

Network Security - Efficiency of Virtualized Vs Containerized Firewalls



CYBERSECURITY™
MADE IN EUROPE

Introductions



Nils Undén

CTO

nils.unden@clavister.com



Mattias Fredriksson

Product Owner

mattias.fredriksson@clavister.com



Dave Cremins

Cloud Software Architect

dave.cremins@intel.com

Agenda

- NGFW Deployment Options – virtual/VNF and container/CNF
- Design Considerations
- Clavister NetShield – Performance Measurements from Intel Partner Alliance (IPA) Lab:
 - Virtual/VNF Vs Container/CNF running on Intel 3rd Gen infrastructure

Journey towards Cloud-native



Classic network appliance approach

- Bespoke software and hardware appliances

Virtual network functions (VNFs)


- Virtual Appliances
- Hypervisor

Cloud-native network functions (CNFs)

- Containerized microservices
- Kubernetes (Cloud OS)

Commercial-off-the-shelf Hardware powered by Intel® Architecture

Webinar Focus!



- Q: What is "the right model"?
- A: It depends...

Clavister Portfolio *Securing identity, device, vehicle, network and cloud*




SECURITY MANAGEMENT & ANALYTICS
Centralised configuration and analytics



NETWORK SECURITY




SMB & Branch Office



[NetWall 100 Series](#)
[NetWall 300 Series](#)

Campus & Distributed Office



[NetWall 300 Series](#)
[NetWall 500 Series](#)
[NetShield 300 Series](#)
[NetShield 500 Series](#)

Datacenter & Service Providers



[NetWall 6000 Series](#)
[NetShield 6000 Series](#)
[NetShield 9000 Series](#)



[VNF](#)
[CNF](#)



CLOUD SERVICES

Security-as-a-service capabilities – delivered from European cloud platform

An aerial photograph of a suspension bridge spanning a large body of water. The bridge features two tall, slender towers and a network of cables supporting the deck. The surrounding landscape includes green hills and some industrial or construction sites in the distance. A semi-transparent dark rectangular overlay covers the central portion of the image, containing the text 'Design considerations'.

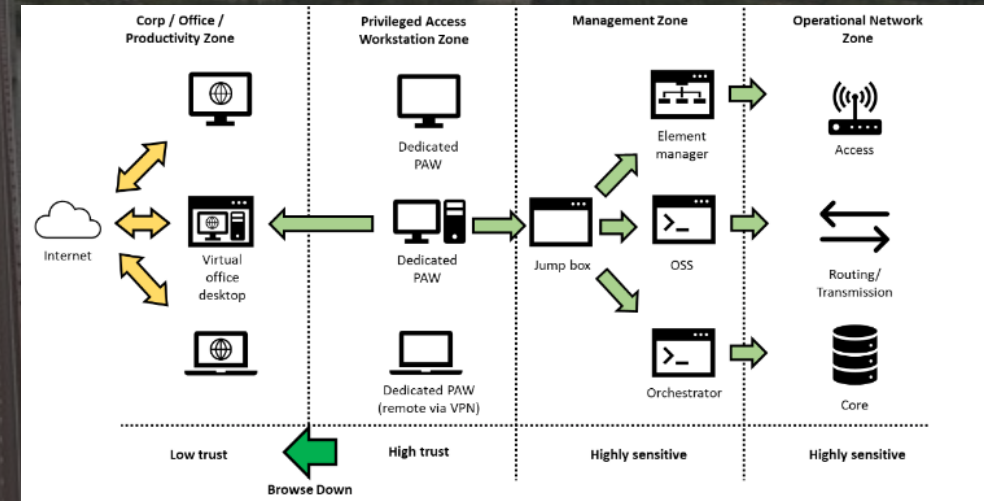
Design considerations

Zero Trust

- What is Zero Trust?
 - No implicit trust granted to assets or user accounts based solely on their location
 - Never trust, always verify
 - Protect resources rather than network segments

