



Disaster Recovery with NSX and SRM

Recovering NSX backed Data Centers utilizing SRM (Site Recovery
Manager)

NSBU
UPDATED 4/21/2016

VERSION 3.2

Revision History

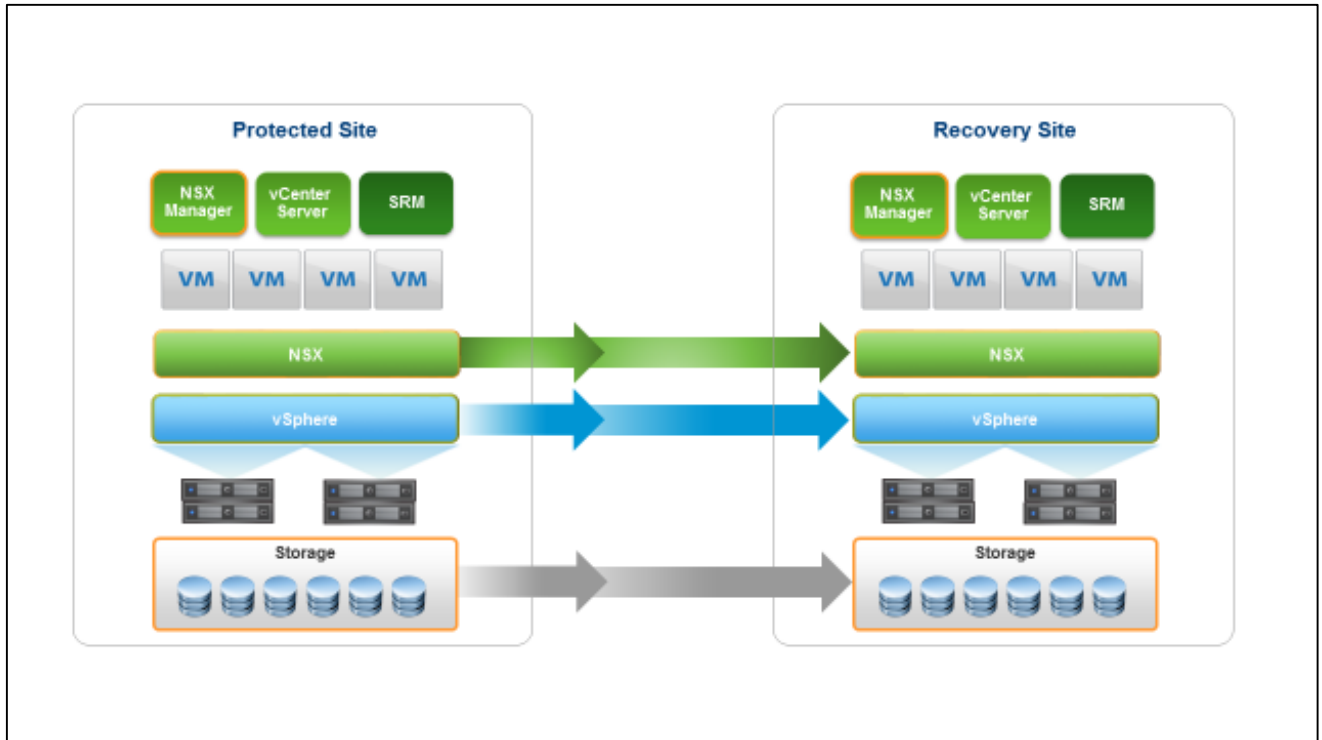
Version	Date	Updates
3.2	4//21/2016	Initial Publication

Table of Contents

Intended Audience	3
Overview	4
Challenges with existing DR Solutions	5
Disaster Recovery Scenarios.....	6
Requirements	7
Initial Configuration.....	12
Solution	16
Initial Set-up	16
Controlling Ingress and Egress Traffic.....	17
Planned Migration/Partial Application Failover	19
Full Application Failover	21
Selecting the N-S Routing Design	23
Handling Full Site Failure.....	24
Other Scenarios: Failback, Re-Protect	26
Ingress Traffic Management using GSLB (Optional)	26
Additional Design Considerations	28
NSX and SRM Integration.....	31
Recovering NSX Components – Detailed Steps.....	32
Recovering NSX Components – Failback Procedure	42
Recovering Control VM.....	43
References	46

Intended Audience

This document is targeted and written for virtualization architects and administrators with experience in designing, deploying and troubleshooting technologies such as VMware vSphere, Site Recovery Manager and NSX along with standard storage technologies.



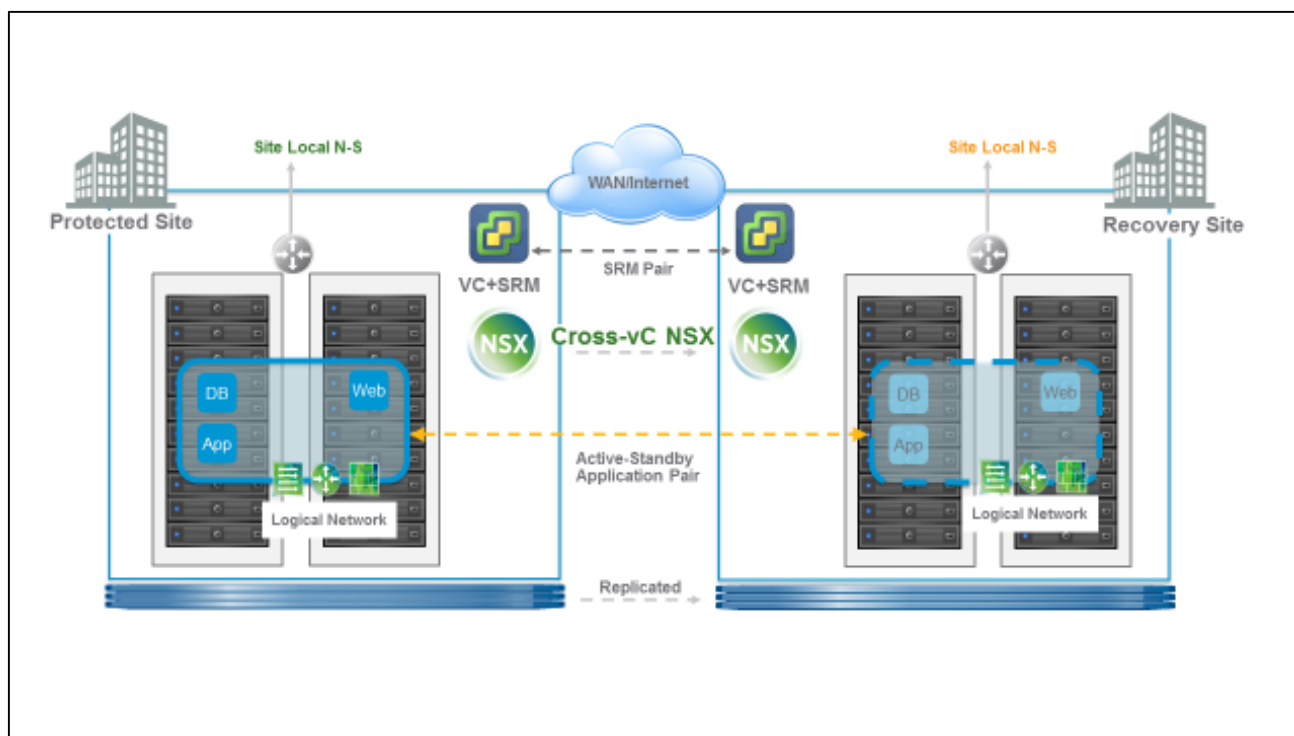
Overview

IT organizations require a methodology for replicating and recovering workloads from a primary site to a recovery site in an event of a disaster or an unplanned outage. To facilitate and automate this recovery process of workloads VMware has products such as Site Recovery Manager (SRM) and vSphere Replication that can automate and orchestrate the recovery process during a failure from a primary site to a recovery site. SRM recovers replicated Virtual Machines in a Secondary Data Center and can perform network mapping (and re-mapping) between the Primary and Secondary locations so that Virtual Machines that are recovered can be reconnected to a L2 Network. These networks can be a VLAN backed Distributed Virtual Port Group (dvPG) or a NSX Logical Switch.

NSX and network virtualization enhances the Disaster Recovery (DR) solution by not only preserving L2 but also recovering the entire logical network topology at the recovery site. NSX also adds API based automation at the networking layer to further improve Recovery Point Objective (RPO) and Recovery Time Objective (RTO). Combining NSX with a SRM based DR design dramatically simplifies the recovery of vital networking services in the secondary location including Logical Switches, Distributed Logical Routers, and Distributed Firewall (DFW) Rules. This document will explain the process of recovering workloads running at Protected and Recovery sites backed by NSX virtual networks.

NSX supports seamless spanning of network and security policies across multiple sites through the use of the Cross-VC NSX feature introduced in NSX 6.2. The DR solution described in this solution brief is based on this Cross-VC NSX feature; however, a DR solution can also be built without leveraging Cross-VC NSX using an external replication/synchronization mechanism (such as vRO) to recreate Logical Networks and Security between NSX instances across the two sites.

Cross vCenter NSX greatly simplifies the process. Deployment consists of Universal Logical Switches, Universal Distributed Logical Router, and Universal Distributed Firewall. These universal objects facilitate the creation of a single unified logical network (L2, L3, DFW) across the protected and recovery site; the application can failover and recover seamlessly without the need for manually re-creating the network on the recovery site or manually mapping/re-mapping IP addresses.



Challenges with existing DR Solutions

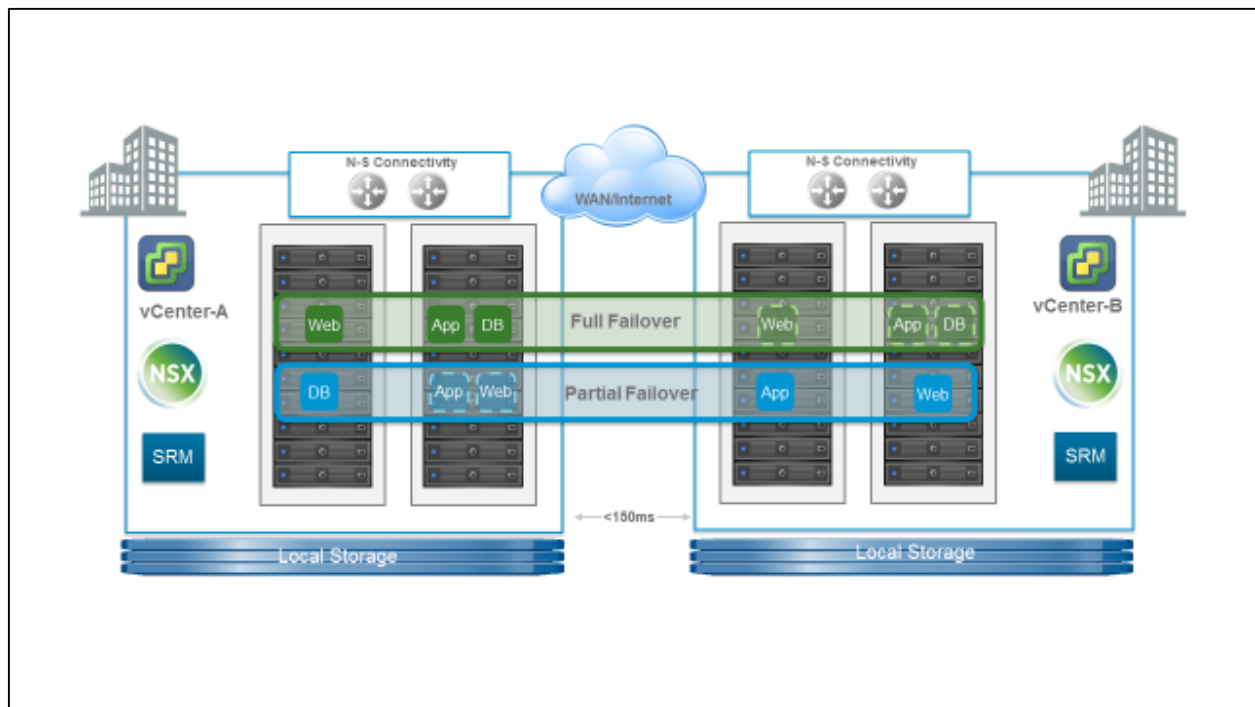
The way Disaster Recovery is done today remains a challenging and complex task. This is primarily due to several underlying issues associated with recovering an application at another physical location with all the associated compute, storage and networking constructs. Existing solutions such as SRM faithfully recover Compute and Storage reliably in an automated fashion, but networking continues to be tied to underlying physical infrastructure. This tie-in with the physical infrastructure causes the following operational and deployment challenges:

- Replicating Layer 2 and Layer 3 network topology across different physical devices at the target and recovery sites
- Replicating security policies across the two data centers
- Reliance on physical IP address in creating policies
- Re-mapping IP address leads to re-writing the associated policies
- Preserving IP address across the data center requires the use of L2 Extension technologies such as Cisco Overlay Transport Virtualization (OTV), or industry standard Virtual Private LAN Services (VPLS) configured generally on the WAN Router – thus, the application's Layer-2 domain is exposed from the rack to the Data Center WAN edge

Disaster Recovery Scenarios

Disaster Recovery solution with NSX is designed and tested to support the following failure and recovery conditions:

- **Partial Application Failover** – Only a part of application failed over from Protected to Recovery site. The application components on the Protected and Recovery site continue to function and communicate as before.
- **Full Application Failover** – Entire application failed over from the Protected site to the Recovery site .
- **Site Failure** – The entire site has failed including NSX components at the protected site. The application and NSX components are recovered at the recovery site.



