VMware NSX-T Data Center 2.4.3 Release Notes

VMware NSX-T Data Center 2.4.3 | 07 November 2019 | Build 15008150

Check regularly for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

- Compatibility and System Requirements
- API and CLI Resources
- Revision History
- Resolved Issues
- Known Issues

Compatibility and System Requirements

For compatibility and system requirements information, see theNSX-T Data Center Installation Guide.

API and CLI Resources

See code.vmware.com to use the NSX-T Data Center APIs or CLIs for automation.

The API documentation is available from the **API Reference** tab. The CLI documentation is available from the **Documentation** tab.

Document Revision History

07 November 2019. First edition.

Resolved Issues

- **Fixed Issue 2294410 - Some Application IDs are detected by the L7 firewall.**
  The following L7 Application IDs are detected based on port, not application: SAP, SUNRPC, and SVN. The following L7 Application IDs are unsupported: AD_BKUP, SKIP, and AD_NSP.

- **Fixed Issue 2295470 - Firewall filters not present after migrating to NSX-T from NSX for vSphere.**
  If services are used in many firewall rules, frequent updates on services may be caused during the migration process. As a result, firewall filters are not installed on the ESXi host. This may cause traffic disruption.
- **Fixed Issue 2314537 - Connection status is down after vCenter certificate and thumbprint update.**
  No new updates from vCenter sync with NSX and all on-demand queries to fetch data from vCenter will fail. Users cannot deploy new Edge/Service VMs. Users cannot prepare new clusters or hosts added in the vCenter. Log locations: /var/log/cm-inventory/cm-inventory.log and /var/log/proton/nsxapi.log on the NSX Manager node.
- **Fixed Issue 2331683 - Add-Load-balancer form on Advance UI not showing updated capacity of version 2.4.**
  When add-load-balancer form is opened, the form-factor-capacity shown on the Advance UI is not updated as per 2.4 version. The capacity shown is from the previous version.
- **Fixed Issue 2424081 - Redistributed routes missing intermittently.**
  Network traffic may be impacted as route to T1 is unavailable for some time.
- **Fixed Issue 2424293 - After VM is migrated to a new host, the VIF connected to a different logical port, distributed firewall rules fail to get applied.**
  VM IP address wasn't realized because UI shows it is duplicated to the IP address on a deleted port.
- **Fixed Issue 2424372 - vIDM fails to provide access token to NSX-T due to improper request made by NSX-T reverse proxy to VIDM.**
  Race condition in reverse-proxy whereby the read logic was reading before the writer (NAPI) had completed writing the file to disk.
- **Fixed Issue 2424386 - Static routes configured on T0 router fail to get installed successfully.**
  VRF for the T0 router won't show any configured static routes. No new configuration is realized.
- **Fixed Issue 2424402 - NSX-T Host upgrade may fail due to insufficient space in /tmp.**
  Host upgrade fails since /tmp does not have sufficient space.
- **Fixed Issue 2424852 - Hosts using CA signed certificate may fail upgrade with error, "Unexpected error while upgrading upgrade unit: Invalid host thumbprint".**
  Host upgrade fails.
- **Fixed Issue 2424855 - NSX Backup operation may fail or can take a long time to complete.**
  Cluster backup can take more than 30 minutes and then fail.
- **Fixed Issue 2327494 - NSX-T host upgrade may fail if SNMP is enabled on the ESXi host.**
  NSX-T upgrade fails if SNMP is enabled on ESX.

