Backup/DR Services - Nutanix Disaster Recovery (Policy Driven DR / Run Books)

[PDF generated August 10 2026. For all recent updates please see the Nutanix Bible releases notes located at https://nutanixbible.com/release_notes.html. Disclaimer: Downloaded PDFs may not always contain the latest information.]

Test Drive

For those who are interested in getting hands on, take it for a spin with Nutanix Test Drive!

https://www.nutanix.com/test-drive-disaster-recovery

The Nutanix Disaster Recovery feature set provides policy driven backup, DR and run book automation services configured via Prism Central (PC). This capability builds upon and extends the native DR and replications features that have been available in AOS and configured in PE for years. For more information on the actual back-end mechanism being leveraged for replication, etc. refer to the 'Backup and Disaster Recovery (DR)' section in the 'AOS' section.

Supported Configurations

The solution is applicable to the configurations below (list may be incomplete, refer to documentation for a fully supported list):

Core Use Case(s):

- Policy based backups and replication
- DR run book automation
- DRaaS (Managed Service Providers)

Management interfaces(s):

- Prism Central (PC)

Supported Environment(s):

- On-Prem:
  - AHV
  - ESXi
- Cloud:
  - NC2 (AWS & Azure)
  - MST (Pilot Light or Zero Compute)

Upgrades:

- Part of AOS

Compatible Features:

- AOS BC/DR features

Key terms

The following key terms are used throughout this section and defined in the following:

- Recovery Point Objective (RPO)
  - Refers to the acceptable data loss in the event of a failure. For example, if you want an RPO of 1 hour, you'd take a snapshot every 1 hour. In the event of a restore, you'd be restoring data as of up to 1 hour ago. For synchronous replication typically an RPO of 0 is achieved.

- Recovery Time Objective (RTO)
  - Recovery time objective. Refers to the period of time from failure event to restored service. For example, if a failure occurs and you need things to be back up and running in 30 minutes, you'd have an RTO of 30 minutes.

- Recovery Point
  - A restoration point aka snapshot.

Implementation Constructs

Within Nutanix Disaster Recovery, there are a few key constructs:

Protection Policy

- Key Role: Backup/Replication policy for assigned categories
- Description: A protection policy defines the RPO (snapshot frequency), recovery location (remote cluster), snapshot retention (local vs. remote cluster), and associated categories. With Protection Policies everything is applied at the category level (with a default that can apply to any/all). This is different from Protection Domains where you have to select VM(s).

The following image shows the structure of the Protection Policy:

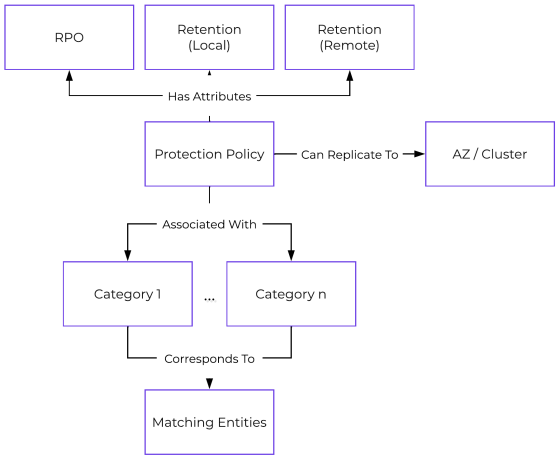

DR - Protection Policy

Recovery Plan

- Key Role: DR run book
- Description: A Recovery Plan is a run book that defines the power on sequencing (can specify categories or VMs/VGs) and network mapping (primary vs. recovery and test failover / failback). This is most synonymous with what people would leverage SRM for. NOTE: a Protection Policy must be configured before a Recovery Plan can be configured. This is necessary as the data must exist at the recovery site in order for it to be recovered.