Requirements

This section outlines the configuration requirements that are needed for the NSX and SRM based Disaster Recovery solution. Most of the requirements are the standard requirements and prerequisites for a vSphere and SRM deployment. The basic NSX deployment outline is augmented with Cross-VC NSX. Please refer to the NSX Design Guide and Cross-VC NSX Deployment Guide for additional details.

Software Requirements Summary

Product	Version	Notes
vSphere	6.0	NSX 6.2 supports vSphere 5.5, however, the DR solution requires support for Cross-VC NSX which is supported with vSphere 6.0 only
NSX	6.2	NSX 6.2 is required for the DR solution outlined in this document
SRM	6.0 6.1	SRM 6.1 is recommended for the NSX based DR solution due to additional integration. However, the solution can be implemented with SRM 6.0 as well. Port groups will not be automatically identified for network mapping in SRM 6.0 (a feature introduced in SRM 6.1)

ESX and vCenter

To ensure independent failure domain for the management plane, separate vCenter instances need to be installed on the Protected and Recovery sites and have clusters of ESX hosts configured. The recommended deployment for an environment running NSX is to have two separate vCenter servers at each site, one for the management clusters and a second one for the compute and edge clusters. This requirement is in-line with NSX recommended guidelines for standard deployments.

SRM

This design doesn't impose any additional restriction for SRM deployment. The SRM requirements remain in-line with the existing recommendation. The only additional requirement is all the Port Groups for the Protected Workloads must be backed by NSX Universal Logical Switches. No network remapping is required in order to deploy this solution.

Storage Replication

This is a basic requirement for any SRM deployment. SRM works with either array-based replication or vSphere replication. Array-based replication requires an SRA (Storage Replication Adapter) to be installed on the SRM server so that it can communicate to the Storage Array. vSphere Replication requires deployment of vSphere Replication Appliances at both sites.

NSX Requirements

Separate NSX manager appliances need to be deployed at each site and linked with the same vCenter server that SRM is linked to. This solution requires **Cross-VC NSX support (NSX 6.2)** where one NSX Manager is deployed in a Primary role and the other in a Secondary role. The Universal Controller Cluster (UCC) is deployed at the protected site with the Primary NSX Manager.

It is beyond the scope of this document to outline NSX 6.2 deployment; please refer to Cross-VC NSX documentation for additional details. Following is a brief overview of Cross-VC NSX deployment:

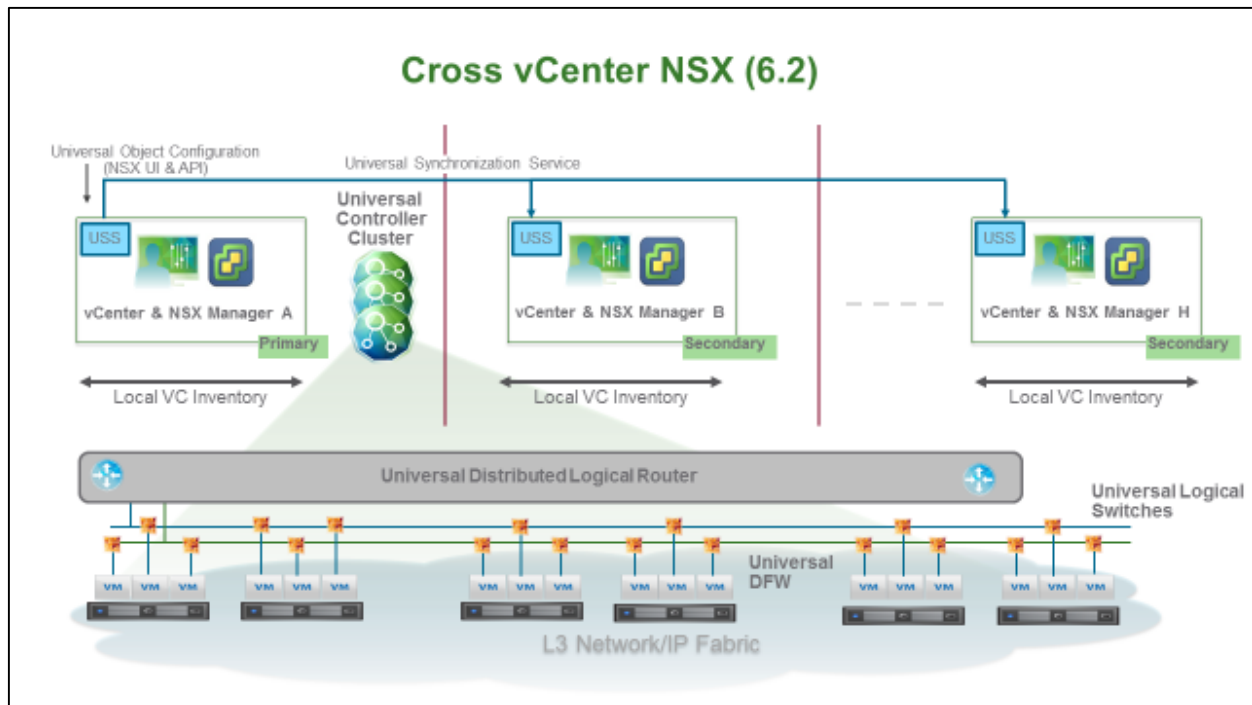
- Cross-VC NSX expands logical networking and security across multiple vCenter domains/sites (Layer 2, Layer 3 and Distributed Firewall)
- Cross-VC NSX supports Universal Logical Objects – Universal Logical Switch, Universal Distributed Logical Router, Universal DFW Rules can be created on the Primary NSX Manager and synchronized across all Secondary NSX Manager(s)
- Cross-VC NSX deployment has one Primary NSX Manager and multiple Secondary NSX Managers
- For each Cross-VC NSX deployment there will be **only one** Controller Cluster called the Universal Controller Cluster (UCC) which will control both the Universal and local logical networks
- Universal logical objects are created/deleted/updated only on the Primary NSX Manager but are readable on the Secondary NSX Manager
- Universal Synchronization Service synchronizes Universal objects across vCenter immediately as well as on demand using force synchronization
- Cross vCenter NSX supports Universal DFW rules; any rules created in the Universal section are synchronized across all NSX Managers

NSX Requirements (Continued)

- Universal DFW rules can only use IP Sets, MAC Sets, and Security Groups containing IP Sets/MAC Sets
- Universal DLR with Local Egress requires two control VMs to be deployed to peer with the Edge Services Gateway (ESG) at each site – this facilitates site specific egress
- Upon failure of a Primary NSX Manager and/or Universal Controller Cluster, the Cross vCenter NSX can be recovered without any loss of configuration data
- Upon failure of a primary NSX manager, a secondary NSX Manager can be promoted to primary NSX Manager and a new Universal Controller Cluster can be deployed

NOTE:

If you have deployed a set of two vCenter servers for management, compute and edge, you will need to link both SRM and NSX to the compute / edge vCenter servers. If your configuration is a single vCenter server at each site, link the SRM and NSX components to the vCenter server that is on the respective supporting site.



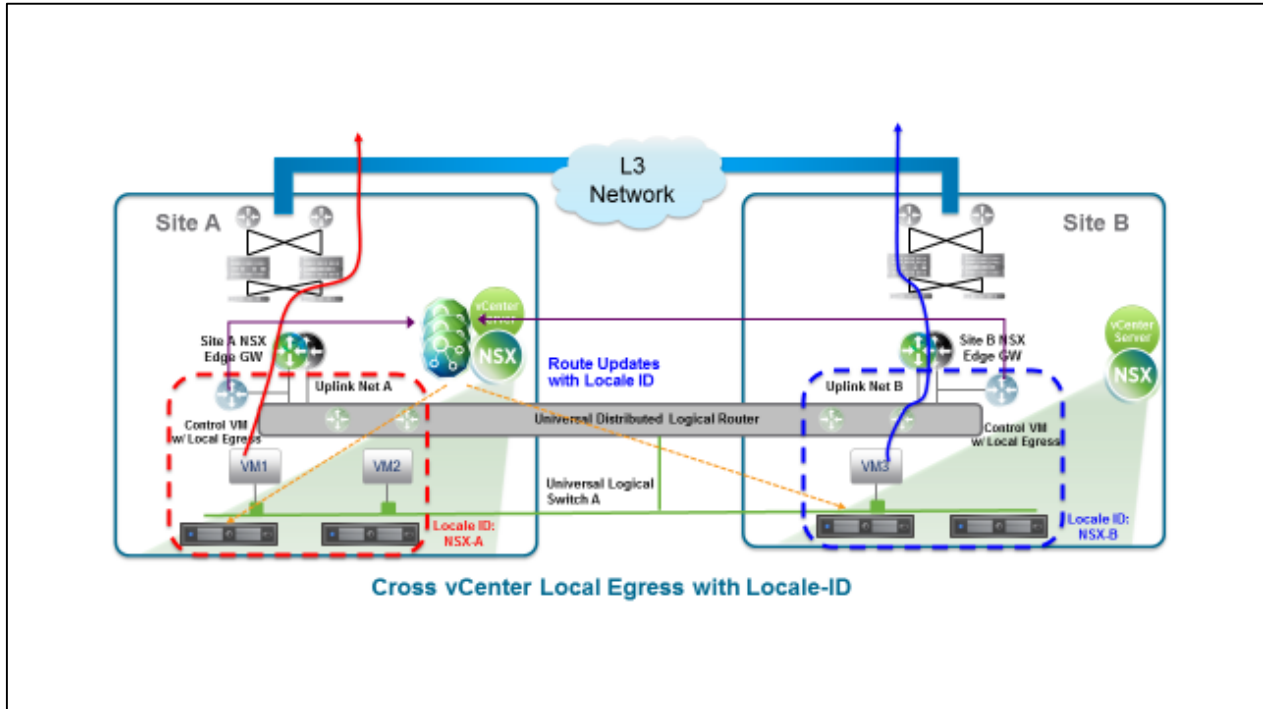
NSX 6.2 and Site Egress Traffic Localization

In order to avoid traffic trombone and additional latency induced in Data Center Interconnect (DCI) traversal, sometime it is desirable to have North-South (N-S) traffic egress from the same site where the workload resides. To address this scenario NSX 6.2 introduced the Egress Localization feature that allows all egress traffic to egress out of a specific site based on a pre-configured Locale-ID. A detailed discussion of how Egress Localization is achieved in NSX is beyond the scope of this document. Egress Localization in context of DR should be treated as a mechanism that allows controlling how the traffic Egresses for workloads; it can be controlled on a per Host or per Cluster or per DLR instance. Some salient points regarding the feature are below:

- By default, Locale ID is set to the NSX Manager UUID. The Locale ID setting is ignored unless Local Egress is enabled. For DR, Local Egress must be enabled
- Once Local Egress is enabled, Locale-ID can be set on per Host or per Cluster or per DLR. For Disaster Recovery we will set Locale-ID on per Cluster (or Host) level
- For each Universal DLR deployed, two control VMs will be deployed. One at the Protected and one at the Recovery site. Each control VM will peer with the Edge

Services Gateway (ESG) at the local site to learn N-S routes specific to the site via Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP)

- DLR running on the hosts in the Protected and Recovery sites will get the route updates from the controller matching the configured Locale-ID. Updating Locale-ID to redirect egress traffic during different DR events will be discussed in the next section. The NSX based DR design leverages Locale-ID as a mechanism to control N-S traffic going from protected workloads (Egress)



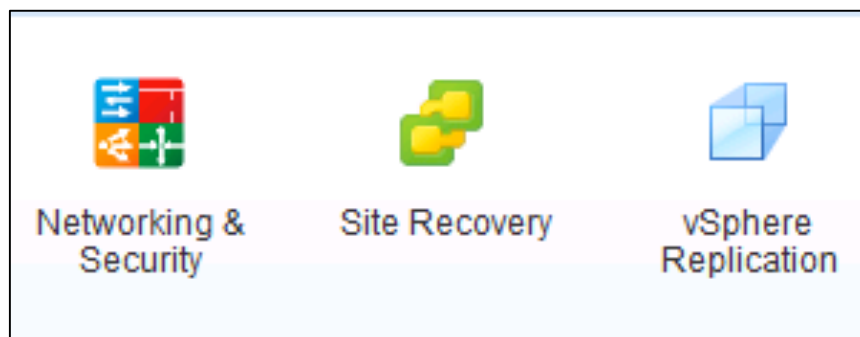
Initial Configuration

Please refer to the NSX 6.2 Documentation, NSX Design Guide, and Cross-VC NSX Deployment Guide before proceeding ahead with the design and configuration.

Configuration Steps

For this solution, the underlay physical network configuration, host preparation and connectivity with the N-S physical routers should be in-line with the existing best practices for NSX deployment. The steps below outline Cross-VC NSX deployment for DR.

