NSXICM4 Notes & Errata

Part Number EDU-EN-NSXICM4-LECT (28-NOV-2022)

Version	Date	Author(s)	Description of Change
01	September 20, 2023	Fabrizio de Luca	Initial Release
02	October 3, 2023	Fabrizio de Luca	Added an entry for slide 4-21

Module 2

2-9 - NSX Features (1)

MORE DETAILS —> *DPU-Based Acceleration* <— feature announcement can be found at https://blogs.vmware.com/networkvirtualization/2022/08/announcing-dpu-based-acceleration-for-nsx.html/. Solution page is located at https://www.vmware.com/solutions/smartnic-dpu-based-acceleration.html

2-12 - High-Level Architecture of NSX

WRONG —> *Consumption plane: Although the consumption plane is not part of NSX, it provides integration into any CMP through the REST API and integration with VMware cloud management planes, such as* ***vRealize Automation*** <— the product has been renamed as **Aria Automation**.

2-28 - Data Plane Components

BOTCHED / INCOMPLETE —> *Provides stateful and gateway services* <— the author probably meant to say: "**Provides stateful and stateless gateway services**".

Module 3

3-6 - Implementing NSX in vSphere

BOTCHED / INCOMPLETE DESCRIPTION —> ***Preconfigure transport nodes****, including transport zones, IP pools, and uplink profiles* <— this step is about **configuring transport nodes prerequisites**, not the transport nodes themselves.

TYPO —> *Deploy the NSX Edge* ***node*** *(VM or bare metal)* <— despite you can deploy just a single Edge node, this is not recommended – especially in a production environment – therefore the sentence shall refer to multiple **nodes**.

3-19 - Management Cluster Status: GUI (2)

WRONG —> *On the* ***Overview*** *page in the NSX UI* <— this is actually the **System > Appliances > NSX Manager** page.

3-28 - NSX Manager Policy and Manager Views

BOTCHED / MISLEADING —> **Use the Manager mode UI** in the following instances: ***When NSX is integrated with*** *cloud management platforms, for example,* ***VMware Aria Automation (former name: vRealize Automation)*** […] <— despite being still possible using the Manager API when connecting VMware Aria Automation to NSX (there is a radio button option in the NSX Cloud Account settings in Aria Automation Assembler, formerly vRealize Automation Cloud Assembly), **the default option is to use the NSX Policy API**.

3-34 - About the System Tab

WRONG —> *The* ***Overview*** *page shows the number and details of the management nodes and the cluster* <— this is the **System > Appliances > NSX Manager** page.

3-45 - About IP Address Pools

BOTCHED / IMPROVABLE —> ***Each transport node has a tunnel endpoint (TEP)*** <— the author should have better explained – as on the other hand stated in a later sentence – that each transport node has **one or more tunnel endpoints (TEPs)**; alternatively, they should have stated that each transport node has **at least one tunnel endpoint (TEP)**.

WRONG <— in the graphic, within the **Set Subnets** dialog box screenshot in the lower right corner of the slide, the **ADD SUBNET > IP Block** option is highlighted, yet configuring an IP pool for the TEPs requires one or more **IP Ranges** to be configured.

BOTCHED / IMPROVABLE —> ***Each transport node has a TEP*** <— same as in the first note for this slide.

3-47 - About Transport Zones (2)

BOTCHED / MISLEADING —> ***A transport node can belong to multiple transport zones: one overlay transport zone and multiple VLAN transport zones*** <— the "one overlay

transport zone" limit **only applies to NSX Edge transport nodes**: **"an NSX Edge can belong to one overlay transport zone and multiple VLAN transport zones".**
[Source: https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-53295329-F02F-44D7-A6E0-2E3A9FAE6CF9.html]

