The Four Horsemen Of the Virtualization Apocalypse

Geekin' :: VirtSec Style

- Setup & Context
- x86 Virtualization Overview in 90 Seconds
- Virtual Networking Architecture
- VirtSec Solutions Landscape
- The Four Horsemen
- Wrap-Up

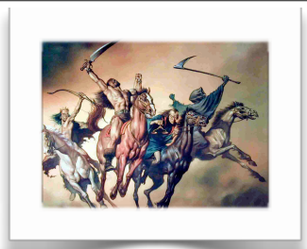

Status Quo = FAIL?

Reality Bites

Replicating many highly-available security applications and network topologies in virtual switches doesn't work

"It ain't all rainbows and unicorns..."

Bumpy Road Ahead

"Everything's Under Construction..."

Monolithic security vendor virtual appliances are the virtualization version of the UTM argument

If It Ain't Fixed, Don't Break It

"She just don't run like she used to..."

Virtualized Security can seriously impact performance, resiliency and scalability

PacketFilter

Penny Wise & Pound Foolish

Virtualizing security will not save you money, it will cost you more

"Money for nuthin' and my chips for free..."

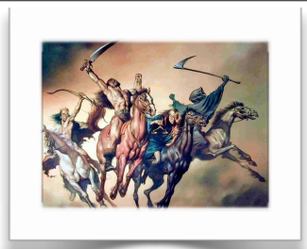

Where To Start?

- Setup & Context
- x86 Virtualization Overview in 90 Seconds
- Virtual Networking Architecture
- VirtSec Solutions Landscape
- The Four Horsemen
- Wrap-Up

x86 Virtualization Overview

From This To This

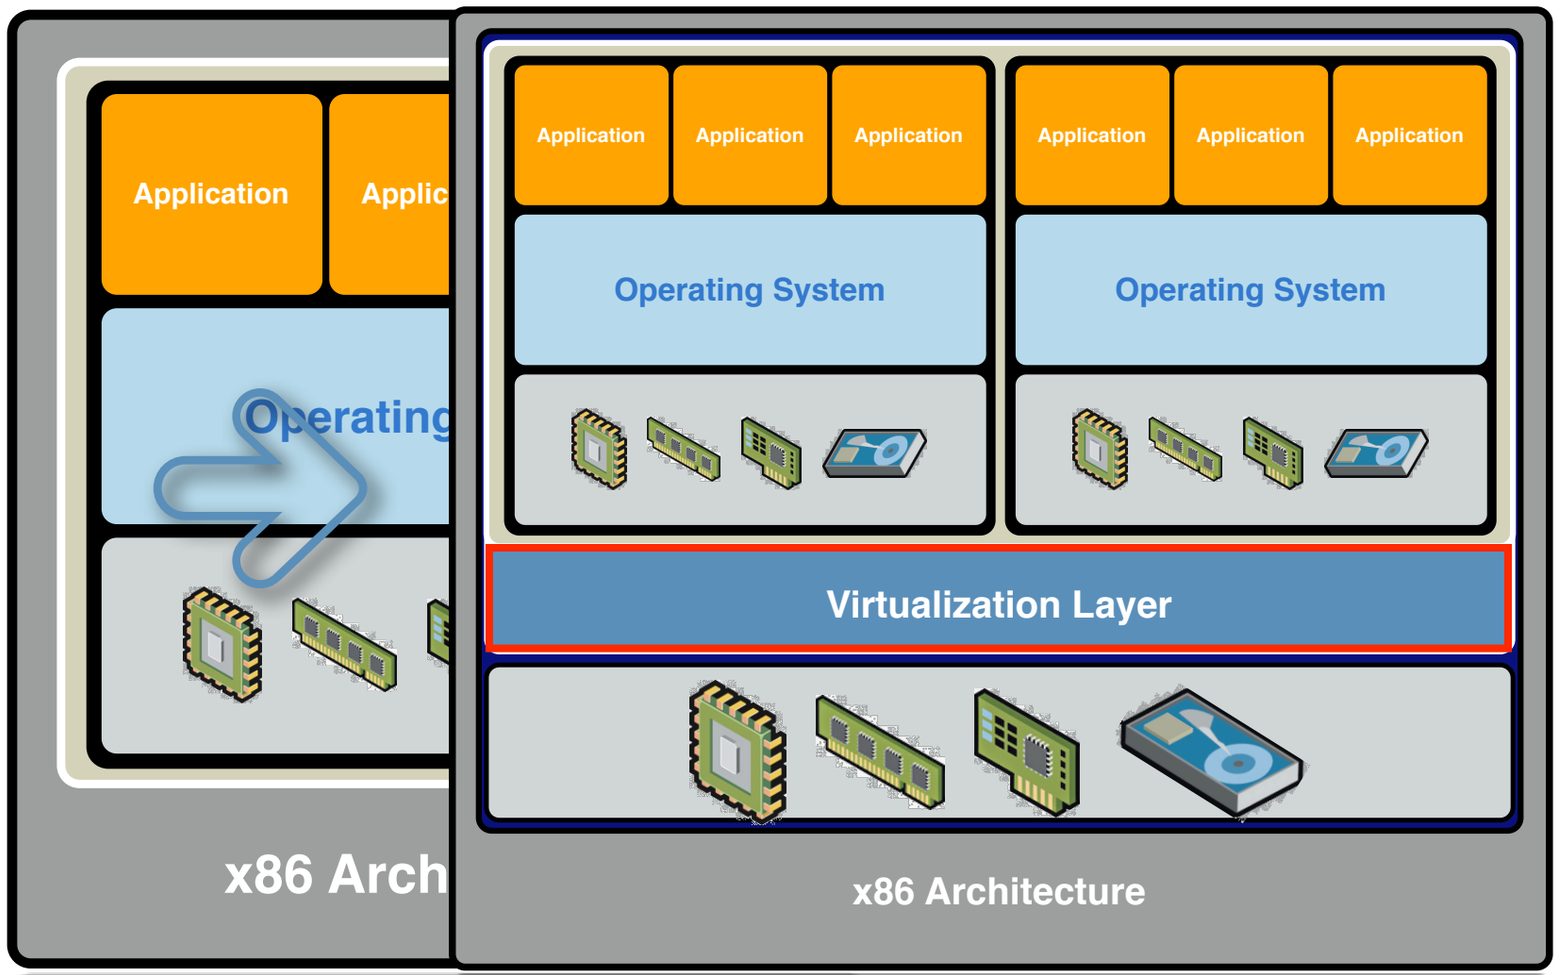

*Represents "Type 1" or Bare Metal Virtualization

x86 Hierarchical Protection Domains/Rings

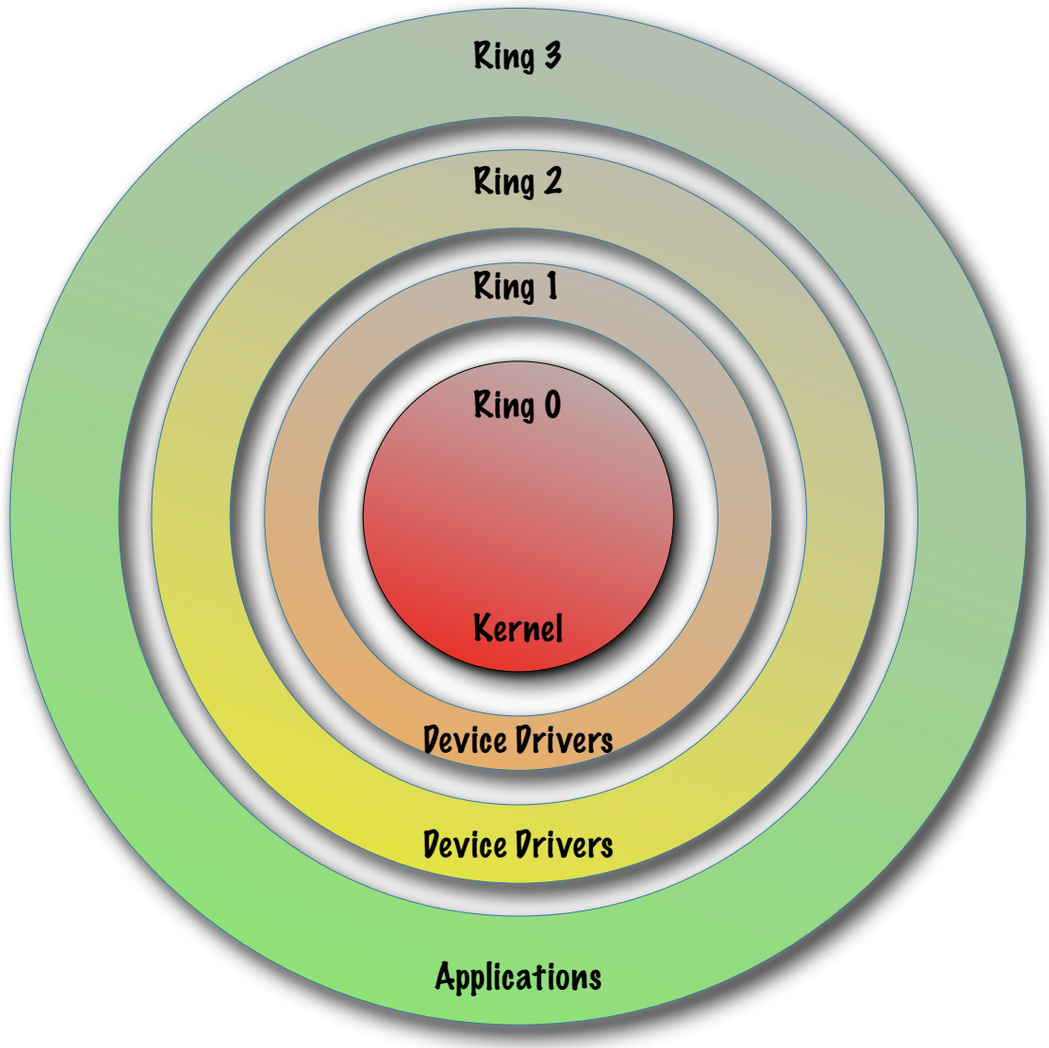

Most Privileged

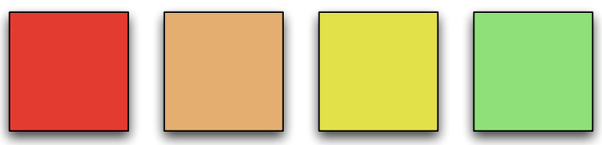

Least Privileged

Adapted from: http://en.wikipedia.org/wiki/Supervisor_mode

x86 Protection Ring Compression For Dummies

Virtualized: Physical/ Software Own-Virtualized

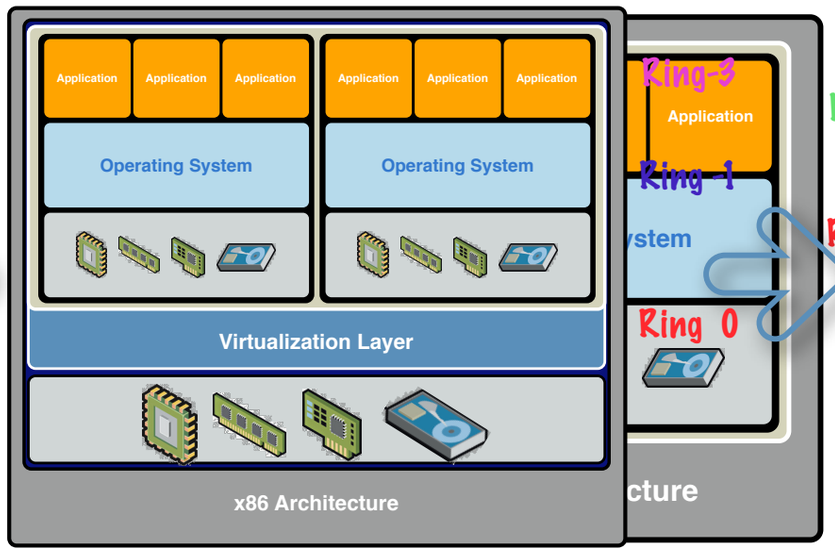

- The Guest OS is de-privileged into Ring-1 and the VMM takes its place in Ring 0
- The Guest OS still thinks it is running in Ring 0 with all the privileges thereof
- Can cause issues/conflicts due to contention for the 17 x86 privileged platform control instructions

Virtualized: Hardware Assisted

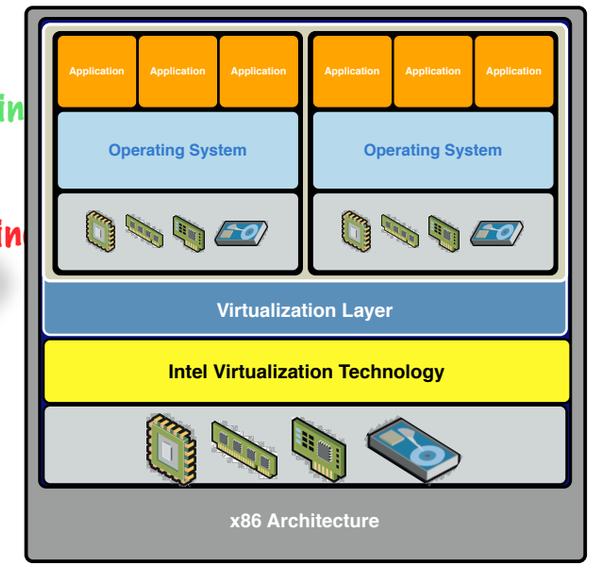

- In this example, Intel VT provides the VMM with an exclusive privileged level where it resides and executes
- The Guest OS is not de-privileged and is running in Ring 0
- Context switching between VMM and Guest OS's are hardware supported

*There is also para-virtualization, not covered here...

Hypervisors Are a Disruptive Commodity

vmware®

CITRIX®

Microsoft®

IBM®

Sun microsystems®

Parallels™

ORACLE®

Xen®

*Yes, there are others, but these have pretty logos...

...and they're showing up everywhere

No One RingO To Rule Them All!

Which means:

- You will likely end up with 4-5 virtualization platforms/VMM's spread out across the horizon of your enterprise

The key differentiators?

- Management, integration, extensibility and security

We need open standards for solution interoperability

- If you have issues with the "simple complexity" of a single virtualization platform, imagine when you have many

Debating Virtualization & Security

Many debates and much ado stems from the inability to distinguish between three fundamental concerns:

- Securing Virtualization
- Virtualizing Security
- Security Via Virtualization

Separate the technical, architectural, and philosophical from the functional, operational and organizational

Time For Sublime Design

- Setup & Context
- x86 Virtualization Overview in 90 Seconds
- Virtual Networking Architecture
- VirtSec Solutions Landscape
- The Four Horsemen
- Wrap-Up

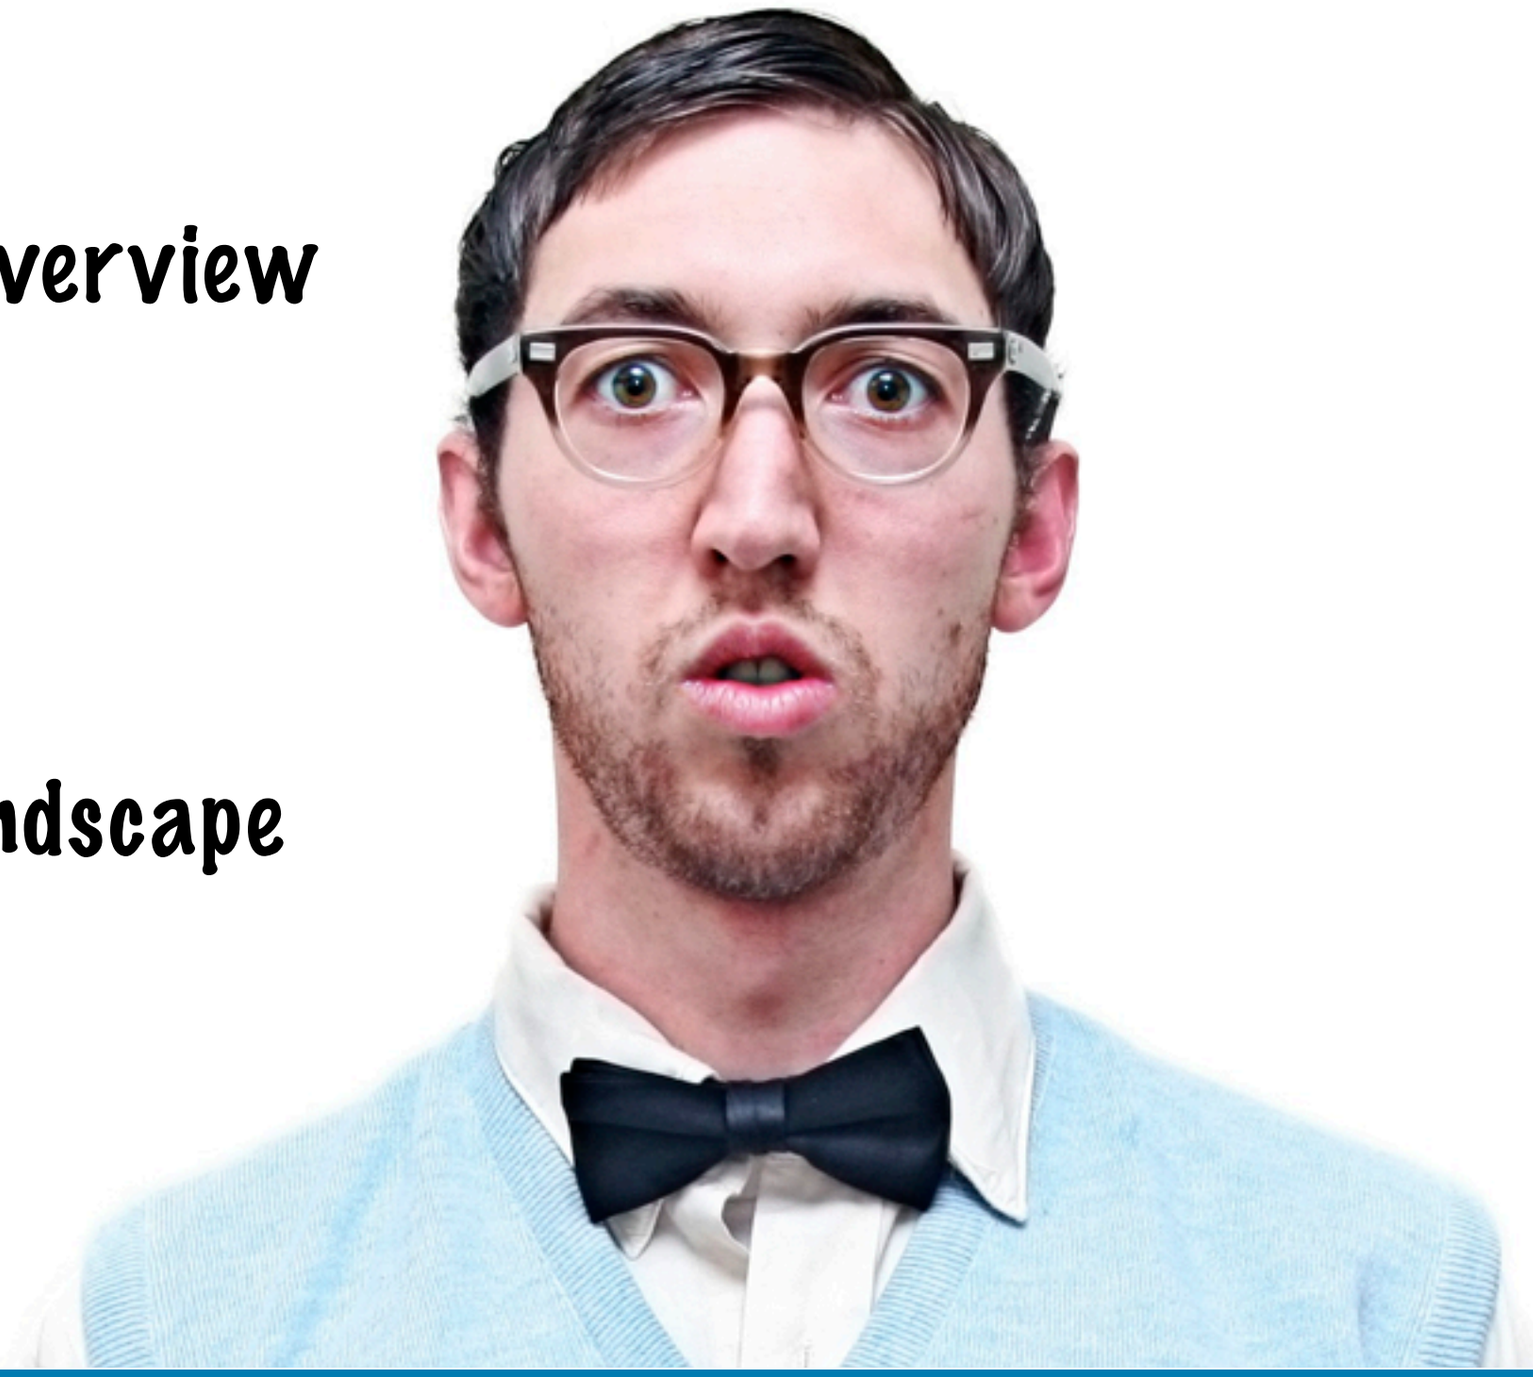

Caveats

- This presentation uses VMware ESX as my virtualization platform example featuring data networking only; storage is a whole other universe of security fun...
- It's true you can achieve very robust/resilient integrated network and virtual infrastructure designs, but the moment you try and integrate security...not so much...
- There are far too many dirty little secrets and unspoken truths regarding implementing VirtSec today; we're going to talk about them here

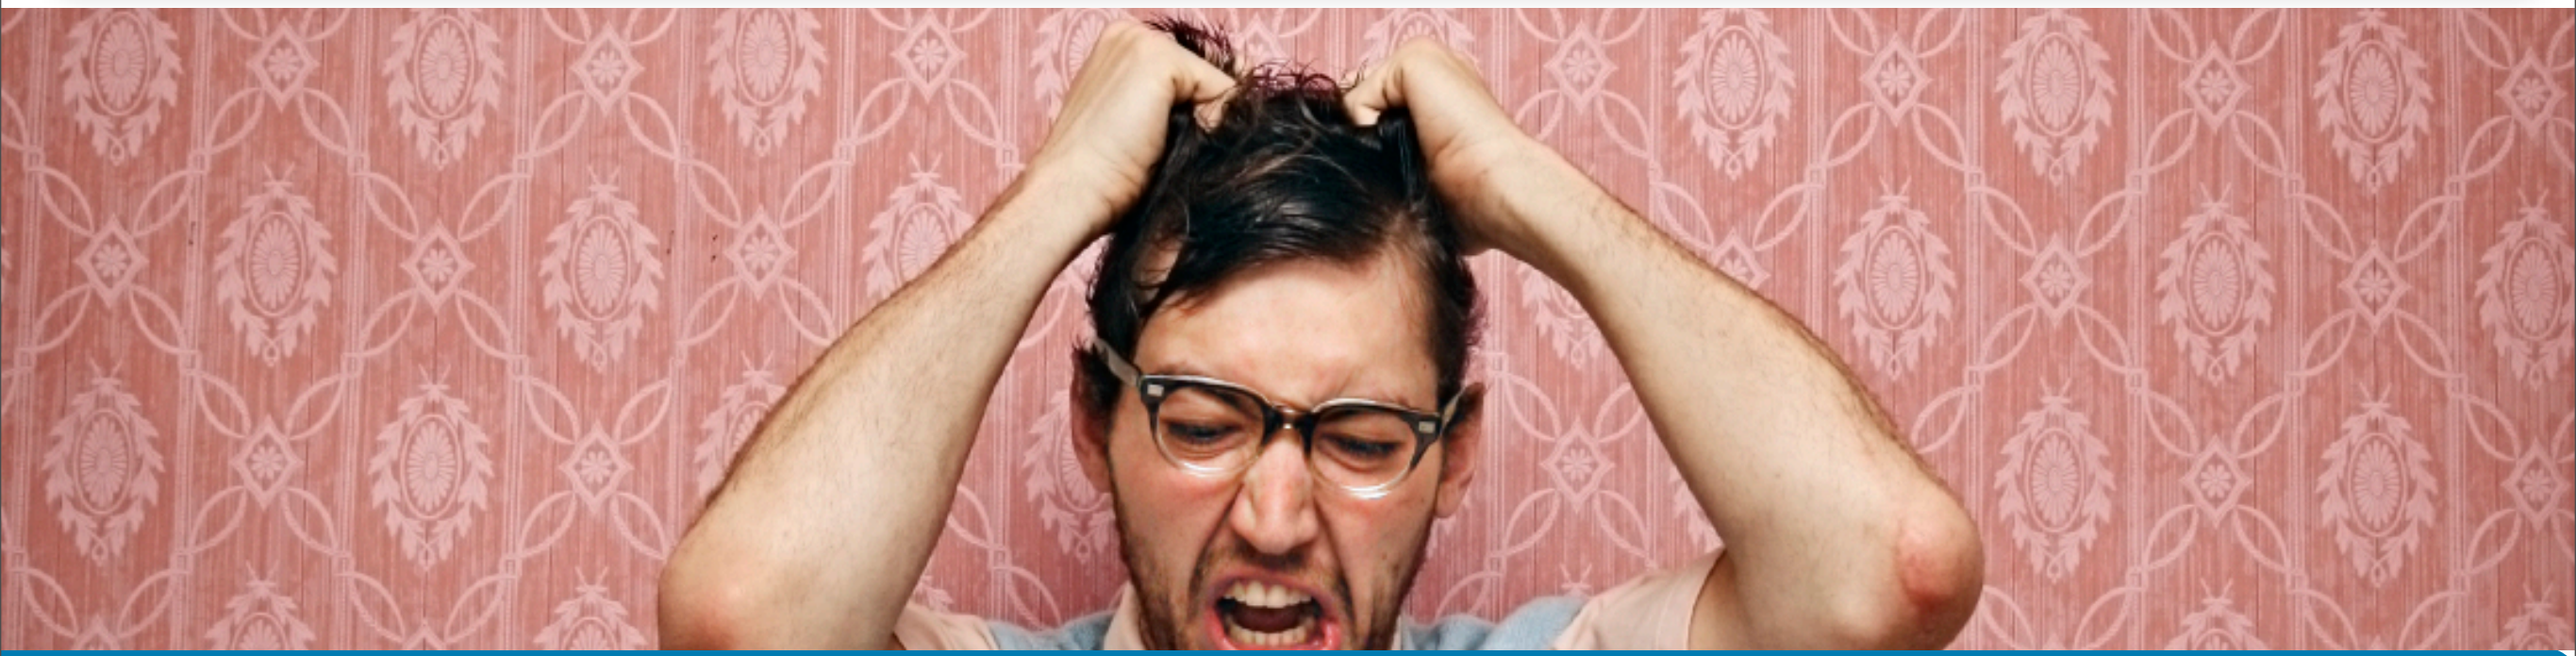

But dOOd, What About Virtualization Malware!?

There are many really interesting topics to discuss here:

- Hypervisor Malware & Hyperjacking
- Exploiting the virtualization chipsets for fun and profit
- Hardware/Firmware abuse
- Control channel manipulation

I'm neither qualified or motivated to talk about these topics and we've got much more profound and fundamental sets of issues to discuss.

There's an entire track dedicated to this stuff. Go there.