- Two independent sites with Layer 3 network connectivity between them would serve as an underlying fabric for the NSX backed overlay (logical networks). The fact that there is no other requirement from the underlying physical infrastructure between the sites greatly simplifies the deployment since there is no need to provide L2 stretching and other WAN connectivity options to preserve the application IP address and associated security policies.
- Each site will have the following.
 - One or more vCenter Servers running a supported version of vCenter and ESXi
 - A number of ESXi servers in a datacenter managed by this vCenter Server where Virtual Machines will run.
 - A SRM 6.0/6.1 Server paired and registered with the vCenter Servers at the Protected and Recovery Sites
 - SRM Configuration based on standard best practices







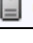
- A **Primary** NSX Manager (NSX 6.2) paired and registered with the vCenter Server that manages the ESXi hosts at the Protected Site

Primary NSX Manager

NSX Manager: 192.168.0.15 (Role: Primary)

NSX Component Installation on Hosts

 Actions


Clusters & Hosts	Installation Status
▼  Edge	✓ 6.2.0
 esx4.mylab.vmware.com	✓ 6.2.0
▼  Compute	✓ 6.2.0
 esx3.mylab.vmware.com	✓ 6.2.0


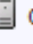

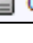
- A **Secondary** NSX Manager (NSX 6.2) paired and registered with the vCenter Server that manages the ESXi hosts at the Recovery Site

Secondary NSX Manager

NSX Manager: 192.168.2.15 (Role: Secondary)

NSX Component Installation on Hosts

 Actions

Clusters & Hosts	Installation Status
▼  Recovery-Edge	✓ 6.2.0
 dr-esx1.mylab.vmware.com	✓ 6.2.0
▼  Recovery-Compute	✓ 6.2.0
 dr-esx3.mylab.vmware.com	✓ 6.2.0

- A NSX Universal Controller Cluster (UCC) of 3 nodes deployed at the Protected Site; all three nodes MUST be deployed at the same physical location

NSX Controller nodes						
Actions						
Controller IP Address	ID	Status	Software Version	NSX Manager		Managed By
192.168.0.31	controller-2	Normal	6.2.44098		192.168.0.15	192.168.0.15
192.168.0.30	controller-1	Normal	6.2.44098		192.168.0.15	192.168.0.15
192.168.0.32	controller-3	Normal	6.2.44098		192.168.0.15	192.168.0.15

- A Universal Transport Zone encompassing Hosts (or Clusters) at the Protected and Recovery site is required to create Universal Logical Switches (ULS) spanning the sites. ULS preserves the logical L2 (and the underlying subnet) of the Protected application

Universal Transport Zone/Universal Logical Switches

NSX Manager: 192.168.0.15 (Role: Primary)					
Actions					
Name	Status	Transport Zone	Scope	Segment ID	
Universal-Web-LS	Normal	Universal Transport Zone	Universal	900001	
Universal-Transit-LS-Recov...	Normal	Universal Transport Zone	Universal	900005	
Universal-Transit-LS-Protect...	Normal	Universal Transport Zone	Universal	900004	
Universal-SW-01	Normal	Universal Transport Zone	Universal	900000	
Universal-DB-LS	Normal	Universal Transport Zone	Universal	900003	
Universal-App-LS	Normal	Universal Transport Zone	Universal	900002	

- Each site (Protected and Recovery) will have its own set of ESGs for the N-S connectivity – this is required for **Local Egress** and to redirect traffic independently from either Protected or Recovery site

Protected Site (UDLR and ESG)

NSX Edges		
NSX Manager: 192.168.0.15 (Role: Primary)		
Id	Name	Type
edge-528b0e3d-ed46-41da-a469-63ec8e572...	Universal-DLR-Protected	Universal Distributed Router
edge-1	Perimeter-Protected	NSX Edge

Recovery Site (UDLR and ESG)

NSX Edges		
NSX Manager: 192.168.2.15 (Role: Secondary)		
Id	Name	Type
edge-528b0e3d-ed46-41da-a469-63ec8e572...	Universal-DLR-Protected	Universal Distributed Router
edge-1	Perimeter-Recovery-SiteB	NSX Edge

- At each site, a Control VM is deployed and local ESG is deployed. This allows each site (Protected and Recovery) to learn and advertise N-S route independently. OSPF is shown in the example below but internal BGP (iBGP) can be used as well.

Protected Site (UDLR Routing)

Routing Configuration :		OSPF Configuration :					
Locale ID :	4225C00E-7A85-FF68-7917-8BAE8C3D9616	Status :	✓ Enabled				
		Protocol Address :	10.114.220.27				
		Forwarding Address :	10.114.220.26				
Area Definitions :		Area to Interface Mapping :		Route Redistribution table :			
Area ID	Type	Interface	Area ID	Learner	From	Prefix	Action
51	Normal	Uplink-to-Protected-Universal-Transit-LS	51	OSPF	Connected	Any	Permit

Recovery Site (UDLR Routing)

Routing Configuration :		OSPF Configuration :					
Locale ID :	42344FE6-5CED-C58A-994A-1D20A1F5BC83	Status :	✓ Enabled				
		Protocol Address :	10.114.220.35				
		Forwarding Address :	10.114.220.34				
Area Definitions :		Area to Interface Mapping :		Route Redistribution table :			
Area ID	Type	Interface	Area ID	Learner	From	Prefix	Action
51	Normal	Uplink-to-Recovery-Universal-Transit-LS-SiteB	51	OSPF	Static routes	Any	Permit

Solution

As outlined in the physical design, the solution leverages NSX 6.2 universal objects to create a single unified logical network that exists at both the Protected and Recovery site. Protected applications are placed on a Universal Logical Switch which is connected to a Universal Distributed Logical Router. Security policies are configured using Universal Distributed Firewall rules. This set-up ensures that the entire Logical Network (L2, L3, and Firewall) for the application spans seamlessly across both sites. When the application fails over, it gets placed on the **same** logical network as was on the recovery site and ensures the same consistent security policy via Universal Distributed Firewall Rules.

The logical objects are created only once on the Primary NSX Manager and are immediately synchronized to the Secondary NSX Manager(s); this leads to guaranteed correctness (since no manual intervention is needed) while the application's networking and security is built automatically without modifying the underlying physical infrastructure. Since the entire configuration is synchronized from Primary to Secondary NSX Manager(s) – loss of Primary NSX Manager (and associated Universal Controller Cluster) doesn't lead to any loss on Logical Network and Security configuration for the application.

Universal Logical Switching across the two sites over Layer-3 physical fabric provides IP address preservation for the application while maintaining a stable and scalable physical underlay – allowing the application to failover without any remapping to a new IP subnet at the recovery site. The Universal DFW based security policies in place at the Recovery site will remain in force since the application IP address is fully preserved.

Initial Set-up

The initial set-up outlined below, is built upon the physical components outlined earlier. Following is a brief outline of the initial set-up.

- SRM and vCenter deployed at the Protected and Recovery site
- Primary NSX Manager and UCC deployed at the Protected site and Secondary NSX Manager deployed at the Recovery site
- A Universal Transport Zone (UTZ) created at the Primary site. The Hosts and Clusters in the Universal Transport Zone are added at each NSX Manager; for example, protected hosts/clusters get added at the Primary NSX Manager and recovery hosts/cluster get added at the Secondary NSX Manager
- Web, App, and DB Universal Logical Switches are created at the Primary NSX Manager (these Universal Logical Switches get automatically synchronized by the Primary NSX Manager to the Secondary NSX Manager)

- A Universal DLR for routing across Web, App and DB tiers of the Application is created at the Primary NSX Manager which will get automatically synchronized to the Secondary NSX Manager

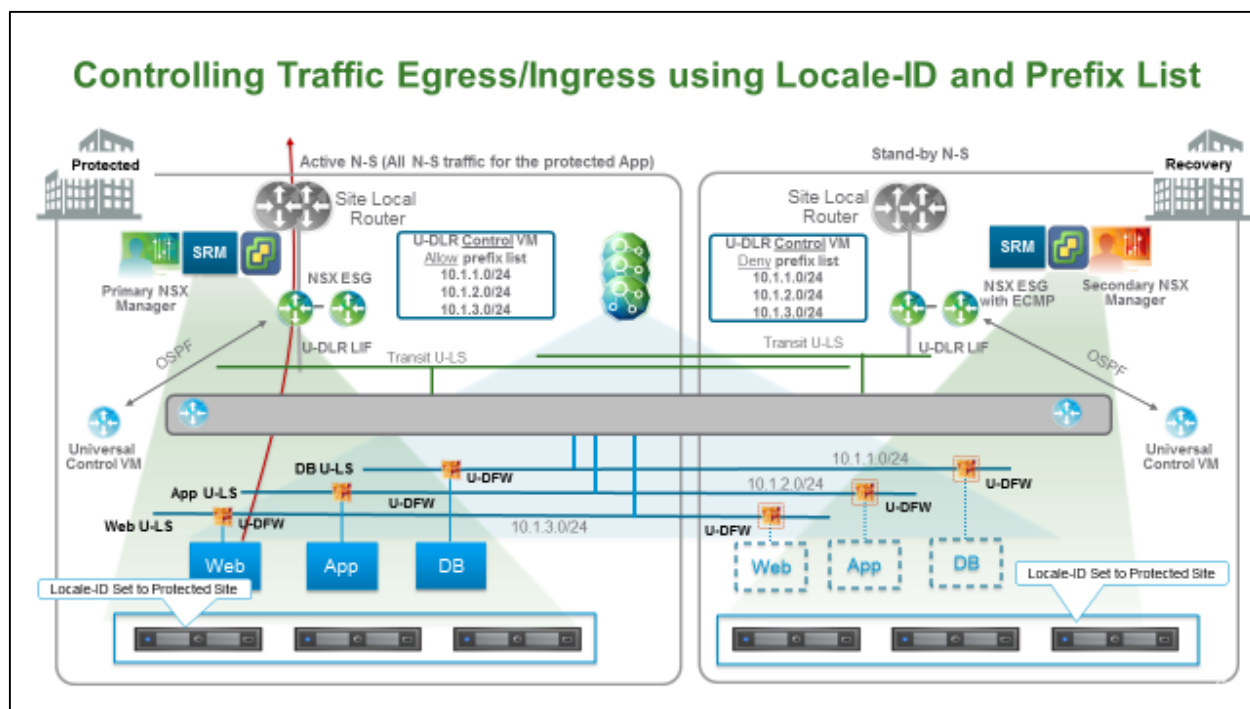
Controlling Ingress and Egress Traffic

There are two approaches to controlling N-S routed traffic in the DR Solution:

a) Locale-ID based N-S traffic redirection

In this design option Universal Distributed Logical Router (UDLR) exist on both primary and recovery site and ensures that East-West routing across application components continue to function without any manual intervention. **Two** UDLR Control VMs are deployed and two UDLR uplink Logical Interface (LIF) are connected to the respective ESG at each site to ensure that North-South routing continues to work after the failover. The use of two UDLR Control VMs allow full **independent** routing adjacency to exist at both sites, and, in the event of a site failure, the other Control VM continues to function. To control the ingress and egress traffic from the Protected site until site failure, a **deny prefix list** is used to restrict the route advertisement from the Recovery site. This will ensure that routes are only advertised from the Protected site and all the traffic will ingress/egress from there.

- Two transit Universal Logical Switches are created for UDLR uplink LIF; the LIF is connected to the site local ESG
- At each site one Control VM for the UDLR is deployed; each Control VM peers with the local ESG via dynamic routing protocol (BGP/OSPF). Since two control VMs are deployed one at Protected and one at the Recovery site, in the event of complete site failure, the recovery site control VM is ready with routing adjacency pre-established with the local ESG
- **Site Egress Traffic:** Local Egress is enabled and the Locale-ID for the hosts at both the Protected and Recovery site is set to the UUID of the Protected site (it MUST be in UUID format and UUID of the Primary NSX Manager is the default)
- **Site Ingress Traffic:** Using a route prefix list, respective subnets (Web/App/DB) are only advertised from the Protected site preventing the recovery site from receiving any inbound traffic (the Web/App/DB prefix advertisement from Recovery site is not advertised because of the deny clause on the prefix at Recovery site),



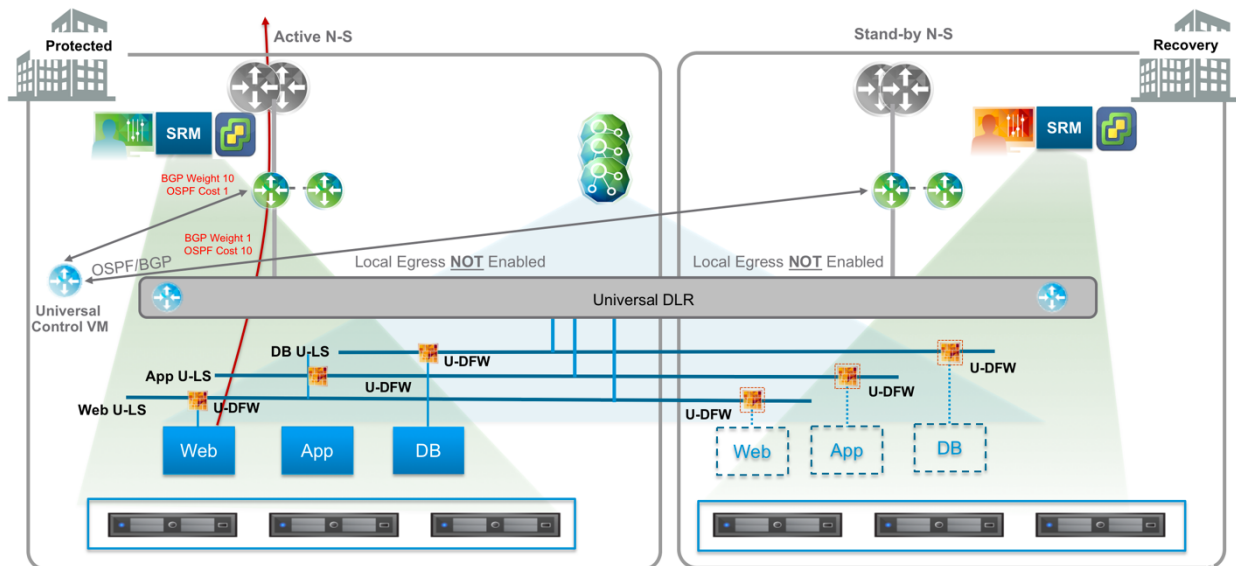
b) Dynamic Routing based N-S traffic redirection

In this design, Universal Distributed Logical Router (UDLR) exists on both Primary and Recovery site, however, **only one** UDLR Control VM is deployed. Only one uplink LIF is connected to all the ESGs at both the sites. Full routing adjacency exist between the Control VM and the ESGs at each site. In this set-up, routes are learned from ESGs from both the sites as well as the NSX backed logical networks advertised from both the ESGs (and sites). To control the egress traffic from the Protected site until site failure, routing metric or cost is utilized (OSPF Link Cost or BGP weight). For OSPF, cost is set lower on the ESG for Recovery site routes; for BGP, BGP weight on the UDLR is utilized and set lower for the BGP adjacency with the Recovery site ESG (BGP attributes should also be configured on the physical network to prefer correct ESG). This will ensure that even through routes are advertised from both Primary and Secondary sites, all the traffic always ingress/egress from the Protected site ESG in normal conditions.

