RELEVANT TECHNOLOGIES
: creating security knowledge

Solve Cloud Packet Mysteries

Laura Taylor, ltaylor@relevanttechnologies.com

March 24, 2014

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Bad Problem: Packets Can Leak from One Tenant to Another

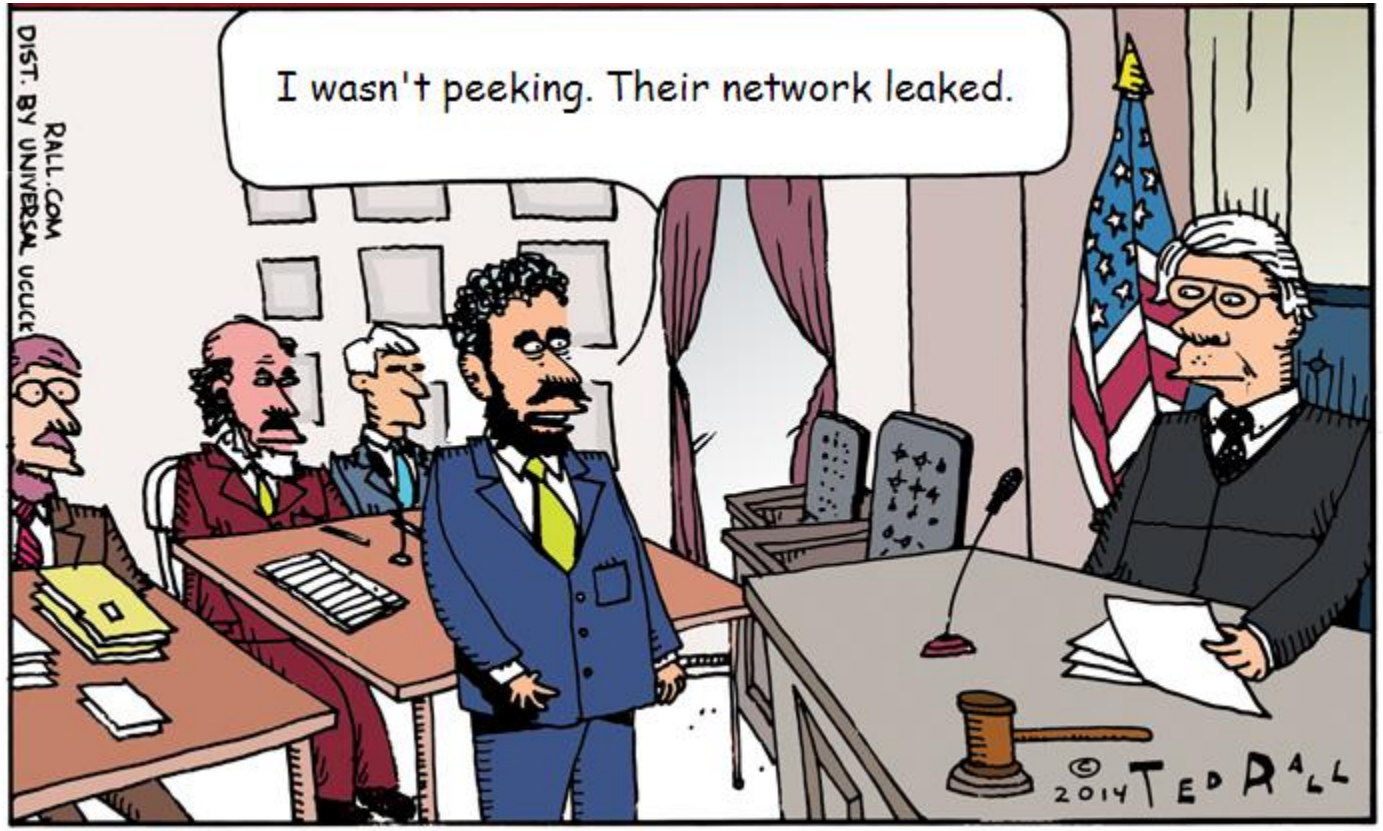

Make it look like they accidentally let their packets leak.