- **Fixed Issue 2416081 - Add capability to increase number of Virtual Servers (from 10 to 20) in a small form factor NSX-T load balancer.**
  Some virtual servers don't work for small load balancer.
- **Fixed Issue 2426255 - Add restart capability for JVM services on NSX-T Manager node.**
  System unusable for any meaningful NSX related network configuration or operation if a critical service component is affected.
- **Fixed Issue 2432859 - FIREWALL-PKTLOG events from cloud VMs are not forwarded by Public Cloud Gateway via syslog.**
  Firewall logs from Cloud VMs running NSX enforced mode are not forwarded via syslog from Public Cloud Gateway (PCG). CLI on PCG to forward logs from VMs do not work.
- **Fixed Issue 2424394 - DHCP packets relayed by NSX-T DR cannot reach more than 10 hops.**
  When DHCP server is more than 10 hops away, the relayed DHCP packets cannot reach the server.
- **Fixed Issue 2441080 - NAT may not be performed when flow-cache enabled.**
  For a given traffic flow undergoing NAT on the edge, some packets of the flow may not be translated as expected. End systems that are accessed via NAT may become inaccessible.
- **Fixed Issue 2444968 - Automatic backups are not generated.**
  Backups are not generated even though automatic backups flag is turned on.
- **Fixed Issue 2448712 - NSX-T host upgrade fails due to insufficient space in bootbank partition.**
  Upgrade fails if there is not enough space in bootbank partition of an ESXi host.
- **Fixed Issue 2424838 - Invalid rules are applied to traffic when a T0 downlink has a T1 SR associated.**
  Traffic switched 2 T1 was subject to service at the T0, when a T1 had an SR.
- **Fixed Issue 2422004 - System generated IP prefix list is not migrated correctly during upgrade from NSX-T 2.3 to 2.4.**
  System generated IP prefix list shows no prefixes for some of the downlink ports. Tier-0 is not advertising downlink IP's correctly causing traffic to go down.
- **Fixed Issue 2415066 - NSX-T Host Upgrade from releases prior NSX-T 2.4.3 to NSX-T 2.5.0 will fail.**
  The vmkernel logs show nsxt-kcp module is doing unloading but can't finish the process.
- **Fixed Issue 2434224 - NSX-T 2.4.0 based Edge Appliance may experience a impact to datapath due to segmentation fault caused by exhaustion of flowcache mask.**
  Active edge will be failed over to standby edge.
- **Fixed Issue 2387301 - LACP status would remain down on physical switches when integrated with NSX-T prepared host configured for LACP.**
  Loss of connectivity across VMs and Service connections (Service Routers, etc).

- **Fixed Issue 2448254 - Intermittent network connectivity loss to VMs in multi-tenant environments with overlapping IP subnets.**
  Network connectivity loss.

- **Fixed Issue 2425861 - NSX-T /api/v1/edge-clusters/<id>/state API returns "in_sync" realization state, which is not documented as an expected state.**
  Edge cluster's state and status metrics are not collecting as online metrics.

- **Fixed Issue 2438550 - Fix false positive alert "CorfuDB is disconnected, set Cluster Status Down."**
  Exceptions seen in the log but the exceptions are non fatal.

Known Issues

The known issues are grouped as follows.

- General Known Issues
- Installation Known Issues
- NSX Manager Known Issues
- NSX Edge Known Issues
- Logical Networking Known Issues
- Security Services Known Issues
- Load Balancer Known Issues
- Solution Interoperability Known Issues
- Operations and Monitoring Services Known Issues
- Upgrade Known Issues
- API Known Issues
- NSX Cloud Known Issues

General Known Issues

- **Issue 2389109 - BGP/Routing does not work on T0-SR if Edge hostname starts with a number.**
  BGP/Routing does not work on T0-SR if Edge hostname starts with a number and configuration is not pushed to the routing stack. This is a known limitation.

  Workaround: Use the CLI to change the hostname so a number is no longer the first digit in the hostname. Enable and disable maintenance mode on the Edge node to effect the change.

- **Issue 2239365 - "Unauthorized" error is thrown.**
  This error may result because the user attempts to open multiple authentication sessions on the same browser type. As a result, login will fail with above error and cannot authenticate. Log location: /var/log/proxy/reverse-proxy.log /var/log/syslog

  Workaround: Close all open authentication windows/tabs and retry authentication.

- **Issue 2252487 - Transport Node Status is not saved for BM edge transport node when multiple TN is added in parallel.**
  The transport node status is not shown correctly in MP UI.

Workaround:

1. Reboot the proton, all transport node status can be updated correctly.
2. Or, use the API https://<nsx-manager>/api/v1/transport-nodes/<node-id>/status?source=realtime to query the transport node status.

- **Issue 2256709 - Instant clone VM or VM reverted from a snapshot loses AV protection briefly during vMotion.**
  Snapshot of a VM is reverted and migrates the VM to another host. Partner console doesn't show AV protection for migrated instant clone VM. There is a brief loss of AV protection.

  Workaround: None.

- **Issue 2261431 - Filtered list of datastores is required depending upon the other deployment parameters.**
  Appropriate error on UI seen if incorrect option was selected. Customer can delete this deployment and create a new one to recover from error.