- Only one transit Universal Logical Switch is created for UDLR uplink LIF, connecting to the ESGs of both the sites
- Only one Control VM for the UDLR is deployed peering with ESGs from both the sites via dynamic routing protocols (OSPF/BGP). In the event of complete site failure, the Control VM needs to be redeployed at the Recovery site and the routing adjacency has to be re-established
- Local Egress is NOT enabled and the Locale-ID setting is NOT required

- Set OSPF/BGP cost/attributes so that the Protected site is always the preferred route for ingress and egress traffic
- **Site Egress Traffic:** The NSX Controller will install only the best routes (lowest cost) learned by the Control VM in the DLR instance on each host. In this case it will be via the ESG from the protected site. This will cause all the traffic to egress from the protected site
- **Site Ingress Traffic:** The upstream routers connected to ESG on both the sites will learn the routes, however, the upstream router from the Protected site will receive all the ingress traffic since that is the best cost route to the NSX backed logical networks

Controlling Traffic Egress/Ingress Using Link Cost/Weight



Planned Migration/Partial Application Failover

Most enterprise designs require that Disaster Recovery design should also be able to handle other frequent events such as planned migration for maintenance or partial failover of the application (as opposed to full failover). A partial failover of an application not only involves recovering the application reliably but also involves ability to connect to the components that are still running at the Protected site (cross-site networking). Cross-VC NSX enables both Cross-VC Networking and Security across the sites and pre-created universal logical networks for instantaneous application failover support. Following are the steps involved in Partial Failover via SRM.

- Execute the partial failover and recovery plan from SRM; in this case failing over Web and App VM while keeping the DB VM at the Protected site

Site Recovery

Protection Groups 2

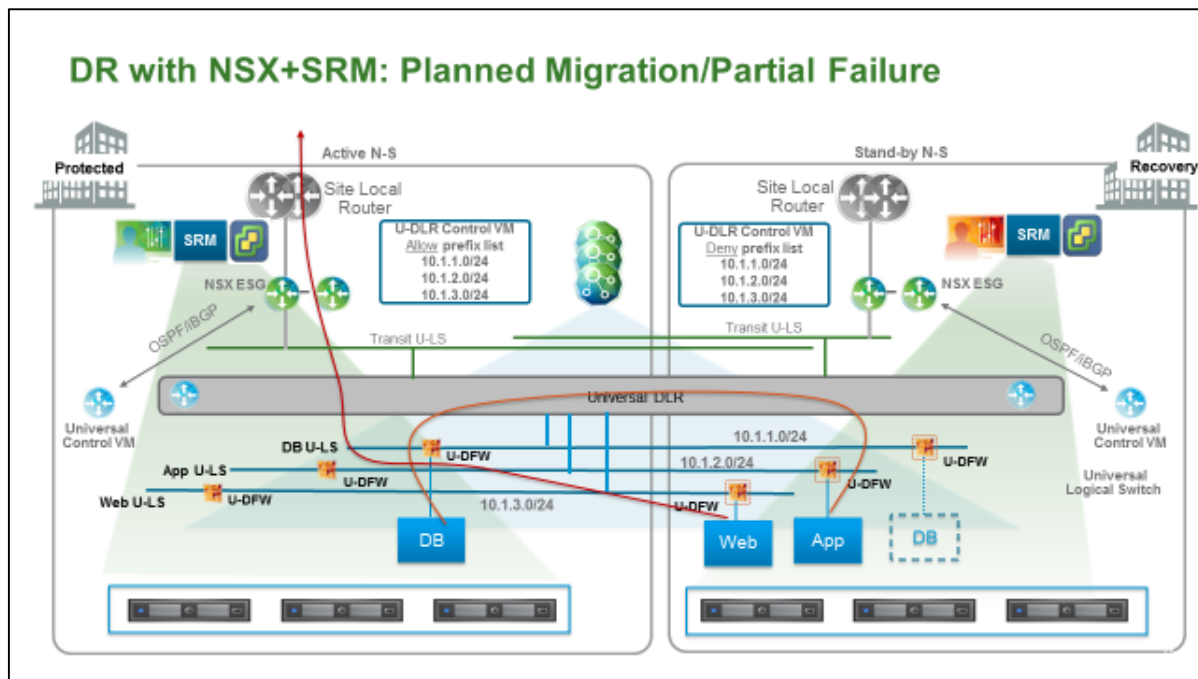
Database

WEB and APP

Recovery Step

1. Synchronize Storage
2. Restore Recovery site hosts from standby
3. Suspend Non-critical VMs at Recovery Site
4. Create Writeable Storage Snapshot
5. Configure Test networks
6. Power On Priority 1 VMs
7. Power On Priority 2 VMs
8. Power On Priority 3 VMs
9. Power On Priority 4 VMs
10. Power On Priority 5 VMs

- Zero Touch Partial Migration:** From the NSX perspective there is no change needed to execute the partial failover (or planned migration) since both the Logical Networks and cross-site networking works as outlined in the initial set-up
- Ingress/Egress Traffic:** The traffic continues to ingress and egress from the Protected site for both the designs – Locale ID based Traffic Ingress/Egress (shown below) and Dynamic Routing based Ingress/Egress (discussed earlier but not shown below)



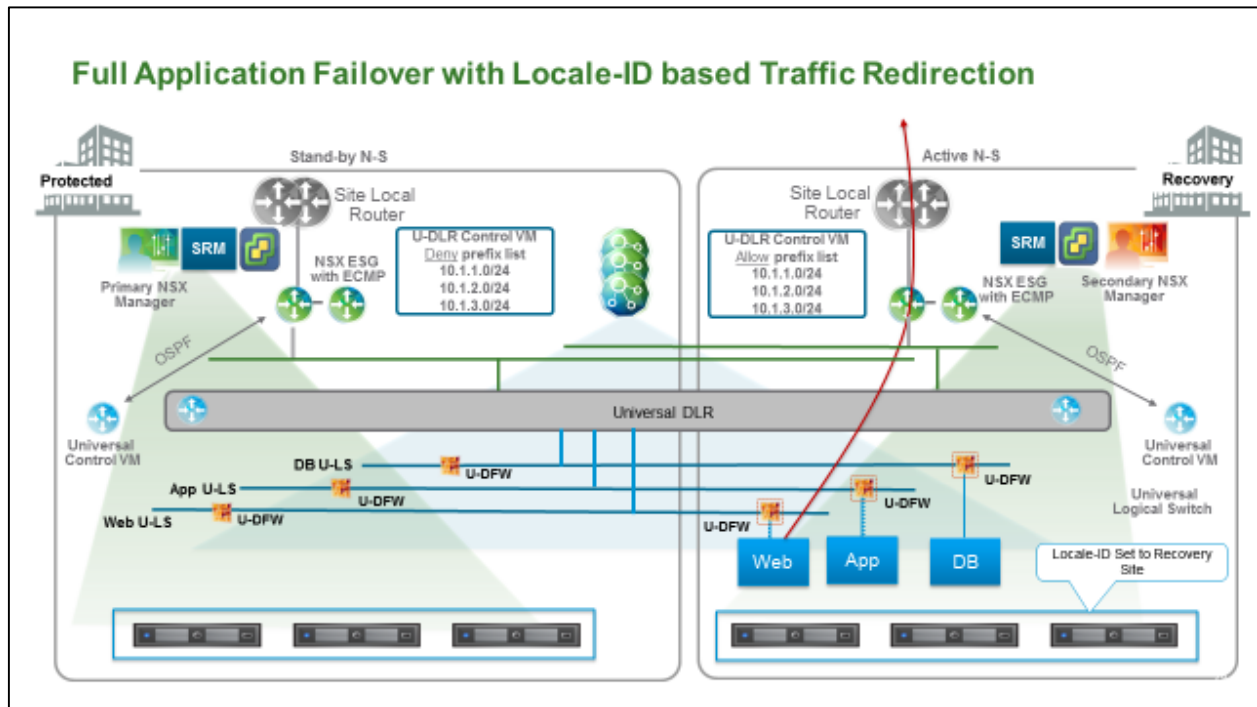
Full Application Failover

Full application failover includes moving over the entire Application (Web, App and DB) from the Protected to the Recovery site. For per application failover designs, the Locale-ID based N-S traffic redirection is a preferred approach; following below are the steps involved in full application failover.

- Execute the full application failover plan from SRM; in this case, failing over Web, App, and DB VMs
- From the NSX perspective this involves traffic ingress/egress from the recovery site

Application failover with Locale-ID based N-S traffic redirection

- **Site Egress Traffic:** To achieve traffic egress from the Recovery site, Locale-ID of the Recovery site host/cluster is updated to UUID of the Secondary NSX Manager (can be done via API or UI)
- **Site Ingress Traffic:** Ingress traffic is controlled by advertising the prefixes from Recovery site and withdrawing from the Protected site. This is accomplished using Allow/Deny clause on the prefix list (can be done via UI or API)

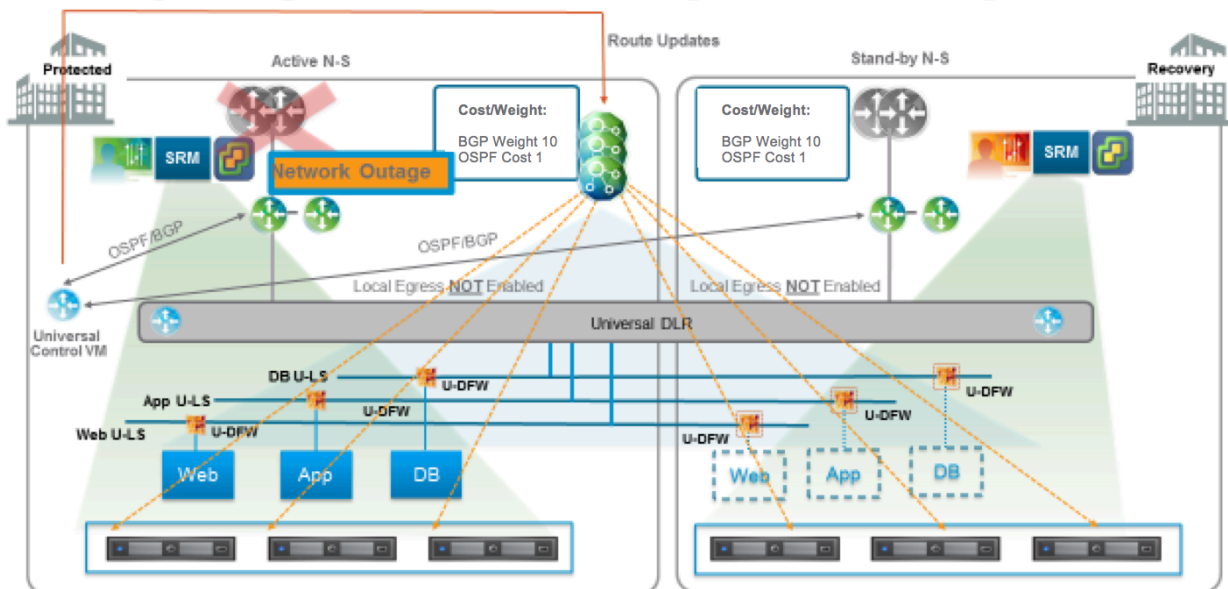


Dynamic Routing based N-S traffic redirection

This approach does NOT provide per-subnet granularity that Locale-ID based failover provides and hence not recommended for granular per application failover; see the next section regarding how to select the optimal N-S routing design for your deployment. This approach is recommended when it is required to automate the application (or multiple applications) failover in response to upstream routing failure. For dynamic routing based failover, the Egress ESG or other upstream routers must failover to trigger a re-routing via the recovery site.

