



Deploying iSCSI High Availability in 2-Node XAPI  
Based Xen Virtualization Environments with HA-lizard



**Installation and Reference Guide  
High Availability iSCSI Add-On  
for HA-Lizard**

Version 1.4



The information in this document and any product or service specifications referred to herein are subject to change without notice.

XenServer, XenCenter, Xen Cloud Platform and XCP are registered trademarks or trademarks of Citrix System, Inc and Xen.org  
DRBD and Distributed Replicated Block Device are registered trademarks or trademarks of LINBIT HA-Solutions GmbH

No part of this document may be reproduced, copied, altered or transmitted in any form or by any means, electronic, mechanical or otherwise for any purpose whatsoever, without the express written permission of the Copyright owner.

The information provided in this document is intended as a guide only and is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

No support is provided as part of the information in this document or any related software. Contact the project sponsor, Pulse Supply ([www.pulsesupply.com](http://www.pulsesupply.com)), for details on support offerings.

Document Copyright © 2014 Pulse Supply  
All rights reserved.

**IMPORTANT**

#####

!!! HA-Lizard iSCSI-HA is free software: you can redistribute it and/or modify  
!!! it under the terms of the GNU General Public License as published by  
!!! the Free Software Foundation, either version 3 of the License, or  
!!! (at your option) any later version.

!!!  
!!! HA-Lizard iSCSI-HA is distributed in the hope that it will be useful,  
!!! but WITHOUT ANY WARRANTY; without even the implied warranty of  
!!! MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
!!! GNU General Public License for more details.

!!!  
!!! You should have received a copy of the GNU General Public License  
!!! along with HA-Lizard iSCSI-HA. If not, see <<http://www.gnu.org/licenses/>>.

#####



## Table of Contents

1. iSCSI-HA Add-on for XCP and XenServer.....	5
• Purpose .....	5
• Requirements .....	6
2. Installation .....	7
• Installing .....	7
• Removing.....	7
3. Configuring.....	8
• Configuration Parameters .....	8
- HA Monitor Configuration .....	8
- MONITOR_MAX_STARTS.....	8
- MONITOR_KILLALL.....	8
- MONITOR_DELAY.....	8
- MONITOR_SCANRATE .....	8
- DRBD_RESOURCES.....	8
- ISCSI_TARGET_SERVICE .....	8
- DRBD_VIRTUAL_IP and DRBD_VIRTUAL_MASK .....	9
- DRBD_INTERFACE .....	9
- ENABLE_LOGGING.....	9
- Email Alert Settings .....	9
4. Managing iSCSI-HA Services .....	10
• iSCSI-HA System Service.....	10
- iSCSI-HA Watchdog Service .....	10
- Disabling the Watchdog Service .....	10



5. iscsi-cfg CLI Tool.....	11
• System Logging .....	11
• Viewing Configuration Parameters .....	11
• Viewing iSCSI-HA Status .....	11
• Operating in Manual Mode .....	12
- Enabling Manual Mode .....	13
- Disabling Manual Mode .....	13
- Becoming the Storage Primary Role.....	13
- Becoming the Storage Secondary Role.....	13
• Manual Mode Best Practices.....	13
- Existing Manual Mode Cleanly .....	14
- Entering Manual Mode .....	14
- Example Maintenance Operations.....	15
6. Miscellaneous.....	16
• Dependencies and Compatibility.....	16
• Security and Ports.....	16
• Support.....	16

