




Citrix XenServer: High Availability Quick Start Guide

www.citrix.com





Contents

- What is XenServer High Availability? 3
- Enabling High Availability 3
 - Increasing the Maximum Failure Capacity for a Pool 7
- Additional Reading 8



What is XenServer High Availability?

With High Availability enabled, if one of your hosts fails, its VMs restart automatically on other hosts in the same pool. If the pool master fails, High Availability automatically selects a new host as the master from any host in the pool.

High Availability works by creating a failure plan (that is, by calculating how many hosts can be restarted based on the priorities you set). The number of hosts that can be restarted is based on the available resources (CPU, memory) in the pool. As you specify the restart priority for VMs, XenServer evaluates the resources required to start each VM. When there are not enough resources to restart all the VMs set to be restarted, the pool reaches its Maximum failure capacity.

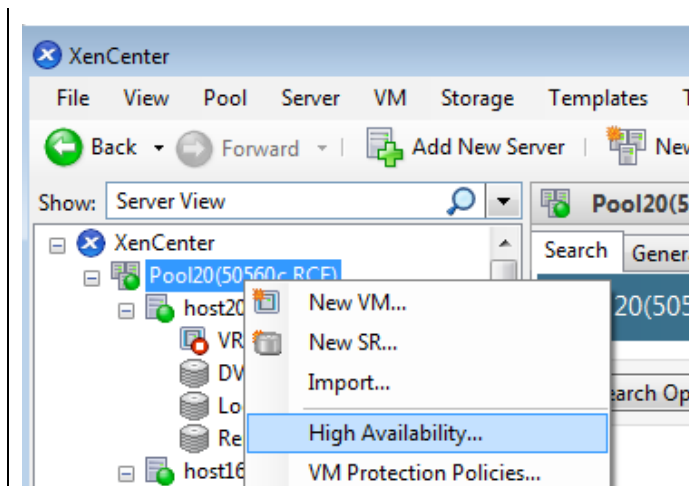
Note: To configure High Availability, all hosts in the pool must have licenses for Citrix XenServer Advanced Edition or higher. To learn more about the features available in different XenServer Editions, visit www.citrix.com/xenserver/editions.

High Availability requires the pool use remote storage, and Citrix recommends bonding the pool's management interface and multipathing storage for the heartbeat storage repository (SR). High Availability is explained in more detail in the *XenServer Administrator's Guide*.

Note: If Role-Based Access Control (RBAC) is enabled, users with permissions lower than Pool Operator will not be able to start VMs in some situations. For more information, see the *XenServer Administrator's Guide*.

