

SHARKFEST 2015

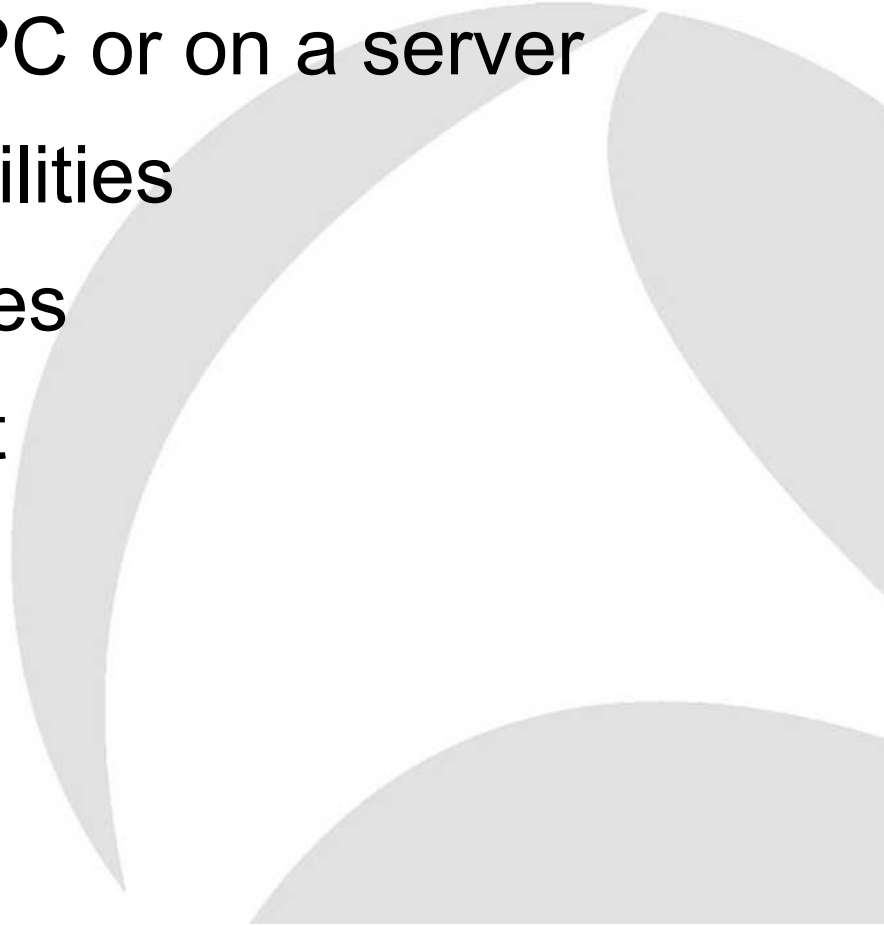
WIRESHARK DEVELOPER AND USER CONFERENCE



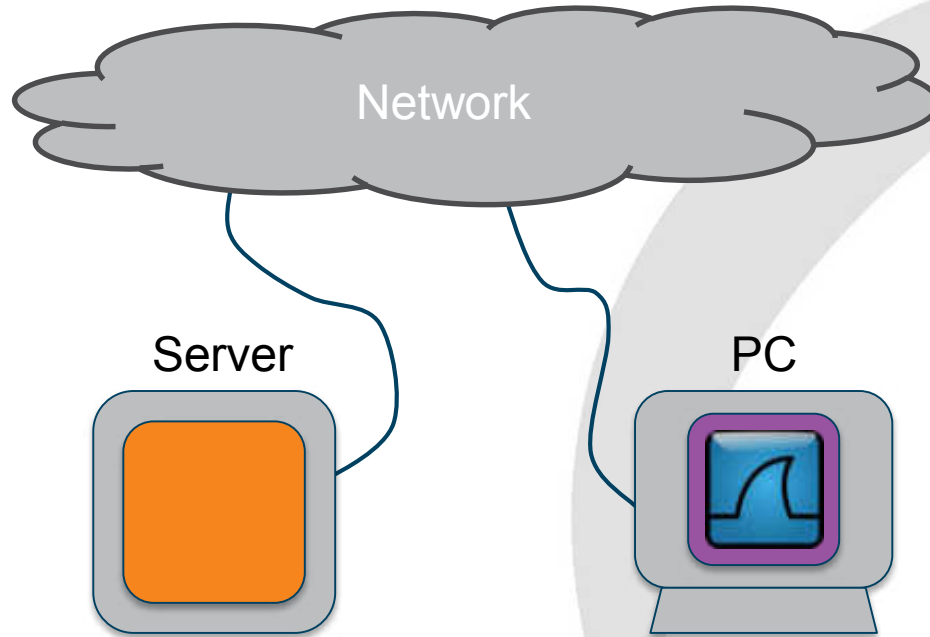
Packet Capture Techniques

Paul Offord, Advance7

Groups of capture techniques

- Directly from the user PC or on a server
 - Based on switch capabilities
 - Via purpose-built devices
 - In a virtual environment
- 

On The Client: Topology



Wireshark executables

Start Wireshark

| | | | | | | |
|-------------------|------|-----------|-----------|-------|-------------------------------|--|
| [-] iexplore.exe | 0.03 | 33,896 K | 36,776 K | 11664 | Internet Explorer | Microsoft Corporation |
| [-] iexplore.exe | 2.70 | 626,424 K | 588,384 K | 16260 | Internet Explorer | Microsoft Corporation |
| [+] firefox.exe | 5.14 | 654,512 K | 648,940 K | 14964 | Firefox | Mozilla Corporation |
| [-] OUTLOOK.EXE | 0.51 | 137,504 K | 203,964 K | 6036 | Microsoft Outlook | Microsoft Corporation |
| [-] WINWORD.EXE | 0.05 | 60,784 K | 56,312 K | 2352 | Microsoft Word | Microsoft Corporation |
| [-] POWERPNT.EXE | 0.04 | 149,784 K | 178,268 K | 3288 | Microsoft PowerPoint | Microsoft Corporation |
| [+] procexp.exe | | 3,040 K | 8,776 K | 17256 | Sysinternals Process Explorer | Sysinternals - www.sysinternals.com |
| [-] Wireshark.exe | 0.01 | 70,732 K | 75,820 K | 15680 | Wireshark | The Wireshark developer community, ... |
| [-] AppEx.exe | 0.36 | 2,968 K | 1,372 K | 5668 | Alps Printing-device Driver f | Alps Electric Co., Ltd. |

Start a capture

| | | | | | | |
|----------------------|--------|-----------|-----------|-------|-------------------------------|--|
| [-] iexplore.exe | 0.03 | 33,820 K | 36,784 K | 11664 | Internet Explorer | Microsoft Corporation |
| [-] iexplore.exe | 1.23 | 626,328 K | 588,448 K | 16260 | Internet Explorer | Microsoft Corporation |
| [+] firefox.exe | 2.35 | 658,132 K | 652,612 K | 14964 | Firefox | Mozilla Corporation |
| [-] OUTLOOK.EXE | 0.55 | 137,496 K | 203,952 K | 6036 | Microsoft Outlook | Microsoft Corporation |
| [-] WINWORD.EXE | < 0.01 | 60,784 K | 58,052 K | 2352 | Microsoft Word | Microsoft Corporation |
| [-] POWERPNT.EXE | 0.02 | 150,512 K | 179,052 K | 3288 | Microsoft PowerPoint | Microsoft Corporation |
| [+] procexp.exe | | 3,040 K | 8,776 K | 17256 | Sysinternals Process Explorer | Sysinternals - www.sysinternals.com |
| [-] Wireshark.exe | 0.44 | 86,508 K | 78,840 K | 15680 | Wireshark | The Wireshark developer community, ... |
| [-] dumpcap.exe | 0.01 | 4,660 K | 7,804 K | 9392 | Dumpcap | The Wireshark developer community |
| [-] SnippingTool.exe | 0.27 | 5,060 K | 14,472 K | 16796 | Snipping Tool | Microsoft Corporation |
| [-] AppEx.exe | 0.06 | 2,968 K | 1,372 K | 5668 | Alps Printing-device Driver f | Alps Electric Co., Ltd. |

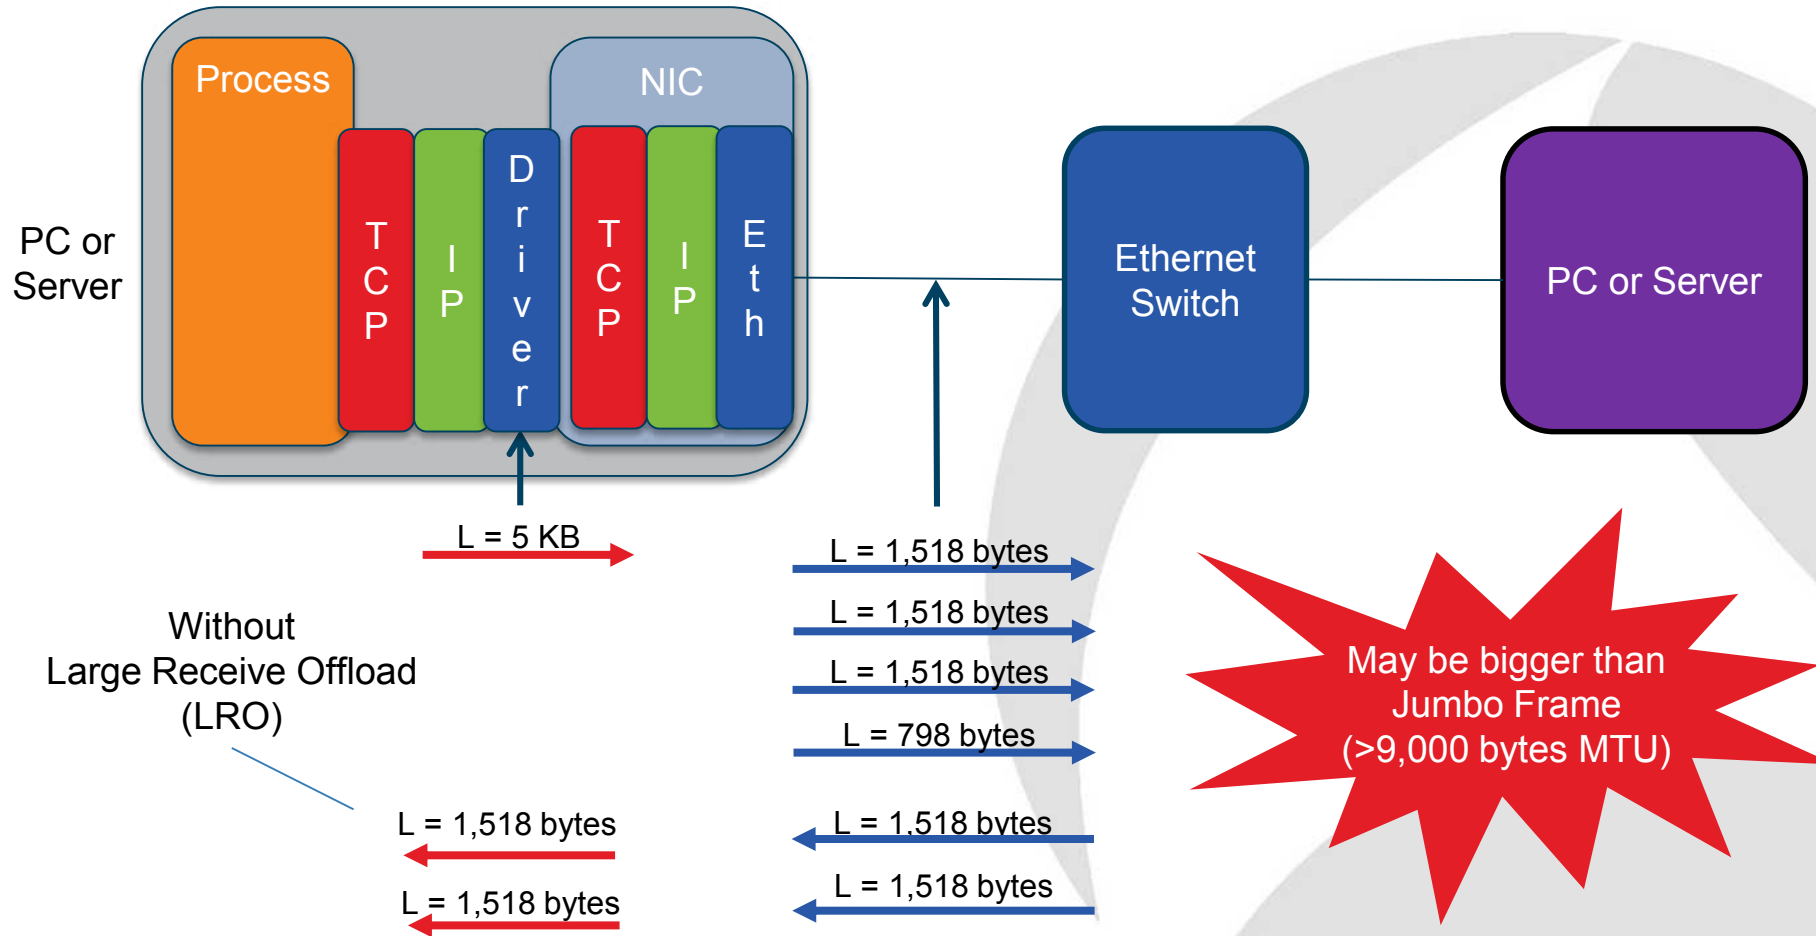
Very Large Frames

The image shows a Wireshark network traffic capture window. The main pane displays a list of network packets. Two packets are highlighted with red circles: packet 4342 and packet 457. Packet 4342 is a TCP segment of a reassembled PDU with a length of 4342 bytes. Packet 457 is an application data packet with a length of 17206 bytes. The status bar at the bottom indicates that frame 457 is 17206 bytes on wire (137648 bits), 17206 bytes captured (137648 bits). The status bar also shows the source and destination of the frame: Ethernet II, Src: Vmware_Bb:62:1e (00:50:56:8b:62:1e), Dst: All-HSRP-routers_00 (00:00:0c:07:ac:00).

