

# DELL EMC VxRAIL™ vSAN STRETCHED CLUSTERS PLANNING GUIDE

## **ABSTRACT**

This planning guide provides best practices and requirements for using stretched clusters with VxRail appliances.

January 2019

# TABLE OF CONTENTS

- Intended Use and Audience ..... 3
- Overview ..... 3
  - vSphere & vSAN ..... 3
  - Fault Domains ..... 4
  - VxRail Cluster Nodes ..... 4
    - VxRail Cluster Deployment Options..... 4
  - Witness Host ..... 4
- VxRail Cluster Requirements ..... 4
  - vCenter Server Requirements ..... 5
    - Customer Supplied vCenter Server Requirements..... 6
- Networking & Latency ..... 6
  - Layer 2 and Layer 3 Support ..... 6
  - Supported Geographical Distances ..... 7
  - Data Site to Data Site Network Latency ..... 7
  - Data Site to Data Site Bandwidth ..... 7
  - Data Site to Witness Network Latency ..... 7
  - Data Site to Witness Network Bandwidth ..... 8
  - Inter-site MTU consistency..... 8
  - Connectivity ..... 8
- Conclusion ..... 8
- Appendix A: VxRail Stretched Cluster Setup Checklist..... 9
- Appendix B: VxRail Stretched Cluster Open Port Requirements .....10

## Intended Use and Audience

This guide is intended for customers, Dell EMC and Business Partner Sales teams, and implementation professionals to understand the requirements for Stretched Cluster support with the Dell EMC VxRail Appliance. Services from Dell EMC or an Authorized VxRail Services Partner are required for implementation of Stretched Clusters.

This document is not intended to replace the implementation guide or to bypass the service implementation required for Stretched Clusters. Customers who attempt to set-up Stretch Clusters on their own will invalidate support.

Additionally, customers will need to contact support to facilitate upgrades. Customer driven upgrades of VxRail Stretched Cluster implementations are not permitted.

## Overview

This planning guide provides best practices and requirements for using stretched cluster with a VxRail Appliance. This guide assumes the reader is familiar with the [vSAN Stretched Cluster Guide](#). This guide is for use with a VxRail Appliance only.

The vSAN Stretched Cluster feature creates a stretched cluster between two geographically separate sites, synchronously replication data between sites. This feature allows for an entire site failure to be tolerated. It extends the concept of fault domains to data center awareness domains.

VxRail 4.5.070 introduced vSAN 6.6 which includes local site protection and site affinity for Stretched Clusters allowing unbalanced configurations. The following is a list of the terms used for vSAN Stretched Clusters:

- *Preferred/Primary site* – one of the two data sites that is configured as a vSAN fault domain.
- *Secondary site* – one of the two data sites that is configured as a vSAN fault domain.
- *Witness host* – a dedicated ESXi host or vSAN witness appliance that is host to the witness component that coordinates data placement between the preferred and secondary site and assists in the failover process. This is the third fault domain.

The vSAN Storage Policies that impact the VxRail Cluster configuration are:

- *Primary Failures to Tolerate (PFTT)<sup>1</sup>/Failures to Tolerate (FTT)* – for stretched clusters this rule has two possible values: 0 ensures protection on a single site; 1 enables protection across sites.
- *Secondary Failures to Tolerate (SFTT) (only applicable starting with vSAN 6.6/VxRail 4.5.070)* – the rule that defines the number of host and device failures that a virtual machine object can tolerate in the local site. Possible values: 0,1,2,3.
- *Failure Tolerance Method*- either RAID-1 (mirroring) used when performance is important or starting with vSAN 6.6/VxRail 4.5.070, RAID-5/6 (erasure coding) used when capacity is important. For stretched clusters, this only applies to the Secondary Failures to Tolerate setting. This is the local file protection mode.
- *Affinity (only applicable starting with vSAN 6.6/VxRail 4.5.070)* - this policy is applicable when PFTT is set to 0. It is set to Preferred or Secondary to determine which sites stores the vSAN object.

## vSphere & vSAN

For vSAN stretched cluster functionality on VxRail, vSphere Distributed Resource Scheduler (DRS) is required. DRS will provide initial placement assistance, and automatically migrate virtual machines to their corrected site in accordance with the Host/VM affinity rules. It can also help locate virtual machines to their correct site when a site recovers after a failure.

