



We will begin shortly...

VCF Monetisation & TCO: Lunch and Learn



15th May 2025 | 12:00 PM- 2:00 PM AEST



In Person Venue - VMware by Broadcom Sydney Office, 8/175 Pitt St, Sydney NSW 2000
Online Attendance - Microsoft Teams Webinar



HYBRID EVENT

VCF Monetisation & TCO: Lunch and Learn



15th May 2025 | 12:00 PM- 2:00 PM AEST



In Person Venue - VMware by Broadcom Sydney Office, 8/175 Pitt St, Sydney NSW 2000

Online Attendance - Microsoft Teams Webinar

Meet The Speakers



Gary Marshall

Cloud Strategist (APJ) | VCSP Solutions
VCF Division, VMware By Broadcom



Phi Dao

Senior Technology Advisor
Crayon ANZ



Disclaimer

- This information is confidential. As part of your registration, you have either agreed to treat this information as confidential or agreed to a separate non-disclosure agreement.
- Certain information in this presentation may outline Broadcom's general product direction and may contain product features or functionality that are currently under development
- This presentation shall not serve to (i) affect the rights and/or obligations of Broadcom or its licensees under any existing or future license agreement or services agreement relating to any Broadcom software product; or (ii) amend any product documentation or specifications for any Broadcom software product.
- This presentation is based on current information and resource allocations and is subject to change or withdrawal by Broadcom at any time without notice.
- The development, release and timing of any features or functionality described in this presentation remain at Broadcom's sole discretion.
- Notwithstanding anything in this presentation to the contrary, upon the general availability of any future Broadcom product release referenced in this presentation, Broadcom may make such release available to new licensees in the form of a regularly scheduled major product release.
- Such release may be made available to licensees of the product who are active subscribers to Broadcom maintenance and support, on a when and if-available basis.
- The information in this presentation is not deemed to be incorporated into any contract.


Agenda

- Enablement Objective
- CSP Transition to VCF 9
- Infrastructure Optimizations
 - Host / Core
 - Networking
 - Storage
- License Portability (BYoS)
- Monetization

Outcomes

VCF Solutions Accelerate CSPs GTM

Enabling our CSPs to deliver differentiated solutions to market



CSP Solutions

- Differentiated
- Vertically Aligned
- Professional & Managed Services



Optimization and Monetization of VCF-based Services

Enablement Outcomes

- Objective
 - Enable you to deliver differentiated cloud services through the adoption of the full VCF stack and the enhancement or development of new VCF-powered offerings
- Outcome
 - Identify 2 to 3 candidate services from the VCF Solution Catalogue for each of **Optimization** and discovery and deeper investigation for **Monetization**



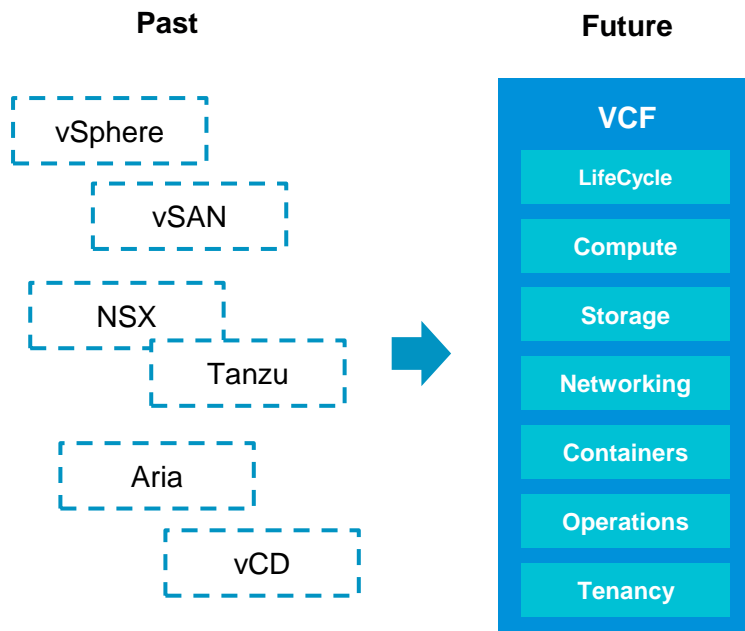
**Accelerating
CSP Growth**

CSP Transition to VCF 9.x


Unifying The VMware Technology Stack

Simplification through a unified technology stack & product offering

Unified Products...



...Simplified Consumption



CSP Transition to VCF 9

Moving to VCF will be a transitional journey

Phase 1

Product Enablement
& Evaluation



Gain a full understanding of VCF 9.x capabilities, features and design considerations

Phase 2

Transition Readiness
Planning



Assess your infrastructure against new VCF 9.x capabilities to develop your migration plans, design / operational changes and overall migration readiness

Phase 3

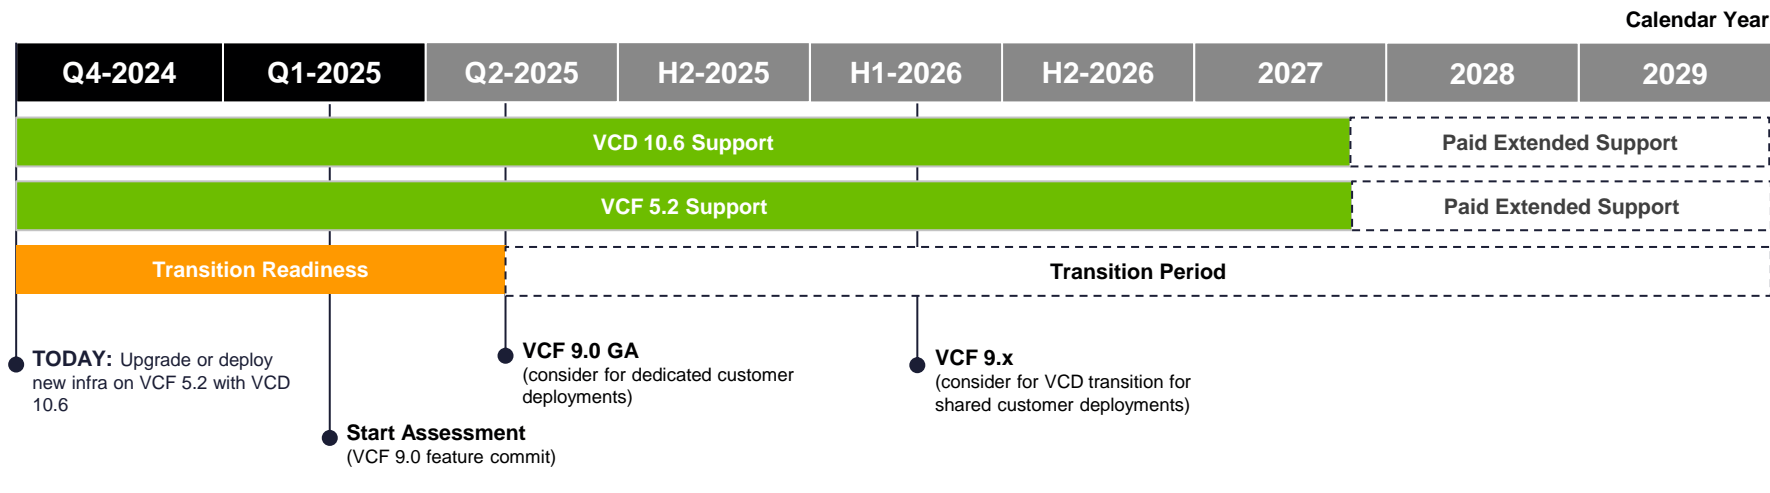
Transition Execution



Execute your transition to VCF 9.x using in-house resources or Broadcom Professional Services

CSP Transition to VCF 9

Current guidance for transitioning from VCF 5.x / VCD 10.6 to VCF 9.x *

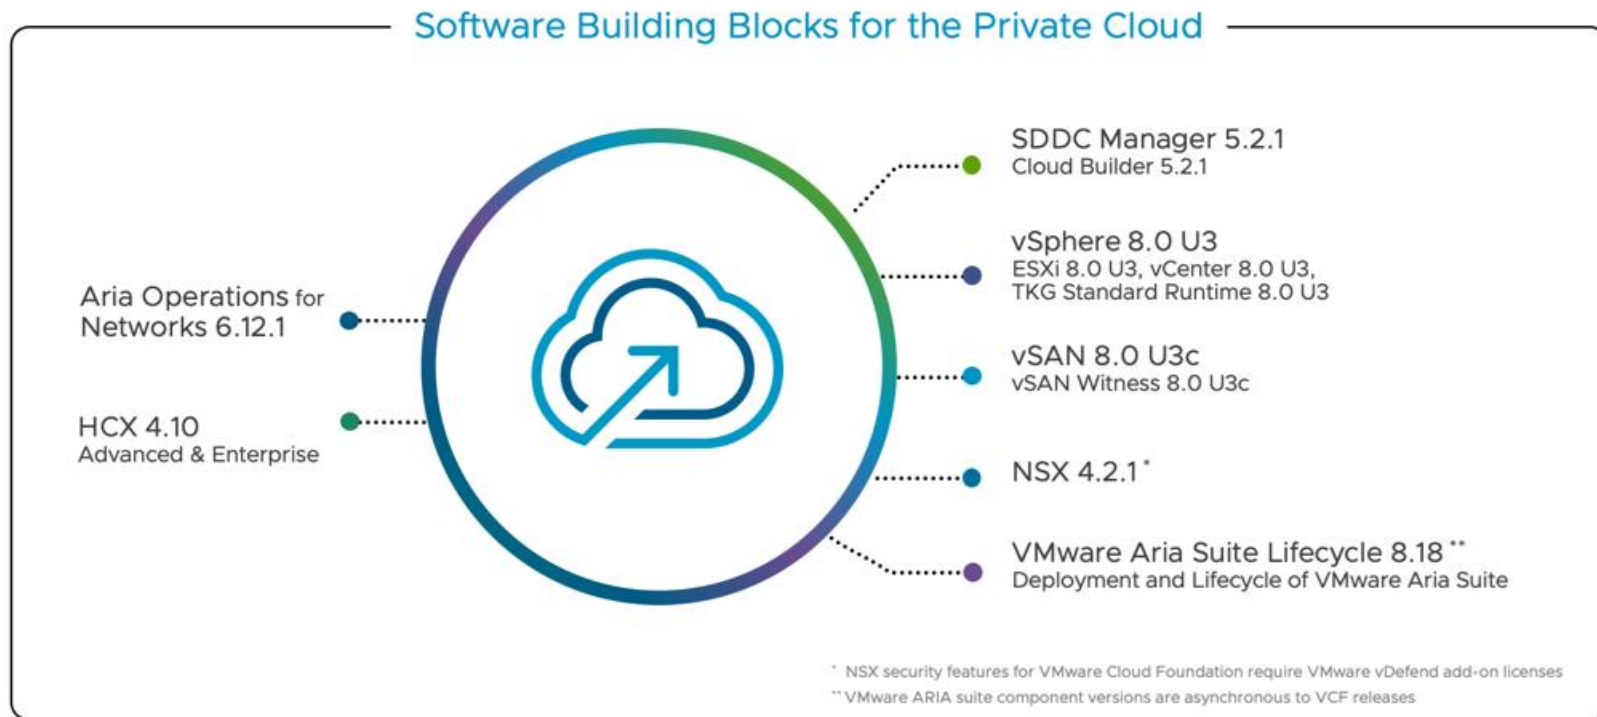


Key Considerations

- Initial 9.0 GA likely more suitable for **dedicated customer** deployments, with following releases providing more mature feature parity to transition from VCD for **shared** environments
- CSPs can start their assessments once VCF 9.0 GA release features are locked & committed
- Transition starting point for CSPs will vary based on initial VCF 9.x release features, combined with each CSPs unique requirements and use cases
- Extended support options are available to assist with extended transition duration

Step-1 : Upgrade to VCF 5.2.1 BOM

VMware Cloud Foundation 5.2.1 contains the following VMware products



Step - 2 : Plan for VCF 9.x

Preposed Plan

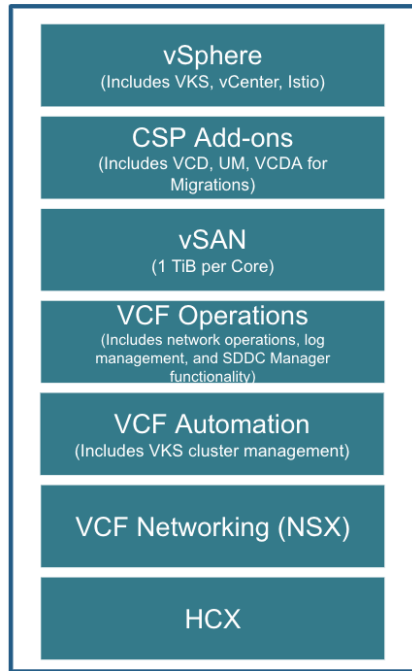
- Apply the VCD patch to enable compatibility with vSphere 9 and NSX 9
- Deploy the VCF Installer and import the existing vSphere 8.x and NSX 4.2.x environment
- Upgrade vSphere and NSX to version 9
- At this stage, the CSP will continue operating VCD 10.6.x on top of vSphere 9 and NSX 9
- Continue to evaluate VMware Cloud Foundation Automation (VCFA) and if good with feature then plan for switch over from vCD to VCFA

Infrastructure Optimizations

Significant Value to Unlock with VCF

Focused optimization guidance to accelerate CSP growth & profitability

VMware Cloud Foundation



The Challenge

- Unit of measurement was vRAM
 - Environment not optimized for the current license model (CPU Cores)
- Historically additional features and capabilities that are now bundled with VCF came at additional (significant) cost
 - Limited uptake of value-added features

Optimization Approach

- Develop and execute a plan to optimize the architecture for Core licensing
- Leverage VCF functionality to driving both internal *operational efficiency* and further *monetization opportunities* with end-customers
- Notable inclusions are Aria Suite Enterprise (VCF Operations, VCF Automation, Aria Operations for Networks Enterprise (VRNI), VCF Networking (NSX), HCX, vSAN and SDDC Manager

Why Optimise? Commit Contracts are Already in Place ...


Significant changes from vRAM license deployments and features to core-based VCF licensing




Optimization: Host & Core

Optimizing Your Licensing Usage & Consumption

Most common host infrastructure challenges for CSPs



Hosts w/ <16-pCores



Hosts w/ Excessive Cores




Unbalanced vCPU to pCore




Unbalanced Hosts



Shift from vRAM to Core




Excessive Spare Hosts



Idle DR Sites

VCF Host CPUs with <16 pCores




The Challenge

- VCF license applies a minimum of 16-pCores per CPU
- Older or legacy CPUs with 4-/8-/12-pCores will now be charged at the minimum rate of 16-pCores per CPU
- Artificially charging for pCores that are not physically present, driving up license costs and commit requirements

Optimization Approach

- Audit hosts to capture the distribution of pCores per CPU
- Retire hosts with <16-pCores and consolidate workloads onto hosts with higher pCore-counts to ensure the most efficient availability of CPU infrastructure and resulting allocation of licensing requirements
- Some workload requirements have specific pCore requirements for <16-pCore CPUs (for example: database hosts), and for these you need to calculate which is the better benefit – lower cores for DB licensing, for example, vs minimum VCF requirements

Low / Unbalanced vCPU to pCore Contention




The Challenge

- Artificially increased VCF licensing costs due to inefficient use of pCores
- Mix of different workloads (Test / Dev, Production, Performance) consolidated on a single cluster leading to inefficient and unbalanced utilization of pCores

Optimization Approach

- Audit hosts to capture the average cluster & host-level contention
- Evaluate the purpose of a cluster – Test / Dev Workloads may operate at an average contention of 8vCPU:1pCore, Production Workloads 4vCPU:1pCore and Performance Workloads at 2vCPU:1pCore
- Remove under-utilized hosts from the cluster and consolidate workloads onto fewer hosts to achieve desired ratio
- Evaluate the benefit of having dedicated clusters for specific workload requirements to ensure efficient use of pCores as well as better meeting of defined SLA requirements
- CPU readiness metric can be captured and presented to tenants via an Aria Ops Dashboard, improving demonstration of meeting a mature and advanced SLA metric

Shift from vRAM to Cores



The Challenge

- With the prior vRAM-based licensing model, CSPs have built offerings within their service catalogue using a 100% vRAM-based cost model and chargeback pricing

Optimization Approach

- Evaluate the existing service catalogue and re-design / re-architect / re-define as required, recognizing that consumed pCores is now the driving element in the software licensing component rather than consumed vRAM
- Changes in service catalogues may be required:
 - *With dedicated environments, charging on an entire physical host / cluster may be the best approach (including spares), although a VM-based service catalogue is certainly still feasible;*
 - *For shared environments, a VM-based service catalogue needs to be re-calculated with VM design focus shifted from vRAM to vCPU and the level of contention and sparing in the environment*