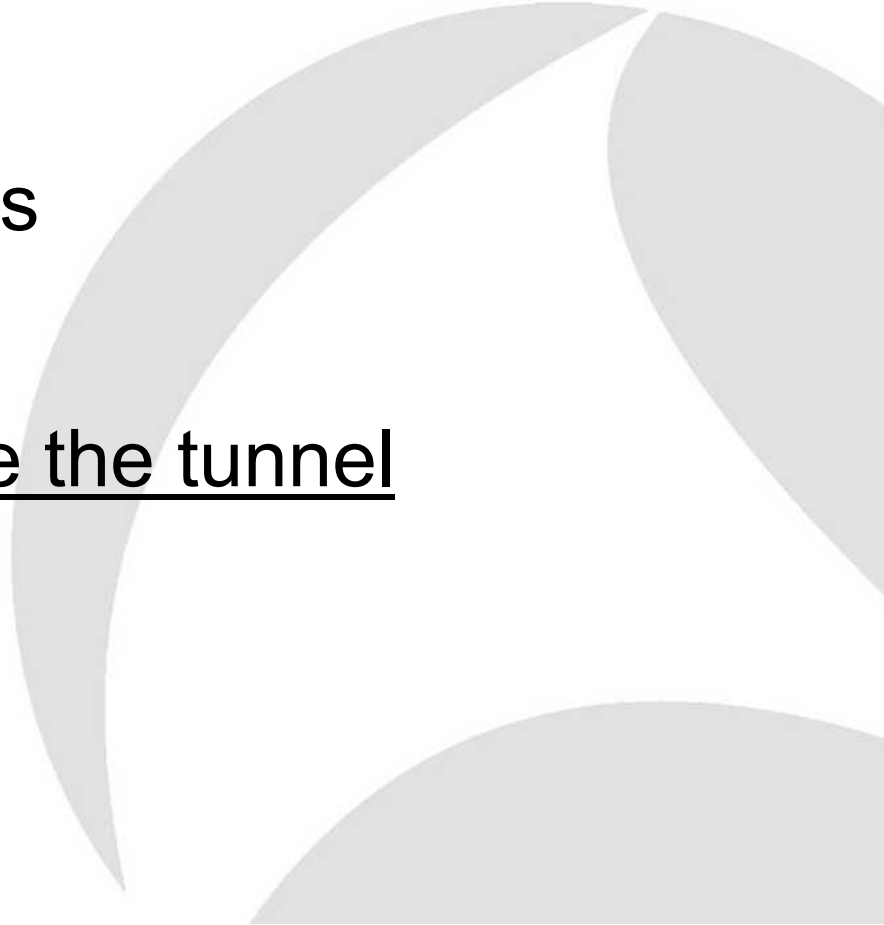
| No. | Time | Delta | Source | Destination | Protocol | Src Port | Dst Port | Length | Info |
|-----|-----------------|----------|---------------|---------------|----------|----------|----------|--------|--|
| 410 | 14:25:08.441961 | 0.000000 | 10.160.58.121 | 10.60.25.175 | TLSv1 | 53861 | 443 | 590 | Application Data |
| 411 | 14:25:08.441983 | 0.000022 | 10.160.58.121 | 10.60.25.175 | TLSv1 | 53861 | 443 | 376 | Application Data |
| 412 | 14:25:08.442037 | 0.000054 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 54 | https-53861 [ACK] Seq=4049927767 Ack=2437570878 Win=512 Len=0 |
| 426 | 14:25:08.620193 | 0.178156 | 10.60.25.175 | 10.160.58.121 | TLSv1 | 443 | 53861 | 507 | Application Data |
| 438 | 14:25:08.817954 | 0.197761 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437570878 Ack=4049928240 Win=16368 Len=0 |
| 442 | 14:25:09.121080 | 0.303126 | 10.160.58.121 | 10.60.25.175 | TLSv1 | 53861 | 443 | 590 | Application Data |
| 443 | 14:25:09.121084 | 0.000004 | 10.160.58.121 | 10.60.25.175 | TLSv1 | 53861 | 443 | 328 | Application Data |
| 444 | 14:25:09.121159 | 0.000075 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 54 | https-53861 [ACK] Seq=4049928240 Ack=2437571688 Win=512 Len=0 |
| 445 | 14:25:09.125497 | 0.004338 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 4342 | TCP segment of a reassembled PDU |
| 446 | 14:25:09.126419 | 0.000922 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049932528 Win=16482 Len=0 |
| 447 | 14:25:09.126452 | 0.000033 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 6486 | [TCP segment of a reassembled PDU] |
| 448 | 14:25:09.126998 | 0.000546 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049938424 Win=16482 Len=0 |
| 449 | 14:25:09.127034 | 0.000016 | 10.60.25.175 | 10.160.58.121 | TLSv1 | 443 | 53861 | 8094 | Application Data |
| 450 | 14:25:09.127304 | 0.000290 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049939496 Win=16482 Len=0 |
| 451 | 14:25:09.127318 | 0.000014 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 2198 | [TCP segment of a reassembled PDU] |
| 452 | 14:25:09.127517 | 0.000199 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049944856 Win=16482 Len=0 |
| 453 | 14:25:09.127532 | 0.000015 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 7558 | [TCP segment of a reassembled PDU] |
| 454 | 14:25:09.128243 | 0.000711 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049950216 Win=16482 Len=0 |
| 455 | 14:25:09.128245 | 0.000002 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049955576 Win=15812 Len=0 |
| 456 | 14:25:09.128247 | 0.000002 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049966648 Win=15544 Len=0 |
| 457 | 14:25:09.128271 | 0.000024 | 10.60.25.175 | 10.160.58.121 | TLSv1 | 443 | 53861 | 17206 | Application Data |
| 458 | 14:25:09.129323 | 0.003052 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049962008 Win=14204 Len=0 |
| 459 | 14:25:09.129349 | 0.000026 | 10.60.25.175 | 10.160.58.121 | TLSv1 | 443 | 53861 | 7558 | Application Data |
| 460 | 14:25:09.130812 | 0.001463 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049967368 Win=12864 Len=0 |
| 461 | 14:25:09.130851 | 0.000039 | 10.60.25.175 | 10.160.58.121 | TCP | 443 | 53861 | 7558 | [TCP segment of a reassembled PDU] |
| 462 | 14:25:09.131414 | 0.000563 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049972728 Win=11524 Len=0 |
| 463 | 14:25:09.131432 | 0.000018 | 10.60.25.175 | 10.160.58.121 | TLSv1 | 443 | 53861 | 7558 | Application Data |
| 464 | 14:25:09.132201 | 0.000769 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049978088 Win=10184 Len=0 |
| 465 | 14:25:09.132203 | 0.000002 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049983448 Win=8844 Len=0 |
| 466 | 14:25:09.132206 | 0.000003 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049988808 Win=7504 Len=0 |
| 467 | 14:25:09.132308 | 0.000102 | 10.160.58.121 | 10.60.25.175 | TCP | 53861 | 443 | 60 | 53861->https [ACK] Seq=2437571688 Ack=4049994168 Win=6164 Len=0 |

Frame 457: 17206 bytes on wire (137648 bits), 17206 bytes captured (137648 bits)
Ethernet II, Src: Vmware_Bb:62:1e (00:50:56:8b:62:1e), Dst: All-HSRP-routers_00 (00:00:0c:07:ac:00)

TCP Segmentation Offload




On The Client: Advantages

- Easy to achieve
 - Zero disruption to services
 - Capture wireless traffic
 - Capture VPN traffic inside the tunnel
- 


On The Client: Considerations

- TCP Seg. Offload can be confusing
- Disk contention may cause lost packets
- Potential performance hit when saving to C:
 - Page files, EXEs, DLL, Memory Mapped Files
 - Consider USB drive
- Use dumpcap for long-term captures


On The Server: Topology



Discovering unknown interactions



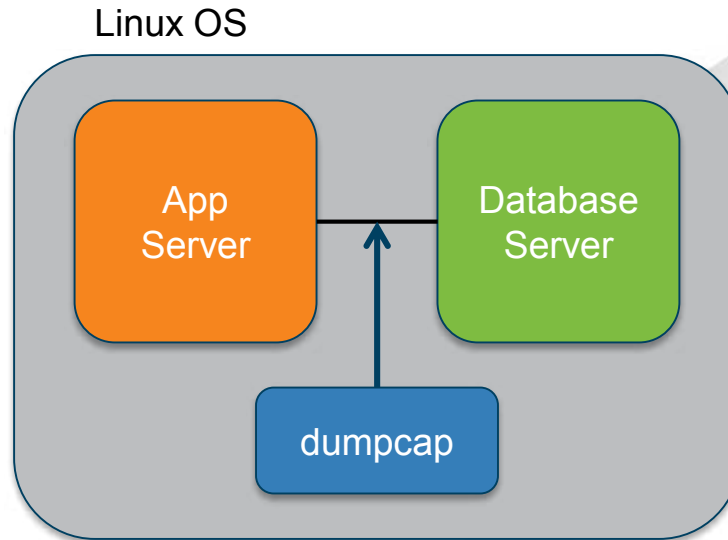
On The Server: Advantages

- Relatively easy to achieve
 - Minimal disruption to services
 - Change Request probably needed
 - All client traffic visible
 - All interactions with other services visible
 - Blade and VM east-west traffic visible
- 

On The Server: Considerations

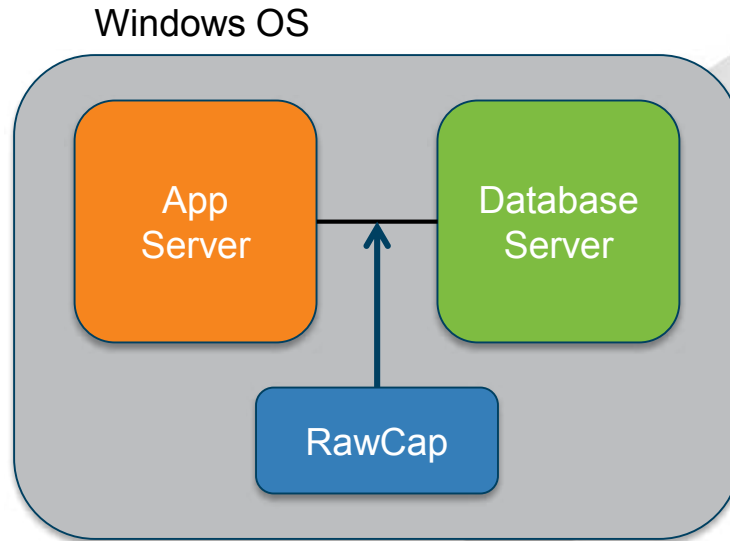
- TCP Seg. Offload can be confusing
- Volume of data higher than client-side capture
- Save to a dedicated volume
 - Not to C: drive, database log vols, etc.
 - USB drives work well
- Use dumpcap not tshark or Wireshark
- Care needed when teaming used
- Intra-OS tracing not possible on Windows
 - Loopback adapter not the same as Linux

Via loopback



Source and destination IP address is 127.0.0.1
– use TCP port number to determine packet direction

