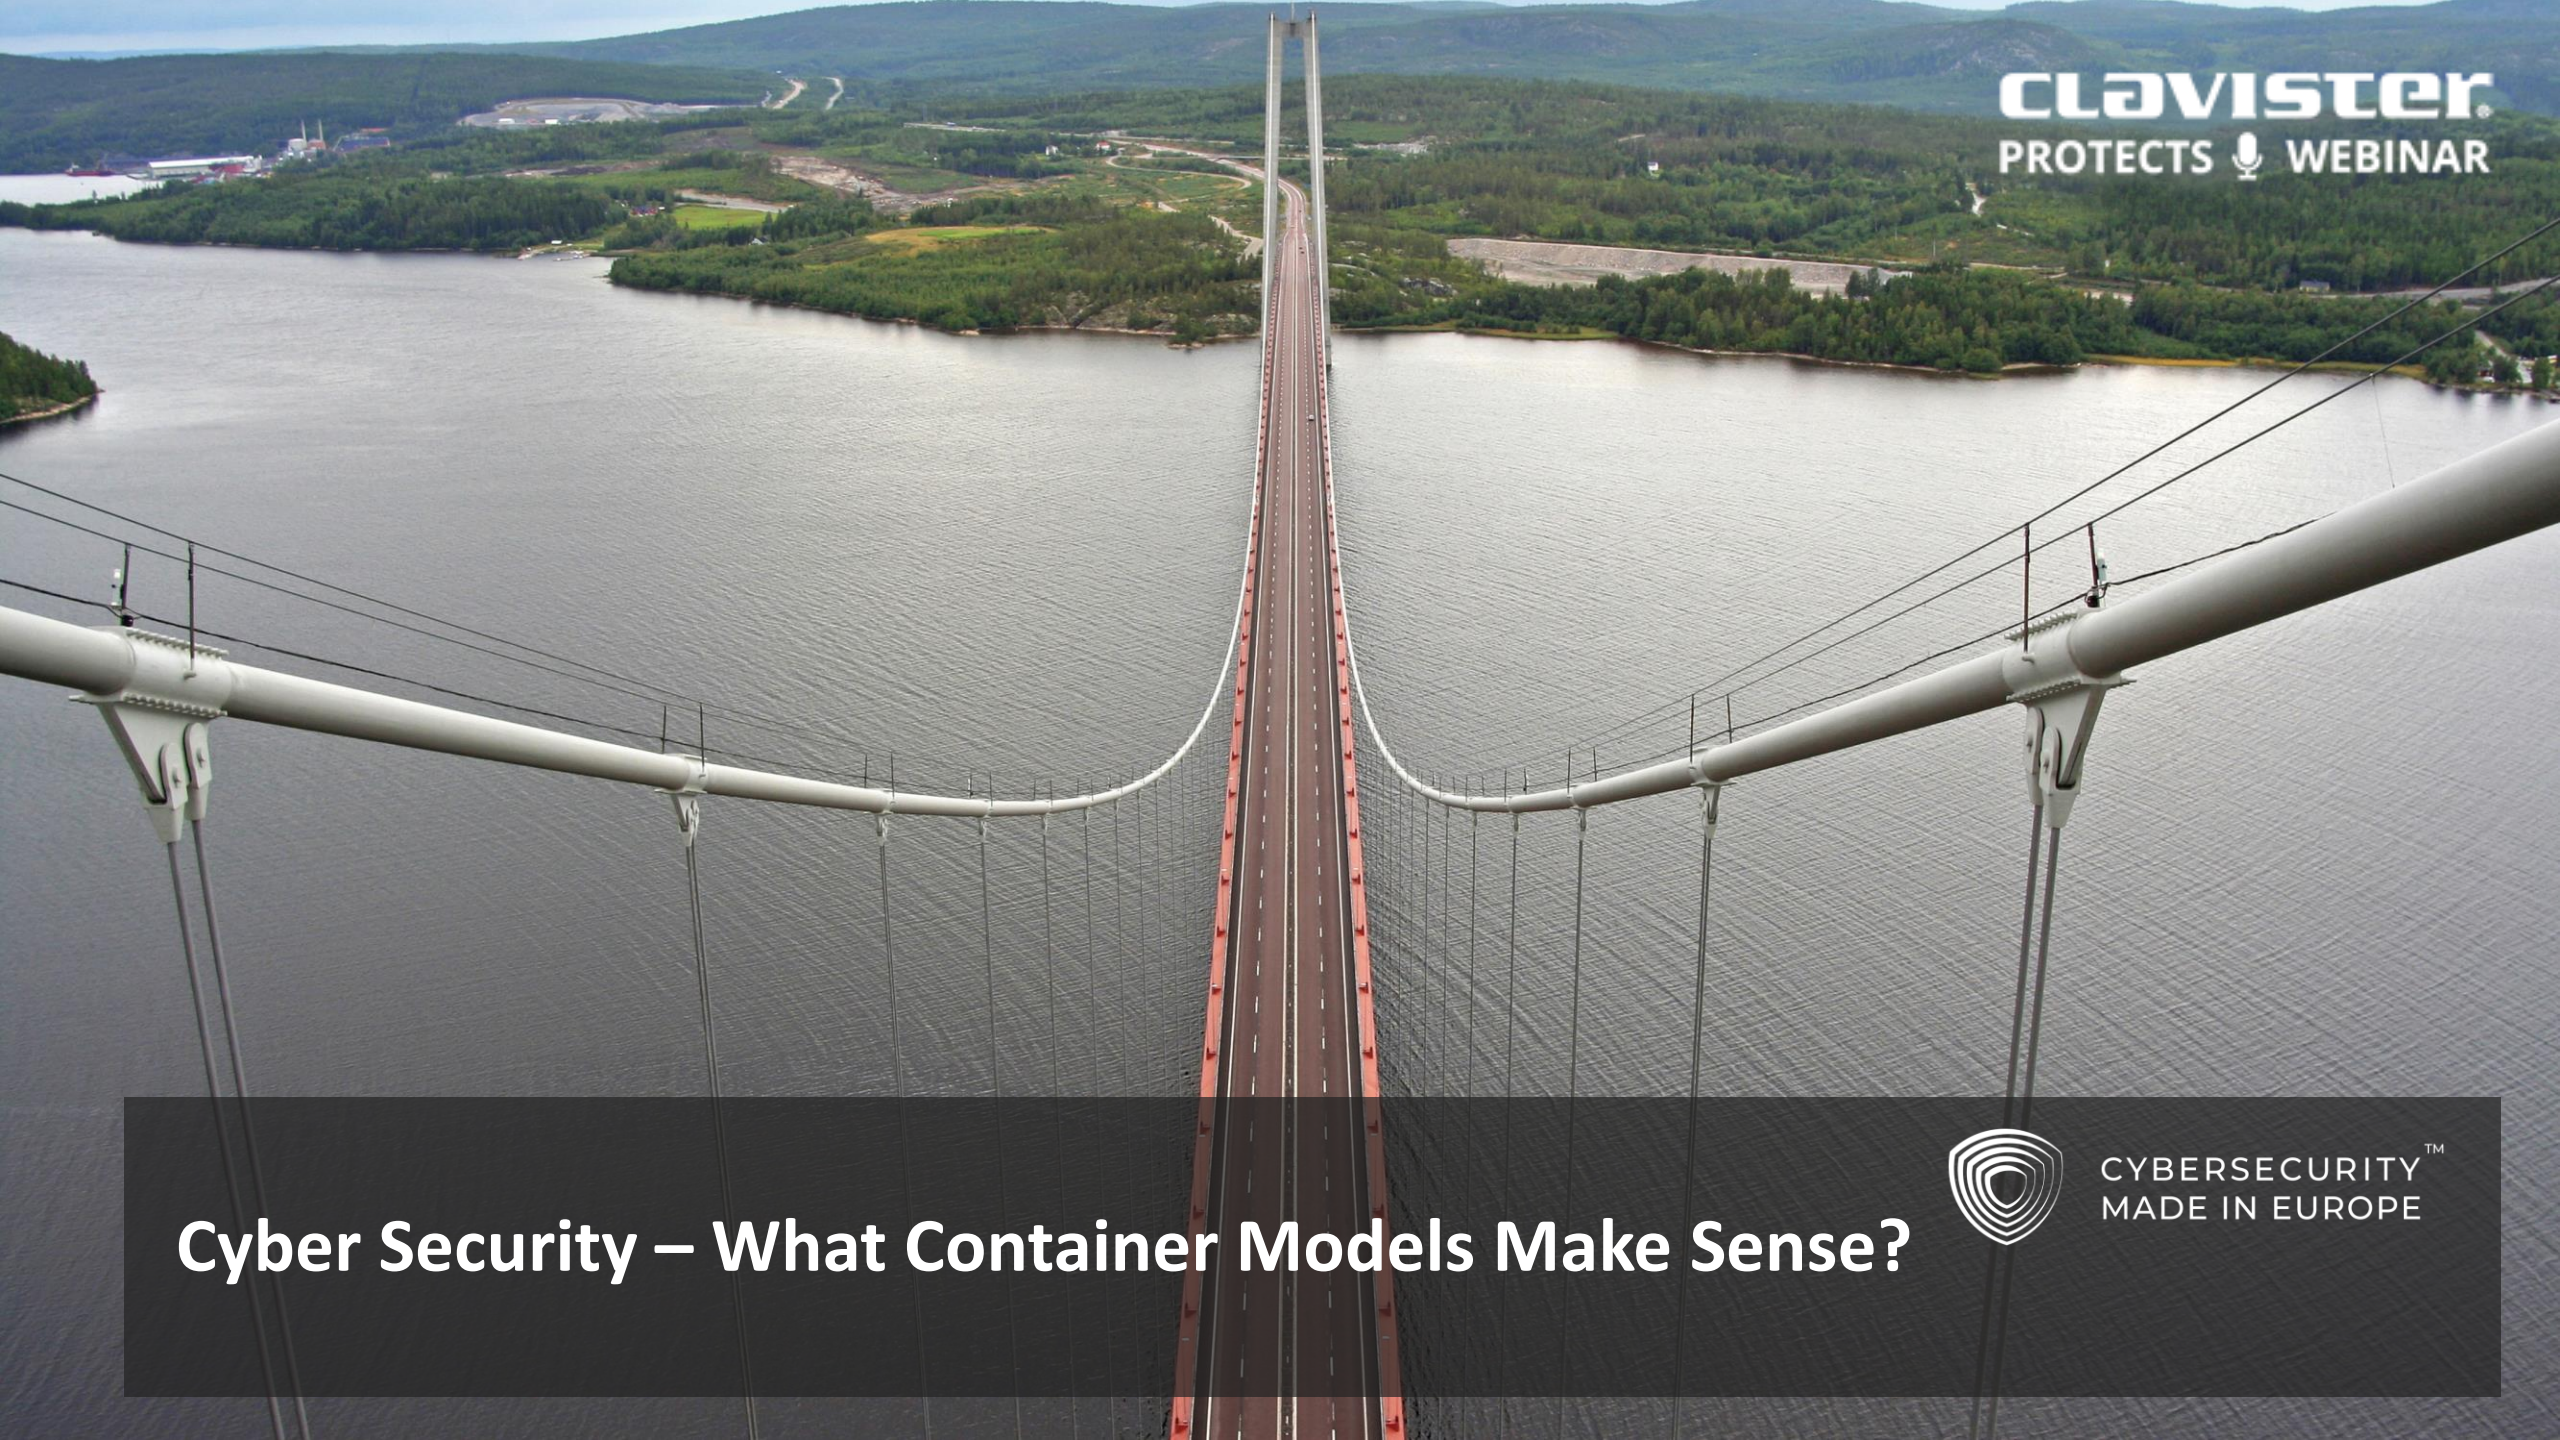

Cyber Security – What Container Models Make Sense?

