



aCloud

Stretched Cluster Configuration Guideline

Version 5.8.7R1



Change Log

Date	Change Description
May 2, 2019	Version 5.8.7R1 document release.

CONTENT

Chapter 1 Requirement	1
1.1 Deployment Requirement.....	1
1.2 Topology	1
Chapter 2 Create Virtual Datastore.....	2
2.1 Create stretched datastore	2
Chapter 3 Run VM on fault domain	7
3.1 Create new virtual machine	7

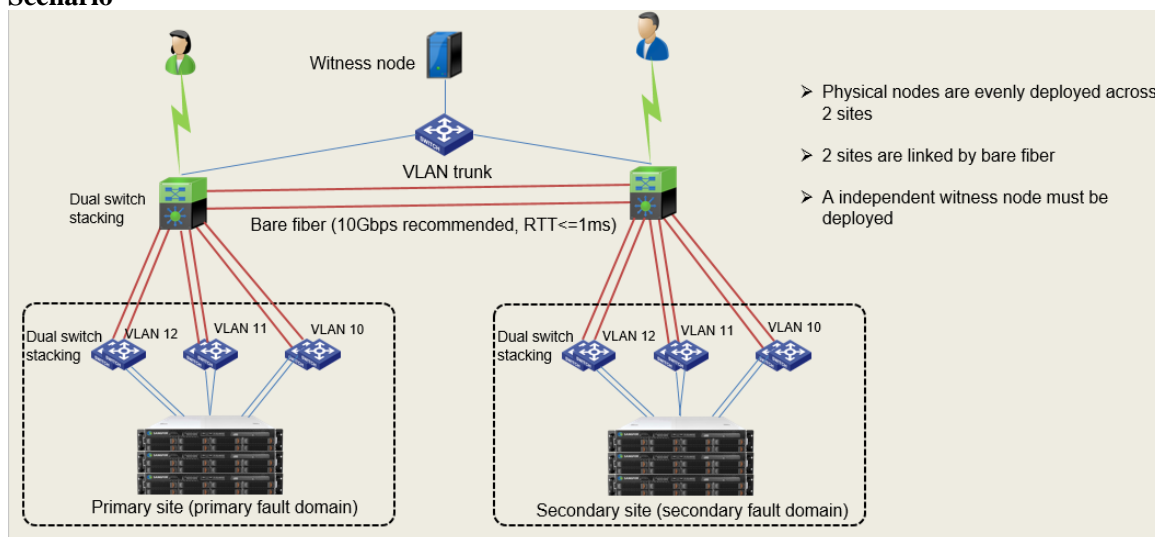
Chapter 1 Requirement

1.1 Deployment Requirement

- **Deployment Scale**
 - Only 1 stretched virtual volume is supported by 1 cluster
 - Starts with 4 nodes, maximum 24 nodes
- **Site link**
 - 10Gbps bare fiber is recommended, $RTT \leq 1\text{ms}$
 - 1Gbps is supported for small deployment (4-6 nodes)
- **Witness link**
 - Recommended $RTT \leq 1\text{ms}$ (fluctuation within 5ms is accepted)
 - 100Mbps is recommended
- **Supported scenarios**
 - Different floors inside the same building
 - 2 adjacent buildings
 - Inside the same campus
- **Scenarios to be supported**
 - 3 centers across 2 cities
 - Metropolitan active-active

1.2 Topology

Scenario



Chapter 2 Create Virtual Datastore

2.1 Create stretched datastore

1. Log in to the aCloud web console and go to [Storage] – [Virtual Storage] – [New]
Select **Stretched datastore** and click “Next” to proceed next page.

Create Virtual Datastore

1 Basics — 2 Select Node — 3 Specify Fault Domains — 4 Use of Disk — 5 Confirm

Virtual Datastore Name:

Type:

☐ Ordinary datastore

☒ **Stretched datastore**

A stretched datastore is applicable to active-active data center scenario where nodes are deployed in two server rooms. Associate the nodes with the stretched datastore and add nodes in server rooms to primary and secondary fault domain respectively.

Data Replicas:

☒ **2 replicas**

All data has two replicas stored on two fault domains respectively. If one server room (fault domain) encounters failure, there is a replica in the other server room so that the business will not be interrupted.