Virtual Networking Architecture

Virtual Networking

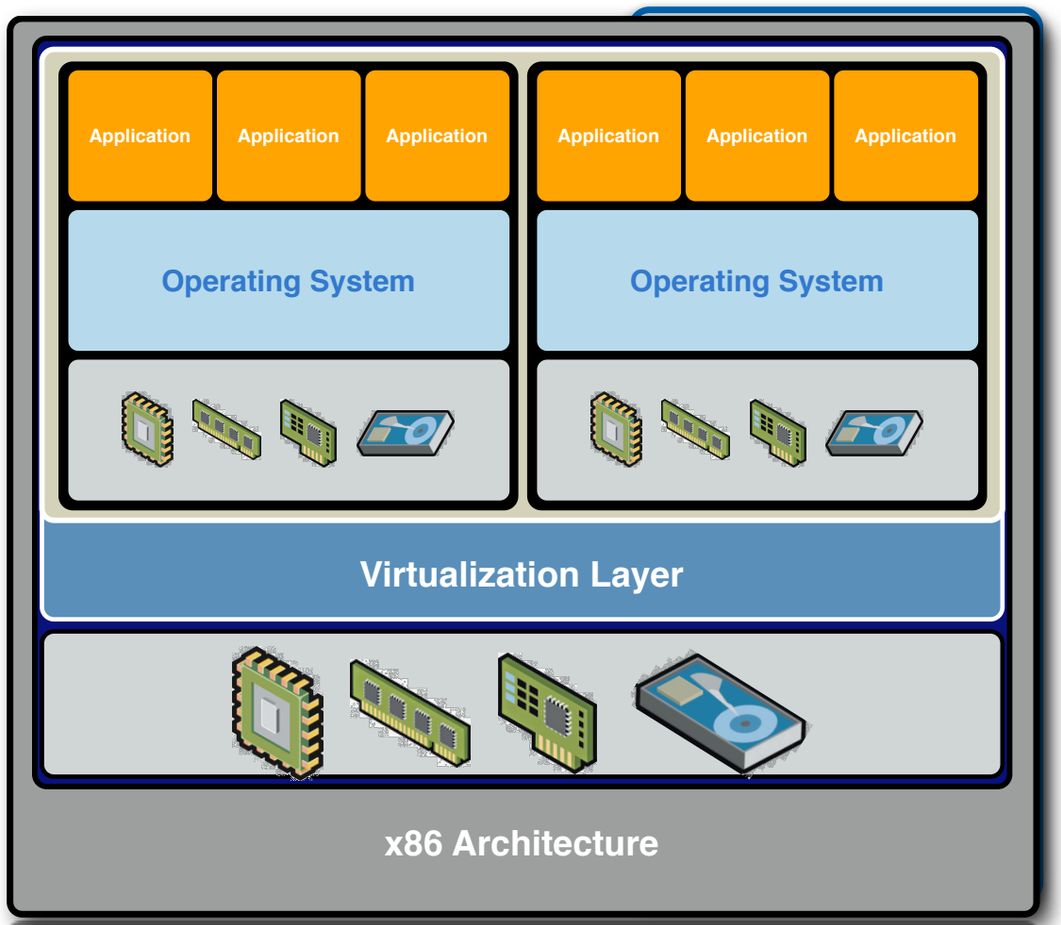

Virtual Switch

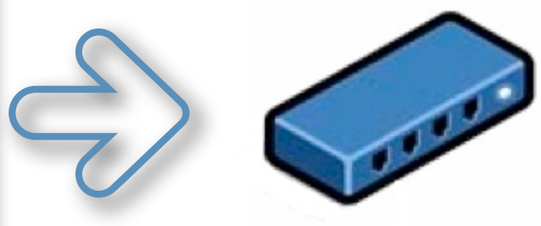

*Not shown for clarity: Service Console/VMKernel/Storage Networking

Virtual Switch Defined

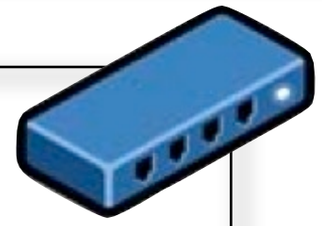

A Virtual Switch:

- Is a software-based networking construct that runs in the virtualization platform's kernel
- Purposely-designed layer-2 (L2) switch which is loaded dynamically at runtime with functional modules such as:
  - Core L2 forwarding engine
  - VLAN tagging, stripping & filtering
  - L2 security, checksum and segmentation offload
- Some features normally found in physical L2 switches are not present by design to provide for integrity, isolation and secure connectivity (no STP, VTP, ISL, etc...)

Virtual Switch Visualized

Abstracted Model

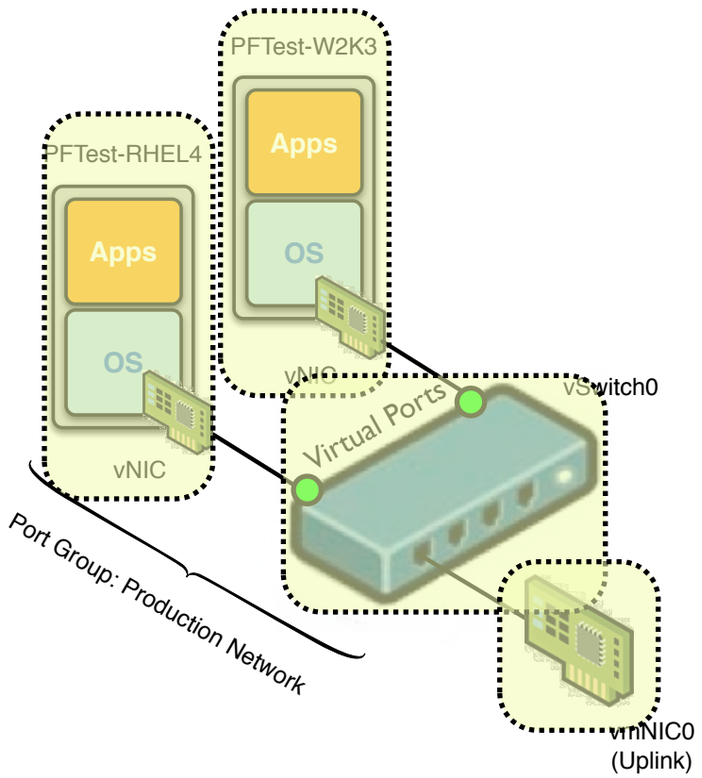

VMware Virtual Infrastructure Client View

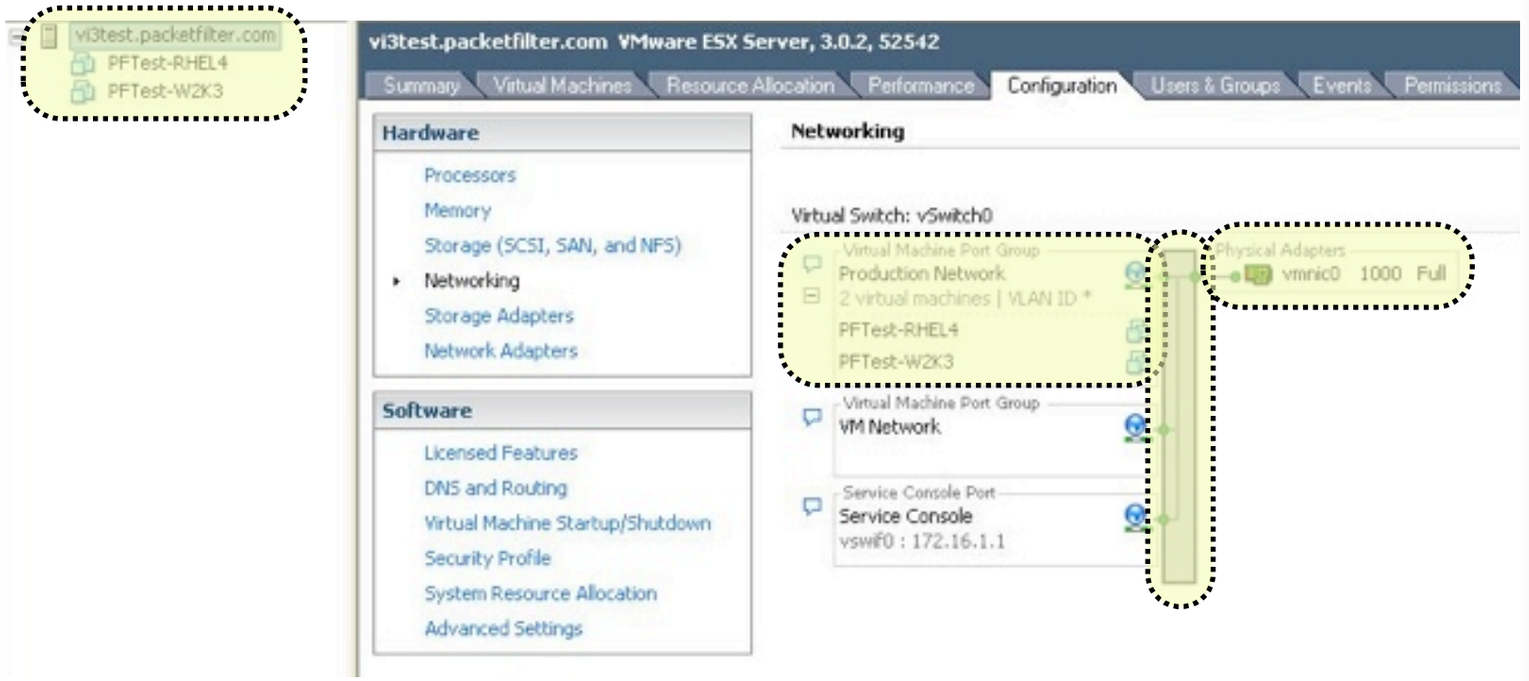

* I purposely left off the VMotion and Service Console networks in the model for clarity

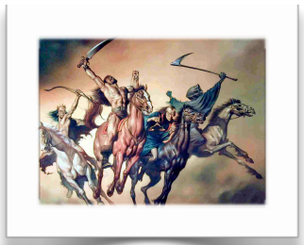

vSwitch Correctness

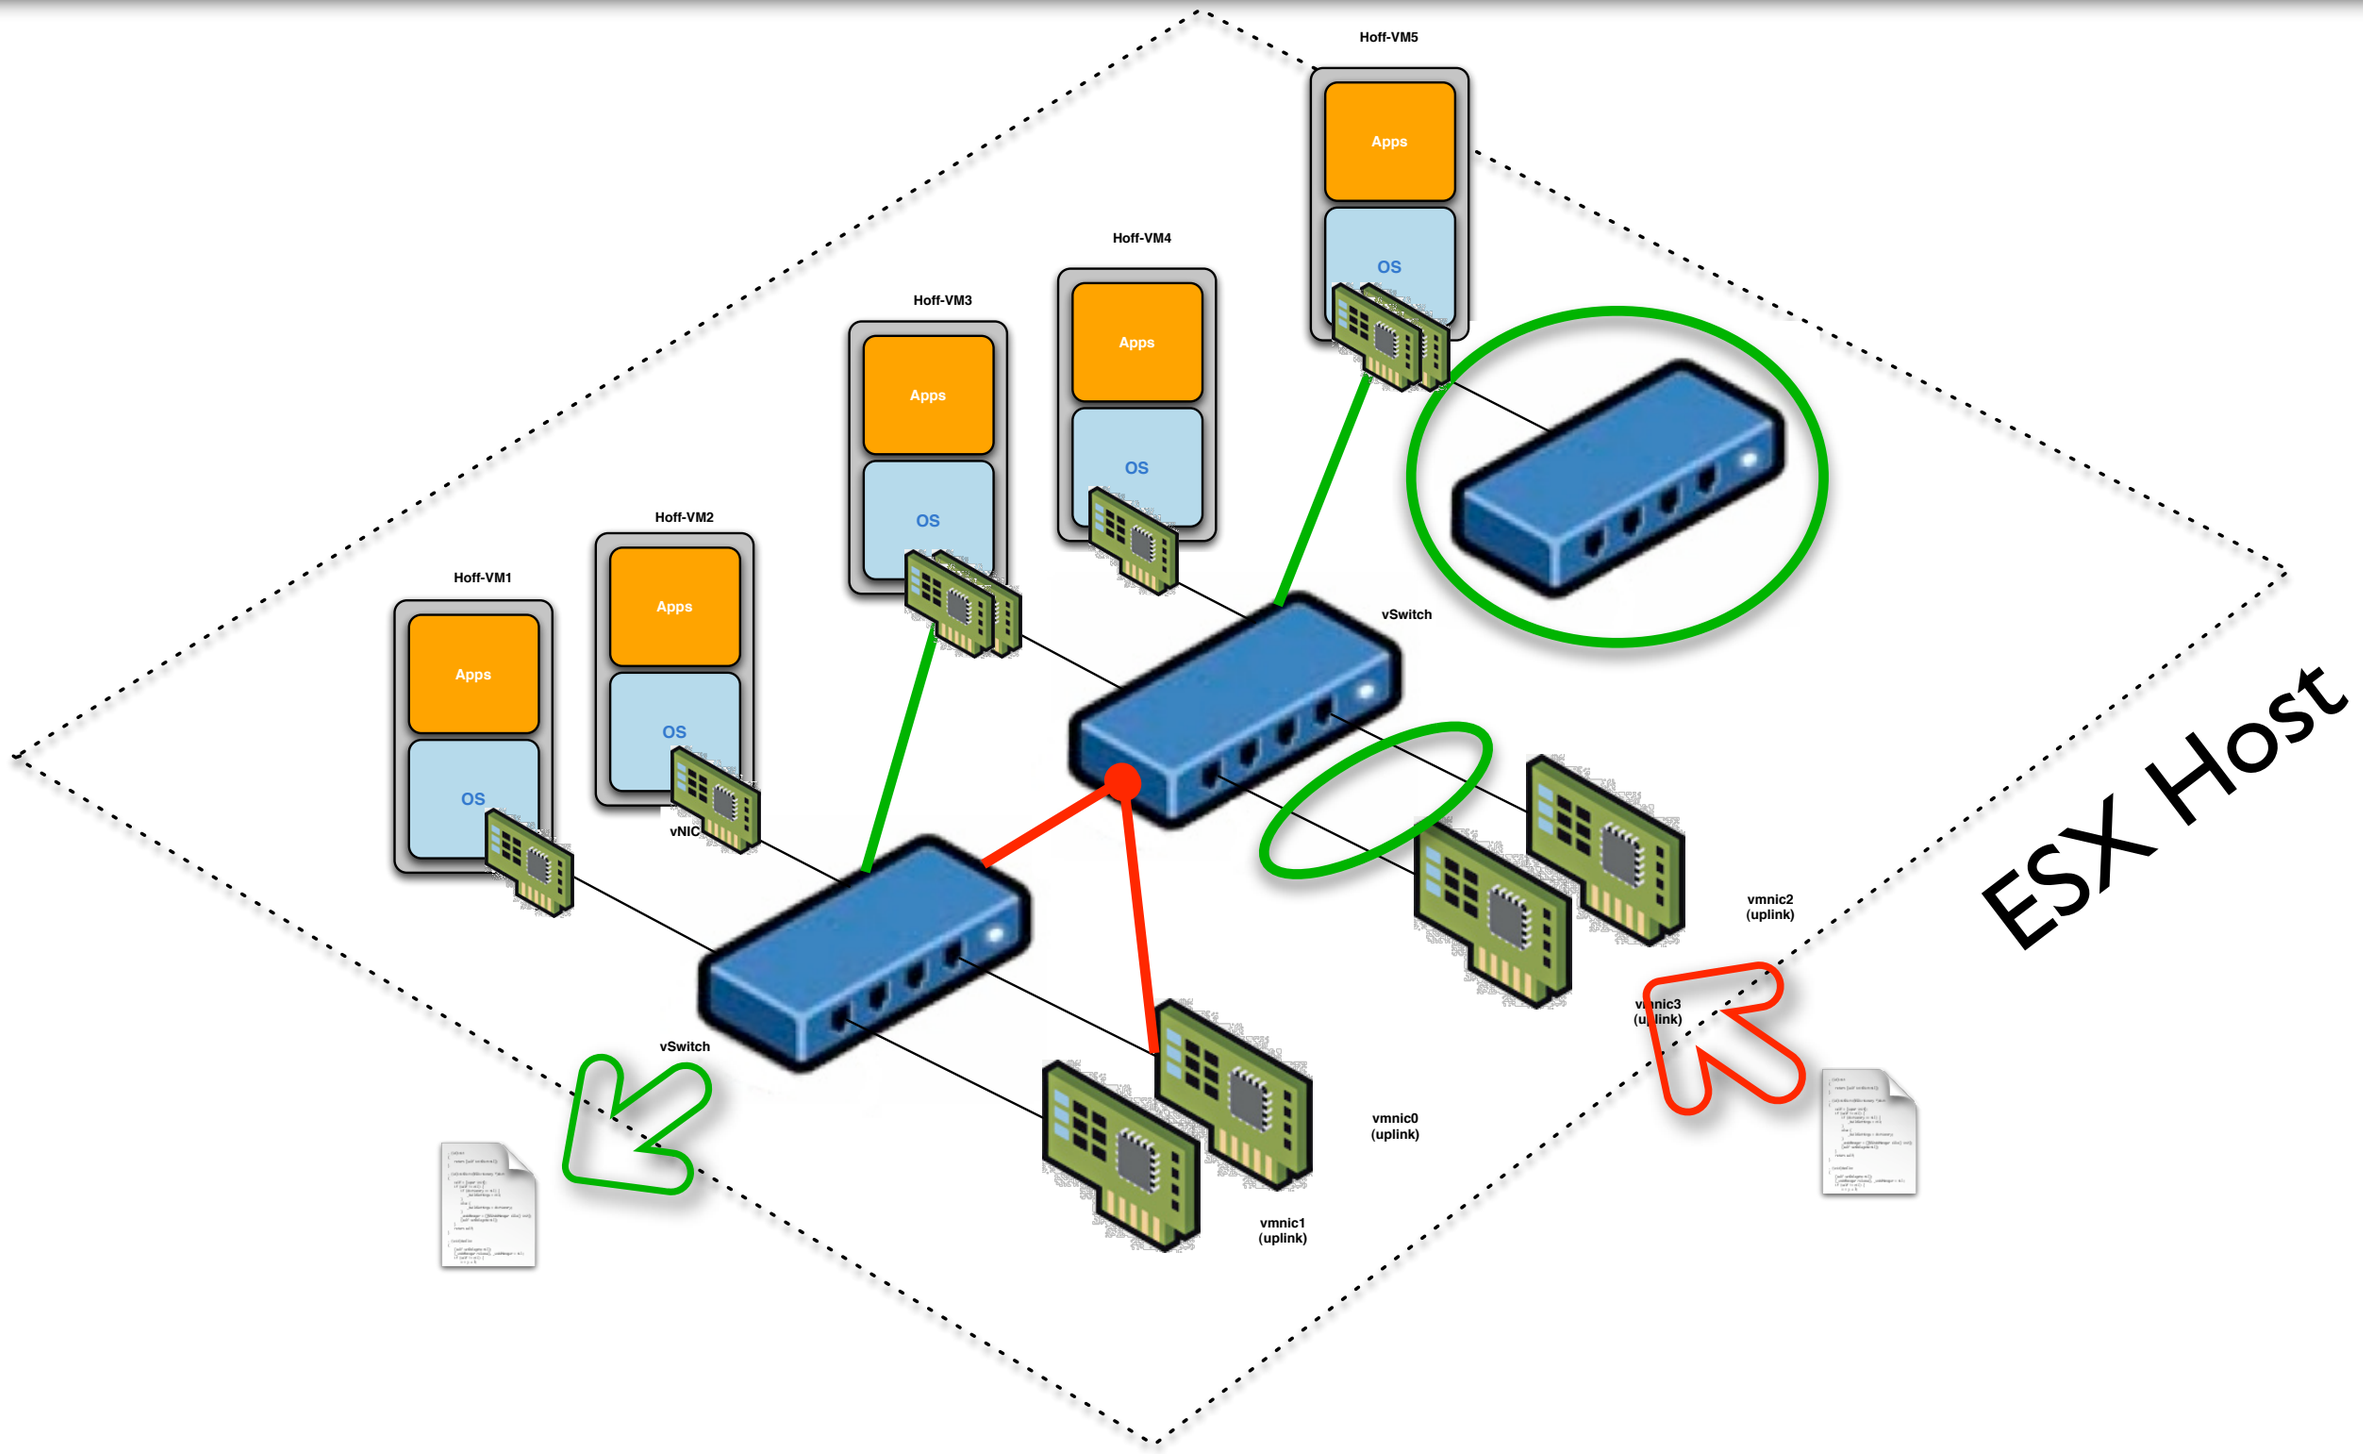

Comparing vSwitches to pSwitches

Similar

- It's a basic Layer-2 switch
- vSwitches maintain MAC forwarding tables & perform frame destination lookup and forwarding
- vSwitches support VLAN segmentation per port (access/trunk)
- Supports copying packets to a mirror port via promiscuous mode

Dissimilar

- Cannot cascade vSwitches
- vSwitches do not learn from the network to populate forwarding tables; no learning of unicast addresses and no IGMP snooping to learn multicast group membership
- vSwitches make private copies of frame data used to make forwarding or filtering decisions
- Frame data is carried outside the frame as it passes through the virtual switch
- vSwitches have no dynamic trunking protocol support such as STP and therefore enforce a single-tier network topology

vSwitch Security Options

vSwitches offer some nifty security features:

- Configure promiscuous mode (per portgroup) for selective mirroring
- MAC Address changes prevents VM's from changing/ spoofing their MAC addresses
- Can restrict "forged transmissions" that would potentially allow VM's to send traffic from nodes other than themselves

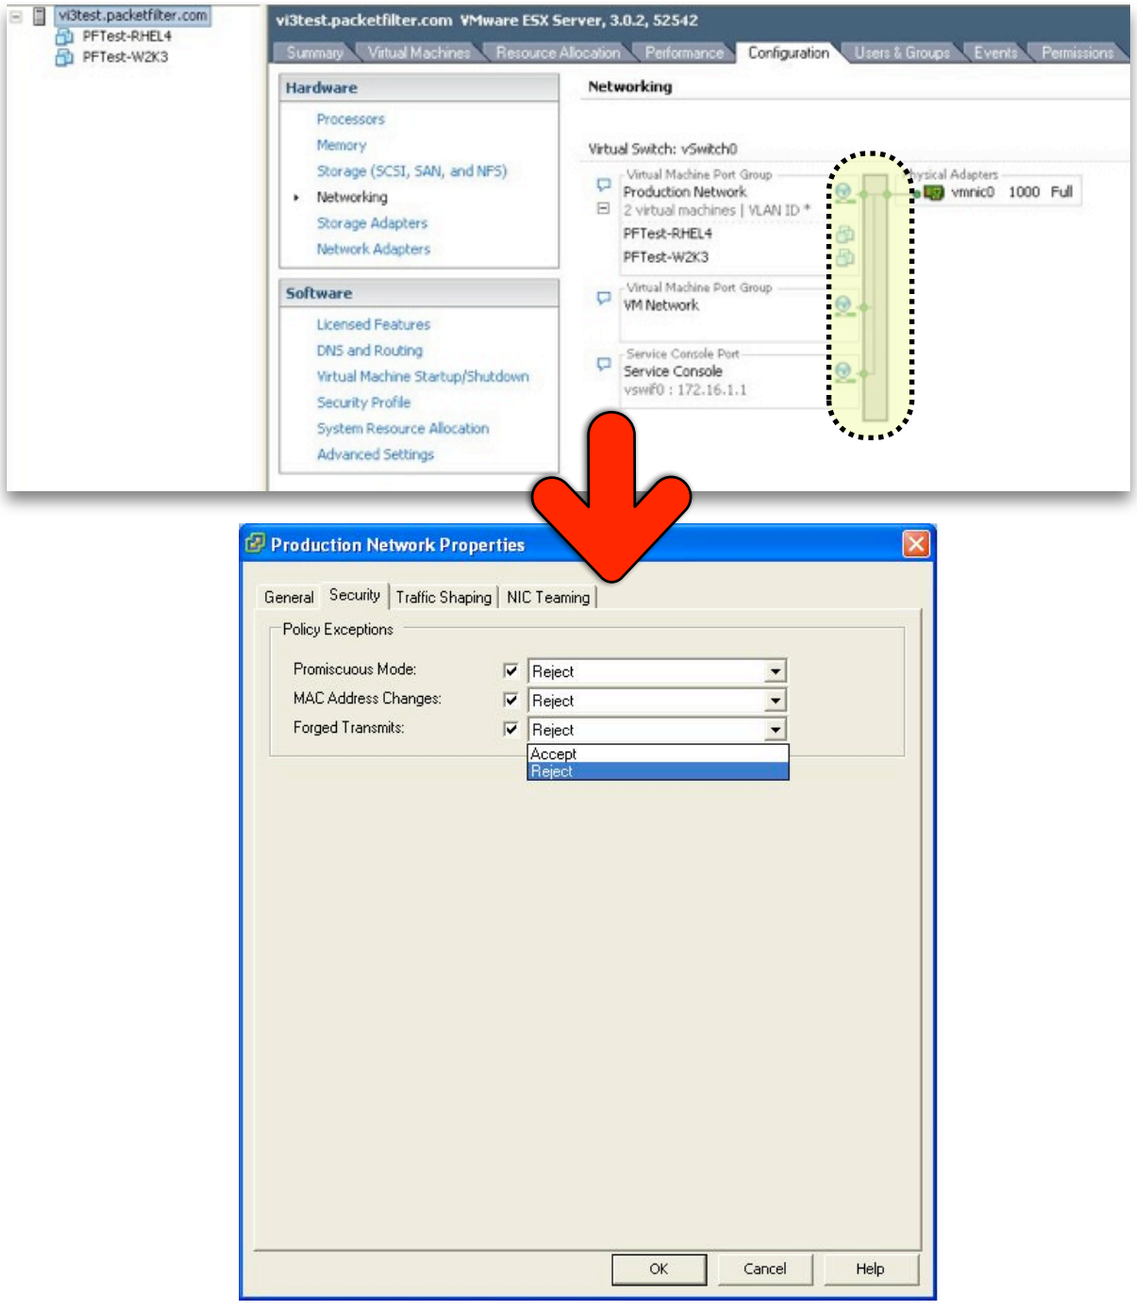

You're Making Me All Weepy!