Introductions

Nils Undén

CTO

Nils.unden@clavister.com

Mattias Fredriksson

Product Owner

Mattias.fredriksson@clavister.com

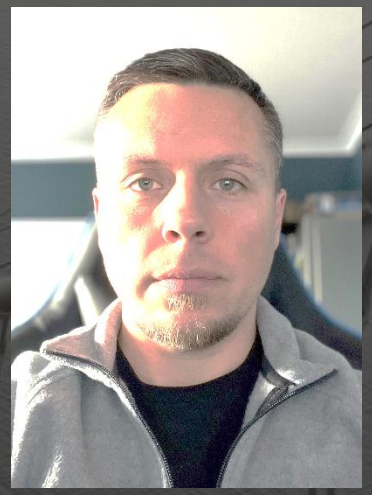

Dave Cremins

Cloud Software Architect

dave.cremins@intel.com

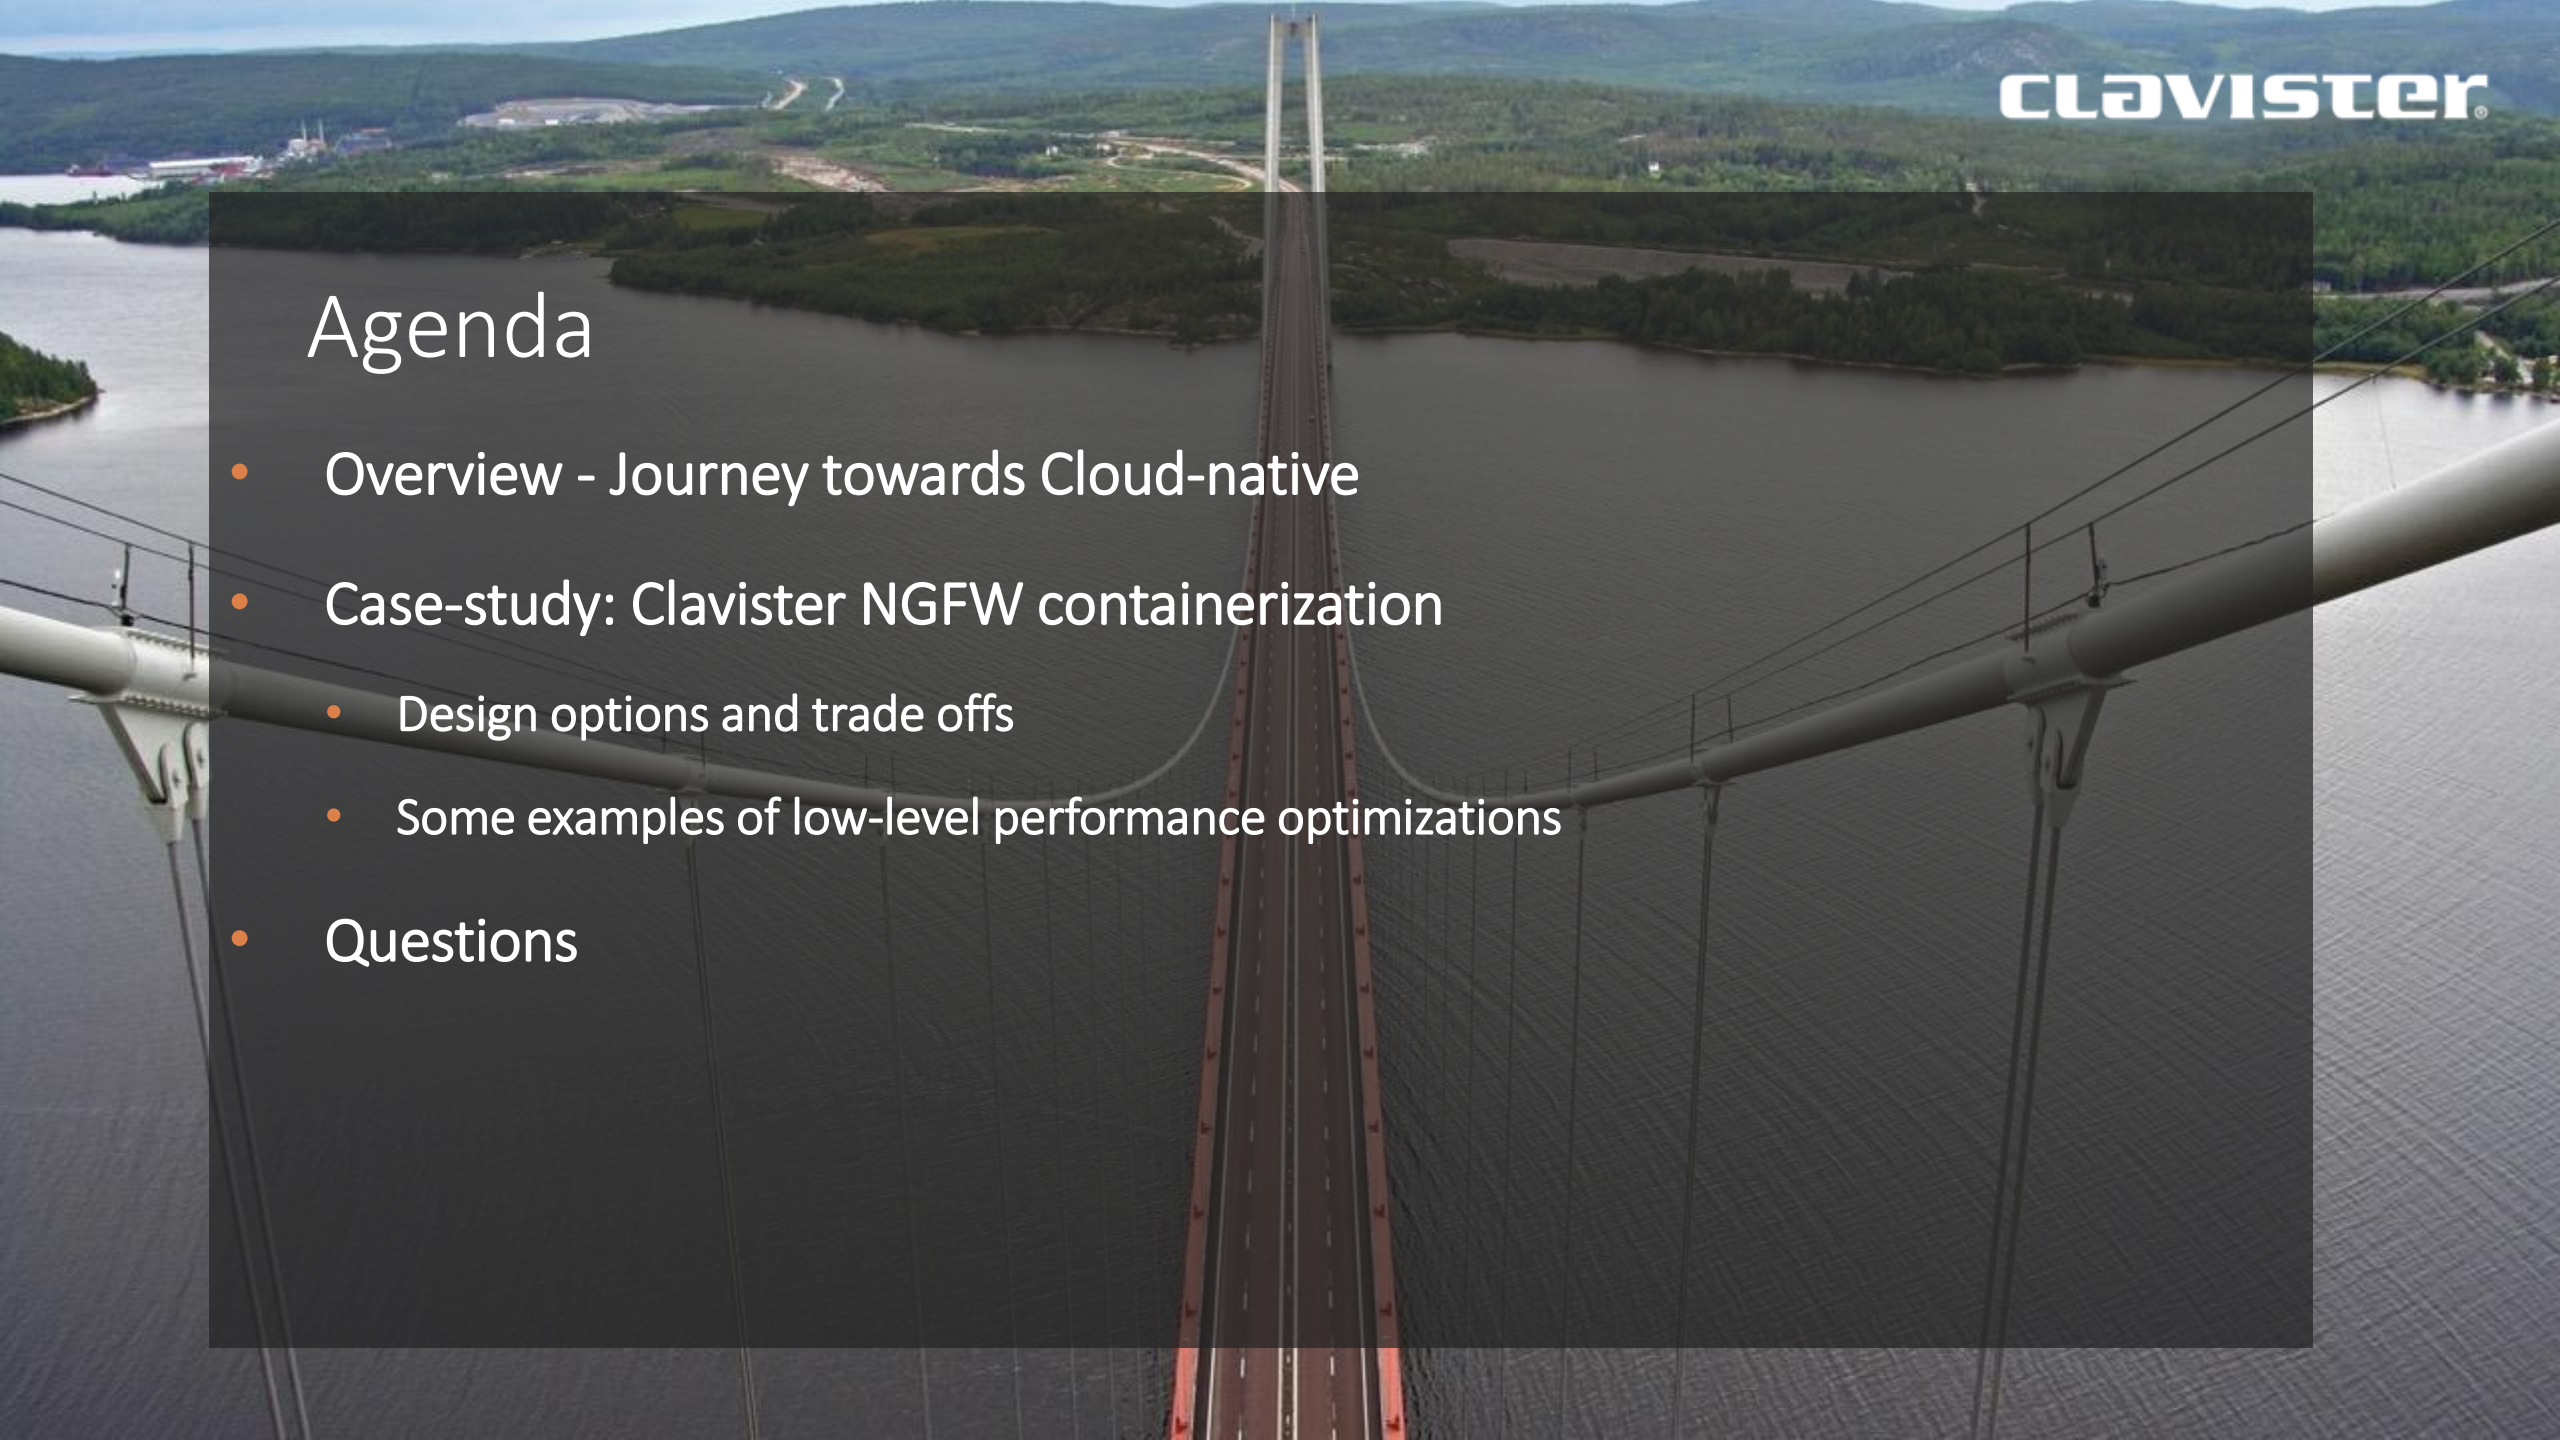

Agenda

- Overview - Journey towards Cloud-native
- Case-study: Clavister NGFW containerization
  - Design options and trade offs
  - Some examples of low-level performance optimizations
- Questions

Journey towards **Cloud-native**

Traditional	Virtual	Cloud-native
Classic network appliance approach	Virtual network functions (VNFs)	Cloud-native network functions (CNFs)
• Bespoke software and hardware appliances	• Virtual Appliances • Hypervisor • Intel® COTS HW	• Containerized microservices • Kubernetes (Cloud OS) • Intel® COTS HW