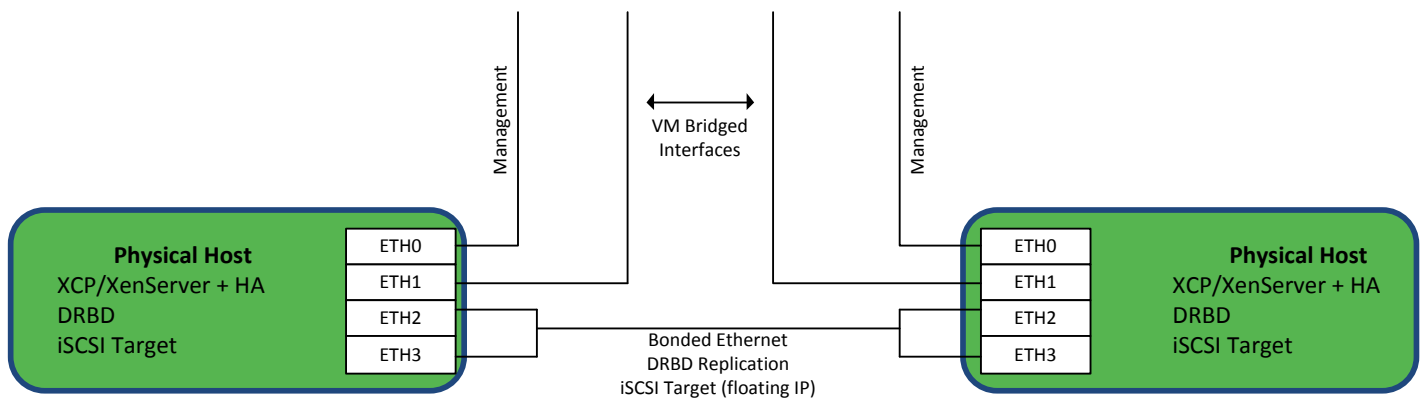
# 1. iSCSI-HA Add-on for XCP and XenServer

- Purpose**

iSCSI-HA is an add-on module for two node pools utilizing Xen Cloud Platform (XCP) or XenServer virtualization environments. It is intended to build highly available two-node clusters with local storage without limiting pool advanced functionality such as live migration.

Generally, this is achieved with separate iSCSI/SAN and Xen clusters which require a minimum of four physical servers and a pair of redundant Ethernet switches to reach an adequate level of fault tolerance. This may not be the most efficient use of hardware for small cluster applications. The goal of iSCSI-HA is to provide a simple framework for building compact, highly available pools utilizing XenServer or Xen Cloud Platform with just two physical hosts.

iSCSI-HA requires DRBD for block replication of storage and an iSCSI target framework such as TGT. A sample pool design/diagram is shown below for a highly available two node pool.



In this example, the DRBD and iSCSI interface is provided via direct attachment (no Ethernet switches) on a bonded Ethernet link. This approach greatly eliminates the possibility of a split brain scenario since there are no networking devices interconnecting the hosts in any way. Additionally, utilizing a bonded Ethernet link further eliminates the possibility of communication interruption between the hosts.

The iSCSI-HA add-on does not make any decisions or employ any logic relating to cluster management and the roles of the hosts. It relies on an external HA tool such as open source HA-Lizard or XenServer HA licensing. The iSCSI-HA add-on relies on the external HA logic to ensure that a pool Master is always available. Based on this, iSCSI-HA will assign a single shared/floating IP address to the Master and promote DRBD resources to follow the floating IP. The slave host will be in a demoted state at all times. In the event of



a change of pool roles, iSCSI-HA will automatically detect the new roles and promote the new Master as the iSCSI target and demote the former Master to slave/standby mode.

- **Requirements**

- XCP version 1.6 or XenServer version 6.1/6.2
- DRBD Version 8.3.x
- iSCSI target – TGT
- Pool HA – Open Source HA-Lizard 1.6.41.4 or newer

iSCSI HA features provided:

- Support for shared/floating virtual IP address to export iSCSI LUNs
- Automatic promoting of DRBD resources on pool Master
- Automatic demoting of DRBD resources on pool Slave
- Manual mode allows for host upgrade and reboots with no storage downtime.
- Management of iSCSI service
- Extensive Logging capabilities to system log file
- Email alerting
- Dynamic iSCSI target selection auto-selects roles
- No changes to existing pool configuration required. All logic is external.
- **Minimal dependencies** – does not compromise pool stability or introduce complex SW packages. Designed to work with the resident packages on a standard XCP/XenServer host.