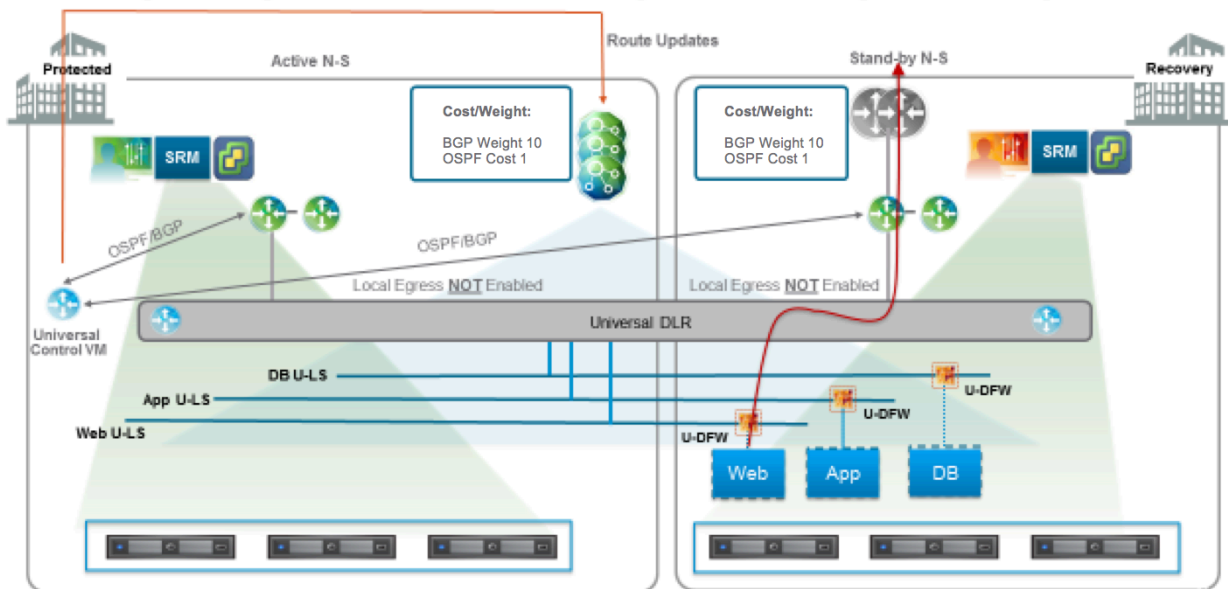
- **Site Egress Traffic:** In case of an uplink failure, ESG will withdraw the routes triggering a route update via Control VM which would in turn trigger route updates in each DLR instance running on the host. The Controller will update with the N/S route via ESG at the recovery site since it is the only available N-S route (even with higher cost)
- **Site Ingress Traffic:** The Ingress traffic will also route from the Recovery site once ESG at Protected site adjacencies are lost due to local or site failure. The traffic will be re-routed by the upstream routers in response to a network/link failure at the protected site. (Completely automatic via dynamic routing)

Traffic Egress/Ingress with Link Cost/Weight: Network Outage Scenario



In the scenario below, the application (Web, App, DB) is shown migrated to recovery site, although it is not necessary. The recovery plan, in response to an upstream network failure, can move the application (Web, App, DB) or keep it at the protected site.

Traffic Egress/Ingress with Link Cost/Weight: Redirecting to Recovery Site



Selecting the N-S Routing Design

As noted below, both the failover approaches have their own benefits and downsides. We recommend failover using dynamic routing whenever possible, because it's a simpler and an automated approach which requires no user intervention to redirect traffic in response to an upstream network failure. The Locale-ID based failover is recommended when there is a need to redirect traffic more granularly in response to per application failure (while the network is fully intact). Locale-Id (with prefix list) based failover require manual steps (that can be automated) to update the Locale-ID and prefix list in response to an application failover as discussed in earlier sections.

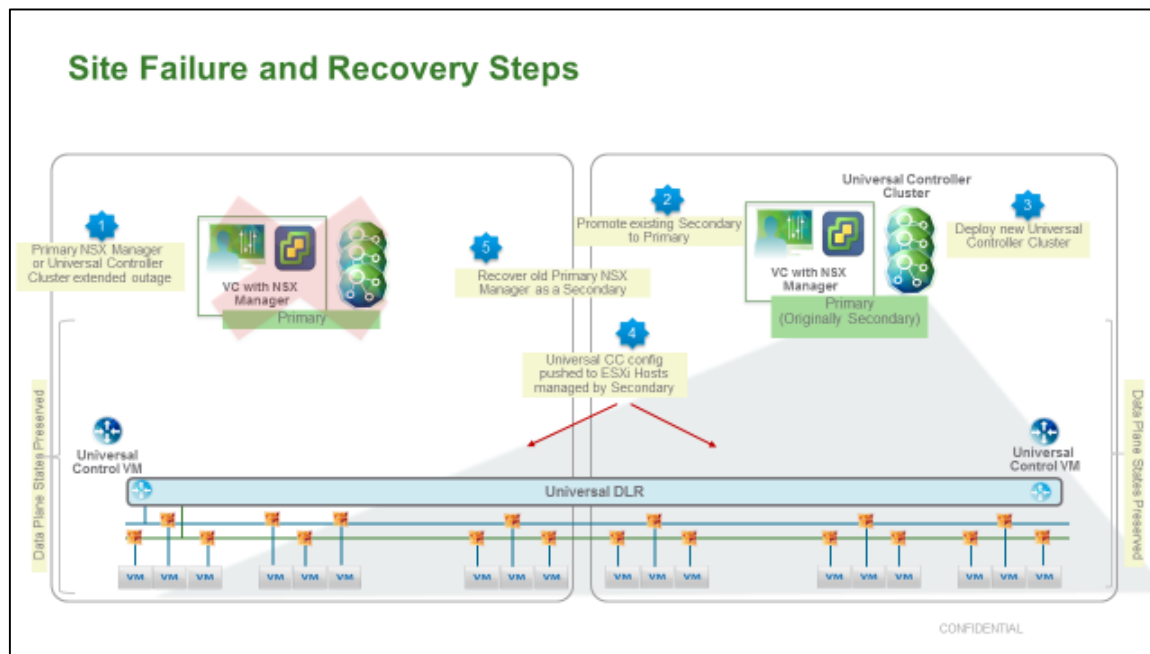
Pros & Cons	Benefits	Downside
Deployment		
Failover using Dynamic Routing	<ul style="list-style-type: none"> ■ Dynamic Routing detects ESG/upstream failures ■ Automatic network failover (via Routing Updates) 	<ul style="list-style-type: none"> ■ Purely Active/Passive for North-South traffic ■ Requires Control-VM redeployment in the event of an infrastructure failure
Failover using Local Egress	<ul style="list-style-type: none"> ■ Granular/Controlled Failover (per network) ■ Provides opportunities for automation triggered by compute operations/DR Orchestration ■ Control VM doesn't need to be redeployed (already provisioned at recovery site) 	<ul style="list-style-type: none"> ■ Requires Additional Configuration (Control VM, Local Egress) ■ Not fully automatic in the event of upstream network or site failures ■ Route redistribution or Locale-ID changes required to failover routing

Handling Full Site Failure

Complete site failure involves running the recovery plan for all the applications and recovering NSX components that are needed for continued function of Cross-VC NSX deployed across the sites. A detailed discussion of NSX component recovery is beyond the scope of this document, please refer to Cross-VC NSX Deployment guide. A brief outline of NSX component recovery in the event of complete site failure is provided below.

- Run the SRM Recovery plan for site failure
- The solution assumes that NSX Primary Manager and Universal Controller Cluster at the Recovery site are inaccessible
- As a part of SRM Recovery the VMs will be placed in the right port groups which are indeed backed by VXLAN backed Universal Logical Switches; loss of NSX components at the Protected site will NOT impact placement of recovered VMs on the appropriate port groups
- The Universal Logical Switch, Universal DLR and Control VM at the Recovery site will be intact even though Primary NSX Management/Universal Controller Components are lost. There is also complete separation of control and data plane; NSX data plane runs in the hypervisor/host and will continue to function even without a Controller or Manager
- The control plane and data plane separation allows recovery of NSX components at the Secondary site while keeping the **existing** data plane state intact (the new states or changes in the physical topology will not get realized in the data plane)

Steps for Recovering NSX Components (NSX Manager and Controller)



Initial NSX Component Recovery:

Step 1. From the Recovery site, disconnect the Secondary NSX Manager from the primary NSX manager

Step 2. Make the Secondary NSX Manager role Primary

Step 3. Re-deploy the controller cluster on the Recovery site from the newly promoted NSX Manager

Step 4. Sync Controller State

Step 5. If Local Egress is used, Universal Control VM already exists, otherwise Universal Control VM needs to be redeployed

Once Initial Primary Site Recovers:

Step 6. Once the initial protected site has been recovered, the respective NSX Manager can be added as a Secondary to the new Primary NSX Manager. The first step is to force the removal of the old registered secondary NSX Manager on the original primary site's NSX Manager.

Step 7. Demote the Protected site/original primary site NSX Manager to transit role

Step 8: Delete UCC nodes at the Protected site/original primary site

Step 9. Also, depending on deployment model, if needed, delete the Control VM at the Protected site as it has been successfully redeployed at the new primary site.

Step 10. Add the Protected site NSX Manager that is currently in transit mode as a Secondary NSX Manager to the newly promoted Primary NSX Manager at the Recovery site

At this point, VC components are all functional and working; there are two NSX Managers deployed and the Primary and Secondary roles have been set. The Universal Controller Cluster is displayed in a normal state.

Other Scenarios: Failback, Re-Protect

There are a few other SRM operations such as Failback and Re-Protect that are not described in the solution above. From NSX perspective these scenarios are merely an extension of different sections described above.

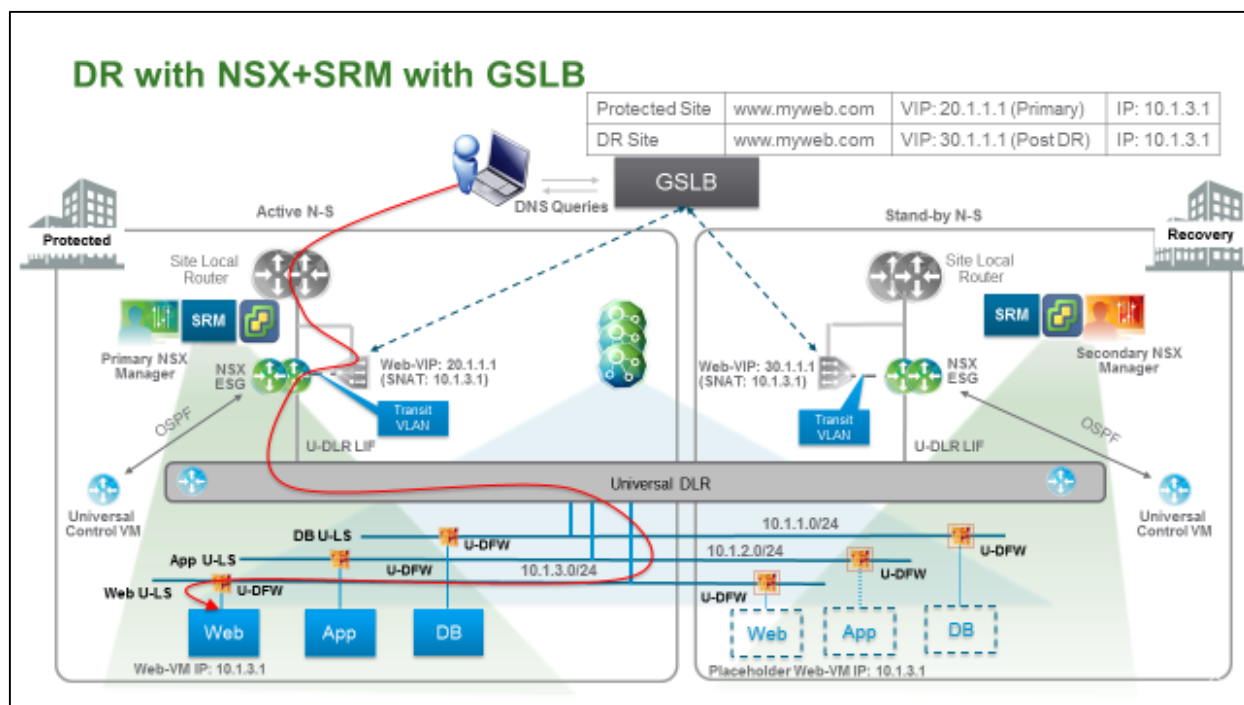
- **SRM Failback Operation:** This operation is merely the failover operation described above in the reverse direction. The universal logical networks continue to function the same way regardless of the direction of application failover (Protected to Recovery or Recovery to Protected)
- **SRM Re-Protect Operation:** This operation is similar to initial-set up described above, except Primary NSX Manager and Universal Controller Cluster reside at the Recovery site

Ingress Traffic Management using GSLB (Optional)

While the failover and recovery solution with NSX is sufficient to address most of the designs, there are scenarios where it is preferred to redirect traffic to Protected and Recovery sites using Global Site Load Balancer (GSLB). Traffic redirection based on GSLB is another way to redirect traffic into a site (similar to Locale-ID and dynamic routing based designs described earlier). It is beyond the scope of this document to outline a complete GSLB solution; it is highly recommended to refer to GSLB vendor's documentation for design details of such a deployment. Following below is a brief outline of how GSLB based traffic redirection can be achieved in an NSX based DR solution.

