PROFITAP

WHITE PAPER:

IN-LINE TAPPING IN THE DATA CENTER

WIRESHARK HEROES SERIES

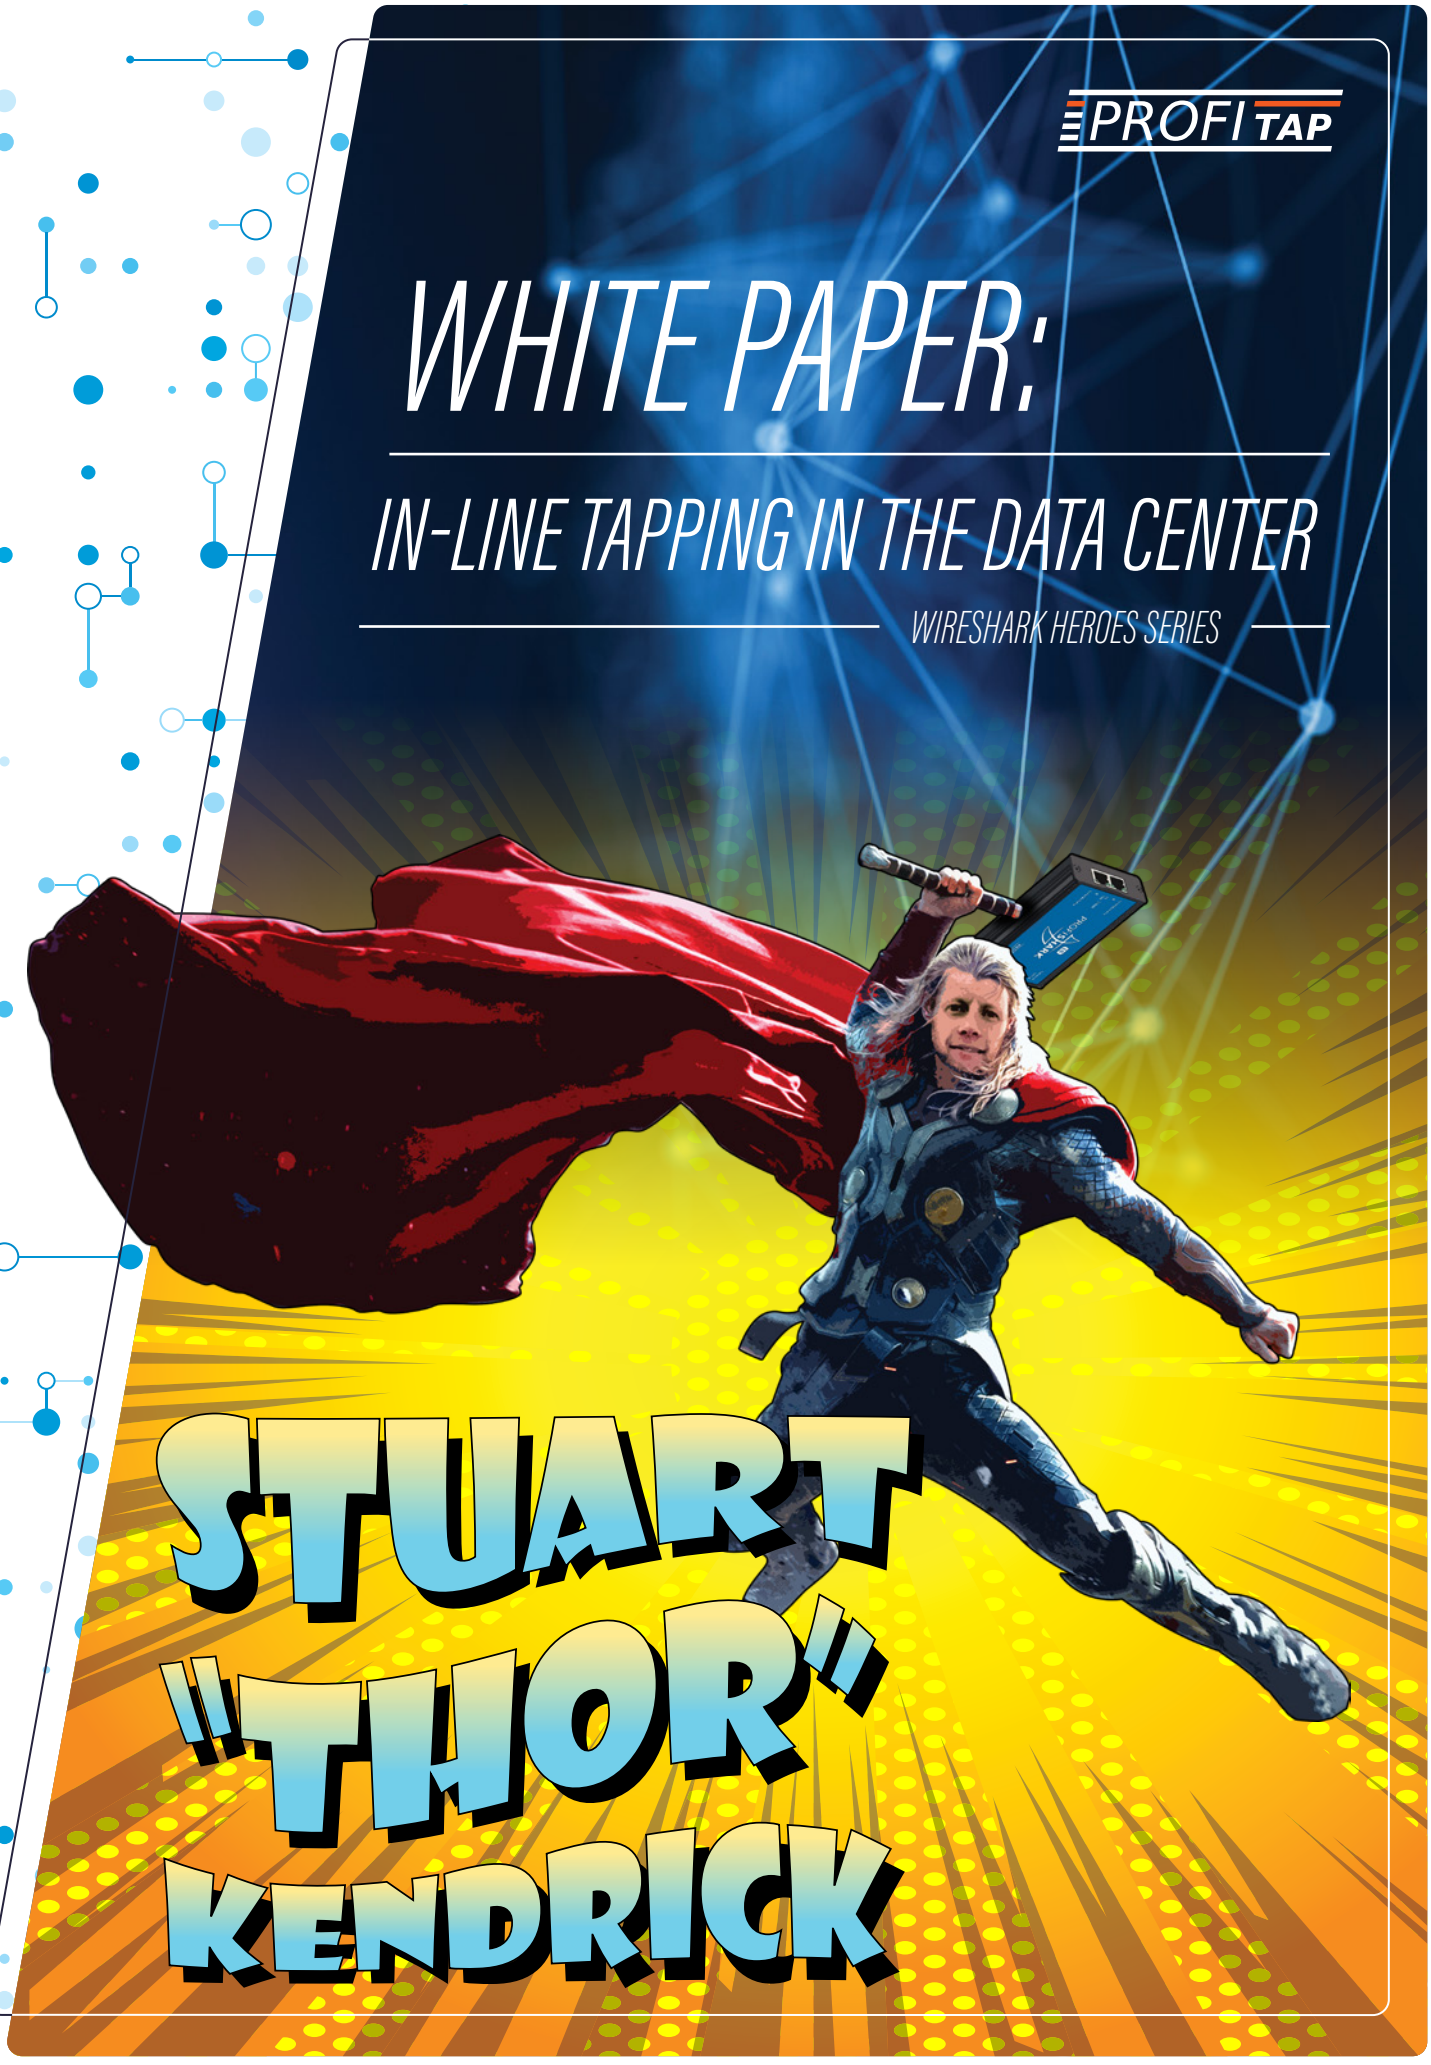

VISIT WWW.PROFITAP.COM

STUART KENDRICK

SYSTEMS ENGINEER AT ALLEN INSTITUTE

Has functioned as both ITIL Problem Manager and Problem Analyst, provided 3rd tier support, and contributed to design efforts. Stuart writes and maintains an enterprise network and device management and monitoring application (the Netops Toolkit). He specializes in transport, monitoring, and packet analysis, provide mentoring and communication training, teaches Root Cause Analysis workshops, and coordinates the efforts of multiple groups interacting with multiple vendors to solve problems or design solutions.