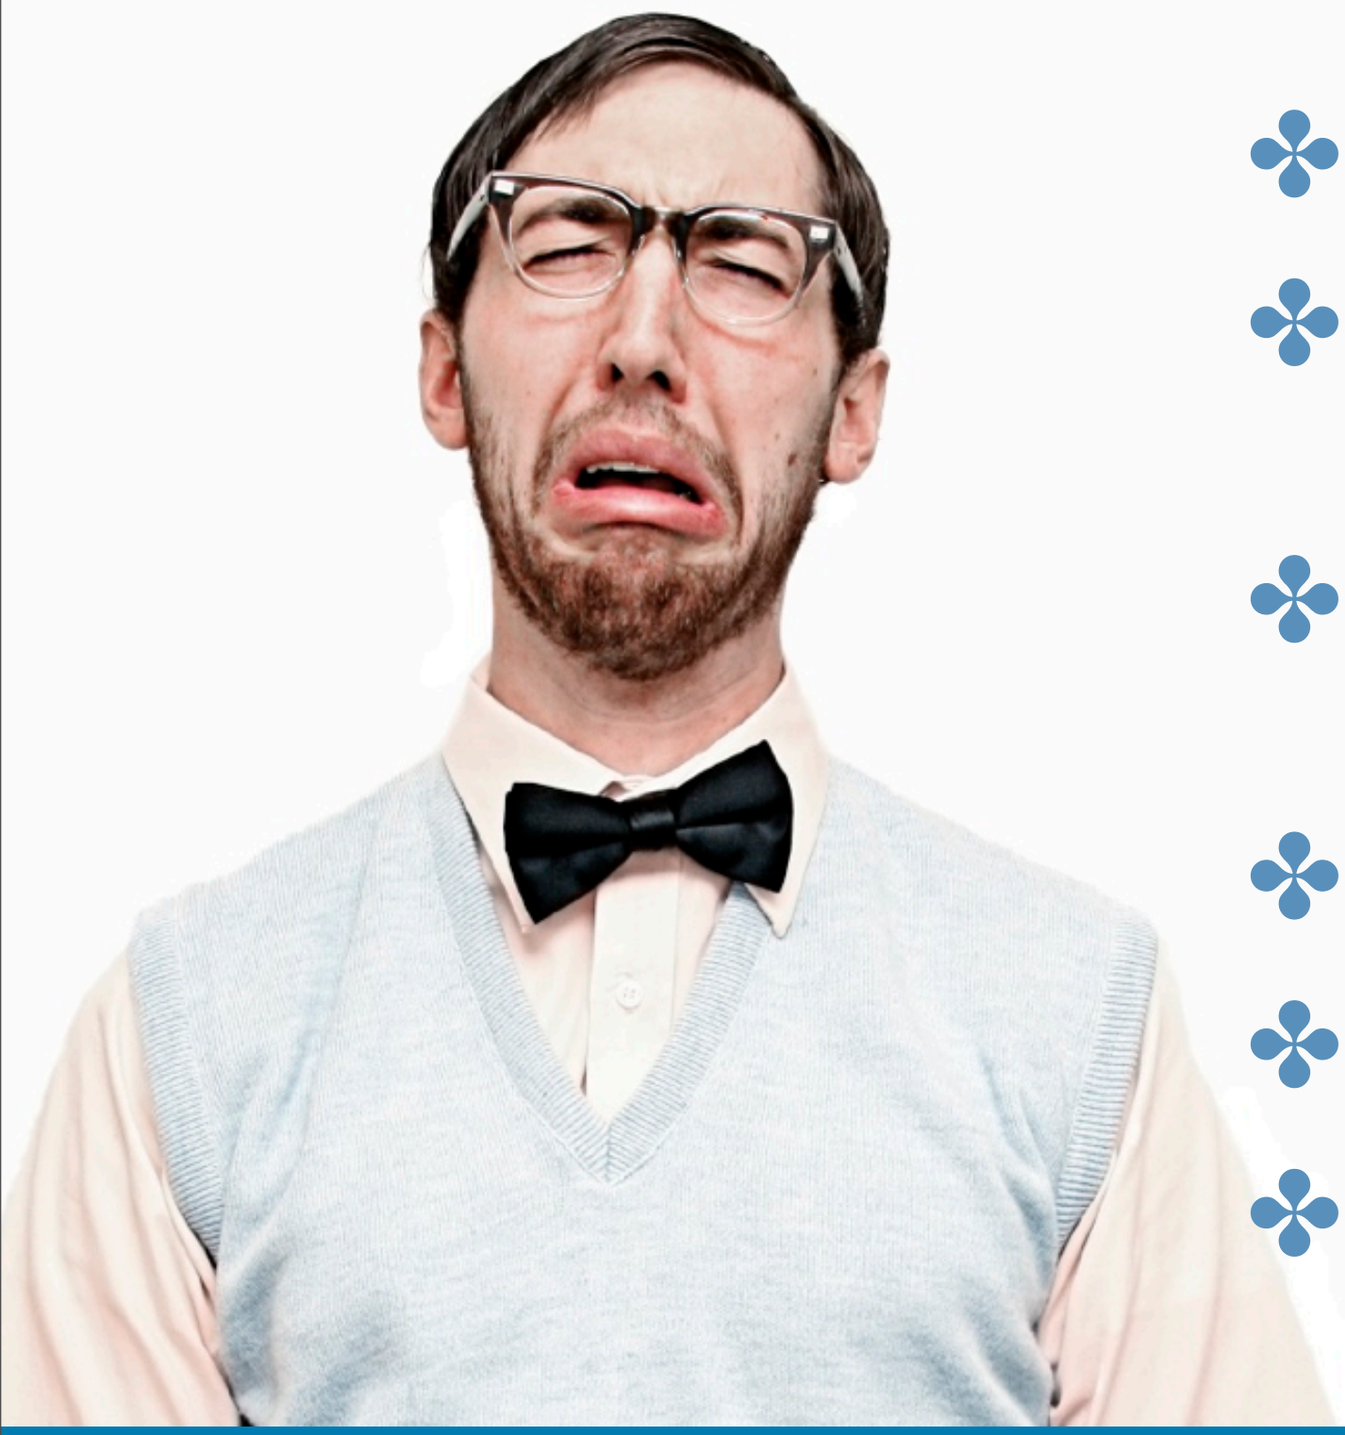

- Setup & Context
- x86 Virtualization Overview in 90 Seconds
- Virtual Networking Architecture
- VirtSec Solutions Landscape
- The Four Horsemen
- Wrap-Up

VirtSec Technology Landscape

- Evolving solutions from existing players as well as emerging startups & the virtualization platform providers
- You will need to invest differently in order to effectively manage risk in a virtualized environment
- The next 12-18 months will be difficult due to the gold rush effect
- There is (still) no silver bullet, just a lot of silver buckshot