Excessive Spare VCF Hosts




The Challenge

- CSPs are now limited in maintaining adequate spare host capacity due to the change from vRAM to Core based licensing
- Additional VCF licensing requirements further restrict spare host capacity:
 1. Any host within a cluster that has an active VM (of any size) will be charged for the duration that the VM is active
 2. Any host that is deployed as an active vSAN node is chargeable (even if it does not have an active VM operating on it)

Optimization Approach

- Evaluate spare host capacity on a per cluster basis and reduce where practical (whilst maintaining SLA commitments)
 1. Example: effectively $n+3$ host spares are being maintained and only $n+1$ is required, remove 2 active hosts from the cluster to ensure that active VMs do not float across all $n+3$ active hosts in the cluster when that additional capacity is not required
- Activate real-time and transparent capacity planning with Aria Operations to optimize spare host capacity
- Re-size per host vSAN requirements to reduce host count to achieve optimal density and sparing as required

Idle DR Sites With VCF Hosts



The Challenge

- Under-utilized hosts at a DR site were previously charged only for the consumed vRAM
- Per-host DR appliance (or proxy VM) causes an entire host to be charged for the duration of the month that the VM is running
- Hosts in an active vSAN cluster are always charged (even without active workloads)

Optimization Approach

- Under the VCF license construct, hosts without an active (powered on) workload will not be charged
- Examine RTO/RPO requirements to ensure only the minimum active hosts necessary to achieve SLAs are activated and that the service is adequately charged to the tenant
- Where applicable, re-architect DR solutions to best map to efficient allocation of licensed resources (e.g. utilisation of VLR / VCDA)
- Maintain minimum number of active hosts in a vSAN cluster and spares in maintenance mode, ensuring rebuild time when spares are activated meets the defined SLA requirements
- Where application, during a DR event, activated spares attract pro-rated charging, and at overage rates (on-demand)

Optimization: Networking

Optimizing Your Licensing Usage & Consumption

Most common network infrastructure challenges for CSPs



NSX Edition / Feature Complexity




Firewall Licensing Scalability



Load-Balancer Licensing Scalability



NSX Edition / Feature Complexity



The Challenge

- Prior NSX product packaging required CSPs to consume from a broad portfolio of different NSX product editions with different features and capabilities
- This created product and pricing complexities for CSPs to deliver capabilities within their service offerings that scaled well and catered to per-tenant requirements

Optimization Approach

- CSPs can now utilize all switching, routing and basic firewall features and capabilities of VCF Networking (NSX) under a single product and pricing model. This includes but not limited to:
 - Switching / NAT
 - IPv4 Static & Dynamic Routing
 - IPv6 Static & Dynamic Routing
 - ECMP
 - VRF Tier-0 Gateway
 - Ethernet VPN
 - L2 & L3 VPN
 - Multi-vCenter Networking & Security
 - Federation
- CSPs can also consume Aria Operation for Network Monitoring and HCX Enterprise to drive additional value and capabilities for their customers and increase revenue and profitability

Firewall Licensing Scalability



The Challenge

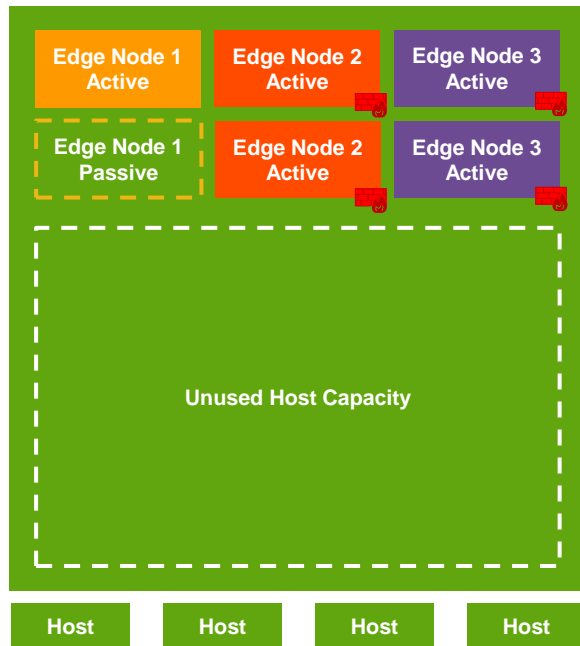
- Historically, NSX Advanced Firewall features (eg. DFW, GFW, ATP for DFW & GFW) could be licensed within a cluster at a far more granular level, providing CSPs better flexibility in managing licensing costs on a per-tenant basis, vs all-or-nothing
- The new ANS product packaging requires licensing across an entire cluster with all cores in the cluster licensed for these advanced features. This minimizes flexibility for CSPs who deploy shared-infrastructure where per-tenant granularity is critical, thus causing significant increase in licensing costs

Optimization Approach


- Evaluate which network security features are required on a per-tenant basis, and migrate tenant workloads to clusters with appropriately licensed features
- Determine what features or how much FW licensing functionality is required
- Deploy dedicated clusters for the specific level of functionality required to enhance flexibility in providing advanced capabilities on a per-tenant basis:
 1. Cluster dedicated to base NSX features only (included in standard VCF license)
 2. Separate cluster dedicated to workloads requiring vDefend features
 3. Separate cluster dedicated to workloads requiring vDefend ATP features (such as IDS/IPS)

vDefend Firewall Cost Optimization

Example #1: Dedicated Cluster for Edge Nodes



VCF Edge Cluster




VCF Workload Cluster

Total Licensing Cost: **\$104,960**

- 8 VCF Hosts w/ 32 cores each
 - $8 \times 32 = 256$ VCF core licenses
 - $256 \times \$350 = \$89,600$
- 6 Edge Node VMs w/ 8 vCPU each
 - 2 A/P Edge Nodes (no security configuration)
 - 4 A/A Edge Nodes (licensed with vDefend GFW)
 - $4\text{GW} \times 8\text{vCPU} \times 4\text{FW cores} = 128$ FW core licenses
 - $128 \times \$120 = \$15,360$

Load-Balancer Licensing Scalability



The Challenge

- CSP environments scale exponentially faster than enterprise customer environments due to the “service provider” business model where CSPs are operating, managing and delivering infrastructure and services to hundreds or thousands of customers (tenants).
- This exponential scale can create significant increases in licensing costs when leveraging a 1:1 deployment model for Avi Load-Balancer (e.g. 1 appliance per customer / tenant)


Optimization Approach

- CSPs can leverage a “shared” design model for deploying Service Engines to help reduce exponential increases in licensing costs
- Shared design model offers options for **Active / Standby** or **Active / Active** deployment designs, as well as a new **Bandwidth-based Charging Model** becoming available. This provides CSPs with flexibility to meet different customer (tenant) requirements while maintaining optimal license costs


Pricing Options

Per Year List Price (US\$)

Dedicated - \$11,390



Shared - \$1265




One Service Engine can have 9 tenants
9 vNICs per SE

New packaging option ...

Shared with Bandwidth Limits

200Mbps - \$442 (5 cents per hour)
25Mbps - \$252 (3 cents per hour)



Single Tenant can be 200 or 25Mbps

Avi Load Balancer Licensing Optimization Example

Sizing / Pricing for variable LB utilization - Dedicated vs Shared

Dedicated

# of Tenants	Dedicated Model (A/S SE-Group per tenant)	Total SU
630 - Low Perf (70%)	1260	1260 SU
180 - Medium Perf (20%)	360	720 SU
90 - High Perf (10%)	180	720 SU
Total (List)		\$15,376,500

Right-sizing (2,700 SU) vs. 1-size-fits-all (3,600 SU)

Performance	SU per SE
Low	1
Medium	2
High	4

Shared

# of Tenants	Shared Model 1 (A/S)	Service Units	Shared Model 2 N (1) + M (1)	Service Units	Shared Model 3 (Mixed 1/3 rd A/S, 2/3 rd N+M)	Service Units
900	L 630/9 * 2 = 140 M 180/9 * 2 = 40 H 90/9 * 2 = 20 Total SE = 200	L 140 (140*1) + M 80(40*2) + H 80(20*4) Total SU = 300	L 630/9 + 25% = 88 M 180/9 + 25% = 25 H 90/9 + 25% = 13 Total SE = 126	L 88 (88*1) + M 50 (25*2) + H 52 (13*4) Total SU = 190	300(A/S) 600 (N+M) L 210/9*2=48 L 420/9 + 25%= 59 M 60/9*2 =14 M 120/9 +25% = 19 H 30/9*2 = 8 H 60/9 + 25% = 9 L 107 (48+59) , M 33 (14+19), H 17 (8+9) Total SE = 157	L 107 (107*1)+ M 66 (33*2) + H 68 (17*4) Total SU = 241
Total (List)	\$1,708,500		\$1,082,050		\$1,372,495	

N+M Architecture vs. A/S HA

Network Infrastructure Optimizations

CSP guidance for optimal Application Connectivity, Security & Performance

VCF Networking

Increase utilization of new **included** features:

- Switching / NAT
- IPv4 Static & Dynamic Routing
- IPv6 Static & Dynamic Routing
- ECMP
- VRF Tier-0 Gateway
- Ethernet VPN
- L2 & L3 VPN
- Multi-vCenter Networking & Security
- Federation

Includes additional networking technologies:

- Aria Network Insight Enterprise
- HCX Enterprise

vDefend Firewall

Split Edge node clusters into firewall-enabled and non-firewall to optimize GFW charges

Right-size Edge node cluster that are firewall-enabled

Move bare-metal edge nodes to VM to avoid full VCF licensing per BM-host

Avi Load Balancer

Right-size Service Engines based on throughput requirements

Leverage N+M High-Availability mode that's available in ALB Enterprise and optimize use of spare Service Engines compared to Active-Standby HA

Utilize global or per-tenant auto-scaling with policy limits to deploy additional Service Engines only when required

Manage license allocation across the fleet by allowing controllers to pull licenses from Avipulse as needed

Optimization: Storage

vSAN is a Critical & High-value Component in VCF

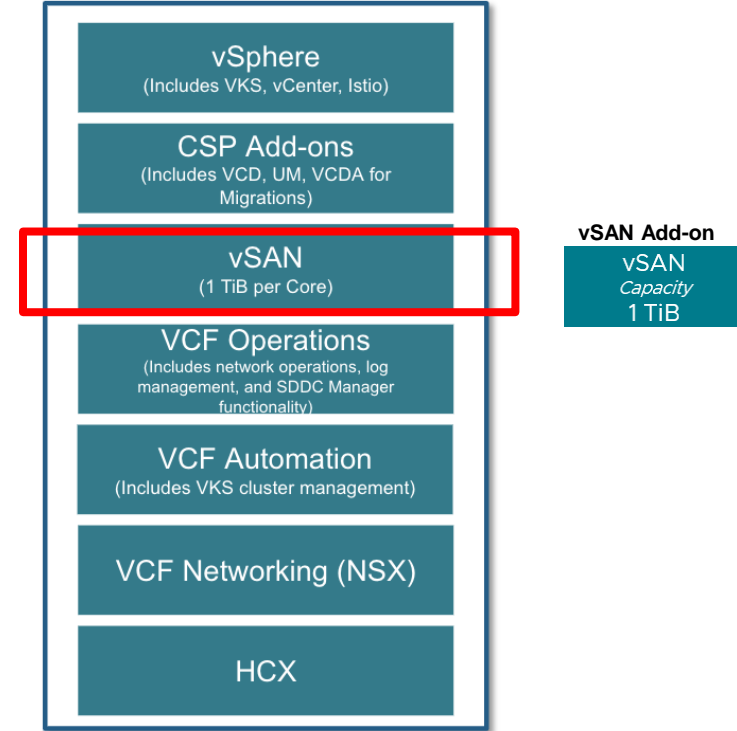
vSAN Enterprise is now included in VCF

Retired VCPP Program

Complete Product List	Points / Month	Charged Per
Storage and Availability		
VMware vSAN Standard	0.02	GB per month for Used Capacity
VMware vSAN Advanced	0.03	GB per month for Used Capacity
VMware vSAN Enterprise	0.04	GB per month for Used Capacity



VMware Cloud Foundation




vSAN vs Traditional Arrays

Remove hardware dependencies, decrease TCO and increase agility with vSAN



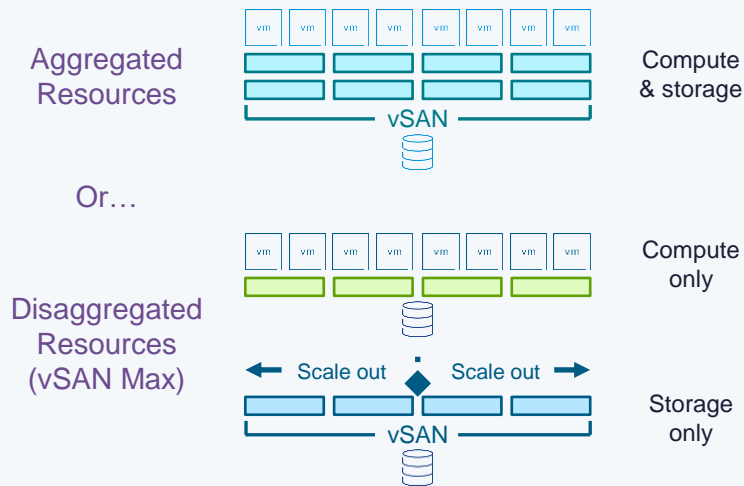
Legacy Three-Tier

- Proprietary storage hardware
- Limited scalability
- Siloed management
- Limited to Datacenter



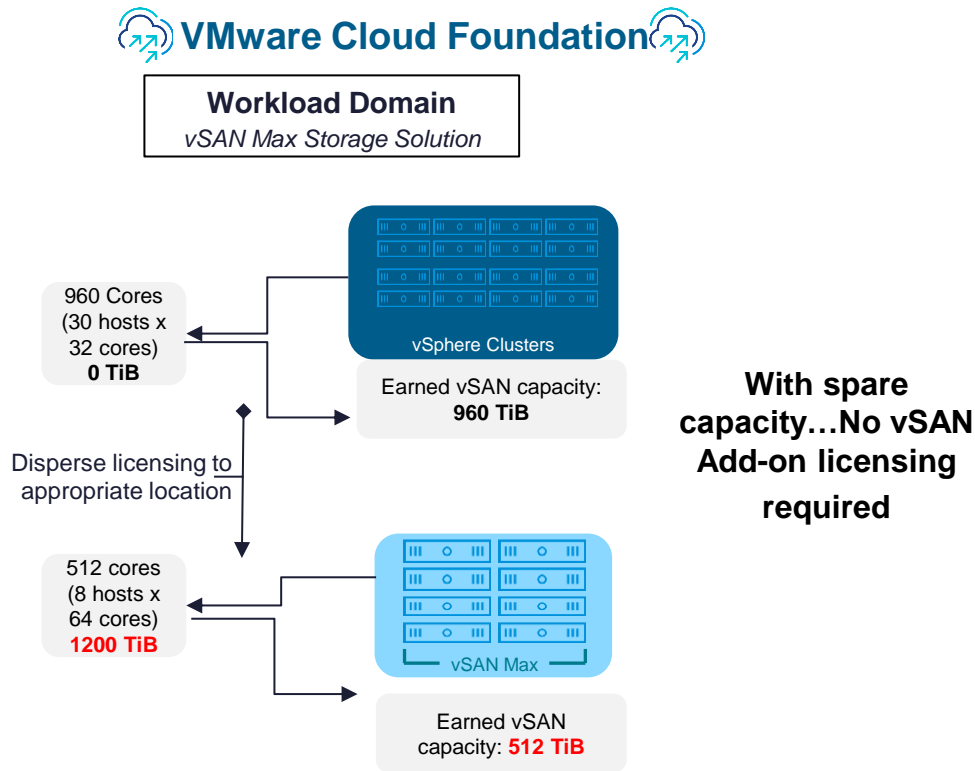
Hyperconverged

- Commodity server hardware
- Distributed scale-out storage
- Common management
- Datacenter or Edge



VCF Storage Cluster (vSAN Max) licensing scenario

Example VCF Compute and vSAN Max cluster deployment



Customer requirements

Workload clusters

- 10 x VCF Compute clusters - 3-nodes each
- 1 x vSAN Max Clusters - 8-node cluster
- **1200 TiB** of raw storage
- However earned = **512 TiB**

VCF License allocation

Total Storage Required: **1200 TiB**

- Capacity included with purchased VCF Cores: $960 + 512 = \sim 1500$ TiB

vSAN capacity earned from licensing VCF Compute clusters aggregated with vSAN Max cluster.