<sup>1</sup> Prior to vSAN 6.6 (a component of VxRail 4.5.70), this was referred to as Failures to Tolerate (FTT).

## Fault Domains

Fault domains (FD) provide the core functionality of vSAN Stretched Cluster. The maximum number of fault domains in a vSAN Stretched Cluster is 3. The first Fault Domain can be referred as “Preferred” data site, the second Fault Domain can be referred as “Secondary” data site, and the third Fault Domain is the witness host site. It is important to keep utilization per data site below 50% to ensure proper availability should either the Preferred or Secondary site go offline.

## VxRail Cluster Nodes

vSAN Stretched Clusters are deployed across 2 sites in an Active/Active configuration. An identical number of ESXi hosts is required prior to vSAN 6.6/VxRail 4.5.070 to ensure a balanced distribution of resources. Starting with vSAN 6.6/VxRail 4.5.070, unbalanced configurations are supported; however, it is a best practice to have an identical number of ESXi hosts across the 2 sites. VM/Host Affinity rules must be set for an unbalanced configuration.

Each data site is configured as a Fault Domain. An externally available third site houses a Witness appliance, which makes up the third Fault Domain.

## VxRail Cluster Deployment Options

A Customer must plan the VxRail Stretched Cluster deployment prior to installation. Depending on the number of nodes in the VxRail Cluster, a customer can:

- Deploy up to 16 nodes, 8 per site, on initial deployment or
- Deploy the minimum number of nodes per site for initial deployment and then scale out additional nodes either at installation or during the VxRail Stretched Cluster life cycle.

## Witness Host

Each vSAN Stretched Cluster configuration requires a Witness host. The Witness must reside on a third site that has independent paths to each data site. While the Witness host must be part of the same vCenter as the hosts in the data sites, it must **not** be on the same cluster as the data site hosts. The Witness ESXi OVA is deployed using a virtual standard switch (vSS).

A vSAN Witness Appliance, or a physical host, can be used for the Witness function. The vSAN Witness Appliance includes licensing, while a physical host would still need to be licensed accordingly.

**NOTE:** If you are using the Witness host OVA file it comes with a license. Thus it will not consume a vSphere license. However, if you are using a physical host, it will require a vSphere license.

## VxRail Cluster Requirements

This section describes the requirements necessary to implement vSAN stretched clusters in a VxRail Cluster.

- The VxRail Cluster must be deployed across 2 physical sites in an Active/Active configuration.
- The VxRail Cluster must be VxRail 3.5 release or higher.
- For VxRail 3.5, 4.0 and 4.5.0, each data site must have an identical number of nodes.
- Starting with vSAN 6.6/VxRail 4.5.070, we recommend each data site have an identical number of nodes, but it is not required.
- Failure Tolerance Method of RAID-5/6, available starting with vSAN6.6/VxRail 4.5.070, the configuration must be all-flash.
- The maximum supported configuration is 15+15+1 (30 nodes+1 witness).
- The minimum number of nodes is dependent on the VxRail Version and Stretched Cluster configuration. See Table 1.

VxRail Version		Minimum Nodes Preferred Site + Secondary Site + Witness
VxRail 3.5		4 + 4 + 1
VxRail 4.0.x and 4.5.0		3 + 3 + 1
VxRail 4.5.070 and beyond  <b>NOTE:</b> This is configuration dependent on the values set for PFTT, SFTT, and Failure Tolerance Method.	PFTT = 1; SFTT=1; Failure Tolerance Method=RAID-1 (Mirroring)	3 + 3 + 1
	PFTT = 1; SFTT=2; Failure Tolerance Method=RAID-1 (Mirroring)	5 + 5 + 1
	PFTT = 1; SFTT=3; Failure Tolerance Method=RAID-1 (Mirroring)	7 + 7 + 1
	PFTT = 1; SFTT=1; Failure Tolerance Method=RAID-5/6 (Erasure Coding)	4 + 4 + 1
	PFTT = 1; SFTT=2; Failure Tolerance Method=RAID-5/6 (Erasure Coding)	6 + 6 + 1

**Table 1 VxRail Version Minimum # of Nodes per Site**

- A witness host **must** be installed on a separate site as part of the installation engagement. See Table 2 for version compatibility.