The following image shows the structure of the Recovery Plan:

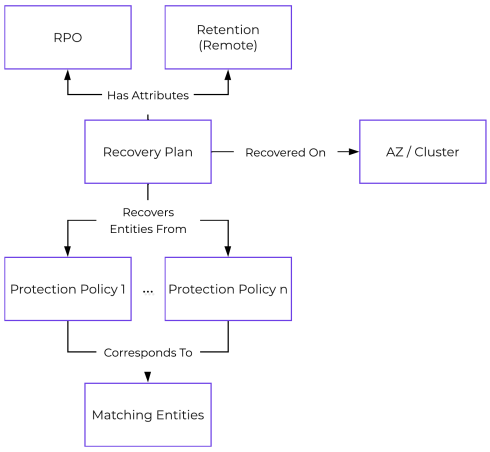

DR - Recovery Plan

Linear Retention Policy

- Key Role: Recovery Point retention policy
- Description: A linear retention policy specifies the number of recovery points to retain. For example, if the RPO is 1 hour and your retention is set to 10, you'd keep 10 hours (10 x 1 hour) of recovery points (snapshots).

Roll-up Retention Policy

- Key Role: Recovery Point retention policy
- Description: A roll-up retention policy will "roll-up" snaps dependent on the RPO and retention duration. For example, if the RPO is 1 hour and your retention is set to 5 days it'll keep 1 day of hourly and 5 days of the latest daily recovery points. The logic can be characterized as follows: If retention is n days, keep 1 day of RPO and n days of latest daily recovery points. If retention is n weeks, keep 1 day of RPO and 1 week of daily and n latest weekly recovery points. If retention is n months, keep 1 day of RPO and 1 week of daily and 1 month of weekly and n months of monthly recovery points. If retention is n years, keep 1 day of RPO, 1 week of daily, 1 month of weekly, 1 year of monthly recovery points and n latest yearly recovery points.

Linear vs. roll-up retention

Use linear policies for small RPO windows with shorter retention periods or in cases where you always need to be able to recover to a specific RPO window.

Use roll-up policies for anything with a longer retention period. They're more flexible and automatically handle snapshot aging / pruning while still providing granular RPOs for the first day.

The following shows a high-level overview of the Nutanix Disaster Recovery constructs:

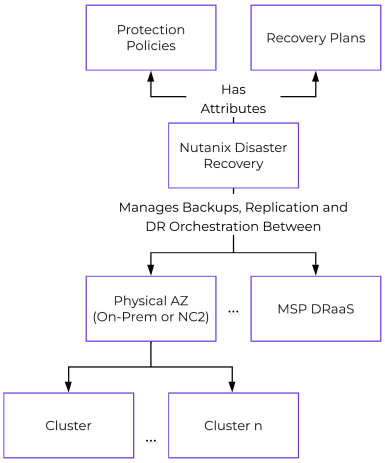

DR - Overview

Usage and Configuration

The following sections cover how to configure and leverage Nutanix Disaster Recovery.

The high-level process can be characterized into the following high-level steps:

1. Connect to Availability Zones (AZs)
2. Configure Protection Policies
3. Configure Recovery Plan(s)
4. Perform/Test Failover & Failback

Connect Availability Zone(s)

The first step is connecting to an AZ which can be a another on-premises AZ or in the cloud on NC2.

In PC, search for 'Availability Zones' or navigate to 'Administration' -> 'Availability Zones':