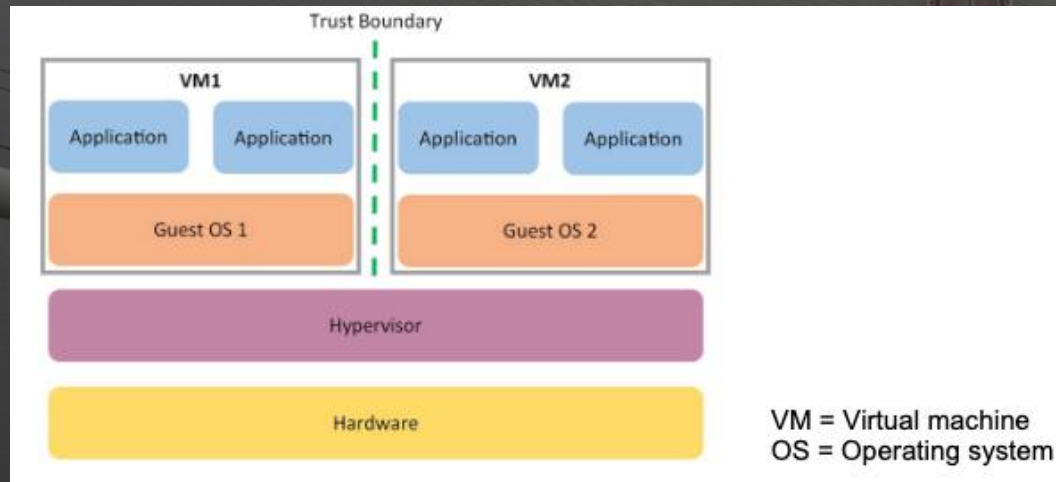
Designing for Zero Trust

- Do not trust even your own network
- Deny by default
- Separation of roles
- Separate the network into different trust domains, with:
 - Firewalls protecting the perimeter of each domain
 - Services network communication protected by the firewall
 - Dedicated hardware