EXPERTISE:

- *Root Cause Analysis*
- *ITIL Problem Management*
- *SNMP-based Management Applications*
- *Packet Analysis*
- *Ethernet/IP Transport*

WWW.SKENDRIC.COM

CONTENT

OVERVIEW

I find a portable in-line tap to be a useful trouble-shooting tool generally – easy to inert in between a desktop station and its network jack, grab a pcap, see what it is happening.

But I even find it useful in the Data Center, and that is the subject of this post.

SCALE-OUT STORAGE

We deploy an Isilon OneFS storage system. From a physical point of view, the Isilon product looks like a bunch of 4 RU servers, sporting 10G (or 40G) Ethernet NICs on their front-side and 40G Ethernet (or, for the older nodes, InfiniBand) NICs on their back-side (all the Nodes talk to each other over the back-side network). The more Nodes you add, the more storage, RAM, cache, and network I/O the system offers. And it scales from hundreds of TB to hundreds of PB.

From a logical point of view, all those nodes present their space inside a single file system. For storage administrators supporting certain applications, this a big win – typical storage products require that you divvy up your total storage into little hunks of tens, hundreds, or occasionally a few thousands of TB. And you are forever shuffling files around from one 'volume' to another, as a given volume runs out of space. In large systems, this chore consumes FTEs; in a OneFS system, this chore doesn't exist – the entire storage space lives inside a single file system. By analogy, consider if your laptop ran OneFS. Every time you ran out of space, what if you just plugged another USB stick into it and poof!, C:\ just got bigger. That's what an Isilon system feels like when you are driving it.

This approach shines for us – in our business, we capture high-resolution images of cells and their interconnections, streaming off custom-built microscopes. Each year, we purchase a few more PB worth of nodes (starting next year, a few more tens of PB), plug them into the Isilon cluster, it mutters to itself for a few hours (OK, sometimes for a few days), and then away we go – more space.

THE CHALLENGE

We had deployed an IP scheme for our cluster without understanding the cluster's demands for IP addresses. For highly-available NFS, the cluster does fine by assigning a single IP address to each node. But for highly-available SMB, the cluster wants several IP addresses per node, for reasons which escape me at the moment. We are at 46 nodes today, planning to add another hundred plus over the next few years. And we were running out of IP addresses. So, we devised a plan to pipe another VLAN into the cluster, an empty /22, and then migrate the cluster into this new subnet.

DC Isilon

gila

Network Design Intent

(1) 10G interface carrying Data Plane traffic
(1) 1G interface carrying management traffic

Configuration

```
gila-1% isi network subnets list

ID                  Subnet              Gateway|Priority    Pools                SC Service
-------------------------------------------------------------------------------------------
groupnet0.subnet0   10.80.104.0/22      10.80.104.1|1       Production           10.80.106.136
groupnet0.subnet1   10.80.100.0/22      10.80.100.1|10      Management           10.80.102.74
groupnet0.subnet2   172.20.0.0/16       172.20.0.1|2        HPC                  172.20.102.136
groupnet0.subnet3   10.80.112.0/22      10.80.112.1|3       Production-Static    10.80.112.15
                                                            Production-Dynamic
-------------------------------------------------------------------------------------------
Total: 4
gila-1%
gila-1% isi network interfaces list

LNN    Name                Status       Owners                          IP Addresses
-------------------------------------------------------------------------
1      10gige-1            Up           groupnet0.subnet0.Production    10.80.106.75
                                        groupnet0.subnet2.HPC           172.20.102.86
                                                                        172.20.102.112
                                                                        172.20.102.124
1      10gige-2            No Carrier   -                               -
1      ext-1               Up           groupnet0.subnet1.Management    10.80.102.88
1      ext-2               No Carrier   -                               -
```

Fine, how hard can it be? We already have (2) VLANs piped into this cluster: just add a third. Well, we tried this, and the entire cluster become inaccessible. The catch with building one big storage system that everyone uses is that ... well, when it quits working, everyone notices. But we'll dance ahead of that sorry moment and focus on the technical side of the issue.

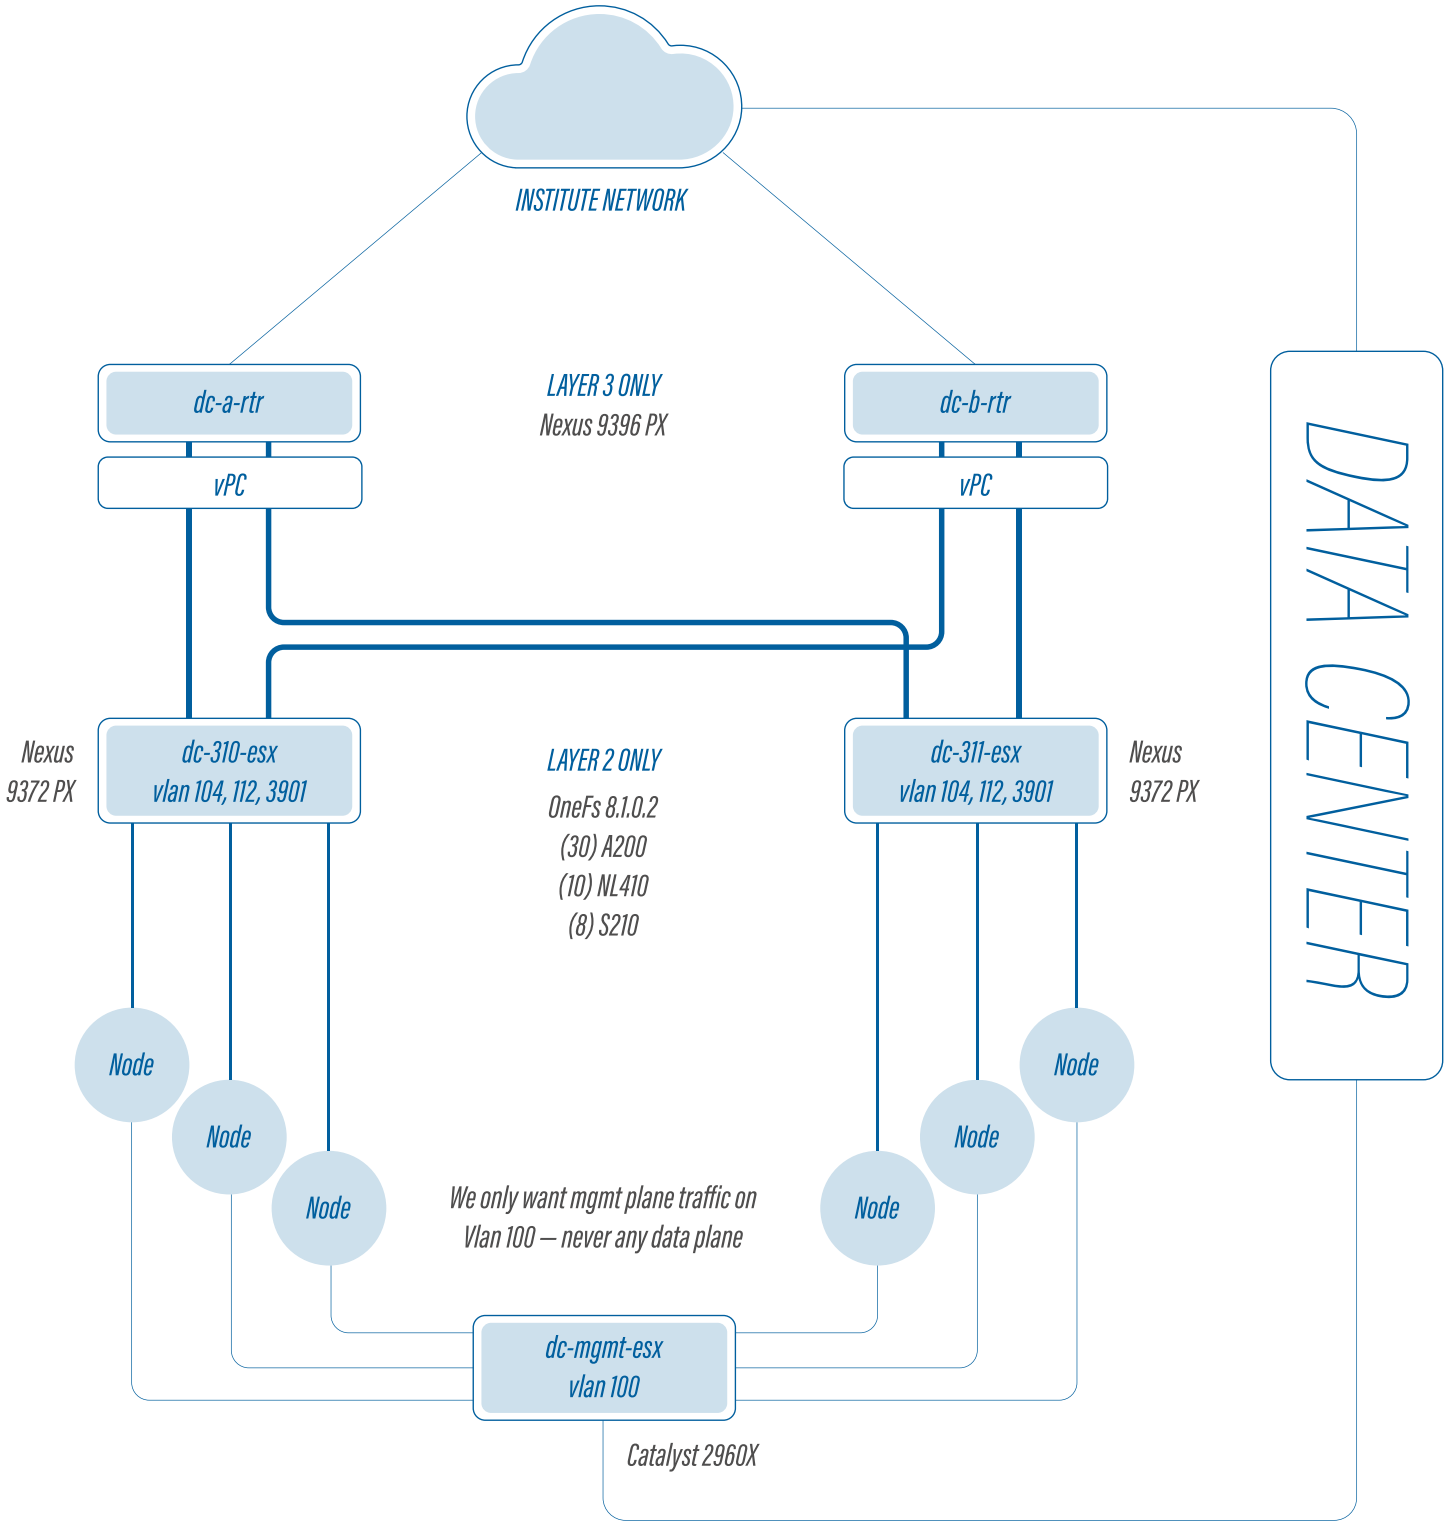

LEGEND

Vlan 100 = 10.80.100.0/22	Mgmt only
Vlan 104 = 10.80.104.0/22	Legacy Data Plane (SMB & NFS
Vlan 112 = 10.80.112.0/22	New Data Plane (SMB & NFS)
Vlan 3901 = 172.20.0.0/16	HPC Data Plane (NFS only)

