

User Guide

PowerChute™ Network Shutdown v5.1

VMware®

990-4595L-001
05/2024

Schneider Electric Legal Disclaimer

The information presented in this manual is not warranted by Schneider Electric to be authoritative, error free, or complete. This publication is not meant to be a substitute for a detailed operational and site specific development plan. Therefore, Schneider Electric assumes no liability for damages, violations of codes, improper installation, system failures, or any other problems that could arise based on the use of this Publication.

The information contained in this Publication is provided as is and has been prepared solely for the purpose of evaluating data center design and construction. This Publication has been compiled in good faith by Schneider Electric. However, no representation is made or warranty given, either express or implied, as to the completeness or accuracy of the information this Publication contains.

IN NO EVENT SHALL SCHNEIDER ELECTRIC, OR ANY PARENT, AFFILIATE OR SUBSIDIARY COMPANY OF SCHNEIDER ELECTRIC OR THEIR RESPECTIVE OFFICERS, DIRECTORS, OR EMPLOYEES BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL, OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS, CONTRACT, REVENUE, DATA, INFORMATION, OR BUSINESS INTERRUPTION) RESULTING FROM, ARISING OUT, OR IN CONNECTION WITH THE USE OF, OR INABILITY TO USE THIS PUBLICATION OR THE CONTENT, EVEN IF SCHNEIDER ELECTRIC HAS BEEN EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SCHNEIDER ELECTRIC RESERVES THE RIGHT TO MAKE CHANGES OR UPDATES WITH RESPECT TO OR IN THE CONTENT OF THE PUBLICATION OR THE FORMAT THEREOF AT ANY TIME WITHOUT NOTICE.

Copyright, intellectual, and all other proprietary rights in the content (including but not limited to software, audio, video, text, and photographs) rests with Schneider Electric or its licensors. All rights in the content not expressly granted herein are reserved. No rights of any kind are licensed or assigned or shall otherwise pass to persons accessing this information.

This Publication shall not be for resale in whole or in part.

Table of Contents

Introduction	1
UPS Configuration	2
Network Configuration	3
UPS Configuration Options	4
Network Management Card Connection	10
Advanced UPS Setups	13
Outlet Group Registration	15
Network Management Card Settings	16
VMware Configuration	17
Standalone VMware Host Details	18
vCenter Server Settings	20
VMware Host Protection	22
Advanced UPS configuration	24
Physical UPS Setup Power Protection	25
Virtualization Settings	26
Virtual Machine Migration	28
Virtual Machine and Virtual Appliance Shutdown/Startup	33
vSAN Settings	38
Host Maintenance Mode	40
vCenter Server Virtual Machine Shutdown	42
vSphere Cluster Services (vCLS)	43
High Availability	44
Virtual Machine Prioritization	45
Active Directory VMware Configuration	54
Hyperconverged Infrastructure Support for VMware	61
Nutanix Support	62
Nutanix Files (Acropolis File Services) Shutdown/Startup	64
Protection Domain Settings	66
Cluster Shutdown/Startup	67
Controller VM Shutdown/Startup	68
CVM/Cluster Details	70
HPE SimpliVity Support	71
OmniStack Virtual Controller Shutdown/Startup	72
Cisco HyperFlex Support	74
Cluster Shutdown/Startup	75
Controller VM Shutdown/Startup	76
Dell VxRail Support	78
VxRail Shutdown and Startup	80
Shutdown Settings	82

UPS Shutdown	83
Shutdown Command Files	84
Shutdown Settings for Advanced UPS Configurations	86
SSH Settings	92
SNMP Configuration.....	96
SNMPv1 Configuration.....	97
SNMPv3 Configuration	98
SNMP Trap Configuration.....	99
SNMP Data Points.....	101
Event Configuration	109
Start-up and Shut-down Scheduling	113
Sequenced Server Shutdown	114
Sample Shutdown Scenarios.....	115
VMware: UPS without Outlet Groups	115
VMware: UPS with Outlet Groups.....	119
VMware with Nutanix Support: UPS without Outlet Groups	124
VMware with Nutanix Support: UPS with Outlet Groups	128
VMware with HPE SimpliVity Support: UPS without Outlet Groups	133
VMware with HPE SimpliVity Support: UPS with Outlet Groups	140
VMware with HyperFlex Support: UPS without Outlet Groups.....	147
VMware with HyperFlex Support: UPS with Outlet Groups.....	154
VMware with Dell VxRail Support: UPS without Outlet Groups.....	161
VMware with Dell VxRail Support: UPS with Outlet Groups.....	169
VMware Shutdown - Single UPS Configuration.....	177
VMware Shutdown - HA Cluster.....	178
VMware Shutdown - vSAN Cluster.....	181
VMware Shutdown - Advanced UPS Configuration	189
VMware Shutdown - VM Prioritization - Single UPS Configuration	190
VMware Shutdown - VM Prioritization - Advanced UPS Configuration	191
Nutanix Shutdown - VM Prioritization - Single UPS Configuration	192
Nutanix Shutdown - VM Prioritization - Advanced UPS Configuration	193
HPE SimpliVity Shutdown - VM Prioritization - Single UPS Configuration	194
HPE SimpliVity Shutdown - VM Prioritization - Advanced UPS Configuration	195
HyperFlex Shutdown - VM Prioritization - Single UPS Configuration	196
HyperFlex Shutdown - VM Prioritization - Advanced UPS Configuration	198
Dell VxRail Shutdown - VM Prioritization - Single UPS Configuration.....	199
PowerChute Events and Logging.....	201
Configurable Events	202
Configurable Environmental Events	207
Non-Configurable Events	208
Configuration (INI) File Events	212

SSH Action Events	214
Java Update Events	215
License Events	216
VMware Virtualization Events.....	218
Hyperconverged Infrastructure Virtualization Events	225
Nutanix Virtualization Events.....	226
HPE SimpliVity Virtualization Events	228
HyperFlex Virtualization Events.....	229
Dell VxRail Virtualization Events	230
Critical Events in a Redundant-UPS Configuration	231
Critical Events in a Parallel-UPS Configuration	232
Logging	233
Logging to Third Party Systems.....	234
KeyStore and TrustStore file for Secure Syslog	236
General.....	237
Communications Settings.....	238
PowerChute Agents.....	239
PowerChute Configuration File.....	240
Resetting your PowerChute username or password	240
Java Update	241
User Interface Session Timeout	242
Check for Updates.....	243
Customer Support	244
Customer Experience Improvement Program (CEIP)	245
License	246
Free/Open Source License Details.....	251
Reset Virtual Appliance Root Password.....	252
Troubleshooting	255
Network Management Card Troubleshooting	256
VMware Troubleshooting.....	257
Nutanix Troubleshooting.....	266
HPE SimpliVity Troubleshooting.....	269
HyperFlex Troubleshooting.....	271
Dell VxRail Troubleshooting.....	273
Browser Troubleshooting.....	276
SSH Actions Troubleshooting.....	277
SNMP Troubleshooting.....	278
General Troubleshooting	279

Introduction

PowerChute™ Network Shutdown (PowerChute) works in conjunction with the UPS Network Management Card (NMC) to provide network-based shutdown of multiple computer systems.

In the case of a UPS critical event, the software performs a graceful, unattended system shutdown before the UPS battery is exhausted. The number of protected systems is limited only by the capacity of the UPS.

View these [Application Notes](#) for detailed information on using PowerChute in specific environments.



After installation, it is essential to configure the software using the PowerChute Setup wizard. This ensures that PowerChute is aware of UPS critical events in order to protect your system.


UPS Configuration

This section contains information on the topics below:

- [Network Configuration](#)
- [UPS Configuration Options](#)
- [Network Management Card Connection](#)
- [Advanced UPS Setups](#)
- [Outlet Group Registration](#)
- [Network Management Card Settings](#)

Network Configuration

PowerChute can use IPv4 or IPv6 to communicate with the Network Management Card(s).

 IPv6 support is available only for Network Management Card firmware 2.6.0.X or higher, and Network Management Card 3.

Select IP

If your computer has more than one IPv4 address you will need to select one of the available addresses. The IP address you select will be registered with the NMC and displayed in the NMC user interface under Configuration - PowerChute Clients.

IPv6 Configuration

If you are using IPv6 to communicate with the NMC(s), each network adapter on your machine will typically have several IP addresses assigned to it. Each adapter will have at least one link-local address and one global unicast address assigned to it.

Use the **Unicast IP Address** drop-down box to specify which address to use. The address type selected in this drop-down box must match the address type that you enter for the NMC(s) on the Network Management Card Connection page. This unicast address will be registered on the NMC(s) and displayed on the PowerChute Network Shutdown Clients page of the NMC.

`fe80::88c8:3d95:bc02:74cc` is an example of a link-local address.

`2001:112:1:0:88c8:3d95:bc02:74cc` is an example of a global unicast address.

Multicast Option

The NMC supports sending communication packets to an IPv6 Multicast address instead of sending unicast packets to each PowerChute agent. To use this, enable the **Multicast** check box and enter an IPv6 Multicast address.

The multicast address that is entered here will be registered on the NMC(s) instead of the unicast address and displayed on the PowerChute Network Shutdown Clients page of the NMC. The NMC(s) will send communication packets to that multicast address.

`FF02::1` is an example of a **multicast address with link-local scope** so that only nodes on the same physical network segment will receive it. If using a link-local unicast address, you must use a multicast address with link-local scope.

`FF0E::1` is an example of a **multicast address with global scope** and the NMC will use its global unicast address to send the packet. If using a global unicast address you must use a multicast address with global scope.

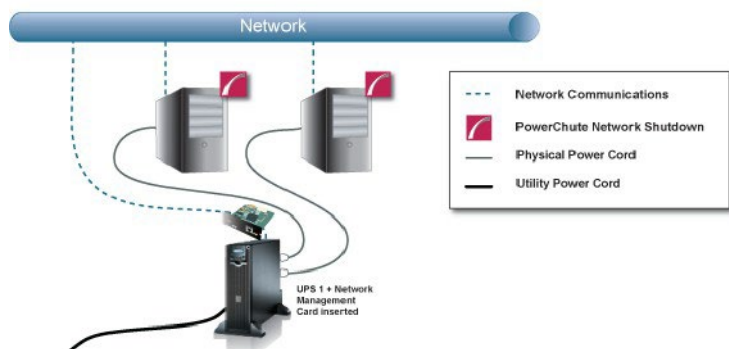
For detailed information, please view "The Communications Process of PowerChute Network Shutdown" [here](#).

UPS Configuration Options

For a detailed overview of which UPS's support each configuration, please view the "PowerChute Network Shutdown Operating Modes and supported UPS Configurations" Application Note [here](#).

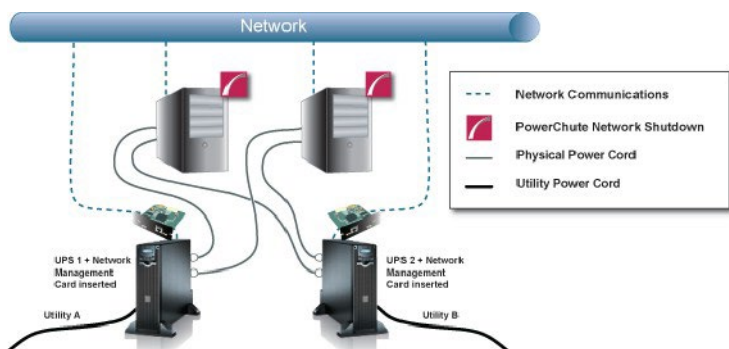
Single-UPS Configuration

Single-UPS Configuration: All servers are protected by a single UPS. The UPS Network Management Card communicates with each server that has PowerChute installed.



Redundant-UPS Configuration


Redundant-UPS Configuration: Two or more UPS's of the same model protect each server. Each UPS can support the server load on its own. All UPS Network Management Cards communicate with each server that has PowerChute installed.



For detailed information, please view "Using PowerChute Network Shutdown in a Redundant-UPS Configuration" Application Note [here](#).

Parallel-UPS Configuration

Parallel-UPS Configuration: Two or more UPS's of the same model protect the load and provide redundancy or increased capacity depending on the load. The UPS outputs are tied together so a single output goes to the load. All UPS Network Management Cards communicate with each server that has PowerChute installed



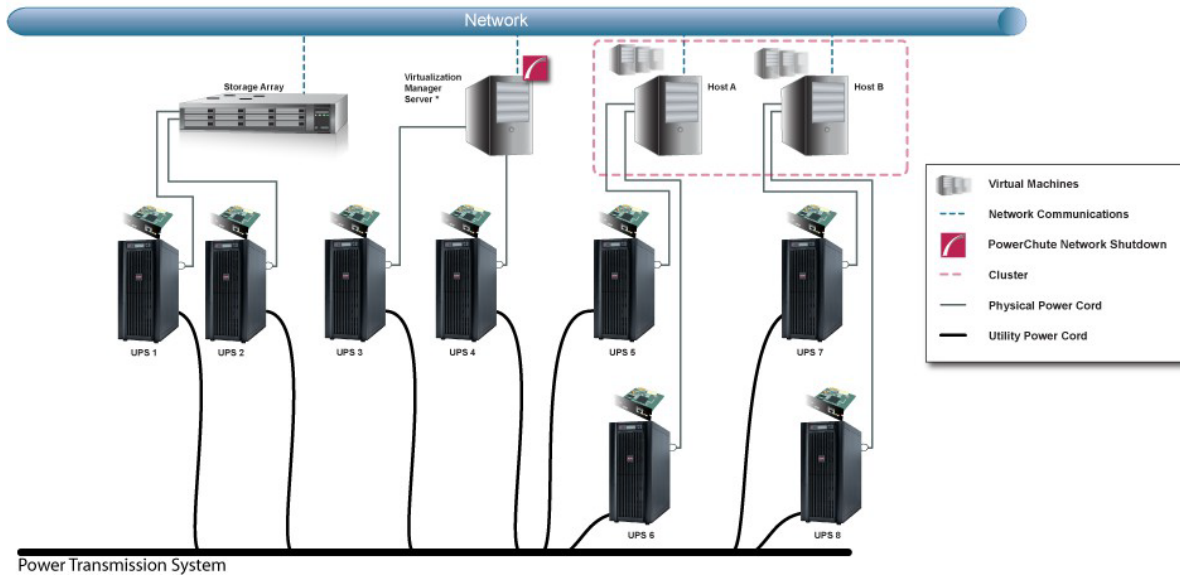
Note: To use the Parallel-UPS configuration, your UPS devices must already be configured to operate in parallel mode.

For detailed information, please view “Using PowerChute Network Shutdown in a Parallel-UPS Configuration” Application Note [here](#).

Advanced UPS Configuration

Standalone VMware hosts and hosts managed by vCenter Server are supported for Advanced UPS Configuration.

Advanced UPS Configuration: PowerChute can monitor both Single UPS's and groups of Redundant UPS's protecting your virtualization environment. If using Redundant UPS groups, redundancy levels can be set on a per group basis e.g. N+1, N+2.



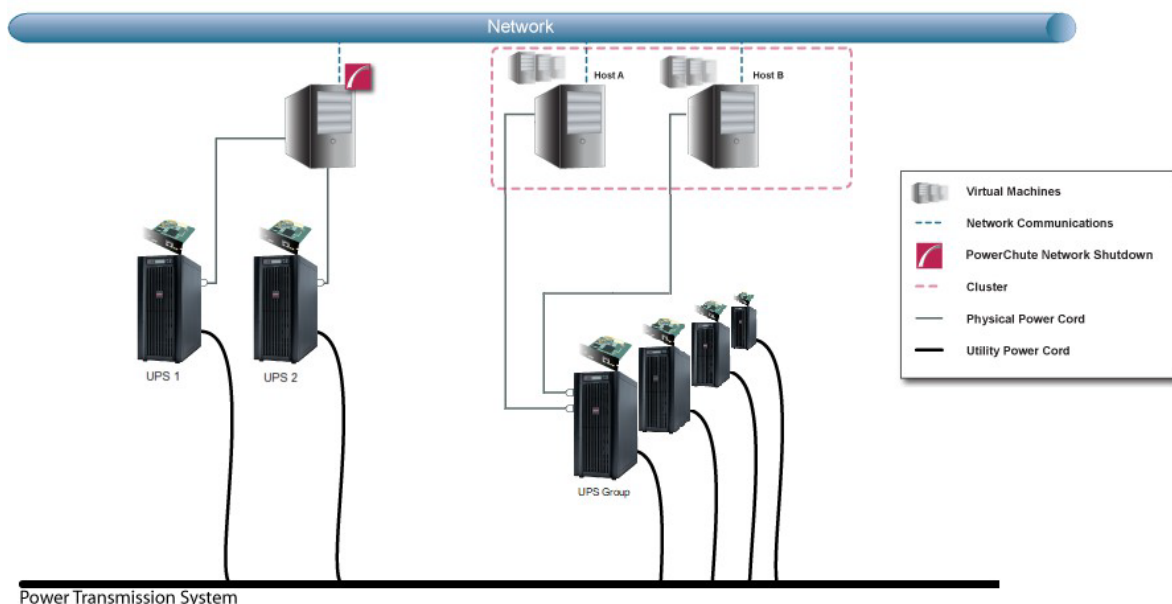
* Virtualization Manager Server is vCenter Server

A single PowerChute Agent can manage all the UPS's in the cluster. Each UPS protects one or more VMware Hosts.

For detailed information, please view the “Using PowerChute Network Shutdown in an Advanced Redundant Setup” Application Note [here](#).

Advanced UPS Configuration - Nutanix Support

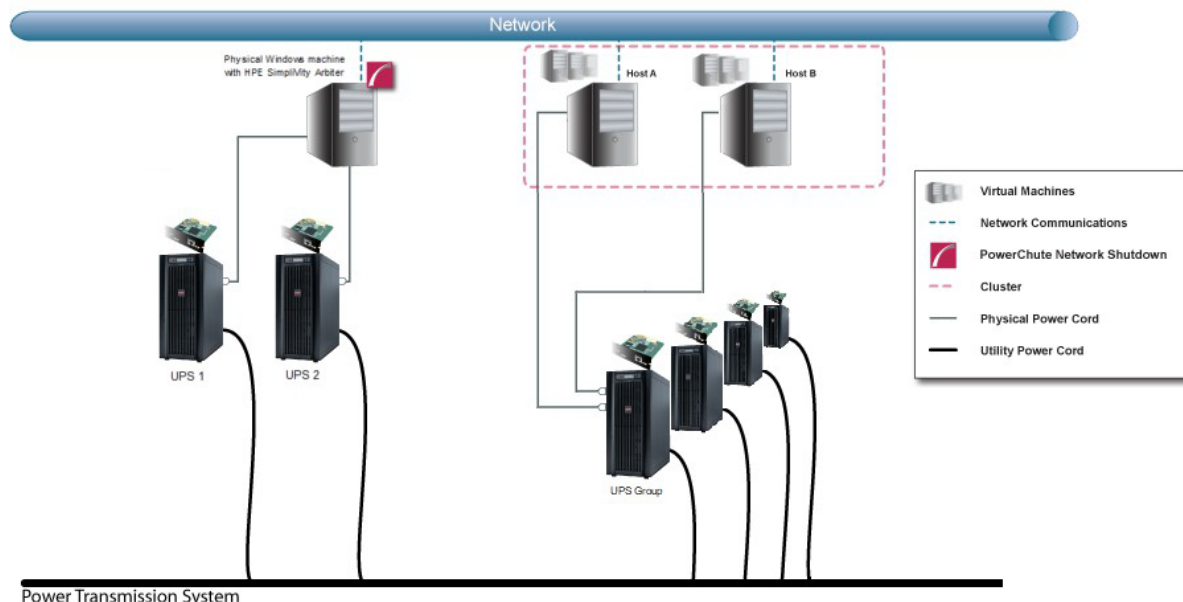
Advanced UPS Configuration: PowerChute can monitor both Single UPS's and groups of Redundant UPS's protecting your virtualization environment. If using Redundant UPS groups, redundancy levels can be set on a per group basis e.g. N+1, N+2.



i In an advanced UPS configuration, all hosts in the Nutanix Cluster must be protected by the same UPS or UPS group.

Advanced UPS Configuration - HPE SimpliVity Support

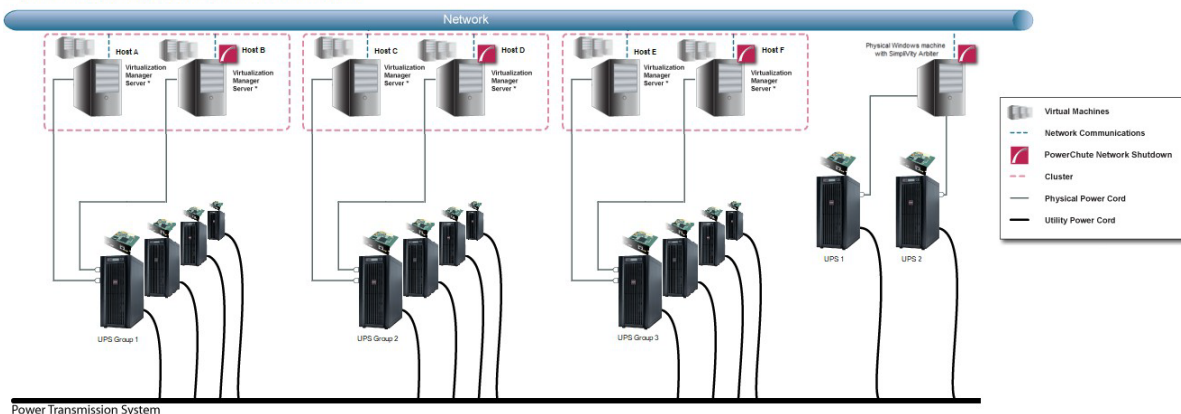
Advanced UPS Configuration: PowerChute can monitor both Single UPS's and groups of Redundant UPS's protecting your virtualization environment. If using Redundant UPS groups, redundancy levels can be set on a per group basis e.g. N+1, N+2.




i In an advanced UPS configuration, there must be one UPS setup protecting the HPE SimpliVity Cluster, and one UPS setup protecting the external physical Windows machine with HPE SimpliVity Arbiter and PowerChute. vCenter Server may be deployed as a VM on the HPE SimpliVity Cluster or on the same physical machine as the Arbiter service and PowerChute. Hosts in the HPE SimpliVity Cluster must be mapped to the same UPS group.

HPE SimpliVity Federation

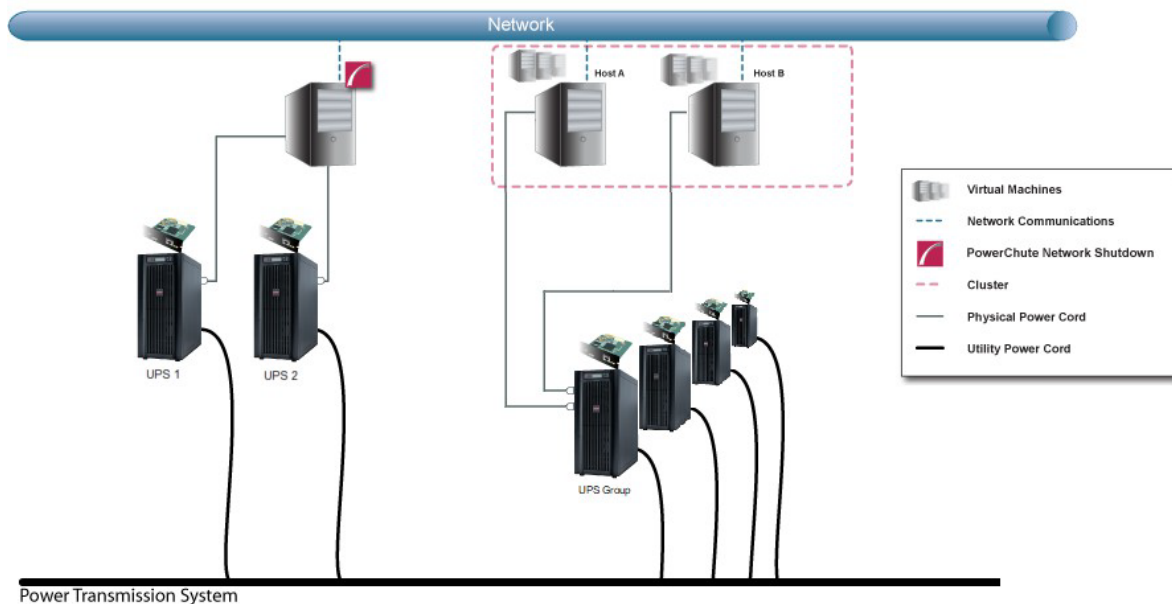
HPE SimpliVity Federation: Multiple HPE SimpliVity Clusters with a PowerChute instance for each cluster, and PowerChute installed on a physical Windows machine outside the cluster with HPE SimpliVity Arbitrator.




 If you have multiple HPE SimpliVity Clusters, it is recommended to have one instance of PowerChute per cluster.

Advanced UPS Configuration - Cisco HyperFlex Support

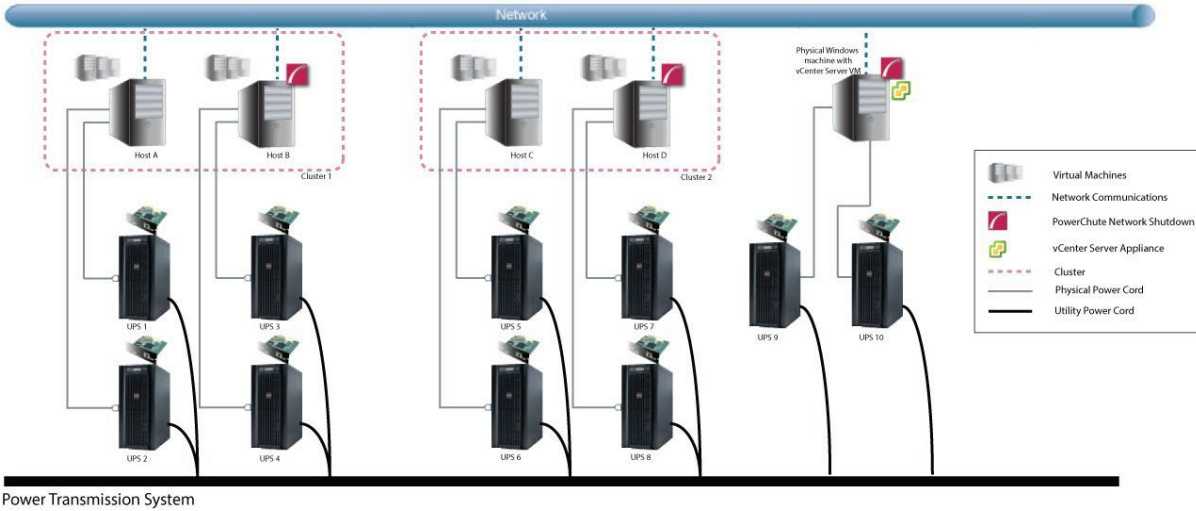
Advanced UPS Configuration: PowerChute can monitor both Single UPS's and groups of Redundant UPS's protecting your virtualization environment. If using Redundant UPS groups, redundancy levels can be set on a per group basis e.g. N+1, N+2.




 In an advanced UPS configuration, all hosts in the HyperFlex Cluster must be protected by the same UPS or UPS group.

Cisco HyperFlex Remote Office / Branch Office (ROBO) Configuration

Cisco HyperFlex Remote Office / Branch Office (ROBO): A 2-node HyperFlex Edge Cluster with a central vCenter Server outside the cluster.

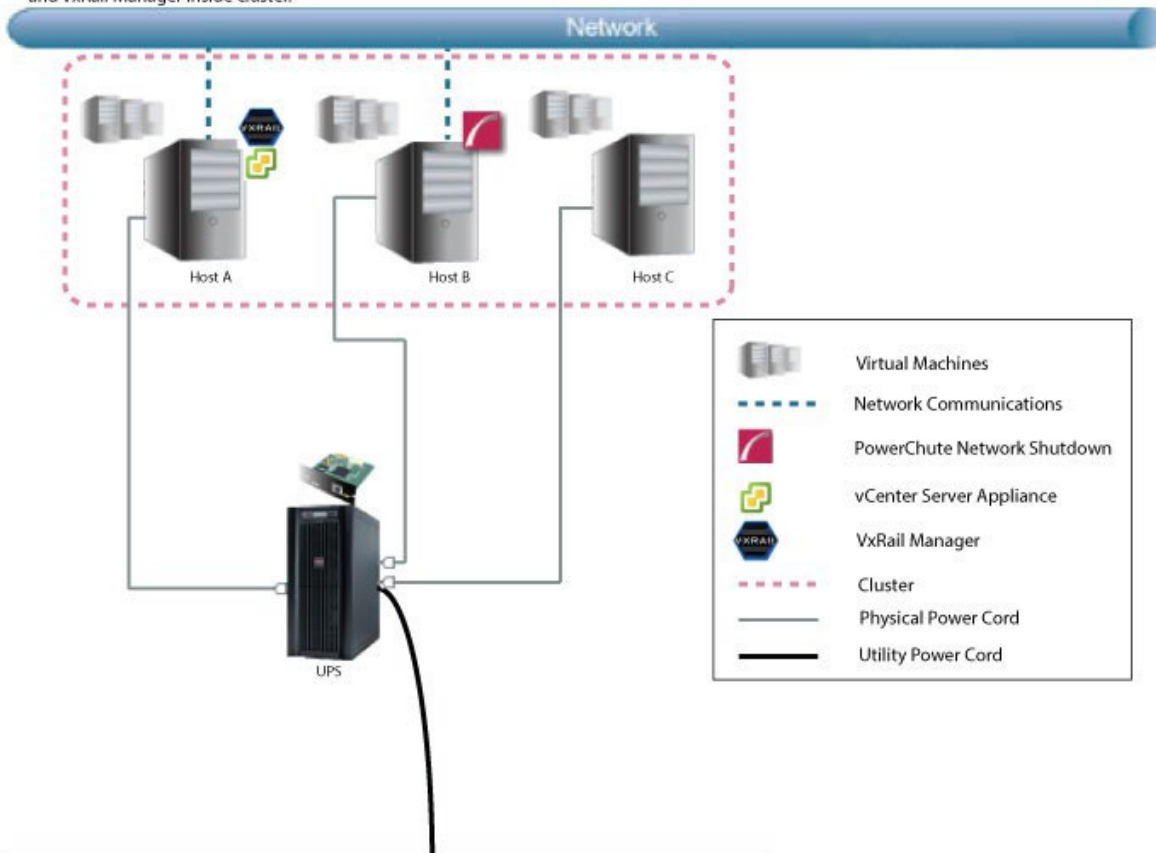


Power Transmission System

 If there are multiple HyperFlex Edge Clusters in a Remote Office / Branch Office (ROBO) configuration, it is recommended to have one instance of PowerChute per cluster.

Dell VxRail Single Configuration

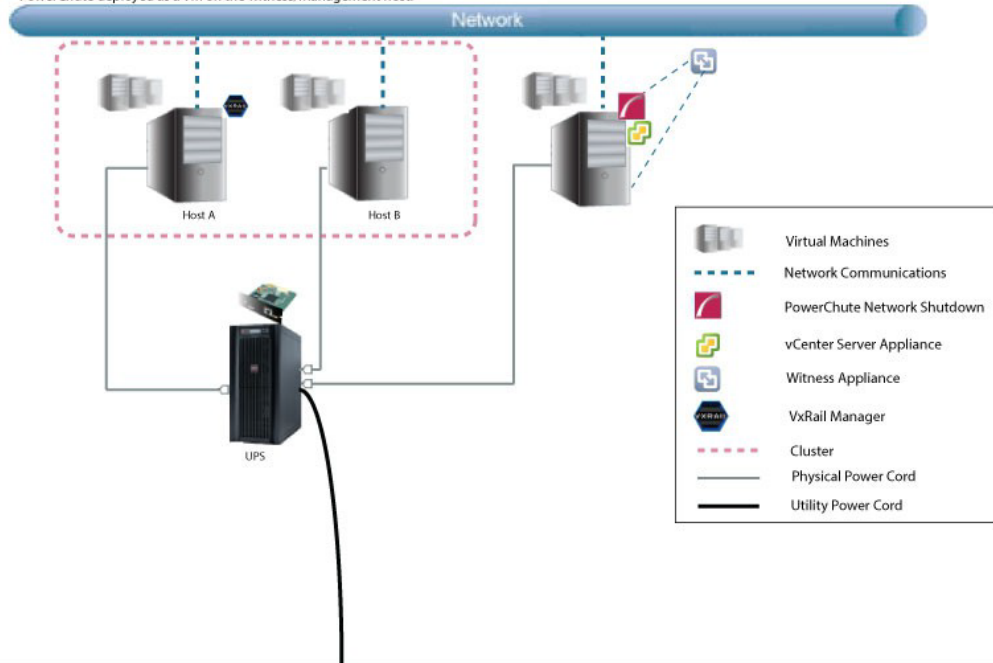
Dell VxRail Single UPS Configuration: A 3-node Dell VxRail Cluster with PowerChute deployed as VM inside cluster, and vCenter Server and VxRail Manager inside cluster.




Power Transmission System

Dell VxRail Stretched Cluster Configuration

Dell VxRail Stretched Cluster Configuration: A 2-node Dell VxRail Cluster with VxRail Manager inside cluster, vCenter Server outside cluster, and PowerChute deployed as a VM on the witness/management host.



Power Transmission System

 Advanced UPS configurations are not supported for Dell VxRail

Network Management Card Connection

Network Management Card 2 firmware v6.8.0 and higher, and Network Management Card 3 firmware v1.1.0.16 and higher uses the HTTPS protocol by default. The default protocol for PowerChute is HTTPS which can be changed to HTTP, although it is not recommended.

The default port is 80 for HTTP, and 443 for HTTPS. Do not change this number unless you changed the port being used by your NMC.

For Redundant and Parallel configurations, you need to enter more than one **IP address** to enable communications with all the relevant NMCs.

For more information on UPS configurations and supported UPS models, view the Application Note "PowerChute Network Shutdown Operating Modes and supported UPS Configurations" [here](#).

Add each IP address using the **+ Add IP Address** button. Enter the IP address of the NMC in the UPS. Click **OK**.

The screenshot shows the "PowerChute Setup: UPS Details" window. It features a "Protocol" dropdown menu set to "https" and a "Port" text box containing "443". Below these is a section titled "Network Management Cards" with a "+ Add IP Address" button. This section contains a list of two IP addresses: "1.1.1.1" and "1.1.1.2". Each IP address has a trash can icon for deletion and a pencil icon for editing.

To edit an IP address, click the  icon. To delete an IP address, click the  icon.

Certificate Management

The registration process will be unsuccessful if you attempt to register an NMC with an untrusted SSL certificate and the "Error registering with the Network Management Card(s)" error message will be displayed. Click the **Fix Issues** button to view the Certificate Management page where you can view the certificate details including the hostname, and expiry date.



You can only add the NMC host certificate to the PowerChute-keystore via this method. You can add CA certificates to the PowerChute-keystore using the keytool command. Please refer to the steps outlined in the "Changing the password for the Java Keystore" topic in the [PowerChute Security Handbook](#) for more information.



If you exit the PowerChute Setup wizard without adding the NMC SSL certificate(s) to the **PowerChute-keystore**, PowerChute will not be able to communicate properly with the NMC. Ensure that you complete the NMC registration process in the PowerChute Setup wizard and add any NMC SSL certificates to the **PowerChute-keystore** if you trust them.

Click the  button to view additional details of a certificate including:

- Version
- Subject
- Issuer
- Serial Number
- Valid From
- Valid To
- Public Key Information
- Signature Algorithm
- Certificate Chain

Check the checkbox and click the **Accept Certificate** button to add the displayed certificate to the **PowerChute-keystore**. If there are multiple certificates listed in the Certificate Chain, you can select them and import them at the same time. When all certificates have been validated, click **Next** to continue with your PowerChute setup. You should only add certificates you recognize and trust to the **PowerChute-keystore**.



PowerChute does not validate certificates added to the **PowerChute-keystore**. PowerChute will continue to trust certificates in the **PowerChute-keystore** after they have expired or been revoked. It is your responsibility to remove these certificates from the **PowerChute-keystore** once they are no longer needed.

Verify Certificates in PowerChute-keystore

You can verify that the displayed certificates were successfully added to the **PowerChute-keystore**:

1. Launch a terminal or command prompt and change the directory to the group1 directory of where PowerChute is installed. If the default location was chosen during installation, this will be:
 - C:\Program Files\APC\PowerChute\group1 for Windows machines
 - opt/APC/PowerChute/group1 for Linux machines
2. Run the command:
 - ..\jre_x64\bin\keytool.exe -list -keystore PowerChute-keystore for Windows machines
 - ../jre_x64/bin/keytool -list -keystore PowerChute-keystore for Linux machines
3. When prompted, enter the password you created for the PowerChute-keystore.



You must manually change the **PowerChute-keystore** password before you can access it. Please refer to the "Changing the password for the Java Keystore" topic in the [PowerChute Security Handbook](#).

4. The certificates and their fingerprints will be displayed.

```

Microsoft Windows [Version 10.0.19043.1766]
(c) Microsoft Corporation. All rights reserved.

C:\Users\Dr. Bunsen Honeydew>cd "%Program Files\APC\PowerChute\group1"

C:\Program Files\APC\PowerChute\group1>keytool --list --keystore PowerChute-keystore
Enter keystore password:

***** WARNING WARNING WARNING *****
* The integrity of the information stored in your keystore *
* has NOT been verified! In order to verify its integrity, *
* you must provide your keystore password. *
***** WARNING WARNING WARNING *****

Keystore type: JKS
Keystore provider: SUN

Your keystore contains 1 entry

apc1c7234.gal.lab, 28 Jun 2022, trustedCertEntry,
Certificate fingerprint (SHA-256): FA:95:E8:8D:F5:84:AC:BF:B5:40:2A:44:DE:D3:23:2C:81:D0:E2:21:6E:F2:8C:A3:8B:0F:8D:13:C
8:DC:55:2D

C:\Program Files\APC\PowerChute\group1>_
  
```



The contents of the **PowerChute-keystore** can be viewed without a password. However, the integrity of the **PowerChute-keystore** cannot be verified without a password.



NMC PowerChute Support

In NMC 2 firmware v6.8.0 and higher and NMC 3 firmware v1.1.0.16 and higher, PowerChute support is disabled by default. This may result in an error while PowerChute registers with the NMC(s). If an error is displayed, log in to the NMC UI and ensure PowerChute support is enabled.

In the NMC Web UI, you must specify a user name and authentication phrase before PowerChute can be enabled. You must also choose the protocol used to communicate with PowerChute (HTTP/HTTPS). **NOTE:** The chosen protocol must be enabled on the NMC before PowerChute communications can be established. For more information, refer to the [NMC User Guide](#).

Advanced UPS Setups

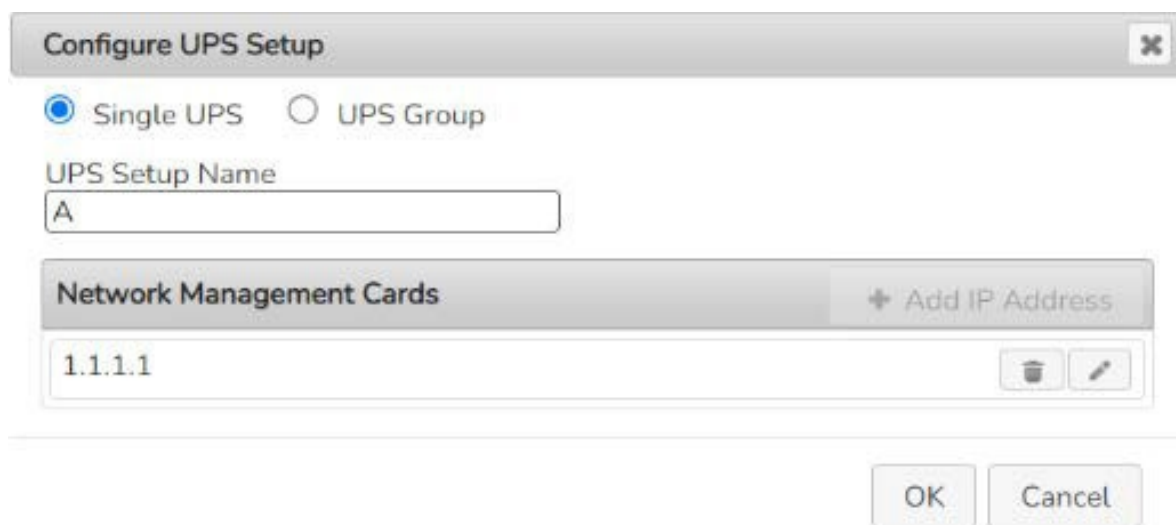
Add UPS Setup

In an Advanced UPS configuration, a single instance of PowerChute Network Shutdown can monitor multiple UPS setups and initiate graceful shutdown of equipment based on different redundancy levels. Each setup can be a single UPS or a UPS group. A single UPS setup is represented by the  icon. A UPS group is represented by the  icon.

For example, one setup may be a group of UPS's that are configured with N+2 redundancy. Another setup may be a single UPS.

On the **UPS Details** page of the Setup Wizard, click the **+ Add UPS(s)** button to create a new setup.

To create a setup with a single UPS, on the Configure UPS Setup dialog choose **Single UPS**:





1. Enter a **UPS Setup Name** (with a maximum of 20 ASCII characters)
2. Click the **+ Add IP Address** button and enter the IP address of the Network Management card in the UPS. Click **OK**.
3. Click **OK** to complete Single UPS Setup.


To create a setup with a group of UPS devices, choose **UPS Group**:

The screenshot shows a dialog box titled "Configure UPS Setup". At the top right is a close button (X). Below the title bar are two radio buttons: "Single UPS" (unselected) and "UPS Group" (selected). Underneath is a text field labeled "UPS Setup Name" containing the character "A". A section titled "Network Management Cards" contains a list of two IP addresses: "1.1.1.1" and "1.1.1.2". To the right of each IP address are two small icons: a trash can (delete) and a pencil (edit). Above this list is a button labeled "+ Add IP Address". At the bottom of the dialog are two buttons: "OK" and "Cancel".

1. Enter a **UPS Setup Name** (with a maximum of 20 ASCII characters)
2. Click the **+ Add IP Address** button and enter the IP address of the Network Management card in the UPS. Click **OK**.
3. Repeat for each of the UPS devices to be added to the UPS group. A minimum of 2 IP addresses is required to set up a UPS Group.
4. Click **OK** to complete Group UPS Setup.

Repeat for each UPS setup required.

To edit a UPS Setup, click the  icon. To delete a UPS setup, click the  icon. Click the **Next** button to go to the next step of the Setup Wizard.


 PowerChute has been tested with a total of 16 NMCs in an advanced configuration. However it is possible to configure for more than 16 NMCs in this configuration.

For detailed information, please view the "Using PowerChute Network Shutdown in an Advanced Redundant Setup" Application Note [here](#).

Outlet Group Registration

If your UPS supports outlet groups you must specify which one the server is being powered by so that PowerChute can monitor it for shutdown events and also issue turn-off commands to that outlet group.

If PowerChute is re-configured to use a different UPS in an Advanced UPS Configuration, you may not be able to map Hosts to the new UPS on the Host Protection page. This only occurs if the new UPS has the same UPS Setup name but a different outlet group configuration to the original one. To reconfigure a PowerChute installation with a new UPS under the same Setup name, use a different UPS Setup name when re-running the Setup Wizard to change the UPS.

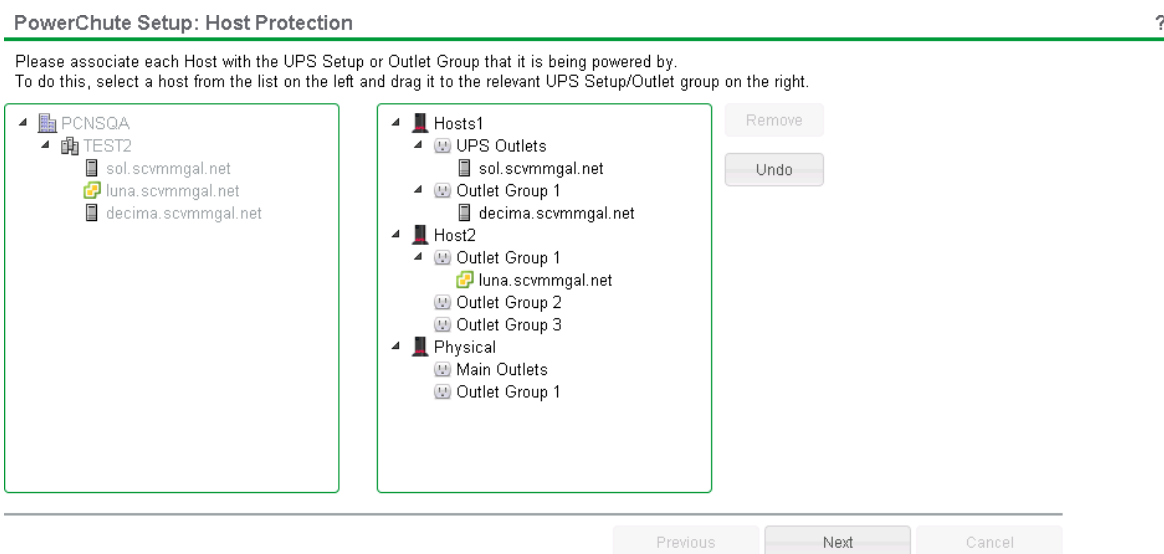
 It is not supported to map UPS outlet groups when PowerChute is configured with Nutanix, HyperFlex, SimpliVity, VxRail, or vSAN.

UPS Shutdown Behavior in Mixed UPS Environments

If your VMware hosts are being powered by a mix of outlet-aware UPS's (e.g. SMX/ SMT) and non-outlet-aware UPS's (e.g. SU/ SUA) in a Redundant UPS Configuration, PowerChute only provides the option to turn off the UPS and not the outlet group.

Your VMware hosts and their virtual machines are still protected if there is a UPS critical event or if the outlet group is commanded to shut down e.g. via the NMC User Interface.

This also applies for the **Advanced UPS configuration** if there are hosts associated with more than one outlet group on the same NMC.



PowerChute Setup - VMware hosts associated with more than one outlet group on the same NMC



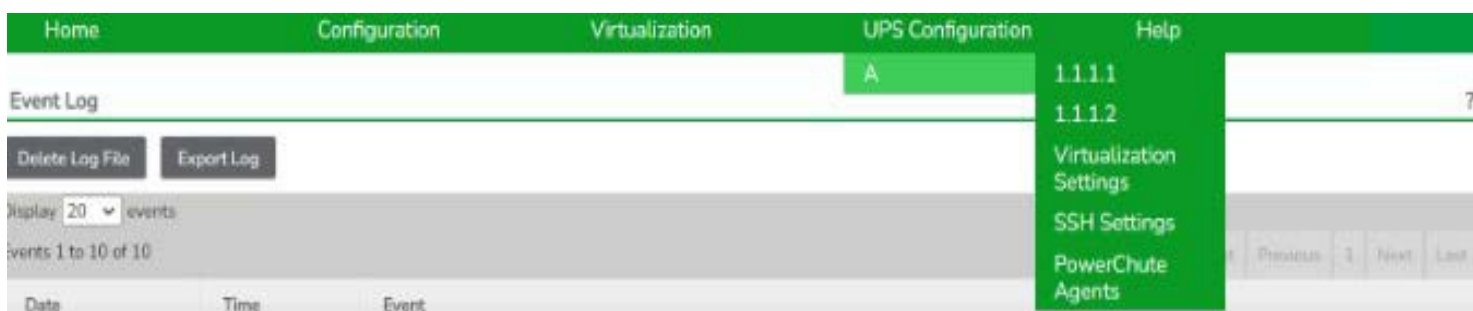
Host Protection - VMware hosts associated with more than one outlet group on the same NMC

Network Management Card Settings

For Single and Redundant and Parallel UPS configurations, the IP address of each NMC that PowerChute is communicating with is displayed under the **UPS Configuration** menu option.



For Advanced UPS configuration, each UPS Setup is displayed as a menu item and the IP address of the NMC(s) with which PowerChute is communicating is displayed under each UPS setup.



Click on the IP address to view the UPS information specific to that NMC. If the NMC has VMware Hosts associated with it, you can edit these settings for a specific NMC on this page. This overrides the global NMC settings configured via the initial PowerChute Setup or via the Shutdown Settings screen.

UPS information displayed includes:

- NMC IP Address
- UPS model name
- UPS configuration

The NMC Host Name from the NMC's DNS settings page under Network - DNS - Configuration is also displayed. This is not the same as the UPS name that can be set under Configuration - UPS General on the NMC.

Clicking the **Launch** button opens the NMC user interface.

VMware Configuration

When VMware® Support has been enabled the options below are displayed:


PowerChute Setup: VMware Configuration

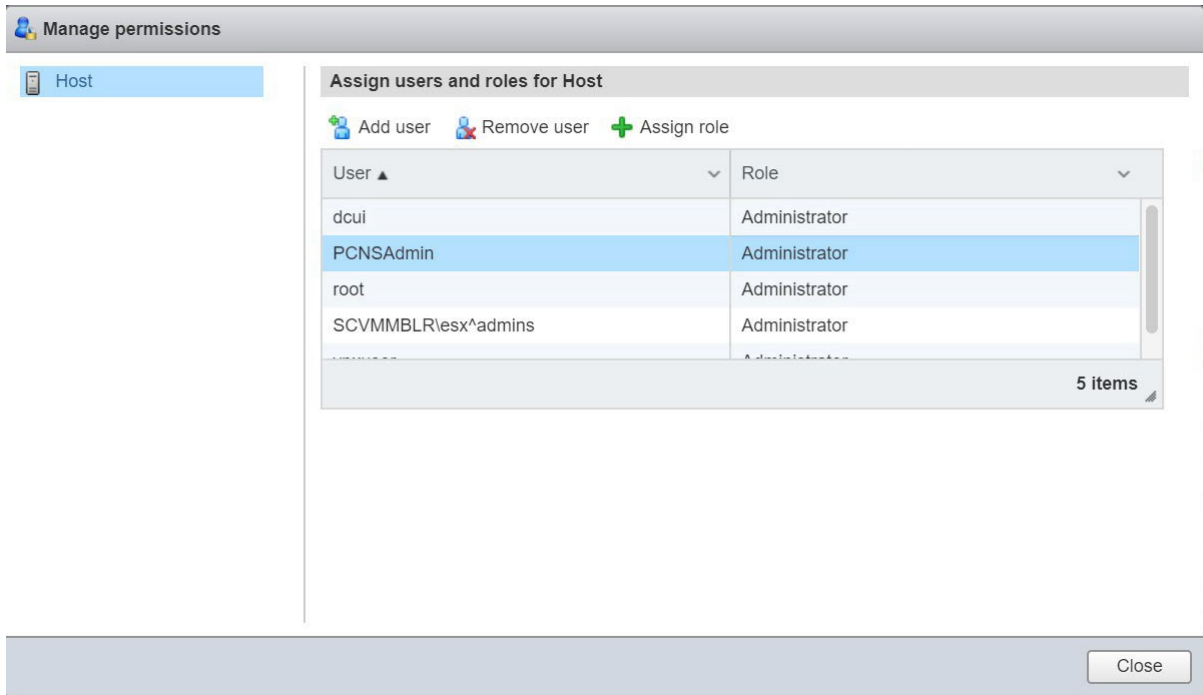
- Standalone VMware Host.
- Host managed by vCenter Server.

- Choose [Standalone VMware Host](#) to protect a single host that is not managed by vCenter Server.
- Choose [Host managed by vCenter Server](#) for HA cluster support and to manage multiple VMware hosts.

Standalone VMware Host Details

When deployed as a virtual appliance or installed on a vMA, PowerChute connects directly to the VMware host to shut it down using the credentials you enter.

 PowerChute must be configured to connect to the VMware host using an account with the Administrator role. This can be a local user account or an Active Directory User Account that is a member of the "ESX Admins" Domain Security Group. The "ESX Admins" group needs to be created via "Active Directory Users and Groups". When an ESXi host is joined to an Active Directory domain, VMware automatically assigns the Administrator role to the "ESX Admins" group.



Manage permissions

Host

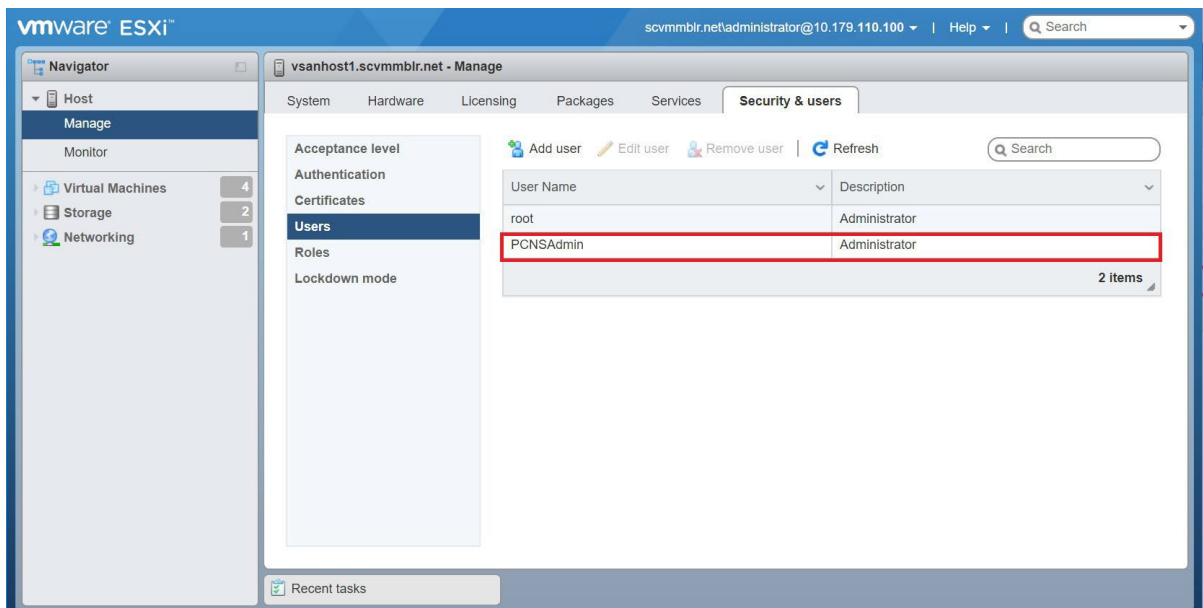
Assign users and roles for Host

+ Add user - Remove user + Assign role

User	Role
dcui	Administrator
PCNSAdmin	Administrator
root	Administrator
SCVMMBLR\esx^admins	Administrator
.....	Administrator

5 items

Close



vmware ESXi™

scvmmblr.net/administrator@10.179.110.100 | Help | Search

vsanhost1.scvmmblr.net - Manage

System Hardware Licensing Packages Services **Security & users**

Acceptance level
Authentication
Certificates
Users
Roles
Lockdown mode

+ Add user Edit user - Remove user Refresh

User Name	Description
root	Administrator
PCNSAdmin	Administrator


2 items

Recent tasks

PowerChute does not directly shut down virtual machines with this configuration. In order to gracefully shut down VMs, you must use the **Virtual Machine Startup/Shutdown** settings for the VMware host in the vSphere Client. This is accessible by clicking **Manage > System > Autostart** under the Host in the vSphere Client.

The screenshot shows the VMware ESXi vSphere Client interface. The main window displays the 'Autostart' settings for the host 'vsanhost1.scvmmblr.net'. The 'Autostart' section is expanded, showing a table of virtual machines with their shutdown behavior, autostart order, and start delay. The 'Stop action' is currently set to 'Power off'.

Virtual machine	Shutdown b...	Autostart order...	Start delay
Windows2016_150	System default	1	120 s
VSANWIN	System default	2	120 s
RHEL_150	System default	3	120 s
VSANRHEL151	System default	4	120 s

 The shutdown action must be changed to "Guest Shutdown" and VMware Tools must be installed on each VM to carry out Guest OS shutdown. If VMware Tools are not installed, the VMs will be powered off.

vCenter Server Settings

PowerChute connects to the vCenter Server to perform VM migration, VM shutdown, vApp shutdown, and VMware host shutdown operations.

It is recommended that you configure an Active Directory user account with the Administrator role for vCenter Server and the VMware hosts being managed by PowerChute.



If Active Directory is not available, it is recommended that you configure a local user account with the Administrator role that exists on vCenter Server and on each of the VMware hosts being managed by PowerChute.

If vCenter Server is running on a VM you must configure an Active Directory account or shared Local User account for host shutdown commands to work correctly. For a [Dell VxRail](#) standard cluster configuration, you must configure an Active Directory account. For more information see [Active Directory VMware Configuration](#).

If the vCenter Server is unavailable when a critical UPS event occurs, PowerChute will still be able to connect directly to the VMware hosts using this Active Directory or shared local user account to shut down VMs and the hosts themselves. VM migration and vApp shutdown and startup are not supported if the vCenter Server is unavailable.

When attempting to connect to vCenter Server, if an untrusted SSL certificate is detected, a dialog to view the certificate will appear. Please verify that the certificate is correct and accept it in order to proceed. The certificate will be added to the PowerChute-keystore.

If the vCenter Server is unavailable in a [Dell VxRail configuration](#), the VxRail API shut down will be unsuccessful and prevent the hosts from being shut down. If vCenter Server is not available, PowerChute cannot provide graceful shut down of the VxRail cluster.

If the vCenter Server is unavailable during shutdown and not available during VM startup, any vSphere Cluster Services (vCLS) VMs in the cluster will not be started and High Availability (HA) cannot be enabled by PowerChute.

The screenshot shows the PowerChute Network Shutdown web interface. At the top, there is a navigation bar with links for Home, Configuration, Virtualization, UPS Configuration, and Help. The user is logged in as 'powerchuteserver' and can click 'Logout'.

The main content area is divided into three sections:

- vCenter Server Details:** This section contains several input fields:
 - vCenter Server Protocol: A dropdown menu set to 'https'.
 - vCenter Server Port: A text input field containing '443'.
 - vCenter Server IP Address/Hostname: A text input field containing '10.1.1.1'.
 - vCenter Server Username: A text input field containing 'UserOne'.
 - vCenter Server Password: A text input field.
 - vCenter Server running on a VM: A checkbox that is currently unchecked.
- PowerChute Access:** This section has a heading and a sub-heading 'Select the access protocol for the user interface.' Below this is a 'Protocol' dropdown menu set to 'HTTPS - Encrypted'.
- PowerChute Security:** This section has a heading and a sub-heading 'The information below is used to log on to PowerChute Network Shutdown and for authentication with the Network Management Card.' It contains three input fields:
 - Administrator Username: A text input field containing 'UserName'.
 - Administrator Password: A text input field containing '<hidden password>'.
 - Administrator Authentication Phrase: A text input field containing '*****'.
 Below these fields are three buttons: 'Check Details', 'Apply', and 'Undo'.

Communications Settings - Main UI - vCenter Server running on a VM option

If vCenter Server is running on a VM, the option **vCenter Server Running on a VM** must be selected so that PowerChute can perform additional validation when trying to locate the vCenter Server VM and its parent host.

- DNS/Hostname resolution issues can lead to a problem where PowerChute cannot correctly identify the vCenter Server VM or its parent Host.
- This can also occur if VMware Tools are not installed and running on the vCenter Server VM.
- This results in vCenter Server VM being shut down too early in the sequence along with the other VMs.

By enabling the checkbox PowerChute will check for these kinds of problems and display a warning message on the Host Protection page or log an event in the Event Log. If vCenter Server is installed on a Physical machine, or on a VM that is not managed by the vCenter Server, this option should not be selected.

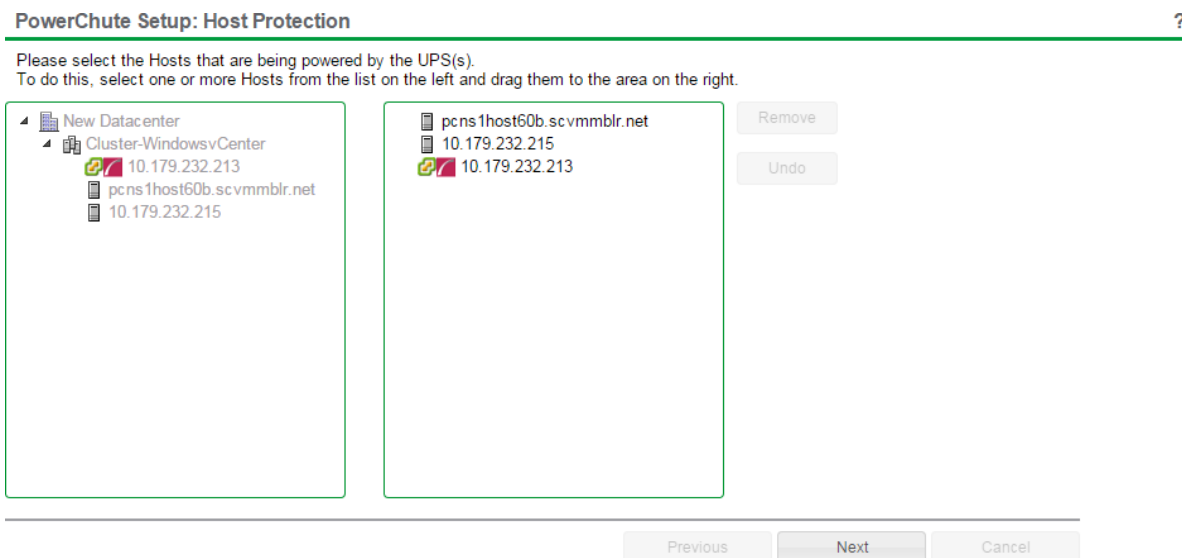
For more information on vCenter Server VM shutdown events see [Virtualization Events](#).

VMware Host Protection


Once connected to vCenter Server, PowerChute displays all of the VMware hosts in the inventory in a tree view (similar to what you see using vSphere client). From this screen you can select the hosts that PowerChute should protect.

Single and Redundant, and Parallel-UPS configurations

If PowerChute is installed on a physical Windows® machine you must specify the VMware hosts in the left-hand panel that you want to protect by dragging them to the right-hand panel of this screen.




When a critical UPS event occurs, PowerChute will shut down VMs and the Hosts in the order that they appear in the right-hand panel. You can change this order by clicking on a host in the right-hand panel and dragging it up or down.