Optimization: Take-aways

Key Take-Aways & Next Steps


- Reviewing your infrastructure for potential optimizations is an **important first step** in preparing your transition to VCF
- Take time to perform a thorough review across all key infrastructure areas including **host compute, networking, storage** and more...
- **Collaborate with your Business, Solution & Products teams** to understand where design or architecture changes may help drive cost optimizations to lower overall operating costs

Key Take-Aways & Next Steps



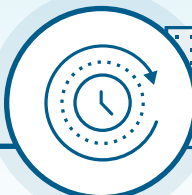
Cost Planning & Optimisation

- VCF component optimization
- Maximize VCF consumption




Business & Offering Strategy

- Target customers
- Solution use cases for public and dedicated cloud



Managed Services Monetization

- GTM development
- MVP development
- Maximize cloud value
- Third-party take out
- Higher margin from higher value & differentiated services



BYoS Opportunity

- Target customers or industry verticals
- Optimized service offering
- Targeted use cases
- Accelerate adoption
- High value deployment & management skills



VCF Add-ons

- Consumption of full-stack VCF
- Productization and market penetration
- Revenue acceleration

License Portability

What is License Portability

How do end customers benefit from License Portability

01

Entitlement

License Portability is an entitlement offered to customers who have purchased the new VCF subscription offering (qualifying version purchased after December 13, 2023)

Investment Protection

02

Licensee

Customers continue to be the licensee of VCF and will have full control of their licenses and compliance requirements outlined in the VMware Cloud Foundation Specific Program Documentation (SPD).

Simplicity

03


Deploy

Deploy and utilize their VCF licenses across a range of compatible environments whether that's in their own data center, a hosting provider, a cloud service provider or a public cloud hyperscaler.

Flexibility

Why License Portability For VCSP

License Portability benefits to CSPs



Removes
Software
Investment
Barriers

End customer can
continue to utilize the
existing VCF purchase

Faster
Onboarding
of New
Workloads

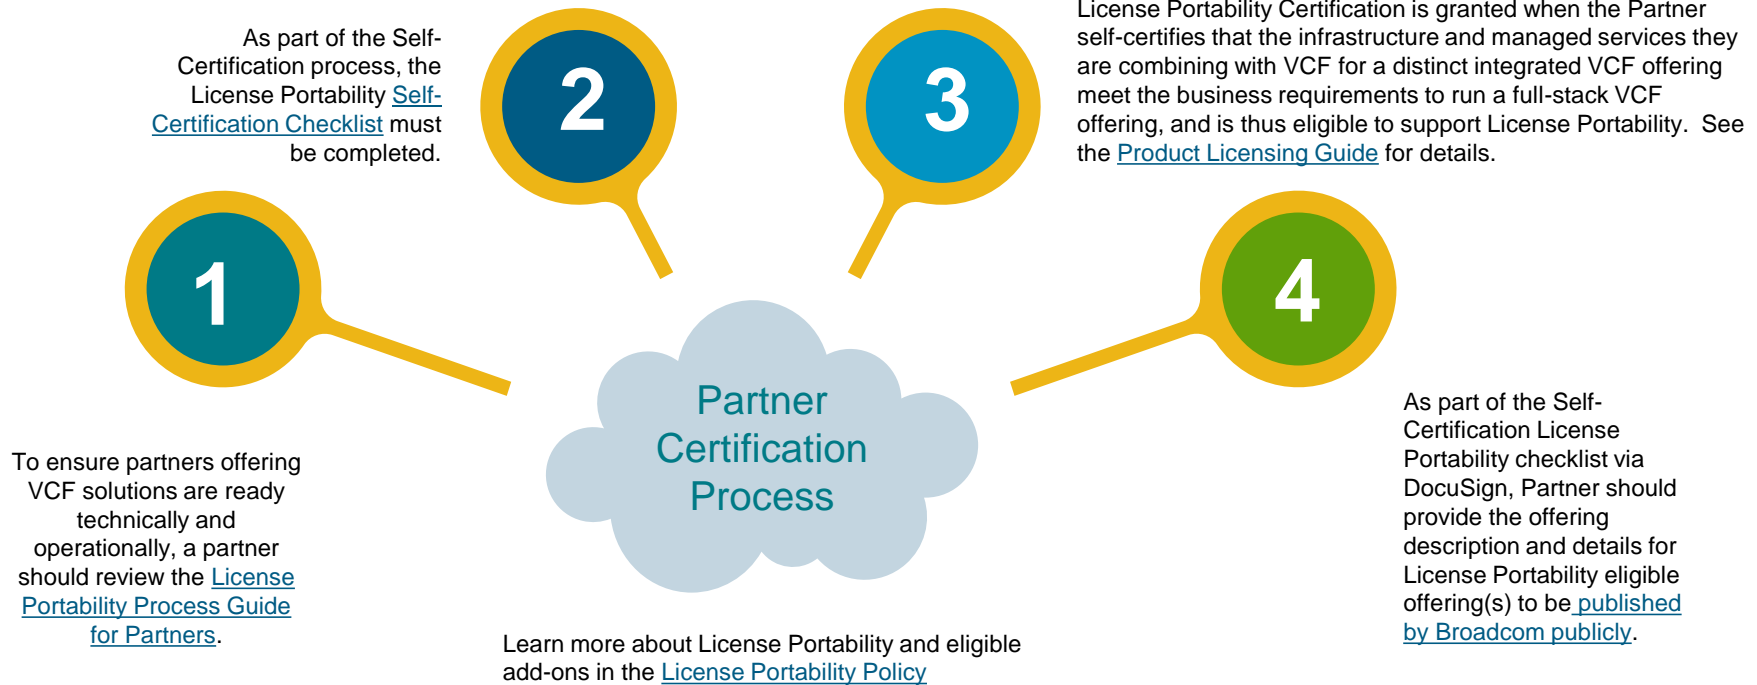
End customer no longer
needs to wait for
subscriptions to end
before migrating to CSP

Sell
Additional
Managed
Services

CSPs can layer additional
managed services onto
customer's VCF cloud

License Portability Process for CSPs

How do CSPs get Certified for License Portability



CSP Monetization with **VCF Managed Services**

Agenda

- Monetization Opportunity With VCF
- VCF Platform
- Storage Services
- Networking, Load Balancing & Security Services
- Kubernetes Services
- Cloud & Network Monitoring
- Sovereign Cloud
- Partner Enablement Services

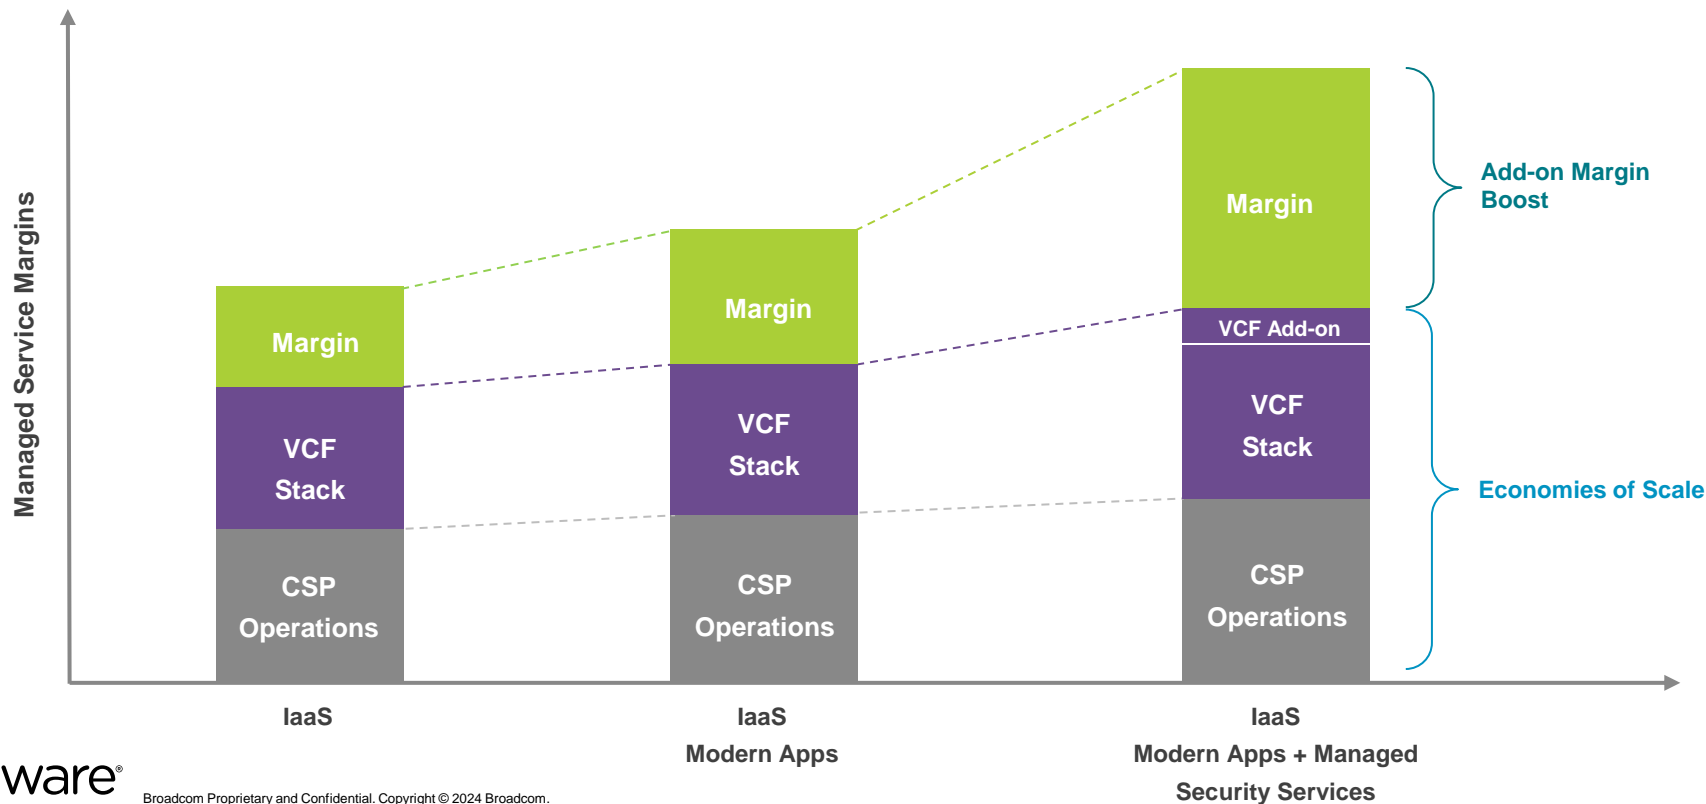
Monetizable Services with VCF

Grow beyond IaaS to generate revenue with additional Managed Services



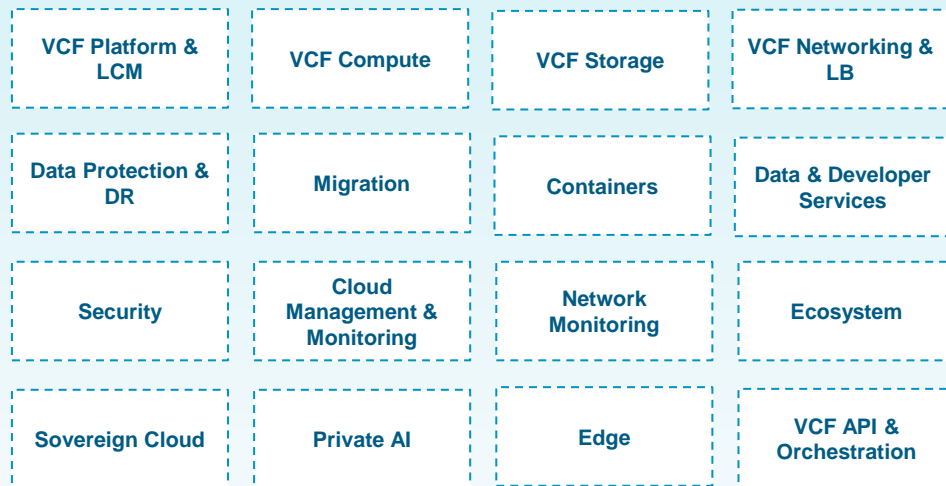
VCF Economies of Scale Enable CSPs to Achieve Higher Margins

Integrated interoperable components lower CSP operating costs



VCF Solutions Accelerate CSPs GTM

Enabling our CSPs to deliver differentiated solutions to market



CSP Solutions

- Differentiated
- Vertically Aligned
- Professional & Managed Services



Your Journey to Monetization with VCF

Solution Kits help accelerate the Productization process



Market Trends
& Demand



Technology,
Service &
Monetization
Models



Architecture,
Deployment &
Onboarding
Framework



Support, SLA &
Lifecycle Operational
Frameworks



Marketing, GTM &
Sales Enablement

Time-to-Market (Productization)