- GSLB based solution works like a DNS with ability to monitor health of a site or Virtual IP of an application running on a site
- No changes are necessary in the logical IP addressing and connectivity in the NSX based DR design described earlier, the IP Address(s) of the Web (or any Internet facing) application will have a corresponding public Virtual IP (VIP) served by the site load balancer (SLB)

- The Web VM default gateway continues to be the NSX DLR, the N-S traffic from Web VM goes out via the NSX ESG; on the same host as the ESG, VXLAN traffic is also terminated. The traffic then subsequently traverses SLB which is advertising Web-VIP to the ESG (static route or dynamic routing) over Transit VLAN (the SLBs do **NOT** require VXLAN termination capabilities in this design)
- One or more Site Load Balancers (Active/Standby) are deployed at both the Protected and Recovery site serving the Public VIP of the Internet facing application and providing the Source Network Address Translation (SNAT) functionality
- The SLBs are deployed parallel to the NSX ESG with a common transit VLAN, the VXLAN is already terminated on the same host as the ESG is on before the packet hits the SLB (does NOT require any VXLAN Termination on the SLB)
- The Public VIP of the Web VM is advertised from both the sites, but GSLB is configured (Active/Standby) to resolve DNS entry of the Protected Site (20.1.1.1) to the incoming Client requests
- Upon complete application failover GSLB will resolve to Recovery Site DNS Entry (30.1.1.1), the failure detection and redirection mechanism depends on GSLB health check mechanism
- Unlike the route advertisement based approach this approach doesn't involve prefix lists or enabling/disabling UDLR LIF; the VIP is always advertised and GSLB is responsible for redirecting to the right location
- Except for Internet facing applications with VIP on SLB, all the other N-S traffic goes through NSX ESG and doesn't traverse SLB; this keeps the N-S traffic flow unchanged for all the other non-Internet facing applications

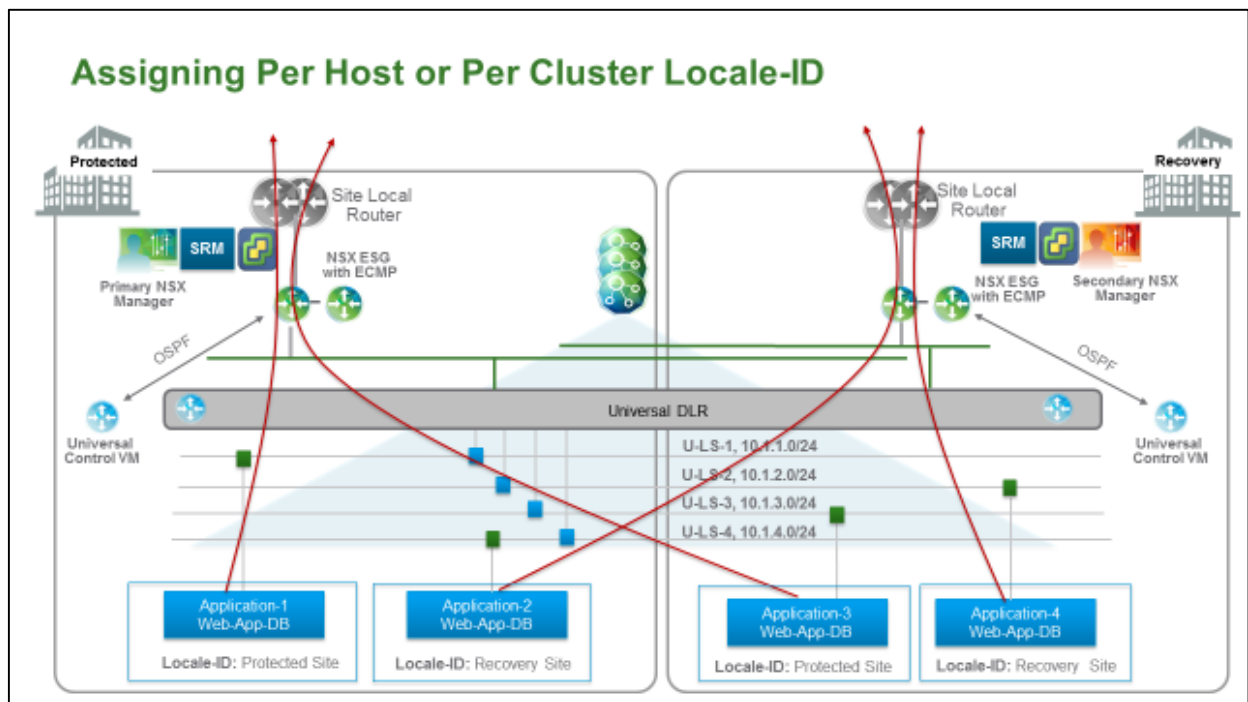


Additional Design Considerations

Following below is a summary of additional design considerations that should be kept in mind for a DR deployment with Cross-VC NSX.

- Some of the caveats associated with Cross-VC NSX impact the Disaster Recovery solution outlined above (Please refer to NSX 6.2 documentation)
- To support Cross-VC NSX and resulting VXLAN traffic between the two sites, a MTU size of 1600 is required for this solution
- The maximum latency between the sites must be under 150ms RTT
- The Universal Logical Switches cannot be bridged to physical workloads either using NSX bridging or third party VTEPs
- Universal DFW firewall rules can only use MAC Sets, IP Sets, and Security Groups containing MAC/IP Sets (Security Tags, VM Names etc. cannot be used)
- Third Party/Partner Services insertion is NOT supported for Universal Objects
- End point security services from the partners is NOT supported on Universal Logical Entities

- The route advertisement and GSLB based approach to control Ingress traffic are not mutually exclusive; the approaches will likely co-exist in most designs. Some North/South traffic will be protected via GSLB and the rest with route advertisements
- The maximum granularity of Locale-ID is on a per host basis (can be per Cluster as well); in a scenario where multiple applications are failing over on the same host (or Cluster if the Locale-ID is assigned at the Cluster level), all the applications MUST share a “locale” where all the N-S traffic will egress (this could be either Recovery or Protected site ESG, but cannot be both for the applications sharing host or a cluster with a single Locale-ID assignment)
- In the example below since Application 1 and Application 2 at the protected site have different Locale-IDs, these applications cannot share a host (or a cluster if the Locale-ID is assigned at a cluster level). Similarly, Applications 3 and Applications 4 cannot share a host (or a cluster if the Locale-ID is assigned at a cluster level) due to different Locale-ID requirements for each application



Run book Automation

The NSX and SRM based DR solution provides several avenues of run book automation using NSX REST API and SRM APIs.






















- The steps outlined to control traffic egress after application failover can be automated by using an NSX REST API based update to Locale-ID
- NSX component recovery described above can also be similarly automated by leveraging the NSX REST API
- SRM operations like Failover, Migration, Re-Protect and Failback can also be automated by using NSX REST API and SRM API in conjunction

NOTE:

It is beyond the scope of this document to outline a detailed DR Runbook Automation procedure. However, workflow automation tools such as VMware vRO and leveraging REST APIs are right steps towards further automation of the DR process outlined above.

NSX and SRM Integration

SRM 6.1 supports automatic network mapping, which allows the Protected site logical network to automatically map to the Recovery site logical network if the port groups are backed by the same Universal Logical Switch on both the sites. This automatic mapping of networks in SRM facilitates faster initial set-up and ability to build a large scale recovery plan without manually mapping each protected network. As shown below, Web, App and DB Universal Logical Switches (among others) are automatically mapped for failover/recovery.

comp-vc6.mylab.vmware.com		Actions
Summary	Monitor	Manage
Related Objects		
Network Mappings	Folder Mappings	Resource Mappings
Storage Policy Mappings	Placeholder Datastores	Advanced Settings
Permissions		
		
comp-vc6.mylab.vmware.com	dr-comp-vc6.mylab.vmware.com - Recovery Network	Reverse Mapping Exists
 VLAN-1020-DVS	 DR-VLAN-1010	Yes
 VLAN-3-DVS	 DR-VLAN3-DVS	Yes
 VLAN142-DVS	 DR-VLAN12-DVS	Yes
 VM Network	 VM Network	Yes
 vxx-dvs-18-universalwire-1-sid-900000-Universal-SW-01	 vxx-dvs-18-universalwire-1-sid-900000-Universal-SW-01	Yes
 vxx-dvs-18-universalwire-2-sid-900001-Universal-Web-LS	 vxx-dvs-18-universalwire-2-sid-900001-Universal-Web-LS	Yes
 vxx-dvs-18-universalwire-3-sid-900002-Universal-App-LS	 vxx-dvs-18-universalwire-3-sid-900002-Universal-App-LS	Yes
 vxx-dvs-18-universalwire-4-sid-900003-Universal-DB-LS	 vxx-dvs-18-universalwire-4-sid-900003-Universal-DB-LS	Yes
 vxx-dvs-18-universalwire-5-sid-900004-Universal-Transi...	 vxx-dvs-18-universalwire-5-sid-900004-Universal-Transi...	Yes
 vxx-dvs-18-universalwire-6-sid-900005-Universal-Transi...	 vxx-dvs-18-universalwire-6-sid-900005-Universal-Transi...	Yes

Recovering NSX Components – Detailed Steps

Initial State (Normal Function): All Cross-VC NSX Components are functional. The two NSX managers deployed and the Primary and Secondary roles have been set. The controller cluster is displayed in a normal state.

The screenshot displays the VMware vSphere Web Client interface. The left-hand 'Navigator' pane shows the 'NSX Managers' folder selected under 'Networking & Security Inventory'. The main content area is titled 'Installation' and contains two tables.

NSX Managers Table:

NSX Manager	Role	IP Address	vCenter	Version
10.100.1.72	Primary	10.100.1.72	10.100.1.71	6.2.2.3604087
10.200.1.72	Secondary	10.200.1.72	10.200.1.71	6.2.2.3604087

NSX Controller nodes Table:

Controller Node	NSX Manager	Managed By	Status	Peers	Software Version
10.100.1.75 controller-7	10.200.1.72	10.100.1.72	✓ Connected		6.2.46427
10.100.1.74 controller-8	10.200.1.72	10.100.1.72	✓ Connected		6.2.46427
10.100.1.73 controller-9	10.200.1.72	10.100.1.72	✓ Connected		6.2.46427
10.100.1.75 controller-12	10.100.1.72	10.100.1.72	✓ Connected	■ ■	6.2.46427
10.100.1.73 controller-10	10.100.1.72	10.100.1.72	✓ Connected	■ ■	6.2.46427
10.100.1.74	10.100.1.72	10.100.1.72	✓ Connected	■ ■	6.2.46427

Site Failure Event: In the event of a site failure and the Primary NSX Manager and Universal Controller Cluster is not available, the logical networking elements are set to a read only mode and no new logical components can be modified or deployed from the NSX managers.

Recovery Step 1:

From the Secondary NSX Manager, disconnect from the Primary NSX Manager

The screenshot shows the VMware vSphere Web Client interface. The left sidebar contains a 'Navigator' pane with 'Networking & Security' expanded to 'NSX Managers'. The main content area shows the 'NSX Managers' configuration page. A table lists the NSX Managers, with one entry at IP 10.200.1.72. Below it, the 'NSX Controller nodes' table is visible, showing three nodes: controller-1 (10.100.1.75, Connected), controller-2 (10.100.1.74, Disconnected), and controller-3 (10.100.1.73, Connected). The 'Disconnected' status for controller-2 indicates it is not connected to the primary NSX Manager.

Click YES when prompted.

10.200.1.72 - Disconnect from Primary NSX Manager

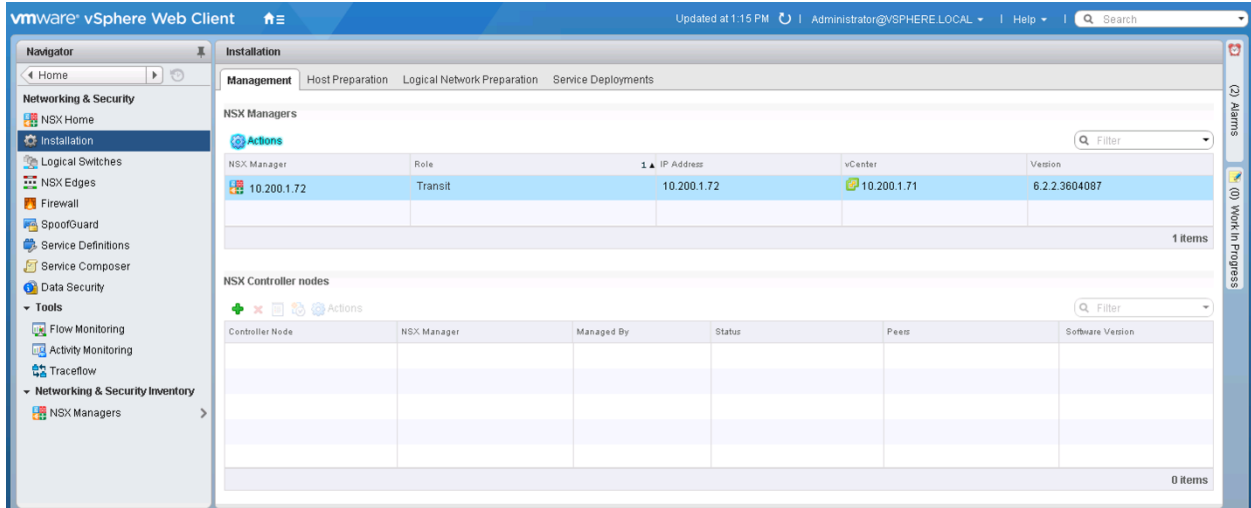


This operation will disconnect 10.200.1.72 from its primary NSX Manager.
Do you want to continue?