VxRail Version	Witness Host OVA Version
VxRail v3.5	OVA Version 6.2
VxRail v4.0.x	OVA Version 6.2
VxRail v4.5.x	OVA Version 6.5
VxRail v4.7.x	OVA Version 6.7

**Table 2 VxRail/Witness Host OVA Compatibility Chart**

## vCenter Server Requirements

Prior to VxRail 4.5.200, only a Customer Supplied vCenter can be used for stretched clusters. Starting with VxRail 4.5.200, either a VxRail vCenter Server or a Customer Supplied vCenter Server can be used for stretched clusters. However, please refer to the [VxRail vCenter Server Planning Guide](#) for caveats of using VxRail vCenter Server.

Customer Supplied vCenter Server Appliance is the recommended choice.

## Customer Supplied vCenter Server Requirements

The following are the Customer Supplied vCenter Server requirements:

- The customer must provide the vSphere Enterprise Plus license.
- The Customer Supplied vCenter Server version must be in the [VxRail and external vCenter interoperability matrix](#). In addition, the ESXi version of the cluster hosting the Customer Supplied vCenter must be identical to the ESXi host version of the VxRail Cluster. Check the VxRail Release Notes for to determine the proper version numbers.
  - VxRail 3.5 and vSphere 6.0, version details can be found in [VxRail Appliance Software 3.5 Release Notes](#).
  - VxRail 4.0.x and vSphere 6.0, version details can be found in [VxRail Appliance Software 4.0.x Release Notes](#).
  - VxRail 4.5.x and vSphere 6.5, version details can be found in [VxRail Appliance Software 4.5.x Release Notes](#).
  - VxRail 4.7.x and vSphere 6.7, version details can be found in [VxRail Appliance Software 4.7.x Release Notes](#).

To join the Customer Supplied vCenter Server you will need:<sup>2</sup>

- Know whether your Customer Supplied vCenter Server has an embedded or non-embedded Platform Services Controller. If the PSC is non-embedded, you will need the PSC FQDN.
- Know the Customer Supplied vCenter Server FQDN.
- Know the Customer Existing Single Sign-on domain (SSO) (For example vsphere.local)
- Create or select a datacenter on the Customer Supplied vCenter Server for the VxRail Cluster to join.
- Specify the name of the cluster that will be created by VxRail in the selected datacenter when the cluster is built. It will also be the name of the distributed switch. This name must be unique and not used anywhere in the datacenter on the Customer Supplied vCenter Server.
- Verify the customer DNS server can resolve all VxRail ESXi hostnames prior to deployment.
- Create or re-use<sup>3</sup> a VxRail management user and password for this VxRail cluster on the Customer Supplied vCenter Server. The user created must be:
  - Created with no permissions
  - Created with no roles assigned to it
- (Optional) Create a VxRail admin user and password for VxRail on the Customer Supplied vCenter Server.

## Networking & Latency

### *Layer 2 and Layer 3 Support*

A stretched cluster in VxRail requires Layer 2 connectivity between data sites. Connectivity between the data sites and the witness must be Layer 3. Figure 1 illustrates a supported configuration.

<sup>2</sup> See the [DELL EMC VxRail vCenter Server Planning Guide](#) for more detailed information.

<sup>3</sup> If a previous VxRail Cluster has been deployed on the Customer Supplied vCenter Server, the VxRail Management User can be re-used if the customer chooses.