1G Ethernet

10G Ethernet

40G Ethernet

THE CHANGE

BEFORE:

OneFS lets you configure the cluster's view of the world via a GUI or via a CLI. Once you make a change to the cluster configuration, OneFS then propagates that change to each node for you.

From a networking point of view, here is what the cluster looked like before we tried to pipe the new VLAN (V112) into the cluster.
Focus on the Purple Lines.

```
gila-1 02:54:33% isi network subnets view groupnet0.subnet0
isi network subnets view groupnet0.subnet0
               ID:  groupnet0.subnet0
             Name:  subnet0
        Groupnet:  groupnet0
           Pools:  Production
      Addr Family:  ipv4
        Base Addr:  10.80.104.0
             CIDR:  10.80.104.0/22
      Description:  Production
        Dsr Addrs:  -
          Gateway:  10.80.104.1
 Gateway Priority:  1
              MTU:  1500
        Prefixlen:  22
          Netmask:  255.255.252.0
  Sc Service Addr:  10.80.106.136
  Sc Service Name:  Sc Service Name:
     VLAN Enabled:  False
          VLAN ID:  104
  gila-1 02:58:25%


gila-1 02:58:25% isi network subnets view groupnet0.subnet2
isi network subnets view groupnet0.subnet2
               ID:  groupnet0.subnet2
             Name:  subnet2
        Groupnet:  groupnet0
           Pools:  HPC
      Addr Family:  ipv4
        Base Addr:  172.20.0.0
             CIDR:  172.20.0.0/16
      Description:  HPC
        Dsr Addrs:  -
          Gateway:  172.20.0.1
 Gateway Priority:  100
              MTU:  1500
        Prefixlen:  16
          Netmask:  255.255.0.0
  Sc Service Addr:  172.20.102.136
  Sc Service Name:
     VLAN Enabled:  True
          VLAN ID:  3901
  gila-1 02:59:16%
```

```
gila-1 02:59:16% isi network subnets view groupnet0.subnet3
isi network subnets view groupnet0.subnet3
              ID:   groupnet0.subnet3
            Name:   subnet3
        Groupnet:   groupnet0
           Pools:   Production-Static, Production-Dynamic
     Addr Family:   ipv4
       Base Addr:   10.80.112.0
            CIDR:   10.80.112.0/22
     Description:   Production static/dynamic
       Dsr Addrs:   -
         Gateway:   10.80.112.1
Gateway Priority:   Gateway Priority: 3
             MTU:   1500
       Prefixlen:   22
         Netmask:   255.255.252.0
 Sc Service Addr:   10.80.112.15
 Sc Service Name:
    VLAN Enabled:   False
         VLAN ID:   112
 gila-1 02:59:33%
```

And here is what the switch ports looked like:

```
interface Ethernet1/2
   description Isilon
   switchport mode trunk
   switchport trunk native vlan 104
   switchport trunk allowed vlan 104,3901
   spanning-tree port type edge
   spanning-tree guard root
   mtu 9216
   storm-control broadcast level 1.00
   storm-control multicast level 1.00
   storm-control action shutdown
   storm-control action trap
```

After:

So we come along and enable VLAN tagging on V112. Again, focus on the purple lines.

```
gila-2 02:54:07% isi network subnets modify groupnet0.subnet0
--vlanenabled=true
isi network subnets modify groupnet0.subnet0 --vlan-enabled=true
isi network subnets modify groupnet0.subnet3 --vlan-enabled=true
gila-2 02:54:28% isi network subnets modify groupnet0.subnet3
--vlanenabed=true
In-Line Tapping in the Data Center 6 Created: 2018-05-19
Stuart Kendrick Updated: 2018-05-19
isi network subnets modify groupnet0.subnet3 --vlan-enabled=true
gila-2 02:54:29%
```

```
gila-1 02:53:56% isi network subnets view groupnet0.subnet0
isi network subnets view groupnet0.subnet0
              ID:  groupnet0.subnet0
            Name:  subnet0
        Groupnet:  groupnet0
           Pools:  Production
     Addr Family:  ipv4
       Base Addr:  10.80.104.0
            CIDR:  10.80.104.0/22
     Description:  Production
       Dsr Addrs:  -
         Gateway:  10.80.104.1
Gateway Priority:  1
             MTU:  1500
       Prefixlen:  22
         Netmask:  255.255.252.0
 Sc Service Addr:  10.80.106.136
 Sc Service Name:
    VLAN Enabled:  True
         VLAN ID:  104

gila-1 02:54:30% isi network subnets view groupnet0.subnet3
isi network subnets view groupnet0.subnet3
              ID:  groupnet0.subnet3
            Name:  subnet3
        Groupnet:  groupnet0
           Pools:  Production-Static, Production-Dynamic
     Addr Family:  ipv4
       Base Addr:  10.80.112.0
            CIDR:  10.80.112.0/22
     Description:  Production static/dynamic
       Dsr Addrs:  -
         Gateway:  10.80.112.1
Gateway Priority:  3
             MTU:  1500
       Prefixlen:  22
         Netmask:  255.255.252.0
 Sc Service Addr:  10.80.112.15
 Sc Service Name:
    VLAN Enabled:  True
         VLAN ID:  112
 gila-1 02:54:33%


interface Ethernet1/2
 description Production and HPC
 switchport mode trunk
 switchport trunk native vlan 104
 switchport trunk allowed vlan 104,112,3901
 spanning-tree port type edge
 spanning-tree guard root
 mtu 9216
 storm-control broadcast level 1.00
 storm-control multicast level 1.00
 storm-control action shutdown
 storm-control action trap
```

I use my favorite tactical monitoring tool, mass-ping[1], to watch the cluster during the change … quickly see that things are going south … and we back out.

```
root@vishnu:/home/netops/rpts/mass-ping/Isilon/Enable-VLANs/2018-03-22#
massping
-s yes -f /home/netops/etc/dc-isilon-gear -n enable-vlan-tagging-2 -m .
-c "Enable Vlan Tagging 2"
Sanity check...
Identifying live hosts...

Beginning with 60 live addresses
Starting: Thursday March 22, 2018 at 02:52:19
Pinging targets every 1 seconds with timeout 0.2 seconds, running for 10
minutes, hit Ctrl-C to cancel...
60 60 60 60 60 60 60 60 60 60 15 15 17 20 24 27 28 29 29 29 29 29 24 29 29 29 29
29 30 30 30 30 28 28 28 29 31 32 34 36 37 39 40 42 44 45 47 49 51 52 54 56 57 5
8 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60
60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60
60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 15 17 21 2
5 28 28 25 29 29 28 29 29 29 29 29 29 29 30 30 30 30 30 28 28 29 31 32 34 36 37
38 40 42 44 46 47 49 51 53 54 56 58 58 58 58 60 60 60 60 60 60 60 60 60 60 60 60
60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 6
0 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60
60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60
60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 60 6
[…]
# Title: Mass Ping Report
#
# Institution: Widgets International
#
# Date of Report: Thursday March 22, 2018 at 02:57:20
#
# Description: This report portrays pings hit and missed
#
# Active: 60
#
# Title; Enable Vlan Tagging 2
#
# Errors:
#
# Questions: If you have questions or comments regarding this
# report, please mail them to xxx.
#
# target                hits             misses
# ---------------------------- ----- ------
 gila-01                295              6
 gila-02                295              6
 gila-03                247              54
 gila-04                296              5
 gila-05                296              5
 gila-06                294              7
 gila-07                295              6
 gila-08                294              7
 gila-09                291              10
 gila-10                291              10
 gila-11                291              10
```

[1] *mass-ping pings a bunch of IP addresses, going to great effort to ping each one exactly once/second. It gives you a simple CLI display as to how many of those addresses are returning pings each second … and after it has finished running, it produces both a textual and a graphical report.*

```
gila-12             298             3
gila-13             295             6
gila-14             277             24
gila-15             277             24
gila-16             250             51
gila-17             251             50
gila-18             251             50
gila-19             239             62
gila-20             234             67
gila-21             237             64
gila-22             235             66
gila-23             232             69
gila-24             222             79
gila-25             235             66
gila-26             236             65
gila-27             230             71
gila-28             218             83
gila-29             245             56
gila-30             226             75
gila-31             242             59
gila-32             222             79
gila-33             235             66
gila-34             236             65
gila-35             228             73
gila-36             239             62
gila-37             237             64
gila-38             233             68
gila-39             244             57
gila-40             224             77
gila-41             239             62
gila-42             239             62
gila-43             223             78
gila-44             222             79
gila-45             220             81
gila-mgt-01         300             1
gila-mgt-02         301             0
gila-mgt-03         256             45    # This is weird - I'm ignoring
                                          # it for now
gila-mgt-04         301             0
gila-mgt-05         301             0
gila-mgt-06         301             0
gila-mgt-07         301             0
gila-mgt-08         301             0
gila-mgt-09         301             0
gila-mgt-10         301             0
gila-mgt-11         301             0
gila-mgt-12         301             0
gila-mgt-13         300             1
gila-mgt-14         300             1
gila-mgt-15         300             1

Ending /opt/local/script/mass-ping
root@vishnu:/home/netops/rpts/mass-ping/Isilon/Enable-VLANs/2018-03-22#
```

I don't have an explanation for why some nodes missed more pings than others. Nor why that single management address missed so many pings – I predicted that the management addresses would be unaffected by this event, since they live on separate NICs attached to a separate physical network. The pain we experienced is perhaps easier to see in a graphical view of mass-ping output:

Mass-Ping: Enable Vlan Tagging 2

Title	Mass-Ping: Enable Vlan Tagging 2
Invocation	mass-ping -i 1 -w 600 -t 0.2 -f /home/hetops/etc/dc-isilon-gear
Details	Run from vishnu on 2018-03-22 at 02-52-18 by root
Errors	
Node count	60
Time count	301

ANALYZE

What is going on? Well, my first thoughts turned, of course, to VLAN tagging – are the switch and the Nodes disagreeing on which frames to tag?

