

From Core to Edge: A Unified Approach for Sovereign Data and Agentic AI

Tomas Michaeli

Technical Strategy Lead, Nord & CEMEA

tomas.michaeli@broadcom.com

Sep-2026



What are the prominent Edge use cases?

Retail



- Loss & fraud prevention
- Personalized Shopping
- Immersive Retailing
- Improve compliance & safety

Manufacturing



- Predictive Maintenance
- AR/VR Enabled Remote Operations
- Maximize production uptime and operating efficiency

Logistics/Transportation



- Real-time inventory tracking
- Improved fleet operations
- Safety & security
- Efficiency improvement

Energy




- Increase production visibility and efficiency
- Virtualized Protection and Automation functions
- Maintain regulatory compliance

What are the key customer challenges at the Edge?



Limited Network

High latency, low bandwidth, extended outages, outbound only, NAT



Limited Footprint

Smaller compute capacity with long hardware refresh cycles, edge devices support



Scale (Sites)

1000s of distributed near-identical sites, geographically distributed



Lack of Local IT

Zero to limited local IT skill, truck roll is both expensive and time consuming




Privacy & Security

Need to protect data and IP from cyber and local risks

Edge Topologies


Connected across WAN and Disconnected



P P = Container on ESX

Edge Design - Without VCF Full Stack


VCF Ops with vCenter for distributed edge



- VCF Ops (mandatory) used for central monitoring and licensing
- vCenter manages multiple distributed edge sites
- Support for single host and multi-host cluster deployment
- **300ms latency** between vCenter and ESX hosts FROM VCF 9.1.1
- Deploy VCF Ops and VC manually or using VCF Installer with VVF option and entitle with VCF Edge license
- Deploy VCF Ops Cloud Proxy can work with link disruption, needed for EDGE with unstable connectivity

With VCF full stack

Centrally Managed Edge Compute Clusters in a Workload Domain



- Centralized management domain
- vCenter and NSX deployed in WLD domain at DC
- Create clusters for each site within same WLD domain
- Support for 2-Node cluster with external storage and 3-Node vSAN clusters
- Support for mixed storage environments

Multiple VCF Edge 9.1 Design Patterns

Use-case driven designs for flexibility

VCF Full Stack

- [Centrally Managed Edge Compute Clusters in a WLD](#)
- [Segregated Edge deployment across Multi-region WLDs](#)
- [Locally Managed Edge](#)
- [VCF Fleet with Multi-Region, Distributed Edge](#)
- [Dark Site Edge](#)
- [vSphere Supervisor at edge in a VCF Instance with Multiple Domains](#)
- [Supervisor Zonal Separation at Edge](#)
- [Single host vSphere Supervisor at edge in a VCF instance](#)
- [Leveraging Existing VCF Operations for Far Edge Expansion](#)


Without VCF Full Stack

- [VCF Operation with vCenter for Distributed Edge](#)
- [Thin Edge with vSphere Supervisor on Two-Node vSAN](#)
- [vSAN Witness Placement for Distributed Edge](#)

[VMware Ports & Protocols](#)

(<https://ports.broadcom.com/network-diagrams/VMware-Cloud-Foundation>)

Zero Touch Provisioning Edge Deployments at Scale




Zero Touch Provisioning

Benefits

- **Automatically bootstrap ESX on bare metal servers using UEFI with no manual intervention.**
- **Deploy Hundreds of Sites in Parallel**
- **Eliminate Truck Rolls and Manual Setup**
- **Standardize Every Edge Location from Day One**

Accelerate and Scale Fleet Maintenance for Edge



Double Fleet Capacity:

Supports **5,000 ESX** hosts per instance, doubling the current 2,500 host limit.

4x Parallel Upgrade Scale:


Parallel upgrade capacity has increased 4x from 64 clusters to **256 clusters** ensuring fleet remains up to date and secure.