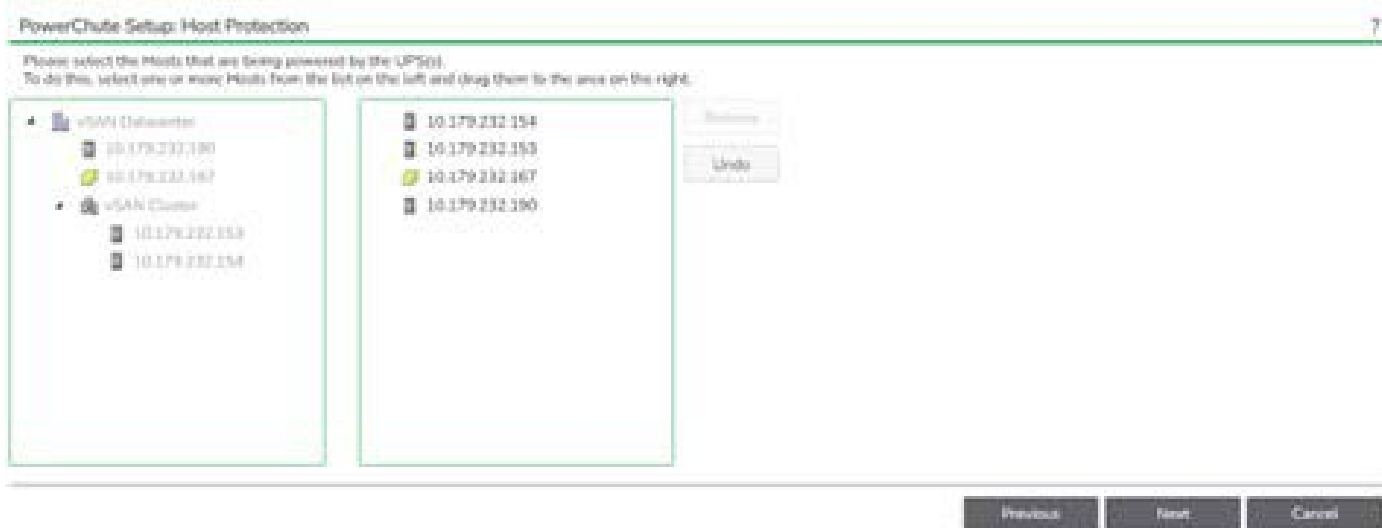
 If vCenter Server or PowerChute is running on a VM they are shut down after the other hosts, irrespective of the order that they appear in the Host Protection user interface. The other ESXi hosts are shut down in the order that they appear in the UI.


If PowerChute is installed on the vMA or deployed as a Virtual Appliance, the right-hand panel is automatically populated with all hosts in the same cluster.



The host running the PowerChute VM is automatically listed last and can be identified by the  logo.

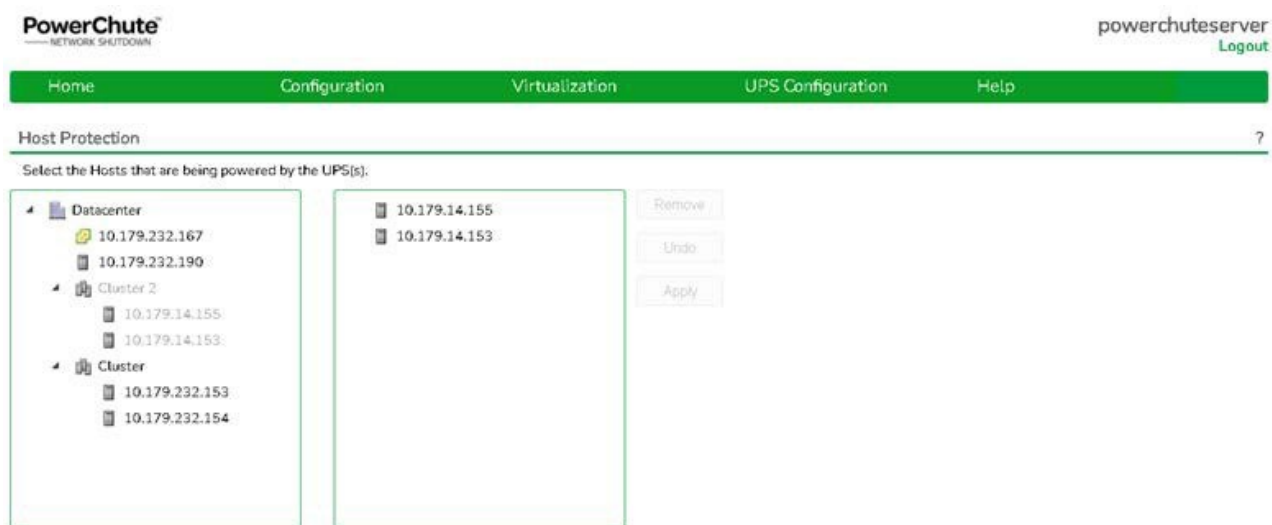
If vCenter Server is running on a VM on one of the ESXi hosts in the cluster it can be identified by this logo .

When you select Target Hosts for protection, you will be prompted to import the vCenter Server Root CA certificates if the certificates are not signed by a trusted party. To proceed with setup, view all certificates and confirm their validity.

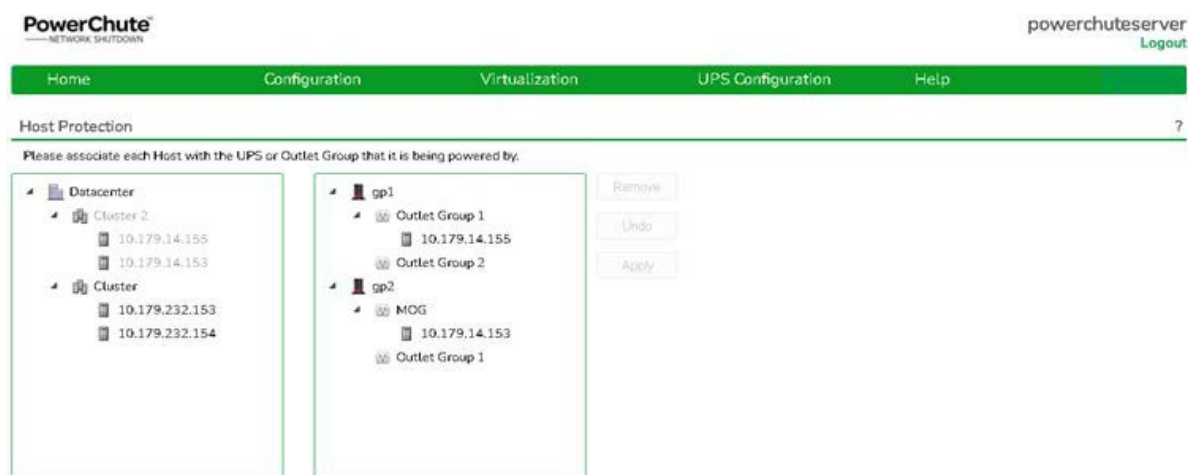


 If vCenter is running on a VM on one of the ESXi hosts and none of the Hosts are highlighted with the logo, or the wrong Host is highlighted, this indicates there is a configuration issue in vSphere setup that will prevent PowerChute from shutting down the vCenter Server VM correctly. For more information see [VMware Troubleshooting: vCenter Server VM Shutdown](#).


If vCenter Server and PowerChute are running on a VM on one of the ESXi hosts in the cluster it can be identified by these logos together:  





Advanced UPS configuration





The VMware hosts are powered separately by one or more UPS(s). The right-hand panel shows the UPS(s) that PowerChute is registered with. You need to associate each VMware host in the left-hand panel with the UPS by which it is powered. To do this, drag each host to the UPS/UPS group in the right-hand panel.

 If you have a Nutanix Cluster with VMware as your hypervisor, in an advanced configuration, all hosts in the Nutanix Cluster must be protected by the same UPS or UPS group. See [Nutanix Support](#) for more information on Nutanix.

 In PowerChute v5.1, all Cluster Hosts must be part of the same UPS setup and **Execute Virtualization Shutdown on Hosts in all UPS Setups** should be enabled on the PowerChute host or physical UPS group (if PowerChute is installed on physical Windows machine) when vSphere Cluster Services (vCLS) is present.

 Adding hosts from different vSAN datastores to the same advanced group is not a valid configuration.

 Advanced configuration is not supported for Dell VxRail.

 If the IP address/Hostname of the vCenter Server or any of the VMware hosts is changed it will be necessary to re-associate the hosts with the UPS's. When this occurs the following event is logged in the PowerChute event log:

Host(s) [Hosts] no longer exist in the vCenter Server Inventory. Please open the Host Protection page and re-select the Hosts that should be protected.

Physical UPS Setup Power Protection

In Advanced UPS Setups, PowerChute can monitor UPS's which are powering equipment outside a VMware cluster (e.g. a Storage Array Device or a physical server machine running PowerChute).

For more information see [Shutdown Settings for Advanced UPS Setups](#).



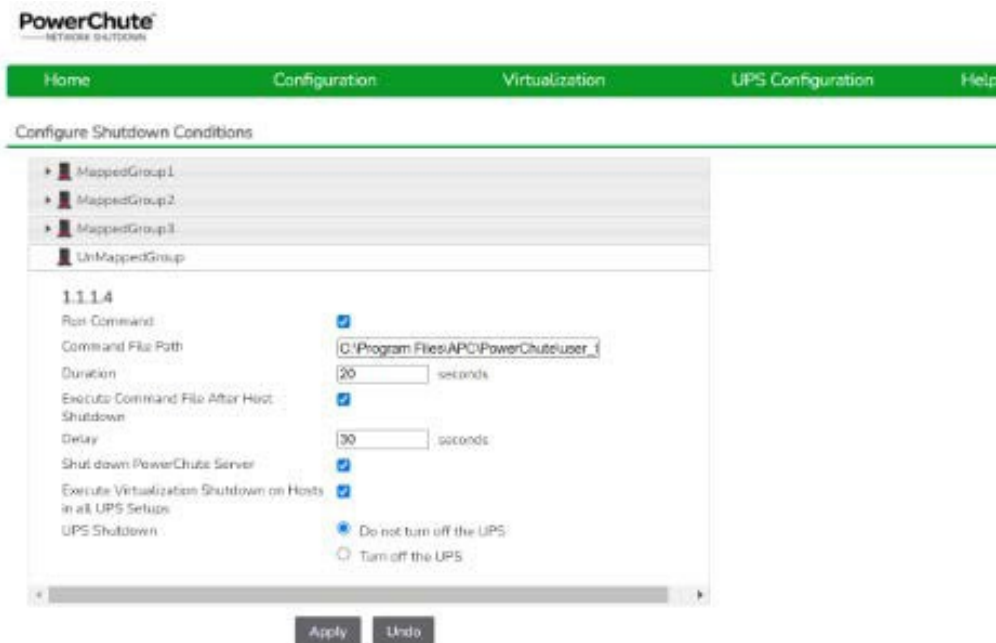
On the **VMware Host Protection** page, do not link VMware hosts with the UPS devices that are powering the physical equipment.

The following additional options will then be displayed on the Shutdown Settings page:

- **Shutdown PowerChute Server** - This is enabled by default and is used to gracefully shut down the physical machine running PowerChute. This option can be disabled if the UPS is powering a Storage Array Device.
- **Execute Virtualization Shutdown on Hosts in All UPS Setups** - This triggers a shutdown sequence using the actions configured on the [Virtualization Settings](#) page. This option should be enabled for all UPS Setups that are powering physical equipment.




If **Execute Virtualization Shutdown on Hosts in all UPS Setups** is enabled in the UI or PowerChute configuration file ("ShutdownVirtualHost" setting), all protected hosts in **all** UPS setups will be shut down when a critical event occurs on that UPS Setup.

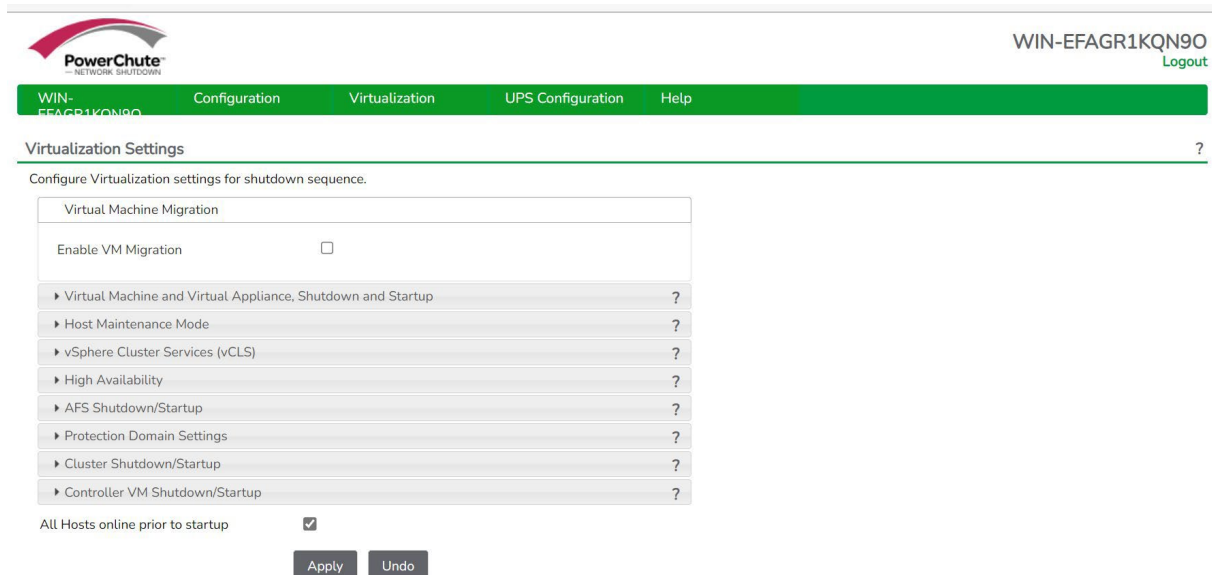


PowerChute User Interface - Configure Shutdown Conditions

Virtualization Settings

Virtual machine (VM) settings like VM Migration and VM Shutdown can be configured in Virtualization Settings.

 If hyperconverged infrastructure (HCI) support is enabled, additional sections will appear in the Virtualization Settings screen. For more information see [Hyperconverged Infrastructure Support](#).



PowerChute
— NETWORK SHUTDOWN

WIN-EFAGR1KQN90
Logout

WIN-EFAGR1KQN90 Configuration Virtualization UPS Configuration Help

Virtualization Settings ?

Configure Virtualization settings for shutdown sequence.


Virtual Machine Migration	
Enable VM Migration	<input type="checkbox"/>
▶ Virtual Machine and Virtual Appliance, Shutdown and Startup	?
▶ Host Maintenance Mode	?
▶ vSphere Cluster Services (vCLS)	?
▶ High Availability	?
▶ AFS Shutdown/Startup	?
▶ Protection Domain Settings	?
▶ Cluster Shutdown/Startup	?
▶ Controller VM Shutdown/Startup	?

All Hosts online prior to startup


Apply Undo


For more information see:

- [Virtual Machine Migration](#)
- [Virtual Machine and Virtual Appliance Shutdown/Startup](#)
- [vSAN Settings](#)
- [Host Maintenance Mode](#)
- [vCenter Server Virtual Machine Shutdown](#)
- [vSphere Cluster Services \(vCLS\)](#)
- [High Availability](#)

 vCenter Server VM Shutdown Duration is the shutdown duration given to the Virtual Machine on which vCenter Server is running. This is configurable and is set to 240 seconds by default.

If vCenter Server is running on a VM and this option is not displayed this indicates a configuration issue with vSphere. For more information see [VMware Troubleshooting](#).

 If the PowerChute virtual appliance is deployed on a vSAN host, the **Skip Maintenance Mode** checkbox must be selected.

 PowerChute will use the **No Action** flag when placing vSAN hosts into maintenance mode when all Hosts in the vSAN cluster are being shut down.

Virtualization settings in Advanced UPS Configuration

In an Advanced UPS Configuration, settings entered on this page will be applied to all UPS Setups if the checkbox **Apply Virtualization Settings to all UPS Setups** is selected. This checkbox is selected by default. If the **Apply Virtualization Settings to all UPS Setups** checkbox is disabled, you can configure different virtualization settings for individual UPS groups.

The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes Home, Configuration, Virtualization, UPS Configuration, and Help. The page title is "Virtualization Settings" with a help icon. Below the title, it says "Configure Virtualization settings for shutdown sequence." The main content area is titled "Virtual Machine Migration" and contains the following settings:

- Enable VM Migration:
- Duration: 120 seconds
- Virtual Machine and Virtual Appliance, Shutdown and Startup: ?
- vSAN Settings: ?
- Host Maintenance Mode: ?
- vSphere Cluster Services (vCLS): ?
- High Availability: ?

Below the migration settings, there are two checkboxes:

- Apply Virtualization Settings to all UPS Setups:
- All Hosts online prior to startup:

At the bottom of the form are "Apply" and "Undo" buttons.

If you have applied settings to individual UPS Setups you should uncheck this option to prevent them from being overwritten.

The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes Home, Configuration, Virtualization, UPS Configuration, and Help. The page title is "Virtualization Settings" with a help icon. Below the title, it says "Configure Virtualization settings for shutdown sequence." The main content area is titled "Virtual Machine Migration" and contains the following settings:

- Enable VM Migration:
- Duration: 300 seconds
- Select Target Hosts for Migration:
- Virtual Machine and Virtual Appliance, Shutdown and Startup: ?
- vSAN Settings: ?
- Host Maintenance Mode: ?
- vSphere Cluster Services (vCLS): ?
- High Availability: ?

Below the migration settings, there are two checkboxes:

- Apply Virtualization Settings to all UPS Setups:
- All Hosts online prior to startup:

At the bottom of the form are "Apply" and "Undo" buttons.

Individual UPS Setup Settings applied

Virtual Machine Migration

NOTES:



- It is not recommended to enable Virtual Machine Migration when all nodes in the cluster are being shut down, except when Delay Maintenance Mode is disabled and VMware DRS is set to fully automated.
- It is not recommended to enable Virtual Machine Migration when [VxRail support](#) is enabled.

If you enable Virtual Machine migration, use the **Duration** field to set the time allowed for the VMs to migrate to another healthy Host in the Cluster. VMs will not be migrated to Hosts that are powered off, in a disconnected state, in maintenance mode, or affected by a UPS critical event.

Custom Target Host Migration

By default, PowerChute will migrate VMs to any available Host in the same cluster. To control where VMs get migrated to, enable the **Select target host for Migration** option in the Virtualization Settings page of the PowerChute Setup or on the Virtualization Settings page in the main UI.

PowerChute Setup: Virtual Machine Settings

Virtual Machine Migration ?

Enable VM Migration

Duration seconds

Customize target host for Migration

PCNSLAB

- SES
 - 10.216.252.22
 - 10.216.252.122
- PCNSDEV
 - 10.216.254.29
 - 10.216.252.146

10.216.252.146

10.216.252.22

Remove

Virtual Machine and Virtual Appliance, Shutdown and Startup ?

Host Maintenance Mode ?

vCenter Server VM Shutdown ?

vSphere Plugin ?

All Hosts online prior to startup

Previous Next

Select target host for Migration - PowerChute Setup Single/Redundant/Parallel UPS Configuration



In Advanced configurations, Select target host for Migration is not available in the Setup wizard, to prevent all UPS Setups using the same set of Target Hosts. It is defined in Virtualization Settings of each individual UPS Setup instead.


The screenshot shows the 'Virtualization Settings' page in the PowerChute interface. The 'Virtual Machine Migration' section is active, with 'Enable VM Migration' checked and a duration of 120 seconds. Under 'Select Target Hosts for Migration', the host 10.179.232.190 is selected from a list that includes Datacenter (10.179.232.167, 10.179.232.190), Cluster 2 (10.179.14.155), and Cluster. Below this, 'All Hosts online prior to startup' is also checked. The 'Apply' and 'Undo' buttons are visible at the bottom.

Select target host for Migration - Single/Redundant/Parallel Configuration

This screenshot is similar to the previous one but shows a different configuration. 'Enable VM Migration' is checked, and the duration is set to 300 seconds. The same target host, 10.179.232.190, is selected from the list. The 'Apply' and 'Undo' buttons are at the bottom.

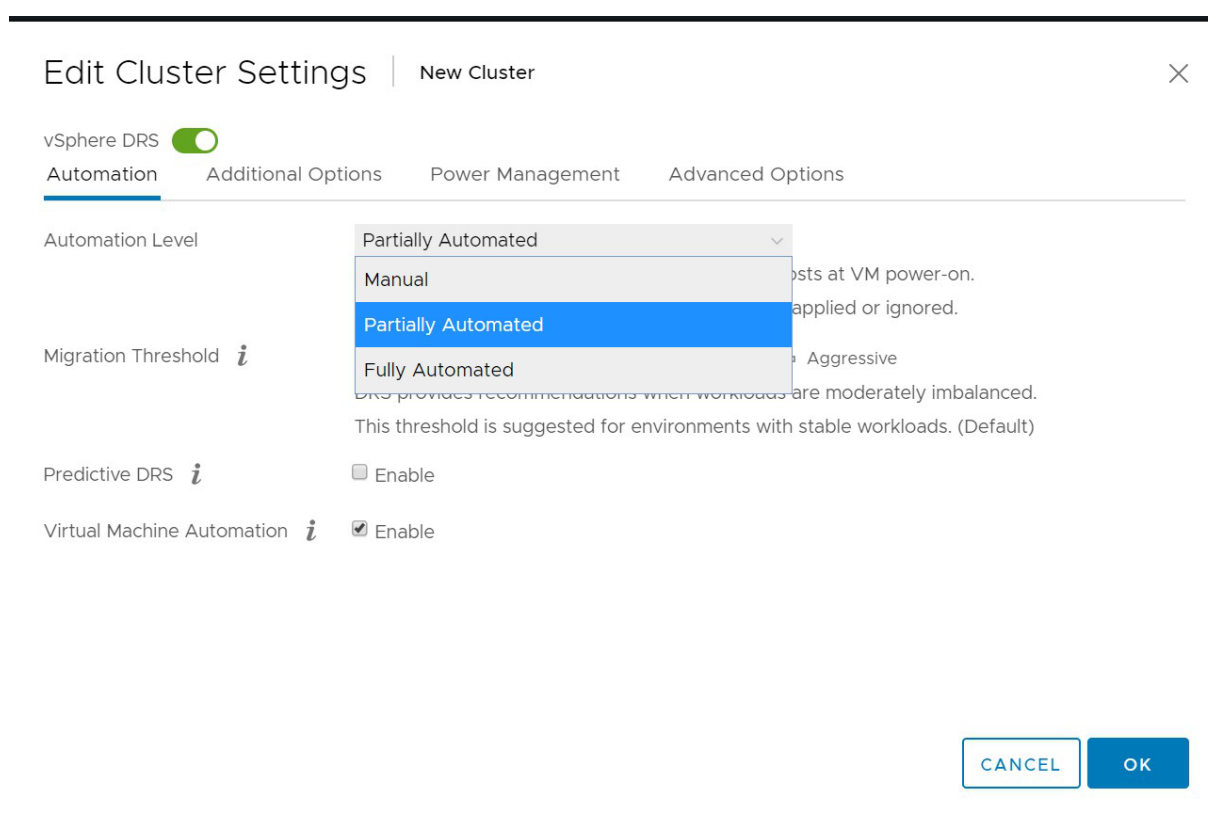
Select target host for Migration - Advanced Configuration

The left pane shows all available hosts in the datacenter. To specify a host to which the VMs should migrate in the event of a UPS critical event, drag and drop the host to the right pane. Using this option will allow VMs to migrate to hosts available in other clusters in the datacenter or Standalone Hosts in the inventory.

 Do not enable custom target host migration for Hosts that are part of a Cluster that has DRS enabled and set to fully automated, as DRS rules will take precedence and VMs may not be migrated as specified. See VM Migration using DRS below for more information.

VM Migration using DRS

If VMware DRS is enabled and set to fully automated for the Cluster, PowerChute will start a maintenance mode task on the host when a UPS critical event occurs if **Delay Maintenance Mode** and **Skip Maintenance Mode** are disabled - these are enabled by default when [HCI support](#) is enabled and for vSAN. This allows DRS to migrate the VMs to other hosts in the cluster.



Edit Cluster Settings | New Cluster

vSphere DRS

Automation | Additional Options | Power Management | Advanced Options

Automation Level: Partially Automated (selected), Manual, Partially Automated, Fully Automated

Migration Threshold: *i* Partially Automated

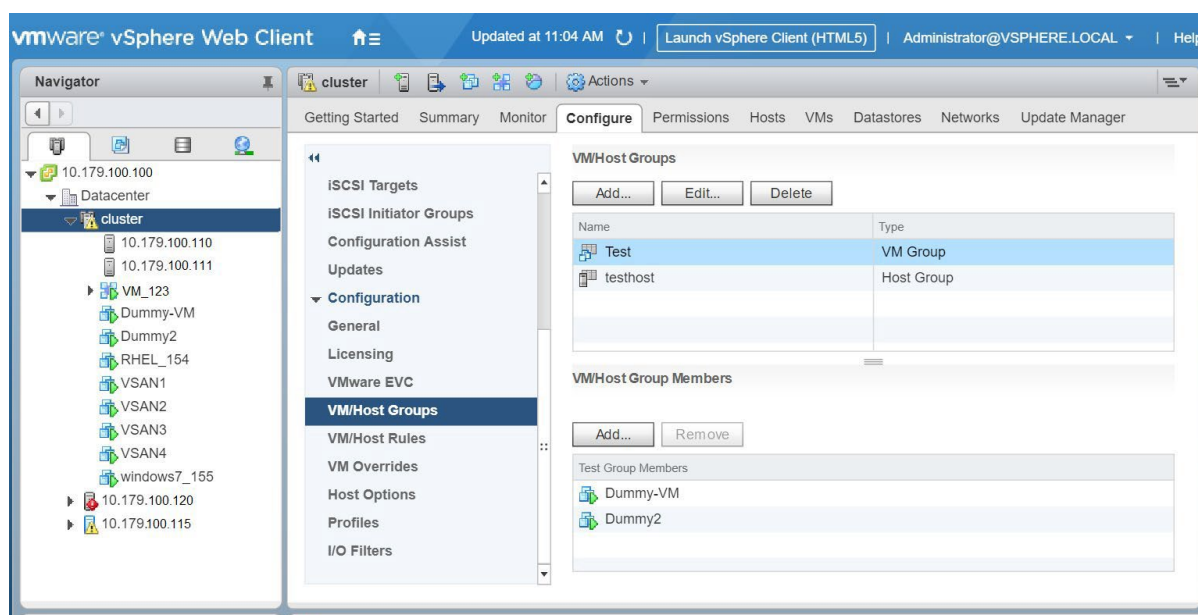
Predictive DRS: *i* Enable

Virtual Machine Automation: *i* Enable

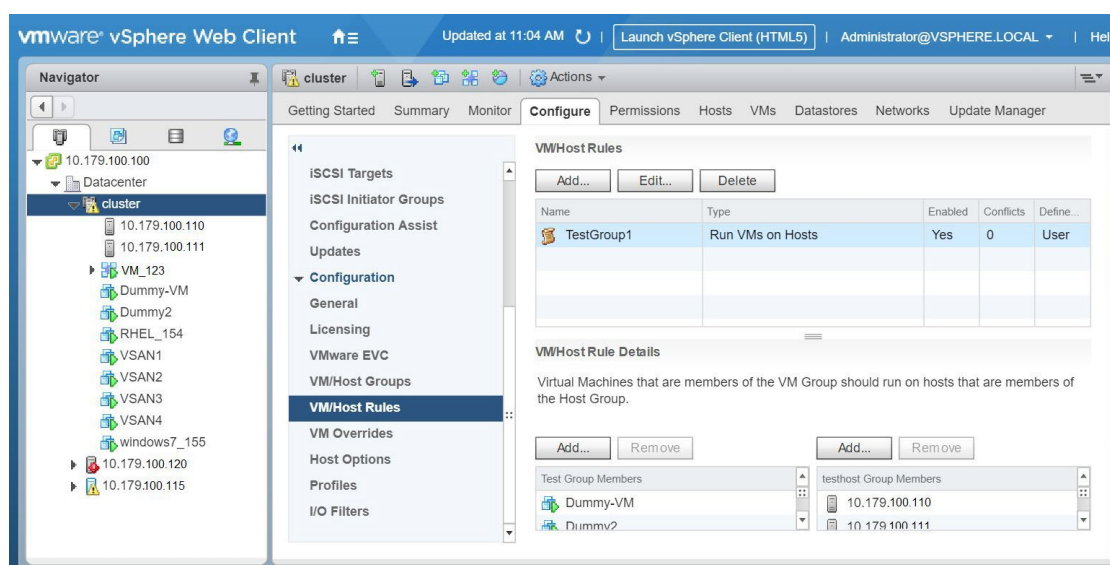
CANCEL OK

DRS is enabled in the Cluster Settings dialog of vSphere. To access Cluster settings, click on a cluster, select **Configure**, click **vSphere DRS**, and click **Edit**.


DRS Rules can be configured to control to which Hosts VMs are migrated:



1. Click on the Cluster, select **Configure**, and select **VM/Host Groups**. Click the **Add** button.



2. Specify a rule for the VM and Host DRS groups. For example, to prevent VMs from being migrated to Hosts in the Host DRS Group, select **Must Not run on hosts in group**.

 If DRS is enabled and set to fully automated and Delay Maintenance Mode is disabled, VM Migration must be enabled in PowerChute with a VM migration duration set, in order to allow Virtual Machines to migrate successfully. See [VMware Troubleshooting](#).

For more information on DRS, see "VMware Distributed Resource Scheduler (DRS)" at [VMware.com](#).

When the duration time elapses PowerChute continues to the next step in the sequence.

DRS and Host Maintenance Mode

If **Delay Maintenance Mode** is enabled in the Virtualization Settings page, and DRS is enabled and set to **Fully automated** when a critical event occurs, PowerChute will set DRS to **Partially automated** before VMs and vApps are shut down. This is done to prevent DRS from migrating VMs back to a critical host. DRS is set back to **Fully automated** when the critical event(s) have been resolved.

VM Migration without DRS

VM migration is also supported without DRS. In this instance, PowerChute will start a maintenance mode task on the host and migrate the VMs to other available hosts during the duration time specified.

If all VMs have been migrated before the duration time has elapsed, PowerChute will wait until the remaining time has passed before proceeding with the next step in the sequence. If the duration time is not long enough, any remaining VMs may not be migrated. These VMs will be shut down gracefully if VM Shutdown is selected as the next step in the sequence.

Virtual Machine and Virtual Appliance Shutdown/Startup

To ensure graceful Virtual Machine (VM)/Virtual Appliance (vApp) shutdown due to a UPS critical event, each VM/vApp must have VMware Tools installed. The **Duration** field is the time allowed for all the VMs/vApps to gracefully shut down.

If the VMs/vApps are shut down before the Duration time, PowerChute waits until this time has elapsed before proceeding to the next step in the sequence.



Using the **Duration** field, you must allow sufficient time for all your VMs/vApps to gracefully shut down before the hosts are commanded to shut down.

The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes Home, Configuration, Virtualization, UPS Configuration, and Help. The main content area is titled "Virtualization Settings" and contains a section for "Configure Virtualization settings for shutdown sequence".

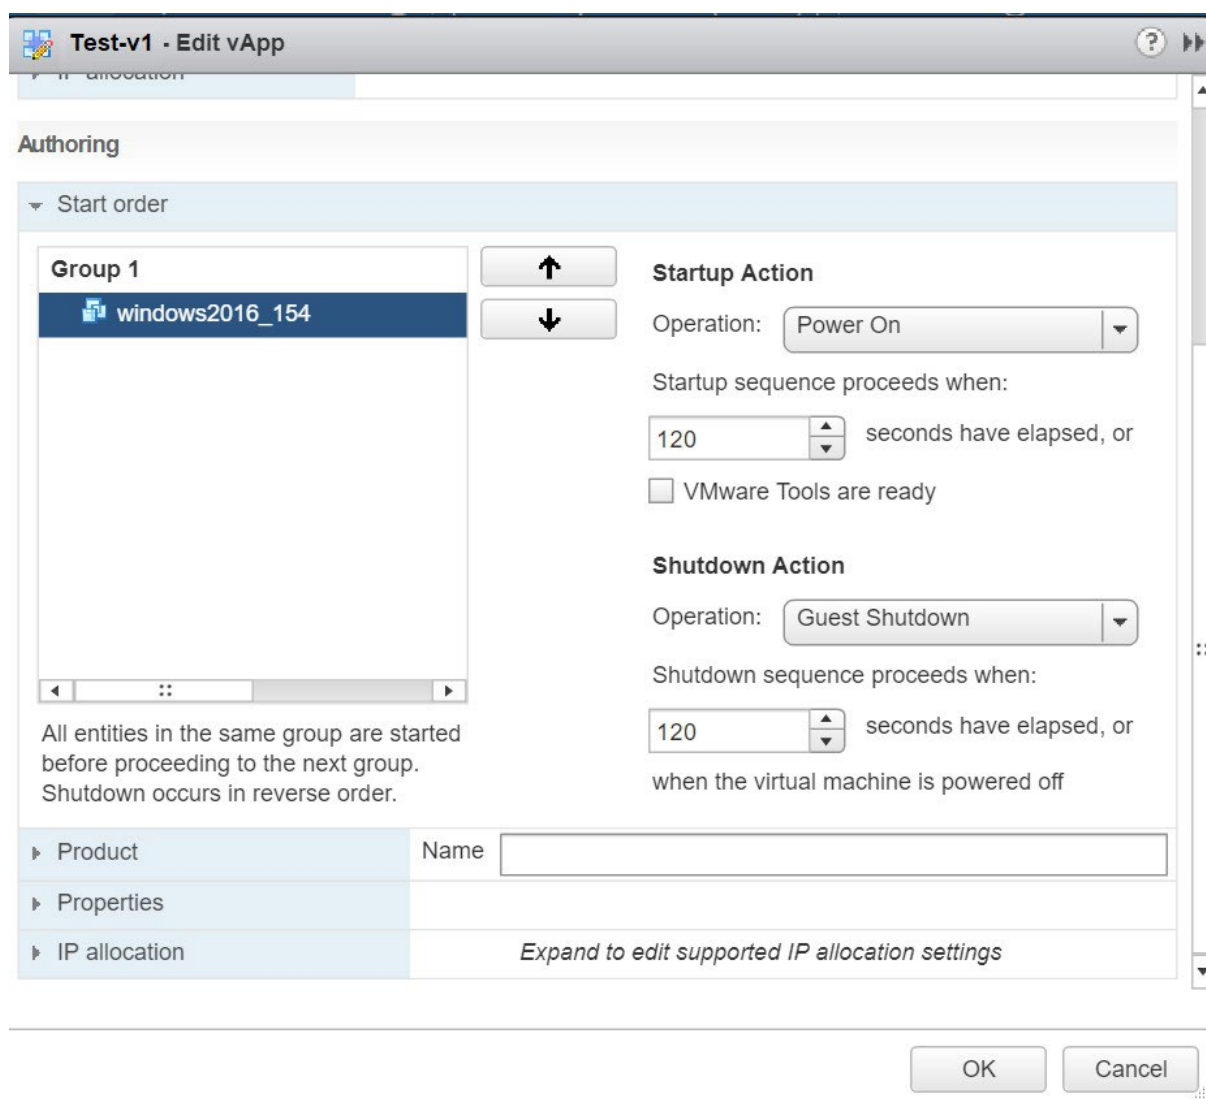
Under the "Virtual Machine Migration" section, the following settings are visible:

- Virtual Machine and Virtual Appliance, Shutdown and Startup
- VM/vApp Shutdown:
- Force vApp Shutdown:
- Shutdown Duration: 120 seconds
- VM/vApp Startup:
- Startup Duration: 240 seconds

Below these settings, there are several expandable sections: vSAN Settings, Host Maintenance Mode, vCenter Server VM Shutdown, vSphere Cluster Services (vCLS), and High Availability. At the bottom, there are two checkboxes: "Apply Virtualization Settings to all UPS Setups" (checked) and "All Hosts online prior to startup" (checked). "Apply" and "Undo" buttons are located at the bottom right of the settings area.

For vApps, the **Shutdown Action - Operation** must be changed from **Power Off** to **Guest Shutdown** to ensure that the VMs in the vApp are shut down gracefully. This can be edited through vCenter Server. Right-click the vApp in the left-hand panel of vCenter Server and click on **Edit Settings**.


By default, there is a 120 second delay between shutting down each VM in a vApp. The next VM in the vApp will not be shut down until this delay has elapsed or the current VM is powered off. This should be factored into the PowerChute VM/vApp Shutdown duration.



The **Force vApp Shutdown** check box is selected by default to ensure that, when VMs in the vApp are on different hosts, they are still gracefully shut down even if some hosts are not impacted by the UPS critical event.

If this checkbox is disabled, vApp shutdown will be skipped if some hosts are still not impacted. In addition, only the vApp VMs on the impacted hosts will be shut down. Any VMs running on hosts not impacted by the UPS event will remain running. This may cause issues if there are dependencies between the VMs in the vApp.

If vCenter Server is running on a VM and not added to a priority group, it is shut down once all other VMs/vApps have been shut down. If the vCenter Server VM is added to a priority group, it will be shut down with the other VMs in its priority group. There is a separate duration for shutting down the vCenter Server VM, which is configurable. In a [Dell VxRail configuration](#), the vCenter Server VM cannot be added to a priority group as it is required for the NMC to authenticate with VxRail Manager. PowerChute shuts itself down after sending the request to the NMC to shut down the VxRail cluster following this delay. If PowerChute is running on a VM, it does not get shut down as it is needed to shut down the VMware hosts. The PowerChute VM will be powered off when the VMware host on which it is running is shut down. If HA is enabled, the PowerChute VM will be restarted automatically once there is a healthy host available in the cluster. See [HA Admission Control](#).

 If vCenter Server is offline or unavailable when a critical event occurs, PowerChute will attempt to connect directly to the VMware hosts to shut down the VMs. To do this, an Active Directory account, or a shared local user account with the Administrator role must exist on vCenter Server and be configured separately on each VMware host. For more information see [Active Directory VMware Configuration](#).

In a [Dell VxRail configuration](#), PowerChute cannot provide graceful shutdown of the VxRail cluster if vCenter Server is offline.

Shutting down a primary vApp with nested vApps

During the PowerChute shutdown sequence, vApps are shut down in no particular order in accordance to their priority group. To shut down a primary vApp (a vApp containing vApps) and exclude the nested vApps in the primary vApp, you can use the [VM Prioritization](#) screen:

1. In the VM Prioritization screen, add the primary vApp to a priority group (for example, Medium).
2. Add the nested vApps to a lower priority group (for example, Group 1).
3. In the **Set VM Shutdown Duration** section of the VM Prioritization screen, provide a sufficient duration for the medium priority group to gracefully shut down. For example, set the shutdown duration for the medium priority group to 240 seconds.
4. Set the shutdown duration for the Group 1 priority group to 0 seconds. If you set a 0 second shutdown duration for a priority group, PowerChute will not shut down the VMs/vApps in this priority group.

Re-starting after a shutdown

Selecting **Enable VM/vApp Startup** re-starts any VMs/vApps that were shut down when a UPS critical event has been resolved and the VMware Hosts are powered on. PowerChute first checks that the VMware host is powered on and connected to vCenter Server.


In **Single and Redundant**, and **Parallel** UPS configurations, the option **Wait for all Hosts Online** is selected by default. When selected, PowerChute waits until all hosts in the cluster are back online before starting the VMs/vApps. The VMs/vApps are started on each host in parallel.


If vCenter Server is running on a VM and it was shut down by PowerChute, and in a [Dell VxRail stretched cluster configuration](#), PowerChute waits until its host comes back online before powering that VM on. PowerChute will then wait for vCenter Server to start before it starts the rest of the VMs/vApps.

Disable option **Wait for all Hosts Online** to allow PowerChute to attempt to start VMs/vApps on VMware hosts as they become available. If the vCenter Server VM Host is available when other Hosts are online, PowerChute will attempt to start the vCenter Server VM and then start VMs/vApps on other hosts.

If vCenter Server was offline or unavailable when VMs/vApps were shut down, PowerChute will start the VMs/vApps by connecting directly to the VMware hosts when they come back online after the UPS critical event has been resolved.

In **Advanced** UPS configurations, the option **Wait for all Hosts Online** is also selected by default. Where vCenter Server is running on a VM that is shut down by PowerChute, PowerChute waits for the vCenter Server host to come back online and then starts the vCenter Server VM, before attempting to start VMs/vApps on other hosts. If the **Wait for all Hosts Online** option is disabled, PowerChute will attempt to start VMs/vApps on VMware hosts as they become available.

 The **All Hosts online prior to startup** checkbox must be selected for Dell VxRail and single, redundant, and parallel UPS configurations.

 If PowerChute is running on a VM, and in a [Dell VxRail stretched cluster configuration](#), the host Witness host will remain in maintenance mode when it re-starts after a critical event is resolved until the PowerChute VM is started when its host comes online.

PowerChute will then take the hosts out of maintenance mode and start the VM/vApp.

In the PowerChute configuration file (`pcnsconfig.ini`), the `"delay_before_vmstartup"` setting allows you to delay VM/vApp start up until the user-specified duration has elapsed. This setting can be used to ensure that cluster/vSAN services have successfully started before attempting to start up the VMs/vApps in the cluster. This setting can be found in the `[HostConfigSettings]` section of the configuration file and the default value is 0. This setting is only relevant in a [Dell VxRail stretched cluster configuration](#).

Unsuccessful VM/vApp Startup

If VM/vApp startup is unsuccessful, PowerChute will wait the time specified in the **VM/vApp Startup Duration** field before attempting to restart the VMs/vApps, and the below entries are logged to the Event Log:

- Attempting to power on VMs on Host [Host] that did not start.
- Attempting to power on vApp [vApp] in datacenter [Datacenter] that did not start.

NOTES:

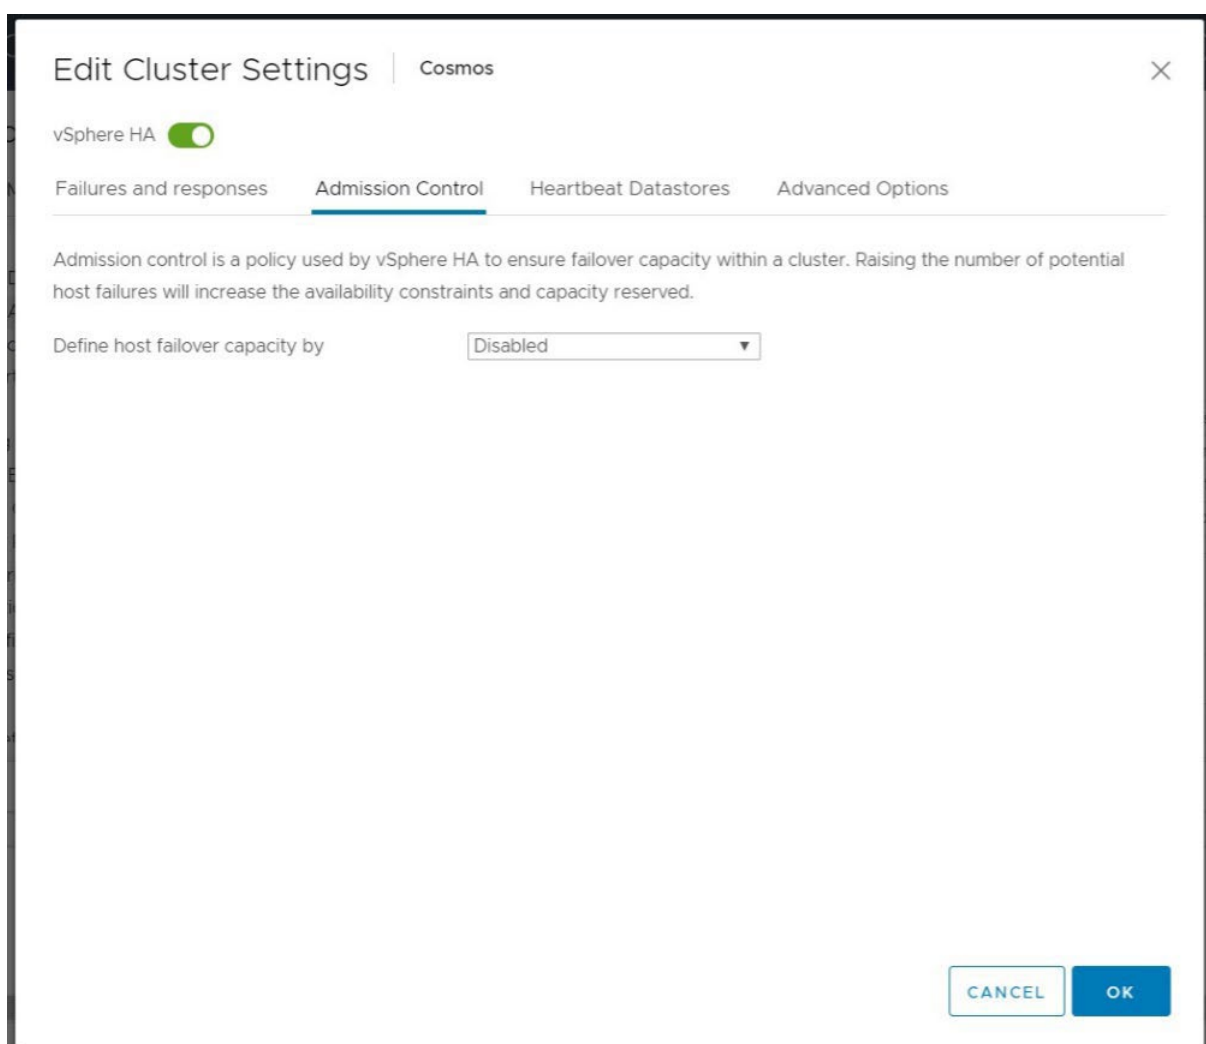


- The **VM/vApp Startup Duration** will vary if **VM Prioritization** is enabled.
- vApp startup will not be attempted if vCenter Server is offline or until vCenter Server is available. For example, when vCenter Server is deployed as a VM, VMs will get started before vApps.
- vApp names in a datacenter must be unique.

If VM/vApp startup is unsuccessful on the second attempt, no further startup attempts are made.

HA Admission Control

HA Admission Control is enabled by default in vSphere. To change Admission Control settings, click on a Cluster in the Inventory, select **Configure**. Click **vSphere Availability**, and click **Edit**.



If PowerChute is installed on a Virtual Machine, it may not be restarted automatically when its Host is powered on after a critical event has been resolved. This can occur if HA Admission Control is enabled or if the Admission Control policy being enforced prevents it.

To allow the PowerChute VM to get started automatically by HA, disable HA Admission Control or modify the Admission Control Policy to allow the PowerChute VM to start. For more information on HA Admission control settings please refer to VMware documentation.

When Admission Control is disabled, HA will attempt to automatically start the VM on which PowerChute is running and PowerChute can begin to monitor associated UPS devices and automatically restart the VMs that it shut down.

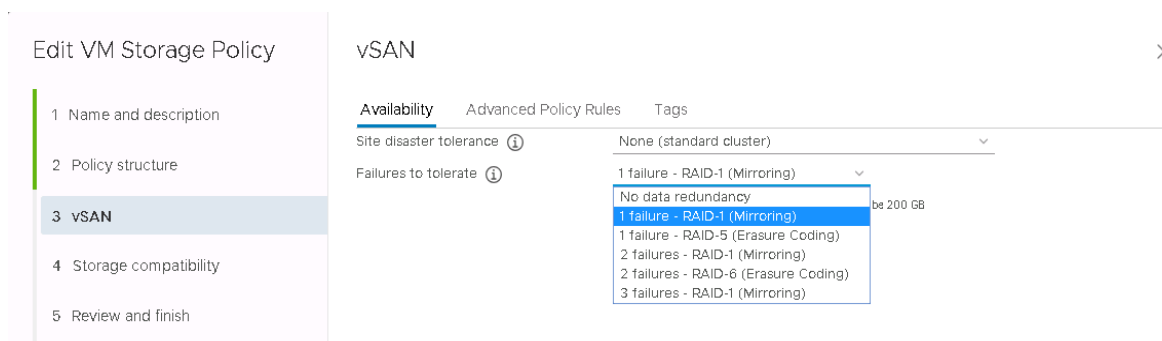
To troubleshoot VM/vApp Startup issues see [VMware Troubleshooting](#).

vSAN Settings

Fault Tolerance Threshold (FTT) support is disabled in PowerChute by default in a vSAN environment. FTT support is only supported in an Advanced UPS configuration where there are vSAN hosts in more than one UPS Setup. If you enable FTT support, PowerChute will take the FTT Level value specified into consideration when starting a maintenance mode task. This value is linked to the number of critical Advanced UPS Setups. The FTT Level configured here should match the Fault Tolerance Threshold in the Storage Policy applied to the vSAN datastore. The default vSAN Storage Policy uses FTT=1.



If PowerChute is deployed as a VM inside the vSAN Cluster, FTT support is not available.



Each PowerChute installation can support 1 vSAN Cluster if FTT is enabled (in an Advanced UPS configuration).



There should be 1 Advanced UPS setup per Fault Domain if FTT support is required.

If the number of critical Advanced UPS Setup(s) is less than or equal to the FTT Level, PowerChute starts a maintenance mode task using the Ensure accessibility flag, and only the critical host(s) will get shut down. This means that VMs can remain running after being migrated to another host in the vSAN Cluster. PowerChute will not disable High Availability and vCLS.

If the number of critical Advanced UPS Setup(s) is greater than the FTT Level, PowerChute will issue a maintenance mode command using the No Action flag. PowerChute will not disable High Availability and vCLS.

In addition to checking for critical events active on other Advanced UPS Setup(s) hosts, PowerChute checks that the hosts are healthy. A host is considered being in an unhealthy state if it is not reachable, not responding, not powered on, or in maintenance mode.

Single UPS Setup(s) (containing one host) in an unhealthy state will be considered when calculating the FTT Level when a critical event occurs. For example:

- Fault Tolerance Threshold is enabled with FTT Level of 1
- PowerChute configured with 4 single UPS Setups
- Critical event occurs on UPS Setup 1
- Host in UPS Setup 2 is in an unhealthy state
- FTT is exceeded as the number of UPS Setups with issues is greater than the FTT Level of 1

The **Shut down All Cluster VMs** checkbox is selected by default if FTT is enabled. It is not recommended to disable this setting. If the number of critical Advanced UPS Setup(s) is greater than the FTT Level, PowerChute will shut down all VMs in the Cluster and place all critical hosts in maintenance mode using the No Action flag before shutting them down. VMs/vApps running on hosts that are not directly impacted by the critical UPS event or power outage will also be shut down if they reside on the vSAN datastore. These hosts will not be placed into maintenance mode or shut down.

The **vSAN Synchronization Duration** field specifies a duration, in seconds, that PowerChute will allow any active vSAN data re-synchronization tasks to complete before shutting down the host if Delay Host Maintenance Mode is enabled. For more information, see [Host Maintenance Mode](#).

Re-starting VMs after a shutdown

The **All Hosts online prior to startup** checkbox in the Virtualization Settings page determines how vSAN Cluster VMs are powered on during the startup sequence.

- If this checkbox is selected, PowerChute will wait for all UPS critical events to be resolved before starting up VMs.
- If this checkbox is unselected, PowerChute will attempt to power on VMs if the FTT Level is sufficient to power them on.



The All Hosts online prior to startup checkbox should be selected for single, redundant, and parallel UPS configurations.

vSAN Cluster Preparation and Recovery Durations

Cluster Preparation – Duration of time for the preparation of the vSAN cluster for shutdown.

Cluster Recovery – Duration of time for the recovery of the vSAN cluster on a startup.


- Default setting of 120 seconds for both fields.
- The values are saved to the `pcnsconfig.ini` file under the “`vsan_prepare`” and “`vsan_recovery`” keys.
- These options are only available for vSAN configurations.

PowerChute Setup: Virtualization Settings

▶ Virtual Machine Migration	?
▶ Virtual Machine and Virtual Appliance, Shutdown and Startup	?
vSAN Settings	
vSAN Synchronization Duration	<input type="text" value="120"/> seconds
vSAN Cluster Preparation Duration	<input type="text" value="120"/> seconds
vSAN Cluster Recovery Duration	<input type="text" value="120"/> seconds
◀	
▶ Host Maintenance Mode	?
▶ vCenter Server VM Shutdown	?
▶ vSphere Cluster Services (vCLS)	?
▶ High Availability	?

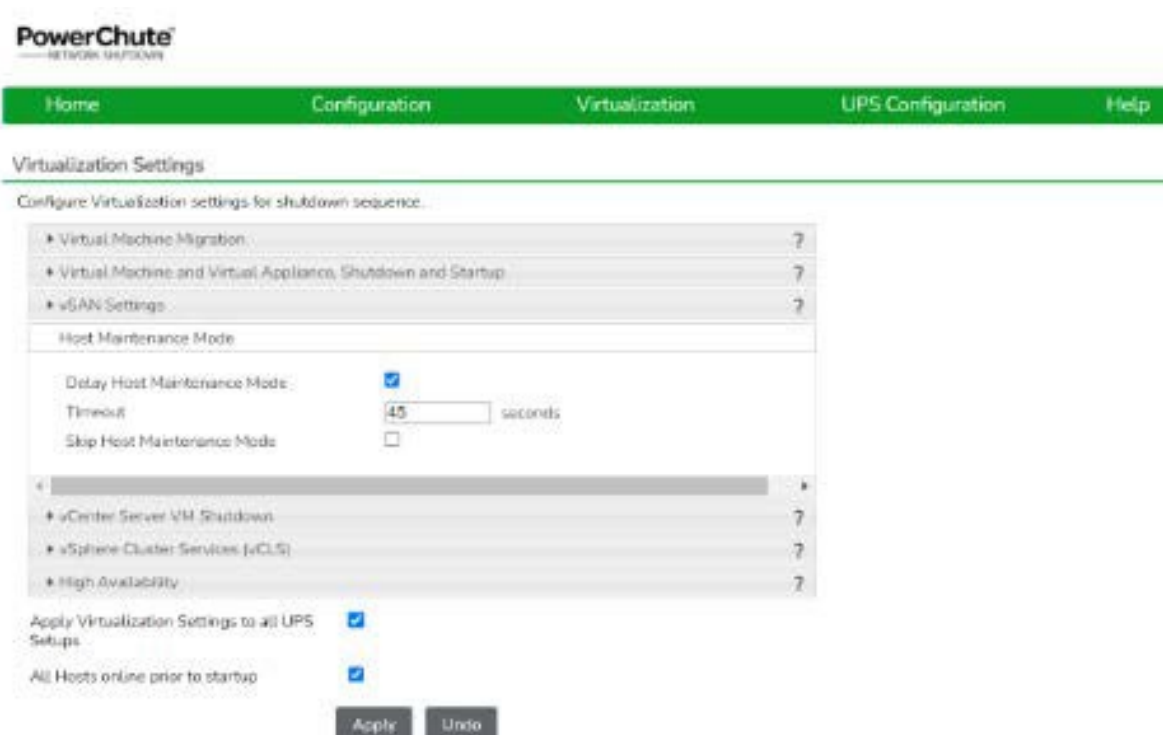
All Hosts online prior to startup

Host Maintenance Mode

 **NOTE:** In a [Dell VxRail configuration](#), Host Maintenance Mode is only relevant for a stretched cluster configuration.


If **Delay Maintenance Mode** is enabled, PowerChute starts a maintenance mode task later in the shutdown sequence, after all Virtual Machines and vApps have shut down. This option should be enabled for Clusters that have DRS set to fully automated to prevent unnecessary VM migration if all hosts in the Cluster are being shut down.


Delay Maintenance Mode should be enabled in a stretched Dell VxRail cluster configuration.



The **Timeout** field allows you to set a value, in seconds, that PowerChute will wait for the maintenance mode task to complete before shutting down the host / Witness Host and moving to the next host in the sequence then shutting down the Management Host in a Dell VxRail stretched cluster. The default value is 15 seconds. In a [Dell VxRail stretched cluster configuration](#), PowerChute will wait for the maintenance mode task to complete before shutting down the Witness Host and then shutting down the Management Host.

In a vSAN environment, the **Delay Maintenance Mode** checkbox is selected by default, and this option should be left selected. This checkbox is also selected by default, and should be left selected, if [HCI support](#) is enabled.

 **NOTE:** It is not recommended to unselect both the **Delay Maintenance Mode** and **Skip Maintenance Mode** checkboxes in an HCI environment.

 **NOTE:** In a vSAN environment, it is not recommended to put all hosts into maintenance mode at the same time.



NOTE: If PowerChute is deployed as a VM on a host in the cluster, the host cannot enter maintenance mode as it must remain powered on to send a shutdown command to its host. This will result in the following message in the Event Log, but this is expected behavior:

"Unable to start Maintenance Mode on Host [Host] as VMs are still powered on. Verify that sufficient time has been configured for VM/vApp/vCSA VM shutdown durations."

If **Delay Maintenance Mode** is enabled for a UPS Setup, hosts will be placed into maintenance mode before being shut down if all VMs are powered off. If the host is part of a vSAN Cluster, PowerChute will check for Active Synchronization in progress and wait for that to complete before shutting down the host. It is not recommended to proceed with host shutdown while Active Synchronization is in progress. See [Maintenance Mode and Data Re-Synchronization Retry](#).

In a vSAN Stretched Cluster, Witness and Management hosts will get put into maintenance mode and shut down after the vSAN Cluster hosts.



NOTE: Witness and Management hosts must be located outside the vSAN/Dell VxRail Cluster.

The value set in the **Timeout** field is the delay in between PowerChute placing each host into maintenance mode and shutting it down. The **Timeout** field is also the delay in between retrying an attempt to put a host into maintenance mode (with retry limit). PowerChute places hosts into maintenance mode and shuts the hosts down in parallel after the initial **Timeout** delay has elapsed for each host.

DRS and Host Maintenance Mode

If DRS is enabled and set to **Fully automated** when a critical event occurs, PowerChute will set DRS to **Partially automated** before VMs and vApps are shut down. This is done to prevent DRS from migrating VMs back to a critical host. DRS is set back to **Fully automated** when the critical event(s) have been resolved.

Maintenance Mode and Data Re-Synchronization Retry

In the PowerChute configuration file (`pcnsconfig.ini`), the "vsan_synch_retry_time" setting allows you to set a value (the default is 10) for retry attempts. This setting is used when retrying to put a host into maintenance mode, and when waiting for data re-synchronization on a vSAN host prior to shutting it down.


PowerChute will re-attempt to put a host into maintenance mode with the **Host Maintenance Mode Timeout** duration in between each attempt, until the task is successful, or the retry limit has been reached.

If PowerChute detects that data re-synchronization is active on a host before shutting it down, PowerChute will wait the **vSAN Synchronization Duration** and re-check until data re-synchronization is no longer active, or the retry limit has been reached.



NOTE: Putting a vSAN host into maintenance mode can trigger a data re-synchronization on the host. In this event, PowerChute will wait until the data re-synchronization is complete (with retry limit) before placing the host into maintenance mode and shutting it down.

vCenter Server Virtual Machine Shutdown

 **NOTE:** In a [Dell VxRail configuration](#), vCenter Server VM Shutdown is only relevant for a stretched cluster configuration.

PowerChute will wait the time specified in the **vCenter Server VM Shutdown Duration** field to gracefully shut down the vCenter Server VM before proceeding with the next step in the shutdown sequence.

NOTES:

- The duration set here will not be applied if the vCenter Server VM is added to a priority group when [VM Prioritization](#) is enabled.
- If VM Prioritization is not enabled, the vCenter Server VM will be the last VM to be shut down before the PowerChute VM.
- If VM Prioritization is enabled and the vCenter Server VM is added to a priority group, it will be shut down with the other VMs in its priority group.
- The vCenter Server VM cannot be added to a priority group in a [Dell VxRail configuration](#).



PowerChute
— NETWORK SHUTDOWN

Home

Configuration

Virtualization

UPS Configuration

Help


Virtualization Settings

Configure Virtualization settings for shutdown sequence.

Virtual Machine Migration	?
Virtual Machine and Virtual Appliance, Shutdown and Startup	?
vSAN Settings	?
Host Maintenance Mode	
Delay Host Maintenance Mode	<input checked="" type="checkbox"/>
Timeout	<input type="text" value="45"/> seconds
Skip Host Maintenance Mode	<input type="checkbox"/>
vCenter Server VM Shutdown	
vSphere Cluster Services (vCLS)	?
High Availability	?
Apply Virtualization Settings to all UPS Setups	<input checked="" type="checkbox"/>
All Hosts online prior to startup	<input checked="" type="checkbox"/>
<input type="button" value="Apply"/> <input type="button" value="Undo"/>	


vSphere Cluster Services (vCLS)

 **NOTE:** This section is not applicable in a [Dell VxRail configuration](#) as the VxRail Manager handles vCLS.

 If PowerChute is configured with the vCenter Server user account, it cannot enable vSphere Cluster Services (vCLS) if it is not assigned the Administrator role in vCenter.

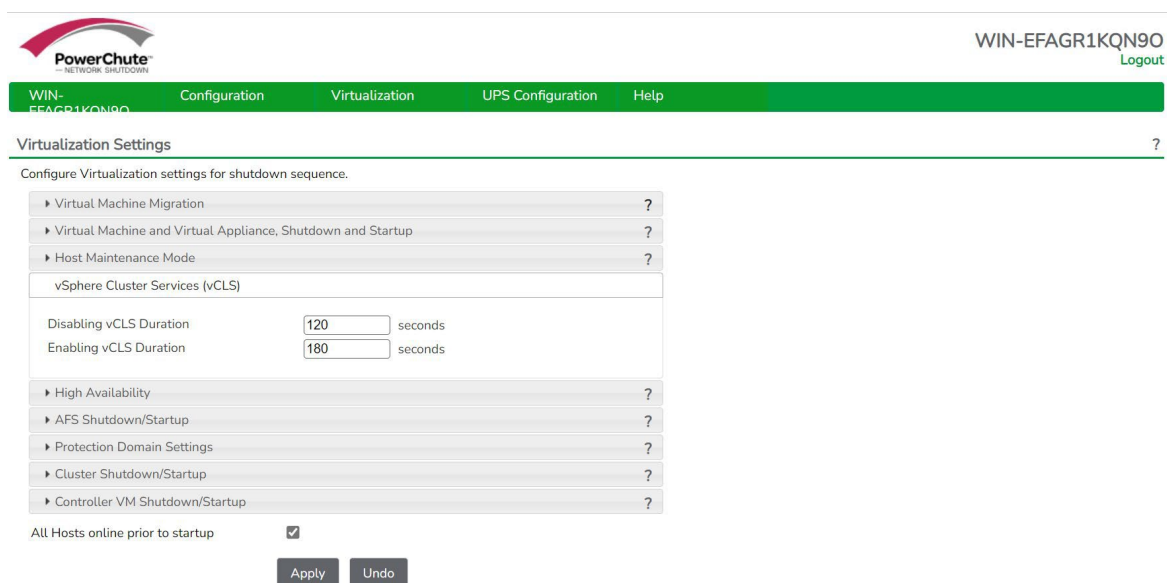
vSphere Cluster Services (vCLS) is a new feature in vSphere 7.0 Update 1 that ensures cluster services, such as vSphere DRS and vSphere High Availability (HA), are available to maintain the resources and health of the workloads running in the clusters independent of the vCenter Server instance availability. PowerChute v5.1 introduces support for disabling vCLS during the shutdown sequence. This is so that the vCLS VMs do not remain powered on and prevent hosts from entering maintenance mode prior to shut down.