Threat Models

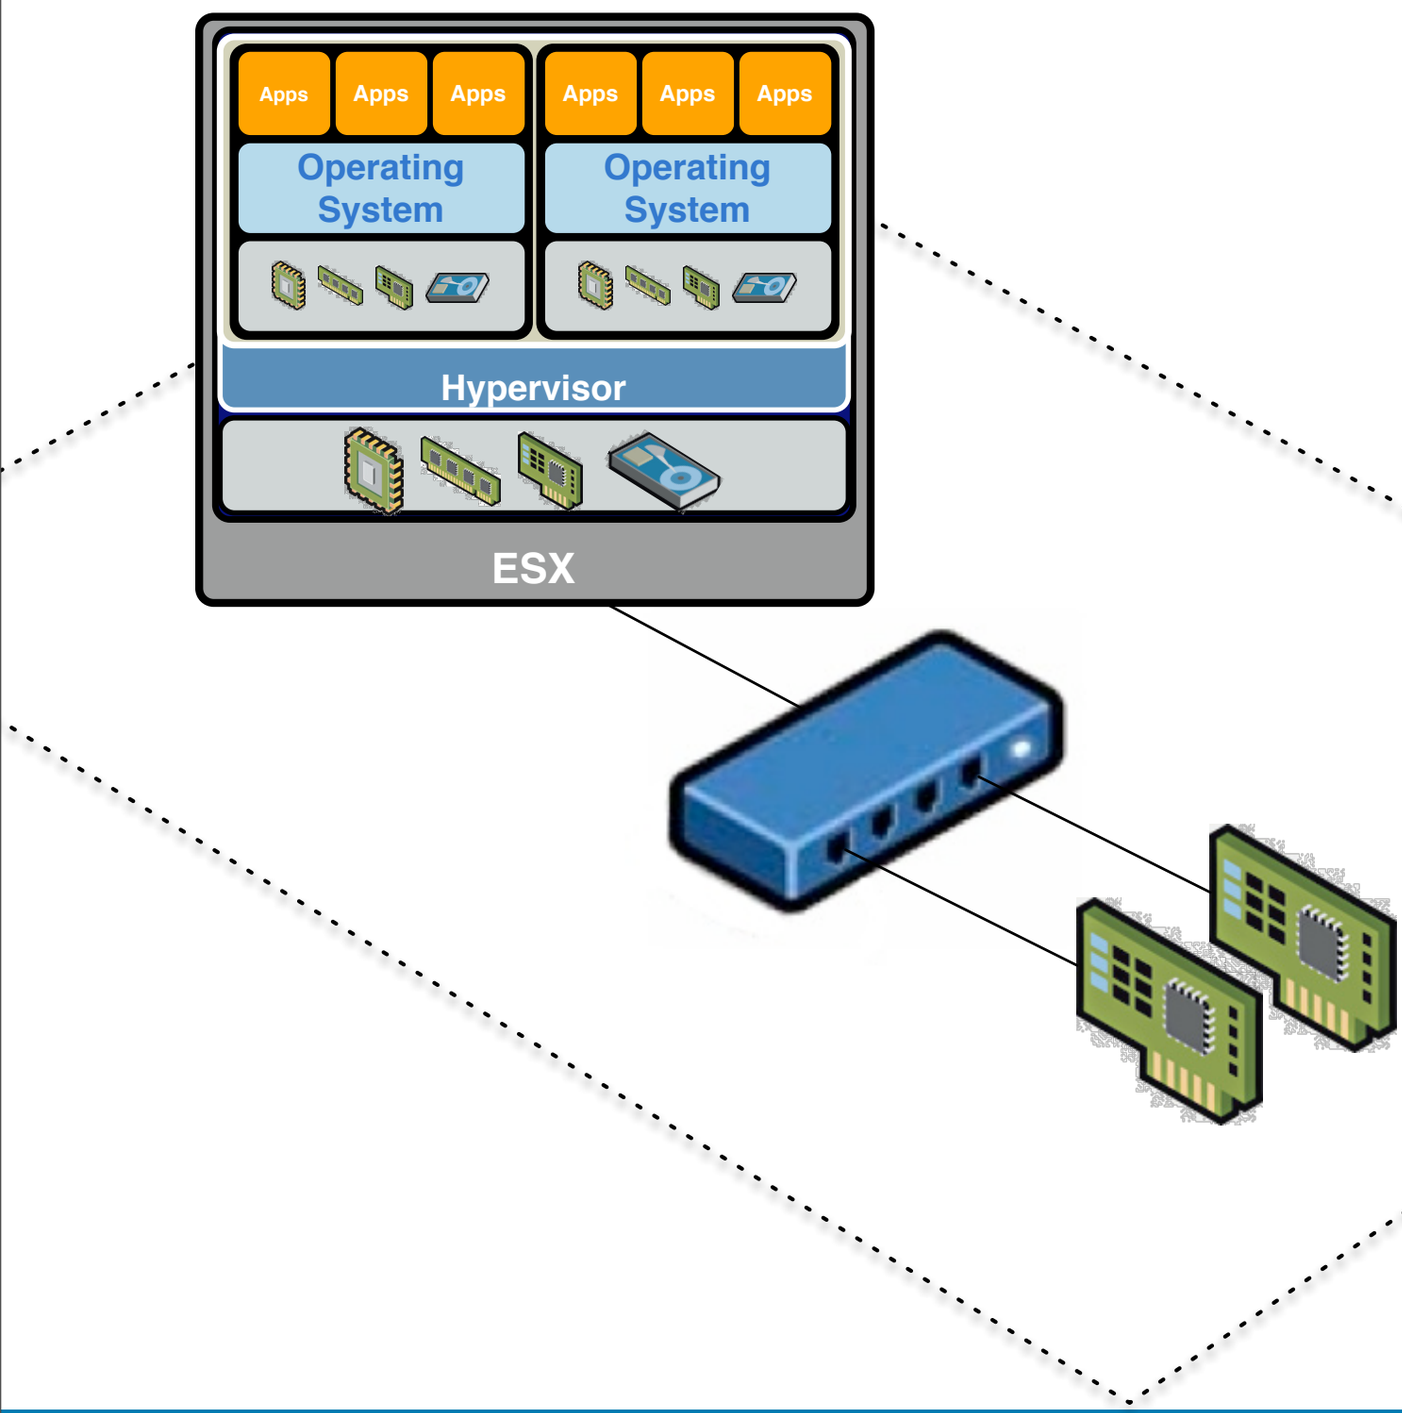

1. Guest to Guest
2. Guest to Host
3. Guest to Self
4. External to Host
5. External to Guest

Threat Models

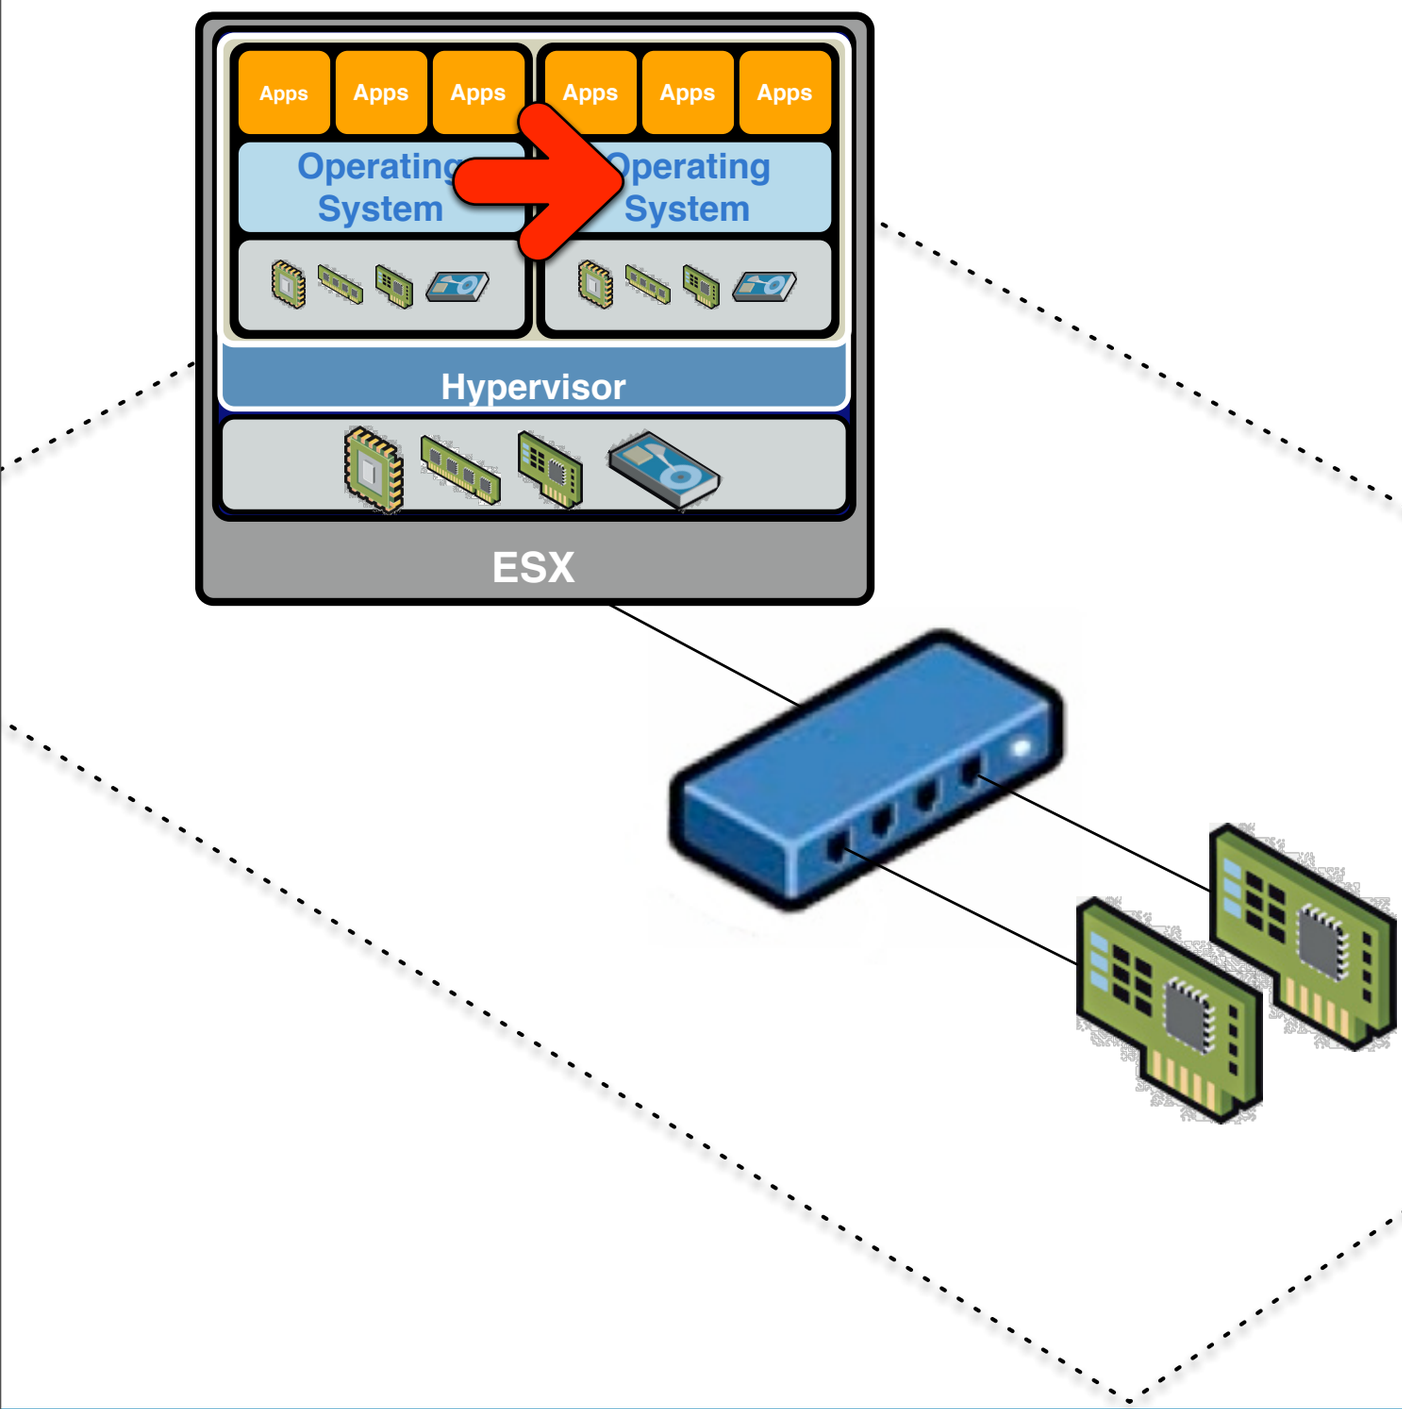

1. Guest to Guest
2. Guest to Host
3. Guest to Self
4. External to Host
5. External to Guest

Threat Models

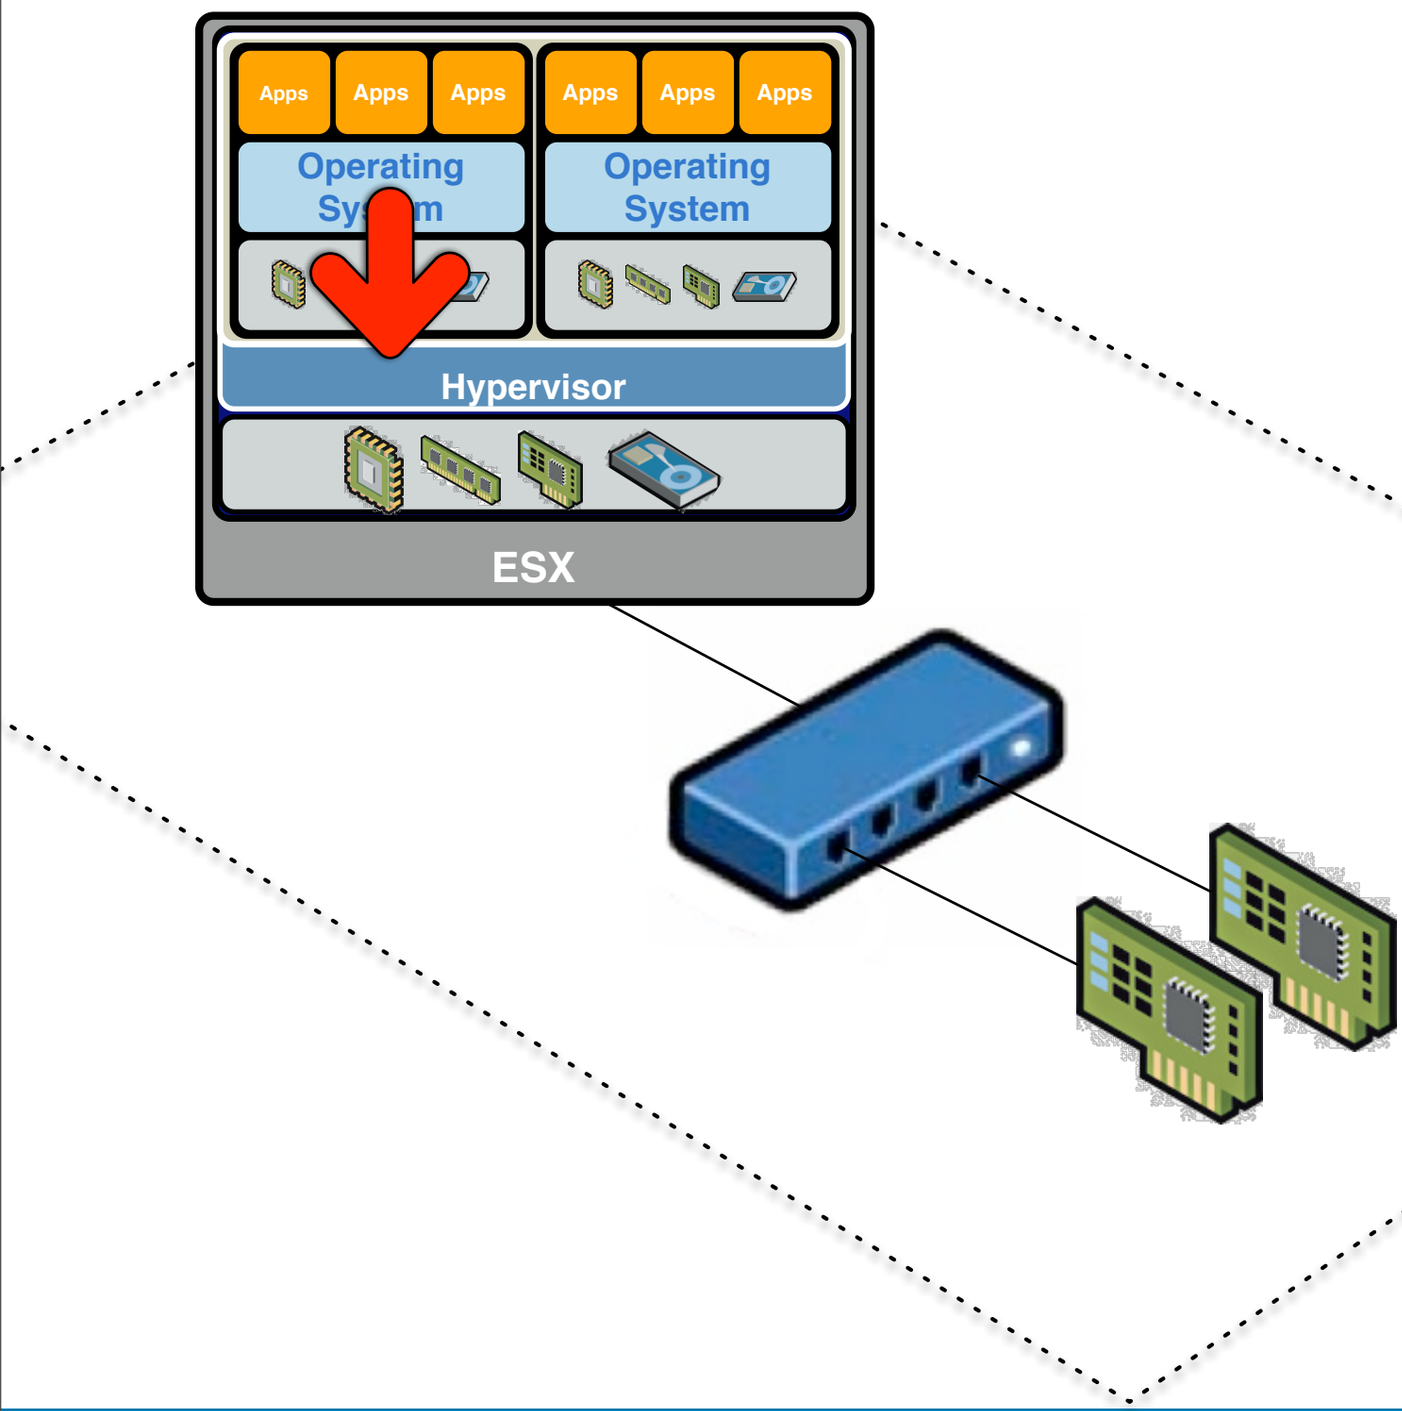

1. Guest to Guest
2. Guest to Host
3. Guest to Self
4. External to Host
5. External to Guest

Threat Models

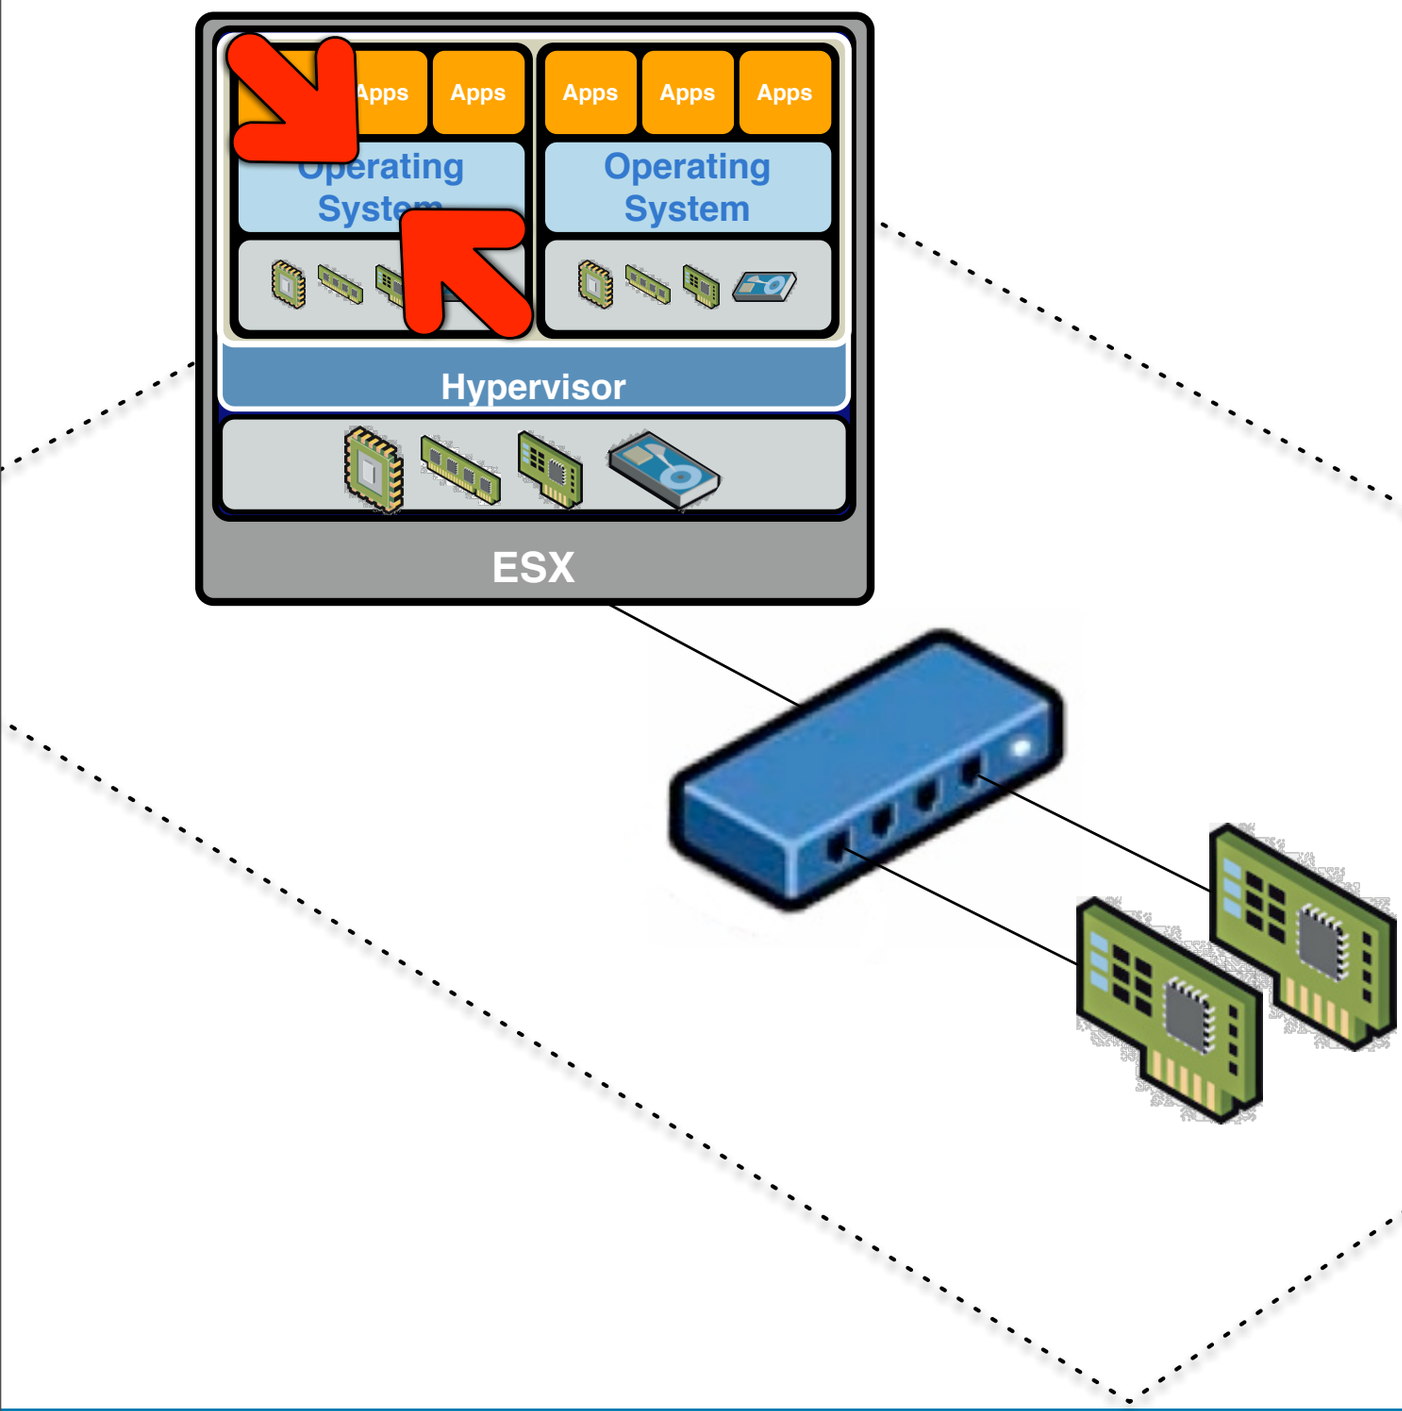

1. Guest to Guest
2. Guest to Host
3. Guest to Self
4. External to Host
5. External to Guest

Threat Models

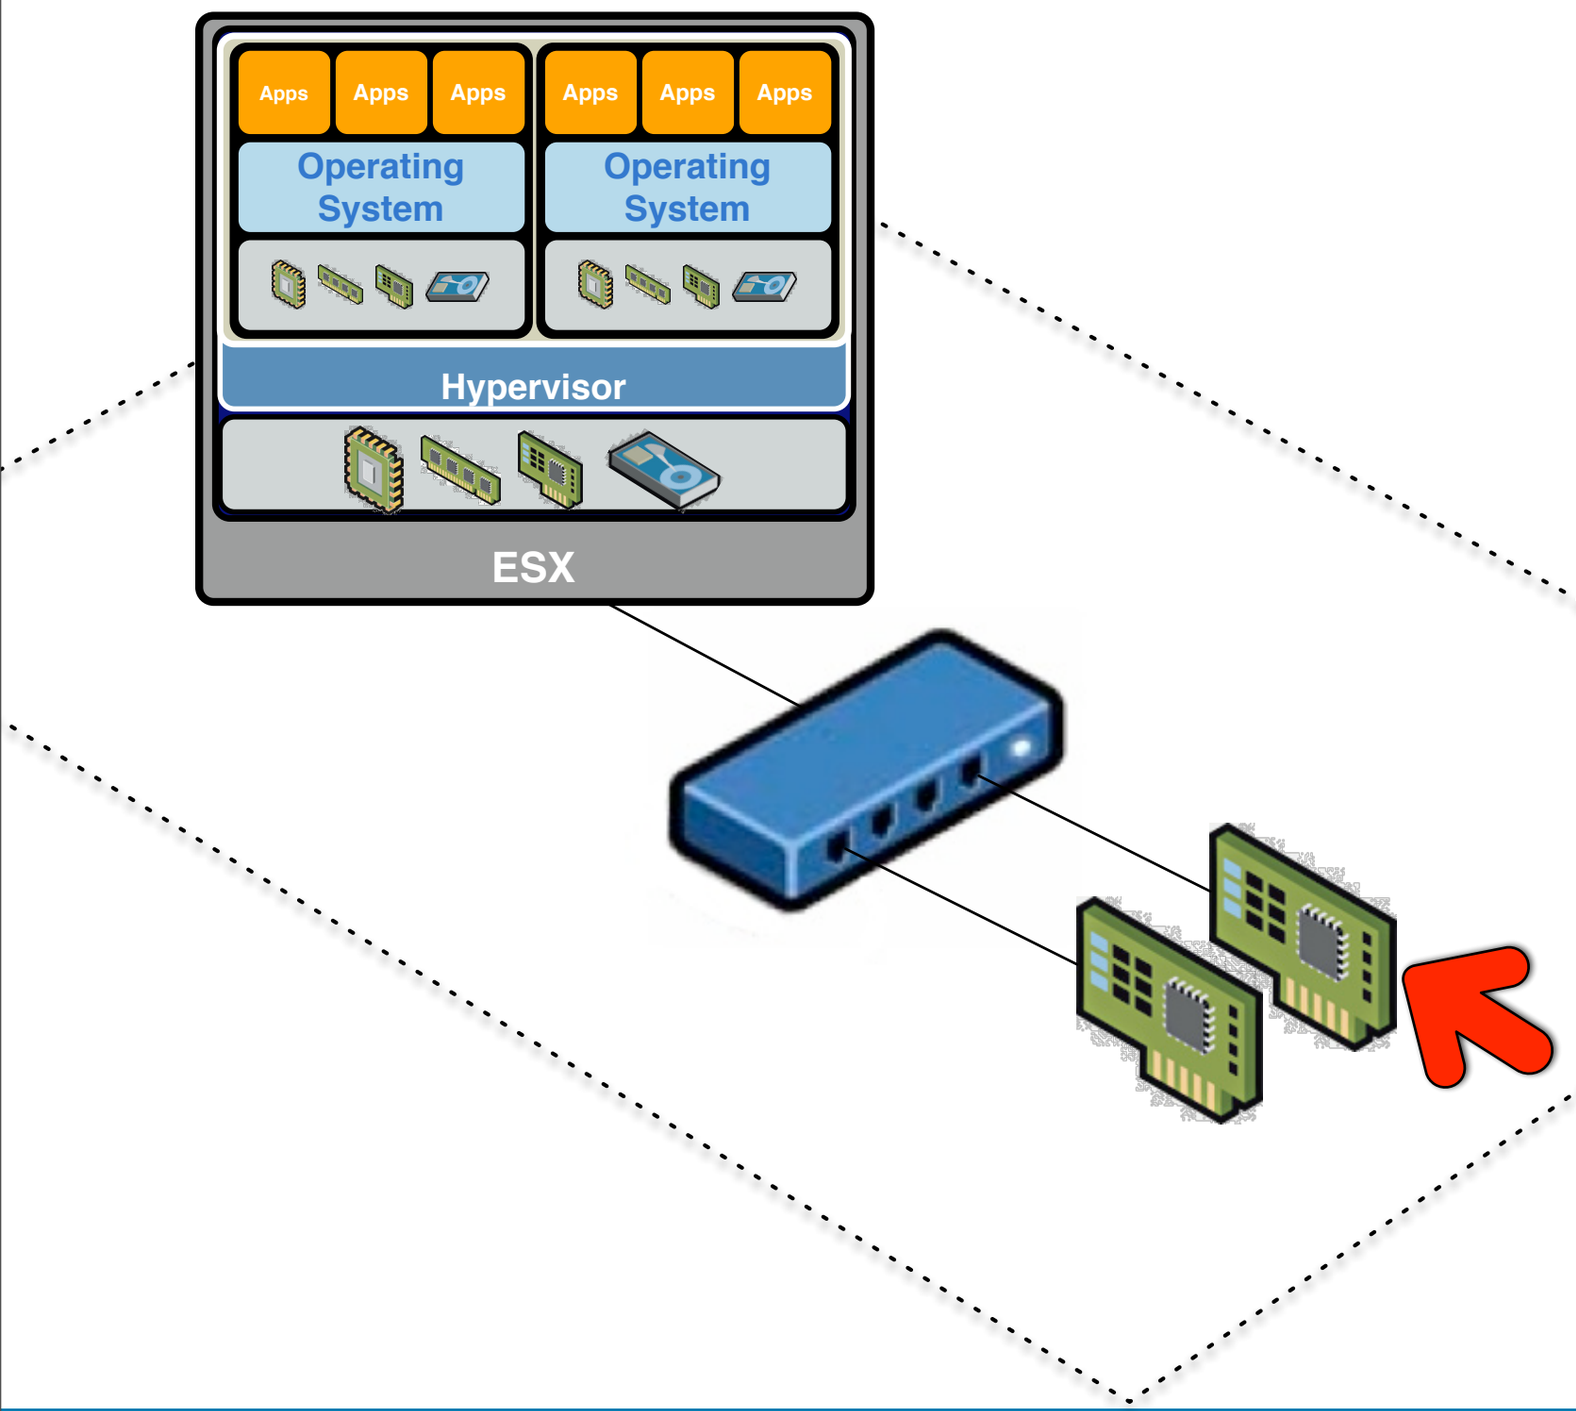

1. Guest to Guest
2. Guest to Host
3. Guest to Self
4. External to Host
5. External to Guest

Threat Models

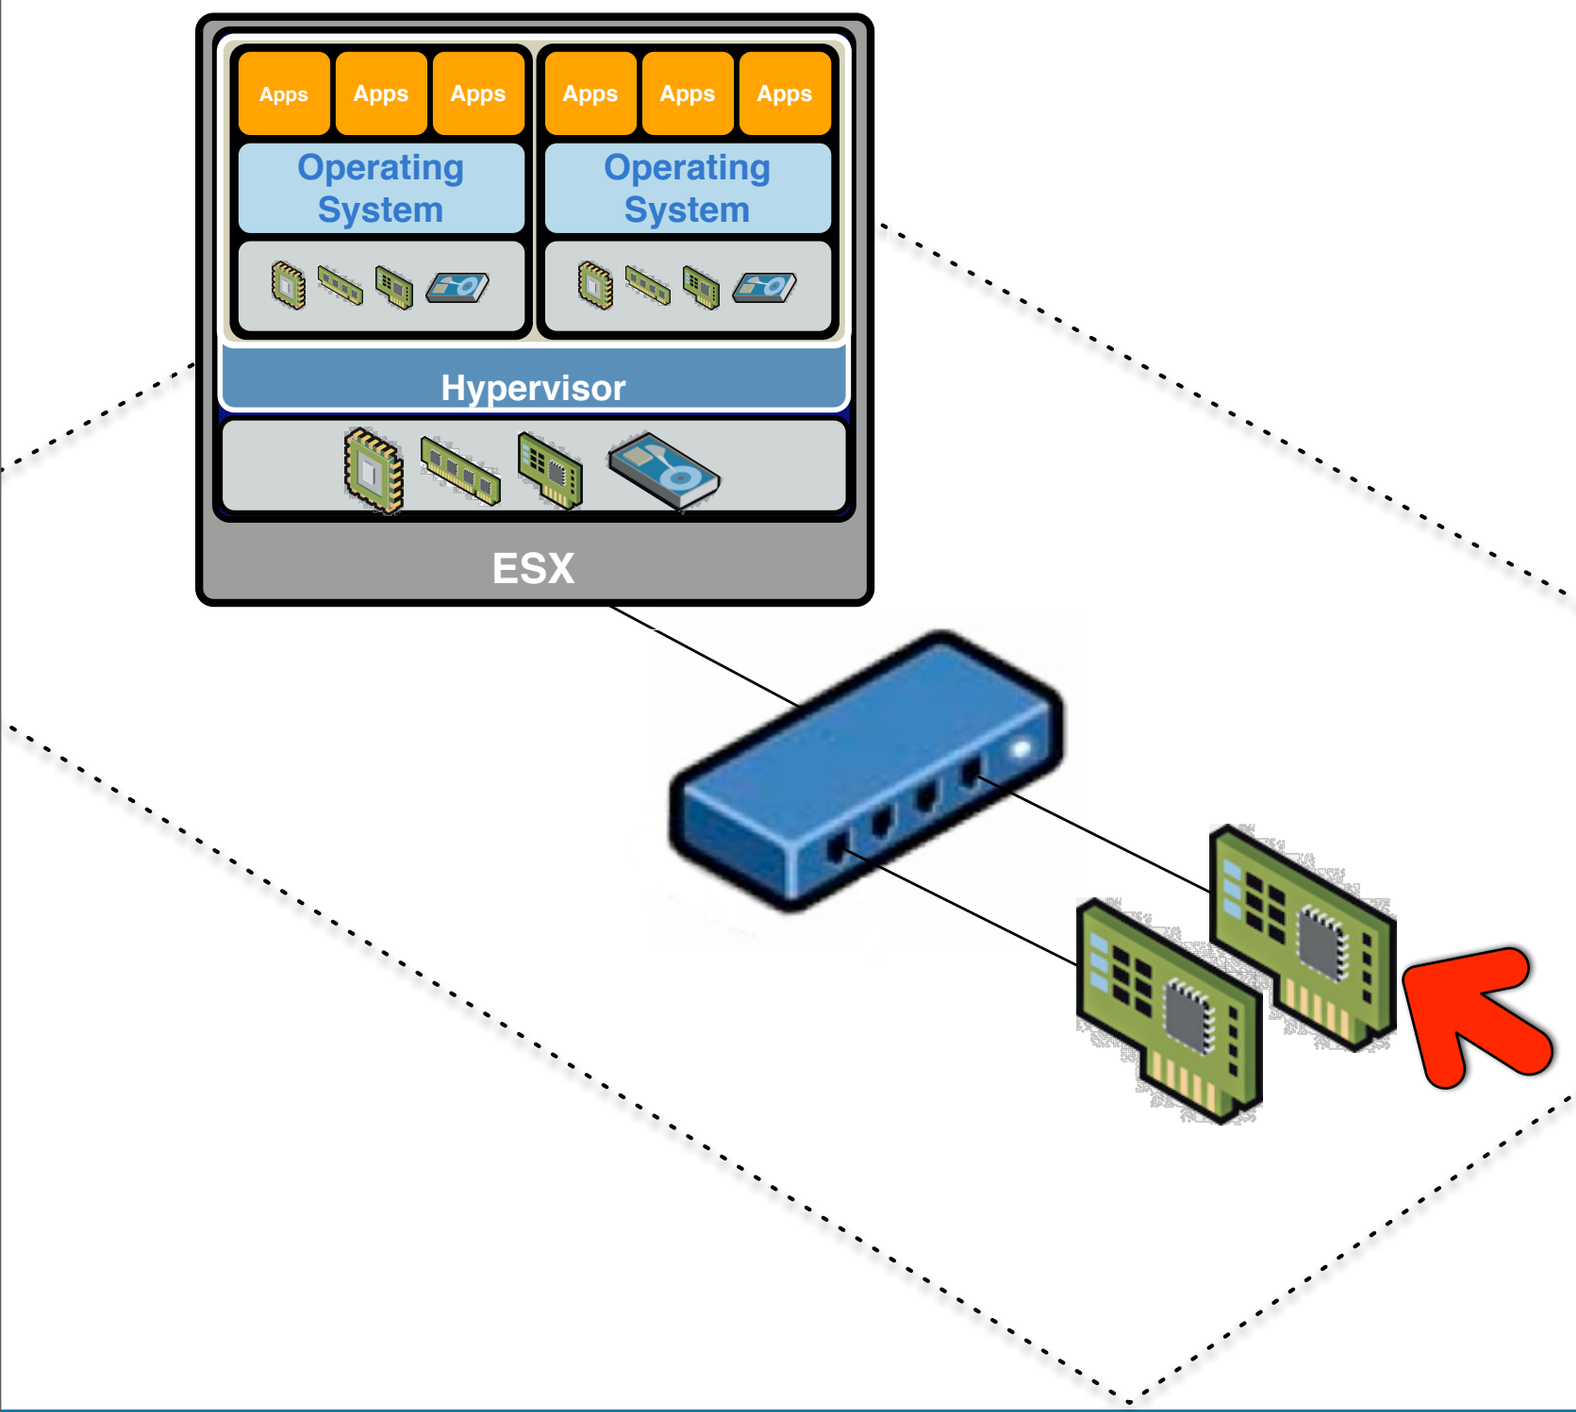

1. Guest to Guest
2. Guest to Host
3. Guest to Self
4. External to Host
5. External to Guest

VirtSec Examples: No Controls (?)

No Security

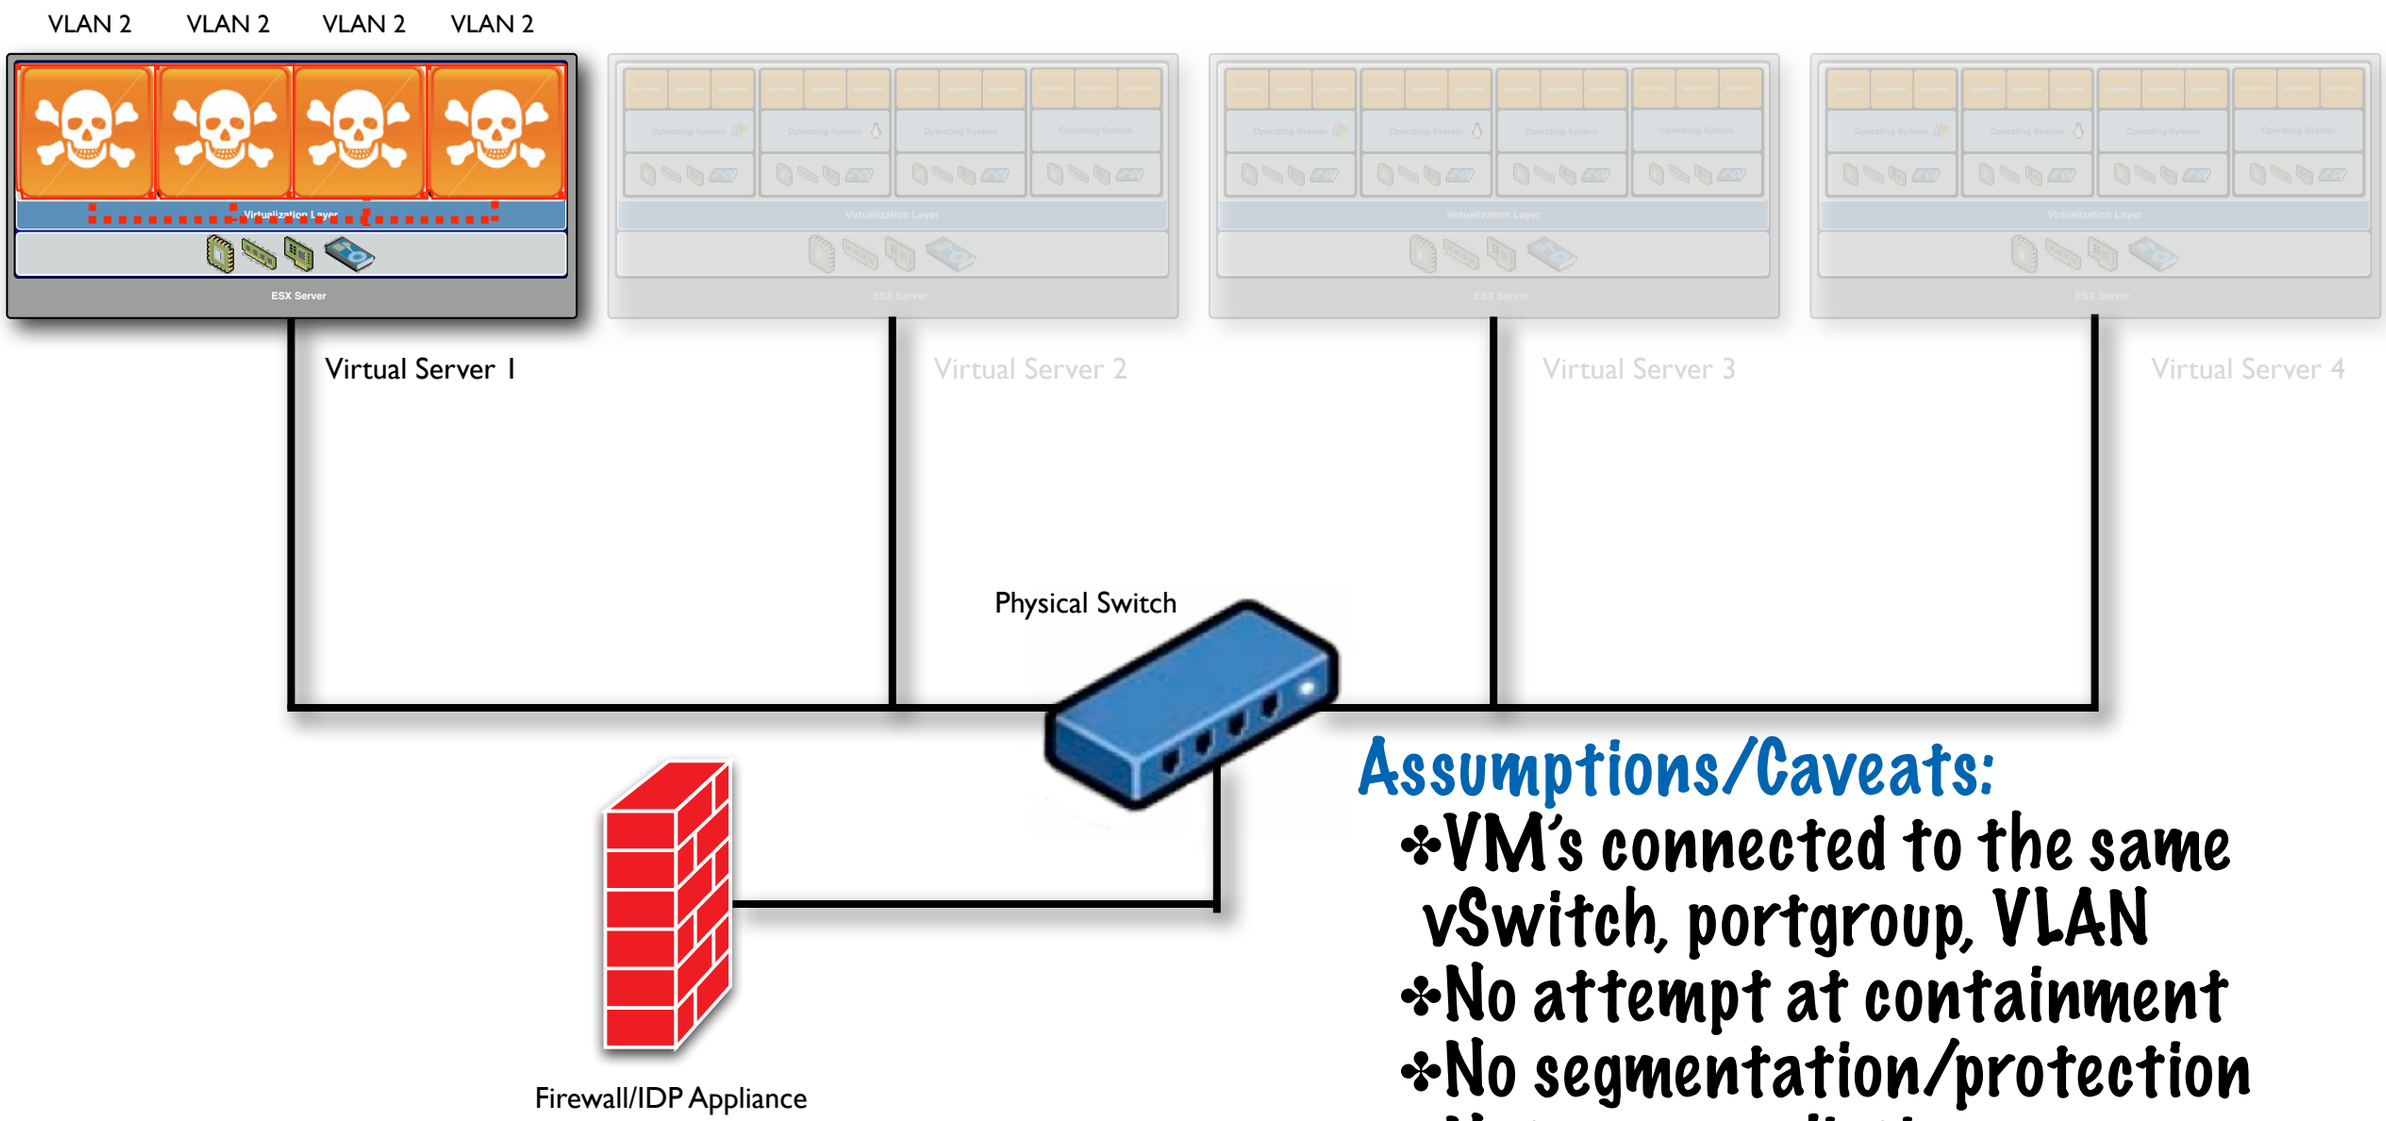

Assumptions/Caveats:

- VM's connected to the same vSwitch, portgroup, VLAN
- No attempt at containment
- No segmentation/protection
- Not very realistic...

VirtSec Examples: Software On the VM

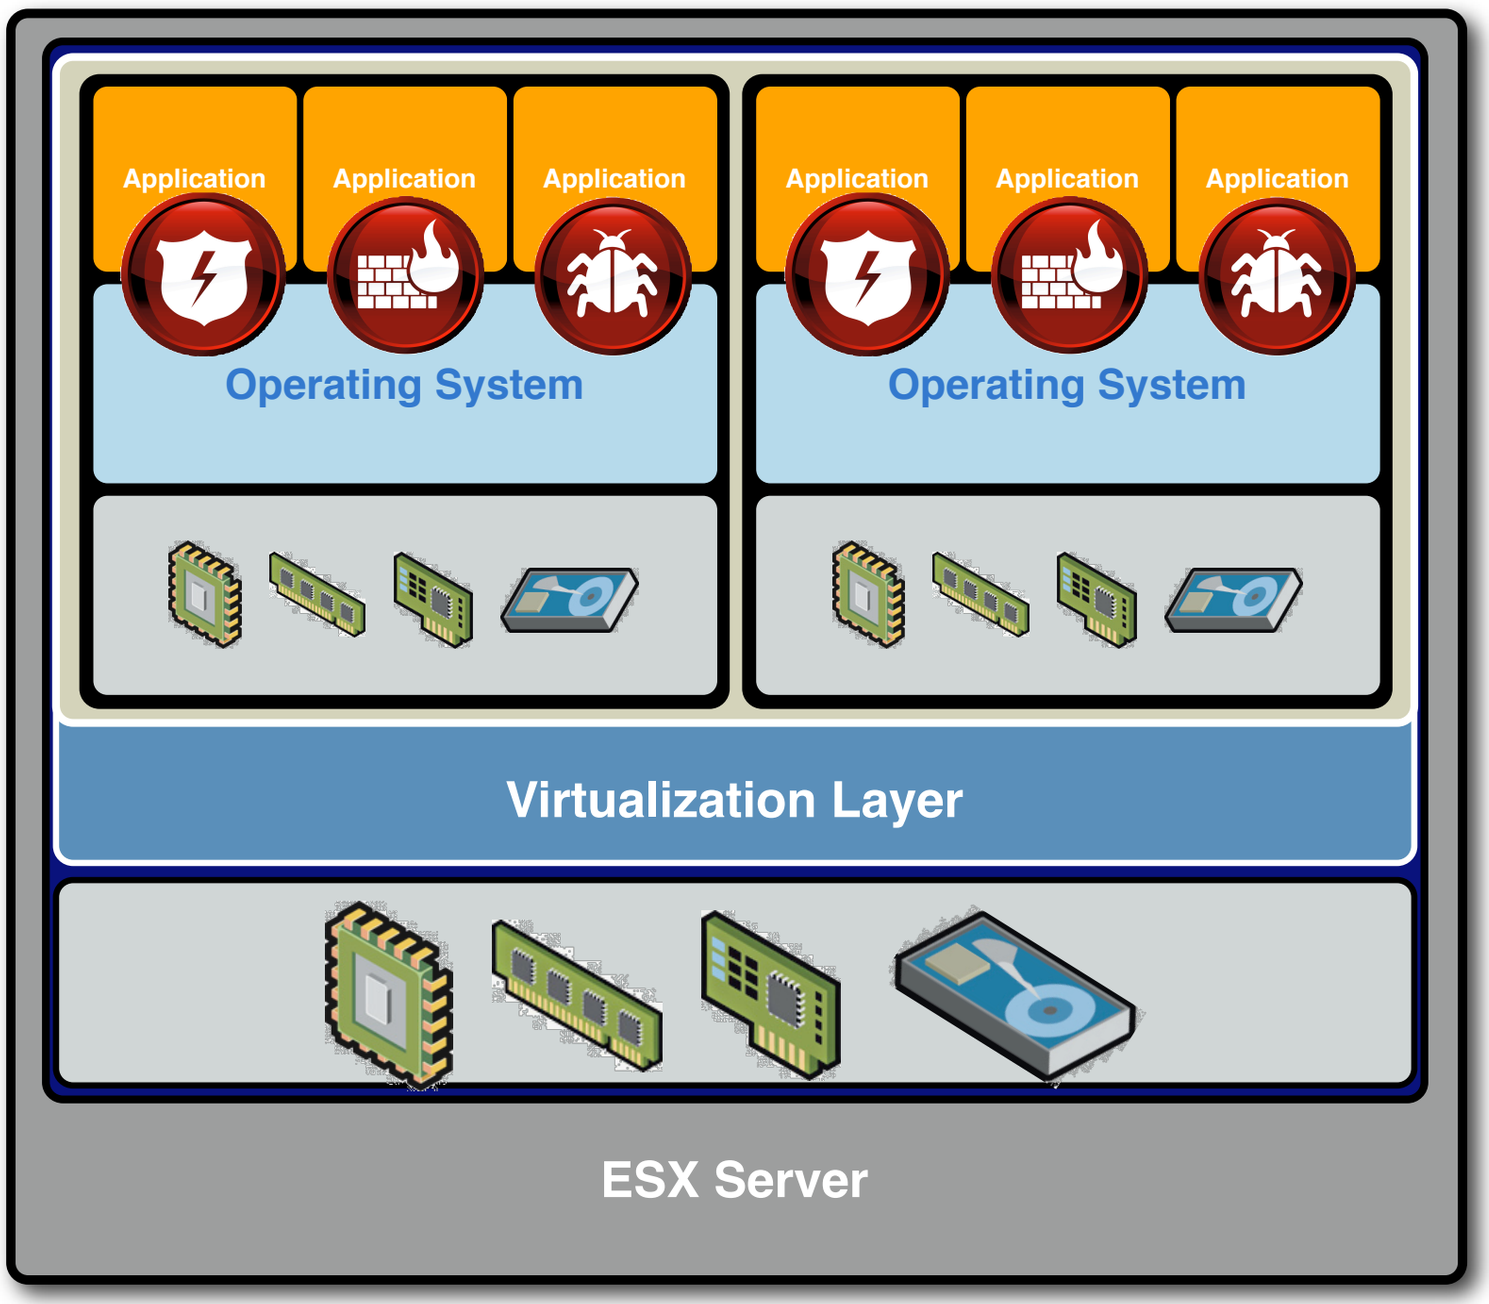

Most anything you run today in your conventional environments will work here...