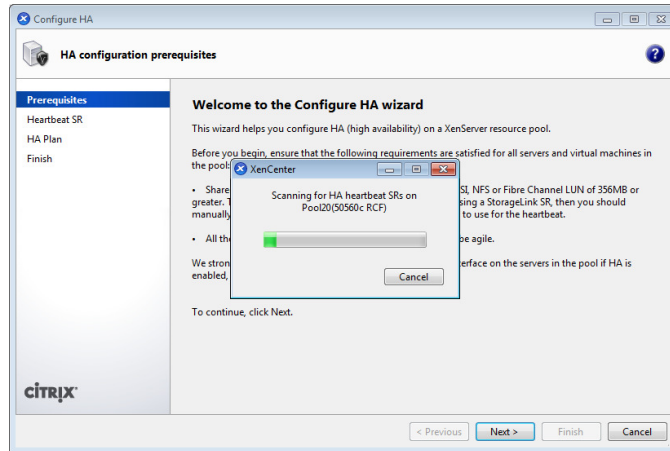
Enabling High Availability

1. Right-click the pool in the **Resources** pane, and then click **High Availability** on the shortcut menu.

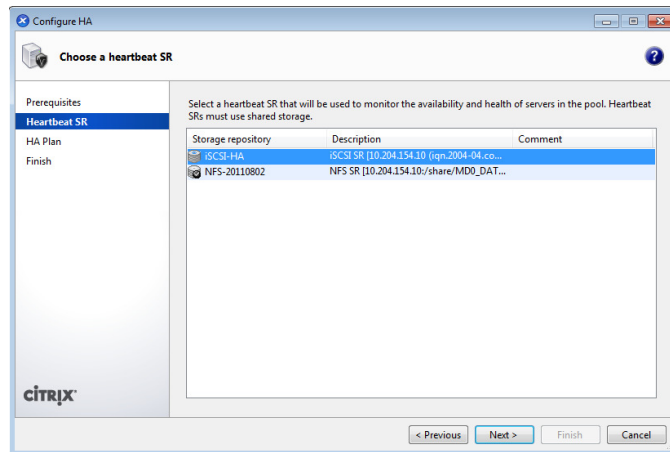


2. Click **Next** on the first page of wizard to scan the pool for a shared iSCSI, Fibre Channel, or NFS LUN that can be used as the pool's heartbeat SR.

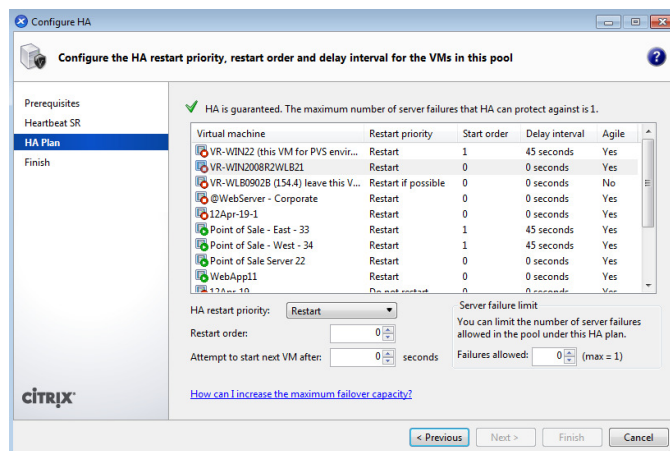
Note: If no suitable SR is found, you must configure appropriate new storage before you can continue.



3. On the **Heartbeat SR** page, choose an SR from the list and then click **Next**.



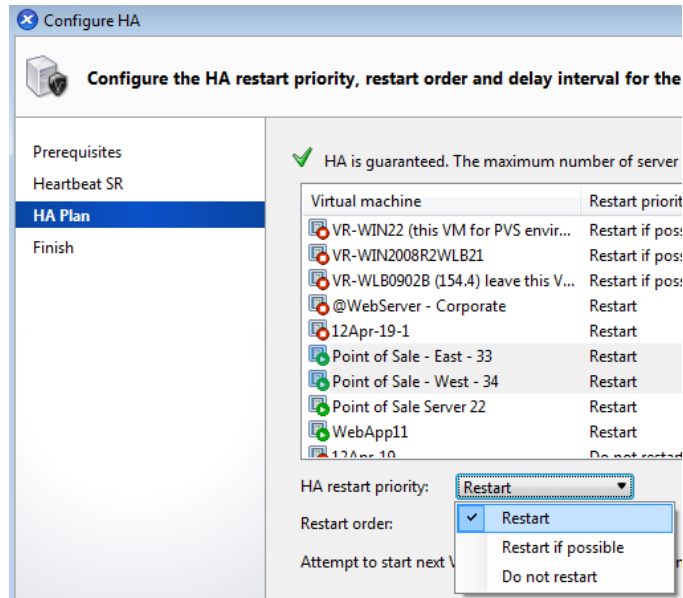
4. In the **Configure HA** page, shown across, configure the restart priority for each VM (by following steps a to c on the following page):




a. Select one or more VMs in the list and set the priority from the **HA restart priority** menu:

- **Restart.** VMs are guaranteed to be restarted if sufficient resources are available in the pool.
- **Restart if possible.** XenServer attempts to restart **Restart if possible** VMs only after it restarts all higher-priority VMs. If XenServer cannot restart **Restart if possible** VMs on its first attempt, it does not retry.