RawCap



Produces PCAP files


IPv4 only

Windows 2008 r2 onwards

**Time for
Questions**



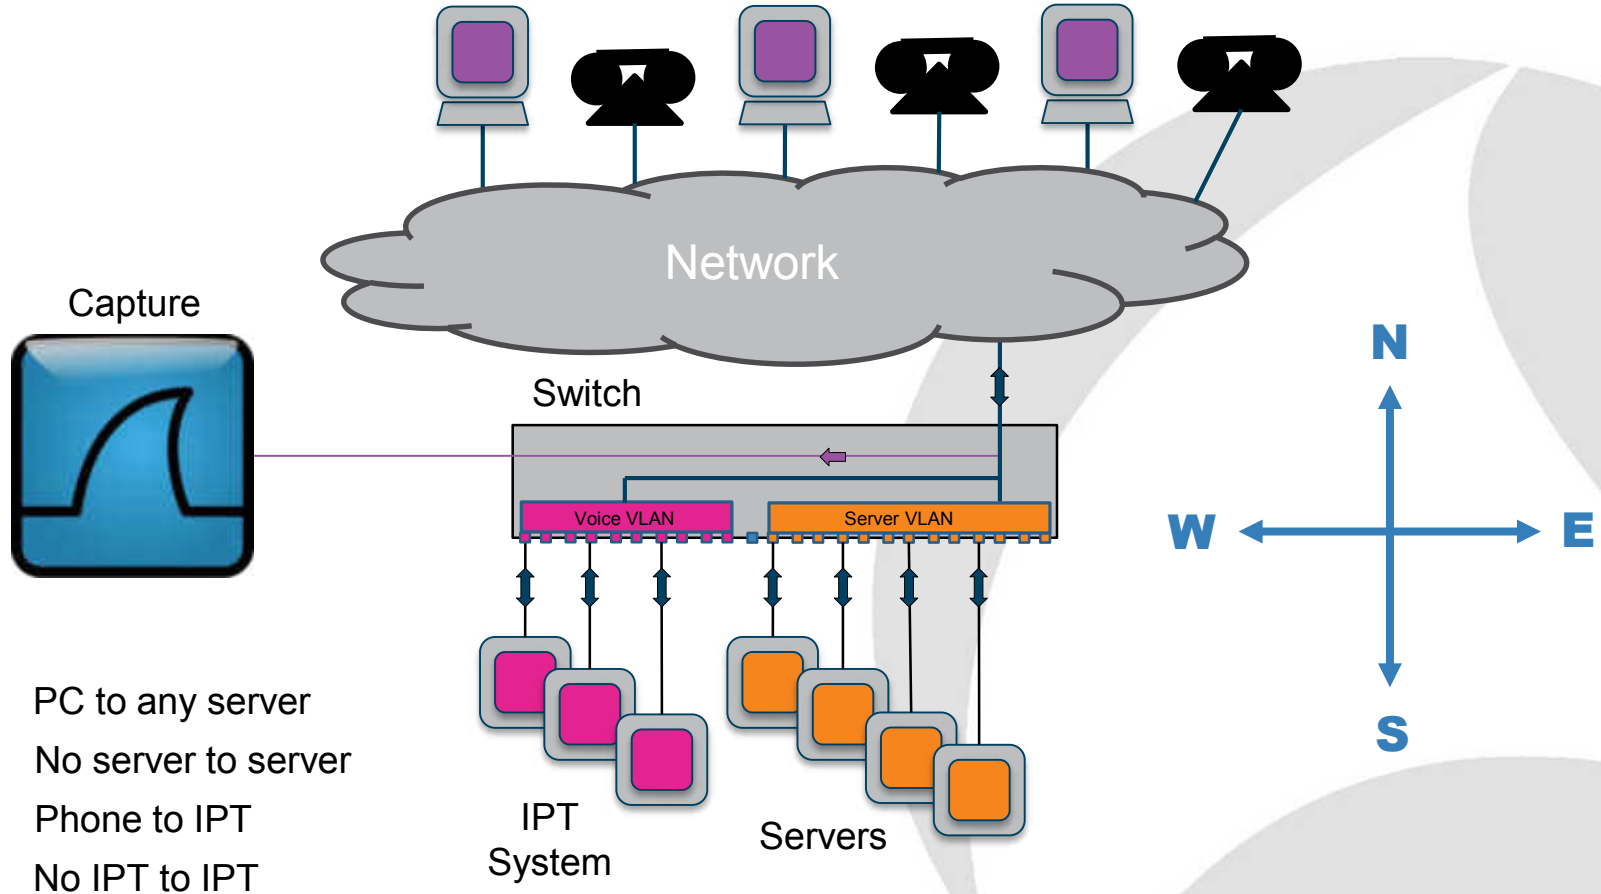
SPAN-Monitor-Mirror: Topology




Cisco – **monitor session ...**

Juniper - **set ethernet-switching options analyzer ...**

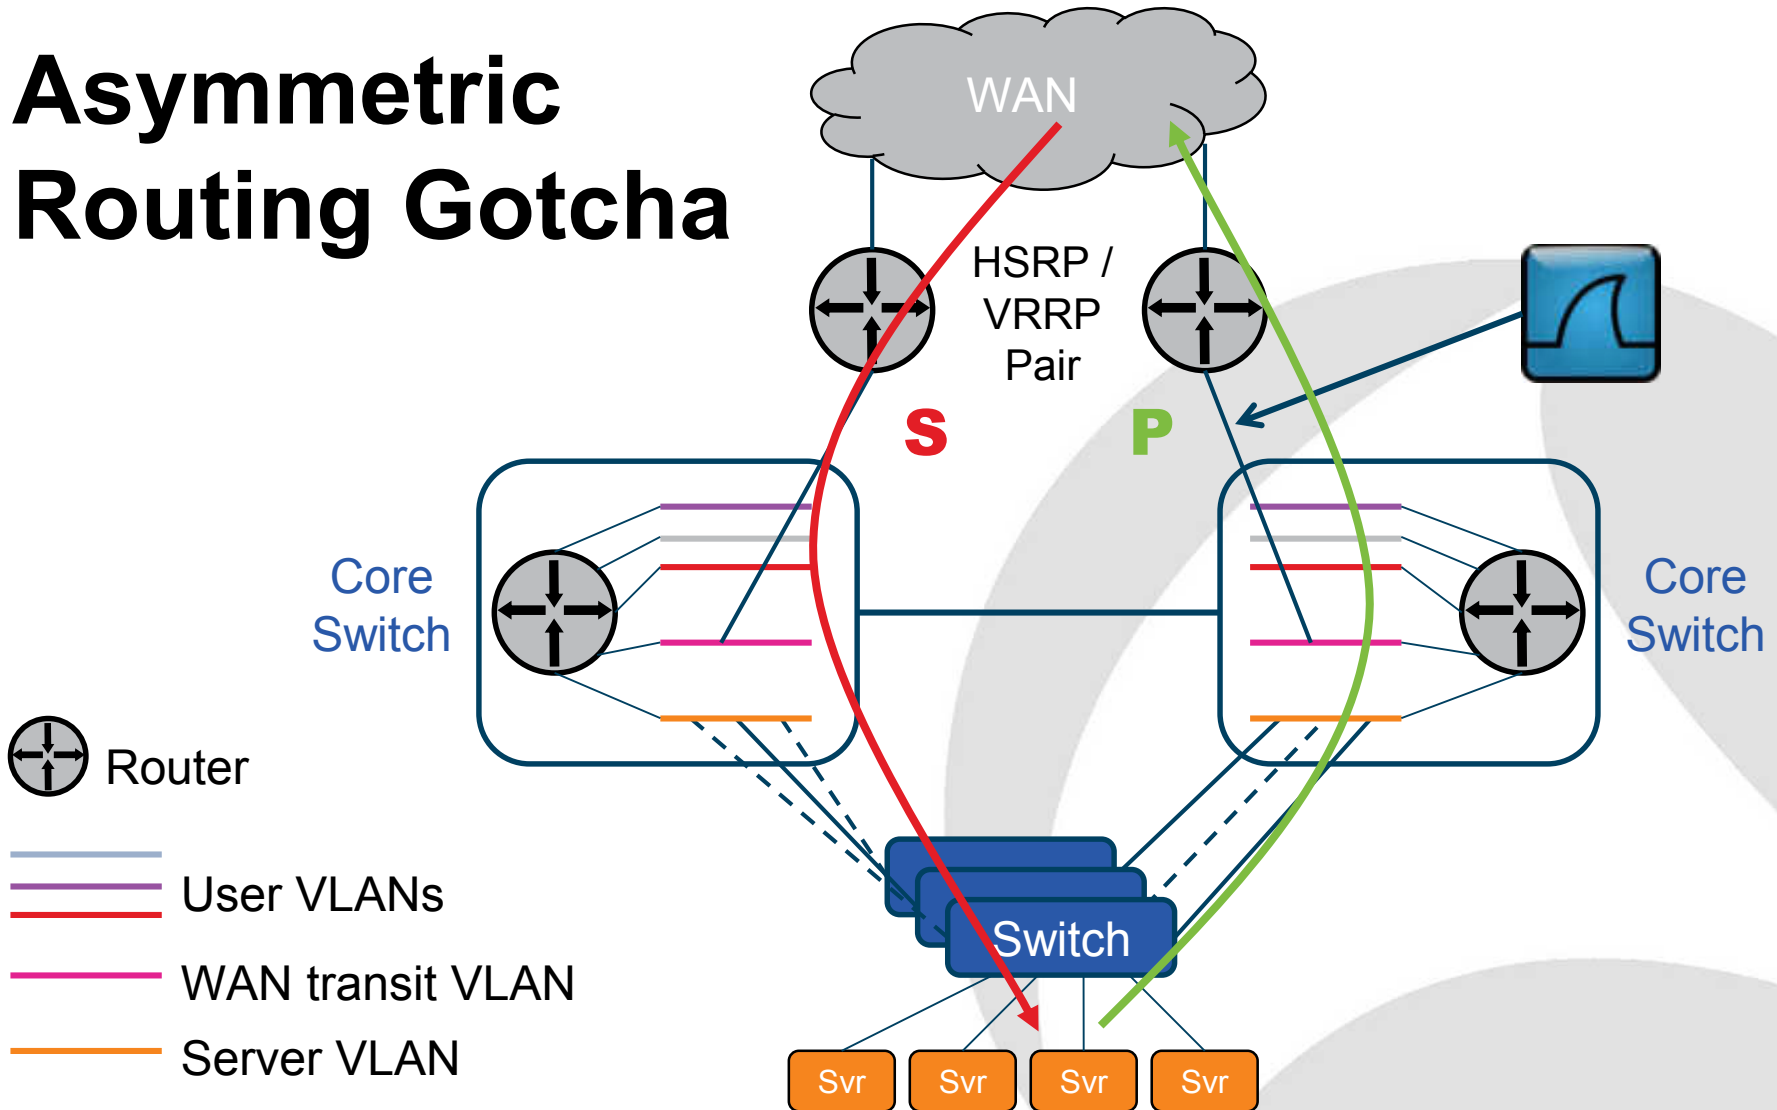
SPAN-Monitor-Mirror: Uplink



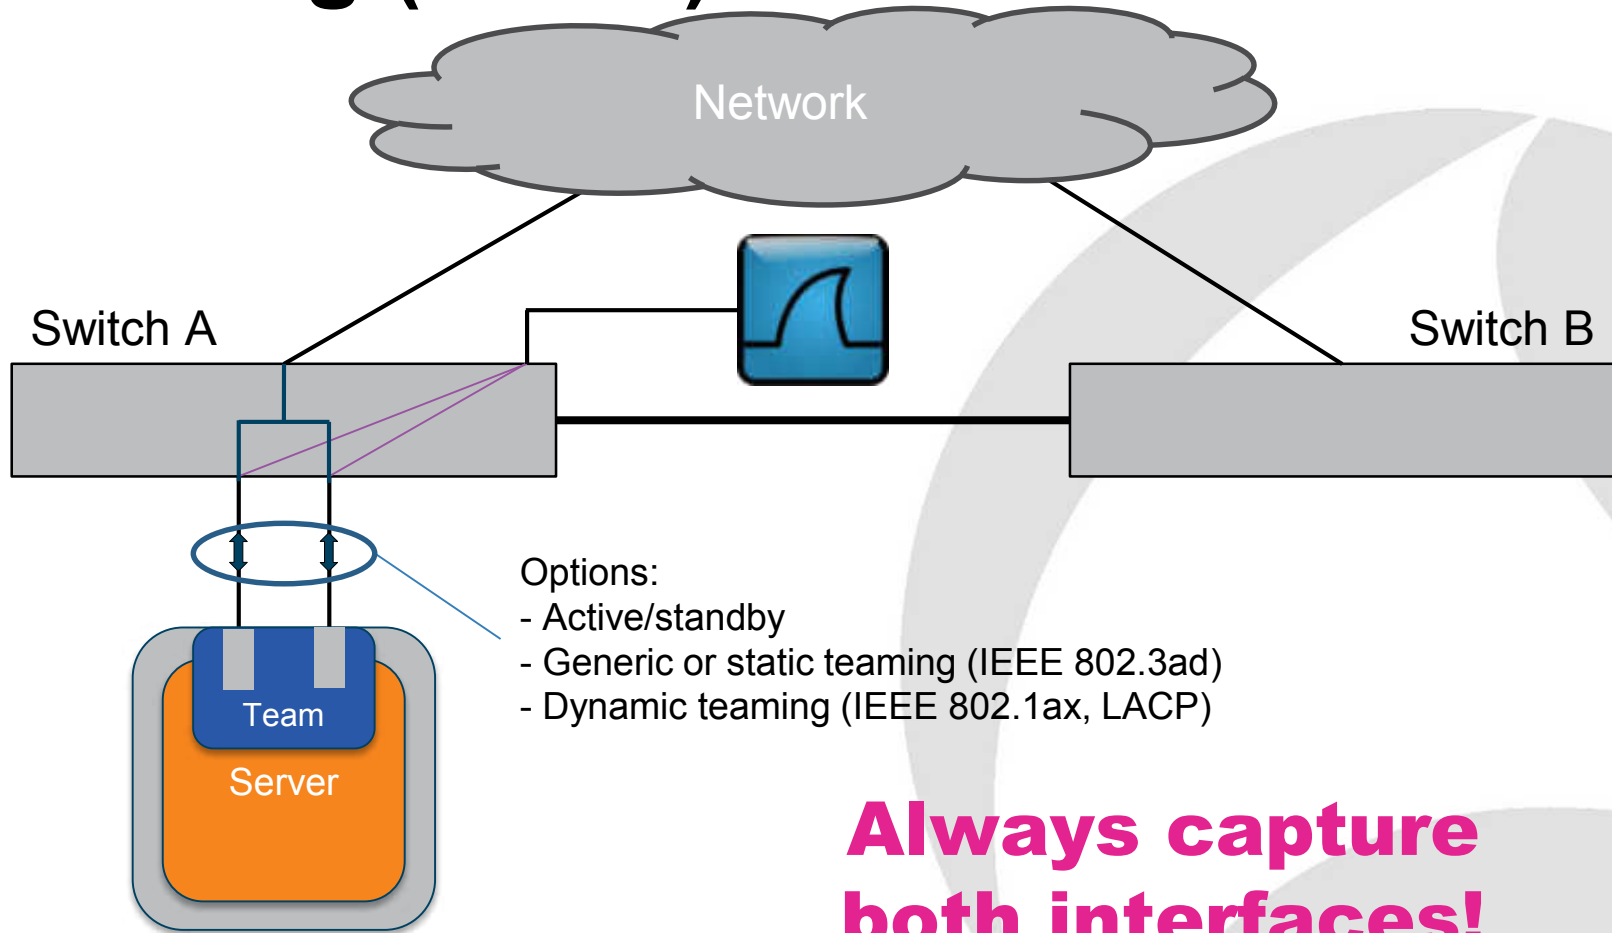
SPAN-Monitor-Mirror: VLAN



Asymmetric Routing Gotcha