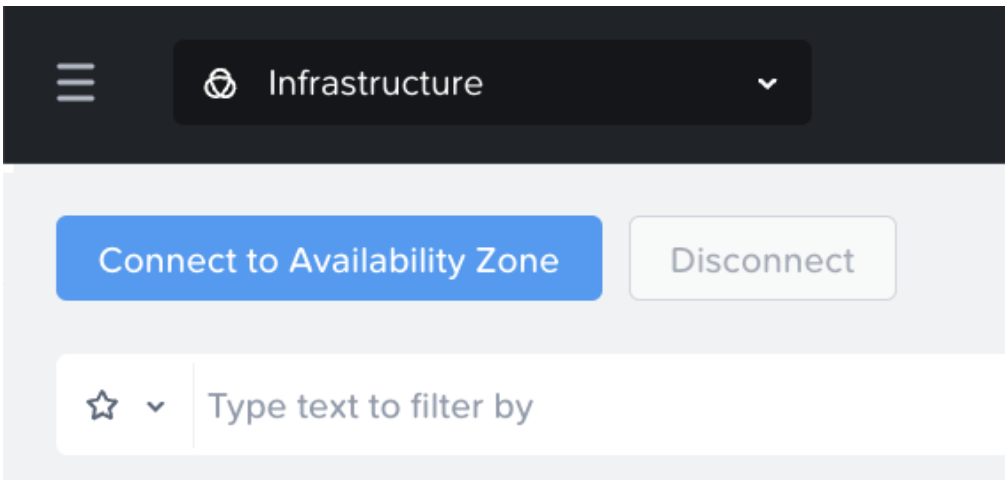

DR - Connect to Availability Zone

Click on 'Connect to Availability Zone' and select the AZ Type ('Xi' or 'Physical Location' aka PC instance):

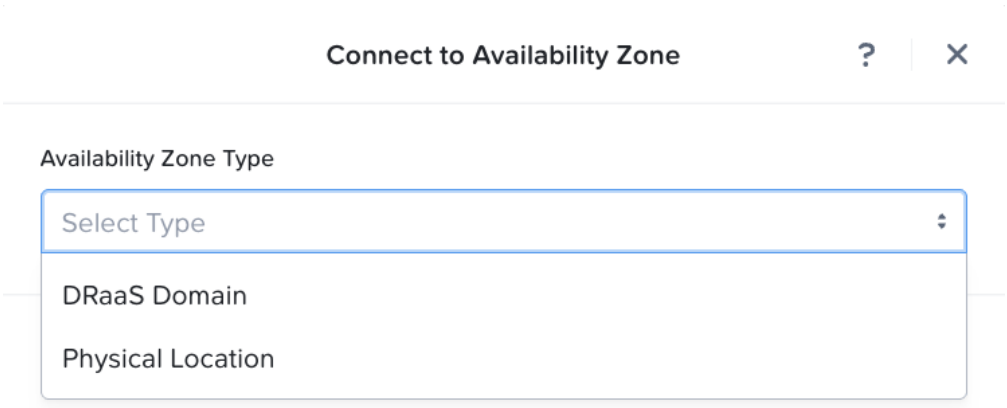

DR - Connect to Availability Zone

Input credentials for PC or Xi and click 'Connect':

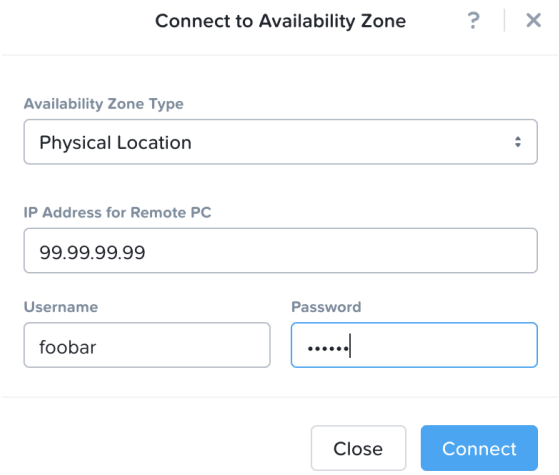

DR - Connect to Availability Zone

The connected AZ will now be displayed and be available.

Configure Protection Policies

In PC, search for 'Protection Policies' or navigate to 'Policies' -> 'Protection Policies':

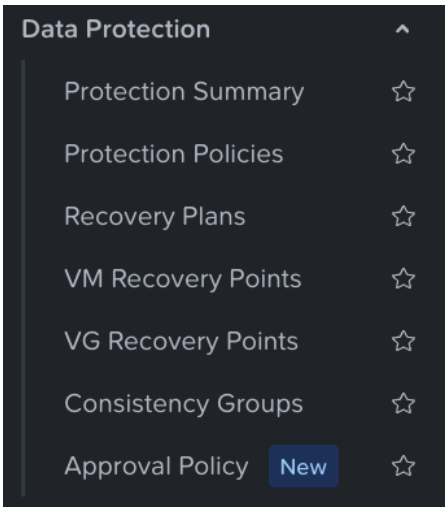

DR - Protection Policies

Click on 'Create Protection Policy':