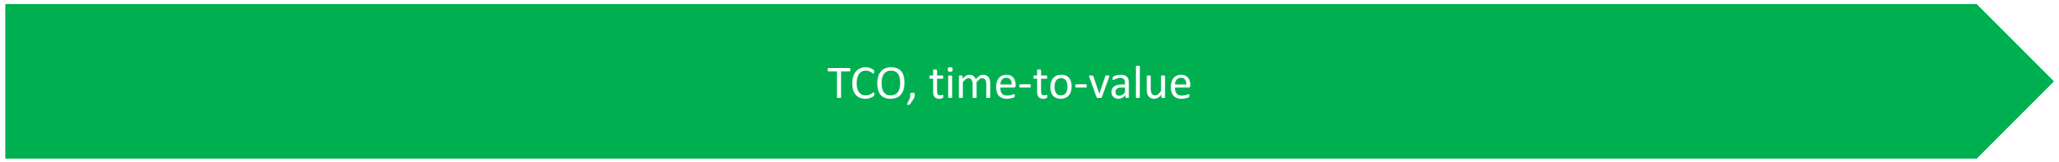

- Q: What is ”the right model”?
- A: It depends…

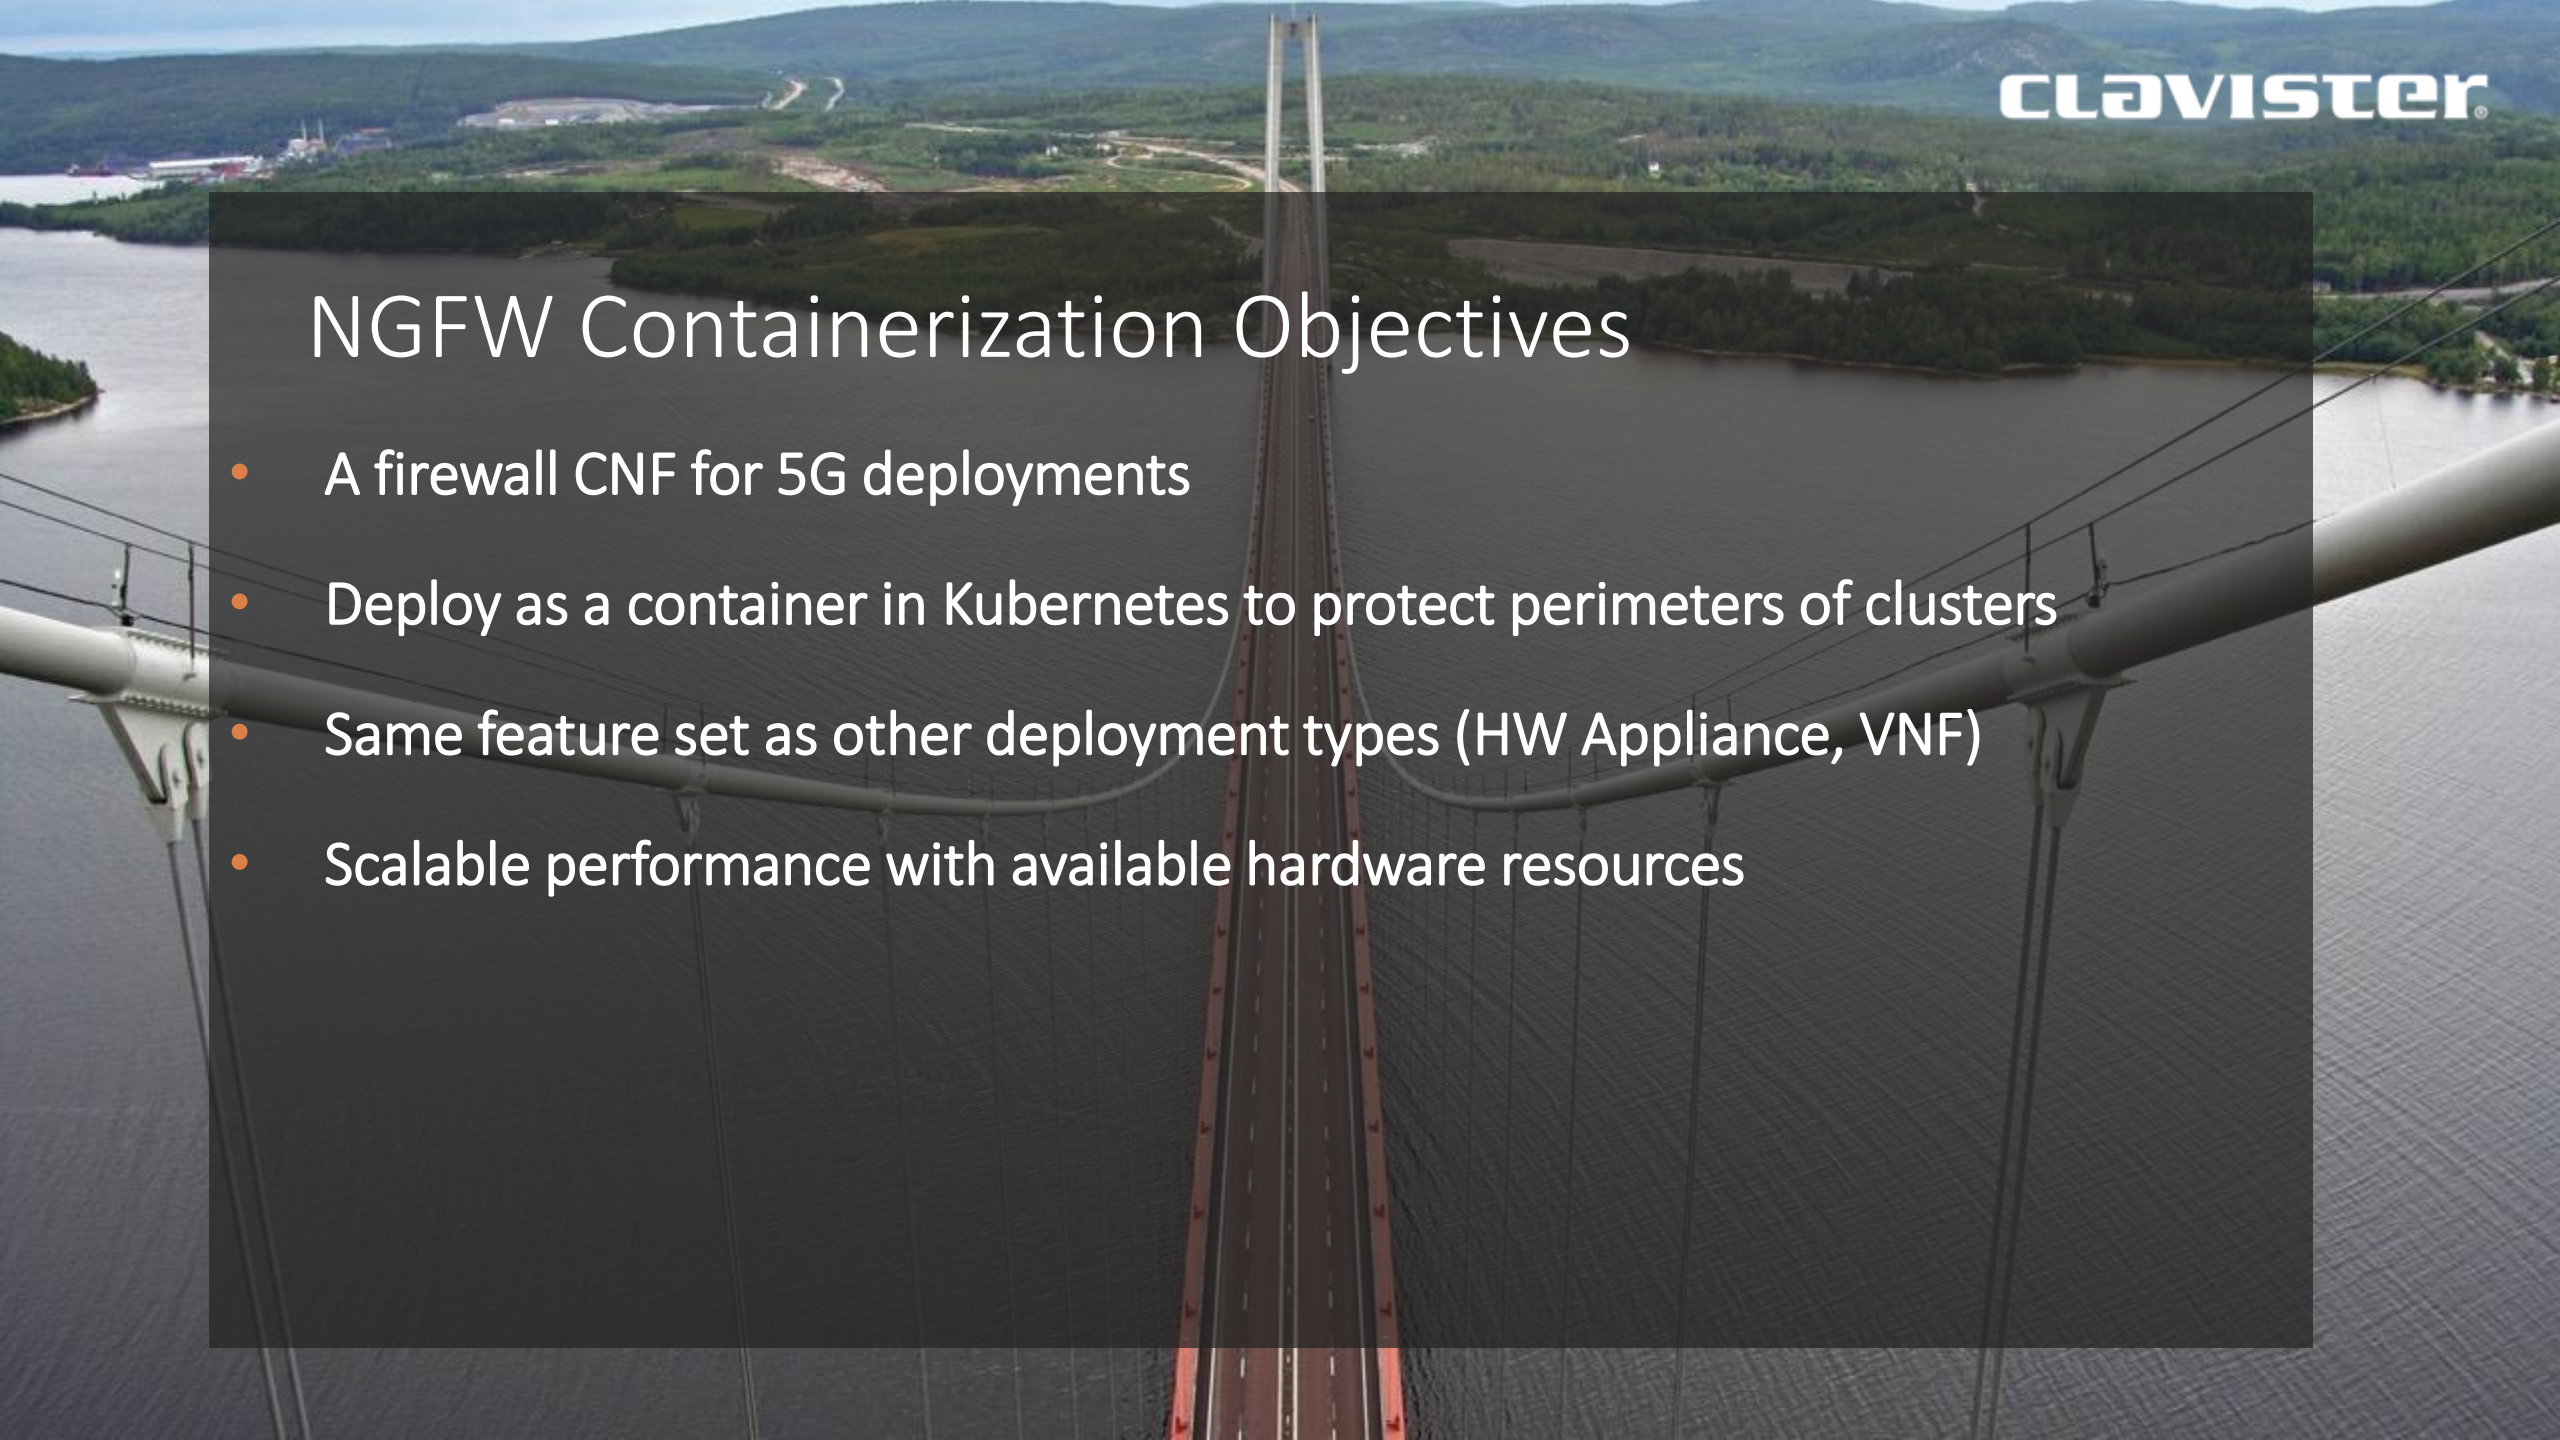

NGFW Containerization Objectives

- A firewall CNF for 5G deployments
- Deploy as a container in Kubernetes to protect perimeters of clusters
- Same feature set as other deployment types (HW Appliance, VNF)
- Scalable performance with available hardware resources