  Workaround: Select shared datastore if you are creating a clustered deployment.

- **Issue 2274988 - Service chains do not support consecutive service profiles from the same service.**
  Traffic does not traverse a service chain and it gets dropped whenever the chain has two consecutive service profiles belonging to the same service.

  Workaround: Add a service profile from a different service to make sure that no two consecutive service profiles belong to the same service. Alternatively, define a third service profile which will perform the same operations of the two original ones concatenated, then use this third profile alone in the service chain.

- **Issue 2275285 - A node makes a second request to join the same cluster before the first request is complete and the cluster stabilized.**
  The cluster may not function properly and the CLI commands get cluster status, get cluster config could return an error.

  Workaround: Do not issue any new join command within 10 minutes to join the same cluster after the first join request.

- **Issue 2275388 - Loopback interface/connected interface routes could get redistributed before filters gets added to deny the routes.**
  Unnecessary routes updates could cause the diversion on traffic for few seconds to min.

  Workaround: None.

- **Issue 2275708 - Unable to import a certificate with its private key when the private key has a passphrase.**
  The message returned is, "Invalid PEM data received for certificate. (Error code: 2002)". Unable to import a new certificate with private key.

  Workaround:

    1. Create a certificate with private key. Do not enter a new passphrase when prompted; press Enter instead.
    2. Select "Import Certificate" and select the certificate file and the private key file.

  Verify by opening the key-file. If a passphrase was entered when generating the key, the second line in the file will show something like "Proc-Type: 4,ENCRYPTED".

  This line is missing if the key-file was generated without passphrase.

- **Issue 2277742 - Invoking PUT https://<MGR_IP>/api/v1/configs/management with a request body that sets publish_fqdns to true can fail if the NSX-T Manager appliance is configured with a fully qualified domain name (FQDN) instead of just a hostname.**
  PUT https://<MGR_IP>/api/v1/configs/management cannot be invoked if a FQDN is configured.

  Workaround: Deploy the NSX Manager using a hostname instead of a FQDN.

- **Issue 2279249 - Instant clone VM loses AV protection briefly during vMotion.**
  Instant clone VM migrated from one host to another. Immediately after migration, eicar file is left behind on the VM. Brief loss of AV protection.

  Workaround: None.

- **Issue 2292116 - IPFIX L2 applied to with CIDR-based group of IP addresses not listed on UI when group is created via the IPFIX L2 page.**
  If you try to create a group of IP addresses from Applied to dialog and enter wrong IP address or CIDR in the Set Members dialog box, those members are not listed under groups. You have to edit that group again to enter valid IP addresses.

  Workaround: Go to the groups listing page and add IP addresses in that group. Then that group can start populating in the Applied to dialog.

- **Issue 1957072 - Uplink profile for bridge node should always use LAG for more than one uplink.**
  When using multiple uplinks that are not formed to a LAG, the traffic is not load balanced and might not work well.

  Workaround: Use LAG for multiple uplinks on bridge nodes.

- **Issue 1970750 - Transport node N-VDS profile using LACP with fast timers are not applied to vSphere ESXi hosts.**
  When an LACP uplink profile with fast rates is configured and applied to a vSphere ESXi transport node on NSX Manager, the NSX Manager shows that the profile is applied successfully, but the vSphere ESXi host is using the default LACP slow timer. On the vSphere hypervisor, you cannot see the effect of lacp-timeout value (SLOW/FAST) when the LACP NSX managed distributed switch (N-VDS) profile is used on the transport node from the NSX manager.

  Workaround: None.

- **Issue 2268406 - Tag Anchor dialog box doesn't show all tags when maximum number of tags are added.**

Tag Anchor dialog box doesn't show all tags when maximum number of tags are added, and cannot be resized or scrolled through. However, the user can still view all tags in the Summary page. No data is lost.

Workaround: View the tags in the Summary page instead.

- **Issue 2310650 - Interface shows "Request timed out" error message.**
  Multiple pages on interface shows the following message: “Request timed out. This may occur when system is under load or running low on resources”

  Workaround: Using SSH, log into the NSX Manager VM and run the "start search resync manager" CLI command.

- **Issue 2320529 - "Storage not accessible for service deployment" error thrown after adding third-party VMs for newly added datastores.**
  "Storage not accessible for service deployment" error thrown after adding third-party VMs for newly added datastores even though the storage is accessible from all hosts in the cluster. This error state persists for up to thirty minutes.