Naturally, the network person thinks they have configured the Nexus switch correctly and the storage person thinks they have configured the Isilon Node correctly.

Time to grab a pcap of the traffic a node and a switch are exchanging, during a repeat of this change. I could of course run tcpdump on the Isilon nodes and SPAN a port on the Nexus switches, in order to capture pcaps. However, I have had only intermittent success in capturing VLAN tags using these methods. Some switches strip out VLAN tags before forwarding frames to a SPAN port; and some NIC drivers strip off VLAN tags before forwarding them to libpcap for tcpdump (or dumpcap or Wireshark) to grab.[2]

So instead, I pulled out an in-line tap – in my case, a ProfiShark 10G. This cute little box has (2) SFP+ inputs and (1) USB 3.1 output. I insert the ProfiShark in-line with the Isilon Node.

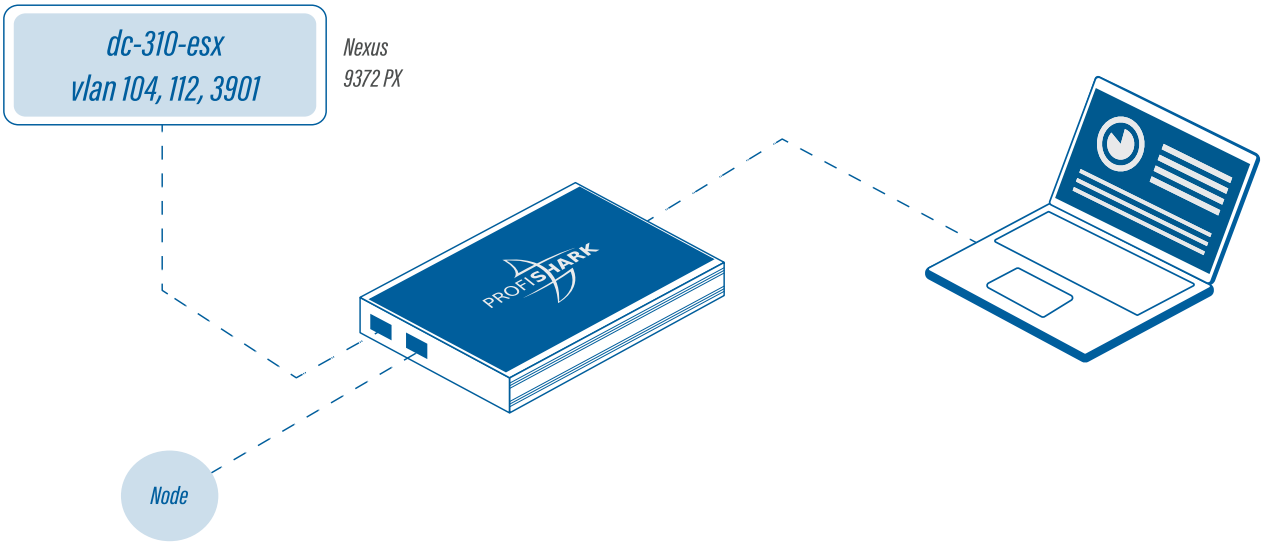

Next, we see the blue ProfiShark 10G unit sitting in top of a stack of Isilon Nodes, operated by the laptop visible at the bottom.

[2] *As an aside, both these methods also stumble when faced with link-local traffic, like LACP and UDLD Hellos; analyzing problems with those protocols also wants an in-line tap.*

Then I go home and wait for the next outage window – I will RDP into my laptop, load Wireshark, and capture on traffic flowing between the Node and the Swtich.

Now, experienced analysts will note that I'm skating over several issues here. First, yes, I did isolate this Node when I installed the ProfiShark. I had an advantage here – OneFS is a distributed system, meaning that a Node can go down, and the end-users don't notice – OneFS dynamically redistributes client connections to other Nodes. So I could do this in the middle of the business day.

In addition, when I start capturing during the next outage window, no way can I capture line-rate 10G traffic – the laptop's hard disk would be overwhelmed, and the resulting pcap would be incomplete. Ideally, I would use a high-end capture engine which can, in fact, capture at linerate 10G. Yes, that's true. And sometimes these nodes are, in fact, running close to line rate 10G. However, for this analysis, I don't care – I just need to see some frames from each direction, in order to assess their tagging. And, in general, I find that most of my servers aren't pushing anywhere near line-rate, and this USB / laptop-hard-disk scheme functions just fine, capturing *all* frames.

As an aside, you can use Wireshark to capture frames from a ProfiShark. Or, you can fire up the heavy client which ProfiTap bundles with their hardware. The Capture screen from that application portrayed below – notice the 'Dropped' count in the lower-left hand corner: this tells you if any frames were, in fact, dropped during this capture session.

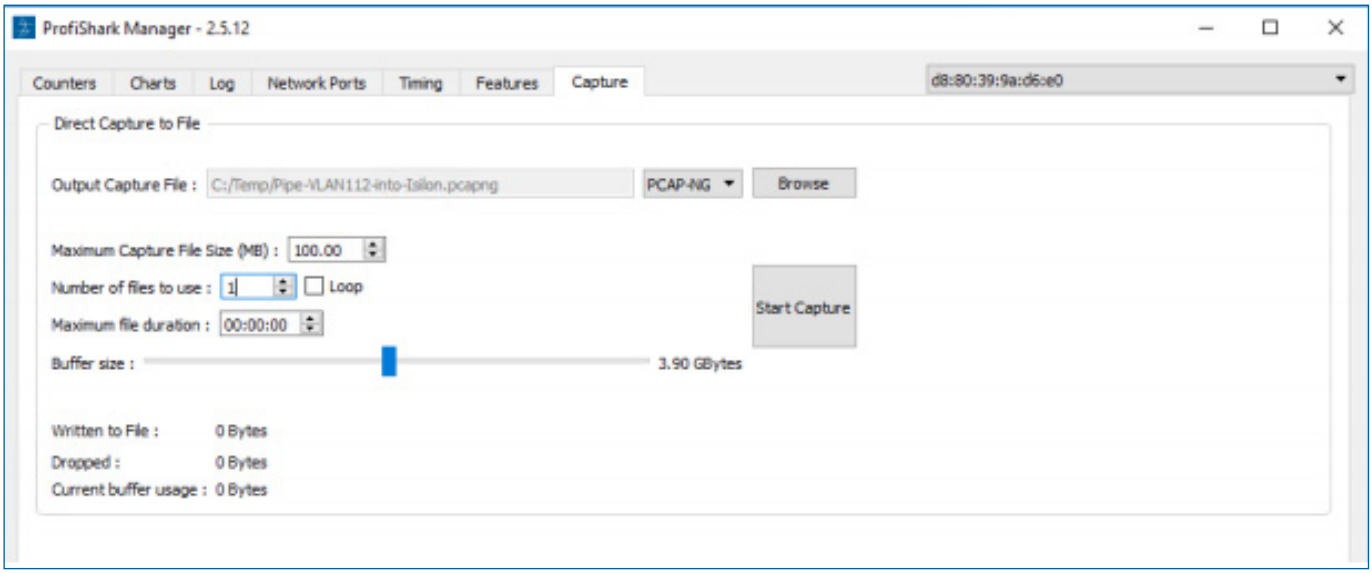

See the Appendix for more screen shots taken from this application.

Anyway, so during the next outage window, we try again, and this time, I capture a pcap.

THE ANSWER

So the Node is tagging VLAN 104, 112, and 3901 frames, while the Switch is tagging only the latter two … more specifically, the Switch is not tagging VLAN 104 frames. [If configured correctly, the Switch would have inserted '104' into the VLAN ID field into those 'TCP Retransmission' frames, see below.]

