VCFDCM52 Notes & Errata

Part Number EDU-EN-VCFDCM52-LEC (06-SEP-2024)

Version	Date	Author(s)	Description of Change
01	September 27, 2024	Fabrizio de Luca	Initial Release
02	October 5, 2024	Pawel Piotrowski	Removed comment on slide 2-9

Module 2

2-5 - About VMware Cloud Foundation

OUT-OF-DATE —> *VMware Cloud Foundation provides a standardized and configured infrastructure for* ***vSphere with Tanzu*** <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

2-6 - VMware Cloud Foundation Software Bill of Materials

INCOMPLETE —> ***You can use VMware Aria Suite Lifecycle to deploy*** *VMware Aria Automation, VMware Aria Operations, VMware Aria Operations for Logs, and Workspace ONE Access* <— **you can install VMware Aria Operations for Networks too**.

MORE DETAILS —> ***Workspace ONE Access*** <— **"While VMware Identity Manager has been officially renamed VMware Workspace ONE Access, VMware Aria Suite Lifecycle uses the VMware Identity Manager 3.3 clustered environment. VMware Aria Suite Lifecycle does not support VMware Workspace ONE Access 20.x and later"**.

[Source: https://docs.vmware.com/en/VMware-Aria-Suite-Lifecycle/8.18/lifecycle-install-upgrade-manage/GUID-5E1CB756-CE86-430D-89C0-DE3831C33738.html]

2-7 - VMware Cloud Foundation Architecture

OUT-OF-DATE (graphic and notes) —> ***vSphere with Tanzu*** *transforms vSphere into a platform for running Kubernetes workloads natively* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

MORE DETAILS —> *Actual virtualized workloads, such as applications and services, run in* ***VI workload domains*** <— **VI** stands for **Virtual Infrastructure**.

2-11 - vSphere Overview

BOTCHED —> *vCenter: Service used to manage multiple ESXi hosts and* ***provide*** *vSphere vMotion and* ***vSphere High Availability functionalities*** <— actually, **vCenter Server does not PROVIDE vSphere HA**: it rather **CONFIGURES** vSphere HA. **All vSphere HA functionalities – from failure detection to workloads failover coordination – are provided by the Fault Domain Manager (FDM) agents running in the ESXi hosts**. Once configured, vSphere HA can provide all its functionalities even if the vCenter Server is involved in the failure. On top of that, why mentioning just those two functionalities and not – in example – **vSphere Distributed Resource Scheduler (DRS)** and others? IMHO, this bullet list item shall list them all (not practical) or none.

MORE DETAILS —> ***Changing generated settings, such as NTP, DNS, or VM names, might result in unpredictable results*** <— starting with VCF 5.2 VMware added support for: **"Out of Band Changes from vCenter: Out of Band changes from vCenter can be manually synced with SDDC Manager. This includes inventory changes (for example, adding a host to a cluster) and object name changes (for example, datacenter name, datastore name, port group name)".**
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/rn/vmware-cloud-foundation-52-release-notes/index.html#What's%20New]
For instructions about using the **VCF Import Tool** with the **sync** parameter click the link below.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-D871245E-5156-4741-A6FA-9A6C03B04315.html]

2-12 - vSAN Overview

INCOMPLETE —> ***Support for both OSA and ESA architectures*** <— starting with VCF 5.2 VMware introduced the following enhancements:

- **"vSAN Max support: vSAN Max is a disaggregated storage offering which enables petabyte scale storage-only clusters. vSAN Max is powered by ESA as the underlying storage platform, which is a high-performance file system that can scale up to high densities with no penalty to performance. ESA also provides other benefits such as built-in, efficient, scalable snapshots, and low overhead data services such as encryption and compression".**
- **"vSAN ESA Stretched Cluster: VCF users can now configure ESA Stretched Cluster in vSAN Ready Nodes. It enables customers to take the concept of fault domains to protect an environment spanning two physical sites from downtime in the event of a site failure".**

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/rn/vmware-cloud-foundation-52-release-notes/index.html#What's%20New]

2-13 - NSX Overview

OUT-OF-DATE —> *NSX also supports modern applications through integration with* ***vSphere with Tanzu*** <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.
[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

2-14 - vSphere with Tanzu Overview

OUT-OF-DATE —> ***vSphere with Tanzu*** *provides the ability to run Kubernetes workloads directly on ESXi hosts and to create upstream Kubernetes clusters in dedicated resource pools* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

2-15 - VMware Aria Suite Overview

OUT-OF-DATE —> **VMware Aria Log Insight**: You deploy **VMware Aria Log Insight** to get real-time log file management and log file analysis for deep operational visibility and fast troubleshooting across hybrid cloud environments <— AND —> ***VMware Aria Network Insight****: You deploy* ***VMware Aria Network Insight*** *to deliver intelligent operations for software-defined networking and security, helping you build an optimized, highly available, secure network infrastructure across multicloud environments* <— **starting with version 8.12 – back in 20th April 2023 – all VMware vRealize Suite products have been renamed to the following:**

- **VMware Aria Suite Lifecycle**
- **VMware Aria Automation**
- **VMware Aria Operations**
- **VMware Aria Operations for Logs**
- **VMware Aria Operations for Networks**

[Source: https://docs.vmware.com/en/VMware-Aria-Suite-Lifecycle/8.12/rn/vmware-aria-suite-lifecycle-812-release-notes/index.html#What's%20New]

2-18 - VMware vSphere Foundation Add-Ons

MORE DETAILS <— get more details about:
VMware Tanzu: https://tanzu.vmware.com
VMware Live Recovery: https://www.vmware.com/products/cloud-infrastructure/live-recovery
VMware Avi Load Balancer: https://www.vmware.com/products/cloud-infrastructure/avi-load-balancer

2-19 - VMware Cloud Foundation Licensing

MORE DETAILS <— get more details about:
VMware HCX: https://www.vmware.com/products/cloud-infrastructure/hcx
VMware Data Services Manager: https://www.vmware.com/products/cloud-infrastructure/data-services-manager

2-20 - VMware Cloud Foundation Add-Ons

MORE DETAILS <— get more details about VMware vDefend solutions at:
VMware vDefend Distributed Firewall: https://www.vmware.com/products/cloud-infrastructure/vdefend-distributed-firewall
VMware vDefend Gateway Firewall: https://www.vmware.com/products/cloud-infrastructure/vdefend-gateway-firewall
VMware vDefend Advanced Threat Prevention: https://www.vmware.com/products/cloud-infrastructure/vdefend-advanced-threat-prevention

VMware Private AI Foundation: https://www.vmware.com/products/cloud-infrastructure/private-ai-foundation-nvidia

Module 3

3-7 - VMware Cloud Foundation Architectures

OUT-OF-DATE —> ***vSphere with Tanzu*** *and Tanzu Kubernetes Grid are supported* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

OUT-OF-DATE —> *Deployments of* ***Horizon workload domain***, ***vSphere with Tanzu***, *Tanzu Kubernetes Grid, and Tanzu Kubernetes Grid Integrated are supported* <— besides the note about **vSphere with Tanzu**, does it still make sense talking about **Horizon** VDI considering that the product is no more VMware's since the EUC Division has been divested from Broadcom to become a new company called **Omnissa**, under KKR (https://www.kkr.com) ownership?

[Source: https://www.omnissa.com/introducing-omnissa-the-former-vmware-end-user-computing-business/]

3-11 - Management Domain Design Decisions

TYPO —> *When deciding on the vCenter Server Appliance size, consider whether you want to deploy a consolidated architecture and plan to scale the design, or (for either business reasons or technical requirements) you plan to* ***maintain the consolidated design with more than 100 hosts*** <— this is likely a typo and the author presumably meant to say **more than 10 hosts**. Assumption is made based on the VMware best practice – also referred in this same slide notes – that suggests that once a consolidated design has reached 8 or 10 nodes (that is, ESXi hosts), customers shall consider moving to the standard design.

3-24 - External Certificates During Bring-Up: Replacing Default Certificates

BOTCHED <— the procedure description contains some inaccuracies and a few missing steps.