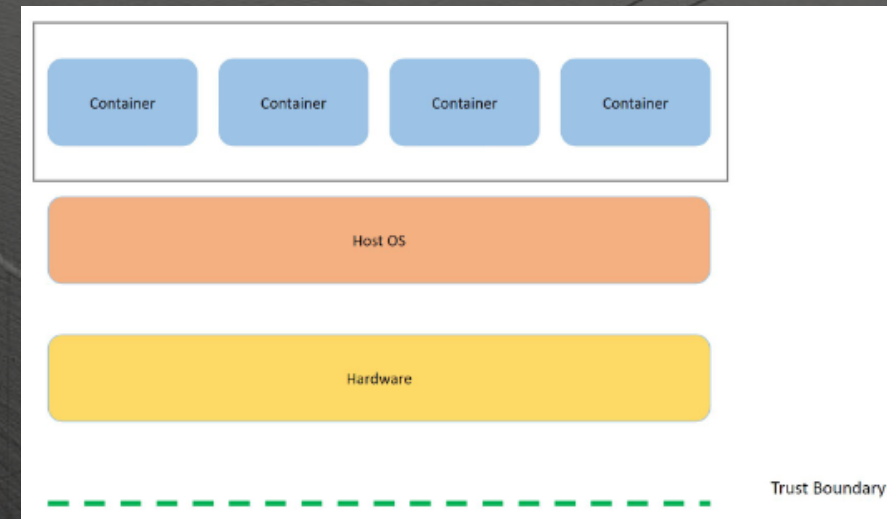


Trust Domains and performance

- Regulatory requirements are evolving – now policy in certain regions*



VNF



CNF

Security

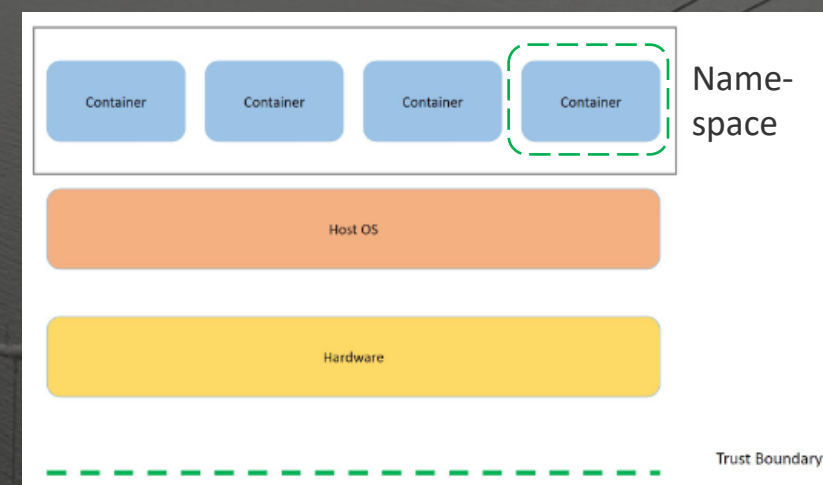


Cost Efficiency

*) https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1057446/Draft_telecoms_security_code_of_practice_accessible.pdf

Trust Domains - continued

- Namespace separation
 - Create separate namespaces for containers to prevent privilege-escalation attacks from within containers
 - Re-map users to run with less privilege on the host, outside of containers

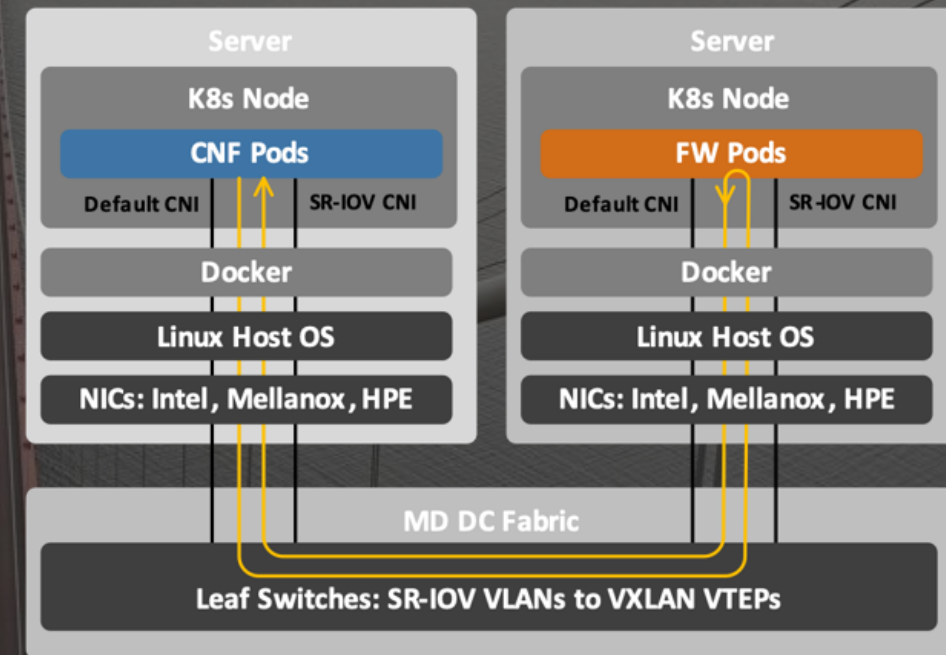


CNF

Clavister NetShield – Performance Measurements from Intel Partner Alliance (IPA) Lab