Benefits

- **Minimize Downtime:**
Reduces maintenance downtime through automated pre-checks and orchestration
- **Drive Fleet Wide Consistency:**
Enforces fleet-wide compliance by automatically aligning every site with versioning policies.
- **Centralized, Consistent Lifecycle Management:**
Centralizes lifecycle management through a single control plane, eliminating fragmented site-by-site updates

Easily Accommodate Edge Environments

2-Node vSAN clusters for remote locations



Consists of 2 nodes at remote site

- Store data resiliently


Primary data center runs virtual witness host appliance

- Determines availability during failure
- Responsible for 1 or more remote sites

2 node clusters managed by single vCenter server at primary data center

Lower Edge Server Costs at Scale

with NVMe Memory Tiering



Enhanced NVMe Memory Tiering

Customer Benefits

- Edge impact: 40% memory cost reduction per site × thousands of sites = significant fleet TCO savings
- Software mirroring improves availability at unmanned remote locations
- Proactive visibility into NVMe device health prevents failures without on-site IT response.

Automated Certificate Management at Edge Fleet Scale

Automatic Renewal of vCenter and ESX TLS Certificates



ESX TLS
Certificates



vCenter TLS
Certificates




VMCA Renewal
During Upgrade

Auto-renewal prevents outages caused due to expired certificates across distributed sites

Industrial vSwitch - EDP Dedicated Low Latency Support

Support for OT use-cases



- Advanced capability added within NSX and exposed via vSphere vDS to enable low latency application networking
- Ensures EDP switch latency of less than 500 microseconds in a very prescriptive environment
- Support for Parallel Redundancy Protocol (PRP) and Process Field Network (PROFINET) protocols
- PROFINET – widely used in industrial automation and manufacturing use-cases for supporting real-time applications such as controlling robotic arms
- Provides virtual networking foundation for the most demanding industrial ethernet deployments

Crucial Components for VKS



vSphere
Supervisor




Networking
VDS/ VPC



Load
Balancer


Supervisor is not VKS



- **vSphere Supervisor** is an integrated **Kubernetes-based Declarative API**
- Deploy **Kubernetes Clusters** with **vSphere Kubernetes Service**
- Deploy **Virtual Machines** with **VM Service**
- Add additional services to **extend the capabilities** of the platform and support your apps


VKS Storage Enablement

Map Physical Storage to Storage Policy



VCF Services

Unified Runtimes, Infrastructure Services, and Advanced Capabilities via a Declarative API



Continuous Delivery at the Edge

GitOps patterns



VMware vSphere
Kubernetes Service

Code Blame 31 lines (29 loc) · 805 Bytes

```
1 apiVersion: kustomize.config.k8s.io/v1beta1
2 kind: Kustomization
3
4 resources:
5   - ../base
6
7 patches:
8   - target:
9     group: cluster.x-k8s.io
10    version: v1beta1
11    kind: Cluster
12    patch: |-
13      - op: replace
14        path: /metadata/name
15        value: site2-vks
16      - op: replace
17        path: /spec/topology/variables/0/value
18        value: best-effort-4xlarge
19      - op: replace
20        path: /spec/topology/variables/1/value
21        value: strb-policy
22      - op: add
23        path: /spec/topology/workers/machine
24        value:
25          overrides:
26            - name: volumes
27              value:
28                - name: containerd
29                  mountPath: /var/lib/containerd
30                  capacity: 100Gi
31                  storageClass: strb-policy
```

argo

Applications Settings User Info Documentation

Resource filters

NAME	KINDS	SYNC STATUS	HEALTH STATUS
NAME	KINDS		
		☐ Synced 1	☐ Healthy 9
		☐ OutOfSync 0	☐ Degraded 0
			☐ Missing 0
			☐ Unknown 0

Applications / site2-vks

DETAILS DIFF SYNC SYNC STATUS HISTORY AND ROLLBACK DELETE REFRESH

APP HEALTH Healthy SYNC STATUS Synced to HEAD (aae9707)

Author: xiaoGaoBroadCom <157322203+xiaoGaoBroadCom@us...> Comment: Update kustomization.yaml

120%

vSphere Client

Search in all environments

site2-vks-tkg-cluster-nodepool-dmo4-v7dtb-vctzc-nnn49

Summary Monitor Configure Permissions Datastores Networks Snapshots Updates

Guest OS

Virtual Machine Details

Tags

Notes

Storage Policies

Broadcom,
Icom Inc. and/or its subsidiaries.