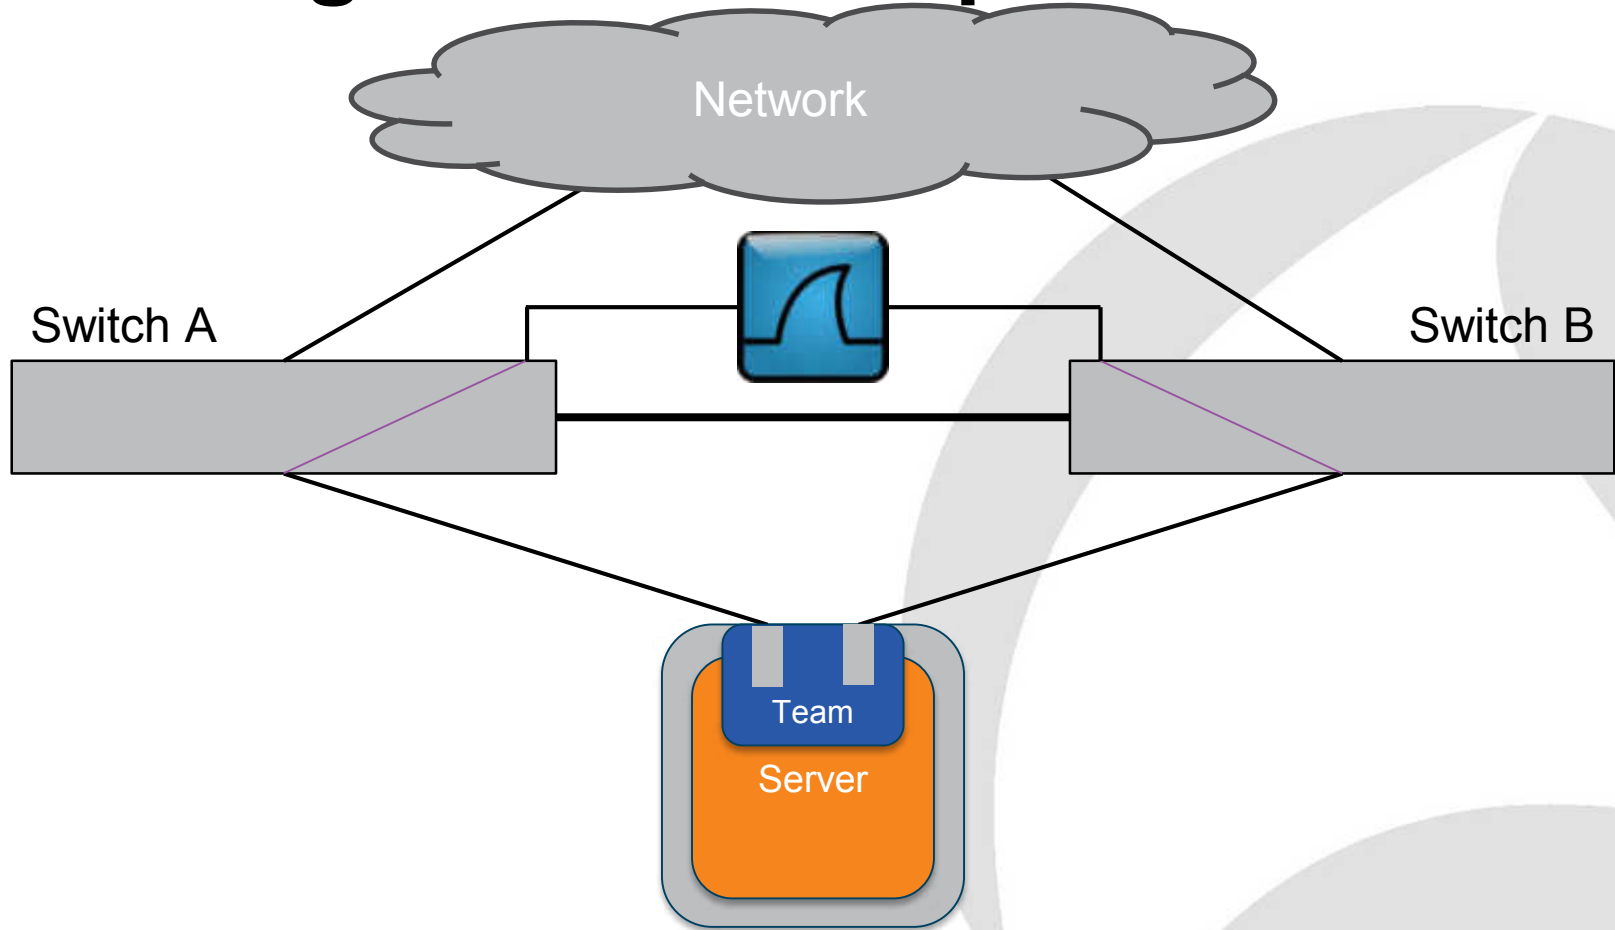


Teaming (LBFO)




**Always capture
both interfaces!**

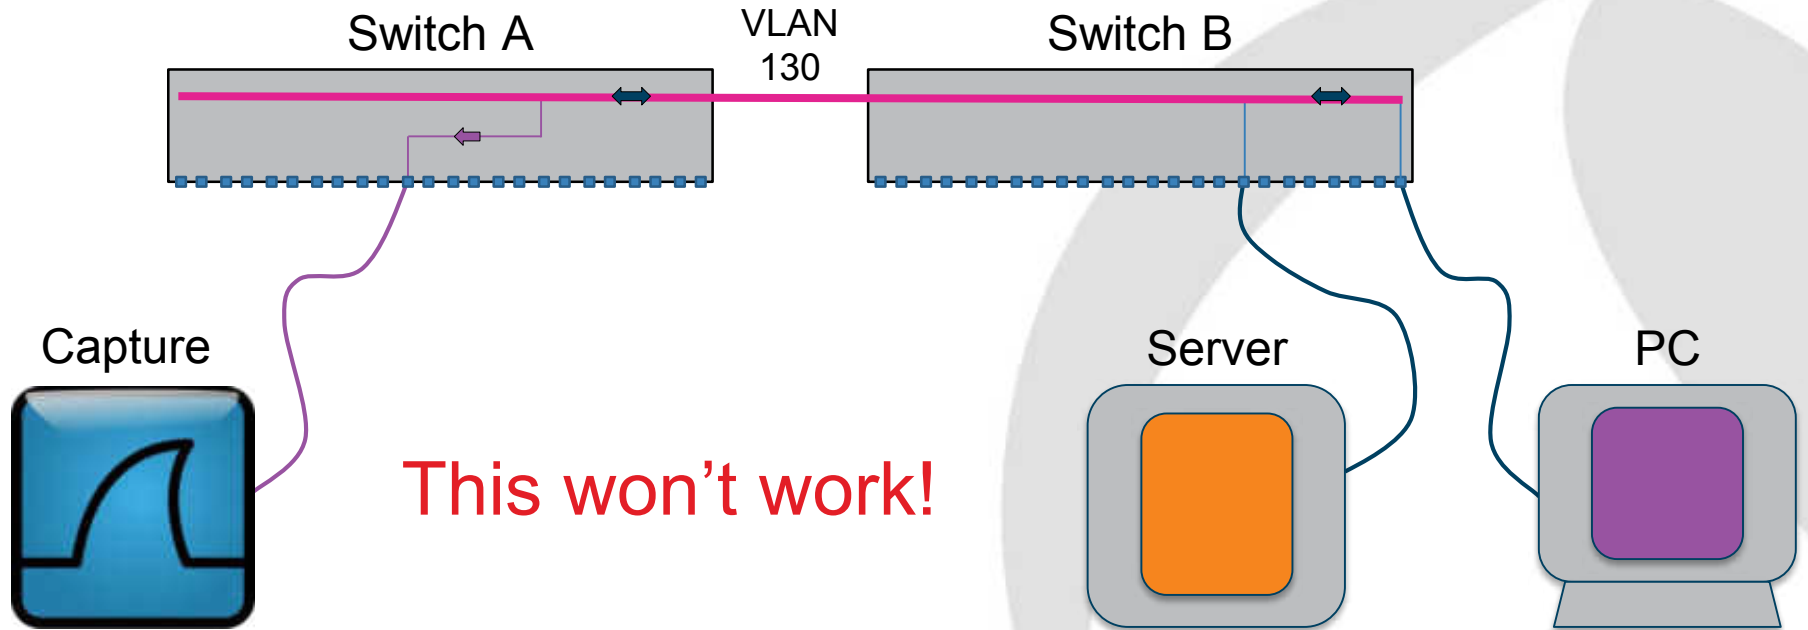
Teaming: Switch independent mode



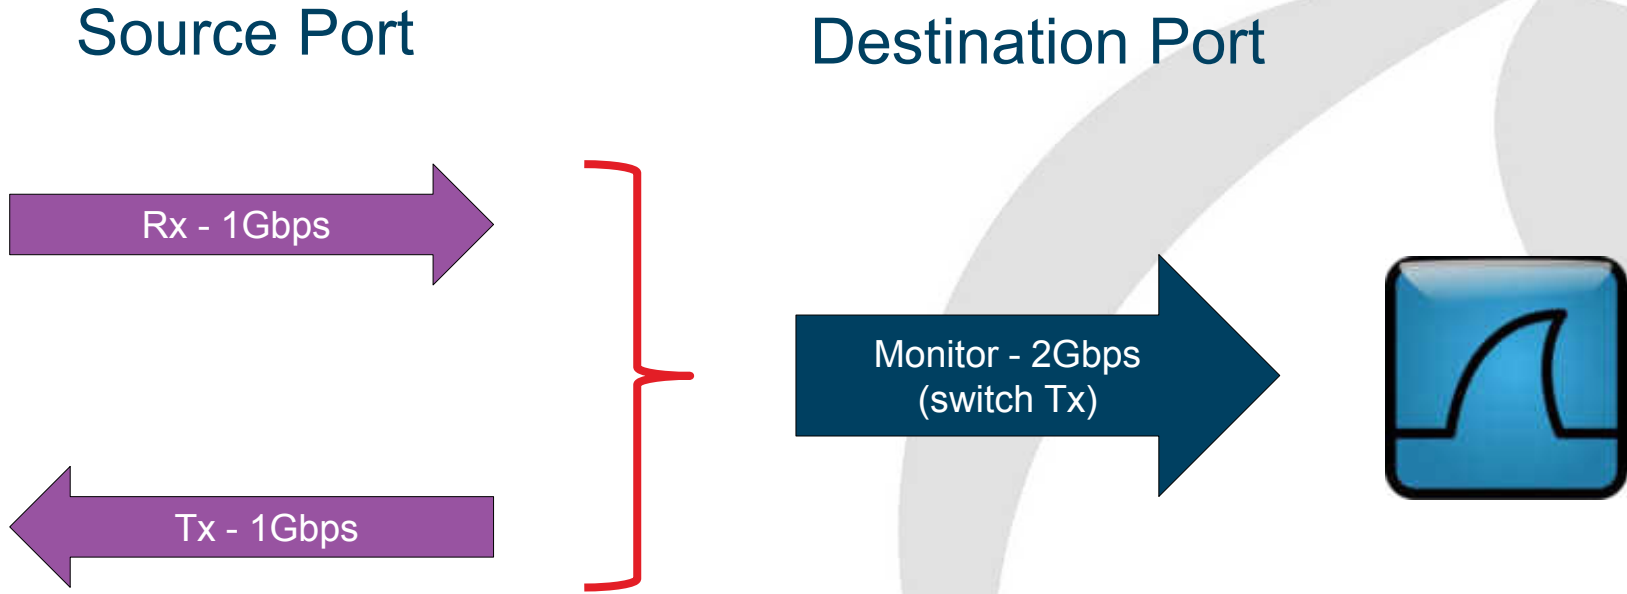
SPAN-Monitor-Mirror: Advantages

- Easy to configure
 - Low risk – non-invasive
 - Multiple sources into one destination
 - Entire VLANs can be monitored
 - Need to monitor on each switch
 - May see duplicates
 - Negligible impact on the switch
- 