- Firewalls
- HIDS
- HIPS
- Anti-virus
- NAC
- Endpoint Assurance
- Patch Management
- Inventory
- Configuration Audit & Control
- Insert $product here

McAfee · Check Point (We Secure the Internet) · symantec. · CISCO
tripwire · Configuresoft · BIGFIX
TREND MICRO · netIQ · Microsoft

Reality Check: The Intra-VM (In)Security Myth

Myth/Security Team Says:

- "Consolidating servers onto the same virtualized host is insecure because you can't secure intra-vm traffic!"

Reality/I ask:

- "When you have two physical servers plugged into the same physical switch in the same VLAN, how do you secure intra-machine traffic?"

Response/Security Team Blushes:

- "Uh, we don't..."

Virtualized

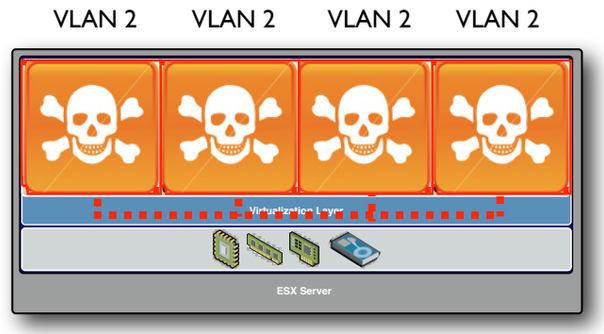

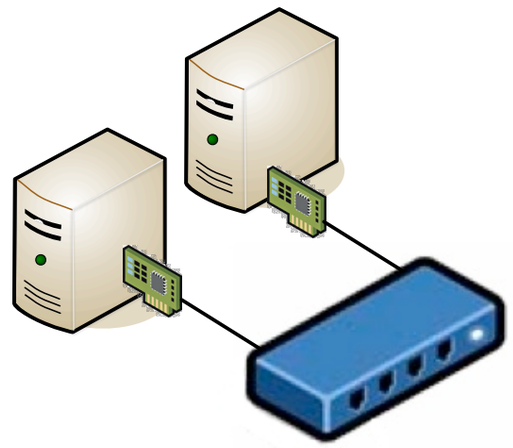

Real World

Doh!

VirtSec Examples: Interacting with External Security

External Security

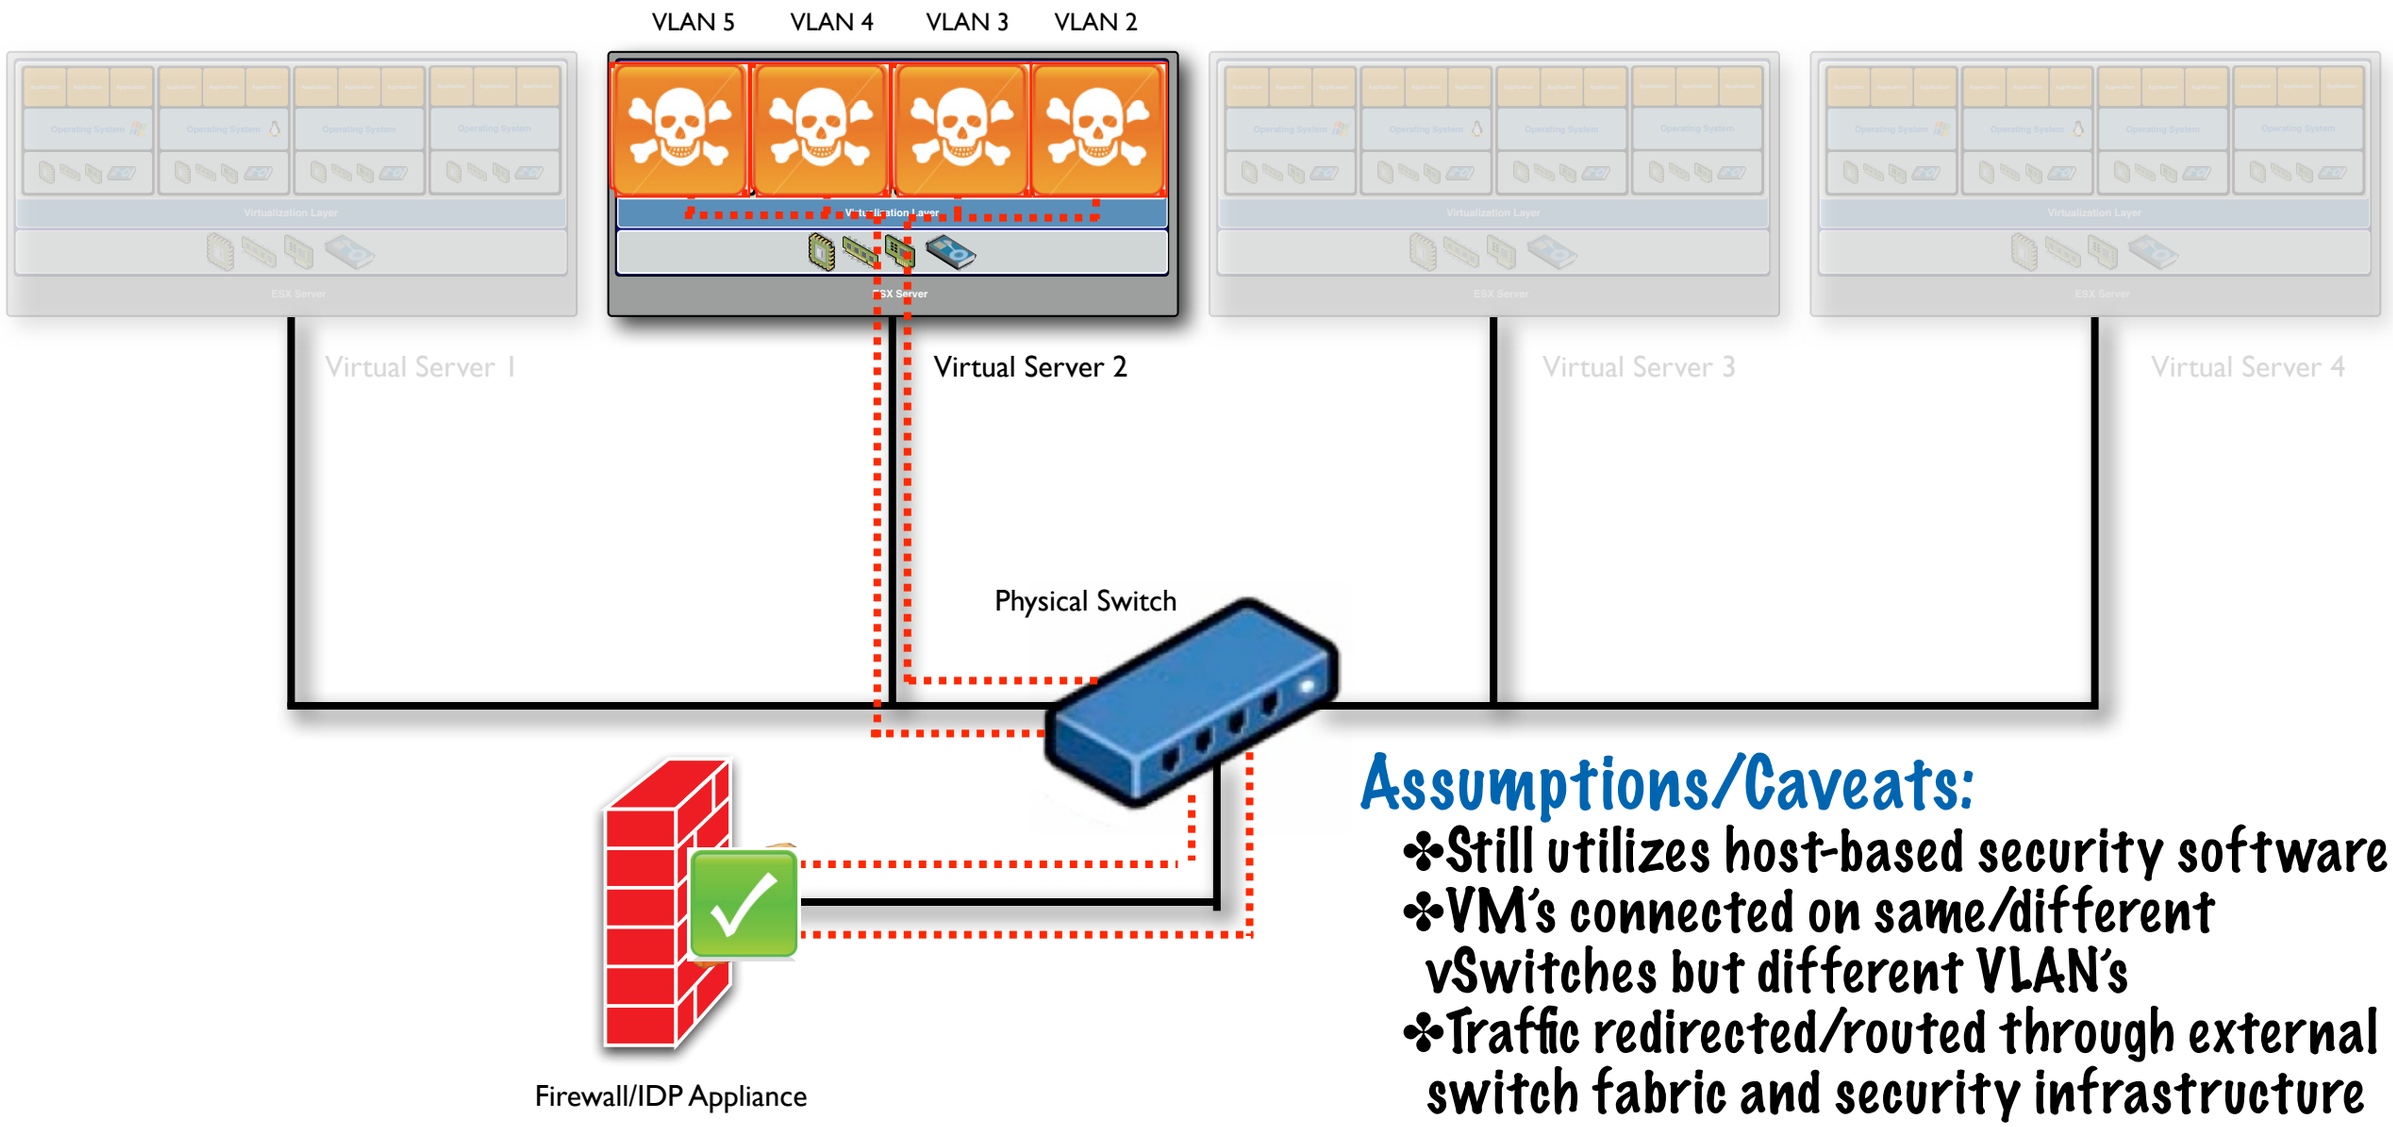

Assumptions/Caveats:

- Still utilizes host-based security software
- VM's connected on same/different vSwitches but different VLAN's
- Traffic redirected/routed through external switch fabric and security infrastructure

VirtSec Examples: Virtual Appliance w/External Security Interaction & VM to VM On Different VLAN/vSwitch

Virtual Appliances

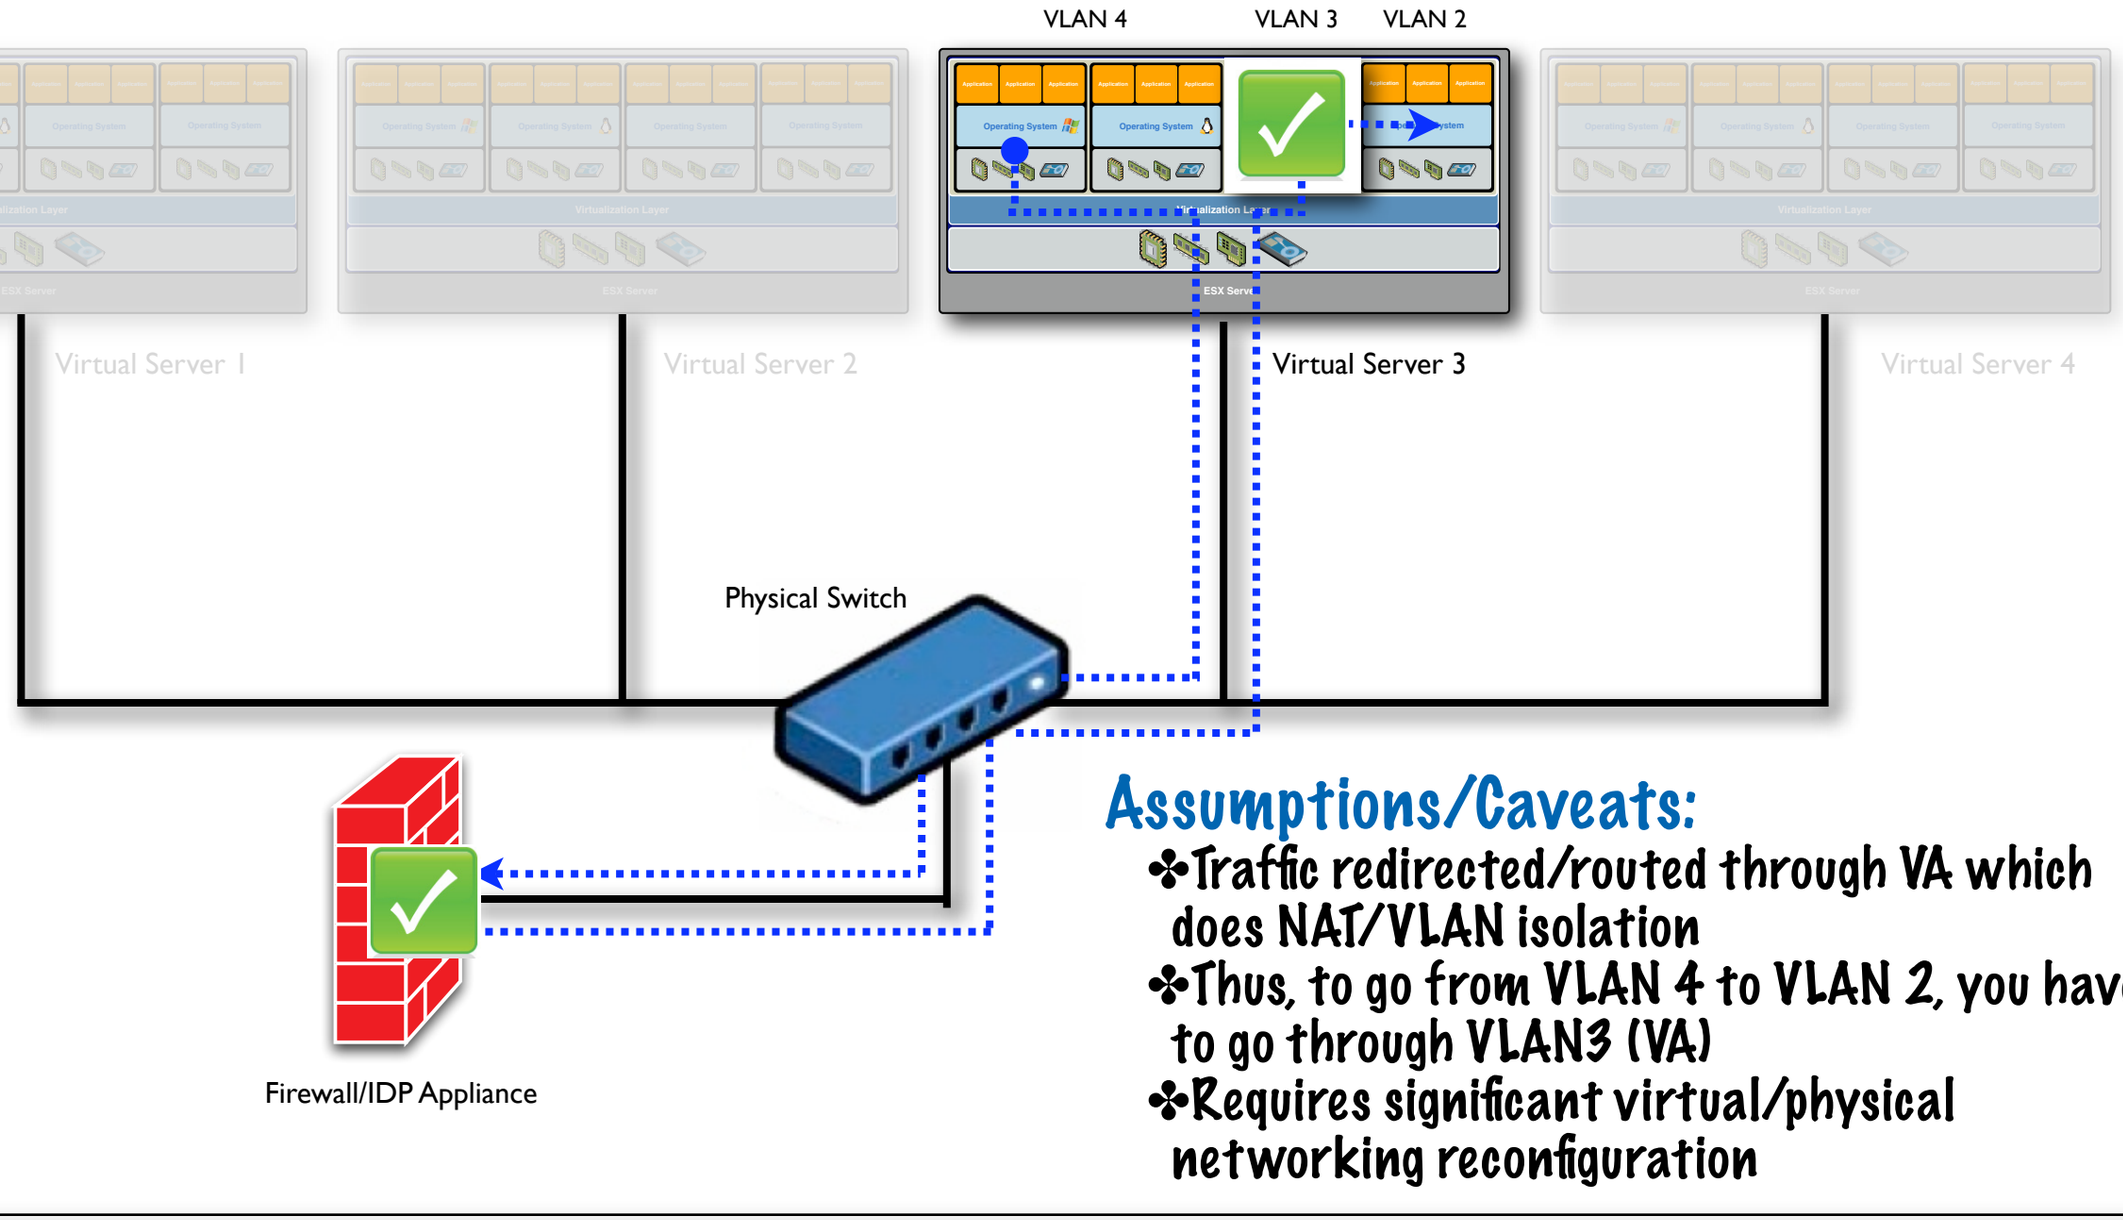

Assumptions/Caveats:

- Traffic redirected/routed through VA which does NAT/VLAN isolation
- Thus, to go from VLAN 4 to VLAN 2, you have to go through VLAN3 (VA)
- Requires significant virtual/physical networking reconfiguration

VirtSec Examples: Virtual Appliance with VM to VM On Different vSwitch/VLAN

Virtual Appliances

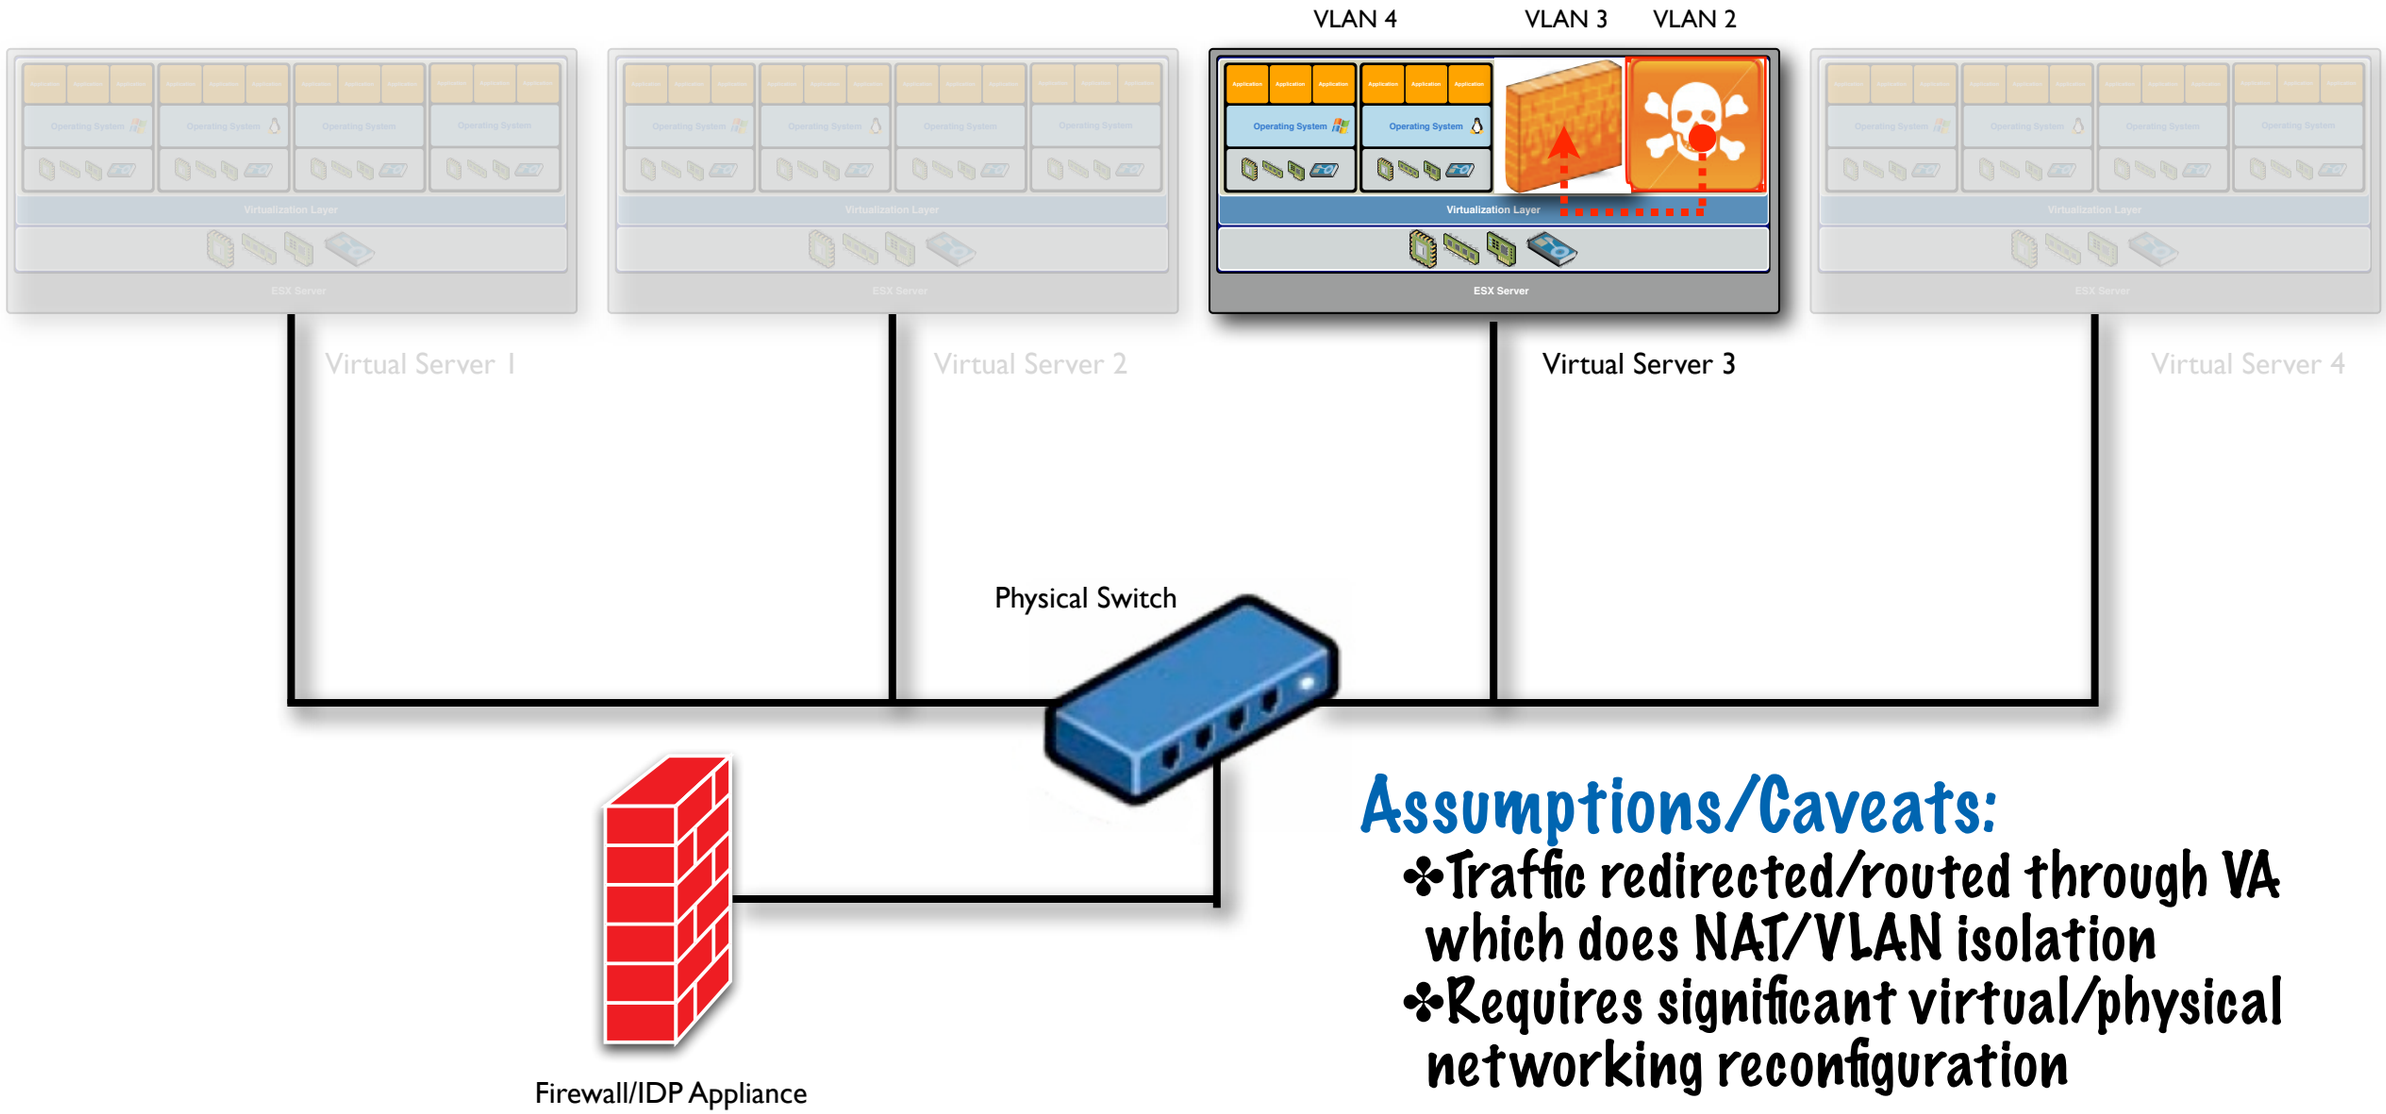

Assumptions/Caveats:

- Traffic redirected/routed through VA which does NAT/VLAN isolation
- Requires significant virtual/physical networking reconfiguration