Next Cancel

2. Select node IP and click “Next” to proceed add to virtual datastore.

Create Virtual Datastore

✓ Basics — 2 Select Node — 3 Specify Fault Domains — 4 Use of Disk — 5 Confirm

Select and Add Node to Virtual Datastore :

ⓘ A stretched cluster should have at least 4 and even number of nodes, each fault domain having the same number of nodes. 4 node(s) selected

<input checked="" type="checkbox"/>	Node Name	Node IP	SSDs	HDDs
<input checked="" type="checkbox"/>	192.168.20.181	192.168.20.181	1	2
<input checked="" type="checkbox"/>	192.168.20.182	192.168.20.182	1	2
<input checked="" type="checkbox"/>	192.168.20.36	192.168.20.36	1	2
<input checked="" type="checkbox"/>	192.168.20.37	192.168.20.37	1	2

Back Next Cancel

- Click on “Add Witness Node”, input witness node IP and password and press OK.

The screenshot shows the 'Create Virtual Datastore' wizard at step 3, 'Specify Fault Domains'. A modal dialog titled 'Add Witness Node' is open, allowing the user to add a witness node. The dialog contains the following fields:

- Node IP: 192.200.19.225
- Username: admin
- Password: (masked with dots)

The dialog has 'OK' and 'Cancel' buttons. In the background, the main window shows a list of nodes with a red arrow pointing to the 'Add Witness Node' button.

- Input name for primary fault domain and secondary fault domain. (eg: Room1 and Room2)

The screenshot shows the 'Create Virtual Datastore' wizard at step 3, 'Specify Fault Domains'. The 'Add Witness Node' button is disabled. The 'Nodes' list on the left shows four selected nodes: 192.168.20.181, 192.168.20.182, 192.168.20.36, and 192.168.20.37. The 'Primary Fault Domain' is set to 'Room1' and the 'Secondary Fault Domain' is set to 'Room2'. Both domains have a 'None' button below them. The wizard has 'Back', 'Next', and 'Cancel' buttons at the bottom.

5. Select nodes and click on right arrow to add to fault domain. Click “Next” to proceed to next steps.

Note: Each fault domain must have same number of node.

The screenshot shows the 'Specify Fault Domains' step of the 'Create Virtual Datastore' wizard. The progress bar at the top indicates steps: Basics (checked), Select Node (checked), Specify Fault Domains (active), Use of Disk (disabled), and Confirm (disabled). The 'Node IP' is 192.200.19.225, and there is an 'Add Witness Node' button. The 'Nodes' list is currently empty, showing 'None'. To the right, there are two fault domain sections. The 'Primary Fault Domain' is 'Room1' and contains two nodes: 192.168.20.181 and 192.168.20.182, both with checked selection boxes. The 'Secondary Fault Domain' is 'Room2' and contains two nodes: 192.168.20.36 and 192.168.20.37, both with checked selection boxes. Red arrows point from the 'Nodes' list area to the right arrow buttons next to each fault domain list. At the bottom are 'Back', 'Next', and 'Cancel' buttons.

6. Verify number of nodes under fault domain and press “OK”.

The screenshot shows the 'Confirm' dialog box for the Virtual Datastore configuration. It contains a warning message: 'Please make sure that the configuration of fault domain is in accordance with actual host deployment in server room. Once the datastore is created, you cannot change the fault domain to which the host belongs. Improper configuration will invalidate active-active data center, causing replicas of a piece of data written to a same data center or other undesired results.' Below the message, it displays the 'Node IP: 192.200.19.230'. The 'Primary Fault Domain' is 'Room1' and lists nodes 192.168.20.181 and 192.168.20.182. The 'Secondary Fault Domain' is 'Room2' and lists nodes 192.168.20.36 and 192.168.20.37. At the bottom are 'OK' and 'Cancel' buttons.

7. Configure use of disk, each node must have cache disk and data disk.

Note: Stretched virtual datastore have to configure witness node **quorum disk**. Each quorum disk require at least 100GB.

Create Virtual Datastore ✕

✓ Basics — ✓ Select Node — ✓ Specify Fault Domains — 4 Use of Disk — 5 Confirm

Expand All Collapse All Edit About Disk Grouping ?