1. Install the base ESXi hosts according to VMware Cloud Foundation documentation.
2. **Validate that the TSM-SSH service is running if you want to connect to the ESXi Shell via SSH.**
3. **Validate that the TSM service is running if you want to connect to the ESXi Shell via the DCUI.**
4. Rename the existing certificate files in the /etc/vmware/ssl directory.
5. Copy the new signed certificate files to the /etc/vmware/ssl directory.
6. Rename the new certificate files to replace the original files (`rui.crt` and `rui.key`).
7. **Restart the host management agents by running the following commands:**
    a. `/etc/init.d/hostd restart`
    b. `/etc/init.d/vpxa restart`

8. **Stop the TSM and/or TSM-SSH service.**

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-deploy/GUID-91824B56-91F3-44FE-B67B-180A50B94717.html]

3-25 - External Certificates During Bring-Up: Customizing the JSON File

MISSING DETAILS —> ***Update the JSON file with the third-party certificate information*** <— the slide doesn't provide any details about how to update the converted JSON file with the third-party certificate information before uploading it to the Cloud Builder appliance VCF deployment wizard. You must "`Update the securitySpec section, choosing Custom for the esxiCertsMode and entering your signing CA chain for certChain`". For an example of the `securitySpec` section formatting, see the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-deploy/GUID-70ECFCCC-8B4D-4BE7-824C-6A6731C9CB06.html]

Module 4

4-10 - Using Identity Providers in VMware Cloud Foundation

MORE DETAILS (NEW in VCF 5.1) —> ***OKTA*** *– External IDP using VMware Identity Service* <— for more details about the configuration procedure, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-779756CD-3FC9-4436-A324-439B8C199515.html?hWord=N4IghgNiBcIPIGsBOYQF8g]

MISSING DETAILS (NEW in VCF 5.2) —> **Microsoft *Entra ID*** *(formerly known as Azure ID)* <— VCF 5.2 has introduced support for Microsoft Entra ID as an Identity Provider (ref. Release Notes). For more details about the configuration procedure, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-F7F02B80-59F3-437B-91A6-AB71E27A50D3.html]

4-11 - About the Online Depot

MORE DETAILS <— for more details about the download of install bundles from either an online or an offline software depot (mutually exclusive), check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-896F8E45-69B2-465B-A378-2CA24AB0EFE9.html]

OUT-OF-DATE —> *In environments where SDDC Manager lacks Internet connectivity, bundles can be downloaded and transferred using the* ***Offline Download Bundle Utility*** <— this tool was used with VCF 4.x; the tool is now called **Bundle Transfer Utility**. For more details about the download procedure, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-lifecycle/GUID-8FA44ACE-8F04-47DA-845E-E0863094F7B0.html]

4-16 - Administering VMware Cloud Foundation

MORE DETAILS <— starting with VCF 5.2 VMware added support for:

"Out of Band Changes from vCenter: Out of Band changes from vCenter can be manually synced with SDDC Manager. This includes inventory changes (for example, adding a host to a cluster) and object name changes (for example, datacenter name, datastore name, port group name)".

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/rn/vmware-cloud-foundation-52-release-notes/index.html#What's%20New]

Despite this feature shall not mean an implicit authorization to make changes from vCenter, **in case SDDC Manager and vCenter Server inventories get out of sync, they can be re-synchronized** using the procedure explained in the online documentation at the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-D871245E-5156-4741-A6FA-9A6C03B04315.html]

4-19 - User-Created Service Accounts

MORE DETAILS —> *You **create service accounts** to provide external applications access to the VMware Cloud Foundation environment through APIs* <— for more details about the procedure to create a **Service (aka Automation) Account**, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-6257AF49-F094-4985-A9E1-CCFDB692979C.html]

4-25 - Rotating Passwords

OUT-OF-DATE —> ***Rotate All: Immediately changes the passwords for all selected accounts*** <— AND —> ***Rotate Now: Immediately changes the passwords for one or more selected accounts*** <— the commands that can be seen in the VCF UI are:

- **ROTATE PASSWORDS**: to rotate passwords only for the select accounts (one or more)
- **ROTATE ALL PASSWORDS**: to rotate all passwords for the selected component

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-28D29FFA-2D81-4781-AD79-85697497D45B.html]

4-26 - Updating and Remediating Passwords

BOTCHED —> ***The UPDATE option updates the password for an account that is in sync with SDDC Manager*** <— the statement doesn't explain who is providing the password: the system (randomly) or the admin (typing it). According to the online documentation, **"You can manually change the password for a selected account. Unlike password rotation, which generates a randomized password, you provide the new password"**.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-815B1682-BF65-4AFA-80FA-A4AB9DFE0853.html]

4-27 - Using the lookup_passwords Command

MORE DETAILS —> *The **SDDC Manager API** is the primary method for retrieving password information* <— we use the **`GET /v1/credentials`** call. For more details, check the link below.

[Source: https://developer.broadcom.com/xapis/vmware-cloud-foundation-api/latest/credentials/]

4-28 - Password Management: Exceptions

BOTCHED —> ***You use the API change the admin password used for the API calls*** <— the sentence doesn't really mean a thing. The author presumably meant to state something similar to what explained in the notes. In example: **You use SSH to access SDDC Manager and run a shell script to change the admin password used for the API calls**.

MORE DETAILS <— despite using the following script is still a viable procedure

`/opt/vmware/vcf/commonsvcs/scripts/auth/set-basicauth-password.sh admin <password>`

VMware's online documentation suggests using the **`PATCH /v1/users/local/admin`** REST API call. For more details, check the link below.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-8BF81022-0DF6-4AC8-8D02-6B15824D98E6.html]

Module 5

5-14 - About VDS

MORE DETAILS —> *The VDS **MTU size must be set to 1,600 bytes or greater** from the vSphere Client to use it with NSX* <— while 1,600 bytes is the bare minimum, an MTU size of 1,700 bytes provides more room for the GENEVE header variable length options used by VMware for security and monitoring purposes. Ideally, jumbo frames shall be set to the maximum size of 9,000 or 9,216 bytes to enable increasing the GENEVE MTU size without needing to change the MTU size of the physical infrastructure.
For more details see VCF Design Guide Requirement ID VCF-NET-REQD-CFG-004.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-5E4E4042-4F97-42BF-9864-AF5E414BE949.html]

5-20 - NSX Manager Cluster Placement

WRONG GRAPHIC <— the vCenter Server Appliance for VI Workload Domain C is missing.

5-22 - Workload Domain Traffic Separation
5-24 - Management Domain Deployment
5-26 - VDS Profile-1: Default Unified Fabric
5-27 - VDS Profile-2: Storage Traffic Separation
5-28 - VDS Profile-3: NSX Traffic Separation

BOTCHED GRAPHIC <— IMHO, the **ESXi Hypervisor Transport Node TEP(s)** shall be placed **UNDER** the NSX Prepared vDS, not **ABOVE**. TEP(s) in ESXi hosts are vmkernel interfaces; hence, they connect to the vSphere VDS with a dvPort, like any other workload. The slide graphic gives students the wrong impression that the TEP(s) are located North the ESXi hosts uplink interfaces, which is absolutely wrong. If the objective was representing the GENEVE Overlay traffic going northbound, then it has to be implemented differently.

5-25 - Dual DPUs Support

MORE DETAILS —> ***VMware Cloud Foundation 5.2 supports Dual DPUs**, in which a single ESXi host can use two DPUs* <— for more details on the **DPU/SmartNIC** technology (formerly known as Project Monterey), check the link below.
[Source: https://core.vmware.com/blog/rise-dpus-infrastructure]

5-29 - Creating a VDS Configuration Using the API

TYPO —> Click **HostCommissionSpeclink** to copy the text that appears in the Value text box and place it in a text editor <— you shall read "**HostCommissionSpec link**".
OUT-OF-CONTEXT —> ***The areas highlighted include the name and the URL.** The ID is in the URL* <— the slide graphic shows the ClusterCreationSpec code, not the **Administration >**

Network Settings page from the VCF UI; hence, there is no Network Pool name nor its ID highlighted in the slide.

5-30 - VI Workload Domain Switch Configuration

OUT-OF-DATE —> ***A single DPU SmartNIC device can be used*** <— **since VCF 5.2, you can deploy a VI Workload Domain that uses a Dual-DPU configuration.** Using the SDDC Manager UI, the procedure is the following:

"To deploy a new VI workload domain that uses ESXi hosts with two data processing units (DPU), you must configure a custom vSphere Distributed Switch (VDS) with the following settings:

- **A single VDS that uses all four DPU-backed nics**
- **Uplinks (uplink1 through uplink4) are mapped to the DPU-backed nics**
- **The NSX network operational mode is set to Enhanced Datapath - Standard**
- **In the NSX uplink profile, uplink-1 and uplink-2 are Active and uplink-3 and uplink-4 are Standby"**

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-E64CEFDD-DCA2-4D19-B5C5-D8ABE66407B8.html#GUID-BCF02F28-844A-4840-9BD2-4E0633FCE178__GUID-D2C76C1F-CCC5-480E-92B3-7394E0180026]

5-34 - Importing vSphere Clusters to VMware Cloud Foundation

TYPO —> *In this scenario, you begin by identifying a vSphere cluster to* ***~~the~~*** *host SDDC Manager* <— ignore the definite article "the".