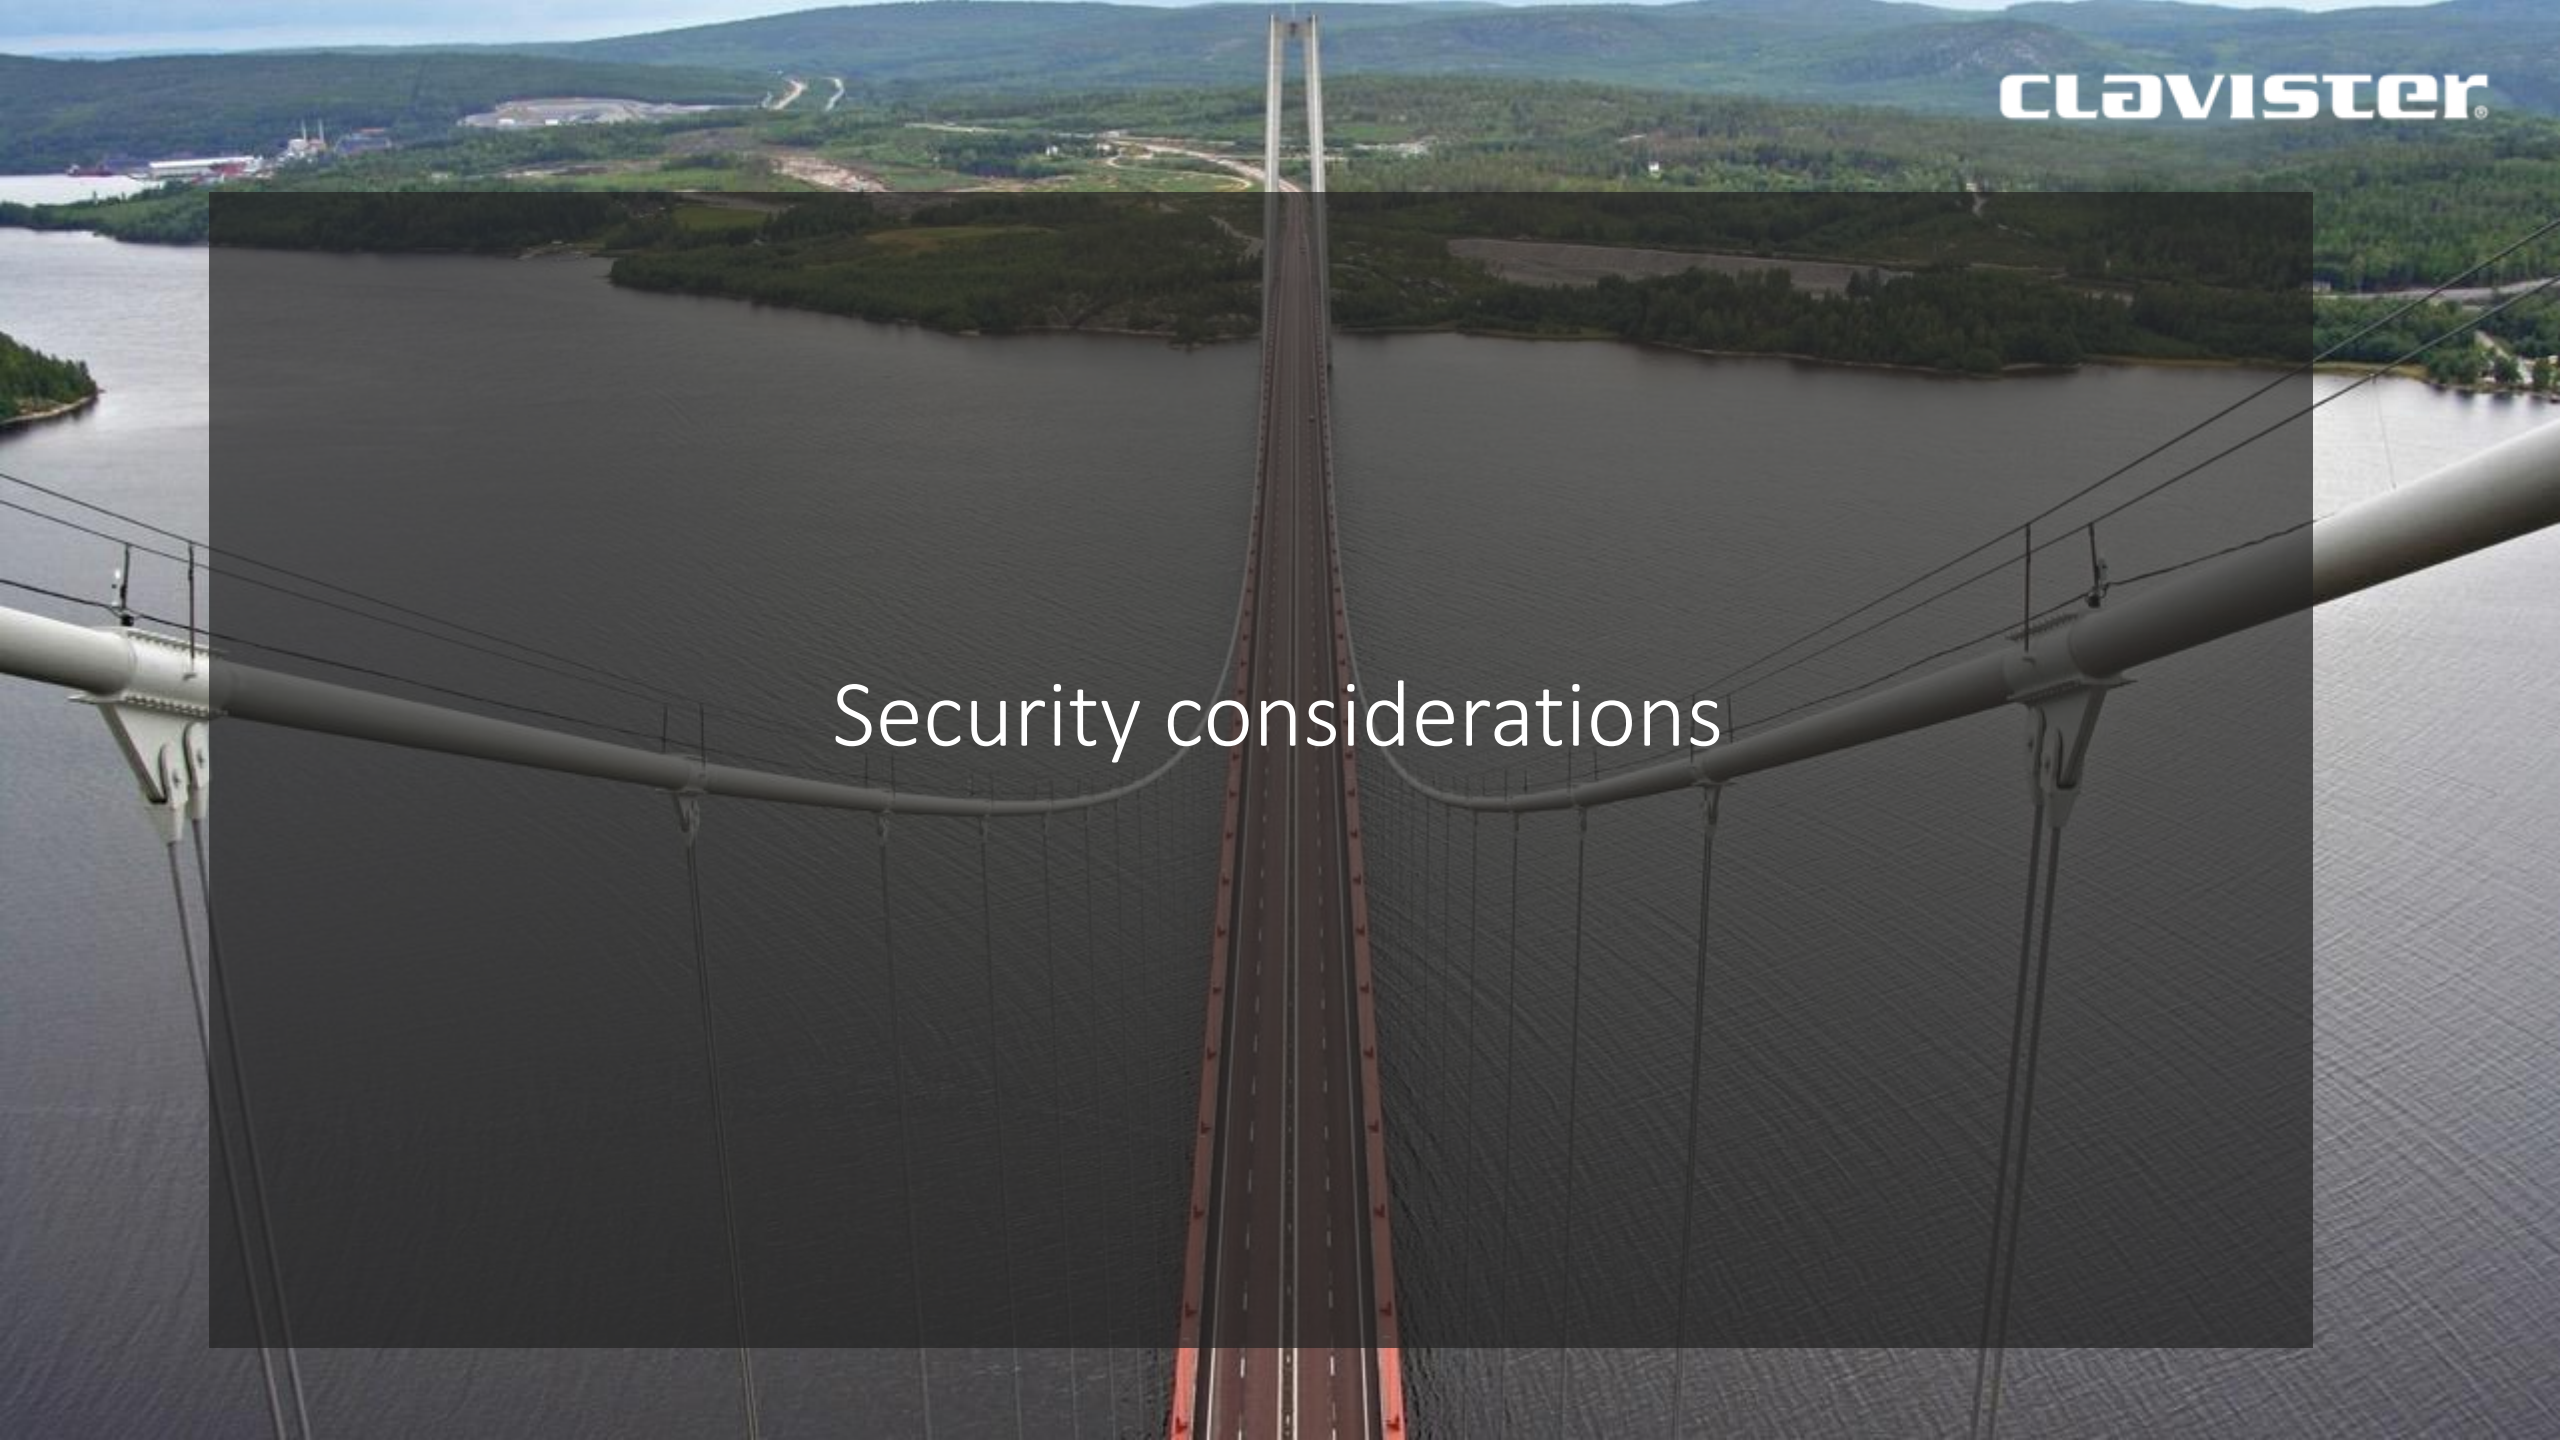

Security considerations

5G Distributed Cloud

- Distributed to improve latency and bandwidth
- End-to-end encryption

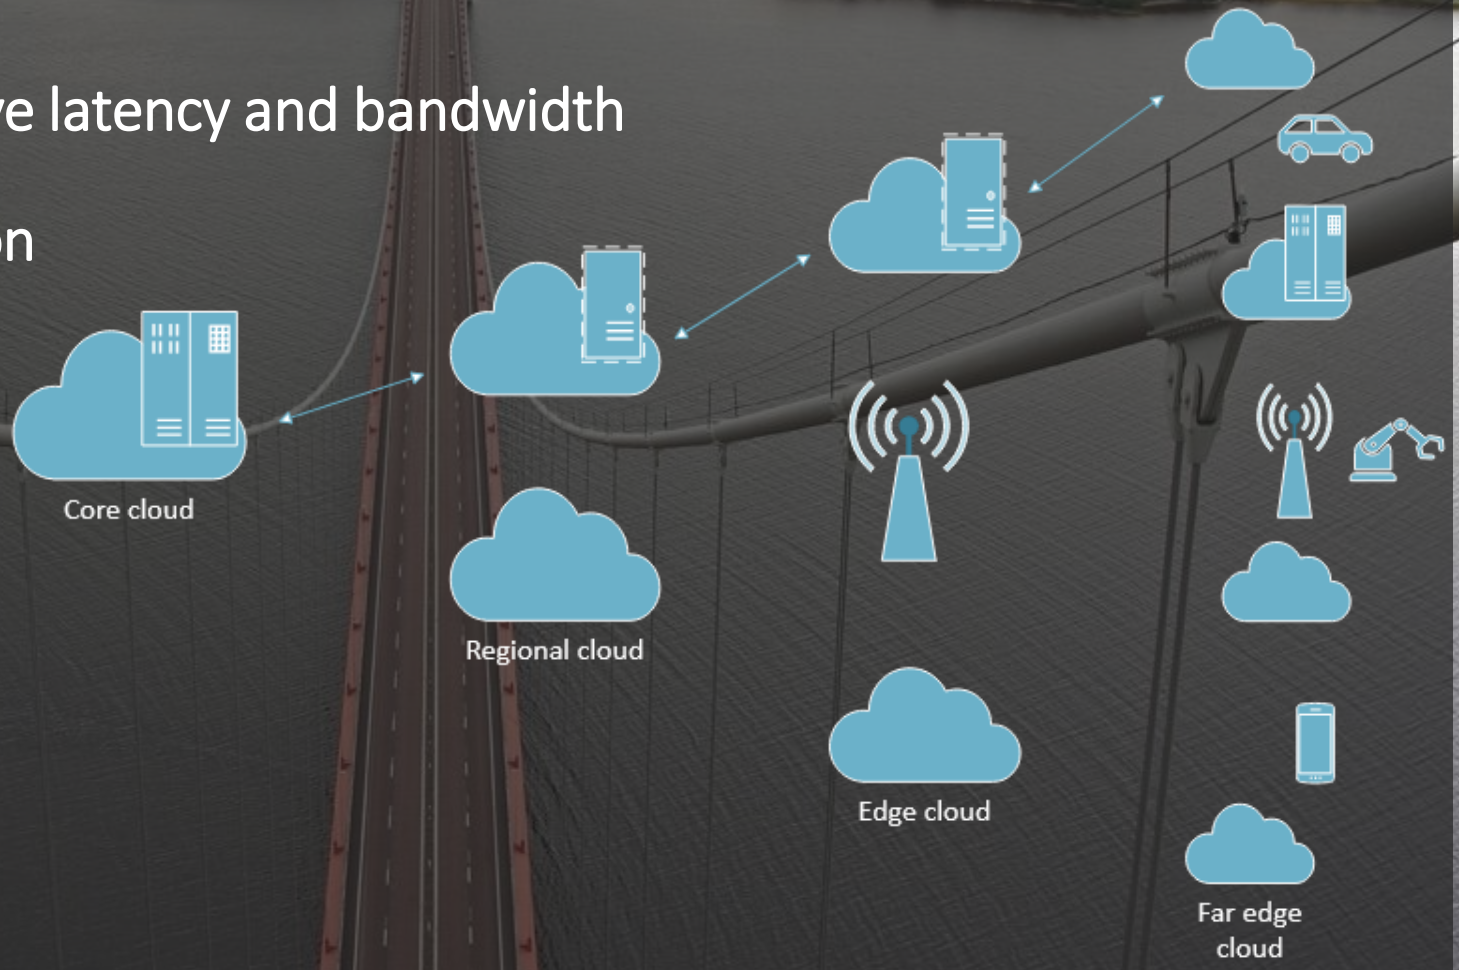


Designing for Zero Trust
What is Zero Trust?
No implicit trust granted to assets or user accounts based solely on their location
Never trust, always verify
Protect resources rather than network segments

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

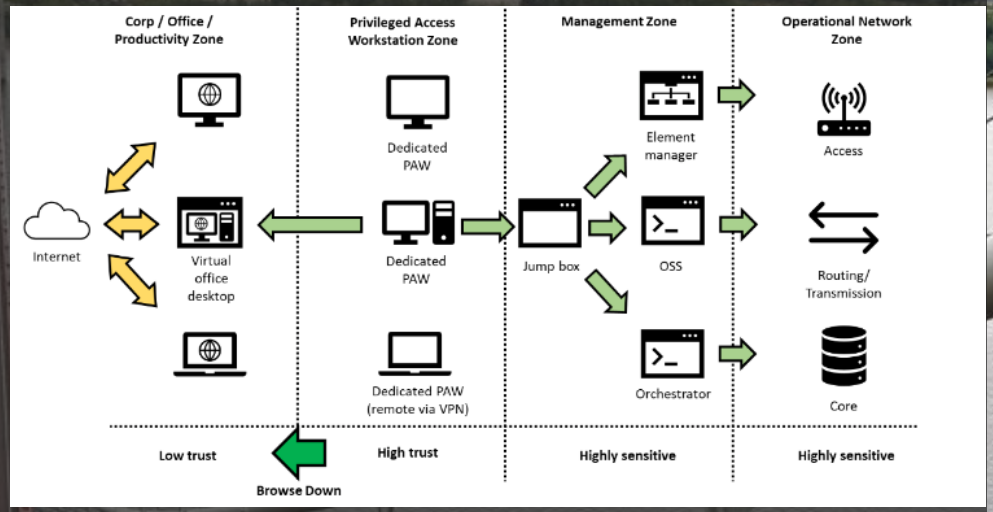

Implementing Trust Domains