Development is well tested and based on Xen Cloud Platform (XCP) version 1.6, XenServer 6.1/6.2 and DRBD 8.3



## 2. Installation

---

An included installer is designed for systems based on Centos DomOs. Manual installation may be required for Debian based DomOs. The following installation instructions will work with a standard ISO install of XCP version 1.6 or XenServer version 6.1.

Before installing:

- DRBD must be installed, configured and running on both hosts
- iSCSI target must be installed, configured and running on both hosts

- **Installing**

- Copy the source tarball into a temporary location (ex. /tmp/)
- Extract its contents and move into the extracted folder  
**`tar -zxvf iscsi-ha-<version>.tgz`**
- Move into the “scripts” folder  
**`cd iscsi-ha-<version>/scripts`**
- Run the installer  
**`./install`**

The installer will check if sendmail packages are installed on the server. These are only required for email alerts. Skip the installation of these packages if email alerting is not required.

Once the installer is completed, iSCSI-HA and watchdog services can be started with the following command:  
**`service iscsi-ha start -w`**

- **Removing**

An uninstall script is provided for Centos based systems. This will completely remove all components from the host.



### 3. Configuring

---

The iSCSI-HA configuration file must be edited to match the pool/environment. The file is located at `/etc/iscsi-ha/iscsi-ha.conf` and can be edited with any text editor such as `vi` or `nano`.

- **Configuration Parameters**

- ***HA Monitor Configuration***

The iSCSI-HA service is run by a monitoring service which runs continuously. The following Monitor settings are used to configure the behavior of the monitor. The provided installer installs and configures the Monitor with default settings that are acceptable in most cases. Once installed, the Monitor will run continuously as a service. Status can be checked with “`service iscsi-ha status`”.

Some values may need to be changed depending on the environment. The Monitor will launch several HA processes in a loop every 10 seconds (`MONITOR_DELAY`). By increasing `MONITOR_DELAY`, it will take longer to detect a change in the pool status and thus longer to move the iSCSI service when needed. Decreasing `MONITOR_DELAY` will more quickly detect failures and recover.

- ***MONITOR\_MAX\_STARTS***

Threshold for when to assume running processes are not responding. Sets how many failed starts to wait before killing any hung processes. Default = 5

- ***MONITOR\_KILLALL***

If `MAX_MONITOR_STARTS` threshold is reached - set whether to kill all iscsi-ha processes. Default = 1  
1 = yes, 0 = no

- ***MONITOR\_DELAY***

Delay in seconds between re-spawning iscsi-ha. Default = 10

- ***MONITOR\_SCANRATE***

Iscsi-ha will not re-spawn unless all current processes are completed. If there are active processes while attempting to start a new iteration, iscsi-ha will wait the number of seconds set here before retrying. Each successive fail will increment a counter (`MONITOR_MAX_STARTS`) that may trigger `KILLALL`. Default = 10

- ***DRBD\_RESOURCES***

Declare the DRBD resource names in a ":" separated list like this `resource1:resource2:resource3`

- ***ISCSI\_TARGET\_SERVICE***

Set the path to the init script for the iSCSI target. Default = `/etc/init.d/tgtd`





- ***DRBD\_VIRTUAL\_IP and DRBD\_VIRTUAL\_MASK***

Declare the virtual "floating" IP address and subnet mask to be used for the iSCSI target NOTE: It is expected that the virtual/floating IP is on the same subnet as other system interfaces. No changes to routing or gateways are made by iscsi-ha the IP set here should work with the routing configuration of the environment.

- ***DRBD\_INTERFACE***

Declare the interface to bind the virtual IP to. Typically xenbr0 or xapi0, xapi1, etc.. Will vary based on environment. Default = xapi0

- ***ENABLE\_LOGGING***

Enable Logging 1=yes, 0=no. Logs are written to /var/log/messages. All log messages are labeled with "iscsi-ha" for easy filtering. View/Filter real time logging with: "tail -f /var/log/messages | grep iscsi-ha"

- ***Email Alert Settings***

- MAIL\_ON: Enable/Disable email alerts. 1 = enabled 0 = disabled
- MAIL\_SUBJECT: Subject Line of email alert
- MAIL\_FROM: The FROM email address used on email alerts
- MAIL\_TO: the email address to send alerts to
- MAIL\_SCREEN\_TIME: time in minutes to suppress duplicate email alerts



## 4. Managing iSCSI-HA Services

---

- **iSCSI-HA System Service**

When installing on a CentOS based Dom0, the installer will install a startup script in /etc/init.d and set it to automatically start each time the server is started. Generally there are no additional steps required in setting up the system service.