Worse Problem:

Malicious code can circumvent tenant isolation methods (Challenge 15)

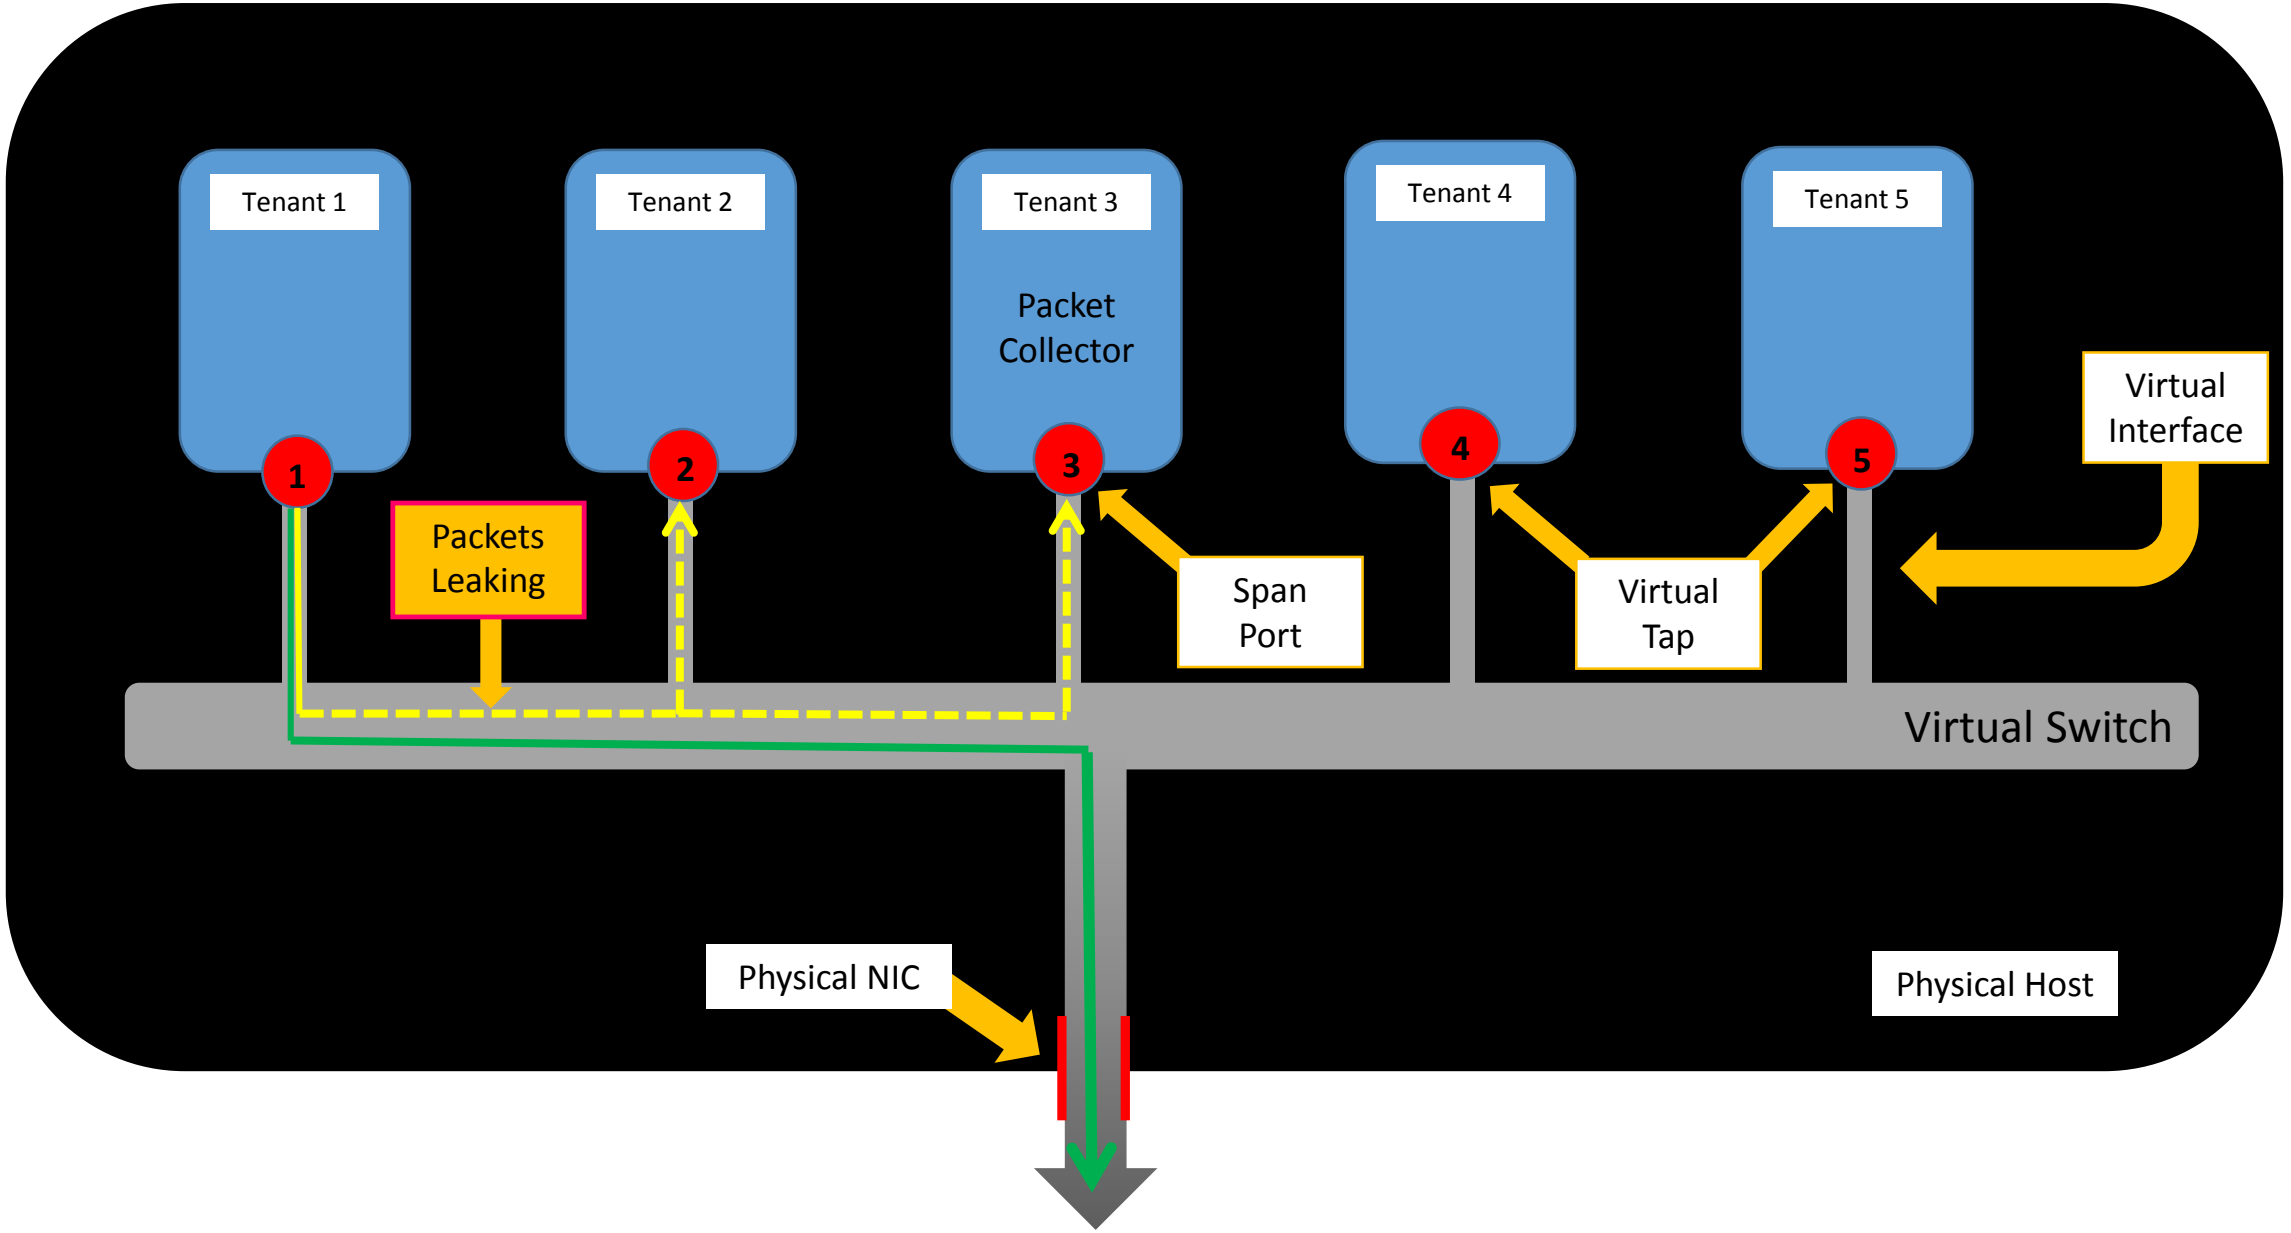

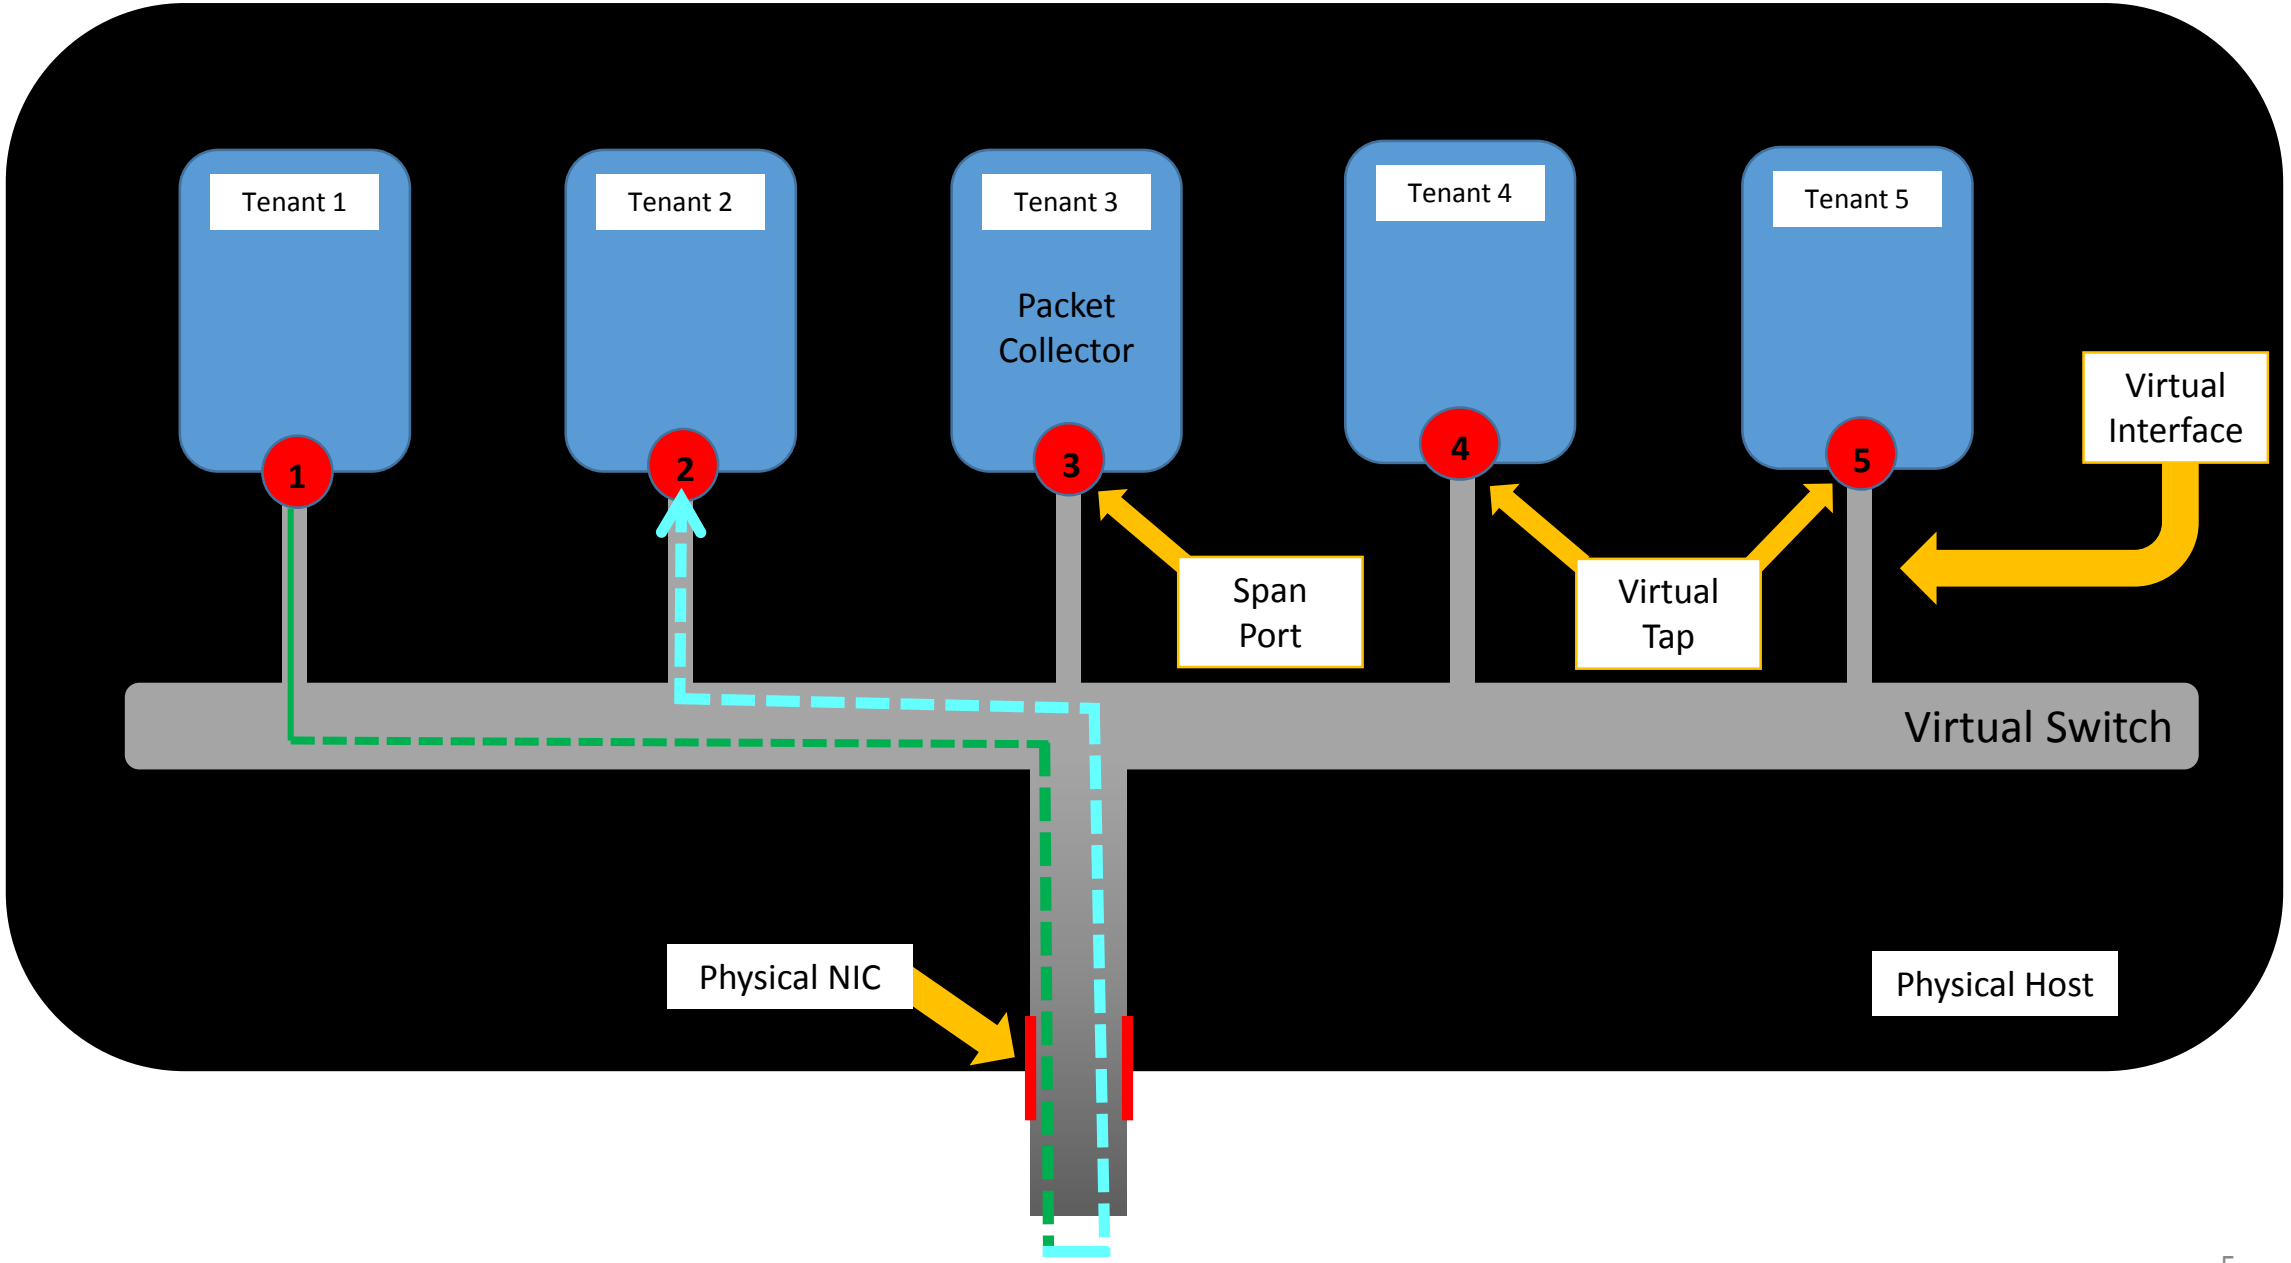

What makes packets leak to other tenants?

- Trojans and virtual machine-based rootkits
- Misconfigured routers
- Misconfigured virtual switches
- Misconfigured live migrations
- Hypervisor abstraction leaks (bugs)
- Network Address Translation limitations
- Switch address tables, unable to scale to network traffic, can overflow leading to flooding of unknown destination frames

The most serious packet leaking scenario?

It is cost prohibitive for network administrators to isolate tenants each on their own network domain -- so tenants are mixed together on shared networks. A common problem is that each tenant may independently assign MAC addresses and VLAN IDs leading to potential duplication of these on the physical network.

Isolating tenants on Layer 3 is not ideal.

Oh dear! What should we do?

But wait. Before we fix this problem, think about how it will effect a forensic investigation?

- In a forensic investigation, even if you have captured the packets using a full packet capture appliance, how will you know which tenant they came from?
- If you can't prove which tenant the packets came from, how can you build a forensics case?

What should CSPs do?

What should CSPs do?

- Test for leaks before initiating tenant provisioning.
- Don’t allow guest virtual machines direct access to the physical hardware.
- Don’t connect the guest virtual machine to the physical network via a bridge/router.
- Maintain the integrity of guest operating systems offered to consumers and protect the kernel using standard procedures for disabling unneeded modules.

How do we do this?