- Regulatory requirements are evolving – may turn into policy*

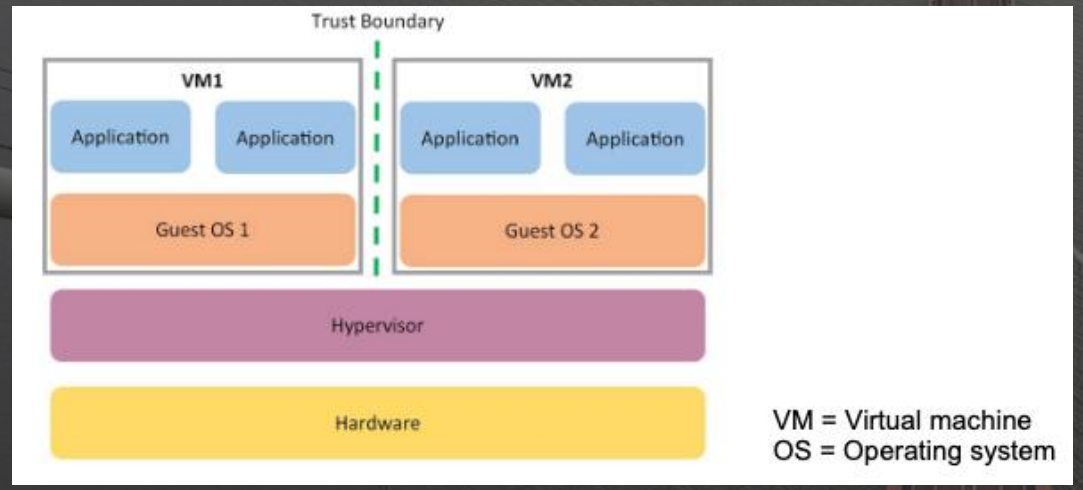

VNF

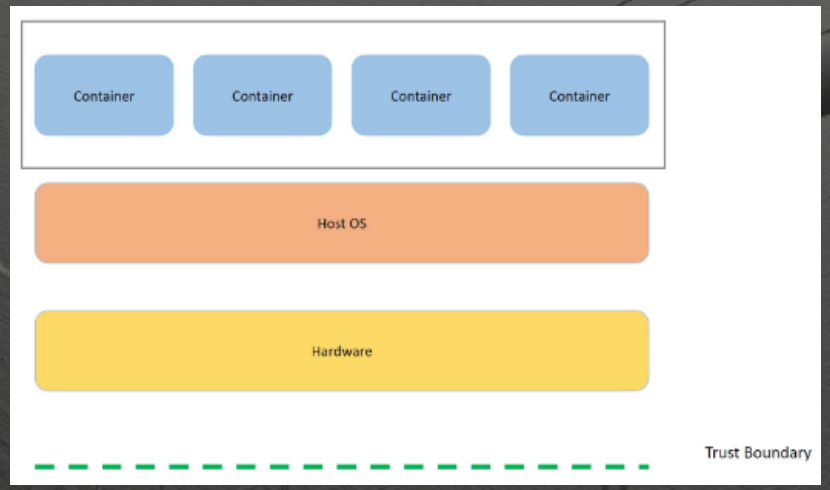

CNF

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

Implementing 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

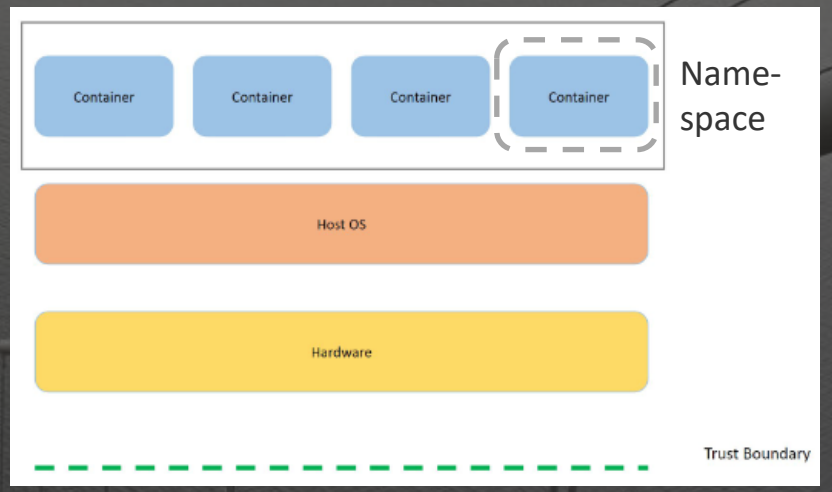

CNF

Security by Design - Networking

- The cluster network is typically used for management
- Additional network attachments depends on the CNF
- Intra-CNF communication:
  - Is done over dedicated subnets on L2/L3 network overlays
  - Always traverse leaf switches