MORE DETAILS —> *You* ***download*** *and deploy the SDDC Manager appliance (OVA) into the cluster and copy the import scripts (Tar archive) to SDDC Manager* <— the VCF 5.2 related software that needs to be downloaded from the Broadcom support portal is:

- VCF SDDC Manager Appliance (VCF-SDDC-Manager-Appliance-5.2.0.0-24108943.ova)
- VCF Import Tool (vcf-brownfield-import-5.2.0.0-24108578.tar.gz)
- VMware Software Install Bundle - NSX_T_MANAGER 4.2.0.0 (bundle-124941.zip)

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-20261403-42CD-42D1-A8FE-CB29816F9825.html]

MORE DETAILS —> *You then use SSH to connect to the SDDC Manager instance to extract the import tarball, and you run the* ***import script with the convert option*** *to convert the vSphere cluster into a management domain* <— the **VCF Import Tool Options and Parameters** are listed in the online documentation page below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-44CBCB85-C001-41B2-BBB4-E71928B8D955.html]

5-35 - Migrating VLAN-Backed Workloads to NSX Overlay

MORE DETAILS —> ***NSX Manager provides a VLAN to Overlay migration tool*** *to help you migrate your workloads from a vSphere VLAN to an NSX overlay network so that you can adopt the full benefits of network virtualization* <— the tool can be access by navigating in the NSX 4.2 UI to **System > Migrate > VLAN-to-Overlay Migration**. Step by step instructions can be found by checking the link below.
[Source: https://docs.vmware.com/en/VMware-NSX/4.2/migration/GUID-5C61A8F6-EF3F-4ACB-9BF8-8C942C143086.html]

5-41 - NSX Edge Cluster Deployment Considerations
BOTCHED —> *The management network and gateway must be reachable from the* ***NSX hosts*** *and NSX Edge instances* <— IMHO, this is too vague, it should be better using **NSX Hypervisor transport nodes**.

5-43 - Specifying the NSX Edge Cluster Use Case
OUT-OF-DATE —> *Select Kubernetes - Workload Management if the NSX Edge cluster is used to provide north-south routing for a* ***vSphere with Tanzu*** *cluster* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.
[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]
BOTCHED —> *This option* ***also*** *enables the NSX Edge cluster to support* ***vSphere infrastructure-as-a-service (IaaS) control plane*** <— **vSphere with Tanzu** (up to version 8.0 Update 2) and **vSphere IaaS Control Plane** (starting with version 8.0 Update 3) – though the technology has been improved and extended with additional functionalities – **are the same thing**.
[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]
MORE DETAILS —> *The NSX Edge large appliance is suitable for environments that require* ***load balancing*** <— **starting with NSX version 4.2.0, VMware has announced a change in the future of the built-in NSX Load Balancer** (which was initially deprecated in version 4.0.1.1):

Entitlement Change for the NSX Load Balancer
In a future major release of NSX, VMware intends to change the entitlement of the built-in NSX load balancer (a.k.a. NSX-T Load Balancer). This load balancer will only support load balancing for Aria Automation, IaaS Control Plane (Supervisor Cluster), and load balancing of VCF infrastructure components.
VMware recommends that customers who need general purpose and advanced load balancing features purchase Avi Load Balancer. Avi provides a superset of the NSX load balancing functionality including GSLB, advanced analytics, container ingress, application security, and WAF.
Existing entitlement to the built-in NSX load balancer for customers using NSX 4.x will remain for the duration of the NSX 4.x release series.

[Source: https://docs.vmware.com/en/VMware-NSX/4.2.0/rn/vmware-nsx-420-release-notes/index.html#Feature%20and%20API%20Deprecations,%20Behavior%20Changes]
WRONG —> *If you require certain services such as **NAT** or VPN, **you must choose active-standby*** <— don't forget that, since NSX v4.0.1.1: **"Stateful Active-Active Edge Services: this release introduces support for stateful services on Tier-0 and Tier-1 gateway in Active-Active HA mode. The following stateful services are supported: L4/L7 Gateway Firewall, URL Filtering, NAT and TLS Inspection"**.
[Source: https://docs.vmware.com/en/VMware-NSX/4.0.1.1/rn/vmware-nsx-4011-release-notes/index.html#What's%20New]
BOTCHED / GENERIC LINK —> *For information about services that require an active-standby configuration, see VMware NSX Administration Guide at **https://docs.vmware.com/en/VMware-NSX/index.html*** <— the direct link to the documentation pages is **https://docs.vmware.com/en/VMware-NSX/4.2/administration/GUID-D33CD33C-7CBE-4414-9109-2584AE80E1E3.html**

5-45 - Configuring NSX Edge Appliances
WRONG —> ***One edge tunnel endpoint (TEP) IP address for each edge node*** *on the edge TEP VLAN* <— IMHO, this shall be better described as **one edge tunnel endpoint (TEP) IP address for each edge node virtual NIC used as an uplink by the edge N-VDS attached to the Overlay Transport Zone**.
Considering the NSX Edge Design Requirements **VCF-NSX-EDGE-REQD-CFG-002** and **VCF-NSX-EDGE-REQD-CFG-004**, and the NSX Edge Design Recommendation **VCF-NSX-EDGE-RCMD-CFG-007** (see the link below "NSX Edge Node Design for VMware Cloud Foundation"), we are rather going to have **two edge TEP IP addresses for each edge node**.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-B6F707E5-E23F-43FB-8396-E0DB4B2371B4.html#nsx-edge-design-requirements-6]
WRONG —> ***One uplink IP address for each edge node*** *on separate uplink VLANs* <— according to the BGP Routing Design Requirements **VCF-NSX-BGP-REQD-CFG-001** and **VCF-NSX-BGP-REQD-CFG-002**, we are rather going to have **two uplink IP addresses for each edge node**.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-FF3FF1C8-BBC8-4502-85A9-20C28D6FD631.html#bgp-routing-design-requirements-1]

5-49 - Adding NSX Edge Cluster Nodes
BOTCHED / NON-SENSE —> *When the Tier-0 service high availability setting is active-standby and you require more than two NSX Edge nodes for services* < AND > *When the Tier-0 service high availability setting is active-active and you require more than eight NSX Edge nodes for services* <— both statements refer to the configuration maximums for a single Tier-0 Gateway – respectively for active-standby (2 nodes) and active-active (8 nodes) high availability modes – and mention the idea

of adding "more than <N> NSX Edge nodes for services". Besides being technically possible – an Edge cluster can have up to 10 Edge nodes assigned -, the author doesn't clarify WHAT will use the additional nodes… they may want to rephrase the concept.

5-50 - Removing NSX Edge Cluster Nodes

BOTCHED / NON-SENSE —> *For active-active configurations, the NSX Edge cluster must contain two or more NSX Edge nodes* ***with two or more Tier-0 routers*** *after the NSX Edge nodes are removed.* <— it simply doesn't make sense: just imagine having a single Tier-0 Gateway with an active-active Edge cluster with, say in example, 4 Edge nodes and you want to remove 1 or 2 of them… according to what stated, you won't be allowed (because you have just 1 Tier-0 router). This rather seems to me a "copy&paste" error: the author may want to rephrase the concept.

5-56 - Connectivity to Physical Layer 3 Devices: Routing Feature Set

WRONG —> *The connectivity between the Tier-0 logical router and the Tier-1 logical router is provided by a router link, which is a /31 subnet within the* ***100.64.0.0/10*** *reserved address space (RFC6598) and is automatically created when deployed.* <— the VMware NSX online documentation always refers to **100.64.0.0/16**, in example:

- NSX Edge Networking Setup: https://docs.vmware.com/en/VMware-NSX/4.2/installation/GUID-370D06E1-1BB6-4144-A654-7AF2542C3136.html
- Add an NSX Tier-0 Gateway: https://docs.vmware.com/en/VMware-NSX/4.2/administration/GUID-7B0CD287-C5EB-493C-A57F-EEA8782A741A.html

Additionally, when checking the NSX API v4.2.0.0, the **transit_subnets** parameter in the body of a `PUT /policy/api/v1/infra/tier-0s/{tier-0-id}` request does refer to **100.64.0.0/16**. **Hence, unless a general mistake has been spread across the documentation, I'm prone to adapt the 100.64.0.0/16 option.**

[Source: https://dp-downloads.broadcom.com/api-content/apis/API_NTDCRA_001/4.2/html/api_includes/method_CreateOrReplaceTier0.html]

5-57 - NSX Edge Node Routing in Active-Active Mode

TYPO —> *With the inter-service router internal BGP (iBGP) routing feature, the service routers (SRs) exchange routing information* ***through iBGP peers*** <— IMHO, the author meant to say **as iBGP peers** (that is, the SRs behave as iBGP peers).

[Source: https://docs.vmware.com/en/VMware-NSX/4.2/administration/GUID-7A0F6D44-4AE8-411F-8D33-2560819699CA.html]

BOTCHED —> *When the Tier-0 gateway is in active-active mode, you can enable inter-service router iBGP routing* ***between the NSX Tier-0 gateways*** <— it's rather **between the** [same, Ed.] **NSX Tier-0 gateway service routers**.

[Source: https://docs.vmware.com/en/VMware-NSX/4.2/administration/GUID-7A0F6D44-4AE8-411F-8D33-2560819699CA.html]

WRONG —> *NSX Manager automatically configures **iBGP peering between the uplink SR components** of the Tier-0 gateway* <— it is actually configured on the T0-GW SR LIF whose type is **internal-routing** (not uplink).
MORE DETAILS —> *This configuration is **transparent** to the administrator* <— yet visible in the routing table and identified with the **isr (Inter-SR)** flag.