3-49 - About VDS

MORE DETAILS —> *The VDS MTU size must be set to 1,600 bytes or greater from the vSphere Client to utilize it with NSX* <— **"Jumbo Frame Support – A minimum required MTU is 1600. However, MTU of 1700 bytes is recommended to address the whole possibility of a variety of functions and future proof the environment for an expanding Geneve header. As the recommended MTU for the NSX is 9000, the underlay network should support this value. This value should be set in the VDS MTU in vSphere and in the Tunnel Endpoint MTU setting in NSX".**
[Source: https://nsx.techzone.vmware.com/resource/nsx-reference-design-guide#a-72-physical-infrastructure-of-the-data-center-A]

3-73 - Next-Generation Infrastructure with DPUs

MORE DETAILS <— DPU-based Acceleration support has been added as follows:

- In **NSX 4.0.1.1**:
  - **"Accelerated networking,**
  - **security (Tech Preview),**
  - **enhanced visibility and compute resources savings".**

[Source: https://docs.vmware.com/en/VMware-NSX/4.0.1.1/rn/vmware-nsx-4011-release-notes/index.html]

- In **NSX 4.1.0**:
  - **"UPT V2 is production ready for NVIDIA Bluefield-2.**
  - **NSX Distributed Firewall (Stateful L2 and L3 firewall)** is available for production deployment with DPU acceleration.
  - **NSX Distributed IDS/IPS (Tech Preview)".**

[Source: https://docs.vmware.com/en/VMware-NSX/4.1.0/rn/vmware-nsx-410-release-notes/index.html]

- In **NSX 4.1.1:**
  - **"NVIDIA BlueField-2 (100Gbps) is supported".**

[Source: https://docs.vmware.com/en/VMware-NSX/4.1.1/rn/vmware-nsx-411-release-notes/index.html]

MORE DETAILS <— the installation procedure is described at https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-F3D2FA09-FE91-4E6A-B484-2EF0DD5BB660.html

Module 4

4-13 - Geneve Header Format

WRONG —> ***Adds an 8-byte UDP header*** <— this was true with the VXLAN protocol used by NSX Data Center for vSphere (now EOL), but with GENEVE – as much better described in the slide notes – **the minimum total header size is 8 bytes, and the maximum is 260 bytes** due to the implementation of the Variable Length Options.

MORE DETAILS —> ***You must use an MTU of 1600 to account for the encapsulation header*** <— **"Jumbo Frame Support – A minimum required MTU is 1600. However, MTU of 1700 bytes is recommended to address the whole possibility of a variety of functions and future proof the environment for an expanding Geneve header. As the recommended MTU for the NSX is 9000, the underlay network should support this value. This value should be set in the VDS MTU in vSphere and in the Tunnel Endpoint MTU setting in NSX".**

[Source: https://nsx.techzone.vmware.com/resource/nsx-reference-design-guide#a-72-physical-infrastructure-of-the-data-center-A]

4-21 - Viewing Configured Segments

WRONG SCREENSHOT <— at this stage, we're still discussing logical segments and haven't introduced logical gateways yet. Hence, the screenshot showing **Web-Segment, App-Segment, and DB-Segment, attached to Tier-1 Gateway "T1-GW-01"** are out-of-context here and belong to the later **Module 5 – NSX Logical Routing**.

4-42 - QoS Segment Profile

BOTCHED GRAPHIC <— when highlighting the **Average Bandwidth**, **Peak Bandwidth** and **Burst Size** fields for **Ingress**, **Ingress Broadcast** and **Egress** network traffic shaping, the author has misplaced the blue and grey rectangles. Elements should be placed as follows:

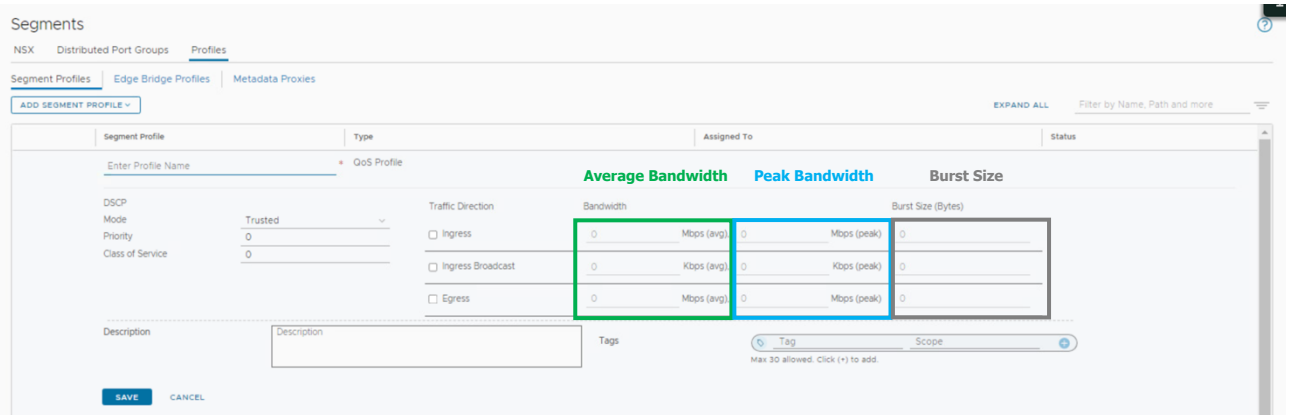

Module 5

5-12 - Edge Nodes and Edge Clusters

BOTCHED / INCOMPLETE —> ***Required to host the Tier-0 gateways*** <— to be more precise, they're **required for hosting the Tier-0 gateways uplink interfaces**.

BOTCHED / INCOMPLETE —> ***Run gateways with centralized and stateful services such as NAT or VPN*** <— services running on Edge gateways can be stateless too, therefore the sentence shall read something like **Run gateways with centralized and stateful services such as NAT or VPN, and stateless services such as BGP, OSPF or Reflexive NAT**.

5-25 - About the NSX Edge Cluster

MORE DETAILS —> *Failure domains can be defined within an edge cluster* <— step by step instructions using the NSX API can be found at https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-FABBBD3C-0928-4E7B-BC30-F04070A76517.html

5-31 - Deploying the NSX Edge Node VM with a Single N-VDS

BOTCHED / MISLEADING —> ***Two TEPs are configured to provide load balancing for the overlay traffic*** <— this statement may incorrectly lead students to believe that, despite they can have up to 4 **`fp-eth#`** interfaces (since NSX-T Data Center 3.2.1, see slide 5-29 comments) in the Edge Node VMs, only 2 TEPs can be configured. Actually, as correctly stated in the slide notes, you **`"use the predefined uplink profile, nsx-edge-multiple-vteps-uplink-profile, to configure the edge node VM with two TEPs"`** (that is, this is a default uplink profile that has been preconfigured by VMware and made available in the product, not the configuration maximum!). Yet, please note that: **`"TEPs are mapped to each of the uplinks configured in the uplink profile for the overlay traffic"`**, therefore if you configure – say – 4 uplinks in a custom uplink profile, you get 4 TEPs in the Edge Node VM. **Hence, the slide graphic simply represents the most common configuration where the ESXi hosts running the Edge Node VMs have two pNICs (typically 10GbE or more)**.

5-40 - Installing NSX Edge on Bare Metal

BOTCHED / INCOMPLETE —> *By default, the following credentials are used: • Root login password: vmware* • ***Password: default*** <— the author forgot some text when copying and pasting from the online documentation. Regarding the second set of credentials (here highlighted in red), the correct statement is "**Admin login password: default**".

[Source: https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-57220956-2AA2-4D79-B0C4-DAC46C28EF69.html]

5-62 - Enabling Route Advertisement in the Tier-1 Gateway

BOTCHED / WRONG TERMINOLOGY —> *Using route advertisement ensures that the networks defined for tenant segments are available for the connected Tier-0 gateway, which can* ***advertise*** *them with the preferred dynamic routing protocol* <— the correct term here is **redistribute**.