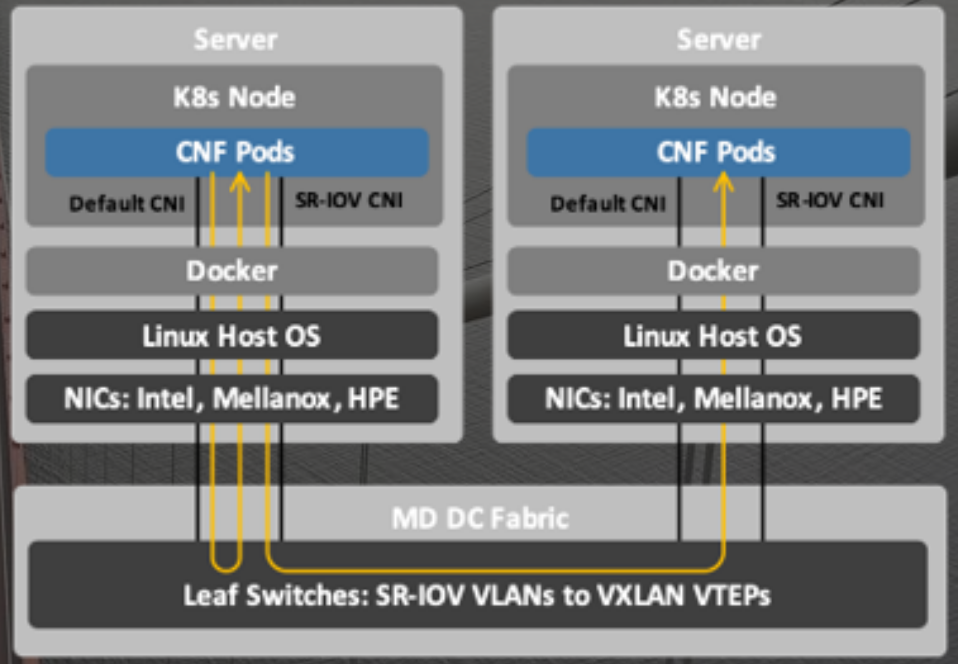

Security by Design - Networking (cont...)

- SR-IOV* VLANs
  used to identify network overlays
- Traffic between egress switches and ingress leaf switches via VXLAN tunnelling
- Interfaces: Intel®, Mellanox®
- CNIs:
  Calico, Cillium, Multus with IPVLAN/SR-IOV

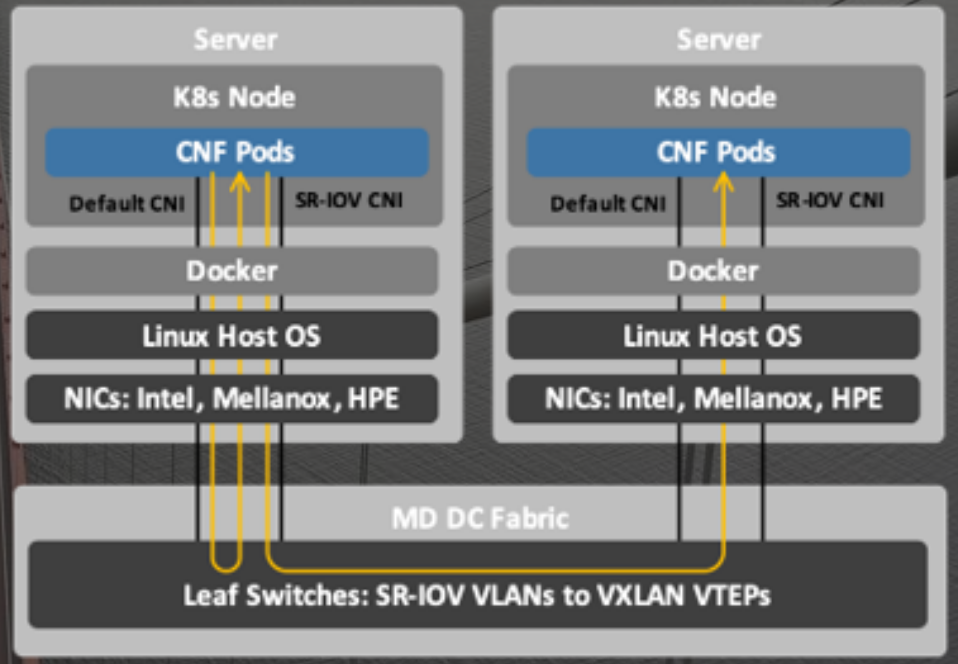

* Single-root input/output virtualization

Clavister NetShield as a container in Kubernetes

NetShield in a 5G deployment

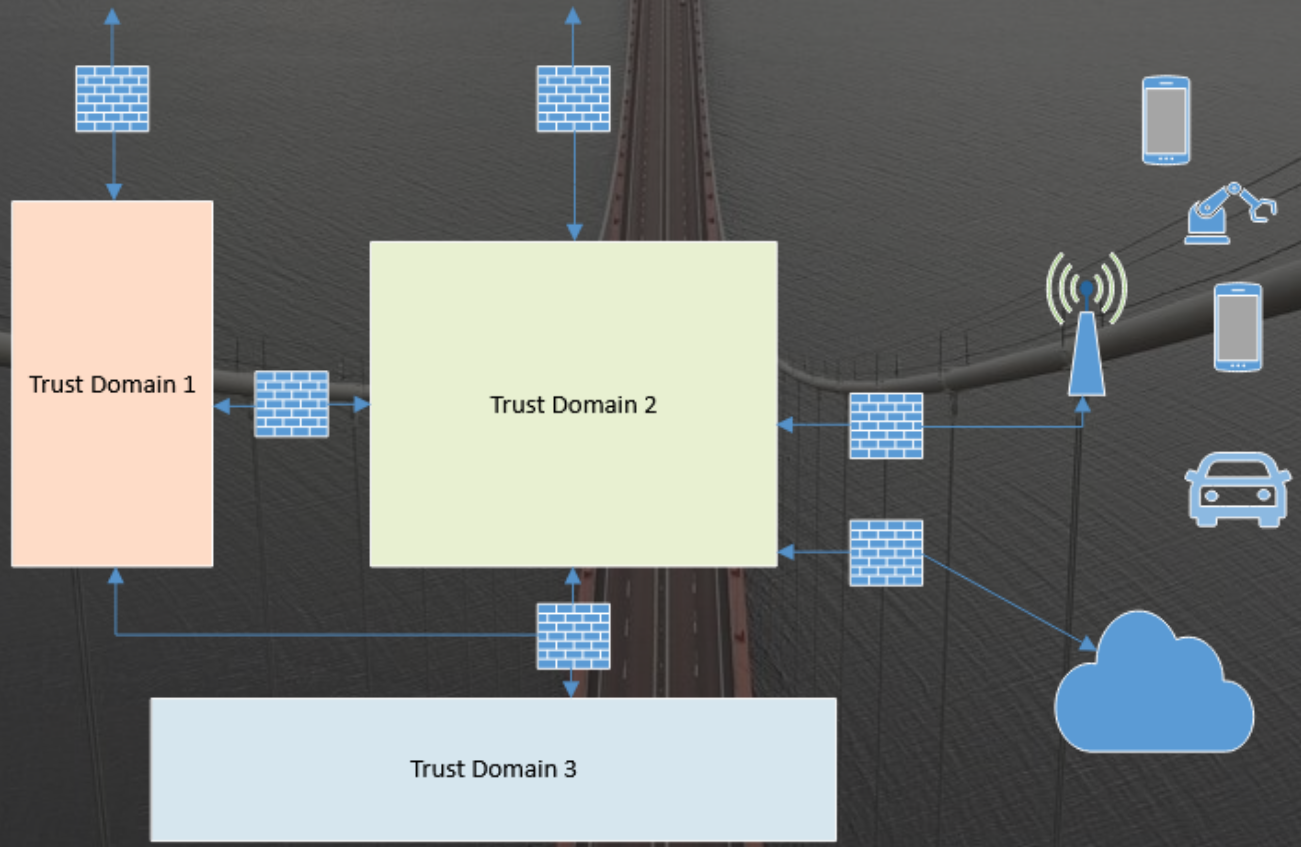

Security By Design - Firewall

- 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.
  - 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.

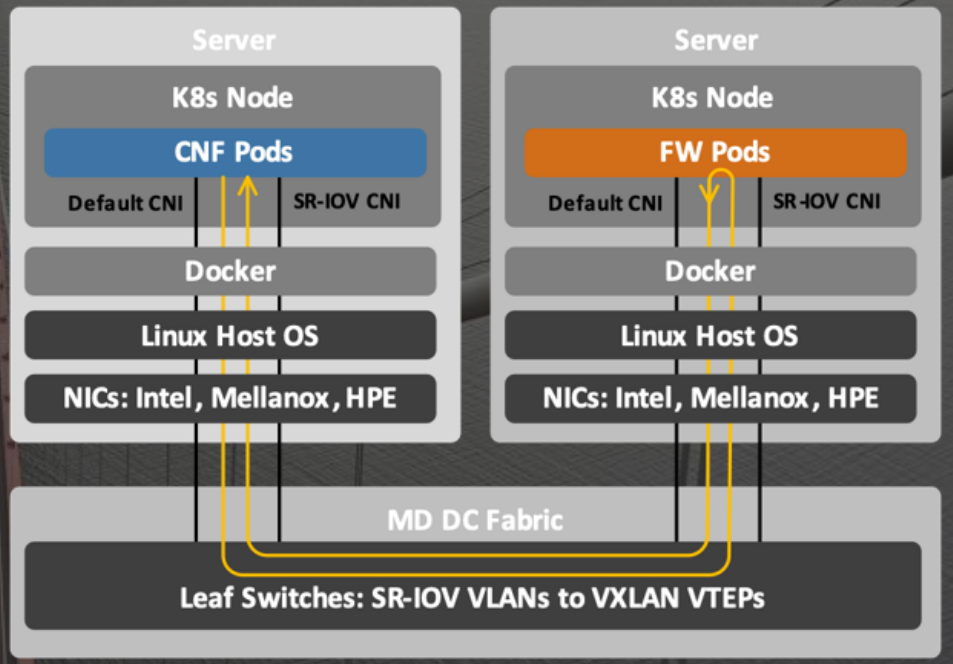

Performance Centric NGFW

- NetShield is based on DPDK – performance scales with number of CPU cores assigned*
  - Based on previous test run together with Intel in 2022.
  - NetShield was running as a KVM VM during that test.
  - Intel® Xeon® Gold 6338N processors running at 2.20 GHz.
  - Intel® Ethernet Network Adapter E810-C 100Gbps Dual-port NICs.
- Static policies for CPU and memory
- Set QoS class to guaranteed
- Interfaces
  - SR-IOV interfaces - enabling high bandwidth NIC:s
  - Af-packet/af-xdp interfaces - an alternative to SR-IOV interfaces if bandwidth requirements are low