If vCLS VMs are detected running on the cluster, PowerChute will automatically disable vCLS, which powers off and deletes the vCLS VMs.

 PowerChute detects vCLS VMs in the cluster beginning with "vCLS". It is recommended you do not name any user VMs in your cluster to "vCLS" to avoid confusion.

You can specify durations for how long PowerChute will wait for vCenter to enable and disable vCLS on your cluster. The **Disable vSphere Cluster Services (vCLS) Duration** value should account for the time it takes for vCenter to power off all the vCLS VMs and delete them. The default value is 180 seconds.

The **Enable vSphere Cluster Services (vCLS) Duration** value should account for the time it takes for new vCLS VMs to be deployed and started up. The default value is 180 seconds.



PowerChute
Network Shutdown

WIN-EFAGR1KQN90
Logout

WIN-EFAGR1KQN90 Configuration Virtualization UPS Configuration Help

Virtualization Settings ?

Configure Virtualization settings for shutdown sequence.

- Virtual Machine Migration ?
- Virtual Machine and Virtual Appliance, Shutdown and Startup ?
- Host Maintenance Mode ?
- vSphere Cluster Services (vCLS)
 - Disabling vCLS Duration: 120 seconds
 - Enabling vCLS Duration: 180 seconds
- High Availability ?
- AFS Shutdown/Startup ?
- Protection Domain Settings ?
- Cluster Shutdown/Startup ?
- Controller VM Shutdown/Startup ?

All Hosts online prior to startup

Apply Undo

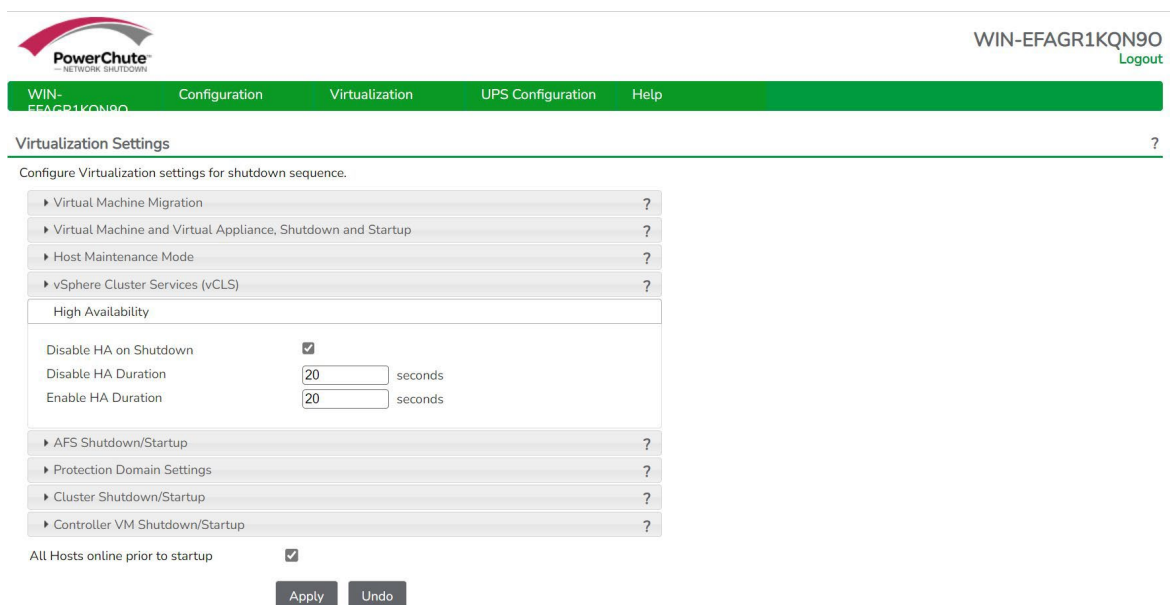
High Availability



NOTE: Disabling High Availability is not applicable in a [Dell VxRail configuration](#).

vCenter High Availability (HA) protects vCenter Server against host and hardware failures and helps you significantly reduce downtime. VMware recommends disabling HA when the entire cluster is being shut down. You can specify this using the **Disable HA on Shutdown** checkbox.

The **Disable HA Duration** field specifies the duration PowerChute will wait, in seconds, to disable HA on vCenter. The **Enable HA Duration** field specifies the duration PowerChute will wait to enable HA during the start up sequence. The default value for both of these fields is 20 seconds.



The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes the PowerChute logo, the user ID 'WIN-EFAGR1KQN90', and a 'Logout' link. Below the navigation bar, the 'Virtualization' tab is selected. The main content area is titled 'Virtualization Settings' and contains a section for 'High Availability' with the following settings:

- Disable HA on Shutdown:**
- Disable HA Duration:** 20 seconds
- Enable HA Duration:** 20 seconds

Other settings visible include 'All Hosts online prior to startup' (checked), and 'Apply' and 'Undo' buttons at the bottom.

Virtual Machine Prioritization

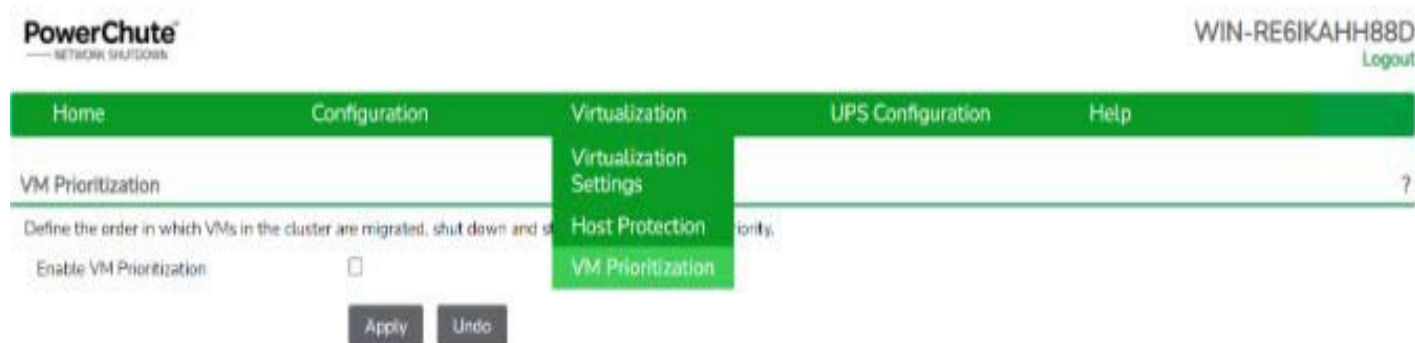
Use Virtual Machine Prioritization settings to specify the order in which VMs migrate, shut down and power on.

VM Prioritization is only available for hosts managed by vCenter Server.



NOTE: If vCenter Server is unavailable at the time of the shutdown, PowerChute will still use the order determined by the priority group to shut down VMs when it connects to each ESXi host. vCenter Server must be available for [Dell VxRail cluster shut down](#).

VM Prioritization is configured in the main PowerChute interface and is disabled by default.



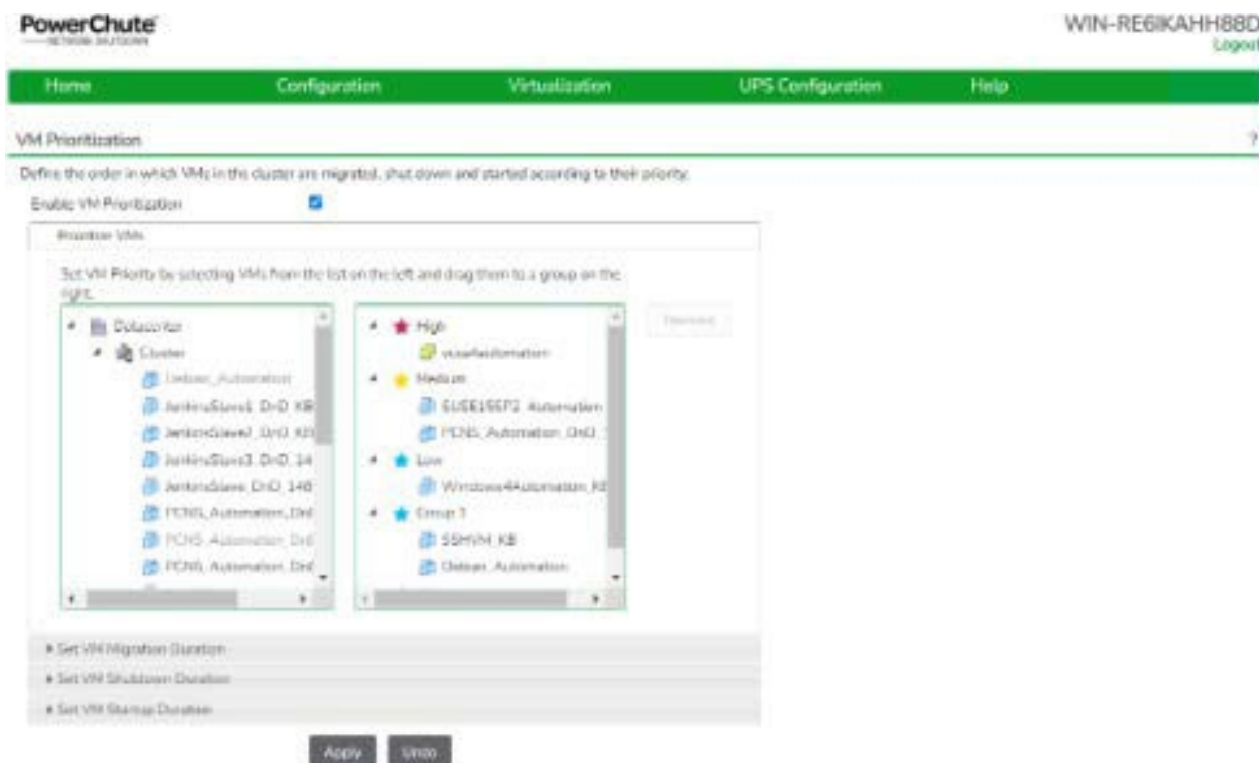
Enable VM Prioritization screen

To enable VM Prioritization, select the **Enable VM Prioritization** checkbox. Four options display:

- **Prioritize VMs**
- **Set VM Migration Duration**
- **Set VM Shutdown Duration**
- **Set VM Startup Duration**

Prioritize VMs


Virtual Machines can be grouped into five priority groups – High, Medium, Low, Group 1 and Group 2. When VM Prioritization is enabled, an **inventory view** of the datacenter, clusters, vCenter Server Appliance (VCSA), VMs and vApps appears on the left. On the right, the High, Medium, Low, Group 1 and Group 2 priority groups are listed.

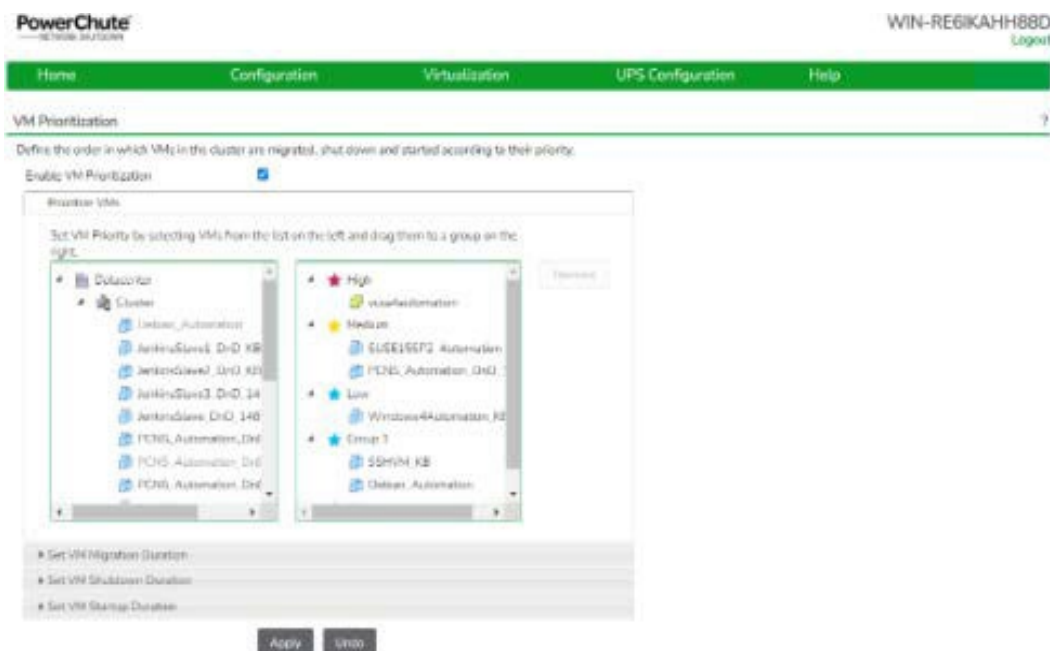


Datacenter inventory view and VM Priority Groups

You can assign a VM/vApp to a priority group by clicking on a VM/vApp on the left hand side and dragging it to a priority group on the right.

Move VMs/vApps between priority groups by dragging them from one group to another. To remove a VM/vApp from a priority group, select the VM/vApp and click the **Remove** button. Any VM/vApp in the inventory that is not assigned to a priority group is considered to be **Un-prioritized**.

 The vCenter Server VM can be added to a priority group if there are special VMs that need to be shut down after vCenter Server and started before it. For example, Platform Services Controller VM, Active Directory Controllers.



The inventory view on the left is populated with VMs that are part of the same HA cluster as any Host protected by a UPS specified on the Host Protection page. VMs on Standalone Hosts managed by vCenter are also present.

NOTES:

- If a HA Cluster contains a protected host, all of its VMs are displayed.
- If a HA Cluster does not contain a protected host, its VMs are not displayed.
- The vCenter Server VM should be added to a priority group on its own, and have no other VMs in its priority group.
- It is not supported to add vApps to a higher priority than the vCenter Server VM as the vCenter Server Appliance needs to be online to interact with the vApps.
- PowerChute does not track changes that are made in the vCenter Server inventory to VMs that have been added to a priority group. If a VM is updated in the vCenter Server Inventory (e.g. renamed, moved), you must manually update the priority group to reflect this change. Prioritized VMs that are not found when a critical event occurs will be ignored when VM operations such as VM Migration, Startup and Shutdown are performed.
- vCLS VMs (VMware v7.0.1 and above) should not be displayed in the inventory view.
- In a [Dell VxRail configuration](#), the vCenter Server VM will not be displayed in the inventory view in a stretched cluster configuration.
- In a [Dell VxRail configuration](#), VxRail Manager and ESRS VMs should not be displayed in the inventory view.
- It is not recommended to rename ESRS VMs in a [Dell VxRail configuration](#). For example, "ESRS_VE".

For PowerChute to correctly identify clusters, VMs and vApps in a datacenter, their names must be unique. It is recommended that a separate instance of PowerChute is used for each datacenter.

If one instance of PowerChute is used to protect hosts in multiple datacenters, it will not be able to identify VMs, vApps or clusters that have the same name, in different datacenters.



In each datacenter in which PowerChute is used to protect hosts, make sure that VMs, vApps and Clusters each have a unique name.

In a [Dell VxRail configuration](#), it is recommended that a separate PowerChute instance is deployed for each VxRail cluster.

Setting Priority Group Durations

You can configure durations for each priority group, for VM Migration, VM Shutdown and VM Startup.

For Priority Groups:

- **VM Migration Duration** sets the time allowed for all VMs in the Priority Group to migrate to another healthy host in the cluster. This also represents the delay between migrating VMs for each Priority Group.
- **VM Shutdown Duration** sets the time allowed for all VMs in the Priority Group to shut down gracefully. This also represents the delay between shutting down VMs for each Priority Group.

The **VM Shutdown Duration** can be set to 0 seconds for any priority group to handle any special VMs that need to be shut down later in the shutdown sequence via a command file or SSH action. This is relevant for VMs running on the Management Host in a [Dell VxRail stretched cluster configuration](#).



If you set a 0 second shutdown duration for a priority group, the shutdown action will be skipped for all VMs in that priority group.

- **VM Startup Duration** sets the time allowed for all VMs in the Priority Group to start up. This also represents the delay between starting VMs for each Priority Group.

PowerChute NETWORK SHUTDOWN WIN-REGIKAHH88D Logout

Home Configuration Virtualization UPS Configuration Help

VM Prioritization ?

Define the order in which VMs in the cluster are migrated, shut down and started according to their priority.

Enable VM Prioritization

Priority VMs

Set VM Migration Duration

Set VM Shutdown Duration

Values entered here will override the Duration values in the Virtualization Settings page.

High	30	seconds
Medium	30	seconds
Low	30	seconds
Group 1	30	seconds
Group 2	0	seconds
Un-prioritized	120	seconds

Set VM Startup Duration

Apply Undo

Set VM Operation durations - VM Migration, Shutdown and Startup

You can set durations for the VMs/vApps in the High, Medium, Low, Group 1 and Group 2 priority Groups, and for VMs/vApps that are Un-prioritized. When VM Prioritization is enabled for the first time, High, Medium, Low, Group 1 and Group 2 durations will have a default value of 0.

Durations for un-prioritized VMs/vApps are automatically set to the global duration values for VM Migration, VM Shutdown and VM Startup, as configured on the [Virtualization Settings](#) page.



If vApps are added to a priority group, they will not be shut down if vCenter Server is offline and no duration is specified for Un-prioritized VMs/vApps. To ensure your vApps gracefully shut down, provide a shutdown duration for the Un-prioritized priority group.

VM Migration and VM Startup durations can be configured to have a value of zero, to skip the migration or startup of VMs/vApp when a critical event is resolved. If VM Migration/VM Startup is set to zero for a priority group, VMs/vApps that do not migrate will be shut down, and will not start up following the shutdown, when the critical event is resolved.

A warning will display in the **VM Prioritization** screen if a VM Shutdown duration is set to zero for any priority group.

Global Virtualization Settings and VM Prioritization

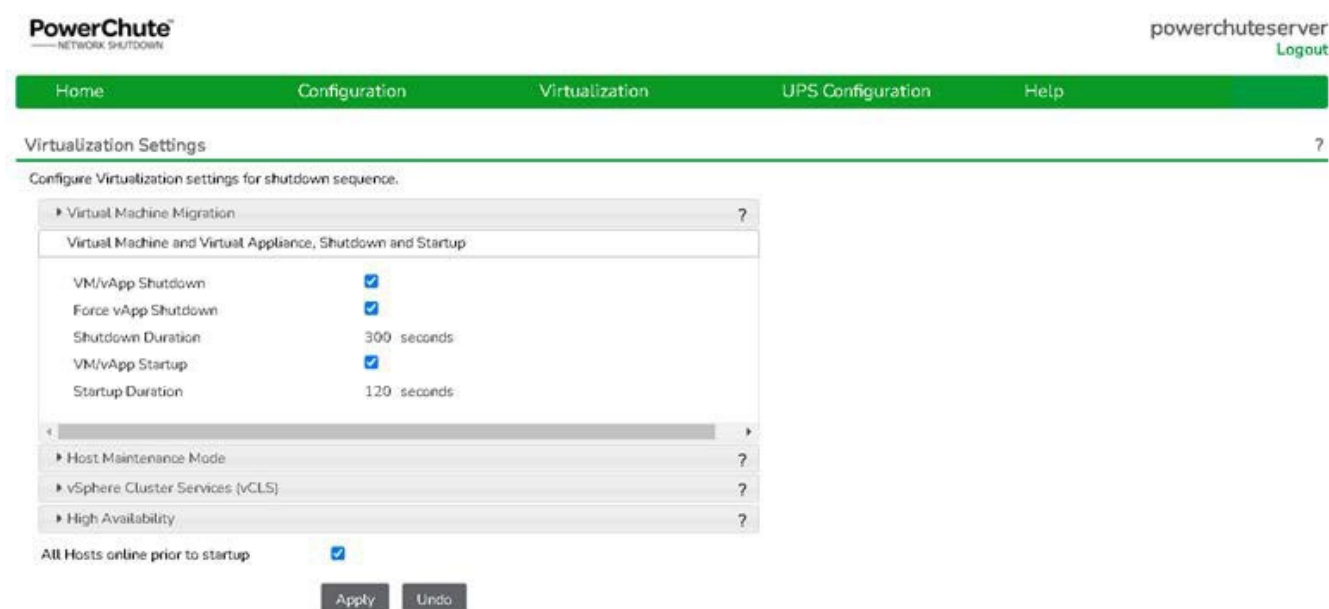
When VM Prioritization is enabled, the VM Migration, VM Startup and VM Shutdown durations set on the Virtualization Settings page automatically match the sum of the respective High/Medium/Low//Group 1/Group 2/Un-prioritized durations set on the VM Prioritization page.

For example, if VM Shutdown durations set on the VM Prioritization screen are as follows:

- High: 90 seconds
- Medium: 60 seconds
- Low: 60 seconds
- Group 1: 30 seconds
- Group 2: 30 seconds
- Un-prioritized: 30 seconds

then the VM Shutdown duration on the Virtualization Settings page is automatically set to:

$90 + 60 + 60 + 30 + 30 + 30 = \mathbf{300 \text{ seconds}}$



Virtualization Setting page - VM Shutdown Duration

If Advanced UPS configuration is used, the durations for each UPS Setup are also changed to reflect the sum of the respective priority group durations for VM Migration, Startup and Shutdown.

On the Virtualization Settings page when VM Prioritization is enabled, VM Migration, Startup and Shutdown durations become read-only and are automatically populated, as described above. If VM Prioritization is disabled, the durations retain the values set by VM Prioritization, but become editable once more.

Prioritized VM Operations Sequence

Operations such as VM Migration, VM Shutdown and VM Startup can be performed for VMs in a priority group.

For VM Migration and VM Startup, the order in which priority groups are processed is as follows:



For VM Shutdown the order in which priority groups are processed is:




The sequences below describe the VM Migration, Shutdown and Startup Sequence for all UPS Configurations.


In this scenario:

- VM Migration is enabled
- DRS is disabled
- Wait for all Hosts Online is enabled
- VM Shutdown is enabled
- VM Startup is enabled

Priority Group VM Migration




A UPS Critical event occurs and PowerChute starts a maintenance mode task on each protected host. First the High priority group VMs migrate, in parallel. When all High priority Group VMs have migrated, the Medium priority VMs migrate in parallel, followed by the Low, Group 1 and Group 2 priority groups, and finally the Un-prioritized VMs migrate.

 PowerChute will proceed to migrate the next group of VMs/vApps when the duration for the priority group has elapsed.

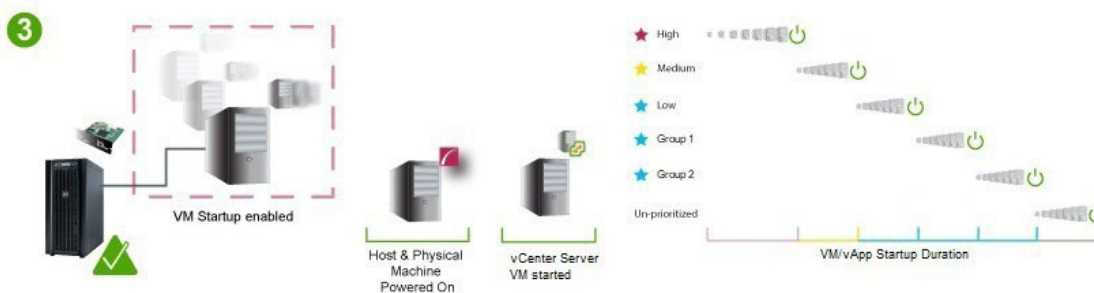
When the VM migration duration elapses, any VMs that have not been migrated will be shut down in the VM Shutdown sequence.

Priority Group VM Shutdown



PowerChute begins sequenced VM/vApp Shutdown. First, the un-prioritized VMs/vApps are shut down sequentially. After the duration for un-prioritized VMs/vApps elapses, the Group 2 priority VMs/vApps are shut down, followed by the Group 1 priority VMs/vApps, the Low priority VMs/vApps, the Medium priority VMs/vApps and finally the High priority VMs/vApps are shut down. When all priority group VMs/vApps have been shut down, PowerChute shuts down the vCenter Server VM if vCenter Server is running on a VM and the vCenter Server VM is not part of a priority group. PowerChute then shuts down the hosts and the physical PowerChute server. In a [Dell VxRail configuration](#), PowerChute sends a request to the NMC to shut down the VxRail cluster after a configurable delay. PowerChute then shuts itself down.


Priority Group VM Startup



When the UPS critical event is resolved and the hosts and physical machine running PowerChute are powered back on, the hosts are taken out of maintenance mode and the vCenter Server VM is started if the vCenter Server VM has not been added to a priority group. In a [Dell VxRail configuration](#), VxRail handles taking the hosts out of maintenance mode and re-starting vCenter Server VM and VxRail Manager VMs. PowerChute begins to start the High priority VMs/vApps with respect to their startup duration. When the High priority startup duration elapses, the Medium priority VMs/vApps are started, followed by the Low priority VMs/vApps, the Group 1 priority VMs/vApps, the Group 2 priority VMs/vApps, and finally, the un-prioritized VMs/vApps are started.

Priority Group VM Operations in Advanced UPS Configuration

In Advanced UPS configuration, ESXi hosts and equipment may be protected by separate UPS devices.



In this example, UPS 2 may experience a critical event at a different time to UPS 3. If the UPS Setups for Host A and B have VM Prioritization enabled, the VM operations that are performed on the priority groups of the two hosts will occur at different times. For example, if UPS 2 experiences a critical event 2 minutes before UPS 3, the High priority VMs on Host B may be shut down at the same time as the Low priority VMs on Host A. VM operations on priority groups occur at the host level.

This also applies in a multi-site scenario where hosts in two geographical locations, Site A and Site B, are powered by two different UPS devices. As the critical event may not occur at the same time at each location, if Site A experiences a critical event first, it is possible to have High priority VMs shutting down in site B at the same time as Low priority VMs shut down at site A.

Active Directory VMware Configuration

In the event that vCenter Server is unavailable, it is recommended to configure an Active Directory account that can be used to connect directly to the VMware hosts to perform shutdown actions. vCenter Server must be available for Dell VxRail cluster shutdown.

This topic outlines how to configure an Active Directory account that can be used to connect directly to the VMware hosts to perform shutdown actions. This topic is only relevant in a stretched cluster configuration.

1. In Active Directory Users and Groups create a group called **ESX Admins** and add your user(s) to the group.



When using Active Directory VMware provides a default AD Group account called "ESX Admins". This group is automatically added to each ESXi host joined to the domain and is granted administrator rights by default.

2. Add Active Directory as an Identity Source in VMware Single Sign On using the vSphere Web Client.
3. Log in to vCenter Server using the vSphere Web Client via a browser - **https://<your_vcenter_ip>** - using default vCenter Server administrator account – **administrator@vsphere.local**.
4. Navigate to **Administration > Single Sign On > Configuration** and then on the Identity Sources tab.

Name	Server URL	Type	Domain	Alias
--	--	System Domain	vsphere.local	--
--	--	Local OS (Default)	localos	--
scvmmblr.net	ldap://10.179.232.211:389	ActiveDirectory	scvmmblr.net	--

5. Click on the **+** symbol to add a new identity source.
6. Select **Active Directory as a LDAP Server**.
7. Enter the domain details; e.g. - **testdomain.com**
 - a. **Name:** testdomain
 - b. **Base DN for Users:** CN=Users, DC=testdomain, DC=com
 - c. **Domain Name:** testdomain.com
 - d. **Alias:** testdomain
 - e. **Base DN for Groups:** CN=Users, DC=testdomain, DC=com

f. **Primary Server URL:** domaincontroller.testdomain.com

g. **Username:** testdomain\domainuser

Add Identity Source ✕

Identity source type: Active Directory over LDAP

Name: scvmmlr.net

Base distinguished name for users: dc=scvmmlr,dc=net

Base distinguished name for groups: dc=scvmmlr,dc=net

Domain name: scvmmlr.net

Domain alias:

Username: administrator@scvmmlr.net

Password:

Connect To:
 Any domain controller in the domain
 Specific domain controllers

Primary Server URL:

Secondary Server URL:

SSL Certificates: BROWSE

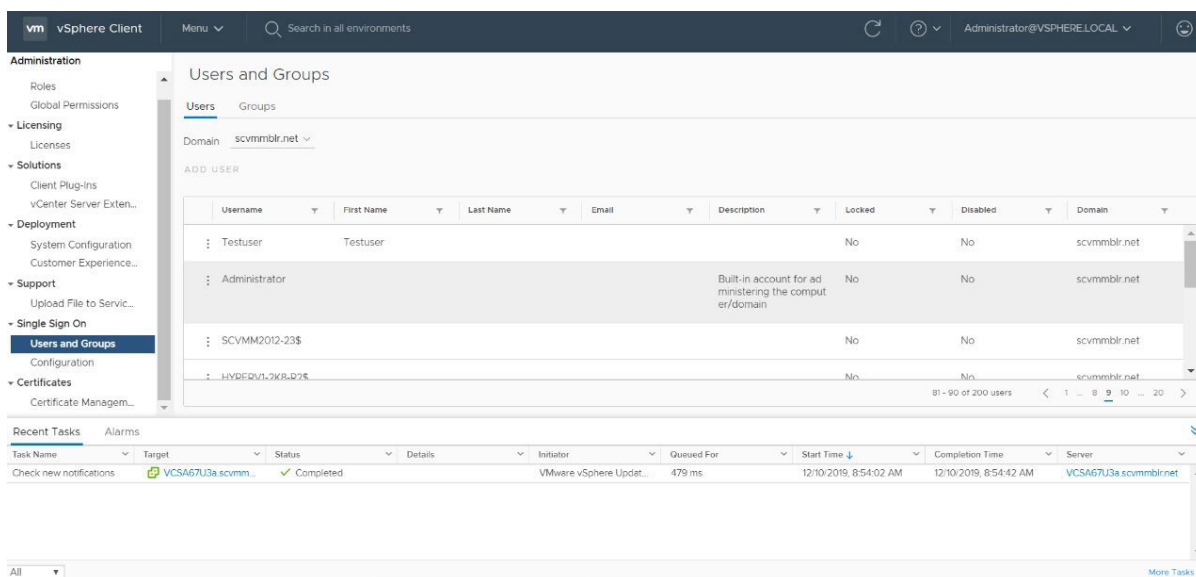
CANCEL
ADD

8. Click OK.

9. Log into vCenter using the vSphere client and navigate to **Menu > Administration > Single Sign On > Users and Groups**.

10. Click on the + symbol and create a user.

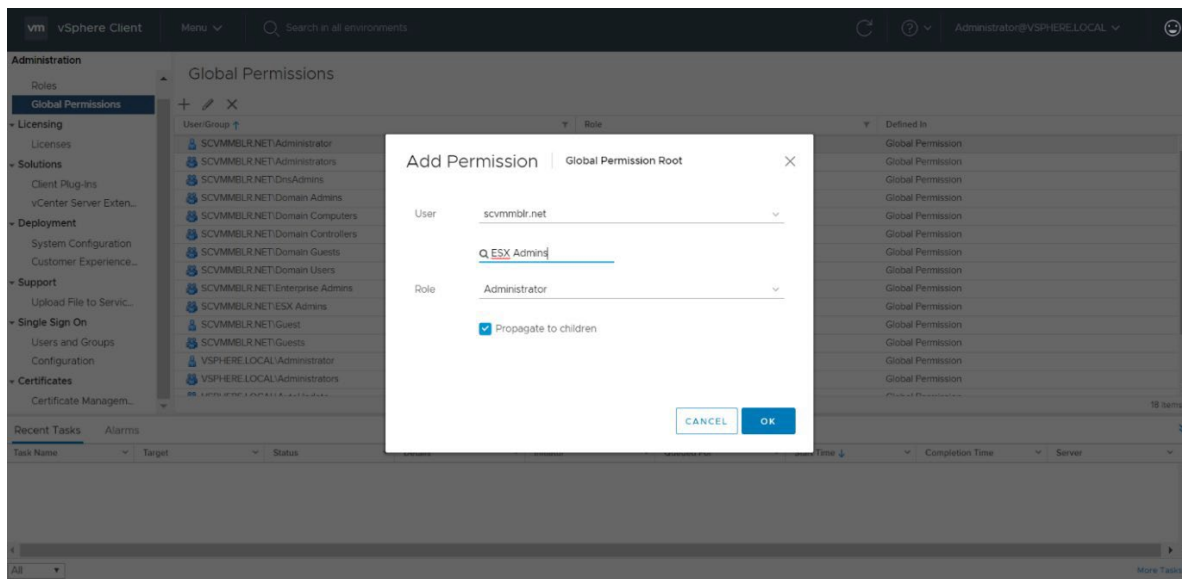
11. Click OK.



12. Navigate to **Menu > Administration > Global Permissions**.



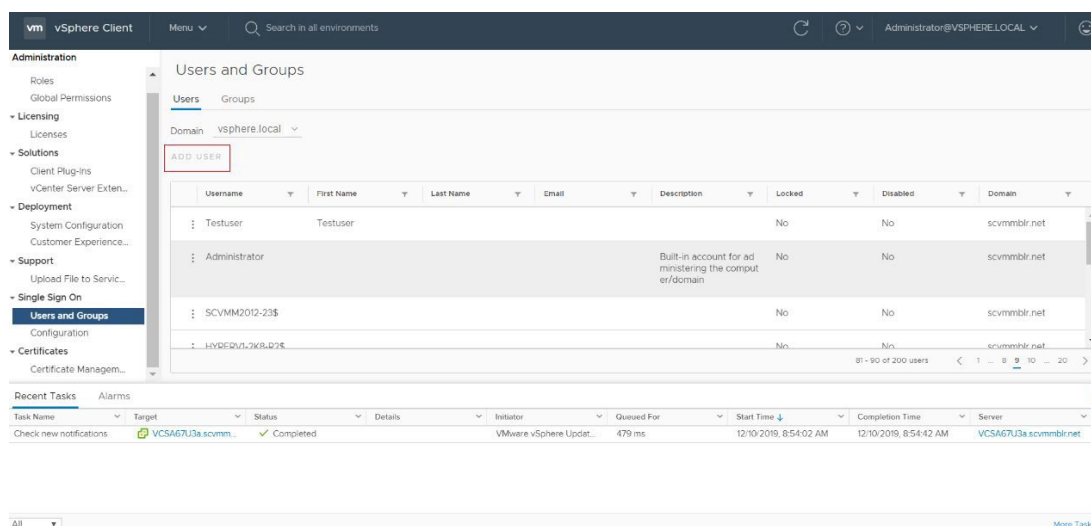
13. Click on the symbol. Select User and change the Assigned Role to Administrator.
14. Confirm "Propagate to children" is selected.
15. Click OK.
16. Add all required roles for the User.



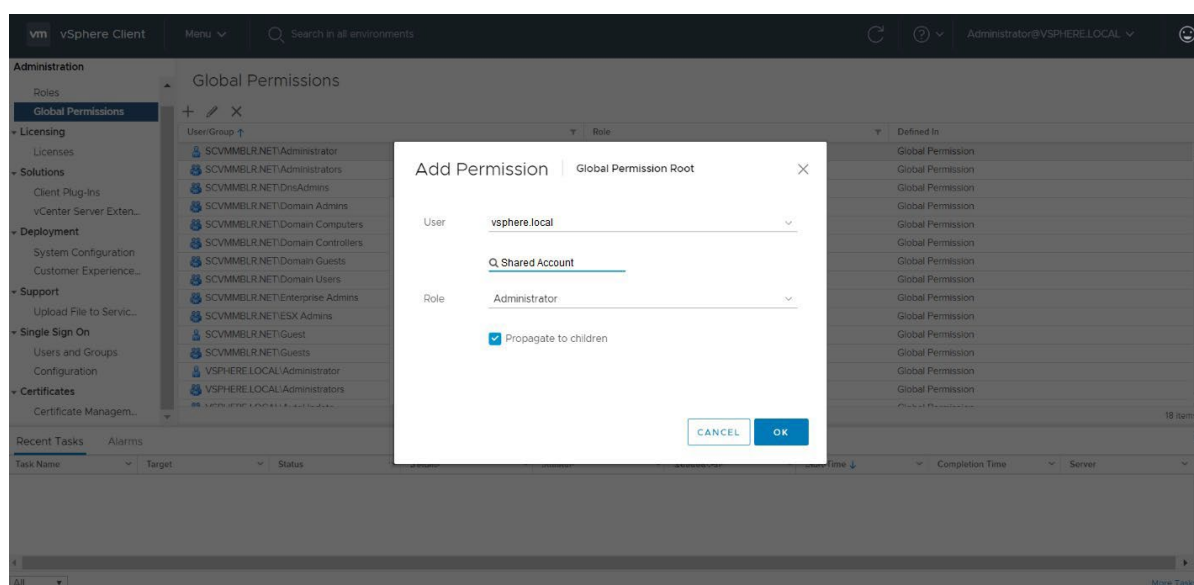
Shared Local Account for vCenter Server and VMware hosts

Create Shared Local Account on vCenter Server - not using local OS

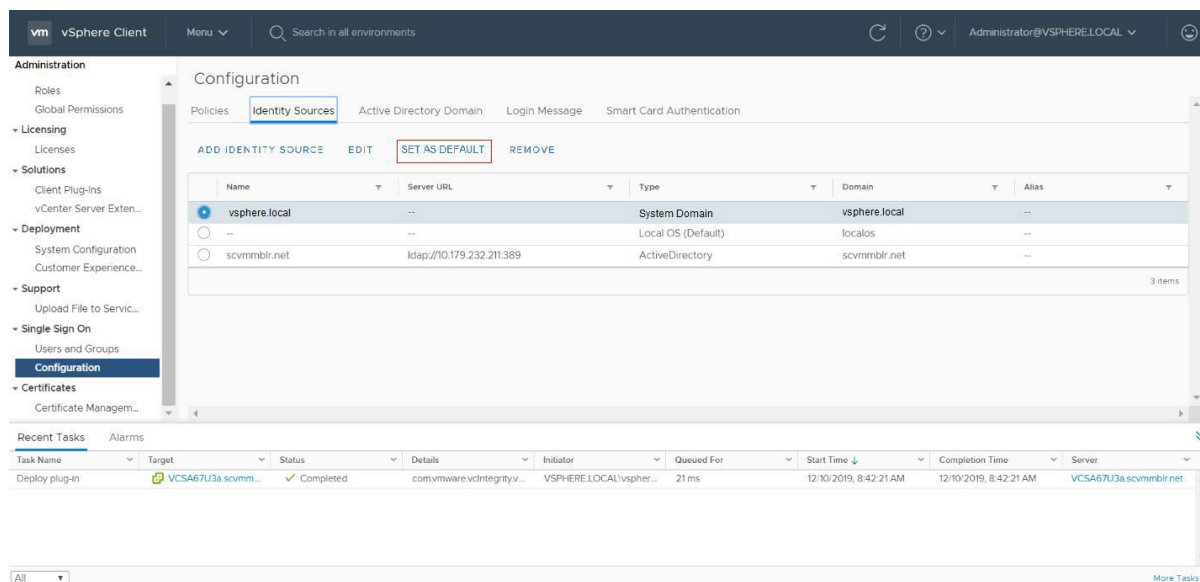
1. In the event that vCenter Server is unavailable a shared account needs to be configured that can be used to connect directly to the VMware hosts to perform shutdown actions. vCenter Server must be available for Dell VxRail cluster shutdown. A shared account needs to be configured that can be used to connect directly to the VMware hosts to perform shutdown actions.
2. Log in to vSphere Web Client as **administrator@vsphere.local**.
3. Navigate to **Menu > Administration > Single Sign On > Users and Groups**.
4. Select the vsphere.local domain and click **Add User**.



5. Enter your vCenter username and password, and click **OK**.
6. Navigate to **Menu > Administration > Global Permissions**.
7. Click on the **+** symbol. Select the User created in Step 4 and change the Assigned Role to Administrator.
8. Confirm "Propagate to children" is selected.
9. Click **OK**.



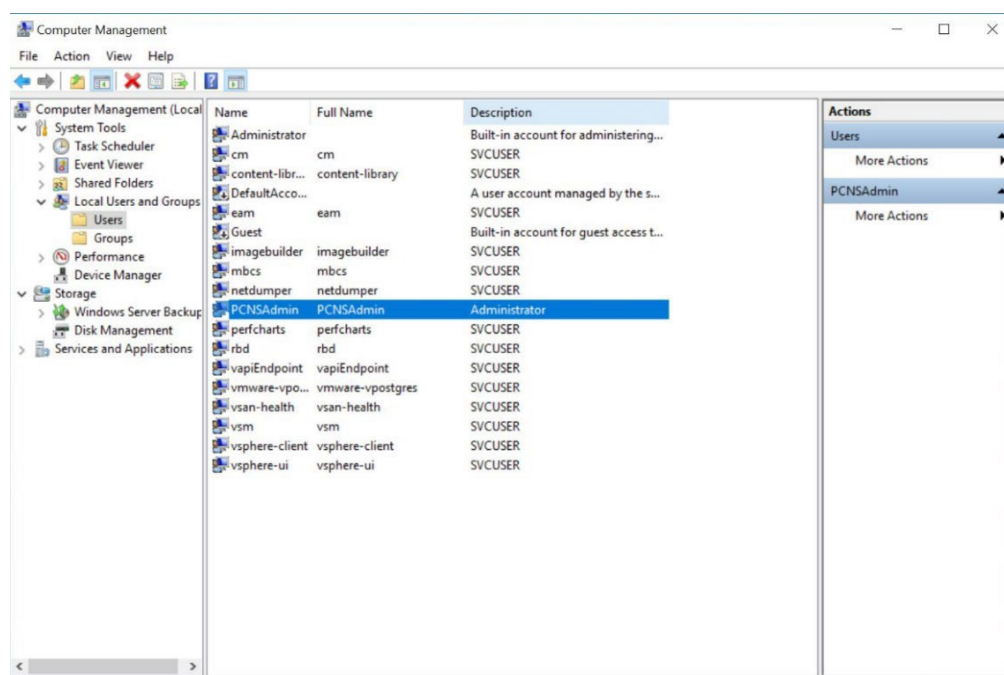
10. Navigate to **Single Sign On > Configuration > Identity Sources**. Select the vsphere.local domain where the new User was added and set this as the default domain.



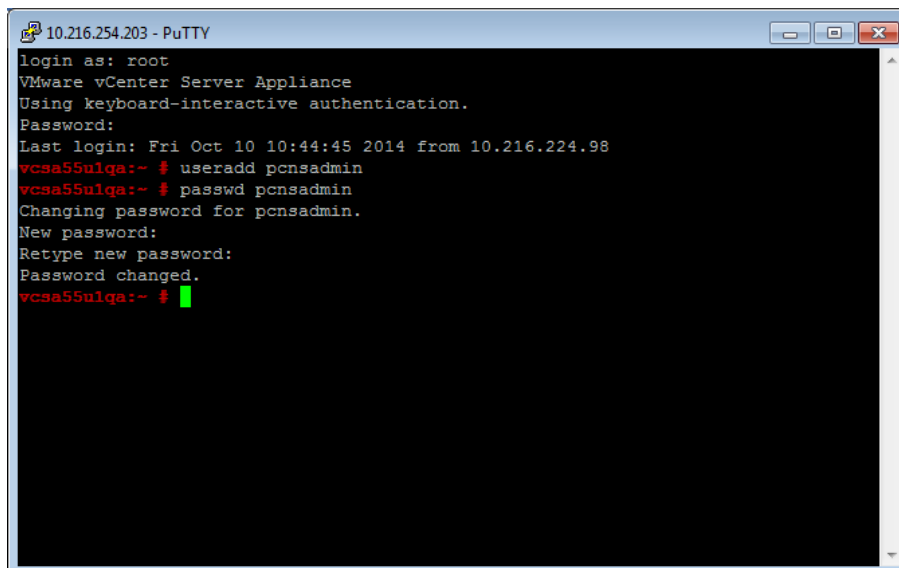
11. In PowerChute, when adding the username, do not include the @vsphere.local suffix. This is not required as ESXi host usernames cannot include @vsphere.local.

Create Shared Local Account on vCenter Server - using local OS

1. In the event that vCenter Server is unavailable a shared account needs to be configured that can be used to connect directly to the VMware hosts to perform shutdown actions. vCenter Server must be available for Dell VxRail cluster shutdown. A shared account needs to be configured that can be used to connect directly to the VMware hosts to perform shutdown actions.
2. If Active Directory is not available then a local user account can be added to vCenter Server.
3. An account with the same name and password then needs to be added to each ESXi host.
4. Log in to vCenter Server machine and add a user via Computer Management -> Local Users and Groups for Windows.

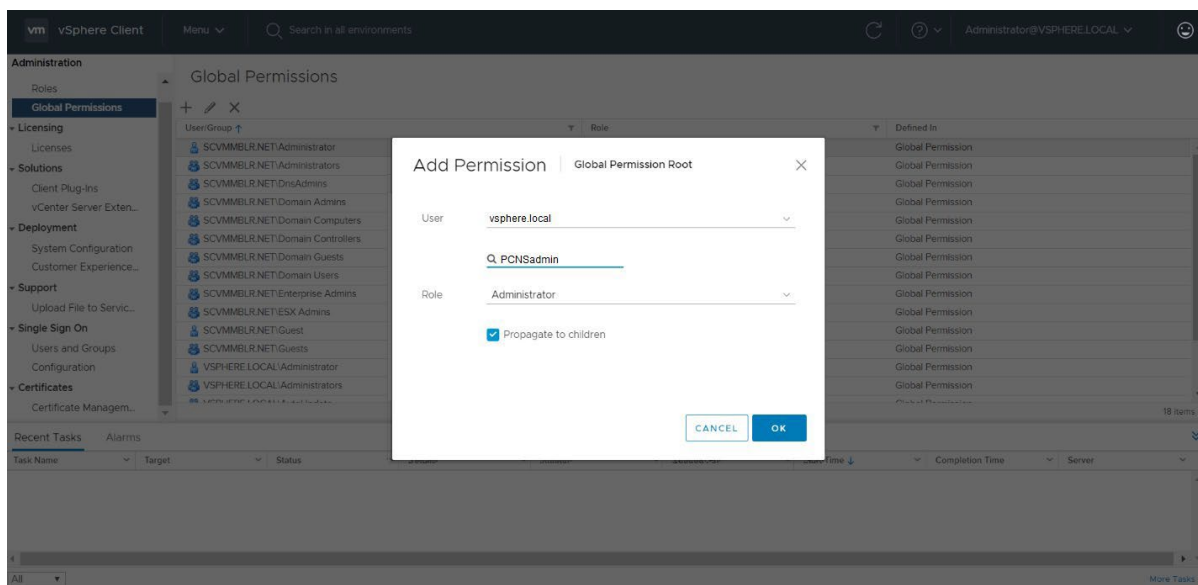


On Linux/vCenter Server Appliance use the terminal commands “useradd” and “passwd”.



```
10.216.254.203 - PuTTY
login as: root
VMware vCenter Server Appliance
Using keyboard-interactive authentication.
Password:
Last login: Fri Oct 10 10:44:45 2014 from 10.216.224.98
vcsa55ulqa:~ # useradd pcnsadmin
vcsa55ulqa:~ # passwd pcnsadmin
Changing password for pcnsadmin.
New password:
Retype new password:
Password changed.
vcsa55ulqa:~ # █
```

5. Log in to vCenter Server using the vSphere Client and navigate to **Menu > Administration > Global Permissions**.
6. Click on the **+** symbol. Select vsphere.local under User, select the User that was added in step 4 and change the Assigned Role to Administrator.
7. Confirm "Propagate to children" is selected.
8. Click OK.



Shared Local Account for vCenter Server and VMware hosts

Create Shared Local Account on each VMware host

1. Log in to the ESXi Client and navigate to **Manage > Security and Users > Users**.
2. Click **Add User**.
3. Enter your username and password, and click **OK**.
4. Right-click the Host in the inventory and click **Permissions**.
5. Click **Add User**.
6. Select the User created in step 2 and change the Assigned Role to Administrator.
7. Click **OK**.



A shared local account should be used when vCenter Server is running on a VM and Active Directory is unavailable.

Hyperconverged Infrastructure Support for VMware

If you have a supported hyperconverged infrastructure (HCI) solution, you can enable HCI support for VMware in the PowerChute Setup wizard. Select a HCI solution from the **Hyperconverged Infrastructure Support** drop-down list in the **vCenter Server Details** screen and click **Next**.

PowerChute Setup: vCenter Server Details

vCenter Server Protocol: https

vCenter Server Port: 443

vCenter Server IP Address/Hostname: 10.1.1.1

vCenter Server Username: administrator@vsphere.local

vCenter Server Password:

vCenter Server running on a VM:

Hyperconverged Infrastructure Support: None (selected)

Options in dropdown: None, HyperFlex, Nutanix, SimpliVity, VxRail

Buttons: Previous, Next, Cancel

For more information see the PowerChute [Installation Guide](#) on the APC website, and the below topics:

- [Nutanix Support](#)
- [HPE SimpliVity Support](#)
- [Cisco HyperFlex Support](#)
- [Dell VxRail Support](#)

Nutanix Support

If you have a Nutanix™ Cluster with ESXi as your hypervisor, you can enable Nutanix support for PowerChute in the PowerChute Setup wizard.

Enable Nutanix Support

To enable Nutanix support, open the PowerChute Setup wizard and follow the steps below.

The screenshot shows the 'PowerChute Setup: vCenter Server Details' window. The 'Hyperconverged Infrastructure Support' dropdown menu is open, with 'Nutanix' selected. Other fields are filled with: vCenter Server Protocol: https; vCenter Server Port: 443; vCenter Server IP Address/Hostname: 10.1.1.1; vCenter Server Username: administrator@vsphere.local; vCenter Server Password: [masked]; vCenter Server running on a VM: [checked]; and Hyperconverged Infrastructure Support: Nutanix. Navigation buttons 'Previous', 'Next', and 'Cancel' are visible at the bottom right.

1. Select **Nutanix** from the **Hyperconverged Infrastructure Support** drop-down list in the **vCenter Server Details** screen and click **Next**.
2. In the **CVM/Cluster Details** screen, enter your Nutanix credentials to connect to your Controller Virtual Machine or Cluster. The recommended configuration is to connect to your Nutanix Cluster.

You can authenticate the connection by entering either the CVM/Cluster IP address and a password, or a SSH key file path. If both authentication options are specified, PowerChute will use the password to connect to the Controller VM/Cluster.



NOTE: You must use the "nutanix" user account credentials to connect to the Cluster/CVM. You cannot use Prism user account credentials to connect.

The SSH key file, if configured, must be located in the user_files folder, or a sub-folder, in the PowerChute installation directory. If the default installation directory was chosen during installation, this location will be:

- C:\Program Files\APC\PowerChute\user_files for Windows systems
- /opt/APC/PowerChute/user_files/ for Linux systems



NOTE: If you are upgrading to PowerChute v4.4, any SSH key files used in a previous version of PowerChute must be manually added to the user_files folder. If the key files are not added to the directory, you will see the below error in the Event Log:

```
ERROR: The ini contains an invalid value for ssh_keyfile_path
in nutanixClusterDetails.
```

It is strongly recommended you validate your configuration after an upgrade.

3. When you have entered your credentials, click **Next** to connect to your Controller VM/Cluster.

4. If the connection to your Controller VM/Cluster is successful, you will be directed to the **Virtualization Settings** screen where you can configure your Nutanix settings. These settings can later be edited in the [Virtualization Settings](#) screen in the PowerChute UI.

Nutanix Virtualization Settings

When Nutanix support is enabled, additional sections will appear in the Virtualization Settings screen in the PowerChute UI.


For more information on these sections, see:


- [Controller VM Shutdown/Startup](#)
- [AFS Shutdown/Startup](#)
- [Protection Domain Settings](#)
- [Cluster Shutdown/Startup](#)

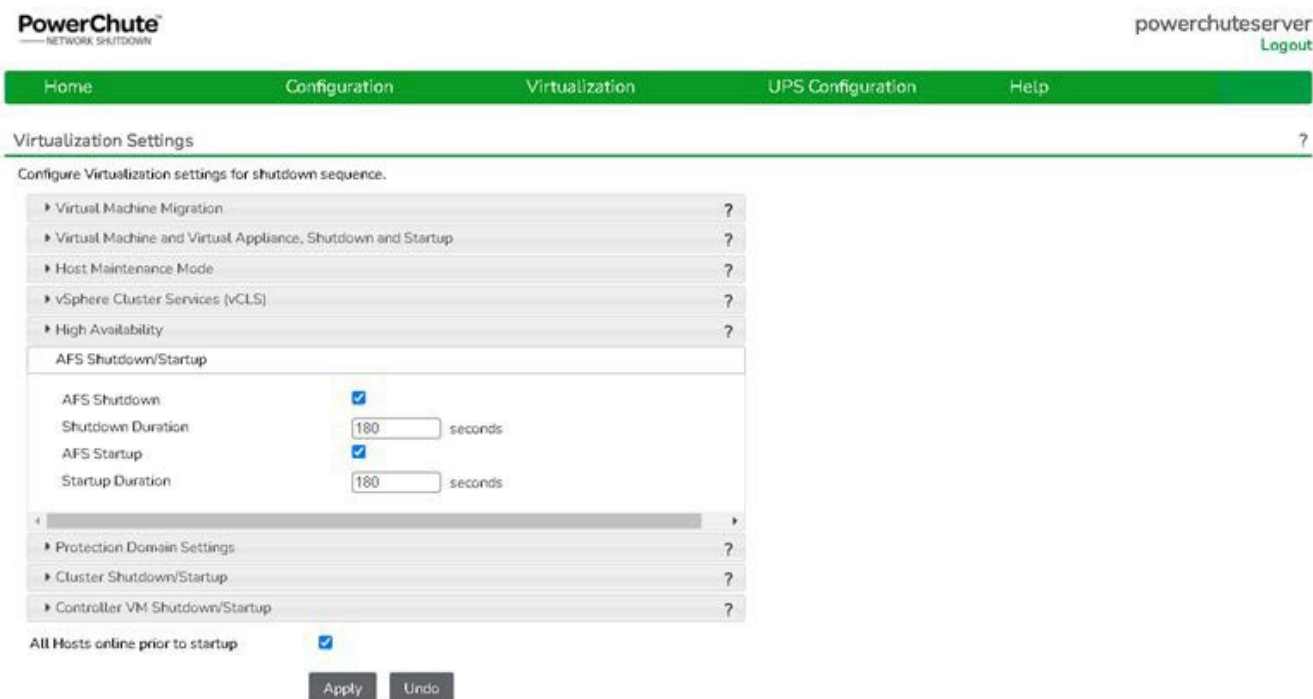
Nutanix Files (Acropolis File Services) Shutdown/Startup

The **AFS Shutdown** checkbox is selected by default, and this option should be left selected to allow graceful Cluster shutdown. In the event of a shutdown, PowerChute stops the AFS service, which shuts down the AFS VMs. This step occurs after the User VMs are gracefully shut down.

The **Shutdown Duration** field is the time allowed for the AFS service and VMs to gracefully shut down and start up following a UPS critical event.

 Using the **Duration** field, you must allow sufficient time for all your AFS VMs to gracefully shut down. If the AFS VMs are not shut down, this may prevent the Cluster and Controller VMs from shutting down.

 If vCenter Server is unavailable during shutdown, an alternative shutdown process is used. In this case, PowerChute Network Shutdown must be configured with passwords for CVMs as SSH keys cannot be used to connect to AFS VMs. AFS VMs must have the same password as CVMs to permit the alternative shutdown process.



The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes Home, Configuration, Virtualization, UPS Configuration, and Help. The user is logged in as 'powerchuteserver'. The 'Virtualization Settings' page is active, displaying a list of configuration options for shutdown sequence. The 'AFS Shutdown/Startup' section is expanded, showing the following settings:

Setting	Value
AFS Shutdown	<input checked="" type="checkbox"/>
Shutdown Duration	180 seconds
AFS Startup	<input checked="" type="checkbox"/>
Startup Duration	180 seconds

Below the settings, there is a checkbox for 'All Hosts online prior to startup' which is also checked. At the bottom of the configuration area, there are 'Apply' and 'Undo' buttons.

AFS Shutdown/Startup

Re-starting after a shutdown


Selecting the **AFS Startup** checkbox re-starts any AFS VMs that were shut down when a UPS critical event has been resolved and the Nutanix Hosts are powered on.

If the AFS VMs are in the process of being started when a critical event occurs, PowerChute waits for the **Startup Duration** time to elapse before shutting down the VMs. This is to ensure that the VMs are shut down gracefully.

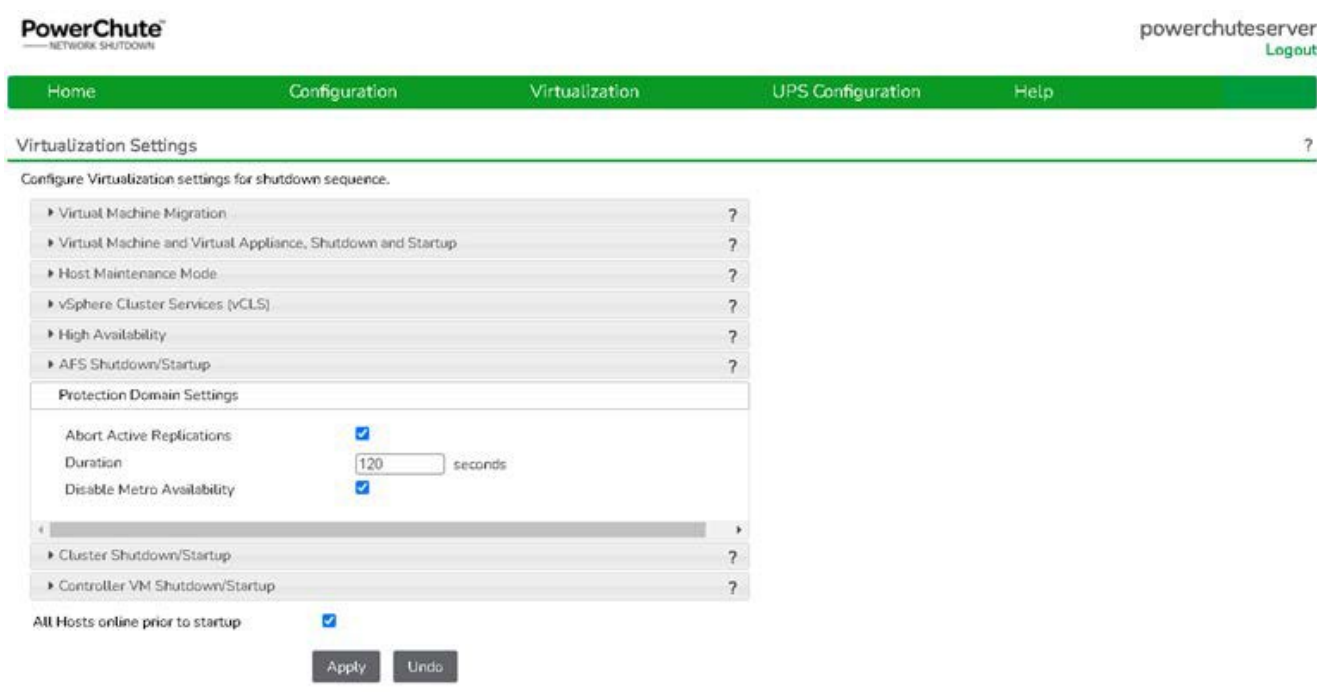
Protection Domain Settings

A protection domain is a collection of VMs that are backed up or replicated on a schedule to recover data. The **Abort Active Replications** checkbox is selected by default, and this option should be selected if your protection domain is configured to replicate on a schedule.

If enabled, PowerChute will wait the time specified in the **Duration** field before aborting any active protection domain replications in the event of a critical UPS event.

 **NOTE:** The **Duration** field will only be taken into consideration if there are active replications when PowerChute reaches this step in the shutdown sequence. If there are no active replications this duration will not be included in the shutdown sequence.

Metro Availability spans a datastore across two sites (local and remote) and synchronously replicates data between the two sites. If the **Disable Metro Availability** checkbox is selected, Metro Availability will be deactivated in the event of a shutdown. This option should be selected if Metro Availability is configured on your Cluster.



The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes Home, Configuration, Virtualization, UPS Configuration, and Help. The current page is titled "Virtualization Settings" and contains a section for "Protection Domain Settings".

Under "Protection Domain Settings", the following options are visible:

- Virtual Machine Migration ?
- Virtual Machine and Virtual Appliance, Shutdown and Startup ?
- Host Maintenance Mode ?
- vSphere Cluster Services (vCLS) ?
- High Availability ?
- AFS Shutdown/Startup ?
- Protection Domain Settings**
 - Abort Active Replications
 - Duration seconds
 - Disable Metro Availability
- Cluster Shutdown/Startup ?
- Controller VM Shutdown/Startup ?

At the bottom of the settings section, there is a checkbox for "All Hosts online prior to startup" which is checked. Below this are "Apply" and "Undo" buttons.

Protection Domain Settings

Cluster Shutdown/Startup

Cluster Shutdown occurs after PowerChute has completed the previous steps in the shutdown sequence: User VM Shutdown, AFS Shutdown, and abort active replications.

The screenshot shows the PowerChute Network Shutdown web interface. The top navigation bar includes Home, Configuration, Virtualization, UPS Configuration, and Help. The 'Virtualization Settings' page is active, displaying a list of settings for the shutdown sequence. The 'Cluster Shutdown/Startup' section is expanded, showing two input fields: 'Shutdown Duration' set to 180 seconds and 'Startup Duration' set to 120 seconds. Below these fields, there is a checkbox for 'All Hosts online prior to startup' which is checked. At the bottom of the section are 'Apply' and 'Undo' buttons.

Cluster Shutdown/Startup

If the Cluster is shut down before the Duration time, PowerChute waits until this time has elapsed before proceeding to the next step in the sequence (shutting down the Controller VMs).



Using the **Shutdown Duration** field, you must allow sufficient time for your Cluster to gracefully shut down before the Controller VMs are commanded to shut down.