VirtSec Examples: Virtual Appliance with VM to VM On Same vSwitch/VLAN

Virtual Appliances

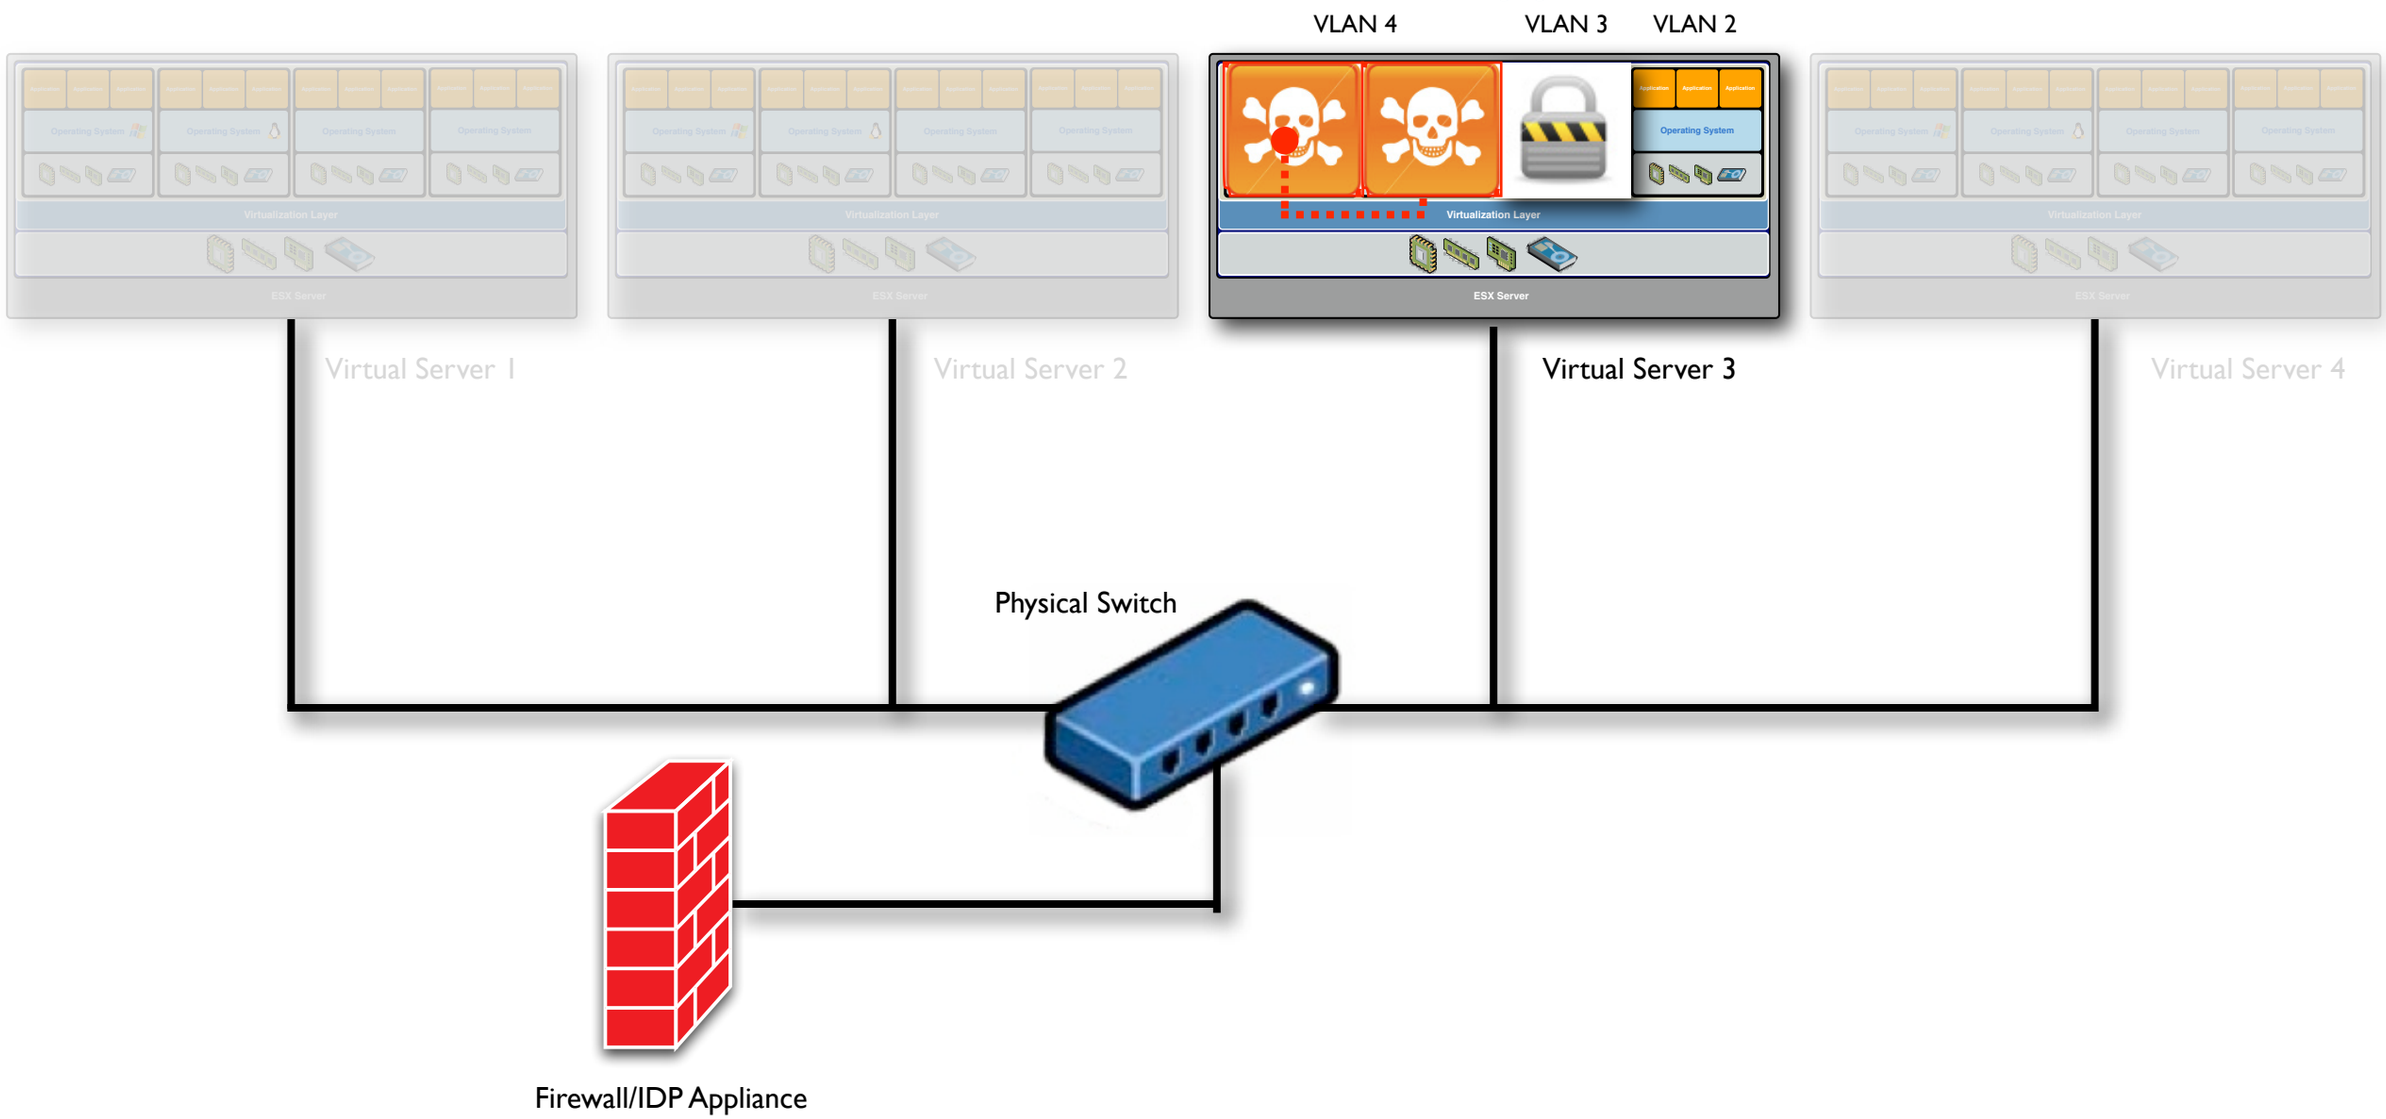

VirtSec Examples: Virtual Appliances

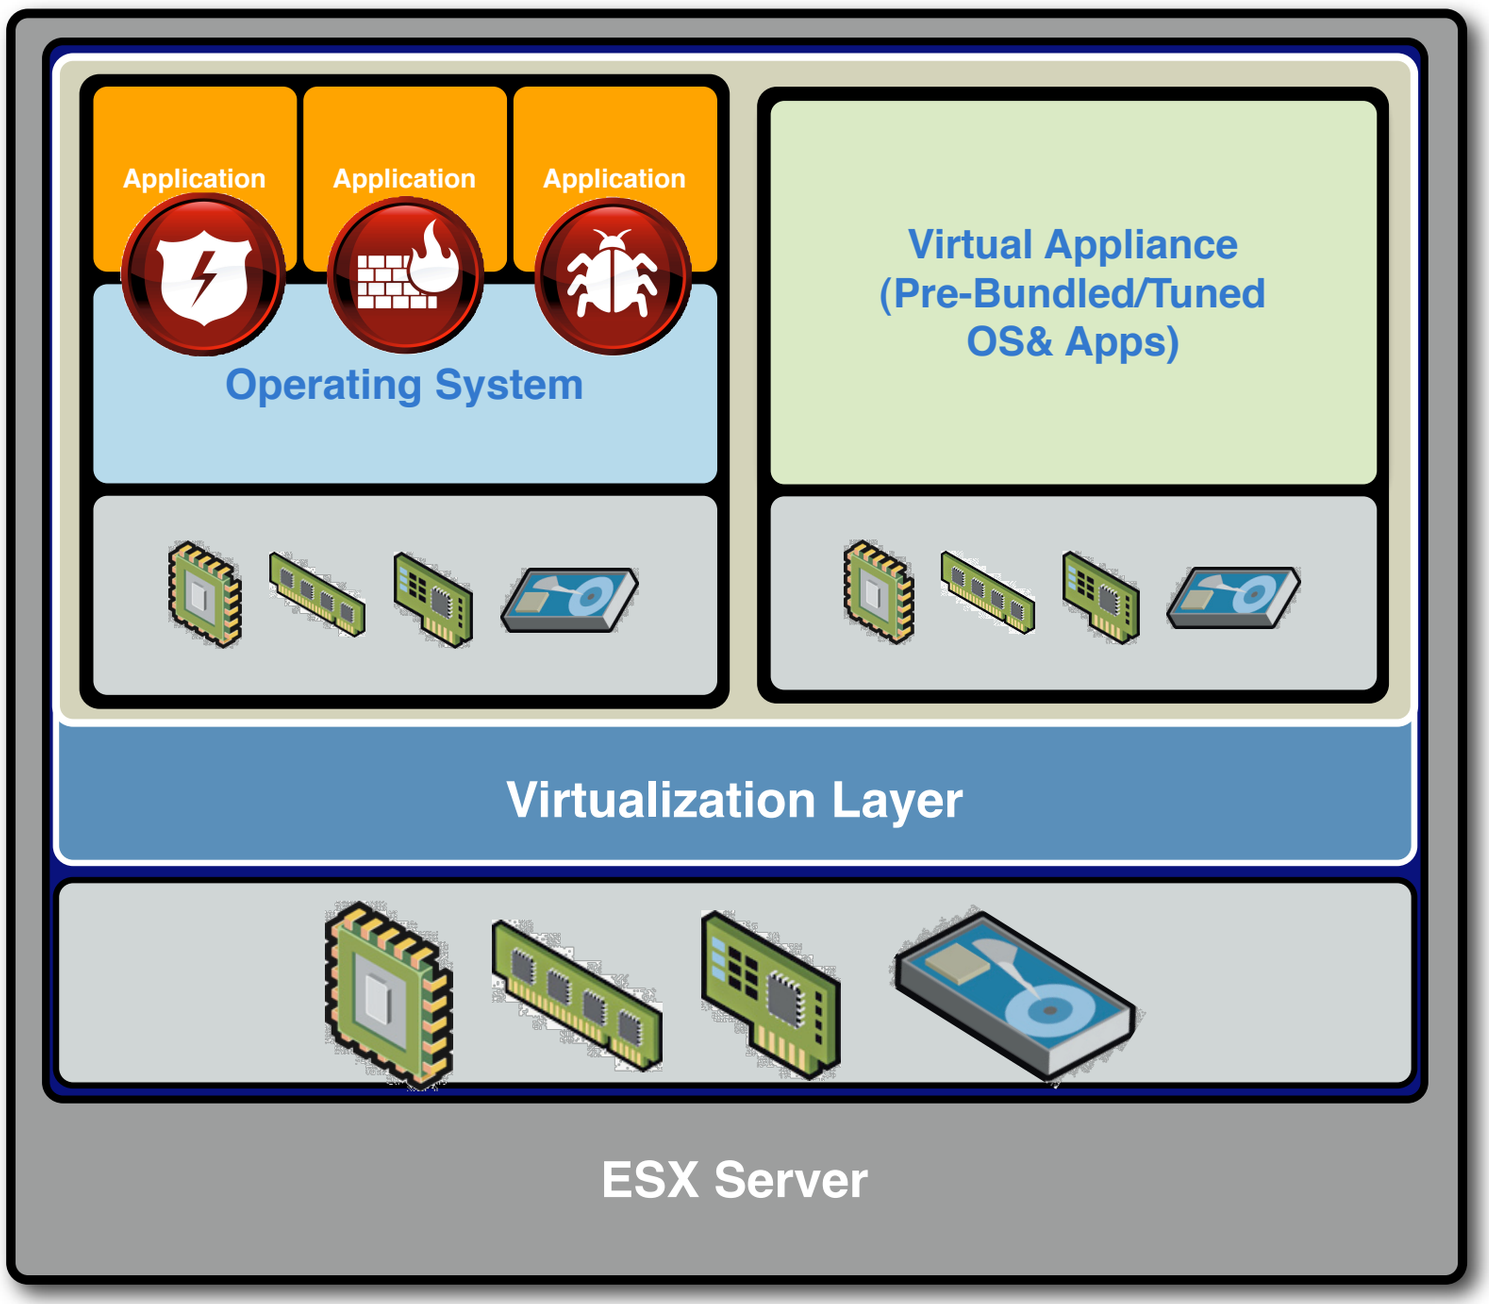

- The trick is forcing the traffic through the virtual appliances (if prevention is required) versus merely monitoring via SPAN for detection/monitoring
- ARP Spoofing/TCP RST's are also used by some vendors for VM's in the same portgroup/VLANs
- Requires careful (and potentially extensive) virtual networking configuration
- Doesn't protect against VM's in the same VLAN
- Does not directly protect the Hypervisor
- HA/Resilience an issue
- Consumes Host Resources
- There are lots of gotchas currently which VMsafe API's will certainly help with...

ALTOR networks
BlueLane
REFLEX SECURITY
catbird
Montego Networks
Secure Switching for Virtual Environments
Apani

Doh!

VirtSec Examples: VMM/ISV API's

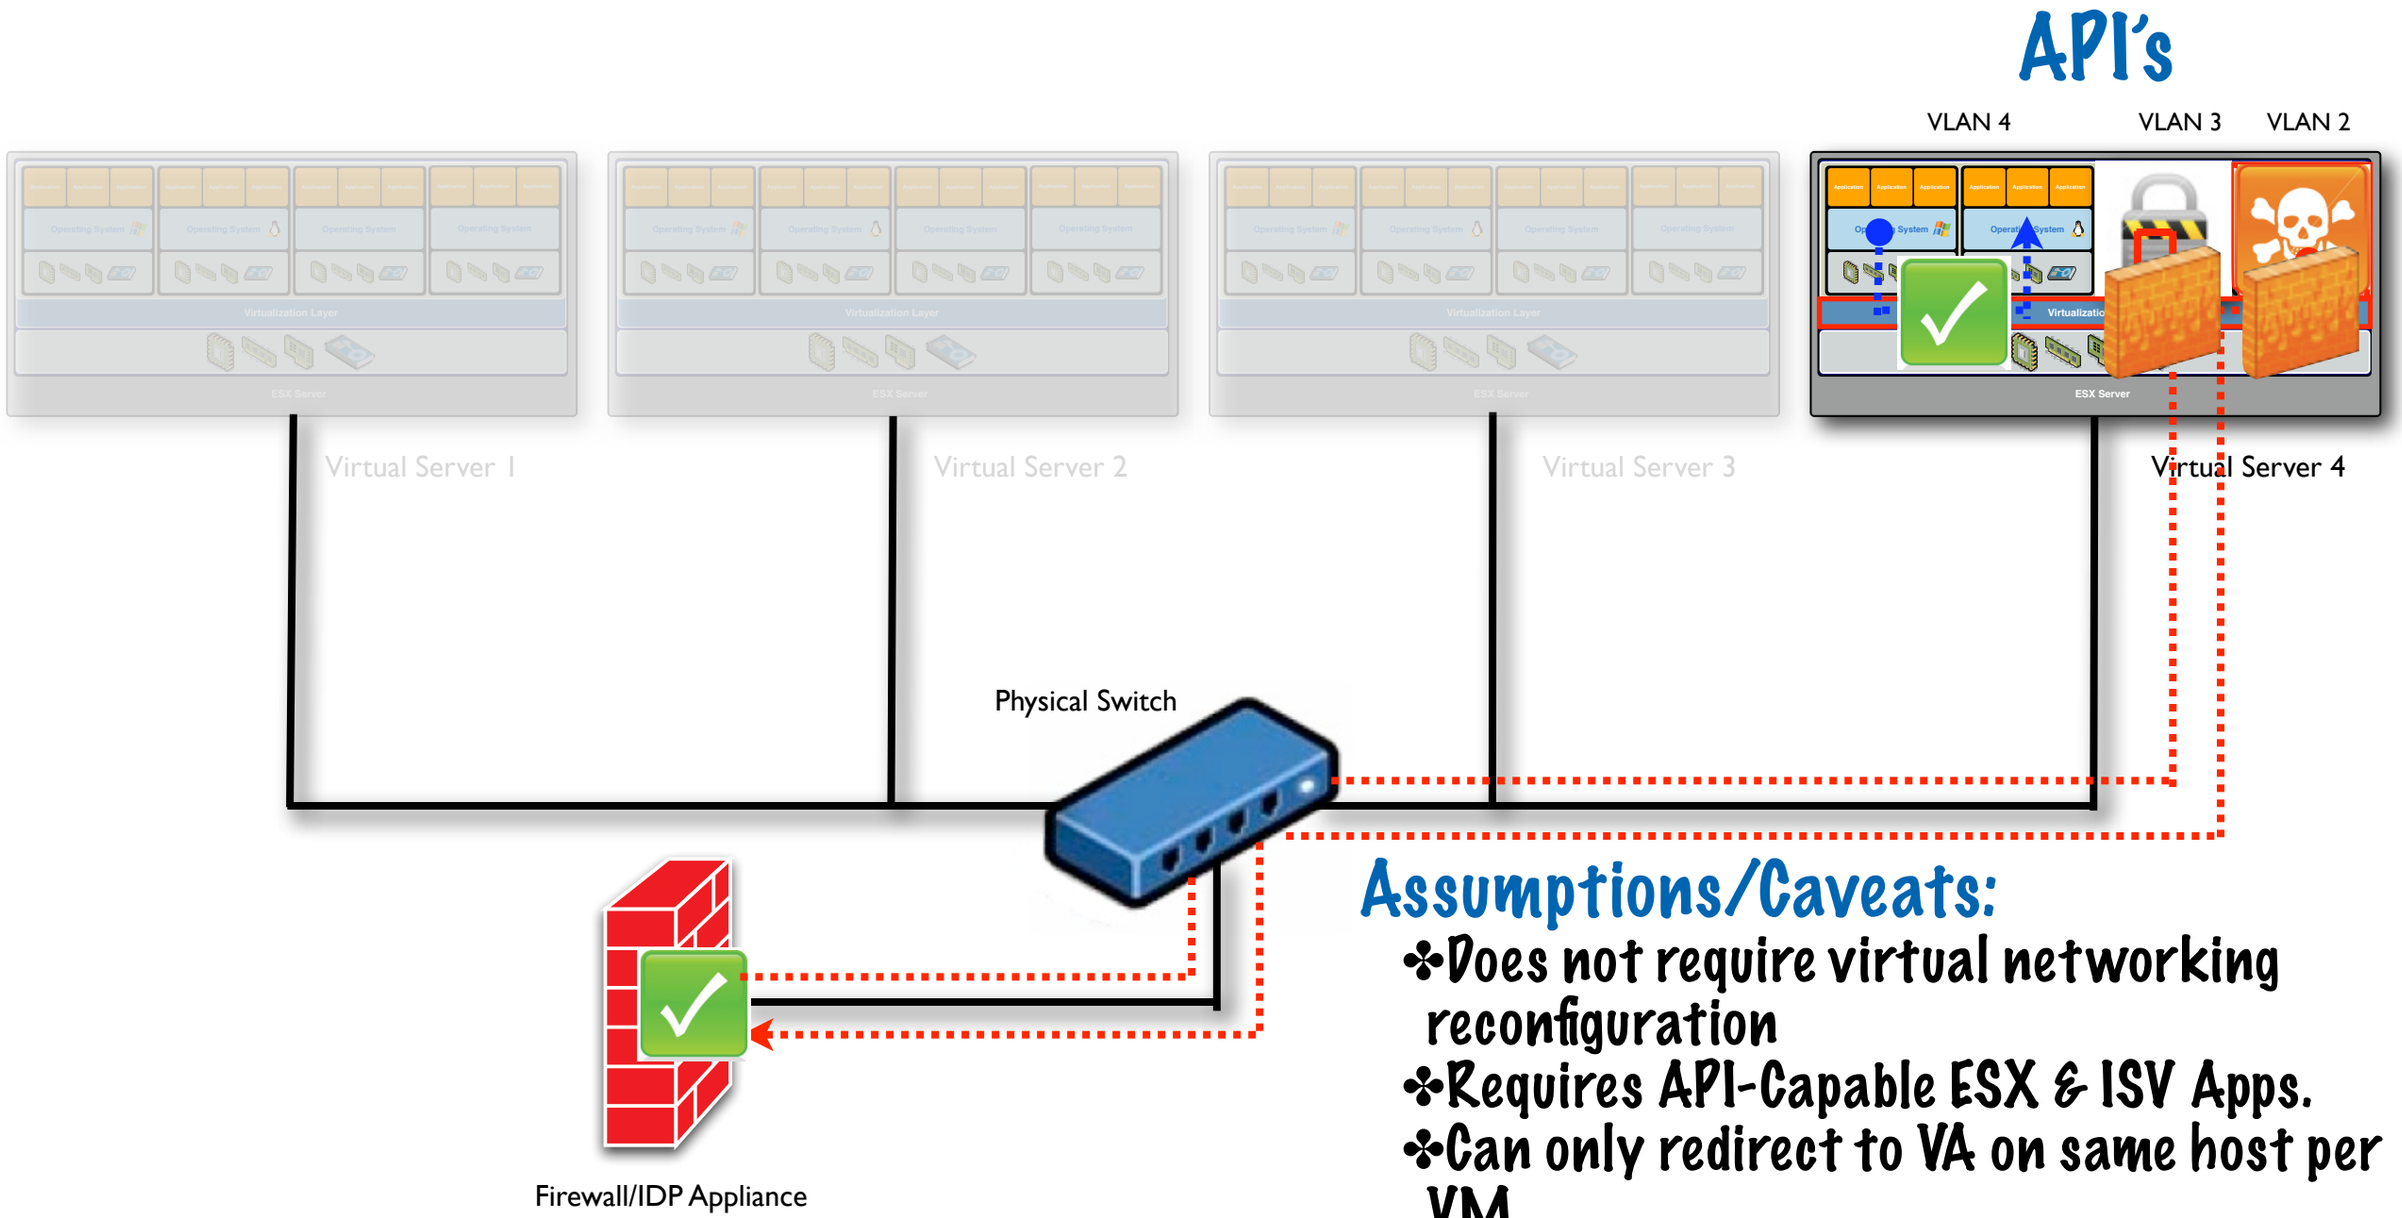

Assumptions/Caveats:

- Does not require virtual networking reconfiguration
- Requires API-Capable ESX & ISV Apps.
- Can only redirect to VA on same host per VM

VirtSec Examples: VMM/ISV API's

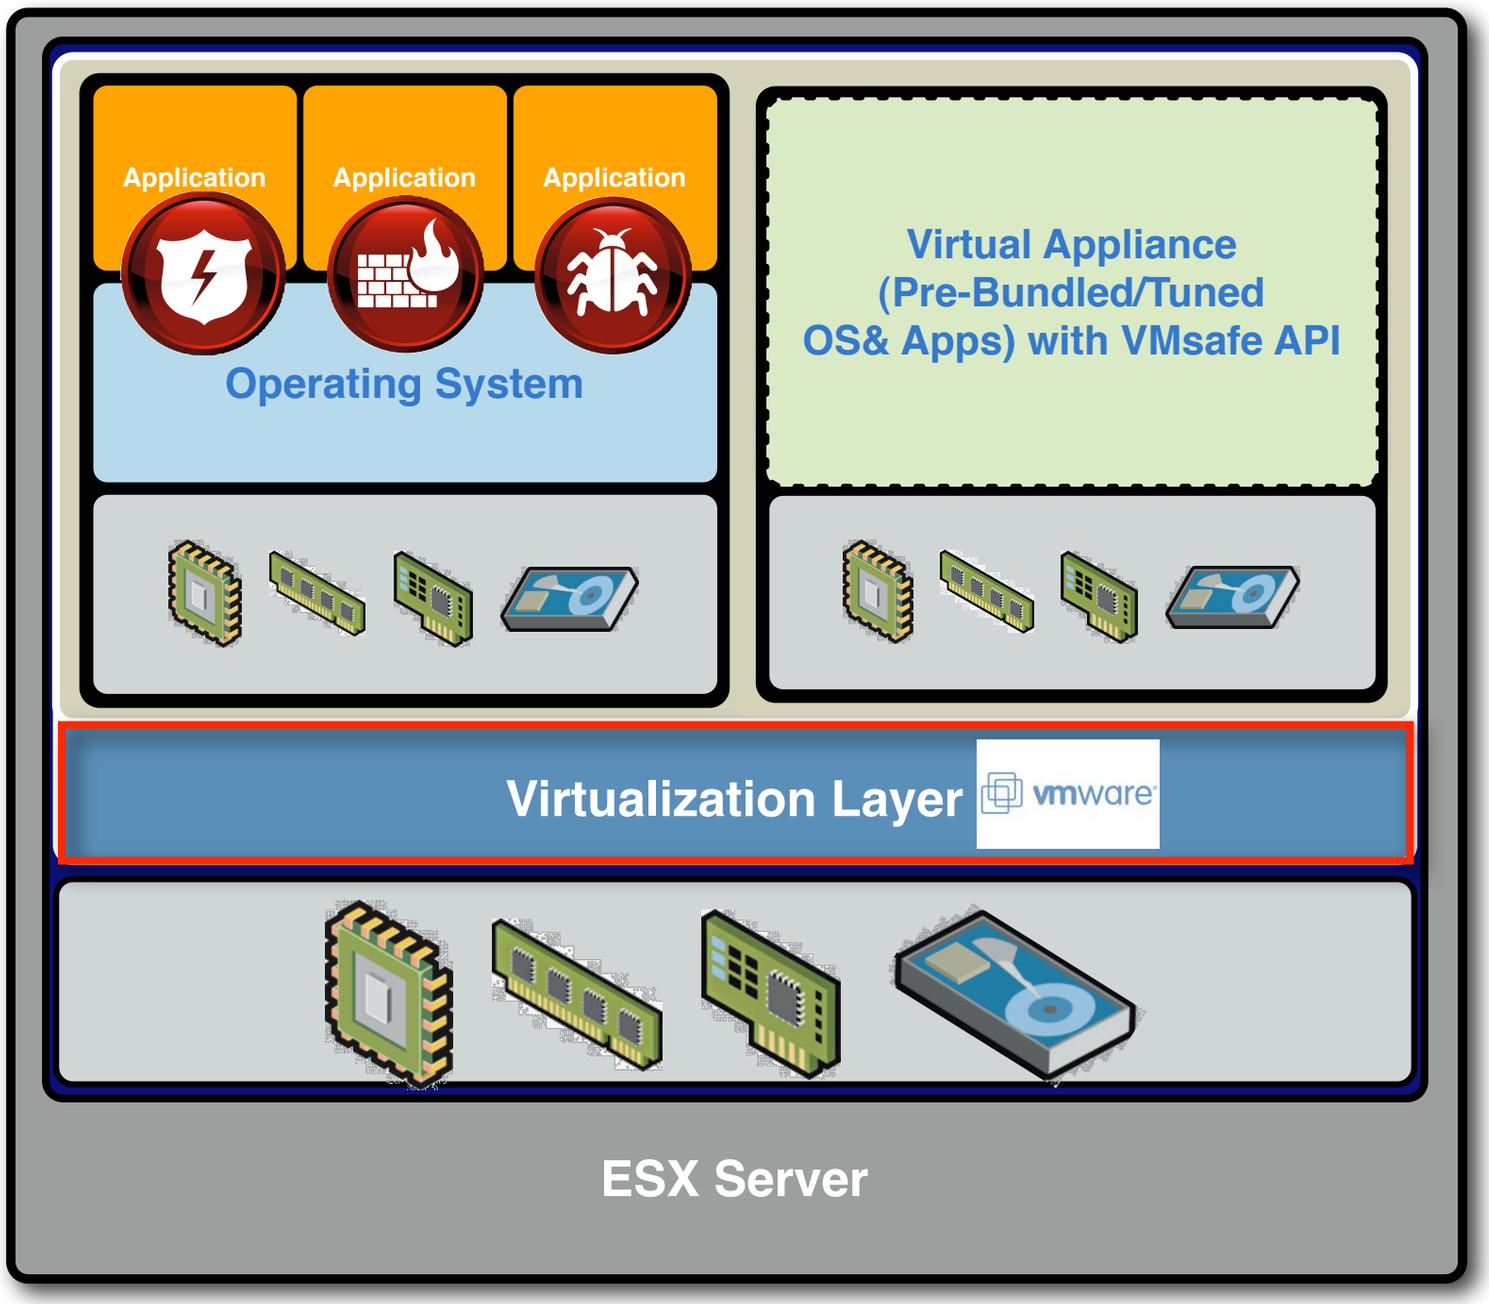

VMware VMsafe:

Security solutions built with VMware VMsafe will provide customers better granularity, visibility, correlation and scalability in virtual machine deployments.