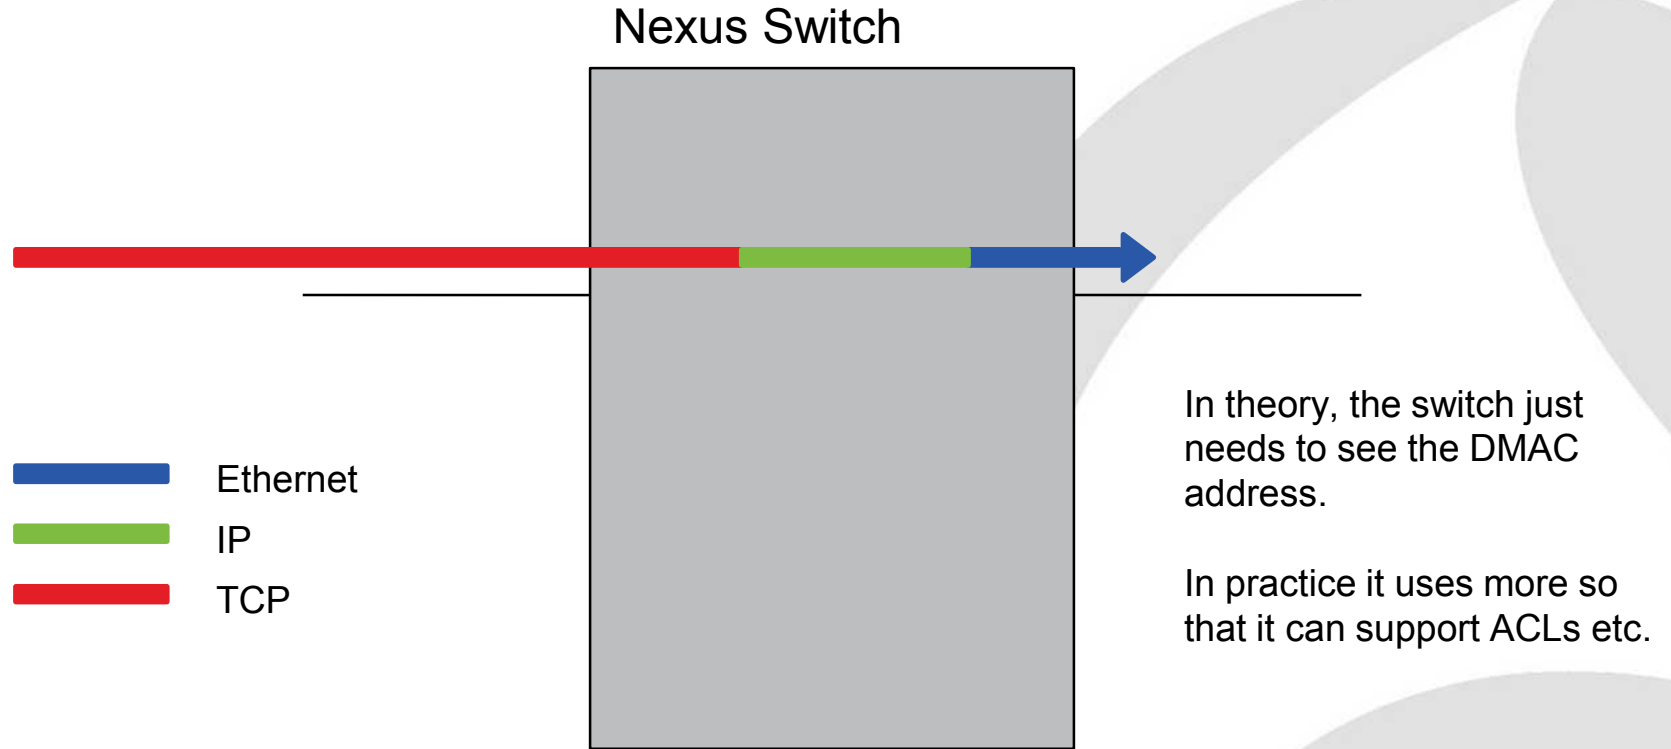
SPAN-Monitor-Mirror: VLAN Gotcha



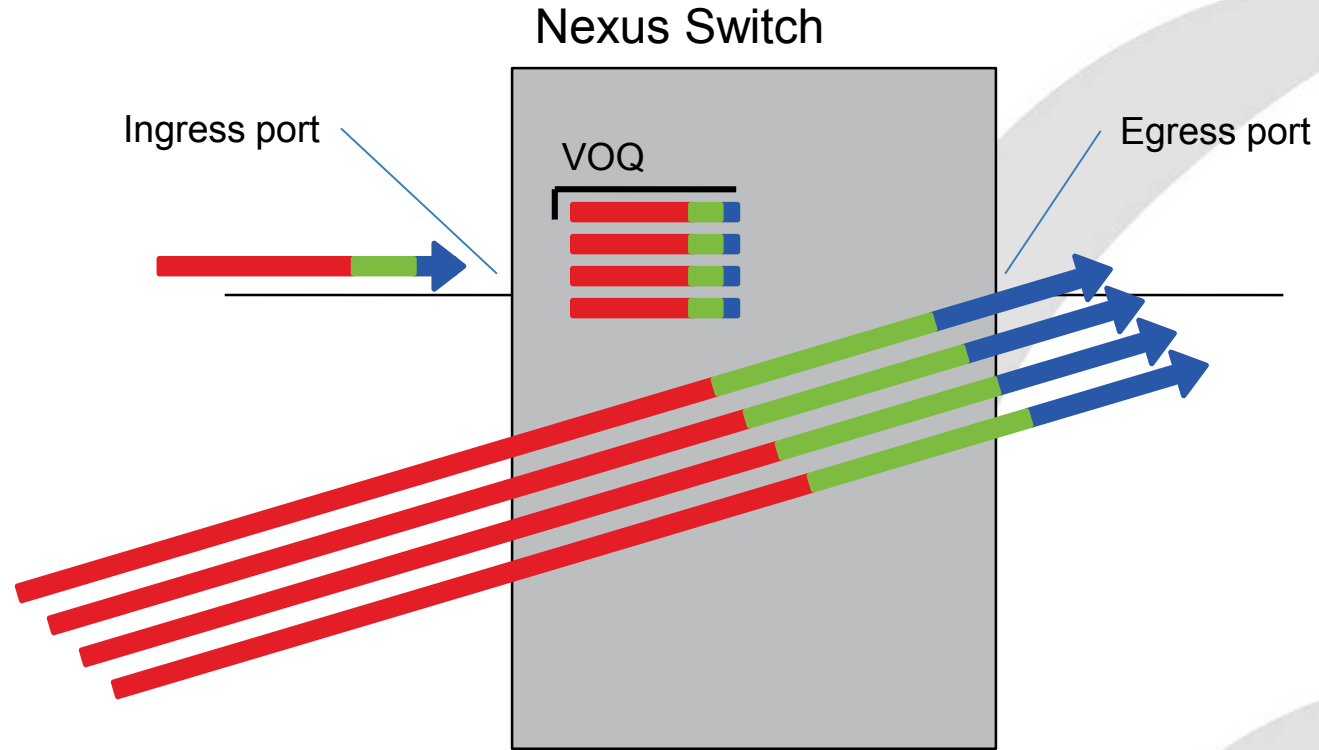
SPAN-Monitor-Mirror: 2-into-1



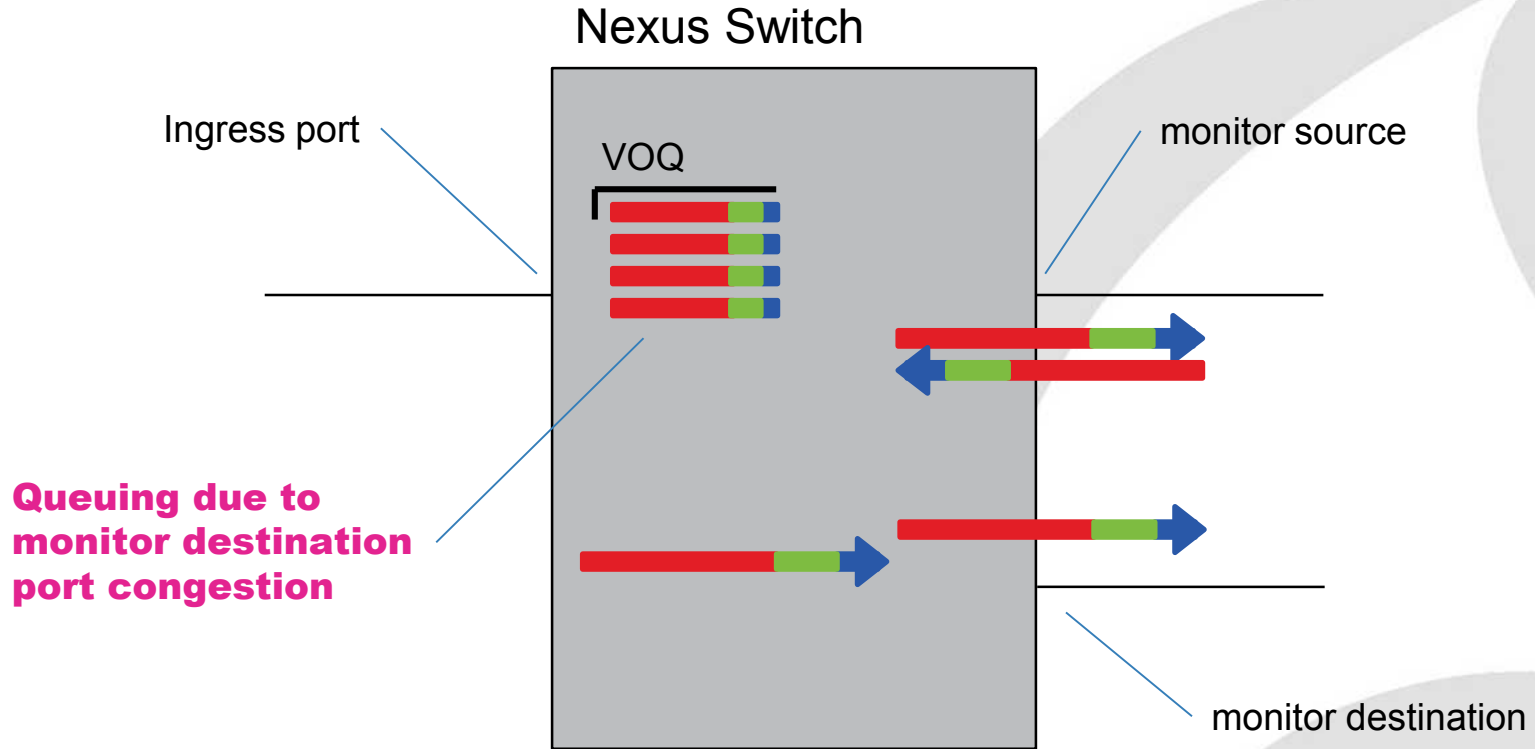
Cut-through switches



Virtual Output Queues



Cisco Nexus Back-pressure Gotcha



SPAN Rate Limiting

Configuring the Rate Limit for SPAN Traffic

By configuring a rate limit for SPAN traffic to 1Gbps across the entire monitor session, you can avoid **impacting the monitored production traffic**. For Nexus 5000 series switches:

- When spanning more than 1Gbps to a 1 Gb SPAN destination interface, SPAN source traffic will not drop.
- When spanning more than 6 Gbps (but less than 10Gbps) to a 10Gb SPAN destination interface, the SPAN traffic is limited to 1Gbps even though the destination/sniffer is capable of 10Gbps.

On the Nexus 5500 series, SPAN traffic is rate-limited to 1Gbps by default so the **switchport monitor rate-limit 1G** interface command is not supported. Also, to avoid **impacting monitored production traffic**:

- SPAN is rate-limited to 5 Gbps for every 8 ports (one ASIC).
- RX-SPAN is rate-limited to 0.71 Gbps per port when the RX-traffic is mirrored to a destination interface.



Different rules and
commands for
Nexus 7000


SPAN-Monitor-Mirror: Considerations

- Overload of the monitor destination
- Back-pressure on source port (Cisco Nexus)
 - Alleviated using source rate limiting
- Limited number of monitor sessions
- Requires a spare switch port for destination
- Makes and models vary -> review first

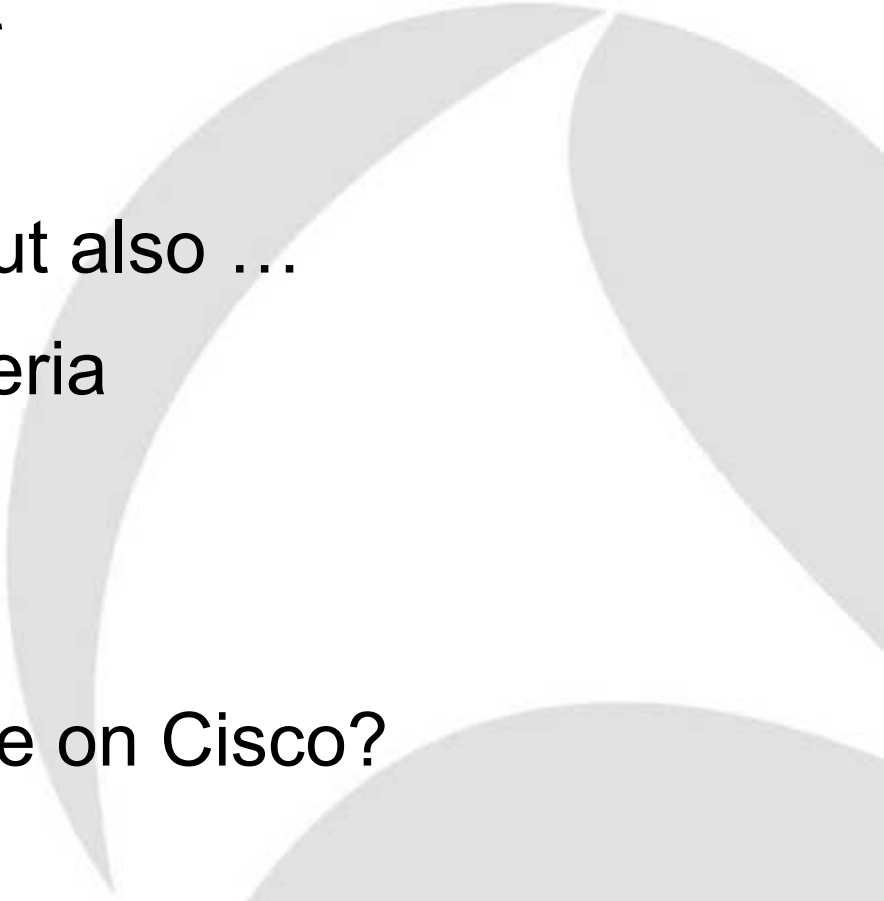
Cisco ACL / VACL: Topology

ACL – Access Control List

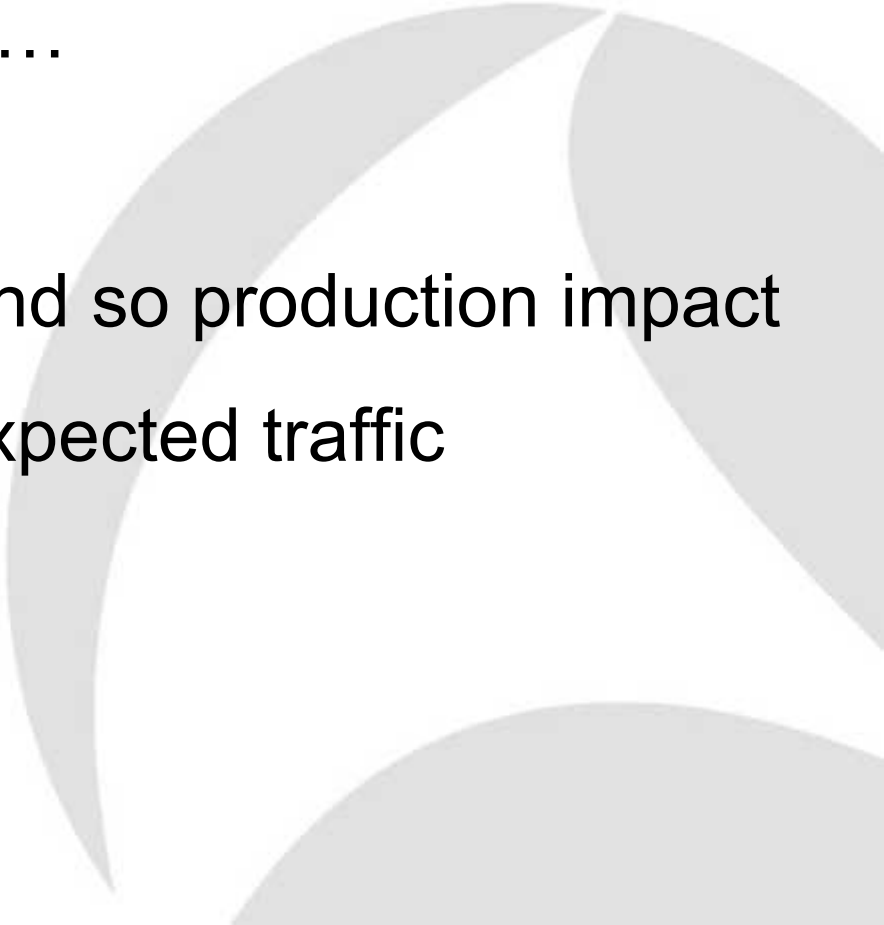
VACL – VLAN Access Control List



Cisco ACL / VACL: Advantages

- VACL Capture on Catalyst
 - ACL Capture on Nexus
 - Similar to monitor/mirror but also ...
 - Wide range of monitor criteria
 - IP addresses, port numbers, etc.
 - Helps avoid destination overload
 - More sessions possible
 - Is this the future for capture on Cisco?
- 

Cisco ACL / VACL: Considerations


- As per monitor/mirror plus ...
 - Complicated to configure
 - Greater risk of a mistake and so production impact
 - Risk of not capturing the expected traffic
- 