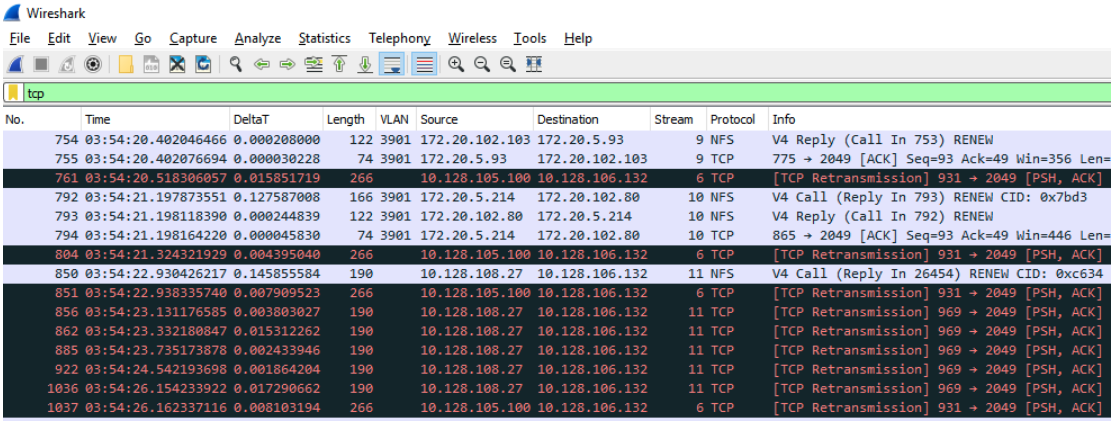

No.	Time	DeltaT	Length	VLAN	Source	Destination	Stream	Protocol	Info
754	03:54:20.402046466	0.000208000	122	3901	172.20.102.103	172.20.5.93	9	NFS	V4 Reply (Call In 753) RENEW
755	03:54:20.402076694	0.000030228	74	3901	172.20.5.93	172.20.102.103	9	TCP	775 → 2049 [ACK] Seq=93 Ack=49 Win=356 Len=
761	03:54:20.518306057	0.015851719	266		10.128.105.100	10.128.106.132	6	TCP	[TCP Retransmission] 931 → 2049 [PSH, ACK]
792	03:54:21.197873551	0.127587008	166	3901	172.20.5.214	172.20.102.80	10	NFS	V4 Call (Reply In 793) RENEW CID: 0x7bd3
793	03:54:21.198118390	0.000244839	122	3901	172.20.102.80	172.20.5.214	10	NFS	V4 Reply (Call In 792) RENEW
794	03:54:21.198164220	0.000045830	74	3901	172.20.5.214	172.20.102.80	10	TCP	865 → 2049 [ACK] Seq=93 Ack=49 Win=446 Len=
804	03:54:21.324321929	0.004395040	266		10.128.105.100	10.128.106.132	6	TCP	[TCP Retransmission] 931 → 2049 [PSH, ACK]
850	03:54:22.930426217	0.145855584	190		10.128.108.27	10.128.106.132	11	NFS	V4 Call (Reply In 26454) RENEW CID: 0xc634
851	03:54:22.938335740	0.007909523	266		10.128.105.100	10.128.106.132	6	TCP	[TCP Retransmission] 931 → 2049 [PSH, ACK]
856	03:54:23.131176585	0.003803027	190		10.128.108.27	10.128.106.132	11	TCP	[TCP Retransmission] 969 → 2049 [PSH, ACK]
862	03:54:23.332180847	0.015312262	190		10.128.108.27	10.128.106.132	11	TCP	[TCP Retransmission] 969 → 2049 [PSH, ACK]
885	03:54:23.735173878	0.002433946	190		10.128.108.27	10.128.106.132	11	TCP	[TCP Retransmission] 969 → 2049 [PSH, ACK]
922	03:54:24.542193698	0.001864204	190		10.128.108.27	10.128.106.132	11	TCP	[TCP Retransmission] 969 → 2049 [PSH, ACK]
1036	03:54:26.154233922	0.017290662	190		10.128.108.27	10.128.106.132	11	TCP	[TCP Retransmission] 969 → 2049 [PSH, ACK]
1037	03:54:26.162337116	0.008103194	266		10.128.105.100	10.128.106.132	6	TCP	[TCP Retransmission] 931 → 2049 [PSH, ACK]

Aha! So, if the Node wants to tag VLAN 104 frames but the switch does not do so, then the Node is discarding incoming (untagged) frames. And that pretty well breaks things.

What is going on? Let's look again at the switch port configuration:

```
interface Ethernet1/2
 description Production and HPC
 switchport mode trunk
 switchport trunk native vlan 104
 switchport trunk allowed vlan 104,112,3901
 spanning-tree port type edge
 spanning-tree guard root
 mtu 9216
 storm-control broadcast level 1.00
 storm-control multicast level 1.00
 storm-control action shutdown
```

What do these two lines do?

```
 switchport trunk native vlan 104
 switchport trunk allowed vlan 104,112,3901
```

Well, we thought they told the switch:

1. Allow frames for VLANs 104, 112, and 3901 onto this port, tagging whatever you transmit
2. And if you receive an untagged frame, accept it and tag it with '104'

But it turns out that it really means:

1. Allow frames for VLANs 104, 112, and 3901 onto this port, tagging whatever you transmit (but see caveat below)
2. When you receive an untagged frame, tag it with '104'
3. And when you transmit a frame arriving from VLAN 104, strip off its tag and then transmit it

And that characteristic #3 was breaking things – the Isilon Node did not have a similar concept of 'native VLAN', and thus discarded untagged (subnet 10.80.104.0/22) frames. Most protocols no worky when one side is tossing all the traffic you send it.

Now, you could argue that if we had a smarter network person, we wouldn't have had to capture a pcap – a smarter network person would have understood the 'native VLAN' concept better, would have seen the mis-interaction with how we were configuring the Isilon node, and would not have made this error in the first place. Heck, a smarter storage person would have picked this up. And I agree.
But we are a small shop, none of us are specialists … we are all generalists … we just aren't that smart. I like working here – I get to do lots of things … but there's no doubt that, as a result, I also get to feel incompetent most days. There are pros and cons to working in small shops versus working in big shops.
So that's my story. Quick, cheap, easy-to-deploy, portable in-line tapping in the Data Center: it is a good thing.

APPENDIX

ProfiShark Manager

Profitap bundles a heavy client called ProfiShark Manager with their gear. You don't have to use it – you can use your favorite analyzer (e.g. Wireshark, many others) to capture pcaps. But the heavy client does offer some neat screens, which I illustrate below.
The discerning reader will notice that the screen shots below are taken from ProfiShark Manager plugged into a ProfiShark 1G+ (the '+' means 'GPS equipped') and not the ProfiShark 10G which I used for the analysis described in this document. Aside from the GPS screen, the ProfiShark Manager GUI offers the same features across all models.

Counters:

A classic speed dial display, gives you quick insight into how full the pipe is.

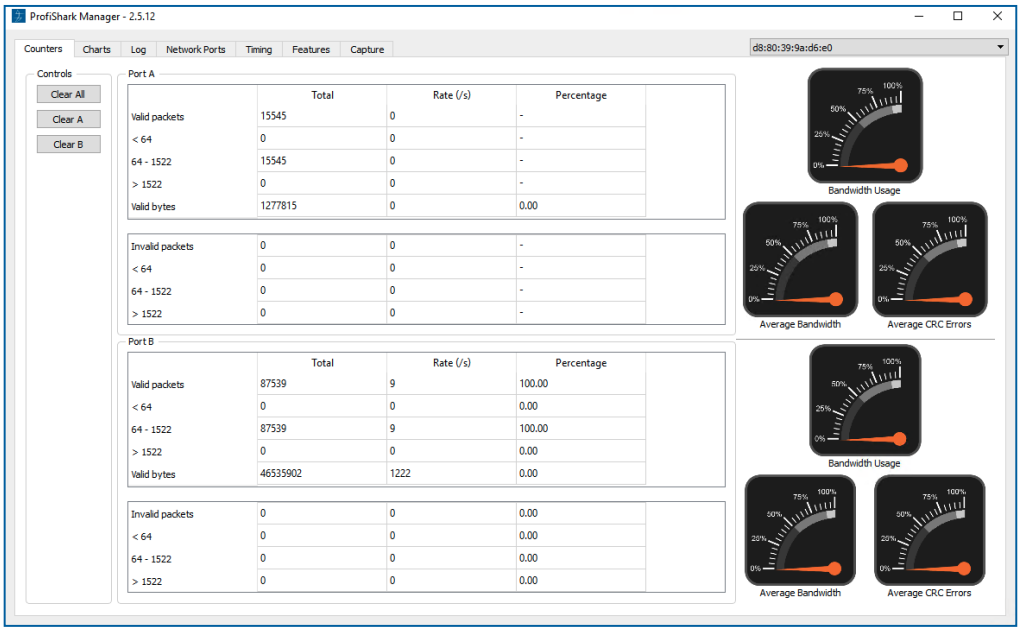

Charting:

Gives you a time line feel for flavors of frames and their rates.

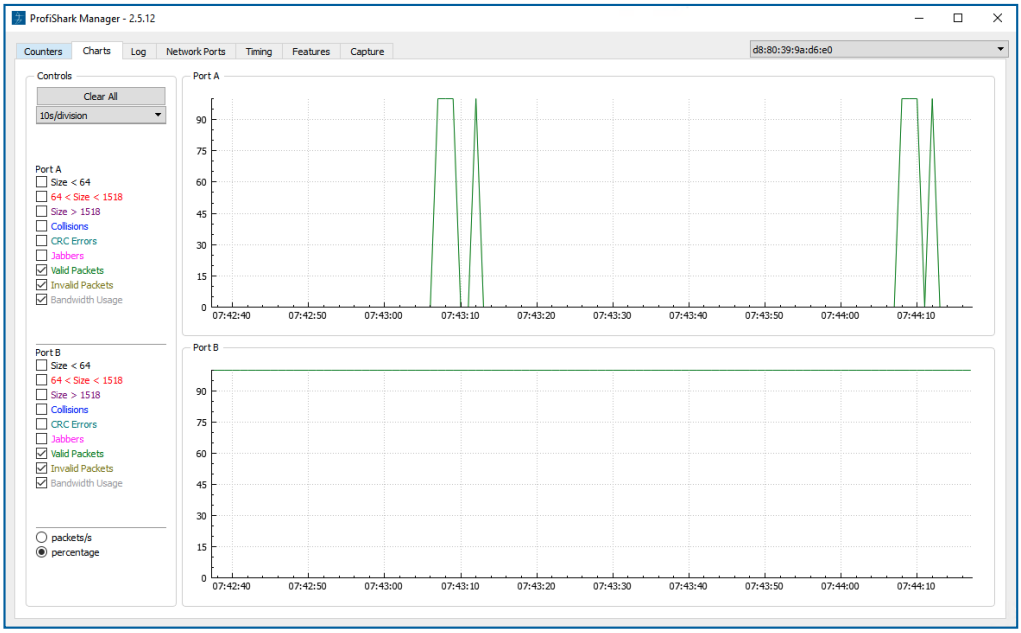

Logging:

Log Bandwidth and CRC events. I have found this useful when I want to log the precise date/time when broadcast storms flood the wire, for later correlation with other events. And to rule-out physical layer errors (no sign of CRCs, for example).