The service can be checked, stopped or invoked with the “*service*” command which manages System V init scripts. The following arguments are supported:

- “service iscsi-ha start”: starts the service
- “service iscsi-ha stop”: stops the service
- “service iscsi-ha restart”: restarts the service
- “service iscsi-ha status”: reports running status of the service

By default a watchdog service is installed and started when installing with the included installer.

Important Note: Stopping the iSCSI-HA service while the watchdog service is running will be ineffective as the watchdog will restart the iSCSI-HA service within a few seconds after a stop. The iSCSI-HA init script can be invoked with a “-w” option to also start or stop the watchdog service with the iSCSI-HA service. The -w option can be used as follows:

- “service iscsi-ha start -w”: starts the service and watchdog
- “service iscsi-ha stop -w”: stops the service and watchdog
- “service iscsi-ha start -w”: starts the service and watchdog
- “service iscsi-ha status -w”: reports running status of the service and watchdog

- ***iSCSI-HA Watchdog Service***

A watchdog service is installed by default and can be started/stopped via the main service init script or can be individually managed with:

- “service iscsi-ha -watchdog start”: starts the service
- “service iscsi-ha -watchdog stop”: stops the service
- “service iscsi-ha -watchdog status”: reports running status of the service

The default watchdog interval for checking the iSCSI-HA service is 10 seconds. This can be changed by editing the variable “WATCH\_INTERVAL” in /etc/init.d/ iscsi-ha -watchdog.

- ***Disabling the Watchdog Service***

The watchdog service can be disabled by invoking the following commands: “service iscsi-ha-watchdog stop” and “chkconfig iscsi-ha-watchdog off”



## 5. iscsi-cfg CLI Tool

---

A command line tool is provided as part of the iscsi-ha package. The tool can be called with:

```
iscsi-cfg
```

A sample output with command line monitoring arguments is shown below:

```
iSCSI-HA Monitoring Tool: Add-on for HA-Lizard: XenServer/XCP High Availability
Usage: iscsi-cfg <action>

Available actions:
<log>:          Watch iSCSI-HA log file output in real time
<get>:          Lists all iSCSI-HA configuration parameters
<manual-mode-enable>: Enter manual mode - required to manually select roles
                  Allows for manually moving iSCSI target to desired host
                  Used to manage rolling updates and server reboots
                  with no VM downtime.
<manual-mode-disable>: Exit manual mode - automatic selection of roles enabled
                       Operation returns to normal - iSCSI-HA manages roles
<become-primary>:  Manually promotes host to primary role regardless of
                       role of the host in the pool master/slave. Only works
                       when operating in manual mode.
<become-secondary>: Manually demotes host to secondary role regardless of
                       role of the host in the pool master/slave. Only works
                       when operating in manual mode.
<status>:         Displays the iSCSI-HA operational status
```

- **System Logging**

A live view of the system logs generated by iSCSI-HA is available by invoking:

```
iscsi-cfg log
```

- **Viewing Configuration Parameters**

A listing of configuration parameters for the local host is available by invoking:

```
iscsi-cfg get
```

- **Viewing iSCSI-HA Status**

The iSCSI-HA service is responsible for managing:

- DRBD Running State
- DRBD Resource State (primary/secondary)
- iSCSI Target (TGT) Running State
- Floating IP Address



The status of each of these can be viewed in real time by invoking:

```
iscsi-cfg status
```