5-63 - Configuring Route Redistribution on the Tier-0 Gateway

MORE DETAILS —> *Redistribution into OSPF happens with the following considerations:*

- *Routes are redistributed as OSPF E2 route type (N2 in NSSA areas).*
- *Redistribution of OSPF E1 routes (N1 in NSSA areas) is not supported.*
- *The OSPF cost of the redistributed routes is always 20.*

<— **"These route types are all external routes. In other words, routes we redistributed into OSPF. Here is the difference between type1 and type2:**

- **Type1 routes add the cost of each OSPF link to the route.**
- **Type2 routes only use the redistributed cost and don't add the cost of OSPF links. In other words, the cost remains the same throughout the OSPF domain".**

[Source: https://notes.networklessons.com/ospf-e1-e2-n1-n2-routes]

In the above example, then, Type2 routes (E2 and N2) will always have OSPF cost 20.

5-95 - Active-Active HA Mode

MORE DETAILS —> *NSX version 4.0.1 adds support for stateful services such as NAT 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]

Module 7

7-43 - Enable or Disable Malicious IP Feeds

MORE DETAILS —> *You can enable or disable the download of Malicious IP Feeds from **NTICS** in the NSX UI* <— "**The NSX Threat Intelligence Cloud Service (NTICS) is a cloud-based service, which provides threat feeds and intelligence for on-premises NSX deployments to connect with and pull data to update NSX security features. When enabled via an on-premises NSX environment, NTICS can be used by a customer to authenticate NSX deployments and provide information used to detect a potential security threat**".

[Source: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/docs/vmw-datasheet-nsx.pdf]

7-47 - Configuring Rules to Block Malicious Ips

BOTCHED / MISLEADING —> ***The configured rule drops any traffic from any internal source going to any destination IP address that is included in the group named Custom Malicious IP Group*** <— in the slide graphic, the **Custom Malicious IP Block** policy is added to the **Application** section, yet – in the **Infrastructure** section – there is still the **Default Malicious IP Block Rules** policy containing two default rules that may be blocking (by default in a greenfield deployment) all traffic to and from known malicious IP addresses. If **Default Malicious IP Block Rules** are enabled, then the custom policy in the **Application** section depicted in the slide graphic will never be evaluated.

7-77 - Configuring Gateway Firewall Policy Settings

WRONG —> ***By default, the firewall does not enforce the need to see a three-way handshake and will pick up sessions that are already established. TCP Strict can be enabled per section to turn off midsession pickup*** <— **TCP Strict is enabled by default on any policy created in any Tier-0 or Tier-1 Gateway**.

[Source: empiric observation]

Module 8

8-13 - IDS/IPS Signature Curation

MORE DETAILS —> *The NSX IDS curator engine combines the IDS signatures from Trustwave, Secureworks, and **Lastline** into a single signature set, which it pushes, as an NSX IDS bundle, to NSX Threat Intelligence Cloud* <— Lastline was a pioneer in anti-malware research and AI-powered network detection and response, and on June 18, 2020 VMware closed its acquisition.
[Source: https://blogs.vmware.com/security/2020/06/lastline.html]
MORE DETAILS —> ***MITRE ATT&CK** tactics and techniques* <— **ATT&CK is largely a knowledge base of adversarial techniques** — a breakdown and classification of offensively oriented actions that can be used against particular platforms, such as Windows. Unlike prior work in this area, **the focus isn't on the tools and malware that adversaries use but on how they interact with systems during an operation**.
ATT&CK organizes these techniques into a set of tactics to help explain to provide context for the technique. Each technique **includes information that's relevant to both a red team or penetration tester for understanding the nature of how a technique works and also to a defender for understanding the context surrounding events or artifacts generated by a technique in use**.
Tactics represent the "why" of an ATT&CK technique. The tactic is the adversary's tactical objective for performing an action. Tactics serve as useful contextual categories for individual techniques and cover standard, higher-level notations for things adversaries do during an operation, such as persist, discover information, move laterally, execute files, and exfiltrate data.
Techniques represent "how" an adversary achieves a tactical objective by performing an action. For example, an adversary may dump credentials to gain access to useful credentials within a network that can be used later for lateral movement. Techniques may also represent "what" an adversary gains by performing an action. This is a useful distinction for the Discovery tactic as the techniques highlight what type of information an adversary is after with a particular action. There may be many ways, or techniques, to achieve tactical objectives, so there are multiple techniques in each tactic category.
[Source: https://medium.com/mitre-attack/att-ck-101-17074d3bc62]
[Source: https://attack.mitre.org]

8-27 - Configuring North-South IDS/IPS Rules

TYPO —> *You create IDS policies and rules, which are shared between multiple gateways, by navigating to **Security > Policy Management > IDS/IPS & Malware Prevention > All Shared Specific Rules*** <— the actual tab label is **All Shared Rules**.