You can only assign VMs that are not agile a **Restart if possible** restart priority since they are tied to one host.



b. In the **Restart order** box, specify the order individual VMs will be started during the recovery operation. VMs set to 0 (zero) will be started first.



c. (Optional.) In the **Attempt to start next VM after** box, enter a value in seconds to specify how long for XenServer is to wait after starting the VM before attempting to start the next group of VMs in the startup sequence (that is, VMs with a higher start order value).

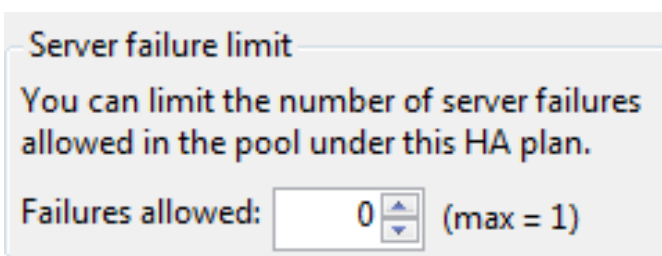


Note: Typically, this setting is used when you want to have one VM start after a particular service starts in another VM (for example, a web server in that depends on a database).

5. In **Server failure limit**, set the number of server failures to allow in this High Availability plan.

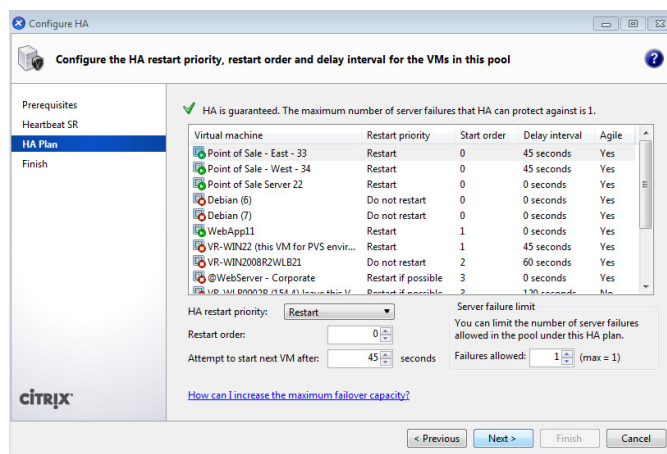
This value specifies the number of server failures that you want to allow in the pool, within the High Availability plan.

This value should be less than or equal to the maximum failure capacity for the pool, shown in the screen capture as **(max = 1)**. For example, to make the **Next** button available, you may need to enter **1** or **0**. If you enter **2**, an error appears.

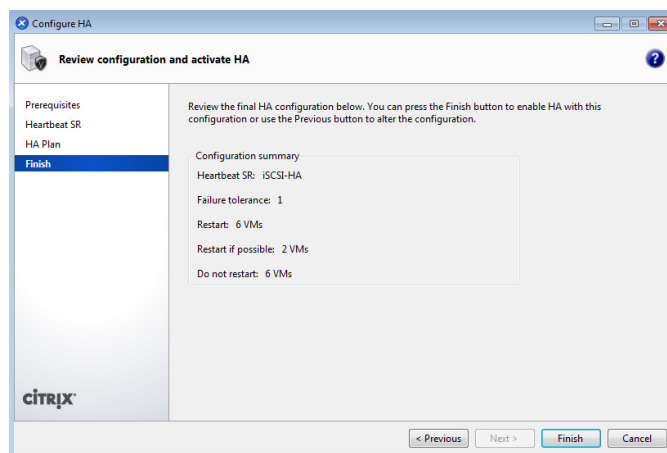


Note: If **max = 0** (zero), the pool is overcommitted. You cannot continue to the next wizard page without either adjusting the HA restart priorities or making more resources available within the pool; see [Increasing the Maximum Failure Capacity for a Pool](#).