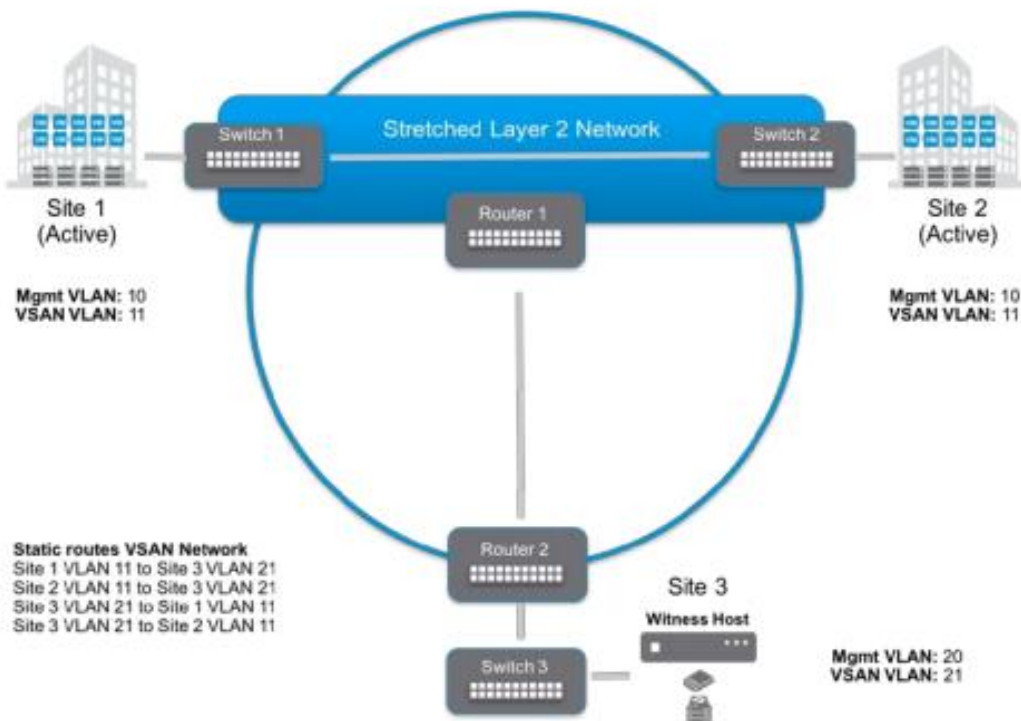


Figure 1. VxRail Supported Topology

## Supported Geographical Distances

For vSAN Stretched Clusters, support is based on network latency and bandwidth requirements, rather than distance. The key requirement is the actual latency numbers between sites.

## Data Site to Data Site Network Latency

Latency or RTT (Round Trip Time) between sites hosting virtual machine objects should not be greater than 5 msec (< 2.5 msec one-way).

## Data Site to Data Site Bandwidth

Bandwidth between sites hosting virtual machine objects will be workload dependent. For most workloads, VMware recommends a minimum of 10Gbps or greater bandwidth between sites.

## Data Site to Witness Network Latency

In most vSAN Stretched Cluster configurations, latency or RTT (Round Trip Time) between sites hosting VM objects and the witness nodes should not be greater than 200 msec (100 msec one-way).

The latency to the witness is dependent on the number of objects in the cluster. VMware recommends that on vSAN Stretched Cluster configurations up to 10+10+1, a latency of less than or equal to 200 milliseconds is acceptable, although if possible, a latency of less than or equal to 100 milliseconds is preferred. For configurations that are greater than 10+10+1, VMware recommends a latency of less than or equal to 100 milliseconds is required.

## *Data Site to Witness Network Bandwidth*

Bandwidth between sites hosting VM objects and the witness nodes are dependent on the number of objects residing on vSAN. It is important to size data site to witness bandwidth appropriately for both availability and growth. A standard rule of thumb is 2Mbps for every 1000 objects on vSAN.

## *Inter-site MTU consistency*

It is important to maintain a consistent MTU (maximum transmission unit) size between data nodes and the witness in a Stretched Cluster configuration. Ensuring that each VMkernel interface designated for vSAN traffic, is set to the same MTU size will prevent traffic fragmentation. The vSAN Health Check checks for a uniform MTU size across the vSAN data network, and reports on any inconsistencies.

## *Connectivity*

- Management network: connectivity to all 3 sites
- VM network: connectivity between the data sites (the witness will not run virtual machines that are deployed on the vSAN cluster)
- vMotion network: connectivity between the data sites (virtual machines will never be migrated from a data host to the witness host)
- vSAN network: connectivity to all 3 sites

## *Witness Traffic Separation*

With the release of vCenter Server 6.7u1, it has become possible to avail of [Witness Traffic Separation \(WTS\)](#) for VxRail Stretched Cluster deployments. This allows an alternate VMkernel interface to be designated to carry traffic destined for the Witness rather than the vSAN tagged VMkernel interface. This feature allows for more flexible network configurations by allowing separate networks for node-to-node and node-to-witness traffic. From a routing perspective, this will allow two independent subnets/routes to be advertised from each Data Node site to the Witness site.