- Use Single Root I/O Virtualization (SR-iOV)
- Use Virtual eXtensible Local Area Network (VXLAN) overlay networks (new IETF standard)

What is SR-iOV?

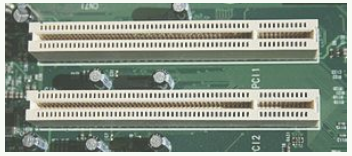
A PCI Connector

- SR-iOV (Single Root I/O Virtualization)
- A PCI SIG standard (www.pcisig.com)
- Modern physical NIC is virtualized at the PCI level
- Virtual machine binds directly to a virtual function within the NIC bypassing the hypervisor vswitch entirely
- One to one assignment from virtual function to virtual switch
- Cisco calls this Data Center Ethernet (DCE)
- Most of the rest of the SAN industry (and IBM) call this Converged Enhanced Ethernet (CCE)

SR-iOV Illustrated

1) Packet arrives at NIC
2) Packet sent to Layer 2 sorter
3) Packet queued to target virtual function
4) Direct memory access operation (DMA) initialized
5) DMA operation completes, packet now in vm memory
6) NIC fires interrupt indicating packet has arrived
7) VMM fires interrupt to vm to announce arrival

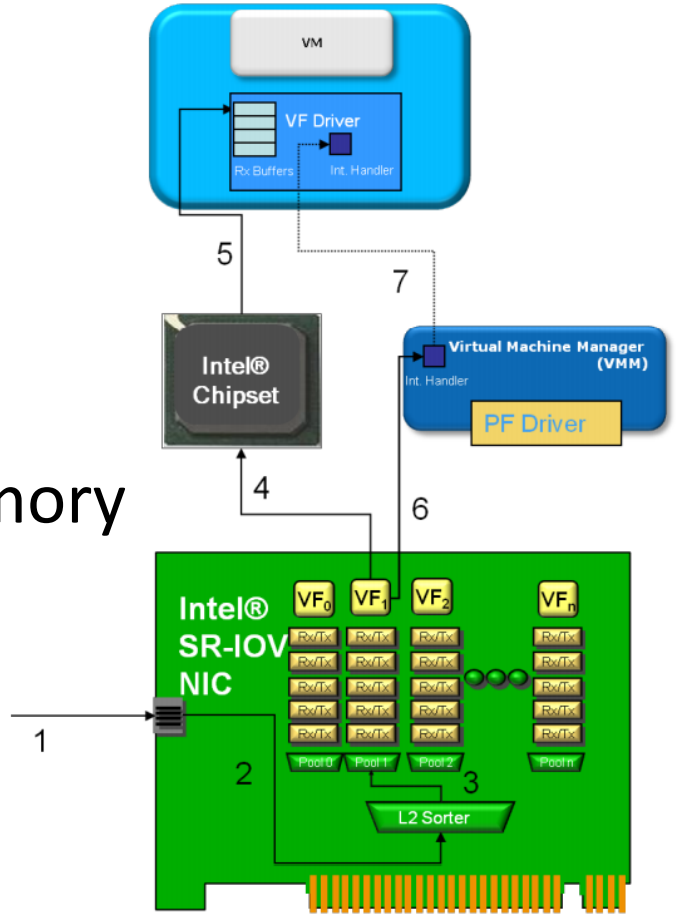

Source: Intel

What is VXLAN?

- Virtual eXtensible Local Area Network
- An overlay network and a new IETF standard
- Enables hypervisor hide references to local resources by using encapsulation
- Encapsulates the packets at the point in time when the guest frames leave the guest network stack and enter the hypervisor -- prior to sending the layer 2 frames.