The information displayed will be relative to the role of the host within the XenServer pool. Below examples show the output from both the pool Master and Slave nodes.

```
*****
iSCSI-HA Status: Version: 1.2.11 iscsi-ha (pid 1145 1141) is running...
Last Updated: Sun Jul 14 20:54:36 EDT 2013

HOST ROLE:           MASTER
DRBD ROLE:           iscsi1=Primary
DRBD CONNECTION:    iscsi1 in Connected state
ISCSI TARGET:       tgttd (pid 10093 10091) is running...
VIRTUAL IP:         10.10.10.3 is local
*****
Control + C to exit
```

```
*****
iSCSI-HA Status: Version: 1.2.11 iscsi-ha (pid 19960 19953) is running...
Last Updated: Sun Jul 14 20:55:31 EDT 2013

HOST ROLE:           SLAVE
VIRTUAL IP:         10.10.10.3 is not local
ISCSI TARGET:       tgttd is stopped
DRBD ROLE:         iscsi1=Secondary
*****
```

The status will not be displayed if the iSCSI-HA service is stopped or a system failure is preventing it from running correctly. In this case, the status output will look like the below example.

```
Sun Jul 14 20:59:30 EDT 2013
#####
## iSCSI-ha Status cache is more than 33 seconds old
## Tolerance = 15 seconds. Check status
#####

Possible causes: 1) iSCSI-HA is actively failing over to standby host
                   this can take up to 60 seconds or more depending on settings

                   2) Service is not running. try 'service iscsi-ha status -w' to check running status
                      If the service is not running, try 'service iscsi-ha start -w' to start the service

Control + C to exit
```

- **Operating in Manual Mode**

Under normal operating conditions, iSCSI-HA operates in a completely automatic and dynamic fashion such that the system can automatically recover from most failures with no user intervention. There are cases, however, where automation must be disabled so that hosts can be safely rebooted or upgraded with zero VM downtime. This is important since under normal operating conditions the storage will *always* be exposed on the pool master. Manual mode allows administrators to disable the iSCSI-HA logic so that storage can be moved between hosts as needed.



For example: Assuming a pool Master must be rebooted as part of an upgrade, an administrator would enable manual mode on both hosts and then safely move the VMs and the storage to the pool slave with no downtime. **Important Note: High Availability must also be disabled in this case before rebooting a host.**

#### - **Enabling Manual Mode**

The following command is used to enable manual mode on a host:

```
iscsi-cfg manual-mode-enable
```

**CAUTION: REBOOTING A HOST THAT IS IN THE PRIMARY ROLE WHILE IN MANUAL MODE WILL CAUSE VMs TO LOSE CONNECTIVITY TO THE STORAGE REPOSITORY. USE CAUTION AND ONLY RESTART A HOST THAT IS IN SECONDARY MODE MAKING SURE THAT THE PEER IS PRIMARY**

**Important Note: This should be applied to both hosts in the pool to avoid any possible contention**

#### - **Disabling Manual Mode**

The following command is used to disable manual mode and return the pool to automatic selection of roles:

```
iscsi-cfg manual-mode-disable
```

**Important Note: Remember to re-enable HA if it was disabled**

#### - **Becoming the Storage Primary Role**

Storage can be manually promoted to master mode while iSCSI-HA is in manual mode. This can only be done on a single host (iSCSI-HA does not support dual primary mode). A warning will be displayed if an attempt to promote a host is made while the peer host is already in the primary storage role. The following command is used to promote a host to primary mode:

```
iscsi-cfg become-primary
```

#### - **Becoming the Storage Secondary Role**

Storage can be manually demoted to the secondary role while iSCSI-HA is in manual mode. This can be done on any host in the pool with no restrictions. When swapping primary/secondary roles, it is first necessary to put both hosts into the secondary role and then promote the desired host to the primary role. The following command is used to demote a host to the secondary role:

```
iscsi-cfg become-secondary
```

#### • **Manual Mode Best Practices**

Manual mode is intended to provide a means for managing a 2-node virtualization/storage cluster manually. This is especially important for environments and applications which require near 100% uptime. When utilizing manual mode, the storage can be exposed on either host as needed. Data replication continues to operate as usual when moving storage between hosts. All iSCSI-HA operations (eg. promoting/demoting storage hosts) are completely transparent to XenServer and any underlying VMs.



iSCSI-HA employs several conditional checks while operating in manual mode. These checks ensure that a user cannot destabilize a system or assert a role which is disallowed for a particular host.