Yes

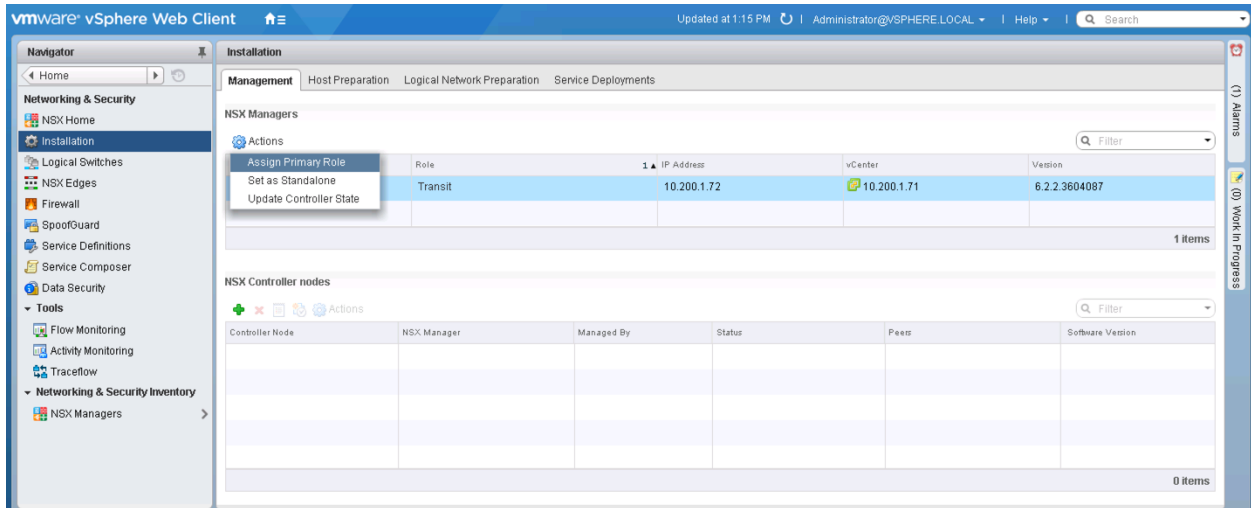
No

The NSX Manager's new role will become **Transit**



Recovery Step 2

Now assign the primary role to the NSX manager marked as **Transit**



Click **YES** when prompted.

10.200.1.72 - Assign Primary Role



This operation will assign primary role to 10.200.1.72.
Do you want to continue?

Yes No

Recovery Step 3

Re-deploy the Universal Controller Cluster at the Recovery site. Below is a screenshot of deploying one controller. This would be repeated until all three controllers are deployed.

Add Controller ?

NSX Manager: * 10.200.1.72

Datcenter: * SanJose

Cluster/Resource Pool: * Edge Cluster

Datastore: * EMC_VNX_1-1

Host: 10.200.1.50

Folder

Connected To: * Edge_Mgmt [Change](#) [Remove](#)

IP Pool: * NSX Controllers [Select](#)

Password: *

Confirm password: *

[OK](#) [Cancel](#)

Below is a screenshot once all three controllers are deployed.

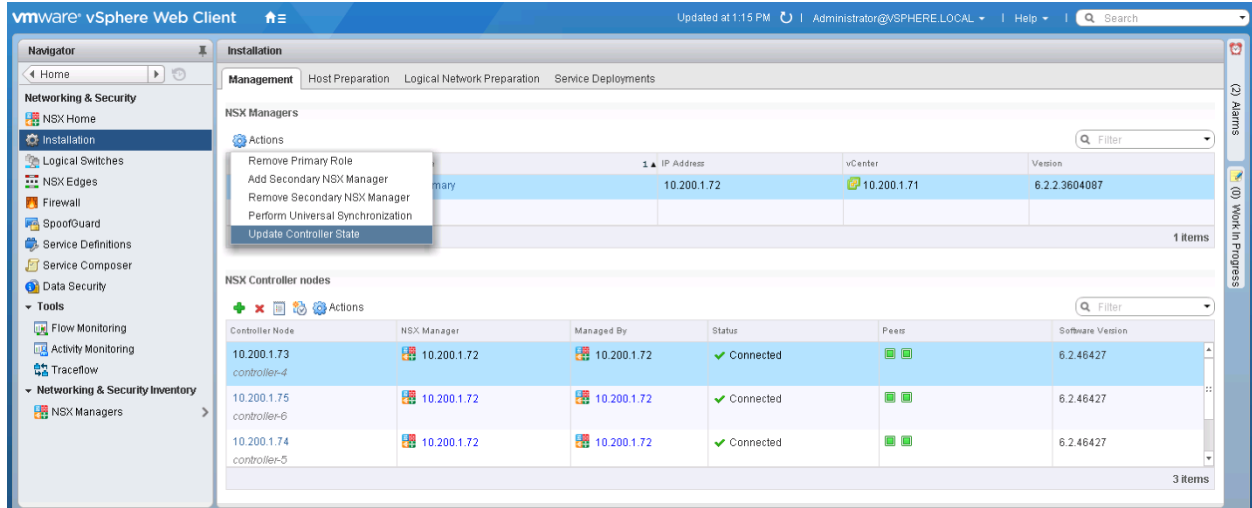
The screenshot shows the VMware vSphere Web Client interface. The left sidebar contains the 'Navigator' with 'Networking & Security' expanded to 'NSX Home' and 'Installation'. The main content area is titled 'Installation' and shows 'Management' with tabs for 'Host Preparation', 'Logical Network Preparation', and 'Service Deployments'. Under 'NSX Managers', there is a table with one entry: NSX Manager 10.200.1.72, Role Primary, IP Address 10.200.1.72, vCenter 10.200.1.71, and Version 6.2.2.3604087. Below this is a section for 'NSX Controller nodes' with a table showing three nodes: 10.200.1.73 (controller-4), 10.200.1.75 (controller-6), and 10.200.1.74 (controller-5). All three nodes are managed by 10.200.1.72 and have a status of 'Connected'. The software version for all is 6.2.46427.

NSX Manager	Role	IP Address	vCenter	Version
10.200.1.72	Primary	10.200.1.72	10.200.1.71	6.2.2.3604087

Controller Node	NSX Manager	Managed By	Status	Peers	Software Version
10.200.1.73 controller-4	10.200.1.72	10.200.1.72	✓ Connected	■ ■	6.2.46427
10.200.1.75 controller-6	10.200.1.72	10.200.1.72	✓ Connected	■ ■	6.2.46427
10.200.1.74 controller-5	10.200.1.72	10.200.1.72	✓ Connected	■ ■	6.2.46427

Recovery Step 4

Update Controller State

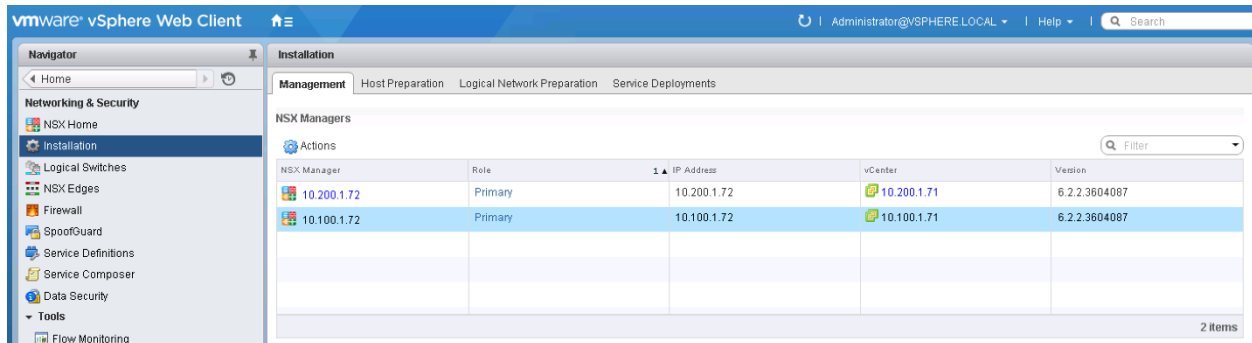


Recovery Step 5

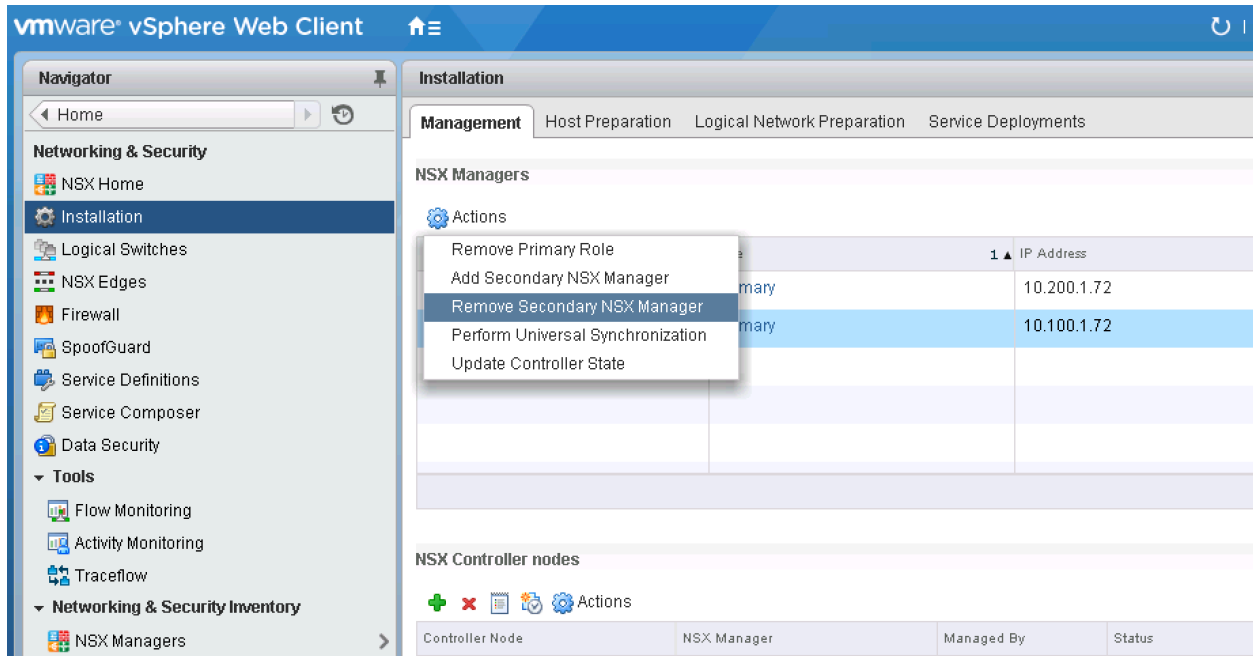
If running a deployment model leveraging routing for Egress failover via one Control VM, please follow the steps outlined in [Recovering Control VM](#) section to recover the Control VM. These steps are NOT needed in a Local Egress (with Locale-ID) deployment where a Control VM is deployed at both sites.

Recovery Step 6 (Once the Primary Site recovers)

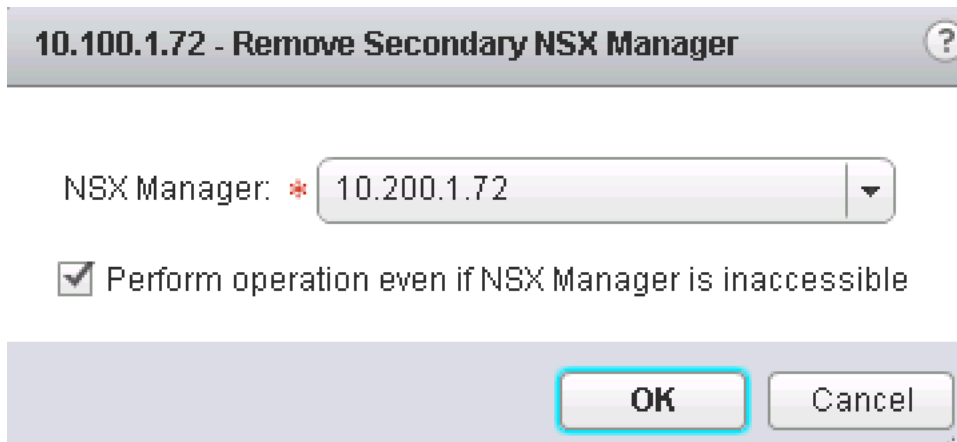
You will initially see that both NSX Managers are assigned the Primary role.



Force the removal of the registered secondary NSX Manager on the Protected site NSX Manager before removing the primary role.