Re-starting after a shutdown

The time specified in the **Startup Duration** field will be the time, in seconds, PowerChute will wait after issuing the Cluster start command before proceeding with the next steps in the startup sequence. The Cluster is re-started after the Controller VMs are powered on.

Controller VM Shutdown/Startup

A Controller Virtual Machine runs on each node in a Nutanix cluster and is responsible for managing the Nutanix Cluster services.

Controller VM Shutdown is enabled by default. In the PowerChute shutdown sequence, the Controller VM is shut down after all other VMs in the Cluster, and the Nutanix Cluster itself are shut down.

The **Duration** field is the time allowed for all the Controller VMs to gracefully shut down. If the Controller VMs are shut down before the Duration time, PowerChute waits until this time has elapsed before proceeding to the next step in the sequence.



Using the **Shutdown Duration** field, you must allow sufficient time for all your Controller VMs to gracefully shut down before the hosts are commanded to shut down.

PowerChute NETWORK SHUTDOWN

powerchuteserver Logout

Home Configuration Virtualization UPS Configuration Help

Virtualization Settings ?

Configure Virtualization settings for shutdown sequence.

- Virtual Machine Migration ?
- Virtual Machine and Virtual Appliance, Shutdown and Startup ?
- Host Maintenance Mode ?
- vSphere Cluster Services (vCLS) ?
- High Availability ?
- AFS Shutdown/Startup ?
- Protection Domain Settings ?
- Cluster Shutdown/Startup ?

Controller VM Shutdown/Startup

Shutdown Duration seconds

Controller VM Startup

Startup Duration seconds

All Hosts online prior to startup

Apply Undo

Controller VM Shutdown/Startup

Re-starting after a shutdown

If HA is enabled in a VMware environment and the **Controller VM Startup** checkbox is selected, PowerChute will re-start the Controller VMs when the Nutanix Hosts are powered on. PowerChute first checks that the host is available.

If HA is disabled, the VMware Virtual Machine Shutdown/Startup feature should be used to re-start the Controller VMs when the hosts are powered on.



It is recommended to wait at least 5 minutes after Controller VM startup to ensure all services are running before re-starting the Cluster. You should account for this using the **Startup Duration** field.

CVM/Cluster Details

This screen allows you to edit the connection details for your Nutanix Controller VM or Cluster.

You can authenticate the connection to your Controller VM or Cluster using one of the following options:

- Enter your CVM/Cluster IP address and CVM/Cluster password.
- Enter your CVM/Cluster IP address, the SSH key file path and its passphrase, if available.

If you wish to change the authentication method used, PowerChute will erase your previous authentication details. For example, if you connected to your Nutanix Cluster using the password option in the PowerChute Setup wizard and you want to use a SSH key file instead, when the **Path to SSH Key File** field is edited and the changes are applied, PowerChute will erase your configured passwords.



NOTE: You must use the "nutanix" user account credentials to connect to the Cluster/Controller VM. You cannot use Prism user account credentials to connect.

The SSH key file, if configured, must be located in the user_files folder, or a sub-folder, in the PowerChute installation directory. If the default installation directory was chosen during installation, this location will be:

- C:\Program Files\APC\PowerChute\user_files for Windows systems
- /opt/APC/PowerChute/user_files/ for Linux systems



NOTE: If you are upgrading to PowerChute v4.4, any SSH key files used in a previous version of PowerChute must be manually added to the user_files folder. If the key files are not added to the directory, you will see the below error in the Event Log:

```
ERROR: The ini contains an invalid value for ssh_keyfile_path in nutanixClusterDetails.
```

It is strongly recommended you validate your configuration after an upgrade.

PowerChute Setup: Controller VM/Cluster Details ?

Cluster Details are required to perform graceful shutdown of the Cluster.

Controller VM/Cluster IP Address	<input type="text" value="10.216.249.29"/>
Controller VM/Cluster Password	<input type="password"/>
Path to SSH Key File	<input type="text"/>
SSH Key Passphrase	<input type="password"/>

Password or Path to SSH Key File required.

CVM/Cluster Details

HPE SimpliVity Support

If you have an HPE SimpliVity Cluster with ESXi as your hypervisor, you can enable HPE SimpliVity support for PowerChute in the PowerChute Setup wizard.

Enable HPE SimpliVity Support

To enable HPE SimpliVity support, open the PowerChute Setup wizard and follow the steps below.

The screenshot shows the 'PowerChute Setup: vCenter Server Details' window. The 'Hyperconverged Infrastructure Support' dropdown menu is open, with 'VxRail' selected. The other fields are: vCenter Server Protocol (https), Accept Untrusted SSL Certificates (checked), vCenter Server Port (443), vCenter Server IP Address/Hostname (10.1.1.1), vCenter Server Username (administrator@vsphere.local), vCenter Server Password (masked), and vCenter Server running on a VM (checked). Navigation buttons for Previous, Next, and Cancel are visible at the bottom right.

1. Select **HPE SimpliVity** from the **Hyperconverged Infrastructure Support** drop-down list in the **vCenter Server Details** screen and click **Next**.
2. In the **HPE SimpliVity Details** screen, enter your HPE SimpliVity credentials to connect to your Cluster. The default user name is "svtcli", and it is not recommended to change this.
3. Click **Next** to validate the connection.
4. If the connection to your HPE SimpliVity Cluster is successful, you will be directed to the **Virtualization Settings** screen where you can configure your HPE SimpliVity settings. These settings can later be edited in the [Virtualization Settings](#) screen in the PowerChute UI.

HPE SimpliVity Virtualization Settings

When HPE SimpliVity support is enabled, the [OmniStack Virtual Controller Shutdown/Startup](#) section will appear in the Virtualization Settings screen in the PowerChute UI.

OmniStack Virtual Controller Shutdown/Startup

An OmniStack Virtual Controller (OVC) is deployed on each HPE SimpliVity node and is responsible for managing the HPE SimpliVity Cluster.

OVC Shutdown is enabled by default. In the PowerChute shutdown sequence, the OVC is shut down after all other VMs in the Cluster. The first OVC VM is shut down using the "Shut Down Safe" option in vSphere, which triggers a Wait for Storage HA Compliance task. When this task is complete, the OVC VM is shut down. Subsequent OVC VMs are shut down using the "Force Virtual Controller Shutdown" command.



All OmniStack Virtual Controllers in the HPE SimpliVity Cluster must be powered on and accessible during PowerChute setup to ensure they are shut down and started by PowerChute as expected.

The **Duration** field is the time allowed for all the OVCs to gracefully shut down. If the OVCs are shut down before the Duration time, PowerChute waits until this time has elapsed before proceeding to the next step in the sequence.



Using the **Shutdown Duration** field, you must allow sufficient time for all your OVCs to gracefully shut down before the hosts are commanded to shut down.

PowerChute[™]
—NETWORK SHUTDOWN

powerchuteserver
Logout

Home Configuration Virtualization UPS Configuration Help

Virtualization Settings ?

Configure Virtualization settings for shutdown sequence.

- Virtual Machine Migration ?
- Virtual Machine and Virtual Appliance, Shutdown and Startup ?
- Host Maintenance Mode ?
- vSphere Cluster Services (vCLS) ?
- High Availability ?

OVC Shutdown/Startup

Shutdown Duration seconds

OVC Startup

Startup Duration seconds

All Hosts online prior to startup

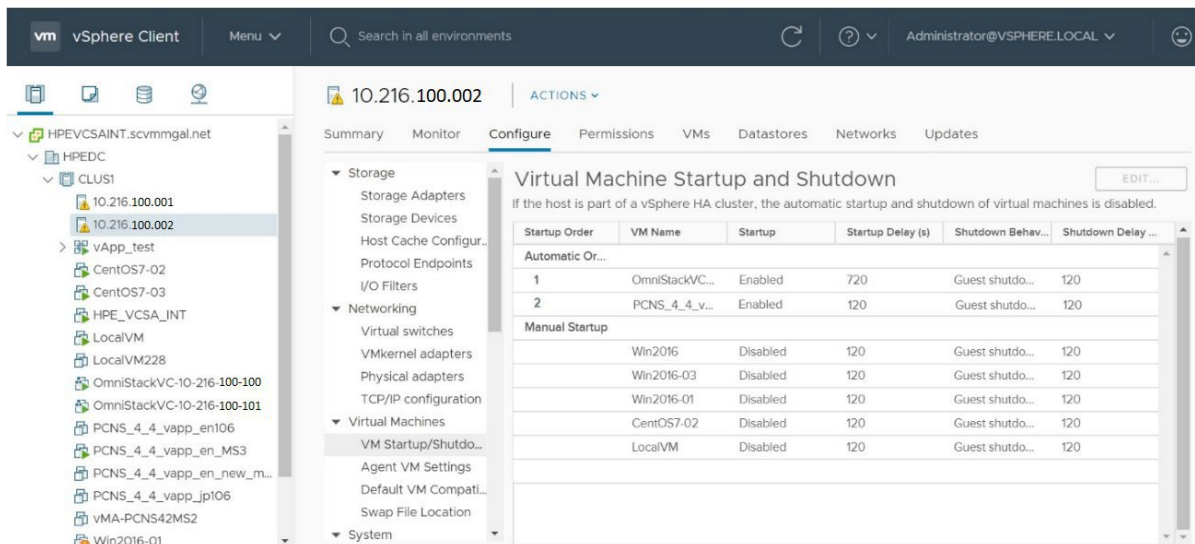
Apply Undo

OVC Shutdown/Startup

Re-starting after a shutdown

If the **Skip Maintenance Mode** checkbox is selected in PowerChute, the OmniStack Virtual Controller VMs will get started automatically using the VMware VM Shutdown/Startup configuration for each host. PowerChute will not start the OmniStack Virtual Controller VMs, but PowerChute will still wait the configured OVC **Startup Duration** before proceeding with User VM startup

If PowerChute is deployed as a VM and VMware HA is disabled for the cluster, the PowerChute VM will be configured to start after the OmniStack Virtual Controller VMs using the VM Shutdown/Startup option in vSphere.



The screenshot shows the vSphere Client interface for a host with IP 10.216.100.002. The left sidebar shows a tree view of the environment, including a cluster named CLUST1 and several VMs. The main pane is titled "Virtual Machine Startup and Shutdown" and contains a table with the following data:

Startup Order	VM Name	Startup	Startup Delay (s)	Shutdown Behav...	Shutdown Delay ...
Automatic Or...					
1	OmniStackVC...	Enabled	720	Guest shutdo...	120
2	PCNS_4_4_v...	Enabled	120	Guest shutdo...	120
Manual Startup					
	Win2016	Disabled	120	Guest shutdo...	120
	Win2016-03	Disabled	120	Guest shutdo...	120
	Win2016-01	Disabled	120	Guest shutdo...	120
	CentOS7-02	Disabled	120	Guest shutdo...	120
	LocalVM	Disabled	120	Guest shutdo...	120



It is recommended to wait at least 5 minutes after OVC startup to ensure all services are running. You should account for this using the **OVC Startup Duration** field.

Cisco HyperFlex Support

If you have a Cisco HyperFlex Cluster with ESXi as your hypervisor, you can enable HyperFlex support for PowerChute in the PowerChute Setup wizard.

Enable Cisco HyperFlex Support

To enable HyperFlex support, open the PowerChute Setup wizard and follow the steps below.

The screenshot shows the 'PowerChute Setup: vCenter Server Details' window. The 'Hyperconverged Infrastructure Support' dropdown menu is open, with 'VxRail' selected. The other fields are: vCenter Server Protocol (https), Accept Untrusted SSL Certificates (checked), vCenter Server Port (443), vCenter Server IP Address/Hostname (10.1.1.1), vCenter Server Username (administrator@vsphere.local), vCenter Server Password (masked), and vCenter Server running on a VM (checked). Navigation buttons for Previous, Next, and Cancel are visible at the bottom right.

1. Select **HyperFlex** from the **Hyperconverged Infrastructure Support** drop-down list in the **vCenter Server Details** screen and click **Next**.
2. In the **HyperFlex Details** screen, enter your HyperFlex credentials to connect to your Cluster.
3. Click **Next** to validate the connection.
4. If the connection to your HyperFlex Cluster is successful, you will be directed to the **Virtualization Settings** screen where you can configure your HyperFlex settings. These settings can later be edited in the [Virtualization Settings](#) screen in the PowerChute UI.

When attempting to connect to HyperFlex Cluster, if an untrusted SSL certificate is detected, a dialog to view the certificate will appear. Please verify that the certificate is correct and accept it in order to proceed. The certificate will be added to the PowerChute-keystore.

Cisco HyperFlex Virtualization Settings

When HyperFlex support is enabled, additional sections will appear in the Virtualization Settings screen in the PowerChute

UI. For more information on these sections, see:

- [Cluster Shutdown/Startup](#)
- [Controller VM Shutdown/Startup](#)

Cluster Shutdown/Startup

Cluster Shutdown occurs after PowerChute has completed the shutdown of User VMs and the vCenter Server VM, if applicable, and before the Controller VMs are shut down.

PowerChute™
NETWORK SHUTDOWN

powerchuteserver
Logout

Home Configuration Virtualization UPS Configuration Help

Virtualization Settings ?

Configure Virtualization settings for shutdown sequence.

- ▶ Virtual Machine Migration ?
- ▶ Virtual Machine and Virtual Appliance, Shutdown and Startup ?
- ▶ Host Maintenance Mode ?
- ▶ vSphere Cluster Services (vCLS) ?
- ▶ High Availability ?
- Cluster Shutdown/Startup
- ▶ Controller VM Shutdown/Startup ?

Shutdown Duration seconds

Startup Duration seconds

All Hosts online prior to startup

Apply Undo

Cluster Shutdown/Startup

If the Cluster is shut down before the Duration time, PowerChute waits until this time has elapsed before proceeding to the next step in the sequence (shutting down the Controller VMs).



Using the **Shutdown Duration** field, you must allow sufficient time for your Cluster to gracefully shut down before the Controller VMs are shut down. It is recommended that you manually test the time needed to stop your Cluster, and specify that as the **Shutdown Duration**.

In the PowerChute configuration file (`pcnsconfig.ini`), the "`cluster_ops_retries`" setting allows you to set a value (the default is 10) for retry attempts. PowerChute will re-attempt to shut down and start up the HyperFlex Cluster with the **Shutdown/Startup Duration** in between each attempt, until the task is successful, or the retry limit has been reached.

Re-starting after a shutdown

The time specified in the **Startup Duration** field will be the time, in seconds, PowerChute will wait after issuing the Cluster start command before proceeding with the next steps in the startup sequence. The Cluster is re-started after the Controller VMs are powered on.

Controller VM Shutdown/Startup

A Controller Virtual Machine is installed on each HyperFlex node in a cluster and is responsible for managing the HyperFlex cluster storage.

Controller VM Shutdown is enabled by default. In the PowerChute shutdown sequence, the Controller VM is shut down after all other VMs in the Cluster, and the HyperFlex Cluster itself are shut down.

The **Duration** field is the time allowed for all the Controller VMs to gracefully shut down. If the Controller VMs are shut down before the Duration time, PowerChute waits until this time has elapsed before proceeding to the next step in the sequence.



Using the **Shutdown Duration** field, you must allow sufficient time for all your Controller VMs to gracefully shut down before the hosts are commanded to shut down.

PowerChute™
— NETWORK SHUTDOWN

powerchuteserver
Logout

Home Configuration Virtualization UPS Configuration Help

Virtualization Settings ?

Configure Virtualization settings for shutdown sequence.

- ▶ Virtual Machine Migration ?
- ▶ Virtual Machine and Virtual Appliance, Shutdown and Startup ?
- ▶ Host Maintenance Mode ?
- ▶ vSphere Cluster Services (vCLS) ?
- ▶ High Availability ?
- ▶ Cluster Shutdown/Startup ?

Controller VM Shutdown/Startup

Shutdown Duration	<input type="text" value="120"/>	seconds
Controller VM Startup	<input checked="" type="checkbox"/>	
Startup Duration	<input type="text" value="300"/>	seconds

All Hosts online prior to startup

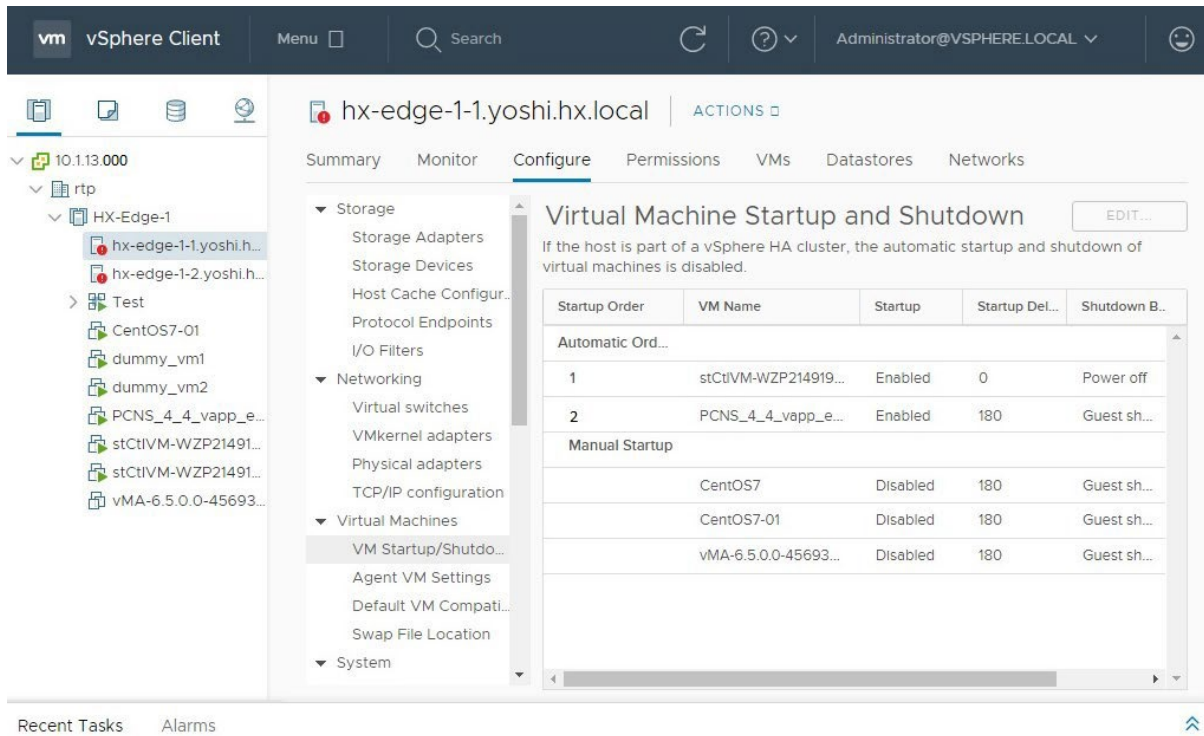
Apply Undo

Controller VM Shutdown/Startup

Re-starting after a shutdown


If the **Skip Maintenance Mode** checkbox is selected in PowerChute, the Controller VMs will get started automatically using the VMware VM Shutdown/Startup configuration for each host. PowerChute will not start the Controller VMs, but PowerChute will still wait the configured Controller VM **Startup Duration** before proceeding with User VM startup

If PowerChute is deployed as a VM and VMware HA is disabled for the cluster, the PowerChute VM will be configured to start after the Controller VMs using the VM Shutdown/Startup option in vSphere.




The screenshot shows the vSphere Client interface for a host named 'hx-edge-1-1.yoshi.hx.local'. The 'Configure' tab is selected, and the 'Virtual Machine Startup and Shutdown' configuration is displayed. The configuration is for a host that is part of a vSphere HA cluster, where automatic startup and shutdown of virtual machines is disabled. A table lists the startup order, VM name, startup status, startup delay, and shutdown behavior for several VMs.


Startup Order	VM Name	Startup	Startup Del...	Shutdown B...
Automatic Ord...				
1	stCtIVM-WZP214919...	Enabled	0	Power off
2	PCNS_4_4_vapp_e...	Enabled	180	Guest sh...
Manual Startup				
	CentOS7	Disabled	180	Guest sh...
	CentOS7-01	Disabled	180	Guest sh...
	vMA-6.5.0.0-45693...	Disabled	180	Guest sh...

 It is recommended to wait at least 5 minutes after Controller VM startup to ensure all services are running before re-starting the Cluster. You should account for this using the **Controller VM Startup Duration** field.

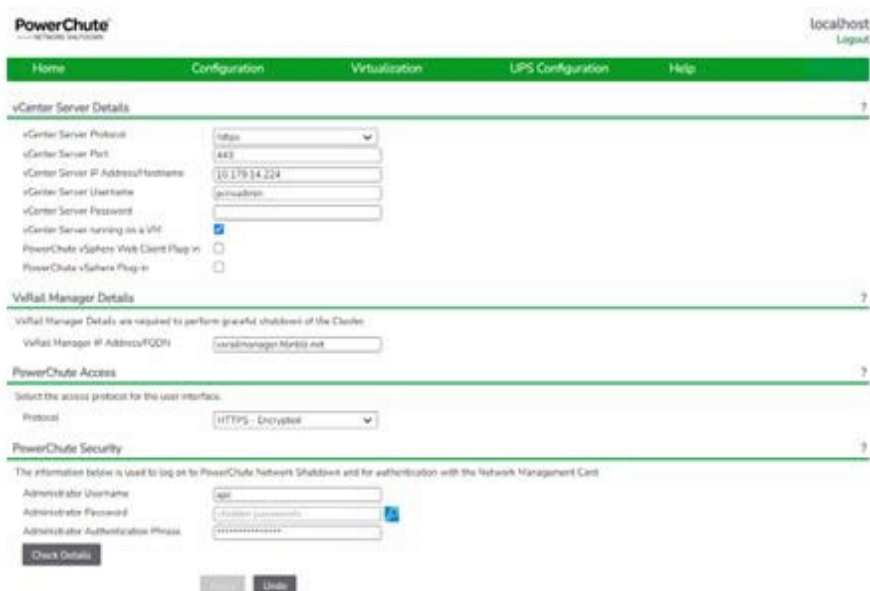
Dell VxRail Support

If you have a Dell VxRail Cluster with ESXi as your hypervisor, you can enable Dell VxRail support for PowerChute in the PowerChute Setup wizard.

 Dell VxRail support is available with the Network Management Card 3. Please refer to the [Operating System and Compatibility Chart](#) for the supported NMC 3 firmware version number.

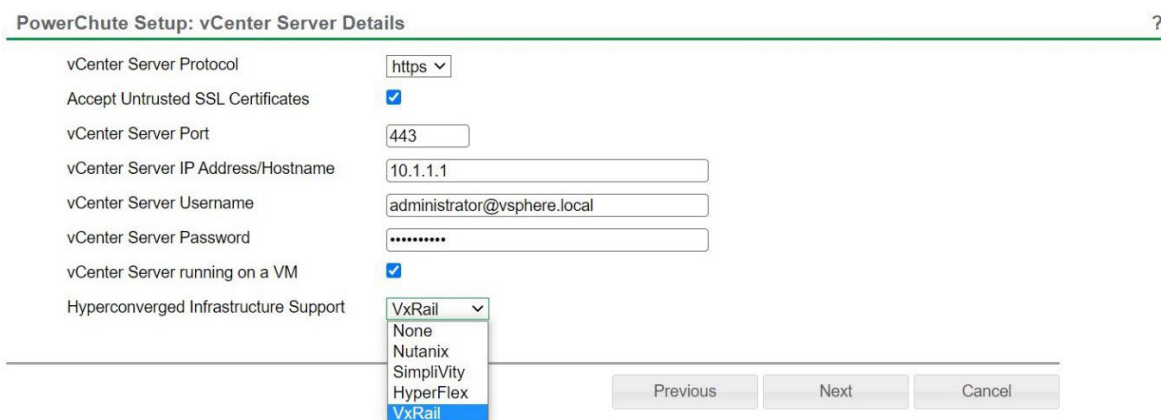
 It is not supported to select HTTP as the communication protocol to communicate with the NMC for Dell VxRail configuration. HTTPS must be selected.

When attempting to connect to VxRail Manager, if an untrusted SSL certificate is detected, a dialog to view the certificate will appear. Please verify that the certificate is correct and accept it in order to proceed. The certificate will be added to the PowerChute-keystore.



Enable Dell VxRail Support

To enable Dell VxRail support, open the PowerChute Setup wizard and follow the steps below.



1. Ensure **VxRail** is selected in the **Hyperconverged Infrastructure Support** drop-down list in the **vCenter Server Details** screen and click **Next**.
2. In the **VxRail Manager Details** screen, enter your VxRail Manager VM IP address or hostname to continue.

NOTES:

- PowerChute requires the VxRail Manager to be available to validate your credentials and configuration.
 - For a standard cluster configuration, a vCenter Server administrator account credentials must be provided.
 - In a stretched cluster configuration, PowerChute will shut down the Witness and Management hosts by directly connecting to them after the vCenter Server VM is shut down. You must configure a shared local user account with the Administrator role on vCenter Server and each VMware host and provide its credentials. For more information, see the [Active Directory VMware Configuration](#) topic.
3. Click **Next** to validate the connection.
 4. If the connection to your Dell VxRail Manager VM/Cluster is successful, you will be directed to the **Virtualization Settings** screen where you can configure your Dell VxRail settings. These settings can later be edited in the [Virtualization Settings](#) screen in the PowerChute UI.

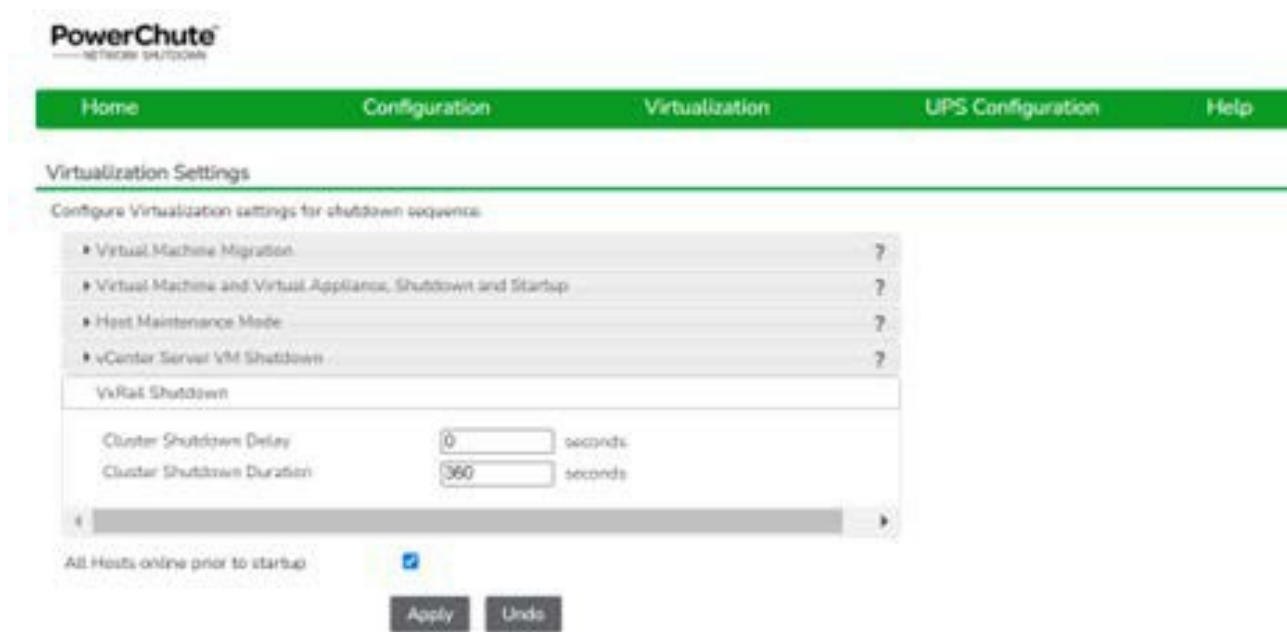
Dell VxRail Virtualization Settings

When Dell VxRail support is enabled, an additional section will appear in the Virtualization Settings screen in the PowerChute UI. For more information on this section, see:

- [VxRail Shutdown and Startup](#)


VxRail Shutdown and Startup


VxRail Cluster Shutdown occurs after PowerChute has completed the shutdown of User VMs. This is initiated by the NMC using the cluster shutdown REST API call via VxRail Manager.




VxRail Shutdown and Startup

The **Cluster Shutdown Delay** is the delay used by the NMC to allow sufficient time for the PowerChute VM to shut down, before the NMC issues a shutdown command to the VxRail Manager. The default value is 120 seconds.

 The **Cluster Shutdown Delay** should be set to 0 seconds for VxRail stretched cluster configurations.

 When a shutdown command file and/or SSH action is configured in a standard cluster configuration, the **Cluster Shutdown Duration** provided must be sufficient enough to allow the command file/SSH action to run and the PowerChute VM to shut down. If a sufficient **Duration** is not provided, the VxRail cluster will not shut down successfully.

Using the **Cluster Shutdown Duration** field, you must allow sufficient time for your Cluster to gracefully shut down once the VxRail Manager has been requested to stop the Cluster. It is recommended that you manually test the time needed to stop your Cluster, and specify that as the **Shutdown Duration**. The default value is 300 seconds.

 In a VxRail stretched cluster configuration, PowerChute waits for the **Cluster Shutdown Duration** to elapse before proceeding with the next step of the shutdown sequence. If the duration provided is not sufficient, cluster shutdown may be unsuccessful.

Re-starting after a shutdown

Depending on your configuration, PowerChute is restarted after a shutdown by one of two methods. If the PowerChute VM and vCenter Server VM are deployed inside the VxRail cluster, PowerChute restarts via a scheduled task after vCenter Server starts. For all other configurations, PowerChute is restarted via [VMware's auto-start policy](#).



NOTE: When the Dell VxRail cluster is restarted, the vSAN datastore is unavailable until all the cluster services are running. In a stretched cluster configuration, it is highly recommended that a value is set for `delay_before_vmstartup` in the `[HostConfigSettings]` section of the PowerChute configuration file. This setting ensures that PowerChute restarts User VMs when the vSAN datastore is accessible.

Scheduled Task

Using this method, PowerChute registers a scheduled task during shutdown to start the PowerChute VM after vCenter Server with a given delay. The delay can be configured in the PowerChute configuration file (`pcnsconfig.ini`) using the `"scheduled_start_delay"` setting in the `VxRailCluster` section.

Once set, the scheduled task can be viewed in vCenter by selecting the PowerChute VM, and navigating to **Configure > Scheduled Tasks**.

If the VxRail cluster services have not fully started, this may be unsuccessful and the PowerChute VM will need to be started manually.

Auto-Start Policy

Using this method, PowerChute registers itself with the Virtual Machine Startup and Shutdown policies on vCenter and the PowerChute VM is restarted with its host. The Virtual Machine Startup and Shutdown policies can be viewed in vCenter by selecting a cluster host, clicking **Configure**, and navigating to **Virtual Machines > VM Startup/Shutdown**.



The auto-start policy will be enabled if PowerChute is deployed as a VM on any VMware setup.



Ensure that the auto-start feature is enabled in ESXi. If it is not enabled in ESXi, PowerChute may not successfully restart after a shutdown. To enable auto-start in ESXi, click on a cluster host, click **System**, click **Autostart**, and set **Enabled** to "Yes".

NOTE: This feature cannot be enabled when HA is enabled.

VxRail manages restarting the VxRail cluster and system VMs. VMware auto-start will attempt to power on the PowerChute VM when it's host reboots. When PowerChute is deployed inside the VxRail cluster this will not work as the hosts will be in maintenance mode after they are powered on. For details on how to have PowerChute automatically power on in this scenario, see Knowledge base article [FAQ000258432](#).

The startup delay has a default value of 0 seconds. This can be configured in the PowerChute configuration file (`pcnsconfig.ini`) using the `"autostart_delay"` setting in the `VxRailCluster` section.



When Dell VxRail and Management hosts are powered on at the same time in a stretched cluster configuration, the datastore may become inaccessible and cluster startup may not work as expected. To prevent this issue, ensure that the Management host is started first followed by the VxRail hosts so the Witness host and vCenter Server are available when the VxRail hosts start up. For more information on how this can be achieved, see the "Cannot start Cluster" section in the [Dell VxRail Troubleshooting](#) topic.

Shutdown Settings

The Shutdown Settings page enables you to configure UPS turnoff and the shutdown command files.

- [UPS Shutdown](#)
- [Shutdown Command Files](#)
- [Shutdown Settings for Advanced UPS Configurations](#)

UPS Shutdown

The default setting is **Do not turn off the UPS**.

You can select **Turn off the UPS** if you want to preserve battery power. Some UPS's do not support UPS turnoff through PowerChute or the NMC. For these models, it can only be done at the UPS itself. Please check your UPS documentation to ensure your model supports UPS turnoff.

If your UPS has Switched Outlet Groups, then the **Turn off the UPS Outlet Group** option enables you to turn off the outlet group that supplies power to the PowerChute protected server after a critical event occurs.



The default behavior for most UPS's if they are turned off following an on-battery shutdown is that they will turn on again once input power is restored.

The **On-Battery Shutdown Behavior** setting can be found in the NMC under Configuration – Shutdown where you can change the behavior to Turn off and Stay off if required.

Turn Off Single UPS On Battery in a Redundant-UPS Configuration



This is not available for an Advanced UPS Configuration that contains UPS Setups with Redundant UPS devices.

In a Redundant UPS configuration you have the option to turn off one of the UPS's after it has switched to battery power. This is designed to prolong the battery life and preserve the battery power of the UPS. If using this feature on a UPS that supports outlet groups the option "Turn off the UPS" should be enabled.

The load is still protected by the other UPS in the configuration.

After the specified delay, PowerChute will issue a command to gracefully turn off the UPS.



If one UPS is on battery and another UPS switches to battery before the configured delay for Single UPS turn off has elapsed, then the first UPS will not be turned off.



If the shutdown action is enabled for the **On Battery** event, a Multiple Critical event condition will occur if a second UPS switches to battery power (after the first UPS has been commanded to turn off by PowerChute). When this occurs the shutdown sequence will start after 10 seconds.

Shutdown Command Files

A Shutdown Command File can be configured to run if a UPS critical event is triggered. In a [Dell VxRail standard cluster configuration](#), the command file is executed after PowerChute sends the NMC request to shut down the VxRail cluster, before its local OS shutdown. In a stretched cluster configuration, the command file is executed after the NMC shuts down the VxRail cluster, before the Witness host is shut down.

Full path to command file: You must specify the full path name of the command file, including the disk drive or volume name. For Linux and Unix systems, the file should execute permissions of **chmod +x [command file name]**.

The command file must be located in the user_files folder, or a sub-folder, in the PowerChute installation directory. If the default installation directory was chosen during installation, this location will be:

- C:\Program Files\APC\PowerChute\user_files for Windows systems
- /opt/APC/PowerChute/user_files/ for Linux systems



NOTE: If you are upgrading to PowerChute v5.1, any command files used in a previous version of PowerChute must be manually added to the user_files folder. If the command files are not added to the directory, you will see the below errors/warnings for the ShutdownStarting event, and all events configured to run a command file, in the Event Log:

```
ERROR: Event ShutdownStarting is enabled for command file execution,
but an invalid value for shutdownCommandFile is specified.
```

```
WARNING: Disabling command file execution for event ShutdownStarting
due to bad parameters. Please validate the configuration.
```

It is strongly recommended you validate your configuration after an upgrade.

Duration: Enter the number of seconds that the shutdown command file requires to execute.



NOTE: For Advanced UPS Configurations, if there are different command files configured for each UPS Setup, PowerChute may need to wait for all command files to finish executing before proceeding with the final steps in the shutdown sequence. This is dependent on the timing that UPS critical events occur on each UPS Setup. PowerChute will automatically increase the Outlet Group Power Off delay or Maximum Required delay (non-outlet aware UPS) to include the combined total of the shutdown command file durations for each UPS Setup. This can impact the runtime available on the UPSs during a shutdown. To accommodate this, set the low battery duration on the UPSs accordingly.

Execute Command File After Host Shutdown: Enable this option to execute the command file following Host Shutdown. This option is only available when PowerChute is installed on a physical machine, and not on a VM.

Delay: Enter the number of seconds that the Host requires to shut down, before the command file is executed.



You must determine the time required for your command file to execute. PowerChute cannot determine whether the command file has completed, so it will wait only the amount of time entered before triggering an operating system shutdown.

The command file runs using the local system account. For Linux/Unix the command file must be executed with root privileges. PowerChute cannot execute programs that require interaction with the desktop; only command line enabled programs are supported.




Shutdown Settings for Advanced UPS Configurations

With Advanced UPS configurations, PowerChute can monitor multiple UPS setups, including single UPS devices and groups of redundant UPS devices that you have created (see [“Advanced UPS Setups”](#)).

For each setup, you need to specify the following:

Field	Description								
Number of UPS's required to power load	<p>Set this value to the minimum number of UPS's that must be available to support the equipment that is being powered by the UPS's in the setup.</p> <p>The value set here will be subtracted from the total number of UPS's in the setup and used to calculate the number of additional (redundant) UPS's.</p> <p>In redundancy terminology, this is the N in N+x.</p> <p>This setting is not displayed for UPS Setups with a Single UPS device.</p>								
Number of additional (redundant) UPS's	<p>This will appear in a setup with more than one UPS. It represents the number of extra UPS's in the setup. This option is associated with the number of UPS critical events required to trigger shutdown:</p> <table border="1" data-bbox="507 1131 1241 1509"> <thead> <tr> <th data-bbox="507 1131 874 1279">Redundancy level</th> <th data-bbox="879 1131 1241 1279">No. of critical events that will trigger a shutdown sequence</th> </tr> </thead> <tbody> <tr> <td data-bbox="507 1285 874 1352">N+1</td> <td data-bbox="879 1285 1241 1352">2</td> </tr> <tr> <td data-bbox="507 1359 874 1426">N+2</td> <td data-bbox="879 1359 1241 1426">3</td> </tr> <tr> <td data-bbox="507 1433 874 1500">N+3</td> <td data-bbox="879 1433 1241 1500">4</td> </tr> </tbody> </table> <p>Multiple critical events occurring on the same UPS does not impact the above table values.</p> <p>In redundancy terminology, this is the x in N+x.</p> <p>This setting is not displayed for UPS Setups with a Single UPS device.</p>	Redundancy level	No. of critical events that will trigger a shutdown sequence	N+1	2	N+2	3	N+3	4
Redundancy level	No. of critical events that will trigger a shutdown sequence								
N+1	2								
N+2	3								
N+3	4								
Total number of UPS's in Setup	This is the total of the above two rows and is calculated automatically.								

Run Command	<p>When a shutdown sequence is triggered you can configure PowerChute to execute a command file.</p> <p>Note: If the same command file is configured for each setup and a shutdown sequence is triggered for more than one setup at the same time, the command file is only executed once.</p> <p>See Shutdown Command Files.</p>								
Shut down PowerChute Server	<p>This is enabled by default and is used to gracefully shut down the physical machine running PowerChute.</p> <p>This option should be disabled if the PowerChute machine is not being powered by the UPS's in a particular setup, and if it is being used to remotely shut down other servers/ equipment.</p> <p>This option is not available if PowerChute is installed on vMA or deployed as a virtual appliance.</p>								
Shut down if Redundancy lost	<p>If this option is enabled, when the number of UPS critical events is the same as the number of additional (redundant) UPS's, a shutdown sequence will be triggered.</p> <p>This option is associated with the number of UPS critical events required to trigger shutdown:</p> <table border="1" data-bbox="507 1122 1240 1503"> <thead> <tr> <th data-bbox="512 1128 874 1272">Redundancy level</th> <th data-bbox="879 1128 1235 1272">No. of critical events that will trigger a shutdown sequence</th> </tr> </thead> <tbody> <tr> <td data-bbox="512 1279 874 1352">N+1</td> <td data-bbox="879 1279 1235 1352">1</td> </tr> <tr> <td data-bbox="512 1359 874 1433">N+2</td> <td data-bbox="879 1359 1235 1433">2</td> </tr> <tr> <td data-bbox="512 1440 874 1503">N+3</td> <td data-bbox="879 1440 1235 1503">3</td> </tr> </tbody> </table> <p>Multiple critical events occurring on the same UPS does not impact the above table values.</p> <p>This option is not shown if there are no additional (redundant) UPS's. For example, this option will not appear if the number of UPS's required to power the load is the same as the total number of UPS's in the group.</p>	Redundancy level	No. of critical events that will trigger a shutdown sequence	N+1	1	N+2	2	N+3	3
Redundancy level	No. of critical events that will trigger a shutdown sequence								
N+1	1								
N+2	2								
N+3	3								
UPS Shutdown	<p>Use this option to set the required UPS behavior after connected equipment or servers are gracefully shut down. For more information see UPS Shutdown.</p>								

<p>Execute Virtualization Shutdown on Hosts in all UPS Setups</p>	<p>Use this option to trigger the actions enabled on the Virtualization Settings page for ESXi hosts that are linked to UPSs/UPS Setups: VM migration/VM Shutdown/vApp shutdown followed by ESXi host shutdown.</p> <p>This option is available only in a configuration in which a UPS setup is powering something other than a virtual host (e.g. a storage array) - it is enabled by default for a Physical UPS setup.</p> <table border="1" data-bbox="507 495 1228 797"><tr><td data-bbox="513 504 651 788"></td><td data-bbox="657 504 1222 788"><p>If Execute Virtualization Shutdown on Hosts in all UPS Setups is enabled in the UI or PowerChute configuration file ("ShutdownVirtualHost" setting), all protected hosts in all UPS setups will be shut down when a critical event occurs on that UPS Setup.</p></td></tr></table>		<p>If Execute Virtualization Shutdown on Hosts in all UPS Setups is enabled in the UI or PowerChute configuration file ("ShutdownVirtualHost" setting), all protected hosts in all UPS setups will be shut down when a critical event occurs on that UPS Setup.</p>
	<p>If Execute Virtualization Shutdown on Hosts in all UPS Setups is enabled in the UI or PowerChute configuration file ("ShutdownVirtualHost" setting), all protected hosts in all UPS setups will be shut down when a critical event occurs on that UPS Setup.</p>		

SSH Settings

PowerChute Network Shutdown can be configured to execute commands on a remote host via an SSH connection. To create an SSH action, click **Add Action** and configure the following:

1. **Name:** A unique name for each SSH action of a length less than or equal to 255 ASCII characters.
2. Configure one of the following authentication methods:
 - **User Name** and **Password:** Enter the user name and password to connect to the remote host.
 - **User Name, SSH Key File Path** and **SSH Key File Password:** Specify the path to a shared SSH key. This option requires you to generate an SSH key and copy it to your target systems. See [SSH Command File Location](#).
3. **IP Addresses/FQDN** and **Port:** The IP address or Fully Qualified Domain Name (FQDN) and port of the target SSH component.
4. **Path to SSH command file:** You must specify the full path name of the command file, including the disk drive or volume number. See [SSH Command File Location](#).
5. **SSH Action Delay:** Enter the amount of time, in seconds, that PowerChute will wait before connecting to the remote host and begin sending commands. The default value is 0.
6. **SSH Action Duration:** Enter the amount of time, in seconds, for the SSH action to complete before proceeding with the rest of the shutdown sequence.



Using the **SSH Action Duration** field, you must allow sufficient time for all your SSH actions to complete.

7. **Execute SSH Action:**
 - **On Startup:** Execute the SSH command file when the PowerChute service re-starts. In an advanced configuration, the SSH command file is executed when the critical UPS event is resolved.
 - **Before VM Shutdown:** Execute the SSH command file before the User VM Migration and Shutdown Step.
 - **Before Host Shutdown:** Execute the SSH command file before host shutdown.
 - **After Host Shutdown:** Execute the SSH command file after host shutdown.
8. **Enable SSH Action:** Allows you to enable or disable the configured SSH action. This checkbox is enabled by default when a new SSH action is created.

Name	BeforeHostShutdown	
User Name	root	
Password	*****	
SSH Key File Path		
SSH Key File Password		
IP Address/Hostname	10.179.232.189	
Port	22	
Path to SSH command file	C:\Program Files\APC\PowerChute\us	
SSH Action Delay	10	seconds
SSH Action Duration	30	seconds
Execute SSH Action	Before Host Shutdown	
Enable SSH Action	<input checked="" type="checkbox"/>	

NOTES:

- PowerChute takes the command file provided and passes it line-by-line to the remote host over an SSH connection. As a result, incomplete lines may be interpreted incorrectly by the remote host. You must ensure that your SSH command file contains complete lines and commands so the remote host can interpret the file correctly.
- The line ending style of the command file must match that of the PowerChute target host operating system. For example, a command file configured on PowerChute running on a Windows host must contain Windows style text line endings.
- Recognized command prompts are:
 - \$ (Linux)
 - # (Linux admin/root)
 - > (Windows, or RPDU)
- Custom command prompts can be added via the PowerChute configuration file (`pcnsconfig.ini`) by adding the `ssh_prompt_regex` setting to the `[SSHAction]` section. For example: to add a custom command prompt of `~`, add `ssh_prompt_regex = \~\s`.
- The PowerChute Event Log only displays that an SSH action has completed. The Event Log does not show if the SSH action has completed successfully or not.
- If a value is specified in the **SSH Action Delay** field, the Event Log does not log that an SSH action is running with a configured delay.



SSH Command File Location

The SSH key file, if configured, and command file must be located in the `user_files` folder, or a sub-folder, in the PowerChute installation directory. If the default installation directory was chosen during installation, this location will be:

- `C:\Program Files\APC\PowerChute\user_files` for Windows systems
- `/opt/APC/PowerChute/user_files/` for Linux systems

NOTE: If you are upgrading to PowerChute v5.1, any key files and/or command files used in a previous version of PowerChute must be manually added to the `user_files` folder. If the files are not added to the directory, you will see the below errors in the Event Log:





```
ERROR: The ini contains an invalid value for ssh_command_file_path in  
section SSHAction0.
```


```
ERROR: The ini contains an invalid value for ssh_keyfile_path in  
section SSHAction0.
```

It is strongly recommended you validate your configuration after an upgrade.

SSH Settings in an Advanced UPS Configuration

In an Advanced UPS configuration, SSH actions can be enabled and disabled for each UPS setup.

The  symbol indicates that an SSH action is enabled in the main SSH Settings screen. To run an SSH action for a particular UPS setup, enable the checkbox next to the  symbol.

SSH actions that are not enabled do not display the  symbol. These disabled actions will not be executed if enabled for a UPS setup.



SNMP Configuration

PowerChute Network Shutdown can be configured to communicate via Simple Network Management Protocol (SNMP), and can be discovered via SNMP by Network Management tools, such as StruxureWare Data Center Expert. Using SNMP, you can query and configure PowerChute settings, and generate SNMP traps for UPS critical events and lost communication events.

SNMPv1 and SNMPv3 are supported by PowerChute Network Shutdown. IPv4 and IPv6 are both supported. Go to **SNMP Settings** in the web user interface to complete the configuration and make PowerChute accessible via SNMP. It is not necessary to re-start the PowerChute service when enabling SNMP or making SNMP configuration changes via the web user interface. PowerChute configuration changes via SNMP are logged to the Event Log.

Enter the **SNMP Discovery Port**. The default value of 161 is automatically populated, but this can be edited if this port is already in use. The Port number availability is automatically checked, and if it is not available, a new port number must be entered.



See:

- [SNMPv1 Configuration](#)
- [SNMPv3 Configuration](#)
- [SNMP Trap Configuration](#)
- [SNMP Data Points](#)
- [SNMP Troubleshooting](#)


SNMPv1 Configuration


Select **Enable SNMPv1 access** to configure the User Profiles required to communicate via SNMPv1. Select **Add Profile** and configure:

1. **Community Name:** The Community Name is sent with each SNMP request to obtain access to a device. The maximum length is 15 ASCII characters.
2. **NMS IP/Host Name:** The IP address, Host Name or Fully Qualified Domain Name of the Network Management System (NMS). An NMS is software that is used to manage software and hardware components on the network. It can be used to manage PowerChute via SNMP by issuing SNMP GET and SET commands. The default value of 0.0.0.0 permits access from any NMS.
3. **Access Type:**
 - **Disable:** No SNMP GET or SET requests are permitted.
 - **Read:** Only SNMP GET requests are permitted.
 - **Read/Write:** SNMP GET and SET requests are permitted.

To edit an existing SNMPv1 user profile, click the  button. To delete an SNMPv1 user profile, click the  button.

Click **Apply** to save the SNMPv1 configuration.

 Certain Network Management Systems require the SNMP Engine ID to communicate via SNMP. The **SNMP Engine ID** is displayed on the SNMP Settings page of the PowerChute user interface.

 SNMPv1 is less secure than SNMPv3. SNMPv1 does not provide encryption or authentication, and the Community Name is sent over the network in plain text. To use encryption and authentication with SNMP, [configure SNMPv3 settings](#).

SNMPv3 Configuration

Select **Enable SNMPv3 access** to configure the SNMPv3 settings. Select **Add Profile** and configure:

1. **User Name:** In SNMPv3, all GET and SET requests and SNMP Traps are matched to a user profile by the User Name. Enter a user name of a length less than or equal to 32 ASCII characters.
2. **Authentication Protocol:** Select MD5, SHA-1 or SHA-2 (SHA256 or SHA512) protocol. It is recommended to use an SHA-2 protocol, if the NMS supports it.



It is not recommended to use the MD5 protocol.

3. **Authentication Passphrase:** Enter an authentication password for the protocol selected, of 8-32 ASCII characters.
4. **Privacy Protocol:** Select AES-128, AES-192*, AES-192 Ex†, AES-256* , AES-256 Ex†, or DES. It is recommended to use the AES-256 protocol, if the NMS and PowerChute JRE support it.
5. **Privacy Passphrase:** Enter a privacy password for the encryption protocol selected, of 8-32 ASCII characters.
6. **Access Type:**
 - **Disable:** No SNMP GET or SET requests are permitted.
 - **Read:** Only SNMP GET requests are permitted.
 - **Read/Write:** SNMP GET and SET requests are permitted.

To edit an existing SNMPv3 user profile, click the  button. To delete an SNMPv3 user profile, click the  button.

Click **Apply** to save the SNMPv3 configuration.



Certain Network Management Systems require the SNMP Engine ID to communicate via SNMP. The **SNMP Engine ID** is displayed on the SNMP Settings page of the PowerChute user interface.

SNMP Trap Configuration

You can specify the device(s) that receive the SNMP traps generated by PowerChute for UPS critical and lost communication events.

To configure a Trap Receiver, select **Add Trap Receiver** and configure:

1. **Enable:** Select the checkbox to enable the Trap Receiver.
2. **NMS IP/Host Name:** The IP address, Host Name or Fully Qualified Domain Name of the NMS.
3. **Port:** The port on which the NMS will listen for incoming traps. The default port number is 162.
4. **SNMPv1:** Select this if you want to send the traps via SNMPv1.
 - **Community Name:** Enter the Community Name of the SNMPv1 user profile to be used as an identifier when SNMPv1 traps are sent to this receiver.
5. **SNMPv3:** Select this if you want to send the traps via SNMPv3.
 - **User Name:** Select the user name of the SNMPv3 user profile to be used as an identifier when SNMPv3 traps are sent to this receiver.

Click the **SNMP Trap Receiver Test** to send a test trap to the configured Trap Receiver. Check the Trap Receiver to ensure that the test trap was received.

To edit an existing SNMP Trap Receiver, click the  button. To delete an SNMP Trap Receiver, click the  button.

UPS Critical Events

PowerChute sends SNMP traps to the configured Trap Receiver(s) upon the following events:

- **PowerChute Critical Event triggers a Shutdown**
When a critical event (such as On Battery) occurs and a Shutdown is triggered, PowerChute sends an SNMP trap detailing the **Event Name**, **UPS Setup** (for advanced configurations), and **Affected Virtual Hosts** (if Virtualization support is enabled).
- **PowerChute Critical Event Resolved**
If the option to **Send Trap when condition is cleared** is enabled, when a PowerChute Critical Event which triggered a Shutdown is resolved, PowerChute sends an SNMP trap to the configured NMS.

Lost Communication Events

PowerChute sends SNMP traps to the configured Trap Receiver(s) upon the following events:

- **Network Communications Lost**
If PowerChute cannot communicate with the Network Management Card of the UPS, a trap is sent to the configured Trap Receiver.
- **UPS Communications Lost**
If the Network Management Card cannot communicate with the UPS, a trap is sent to the configured Trap Receiver.

If the option to **Send Trap when condition is cleared** is enabled, the following traps are sent:

- **Network Communications Lost Resolved**
If PowerChute regains communication with the Network Management Card of the UPS, a trap is sent to the configured Trap Receiver.
- **UPS Communications Lost Resolved**
If communication is regained between the NMC and the UPS, a trap is sent to the configured Trap Receiver.

Other Events

- **Software Update Available Trap**

When the PowerChute Auto Update functionality detects that there is a new update available; a trap is sent to the configured Trap Receiver.


- **PowerChute Test Trap**

When configuring a Trap Receiver, a test trap can be sent to determine if the Trap Receiver is receiving the traps. See [SNMP Trap Receiver Test](#).

See [SNMP Data Points > PowerChute Traps](#) for more information on PowerChute SNMP Trap OIDs.

Configuring SNMP Trap Notification Settings

To configure the settings for UPS Critical Event or Lost Communication traps:

1. Go to **SNMP Settings > SNMP Traps**
2. Click on the  icon next to **UPS Critical Events** or **Lost Communication Events**
3. Select the **Enable** checkbox to enable traps for critical events.
4. **Delay**: Specify the length of time that Event must persist before a trap is sent. If the Event is cleared before this time, no trap is sent.
5. **Repeat Interval**: Specify the time interval in seconds that the trap is re-sent.
6. Select:
 - **Repeat until condition clears** if you want the trap to be sent at the repeat interval until the Event is cleared.
 - **Repeat X times** to specify the number of times the trap will be sent when the Event occurs.
7. Select **Send Trap when condition is cleared** to be notified when the Event is cleared.

Note: If the PowerChute server is shut down due to a UPS Critical Event, no clearing Trap will be sent to the NMS.

SNMP Data Points

The tables below describe the PowerChute configuration details that are available for SNMP polling and/or configuration.

PowerChute Identity Information

Object Identifier Name	Access	Description
pcnshostname	read-only	The hostname of the PowerChute instance.
pcnsVersion	read-only	The version of PowerChute installed.
pcnsOS	read-only	The version of the Operating System upon which PowerChute is installed.
pcnsJavaVersion	read-only	The version of Java upon which PowerChute is running.

PowerChute Networking Settings

Object Identifier Name	Access	Description
pcnsUIProtocol	read-only	The web protocol that is used to connect to the PowerChute web user interface.
pcnsHttpPort	read-only	The port that is used to connect to the PowerChute web user interface.
pcnsHttpsPort	read-only	The port that is used to connect via https to the PowerChute web user interface.
pcnsNetworkConfig	read-only	Configuration of the TCP network: IPv4/IPv6.
pcnsVirtualInstall	read-only	The Virtualization technology for which PowerChute is configured.

PowerChute Network Management Card Settings

Object Identifier Name	Access	Description
pcnsMode	read-only	The configuration of the UPSs that PowerChute is monitoring. See UPS Configuration Options .
pcnsNMCPort	read-only	The port used to connect to all of the Network Management Card(s).
pcnsNMCProtocol	read-only	The web protocol used to connect to all of the Network Management Card(s).
NMC details are contained in an SNMP table named <code>pcnsNmcTable</code> . Each table entry contains:		
pcnsNMCIndex	read-only	The Index of the NMC within the PowerChute setup.
pcnsNMCAddr	read-only	The IP address of the NMC
pcnsNMCOutlet	read-only	The Outlet Group of the NMC on which PowerChute is enrolled.

PowerChute Shutdown Settings

Object Identifier Name	Access	Description
pcnsShutdownCommandFileEnabled*	read-write	Specify if Shutdown a Command File is configured to run if a UPS critical event is triggered. See Shutdown Command Files .
pcnsShutdownCommandFileDelay*	read-write	The number of seconds that a host requires to shut down, before the command file is executed. This setting is applicable to virtualization support only.
pcnsShutdownCommandFile*	read-write	The full path name of the command file, including the disk drive or volume name. See Shutdown Command Files .
pcnsShutdownCommandFileDuration*	read-write	The number of seconds that the shutdown command file requires to execute.
pcnsTurnOffUps*	read-write	The setting to turn off the UPS after performing a graceful shutdown.
pcnsTurnOffSOG*	read-write	The setting to turn off the Outlet Group of the UPS when performing a graceful shutdown.
pcnsRuntimeRemainingThreshold	read-write	This defines a threshold for runtime remaining. When the UPS is running on battery power and the runtime remaining on the UPS drops below the threshold, PowerChute triggers a shutdown sequence. See Sequenced Server Shutdown for more information.

pcnsRuntimeRemainingCmdFileThreshold	read-write	This defines a threshold for runtime remaining. When runtime remaining drops below this threshold, PowerChute executes the command file.
* Note: These OIDs are not available for Advanced UPS Setups. See below for equivalent OIDs for Advanced UPS Setups.		

Shutdown Settings for Advanced UPS Setups

The PowerChute Shutdown Settings for Advanced UPS Setups are contained in an SNMP Table named `pcnsAdvancedShutdownTable`. To SNMP GET or SET a property of a specific Advanced UPS Configuration, you must specify the group number of the UPS Setup in the GET or SET command. To see a list of UPS Setups and their associated numbers, perform an SNMP Get on `.pcnsAdvancedGroupName`

To perform an SNMP GET or SET command on each UPS Setup is, use the format. `[Object Identifier Name].[UPS Setup Number]`

For example, to specify that a command file is configured for the first UPS Setup, the perform an SNMP SET on `.pcnsAdvancedShutdownCommandFileEnabled.1`

Object Identifier Name	Access	Description
pcnsAdvancedGroupNo	read-only	The number of the Advanced UPS Setup.
pcnsAdvancedGroupName	read-write	The name of the Advanced UPS Setup.
pcnsAdvancedGroupNMC	read-only	A list of IP addresses of the NMCs in the Advanced UPS Setup, and the outlets in the associated outlet groups.
pcnsAdvancedShutdownUPSRequiredForLoad	read-write	The number of UPS devices required to power the load.
pcnsAdvancedShutdownCommandFileEnabled	read-write	Specify if a Shutdown Command File is configured to run if a UPS critical event is triggered. See Shutdown Command Files .
pcnsAdvancedShutdownCommandFile	read-write	The full path name of the command file, including the disk drive or volume name. See Shutdown Command Files .
pcnsAdvancedShutdownCommandFileDuration	read-write	The number of seconds that the shutdown command file requires to execute.
pcnsAdvancedShutdownPowerchuteServer	read-write	The setting to gracefully shut down the physical machine running PowerChute. See Shutdown Settings for Advanced UPS Configurations .
pcnsAdvancedShutdownOnLostRedundancy	read-write	If this option is enabled, when the number of UPS critical events is the same as the number of additional (redundant) UPS's, a shutdown sequence will be triggered. See Shutdown Settings for Advanced UPS Configurations .

pcnsAdvancedTurnOffUps	read-write	The setting to turn off the UPS when performing a graceful shutdown.
pcnsAdvancedTurnOffSOG	read-write	The setting to turn off the Outlet Group of the UPS when performing a graceful shutdown.

PowerChute Events

The table below details the OID Names of the [Configurable Events](#) that can be configured via SNMP, and the names of the Events as seen in the PowerChute User Interface.

Object Identifier Name	PowerChute UI Reference
pcnsPowerFailed	UPS On Battery
pcnsPowerRestored	Input Power Restored
pcnsOverload	UPS Overloaded
pcnsOverloadSolved	UPS Overload Corrected
pcnsRunTimExceeded	Runtime exceeded
pcnsRunTimeWithinRange	Runtime is sufficient
pcnsRunTimeBelowThreshold	Runtime remaining below threshold
pcnsRunTimeAboveThreshold	Runtime remaining above threshold
pcnsBatteryDischarged	Battery Discharged
pcnsBatteryChargeInRange	Battery Recharged
pcnsFaultBypassEnabled	Bypass due to hardware error or overload
pcnsBypassEnabled	Maintenance Bypass
pcnsBypassManualEnabled	Bypass ended
pcnsBypassDisabled	Bypass ended
pcnsBypassContactorFailed	Bypass Switch failed
pcnsBypasContactorOk	Bypass Switch replaced

pcnsCommunicationLostOnBattery	Communication lost while on Battery
pcnsCommunicationLost	NMC cannot communicate with the UPS

Object Identifier Name	PowerChute UI Reference
pcnsNetCommunicationLost	PowerChute cannot communicate with the NMC
pcnsCommunicationEstablished	Communication established
pcnsMinRedundancyLost	Minimum Redundancy lost
pcnsMinRedundancyRegained	Minimum Redundancy restored
pcnsParallelRedundancyLost	Parallel Redundancy lost
pcnsParallelRedundancyRegained	Parallel Redundancy restored
pcnsMaxInternalTempExceeded	UPS Temperature Overheated
pcnsMaxInternalTempInRange	UPS Temperature Normal Again
pcnsMinLoadCapabilityLost	Load (kVA) Alarm Violation
pcnsMinLoadCapabilityRegained	Load (kVA) Alarm Violation cleared
pcnsEnvironmentCommunicationEstablished	Communication Established with EMC
pcnsEnvironmentCommunicationLost	Communication Lost with EMC
pcnsTempInRangeProbeX	Temperature Probe X In Range
pcnsTempOutOfRangeProbeX	Temperature Probe X Out Of Range
pcnsHumidityInRangeProbeX	Humidity Probe X In Range
pcnsHumidityOutOfRangeProbeX	Humidity Probe X Out Of Range
pcnsContactFaultX	Contact Zone X Alarm
pcnsContactNormalX	Contact Zone X Normal

For each event you can access:

OID Name	Access	Description
[event name]Desc	read-only	The description of the event.
[event name]EnableLogging	read-write	Enable or disable logging to the event log for this event.
[event name]EnableCommandFile	read-write	Enable or disable command file execution for this event.
[event name]CommandFilePath	read-write	Specify a Command File and full path to be executed upon this event.
[event name]CommandFileDelay	read-write	The number of seconds that a host requires to shut down, before the command file is executed. Note: This is not available for <code>pcnsRunTimeBelowThreshold</code> . See pcnsRunTimeBelowThresholdCommandFileThreshold .

For some events you can access:

OID Name	Access	Description
[event name]EnableShutdown	read-write	Perform a graceful shutdown of the host when this event occurs.
[event name]ShutdownDelay	read-write	The amount of time in seconds that PowerChute should wait before initiating the shutdown process.

For `pcnsRunTimeBelowThreshold` you can access:

OID Name	Access	Description
<code>pcnsRunTimeBelowThresholdCommandFileThreshold</code>	read-write	If Runtime Remaining falls below this threshold, the command file is executed.
<code>pcnsRunTimeBelowThresholdShutdownThreshold</code>	read-write	If Runtime Remaining falls below this threshold, a graceful shutdown of the host is initiated.