Disk Group	Disk	Type	Disk Size	Use of Disk
Group 1	Disk 0	SSD	223.57 GB	Cache disk
	Disk 2	HDD	1.82 TB	Data disk
	Disk 1	HDD	1.82 TB	Data disk

▼ Node: 192.200.19.230 (Witness Node) Quorum Disks: 1 Free: 1

Disk	Type	Disk Size	Use of Disk
6000c290846d590b41c06cf6be258e9a VMware Virtu...	SSD	500 GB	Quorum Disk
6000c29c445ca982b7c1063a3e6bbf9c VMware Virtua...	SSD	500 GB	Free

Back Next Cancel

8. Verify total size of virtual datastore and click “OK” to proceed.

Create Virtual Datastore ✕

✓ Basics — ✓ Select Node — ✓ Specify Fault Domains — ✓ Use of Disk — 5 Confirm

Confirm Configuration of Virtual Datastore (VirtualDatastore1):

7.24 TB Available Space **14.55 TB** Total Space **4** Nodes **2** Replicas

Witness Node:

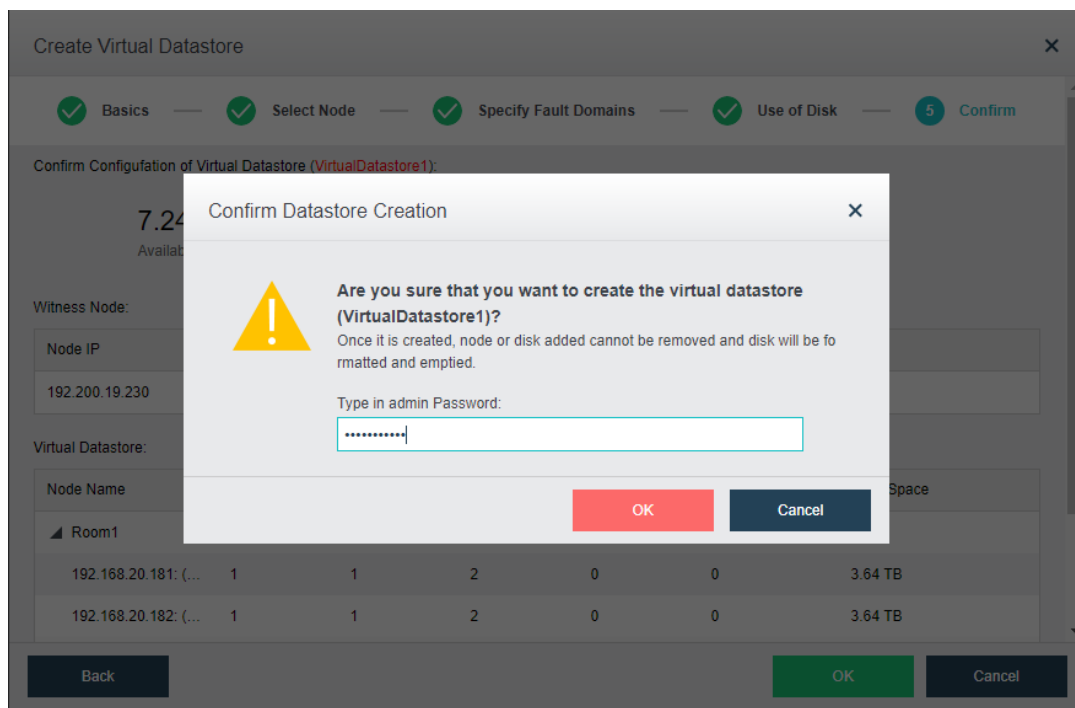
Node IP	Quorum Disk Type	Quorum Disk Capacity
192.200.19.230	SSD	500 GB

Virtual Datastore:


Node Name	Disk Groups	Cache Disks	Data Disks	Spare Disks	Free Disks	Total Space
Room1						
192.168.20.181: (...)	1	1	2	0	0	3.64 TB
192.168.20.182: (...)	1	1	2	0	0	3.64 TB