  Workaround: Retry after thirty minutes. As an alternative, make the following API call to update the cache entry of datastore:
  https://{{NsxMgrIP}}/api/v1/fabric/compute-collections/<CC Ext ID>/storage-resources?uniform_cluster_access=true&source=realtime
  where NsxMgrIP is the IP address of the NSX manager where the service deployment API has failed, and CC Ext ID is the identifier in NSX of the cluster where the deployment is being attempted.

- **Issue 2320855 - New VM security tag is not created if user doesn't click Add/Check button.**
  Interface issue. If a user adds a new security tag to a policy object or inventory and clicks **Save** without first clicking the **Add**/**Check** button next to the tag-scope pair field, the new tag pair is not created.

  Workaround: Be sure to click the **Add**/**Check** button before clicking **Save**.

- **Issue 2328126 - Bare Metal issue: Linux OS bond interface when used in NSX uplink profile returns error.**
  When you create a bond interface in the Linux OS and then use this interface in the NSX uplink profile, you see this error message: "Transport Node creation may fail." This issue occurs because VMware does not support Linux OS bonding. However, VMware does support Open vSwitch (OVS) bonding for Bare Metal Server Transport Nodes.

  Workaround: If you encounter this issue, see Knowledge Article 67835Bare Metal Server supports OVS bonding for Transport Node configuration in NSX-T.

- **Issue 2334442 - User does not have permission to edit or delete created objects after admin user renamed.**
  User does not have permission to edit or delete created objects after admin user is renamed. Unable to rename admin/auditor users.

  Workaround: Restart policy after rename by issuing the command "service nsx-policy-manager restart"

- **Issue 2261818 - Routes learned from eBGP neighbor are advertised back to the same neighbor.**
  Enabling bgp debug logs will indicate packets being received back and packet getting dropped with error message. BGP process will consume additional cpu resources in discarding the update messages sent to peers. If there are large number of routes and peers this can impact route convergence.

  Workaround: None.

- **Issue 2390624 - Anti-affinity rule prevents service VM from vMotion when host is in maintenance mode.**
  If a service VM is deployed in a cluster with exactly two hosts, the HA pair with anti-affinity rule will prevent the VMs from vMotioning to the other host during any maintenance mode tasks. This may prevent the host from entering Maintenance Mode automatically.

  Workaround: Power off the service VM on the host before the Maintenance Mode task is started on vCenter.

Installation Known Issues

- **Issue 1957059 - Host unprep fails if host with existing vibs added to the cluster when trying to unprep.**
  If vibs are not removed completely before adding the hosts to the cluster, the host unprep operations fails.

  Workaround: Make sure that vibs on the hosts are removed completely and restart the host.

NSX Manager Known Issues

- **Issue 2282798 - Host registration may fail when too many requests/hosts try to register with the NSX Manager simultaneously.**
  This issue causes the fabric node to be in a FAILED state. The Fabric node status API call shows “Client has not responded to heartbeats yet”. The /etc/vmware/nsx-mpa/mpaconfig.json file on the host is also empty.

  Workaround: Use the following procedure to recover from this issue.

  1. Use the Resolver option.
  2. Delete the FN from NSX.
  3. Re-add the FN again manually through CLI command "join management-plane".

NSX Edge Known Issues

- **Issue 2283559 - /routing-table and /forwarding-table MP APIs return an error if the edge has 65k+ routes for RIB and 100k+ routes for FIB.**
  If the edge has 65k+ routes for RIB and 100k+ routes for FIB, the request from MP to Edge takes more than 10 seconds and results in a timeout. This is a read-only API and has an impact only if they need to download the 65k+ routes for RIB and 100k+ routes for FIB

using API/UI.

Workaround: There are two options to fetch the RIB/FIB.

  - These APIs support filtering options based on network prefixes or type of route. Use these options to download the routes of interest.
  - CLI support in case the entire RIB/FIB table is needed and there is no timeout for the same.
- **Issue 2204932 - Configuring BGP Peering can delay HA failover recovery.**
  When Dynamic-BGP-Peering is configured on the routers that peer with the T0 Edges and a failover event occurs on the Edges (active-standby mode), BGP neighborship may take up to 120 seconds.

  Workaround: Configure specific BGP peers to prevent the delay.