Choice of Node OS Support for Your Workloads

VKS Cluster Node OS Options

Shipped Out of the Box



Photon OS

Built-in Enterprise Support - No Additional License



Canonical Ubuntu

Built-in Enterprise Support - No Additional License

Custom Node OS Images



Microsoft Windows
Server 2022

Require additional license from vendor



Red Hat Enterprise
Linux

Require additional license from vendor

Features


Bring your own Windows or Red Hat subscription, & build Windows or RHEL based nodes using *imagebaker* shipped with VKS

Benefits

Deploy Windows or RHEL based workloads to VKS without supportability concerns

<https://knowledge.broadcom.com/external/article/408602/increase-the-ephemeral-storage-on-a-vks.html>

Expanded GPU Support for AI Applications at the Edge



**Maximize GPU
Utilization
Across Sites**




**Support
Diverse Edge
Workloads on
One Platform**

Benefits

- **Maximize GPU Utilization Across Sites:**
Share GPUs across multiple VMs/workloads, improving efficiency and reducing hardware costs at distributed locations.
- **Support Diverse Edge Workloads on One Platform:**
Run AI, VM based, and containerized applications together with GPU acceleration on a unified edge platform.
- **Reduce Infrastructure Footprint and Cost:**
Consolidate workloads onto fewer GPU enabled systems, optimizing power, space, and cost in constrained edge environments.

CPU-Based Inference with llama.cpp

Run AI Models efficiently without GPUs



Edge AI

llama.cpp to run quantized AI Models on CPUs for low-cost, local inference

CPU-based execution reduces cost, increases deployment flexibility, and preserves GPUs for large-scale training and inference.

AI Models can run efficiently on CPUs, making them ideal for development, testing, and production use cases where GPUs are not required.

Key Outcomes of VMware Cloud Foundation Edge

Scale, Speed, Efficiency, and Real-Time Intelligence at the Edge



**Scale without
Operational
Complexity
Across
Thousands of
Sites**



**Real Time
Intelligence
at the Edge**



**Lower TCO vs
Fragmented Edge
Solutions**



**Faster Deployment
and Time- to- Value
Across Distributed
Environments**

VCF Edge Use Cases



AI

Preventive Maintenance
CV Assisted Patient
Diagnosis



DC Extension

ROBO
Retail Stores



Real-Time Analytics

Virtual PLCs
Protection Relay in Power
Distribution



Industrial IoT

Automated
Inspection
Inventory
management



Offline Locations

Defense Apps
Oil Rigs

Audi - Edge Cloud 4 Production (EC4P)

Smart factories accelerate the transformation of the global car industry

ABOUT

Audi is one of the world's leading manufacturers of premium cars. The Audi Group includes the Bentley, Lamborghini and Ducati brands

CHALLENGES

- Eliminate downtime to streamline production, ensuring outage prevention and rapid recovery
- Manage complexity with a simple, functional IT landscape

PRODUCTS


VMware Cloud Foundation
VMware Professional Services

SOLUTION

- Consolidated Programmable Logic Controllers (PLCs) into virtual PLCs
- VCF private cloud across many factory use cases
- Automated updates and patches

IMPACT

- Reduced costs, strengthened security while decreasing auto build cycles
- Proactive monitoring and timely updates/patches minimize downtime
- Improved reliability and resilience



Ingolstadt, Germany
Automotive



“Our motto is software, not hardware. VMware Cloud Foundation underpins our EC4P project, virtualizing our production environment with a software defined infrastructure platform.”

Sven Müller, Project Lead, Edge Cloud
for Production, Audi

Resources

- [VCF Edge Design Library](#)
- [Design Blueprint: Single Node with ArgoCD](#)
- [Design Blueprint: Government and Defense Use Case](#)
- [Design Blueprint: Manufacturing Use Case](#)
- [Blog: Operationalizing the Edge](#)
- [Blog: ZTP with VCF Edge 9.1](#)
- [Video - VKS Deployment with VCF](#)
- [Video Demo 1: ArgoCD CI/CD](#)
- [Video Demo 2: Harness CI/CD](#)



THANK YOU!



Adria Forum 2026