5-62 - NSX Edge in a Single-Rack Consolidated Management Domain
BOTCHED —> *North-south traffic takes a maximum of two hops to reach the external physical router. **Traffic passes through the ToR to reach the Tier-1 gateway in the NSX Edge cluster**, which resides in the management cluster. The traffic exits the NSX Edge cluster through the Tier-0 gateway and is sent to the external network through the peered BGP device* <— **the statement takes a specific use-case – when the Tier-1 gateway has been configured for services with an attached Edge Cluster – and makes it look like as if it was a general, always-true, behaviour.** If instead the Tier-1 gateway is configured for distributed routing only, routing between the T1-GW DR and the T0-GW DR components happens locally in the Host Transport Node that runs the source workload, then and only then the traffic heads northbound – through the ToR – to the Edge Node(s) where the active T0-GW SR component(s) is/are located.

5-69 - Application Virtual Networks
BOTCHED / OUT-OF-DATE —> ***Mgmt-RegionA01-VXLAN** is used for VMware Aria Suite workloads that do not require portability between regions: You use this network for VMware Aria Operations for Logs and **VMware Aria Automation Proxy Servers** guided deployment* <— described this way, it may implicitly suggest students that the AVN logical segment names are predetermined, yet their name is whatever has been provided by the administrators in the wizard. Moreover, the segments name suffix **-VXLAN**, mentioned in the slide notes, is a clear leftover from when VCF was using VMware NSX Data Center for vSphere (that has a VXLAN-based overlay).
MORE DETAILS <— the correct mapping between Aria Suite Components and related NSX Segments (using sample segment names) is the following:

Aria Suite component	**AVN Logical Segment**
VMware Aria Operations for Logs	Region-A
VMware Aria Operations	X-Region
Workspace ONE Access	X-Region
VMware Aria Automation	X-Region
VMware Aria Suite Lifecycle	X-Region

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-59E5BEE3-B157-426D-A40C-F21171586863.html#vmware-aria-suite-components-and-nsx-segments-2]

Module 6

6-7 - Workload Domain Types

WRONG —> ***Up to 14 VI workload domains are supported*** <— according to the VMware Configuration Maximums portal:

"Maximum number of Isolated VI Workload Domains per SDDC Manager instance is 24. Maximum number of VI workload domains per SDDC Manager in ELM mode is 14. NOTE: Based on vCenter Server Linked Mode Limitation. Maximum combined limit of Isolated VI workload domains and VI workload domains in ELM mode per SDDC Manager instance is 24".

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=VMware%20Cloud%20Foundation&release=VMware%20Cloud%20Foundation%205.2.0&categories=17-0]

6-8 - About VI Workload Domains

WRONG —> *Principal storage choices include vSAN,* ***NFS*** *v3 and* ***v4.1****, VMFS on FC, or vSphere Virtual Volumes* <— **VCF only supports NFS v3 when used as principal storage.** For more details, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

INCOMPLETE —> ***Supplemental storage choices include NFS v3 and v4.1, VMFS on FC, or vSphere Virtual Volumes*** <— **there are many more options available as supplemental storage.** For more details, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

Better and more complete insights on Principal and Supplemental storage options for VCF can be found by checking the links below.

[Sources:
https://core.vmware.com/blog/principal-and-supplemental-storage-vmware-cloud-foundation-51
https://core.vmware.com/resource/why-vsan-and-vvols-are-best-vmware-cloud-foundation]

6-9 - VI Workload Domain Architecture Models

OUT-OF-DATE —> *You can deploy Workload Management (****vSphere with Tanzu****) in the management domain in VMware Cloud Foundation 4.0.1 and later* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

6-10 - Workload Domain Storage Options

WRONG —> *Principal storage choices include vSAN, **NFS** v3 and **v4.1**, VMFS on FC, or vSphere Virtual Volumes* <— **VCF only supports NFS v3 when used as principal storage.** For more details, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

INCOMPLETE —> ***Supplemental storage choices include NFS v3 and v4.1, VMFS on FC, or vSphere Virtual Volumes*** <— **there are many more options available as supplemental storage.** For more details, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

Better and more complete insights on Principal and Supplemental storage options for VCF can be found by checking the links below.

[Sources:
https://core.vmware.com/blog/principal-and-supplemental-storage-vmware-cloud-foundation-51
https://core.vmware.com/resource/why-vsan-and-vvols-are-best-vmware-cloud-foundation]

WRONG —> ***VMware Cloud Foundation uses the VI workload domain's principal storage as the location for VMs that are required for supporting services on the VI workload domain, such as the NSX Manager cluster, SDDC Manager, and vCenter Server Appliances*** <— whilst it is true for the Management Domain, where all management-related appliances are deployed to the vSAN-based principal storage, **it absolutely wrong and inapplicable to the VI Workload Domains where the storage – principal and/or supplemental – is used for user workloads and related files (i.e.: ISO images).**

TYPO —> *vSAN is the preferred storage type for VMware Cloud Foundation for the advantages **is** provides* <— author for sure meant to say "**it**".

MORE DETAILS —> ***vSAN** simplifies storage provisioning because VMs can be provisioned **using storage policies** that are built into vCenter* <— that is true for every other storage type, provided:

- It does support the vSphere API for Storage Awareness (VASA), such as mid-range and enterprise-class storage arrays do, providing NFS, VMFS and vVOLs datastores.
- vSphere administrators use tag-based VM Storage Policies.

6-11 - Multicluster Support

WRONG —> *You can deploy more than 64 hosts per workload domain when you use multiple clusters in a workload domain. **Each cluster scales to a maximum of 64 nodes**, and you can manage as many VMs and hosts as described in the VMware Configuration Maximums for VMware Cloud Foundation at https://configmax.esp.vmware.com/home* <— **the current configuration maximum for non-vSAN clusters in vSphere 8.x is 96 ESXi hosts per vSphere Cluster** (as it is since vSphere 7.0U1); hence, it should be possible.

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%208.0&categories=2-0]

6-12 - Multiple VI Workload Domains or Multiple Clusters

WRONG —> ***You can add up to 14 VI workload domains. A single cluster in a VI workload domain can scale to 64 hosts*** <— according to the VMware Configuration Maximums portal:

"Maximum number of Isolated VI Workload Domains per SDDC Manager instance is 24. Maximum number of VI workload domains per SDDC Manager in ELM mode is 14. NOTE: Based on vCenter Server Linked Mode Limitation.
Maximum combined limit of Isolated VI workload domains and VI workload domains in ELM mode per SDDC Manager instance is 24".

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=VMware%20Cloud%20Foundation&release=VMware%20Cloud%20Foundation%205.2.0&categories=17-0]

Additionally, **the current configuration maximum for non-vSAN clusters in vSphere 8.x is 96 ESXi hosts per vSphere Cluster** (as it is since vSphere 7.0U1); hence, it should be possible.

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%208.0&categories=2-0]

6-18 - Host Validation

WRONG / OUT-OF-DATE —> ***3. In the Confirm Fingerprint column, select the check mark for each or all hosts*** <— that was the procedure in VCF 4.x. **In VCF 5.x you can confirm individual hosts by clicking the vertical ellipsis besides it and selecting the "Confirm fingerprint" menu option.**

BOTCHED / MISPLACED —> ***5. Remove any hosts for which fingerprints cannot be validated*** <— **IMHO, this step shall be moved after clicking VALIDATE ALL.**

6-24 - Network Pool Sizing

WRONG —> *Each cluster can be expanded to the* ***maximum number of hosts supported by vCenter Server, which is 64 hosts*** <— **the current configuration maximum for non-vSAN clusters in vSphere 8.x is 96 ESXi hosts per vSphere Cluster** (as it is since vSphere 7.0U1); hence, it should be possible.

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%208.0&categories=2-0]

WRONG —> **A vSAN stretched cluster can be expanded to a maximum of 30 hosts** <— current configuration maximum for vSphere 8.x is **40 hosts (20 per site x 2)**.

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%208.0&categories=7-0]

6-33 - Sub-Transport Node Profiles

WRONG GRAPHIC / TYPO <— the TEP VLAN for Sub-Cluster-3 is **VLAN105**, yet this VLAN is already used in Sub-Cluster-2… following the schema for the other two Sub-Clusters, it shall be – say – **VLAN106**.

7-40 - Workload Domain Prerequisites

WRONG —> ***You cannot assign static TEP address pools using VMware Cloud Foundation automated deployment of NSX.*** <— **"If you do not plan to use DHCP, you can use a static IP pool for the NSX host overlay network. The static IP pool is created or selected as part of VI workload domain creation"**.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-4EBF9094-F732-4800-86FF-30C5D88B9758.html]

WRONG —> ***VMware Cloud Foundation creates a DHCP IP address pool in NSX Manager by API during the workload domain creation.*** <— no, it doesn't, it does rather create a **Transport Node Profile** that uses DHCP for **IP Assignment**.