Monetizable Services with VCF

Grow beyond IaaS to generate revenue with additional Managed Services



Solution Basics: GPU-aaS

Available Today with VMware Cloud Director



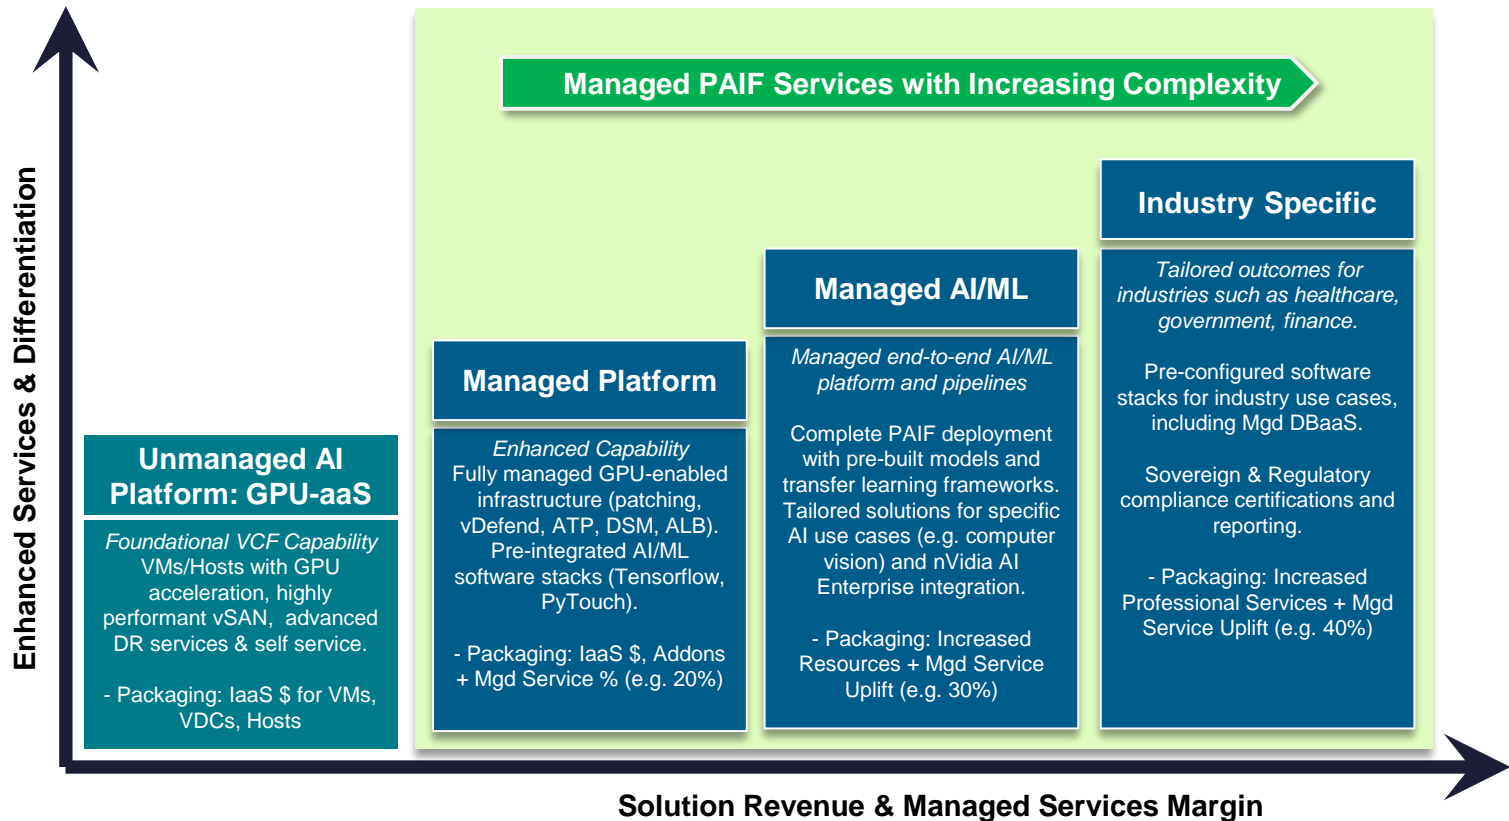
**Coming in
VCF 9.x**



Server

CSPs Delivering a Comprehensive AI Capability

Maturing CSP Solutions to Drive Revenue and Deliver Differentiated Outcomes

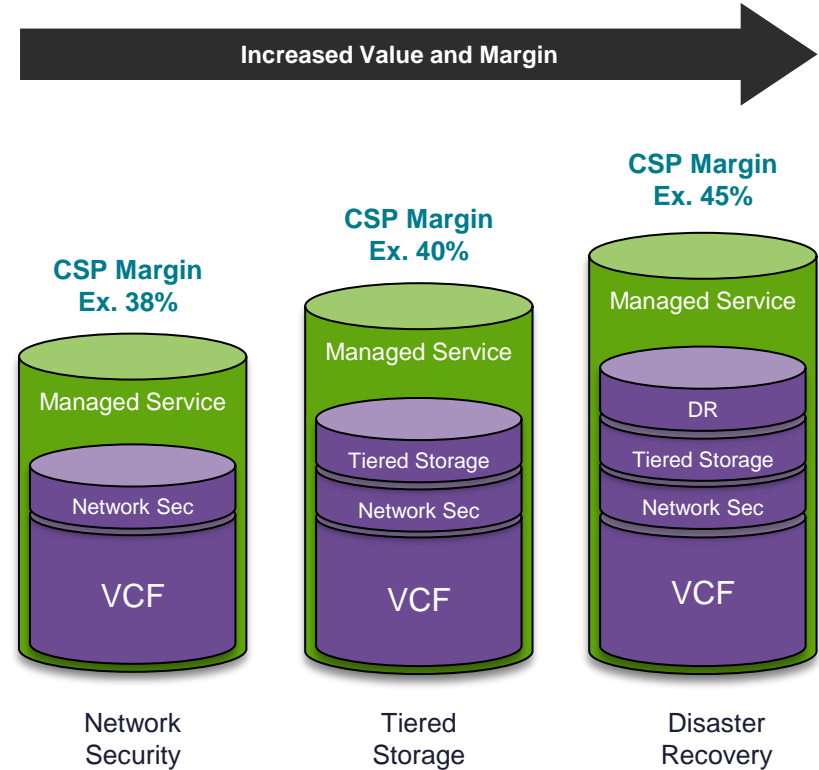



Service Based Strategy / Service Catalog

Creating recurring revenue streams

For this IaaS monetization strategy, the partner should focus on providing modular managed services packages that are aligned with the needs of the customers.

These are examples of some of the packages that the partner can offer.




Up- and Cross-Sell Opportunities

Generate additional Revenue Streams

Solution-centric service catalogue rather vs Infrastructure-led


CSP Margin
Ex. 35%



VCF Base
+ optional package




CSP Margin
Ex. 40%



Additional Features

No VCF add ons required. Hardware updates might be required to support some functionality

CSP Margin
Ex. 40%



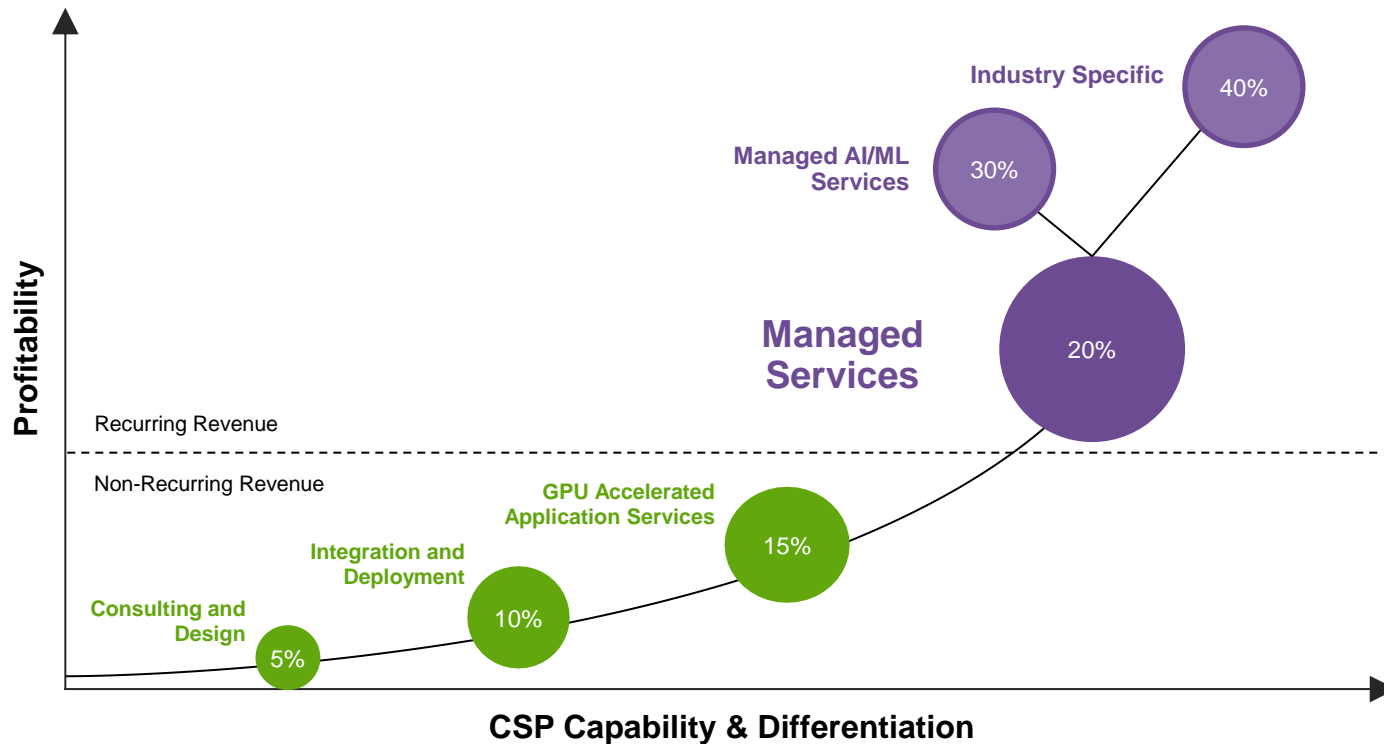
Advanced

These solutions require a mix of add ons and/or third-party licenses

Increased Value and Margin

GPUaaS & Services Sample Profitability Growth

Example Service Catalog



Non-Recurring Revenue


Consulting and Design: Assess customer requirements for GPU-accelerated workloads

Integration and Deployment Services: Integrate GPU infrastructure with VCF software stack

GPU-Accelerated Application Services: Optimize applications to leverage GPU acceleration on VCF

Managed Services: End-to-end lifecycle management of GPU infrastructure on VCF

Key



Service Offering

How to Monetize & Deliver Industry GPUaaS Specific Packages

Example offerings tailored to a variety of unique workload requirements



Healthcare

Provides GPU acceleration for healthcare workloads like medical imaging, genomics, and drug discovery, enabling faster insights and breakthroughs while bolstering data privacy and compliance.



Finance

Can help accelerate financial workloads like risk analysis, fraud detection, and algorithmic trading by leveraging GPU power, while enabling secure multi-party computation and data sharing.



Manufacturing

Can help accelerate manufacturing workloads like predictive maintenance, quality control, and digital twins, while providing confidential computing for protecting proprietary algorithms and data.

Subscription Based Licensing Model


Professional Services

AI-Ready Enterprise-Class Platform



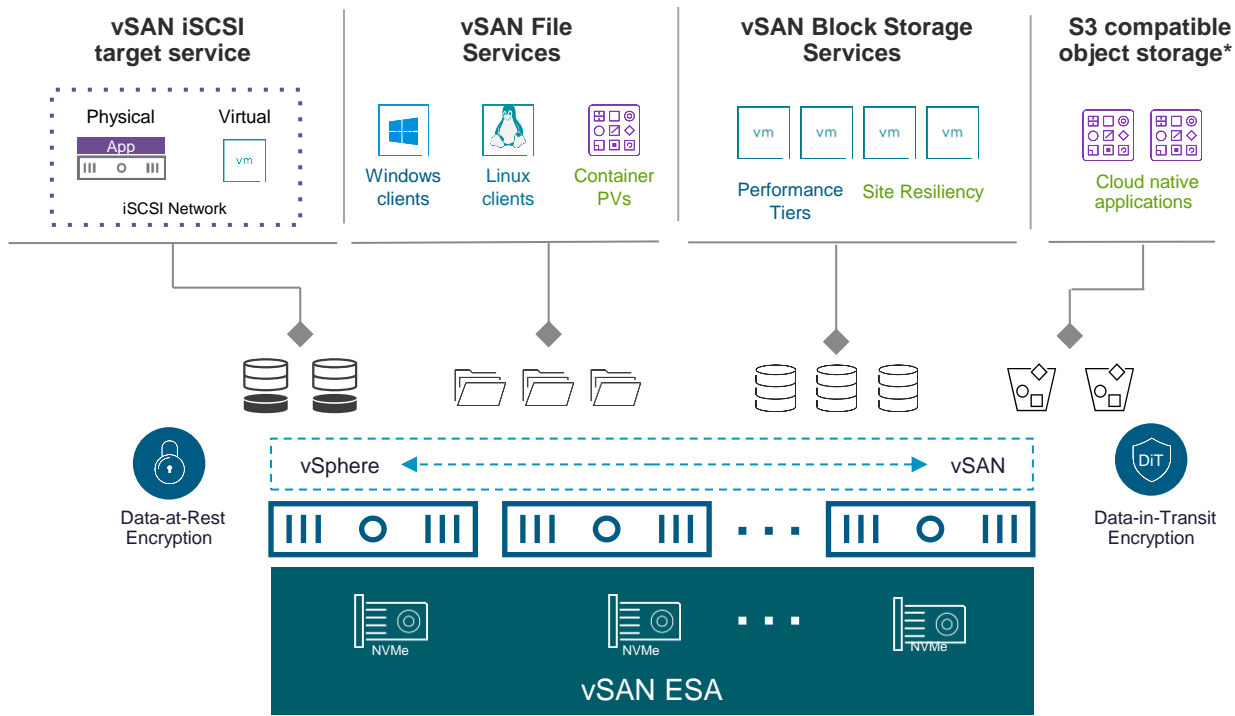
Monetizable Services with VCF

Grow beyond IaaS to generate revenue with additional Managed Services



The Foundation for Storage-as-a-Service with VCF

High-performance Tiered Service Offerings for CSPs

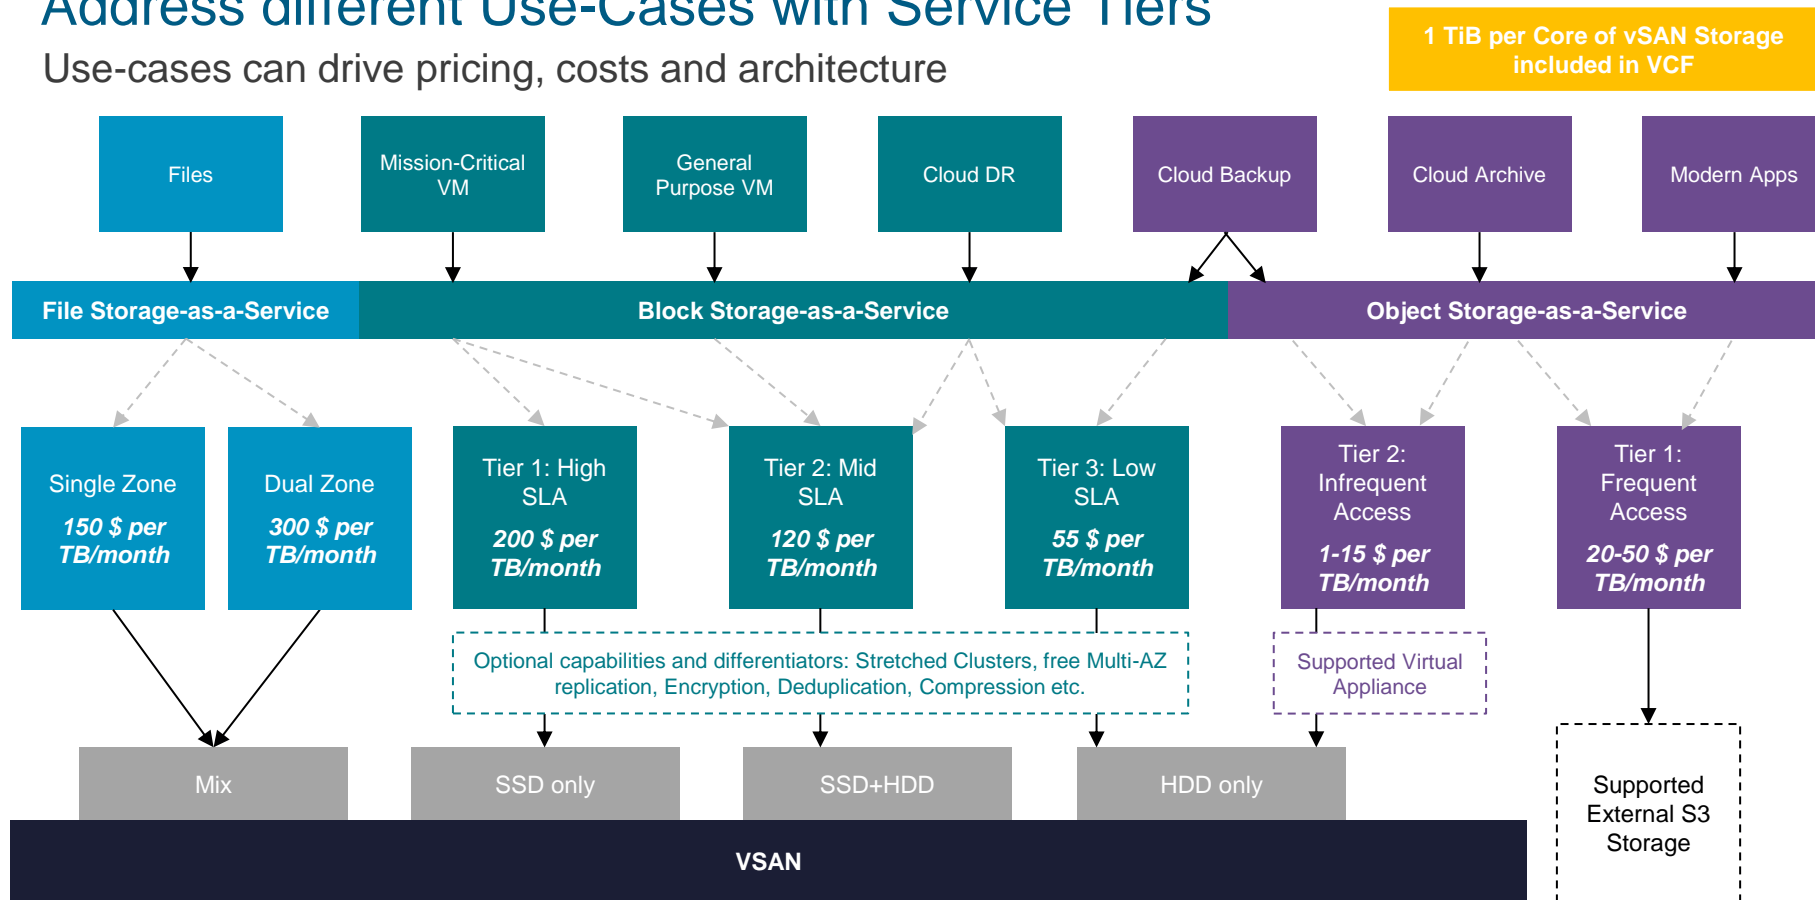


Differentiate Storage Services:

- Block storage at multiple performance tiers, resiliency levels and price points
- SMB and NFS file services
- iSCSI block access for legacy workloads
- S3 compatible object storage*
- Foundational vSAN capabilities meeting different workload requirements and at different price points

Address different Use-Cases with Service Tiers

Use-cases can drive pricing, costs and architecture



LEGO™ Blocks - vSAN Storage Components

Building Blocks: Compute + Storage Clusters

Compute cloud cluster
utilising storage from
dedicated vSAN MAX
cluster

vSAN capacity of 1TiB per
core included in base VCF
license

Licensed vSAN
consumption shared across
clusters

Failure protection of RAID6
assumed

Compression
conservatively estimated at
1.6

Compute
Cluster



+

Storage Cluster
(vSAN Max)



# Drives	6 hosts x 24 drives
Drive Capacity	3.84 TB (3.49 TiB)
Total Raw Capacity	553 TB (503 TiB)
Physical Cores	288 + 1,024 = 1,312
Useable Workload Capacity*	302.5 TiB*
vSAN Licensing	Fully licensed with VCF compute cluster
vSAN Headroom	~500TB spare
Average Cost (2.0x)	6.4c per GiB / month
Average Cost (1.6x)	7.9c per GiB / month

* **Useable Workload Capacity:** Overhead for vSAN operations,
host rebuild, meta-data, vSAN and RAID removed; after
compression.