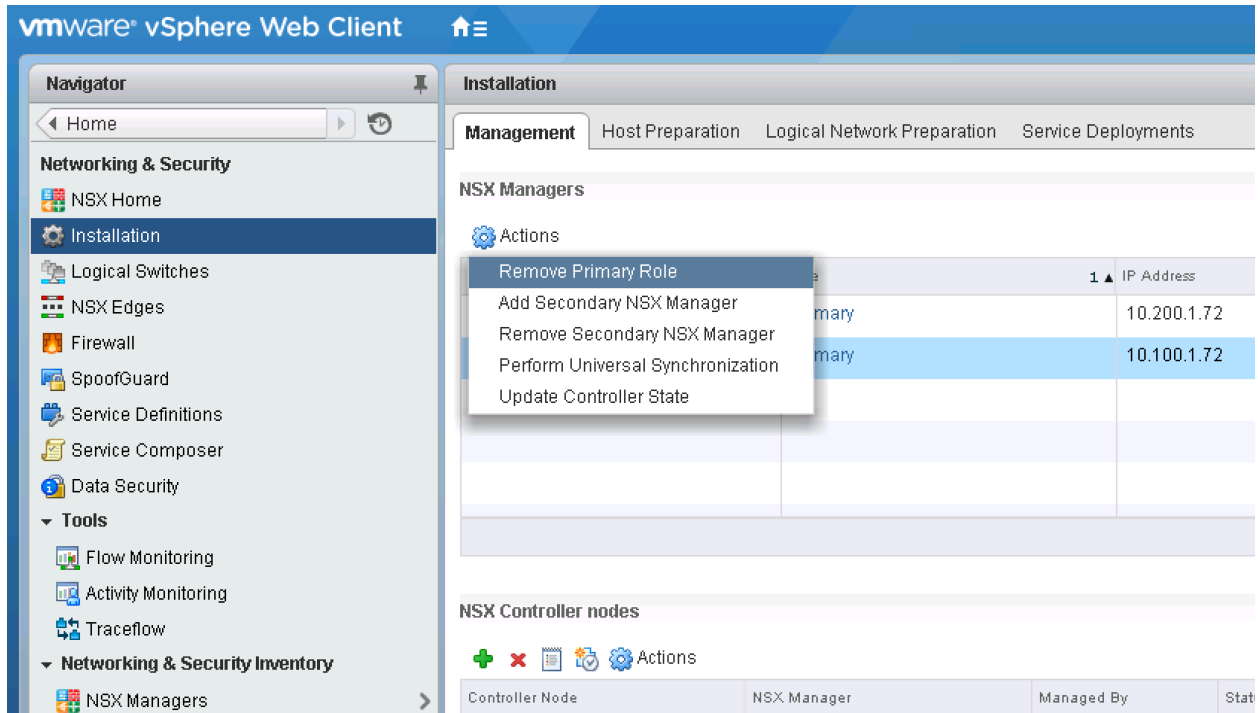


NOTE: Make sure to check the box for Perform operation even if the NSX manager is inaccessible



Recovery Step 7

Demote the original primary site NSX Manager to transit role. Select the original NSX manager and select **Actions** and **Remove Primary Role**



The screenshot shows the VMware vSphere Web Client interface. The left sidebar contains the 'Navigator' with 'Networking & Security' expanded to 'Installation'. The main content area shows the 'Installation' section with a 'Management' tab. Below this is a table of 'NSX Managers'. The table has columns for 'Controller Node', 'NSX Manager', 'Managed By', and 'Stat'. The row for IP address 10.100.1.72 is selected, and a context menu is open over it, with 'Remove Primary Role' highlighted. Below the table is a section for 'NSX Controller nodes' with a table header and an 'Actions' button.

Controller Node	NSX Manager	Managed By	Stat
	primary		
	primary		

Click Yes to proceed with the removal of the primary role assignment.

10.100.1.72 - Remove Primary Role



This operation will remove primary role from 10.100.1.72.
Do you want to continue?

Yes

No

The original Primary NSX Manager is now in transit mode.

The screenshot shows the VMware vSphere Web Client interface. The left sidebar contains the 'Networking & Security' menu with 'Installation' selected. The main content area is titled 'Installation' and has tabs for 'Management', 'Host Preparation', 'Logical Network Preparation', and 'Service Deployments'. Under 'NSX Managers', there is a table with the following data:

NSX Manager	Role	IP Address
10.200.1.72	Primary	10.200.1.72
10.100.1.72	Transit	10.100.1.72

Below this table is the 'NSX Controller nodes' section, which includes a table with columns for 'Controller Node', 'NSX Manager', 'Managed By', and 'Status'.

Recovery Step 8:

Clean up old controllers on the primary site, delete all three of the old controllers at the Protected site. These would get redeployed once you assign the primary role back to the NSX Manager at the Protected site

The screenshot shows a detailed view of the 'NSX Controller nodes' table. A dialog box titled 'Delete Controller' is open, asking 'Do you really want to delete the selected controller?' with 'Yes' and 'No' buttons. The table below has the following data:

Controller Node	NSX Manager	Peers	Software Version
10.200.1.73 controller-24	10.200.1.72	Connected	6.2.46427
10.200.1.74 controller-25	10.200.1.72	Connected	6.2.46427
10.200.1.75 controller-26	10.200.1.72	Connected	6.2.46427
10.100.1.73 controller-10	10.100.1.72	Connected	6.2.46427
10.100.1.74 controller-11	10.100.1.72	Connected	6.2.46427
10.100.1.75	10.100.1.72	Connected	6.2.46427

The 'controller-10' row is highlighted in blue, indicating it is the selected controller for deletion.

Recovery Step 9:

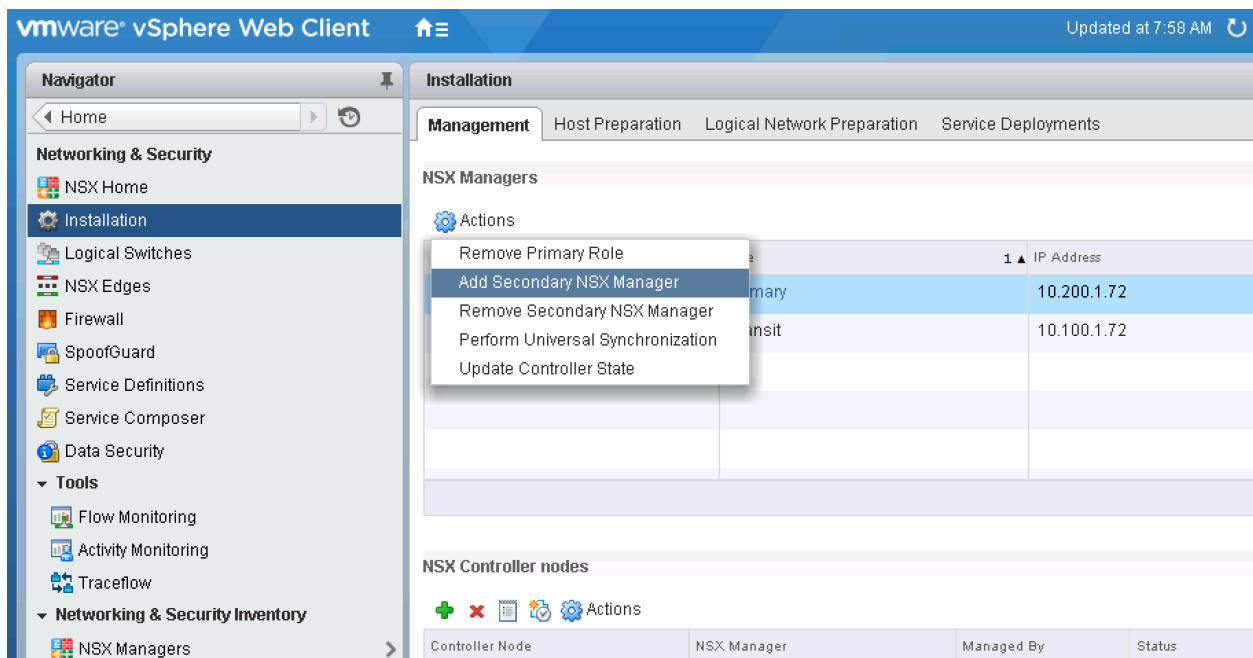
Also, depending on deployment model, if needed, delete the Control VM at the Protected site as it has successfully been redeployed at the new primary site.

Currently, there is a bug that prevents the deletion of the Control VM from the GUI, so a REST API call must be made to do the deletion

DELETE /api/4.0/edges/<edge-id>/appliances/<ha-index>

Recovery Step 10:


Now add the original Primary NSX Manager as a secondary role to the newly promoted Primary NSX Manager by selecting actions and clicking **Add Secondary NSX manager**



The screenshot shows the VMware vSphere Web Client interface. The left sidebar contains the 'Networking & Security' section with 'Installation' selected. The main content area is titled 'Installation' and shows the 'NSX Managers' configuration page. A context menu is open over the 'NSX Managers' table, with 'Add Secondary NSX Manager' highlighted. The table lists two NSX Manager nodes: one with IP address 10.200.1.72 and another with IP address 10.100.1.72. Below the table, there is a section for 'NSX Controller nodes' with a table header including 'Controller Node', 'NSX Manager', 'Managed By', and 'Status'.

Controller Node	NSX Manager	Managed By	Status
	primary		
	nsit		

Enter the NSX Manager credentials.

10.200.1.72 - Add Secondary NSX Manager 

NSX Manager: *

User Name: *


Password: *

Confirm password: *

Accept the SSL certificate to proceed.

Trust Certificate?

NSX Manager 10.100.1.72 presented an SSL certificate with the following thumbprint:

 05:46:8B:DF:03:74:C0:31:AC:0D:CD:8D:BB:12:3B:8C:39:F5:BA:61

Proceed with this certificate?

We can now see the NSX Managers in their newly assigned roles

The screenshot displays the VMware vSphere Web Client interface. The left sidebar shows the 'Networking & Security Inventory' with 'NSX Managers' selected. The main content area is titled 'Installation' and contains two tables.

NSX Managers Table:

NSX Manager	Role	IP Address	vCenter	Version
10.200.1.72	Primary	10.200.1.72	10.200.1.71	6.2.2.3604087
10.100.1.72	Secondary	10.100.1.72	10.100.1.71	6.2.2.3604087

NSX Controller nodes Table:

Controller Node	NSX Manager	Managed By	Status	Peers	Software Version
10.200.1.73 controller-24	10.200.1.72	10.200.1.72	✓ Connected	■ ■	6.2.46427
10.200.1.74 controller-25	10.200.1.72	10.200.1.72	✓ Connected	■ ■	6.2.46427
10.200.1.75 controller-26	10.200.1.72	10.200.1.72	✓ Connected	■ ■	6.2.46427
10.200.1.73 controller-14	10.100.1.72	10.200.1.72	✓ Connected		6.2.46427
10.200.1.74 controller-13	10.100.1.72	10.200.1.72	✓ Connected		6.2.46427
10.200.1.75	10.100.1.72	10.200.1.72	✓ Connected		6.2.46427

All NSX components have now been recovered; the NSX Primary Manager is now running at the Recovery site and the newly deployed Universal Controller Cluster is also running at the recovery site.

Recovering NSX Components – Failback Procedure

Step 1. Delete the Universal Controller Cluster at the secondary site from the newly promoted Primary NSX manager

Step 2. Remove the Secondary NSX manager from the Primary NSX Manager

Step 3. Assign the Primary role to the NSX Manager now in transit mode located at the original Protected site

Step 4. Remove the primary role from the NSX Manager located at the Recovery site and it will be placed into transit mode

Step 5. Add as secondary the newly demoted NSX Manager currently in transit mode located at the Recovery site to the newly promoted Primary NSX Manager located at the Protected site

Step 6. Redeploy the Universal Controller Cluster from the Primary NSX Manager onto the Protected site

Step 7. Update Controller State

Step 8. Depending on deployment model, redeploy the Control VM at the Protected site and delete from Recovery site

Recovering Control VM

This is needed only for deployments leveraging routing for Egress failover using single Control VM. These deployments would require redeploying the Control VM upon site failure as described below. This step is NOT required for Local Egress (with Locale-ID) where Control VM is already deployed at both sites.

Steps for redeploying Control VM:

- First get the UDLR XML response with following **GET**
GET <https://<NSX Manager IP>/api/4.0/edges/<edge-id>> You can use a browser-based client such as RESTClient on Firefox shown below to make the call.

The screenshot shows the RESTClient interface in Firefox. The 'Request' section is expanded, showing a GET method to the URL `https://10.200.1.72/api/4.0/edges/edge-2167d8fa-915d-48b7-a7c1-620906013091`. The 'Headers' section contains two entries: 'Authorization: Basic YWRtaW46Vkl1...' and 'Content-Type: application/xml'. The 'Body' section is empty. The 'Response' section is expanded, showing the following headers:

```
1. Status Code      : 200 OK
2. Cache-Control   : private, no-cache
3. Content-Type    : application/xhtml+xml
4. Date            : Fri, 01 Apr 2016 13:46:22 GMT
5. Expires         : Thu, 01 Jan 1970 00:00:00 GMT
6. Set-Cookie     : JSESSIONID=37A775BE80C37CB584E3C3D3B567CEDE; Path=/; Secure; HttpOnly
7. Transfer-Encoding : chunked
```

There are a few fields that need to be filled out in the body of this XML; the fields that need to be entered are represented as **moref** variables in the below figure. The respective values for these variables can be found via the **Managed Object Browser**: <https://<vCenter IP Address>/mob>

```
<edge>
  <!--remaining edge details snipped-->

  <appliances>
    <appliance>
      <resourcePoolId>$cluster-moref</resourcePoolId>
      <datastoreId>$datastore-moref</datastoreId>
    </appliance>
  </appliances>

  <mgmtInterface>
    <addressGroups>
    </addressGroups>
    <mtu>1500</mtu>
    <connectedToId>$network-moref</connectedToId>
  </mgmtInterface>

  <!--remaining edge details snipped-->
</edge>
```

- Once the XML body is modified, a **PUT** NSX REST API call needs to be made to redeploy the Universal Control VM.
PUT <https://<NSX Manager IP> /api/4.0/edges/<edge-id>>
- Note, the redeployment of Universal Control VM will not recover the routing protocol configuration. This will need to be re-entered or could be inserted as part of the **PUT** NSX REST API call.

References

[NSX Design Guide](#)

[NSX 6.2 Documentation](#)

[Vmware Site Recovery Manager Documentation](#)

[Vmware vRealize Orchestrator Documentation](#)