PowerChute Traps

The table below details the OID Names of the SNMP traps sent by PowerChute for critical and lost communication events.

UPS Critical Events

OID Name	Level	Description
pcnsCriticalEventActive	Severe	PowerChute Network Shutdown has begun a graceful shutdown of the host due to a critical event.
pcnsCriticalEventResolved	Informational	The PowerChute Network Shutdown critical event has been resolved, and graceful shutdown of the host continues.

Lost Communication Events



OID Name	Level	Description
pcnsNetworkComsLost	Severe	PowerChute cannot communicate with the NMC.
pcnsNetworkComsLostResolved	Informational	PowerChute has regained communication with the NMC.
pcnsNMCCComsLost	Severe	The NMC cannot communicate with the UPS.
pcnsNMCCComsLostResolved	Informational	The NMC has regained communication with the UPS.

Other Events

OID Name	Level	Description
pcnsTest	Informational	PowerChute has sent a test trap to the NMS.
pcnsUpdateAvailable	Informational	PowerChute has detected that an update is available.

Event Configuration

When UPS events occur, PowerChute can be configured to log the event, notify users, execute a command file or initiate a system shutdown through the Configure Events screen.

The  symbol indicates that the action is enabled for this event while the  symbol indicates that the action is not enabled.



Descriptions of events are in the [PowerChute Events and Logging](#) sections.

Notifications

PowerChute can send a message to one user or all logged-in users when an event occurs:


Notify all users: For Linux or Unix, all users who are logged onto the server with a terminal prompt open will be notified.

Notify only this user: On Windows, enter the machine name. On Linux or Unix systems, enter the user name. The user will still need to be logged onto the server with a terminal prompt open to be notified.

Repeat Interval: The time interval, in seconds, at which the message will be repeated while the event condition exists. If this field is blank or zero, the message will not be repeated.

Delay (if required): Enter the amount of time in seconds that PowerChute should wait after the event occurs before notifying users. Users will be notified immediately if a shutdown event is triggered.

Event-Driven Command Files

If required, PowerChute can be configured to execute a command file after certain events are triggered. Click the  symbol on the event row and select the **Enable Command File** check box.

Delay: Enter the amount of time in seconds that PowerChute should wait when the event occurs before executing the command file.



If a shutdown command file is also configured, both command files will be executed in parallel.

Full path to command file: You must specify the full path name of the command file, including the disk drive or volume name.

The command file runs using the local system account. PowerChute cannot execute programs that require interaction with the desktop; only command line-enabled programs are supported.

The command file must be located in the `user_files` folder, or a sub-folder, in the PowerChute installation directory. If the default installation directory was chosen during installation, this location will be:

- `C:\Program Files\APC\PowerChute\user_files` for Windows systems
- `/opt/APC/PowerChute/user_files/` for Linux systems



NOTE: If you are upgrading to PowerChute v5.1, any command files used in a previous version of PowerChute must be manually added to the `user_files` folder. If the command files are not added to the directory, you will see the below errors/warnings for all events configured to run a command file in the Event Log. For example:

```
ERROR: Event PowerFailed is enabled for command file execution, but
an invalid value for event_PowerFailed_commandFilePath is
specified.
```

```
WARNING: Disabling command file execution for event PowerFailed due
to bad parameters. Please validate the configuration.
```

It is strongly recommended you validate your configuration after an upgrade.

Shutdown Actions

When the Shutdown Action is enabled for an event, PowerChute treats the event as critical and will trigger a shutdown sequence. Shutdown is not supported for all events: this is indicated by the presence or absence of an icon on the event row.

The **Delay** field is the amount of time in seconds that PowerChute should wait before initiating the shutdown sequence. By default, the **On Battery** event has a delay of 120 seconds, whereas the default for all other events is 0 seconds.

By default, PowerChute will only trigger a shutdown sequence if a low battery condition occurs or the UPS is commanded to turn off. Shutdown cannot be disabled for these events using the PowerChute user interface.

Start-up and Shut-down Scheduling

PowerChute does not directly support the start-up and shut-down scheduling feature as it is available from the Network Management Card. For more information on supported scheduling features, refer to the [NMC User Guide](#).

Sequenced Server Shutdown


The **Runtime Remaining below Threshold** event can be used to sequence the order that your servers shut down during an extended power outage.

This is useful if you have multiple servers powered by the same UPS and you want to extend the runtime for your higher priority servers. It also ensures that lower priority servers are the first to be shut down.

This event will trigger a server shutdown command when the UPS is running on battery power and the runtime has dropped below the threshold configured. You can also configure a command file to execute before shutdown occurs by specifying a higher runtime threshold value for the **Run Command File** event action.

Example

1. You have 3 servers powered by the same UPS. Your lower priority server is Server C while you want to keep Server A running as long as possible.
2. You want Server A to shut down when the UPS protecting it has 10 minutes runtime remaining.
3. You want Server B to shut down when the UPS protecting it has 15 minutes runtime remaining.
4. You want Server C to shut down when the UPS protecting it has 20 minutes runtime remaining.
5. Configure each PowerChute Agent with the following threshold values:
 - Server A – 10 minutes
 - Server B - 15 minutes
 - Server C - 20 minutes
6. Each server is shut down when the runtime remaining drops below the threshold configured.



Sample Shutdown Scenarios

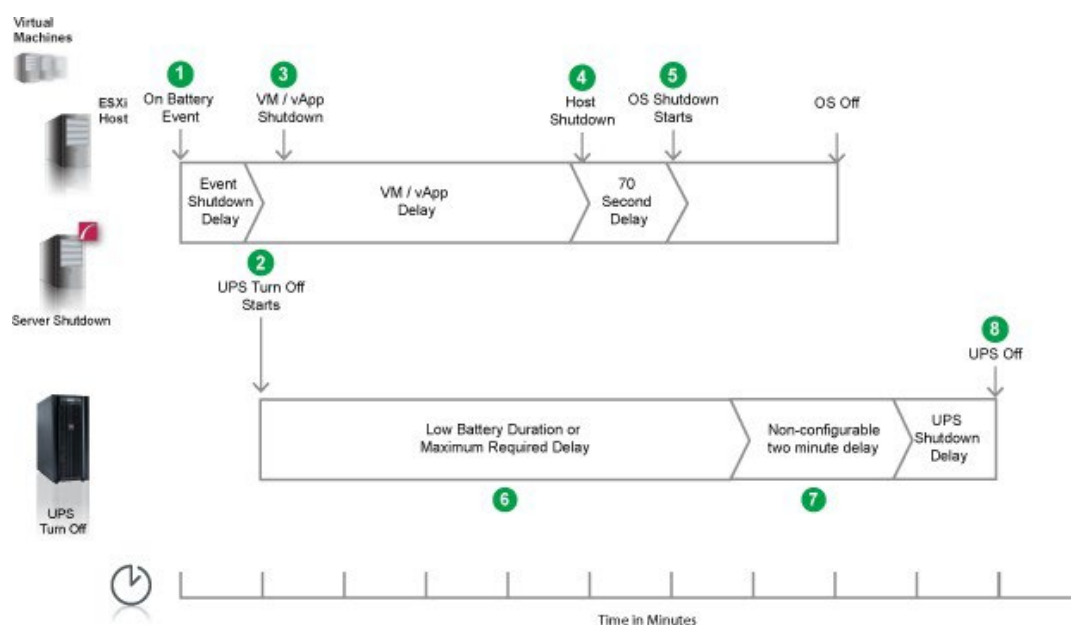
The following scenarios provide examples of how PowerChute and the UPS behave when a shutdown sequence is triggered:

VMware: UPS without Outlet Groups

Example 1: Turn off the UPS enabled, no shutdown command file configured.

PowerChute is installed on a physical machine outside the cluster, configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page. No shutdown command file is configured. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM/vApp Shutdown Duration = 240), VMware Hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. PowerChute issues commands to shut down the VMware hosts.
5. After a 70 second delay the operating system on the physical machine running PowerChute starts to shut down.
6. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

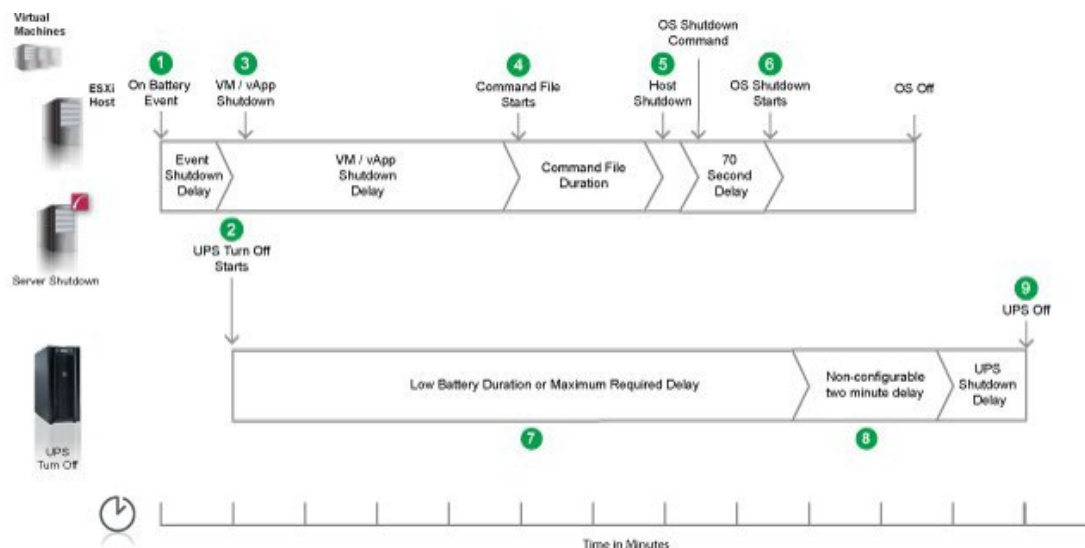
7. After this delay, a further non-configurable two minute delay is counted down.
8. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 7) starts to count down.

Example 2: Turn off the UPS enabled, shutdown command file configured.

PowerChute is installed on a physical machine outside the cluster, and configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page. A shutdown command file is configured. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM/vApp Shutdown Duration = 240), PowerChute starts to execute the shutdown command file.
5. VMware hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. After the duration configured for the shutdown command file has elapsed, PowerChute issues commands to shut down the VMware hosts.
6. An additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
7. The UPS will wait the amount of time indicated by one of the following, whichever is greater:
Low Battery Duration or **Maximum Required Delay**.

These are shown on the Configuration - Shutdown page in the NMC interface.

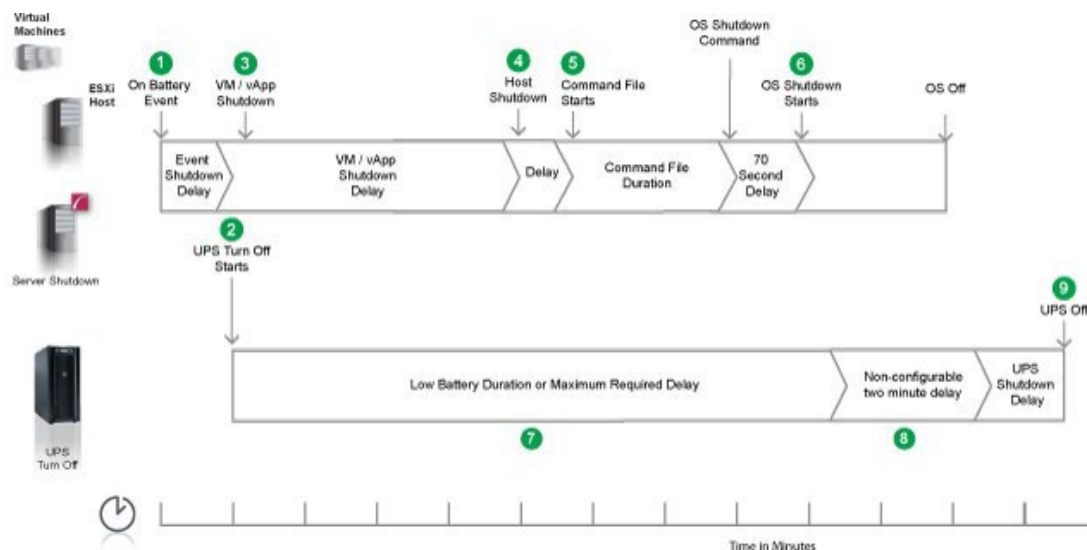
8. After this delay, a further non-configurable two minute delay is counted down.
9. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 8) starts to count down.

Example 3: Turn off the UPS enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled.

PowerChute is installed on a physical machine outside the cluster, and configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM/vApp Shutdown Duration = 240), VMware hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. PowerChute issues commands to shut down the VMware hosts.
5. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
6. After the duration configured for the shutdown command file has elapsed, the OS Shutdown Command is issued and an additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
7. The UPS will wait the amount of time indicated by one of the following, whichever is greater:
Low Battery Duration or Maximum Required Delay.

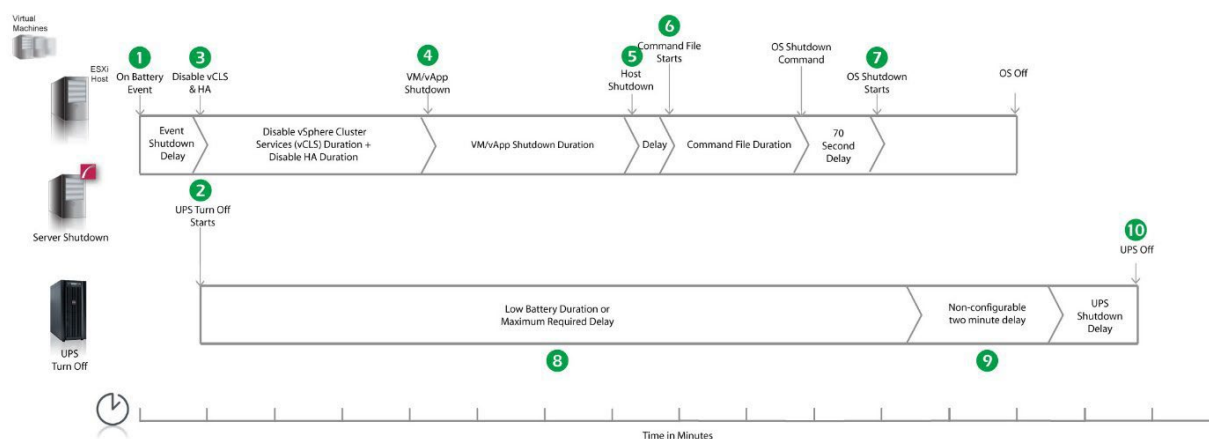
These are shown on the Configuration - Shutdown page in the NMC interface.

8. After this delay, a further non-configurable two minute delay is counted down.
9. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

Example 4: Turn off the UPS enabled, vCLS VMs in cluster, shutdown command file configured, Execute Command File after Host Shutdown enabled.

PowerChute is installed on a physical machine outside the cluster, and configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured. vCLS VMs are detected in the cluster, and the Disable vSphere Cluster Services (vCLS) Duration is 180 seconds. The Disable HA on Shutdown checkbox is selected with a duration of 20 seconds.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute disables vSphere Cluster Services (vCLS) and High Availability (HA) on the cluster.
4. After 3 minutes and 20 seconds (Disable vSphere Cluster Services (vCLS) Duration = 180, Disable HA Duration = 20), PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
5. After 4 minutes (VM/vApp Shutdown Duration = 240), VMware hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. PowerChute issues commands to shut down the VMware hosts.
6. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
7. After the duration configured for the shutdown command file has elapsed, the OS Shutdown Command is issued and an additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:
Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

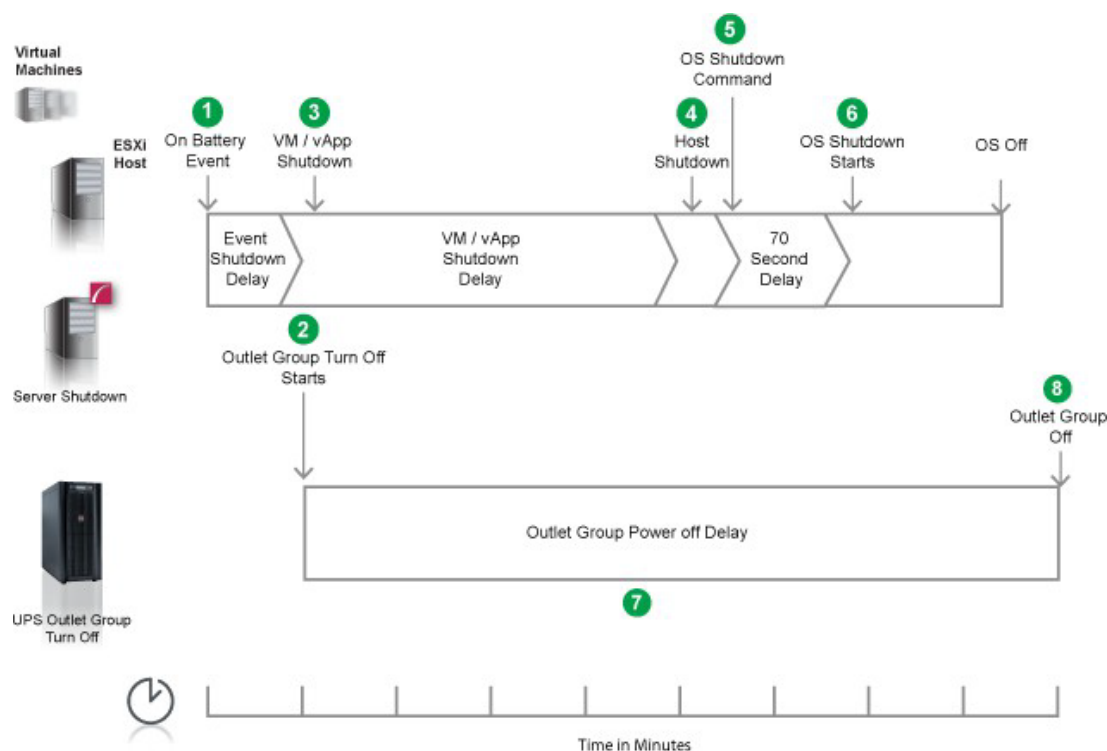
9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface

VMware: UPS with Outlet Groups

Example 1: Turn off the Outlet Group enabled, no shutdown command file configured.

PowerChute is installed on a physical machine outside the cluster, configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the Outlet Group is enabled on the Shutdown settings page. No Shutdown command file is configured. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



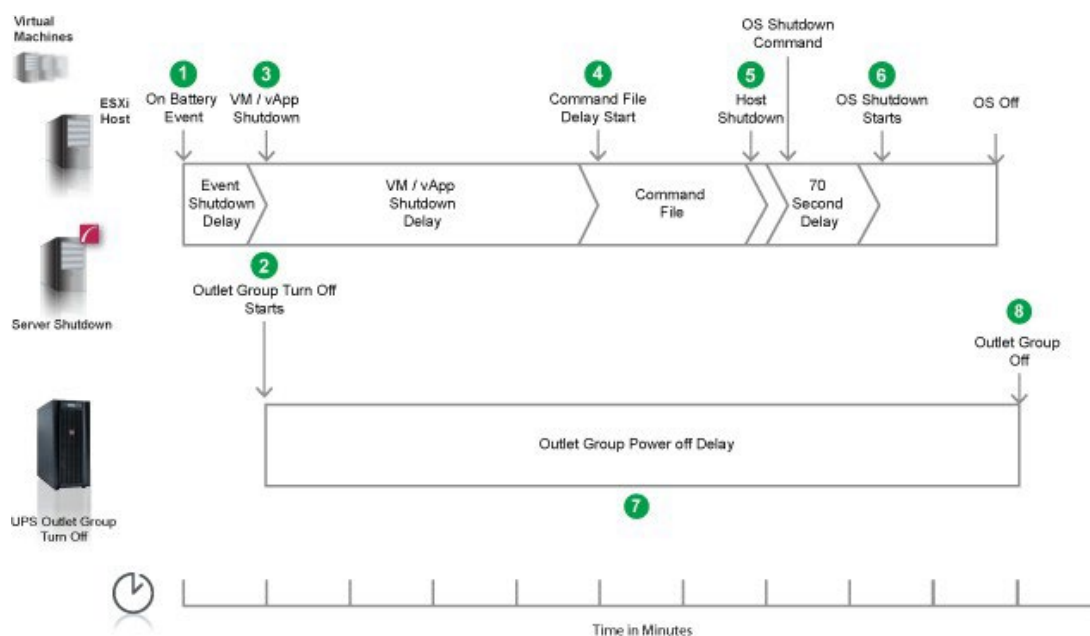
1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS outlet group and the outlet group turn off starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM/vApp Shutdown Duration = 240), VMware Hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. PowerChute issues commands to shut down the VMware hosts.
5. PowerChute issues the operating system shutdown command.
6. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
7. The outlet group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

It is recommended that the outlet group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the outlet group does not turn off before the operating system.

Example 2: Turn off the Outlet Group enabled, shutdown command file configured.

PowerChute is installed on a physical machine outside the cluster, configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the Outlet Group is enabled on the Shutdown settings page. A shutdown command file is configured. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



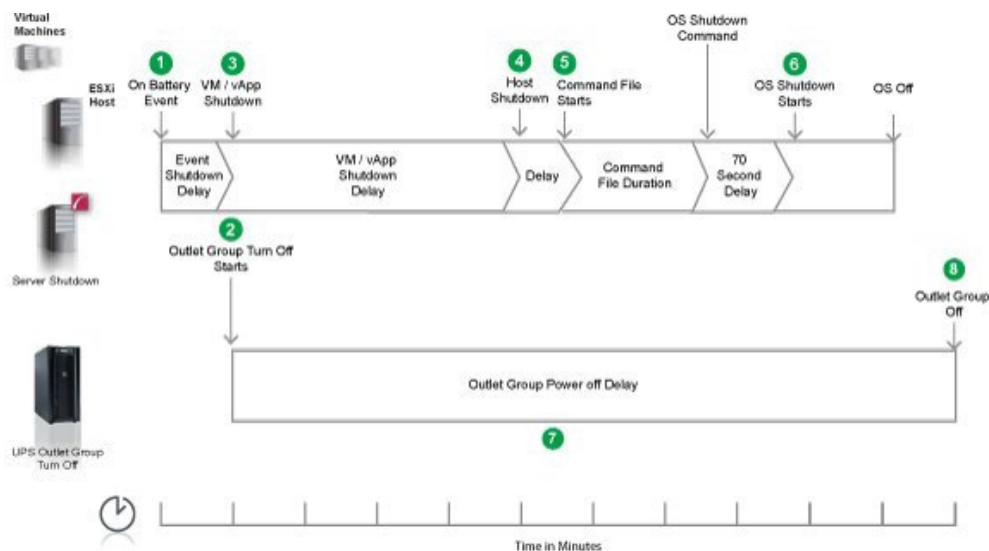
1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the outlet group and the outlet group turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM/vApp Shutdown Duration = 240), PowerChute starts to execute the shutdown command file.
5. VMware Hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. After the duration configured for the shutdown command file has elapsed, PowerChute issues commands to shut down the VMware hosts.
6. An additional 70 second delay is counted down before the operating system starts to shut down.
7. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet groups to turn off before the Main Outlet Group turn off starts.
 - If registered with a Switched Outlet Group only that delay is counted down.

It is recommended that the outlet group **Power Off delay** is configured to allow enough time for the shutdown command file and the operating system shutdown to complete. You should allow extra time to ensure that the outlet group does not turn off before the operating system.

Example 3: Turn off the Outlet Group enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled.

PowerChute is installed on a physical machine outside the cluster, configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the Outlet Group is enabled on the Shutdown settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



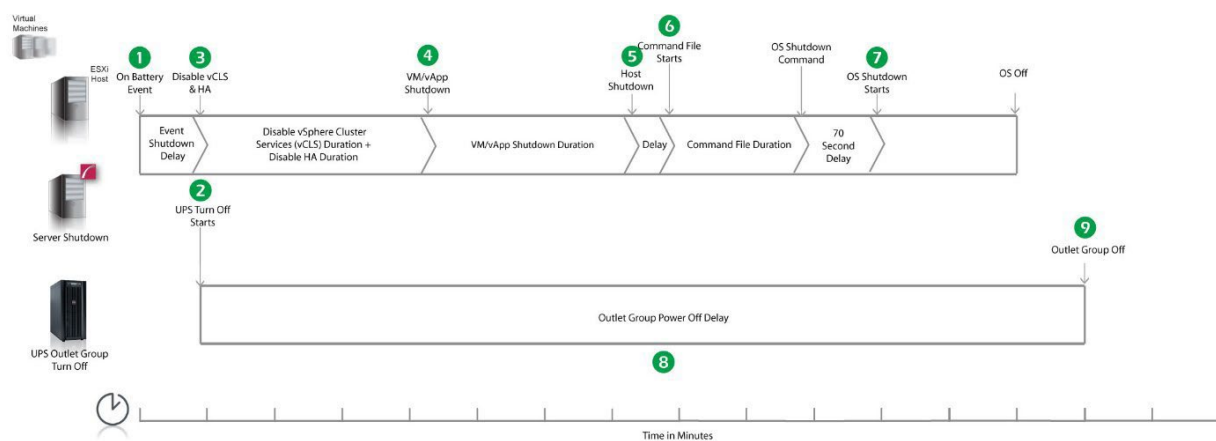
1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the outlet group and the outlet group turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM/vApp Shutdown Duration = 240), VMware Hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. PowerChute issues commands to shut down the VMware hosts.
5. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
6. After the duration configured for the shutdown command file has elapsed, an additional 70 second delay is counted down before the operating system starts to shut down.
7. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet groups to turn off before the Main Outlet Group turn off starts.
 - If registered with a Switched Outlet Group only that delay is counted down.

It is recommended that the outlet group **Power Off delay** is configured to allow enough time for the shutdown command file and the operating system shutdown to complete. You should allow extra time to ensure that the outlet group does not turn off before the operating system.

Example 4: Turn off the UPS enabled, vCLS VMs in cluster, shutdown command file configured, Execute Command File after Host Shutdown enabled.

PowerChute is installed on a physical machine outside the cluster, configured for a Single/Redundant UPS configuration with several VMware Hosts in a HA cluster. The option to Turn off the Outlet Group is enabled on the Shutdown settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. Host Maintenance Mode is disabled. VM/ vApp Shutdown is enabled with 240 second delay configured. vCLS VMs are detected in the cluster, and the Disable vSphere Cluster Services (vCLS) Duration is 180 seconds. The Disable HA on Shutdown checkbox is selected with a duration of 20 seconds.

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute disables vSphere Cluster Services (vCLS) and High Availability (HA) on the cluster.
4. After 3 minutes and 20 seconds (Disable vSphere Cluster Services (vCLS) Duration = 180, Disable HA Duration = 20), PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
5. After 4 minutes (VM/vApp Shutdown Duration = 240), VMware hosts enter Maintenance mode if all VMs are powered off, otherwise the Maintenance mode task is cancelled. PowerChute issues commands to shut down the VMware hosts.
6. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
7. After the duration configured for the shutdown command file has elapsed, an additional 70 second delay is counted down before the operating system starts to shut down.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet groups to turn off before the Main Outlet Group turn off starts.
 - If registered with a Switched Outlet Group only that delay is counted down.

It is recommended that the outlet group **Power Off Delay** is configured to allow enough time for the shutdown command file and the operating system shutdown to complete. You should allow extra time to ensure that the outlet group does not turn off before the operating system.

Recommended Power-Off Delays for Outlet Groups

By default, the outlet group **Power Off Delay** will be the same value as the Low Battery duration configured on the NMC. PowerChute will automatically increase the Power Off Delay for the outlet group it is registered with, if the total shutdown time it needs is greater than the Power Off Delay.

The total shutdown time includes the following values:

- VM Migration delay
- VM Shutdown and Startup Delays
- vApp Shutdown and Startup Delays
- Delay Host Maintenance Mode
- Disable vSphere Cluster Services (vCLS) Duration
- vSAN Recovery Duration
- Disable HA Duration
- vSAN Synchronization Duration (for vSAN configurations)
- Execute Command File after Host Shutdown delay
- Shutdown Command File Duration
- SSH Action Duration
- Built-in delay of 2 minutes (this consists of a 10 second OS shutdown delay and a 60 second OS shutdown duration; rounded up)

NOTE: Delay Host Maintenance Mode



In a non-Advanced configuration, the total will equal the number of protected hosts multiplied by the **Delay Host Maintenance Mode** value.

In an Advanced configuration, the total will equal the number of protected hosts multiplied by the **Delay Host Maintenance Mode** value multiplied by the number of groups containing hosts.



NOTE: vSAN Synchronization Duration

In an Advanced configuration, the total will equal the **vSAN Synchronization Duration** value multiplied by the number of groups containing hosts.

The time required to gracefully shut down your operating system is not covered by the total shutdown time, as PowerChute cannot determine how long it will take to complete.



The Power Off Delay for the outlet group should be long enough for the OS to gracefully shut down. You should add extra time to allow for unforeseen circumstances.

The Low Battery Duration set on the NMC should be equal to or greater than the Power Off Delay for the outlet group.

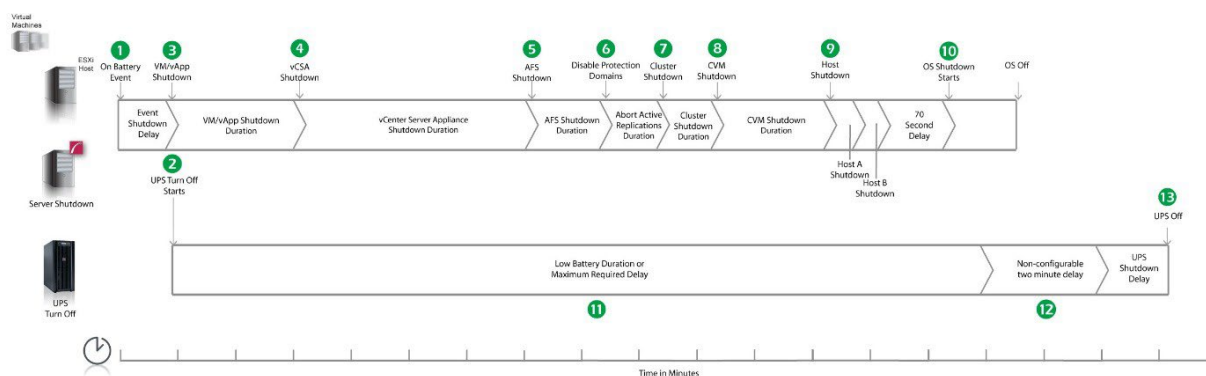
VMware with Nutanix Support: UPS without Outlet Groups

Example 1: Turn off the UPS enabled, VM migration disabled, abort active replications enabled, no shutdown command file or SSH action configured.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. [VM Prioritization](#) is enabled and the vCenter Server Appliance VM is added to a priority group. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- AFS Shutdown Duration = 60 seconds
- Abort Active Replications Duration = 80 seconds
- Cluster Shutdown Duration = 60 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute shuts down Acropolis File Services and the AFS VMs.
6. After 60 seconds (AFS Shutdown Delay), PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
7. After 80 seconds (Abort Active Replications Duration), PowerChute shuts down the Nutanix Cluster.
8. After 60 seconds (Cluster Shutdown Delay), PowerChute issues a command to shut down the Controller VMs.
9. After 120 seconds (Controller VM Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
10. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
11. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

12. After this delay, a further non-configurable two minute delay is counted down.
13. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.
This is configurable on the Configuration - Shutdown page in the NMC user interface.

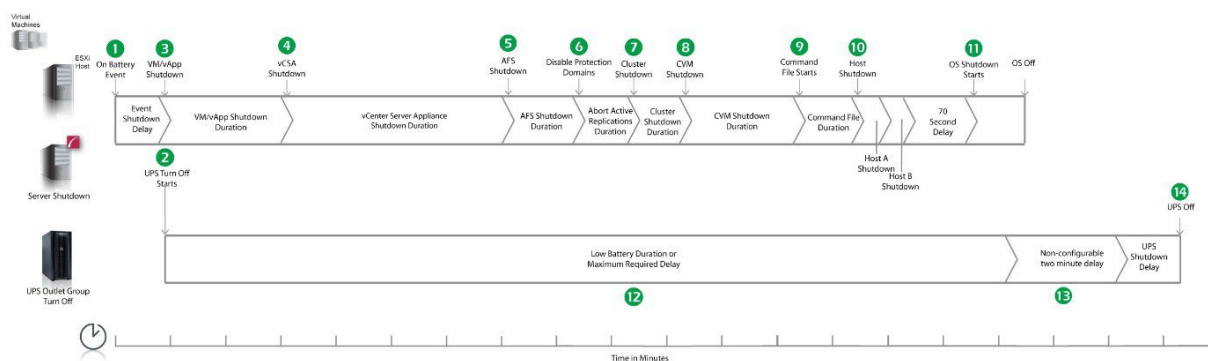
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 12) starts to count down.

Example 2: Turn off the UPS enabled, VM migration disabled, ongoing replication abort delay enabled, shutdown command file configured.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. [VM Prioritization](#) is enabled and the vCenter Server Appliance VM is added to a priority group. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- AFS Shutdown Delay = 60 seconds
- Abort Active Replications Duration = 60 seconds
- Cluster Shutdown Duration = 60 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down Acropolis File Services and the AFS VMs.
6. After 60 seconds (AFS Shutdown Duration), PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
7. After 80 seconds (Abort Ongoing Replications Duration), PowerChute shuts down the Nutanix Cluster.

8. After 60 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
9. After 120 seconds (Controller VM Shutdown Delay), PowerChute starts to execute the shutdown command file.
10. After the delay configured for the command file has elapsed, the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
11. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
12. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

13. After this delay, a further non-configurable two minute delay is counted down.
14. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

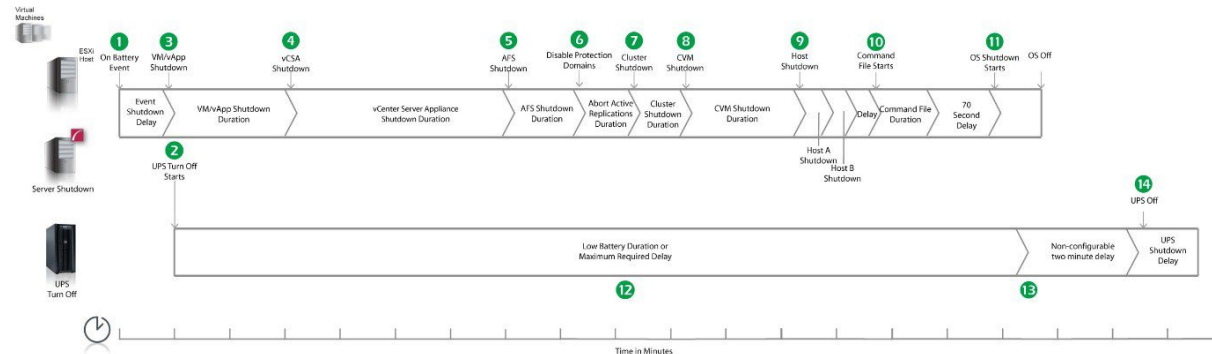
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 13) starts to count down.

Example 3: Turn off the UPS enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. **VM Prioritization** is enabled and the vCenter Sever Appliance VM is added to a priority group. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- AFS Shutdown Duration = 60 seconds
- Abort Active Replications Duration = 80 seconds
- Cluster Shutdown Duration = 60 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down Acropolis File Services and the AFS VMs.
6. After 60 seconds (AFS Shutdown Duration), PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
7. After 80 seconds (Abort Active Replications Duration), PowerChute shuts down the Nutanix Cluster.
8. After 60 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down Controller VMs.
9. After 120 seconds (Controller VM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
10. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
11. After the duration configured for the shutdown command file has elapsed, the OS Shutdown Command is issued and an additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
12. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

13. After this delay, a further non-configurable two minute delay is counted down.
14. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

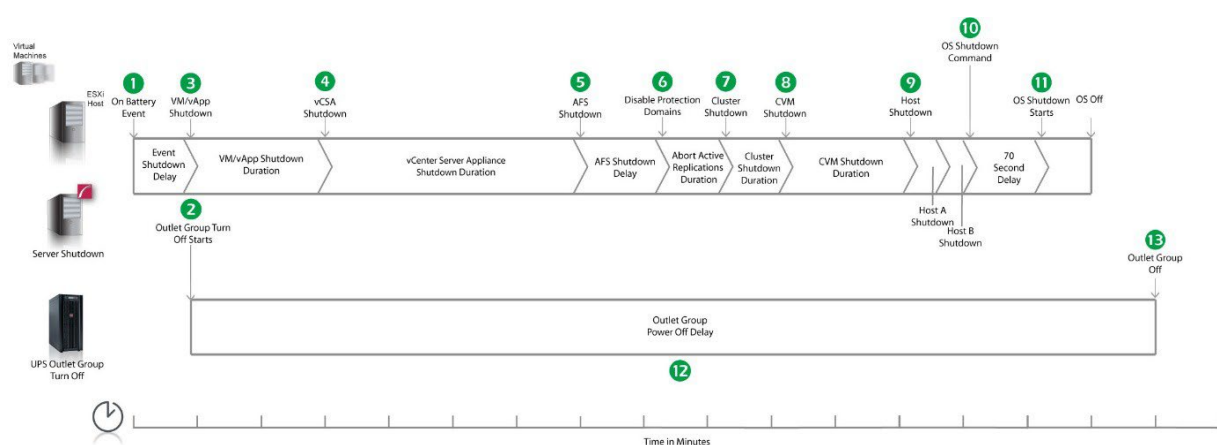
VMware with Nutanix Support: UPS with Outlet Groups

Example 1: Turn off the Outlet Group enabled, AFS shutdown enabled, abort active replications enabled, no shutdown command file or SSH action configured.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. [VM Prioritization](#) is enabled and the vCenter Server Appliance VM is added to a priority group. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- AFS Shutdown Duration = 60 seconds
- Abort Ongoing Replications Duration = 80 seconds
- Cluster Shutdown Duration = 60 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down Acropolis File Services and the AFS VMs.
6. After 60 seconds (AFS Shutdown Duration), PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
7. After 80 seconds (Abort Active Replications Duration), PowerChute shuts down the Nutanix Cluster.
8. After 60 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
9. After 120 seconds (Controller VM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).

10. PowerChute issues the operating system shutdown command.
11. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
12. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

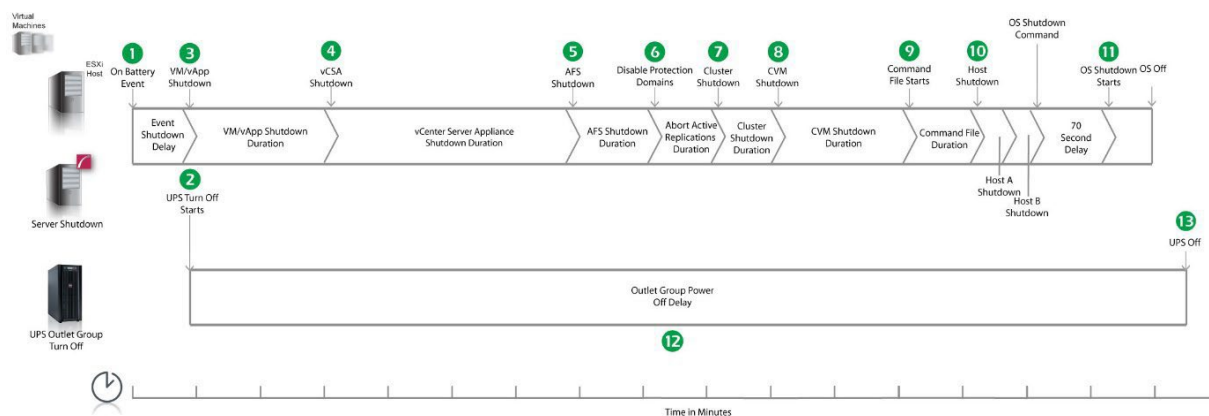
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 2: Turn off the Outlet Group enabled, shutdown command file configured.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. **VM Prioritization** is enabled and the vCenter Sever Appliance VM is added to a priority group. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- AFS Shutdown Duration = 60 seconds
- Abort Ongoing Replications Duration = 80 seconds
- Cluster Shutdown Duration = 60 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down Acropolis File Services and the AFS VMs.
6. After 60 seconds (AFS Shutdown Duration), PowerChute aborts any ongoing VM replications. If enabled in the Protection

Domain Settings page, Metro Availability will also be disabled on your Cluster.

7. After 80 seconds (Abort Active Replications Duration), PowerChute shuts down the Nutanix Cluster.
8. After 60 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
9. After 120 seconds (Controller VM Shutdown Duration), PowerChute starts to execute the shutdown command file.
10. The VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
11. An additional 70 second delay is counted down before the operating system starts to shut down.
12. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

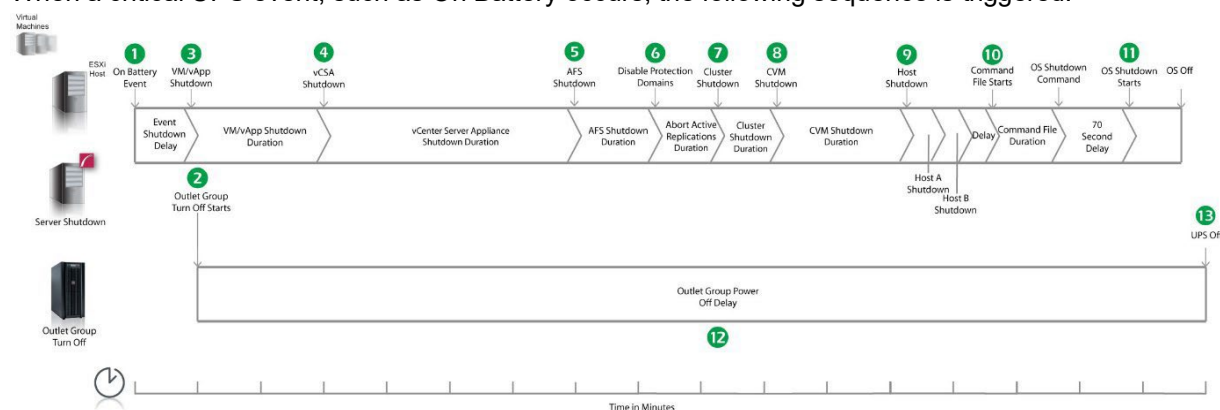
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 3: Turn off the Outlet Group enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. [VM Prioritization](#) is enabled and the vCenter Server Appliance VM is added to a priority group. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- AFS Shutdown Duration = 60 seconds
- Abort Ongoing Replications Duration = 80 seconds
- Cluster Shutdown Duration = 60 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.

2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down Acropolis File Services and the AFS VMs.
6. After 60 seconds (AFS Shutdown Duration), PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
7. After 80 seconds (Abort Active Replications Duration), PowerChute shuts down the Nutanix Cluster.
8. After 60 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
9. After 120 seconds (Controller VM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
10. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
11. After the duration configured for the shutdown command file has elapsed, an additional 70 second delay is counted down before the operating system starts to shut down.
12. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Recommended Power-Off Delays for Outlet Groups

By default, the outlet group **Power Off Delay** will be the same value as the Low Battery duration configured on the NMC. PowerChute will automatically increase the Power Off Delay for the outlet group it is registered with, if the total shutdown time it needs is greater than the Power Off Delay.

The total shutdown time includes the following values:

- VM Migration Duration
- VM/vApp Shutdown and Startup Duration
- Delay Host Maintenance Mode
- AFS Shutdown and Startup Duration
- Abort Active Replications Duration
- Cluster Shutdown and Startup Duration
- Controller VM Shutdown and Startup Duration
- Execute Command File after Host Shutdown delay
- Shutdown Command File Duration
- SSH Action Duration
- Built-in delay of 2 minutes (this consists of a 10 second OS shutdown delay and a 60 second OS shutdown duration; rounded up)

The time required to gracefully shut down your operating system is not covered by the total shutdown time, as PowerChute cannot determine how long it will take to complete.



The Power Off Delay for the outlet group should be long enough for the OS to gracefully shut down. You should add extra time to allow for unforeseen circumstances.

The Low Battery Duration set on the NMC should be equal to or greater than the Power Off Delay for the outlet group.

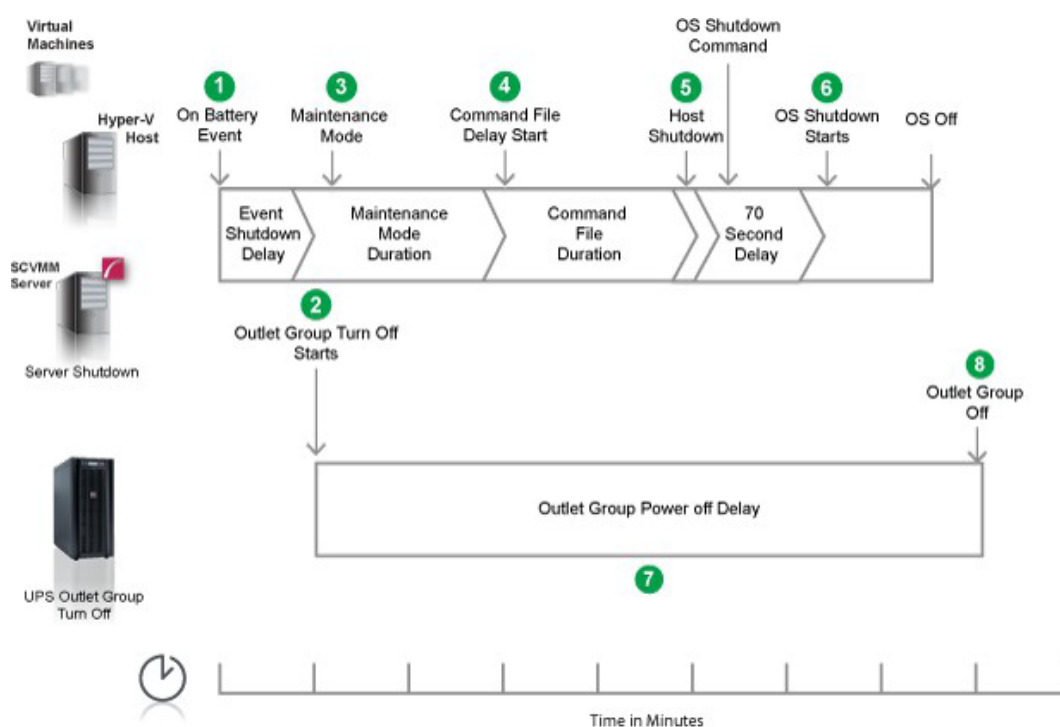
VMware with HPE SimpliVity Support: UPS without Outlet Groups

Example 1a: Turn off the UPS enabled, no shutdown command file or SSH action configured, external vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
5. After 120 seconds (OVC Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
6. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.

7. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

8. After this delay, a further non-configurable two minute delay is counted down.

9. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.

This is configurable on the Configuration - Shutdown page in the NMC user interface.

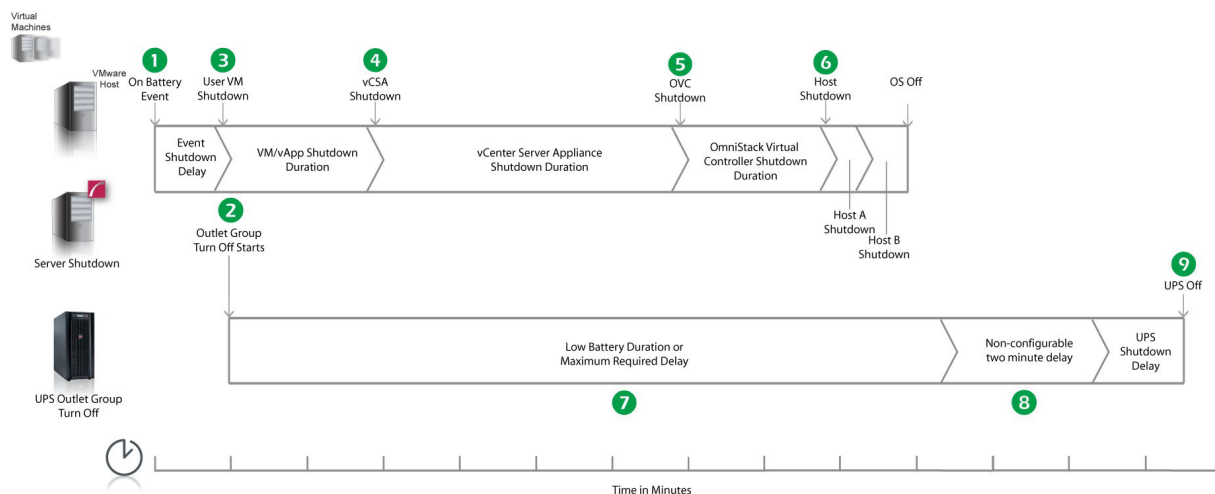
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 8) starts to count down.

Example 1b: Turn off the UPS enabled, no shutdown command file or SSH action configured, internal vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute issues a command to shut down the OmniStack Virtual Controller VMs.
6. After 120 seconds (OVC Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if

all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.

7. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay

These are shown on the Configuration - Shutdown page in the NMC interface.

8. After this delay, a further non-configurable two minute delay is counted down.

9. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.

This is configurable on the Configuration - Shutdown page in the NMC user interface.

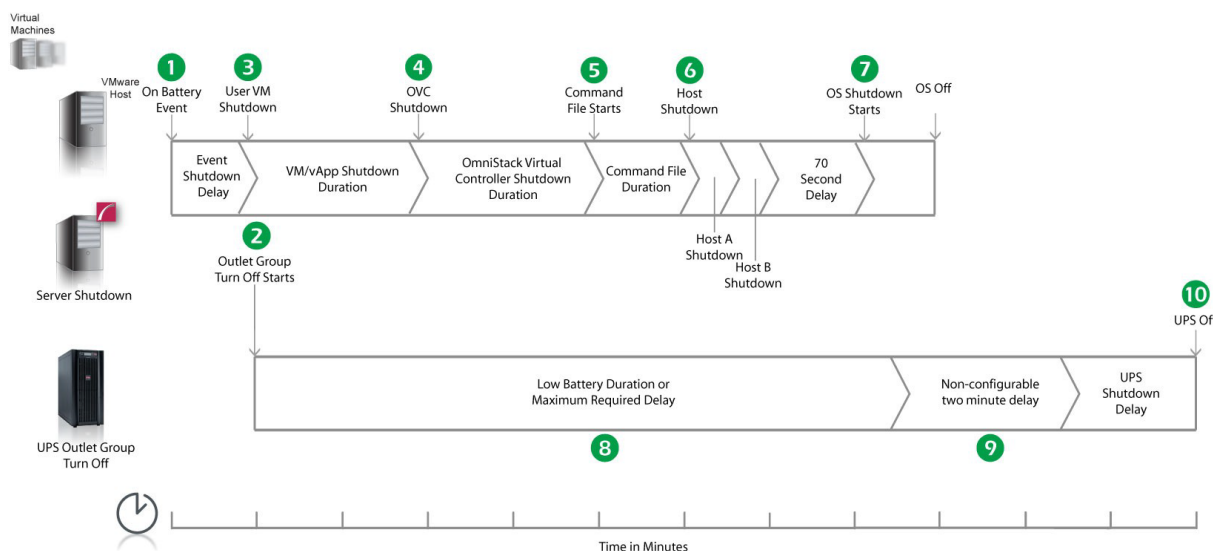
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 2a: Turn off the UPS enabled, shutdown command file configured, external vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
5. After 120 seconds (OVC Shutdown Delay), PowerChute starts to execute the shutdown command file.
6. After the delay configured for the command file has elapsed, the VMware Hosts enter maintenance mode and are shut

down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).

7. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

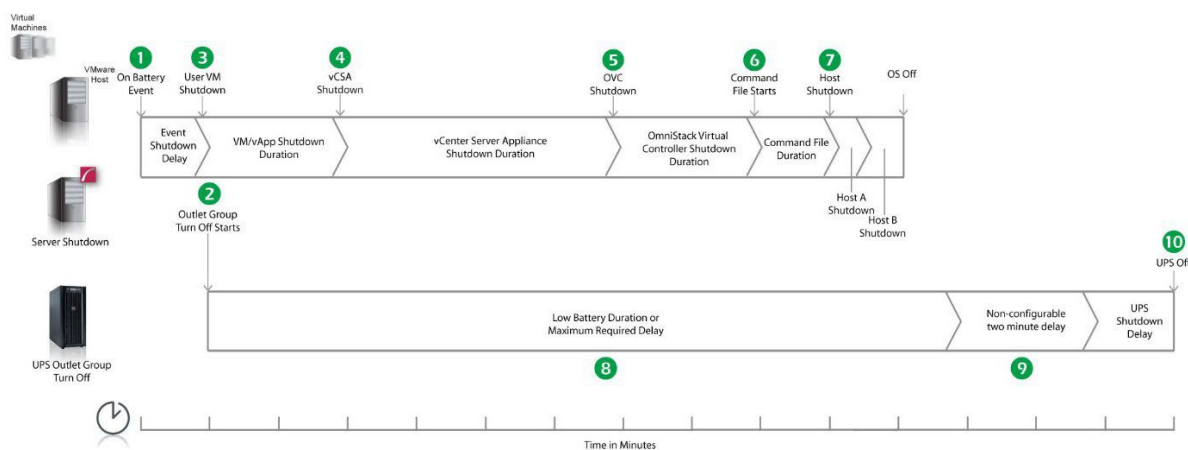
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 9) starts to count down.

Example 2b: Turn off the UPS enabled, shutdown command file configured, internal vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is installed as a VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.

6. After 120 seconds (OVC Shutdown Delay), PowerChute starts to execute the shutdown command file.
7. After the delay configured for the command file has elapsed, the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

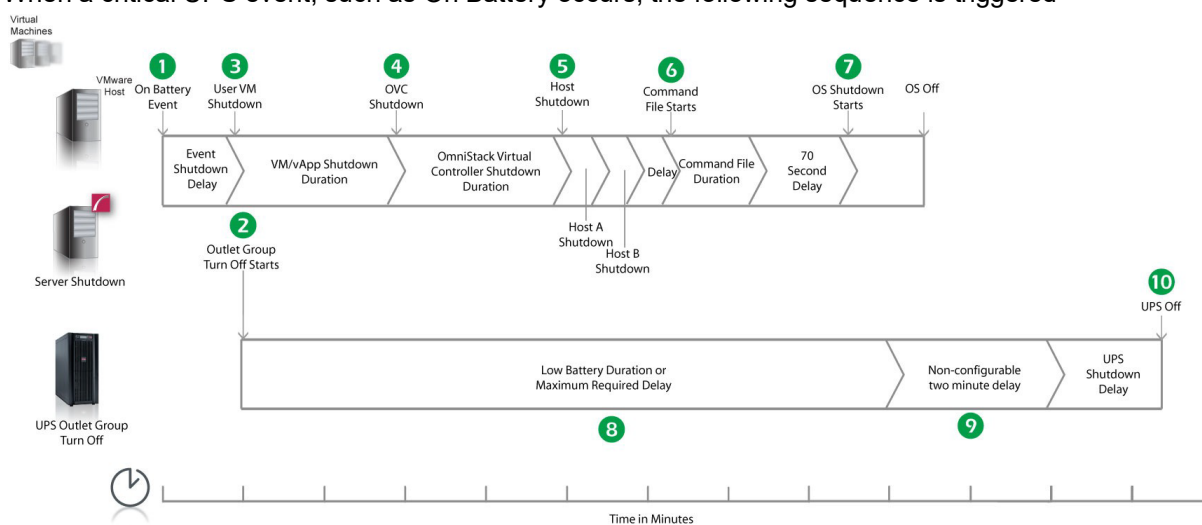
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 3a: Turn off the UPS enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled, external vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.

4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
5. After 120 seconds (OVC Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
6. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
7. After the duration configured for the shutdown command file has elapsed, the OS Shutdown Command is issued and an additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

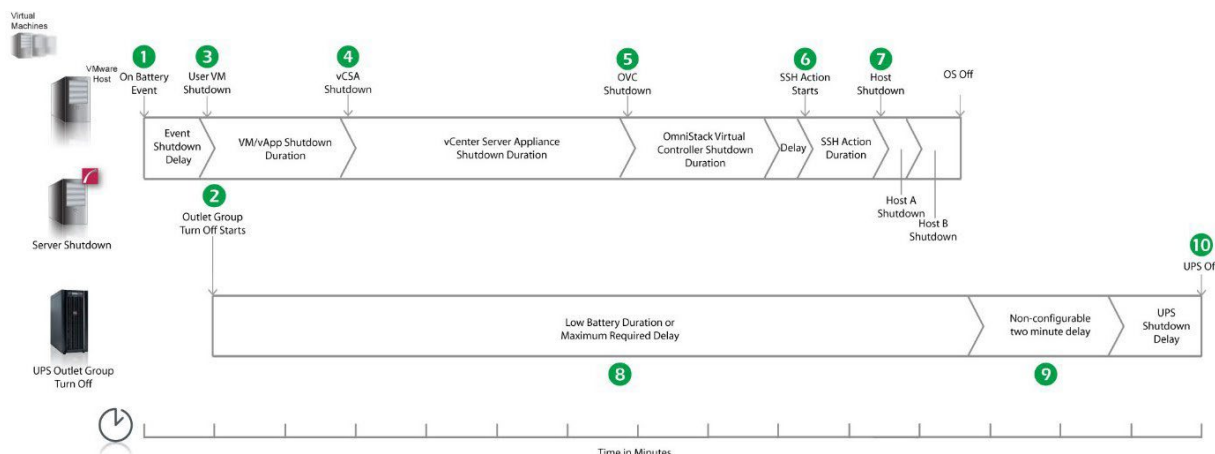
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 9) starts to count down.

Example 3b: Turn off the UPS enabled, SSH action configured, internal vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is installed as a VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
6. After 120 seconds (OVC Shutdown Delay), PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
7. VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or **Maximum Required Delay**.

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.
This is configurable on the Configuration - Shutdown page in the NMC user interface.