- **Issue 2285650 - BGP route tables populated with unwanted routes.**
  When the allowas-in option is enabled as part of the BGP configuration, routes advertised by Edge nodes are received back and installed in the BGP route table. This results in excess memory consumption and routing calculation processing. If higher local preference is configured for the excess routes, this forwarding loop may result in the route table on some routers being populated with redundant routes.

  For example, route X originates on router D, which is adertised to routers A and B. Router C, on which allowas-in is enabled, is peered with B, so it learns route X and installs it in its route table. As a result, there are now two paths for route X to be advertised to router C, resulting in the problem.

  Workaround: You can prevent forwarding loops by configuring the problematic router (or its peer) to block routes being advertised back to it.

Logical Networking Known Issues

- **Issue 2243415 - Customer unable to deploy EPP service using Logical Switch (as a management network).**
  On the EPP deployment screen, the user cannot see a logical switch in the network selection control. If the API is used directly with logical switch mentioned as management network, user will see the following error: "Specified Network not accessible for service deployment."

  Workaround: Deploy using other type of switches like local or distributed.

- **Issue 2288774 - Segment port gives realization error due to tags exceeding 30 (erroneously).**
  User input incorrectly tries to apply more than 30 tags. However, the Policy workflow does not properly validate/reject the user input and allows the configuration. Then Policy then shows an alarm with the proper error message that the user should not use more than 30 tags. At that point the user can correct this issue.

  Workaround: Correct the configuration after the error displays.

- **Issue 2320147 - VTEP missing on the affected host.**

If a LogSwitchStateMsg is removed and added in the same transaction and this operation is processed by the central control plane before management plane has sent the Logical Switch, the Logical switch state will not be updated. As a result, traffic cannot flow to or from the missing VTEP.

Workaround: If you encounter this issue, restart the central control plane.

- **Issue 2327904 - After using pre-created Linux bond interface as an uplink, traffic is unstable or fails.**
  NSX-T does not support pre-created Linux bond interfaces as uplink.

  Workaround: For uplink, use OVS native bond configuration from uplink profile.

- **Issue 2295819 - L2 bridge stuck in "Stopped" state even though Edge VM is Active and PNIC is UP.**
  L2 bridge may be stuck in "Stopped" state even though the Edge VM is Active and the PNIC that backs the L2 bridge port is UP. This is because the Edge LCP fails to update the PNIC status in its local cache, thereby assuming that the PNIC is down.

  Impact to customer:
  Traffic outage for VMs that are reachable by edge l2bridge port

  Workaround: Restart the local-control agent on the affected Edge VM.

- **Issue 2389993 - Route map removed after redistribution rule is modified using the Policy page or API.**
  A route map added to a redistribution rule from the Management Plane interface or API, may be removed if the same redistribution rule is subsequently modified through the Policy page interface or API. This is due to the Policy page interface or API do not support adding route-maps. This can result in advertisement of unwanted prefixes to the BGP peer.

  Workaround: You can restore the route map by returning the management plane interface or API to re-add it to the same rule. If you wish to include a route map in a redistribution rule, it is recommended you always use the management plane interface or API to create and modify it.

Security Services Known Issues

- **Issue 2395334 - (Windows) Packets wrongly dropped due to stateless firewall rule conntrack entry.**
  Stateless firewall rules are not well supported on Windows VMs.

  Workaround: Add a stateful firewall rule instead.

Load Balancer Known Issues

- **Issue 2290899 - IPSec VPN does not work, control plane realization for IPSec fails.**
  IPSec VPN (or L2VPN) fails to comes up if more than 62 LbServers are enabled along with IPSec service on Tier-0 on the same Edge node.

  Workaround: Reduce the number of LbServers to fewer than 62.

- **Issue 2318525 - The next-hop IPv6 routes as the eBGP peer's IP address gets changed to its own IP.**
  In case of eBGP IP4 sessions, advertised IPv4 routes that have their eBGP peer as the next hop, the next hop of the route is NOT changed at the sender side to its own IP address. This works for IPv4, but for IPv6 sessions, the next hop of the route is changed at the sender side to its own IP address. This behavior can result in route loops.

  Workaround: None.
- **Issue 2362688 - If some pool members are DOWN in a load balancer service, the UI shows the consolidated status as UP.**
  When a pool member is down, there is no indication on the Policy UI where the Pool status is green and Up.Workaround:

  Workaround: Workaround: None.