**Time for
Questions**

Blade Enclosure: Front

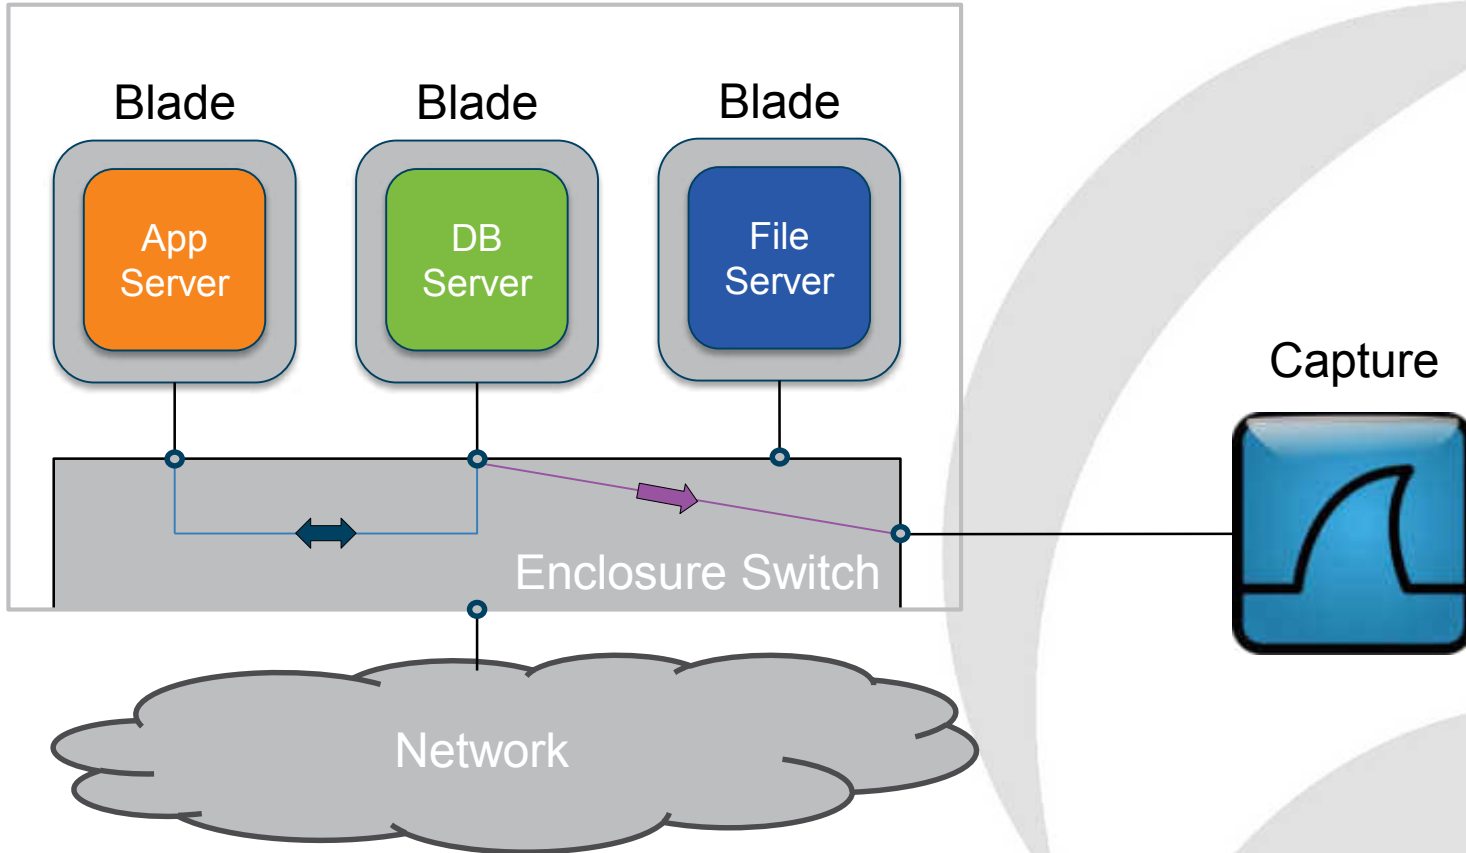


Blade Enclosure: Rear

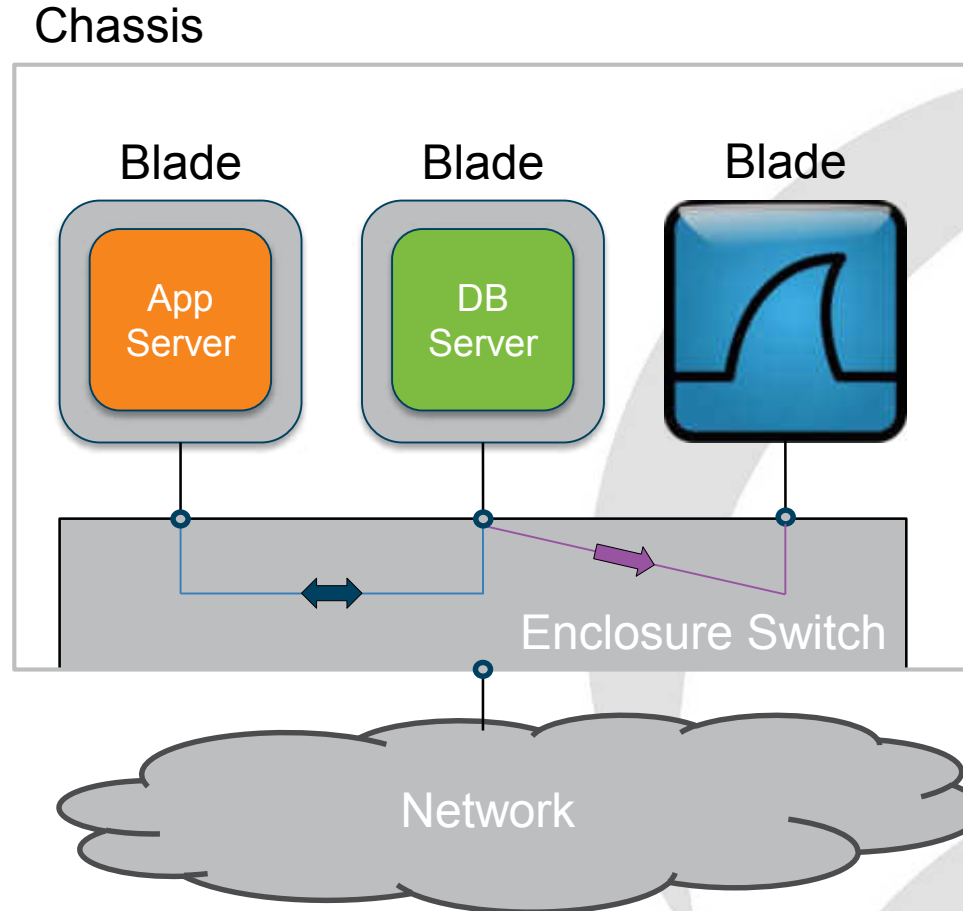


Blade Enclosure: Topology


Chassis



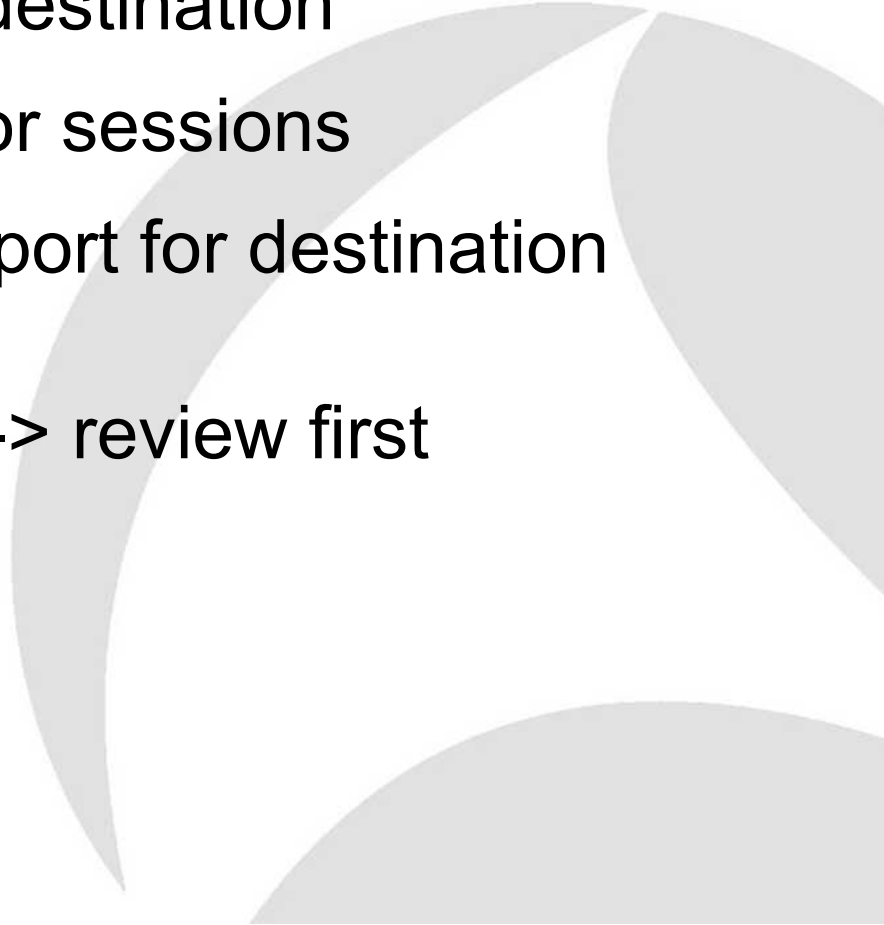
Blade Enclosure Alternative



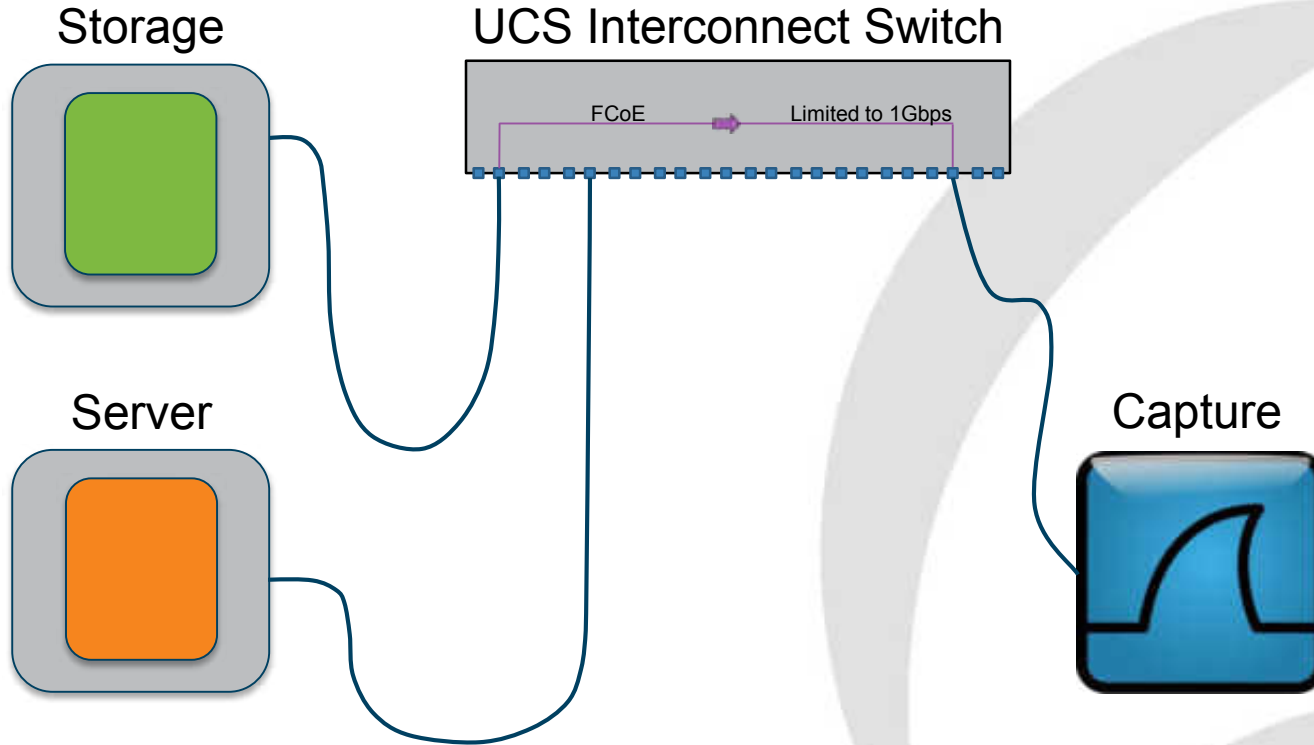
Blade enclosure: Advantages

- Easy to configure
 - Low risk – non-invasive
 - Multiple sources into one destination
 - Often entire VLANs can be monitored
 - Need to monitor on each switch
 - May see duplicates
 - Negligible impact on the switch
- 


Blade enclosure: Considerations

- Overload of the monitor destination
 - Limited number of monitor sessions
 - Requires a spare switch port for destination
 - Often all external ports are in use
 - Makes and models vary -> review first
- 