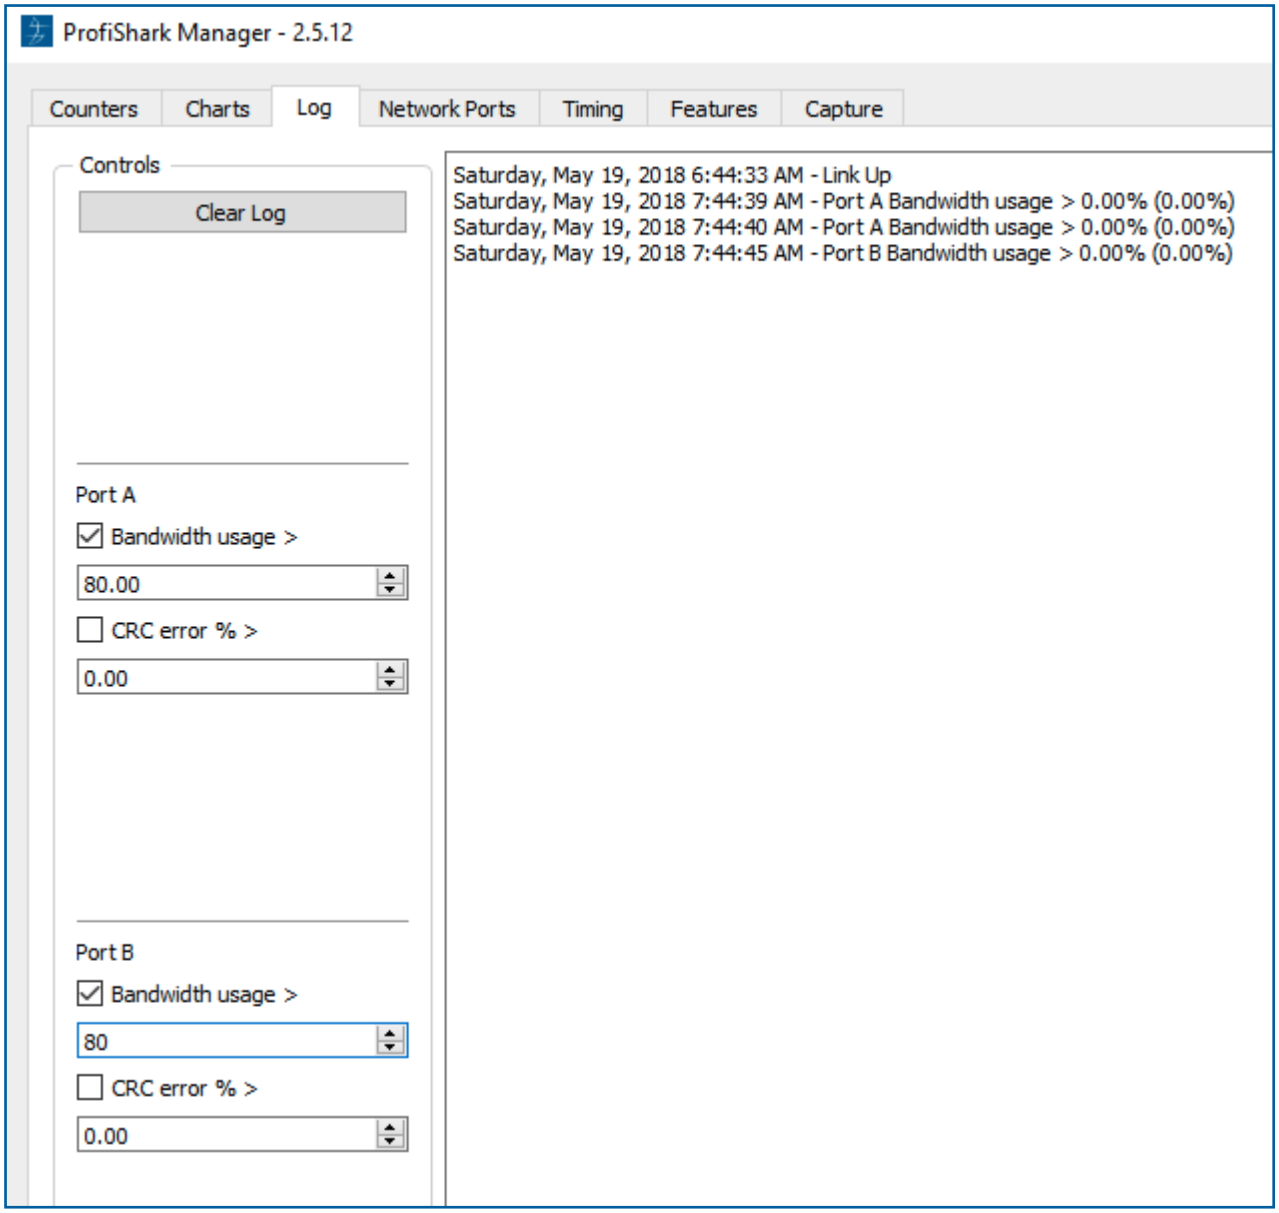

Ethernet Insights:

Glance at the Ethernet-level auto-negotiation parameters: quick way to identify the capabilities of the transceiver you've inserted, without having to Google for its manufacturer specs.

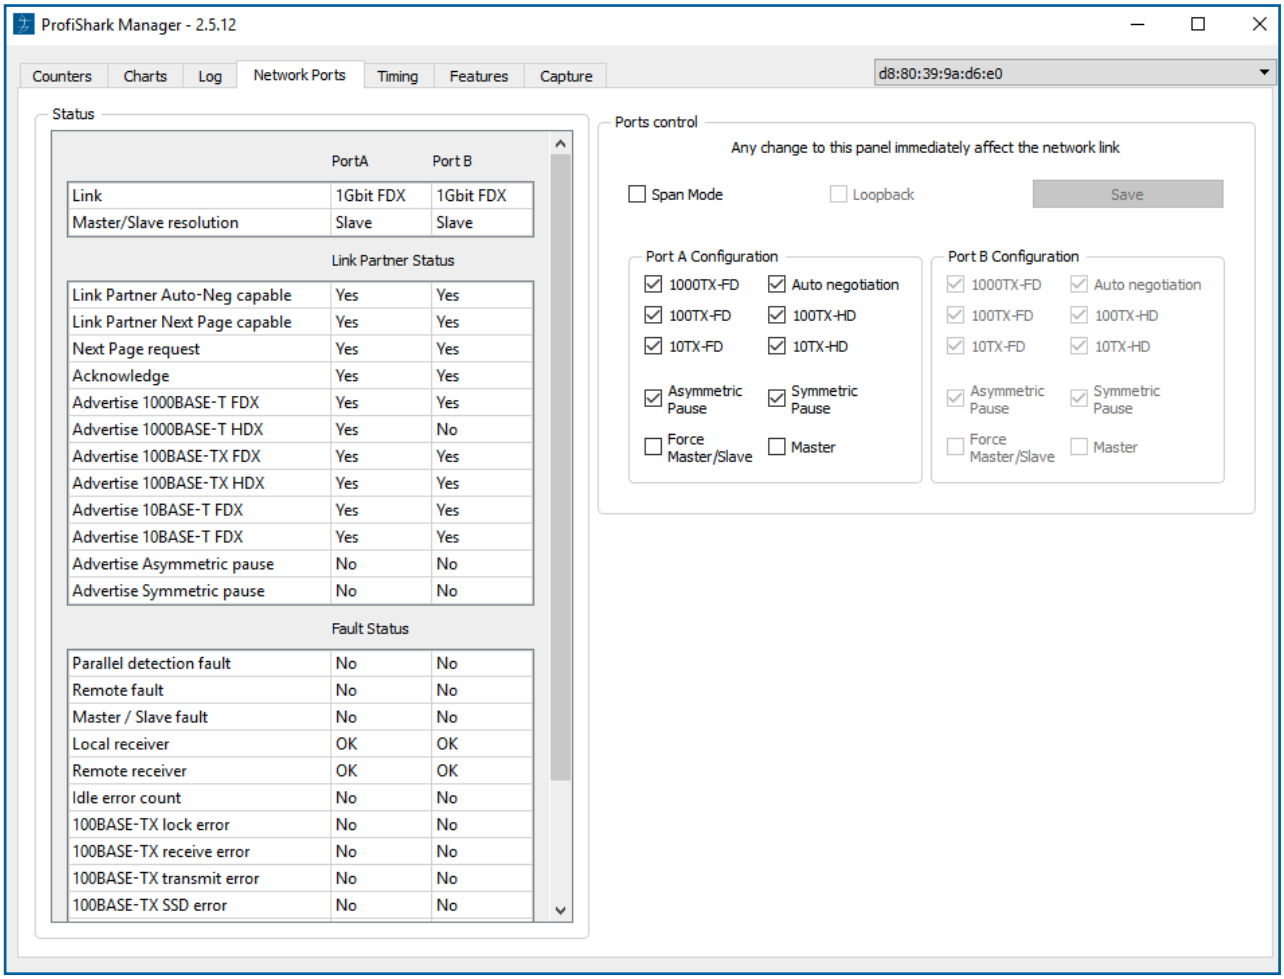

Real-Time Clock:

The '+' models offer a GPS-synchronized real-time clock, which provides highly accurate timestamps in your pcaps.