Storage Benchmarking

Consistency and Performance

Differentiate with simplicity and consistency

Detailed reporting & SLA/SLO capabilities

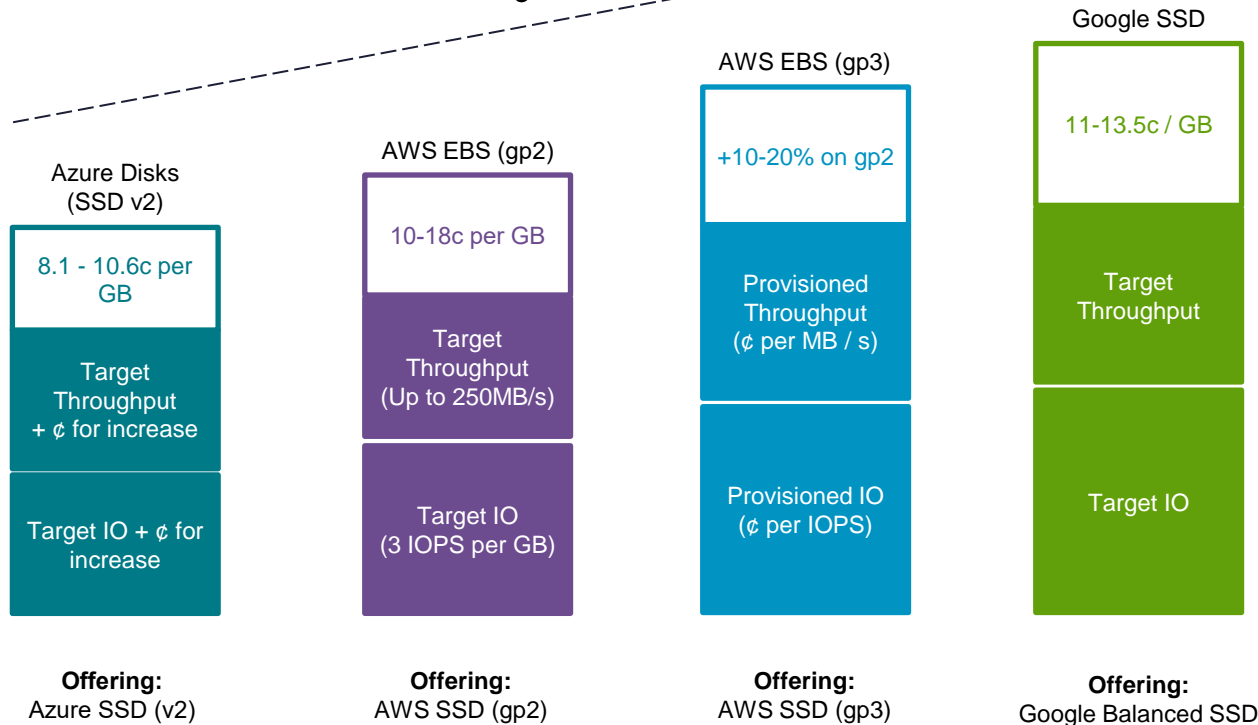
High performance

Hyperscalers can result in complexity and uncertainty

Cost competitive

Google SSD storage highest cost on average with less variation compared to Azure and AWS

Complexity and Cost



Storage Benchmarking w/ TCO Analysis


Consistency and Performance

vSAN TCO displayed with a combined Compute + vSAN Max Clusters & standalone vSAN Max cluster

vSAN compression shown at both “conservative” (1.6x) as well as “aggressive” (2.0x), where the higher value is expected for general cloud workloads

Significant variation on hyperscaler costs according to geo

Significant margin opportunity



Storage Benchmarking w/ Standalone Storage Platform

Consistency and Performance


Hardware + Software
Licensing Costs Only

Operational efficiencies,
network efficiencies, ... not
applied (not a TCO
analysis)

vSAN Max shown to be
cost competitive against
dedicated storage arrays

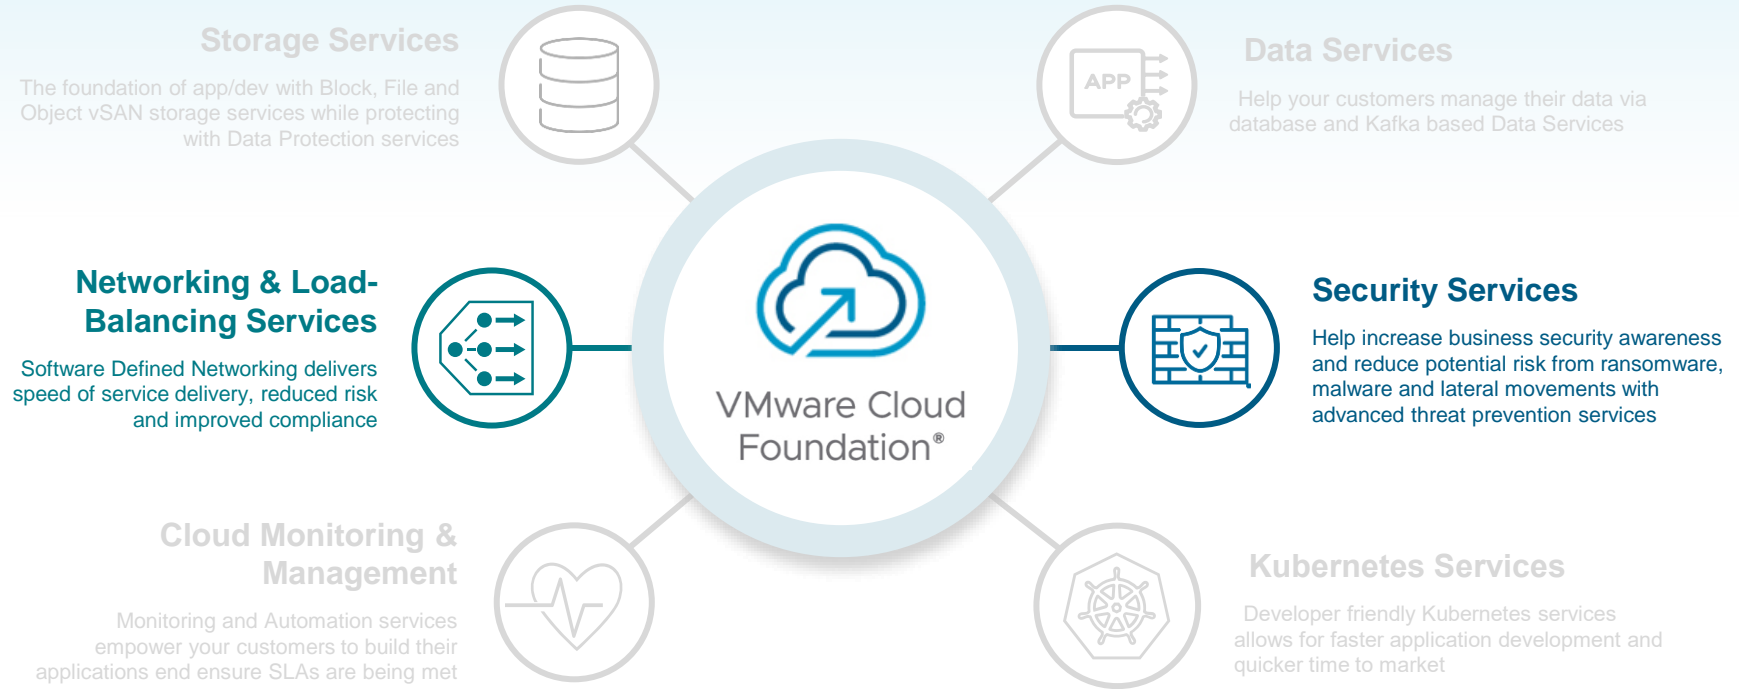
1PiB sizing assumed

Cost estimate summary of useable GiB per month



Monetizable Services with VCF

Grow beyond IaaS to generate revenue with additional Managed Services



Monetizable Services **Add-Ons** with vDefend Firewall and Avi LB

Networking, Security and Load Balancing Services



VMware Cloud
Foundation®

Higher Value Services

Load Balancing

Firewall, IDS/IPS

Routing

Switching

VMs, Containers, Physical Servers

VMware Avi Load Balancer Add-On

- L4 Load Balancing
- L7 Load Balancing
- Load Balancing Policies
- IPAM, DNS
- Global Server Load Balancing
- SSL / TLS
- Container Ingress
- Container LB
- App Health Analytics
- App Latency Analytics
- DDoS Protection
- Web Application Firewall (WAF)
- App Rate Limiting
- IP Reputation

VMware vDefend Advanced Threat Prevention (ATP) Add-On

- Distributed & Gateway IDS/IPS
- Network Traffic Analysis (NTA)
- Malware Prevention incl Sandboxing
- Network Det. & Response
- Gen AI Threat Investigation Co-pilot
- Curated Threat Intelligence feeds

VMware vDefend Firewall Add-On

- Distributed Firewall (L4-L7, Identity, FQDN, Malicious IP Filtering)
- Container Security
- Gateway Firewall (L4-L7, TLS Inspection, FQDN, URL Filtering)
- Security Intelligence (Visualization, Assessment, FW Rule recommendations)

NSX Networking for VCF

- Distributed Routing & Switching
- Overlay Networking
- Static & Dynamic Routing
- IPv4 & IPv6
- Virtual Routing & Forwarding
- EVPN
- Spoofguard
- L2/L3 Multicast
- NAT
- DNS & DHCP Relay
- L2/L3 VPN
- Quality of Service
- Container Networking
- Multi-Site Management
- IPFIX & Port Mirroring
- Object Tagging & Grouping
- Firewall Logging
- Stateless L3 Rules

Monetization Strategy for NSX

Charge based on VCF Add-Ons

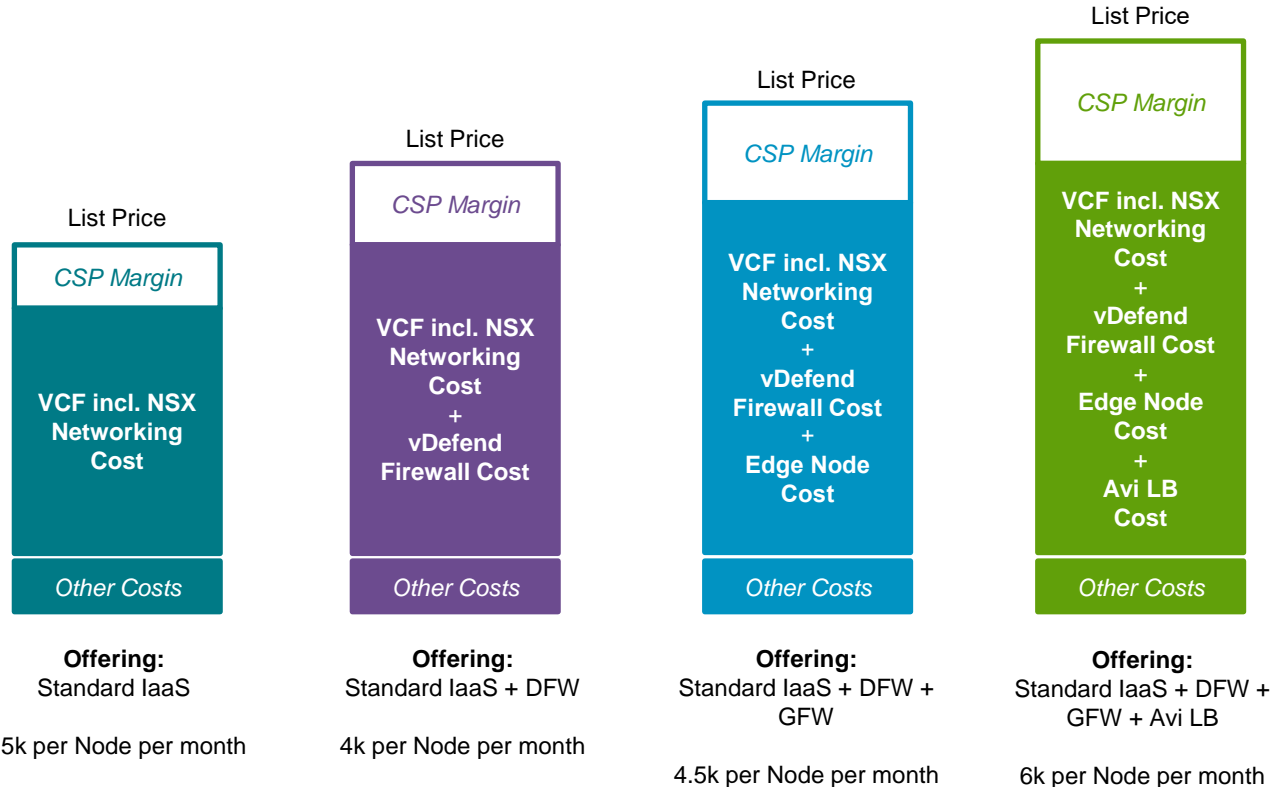
Base charge for VCF NSX networking features

Add-On charges for vDefend Firewall

Add-On charges for Avi LB

Included in IaaS charge or as separately priced item

Less comparable to Hyperscale offerings



CSP Price Comparison and Margin Impact


CSP Solution using VCF + Avi Load Balancer vs. Hyperscaler Solutions

End customer price higher by a factor of up to **13.5x**

Conservative assumptions with only one LB instance, no egress traffic, no additional Load Capacity parameters


Room to optimize Avi LB costs with N+M, less disk performance, customer commit etc.

VMware Avi Load Balancer Container Ingress Cost Comparison - By Workload





Monetizable Services with VCF

Grow beyond IaaS to generate revenue with additional Managed Services



Cloud Management Offerings

Build a complete portfolio of services with VCF Operations



Combine and integrate into a complete **Cloud Management Service Portfolio**

Monitoring as a Service

- Metrics, alarms, dashboards, reports
- Capacity planning, optimization and compliance
- System, App, Platform and Network
- Based on VMware Aria Operations
- Compete with AWS CloudWatch, Azure Monitor, GCP Cloud Monitoring

Logging as a Service

- Centralized Logging, Log Analysis, Dashboards
- Based on Aria Operations for Logs
- Compete with AWS CloudWatch Logs, Azure Monitor Logs, GCP Cloud Logging

Cloud Automation

- Cloud Automation and Orchestration
- IaC and GitOps Capabilities
- Service Catalog and Governance
- Based on Aria Automation
- Compete with AWS OpsWorks, CloudFormation, Azure Automation, GCP Deployment Manager




















Full-Stack Observability

- Application, platform and infra observability
- Traces, events and dependencies
- Based on Aria Operations for Integrations
- Enhanced with Aria Ops for Apps*
- Compete with AWS CloudWatch, X-Ray, Azure Monitor, GCP Cloud Trace