Back OK Cancel

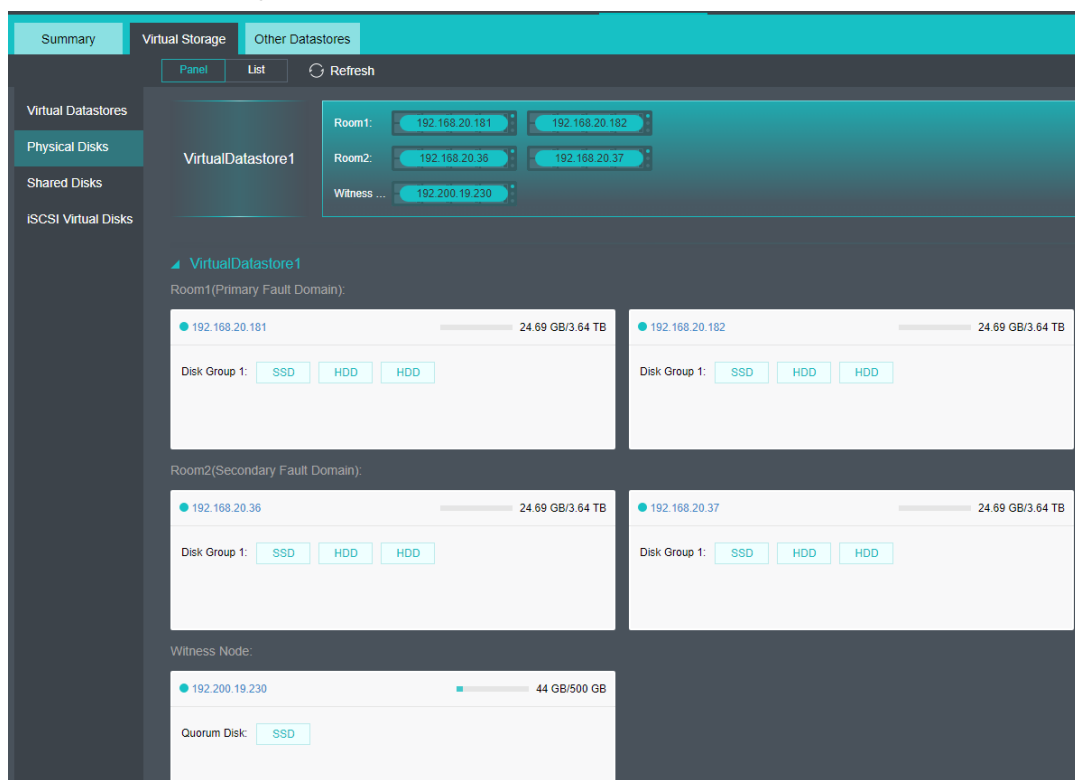
9. Enter password to confirm virtual datastore creation.



10. Diagram below show Stretched datastore currently creating.



Stretched virtual storage overview



Chapter 3 Run VM on fault domain

3.1 Create new virtual machine

Create virtual machine and select primary/secondary fault domain to run.

Create New Virtual Machine

Name:

Group:

Tag:

HA: ☒ Migrate to another node if the node fails [HA Settings](#)

Datastore:

Run on Node:

Guest OS:

Node	IP Address	CPU Usage	Memory Size	Memory Usage
<Auto>	-	-	-	-
<Primary fault domain: Room1>	-	-	-	-
<Secondary fault domain: Room2>	-	-	-	-
192.168.20.181	192.168.20.181	14%	32 GB	28%
192.168.20.182	192.168.20.182	14%	32 GB	30%
192.168.20.36	192.168.20.36	52%	32 GB	32%
192.168.20.37	192.168.20.37	20%	32 GB	32%

☒ Enable NUMA Scheduler [?](#)

☐ Use CPU from host [?](#)

Configuration **Advanced**

Standard: ☐ Low ☒ Typical ☐ High

Processor: 8 core(s)

Memory: 16 GB

Disk 1: 80 GB

CD/DVD 1: None

eth0: Connected To: DefaultEdge

OK **Cancel**

Note: There have 2 option for running method.

Prefer this one: Virtual machine first priority start on selected fault domain. If it failed to start, it will swap to another fault domain.

Require on this one: Virtual machine must start on this fault domain. If it failed to start, it will never start and prompt error message under task logs.

Overview of fault domain node

SANGFOR aCloud VMware vCenter

View By Node [Refresh](#) [New](#) [New Group](#) [Power On](#) [Shut Down](#) [More](#)

2 virtual machine(s) giving alarm [View](#)

Basics	Node	Throughput	IO Speed	Host Resources	Backup
<input type="checkbox"/> Status	VM Name	IP Address	Group		
<input type="checkbox"/> Alarm	centos	192.168.20.124	Default Group		
<input type="checkbox"/> Alarm	centos2	192.168.20.125	Default Group		



SANGFOR

Copyright © SANGFOR Technologies Inc. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of SANGFOR Technologies Inc.

SANGFOR is the trademark of SANGFOR Technologies Inc. All other trademarks and trade names mentioned in this document are the property of their respective holders.

Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied. The information in this document is subject to change without notice. To obtain the latest version, contact the international service center of SANGFOR Technologies Inc