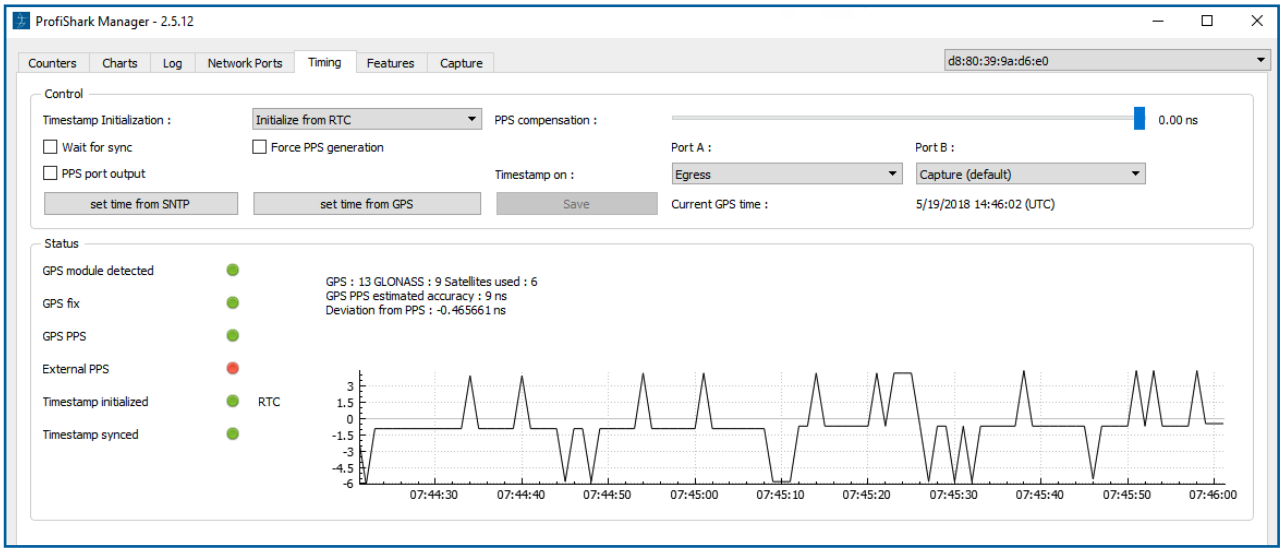

Features:

Optionally enable or disable hardware-level capture features.

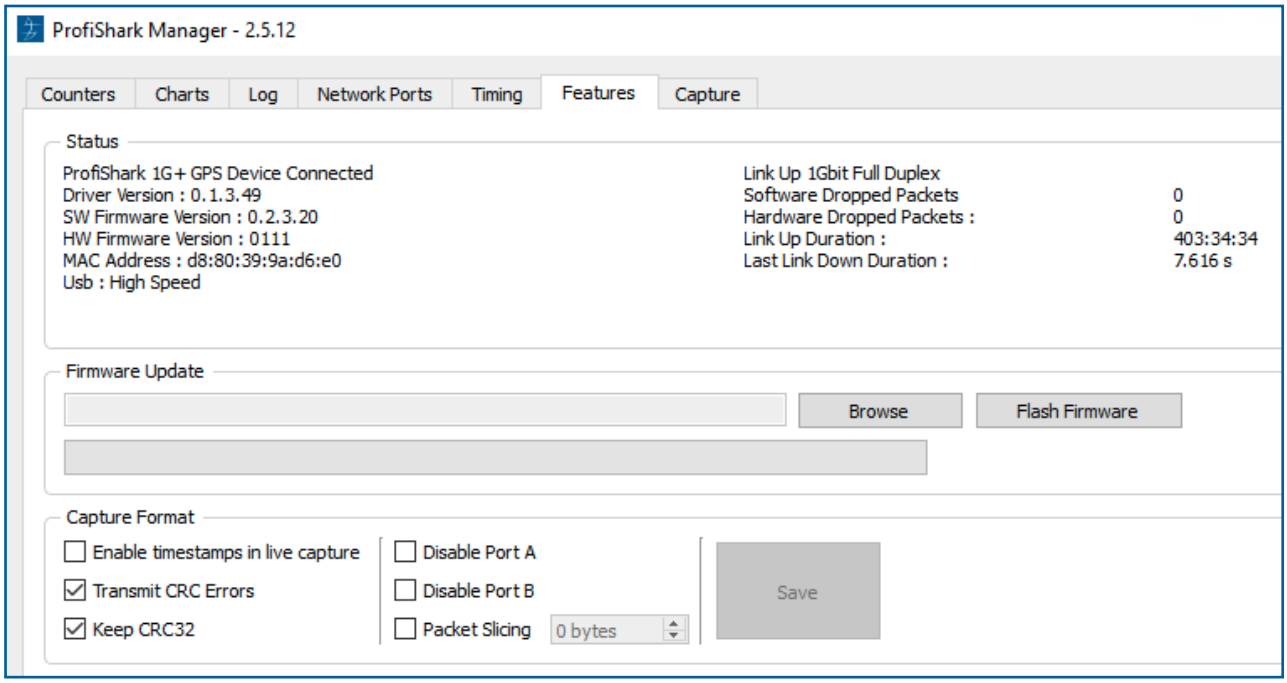

Captures:

And finally the Capture screen.

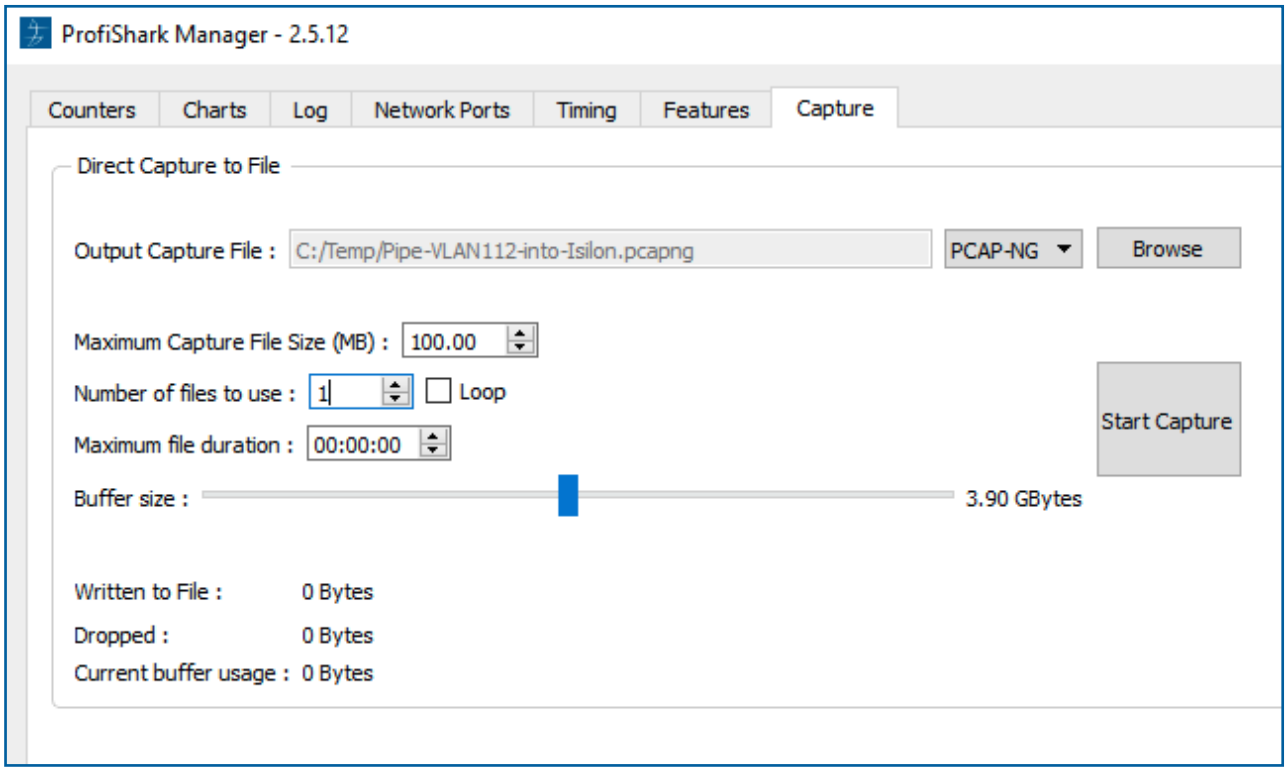

PROFISHARK IN ACTION

Isilon Row

Here we approach the row of Cabinets hosting the Isilon Storage System, with my laptop on a stool and the ProfiShark 10G barely visible above it. The Nexus 9372PX switches are minimally visible at the top of the Cabinets; the large blue LEDs mark the Generation 6 Isilon nodes; the Generation 5 nodes which populate most of these Cabinets aren't visible. And the small blue LEDs mark the vertical plug-strips providing power.

Isilon Cabinets

Walking closer to these Cabinets, we see the mix of Generation 5 and Generation 6 Isilon nodes, the two InfinBand switches which service the Cluster's back-side (those are fed by the bright blue and bright red power cords), plus a somewhat clearer view of the ProfiShark, sitting on top of a stack of Isilon nodes.

Laptop on Stool

Here we see the ProfiShark more clearly.