6-40 - Expanding Workload Domains: Adding Clusters

MORE DETAILS —> *You must use SDDC Manager to add clusters or hosts to existing workload domains.* ***You must not add hosts or clusters using the vSphere Client*** *of the vCenter instance managing the workload resources. If you use the vSphere Client to add compute resources,* ***SDDC Manager is unaware of the change*** *and the hosts are not added to the VMware Cloud Foundation inventory.* ***This omission can cause unpredictable and unwanted behaviour*** *later when you attempt VMware Cloud Foundation workflows, such as upgrades, VI workload domain expansion, VI workload domain deletion, and so on* <— generally speaking this is true and is **always the best practice**. Nevertheless, **"if you make changes to your VMware Cloud Foundation configuration using the vSphere client, you can use the sync workflow in the VCF Import Tool to manually update the SDDC Manager with any out-of-band changes applied from vCenter Server".**

For instructions about using the **VCF Import Tool** with the **sync** parameter click the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-D871245E-5156-4741-A6FA-9A6C03B04315.html]

6-48 - Workload Domain Design Considerations

WRONG —> *vCenter design for a VI workload domain: For this design,* ***you determine the number of vCenter instances in the VI workload domain*** *in addition to their size, networking configuration, cluster layout, redundancy, and security configuration.* <— according to the Configuration Maximums portal, the **"Maximum number of vCenters associated to a VI Workload Domain" is 1**.

[Source:https://configmax.esp.vmware.com/guest?vmwareproduct=VMware%20Cloud%20Foundation&release=VMware%20Cloud%20Foundation%205.2.0&categories=73-0]

Additionally, vCenter HA is not supported in VMware Cloud Foundation (also see slide 6-52).
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-CD0E3596-AB45-442D-B9C8-7BE0BEC0892B.html#GUID-B4380728-190C-428E-83AE-3F7667506948__GUID-36F355B6-D679-47EE-9FF4-EFC576D9582C]

6-54 - Software-Defined Networking Design: Shared NSX Manager Instances
BOTCHED —> ***All VI workload domains can use NSX Edge nodes in the management domain.*** <— phrased this way, this sentence can make students believe that VI Workload Domains can use the Management Domain Edge Cluster(s)' Edge nodes, yet **this is not possible as the NSX fabric created for the Management Domain is not shared with the VI Workload Domains, which are paired with one or more additional NSX fabrics** (depending on whether we have VI Workload Domains or Isolated VI Workload Domains, or a mix of both).
The only possible interpretation of the botched sentence is that **all VI Workload Domains can have their Edge nodes deployed into the Management Domain cluster(s)** or – as the following sentence in the slide notes is explaining – "for scalability, SDDC Manager can be used to deploy an NSX Edge cluster in VI workload domain clusters".

6-57 - Shared Storage Design: Principal and Supplemental Storage
INCOMPLETE <— there are more options for both Principal and Supplementary Storage, for a complete and detailed list, check the link below.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

6-58 - Workload Domain Design Decisions: Storage
BOTCHED / NON-SENSE —> ***You choose between a vSAN datastore or a vSAN cluster as the principal storage method:***

- ***vSAN datastore: If using vSAN as the principal storage in the workload domain, you identify the vSAN policies that each workload requires.***
- ***vSAN cluster: If using vSAN as the principal storage in the workload domain, you determine the level of high availability for the workloads in each workload domain cluster:***
    - ***vSphere HA and admission control***
    - ***Maintenance mode selection: Ensure accessibility or full migration***

<— it doesn't make any sense: there is 1x vSAN datastore per each vSphere Cluster (aka vSAN Cluster when vSAN is enabled). Hence, there is no distinction between the two: one corresponds to the other, and both – obviously – offer the same capabilities as they are the same thing.
If the author(s) meant anything different – in example, "vSAN datastore" implicitly referred to a remotely-mounted vSAN datastore using HCI Mesh / vSAN Max – they shall better rephrase this slide.

6-59 - Workload Domain Design Decisions: ESXi

OUT-OF-DATE —> *For **vSphere with Tanzu**, large NSX Edge nodes are required.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

BOTCHED —> *After completing the application profile assessment, you can use the new vSAN ReadyNode sizer at **https://vsansizer.vmware.com.https://vsansizer.vmware.com/*** <— link is repeated twice.

INCOMPLETE —> ***With the tool, you can perform the following tasks*** <— **you can also choose between vSAN ESA and vSAN OSA (which is a quite important decision as it determines the vSAN architecture!).**

Module 7

7-6 VMware Cloud Foundation Storage Types

INCOMPLETE <— there are more options for both Principal and Supplementary Storage (i.e.: NVMe over RDMA), for a complete and detailed list, check the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

7-19 - Failures to Tolerate

MORE DETAILS —> *You can select a Failures to Tolerate (**FTT**) **value of 0, 1, or 2**.* <— despite **vSAN supports 0|1|2|3** failures to tolerate, **VCF allows for 0|1|2 only** configurations in the VI Workload Domain deployment wizard (see a screenshot in slide 7-20).

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-E64CEFDD-DCA2-4D19-B5C5-D8ABE66407B8.html#procedure-12]

7-27 - Management Domain Deduplication and Compression

OUTDATED SCREENSHOT <— despite it shows the **Enable vSAN Deduplication and Compression** option, this is a pre-VCF5 Deployment Parameters workbook.

7-29 - Enabling Deduplication and Compression on an Existing Cluster

INCOMPLETE <— *Enabling deduplication and compression can significantly affect an active vSAN cluster because vSAN must perform the following tasks:*

1. ***Evacuate the disk group.***
2. ***Format the evacuated disk group.***

<— there is a missing step between steps 1 and 2. According to the online vSAN documentation:

While enabling deduplication and compression, vSAN updates the on-disk format of each disk group of the cluster. To accomplish this change, vSAN:

1. **evacuates data from the disk group,**
2. **removes the disk group,**
3. **and recreates it with a new format that supports deduplication and compression.**

[Source: https://docs.vmware.com/en/VMware-vSphere/8.0/vsan-administration/GUID-C0DBD8E2-A3CD-4E71-A0B3-127D48D89AE0.html]

7-35 - Comparing vSAN Architectures

WRONG —> ***vSAN ESA Compression: off by default*** <— it is quite the opposite: in vSAN ESA, **"the new compression option is enabled by default, but can be**

optionally disabled for workloads though the use of the vCenter Storage Policy Based Management (SPBM) Framework".
[Source: https://core.vmware.com/blog/vsan-8-compression-express-storage-architecture]

7-36 - vSAN ESA Requirements
BOTCHED —> ***Minimum of 32 CPUs*** <— **to say it better, it's 32 CPU cores!**
[Source: https://docs.vmware.com/en/VMware-vSphere/8.0/vsan-planning/GUID-07EFD36A-F844-4E7D-830D-3863E4AA617C.html]

7-39 - vSAN ESA New I/O Engine
TYPO IN THE GRAPHIC <— **vSCS(ESXi)** in place of **vSCSI(ESXi)**
BOTCHED GRAPHIC —> **Compresses [...] Encrypts [...] Checksums [...]** <— **whilst the features listed in the slide are correct, the graphic refers to the vSAN Object Manager that is a layer that sits much lower in the stack than the layer – the vSAN Log-Structured Filesystem (LFS) – that is responsible for compression, encryption, checksum, and many other features (see slide 7-40).**
MORE DETAILS <— technical deep dive on vSAN ESA is available at:

- https://core.vmware.com/vsan-esa
- https://vmwareknowledgebase.com/vmware-vsan-esa-architecture-deep-dive/

7-51 - Parameters for vSAN ESA
MORE DETAILS —> *The user has the option of providing the path to the* ***HCL JSON file manually****.* <— **instructions are provided in the KB article linked below.**
[Source: https://knowledge.broadcom.com/external/article/315556/updating-the-vsan-hcl-database-manually.html]

7-60 - Adding a Cluster to a Workload Domain
MORE DETAILS —> ***vSAN Direct*** *is not supported with vSAN ESA.* <— for more details about what vSAN Direct is and can provide, click the link below.
[Source: https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-with-tanzu-services-workloads/GUID-F7223607-30A5-4B2D-9B06-A55A65FEAA11.html]

7-62 Introducing vSAN Max Support
INCOMPLETE —> *In VMware Cloud Foundation 5.2, customers can use* ***vSAN Max as their primary VMware Cloud Foundation storage****.* <— the author forgot to mention a very important "limitation": **vSAN Max cannot be used as a Principal Storage option for the Management Domain**.
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-design/GUID-2156EC66-BBBB-4197-91AD-660315385D2E.html]

7-68 - Host Design Considerations for vSAN: Boot Device

BOTCHED / NON-SENSE —> *You should boot hosts using the Unified Extensible Firmware Interface (UEFI),* ***but*** *hard disks, CD-ROM drives, SD devices, and USB devices are supported.* <— "***but***" is totally misplaced and misleading. **UEFI is just the more modern firmware alternative to the much older BIOS option. The operating system – in our case vSphere ESXi – still needs to be placed somewhere** (i.e. HDDs/SDDs, SDs, USBs, and even ISO images used to begin an interactive ESXi installation).