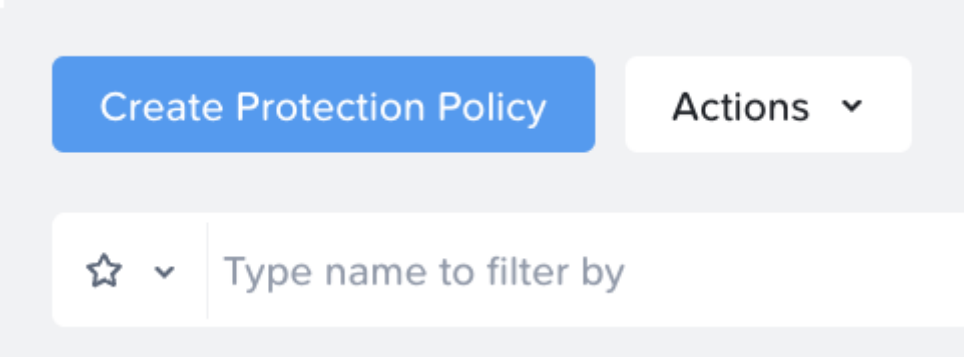

DR - Create Protection Policy

Input the desired name, and select source AZ and cluster then click save. Select recovery location AZ and cluster, then click save:

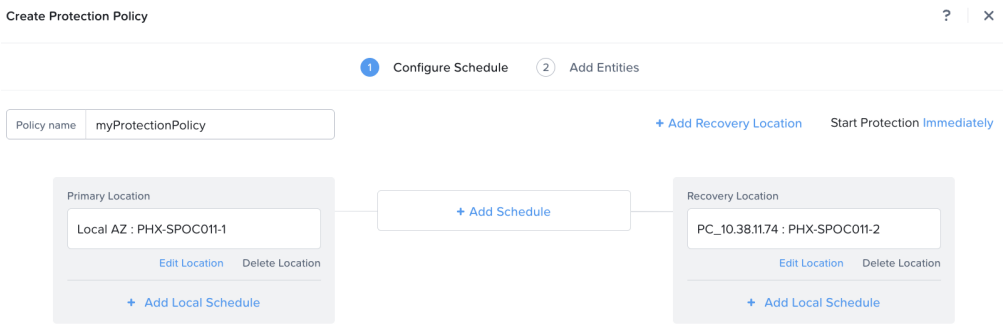

DR - Protection Policy Inputs

Click Add Schedule between the Primary and Recovery locations to add frequency of snapshots and replication, retention and other details.

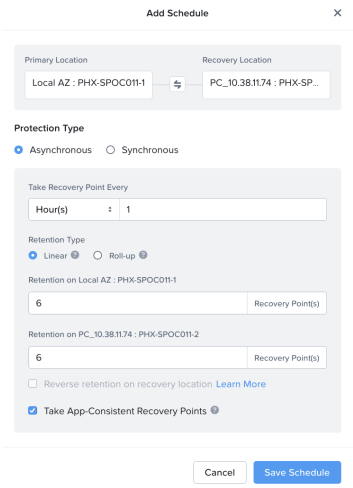

DR - Protection Policy Inputs

NOTE: for DRaaS you don't need select a 'Target Cluster':

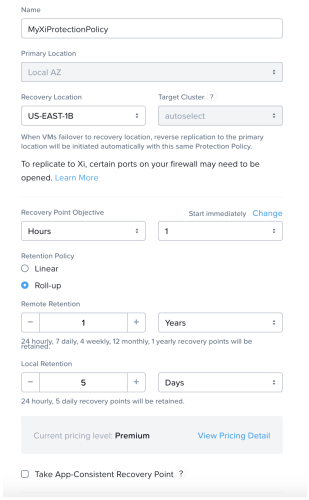

DR - Protection Policy Inputs - DRaaS

Next we'll select the categories for the policy to apply to:

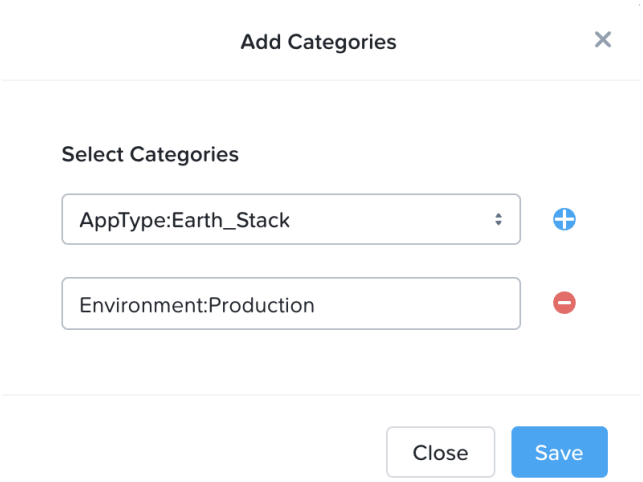

DR - Protection Policy Categories

Click 'Save' and you will now see the newly created Protection Policy:

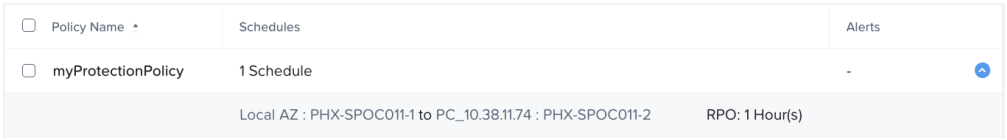

DR - Protection Policies

Configure Recovery Plans

In PC, search for 'Recovery Plans' or navigate to 'Policies' -> 'Recovery Plans':

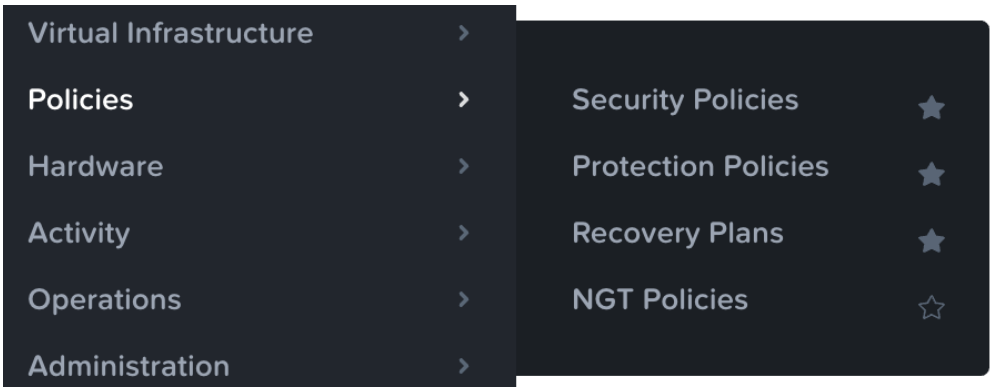

DR - Recovery Plans

On the first launch you will be greeted with a screen to create the first Recovery Plan:

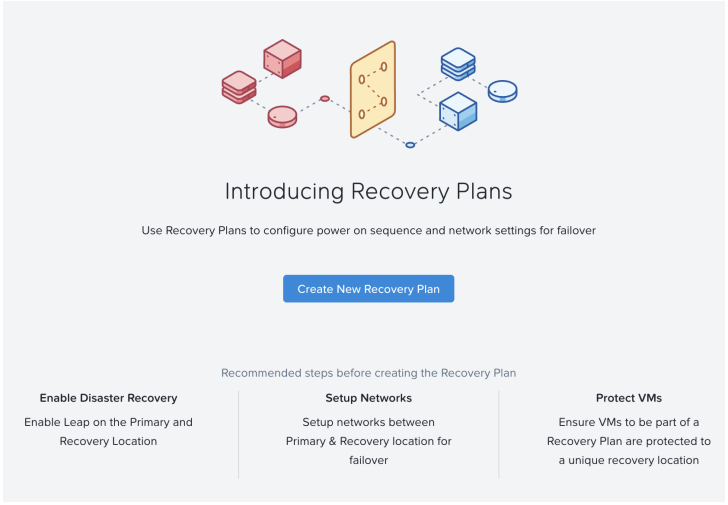

DR - Create Recovery Plan

Select the 'Recovery Location' using the drop down:

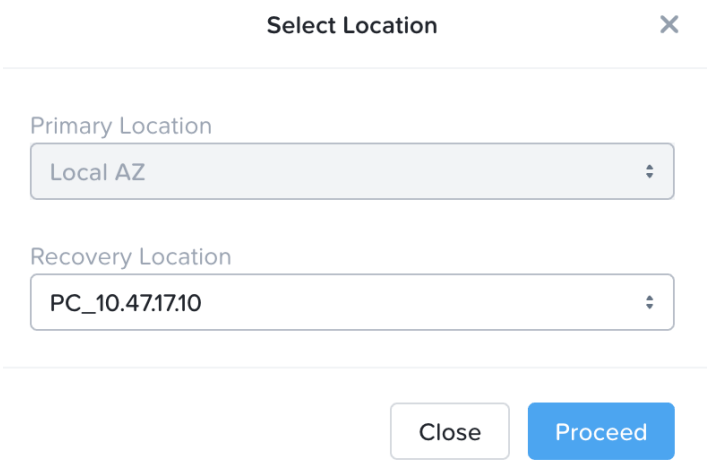

DR - Select Recovery Location

NOTE: This can be either a Xi AZ or Physical AZ (PC with corresponding managed clusters).

Input the Recovery Plan name and description and click 'Next':

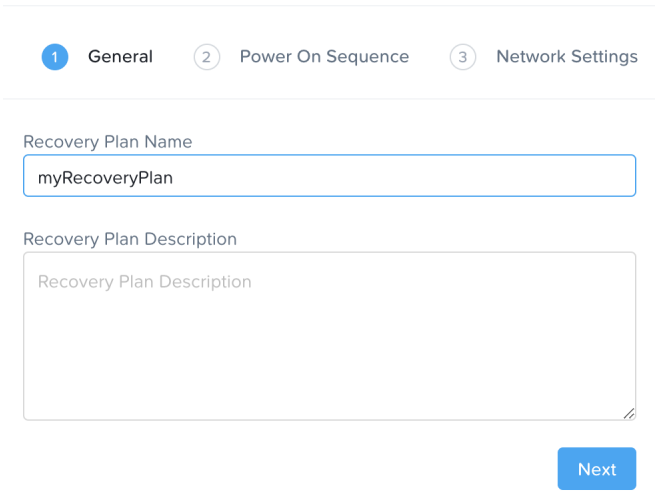

DR - Recovery Plan - Naming

Next click on 'Add Entities' and specify the power on sequence:

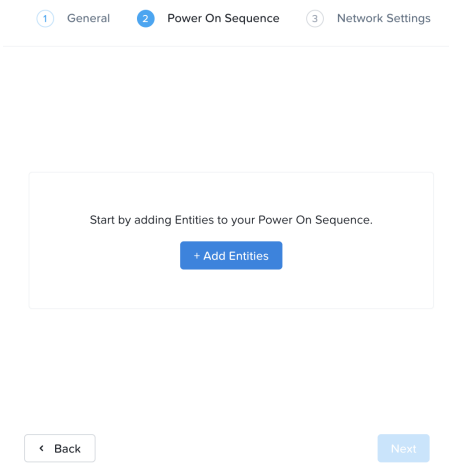

DR - Recovery Plan - Power On Sequence

Search for VMs or Categories to add to each stage:

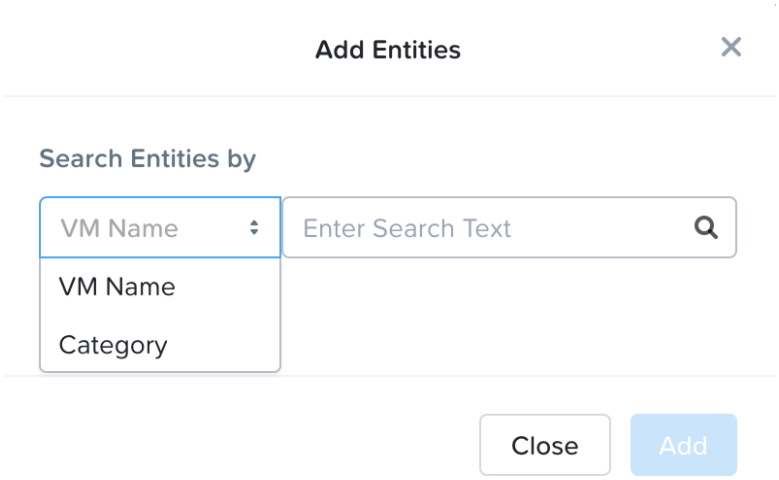

DR - Recovery Plan - Power On Sequence

Once the power on sequence looks good with the stages, click 'Next':

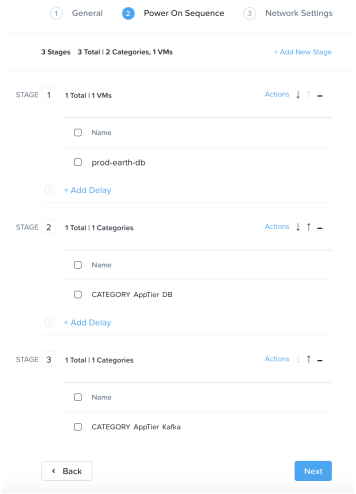

DR - Recovery Plan - Power On Sequence

Power On Sequencing

When determining the power on sequence you will want to stage things as follows:

- Stage 0: Core services (AD, DNS, etc.)
- Stage 1: Services dependent on Stage 0 services, and required for Stage 2 services (e.g. DB Tier)
- Stage 2: Services dependent on Stage 1 services, and required for Stage 3 services (e.g. App Tier)
- Stage 3: Services dependent on Stage 2 services, and required for Stage 4 services (e.g. Web Tier)
- Stage 4-N: Repeat based upon dependencies

We will now map the network between our source and target environments:

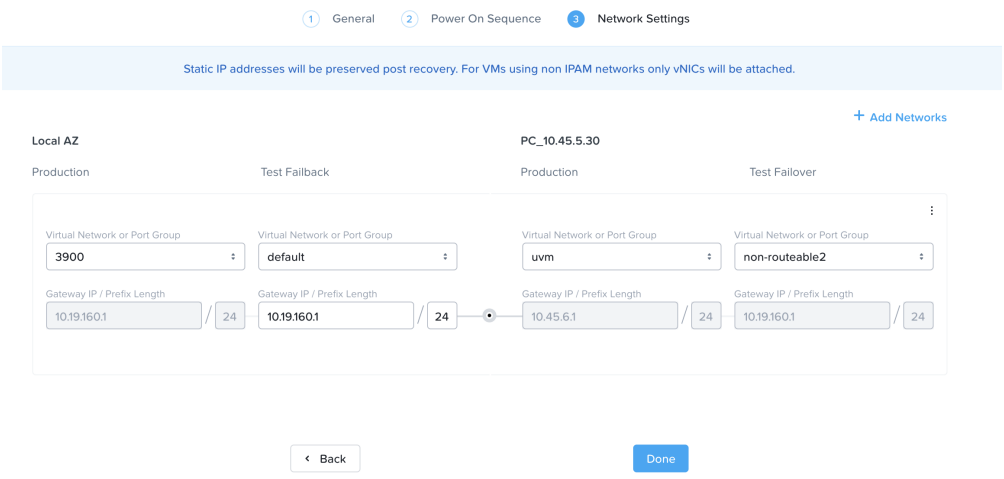

DR - Recovery Plan - Network Mapping

Failover / Failback Networks

In most cases you will want to use a non-routable or isolated network for your test networks. This will ensure you don't have any issues with duplicate SIDs, arp entries, etc.