It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

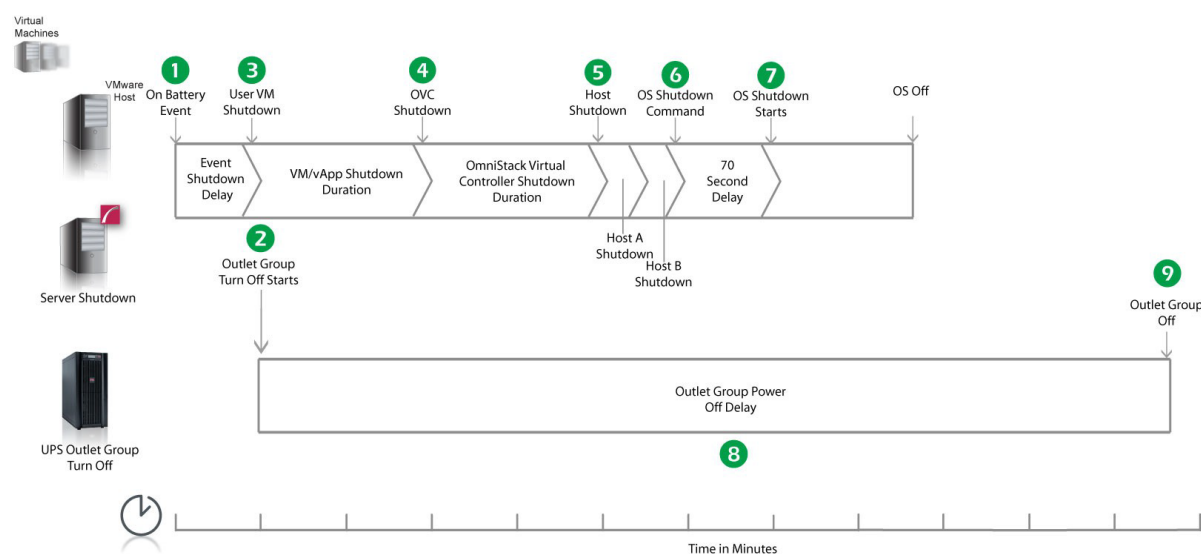
VMware with HPE SimpliVity Support: UPS with Outlet Groups

Example 1a: Turn off the Outlet Group enabled, no shutdown command file or SSH action configured, external vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
5. After 120 seconds (OVC Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
6. PowerChute issues the operating system shutdown command.
7. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

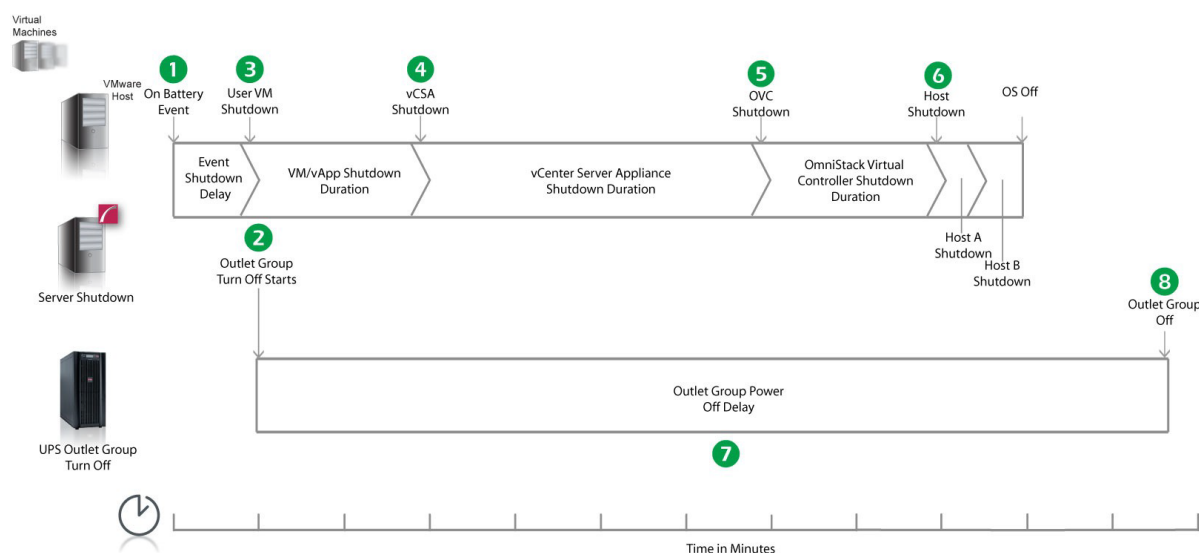
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 1b: Turn off the Outlet Group enabled, no shutdown command file or SSH action configured, internal vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is installed as a VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



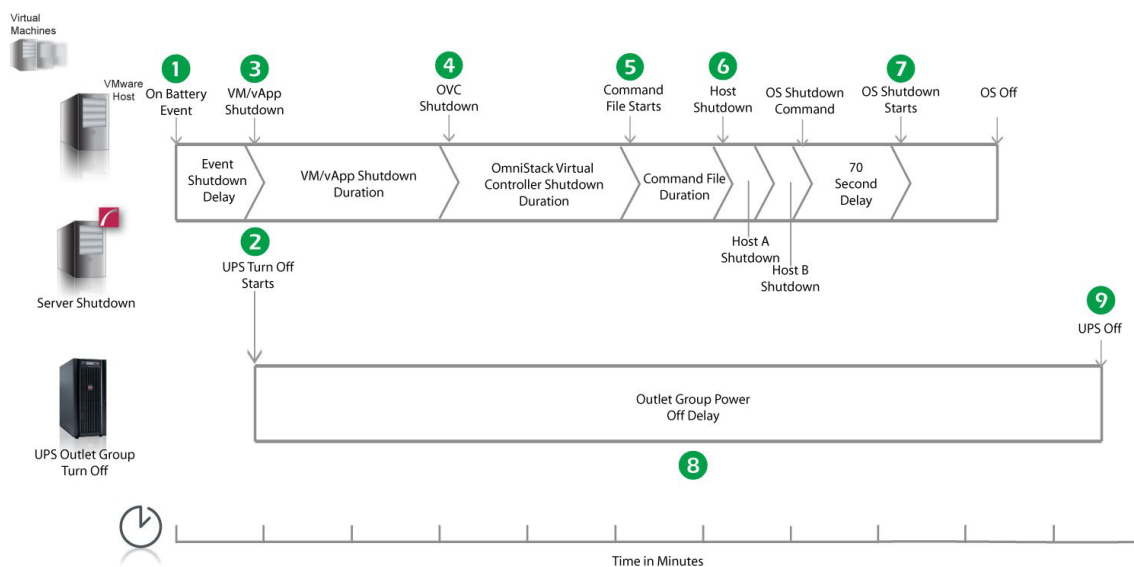
1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
6. After 120 seconds (OVC Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
7. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 2a: Turn off the Outlet Group enabled, shutdown command file configured, external vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
5. After 120 seconds (OVC Shutdown Delay), PowerChute starts to execute the shutdown command file.
6. The VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
7. An additional 70 second delay is counted down before the operating system starts to shut down.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

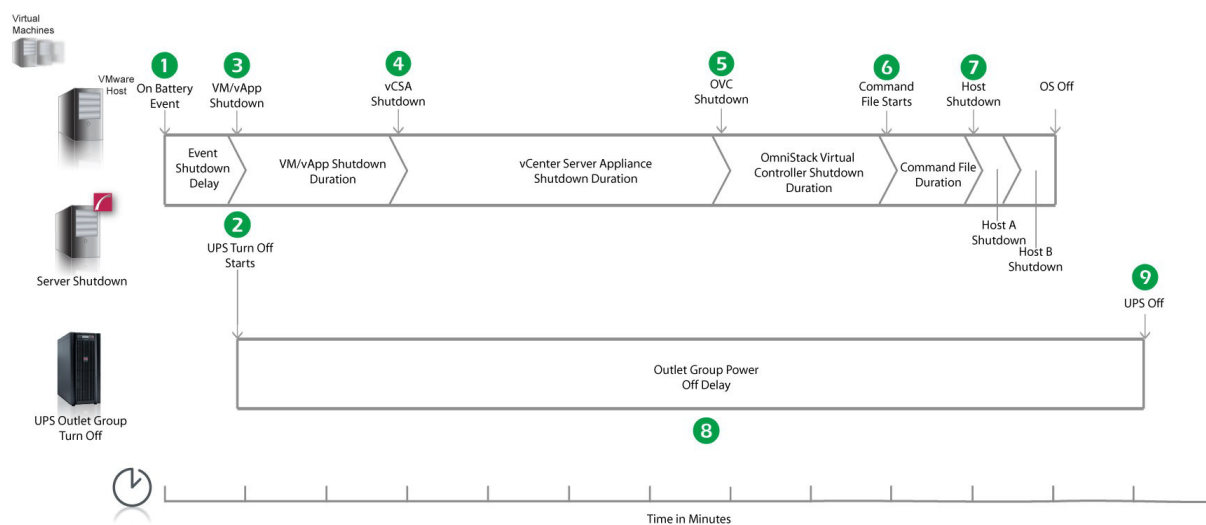
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 2b: Turn off the Outlet Group enabled, shutdown command file configured, internal vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is installed as a VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



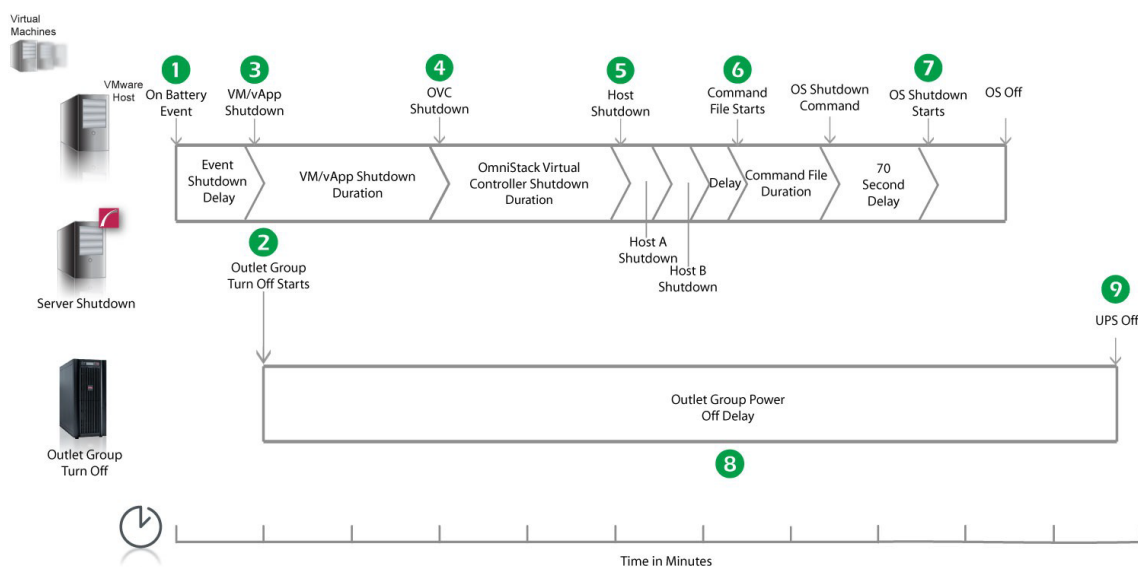
1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
6. After 120 seconds (OVC Shutdown Delay), PowerChute starts to execute the shutdown command file.
7. The VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 3a: Turn off the Outlet Group enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled, external vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
5. After 120 seconds (OVC Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
6. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
7. After the duration configured for the shutdown command file has elapsed, an additional 70 second delay is counted down before the operating system starts to shut down.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.

- If registered with a Switched Outlet Group, only that delay is counted down.

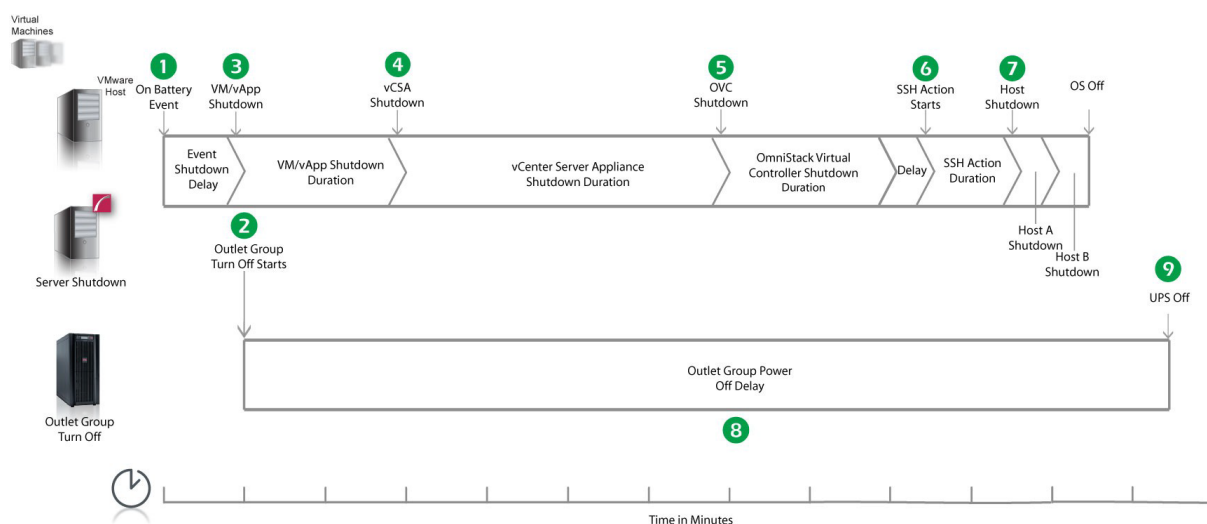
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 3b: Turn off the Outlet Group enabled, SSH action configured, internal vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is installed as a VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- OmniStack Virtual Controller Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute shuts down the HPE SimpliVity OmniStack Virtual Controller.
6. After 120 seconds (OVC Shutdown Delay), PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
7. The VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.

- If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
- If registered with a Switched Outlet Group, only that delay is counted down.

Recommended Power-Off Delays for Outlet Groups

By default, the outlet group **Power Off Delay** will be the same value as the Low Battery duration configured on the NMC. PowerChute will automatically increase the Power Off Delay for the outlet group it is registered with, if the total shutdown time it needs is greater than the Power Off Delay.

The total shutdown time includes the following values:

- VM/vApp Shutdown and Startup Duration
- vCenter Server Appliance Shutdown Duration
- OmniStack Virtual Controller Shutdown and Startup Duration
- Delay Host Maintenance Mode Duration
- Execute Command File after Host Shutdown delay
- Shutdown Command File Duration
- SSH Action Duration
- Built-in delay of 2 minutes (this consists of a 10 second OS shutdown delay and a 60 second OS shutdown duration; rounded up)

The time required to gracefully shut down your operating system is not covered by the total shutdown time, as PowerChute cannot determine how long it will take to complete.



The Power Off Delay for the outlet group should be long enough for the OS to gracefully shut down. You should add extra time to allow for unforeseen circumstances

The Low Battery Duration set on the NMC should be equal to or greater than the Power Off Delay for the outlet group.

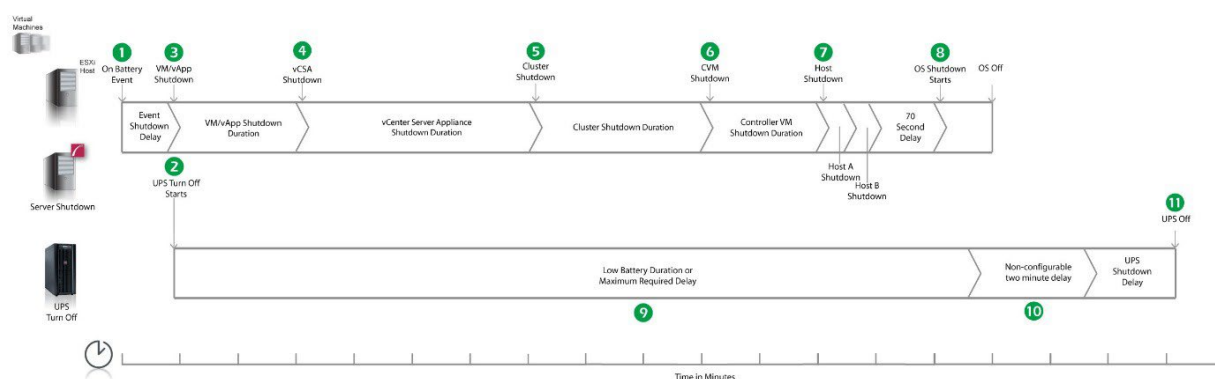
VMware with HyperFlex Support: UPS without Outlet Groups

Example 1a: Turn off the UPS enabled, no shutdown command file or SSH action configured, internal vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Delay), PowerChute shuts down the HyperFlex Cluster.
6. After 180 seconds (Cluster Shutdown Delay), PowerChute issues a command to shut down the Controller VMs.
7. After 120 seconds (Controller VM Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
8. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
9. The UPS will wait the amount of time indicated by one of the following, whichever is greater:
 - Low Battery Duration or Maximum Required Delay**
These are shown on the Configuration - Shutdown page in the NMC interface.
10. After this delay, a further non-configurable two minute delay is counted down.

- The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

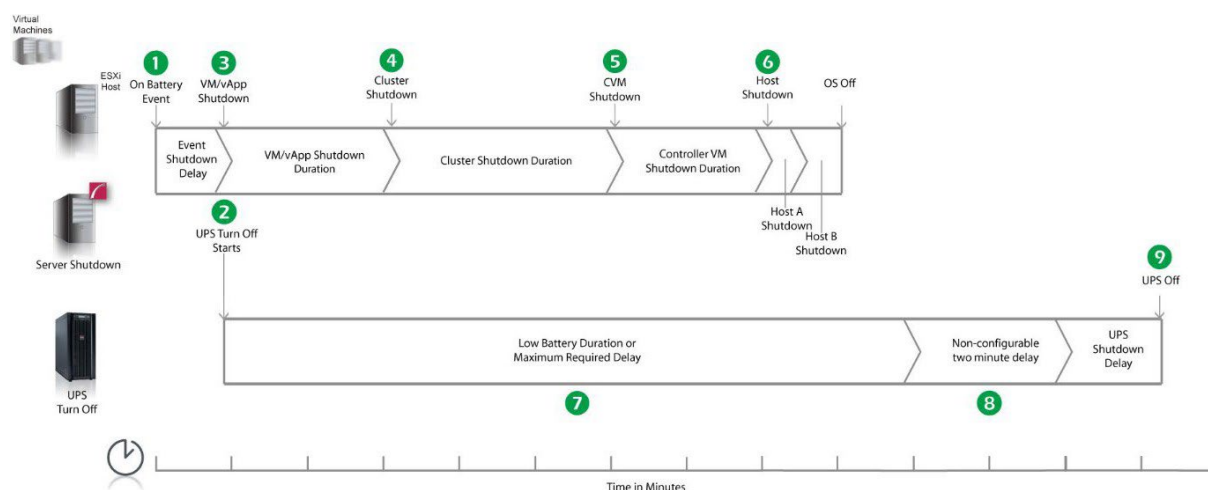
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 10) starts to count down.

Example 1b: Turn off the UPS enabled, no shutdown command file or SSH action configured, external vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an external vCenter Server Appliance. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



- PowerChute reports that the UPS is on battery.
- After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
- PowerChute starts to shut down User VMs and vApps.
- After 2 minutes (VM/vApp Shutdown Delay), PowerChute shuts down the HyperFlex Cluster.
- After 180 seconds (Cluster Shutdown Delay), PowerChute issues a command to shut down the Controller VMs.
- After 120 seconds (Controller VM Shutdown Delay), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
- The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay

These are shown on the Configuration - Shutdown page in the NMC interface.

8. After this delay, a further non-configurable two minute delay is counted down.
9. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.
This is configurable on the Configuration - Shutdown page in the NMC user interface.

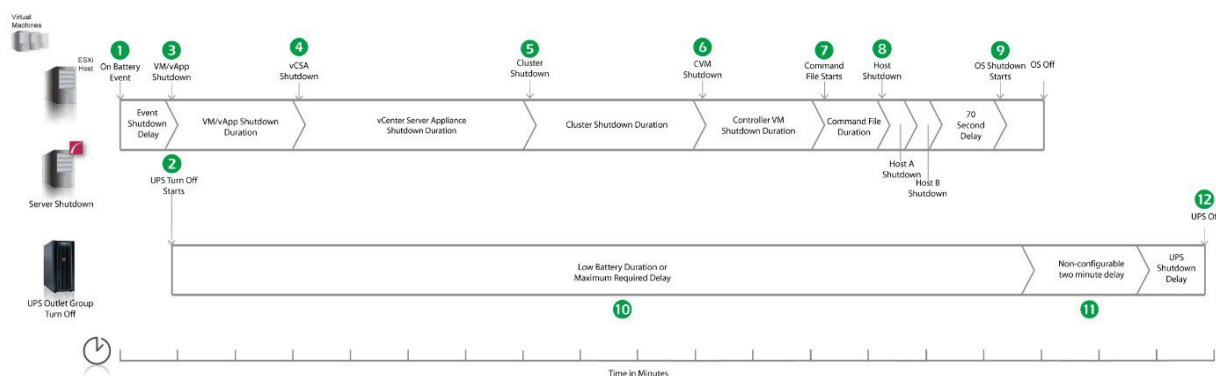
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 2a: Turn off the UPS enabled, shutdown command file configured, internal vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
6. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
7. After 120 seconds (Controller VM Shutdown Delay), PowerChute starts to execute the shutdown command file.
8. After the delay configured for the command file has elapsed, the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).

9. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.
10. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

11. After this delay, a further non-configurable two minute delay is counted down.
12. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.
This is configurable on the Configuration - Shutdown page in the NMC user interface.

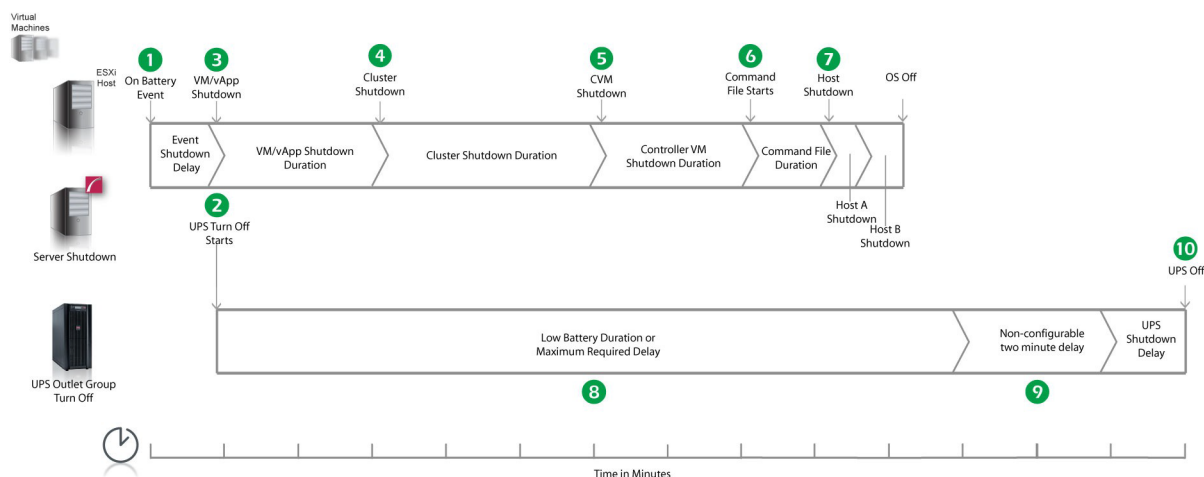
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 11) starts to count down.

Example 2b: Turn off the UPS enabled, shutdown command file configured, external vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an external vCenter Server Appliance. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
5. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
6. After 120 seconds (Controller VM Shutdown Delay), PowerChute starts to execute the shutdown command file.

7. After the delay configured for the command file has elapsed, the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.

8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.

10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.

This is configurable on the Configuration - Shutdown page in the NMC user interface.

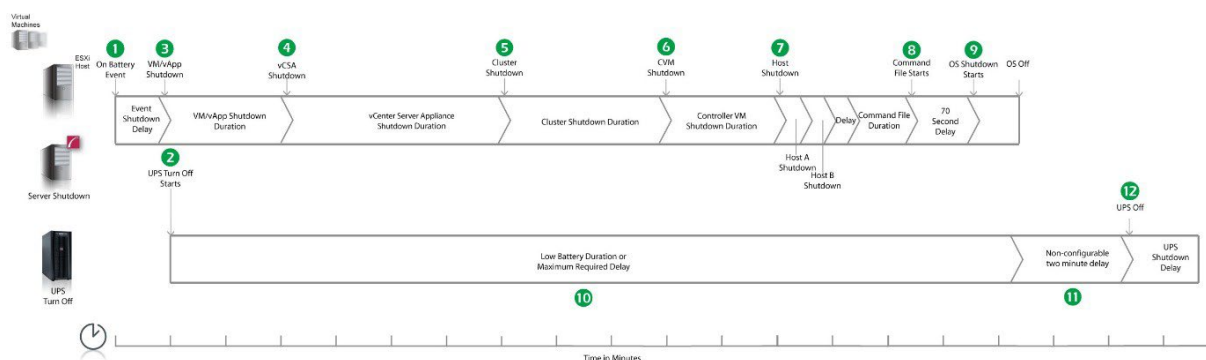
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 3a: Turn off the UPS enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled, internal vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.

6. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down Controller VMs.
7. After 120 seconds (Controller VM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
8. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
9. After the duration configured for the shutdown command file has elapsed, the OS Shutdown Command is issued and an additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
10. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

11. After this delay, a further non-configurable two minute delay is counted down.
12. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

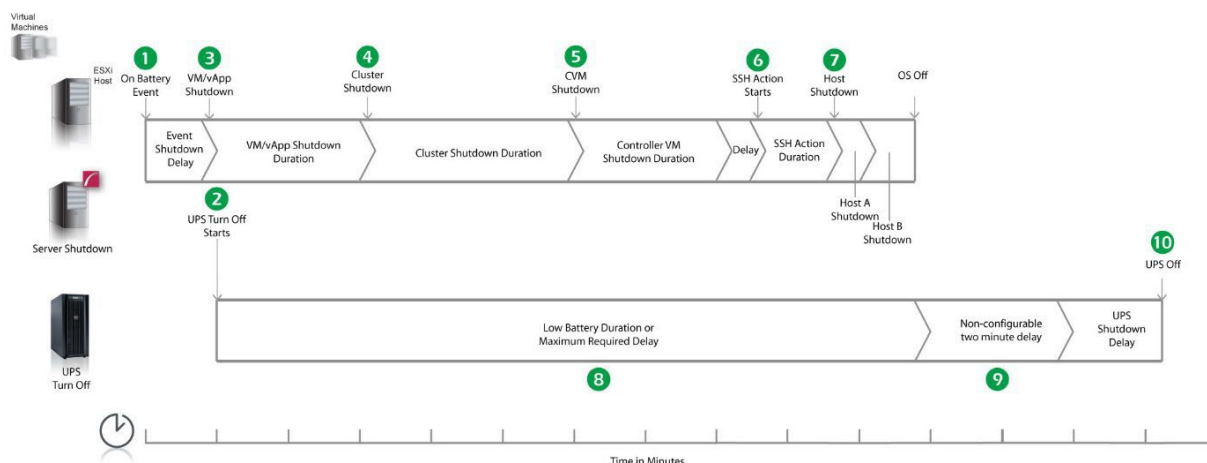
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 11) starts to count down.

Example 3b: Turn off the UPS enabled, SSH action configured, external vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an external vCenter Server Appliance. The option to Turn off the UPS is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
5. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down Controller VMs.
6. After 120 seconds (Controller VM Shutdown Duration), PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
7. VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:
 - Low Battery Duration** or **Maximum Required Delay**.These are shown on the Configuration - Shutdown page in the NMC interface.
9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

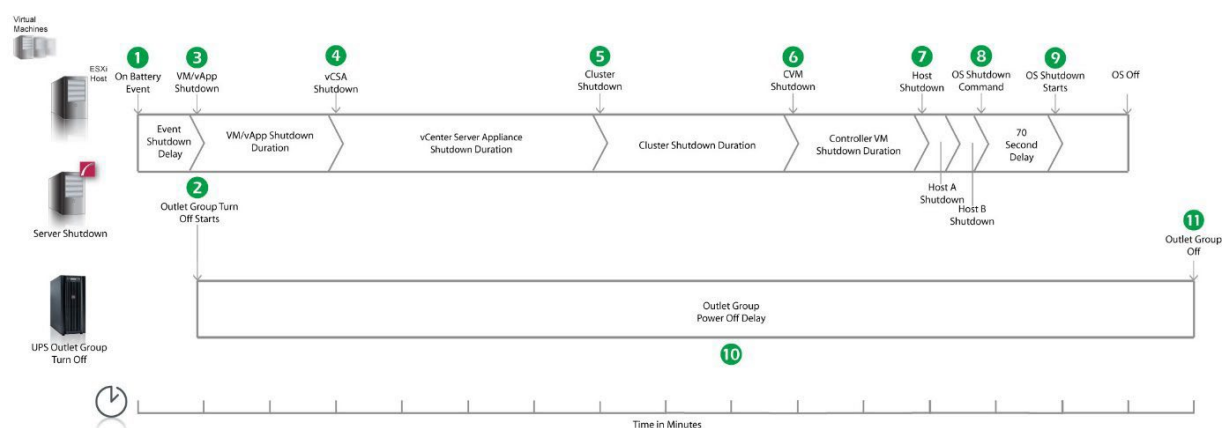
VMware with HyperFlex Support: UPS with Outlet Groups

Example 1a: Turn off the Outlet Group enabled, no shutdown command file or SSH action configured, internal vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
6. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
7. After 120 seconds (CVM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
8. PowerChute issues the operating system shutdown command.
9. After a 70 second delay, the operating system on the physical machine running PowerChute starts to shut down.

10. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.

- If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
- If registered with a Switched Outlet Group, only that delay is counted down.

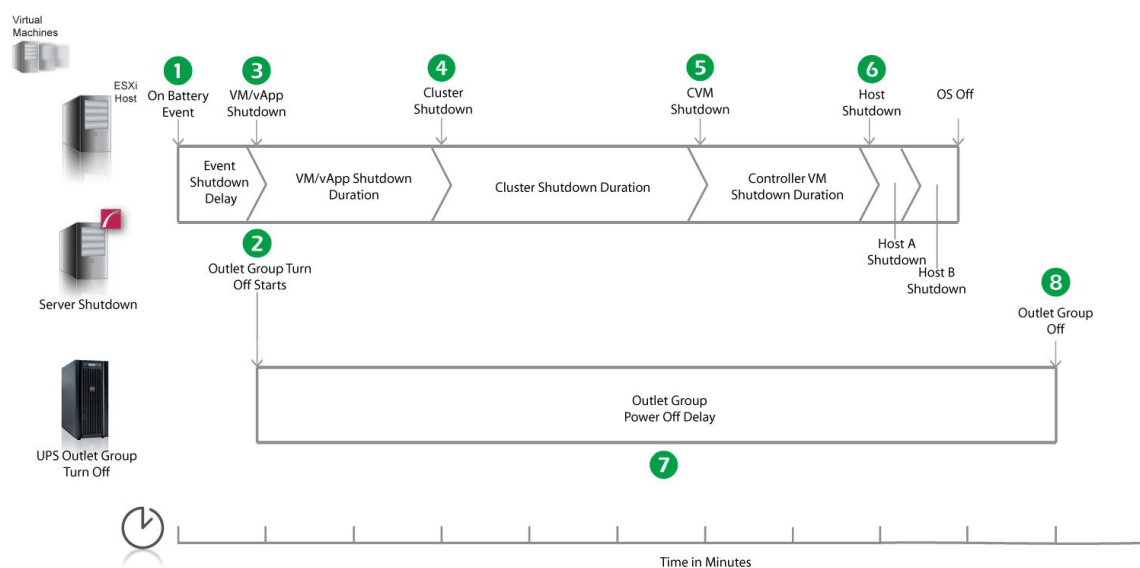
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 1b: Turn off the Outlet Group enabled, no shutdown command file or SSH action configured, external vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an external vCenter Server Appliance. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
5. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.

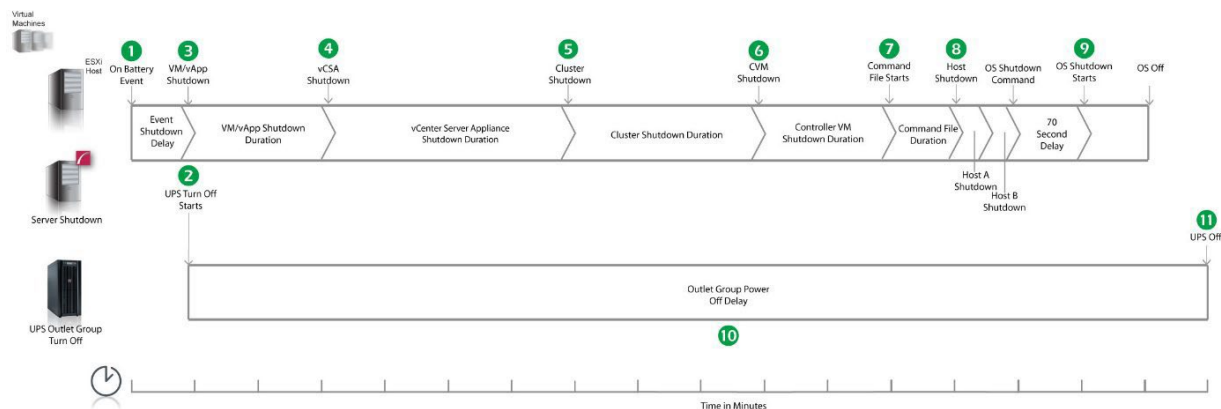
6. After 120 seconds (CVM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
7. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 2a: Turn off the Outlet Group enabled, shutdown command file configured, internal vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
6. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
7. After 120 seconds (CVM Shutdown Duration), PowerChute starts to execute the shutdown command file.

8. The VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
9. An additional 70 second delay is counted down before the operating system starts to shut down.
10. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

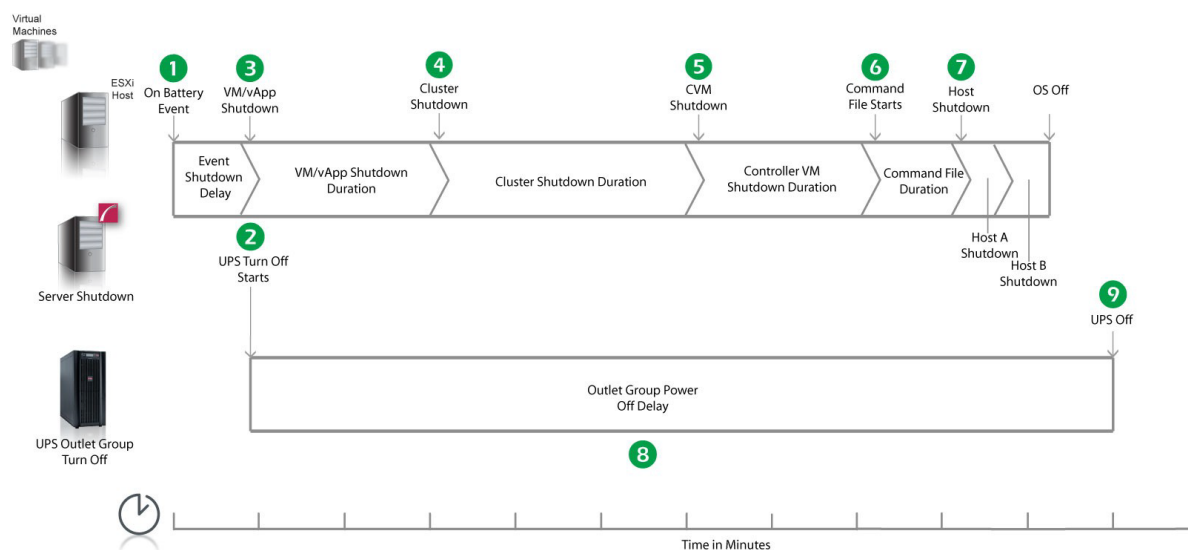
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 2b: Turn off the Outlet Group enabled, shutdown command file configured, external vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an external vCenter Server Appliance. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.

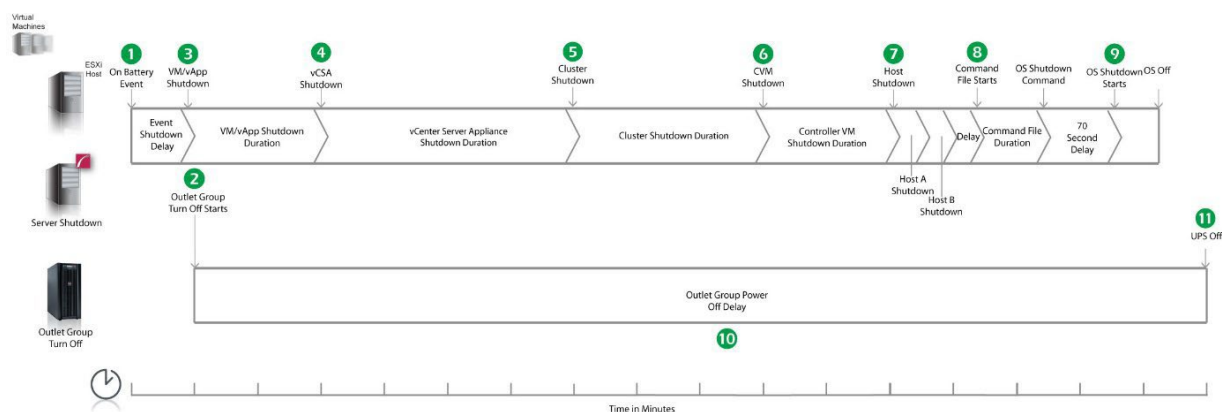
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
5. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
6. After 120 seconds (CVM Shutdown Duration), PowerChute starts to execute the shutdown command file.
7. The VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 3a: Turn off the Outlet Group enabled, shutdown command file configured, Execute Command File after Host Shutdown enabled, internal vCenter Server Appliance, PowerChute on physical Windows machine.

PowerChute is installed on a physical machine outside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A shutdown command file is configured. Execute Command File after Host Shutdown is enabled, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the vCenter Server Appliance.
5. After 240 seconds (vCSA Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
6. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.

7. After 120 seconds (CVM Shutdown Duration), the VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds).
8. Following the 30 second delay configured for the **Execute Command File after Host Shutdown** option, PowerChute starts to execute the shutdown command file.
9. After the duration configured for the shutdown command file has elapsed, an additional 70 second delay is counted down before the operating system starts to shut down.
10. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

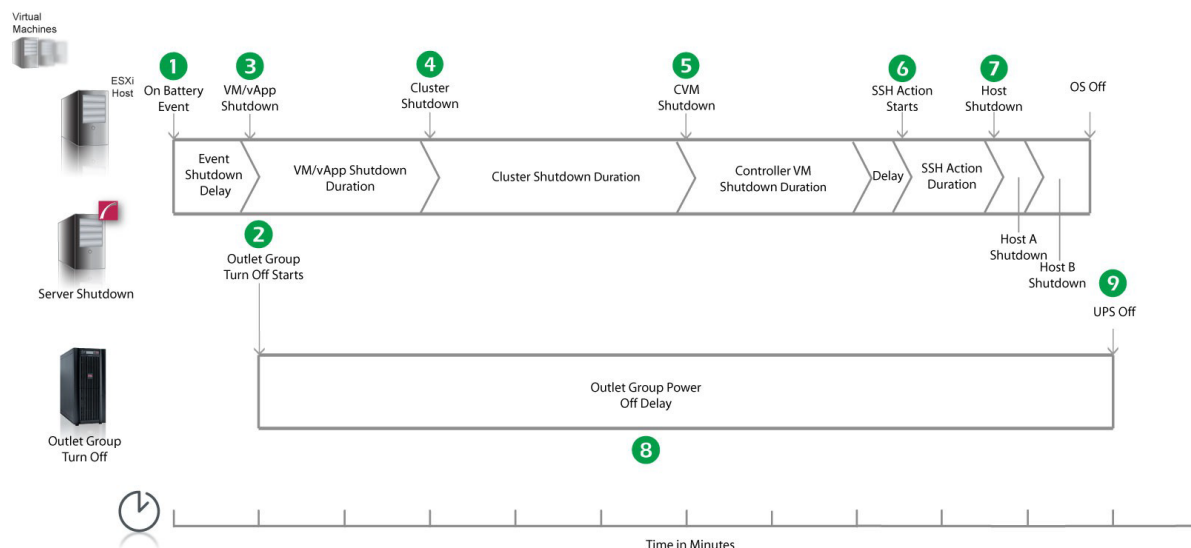
It is recommended that the Outlet Group **Power Off Delay** is configured to allow enough time for the operating system shutdown to complete. You should allow extra time to ensure that the Outlet Group does not turn off before the operating system.

Example 3b: Turn off the Outlet Group enabled, SSH action configured, external vCenter Server Appliance, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration with 2 nodes in a Cluster. HA is enabled on the Cluster and there is an internal vCenter Server Appliance. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page, Delay Host Maintenance Mode is enabled on the Virtualization Settings page. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Duration = 180 seconds
- Controller VM Shutdown Duration = 120 seconds
- Delay Host Maintenance Mode = 30 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute shuts down the HyperFlex Cluster.
5. After 180 seconds (Cluster Shutdown Duration), PowerChute issues a command to shut down the Controller VMs.
6. After 120 seconds (CVM Shutdown Duration), PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
7. VMware Hosts enter maintenance mode and are shut down sequentially if all VMs are powered off, otherwise the maintenance mode task is cancelled. The delay between placing each host into maintenance mode is the value set for Delay Maintenance Mode (30 seconds). The PowerChute host is powered off last.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Recommended Power-Off Delays for Outlet Groups

By default, the outlet group **Power Off Delay** will be the same value as the Low Battery duration configured on the NMC. PowerChute will automatically increase the Power Off Delay for the outlet group it is registered with, if the total shutdown time it needs is greater than the Power Off Delay.

The total shutdown time includes the following values:

- VM/vApp Shutdown and Startup Duration
- Cluster Shutdown and Startup Duration
- vCenter Server Appliance Shutdown Duration
- Controller VM Shutdown and Startup Duration
- Delay Host Maintenance Mode
- Execute Command File after Host Shutdown delay
- Shutdown Command File Duration
- SSH Action Duration
- Built-in delay of 2 minutes (this consists of a 10 second OS shutdown delay and a 60 second OS shutdown duration; rounded up)

The time required to gracefully shut down your operating system is not covered by the total shutdown time, as PowerChute cannot determine how long it will take to complete.



The Power Off Delay for the outlet group should be long enough for the OS to gracefully shut down. You should add extra time to allow for unforeseen circumstances.

The Low Battery Duration set on the NMC should be equal to or greater than the Power Off Delay for the outlet group.

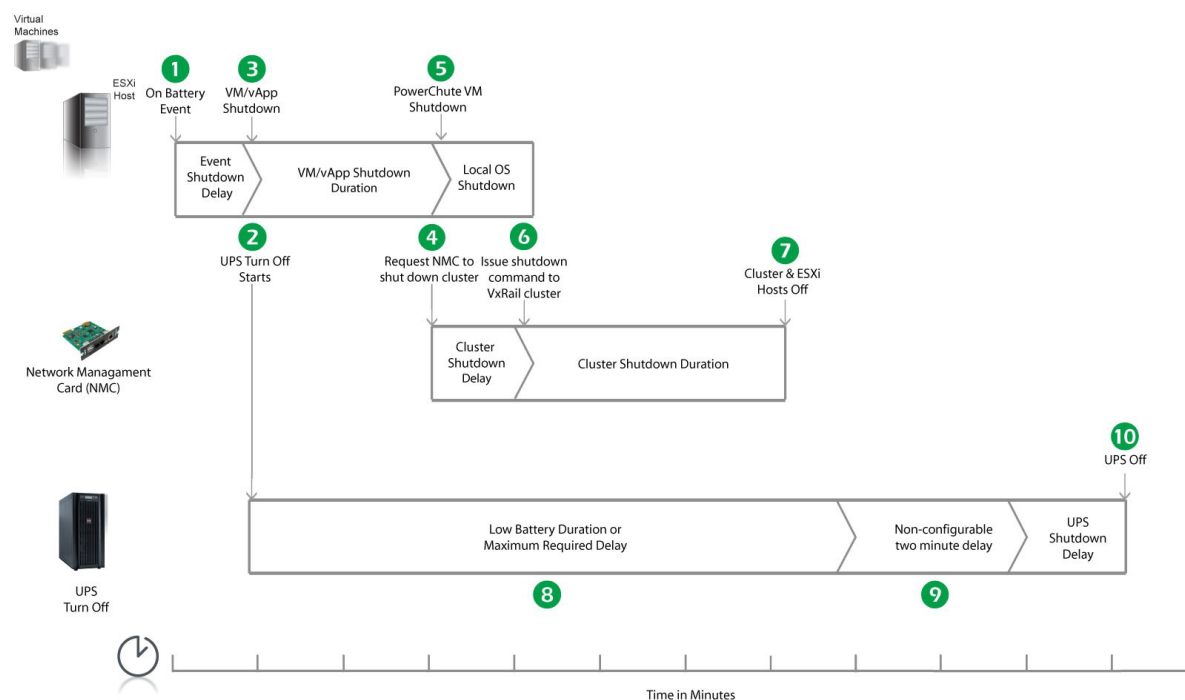
VMware with Dell VxRail Support: UPS without Outlet Groups

Example 1a: Standard Cluster. Turn off the UPS enabled, no shutdown command file or SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM inside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 60 seconds
- Cluster Shutdown Duration = 180 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.
5. PowerChute then starts a local OS shutdown to power itself off. All user VMs must be powered off before cluster shutdown can proceed.
6. After 60 seconds (Cluster Shutdown Delay), the NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM should be off.
7. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down and the UPS will turn off shortly after this.

8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.

10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.

This is configurable on the Configuration - Shutdown page in the NMC user interface.

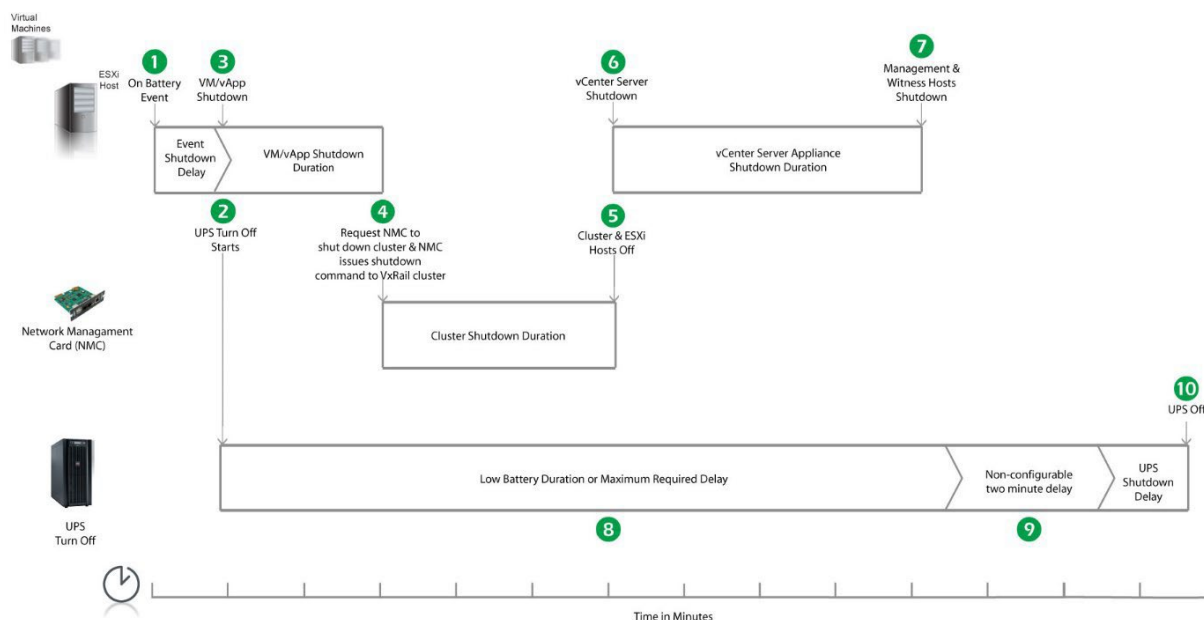
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 1b: Stretched Cluster. Turn off the UPS enabled, no shutdown command file or SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM on the Management host outside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. vCenter Server is deployed as a VM on the Management Host outside the Cluster. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 0 seconds
- Cluster Shutdown Duration = 180 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs should be off.

5. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down.
6. PowerChute shuts down the vCenter Server Appliance.
7. After 240 seconds (vCenter Server Appliance Shutdown Duration), PowerChute shuts down the Management and Witness hosts, including the PowerChute VM. The UPS will turn off shortly after this.
8. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay

These are shown on the Configuration - Shutdown page in the NMC interface.

9. After this delay, a further non-configurable two minute delay is counted down.
10. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

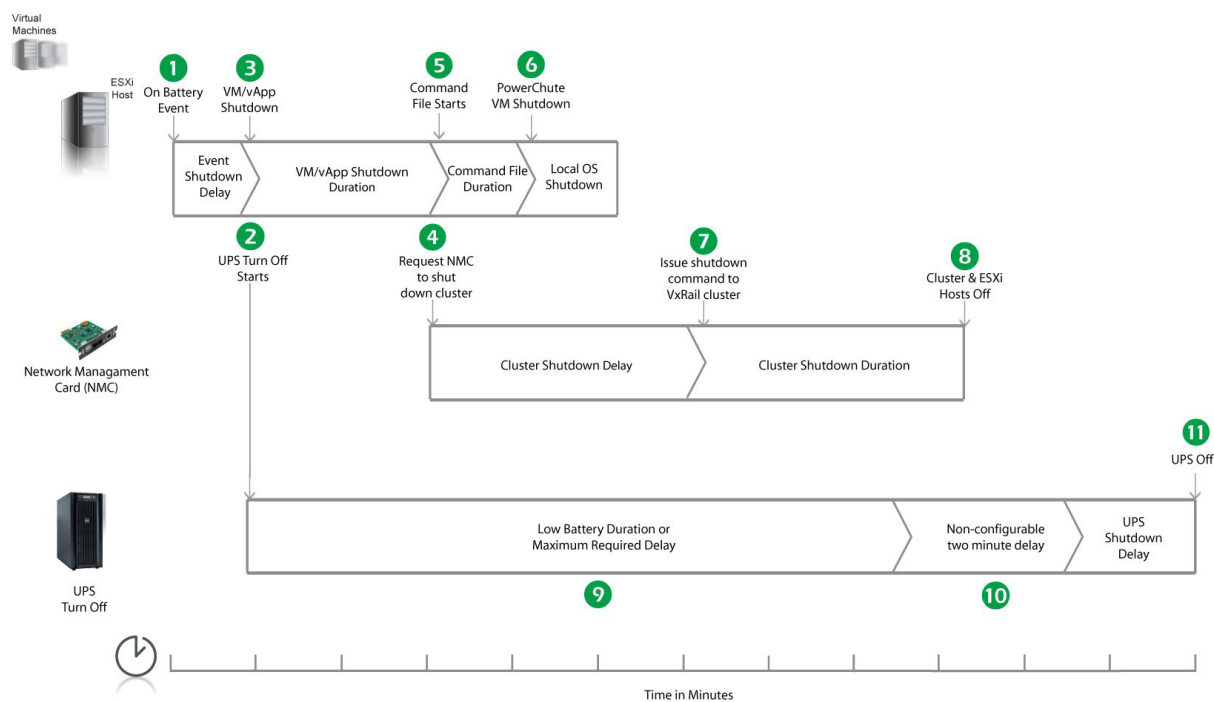
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete. Ideally the operating system should have shut down before the non-configurable two minute delay (step 9) starts to count down.

Example 2a: Standard Cluster. Turn off the UPS enabled, shutdown command file configured, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 180 seconds
- Cluster Shutdown Duration = 180 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.
5. PowerChute starts to execute the shutdown command file.
6. After the delay configured for the command file has elapsed, PowerChute then starts a local OS shutdown to power itself off. All user VMs must be powered off before cluster shutdown can proceed.
7. After 180 seconds (Cluster Shutdown Delay), the NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM should be off.
8. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down and the UPS will turn off shortly after this.
9. The UPS will wait the amount of time indicated by one of the following, whichever is greater:
 - Low Battery Duration** or **Maximum Required Delay**.These are shown on the Configuration - Shutdown page in the NMC interface.
10. After this delay, a further non-configurable two minute delay is counted down.
11. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

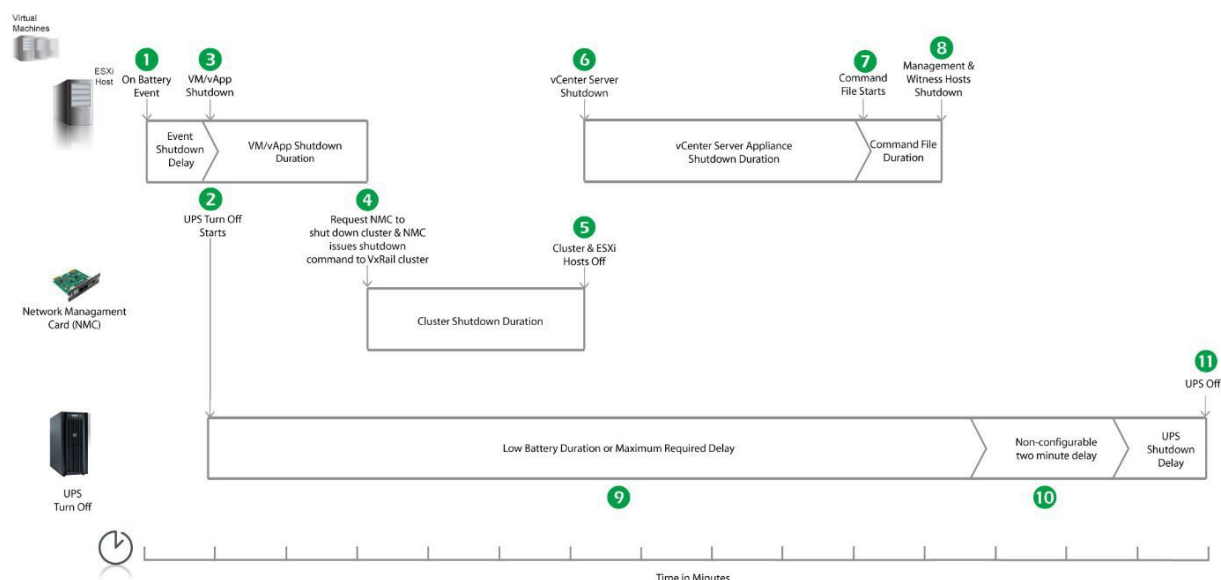
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 2b: Stretched Cluster. Turn off the UPS enabled, shutdown command file configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM on the Management host outside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. vCenter Server is deployed as a VM on the Management Host outside the Cluster. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 0 seconds
- Cluster Shutdown Duration = 180 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs should be off.
5. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down.
6. PowerChute shuts down the vCenter Server Appliance.
7. After 240 seconds (vCenter Server Appliance Shutdown Duration), PowerChute starts to execute the shutdown command file.
8. After the delay configured for the command file has elapsed, PowerChute shuts down the Management and Witness hosts, including the PowerChute VM. The UPS will turn off shortly after this.
9. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

10. After this delay, a further non-configurable two minute delay is counted down.
11. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed. This is configurable on the Configuration - Shutdown page in the NMC user interface.

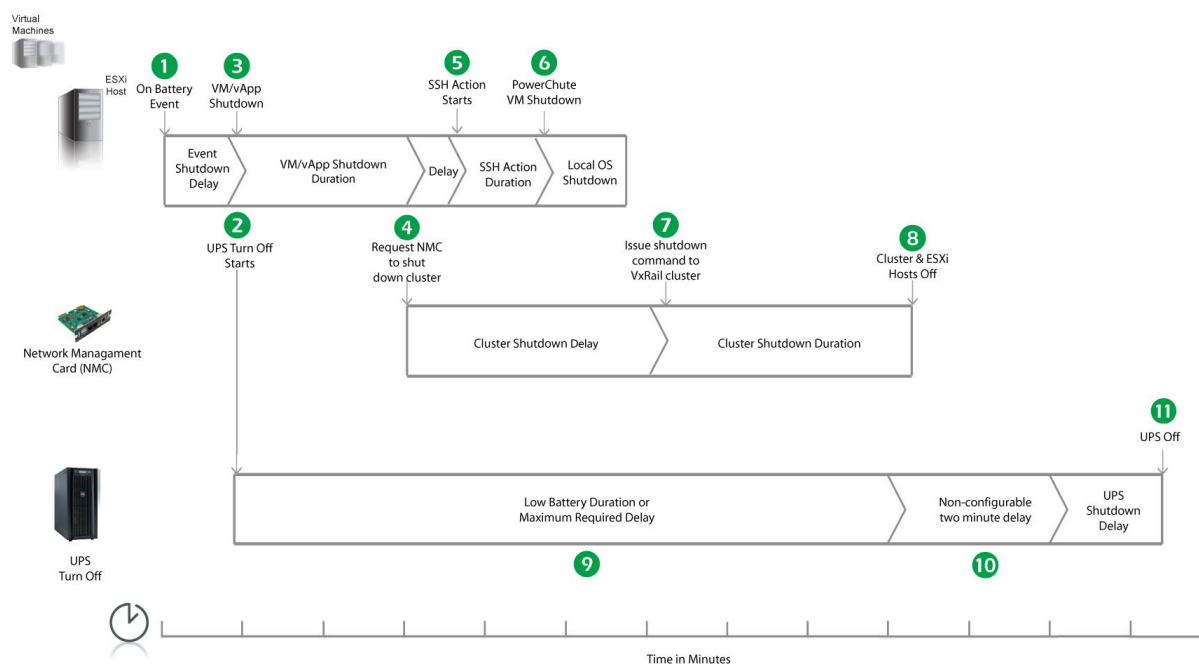
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 3a: Standard Cluster. Turn off the UPS enabled, SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 180 seconds
- Cluster Shutdown Duration = 180 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.
5. PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
6. PowerChute then starts a local OS shutdown to power itself off. All user VMs must be powered off before cluster shutdown can proceed.
7. After 180 seconds (Cluster Shutdown Delay), the NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM should be off.
8. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down and the UPS will turn off shortly after this.

9. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or Maximum Required Delay.

These are shown on the Configuration - Shutdown page in the NMC interface.

10. After this delay, a further non-configurable two minute delay is counted down.

11. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.

This is configurable on the Configuration - Shutdown page in the NMC user interface.

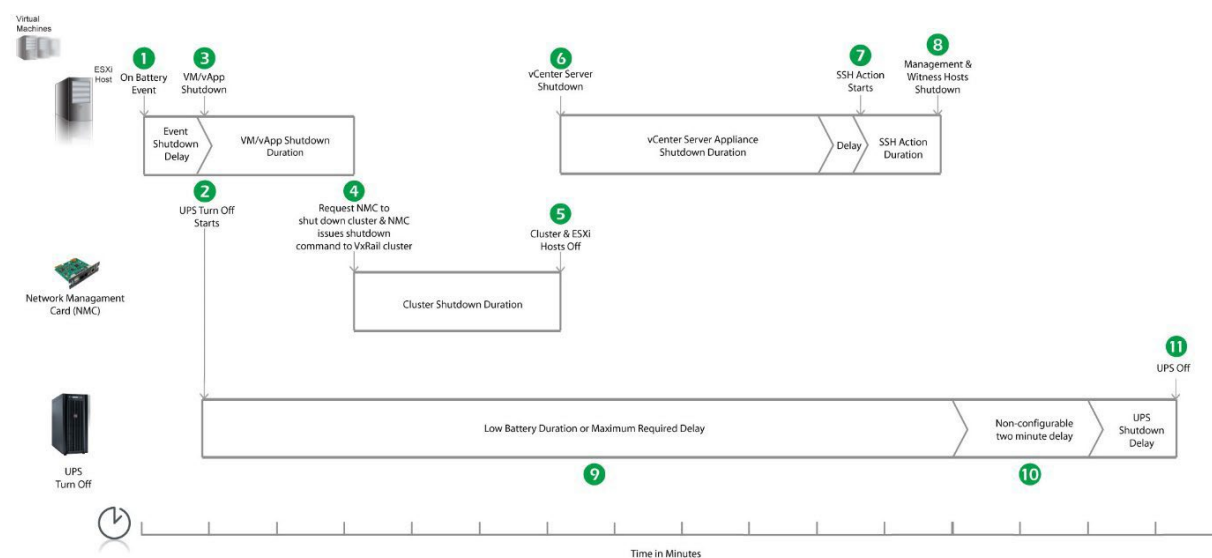
It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

Example 3b: Stretched Cluster. Turn off the UPS enabled, SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM on the Management host outside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. vCenter Server is deployed as a VM on the Management Host outside the Cluster. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 0 seconds
- Cluster Shutdown Duration = 180 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs should be off.

5. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down.
6. PowerChute shuts down the vCenter Server Appliance.
7. After 240 seconds (vCenter Server Appliance Shutdown Duration), PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
8. After the delay configured for the SSH action has elapsed, PowerChute shuts down the Management and Witness hosts, including the PowerChute VM. The UPS will turn off shortly after this.
9. The UPS will wait the amount of time indicated by one of the following, whichever is greater:

Low Battery Duration or **Maximum Required Delay**.

These are shown on the Configuration - Shutdown page in the NMC interface.

10. After this delay, a further non-configurable two minute delay is counted down.
11. The UPS will then turn off after the user-configurable **Shutdown Delay** time has elapsed.
This is configurable on the Configuration - Shutdown page in the NMC user interface.

It is recommended that the **Low Battery Duration** is configured to allow enough time for the Operating System shutdown to complete.

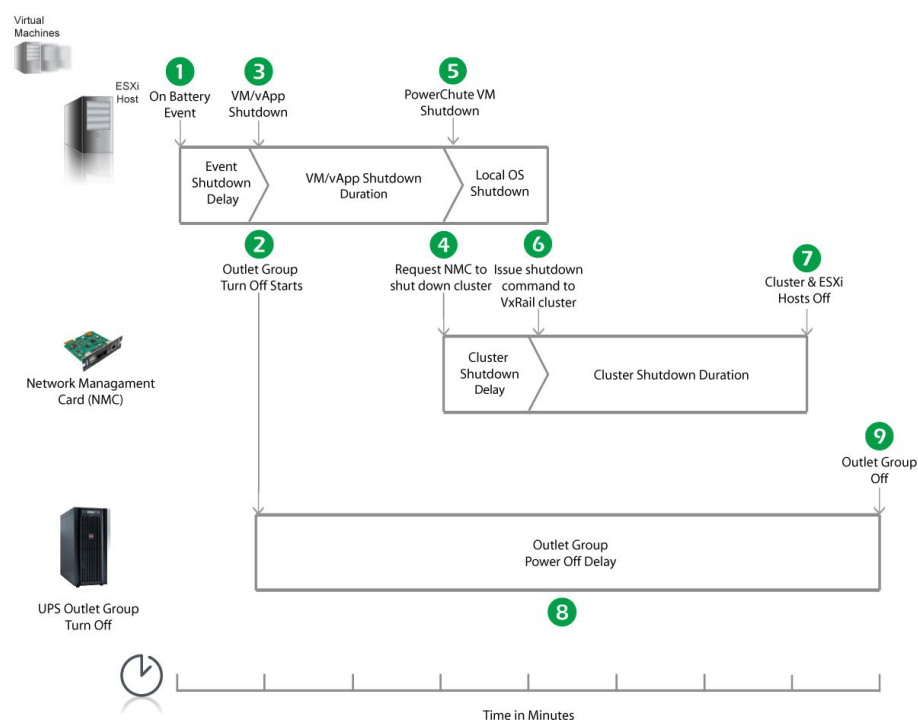
VMware with Dell VxRail Support: UPS with Outlet Groups

Example 1a: Standard Cluster. Turn off the Outlet Group enabled, no shutdown command file or SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 60 seconds
- Cluster Shutdown Duration = 180 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.
5. PowerChute then starts a local OS shutdown to power itself off. All user VMs must be powered off before cluster shutdown can proceed.
6. After 60 seconds (Cluster Shutdown Delay), the NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM should be off.
7. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down and the UPS will turn off shortly after this.

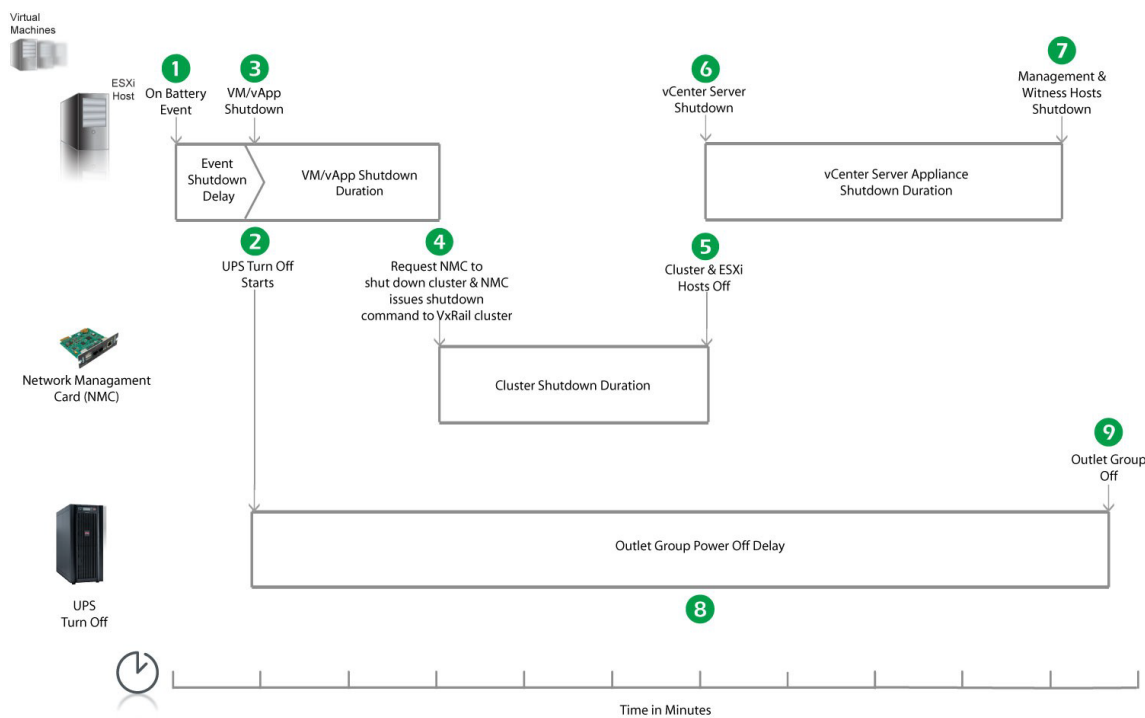
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 1b: Stretched Cluster. Turn off the Outlet Group enabled, no shutdown command file or SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM on the Management host outside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. vCenter Server is deployed as a VM on the Management Host outside the Cluster. No shutdown command file or SSH action configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 0 seconds
- Cluster Shutdown Duration = 180 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Delay), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs should be off.
5. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down.