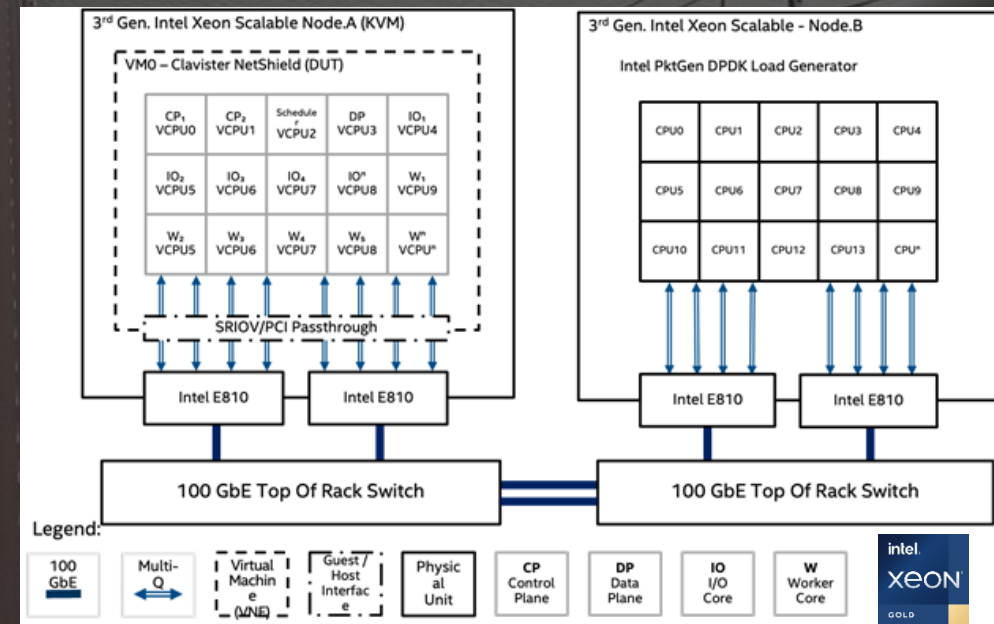
Security By Design - Firewall

- Kubernetes policies prevents communication between pods in a cluster
- NetShield Firewall CNFs:
 - May be deployed on dedicated K8s node or separate cluster.
 - Deployed to protect subnets on network overlays. (Multus CNI / SR-IOV)
 - East/west traffic between pods traverses leaf switch. Pod-to-pod communication on the same node is protected.
 - North/south traffic protected by the firewall.
 - Multiple firewalls can be deployed in parallel on the same cluster.
 - NICs: Multus CNI for SR-IOV, or Af-packet / af-xdp interfaces – if throughput requirements are low



Intel Lab Test Setup*

- Two servers were used in the test
 - 3rd Generation Intel® Xeon Scalable Processor (Intel® Xeon® Gold 6338N Processor 48M Cache, 2.20 GHz)
 - Hyper-Threading was enabled -> 64 vCPUs
 - 2 x 100GbE Intel® Ethernet Network Adapter E810
 - The same setup as was used in VNF tests*
- First server deployed NetShield CNF on Kubernetes and up to 62 vCPUs assigned
- Second server was running PktGen, a traffic generation tool built using DPDK
- The servers were connected via a 100 GbE top of rack switch (200 Gbps max line rate)



Intel Lab Test Setup

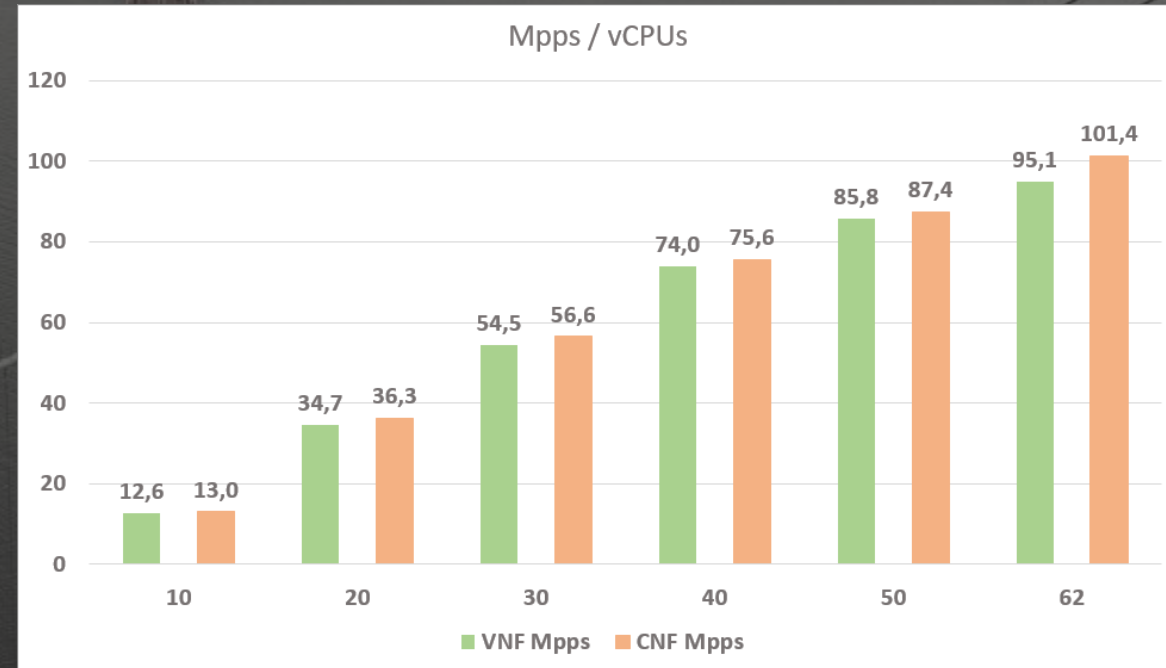
- The NetShield CNF used single root I/O virtualization (SR-IOV)/PCI passthrough to mediate traffic flow
- CPU isolation
 - To avoid workloads contending for available CPU resources
 - Ensures that physical cores are used exclusively by the NetShield Pod
 - Useful for performance sensitive use cases
- Pod specification
 - Static policies for CPU and memory
 - QoS class set to guaranteed