[Source: https://uefi.org]

To be thorough, **the only use case where a local storage device is not required to accommodate ESXi – or more in general, any other OS – is when PXE booting off the network: ESXi is loaded from the network directly into the host memory.**

[Source:http://web.archive.org/web/20231003211409/http://www.pix.net/software/pxeboot/archive/pxespec.pdf]

[Source: https://tools.ietf.org/html/rfc4578]

MORE DETAILS —> ***When USB and SD devices are used for boot devices****, the logs and traces reside in RAM disks that are not persistent during reboots. USB and SD boot device functionality is supported through vSphere 8, including updates.* <— review the KB article with full options and recommendations by clicking the link below.

[Source: https://knowledge.broadcom.com/external/article/317631/sd-cardusb-boot-device-revised-guidance.html]

7-71 - vSAN Sizing: Levels of Overhead

BOTCHED —> *For more information about designing and sizing a vSAN cluster, see vSAN Planning and Deployment at* ***https://docs.vmware.com/en/VMware-vSphere/8.0/vsan-803-planning-deployment-guide.pdfhttps://docs.vmware.com/en/VMware-vSphere/8.0/vsan-803-planning-deployment-guide.pdf****.* <— the link is repeated twice.

Module 8

8-6 - Availability of Key Infrastructure Components

MORE DETAILS / MISLEADING —> ***BGP peers*** *must be in place and available by the time you deploy an NSX Edge cluster for the management domain.* <— stated this way, it may let students believe that using BGP for routing is the only option. According to the online documentation, though, when you deploy an NSX Edge cluster to provide north-south routing and network services to a workload domain, you can define the **"Tier-0 Routing Type: select Static or EBGP to determine the route distribution mechanism for the tier-0 gateway. If you select Static, you must manually configure the required static routes in NSX Manager. If you select EBGP, VMware Cloud Foundation configures eBGP settings to allow dynamic route distribution."**.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-D17D0274-7764-43BD-8252-D9333CA7415A.html#procedure-1]

8-8 - Restoring SDDC Manager Backups

BOTCHED / MISLEADING —> *Use the vSphere Client to* ***deploy a new SDDC Manager OVA Management cluster****.* <— stated this way, it may let students believe that the SDDC Manager can be clustered, instead the SDDC Manager is a single-node deployment. You should read the above-mentioned statement as follows: "**Use the vSphere Client to deploy a new SDDC Manager OVA onto the VCF Management Domain vSphere cluster**".

GENERIC LINK —> *For more information on* ***restoring SDDC Manager backups****, please see the VMware Cloud Foundation Administering Guide at VMware Docs.* <— this is the direct link to the procedure (see the child pages too): **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-3C8BD7CF-E586-4F62-8D85-DCD10F457696.html**

BROKEN LINK —> *For detailed instructions on recovering an entire VMware Cloud Foundation instance from backup, see VMware Cloud Foundation Instance Recovery Guide at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/services/vcf-recovery/GUID-ACBFFD0F-F8CA-47D9-990B-7B0975BCDF5A.html*** <— use instead the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-3C8BD7CF-E586-4F62-8D85-DCD10F457696.html]

8-10 - Restoring NSX Managers

INCOMPLETE —> ***WHOLE SLIDE*** <— the description is way too vague and doesn't cover the possible use cases. **It gives the wrong impression that there is a "one-size-fits-all" approach and this may lead to unnecessary downtimes!**

There are different procedures that may be required to restore the NSX Management node(s), based on the initial state of the cluster after the failure.

First of all, you need to identify the use case:

- **"If all three NSX Manager nodes in an NSX Manager cluster are in a failed state, you begin the restore process by restoring the first cluster node". This procedure is not applicable in use cases when there are operational NSX Manager cluster nodes.**
    - **"If two of the three NSX Manager nodes in the NSX Manager cluster are in a failed state, you begin the restore process by deactivating the cluster".**
    - **"If only one of the three NSX Manager nodes in the NSX Manager cluster is in a failed state, you directly restore the failed node to the cluster".**

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-FCD71AB6-866B-47EB-AAF0-3F4C61743AFA.html]

Check the detailed procedures in the online documentation navigating to the URL below and reviewing all its children pages.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-85ECE776-F5F2-4ABF-B031-8CC6F982403D.html]

8-11 - Using vSphere HA in the Management Domain

WRONG —> *If VMware Tools heart beating fails for a continuous period of **50 seconds**, the VM is reset to ensure that services are restored as quickly as possible.* <— the vSphere HA advanced parameter **das.iostatsinterval** in both Management and VI Workload Domain clusters in VCF 5.2 is set **equal to zero**, disabling the check. According to VMware's online documentation:

"das.iostatsinterval: Changes the default I/O stats interval for VM Monitoring sensitivity. The default is 120 (seconds). Can be set to any value greater than, or equal to 0. Setting to 0 disables the check. Note: Values of less than 50 are not recommended since smaller values can result in vSphere HA unexpectedly resetting a virtual machine".

[Source: https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-availability/GUID-E0161CB5-BD3F-425F-A7E0-BF83B005FECA.html]

8-21 - Stretched Cluster Requirements

MISLEADING —> ***vSphere vMotion VMkernel ports use the vMotion TCP/IP stack instance by default**, and they can therefore use a different gateway than the default TCP/IP stack.* <— **said this way, the sentence may make the students believe that this is always the case.** Instead, **this configuration is applied by default by VCF through the SDDC Manager workflow configuring the ESXi hosts being commissioned.** This includes setting the vMotion TPC/IP stack default gateway based on the Network Pool assigned to the hosts. vSphere ESXi hosts being configured **outside a VCF environment by default get vMotion VMkernel interfaces assigned to the Default TCP/IP stack.** If you want to use the vMotion stack instead, you need to follow the procedure described in the link below.

[Source: https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-vcenter-esxi-management/GUID-30FAA00F-D5F3-475D-820E-5D45517AC18E.html#network-configuration-6]

8-33 - NSX Routing Considerations for Stretched Clusters

BOTCHED / MISLEADING —> *For example,* ***if a segment with the 192.168.21.0/24 network is connected to the Tier-1 router of the NSX Edge cluster, the NSX Edge cluster advertises that network to BGP peers*** *so that traffic can route in to the 192.168.21.0/24 from external sources.* <— stated this way, students may believe that Tier-1 routers run BGP, but they don't run any dynamic routing protocol, nor they can connect to any physical router. **"A tier-1 gateway has downlink connections to segments and uplink connections to tier-0 gateways. You can configure route advertisements and static routes on a tier-1 gateway. Recursive static routes are supported."**.

[Source: https://docs.vmware.com/en/VMware-NSX/4.2/administration/GUID-6244CFD2-4119-4718-BA52-1BC9682A8C6E.html]

That said, you shall read the above-mentioned sentence as something like "For example, **if a segment with the 192.168.21.0/24 network is connected to the Tier-1 router, then:**

- **first, the Tier-1 router has to advertise it to the Tier-0 router;**
- **next, the Tier-0 router has to redistribute it into the BGP protocol;**
- **last, the network is advertised to the BGP peers by NSX Edge cluster (which is where the Tier-0 Service Router component runs the BGP protocol),** *so that traffic can route in to the 192.168.21.0/24 from external sources.*"

10-38 - Expanding a Stretched Workload Domain Cluster

BOTCHED / GENERIC LINK —> *For more information, see the chapter about* ***expanding stretched workload domain clusters*** *in VMware Cloud Foundation in VMware Cloud Foundation Administration Guide at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html****.* <— the direct link to the documentation pages is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-928542C0-65E1-46F1-A62D-E7A106AC593B.html**.

8-39 - Replacing Failed Hosts in a Stretched Workload Domain Cluster

BOTCHED / GENERIC LINK —> *For more information, see the chapter about* ***replacing a failed host in a stretched cluster*** *in VMware Cloud Foundation Administration Guide at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html****.* <— the direct link to the documentation pages is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-92FD3AEE-5B5F-421C-B722-44AE0314D84E.html**.

Module 9

9-13 - VMware Certificate Authority

MORE DETAILS —> ***Provides signed certificates to*** *vSphere, ESXi hosts, SDDC Manager and* ***solutions like NSX****.* <— **"If VMware Cloud Foundation™ (VCF) was used to deploy NSX, the default NSX API and Cluster certificates get replaced with CA certificates signed by the VMware Certificate Authority (VMCA) from vCenter. The API and Cluster certificates might still display in the certificate list, but are not used. Replace the CA-signed certificates using the procedure in the VCF Administration Guide. After you perform the replacement, your NSX Manager stores in the UI contain the API and Cluster certificates, the VMCA CA certificates, and the signed certificates by the third-party organization. From then on, the NSX Manager uses the signed certificate from your organization."**.