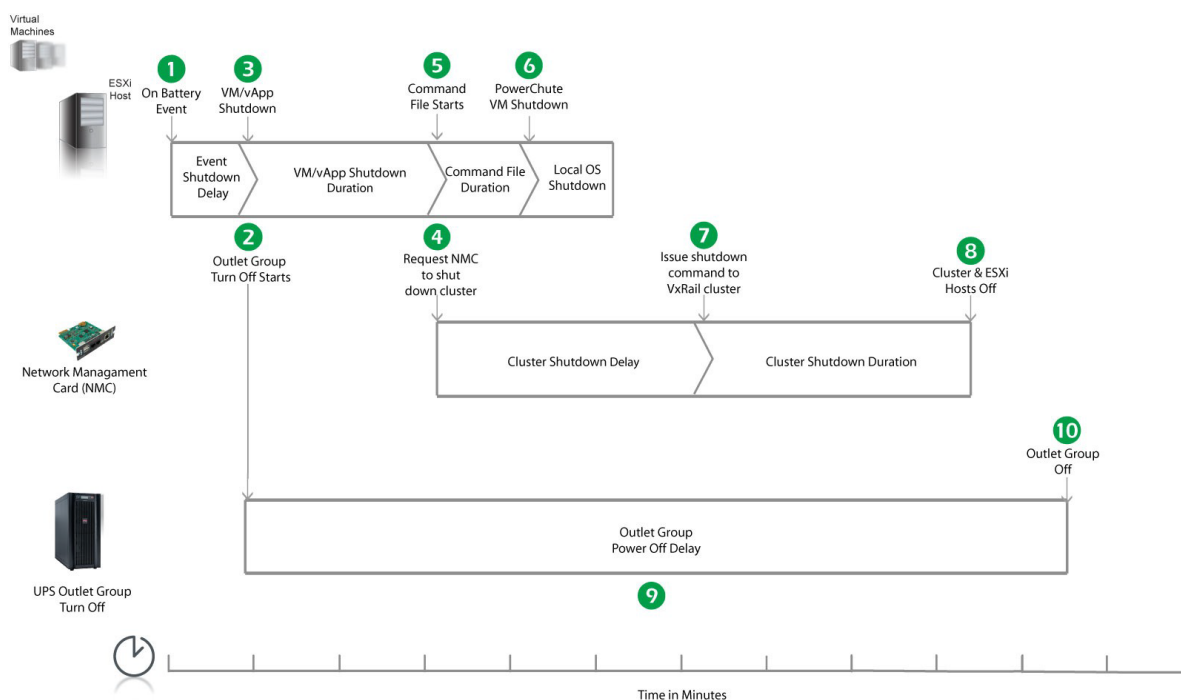
6. PowerChute shuts down the vCenter Server Appliance.
7. After 240 seconds (vCenter Server Appliance Shutdown Duration), PowerChute shuts down the Management and Witness hosts, including the PowerChute VM. The UPS will turn off shortly after this.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 2a: Standard Cluster. Turn off the Outlet Group enabled, shutdown command file configured, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page. A shutdown command file is configured. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 180 seconds
- Cluster Shutdown Duration = 180 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.

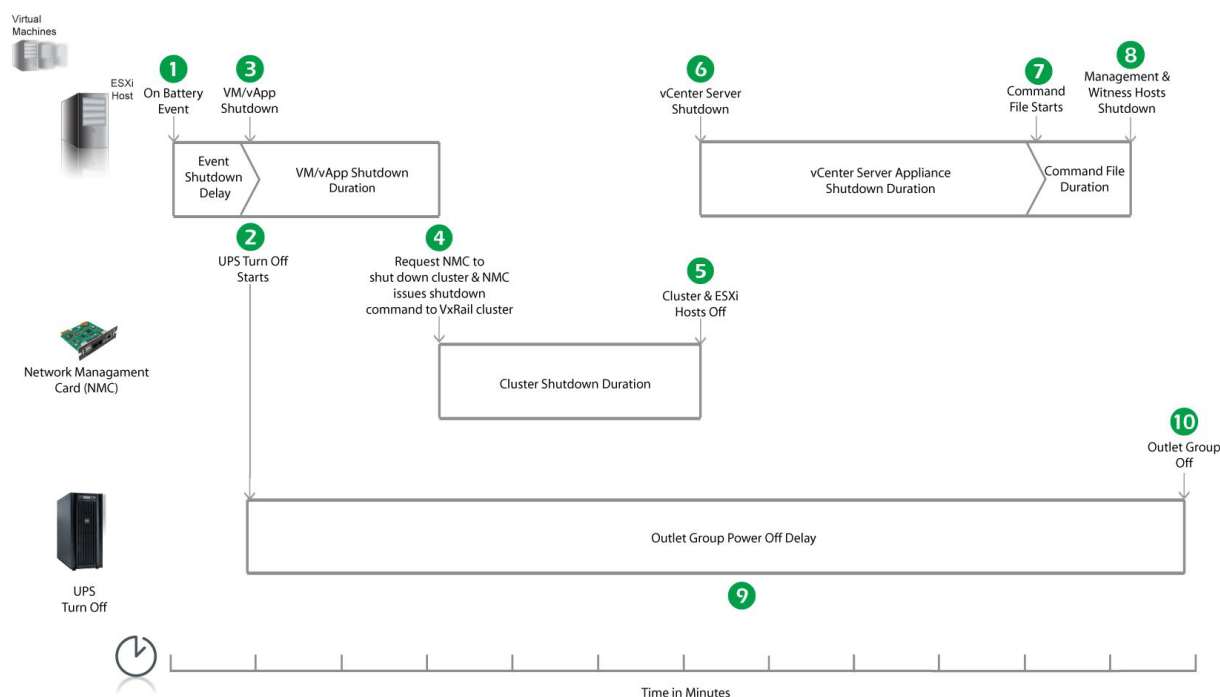
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.
5. PowerChute starts to execute the shutdown command file.
6. After the delay configured for the command file has elapsed, PowerChute then starts a local OS shutdown to power itself off. All user VMs must be powered off before cluster shutdown can proceed.
7. After 180 seconds (Cluster Shutdown Delay), the NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM should be off.
8. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down and the UPS will turn off shortly after this.
9. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 2b: Stretched Cluster. Turn off the Outlet Group enabled, shutdown command file configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM on the Management host outside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. vCenter Server is deployed as a VM on the Management Host outside the Cluster. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 0 seconds
- Cluster Shutdown Duration = 180 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



PowerChute reports that the UPS is on battery.

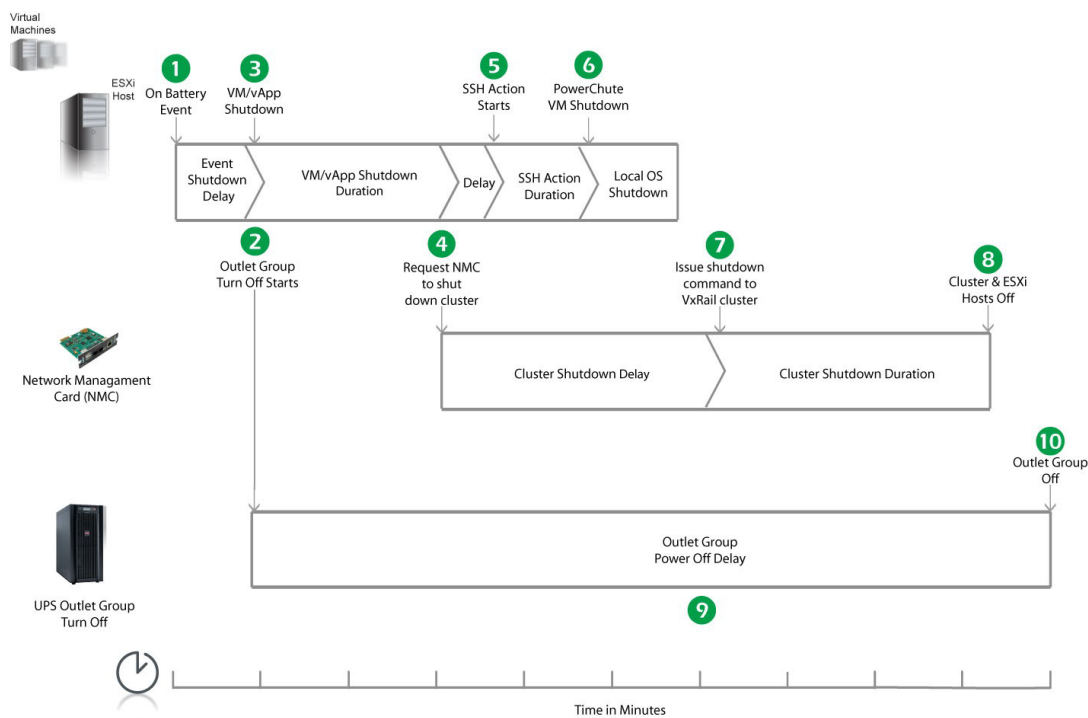
1. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
2. PowerChute starts to shut down User VMs and vApps.
3. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs should be off.
4. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down.
5. PowerChute shuts down the vCenter Server Appliance.
6. After 240 seconds (vCenter Server Appliance Shutdown Duration), PowerChute starts to execute the shutdown command file.
7. After the delay configured for the command file has elapsed, PowerChute shuts down the Management and Witness hosts, including the PowerChute VM. The UPS will turn off shortly after this.
8. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 3a: Standard Cluster. Turn off the Outlet Group enabled, SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as VM inside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the Outlet Group is enabled on the Shutdown Settings page. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 180 seconds
- Cluster Shutdown Duration = 180 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



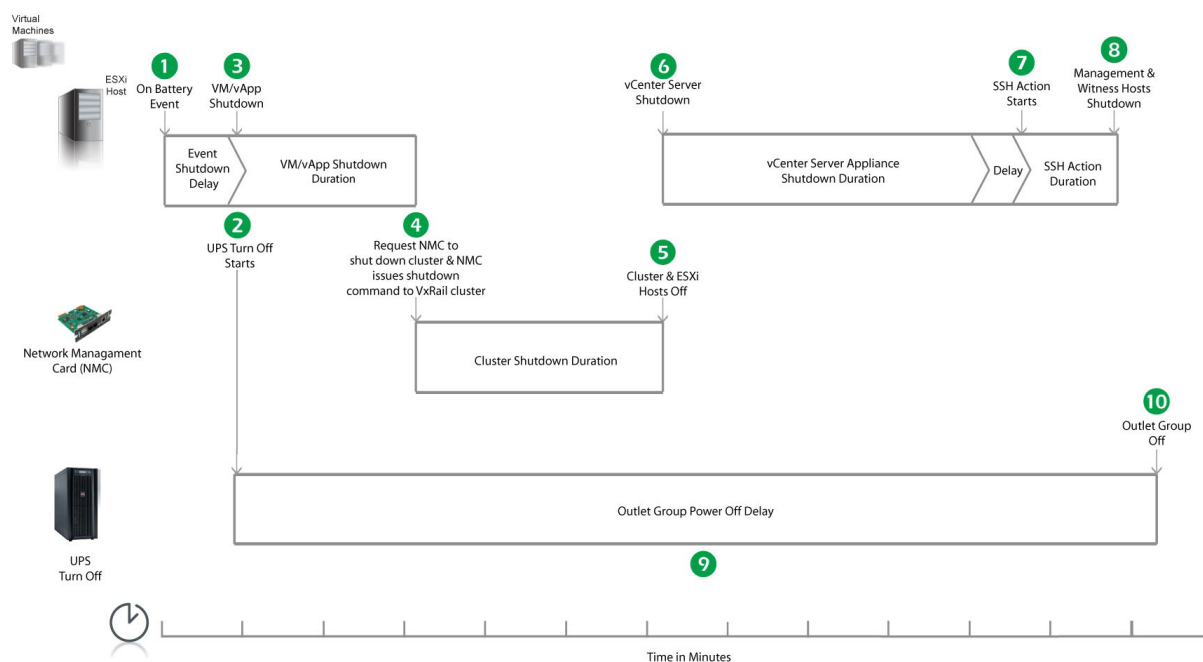
1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.
5. PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
6. PowerChute then starts a local OS shutdown to power itself off. All user VMs must be powered off before cluster shutdown can proceed.
7. After 180 seconds (Cluster Shutdown Delay), the NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM should be off.
8. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down and the UPS will turn off shortly after this.
9. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Example 3b: Stretched Cluster. Turn off the Outlet Group enabled, SSH action configured, PowerChute deployed as a VM.

PowerChute is deployed as a VM on the Management host outside the Cluster, configured for a Single/Redundant UPS configuration. The option to Turn off the UPS is enabled on the Shutdown Settings page. vCenter Server is deployed as a VM on the Management Host outside the Cluster. A SSH action is configured to execute before host shutdown, with a delay of 30 seconds applied. The durations in this example are as follows:

- VM and vApp Shutdown Duration = 120 seconds
- Cluster Shutdown Delay = 0 seconds
- Cluster Shutdown Duration = 180 seconds
- vCenter Server Appliance Shutdown Duration = 240 seconds

When a critical UPS event, such as On Battery occurs, the following sequence is triggered.



1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the Outlet Group and the Outlet Group turn off starts.
3. PowerChute starts to shut down User VMs and vApps.
4. After 2 minutes (VM/vApp Shutdown Duration), PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs should be off.
5. After 180 seconds (Cluster Shutdown Duration), the cluster and ESXi hosts are shut down.
6. PowerChute shuts down the vCenter Server Appliance.
7. After 240 seconds (vCenter Server Appliance Shutdown Duration), PowerChute waits the 30 second delay configured for **Execute SSH Action before Host Shutdown**, and starts to execute the SSH action.
8. After the delay configured for the SSH action has elapsed, PowerChute shuts down the Management and Witness hosts, including the PowerChute VM. The UPS will turn off shortly after this.

9. The Outlet Group will turn off after the **Power Off Delay** (configurable on the **Configuration – Outlet Group** page in the NMC user interface) has elapsed.
 - If registered with the Main Outlet Group, the UPS will wait for any Switched Outlet Groups to turn off before the Main Outlet Group turnoff starts.
 - If registered with a Switched Outlet Group, only that delay is counted down.

Recommended Power-Off Delays for Outlet Groups

By default, the outlet group **Power Off Delay** will be the same value as the Low Battery duration configured on the NMC. PowerChute will automatically increase the Power Off Delay for the outlet group it is registered with, if the total shutdown time it needs is greater than the Power Off Delay.

The total shutdown time includes the following values:

- VM/vApp Shutdown and Startup Duration
- Cluster Shutdown Delay
- Cluster Shutdown Duration
- Delay Host Maintenance Mode (for stretched cluster configurations)
- Shutdown Command File Duration
- SSH Action Duration
- Built-in delay of 2 minutes (this consists of a 10 second OS shutdown delay and a 60 second OS shutdown duration; rounded up)

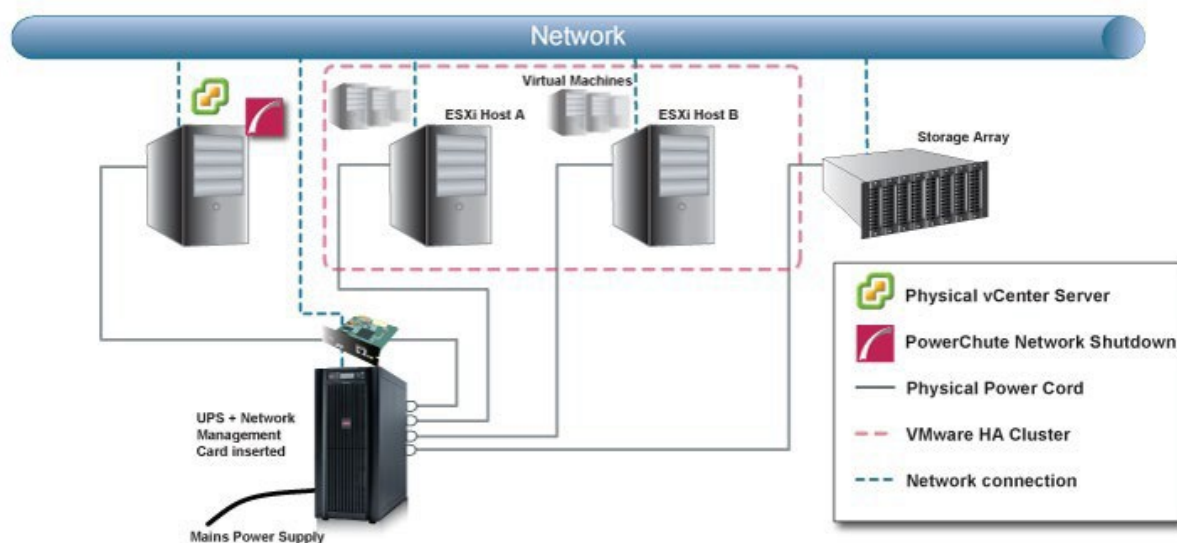
The time required to gracefully shut down your operating system is not covered by the total shutdown time, as PowerChute cannot determine how long it will take to complete.



The Low Battery Duration set on the NMC should be equal to or greater than the Power Off Delay for the outlet group.

VMware Shutdown - Single UPS Configuration

In this example, there are two VMware hosts, a vCenter Server and a storage array being powered by a single UPS. PowerChute is installed on the vCenter Server machine outside the cluster.



The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

1. The UPS has been running on Battery power for x number of seconds.
2. PowerChute issues a command to turn off the UPS and UPS turnoff starts.
3. PowerChute starts a maintenance mode task on the VMware hosts and shuts down the VMs on VMware hosts A and B.
4. PowerChute shuts down the vApp if configured.
5. PowerChute runs the shutdown command file if configured.
6. After the shutdown command file duration has elapsed, PowerChute shuts down the VMware hosts.
7. PowerChute shuts down the vCenter Server Machine.

VMware Shutdown - HA Cluster

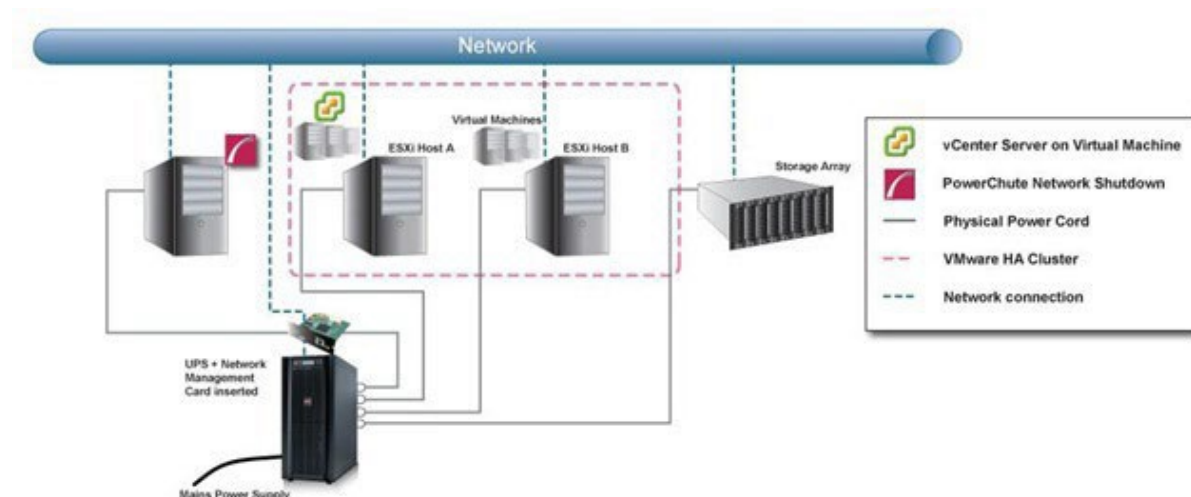
In the following examples, a VMware HA Cluster is protected by a Single, Redundant or Parallel UPS configuration. vCenter Server is running on a virtual machine.

Recommended Deployment

PowerChute can run on a VM in the HA cluster (either installed on the vMA or deployed as a virtual appliance) or be installed on a physical Windows machine outside the cluster. The vCenter Server account configured in PowerChute Network Shutdown must have Administrator permissions on vCenter Server and on each of the ESXi hosts being managed by PowerChute. This can be an Active Directory account or a local user account. For more information see [Active Directory VMware Configuration](#).

Example 1: vCenter Server is running on a VM; PowerChute is installed on a physical Windows machine

- VM & vApp Shutdown enabled with a 120 second delay (i.e. 120 seconds allocated for each action to complete).
- The option to turn off the UPS or Outlet Group is enabled.
- A shutdown command file has been configured with a 120 second duration.



When a critical UPS event, such as UPS on Battery occurs the following sequence is triggered:

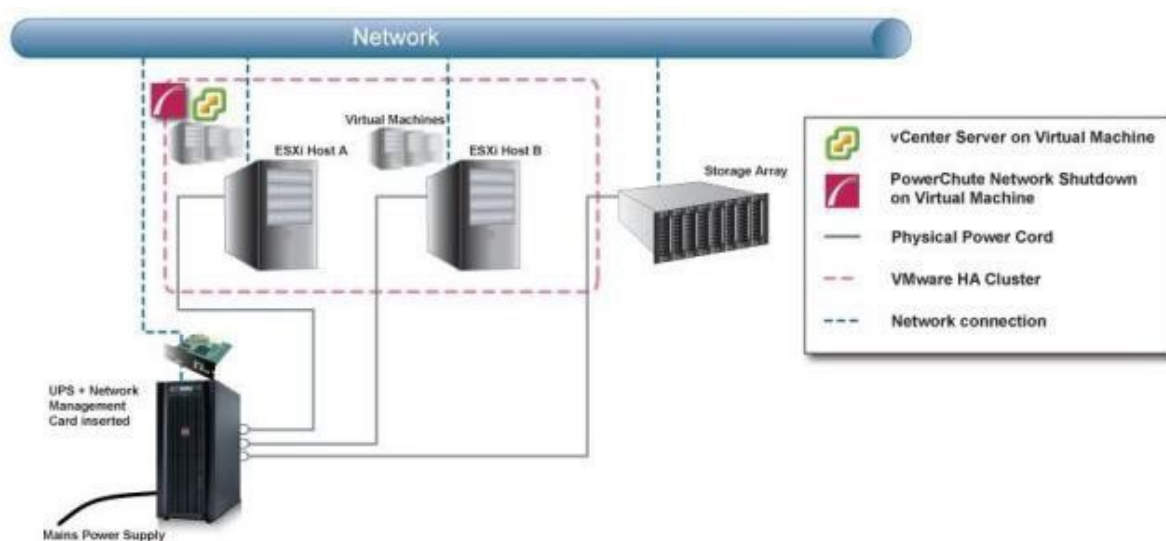
Shutdown Sequence

1. PowerChute reports that the UPS is on battery.
2. Shutdown delay for the On Battery event elapses. PowerChute sends a command to turn off the UPS or Outlet Group.
3. PowerChute starts a maintenance mode task on each Host and then starts VM/vApp shutdown.
4. VM/vApp shutdown durations elapse.
5. PowerChute gracefully shuts down the vCenter Server VM.
6. vCenter VM shutdown duration elapses. PowerChute starts executing the shutdown command file.
7. Shutdown command file duration elapses and PowerChute gracefully shuts down the VMware hosts that are not running the vCenter Server VM.
8. PowerChute shuts down the VMware Host running the vCenter Server VM.

9. OS shutdown sequence starts on the PowerChute physical machine.
10. After a 70 second delay the OS starts to shut down.
11. UPS waits for the duration that is greatest of Low Battery Duration/Maximum Required Delay (Non-Outlet Aware UPS's) or the Outlet Group Power Off Delay.
12. UPS turns off after the user-configurable Shutdown Delay time has elapsed or the Outlet Group turns off after the power off Delay elapses.

Example 2: Both vCenter Server and PowerChute are running on Virtual Machines

- VM & vApp Shutdown enabled with a 120 second delay (i.e. 120 seconds allocated for each action to complete)
- The option to turn off the UPS or Outlet Group is enabled
- A shutdown command file has been configured with a 120 second duration



When a critical UPS event, such as UPS on Battery occurs the following sequence is triggered:

Shutdown Sequence

1. PowerChute reports that the UPS is on battery.
2. Shutdown delay for the On Battery event elapses. PowerChute sends a command to turn off the UPS or Outlet Group.
3. PowerChute starts a maintenance mode task on each host and then starts VM/vApp shutdown.
4. VM/vApp shutdown durations elapse.
5. PowerChute gracefully shuts down the vCenter Server VM.
6. vCenter VM shutdown duration elapses. PowerChute starts executing the shutdown command file.
7. Shutdown command file duration elapses and PowerChute gracefully shuts down the VMware hosts that are not running the vCenter Server or PowerChute VM.
8. PowerChute shuts down the VMware host running vCenter Server VM followed by the host running PowerChute VM.

Note: The Maintenance mode task is cancelled for the Host running PowerChute so HA can attempt to restart the PowerChute VM when its host is powered back on.

9. UPS waits for the duration that is greatest of Low Battery Duration/Maximum Required Delay Non-Outlet Aware UPS's or the Outlet Group Power Off Delay.
10. UPS turns off after the user-configurable Shutdown Delay time has elapsed or the Outlet Group turns off after the power off Delay elapses.

VMware setups with multiple Clusters or Datacenters

If PowerChute is deployed as a virtual appliance we recommend deploying one PowerChute Agent per cluster if your setup has multiple clusters.

In environments where there are multiple clusters or datacenters you can use one copy of PowerChute installed on a physical Windows machine to monitor your hosts. PowerChute should be installed on multiple machines if the datacenters/clusters contain hosts that are in different geographical locations.

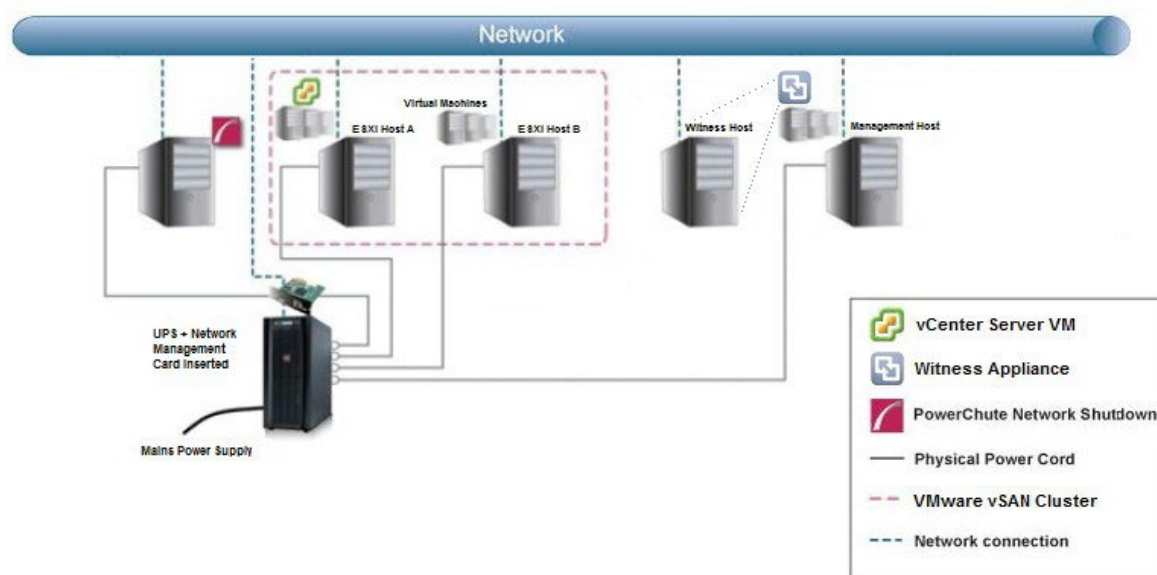
VMware Shutdown - vSAN Cluster

In the following examples, a VMware vSAN Cluster is protected by a Single or Advanced UPS Configuration. vCenter Server is running on a Virtual Machine.

Recommended Deployment

PowerChute can be located on a physical Windows machine outside the vSAN Cluster or deployed as a VM inside the vSAN Cluster. The vCenter Server account configured in PowerChute Network Shutdown must have Administrator permissions on vCenter Server and on each of the ESXi hosts being managed by PowerChute. This can be an Active Directory account or a local user account. For more information see [Active Directory VMware Configuration](#). The vSAN hosts must have SSH access enabled for the vSAN cluster preparation script to run correctly and successfully.

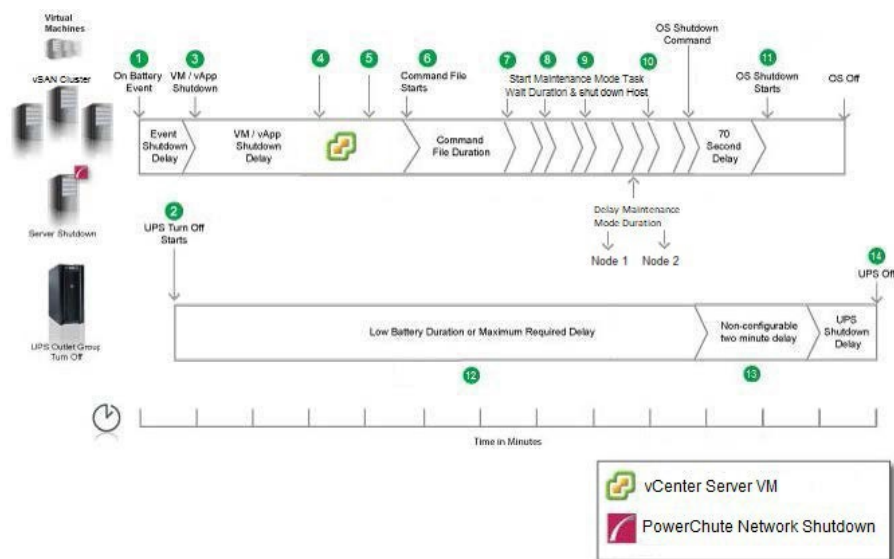
Example 1: Single UPS, 2-Node Stretch Cluster with Witness Appliance and Management Host



Setup

- PowerChute is installed on a physical Windows machine.
- Single UPS configuration.
- [Delay Maintenance Mode](#) is enabled (required for vSAN hosts) with duration set to X seconds.
- vCenter Server is running on a VM inside the vSAN Cluster.
- vSAN Witness Appliance is deployed on the Management Host and added to the inventory as a Host.
- [VM Prioritization](#) is enabled.
- An Active Directory controller VM is added to the High priority group with appropriate duration for shutdown and startup. This VM needs to be shut down after the vCenter Server VM and started before it.
- vCenter Server VM is added to the Medium priority group with appropriate duration for shutdown and startup.
- The option to turn off the UPS or Outlet Group is enabled.
- A shutdown command file has been configured with a 120 second duration.

When a critical UPS event, such as UPS on Battery occurs the following sequence is triggered:



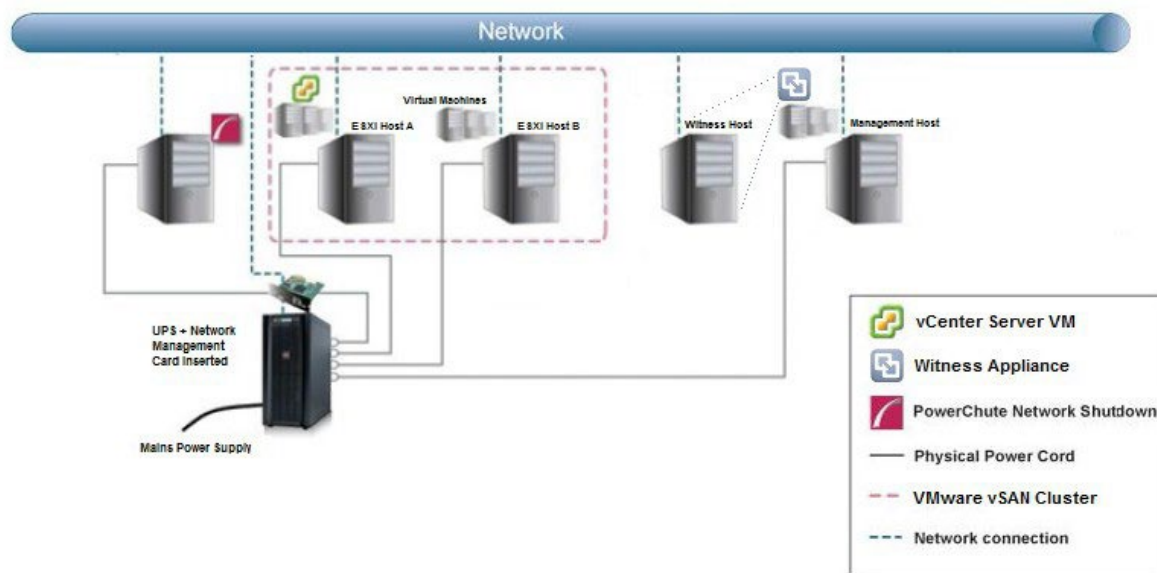
Shutdown Sequence

1. PowerChute reports that the UPS is on battery.
2. Shutdown delay for the On Battery event elapses. PowerChute sends a command to turn off the UPS or Outlet Group.
3. PowerChute starts VM and vApp shutdown on each Host.
4. PowerChute gracefully shuts down the vCenter Server VM in the Medium priority group followed by the Active Directory controller VM in the High priority group during VM Shutdown.
5. VM/vApp Shutdown durations elapse.
6. PowerChute starts executing the shutdown command file.
7. Shutdown command file duration elapses and PowerChute starts a Maintenance Mode task on the first vSAN host, and waits the Delay Host Maintenance Mode Duration. If there is a vSAN synchronization active, PowerChute will wait the specified **vSAN Synchronization Duration** and check if it has completed with respect to the value set for "vsan_synch_retry_time" in the `pcnconfig.ini` file until data re-synchronization is no longer active, or the retry limit has been reached. PowerChute then shuts down the host.
8. Once the Delay Host Maintenance Mode Duration has elapsed, PowerChute will start a Maintenance Mode task on the next vSAN host, and waits the Delay Host Maintenance Mode Duration. If there is a vSAN synchronization active, PowerChute will wait the specified **vSAN Synchronization Duration** and check if it has completed with respect to the value set for "vsan_synch_retry_time" in the `pcnconfig.ini` file until data re-synchronization is no longer active, or the retry limit has been reached. PowerChute then shuts down the host.
9. PowerChute starts a Maintenance Mode task on the Witness Host, waits the Delay Host Maintenance Mode Duration and shuts down the host.
10. PowerChute starts a Maintenance Mode task on the Management Host, waits the Delay Host Maintenance Mode Duration of X seconds and shuts down the host.
11. OS shutdown sequence starts on the PowerChute physical machine. After a 70 second delay, the OS starts to shut down.
12. UPS waits for the duration that is greatest of Low Battery Duration/Maximum Required Delay (Non-Outlet Aware UPS's) or the Outlet Group Power Off Delay.

13. After this delay, a further non-configurable two-minute delay is counted down.
14. UPS turns off after the user-configurable Shutdown Delay time has elapsed or the Outlet Group turns off after the power off Delay elapses.

NOTE: In a vSAN configuration, Witness and Management hosts will get placed into Maintenance Mode and shut down after Cluster hosts if Delay Maintenance Mode is enabled. vSAN Hosts are placed into Maintenance mode using “No data migration” for the vSAN data evacuation mode.

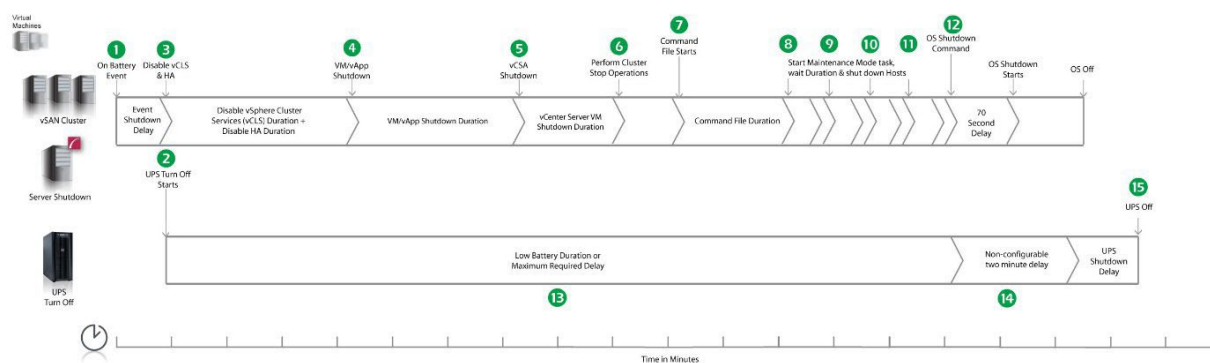
Example 2: Single UPS, 2-Node Stretch Cluster with Witness Appliance and Management Host, vCLS VMs in cluster, Disable HA on Shutdown enabled



Setup

- PowerChute is installed on a physical Windows machine.
- Single UPS configuration.
- Delay Maintenance Mode is enabled (required for vSAN hosts) with duration set to X seconds.
- [Disable vSphere Cluster Services \(vCLS\) Duration](#) set to 180 seconds.
- [Disable HA on Shutdown](#) enabled with a duration of 20 seconds.
- vCenter Server is running on a VM inside the vSAN Cluster.
- vSAN Witness Appliance is deployed on the Management Host and added to the inventory as a Host.
- [VM Prioritization](#) is enabled.
- An Active Directory controller VM is added to the High priority group with appropriate duration for shutdown and startup. This VM needs to be shut down after the vCenter Server VM and started before it.
- vCenter Server VM is added to the Medium priority group with appropriate duration for shutdown and startup.
- The option to turn off the UPS or Outlet Group is enabled.
- A shutdown command file has been configured with a 120 second duration.

When a critical UPS event, such as UPS on Battery occurs the following sequence is triggered:



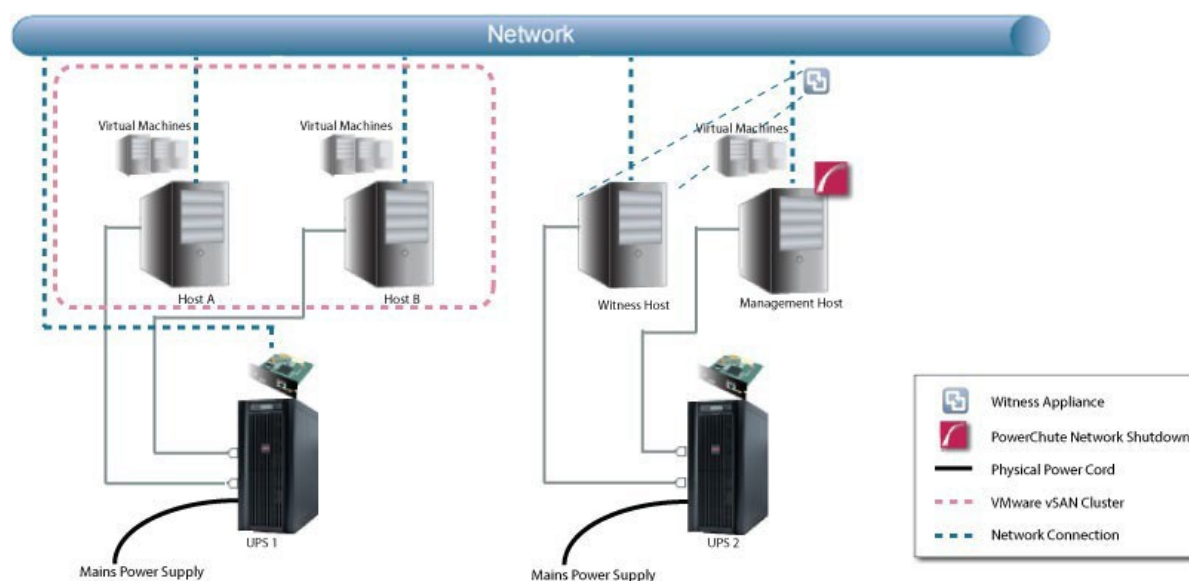
Shutdown Sequence

1. PowerChute reports that the UPS is on battery.
2. Shutdown delay for the On Battery event elapses. PowerChute sends a command to turn off the UPS or Outlet Group.
3. PowerChute disables vSphere Cluster Services (vCLS) and High Availability (HA) on the cluster.
4. After 3 minutes (Disable vSphere Cluster Services (vCLS) Duration = 180, Disable HA Duration = 20), PowerChute starts VM and vApp shutdown on each Host.
5. PowerChute gracefully shuts down the vCenter Server VM in the Medium priority group followed by the Active Directory controller VM in the High priority group during VM Shutdown.
6. PowerChute performs cluster stop operations:
 - a. PowerChute disables cluster member updates by executing the `esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates` SSH command on all protected vSAN nodes.
 - b. PowerChute prepares the vSAN cluster for shutdown by executing the `python /usr/lib/vmware/vsan/bin/reboot_helper.py prepare` command on 1 of the vSAN nodes.
7. PowerChute starts executing the shutdown command file.
8. Shutdown command file duration elapses and PowerChute starts a Maintenance Mode task on the first vSAN host, and waits the Delay Host Maintenance Mode Duration. If there is a vSAN synchronization active, PowerChute will wait the specified **vSAN Synchronization Duration** and check if it has completed with respect to the value set for "vsan_synch_retry_time" in the `pcnsconfig.ini` file until data re-synchronization is no longer active, or the retry limit has been reached. PowerChute then shuts down the host.
9. Once the Delay Host Maintenance Mode Duration has elapsed, PowerChute will start a Maintenance Mode task on the next vSAN host, and waits the Delay Host Maintenance Mode Duration. If there is a vSAN synchronization active, PowerChute will wait the specified **vSAN Synchronization Duration** and check if it has completed with respect to the value set for "vsan_synch_retry_time" in the `pcnsconfig.ini` file until data re-synchronization is no longer active, or the retry limit has been reached. PowerChute then shuts down the host.
10. PowerChute starts a Maintenance Mode task on the Witness Host, waits the Delay Host Maintenance Mode Duration and shuts down the host.
11. PowerChute starts a Maintenance Mode task on the Management Host, waits the Delay Host Maintenance Mode Duration of X seconds and shuts down the host.
12. OS shutdown sequence starts on the PowerChute physical machine. After a 70 second delay, the OS starts to shut down.

13. UPS waits for the duration that is greatest of Low Battery Duration/Maximum Required Delay (Non-Outlet Aware UPS's) or the Outlet Group Power Off Delay.
14. After this delay, a further non-configurable two-minute delay is counted down.
15. UPS turns off after the user-configurable Shutdown Delay time has elapsed or the Outlet Group turns off after the power off Delay elapses.

NOTE: In a vSAN configuration, Witness and Management hosts will get placed into Maintenance Mode and shut down after Cluster hosts if Delay Maintenance Mode is enabled. vSAN Hosts are placed into Maintenance mode using “No data migration” for the vSAN data evacuation mode.

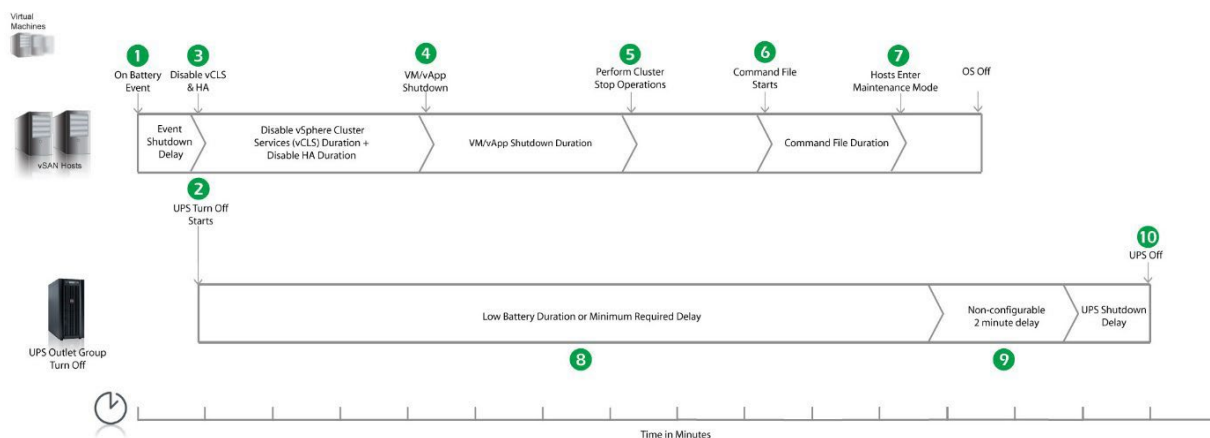
Example 3: Advanced UPS Configuration, 2 Node vSAN Cluster with Witness Appliance and Management Host



Setup

- Advanced configuration containing 2 single UPS setups:
 - UPS Setup 1: vSAN Host A, vSAN Host B.
 - UPS Setup 2: Witness Host and Management Host.
- Delay Maintenance Mode is enabled (required for vSAN hosts) with duration set to X seconds.
- **Execute Virtualization Shutdown on Hosts in all UPS Setups** is enabled on UPS #2 (the PowerChute host) and this setting is disabled on UPS #1.
- PowerChute virtual appliance is running on the Management Host.
- vCenter Server VM is added to the High priority group with appropriate duration for shutdown and startup.
- Fault Tolerance Threshold (FTT) is disabled.
- The option to turn off the UPS or Outlet Group is enabled.
- A shutdown command file has been configured with a 120 second duration.

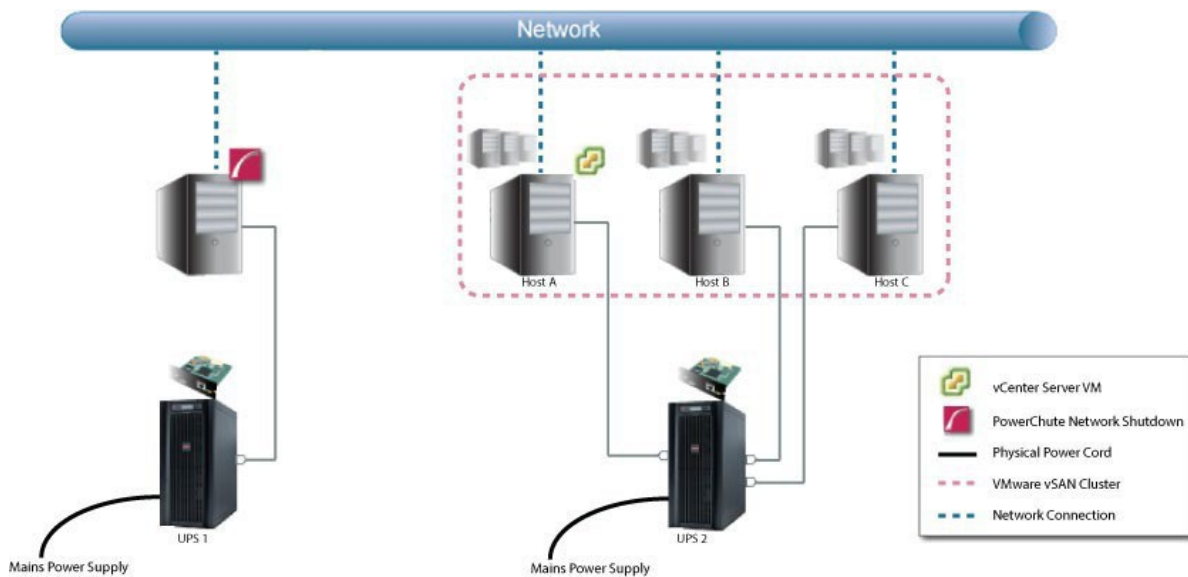
When a critical UPS event, such as UPS on Battery occurs on UPS #1, the following sequence is triggered:



Shutdown Sequence

1. PowerChute reports that the UPS #1 protecting the vSAN Cluster is On Battery.
2. Shutdown delay for the On Battery event elapses. PowerChute sends a command to turn off the UPS or Outlet Group.
3. PowerChute disables vSphere Cluster Services (vCLS) and High Availability.
4. After 3 minutes (Disable vSphere Cluster Services (vCLS) Duration = 180, Disable HA Duration = 20), PowerChute starts VM and vApp shutdown on each Host.
5. After 120 seconds (VM/vApp Shutdown Duration), PowerChute performs cluster stop operations:
 - a. PowerChute disables cluster member updates by executing the `esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates` SSH command on all protected vSAN nodes.
 - b. PowerChute prepares the vSAN cluster for shutdown by executing the python `/usr/lib/vmware/vsan/bin/reboot_helper.py prepare` command on 1 of the vSAN nodes.
6. After 120 seconds (Cluster stop operations delay), PowerChute starts executing the shutdown command file.
7. Shutdown command file duration elapses and all vSAN Cluster Hosts enter Maintenance Mode using No Action as FTT is disabled and the entire cluster is protected by 1 UPS.
8. UPS waits for the duration that is greatest of Low Battery Duration/Maximum Required Delay (Non-Outlet Aware UPS's) or the Outlet Group Power Off Delay.
9. After this delay, a further non-configurable two-minute delay is counted down.
10. UPS turns off after the user-configurable Shutdown Delay time has elapsed or the Outlet Group turns off after the power off Delay elapses.

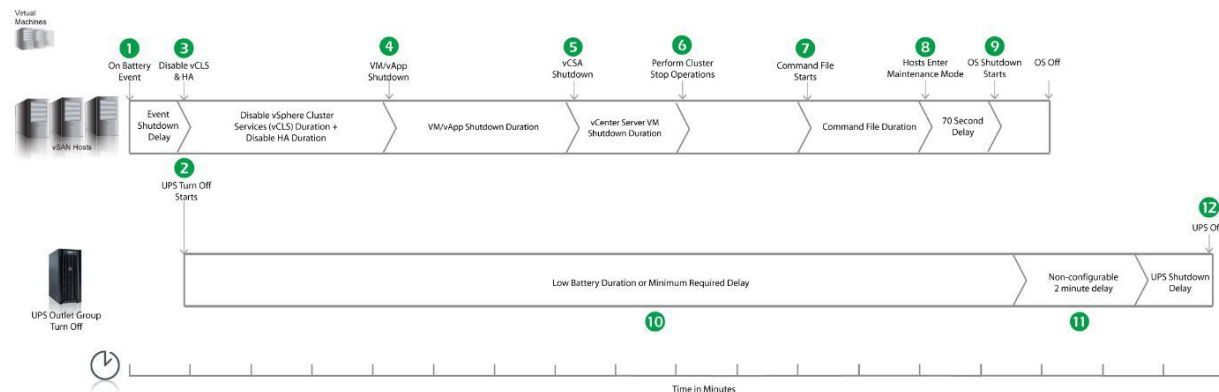
Example 4: Advanced UPS Configuration, 3 Node vSAN Standard Cluster



Setup

- Advanced configuration containing 2 single UPS setups:
 - UPS Setup 1: PowerChute is installed on a physical Windows machine.
 - UPS Setup 2: vSAN Host A, vSAN Host B, vSAN Host C.
- Delay Maintenance Mode is enabled (required for vSAN hosts) with duration set to X seconds.
- **Fault Tolerance Threshold** is disabled.
- vCenter Server is running on a VM inside the vSAN Cluster.
- vCenter Server VM is added to the High priority group with appropriate duration for shutdown and startup.
- The option to turn off the UPS or Outlet Group is enabled.
- A shutdown command file has been configured with a 120 second duration.
- **Execute Virtualization Shutdown on Hosts in all UPS Setups** is enabled on UPS #1 (the PowerChute host) and this setting is disabled on UPS #1.

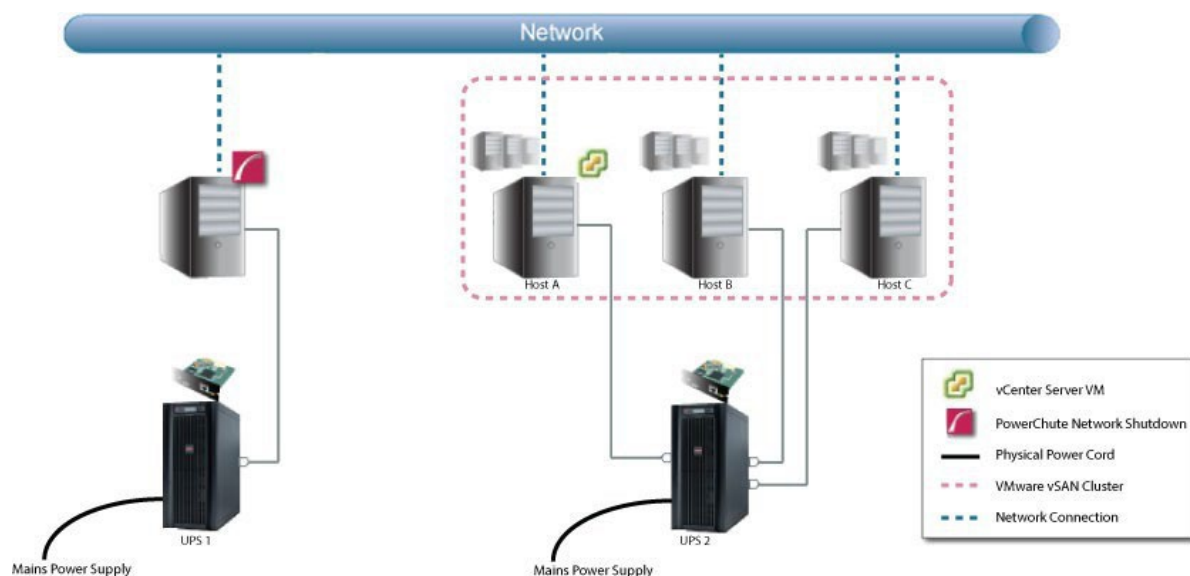
Shutdown Sequence




1. PowerChute reports that the UPS #2 protecting the vSAN Cluster is On Battery.
2. Shutdown delay for the On Battery event elapses. PowerChute sends a command to turn off the UPS or Outlet Group.
3. PowerChute disables vSphere Cluster Services (vCLS) and High Availability.
4. After 3 minutes (Disable vSphere Cluster Services (vCLS) Duration = 180, Disable HA Duration = 20), PowerChute starts VM and vApp shutdown on each Host.
5. After 120 seconds (VM/vApp Shutdown Duration), PowerChute gracefully shuts down the vCenter Server VM in the High priority group.
6. After 60 seconds (vCenter Server VM Duration), PowerChute performs cluster stop operations:
 - a. PowerChute disables cluster member updates by executing the `esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates` SSH command on all protected vSAN nodes.
 - b. PowerChute prepares the vSAN cluster for shutdown by executing the python `/usr/lib/vmware/vsan/bin/reboot_helper.py prepare` command on 1 of the vSAN nodes.
7. After 120 seconds (Cluster stop operations delay), PowerChute starts executing the shutdown command file.
8. Shutdown command file duration elapses and all vSAN Cluster Hosts enter Maintenance Mode using No Action as FTT is disabled and the entire cluster is protected by 1 UPS.
9. The OS Shutdown Command is issued and an additional 70 second delay is counted down before the operating system on the physical machine running PowerChute starts to shut down.
10. UPS waits for the duration that is greatest of Low Battery Duration/Maximum Required Delay (Non-Outlet Aware UPS's) or the Outlet Group Power Off Delay.
11. After this delay, a further non-configurable two-minute delay is counted down.
12. UPS turns off after the user-configurable Shutdown Delay time has elapsed or the Outlet Group turns off after the power off Delay elapses.

VMware Shutdown - Advanced UPS Configuration

Here, UPS #1 is protecting the physical PowerChute machine, and UPS #2 is protecting the 3-node Cluster, and vCenter Sever VM is installed on Host A.



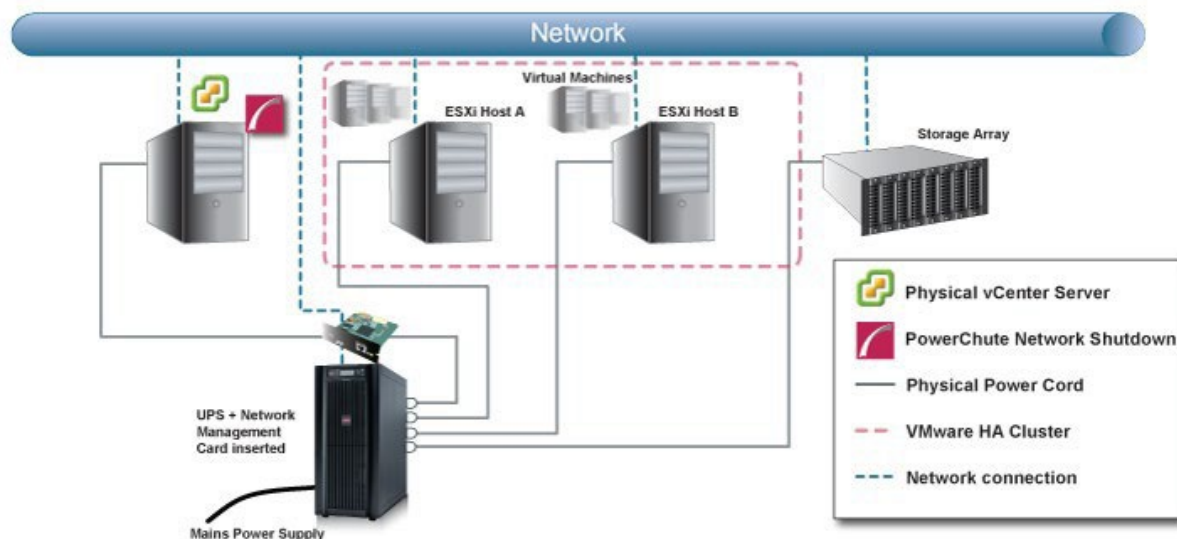
 In PowerChute v5.1, all Cluster Hosts must be part of the same UPS setup and **Execute Virtualization Shutdown on Hosts in all UPS Setups** must be enabled when vSphere Cluster Services (vCLS) is present.

Shutdown Sequence

1. UPS #2 goes on battery.
2. PowerChute issues a command to gracefully turn off UPS #4, if this has been configured.
3. The UPS Critical event is triggered for the three VMware hosts in the cluster.
4. PowerChute starts a maintenance mode task on the three VMware hosts, and shuts down the VMs/vApps on the three VMware hosts.
5. PowerChute shuts down the three VMware hosts.
6. As the physical machine is not affected, PowerChute continues to run.

VMware Shutdown - VM Prioritization - Single UPS Configuration

In this example, there are two VMware hosts, a vCenter Server and a storage array being powered by a single UPS. PowerChute is installed on the vCenter Server machine outside the cluster. [VM Prioritization](#) is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.



The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

1. The UPS has been running on Battery power for x number of seconds. PowerChute starts a maintenance mode task on each protected host.
2. PowerChute begins to shut down the VMs and vApps on VMware hosts A and B in the order in which they are prioritized:

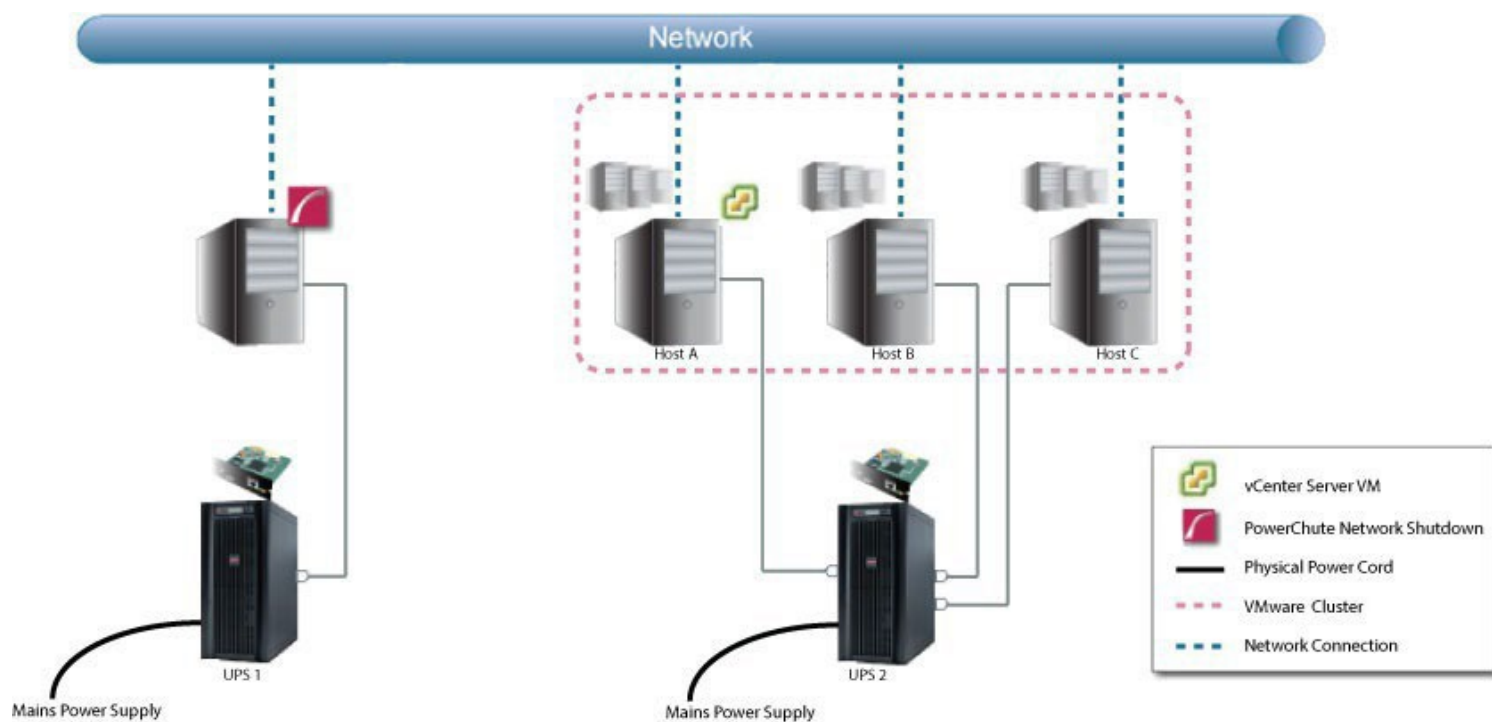
Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. PowerChute sends a shutdown command to all VMs/vApps in each priority group at the same time. The VMs/vApps within each priority group are not shut down in a particular order.

3. PowerChute runs the shutdown command file or SSH action and issues a command to turn off the UPS, if configured.
4. After the shutdown command file or SSH action duration has elapsed, PowerChute shuts down the VMware hosts.
5. PowerChute shuts down the vCenter Server Machine.