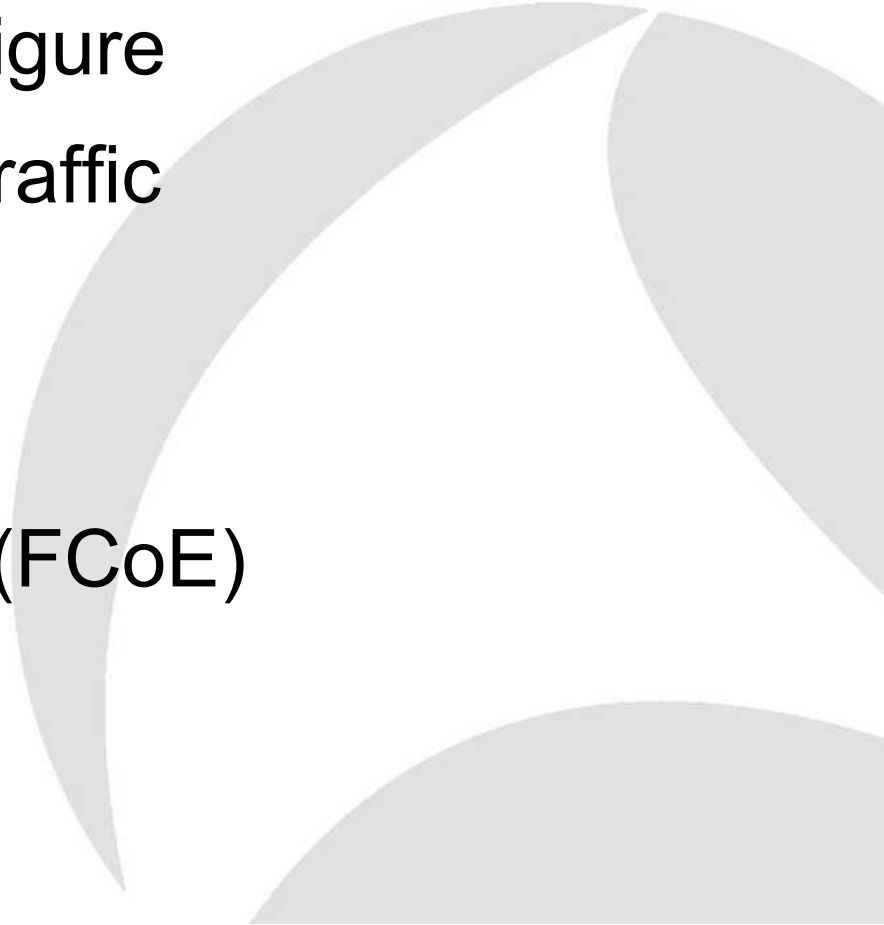
Cisco UCS Fabric Interconnect



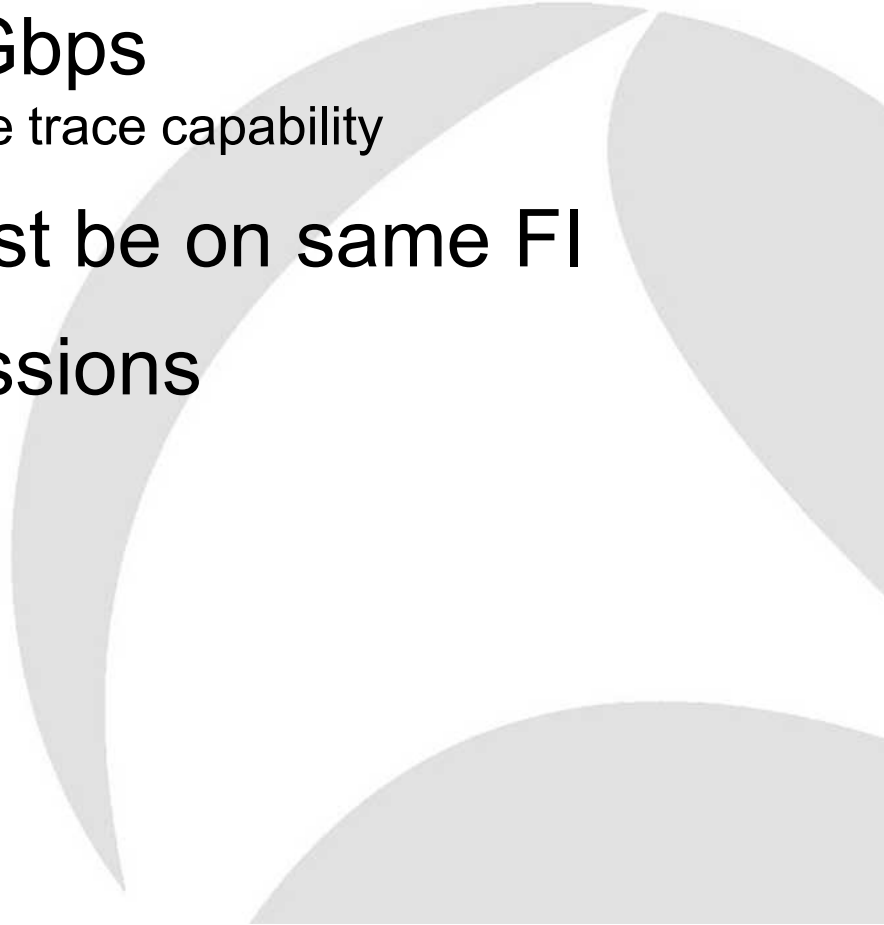
Cisco UCS Fabric Interconnect



UCS Fab Interconnect: Advantages

- Quick and easy to configure
 - Visibility to East-West traffic
 - Monitor multiple source
 - Monitor VLANs
 - Capture storage traffic (FCoE)
- 


UCS Fab Interconnect: Considerations

- Monitoring limited to 1Gbps
 - This probably negates the storage trace capability
 - Monitor src and dst must be on same FI
 - Limit of two monitor sessions
- 


**Time for
Questions**

TAP

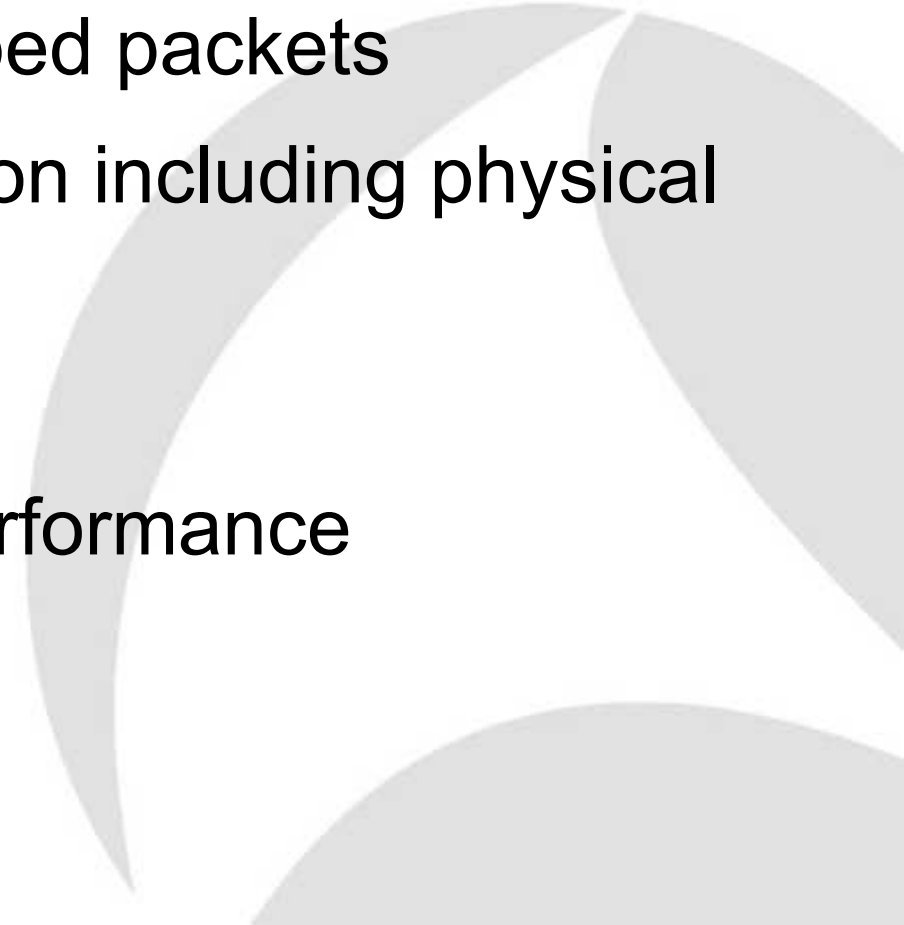
Non-aggregator Tap



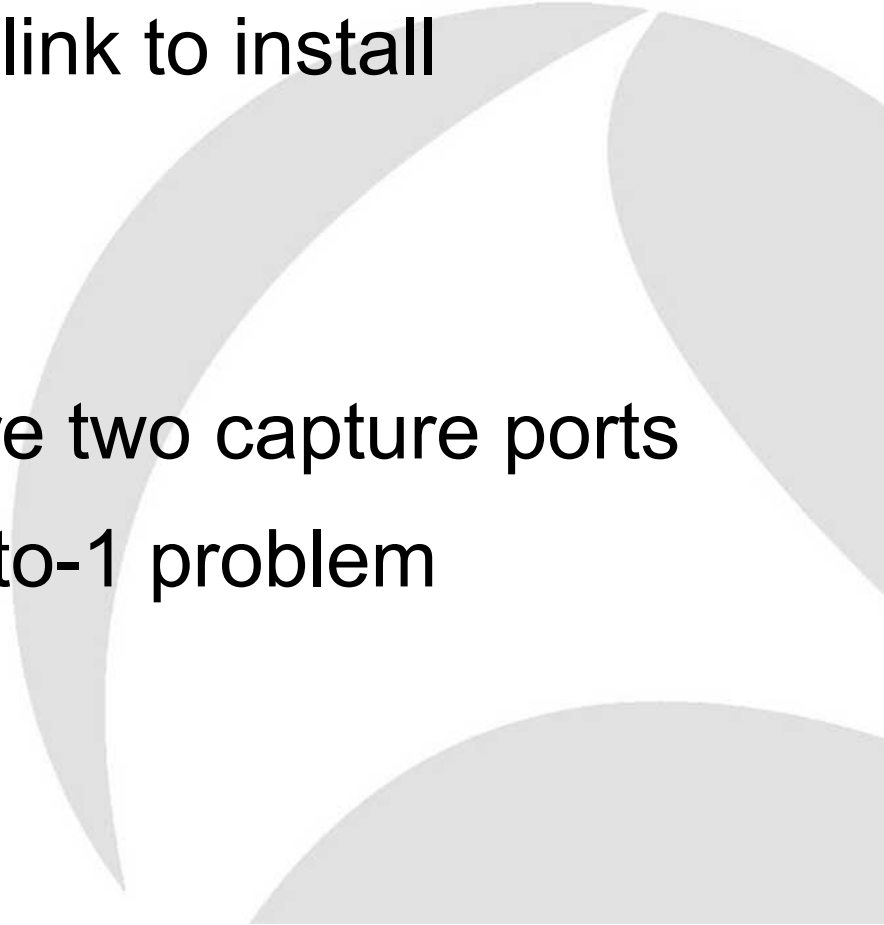
Aggregator Tap



TAP: Advantages

- Reduces risk of dropped packets
 - Captures all information including physical errors
 - Totally passive
 - Will not affect host performance
- 

TAP: Considerations

- Need to break network link to install
 - More expensive
 - Less flexible
 - Non-aggregators require two capture ports
 - Aggregators suffer 2-into-1 problem
- 