[Source: https://docs.vmware.com/en/VMware-NSX/4.2/administration/GUID-CA4DC685-3013-40F4-930D-A10173F8FA25.html]

MORE DETAILS —> ***Does not integrate with the VMware Cloud Foundation certificate Management.*** <— see step by step instructions for installing third-party certificates by clicking on the link below.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-2A1E7307-84EA-4345-9518-198718E6A8A6.html]

9-20 - Configuring Microsoft CA: Requirements

OUTDATED LINK —> *For more information, see the section on* ***creating and adding a Microsoft CA template*** *in VMware Cloud Foundation Administration Guide at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/5.1/vcf-admin/GUID-194B5856-8B2C-40D3-AEC2-DD4324980C05.html****.* <— the updated link is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-194B5856-8B2C-40D3-AEC2-DD4324980C05.html**.

9-21 - Integrating a Microsoft CA with SDDC Manager

WRONG —> ***Template Name: Enter `certsrv` as the template name.*** <— stated this way, may give students the wrong impression that "`certsrv`", as a template name, is mandatory. Actually, **it can be anything the administrator likes, obviously the name must match whatever has been used in the Microsoft CA.**

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-B83D4273-F280-4978-B8BD-63A283F803A9.html]

9-21 - Workflow: Installing Certificates Using an Integrated CA

BOTCHED / INCOMPLETE —> *To install certificates in SDDC Manager using a supported integrated CA:*

1. Select the resource type whose certificates you want to replace.

2. Click GENERATE SIGNED CERTIFICATES.

3. After the certificates generate, click INSTALL CERTIFICATES.

<— regardless of whether you will be using an Integrated Microsoft CA or an Integrated OpenSSL CA, in the list above the creation of a **Certificate Signing Request (CSR) is missing**.

The correct workflow is:

1. **Select the resource type whose certificates you want to replace.**
2. **Click GENERATE CSRS.**
3. **Click GENERATE SIGNED CERTIFICATES.**
4. **Click INSTALL CERTIFICATES.**

The step-by-step procedures can be found at the following URLs:

- **Microsoft CA: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-1C048081-8BF5-437C-BBED-885183B17392.html**
- **OpenSSL CA: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-C08A3777-B438-4E59-80F4-9AA14D71B0B4.html**

BOTCHED / GENERIC LINK —> For additional details on certificate processes, see VMware Cloud Foundation Administration Guide at **https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html**. <— the direct links to step-by-step procedures – for Integrated Microsoft CAs or an Integrated OpenSSL CAs – are the two URLs highlighted in blue right above.

9-24 - Workflow: Installing Certificates Using an External CA

BOTCHED / GENERIC LINK —> For additional details on certificate processes, see VMware Cloud Foundation Administration Guide at **https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html**. <— the process for External CAs (slide 9-24) and Third-Party CAs (slide 9-25) are basically identical. The direct link is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-2A1E7307-84EA-4345-9518-198718E6A8A6.html**

9-25 - Workflow: Installing Third-Party Certificates

BOTCHED / GENERIC LINK —> For additional details on certificate processes, see VMware Cloud Foundation Administration Guide at **https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html**. <— the process for External CAs (slide 9-24) and Third-Party CAs (slide 9-25) are basically identical. The direct link is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-2A1E7307-84EA-4345-9518-198718E6A8A6.html**

Module 10

10-11 - About VMware Aria Suite Lifecycle

INCOMPLETE [text and graphic] —> *You use **VMware Aria Suite Lifecycle to deploy and automatically update** the following **products*** <— VMware Aria Suite Lifecycle can deploy and update **VMware Aria Operations for Networks** too.

[Sources:

- Deploy: https://docs.vmware.com/en/VMware-Aria-Suite-Lifecycle/8.18/lifecycle-install-upgrade-manage/GUID-6CD013DE-A756-4C32-BF03-6731AD0220F2.html
- Update: https://docs.vmware.com/en/VMware-Aria-Suite-Lifecycle/8.18/lifecycle-install-upgrade-manage/GUID-DD5E829A-9E03-4FF7-A8F8-EF5D026520C7.html]

10-27 - Creating vSphere Lifecycle Manager Images

TYPO —> *You can create an image on the vCenter instance for the management domain or a VI workload domain, or **for** an external vCenter instance.* <— IMHO, the author meant to say "**on**".

Module 11

11-9 - VMware Aria Suite Integrations with VMware Cloud Foundation

INCOMPLETE —> *With the VMware Aria Suite Lifecycle appliance deployed, a customer can choose to* ***deploy the VMware Aria Suite products from VMware Aria Suite Lifecycle****.* <— VMware Aria Suite Lifecycle can deploy **VMware Aria Operations for Networks** too.

[Source: https://docs.vmware.com/en/VMware-Aria-Suite-Lifecycle/8.18/lifecycle-install-upgrade-manage/GUID-6CD013DE-A756-4C32-BF03-6731AD0220F2.html]

11-12 - AVN Configuration: Overlay and VLAN Backed

WRONG —> ***For the NSX Edge cluster to be deployed, an upstream router must be configured as a border gateway protocol (BGP) peer to support dynamic routing of the overlay networks. When using overlay networks for dynamic routing, a set of upstream routers must be configured for BGP peering.*** <— according to the VMware online documentation, when deploying an Edge Cluster, you can **"select Static or EBGP to determine the route distribution mechanism for the tier-0 gateway. If you select Static, you must manually configure the required static routes in NSX Manager. If you select EBGP, VMware Cloud Foundation configures eBGP settings to allow dynamic route distribution"**.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.3/vcf-admin/GUID-D17D0274-7764-43BD-8252-D9333CA7415A.html]

BOTCHED —> *For further information on AVN network configuration options, see VMware Cloud Foundation Administration Guide at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html.***
https://docs.vmware.com/en/VMware-Cloud-Foundation/index.htmlhttps://docs.vmware.com/en/VMware-Cloud-Foundation/index.html <— the link is repeated three times.

11-13 - VMware Aria Suite Networking

WRONG —> ***Region-specific network:***
— VMware Aria Operations for Logs
— VMware Workspace ONE Access
— VMware Aria Operations Cloud Proxies <— **VMware Workspace ONE Access is not region-specific and – if a multi-region architecture is used – it must be able to fail over the other region; hence, it gets connected to the cross-region network**.

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.3/vcf-admin/GUID-59E5BEE3-B157-426D-A40C-F21171586863.html#vrealize-suite-components-and-nsx-segments-2]

11-16 - Deployment Workflow

WRONG —> ***One or more NSX Edge nodes*** *are deployed with two-tier routing and network services* <— according to VMware's online documentation:

"A minimum of two NSX Edge nodes is required. NSX Edge cluster creation allows up to 8 NSX Edge nodes if the Tier-0 Service High Availability is Active-Active and two NSX Edge nodes per NSX Edge cluster if the Tier-0 Service High Availability is Active-Standby".

[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-D17D0274-7764-43BD-8252-D9333CA7415A.html]

GENERIC LINK —> *For more information on setting up VMware Aria Suite Lifecycle for VMware Cloud Foundation, see VMware Cloud Foundation Administration Guide at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html****.* <— the direct link to the subject is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-99799810-F752-48BD-A398-180C5E64A581.html**.

11-31 - Deploying Workspace ONE Access with VMware Aria Suite Lifecycle

OUTDATED LINK —> The complete deployment steps are available in the VMware Cloud Foundation admin guide. **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.1/vcf-admin/GUID-249206AC-44B6-4E9A-9A72-8A78CC38716F.html** <— the link to the documentation for the current version is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-249206AC-44B6-4E9A-9A72-8A78CC38716F.html**

11-32 - About VMware Aria Suite Lifecycle Flexible Upgrades

INCOMPLETE —> *From VMware Cloud Foundation 5.0, VMware Aria Suite Lifecycle 8.10 is then used to deploy, update, and patch* ***VMware Aria Suite components, which include VMware Aria Operations for Logs, VMware Aria Operations, and VMware Aria Automation****.* <— VMware Aria Suite Lifecycle can deploy **VMware Aria Operations for Networks** too.

[Source: https://docs.vmware.com/en/VMware-Aria-Suite-Lifecycle/8.18/lifecycle-install-upgrade-manage/GUID-6CD013DE-A756-4C32-BF03-6731AD0220F2.html]

11-48 - Deploying VMware Aria Operations for Logs on VMware Cloud Foundation

INCOMPLETE —> *VMware Aria Suite Lifecycle deploys a three-node cluster for* ***VMware Aria Operations****.* <— the product being discussed is **VMware Aria Operations for Logs**.