- ***Existing Manual Mode Cleanly***

Manual mode cannot be disabled unless both hosts are first restored to their native storage roles. This means that the XenServer pool master must be promoted to the storage master (iscsi-cfg become-master) and the pool slave must be demoted to the storage secondary role (iscsi-cfg become-secondary). With both hosts in their native roles manual mode should be disabled on both hosts with 'iscsi-cfg manual-mode-disable'. This behavior is intentional and ensures that there is no storage downtime when returning the pool to normal (automated) operation.

- ***Entering Manual Mode***

Manual mode is generally used to perform maintenance operations on hosts within a given pool. Although there are no restrictions set when entering manual mode, one should ensure that any pool HA functionality is disabled while operating in manual mode if hosts require a shutdown or restart.



### - **Example Maintenance Operations**

The below example illustrates the steps necessary to perform an update to both the pool master and slave hosts which requires that they be rebooted. The following procedure ensures that there is no downtime on any VMs. For this example it is assumed that the pool master will be rebooted first and then the pool slave.

- 1) Disable HA for the pool (for HA-Lizard 'ha-cfg' can be used to disable HA)
- 2) Enter manual mode on each host ('iscsi-cfg manual-mode-enable')
- 3) Migrate all VMs to the pool slave
- 4) Demote the pool master's storage role ('iscsi-cfg become-secondary')
- 5) Promote the pool slave's storage role ('iscsi-cfg become-primary')  
Important Note: there should be minimal delay between the demote/promote actions to ensure that the VMs experience little storage downtime. Typical delay when manually switching roles is 3 seconds.
- 6) The master host can now be safely worked on (enter maintenance mode, shutdown, rebooted, etc...)
- 7) Rejoin the master host to the pool
- 8) Migrate all VMs from the slave to the pool master
- 9) Demote the pool slave's storage role ('iscsi-cfg become-secondary')
- 10) Promote the pool master's storage role ('iscsi-cfg become-primary')  
Important Note: there should be minimal delay between the demote/promote actions to ensure that the VMs experience little storage downtime. Typical delay when manually switching roles is 3 seconds.
- 11) The slave host can now be safely worked on (enter maintenance mode, shutdown, rebooted, etc...)
- 12) Rejoin the slave host to the pool
- 13) Exit manual mode on both hosts ('iscsi-cfg manual-mode-disable')
- 14) Enable HA for the pool (for HA-Lizard 'ha-cfg' can be used to enable HA)



## 6. Miscellaneous

---

- **Dependencies and Compatibility**

When installing iSCSI-HA onto a default Centos based DomO (XCP or XenServer), all the required tools needed to run iSCSI-HA are resident on the system with the exception of:

- DRBD – version 8.3 required. RPMs are provided in /etc/iscsi-ha/RPM/
- TGT iSCSI Target (can be installed with “yum –enablerepo=base install scsi-target-utils”)

***Package is compatible with XCP version 1.6 and XenServer version 6.1/6.2. Prior releases may work but have not been tested.***

For custom DomO installations, ensure the following tools are available:

xapi and xe toolstack  
/bin/cat  
/bin/awk  
/bin/echo  
/sbin/drbdadm  
/bin/logger  
/sbin/ifconfig  
hostname  
/bin/mail  
/sbin/ip  
/sbin/arping

- **Security and Ports**

- iscsi port 3260 used as the listen port for the iscsi target
- ICMP (ping) is used to check whether the virtual IP is live
- DRBD – ensure that the port numbers specified in drbd.conf are open

- **Support**

- Post a question on the support forum  
<http://www.halizard.com/index.php/forum>
- Contact the project sponsor for paid support options  
<http://www.pulsesupply.com>