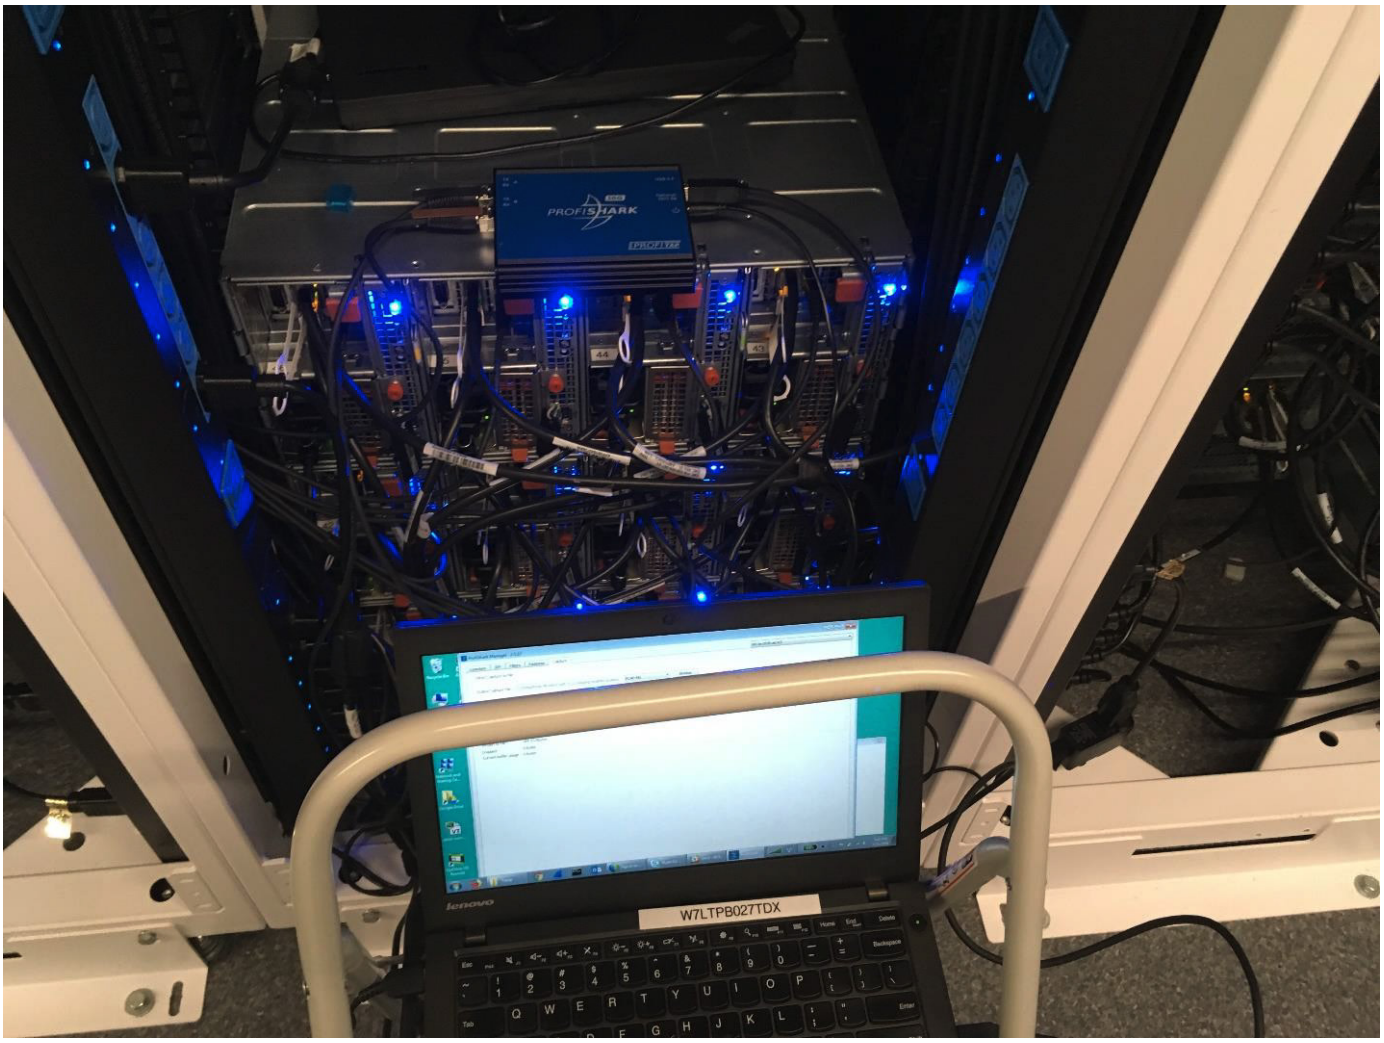

ProfiShark Close-up

And here we focus on the ProfiShark. Inserted into its left-hand side are the (2) 10G twinax cables which place it in-line with one of the Isilon Nodes below. On the right-hand side are inserted the USB cable connecting it to the laptop, along with a power cord attaching it to a wall-wart power supply. This power supply isn't a requirement – the ProfiShark will power itself from the USB link to the laptop. However, I wanted to be able to grab the laptop and walk away with it, leaving the ProfiShark behind. If I did that without first providing external power, then the ProfiShark would go dark and the Isilon Node would be disconnected from the network.

Laptop Closeup

You can use Wireshark to capture from the ProfiShark just fine – but I figure we are all familiar with Wireshark, so I would illustrate here the use of the dedicated ProfiShark Manager application, which also provides a Capture interface. If you have multiple ProfiShark units attached to your laptop, then you select the unit using the top right-hand drop down menu

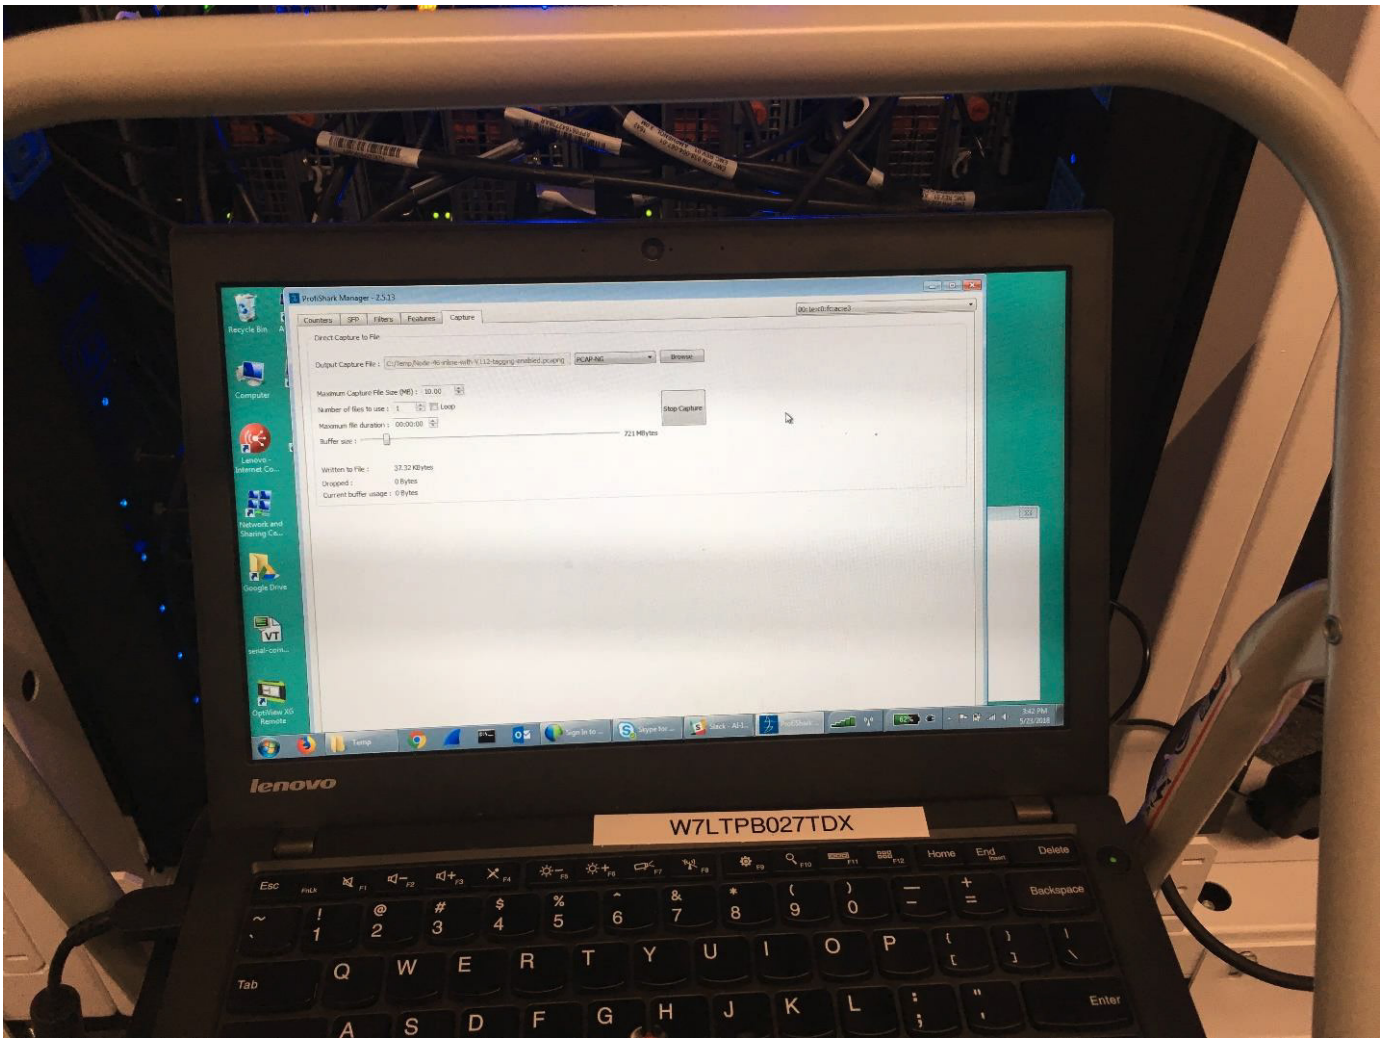

IT ALL STARTS WITH VISIBILITY

PROFITAP

Profitap develops a wide range of state-of-the-art and user-friendly network monitoring tools for both SMEs and the enterprise sector. Our wide range of high-density network TAPs, field service troubleshooters and network packet brokers are extremely performant, providing complete visibility and access to your network, 24/7.

We've been creating monitoring solutions for network analysis and traffic acquisition for more than 33 years. Therefore, we are experts in our field and our award-winning ProfiShark® 1G stands to prove it. This lightweight, advanced and portable network TAP is one the most innovative products on the market.

With more than 1,000 clients from 55 countries, PROFITAP has become a must-have solution for many important businesses, many of which are among Fortune 500 companies.

PROFITAP HQ B.V.
HIGH TECH CAMPUS 9
5656 AE EINDHOVEN
THE NETHERLANDS

sales@profitap.com
www.profitap.com

Profitap

@Profitap

profitap-international