Module 12

12-2 - Importance

MORE DETAILS —> *VMware Private AI Foundation with NVIDIA is a joint Generative AI platform that enables enterprises to run **RAG workflows**, fine-tune and customize large language models (LLMs), and run inference workloads in their data centers to address privacy, choice, cost, performance, and compliance concerns.* <— **"Retrieval-augmented generation (RAG) is a technique for enhancing the accuracy and reliability of generative AI models with facts fetched from external sources".**

[Source: https://blogs.nvidia.com/blog/what-is-retrieval-augmented-generation/]

12-7 - Large Language Model Operational Components

TYPO —> *These models consist of multiple layers of self-attention mechanisms and feed-forward neural networks with billions **or** neurons and parameters that must be trained over terabytes of data.* <— IMHO, the author meant to say "**of**".

12-9 - GPU Virtualization Using NVIDIA GPUs

WRONG —> *A host VIB is installed from the NVIDIA AI Enterprise Suite, allowing GPUs to be shared or partitioned into several GPU instances for running multiple VM and **Tanzu Kubernetes Grid** workloads. VM or **Tanzu Kubernetes Grid** workloads can then be configured to consume GPU resources from a selection of preconfigured profiles.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

12-12 - Shared GPU Modes

WRONG —> *NVIDIA GPUs can be shared with traditional VMs or **Tanzu** worker node VMs in Time-Slicing mode or Multi-Instance GPU (MIG) mode.* <AND> *Manual and automated vMotion (maintenance mode) for traditional VM-based workloads (does not apply to **Tanzu Kubernetes Grid (TKG)** worker node VMs)* <AND> *vSphere HA for VM workloads and non-persistent **TKC** workloads* <AND> *NVIDIA vGPU support of a VM-based workload includes support for traditional vSphere VMs and **Tanzu** worker node VMs. **Tanzu** worker nodes are virtual machines which run Kubernetes container workloads.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

12-13 – Time-Slicing Versus MIG Mode Considerations

WRONG —> *A VM workload can be an ordinary virtual machine or a **Tanzu Kubernetes Grid** worker node VM.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

12-14 - NVIDIA GPU Configuration for VMware Cloud Foundation

WRONG —> *Select the preconfigured NVIDIA vGPU profile to be used with the VM or **Tanzu Kubernetes Grid** workload that requires the vGPU (reboot required).* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

BOTCHED —> *For more in-depth configuration details of a GPU-enabled cluster in VMware Cloud Foundation, see the three-part blog series "Configuring NVIDIA GPU devices for Tanzu Kubernetes Clusters in VMware Cloud Foundation" at **https://blogs.vmware.com/cloud-foundation/2022/04/18/configuring-nvidia-gpu-devices-for-tanzu-kubernetes-clusters-in-vmware-cloud-foundation-part-i/.https://blogs.vmware.com/cloud-foundation/2022/04/18/configuring-nvidia-gpu-devices-for-tanzu-kubernetes-clusters-in-vmware-cloud-foundation-part-i/*** <— the link is repeated twice.

12-17 - Creating VM Class for a TKG Worker Node VM

WRONG —> *To configure a **Tanzu Kubernetes Grid** worker node VM with a GPU, you must create a VM Class.* <AND> *To configure a **Tanzu Kubernetes Grid** worker node VM with an NVIDIA A30, A100, H100, or L40S GPU:* <AND> *The host is then commissioned to a workload domain Cluster configured for Workload Management (**Tanzu**).* <AND> *To configure for **Tanzu**, create a VM class and assign a GPU PCIe as a resource for the data scientist or developer to consume.* <AND> *Commission the host to a workload domain configured for Workload Management (**Tanzu**).* <AND> *When configuring a **Tanzu Kubernetes Grid** worker node VM with a shared or partitioned GPU instance, select the NVIDIA Grid vGPU option.* <AND> *When creating a VM class, the **Tanzu Kubernetes Grid** worker node VM can be configured for either Time-Slicing mode or MIG through the vSphere Client.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

12-18 - NVIDIA GPUDirect RDMA

BOTCHED —> *See Developing a Linux Kernel Module using GPUDirect RDMA at **https://docs.nvidia.com/cuda/gpudirect-rdma/index.html. https://docs.nvidia.com/cuda/gpudirect-rdma/index.html*** <— the link is repeated twice.

12-20 - NVIDIA NVLink

BOTCHED —> *For more information, see the NVIDIA NVLink and NVSwitch page at* ***https://www.nvidia.com/en-au/data-center/nvlink/https://www.nvidia.com/en-us/data-center/nvlink/****.* <— the link is repeated twice.

12-23 Private AI Foundation with NVIDIA Use Cases

WRONG —> *Cloud administrators can provide DevOps engineers with a Private AI Foundation with NVIDIA environment for provisioning production-ready AI workloads on* ***Tanzu Kubernetes Grid*** *clusters on* ***vSphere with Tanzu****.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

12-24 Private AI Foundation with NVIDIA Components

WRONG [graphic & text] —> *Content Library for* ***Tanzu Kubernetes*** *Releases* <AND> ***Tanzu Kubernetes Grid*** *Cluster* <AND> *Supervisor: One or more vSphere clusters enabled for* ***vSphere with Tanzu*** *so that you can use the Kubernetes API to run VMs and containers on vSphere* <AND> *Deployed on a* ***Tanzu Kubernetes Grid*** *cluster* <AND> *You use the NVIDIA Licensing Portal to generate a client configuration token to assign a license to the guest vGPU driver in the deep learning virtual machine and the GPU Operators on* ***Tanzu Kubernetes Grid*** *clusters.* <AND> *Content libraries store the images for the deep learning VMs and for the* ***Tanzu Kubernetes Grid*** *releases.* <— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

12-26 Life Cycle Management for VMware Cloud Foundation GPU-Enabled Workloads

WRONG —> *In a VMware Cloud Foundation* ***Tanzu Kubernetes Grid*** *cluster:*

• GPU-enabled ***Tanzu Kubernetes Grid*** *VMs must be manually powered off before vSphere Lifecycle Manager operations.*

• Re-instantiate a ***Tanzu Kubernetes Grid*** *worker node VM on another host.*

<— starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.

[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

OUT-OF-DATE —> *GPU-enabled hosts are not supported in the management domain as* ***the management domain is configured using vSphere Lifecycle Manager baselines****.* <— according to the VMware's online documentation:

"Enable vLCM Cluster Image: Select Yes to use vSphere Lifecycle Manager images for managing the lifecycle of ESXi hosts in the primary cluster of

management domain. VMware Cloud Builder extracts a vSphere Lifecycle Manager image from the first ESXi host and applies that image to all the hosts in the cluster. The vSphere Lifecycle Manager image is also imported into SDDC Manager (available at Lifecycle Management > Image Management.
Note: vSAN Express Storage Architecture (ESA) requires vSphere Lifecycle Manager images.
Select No to use vSphere Lifecycle Manager baselines for managing the lifecycle of ESXi hosts in the primary cluster of management domain".
[Source: https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-deploy/GUID-89E3C801-FA40-42F1-A7F1-C1B6859589E2.html]
WRONG —> *Life cycle management considerations for **Tanzu Kubernetes Grid** worker node VMs running K8s containers configured with NVIDIA vGPU:* <AND> *Before commencing VMware Cloud Foundation life cycle management, patching, and upgrade operations, **Tanzu Kubernetes Grid** workloads must be manually powered off.* <AND> *If a VMware Cloud Foundation operator does not manually power off a GPU-enabled **Tanzu Kubernetes Grid** workload, the VMware Cloud Foundation life cycle management operation will fail.* <AND> *Assuming GPU resources are available, **Tanzu Kubernetes Grid** workloads can simply be re-instantiated on a different host.* <AND> *vSphere and VMware Cloud Foundation do not support vSphere vMotion of GPU-enabled **Tanzu Kubernetes Grid** worker node VMs configured with persistent storage to another host.* <—
starting with vSphere 8 Update 3, **vSphere with Tanzu** has been renamed as **vSphere IaaS Control Plane**.
[Source: https://core.vmware.com/resource/whats-new-vsphere-update-3-vsphere-iaas-control-plane]

Module 13

13-7 - SoS Tool: Performing Health Checks

BOTCHED / GENERIC LINK —> *For more information about the options available with the SoS tool, see the section about* ***SoS utility options*** *in the* ***VMware Cloud Foundation Operations and Administration Guide*** *at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html****.* <— the direct link to the subject is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-9EBC6D42-B799-4177-9EFF-78E98FDBA0FD.html**

13-9 - SoS Tool: Collecting Log Files

BOTCHED / GENERIC LINK —> *For more information about log bundle collection, see the section about* ***collecting logs*** *for your VMware Cloud Foundation system in the* ***VMware Cloud Foundation Operations and Administration Guide*** *at* ***https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html****.* <— the direct link to the subject is **https://docs.vmware.com/en/VMware-Cloud-Foundation/5.2/vcf-admin/GUID-77C40307-8466-4331-A510-9334C0F6CF32.html**