*) https://networkbuilders.intel.com/solutionslibrary/clavister-netshield-delivers-scalable-performance-up-to-95-mpps1

High Availability

- K8s provides redundancy by design
- When state-synchronization is required to yield seamless failovers
  - Active-Passive HA pair
  - Anti-affinity – two alternatives
    - HA pair runs as a set of Pods in the same Kubernetes cluster – but on different nodes/hardware
    - HA pair runs as a set of Pods on different Kubernetes clusters – on different hardware
  - Multus is needed for additional pod interfaces
  - SR-IOV and an external switch is required

NGFW Manager Overview

- Clavister NGFW manager – InCenter – runs as a pod in the same Kubernetes cluster as the firewalls
- Multiple firewall HA-pairs can run in the same Kubernetes cluster

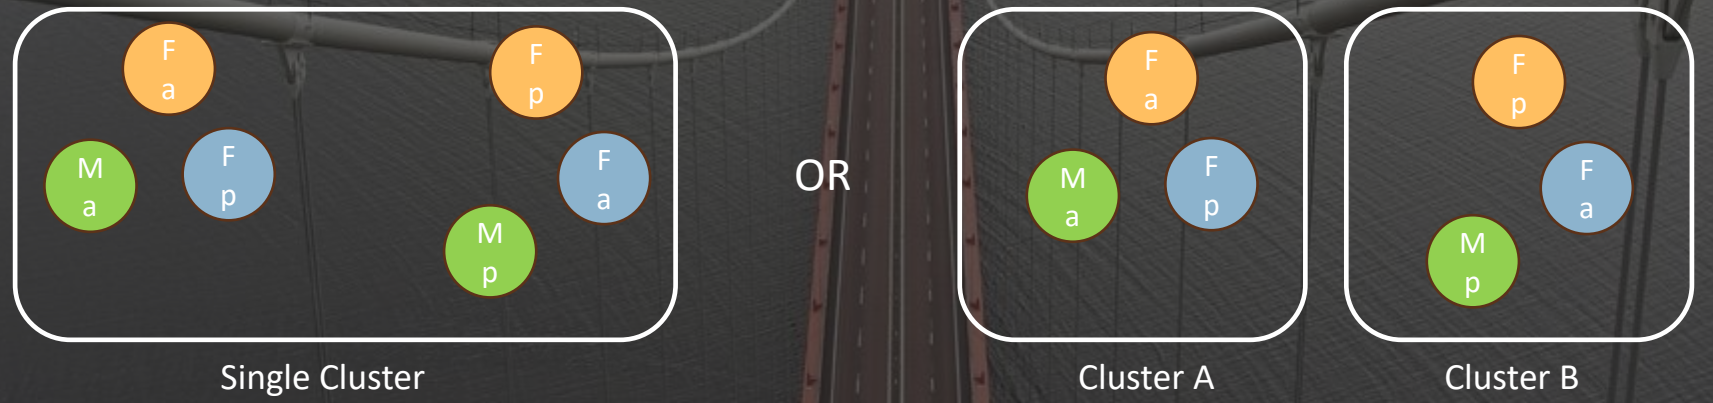

Journey towards **Cloud-native**

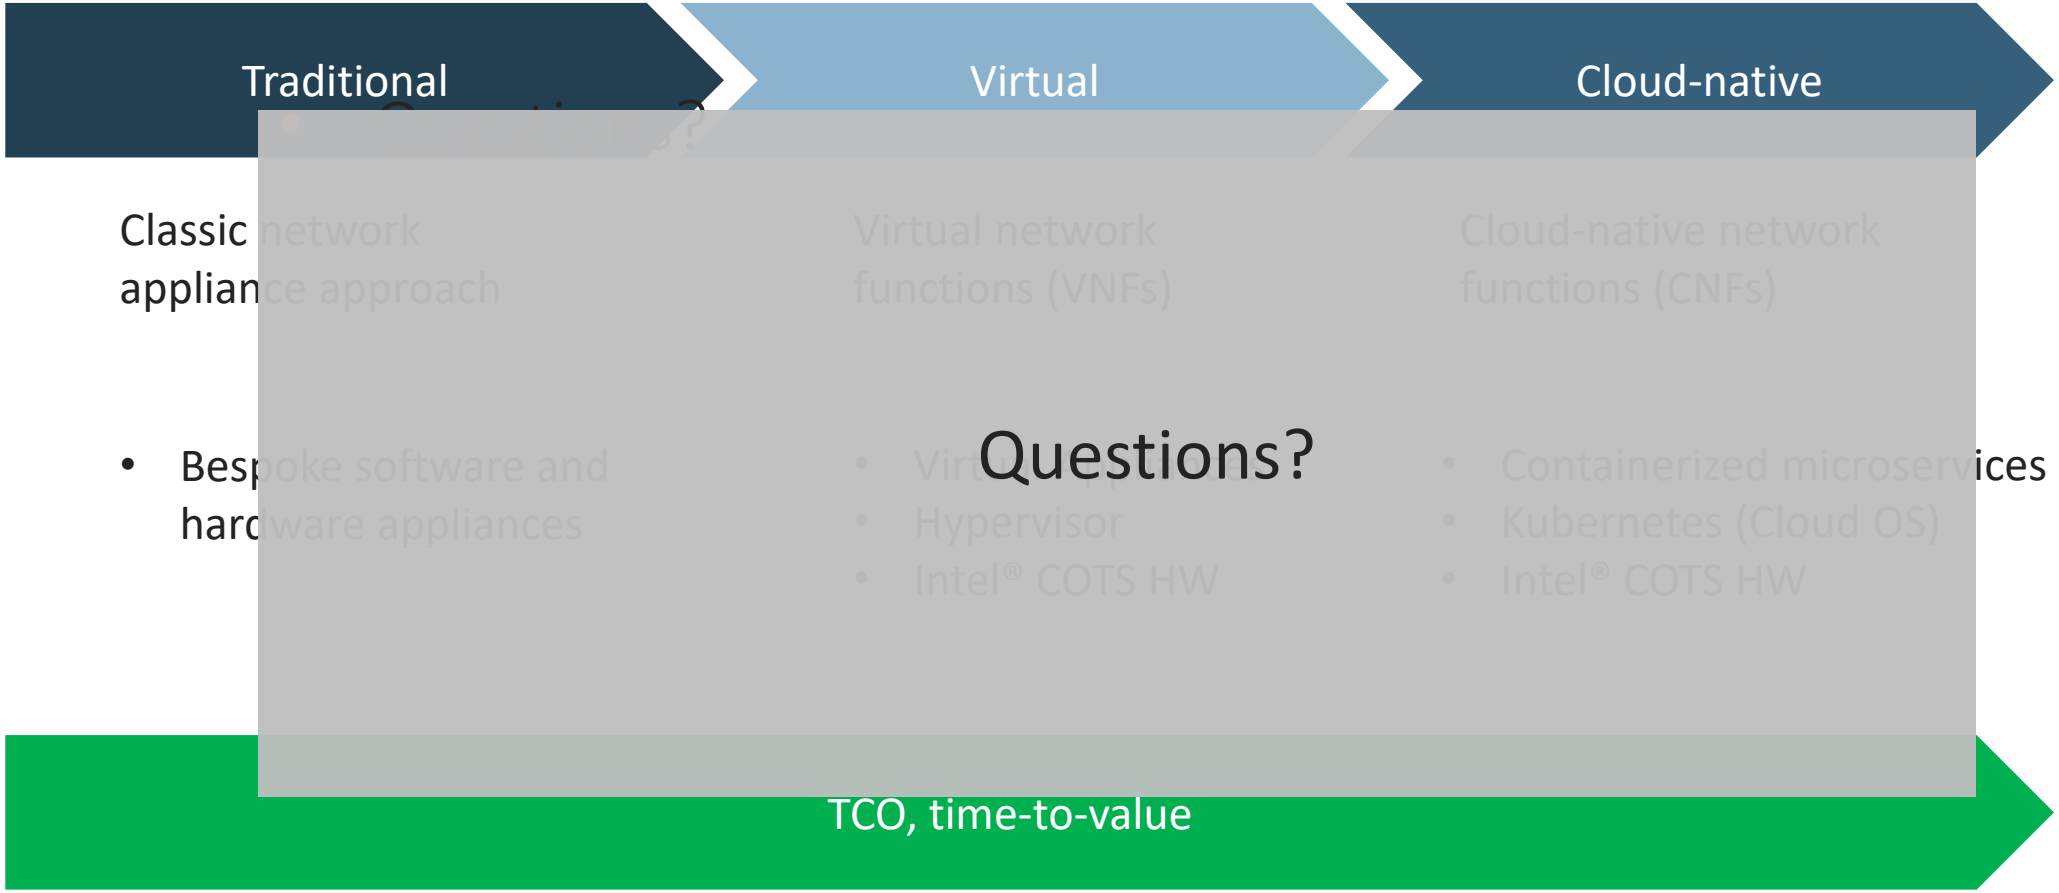

- Q: What is ”the right model”?
- A: It depends…


Thank You!
SECURITY BY
SWEDEN