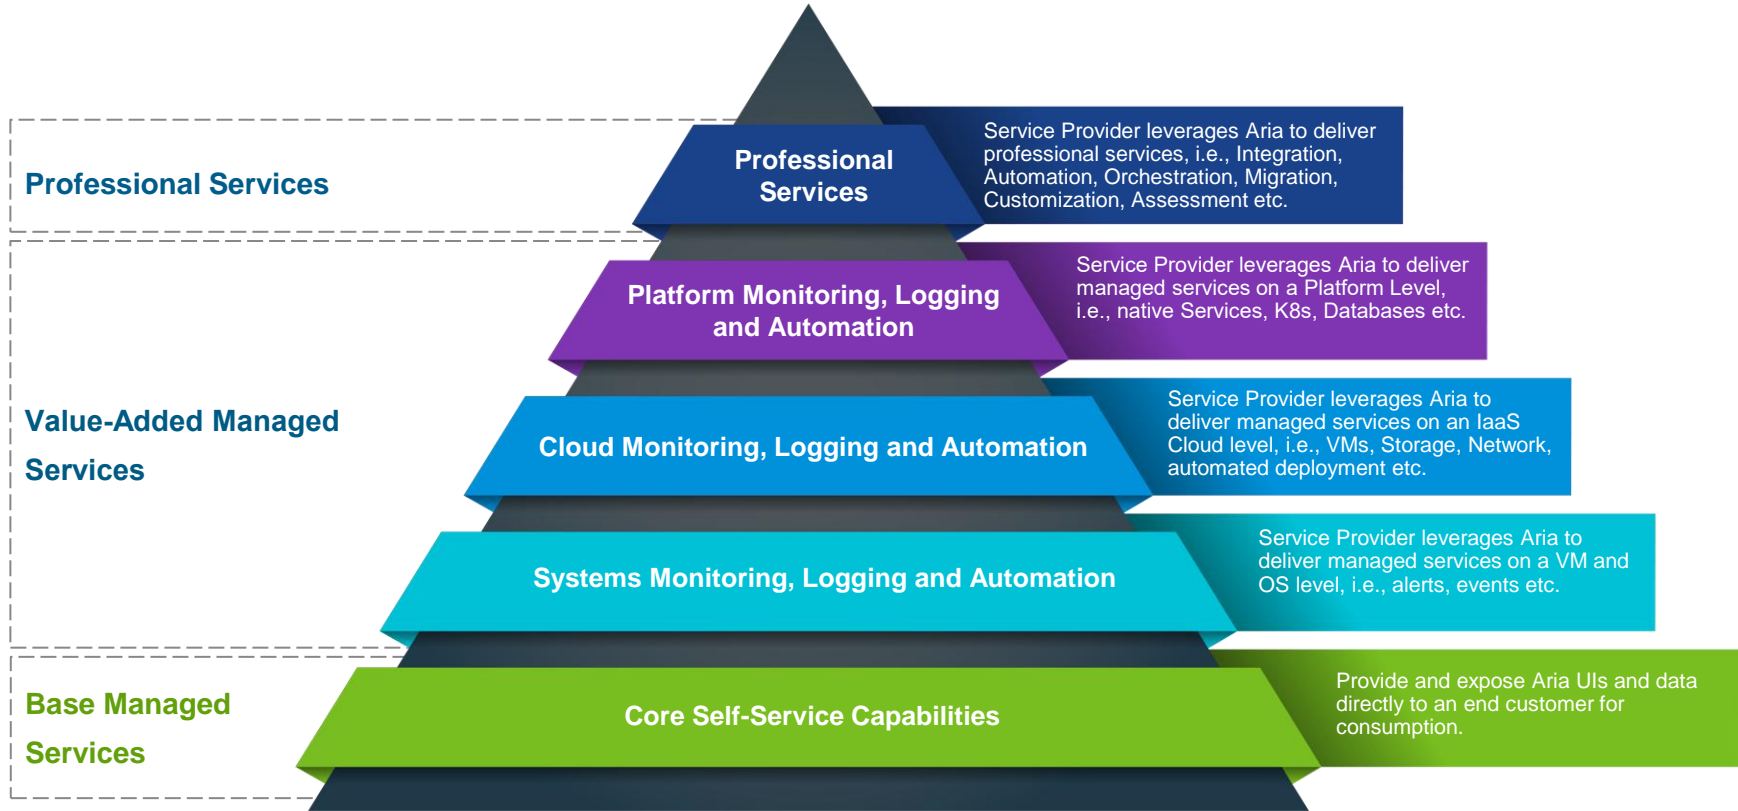
For VMware Cloud Foundation, Public Clouds & Third-Party Components

Easier Offering Portfolio compared to Hyperscalers

Example: Cloud Monitoring and Logging Services

Amazon Web Services	Microsoft Azure	Google Cloud Platform	VMware Service Provider
 AWS CloudWatch Metrics per metric, per API requested metric, per metric stream	 Azure Monitor Metrics per metric query, per metric ingestion	 Google Cloud Monitoring per metric, per API requested metric, per monitoring uptime checks	 VMware Cloud Foundation® Aria Operations Aria Operations for Logs Aria Operations for Networks Aria Operations for Integrations Charging Models: <ul style="list-style-type: none">• per Aria Instance• per Object / VM• granular
 AWS CloudWatch Logs per GB data ingestion, per GB data storage, per GB data scanned, per GB vended logs	 Azure Monitor Logs per GB basic logs, per GB analytics logs in tiers, per log search query, per GB archive, per export, per GB retention etc.	 Google Cloud Logging per GB data ingestion, per GB data storage	
 AWS CloudWatch Dashboards per dashboard	 Azure Dashboards Free	 Google Cloud Trace per trace ingestions	
 AWS CloudWatch Alarms per alarm metric, per metric analyzed, per composite alarm	 Azure Monitor Alert Rules per alert	 Error Reporting Cloud Logging charge	
 AWS CloudWatch Events per million custom event, per million cross-account event	 Azure Monitor Notifications per update event, per E-mails, per push notification, per web hook, per SMS, per voice call	 Cloud Profiler Free	
 AWS CloudWatch Additional charging metrics for insights, network resources, canaries, etc.	 Azure Monitor Additional charging metrics for web tests, analytics, Prometheus etc.	 Google Cloud Operations Suite Varying charges for GKE, Prometheus, Anthos etc.	

Strategies for Monetizing VCF Operations



VCF Operations Pricing Models

VCF Ops Instance (\$4k Cost): \$10k Revenue

50k Objects at \$2: \$100k Revenue

\$10M. Metrics at \$0.3: \$3M. Revenue

per Aria Instance



- Charge per Aria Instance
- Offer different instance sizes based on customer requirements
- Fixed charge with unlimited Objects, Metrics, Alarms, Dashboards, Reports etc.
- Chargeback via Aria Instance Tag
- High revenue and margin opportunity
- Differentiate from granular and unpredictable charges in AWS CloudWatch, Azure Monitor, GCP Operations Suite etc.

per Object / VM



- Charge per managed Object in Aria
- Object can be:
 - vCenter, Cluster, Host, VM
 - Data Store, vSwitch, Port Group
 - Folder, Resource Pool, Namespace
 - Public Cloud resource
 - OS, App and platform instance
 - Infrastructure object
 - Application adapters
- Chargeback via Object Tag where possible or custom query
- Good balance between granularity, complexity and margin

Granular



- Charge on a granular consumption basis
- For Example:
 - per Metric
 - per Alarm
 - per Dashboard
 - per Report
 - per GB Logs
 - per trace
- Chargeback via custom query
- Aligned with hyperscale PnP
- Low initial margin, huge revenue/margin potential at scale

TCO Comparison Example - VCF Operations

Significant Revenue and Margin Opportunity

Per Instance model assumes IaaS VM
Unit Sales Price plus 50% margin


Market potential based on AWS granular
charging model

per Object model assumes mean value
between the other two

In general, per Object model is still
considerable cheaper compared to
former VCPP per OSI price at 38% CSP
discount and 20% margin

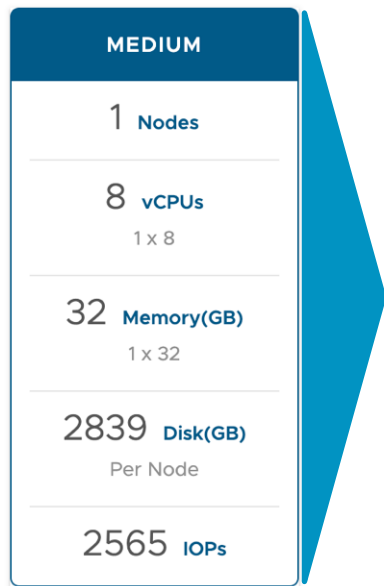
(Note: Logarithmic scale)

Charging Model Revenue Comparison across VCF Operations




TCO Comparison Example - VCF Operations

Significant Revenue and Margin Opportunity



Maximum Metrics: 2,500,000

Cost per VM: **1,200 USD**
2x (HA): **2,400 USD**
25% Margin: **3,000 USD**



Comparable AWS Costs

Tiered price for: 1,545,000 metrics

10,000 metrics x 0.30 USD = 3,000.00 USD

240,000 metrics x 0.10 USD = 24,000.00 USD

750,000 metrics x 0.05 USD = 37,500.00 USD

545,000 metrics x 0.02 USD = 10,900.00 USD

Total tier cost: 3,000.00 USD + 24,000.00 USD + 37,500.00 USD + 10,900.00 USD = 75,400.00 USD

CloudWatch Metrics cost (monthly): 75,400.00 USD

Tiered price for: 1,000 Dashboards

3 Dashboards x 0.00 USD = 0.00 USD

997 Dashboards x 3.00 USD = 2,991.00 USD

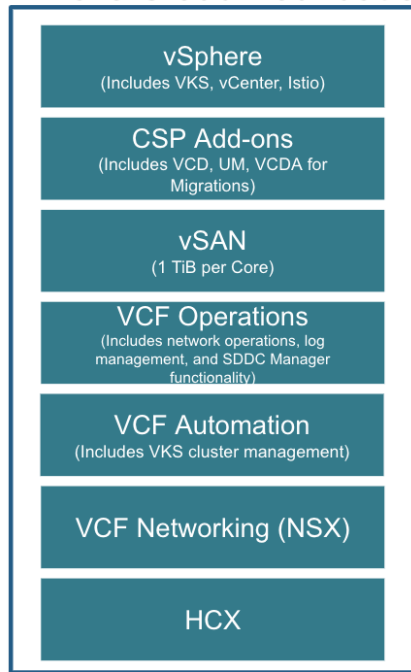
Total tier cost: 0.00 USD + 2,991.00 USD = 2,991.00 USD (Dashboards cost)

CloudWatch Dashboards and Alarms cost (monthly): 2,991.00 USD

Potential Revenue: 80,000 USD

VCF – Putting **VCF Network Operations** (VRNI) to work ...

VMware Cloud Foundation

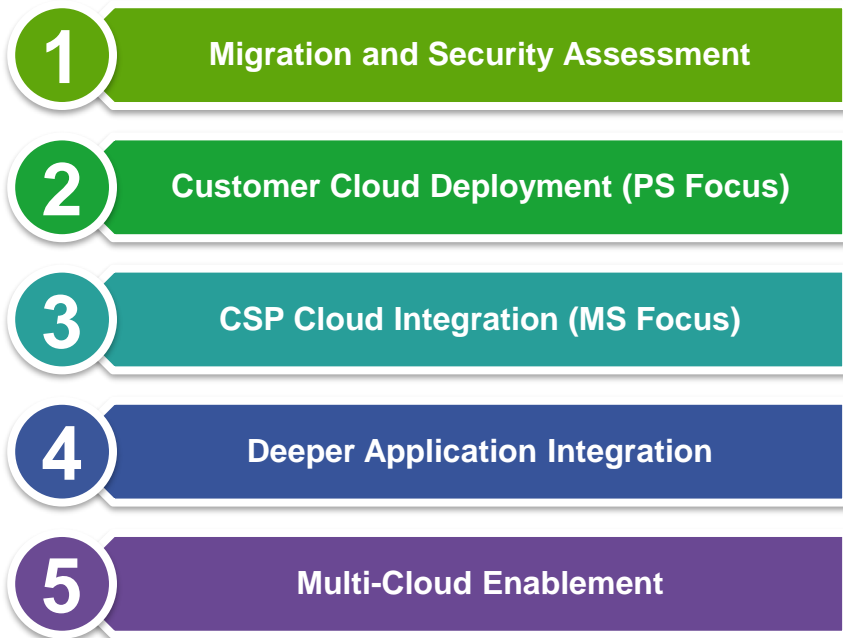


Networking Monitoring Opportunities

- Unlock further value (margin + revenue) across the VCF Deployment Models
- Professional Service opportunities in Private Cloud (on premise with the customer, in CSP DC or at the Edge)
- Ongoing Managed Service opportunities across Private and Public Clouds (CSP + Hyperscaler)
- Networking and Application Visibility
- Inclusive VCF license bundle

Service Implementation Highlights

Realizing Professional & Managed Services Opportunities with Customer Focus

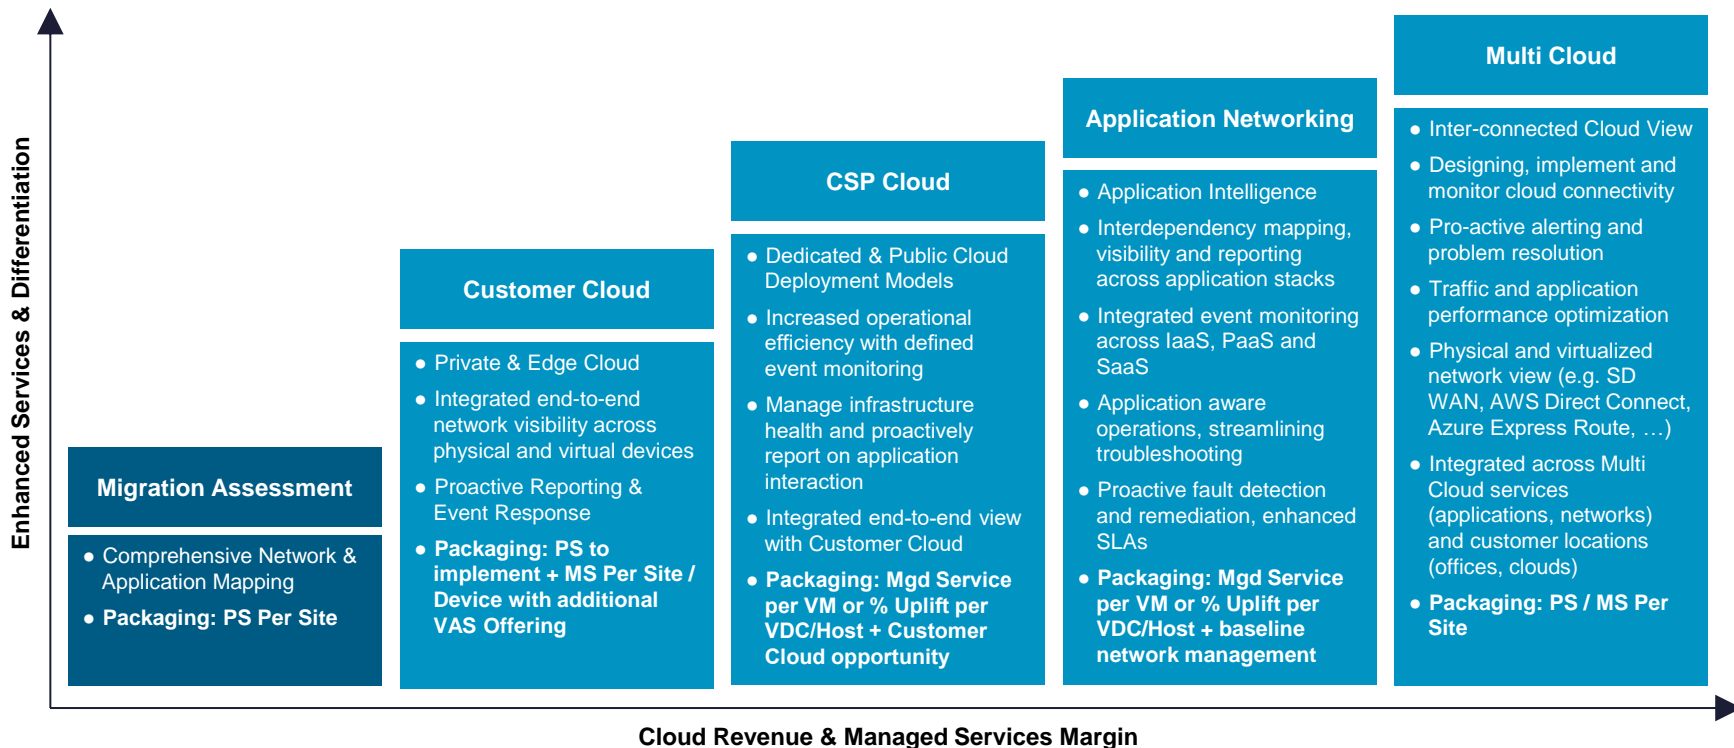


Target Customer

- Customers in a mature cloud maturity stage
- Customer that has complex cloud and network challenges and needs help from a CSP
- Customers looking to move to a OpEx model
- Customers retiring their data centers and increasingly moving to Cloud
- Customer wants to embark on digital transformation strategies
- Customers suffering Cloud pain due to sub-optimal network and application deployments