## *Conclusion*

In short, the vSAN stretched cluster feature is available in VxRail Appliance Release 3.5 and later. It creates a stretched cluster between two geographically separate sites, synchronously replicating data between sites, and enabling enterprise-level availability. The stretched cluster feature allows for an entire site failure to be tolerated, with no data loss and near zero downtime.



## Appendix A: VxRail Stretched Cluster Setup Checklist

Required Reading	<ul style="list-style-type: none"> <li>✓ Read the <a href="#">VMware vSAN Stretched Cluster Guide</a>.</li> <li>✓ Read the <a href="#">VxRail vCenter Server Planning Guide</a>.</li> </ul>
VxRail Version	<ul style="list-style-type: none"> <li>✓ The minimum version is VxRail 3.5</li> <li>✓ No mixed clusters are supported (i.e., VxRail 4.5 and 4.0 in the same cluster)</li> </ul>
vSphere License	<ul style="list-style-type: none"> <li>✓ vSphere Enterprise Plus license is required</li> <li>✓ You cannot reuse the VxRail vCenter Server license on any other deployments.</li> </ul>
Number of Nodes	<ul style="list-style-type: none"> <li>✓ Review Table 1 in this guide for the minimum number of nodes.</li> <li>✓ The maximum supported configuration is 15+15+1 (30 nodes+1 witness).</li> </ul>
Customer Supplied vCenter Server (Recommended choice and required prior to VxRail 4.5.200)	<ul style="list-style-type: none"> <li>✓ Required prior to VxRail 4.5.200.</li> <li>✓ The Customer Supplied vCenter Server version must be in the <a href="#">VxRail and external vCenter interoperability matrix</a>.</li> </ul>
Fault Domains	<ul style="list-style-type: none"> <li>✓ Must have 3 Fault Domains (preferred, secondary, and witness host)</li> </ul>
Network Topology	<ul style="list-style-type: none"> <li>✓ vSAN traffic between the data sites must be Layer 2.</li> <li>✓ vSAN traffic between the witness host and the data sites must be Layer 3.</li> </ul>
Data Site to Data Site Network Latency	<ul style="list-style-type: none"> <li>✓ Latency or RTT between data sites should not be greater than 5 msec. (&lt;2.5 msec one-way)</li> </ul>
Data Site to Data Site Bandwidth	<ul style="list-style-type: none"> <li>✓ A minimum of 10Gbps is required.</li> </ul>
Data Site to Witness Network Latency	<ul style="list-style-type: none"> <li>✓ For configurations up to 10+10+1, latency or RTT less than or equal to 200 msec is acceptable, but 100 msec is preferred.</li> <li>✓ For configuration greater than 10+10+1, latency or RTT less than or equal to 100 msec is required.</li> </ul>
Data Site to Witness Network Bandwidth	<ul style="list-style-type: none"> <li>✓ The rule of thumb is 2Mbps for every 1000 objects on vSAN.</li> </ul>
Inter-site MTU consistency	<ul style="list-style-type: none"> <li>✓ Required to be consistent between data sites.</li> </ul>
Network Ports	<ul style="list-style-type: none"> <li>✓ Review <a href="#">Appendix B</a> for required port connectivity.</li> </ul>

## Appendix B: VxRail Stretched Cluster Open Port Requirements

The following table lists the open port requirements for a VxRail Stretched Cluster.

Description	Connectivity To/From	L4 Protocol	Port
vSAN Clustering Service	vSAN Hosts	UDP	12345, 2345
vSAN Transport	vSAN Hosts	TCP	2233
vSAN VASA Vendor Provider	vSAN Hosts and vCenter	TCP	8080
vSAN Unicast Agent (to Witness Host)	vSAN Hosts and vSAN Witness Appliance	UDP	12321



[Learn more](#) about Dell EMC VxRail Appliances



[Contact](#) a Dell EMC Expert



[View more](#) resources



Join the conversation  
[@DellEMC\\_CI](#)  
with [#VxRail](#)