8-37 - NSX Application Platform Deployment (1)

MORE DETAILS —> *The **Helm Repository and Docker Registry URLs** must point to the registry and repository **hosted by VMware** or to your private Harbor instance* <— **"By default, the Helm Repository text box has the oci://projects.registry.vmware.com/nsx_application_platform/helm-charts value. This value is the public VMware-hosted repository from which the system obtains the packaged NSX Application Platform Helm chart. The Docker Registry path has the projects.registry.vmware.com/nsx_application_platform/clustering value. This value is the public VMware-hosted registry location from which the system obtains the NSX Application Platform docker images".**

[Source: https://docs.vmware.com/en/VMware-NSX/4.0/nsx-application-platform/GUID-A299BF64-C378-4B1F-8FD1-FBDFB10C2D56.html]

8-38 - NSX Application Platform Deployment (2)

MORE DETAILS —> *Interface Service Name (FQDN) & Messaging Service Name* <— a slightly better detailed description can be found on the on-line documentation at https://docs.vmware.com/en/VMware-NSX/4.0/nsx-application-platform/GUID-D54C1B87-8EF3-45B3-AB27-EFE90A154DD3.html

8-56 - ESXi Host Components

MORE DETAILS —> *Beginning with NSX 4.0.1, east-west malware prevention is also supported for Linux VMs. The Guest Introspection thin agent for Linux is available as part of the operating system specific packages (OSPs). The packages are hosted on the VMware packages portal* <— instructions to download and install the packages can be found at: https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-4871C429-CFE6-41C9-86C9-7FCFE9C95EC8.html

8-60 - Setting the Cloud Region

MORE DETAILS —> *The FQDN for the United States cloud is **nsx.west.us.lastline.com**, and the FQDN for the European cloud is **nsx.nl.emea.lastline.com*** <— this is the **NSX Advanced Threat Prevention Cloud** service that results from the acquisition VMware made in 2020 of Lastline (see details and link to the press release on note 8-13).

8-62 - Service Registration

MORE DETAILS —> *Before you can use malware prevention on the transport nodes, you must register the malware prevention service and deploy SVMs on each host* <— the slide doesn't provide any detail about where to find the SVM and the meaning of the **"ovf_url": "<OVF_PATH>"** parameter in the body of the REST API call that you need to use to register the service. The **NSX SVM Appliance** .ova file can be downloaded from https://customerconnect.vmware.com/downloads/info/slug/networking_security/vmware_nsx/4_x

then extract and store all its files on a web server that must be accessible to NSX Manager, all ESXi hosts where you plan to deploy the NSX Malware Prevention SVM, and the VMware vCenter that is registered to NSX.

To download the SVM follow the on-line instructions at:
https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-0A8BF7D8-9C2E-48A5-8219-17C00F1EC13A.html?hWord=N4IghgNiBcIOIFcCmBnALgAgJIDs0CcB7FAByQGM0BLQnDNACyrrAHMk8MAzQ-DAGWYIAHiAC+QA#download-the-ova-file-of-nsx-malware-prevention-service-virtual-machine-8

To configure the body of the registration REST call see:
https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-0A8BF7D8-9C2E-48A5-8219-17C00F1EC13A.html?hWord=N4IghgNiBcIOIFcCmBnALgAgJIDs0CcB7FAByQGM0BLQnDNACyrrAHMk8MAzQ-DAGWYIAHiAC+QA#register-the-nsx-distributed-malware-prevention-service-9

8-64 - Service Deployments

BOTCHED / INCOMPLETE <— **the slide doesn't explain the purpose of the mandatory RSA public key of the SVM**. Looking at the on-line documentation, you can read the following: **"To download log file from the SVM for troubleshooting purposes, read-only SSH access to the NSX Malware Prevention SVM is required. SSH access to the admin user of the SVM is key-based (public-private key pair). A public key is needed when you are deploying the service on an ESXi host cluster, and a private key is needed when you want to start an SSH session to the SVM"**.