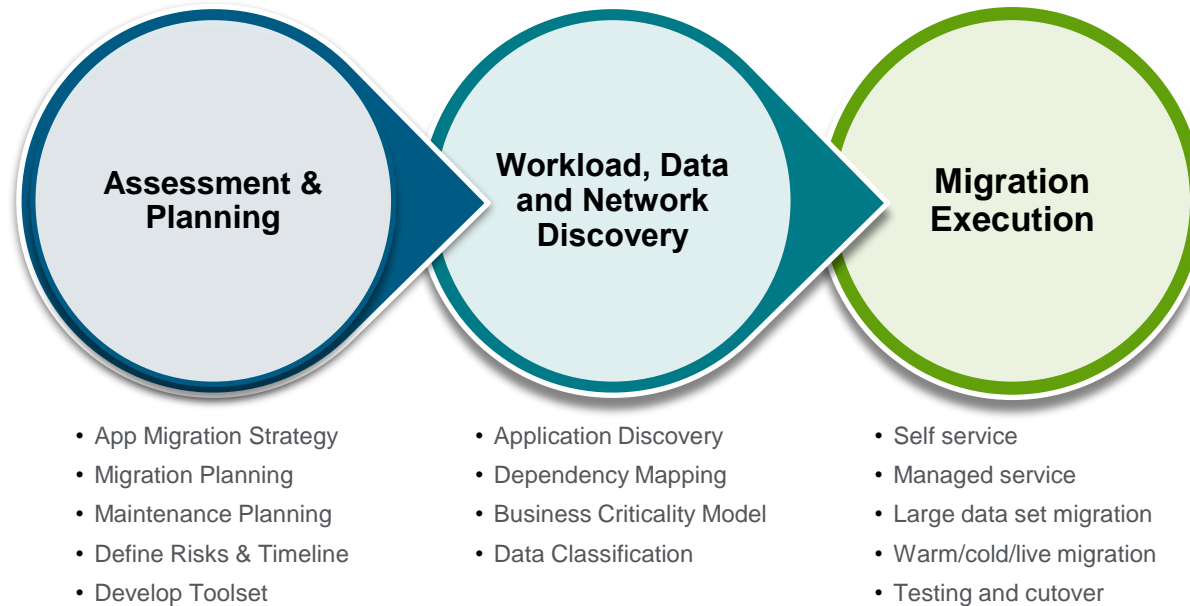
Delivering Comprehensive Network Monitoring Capabilities

Hybrid Cloud Professional Services + Managed Services Opportunities



Network Monitoring - Monetizing Pre-Migration Services

Deliver an end-to-end migration solution for private and public clouds



Monetizable Services with VCF

Grow beyond IaaS to generate revenue with additional Managed Services



Kubernetes as a Service Offering Tiers

Bring Offerings to Market in phased Approach to align with internal Learning Curve




Entry-Level Kubernetes

For Hosted and Managed Private Cloud:

- **Easy Cluster Deployment**
- 99.9% Availability SLA, best effort
- Compete with VMC, hyperscale Compute services and low-tier Kubernetes services


CSP self-service offering to capture Kubernetes and Developer workloads



Mid-Tier Kubernetes

For Managed Private and Public Clouds:

- Easy Cluster Deployment
- **Easy Cluster Lifecycle Management**
- **Enterprise Container Networking**
- **Choice of Kubernetes Management Tools**
- 99.95% Availability SLA, financially backed
- Compete with premium hyperscale Kubernetes services



High-End Kubernetes

For all VCF Deployment Models:

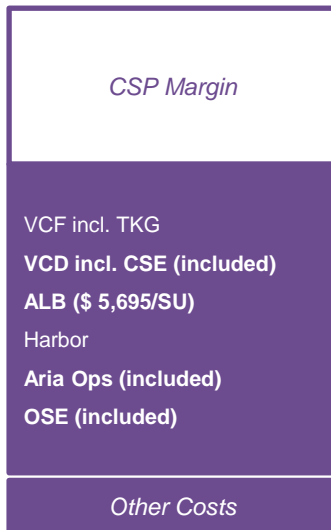
- Easy Cluster Deployment
- Easy Cluster Lifecycle Management
- Enterprise Container Networking
- Choice of Kubernetes Management Tools
- **Multi-Cloud Management**
- **Advanced Policy Management**
- 99.95% Availability SLA, financially backed
- Compete with premium hyperscale Kubernetes services and Add-Ons

Kubernetes as a Service Offering Tiers (BoM)

Cost Planning & Margin Model Examples



Offering:
Entry KaaS



Offering:
Mid-Tier KaaS



Offering:
High-End KaaS

Kubernetes Service Monetization

Market Price Points & Monetization Recommendations

Kubernetes Cluster: Worker & Master Nodes

Single-Node Control Plane:

- Price approx. \$850 per year*
- 4 vCPU, 16 GB vRAM
- Limited to 100 Worker Nodes
- Uptime SLA: 99.9%

Multi-Node Control Plane:

- Price approx. \$5,200 per year*
- 3x 8 vCPU, 32 GB vRAM
- Limited to 250 Worker Nodes
- Uptime SLA: 99.95%

Custom control plane options

Use same Worker pricing as IaaS VMs

+ Additional Networking Services: Load Balancer & Ingress

Per Tenant Advanced Load Balancer:

- More predictable / less granular compared to hyperscale offerings (by throughput, no. of rules, no. of instances)
- Charge for right-sized Service-Engine cluster(s)
- Offer full set of Kubernetes capabilities: L4 LB, L7 Ingress, WAF, GSLB, DDoS etc.

Hard requirement for VCD (managed private cloud and public cloud)

+ Advanced Kubernetes Services: Management, Registry etc.

Container Registry:

- Charge for VM resources (+ optional managed service fee)

Monitoring / Logging:

- Self-Service: Charge for solution resources (Aria / OSS components)
- Managed Service: Charge per instance under management (node, cluster, VM etc.)

Advanced Cluster Management (TMC):

- Self-Service: Charge for solution resources and license (TMC-SM)
- Managed Service: Charge per instance under management (node, cluster, VM etc.)

Kubernetes as a Service Offering Competitive Pricing


Example Comparison: Worker & Master Nodes

EKS, AKS, GKE: Standard K8s Offerings don't charge beyond IaaS VM costs plus a control plane fee

More expensive offerings from Hyperscalers like Azure Premium or RHOS

CSP KaaS is Cost Competitive in the Market and a valuable addition to an overall service catalogue

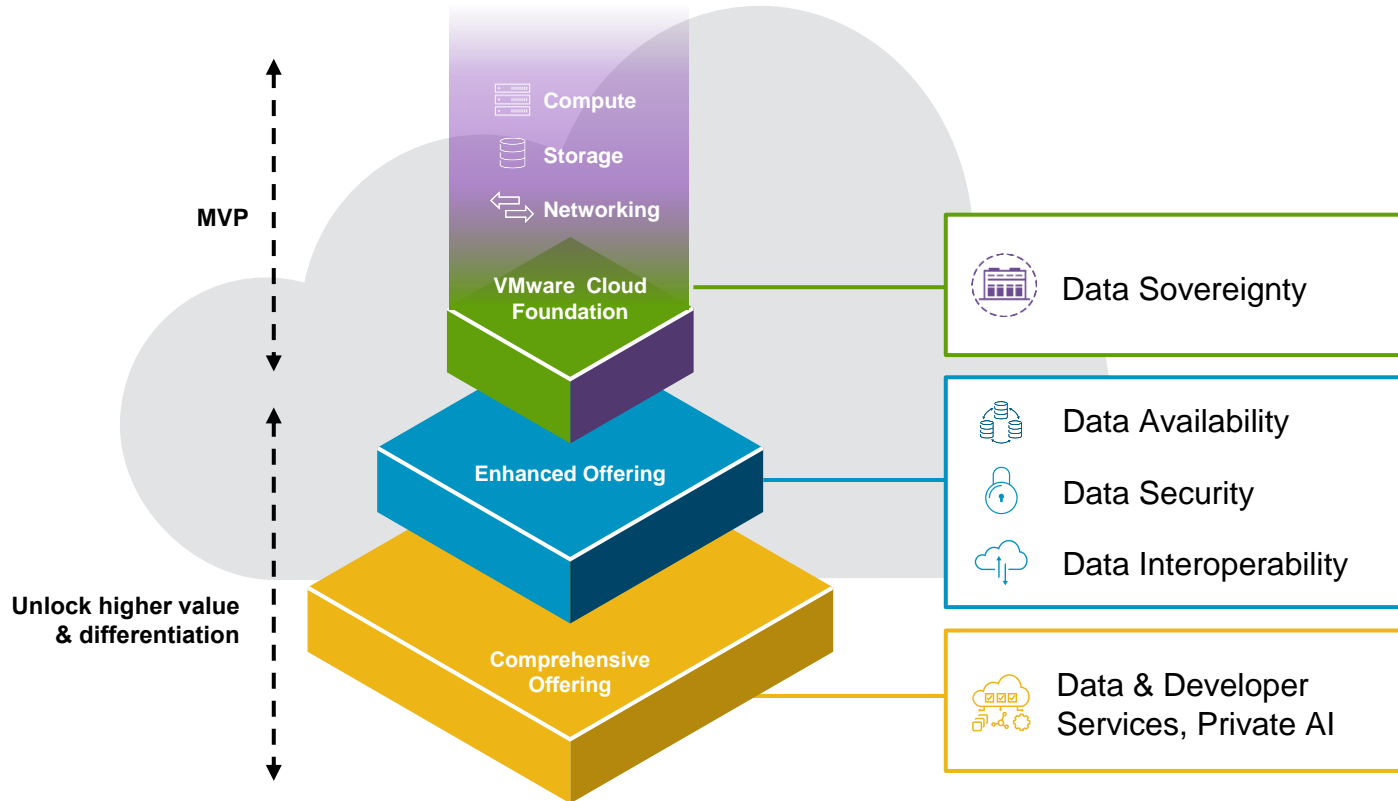
1Y Cost to Customer for running 3 Node (4x16) Kubernetes Clusters



Sovereign Cloud

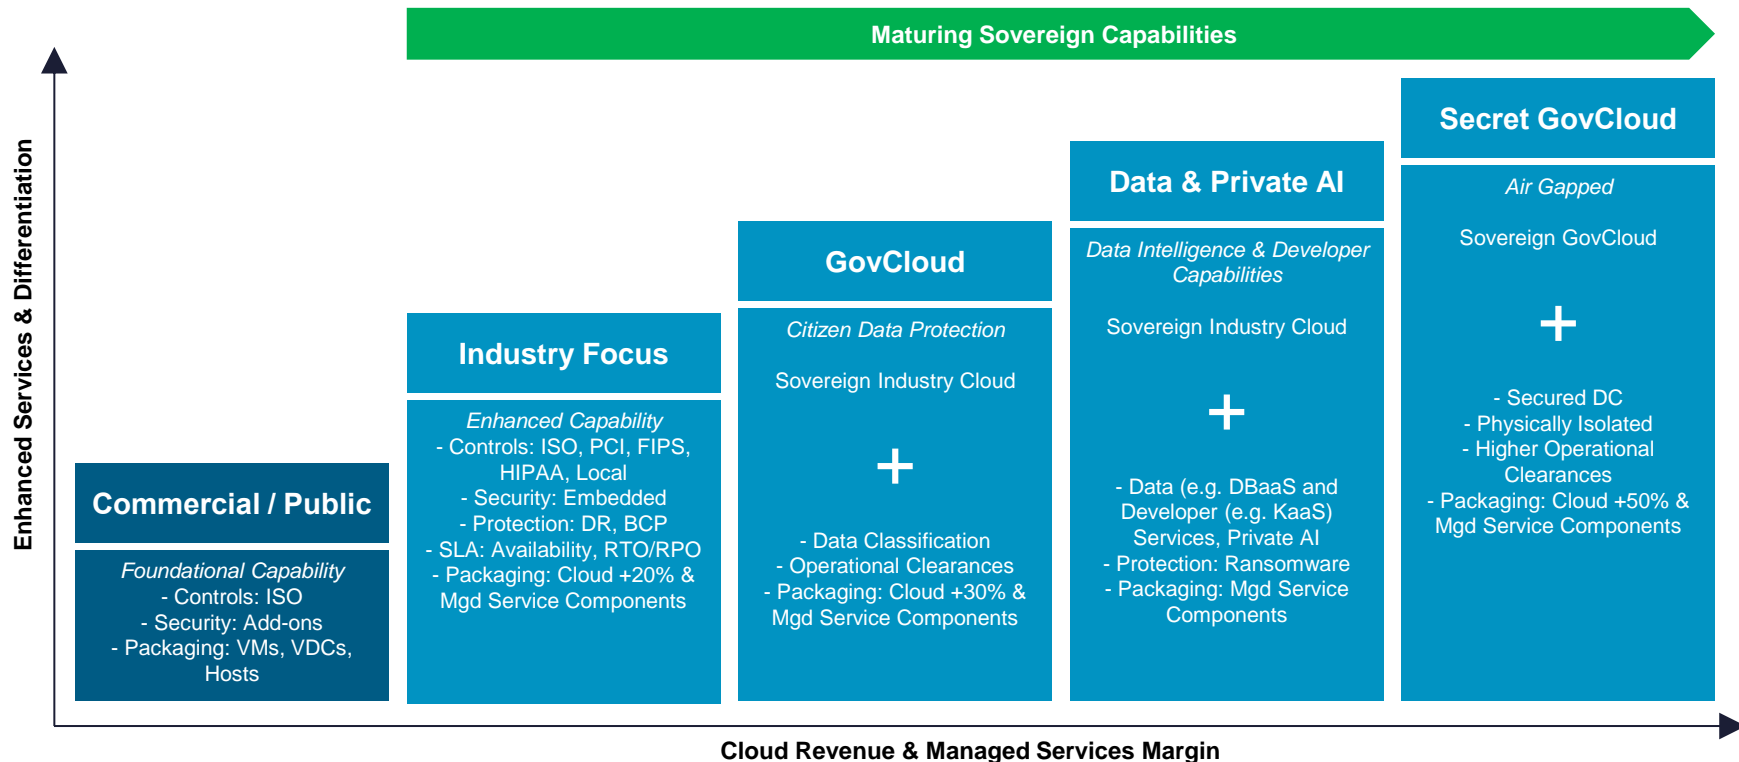
Evolving Sovereign Partner Offerings

From Data Sovereignty to Enhanced Digital Solutions



Delivering a Comprehensive Sovereign Cloud Capability

Managing Workload Risk and Driving Revenue and Margin with Differentiated Sovereign Clouds