Intel Lab Test Setup

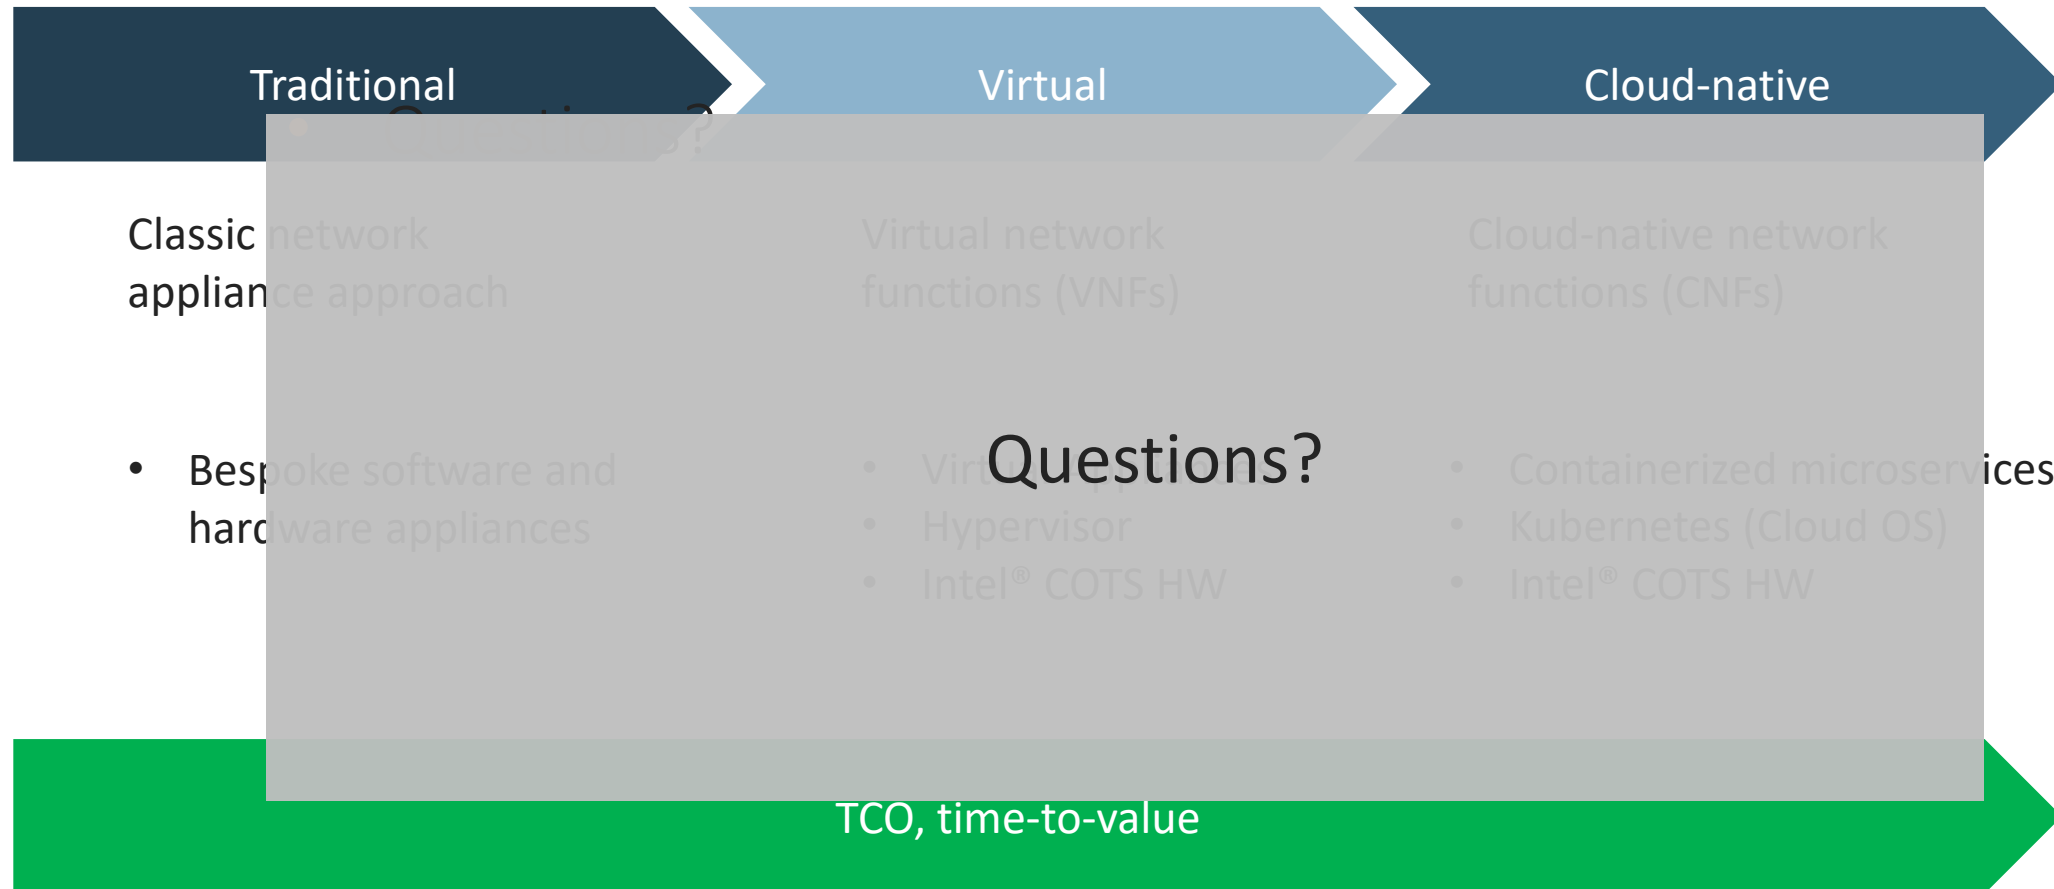
- NetShield CNFs tests run using different amounts of vCPU to visualize performance scaling
 - 10, 20, 30, 40, 50 and 62 vCPUs
 - 2 vCPUs left for the rest of the system
- Pktgen – traffic generation tool based on DPDK
 - RFC 2544 benchmark to test highest possible UDP packet throughput without packet loss
 - Tests run with different packet sizes: 64, 128, 256, 512, 1024 and 1518 bytes
- Test result emphasis on packets-per-second
 - Test setup max line rate was 200 Gbps
 - Results reach line rate at higher packet sizes – same result as we saw when testing VNF throughput

VNF vs CNF – The Lab Results

- NetShield performance scales with number of CPU cores assigned
- Chart shows millions of packets-per-second with 64B packet size
- Test results shows us that performance scales regardless of deployment model – VNF, CNF



Journey towards Cloud-native



- Q: What is "the right model"?
- A: It depends...

Thank You!