[Source: https://docs.vmware.com/en/VMware-NSX/4.0/administration/GUID-0A8BF7D8-9C2E-48A5-8219-17C00F1EC13A.html#generate-publicprivate-key-pair-for-ssh-access-to-svm-3]

8-65 - Service Deployment Validation from the NSX UI

MORE DETAILS —> ***NTICS*** *Reputation Service Unreachable* <— **"The NSX Threat Intelligence Cloud Service (NTICS) is a cloud-based service, which provides threat feeds and intelligence for on-premises NSX deployments to connect with and pull data to update NSX security features. When enabled via an on-premises NSX environment, NTICS can be used by a customer to authenticate NSX deployments and provide information used to detect a potential security threat"**.

[Source: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/docs/vmw-datasheet-nsx.pdf]

8-89 - NSX Intelligence Activation

MORE DETAILS —> *For more information about upgrading from an earlier version of NSX Intelligence to NSX Intelligence 4.0.1 or later […]* <— upgrading the **NSX Application Platform** and its related **NSX features** (i.e., **NSX Intelligence**) requires deploying the **Upgrade**

Coordinator pod to the Tanzu Kubernetes Grid cluster or upstream Kubernetes cluster where the NSX Application Platform is running and then use the NSX UI to drive the upgrade process. [Source: https://docs.vmware.com/en/VMware-NSX/4.0/nsx-application-platform/GUID-A6C69D8C-ED13-48E4-B323-DF8D971EE5FE.html]

Module 9

9-47 - NSX Advanced Load Balancer Feature Edition Comparison (1)

BROKEN LINK —> **https://avinetworks.com/docs/22.1/nsx-license-editions/** <— current editions comparison can be found at **https://docs.vmware.com/en/VMware-NSX-Advanced-Load-Balancer/22.1/Administration_Guide/GUID-B5EC8F3B-A75E-4809-A653-6EBE08CFED81.html**

9-103 - L2 VPN Considerations

WRONG —> An L2 VPN session can extend up to **4094** L2 segments <— in the .PPT slides used by the VMware Certified Instructors, the maximum number of L2 segment per L2 VPN is wrong, the correct configuration maximum is **512**, irrespectively of the Edge Node form factor.

NOTE: The students' eBook has been already fixed.

[Source: https://configmax.esp.vmware.com/guest?vmwareproduct=VMware%20NSX&release=NSX%204.0.1&categories=21-48]

Module 10

10-39 - Object-Based RBAC in a Multitenancy Environment

MORE DETAILS —> *Object-based RBAC is only configurable through NSX API in NSX 4.0.1* <and> *Object-based RBAC is only configurable through NSX API in NSX 4.0.1. NSX UI support is expected in future releases* <— object-based RBAC and, more in general, Multi-Tenancy in the NSX UI are supported starting with version 4.1.0.

[Source: https://docs.vmware.com/en/VMware-NSX/4.1.0/rn/vmware-nsx-410-release-notes/index.html]

Module 11

11-12 - Ownership of Logical Configuration (2)
BOTCHED / DUPLICATED TEXT —> ***Services*** <— in the slide notes bulleted list, the item "Services" is repeated twice.