Enables partners to build security solutions in the form of a virtual machine that can access, correlate and modify data to help control and protect:

- Memory and CPU
- Networking
- Process execution
- Storage

Note:

- Requires re-tooled ISV software & virtualization platforms
- Per-VM policies can only redirect to a VA/VM within the same host
- Coarse triggers
- Dispositions are limited

VirtSec Futures: 3rd Party vSwitches

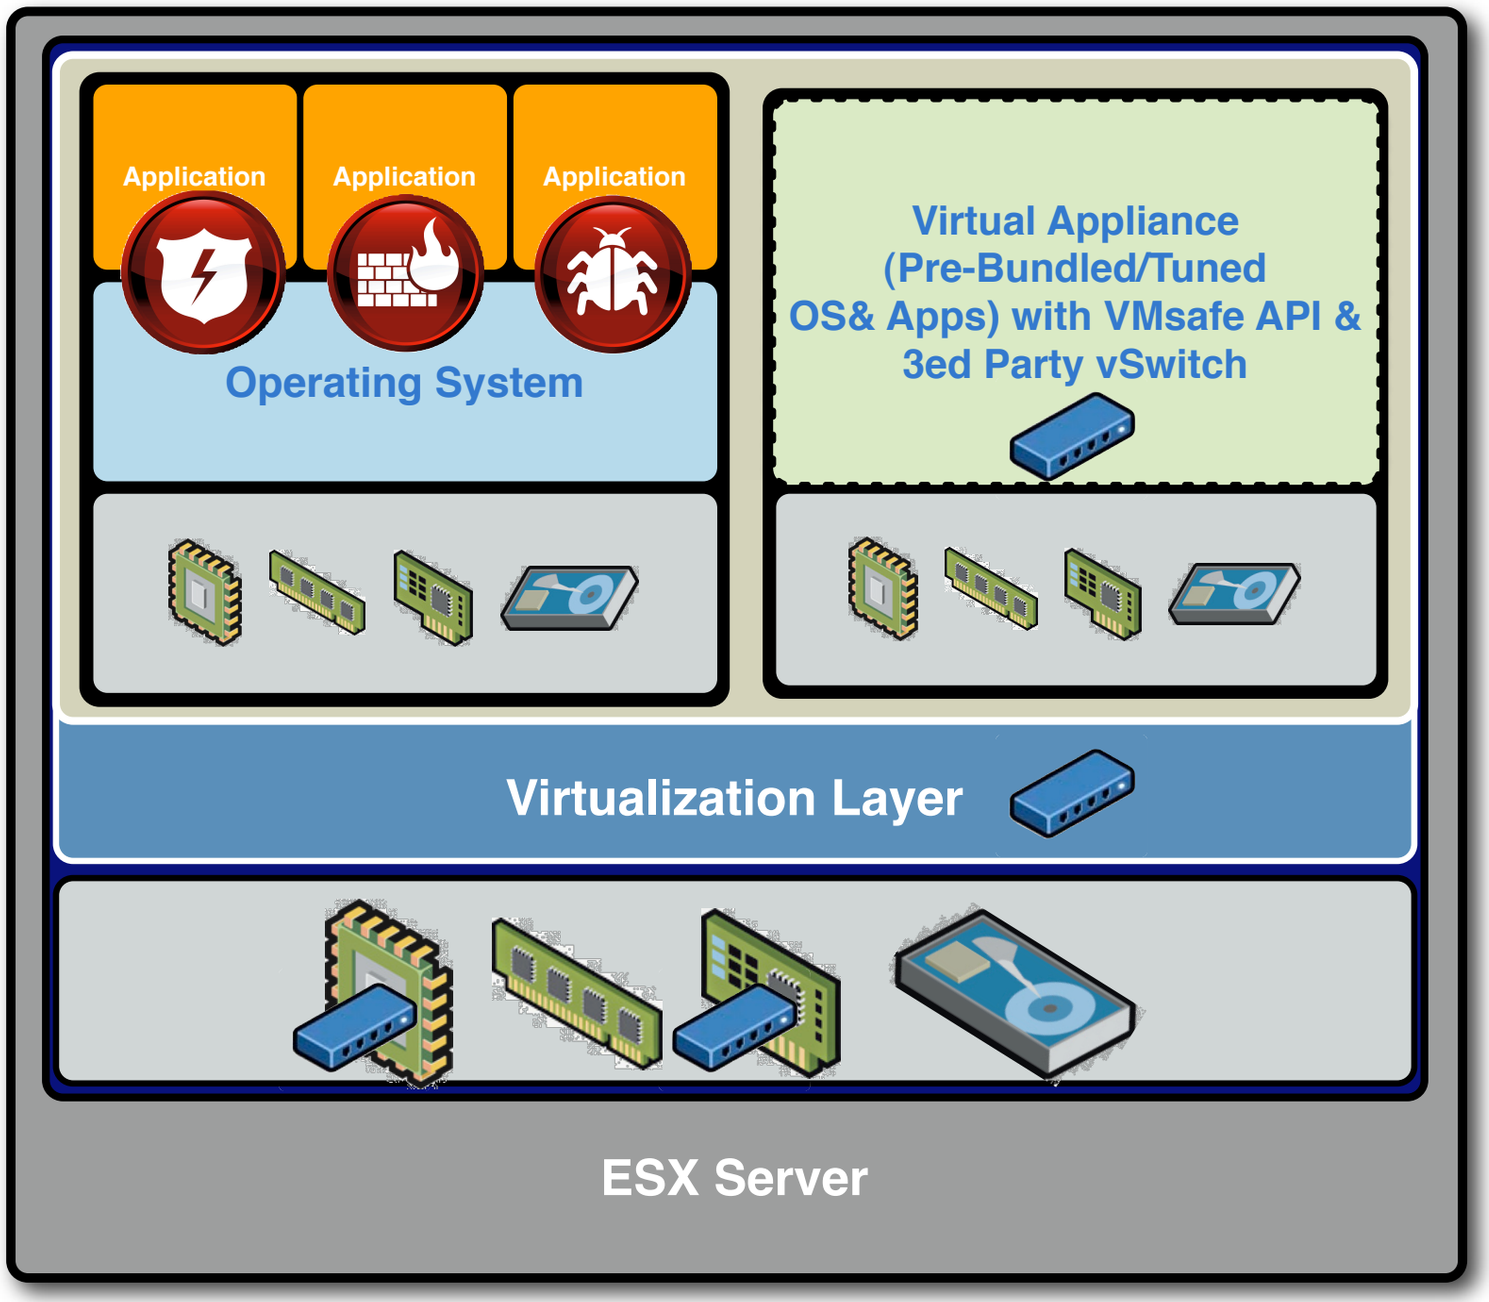

Third Party vSwitches

- Acts as a policy-driven intelligent disposition director to 3rd party security functions
- Allows integration/replication of external software, fabric capabilities and policy
- Consistency in networking capabilities (load-balancing, QoS, L3-7, etc...)
- vSwitches evolving to reside in hardware & software:
  - Hypervisor
  - VA/VM
  - Underlying Virtualization-enabled CPU's
  - In "new" breed of NIC cards

intel · NetEffect · Neterion · Cisco · Montego Networks (Secure Switching for Virtual Environments) · Mellanox Technologies

VirtSec Futures: I/O Virtualization

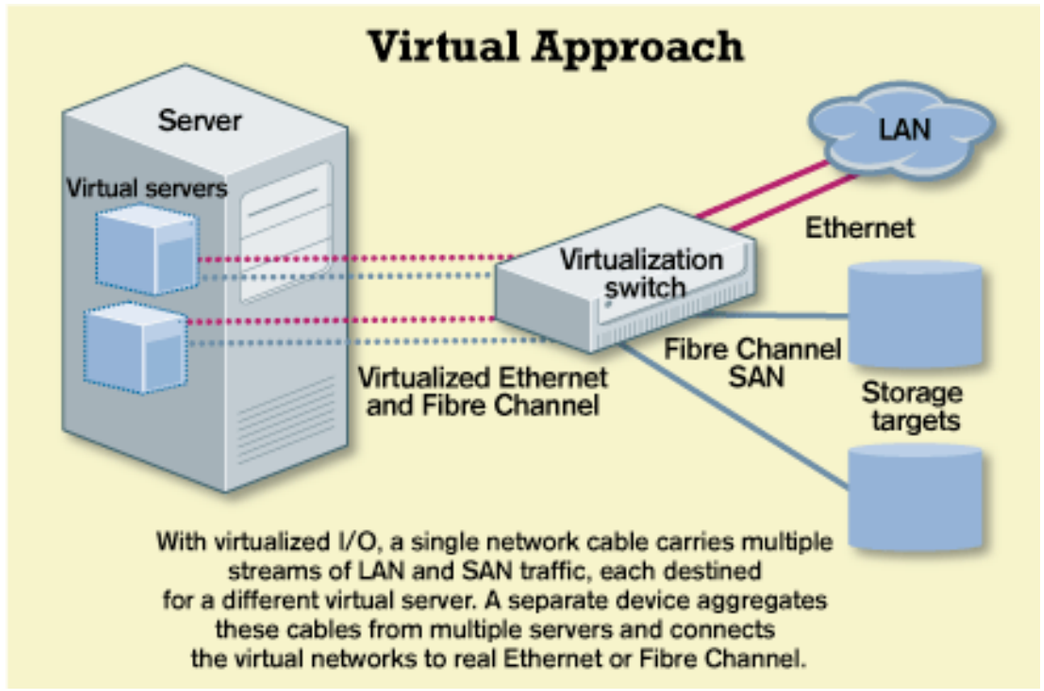

I/O Virtualization:

- Single network connection provides virtualized fabric interconnectivity for LAN & SAN
- Ultimately your VM's run in the switch
- All your VM's are belong to us!

- Cisco 7000 Nexus
- Brocade DXC Backbone
- 3Leaf V-8000 Virtual I/O Server
- Xsigo I/O Director

CISCO

3LEAFSYSTEMS

xsigo systems

BROCADE

My Head's In the Cloud(s)

- Today's virtualization offerings are just the beginning
- Cloud/Grid/Utility computing is real
- SaaS, Amazon's EC2, Clean Pipes are all great examples today of what's coming tomorrow
- How are we going to secure the abstraction of a cloud-based virtualized processes, memory space, I/O?

The End Is Nigh! Run Away!

- Setup & Context
- x86 Virtualization Overview in 90 Seconds
- Virtual Networking Architecture
- VirtSec Solutions Landscape
- The Four Horsemen
- Wrap-Up

Vini, Vidi, Wiki...

- Monolithic security vendor virtual appliances are the virtualization version of the UTM argument
- Virtualized Security can seriously impact performance, resiliency and scalability
- Replicating many highly-available security applications and network topologies in virtual switches don't work
- Virtualizing security will not save you money, it will cost you more

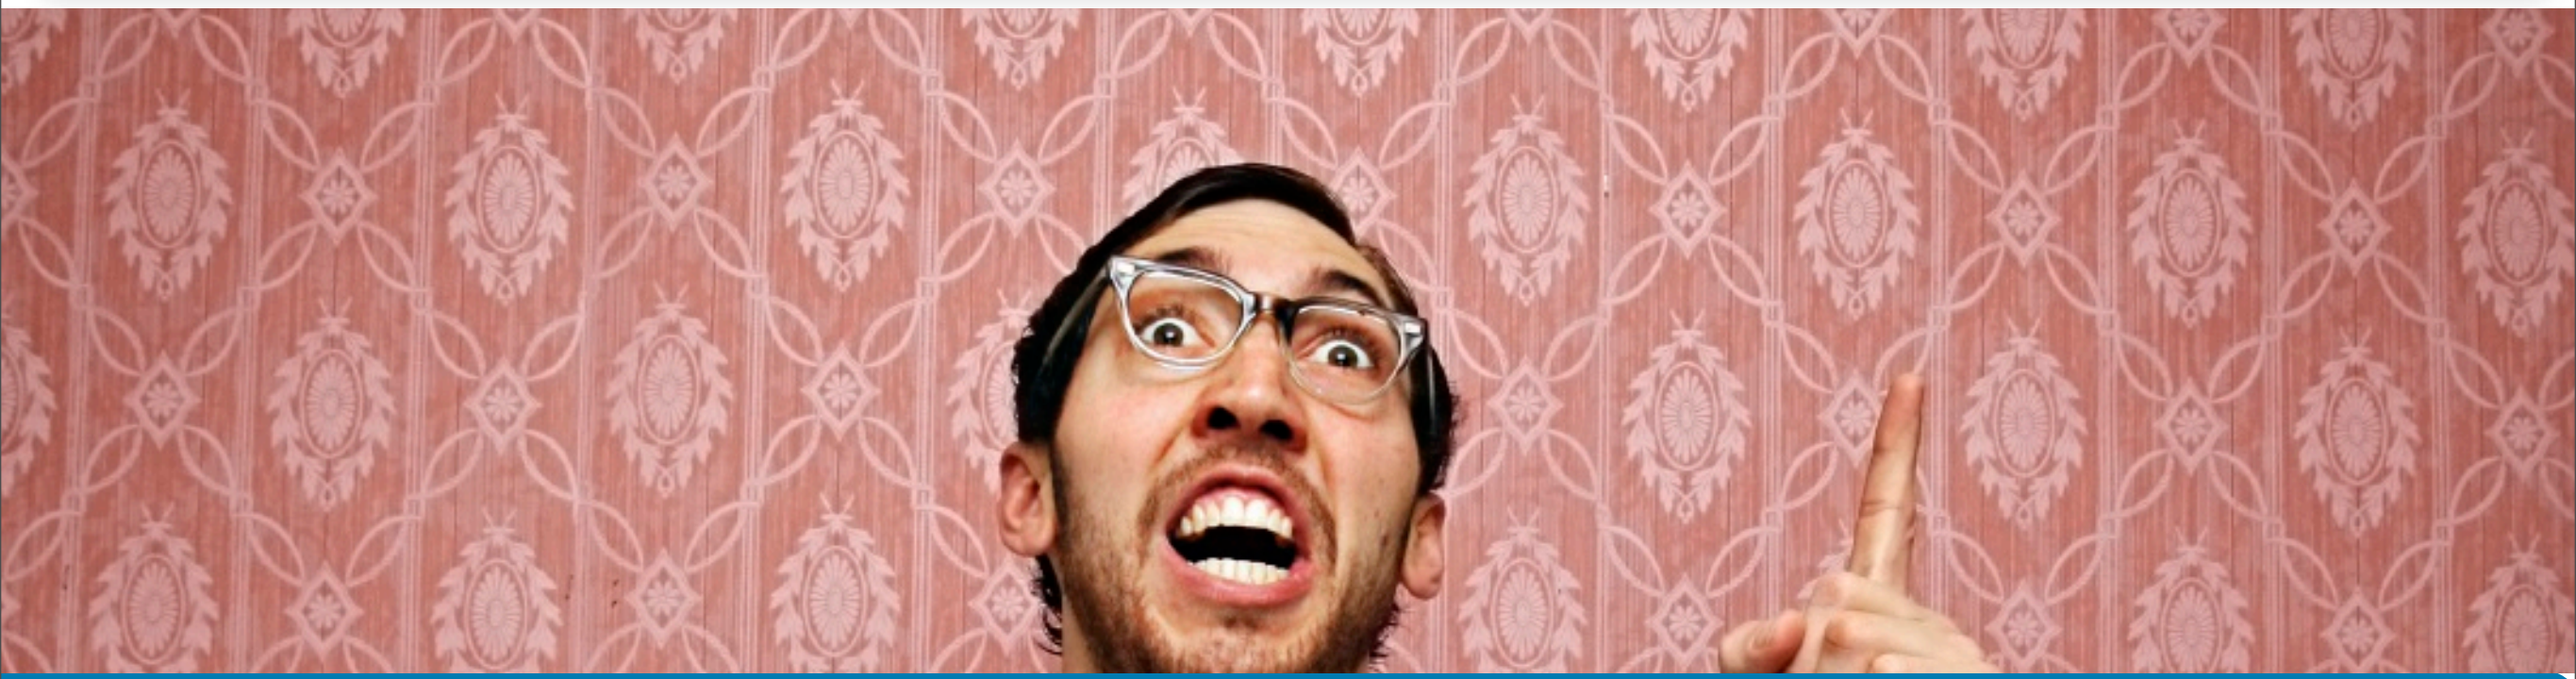

Example: Virtualizing the DMZ*

Typical Screened-Subnet DMZ:

- Trust zones separated by physical controls on separate switches & host groups/ clusters

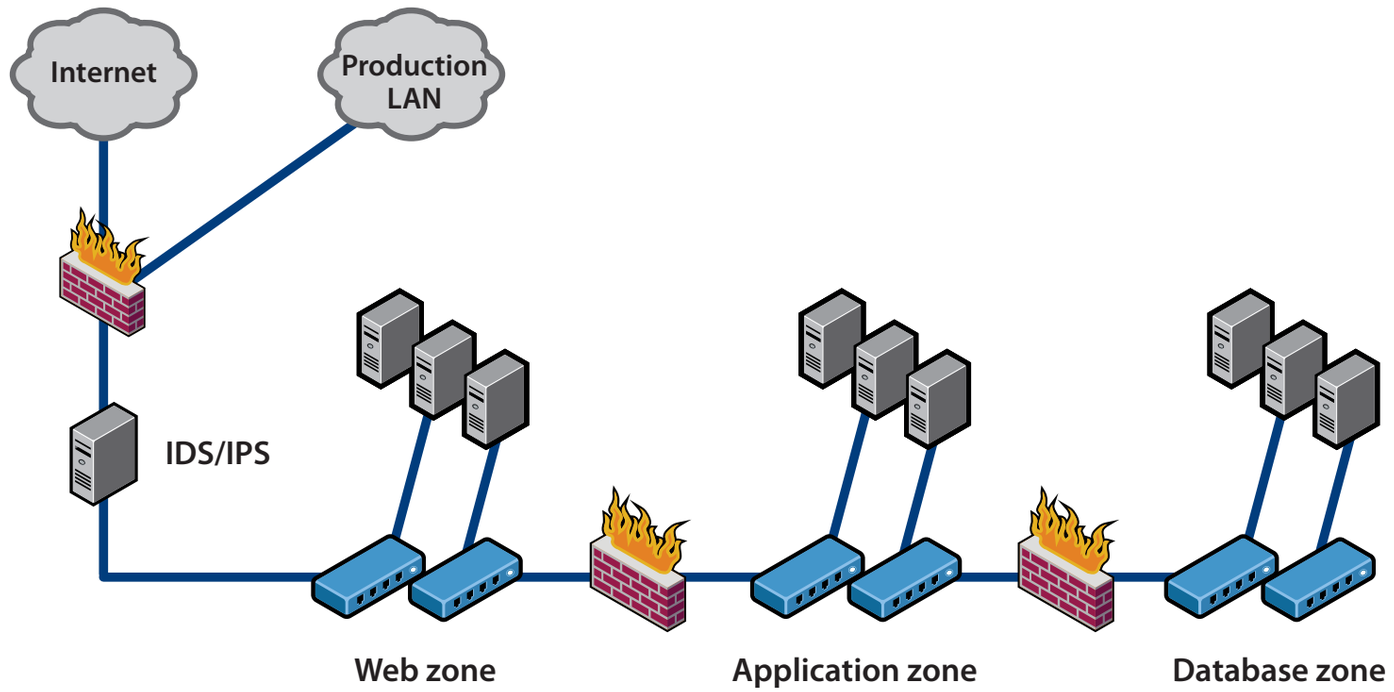

Figure 1 — A typical DMZ in a physical environment

*Images/Concept from VMware Whitepaper: DMZ Virtualization with VMware Infrastructure

Example: Virtualizing the DMZ*

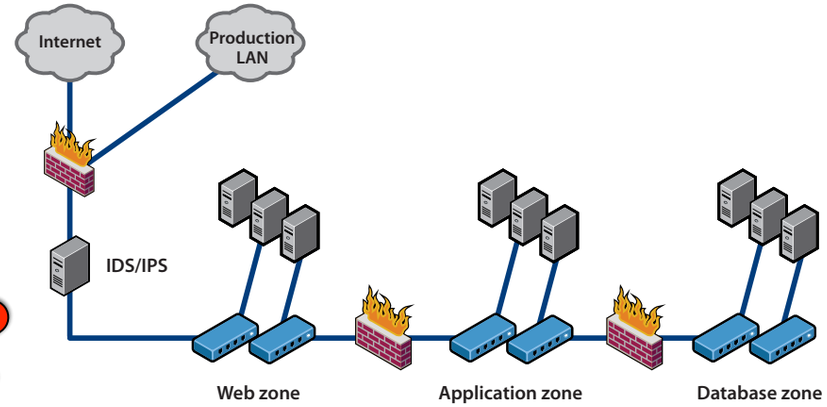

Figure 1 — A typical DMZ in a physical environment

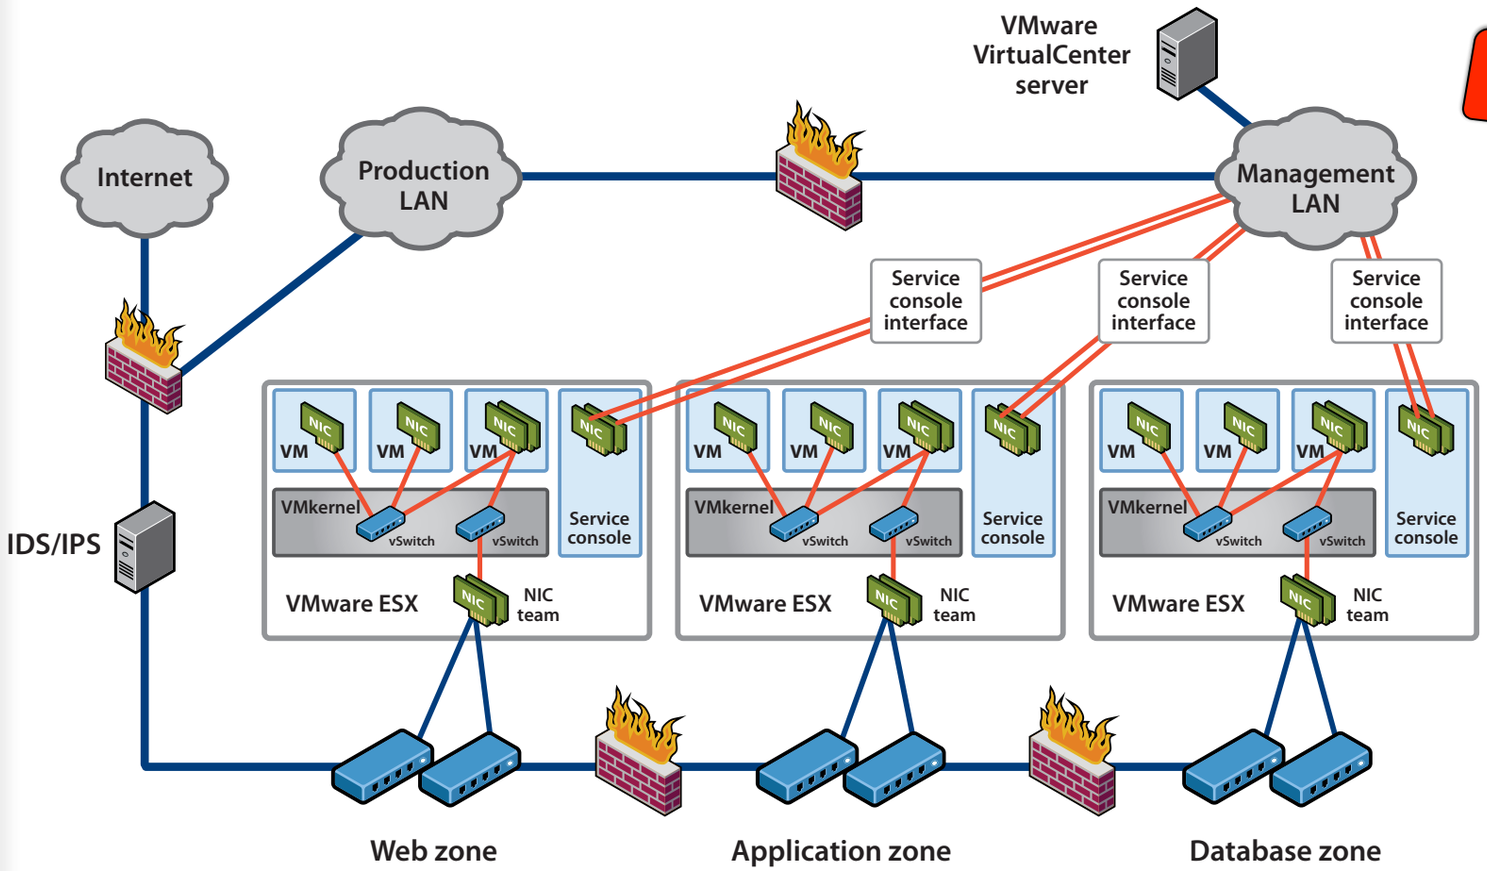

Figure 3 — Partially collapsed DMZ with separate physical trust zones

Partially Virtualized Screened-Subnet DMZ:

- Trust zones separated by both physical controls but VM's on multiple ESX clusters

*Images/Concept from VMware Whitepaper: DMZ Virtualization with VMware Infrastructure

Example: Virtualizing the DMZ*

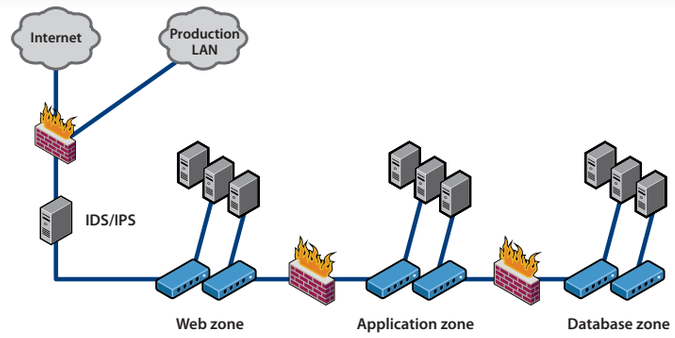

Figure 1 — A typical DMZ in a physical environment

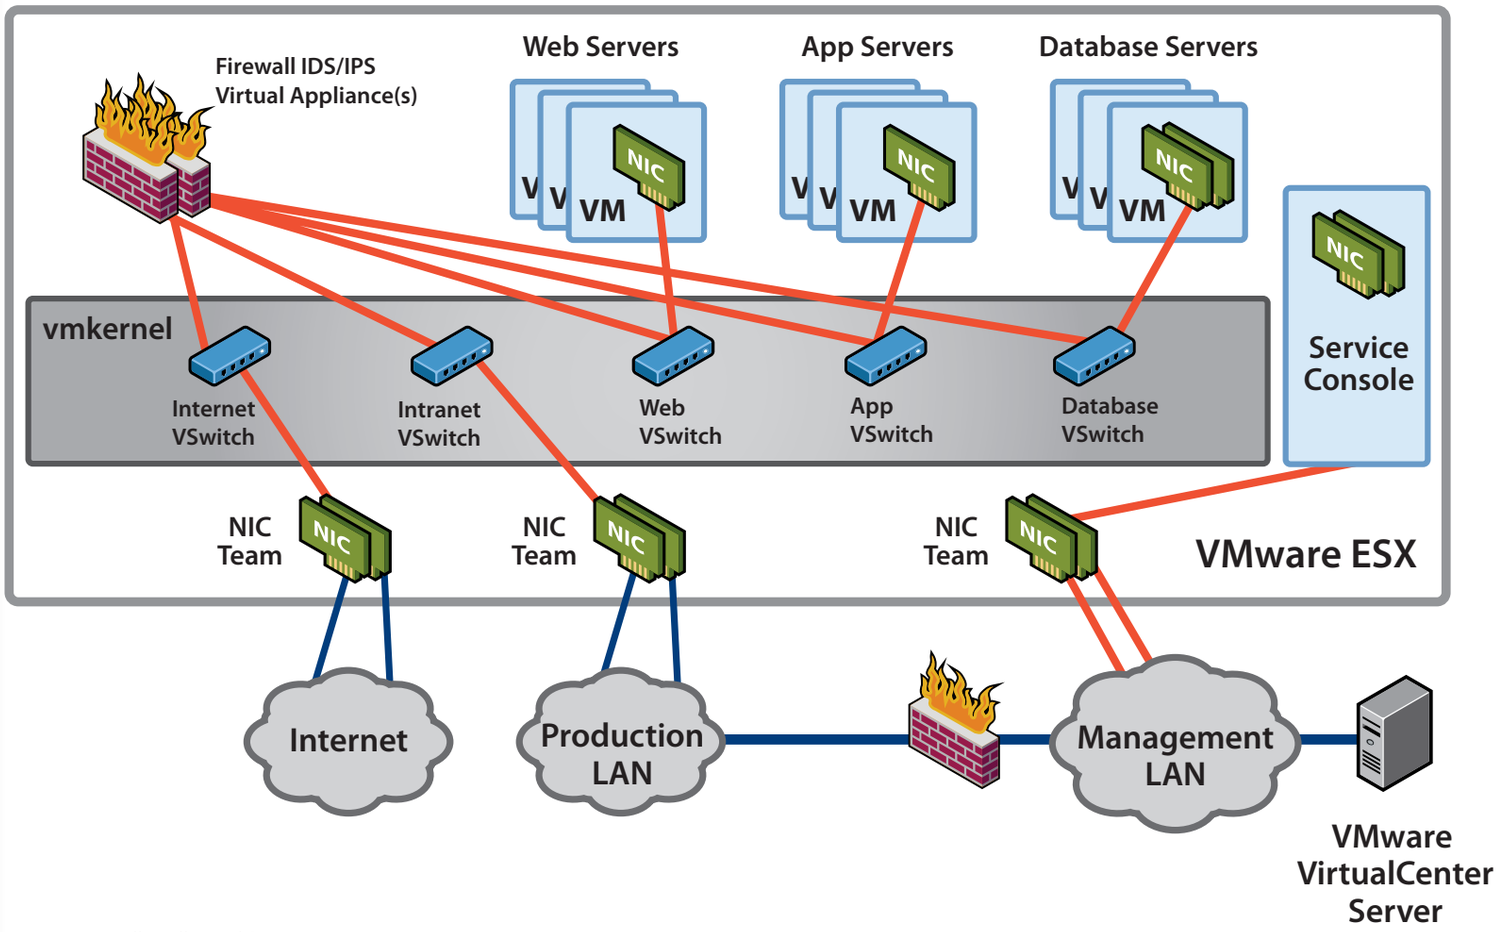

Figure 5 — Fully collapsed DMZ

Completely Virtualized Screened-Subnet DMZ:

- Trust zones separated by virtual controls on a single ESX Cluster

*Images/Concept from VMware Whitepaper: DMZ Virtualization with VMware Infrastructure

Who Are the Four Horsemen?