Accelerating CSP Growth



CSP Support from CSP Cloud Practice & CSP Solution Teams

Accelerating partner growth and innovation


Strategy & Business Guidance



Cost Planning & Optimisation



Monetization, Planning & Pricing Strategies



Business Case Development & Deployment Plan



GTM Plan Development & Execution + Field Sales Enablement

Architecture & Design Guidance



Reference Architectures, Design Blueprints



Deep Technical Partner Activation & Enablement

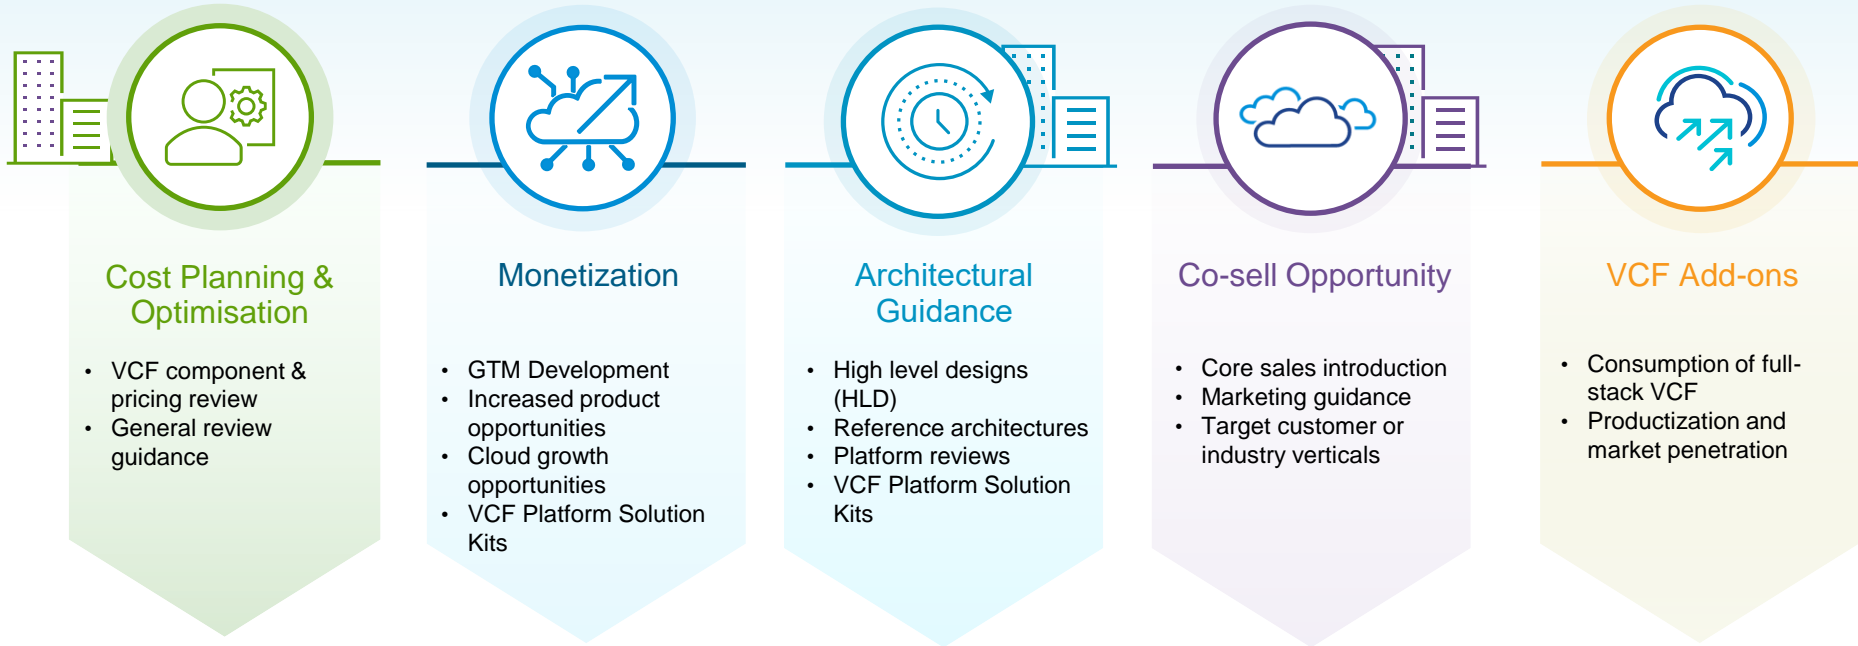


Service Design (HLD) & Offering Development



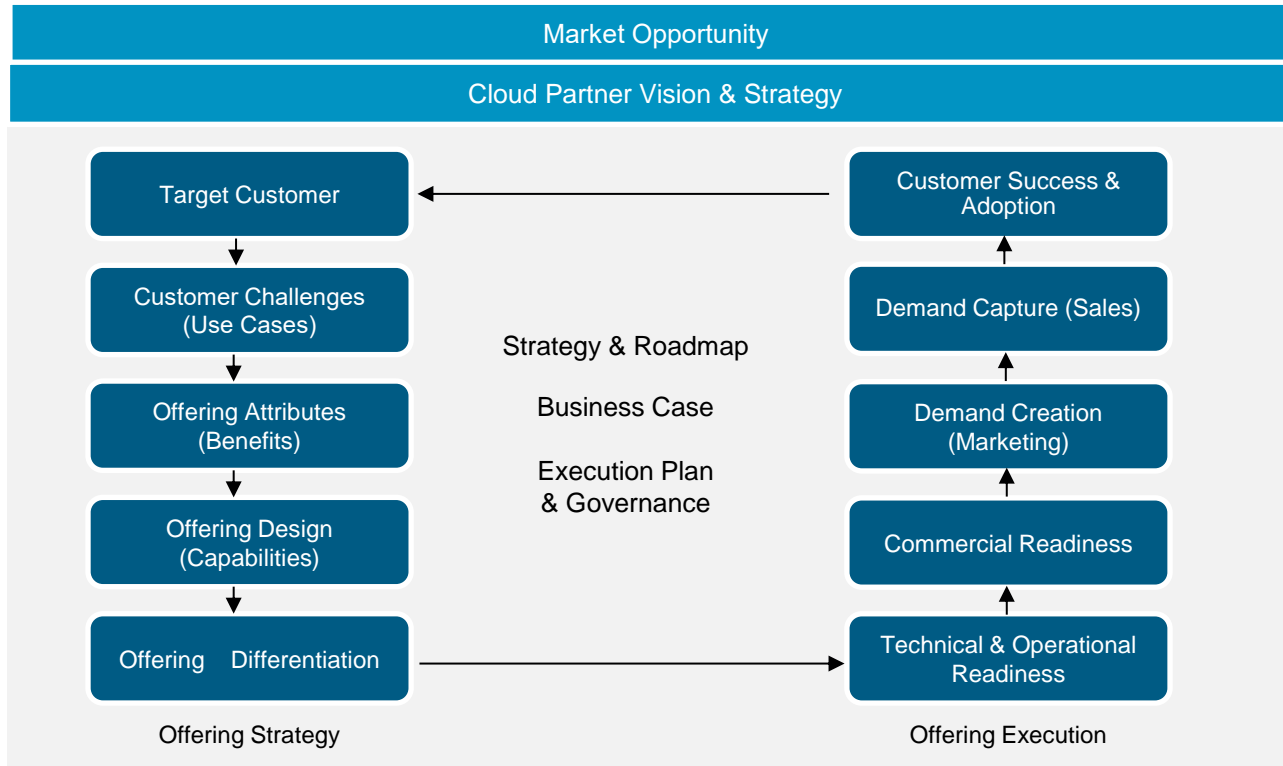
Solution Validation, PoC Support & Execution

Key Engagement Discussions with CSPs



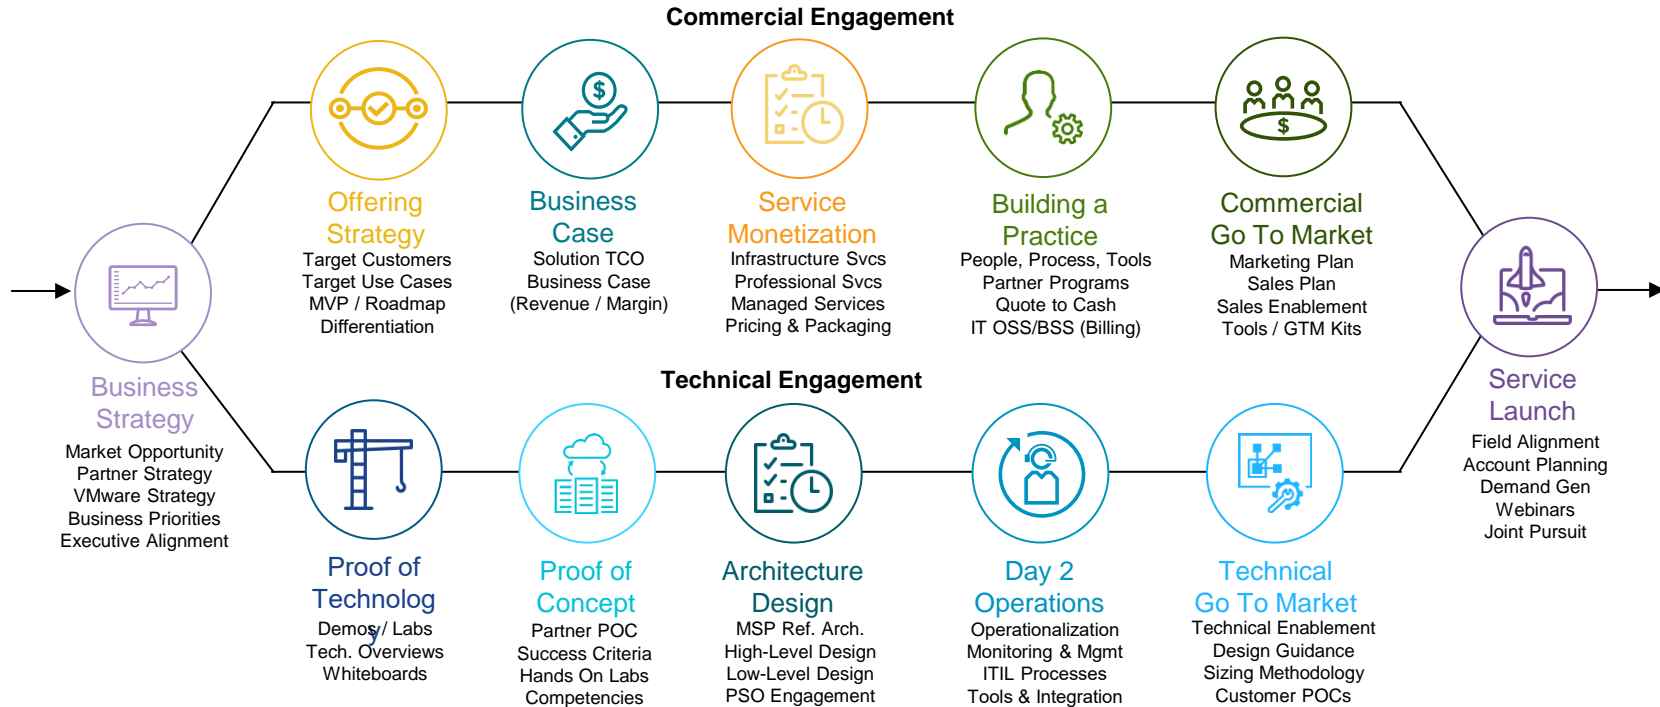
Partner Best Practices Offering Development Framework

Core elements of a partner GTM offering strategy & execution plan



VCSP Acceleration Opportunity

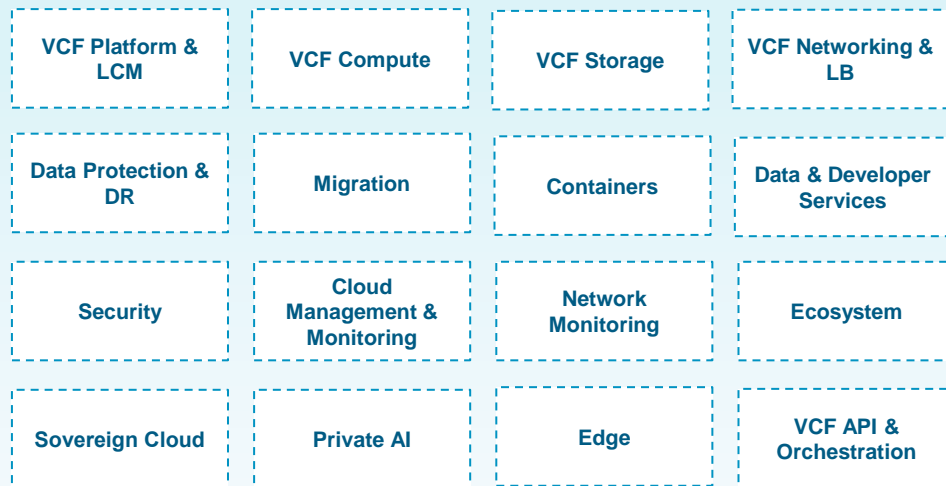
Accelerating partner growth and service creation



Partner Management: PBM & CSA

VCF Solutions Accelerate CSPs GTM

Enabling our CSPs to deliver differentiated solutions to market



CSP Solutions

- Differentiated
- Vertically Aligned
- Professional & Managed Services





Questions?

Book A Meeting



Scan Code Get a Callback or
Contact : phi.dao@crayon.com



Thank You!

Appendix

vSAN Turbocharged


Performance Benchmarking

- Shared under NDA, please do not re-distribute
- Indicative benchmarking results obtained given the criteria outlined in the following slides

Database Workloads



Linear Scalability




Lenovo ThinkAgile VX7531

- 8 node ESA cluster
- 2 Intel Xeon Gold 6348
@2.60GHz, 28 cores each
- 1024GB RAM
- Mellanox 100Gbits/s full duplex
- 6x6.4TB Micron 7450 NVMe

Pure Storage Takeout with vSAN Max


Preliminary
Testing

200 VM'S  Microsoft®
SQL Server®




Replacing Pure
with vSAN ESA

200 VM'S  Microsoft®
SQL Server®



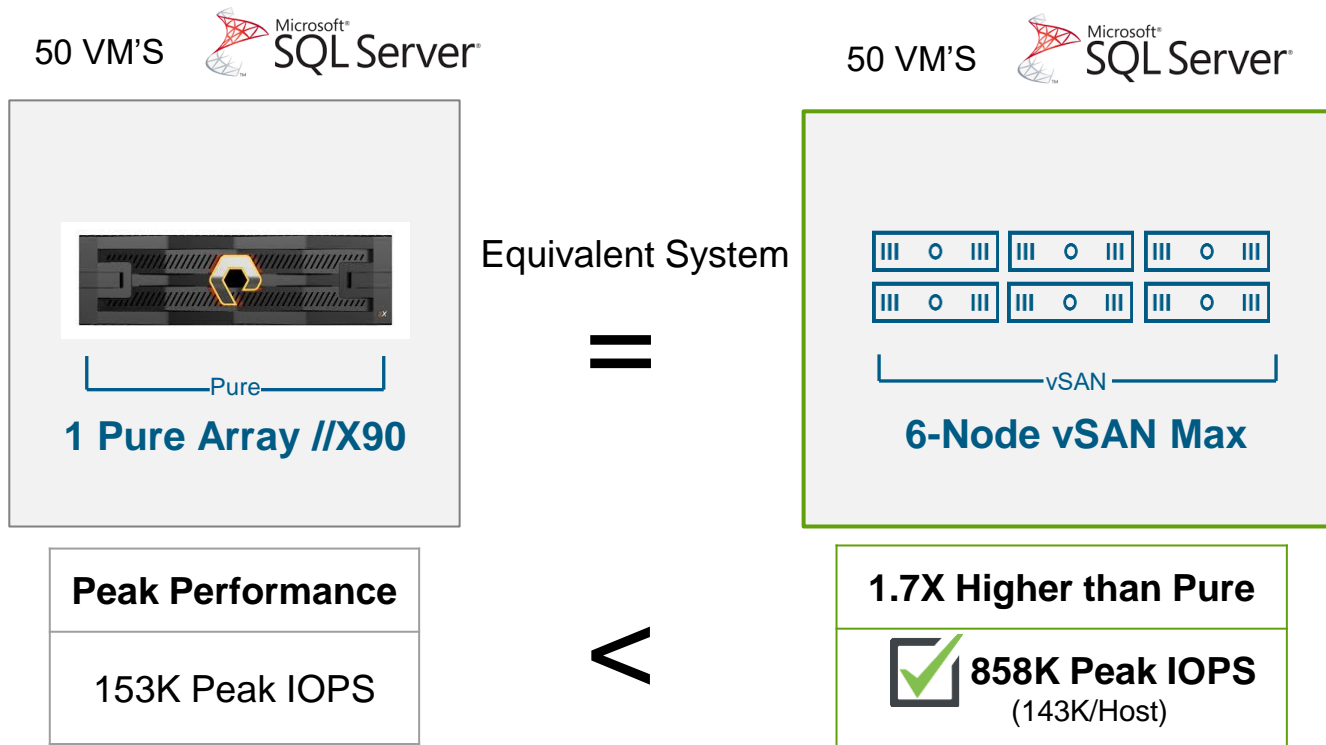
Expected Peak IOPS	Cost Baseline
1.8M IOPS (500K IOPS at 27% load)	\$1594/TiB (normalized)

Beats Pure Perf.	Lowers TCO
 2M IOPS (142K/Host)	 \$1104/TiB (30% lower)

Conclusion: vSAN ESA with 14-nodes can match the max performance of 4 Pure Storage Arrays at 31% lower of the cost per TB.

Customer Example - vSAN Performance is Better than Pure

Preliminary
Testing



vSAN Performance Details: Benchmark Testing

8K 70/30 Random Read Workload - Equivalent to SQL Database Workload

Preliminary
Testing


vSAN ESA & vSAN Storage Cluster – Benchmarks (vSAN 9.0, 25GbE + NVMe)

Metric	vSAN ESA (1 Node)	vSAN ESA (6 Nodes)	vSAN Storage Cluster (1 Node)	vSAN Storage Cluster (6 Nodes)
Peak IOPS (8KB, 70% Reads)	~143K	~858K	~142K	~852K
Latency	sub-ms	sub-ms	sub-ms	sub-ms

Identical Testbed for Apples-to-Apples Comparison

Storage Cluster	HCI Cluster
Hardware configuration: <ul style="list-style-type: none">• vSAN ESA• 6-node cluster• FTT=2 R6• 6 NVMe devices per host• 56 cores/host• 768G RAM per host• 25GbE networking	Hardware configuration: <ul style="list-style-type: none">• vSAN ESA• 6-node cluster• FTT=2 R6• 6 NVMe devices per host• 56 cores/host• 768G RAM per host• 25GbE networking

vSAN vs Pure Storage : Performance Scalability



- **vSAN's Scale-Out Model:**
Predictable, **Linear Growth**
- **Pure Storage's Scale-Up Model:**
Expensive Step Up

Conclusion


- vSAN offers better scalability of **performance** and **capacity** **without** unnecessary **overprovisioning**

Disaggregated Storage



Performance Comparisons - VCF 5.2


Comparing traditional storage with vSAN deployments



- vSAN HCI performance is ~1.8x higher than equivalent traditional storage platforms
- vSAN Storage Cluster performance is ~1.4x higher

Performance Comparisons - VCF 9.0

Comparing traditional storage with vSAN deployments



- VCF 9.0 closes the gap between vSAN Storage Cluster & vSAN HCI
- Performance gap between vSAN Storage Cluster and traditional storage increases to ~1.7x