF point

BFF and range submitted

Reference previous V2.0.2 Emerald Training slides

"Product Family, Best Foot Forward, Test Points and Qualification Ranges" see starting with slide #30

If things go wrong

- Filenames in the SFS2014 distribution are all lowercase or all upper case. This happens on many Unix systems. One can either fight their way through the RockRidge extensions at mount time, or just load the DVD in a Windows system, and copy the contents to a common shared location.
- Netmist load generator stores per process logs on each client in /tmp/netmist_C*.log (or wherever, if specified) These contain details of causes for failures.
- The sfslog.[vda,vdi,swbuild,database].txt files contain a high level summary, with a pointer to which client process failed. (so you know what log to examine in /tmp)

Common failures

- DNS not setup properly. Both hostname to IP address and IP address to hostname need to work for all clients.
- Ensure that 127.0.0.1 hostname is *not* in /etc/hosts
- Password challenges during startup. Failure to setup the ssh keys so that challenges are not needed.
- Hangs at the beginning. Forgot to disable firewalls.
- ENOSPC. Failure to configure sufficient disk space.
- Runs marked INVALID. Load is exceeding server capabilities.

Common failures

- Windows failure to startup. You must have a dedicated Prime that does not present load.
- You must have credentials that give the benchmark permission to access each other via WMI, and permission to access the storage. These are Active Directory configuration issues. Talk to the IT person that maintains the Active directory for your lab.
- Set User in the sfs_rc file to: DOMAIN\accountname and the Password to the account's password.
- If all else fails, there is a SPECsfs2014 User's guide here: https://www.spec.org/sfs2014/index.html#userguide

Common failures

- Contact the SPEC Storage subcommittee support at:
  sfs2014support@spec.org

Q & A

- Q & A session