Network Packet Broker

Filter


Slice

Timestamp



De-dup

Many-to-many



SPAN
Monitor
Mirror

TAP

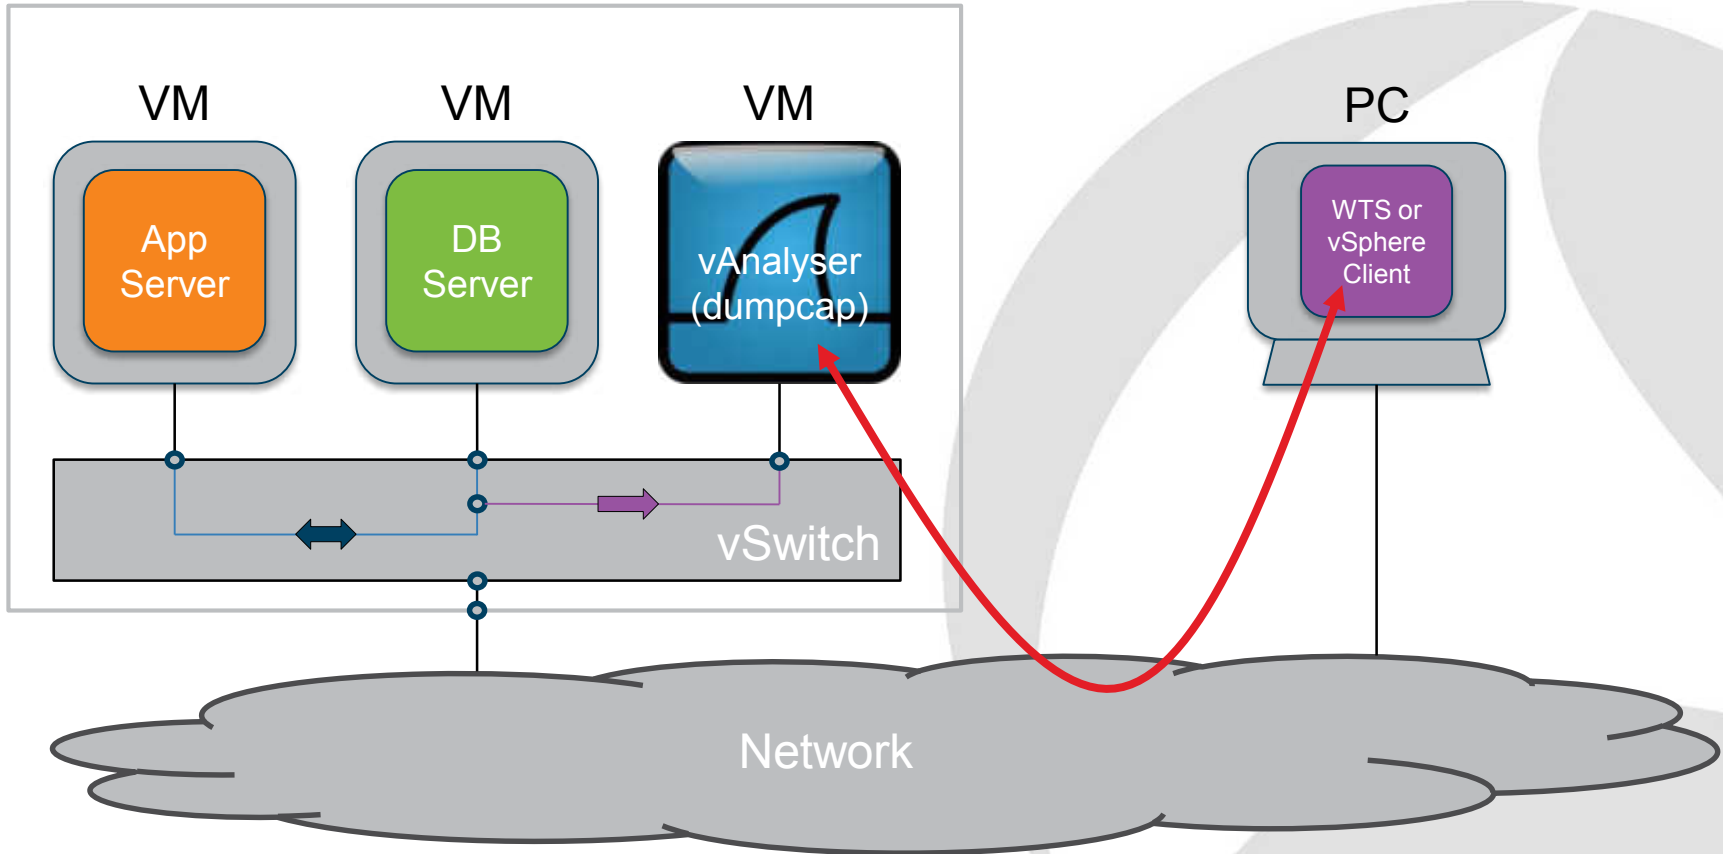
Other
Sources

**Time for
Questions**



ESX vSwitch Promiscuous Mode

ESX Host



Promiscuous Mode: Advantages

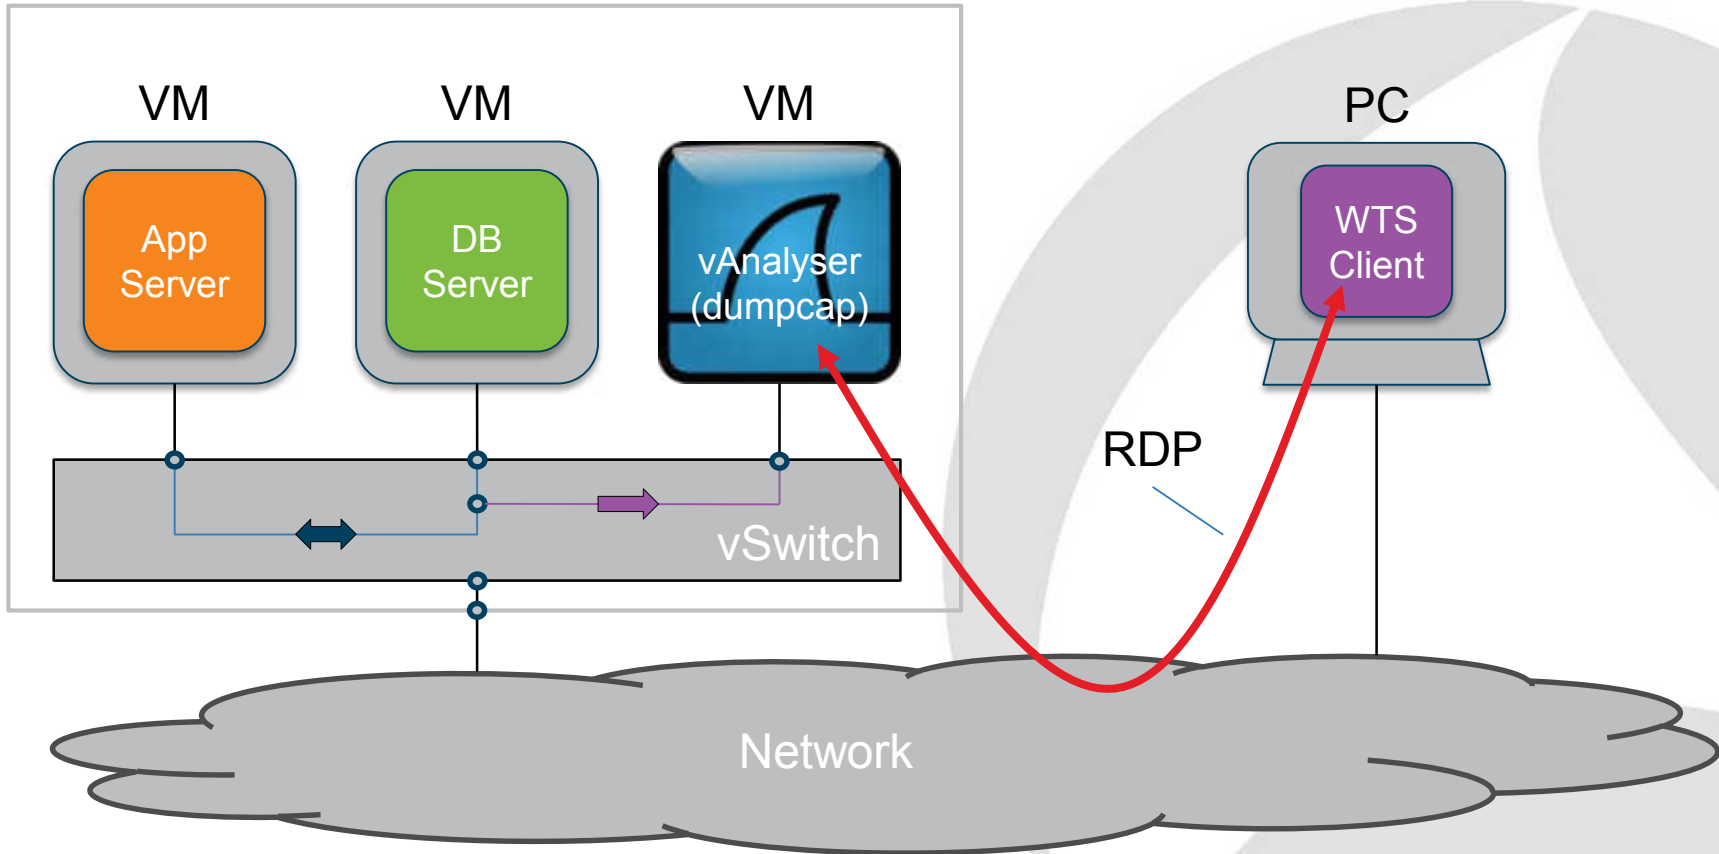
- Minimal disruption to services
 - Change Request probably needed
- Can capture all intra-vSwitch traffic
 - East-West

Promiscuous Mode: Considerations

- vAnalyser VM required
- Care regarding destination of trace data
 - Not to sensitive volumes
- Anecdote that causes high CPU load
 - This has not been our experience
- Capture will not follow vMotioned guest

Hyper-V Monitor Port

Hyper-V



Hyper-V Monitor Port: Advantages

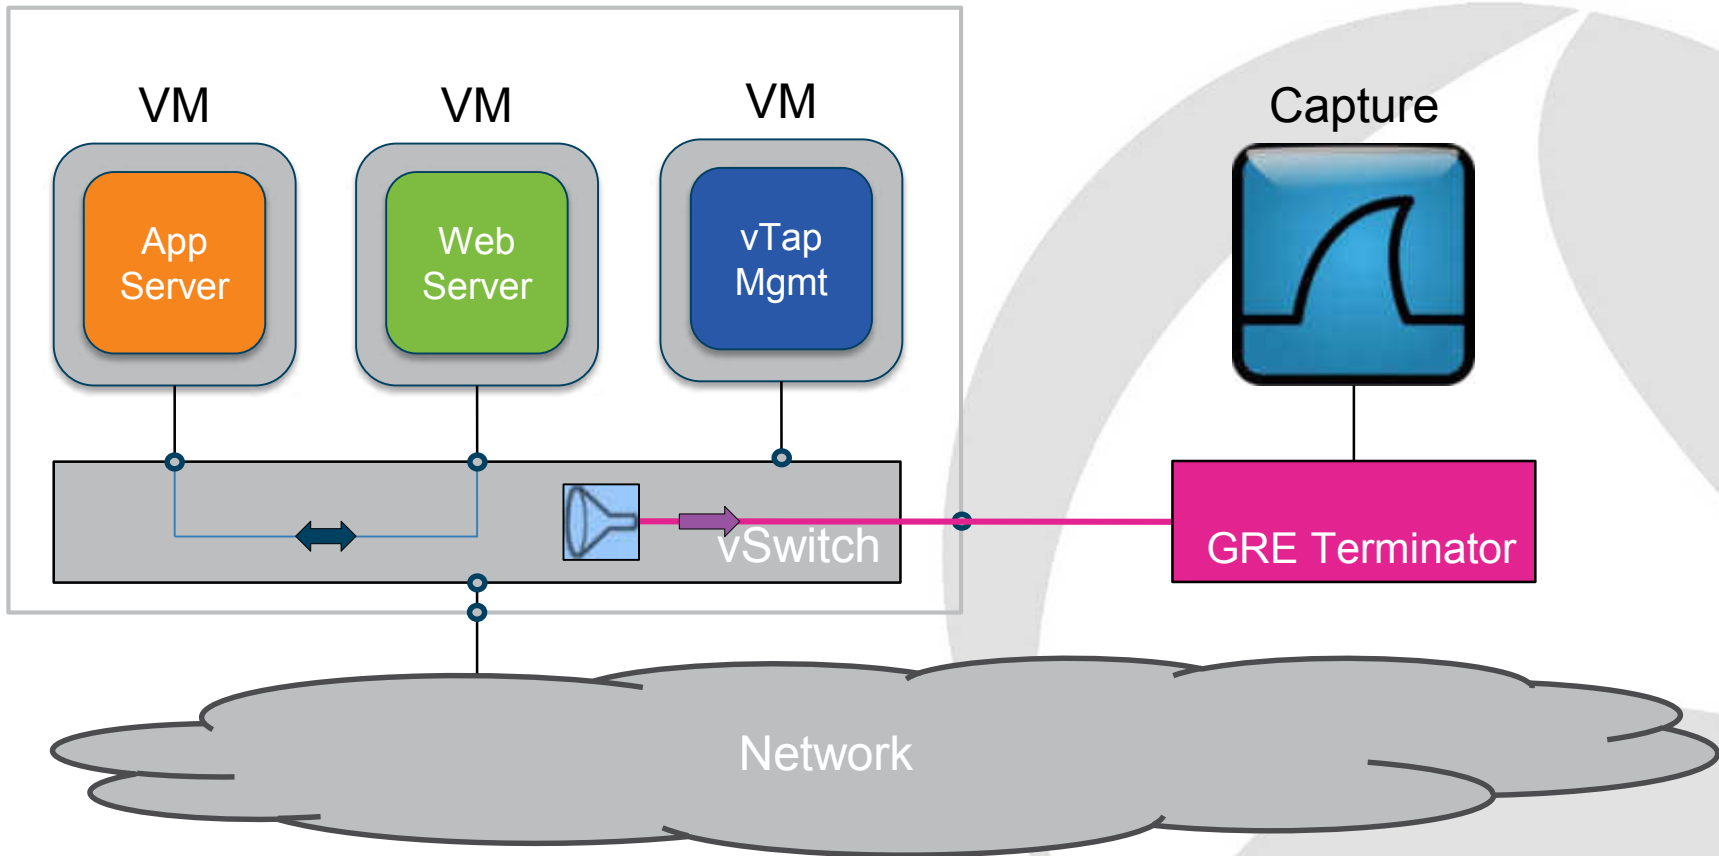
- Similar to monitor/mirror on a physical switch
- Minimal disruption to services
 - Change Request probably needed
- Can capture all intra-vSwitch traffic
 - East-West

Hyper-V Monitor Port: Considerations

- vAnalyser required
- Care regarding destination of trace data
 - Not to sensitive volumes

Ixia Phantom vTap

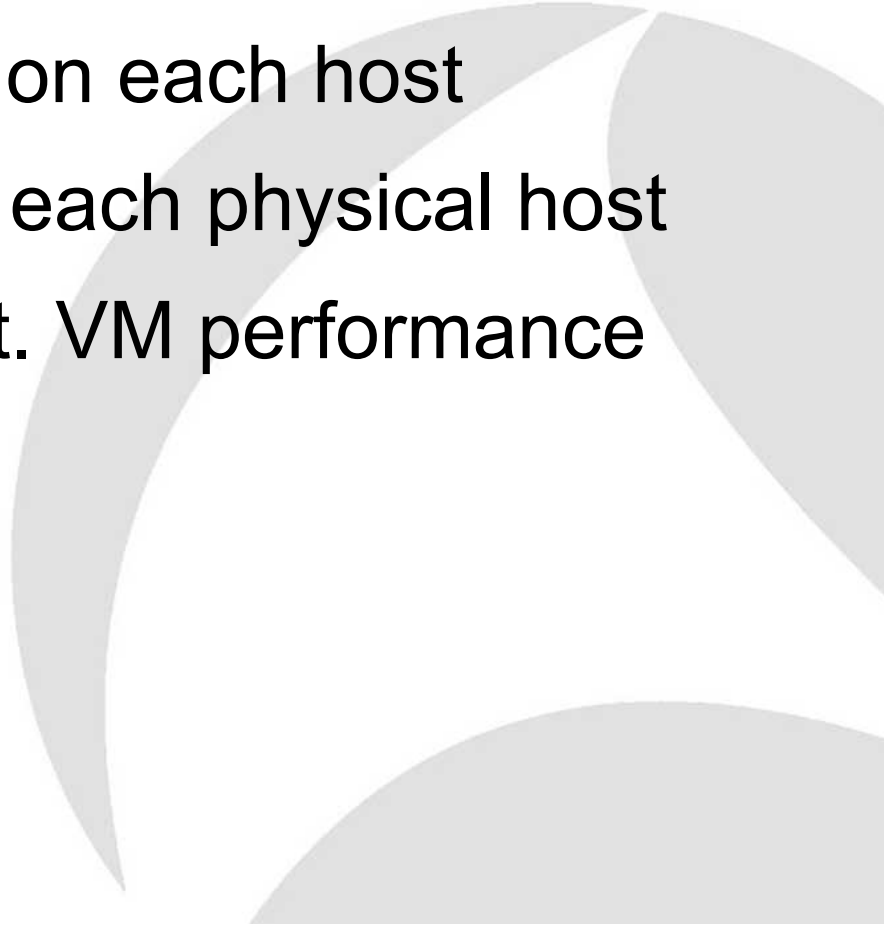
ESX / Hyper-V / XenServer



Ixia Phantom Tap: Advantages

- No software required on VM's
- No impact to VM performance
- vTap can capture all vSwitch traffic
- Or can capture specific traffic
- Works on the leading hypervisors
- Can track a VM thru' ESX vMotion

Ixia Phantom Tap: Considerations

- vTap Management VM on each host
 - Annual subscription for each physical host
 - Sensitive to vTap Mgmt. VM performance
- 

Further information

Paul Offord FBCS CITP
Mobile: +44 1279 211 668
Email: paul.offord@advance7.com
Web: www.advance7.com



LinkedIn Communities

LoveMyTool - Building an Open Community for Network Management and Monitoring

Protocol Analysis, Data Recorder, CALEA, Lawful Intercept, Application Performance, User Experience, Industrial Ethernet, Data Loss Prevention, Deep Packet Inspection, NetFlow, SOX, HIPAA and PCI Compliance, Switching and Routing, Forensics, VoIP, IPTV ... etc.



TribeLab

- Free tutorials
- Free guides
- Free resources