VXLAN Illustrated

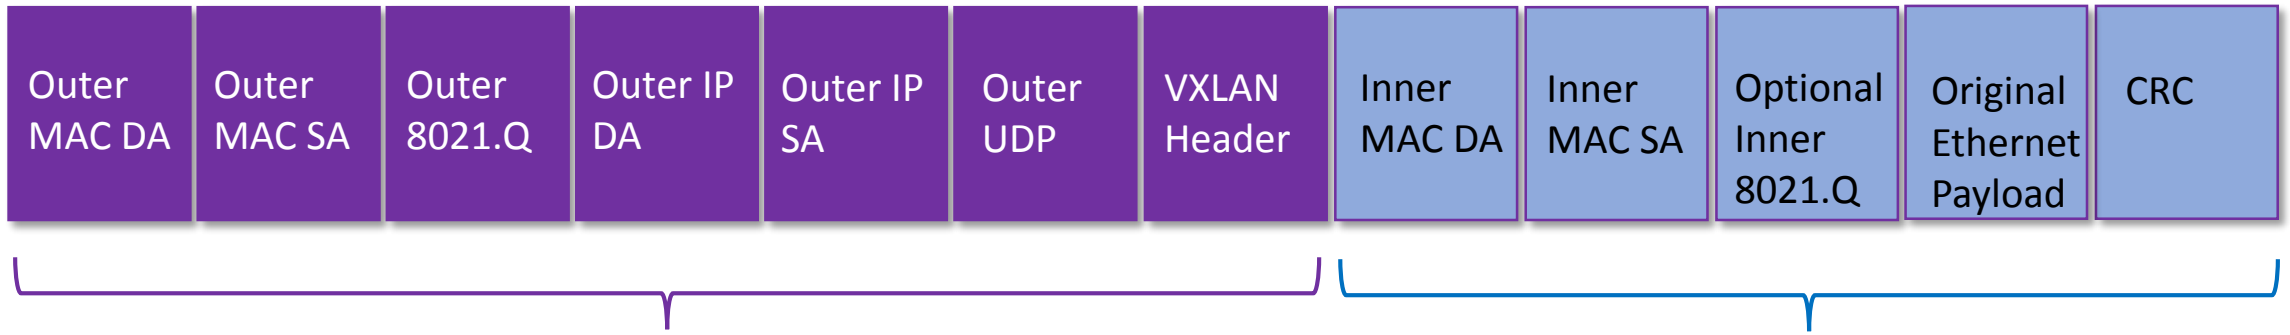

What should tenants do?

What should CSP tenants do?

- Ask your CSP what controls they have put in place to isolate tenants
- Ask your CSP what they did to test for packet leaks between tenants
- Ask CSP if the switches they are using support SR-iOV or VXLAN
- Assume your CSP might not be isolating tenants correctly
  - Put a virtual tap (vtap) on your virtual interface
  - Use the vtap to start collecting packets and look for packet leaks. Your CSP might not be aware of existing packet leaks so tenants might need to watch out for this themselves.
  - Report packet leaks to the CSP as a security incident

The Takeaway

- If you use the wrong technologies, and/or if the cloud is not configured correctly, it could become nearly impossible to prove which tenants the packets came from, therefore making it extremely challenging to build a forensics case.

References

Landau, D. Hadas, M. Ben-Yehuda, *Plugging the Hypervisor Abstraction Leaks Caused by Virtual Networking*, IBM, April 2010

Rutkowska, *Security Challenges in Virtualized Environments*, Invisible Things Lab, April 2008

King, P. Chen, Y. Wang, C. Verbowski, H. Wang, J. Lorch, *SubVirt: Implementing malware with virtual machines*, IEEE Symposium on Security and Privacy, March 2006

Mahalingam, Dutt et al., IETF Draft, *VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks* (Expires 8/3/2014)

Sridharan, et al., IETF Draft, *Network Virtualization using Generic Routing Encapsulation* (Expires 8/14)

Ormandy, An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments, Google, April 2007

Garfinkel, M. Rosenblum, *When Virtual Is Harder Than Real: Security Challenges in Virtual Machine-Based Computing Environments*, Stanford University, August 2005

Hooda, S. Kapadia; P. Krishnan, *Using TRILL, FabricPath, and VXLAN: Designing Massively Scalable Data Centers (MSDC) with Overlays*, Cisco Press, February 6, 2014

PCI-SIG SR-IOV Primer, An Introduction to SR-IOV Technology, Intel, January 2011

Questions?

ltaylor@relevanttechnologies.com