The Four Horsemen of the Apocalypse represent the "...forces of man's destruction described in the Bible in the Book of Revelations" and are "...named after the powers they represent"*

- War
- Pestilence
- Death
- Famine

*Wikipedia

War I Episode 7: Revenge Of the UTM Clones

Monolithic security vendor virtual appliances are the virtualization version of the UTM argument:

- The notion that we will deploy a single vendor/monolithic security VA in each host is silly
- If you're still stuck on "defense in breadth," you're going to deploy more than one security virtual appliances on each host
- UTM performance sucks when you flip all the switches

The VAUTM Conundrum

- How do you ensure that traffic is statefully directed to the appropriate individual in-line security bumps in the stack?

- The more security VA's you add, the less VM's you can service

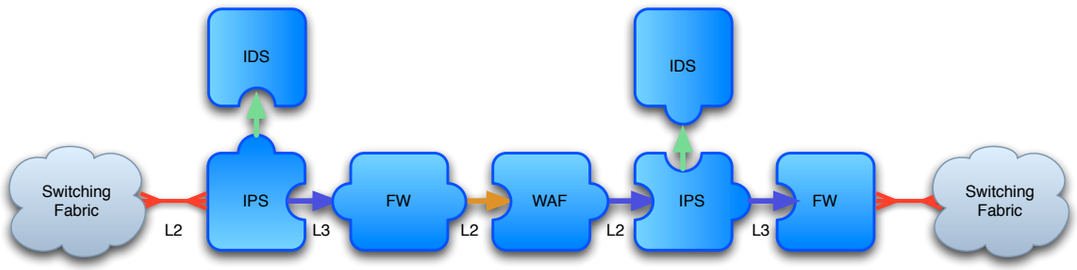

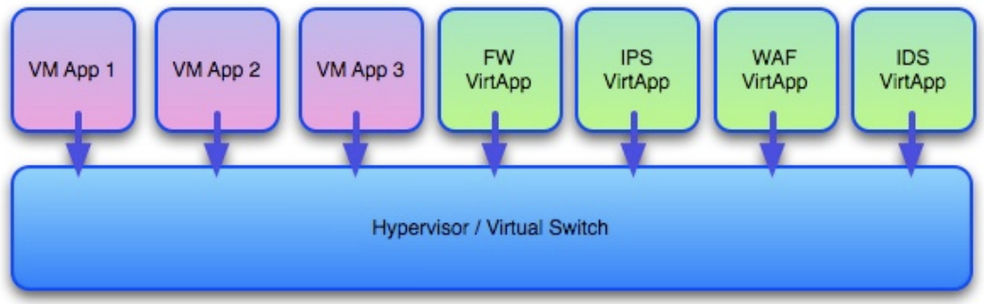

Virtualized Security can seriously impact performance, resiliency and scalability

- Performance overhead of in-line security VA/VMs & API's is extremely difficult to predict
- Today we rely on multiple load-balanced high-performance multi-core COTS H/W or dedicated ASIC/FPGA equipped appliances for acceptable throughput/low latency...
- We're now going to expect that software based VA's which are not optimized or do not utilize paravirtualized drivers to perform the same?
- Security functions are competing for the same resources as the VM's you're trying to protect

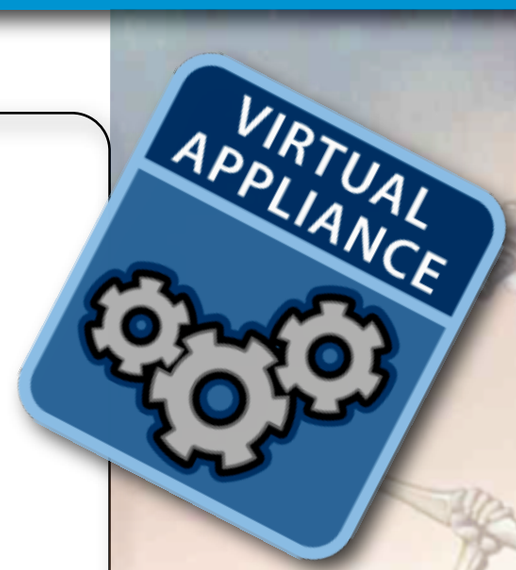

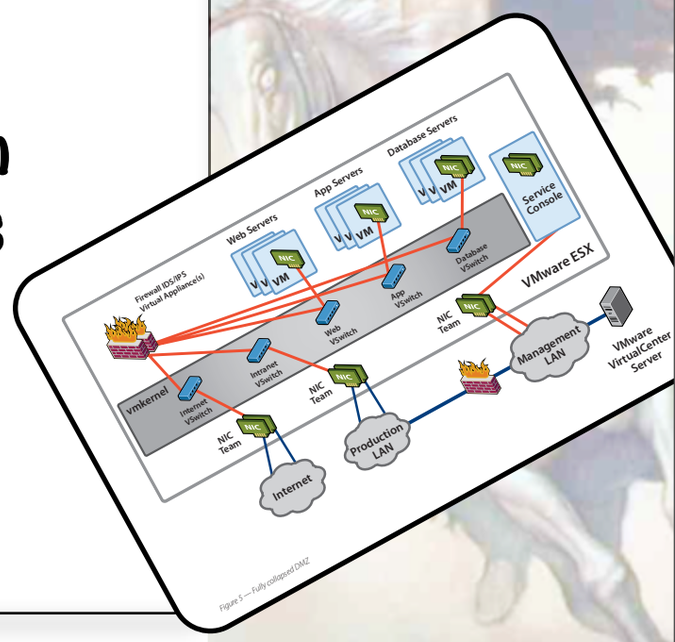

Drinking From the Firehose

- VMware showed tests* with linux-based VM-VM throughput on the same vSwitch of ~2.5Gb/s
- Most dedicated hardware appliances have trouble at those rates at small packets/low latency
- What happens when you try to choke every flow through a non-optimized, software-only virtual appliance in/out of every VM?
- What happens when we add multiple 1Gb/s or 10Gb/s bonded pipes feeding our servers?

Figure 9. Linux Virtual Machine to Virtual Machine TCP Throughput

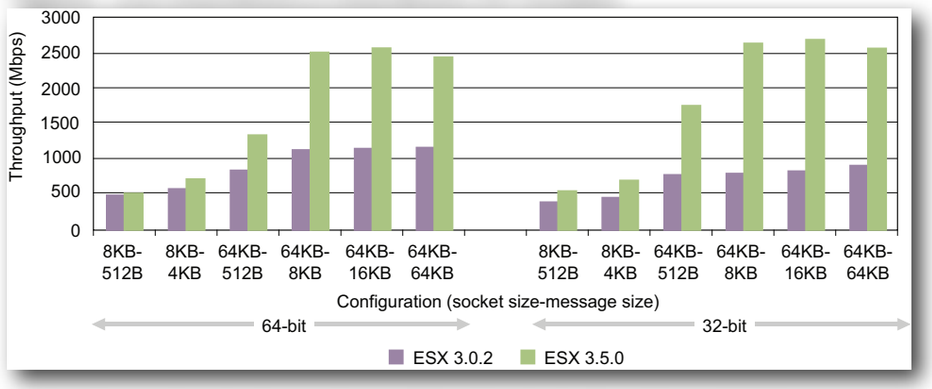

Thus, the virtual machine to virtual machine TCP throughput on ESX Server 3.5 can exceed 2.5 Gbps for some operating systems while speeds of physical networks with 1 Gbps NICs are limited to approximately 950 Mbps.

*Networking Performance VMware® ESX Server 3.5

Public Service Announcement

Every time you deploy a security virtual appliance...

God kills a kitten.

Death I The Network Is The Computer?

Replicating many highly-available security applications and network topologies in virtual switches don't work

- Security applications are incredibly topology sensitive
- Affinity between the physical, logical and policy elements breaks when things move
- It's not that you can't get network-based HA to work, it's the support of the applications and their secret sauce that breaks.
- Most physical appliances use heavily tweaked kernels and drivers which aren't supported natively in virtualization stacks; performance suffers and HA may no longer work
- Failover and HA options for stateful security applications currently suck

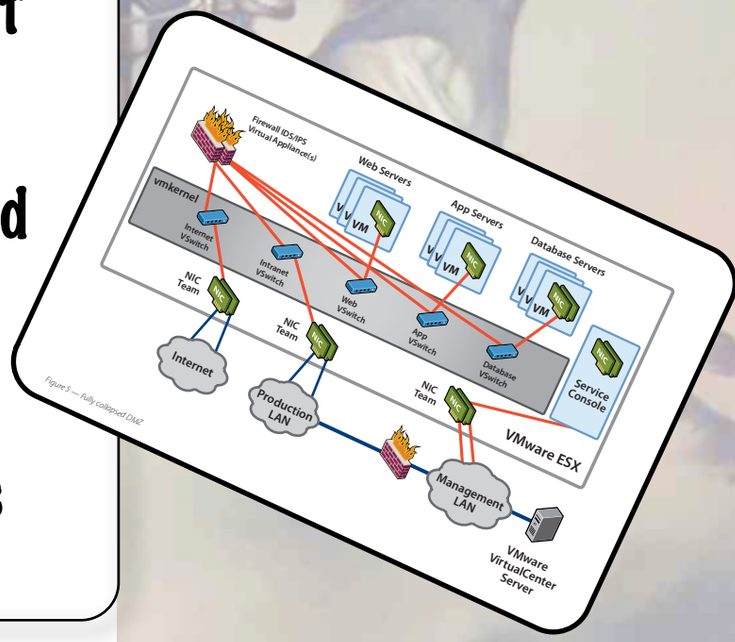

*Image from Cisco Whitepaper: VMware Infrastructure 3 in a Cisco Network Environment

Resilient Network Designs Are Achievable

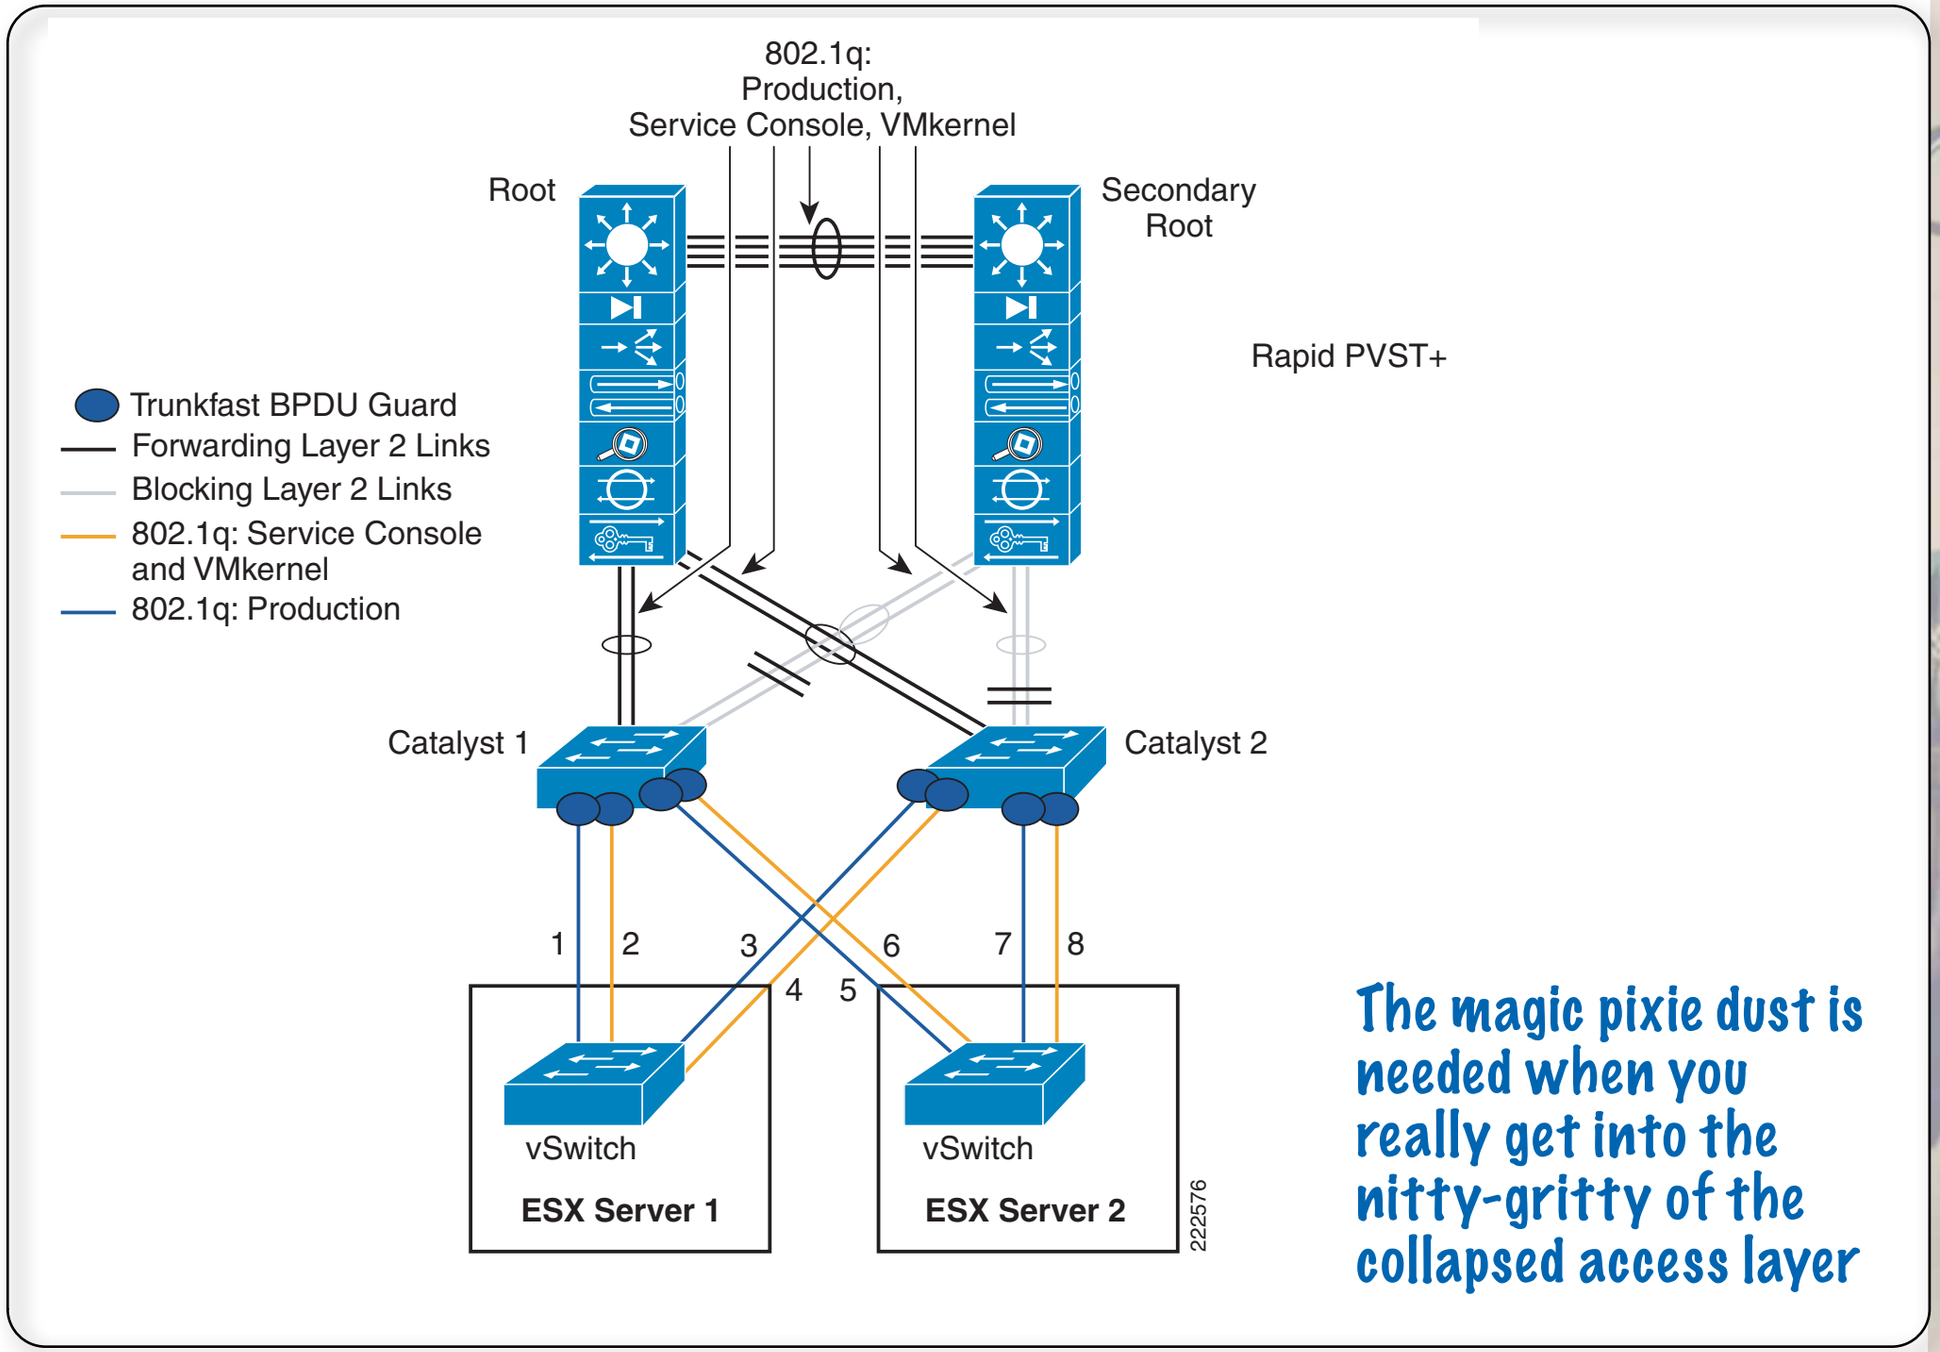

The magic pixie dust is needed when you really get into the nitty-gritty of the collapsed access layer

*Image from Cisco Whitepaper: VMware Infrastructure 3 in a Cisco Network Environment

...But the Security Application Failure Recovery Options Suck

What happens when these security virtual appliances fail?

- Application-level and VMware HA clustering do not take into consideration the network topology sensitivities of security applications
- Security applications and the networking stacks are not stateful and do not exchange telemetry
- Moving the security VA to another box leaves the VM's unprotected or disconnected/isolated on the original
- Failing over an entire cluster-member's inventory of VM's due to the failure of a security component is ludicrous

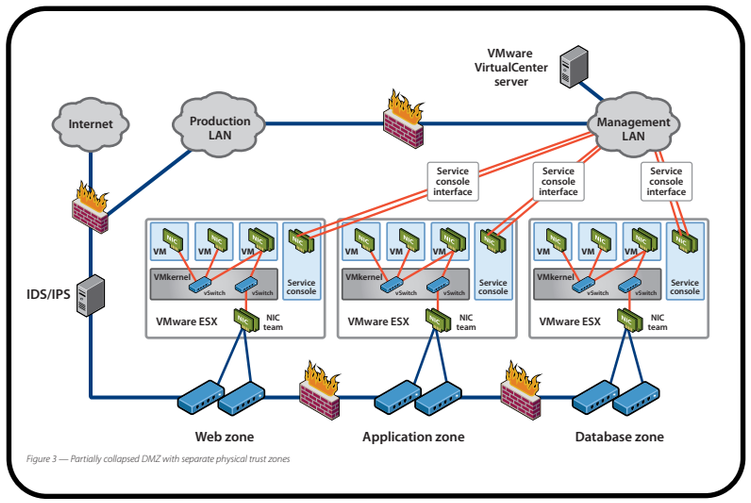

Figure 3 — Partially collapsed DMZ with separate physical trust zones

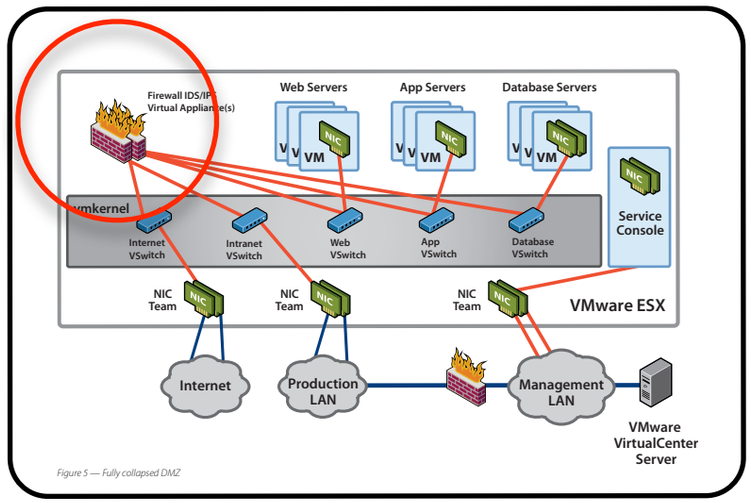

Figure 5 — Fully collapsed DMZ

Run...It's the Fuzz!

I was going to show you a really cool demo with ERNW using their modified L2 Sulley fuzzing framework abusing the HA protocols of a well-known firewall vendor to show you how fragile this stuff is in terms of performance/resiliency, but:

- The version that runs under ESX doesn't utilize VRRP
- The version that runs under ESX doesn't use the native clustering functionality
- In fact, the version that runs under ESX doesn't support HA in <u>any way</u>...

Get used to it, as no VirtSec virtual appliance vendor I've spoken to currently supports native HA/LB in the virtual appliance version of their products!

Doh!

Famine I Spinning VM Straw Into Budgetary Gold

Virtualizing security will not save you money, it will cost you more

- We won't get rid of physical appliances or security line cards in switches, in fact, we'll probably have to buy more...
- We won't get rid of host-based security software
- That means that when we add VirtSec solutions, these solutions & their licenses are cost-additive
- As we add more solutions, we add complexity
- Who's going to administer these solutions?

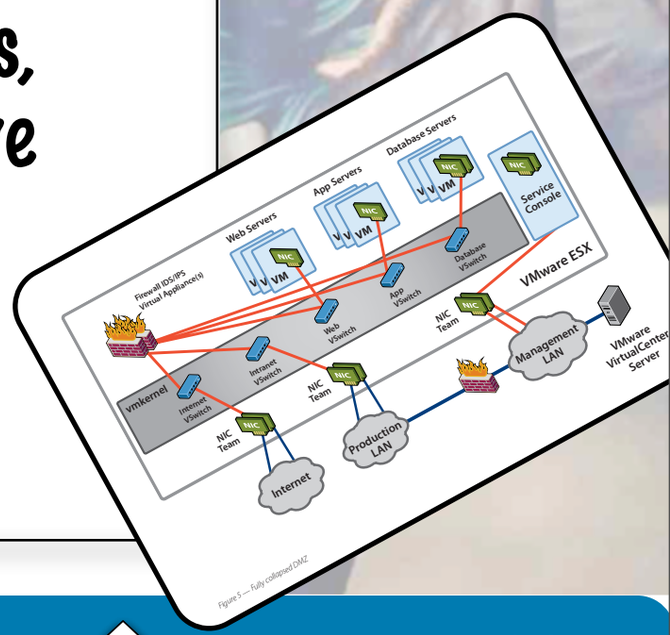

Parting Is Such Sweet Sorrow

- Setup & Context
- x86 Virtualization Overview in 30 Seconds
- Virtual Networking Architecture
- VirtSec Solutions Landscape
- The Four Horsemen
- Wrap-Up

Great, So What Can I Do Today?

Guidelines:

- Use risk assessment and threat modeling
- Design and provision your virtual network very carefully or find someone you can convince to do so
- Provide for zoned/screened VM's that are grouped based on risk/value requiring like policies and failure domains
- Thoroughly evaluate existing and emerging tools to determine value vs. disruption
- Dig deep on performance, HA/LB and scalability on virtual appliances
- Understand what impact VMotion, DRS, HA have on security solutions
- Be able to quantify performance impacts on hosts/networks
- Push virtualization platform providers to reveal roadmaps

Hope Is Not a Strategy, But It Doesn't Hurt

- We need a trust model in hardware & software
- We need affinity between the VM and protection schemes/policies
- Centralized VM registration providing telemetry that controls spin-up, state and mobility capabilities regardless of vendor
- Comprehensive discovery, profiling, dynamic configuration & security management of all VM's -- online or offline
- Intelligent networking capabilities within the virtual switching infrastructure for consistency, visibility and security including integrated virtual network admission control & access Control (vNAC)
- Correlation of telemetry between VM Management and internal/external security planes to tie in virtualization, network and security provisioning/ management into a consolidated single pane of glass

Vini, Vidi, Wiki...

- Monolithic security vendor virtual appliances are the virtualization version of the UTM argument
- Virtualized Security can seriously impact performance, resiliency and scalability
- Replicating many highly-available security applications and network topologies in virtual switches don't work
- Virtualizing security will not save you money, it will cost you more

Thanks For Not Leaving ;)

PacketFilter