6. Click the next VM for which you want to configure High Availability to register the settings. Repeat this process until you configure all desired VMs. Click **Next** when you have finished High Availability plan configuration.



7. On the last page of the wizard, click **Finish** to enable High Availability and close the wizard.



8. After clicking **Next**, while XenServer is enabling HA, you can check the **Logs** tab to see the progress.



Increasing the Maximum Failure Capacity for a Pool

When you first enable High Availability on a pool, XenServer calculates a failure plan based on the resources available at that time. To create a failure plan, XenServer determines if the VMs set to **Restart** can be restarted if the maximum number of server failures occurs (that is, if there would be enough resources available for all the VMs to restart).

If XenServer cannot create a plan in which all VMs with the **Restart** priority can be reliably restarted, the pool is considered overcommitted. The pool can also be overcommitted for reasons such as not enough free memory or changes to virtual disks and networks that affect which VMs can be restarted on which servers.

HA cannot be guaranteed with the specified settings and the current pool resources.

Virtual machine	Restart priority	Start order	Delay interval	Agile
VR-WIN22 (this VM for PVS envir...	Restart	0	0 seconds	Yes
VR-WIN2008R2WLB21	Restart	0	0 seconds	Yes
VR-WLB0902B (154.4) leave this V...	Restart if possible	0	0 seconds	No
@WebServer - Corporate	Restart	0	0 seconds	Yes
12Apr-19-1	Restart	0	0 seconds	Yes
Point of Sale - East - 33	Restart	1	45 seconds	Yes
Point of Sale - West - 34	Restart	1	45 seconds	Yes
Point of Sale Server 22	Restart	0	0 seconds	Yes
WebApp11	Restart	0	0 seconds	Yes
12Apr-19	Do not restart	0	0 seconds	Yes

HA restart priority:

Restart order:

Attempt to start next VM after: seconds

[How can I increase the maximum failover capacity?](#)

Server failure limit
 You can limit the number of server failures allowed in the pool under this HA plan.
 Failures allowed: (max = 1)
 ❌ This number may not exceed the pool's maximum failure capacity.

To increase the maximum failure capacity for a pool, you need to do one or more of the following:

- Reduce the number of VMs set to **Restart** as their restart priority.
- Increase the amount of RAM on your servers or add more servers to the pool to increase its capacity.
- Reduce the amount of memory configured on some VMs.
- Shut down non-essential VMs.



XenServer does not consider VMs set to have the **Restart if possible** restart priority when calculating a failure plan, but it will attempt to restart them one time if a server running these VMs fails.

Note: If you attempt to start or resume a VM and that action would cause the pool to be overcommitted, a warning is displayed in XenCenter. A pool is overcommitted when XenServer cannot calculate a failure plan because there are insufficient resource available to carry out the plan if a host failed. When a warning message appears, you can cancel the operation, or proceed anyway, causing the pool to become overcommitted.

Additional Reading

High Availability Whitepaper. http://support.citrix.com/servlet/KbServlet/download/21018-102-479340/HA_deep_2.pdf

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is the leading provider of virtualization, networking and software as a service technologies for more than 230,000 organizations worldwide. Its Citrix Delivery Center, Citrix Cloud Center (C3) and Citrix Online Services product families radically simplify computing for millions of users, delivering applications as an on-demand service to any user, in any location on any device. Citrix customers include the world's largest Internet companies, 99 percent of Fortune Global 500 enterprises, and hundreds of thousands of small businesses and prosumers worldwide. Citrix partners with over 10,000 companies worldwide in more than 100 countries. Founded in 1989, annual revenue in 2010 was \$1.87 billion.

©2011 Citrix Systems, Inc. All rights reserved. Citrix®, Access Gateway™, Branch Repeater™, Citrix Repeater™, HDX™, XenServer™, XenCenter™, XenApp™, XenDesktop™ and Citrix Delivery Center™ are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.