11-24 - Active Global Manager Configuration
BOTCHED / MISSING SLIDE <— the **Infrastructure Onboarding** process is described by directly starting with the step required to configure the **Active Global Manager Cluster**. IMHO, it would have been better beginning – for the sake of clarity – with a slide showing the deployment process of the **Global Manager Cluster** nodes, so to demonstrate that **"installing a Global Manager appliance is similar to installing an NSX Manager appliance. The only difference is that when you deploy the appliance, you select NSX Global Manager for the role"**.
[Source: https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-E6C5AA1E-2C3C-42D1-B386-6C99B92E5B21.html]

11-26 - Adding Standby Global Manager (2)
BOTCHED / MISSING INFORMATION —> *You provide the Location B GM information* <— the author forgot to mention that in the **FQDN/IP** input field you must "**enter the FQDN or IP address of the Global Manager cluster VIP at the secondary location. Do not enter an individual Global Manager FQDN or IP**".
[Source: https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-963219D2-C772-4918-8EF0-C0B658F53CA2.html#procedure-1]

11-27 - Adding a Location
BOTCHED / MISSING INFORMATION —> *[…] provide the details about NSX Manager at Location A* <— the author forgot to mention that in the **FQDN/IP** input field you must "**enter the FQDN or IP address of the NSX Manager cluster VIP. Do not enter an individual NSX Manager FQDN or IP**".
[Source: https://docs.vmware.com/en/VMware-NSX/4.0/installation/GUID-88B330FB-752D-4FB3-9B1C-B283E92060BC.html#procedure-2]

11-37 - Tier-0 and Tier-1 Gateways: Logical Topologies (1)
TYPO —> *Tier-0 and Tier-1 gateways can still be local to* ***Location Manager*** <— this is not a new object, it's just the **Local Manager Cluster**.

11-38 - Tier-0 and Tier-1 Gateways: Logical Topologies (2)
WRONG —> ***The span of the Tier-0 gateway is equal to or a subset of the span of the Tier-1 gateway*** <— nope, it's exactly the opposite: ***the span of the Tier-1 gateway is equal to or a subset of the span of the Tier-0 gateway***. See also slide 11-39 graphic for a visual example and notes for a description.

11-39 - Tier-0 and Tier-1 Gateways: Logical Topologies (3)
WRONG —> ***In the example, the span of the T0-Stretched and both T1-Stretched gateways are not the same. These connections are not possible*** <— the reason why the T0-Stretched and both the T1-Stretched gateways in the graphic cannot be connected is correctly explained in the slide notes: **"a connection is not possible between T1-Stretched-1 and T1-Stretched-2 to T0-Stretched because T0-Stretched is not stretched to Location C"**. That is, the span of neither T1-Stretched-1 nor T1-Stretched-2 is equal to or a subset of the span of the T0-Stretched gateway. In other words, if you wish, both T1-Stretched-1 and T1-Stretched-2 do exist in a location (Location C) where the T0-Stretched gateway doesn't.

11-40 - Single-Location Tier-0 Gateway Deployments
TYPO —> *The traffic ingresses or egresses from the edge **nodes** with the active Tier-0 gateway* <— when a Tier-0 gateway is deployed with an Edge Cluster in A/S mode the active Edge node is only one, therefore the highlighted word shall read "**node**" (singular).

11-44 - Multilocation T0-Stretched Gateway Modes (1)
WRONG / OUT-OF-DATE —> ***The Tier-0 gateway in the active-active high availability mode does not support stateful NAT. However, stateless NAT can be used*** <— not anymore: starting with NSX 4.0.1, on Tier-0 gateways with A/A HA mode **"the following stateful services are supported: L4/L7 Gateway Firewall, URL Filtering, NAT and TLS Inspection"**. Here NAT means both Source NAT and Destination NAT.
[Source: https://docs.vmware.com/en/VMware-NSX/4.0.1.1/rn/vmware-nsx-4011-release-notes/index.html]

11-45 - Multilocation T0-Stretched Gateway Modes (2)
BOTCHED / MISSING TEXT —> *For the Tier-0 gateway configured in **a primary** setup […]* <— this slide refers to an All Primary Locations design, therefore the sentence shall read: "For the Tier-0 gateway configured in an **all primary** setup […]".
WRONG / OUT-OF-DATE —> ***Tier-0 does not support services in this deployment mode: Stateless NAT can be used*** <— same as on the note 11-44 above.