Solution Interoperability Known Issues

- **Issue 2289150 - PCM calls to AWS start to fail.**
  If you update the PCG role for an AWS account on CSM from*old-pcg-role* to *new-pcg-role*, CSM updates the role for the PCG instance on AWS to *new-pcg-role*. However, the PCM does not know that the PCG role has been updated and as a result continues to use the old AWS clients it had created using *old-pcg-role*. This causes the PCM AWS cloud inventory scan and other AWS cloud calls to fail.

  Workaround: If you encounter this issue, do not modify/delete the old PCG role immediately after changing to new role for at least 6.5 hours. Restarting the PCG will re-initialize all AWS clients with new role credentials.

Operations and Monitoring Services Known Issues

- **Issue 2316943 - Workload unprotected briefly during vMotion.**
  VMware tools takes a few seconds to report correct computer name for VM after vMotion. As a result, VMs added to NSGroups using computer name are unprotected for a few seconds after vMotion.

  Workarround: For groups to be used in DFW rules, use VM name-based criteria instead of computer name-based criteria.

Upgrade Known Issues

- **Issue 2288549 - RepoSync fails with checksum failure on manifest file.**
  Observed in deployments recently upgraded to 2.4. When an upgraded setup is backed up and restored on a fresh deployed manager, the repository manifest checksum present in the database and the checksum of actual manifest file do not match. This causes the RepoSync to be marked as failed after backup restore.

  Workaround: To recover from this failure, perform the following steps:

  1. Run CLI command `get service install-upgrade`
     Note the IP of "Enabled on" in the results.

2. Log in to the NSX manager IP shown in "Enabled on" return of the above command.
3. Navigate to **System > Overview**, and locate the node with the same IP as "Enabled on" return.
4. Click **Resolve** on that node.
5. After the above resolve operation succeeds, click **Resolve** on all nodes from the same interface.
   All three nodes will now show RepoSync status as **Complete**.

- **Issue 2277543 - Host VIB update fails during in-place upgrade with 'Install of offline bundle failed on host' error.**
  This error may occur when storage vMotion was performed on the host before doing an in-place upgrade from NSX-T 2.3.x to 2.4 and hosts running ESXi-6.5P03 (build 10884925). The switch security module from 2.3.x is not get removed if storage vMotion was performed just before the host upgrade. The storage vMotion triggers a memory leak causing the switch security module unload to fail.

  Workaround: See Knowledge Base article 67444Host VIB update may fail when upgrading from NSX-T 2.3.x to NSX-T 2.4.0 if VMs are storage vMotioned before host upgrade.

- **Issue 2276398 - When an AV Partner Service VM is upgraded using NSX, there may be up to twenty minutes of protection loss.**
  When a Partner SVM is upgraded, the new SVM is deployed and old SVM is deleted. SolutionHandler connection errors may appear on the host syslog.

  Workaround: Delete the ARP cache entry on the host after upgrade and then ping the Partner Control IP on the host to solve this issue.

- **Issue 2330417 - Unable to proceed with upgrade for non-upgraded transport nodes.**
  When upgrading, the upgrade is marked as successful even though some transport nodes are not upgraded. Log location: /var/log/upgrade-coordinator/upgrade-coordinator.log.

  Workaround: Restart the upgrade-coordinator service.

API Known Issues

- **Issue 2260435 - Stateless redirection policies/rules are created by default by API, which is not supported for east-west connections.**
  Stateless redirection policies/rules are created by default by API, which is not supported for east-west connections. As a result, traffic is not't get redirected to partners.

  Workaround: When creating redirection policies using the policy API, create a stateful section.

- **Issue 2332397 - API allows creation of DFW policies in nonexistent domain.**
  After creating such a policy on a nonexistent domain, the interface becomes unresponsive when user opens up a DFW security tab. The relevant log is /var/log/policy/policy.log.

  Workaround: Create the domain, with the same ID, on which the policy was created. This permits the validation to succeed.

NSX Cloud Known Issues

- **Issue 2275232 - DHCP would not work for VMs on cloud if DFWs Connectivity_statregy is changed from BLACKLIST to WHITELIST.**
  All the VMs requesting for new DHCP leases would lose IPs. Need to explicitly allow DHCP for cloud VMs in DFW.

  Workaround: Explicitly allow DHCP for cloud VMs in DFW.