VMware Shutdown - VM Prioritization - Advanced UPS Configuration

In this example, separate UPS devices are powering the physical PowerChute server, and 3 VMware hosts in the cluster. [VM Prioritization](#) is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.

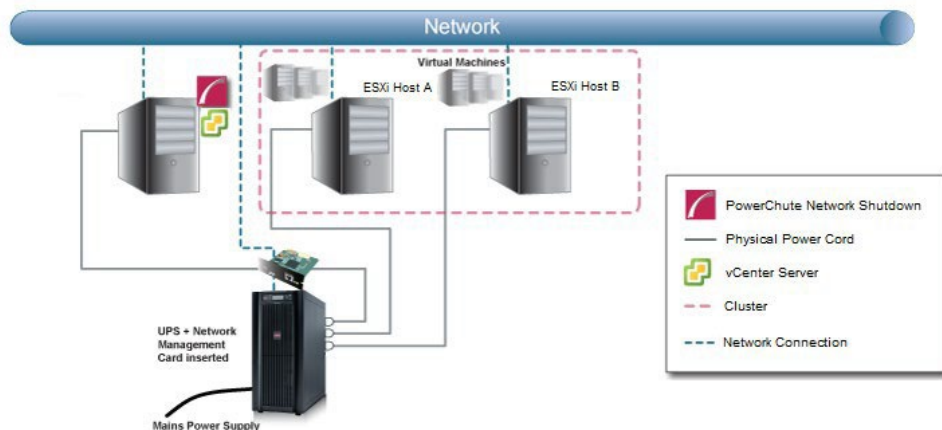


In an advanced UPS configuration in PowerChute v5.1, all Cluster Hosts must be part of the same UPS setup and **Execute Virtualization Shutdown on Hosts in all UPS Setups** should be enabled on the PowerChute host or physical UPS group (if PowerChute is installed on physical Windows machine) when vSphere Cluster Services (vCLS) is present.

If a critical event occurs on UPS #2 protecting the VMware cluster, all of the hosts in the cluster will be shut down and no VM migration or VM prioritization will take place.

Nutanix Shutdown - VM Prioritization - Single UPS Configuration

In this example, there are two VMware hosts, and a vCenter Server being powered by a single UPS. PowerChute is installed on the vCenter Server machine outside the cluster. Shutdown command files/SSH actions are configured to run after host shutdown. Delay Maintenance Mode is enabled, and VM Prioritization is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.



The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

1. The UPS has been running on Battery power for x number of seconds.
2. PowerChute issues a command to turn off the UPS, if configured.
3. PowerChute begins to shut down the VMs and vApps on VMware hosts A and B in the order in which they are prioritized:

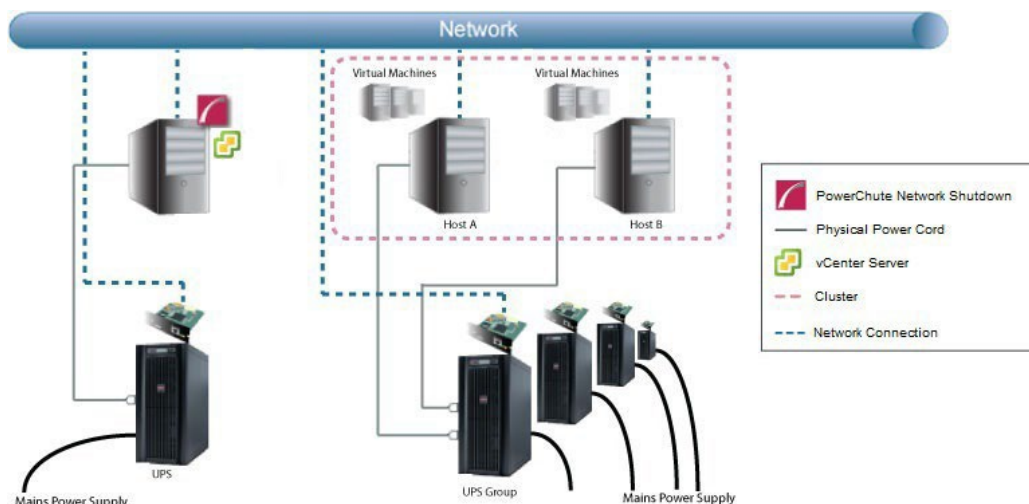
Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. The VMs/vApps within each priority group are not shut down in a particular order.

4. PowerChute shuts down Acropolis File Services.
5. PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
6. PowerChute shuts down the Nutanix Cluster.
7. PowerChute issues a command to shut down the Controller VMs.
8. PowerChute starts a maintenance mode task on each protected host.
9. PowerChute shuts down the VMware hosts.
10. PowerChute runs the shutdown command file or SSH action, if configured.
11. PowerChute shuts down the physical machine running vCenter Server and itself.
12. The UPS turns off after the shutdown time has elapsed.

Nutanix Shutdown - VM Prioritization - Advanced UPS Configuration

In this example, a UPS device is powering the vCenter Server, and a single UPS is powering the two Nutanix hosts in the Cluster. PowerChute is installed on the vCenter Server machine and is monitoring all UPS devices. Shutdown command files/SSH actions are configured to run after host shutdown. Delay Maintenance Mode is enabled, and VM Prioritization is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.



Critical event on UPS Group: Option to shut down virtual hosts is enabled for this UPS. VM Migration and VM Prioritization are enabled.

1. A UPS in the UPS group goes on battery.
2. A UPS critical event is triggered for Host A and Host B.
3. PowerChute issues a command to turn off the UPS, if configured.
4. As there is no healthy host remaining in the Cluster, PowerChute cannot migrate the VMs. VM Migration is skipped and PowerChute proceeds to shut down the VMs/vApps on the critical hosts.
5. VMs/vApps will be shut down in the VM Shutdown sequence, in the order in which they are prioritized:


Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. The VMs/vApps within each priority group are not shut down in a particular order.

6. PowerChute shuts down Acropolis File Services.
7. PowerChute aborts any ongoing VM replications. If enabled in the Protection Domain Settings page, Metro Availability will also be disabled on your Cluster.
8. PowerChute shuts down the Nutanix Cluster.
9. PowerChute issues a command to shut down the Controller VMs.
10. PowerChute starts a maintenance mode task on each protected host.
11. PowerChute shuts down Host A and Host B.

HPE SimpliVity Shutdown - VM Prioritization - Advanced UPS Configuration

In this example, a UPS device is powering the external vCenter Server, and a UPS group is powering the two HPE SimpliVity hosts in the Cluster. PowerChute is installed on the vCenter Server machine and is monitoring all UPS devices. Shutdown command files/SSH actions are configured to run after host shutdown. Delay Maintenance Mode is enabled, and VM Prioritization is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.



Critical event on UPS Group: Option to shut down virtual hosts is enabled for this UPS. VM Prioritization is enabled.

1. A UPS in the UPS group goes on battery.
2. A UPS critical event is triggered for Host A and Host B.
3. PowerChute issues a command to turn off the UPS, if configured.
4. VMs/vApps will be shut down in the VM Shutdown sequence, in the order in which they are prioritized:

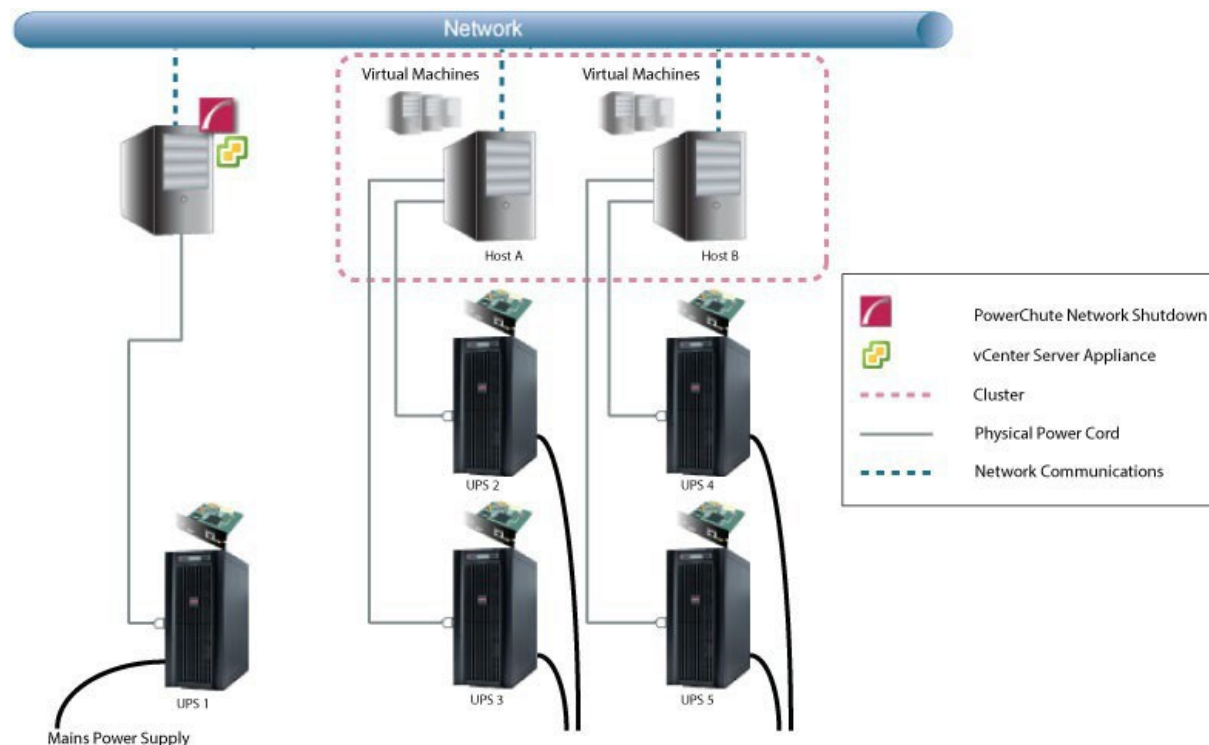
Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. The VMs/vApps within each priority group are not shut down in a particular order.

5. PowerChute issues a command to shut down the Controller VMs.
6. PowerChute starts a maintenance mode task on the hosts.
7. PowerChute shuts down Host A and Host B.
8. PowerChute runs the shutdown command file or SSH action, if configured.
9. The UPS turns off after the shutdown time has elapsed.

HyperFlex Shutdown - VM Prioritization - Single UPS Configuration

In this example, there are two HyperFlex nodes, and an external vCenter Server being powered by a single UPS. PowerChute is installed on the vCenter Server machine outside the cluster. Shutdown command files/SSH actions are configured to run after host shutdown. Delay Maintenance Mode is enabled, and VM Prioritization is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.



The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

1. The UPS has been running on Battery power for x number of seconds.
2. PowerChute issues a command to turn off the UPS, if configured.
3. PowerChute begins to shut down the VMs and vApps on VMware hosts A and B in the order in which they are prioritized:

Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

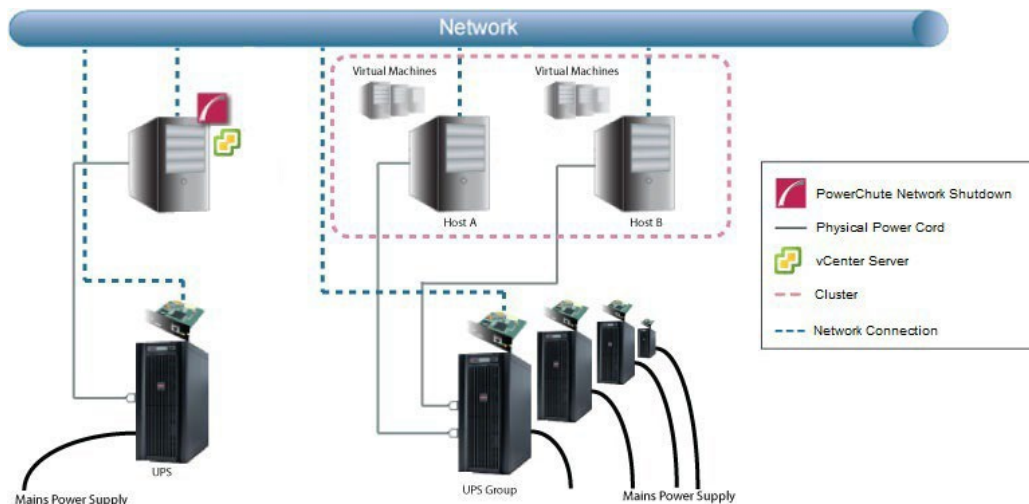
First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. The VMs/vApps within each priority group are not shut down in a particular order.

4. PowerChute shuts down the HyperFlex Cluster.
5. PowerChute issues a command to shut down the Controller VMs.
6. PowerChute starts a maintenance mode task on each protected host.
7. PowerChute shuts down the VMware hosts.

8. PowerChute runs the shutdown command file or SSH action, if configured.
9. PowerChute shuts down the physical machine running vCenter Server and itself.
10. The UPS turns off after the shutdown time has elapsed.

HyperFlex Shutdown - VM Prioritization - Advanced UPS Configuration

In this example, a UPS device is powering the external vCenter Server, and a UPS group is powering the two HyperFlex hosts in the Cluster. PowerChute is installed on the vCenter Server machine and is monitoring all UPS devices. Shutdown command files/SSH actions are configured to run after host shutdown. Delay Maintenance Mode is enabled, and VM Prioritization is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.



Critical event on UPS Group: Option to shut down virtual hosts is enabled for this UPS. VM Prioritization is enabled.

1. A UPS in the UPS group goes on battery.
2. A UPS critical event is triggered for Host A and Host B.
3. PowerChute issues a command to turn off the UPS, if configured.
4. VMs/vApps will be shut down in the VM Shutdown sequence, in the order in which they are prioritized:

Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

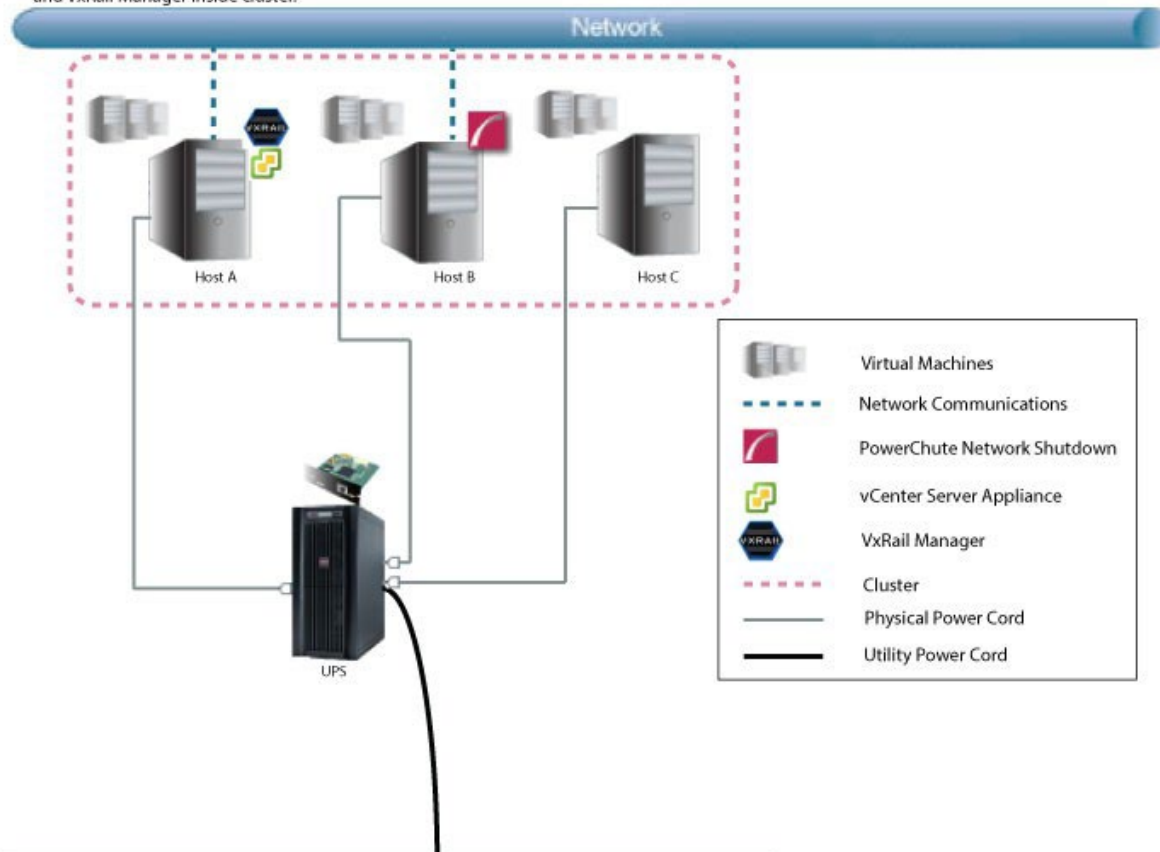
First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. The VMs/vApps within each priority group are not shut down in a particular order.

5. PowerChute shuts down the HyperFlex Cluster.
6. PowerChute issues a command to shut down the Controller VMs.
7. PowerChute starts a maintenance mode task on each protected host.
8. PowerChute shuts down Host A and Host B.
9. PowerChute runs the shutdown command file or SSH action, if configured.
10. The UPS turns off after the shutdown time has elapsed.

Dell VxRail Shutdown - VM Prioritization - Single UPS Configuration

In this example, there are three Dell VxRail nodes, one with PowerChute installed, and one with an internal vCenter Server and VxRail Manager. Shutdown command files/SSH actions are configured to run after host shutdown. VM Prioritization is enabled and VMs/vApps are prioritized into High, Medium, Low, Group 1, Group 2 priority groups.

Dell VxRail Single UPS Configuration: A 3-node Dell VxRail Cluster with PowerChute deployed as VM inside cluster, and vCenter Server and VxRail Manager inside cluster.



Power Transmission System

The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

1. A UPS has been running on Battery power for x number of seconds.
2. PowerChute issues a command to turn off the UPS, if configured.
3. PowerChute begins to shut down the VMs and vApps on VMware hosts A and B in the order in which they are prioritized:

Un-prioritized > ★ Group 2 > ★ Group 1 > ★ Low > ★ Medium > ★ High

First, the un-prioritized VMs/vApps are shut down sequentially. As the duration for un-prioritized VMs/vApps elapses, the Group 2 VMs/vApps are shut down, followed by Group 1 priority VMs/vApps, then Low priority VMs/vApps, and Medium priority VMs/vApps, and finally the High priority VMs/vApps are shut down. The VMs/vApps within each priority group are not shut down in a particular order.

4. PowerChute runs the shutdown command file or SSH action, if configured.
5. PowerChute sends a request to the Network Management Card to shut down the cluster using the VxRail REST API.

6. PowerChute starts a local OS shutdown to power itself off.
7. The NMC issues the VxRail cluster API shutdown command. At this point, all user VMs including the PowerChute VM are off.
8. The cluster and ESXi hosts are shut down.
9. The UPS turns off after the shutdown time has elapsed.

PowerChute Events and Logging

The Event Log displays UPS events that affect PowerChute and the load that it is protecting. Not all UPS events are logged. The log is refreshed automatically every 30 seconds.

By default, event logging is enabled for all configurable and non-configurable PowerChute events. To disable logging of an event, use the **Configure Events** screen.

The **EventLog.txt** file is located in the **group1** folder where PowerChute is installed. When the file reaches 1000 log entries, the oldest third of the file is deleted.

1000 is the default value, but you can change it using the PowerChute Configuration (INI) File. To do this:

1. Stop the PowerChute service/daemon. For more information, see Knowledge Base article [FA290624](#).
2. Locate the **pcnsconfig.ini** file in the **group1** folder where PowerChute is installed and open it using a text editor.
3. In the section **[EventLog]** change the value for `logsize` to the desired value. For example, to change the value to 2000 entries, change `logsize` to:

```
logsize = 2000
```
4. Save the `pcnsconfig.ini` file.
5. Restart the PowerChute service/daemon.



To completely clear the Event Log, use the **Delete Log File** button. Use **Export Log** to download a copy of the Event Log as a text file.

Configurable Events

- Available runtime has been exceeded

For both conditions below, the “total shutdown time” includes the following durations:

- VM migration duration
- VM shutdown and startup duration
- vApp shutdown and startup duration
- vCenter Server VM shutdown duration
- Disable vCLS duration
- Disable HA duration
- AFS shutdown and startup duration - Nutanix only
- Abort Active Replications duration - Nutanix only
- Cluster shutdown and startup duration - Nutanix, HyperFlex, and Dell VxRail only
- Controller VM shutdown and startup duration - Nutanix and HyperFlex only
- OVC shutdown and startup duration - SimpliVity only
- Maintenance mode duration for Dell VxRail stretched cluster
- Shutdown command file duration
- SSH action duration
- Built-in duration of 2 minutes (this consists of a 10 second OS shutdown duration and a 60 second OS shutdown duration; rounded up)

This event occurs with either of the following conditions:

Condition 1.

When the total shutdown time required by PowerChute is greater than the Low Battery Duration minus two minutes configured for the UPS. In the event of a low battery condition, PowerChute will not have enough time to complete the shutdown sequence before the UPS powers off. For example, if the total shutdown time required is 3 minutes and Low Battery Duration is 4 minutes, the Available Runtime has been Exceeded event will be triggered.

Resolution: Increase the Low Battery Duration value on the NMC using Configure - Shutdown or decrease the shutdown durations being used by PowerChute.

Condition 2.

When the shutdown duration configured for the UPS On Battery event plus the total shutdown time required by PowerChute is greater than the Runtime Remaining on the UPS minus two minutes. This condition can be caused by having too great a load on the UPS when the battery is fully charged.

Resolution:

1. Remove some equipment from the UPS to increase the available runtime.
2. Decrease the shutdown duration time for the UPS On Battery event.
3. Decrease the command file execution time using the Shutdown Settings screen.

This event is logged and event actions are carried out even if it occurs on a single UPS in a Redundant or Parallel UPS configuration.

- Available runtime is sufficient

The available UPS Runtime/ Low Battery Duration is sufficient for PowerChute to shut down all equipment gracefully.

- Battery is discharged

The UPS battery runtime has fallen below an acceptable range. If there is a power outage, a low battery condition will occur. This can be caused if the UPS has been operating on battery for an extended time period.

If a Battery Recharged event does not occur within four hours, the UPS may not be charging properly, please contact [APC Customer Support](#).

- Battery has recharged.

The battery runtime of the UPS has returned to within an acceptable range.

- UPS in Bypass due to an internal hardware problem or UPS overload.

The UPS has switched to bypass due to an internal hardware problem or because the UPS is overloaded.

- UPS has switched to bypass in response to the bypass switch at the UPS, typically for maintenance.

A user put the UPS into bypass mode using a hardware switch.

- UPS has switched to bypass in response to the UPS front-panel or a user-initiated software command, typically for maintenance.

The UPS has switched to bypass and cannot protect the load if a power outage occurs. This is a normal condition if maintenance is being performed on the UPS.

If this event occurs when the UPS was not deliberately put into bypass, please contact [Customer Support](#).

- UPS is no longer in Bypass.

The UPS is no longer in a bypass state.

- Bypass switch is not working properly.

The bypass contactor is not operating properly. This will prevent the UPS from being placed in bypass or returning from bypass. Please contact [Customer Support](#).

- Bypass switch has been replaced. Communication has been lost while on battery.

The bypass contactor is now operating properly.

- Communication has been lost while on battery.

PowerChute lost communication while the UPS was on battery and cannot detect a Low Battery condition if the power outage continues. Graceful shutdown cannot be guaranteed.

This occurs when the UPS is on battery and:

- The Management Card cannot communicate with the UPS
or
- PowerChute cannot communicate with the Management Card.

- Network Management Card cannot communicate with the UPS.

Communication between the NMC and the UPS has been lost. Make sure that the NMC is firmly inserted in its slot. This can occur during a firmware upgrade of the NMC.

- PowerChute cannot communicate with the Network Management Card.

Network communication between PowerChute and the NMC has been lost. See [Network Management Card Troubleshooting](#). This can occur during a firmware upgrade of the NMC.

- Communication has been established.

Communication has been established between PowerChute and the NMC.

- UPS has switched to battery power.

The UPS has switched to battery operation due to a power outage. If you can't restore power to the UPS, do the following:

1. If there is no general power outage (i.e. if only this UPS has lost input power), check the building wiring and circuit breakers.
2. If this event occurs occasionally and briefly, check to see if equipment on the same electrical circuit as the UPS uses high power periodically.
3. This event can also be caused by poor power quality (i.e. power fluctuation). Decrease the sensitivity of the UPS through the NMC user interface.
4. If the condition persists, contact an electrician to analyze your utility power.

- UPS is no longer running on battery power or output power has been turned on.

The UPS is no longer running on battery power.

- The load has exceeded the user specified alarm threshold.

The load on your UPS has exceeded the maximum load threshold, set in the NMC user interface. Reduce the load on the UPS or upgrade to a device that can support the existing load.

- The load no longer exceeds the user specified alarm threshold.

The load on your UPS is no longer above the load threshold.

- Minimum redundancy lost.

The UPS has too great a load or there are not enough power modules operational to support the desired redundancy.

Check to see that all power modules are functioning properly and that the redundancy configuration is correct.

If the condition persists, contact [Customer Support](#).

- Minimum redundancy restored.

The UPS can now support the desired redundancy.

- Parallel redundancy lost.

The system has too great a load or there are not enough operational UPS's to support the desired redundancy level.

Check to see that all UPS's are functioning properly and that the redundancy configuration is correct.

If the condition persists, contact [Customer Support](#).

- Parallel redundancy restored.

The Parallel UPS system can now support the desired redundancy.

- The runtime remaining has dropped below the configured threshold while on Battery.

The runtime remaining has dropped below the configured threshold while on battery. You can configure this threshold using the shutdown action on the Configure Events page.

When the UPS is running on battery power and the runtime remaining on the UPS drops below the threshold, PowerChute will trigger a shutdown sequence. See "[Sequenced Server Shutdown](#)" for more information.

- The runtime remaining is now above the configured threshold or input power has been restored.

Occurs when the UPS runtime is greater than the user defined threshold or if the UPS is no longer running on battery power.

- UPS has overheated which can cause damage.

The UPS's internal temperature is too high. Make sure that there is at least one inch of clearance around the UPS, and that the UPS ventilation ports are not blocked. If this condition is not resolved quickly, damage may occur to your UPS.

- UPS is no longer overheated.

The UPS's internal temperature has returned to an acceptable level.

- UPS output overload.

The UPS has sensed a load greater than 100 per cent of its rated capacity. Remove attached equipment from the UPS until the condition is corrected. If this condition happens occasionally and briefly, check to see if some equipment connected to the UPS is using high power periodically (e.g. connected laser printers or photocopiers). If the condition persists, contact [Customer Support](#).

- UPS overload condition has been corrected.

A condition that caused the UPS output overload event to occur has been corrected.

Configurable Environmental Events

The following events are logged and event actions are carried out even if they occur on a single UPS in a Redundant or Parallel UPS configuration.

- Ambient Temperature Out Of Range Probe X.

The temperature exceeds the threshold configured for the Environmental temperature probe.

- Ambient Temperature In Range Probe X.

The temperature no longer exceeds the threshold configured for the Environmental temperature probe.

- Humidity out of Range Probe X.

The humidity exceeds the threshold configured for the Environmental humidity probe.

- Humidity In Range Probe X.

The humidity no longer exceeds the threshold configured for the Environmental humidity probe.

- Communication lost with Environmental Monitor.

PowerChute has stopped receiving data from the Environmental Monitoring Card or the probe has been removed from the Universal I/O (UIO) port on the NMC.

Check to see that the Environmental Monitoring Card is firmly inserted in its slot and has power. Check that environmental monitoring information is accessible through the NMC user interface.

If PowerChute cannot communicate with the NMC you will need to correct that problem first.

- Communication established with the Environmental Monitor.

PowerChute Network Shutdown is receiving data from the Environmental Monitoring Card/Probe.

- Contact X Alarm.

One of the environmental input contacts is in an alarm state. Check in the location being monitored by this contact.

- Contact X Normal.

One of the environmental input contacts has returned to a normal condition.

Non-Configurable Events

- Three unsuccessful logon attempts detected. Temporarily denying logon attempts from machine with IP <IP address>.

There have been three invalid login attempts from a machine with the IP address listed in the event. Further login attempts will be prevented from this machine for two minutes. This is a security measure designed to prevent brute-force login attempts.

- Username was changed by user [User] from IP address [IP address]. New username is [Username].

The Username has been changed by the user at the specified IP address. This is a security feature to notify the user when the Username has been changed.

- Password was changed by user [User] from IP address [IP address].

The password has been changed by the user at the specified IP address. This is a security feature to notify the user when the password has been changed.

- Authentication phrase was changed by user [User] from IP address [IP address].

The authentication phrase has been changed by the user at the specified IP address. This is a security feature to notify the user when the authentication phrase has been changed.

- Low-battery condition occurred.

The runtime remaining on the UPS has dropped below the Low battery duration value while the UPS was on battery.

- UPS Turn off has been initiated.

A graceful shutdown command has been issued to the UPS using the NMC User Interface, the LCD display or by PowerChute. This event is logged for all UPS Configurations.

- PowerChute Network Shutdown version X monitoring started.

The PowerChute Web service has been started.

- Shutdown process started <OS name> will shut down soon.

The operating system has started to shut down in response to a critical UPS event.

- Error: Outlet Group X is turned off for NMC X.

The outlet group that PowerChute is registered with is turned off. This can indicate that PowerChute is not configured for the correct outlet group.

- Warning: Outlet Group X is turning off for NMC X.

The outlet group that PowerChute is registered with is shutting down. A shutdown sequence will be started as a

result.

- No Outlet Group specified. Using outlet group X.

If PowerChute was not registered with an Outlet group during setup it will be automatically registered with the first outlet group on the UPS by default.

- PowerChute is unable to open TCP port [number]. Check that TCP port [number] is free.

PowerChute uses TCP ports 3052 and 6547 for the Web User Interface. This event will be logged if another application is already using either of the above ports.

Use the netstat command to identify which process is using these ports or change the values using the PowerChute Configuration File.

- PowerChute is attempting to open TCP port [number]

PowerChute has begun to open the port it requires for the Web UI.

- PowerChute successfully opened TCP port [number].

PowerChute has successfully opened the port it requires for the Web UI.

- PowerChute is unable to open UDP port 3052. Check that UDP port 3052 is free. This is required for NMC communication.

PowerChute uses UDP port 3052 for communication with the NMC. This event will be logged if another application is already using this port. Use the netstat command to identify which application is using the port. This port cannot be changed.

- PowerChute is attempting to open UDP port 3052.

PowerChute has begun to open the port required for NMC communications.

- PowerChute successfully opened UDP port 3052.

PowerChute was able to open the port it needs for NMC communications.

- PowerChute cannot communicate with Network Management Card [ip_address]

Reported when the PowerChute Agent cannot communicate with the Network Management Card over the network. This could be due to a mismatch in security credentials or a network issue.

- Network Management Card [ip_address] cannot communicate with the UPS.

Reported when the Network Management Card cannot communicate with the UPS. If this issue persists please contact technical support.

- Connection unsuccessful because PowerChute received an untrusted SSL certificate from the NMC [protocol]://[ip_address]

This can occur if registering with an NMC that has HTTPS enabled and is using an SSL certificate that is not signed by a trusted root certification authority.

Ensure that trusted certificates are added to the **PowerChute-keystore** during registration. See [Certificate Management](#) for more information.

- PowerChute received an untrusted SSL certificate from the NMC `https://[ip_address]`.

Occurs when registering with an NMC that has HTTPS enabled if the SSL cert is not signed by a trusted root certification authority.
- PowerChute added a Network Management Card Self-Signed Certificate to the keystore.

PowerChute added an NMC SSL certificate to the **PowerChute-keystore**. See [Certificate Management](#) for more information.
- UPS `[ip_address]` is running on battery power

Reported when one UPS goes on battery in a UPS configuration with multiple UPS's.
- The On Battery UPS is no longer running on Battery power or output power has been turned on.

Reported when one UPS in a UPS configuration with multiple UPS's returns to On Line operation.
- Outlet on UPS is turning off / UPS is turning off.

The advanced option is enabled under UPS Shutdown Settings for a Redundant UPS Configuration and one UPS is on Battery.
- UPS `[ip_address]` has turned off.

Reported when one UPS turns off in a UPS configuration with multiple UPS's.
- The turned off UPS has switched to On Line operation.

Reported when one UPS turns back on in a UPS configuration with multiple UPS's.
- Multiple UPS's have been commanded to turn off / Outlet Group turn off has been initiated on Multiple UPS's.

Reported in a Redundant UPS Configuration.

In n+1 redundancy, this is reported when 2 UPS's turn off.

In n+2 redundancy, this is reported when 3 UPS's turn off.

In n+3 redundancy, this is reported when 4 UPS's turn off.
- Multiple UPS have turned off.

Reported in a Redundant UPS Configuration.

In n+1 redundancy, this is reported when 2 UPS's turn off.

In n+2 redundancy, this is reported when 3 UPS's turn off.

In n+3 redundancy, this is reported when 4 UPS's turn off.

- Multiple Critical Events occurred.

This occurs in a Redundant or Parallel-Redundant UPS Configuration when two different critical UPS events are active.

- Parallel-UPS Configuration not supported at address [ip_address].

One of the UPS devices in a Parallel-UPS configuration has been removed from the Parallel system.

- Turning off UPS [NMC IP Address].

PowerChute has sent a graceful shutdown command to the UPS. This is logged when a critical event occurs and the option to Turn off the UPS is enabled on the Shutdown settings page.

- Turning off outlet [Outlet Name] on UPS [NMC IP Address]

PowerChute has sent a graceful shutdown command to the UPS Outlet group. This is logged when a critical event occurs and the option to Turn off the UPS Outlet Group is enabled on the Shutdown settings page.

- SNMP[version]: New connection by user [User] from [IP Address].

A new user has connected to PowerChute via SNMP. This event is logged the first time a user connects after the PowerChute service restarts, or a SNMP setting is changed.

- SNMP[version]: Unsuccessful connection attempt by user [User] from [IP Address].

PowerChute detected a new user attempting to connect via SNMP. This event is logged the first time a user is unable to connect after the PowerChute service restarts, or a SNMP setting is changed.

NOTE: Some SNMP managers make unsuccessful attempts as part of their connection process. This will be indicated by the user "initial".

- SNMP: Configuration changed by user [User] from [IP Address]. [Config.ini Section].[Config.ini Key] set to [New Value].

A PowerChute setting has been changed by [User] via SNMP.

Configuration (INI) File Events

The table below lists events that may be logged as a result of manual changes to the PowerChute Configuration File.

See [PowerChute Configuration File](#).



Before editing the Configuration file manually you should save a backup copy locally.

- **Error: PowerChute cannot find the configuration file or the backup configuration file. Shutting down.**

PowerChute cannot locate pcsnconfig.ini or pcnsconfig_backup.ini to error.log in the group1 folder where PowerChute is installed. Please re-install PowerChute. If this does not resolve the issue contact APC technical support.
- **Error: The ini file is missing the required [x] section**

A required section is missing or incorrectly named.
- **Error: The ini file is missing [x] key from section [x].**

A required key is missing. Replace the missing key from a backup file.
- **Error: The ini file could not find IP address information in section [x].**

NMC IP addresses are missing from the [NetworkManagementCard] section.
- **Error: The ini contains an invalid value for [x] in section [x].**

An invalid value is present in the file and no previous good value or default is available in the backup file.
- **Error: The ini contains an invalid value for [x] in section [x]. Using {2} instead. Please validate the configuration.**

An invalid value is found in the file but a previous valid value or default value is available in the backup file. This should be checked but no further action may be needed.
- **Error: The key [x] in section [x], did not match the supplied regular expression.**

This can occur if you enter a username value that contains unsupported characters or if you entered a value other than http/https for the key protocol.
- **Error: Could not convert the value of [x] in section [x] to its expected type.**

This can occur if you enter a non-numeric value where a numeric value is expected for example.
- **Error: Event [x] is enabled for command file execution, but an invalid value for [x] is specified**

The command file specified cannot be found.

- The ini file has entries defined outside of a section.

There are extra entries outside of a section that PowerChute does not recognize. These can be deleted.

- The invalid key [x] should be deleted from section [x] in the ini file.

The configuration file contains keys that PowerChute does not recognize. These can be deleted.

- The ini file has detected duplicate values for [x] in section [x].

When this occurs PowerChute will use the first value and this may result in an incorrect value being used e.g. if you enter 2 values for the HTTP port (80 and 8080), PowerChute will use 80 instead of 8080.

- The invalid section [x] should be removed from the ini file.

The configuration file contains a section that PowerChute does not recognize. This can be deleted.

- Disabling command file execution for event [x] due to bad parameters. Please validate the configuration.

This can occur on a Linux/ Unix system if the path to the command file is valid but the file itself does not have execute permissions.

- Username was changed from [Username 1] to [Username 2] via ini file.

This notifies the user that the username has been changed via the ini file, for security purposes.

- Password was changed via ini file.

This notifies the user that the password has been changed via the ini file, for security purposes.

- Authentication phrase was changed via ini file.

This notifies the user that the authentication phrase has been modified via the ini file, for security purposes.

SSH Action Events

- Running SSH Action: [Action].

PowerChute is running the SSH action [Action] on the remote host.

- SSH Action [Action] has already run.

In an advanced UPS configuration, PowerChute runs each SSH action once for each host in the advanced group. This may result in multiple hosts attempting to run the same SSH action at the same time.

This event is shown if a SSH action has already run on a host in the advanced group.

- SSH Action [Action] has completed.

PowerChute has successfully completed executing the SSH action [Action] on the remote host.

- SSH Action [Action] has not completed within the configured duration.

The SSH action [Action] could not complete as insufficient time was configured. Ensure sufficient time is provided in the **SSH Action Duration** field in the [SSH Settings](#) screen for your SSH actions to complete.

Java Update Events

- Updating Java using [Java file].

PowerChute is attempting to update the Java version used by PowerChute to [Java file]. Allow 2-3 minutes for the Java update to complete.

- Java has successfully updated. Restarting PowerChute.

The Java version used with PowerChute has successfully updated. The PowerChute service will restart for changes to take effect.

- Unable to update Java. See error.log for details.

PowerChute was unable to update the Java version used with PowerChute. For more information, see error.log.

License Events

The events below are non-configurable and relate to activating, returning, and expired licenses.

- [Quantity] License(s) successfully activated.

Your PowerChute license(s) have successfully been activated. The **Quantity** associated with your **Activation ID** has been decreased by [Quantity] of licenses activated.

- License quantity has been reached. Purchase additional licenses through your Schneider Electric IT partner.

PowerChute is unable to validate your PowerChute license(s) as the **Quantity** associated with your **Activation ID** has been exceeded. You can purchase additional licenses through your Schneider Electric IT partner.

- Activation ID is incorrect. Please verify that the entered Activation ID is correct.

The provided **Activation ID** is incorrect and your PowerChute license(s) could not be activated. Please verify that the entered Activation ID is correct, and try again.

- Unable to connect to the Cloud Licensing Server. Please verify that you have Internet access and try again.

PowerChute was unable to connect to the Cloud Licensing Server and cannot validate your license or return your license. Please verify that you have Internet access, and try again.

- [Quantity] License(s) successfully returned.

Your PowerChute license(s) have successfully been returned. The **Quantity** associated with your **Activation ID** has been increased by [Quantity] of licenses returned.

- License will expire in [Days] days.

Your PowerChute license will expire in [Days] days. If you do not renew your license before the **Renewal Date**, you must purchase a new license and update your license details in the **License** screen.

- License is expired. Please purchase a new license.

Your PowerChute license has expired as the **Renewal Date** has been reached. PowerChute provides a 60-day grace period after the **Renewal Date** for you to purchase a new license. If you do not purchase a license during the grace period, PowerChute will no longer function.

Your PowerChute license has expired as the **Renewal Date** has been reached. PowerChute provides a 60-day grace period after the **Renewal Date** for you to purchase a new license. If you do not purchase a license during the grace period, PowerChute will no longer function.

- [Days] days remaining in the grace period. PowerChute will no longer function after this.

There are [Days] days remaining in the 60-day grace period. If you do not purchase a license during the grace period, you will no longer have access to configuration options in the PowerChute UI and the PowerChute Setup wizard - except the **License** screen - and shutdowns and startups will be canceled.

If you purchase a new PowerChute license during the grace period, update your license details in the **License** screen for continued PowerChute functionality.

There are [Days] days remaining in the 60-day grace period. If you do not purchase a license during the grace period, you will no longer have access to configuration options in the PowerChute UI and the PowerChute Setup wizard - except the **License** screen - and shutdowns and startups will be canceled.

If you purchase a new PowerChute license during the grace period, update your license details in the **License** screen for continued PowerChute functionality.

- License quantity updated, [Quantity] License(s) successfully activated.

Your license quantity has been updated, and [Quantity] licenses are successfully activated.

Your PowerChute license has expired as the **Renewal Date** has been reached. PowerChute provides a 60-day grace period after the **Renewal Date** for you to purchase a new license. If you do not purchase a license during the grace period, PowerChute will no longer function.

VMware Virtualization Events

The events below are non-configurable and relate to virtualization tasks such as VM shutdown.

- UPS critical event triggered a shutdown sequence on Host [Host]

A critical event has occurred on a UPS associated with [Host]. This will trigger a shutdown sequence using the actions configured on the Virtualization Settings page.
- DRS will attempt to migrate powered on VMs to another Host in the cluster.

When VM migration is enabled in PowerChute and DRS automation level is set to fully automated, PowerChute starts a maintenance mode task on the host and allows DRS to migrate VMs to available hosts in the cluster.
- PowerChute will attempt to migrate powered on VMs to another Host in the cluster.

When VM migration is enabled in PowerChute and DRS is disabled, PowerChute migrates VMs to available hosts in the cluster. If using the Custom Target Host migration option PowerChute will attempt to migrate to those target Hosts.
- Unable to find a suitable Host to migrate VMs from Host [Host].

PowerChute could not find any suitable host within the cluster to migrate the VMs from [Host] to.
- Insufficient time configured to migrate VMs from Host [Host].

PowerChute was unable to migrate VMs from [Host] as insufficient time was provided in the **Duration** field.
- Unable to migrate any VMs from Host [Host].

PowerChute was unable to migrate any VMs from [Host].
- Unable to migrate all VMs from Host [Host].

PowerChute successfully migrated some of the VMs from [Host]. However, not all VMs were migrated.
- No VMs require migration from Host [Host].

All powered on VMs have already been migrated on [Host].
- VMs have been successfully migrated from Host [Host].

A critical event has occurred on a UPS associated with [Host] and PowerChute has successfully migrated virtual machines from it.
- Starting Maintenance Mode task on Host [Host].

PowerChute started a maintenance mode task on [Host].
- Host [Host] has successfully entered Maintenance Mode.

PowerChute has successfully put [Host] into maintenance mode.

- Could not enter Maintenance Mode on Host [Host].

PowerChute was unable to put [Host] into maintenance mode.

- Unable to start Maintenance Mode on Host [Host] as VMs are still powered on. Verify that sufficient time has been configured for VM/vApp/VCSA VM shutdown duration.

A critical event has occurred on a UPS associated with this host and PowerChute was unable to put [Host] into maintenance mode as there are still powered on VMs.

- Maintenance Mode task cancelled on Host [Host] as there are still powered on VMs. Please verify that sufficient time has been configured for VM/vApp/VCSA VM shutdown duration.

A critical event has occurred on a UPS associated with this host and PowerChute has cancelled the maintenance mode task for [Host] as there are still powered on VMs prior to attempting to shut down the host.

- Maintenance Mode Task cancelled on Host [Host] as VCSA VM is shutting down.

PowerChute cancelled the maintenance mode task on [Host] as the vCenter Server Appliance VM is shutting down.

- Maintenance Mode task did not complete within the allowed time, please check your configuration.

A critical event has occurred on a UPS associated with this host and PowerChute could not complete the maintenance mode task for [Host] as the duration configured is insufficient.

- Exit Maintenance Mode on Host [Host].

PowerChute takes the Host out of maintenance mode when a critical event is resolved and the Host is powered back on.

- Migrating VMs on Host [Host] to another Host in cluster [Cluster].

A critical event has occurred on a UPS associated with the [Host] and PowerChute is trying to migrate virtual machines to another host in the [Cluster].

- Migration was not performed because Host [Host] is not part of any cluster.

PowerChute could not migrate the VMs on this host because the host is not part of any VMware cluster.

- Shutting down VMs on Host [Host].

A critical event has occurred on a UPS associated with this host and PowerChute is shutting down its virtual machines.

- Shutting down Host [Host].

PowerChute is now shutting down the host.

- Insufficient time configured to shut down VMs on Host [Host].

PowerChute was unable to shut down VMs from [Host] as insufficient time was provided in the **Shutdown Duration** field

- Shutting down vApp [vApp] in datacenter [Datacenter].
PowerChute is shutting down the specified vApp.
- Shutting down VMs belonging to vApp on Host [Host].
PowerChute is shutting down the VMs that are in the specified vApp.
- vApp shutdown unsuccessful due to timeout, please increase the vApp shutdown duration for a graceful shutdown.
PowerChute was unable to shut down [vApp] as insufficient time was configured for the **Shutdown Duration** field.
- Starting vApp shutdown process.
PowerChute is starting to shut down any vApps.
- Insufficient time configured to startup vApp [vApp]. Startup still in progress.
PowerChute was unable to re-start [vApp] as insufficient time was provided in the **Duration** field.
- No vApp to shut down on Host [Host].
A critical event has occurred on a UPS associated with [Host]. PowerChute has not found any vApp associated with [Host] to shut down.
- Powering on VMs on Host [Host].
A critical event has been resolved and a host has restarted. PowerChute is trying to start up the VMs on that host.
- Powering on vApp [vApp] in datacenter [Datacenter].
A critical event has been resolved and the host has restarted. PowerChute is starting up the specified vApp in the VMware datacenter.
- No Hosts have been associated with a UPS. PowerChute cannot shut down any Hosts or their VMs if a critical UPS event occurs.
PowerChute is advising you to link your VMware hosts to a UPS Setup. See [VMware Host Protection](#).
- Shutting down machine that PowerChute is running on.
The PowerChute machine is shutting down.
- Cannot connect to vCenter Server. PowerChute may not be able to issue commands to Virtual Machines or Hosts.
PowerChute cannot connect to vCenter Server over the network. VM Migration and vApp shutdown cannot be performed when this occurs.

- vCenter Server authentication error. PowerChute may not be able to issue commands to Virtual Machines or Hosts.

PowerChute cannot connect to vCenter Server using its credentials. Check that the username and password entered under [vCenter Server Settings](#) are correct. VM Migration and vApp shutdown cannot be performed when this occurs.

- Cannot connect to host. PowerChute may not be able to issue commands to the Host.

PowerChute cannot connect to the host over the network.

- Host authentication error. PowerChute may not be able to issue commands to the Host.

PowerChute cannot connect to the VMware host using its credentials. Check that the username and password entered under [vCenter Server Settings](#) are correct.

- Shutdown Host unsuccessful for Host [Host].

A critical event has occurred on a UPS associated with [Host] and PowerChute was unable to shut down [Host].

- Shutting down vCenter Server VM Host [Host].

PowerChute is shutting down the Host containing the vCenter Server VM.

- Attempting to power on VMs on Host [Host] that did not start.

PowerChute could not previously restart VMs on the specified host and is now trying again

- Attempting to power on vApp [vApp] in datacenter [Datacenter] that did not start.

PowerChute could not previously restart a specified vApp and is now trying again.

- vApp [vApp] will not be shut down as it contains the Virtual Machine running PowerChute. Please remove the PowerChute from the vApp.

PowerChute will not shut down the specified vApp as PowerChute is installed on a VM in that vApp.

The VM running PowerChute must be removed from the vApp.

- Cannot connect to vCenter Server. PowerChute will not be able to perform VM Migration.

PowerChute cannot connect to vCenter Server in order to perform VM Migration during a shutdown sequence.

- Cannot connect to vCenter Server. PowerChute will not be able to perform vApp Shutdown.

PowerChute cannot connect to vCenter Server in order to perform vApp Shut down during a shutdown sequence.

- Shutdown sequence is already in progress on VMware Hosts.

With an Advanced UPS configuration, when critical events occur on UPS Setups at different times, this event indicates that the shutdown is already in progress so it will not be started again.

For example: if a critical event occurs on a UPS Setup with VMware Hosts linked and later on a critical event occurs on a Physical UPS setup, the shutdown sequence will not be performed twice, as it is already in progress from the critical event on the first UPS setup.

- VM/vApp startup is in progress for Host [Host]. PowerChute will wait for the startup delay to elapse before starting the shutdown sequence.

A critical event has triggered a shutdown sequence, however as either a VM or a vApp startup is already in progress, PowerChute will wait for the specified duration to elapse before continuing with the shutdown sequence. See [Virtual Machine Shutdown/Startup](#).

- vCenter Server is accessible. PowerChute will be able to issue commands to Virtual Machines or Hosts.

PowerChute can now connect to the vCenter Server.

- Host is accessible. PowerChute will be able to issue commands to the Host.

PowerChute can now connect to the host.

- UPS critical event: [Event].

The specified critical event has been occurred on a UPS. This will trigger a shutdown sequence if not other shutdown sequence is currently active.

- UPS critical event: [Event] resolved.

The specified critical event has been resolved.

- UPS critical event: [Event] resolved on Host [Host].

The specified event has been resolved.

- UPS critical event resolved on Host [Host].

The event has been resolved.

- HA enabled for cluster [Cluster]. HA will attempt to re-start PowerChute if the Host on which it is running is shut down.

This is logged when PowerChute previously detected that HA was disabled.

- HA disabled for Cluster [Cluster]. PowerChute will not be re-started automatically by HA if the Host on which it is running is shut down.

The specified cluster is not a high availability cluster. Because of this, if the host containing PowerChute in this cluster is shut down, PowerChute cannot be restarted automatically. See also [HA Admission Control](#).

- vApp [vApp] in datacenter [Datacenter] will not be shut down because one or more VMs are running on a host unaffected by this critical UPS event.

This occurs if **Force vApp shutdown** is disabled.

- vApp [vApp] in datacenter [Datacenter] will not be shut down. The vApp is already powered off.

A critical event caused a shutdown but the specified vApp is already powered off.

- vApp [vApp] will not be shut down as it contains the vCenter Server VM. Please remove vCenter Server VM from the vApp.

PowerChute will not shut down the specified vApp as the vCenter Server VM is installed on a VM in that vApp.

The VM running the vCenter Server VM must be removed from the vApp.

- vCenter Server VM [VM] cannot be gracefully shut down. Please check vCenter Server VM Shutdown duration.

PowerChute is unable to gracefully shut down the VM.

Check that the VM Shutdown Duration is long enough to allow for VMs to be shut down gracefully.

- Shutting down vCenter Server VM [VM].

PowerChute is shutting down the VM running vCenter Server.

- Attempting to start vCenter Server VM [VM].

PowerChute is attempting to start the VM running vCenter Server.

- Host(s) [Host1], [Host2] no longer exist.

This event is logged when Hosts that PowerChute is configured to protect (via the [Host Protection Page](#)) are no longer present in the vCenter Server Inventory. When this occurs the old hosts will be removed and you need to update the PowerChute configuration using the Host Protection page.

- PowerChute cannot locate the vCenter Server VM in the Inventory. See the troubleshooting section in the Online Help.

This event is logged if PowerChute cannot determine which VMware host is running vCenter Server VM. This can indicate a vSphere Configuration issue and will prevent PowerChute from gracefully shutting down the Dell VxRail cluster. vCenter Server VM.

- The vCenter Server VM found in the Inventory is powered off. See the troubleshooting section in the Online Help.

This event is logged if PowerChute has identified the vCenter Server VM in the inventory but the VM is powered off. This can indicate a vSphere Configuration issue and will prevent PowerChute from gracefully shutting down the vCenter Server VM that PowerChute is configured for. the VxRail Cluster from being shut down.

- Detected vSAN Synchronization in progress on Host [Host], waiting duration.

Active data re-synchronization has been detected on the host with Delay Maintenance Mode enabled. PowerChute will wait and check for active data to complete with a retry limit before proceeding with host shutdown.

- Detected vSAN Resynchronization in progress on Host [Host]. Re-try limit reached, proceeding with Host shutdown.

PowerChute will proceed with host shutdown, as the retry limit has been reached.

- FTT exceeded triggered a shutdown of vSAN Cluster VMs on Host [Host].

Fault Tolerance Threshold (FTT) is enabled and has been exceeded. Shut down all Cluster VMs is also enabled and has triggered a shutdown of all VMs on non-critical hosts in the vSAN Cluster.

- Shutting down VMs in vSAN cluster due to FTT exceeded. PowerChute will wait for this to complete before proceeding with shutdown sequence.

A critical event has occurred on a critical host that is currently getting its VMs/vApps shut down in a separate shutdown sequence (as a non-critical host) due to Shut down all Cluster VMs being enabled.

- Starting Maintenance Mode Task on Host [Host] using data evacuation mode: [Mode].

A maintenance mode task has been started on a vSAN host using the data evacuation mode: **No Action** or **Ensure Accessibility**.

- Ensure Accessibility - vSAN data will be reconfigured to ensure storage object accessibility.
- No Action - No special action will be taken regarding vSAN data.

- Attempting to start vSAN Witness VM [VM].

PowerChute is attempting to start Witness Host VM.

- Disabling vSphere Cluster Services (vCLS) on cluster [cluster].

PowerChute is attempting to disable vSphere Cluster Services (vCLS) on the cluster.

- Unable to disable vSphere Cluster Services, vCLS VMs will not be deleted.

PowerChute was unable to disable vSphere Cluster Services (vCLS) within the specified **Disable vSphere Cluster Services (vCLS) Duration**. Ensure that vCenter Server is available.

- vSphere Cluster Services (vCLS) VMs have not gracefully been shut down in the allocated time.

PowerChute was unable to delete the vSphere Cluster Services (vCLS) VMs within the specified **Disable vSphere Cluster Services (vCLS) Duration**.

- Disabling High Availability (HA) on cluster [cluster].

PowerChute is attempting to disable High Availability (HA) on the cluster.

- Enabling vSphere Cluster Services (vCLS) on the cluster [cluster].

PowerChute is attempting to enable vSphere Cluster Services (vCLS) on the cluster.

- Unable to enable vSphere Cluster Services, vCLS VMs will not be created.

PowerChute was unable to enable vSphere Cluster Services (vCLS) and create the vCLS VMs on the host within the specified **Enable vSphere Cluster Services (vCLS) Duration**. Ensure that vCenter Server is available.

- Enabling High Availability (HA) on cluster [cluster].

PowerChute is attempting to enable High Availability (HA) on the cluster.

Hyperconverged Infrastructure Virtualization Events

The events below are non-configurable and relate to virtualization tasks such as cluster shutdown. The events listed here are common across the hyperconverged infrastructure (HCI) environments: Nutanix, HPE SimpliVity, and HyperFlex.

- Shutting down Cluster.

PowerChute is attempting to shut down the Cluster.

- Cluster shut down successfully.

PowerChute successfully shut down the Cluster.

- Cluster cannot be gracefully shut down.

PowerChute was unable to shut down the Cluster.

- Starting Cluster.

PowerChute is attempting to start up the Cluster.

- Cluster started successfully.

PowerChute successfully started up the Cluster.

- Cluster cannot be started.

PowerChute was unable to start up the Cluster.

- Shutting down Controller VM [CVM].

A critical event has occurred on the UPS powering the Cluster and PowerChute is shutting down all Controller Virtual Machines in the Cluster. Controller VMs are shut down after all other VMs in the Cluster, and the Cluster itself are shut down.

- Attempting to start Controller VM [CVM].

A critical event has been resolved and PowerChute is attempting to start up the Controller Virtual Machine [CVM]. This event is triggered after the host is online.

- Controller CVM [CVM] did not shut down gracefully.

PowerChute was unable to gracefully shut down Controller VM [CVM].

- Waiting for Controller VM startup to complete.

A critical event has been resolved and PowerChute is waiting for the Controller Virtual Machines in the Cluster to startup. The Cluster will start up when Controller VM startup is complete.

Nutanix Virtualization Events

The events below are non-configurable and relate to virtualization tasks such as VM shutdown.

- Shutting down Nutanix AFS.

A critical event has occurred on the UPS powering the Nutanix Cluster and PowerChute is shutting down the Acropolis File Server (AFS) VMs in the Cluster.

- Nutanix AFS cannot be gracefully shut down.

PowerChute was unable to shut down the Acropolis File Server (AFS) VMs in the Cluster.

- Starting Nutanix AFS.

A critical event has been resolved and PowerChute is trying to start up the Acropolis File Server (AFS) VMs in the Cluster.

- Nutanix AFS cannot be started.

PowerChute was unable to start up the Acropolis File Server (AFS) VMs in the Cluster.

- No replication in progress for Protection Domain. Proceeding with shutdown sequence.

PowerChute did not detect any active replications for a Protection Domain. This step of the shutdown sequence is skipped, and PowerChute continues to the next step, shutting down the Nutanix Cluster.

- Replication in progress for Protection Domain detected. Waiting the configured delay.

PowerChute detected active replications for a Protection Domain. PowerChute will wait the delay specified in the **Duration** field on the Protection Domain Settings screen before aborting replications.

- Aborting replication in progress for Nutanix Protection Domain.

PowerChute is aborting any active replications of your protection domain.

- Nutanix Protection Domain replications cannot be gracefully aborted.

PowerChute was unable to abort your active protection domain replications.

- Disabling Nutanix Metro Availability.

PowerChute starts a maintenance mode task on the host when a critical UPS event occurs.

PowerChute is attempting to disable Metro Availability.

- Nutanix Metro Availability cannot be disabled.

PowerChute could not disable Metro Availability on the Nutanix Cluster.

- Enabling Nutanix Metro Availability.

PowerChute starts a maintenance mode task on the host when a critical UPS event occurs.

PowerChute is attempting to enable Metro Availability.

- Nutanix Metro Availability cannot be enabled.

PowerChute could not enable Metro Availability on the Nutanix Cluster.

HPE SimpliVity Virtualization Events

The events below are non-configurable and relate to virtualization tasks carried out when shutting down a HPE SimpliVity cluster.

- Shutting down OmniStack Virtual Controller [OVC].

A critical event has occurred on the UPS powering the Cluster and PowerChute is shutting down the OmniStack Virtual Controller [OVC] running the Cluster. The OmniStack Virtual Controller is shut down after all other VMs in the Cluster.

- OmniStack Virtual Controller [OVC] did not shut down gracefully.

PowerChute was unable to gracefully shut down OmniStack Virtual Controller [OVC].

- Insufficient time configured for OVC shutdown. Please increase OVC Shutdown Duration.

PowerChute was unable to shut down the OmniStack Virtual Controller [OVC] as insufficient time was provided in the **OVC Shutdown Duration** field.

- Attempting to start OmniStack Virtual Controller [OVC].

A critical event has been resolved and PowerChute is attempting to start the OmniStack Virtual Controller [OVC]. This event is triggered after the host is online.

- OmniStack Virtual Controller [OVC] has not started.

PowerChute was unable to start the OmniStack Virtual Controller [OVC].

HyperFlex Virtualization Events

The events below are non-configurable and relate to virtualization tasks carried out when shutting down a HyperFlex cluster.

- Cannot connect to HyperFlex Rest service. PowerChute may not be able to issue commands to the cluster.

PowerChute cannot establish a connection with the HyperFlex Rest service and may not be able to perform cluster operations.

- Cluster shut down did not complete in time. PowerChute will wait until the cluster has successfully shut down.

PowerChute could not shut down the cluster within the **Shutdown Duration** specified. PowerChute will wait for the **Shutdown Duration** to elapse before attempting to shut down the cluster again with respect to the `cluster_ops_retries` value in the PowerChute configuration file. For more information, see [Cluster Shutdown/Startup](#).

- Cluster startup did not complete in time. PowerChute will wait until the cluster has successfully started.

PowerChute could not start the cluster within the **Startup Duration** specified. PowerChute will wait for the **Startup Duration** to elapse before attempting to start the cluster again with respect to the `cluster_ops_retries` value in the PowerChute configuration file. For more information, see [Cluster Shutdown/Startup](#).

Dell VxRail Virtualization Events

The events below are non-configurable and relate to virtualization tasks carried out when shutting down a VxRail cluster.

- Sending request to Network Management Card [NMC] to shut down Cluster.

PowerChute has sent a request to the Network Management Card to shut down the cluster using the VxRail REST API.

- Network Management Card [NMC] has initiated Cluster shut down.

The Network Management Card has waited for the specified **Cluster Shutdown Delay** to elapse and is shutting down the cluster.

- Network Management Card [NMC] is unable to initiate Cluster shut down.

The Network Management Card was unable to initiate cluster shut down. Ensure that the VxRail credentials provided are correct and that VxRail Manager is available. For more information, see [Dell VxRail Troubleshooting](#).

Critical Events in a Redundant-UPS Configuration



This topic does not apply to Advanced Configuration with Advanced UPS Setups. For more information about Critical Events in Advanced UPS Setups, please view the “Using PowerChute Network Shutdown in an Advanced Redundant Setup” Application Note [here](#).

PowerChute Network Shutdown considers all UPS devices in a Redundant configuration as one UPS System. Each UPS must be able to support the entire load itself.

PowerChute follows these shutdown rules when it detects critical events:

- 2 identical critical events (such as **Low-Battery Condition Occurred** OR **UPS turn off has been initiated**) occurring in succession on 2 UPS devices cause a shutdown. The shutdown is immediate and no configured delay is counted down.
- 2 identical user-configured critical events such as **UPS On Battery**
- **Low-Battery Condition Occurred** and **PowerChute cannot communicate with the Management Card** which always leads to a shutdown. Prior to the PowerChute shutdown process starting, a 10-second delay is counted. No configured delay time is counted down.

Critical Events in a Parallel-UPS Configuration

In Parallel-UPS configurations, the combined outputs of several UPS devices support the load. With this setup, PowerChute monitors the load as it changes to determine whether the mode of operation is Parallel Capacity or Parallel Redundant.

For example, you are operating in a Parallel Redundant mode (i.e., there are more UPS devices available than are required to provide power to the load) and then you increase the load by adding new servers. PowerChute detects if the mode of operation changes to Parallel Capacity (i.e., all UPS devices in the configuration are now required to provide power to the load). This could cause PowerChute to initiate a shutdown if just one critical event is triggered.

Scenario 1: Three 10kVA UPS devices supporting a 16kVA Server Load (Parallel Redundant)

In this Parallel Redundant configuration, two or more critical events occurring cause PowerChute to trigger a graceful shutdown of the server(s).

PowerChute follows these shutdown rules.

- 2 identical critical events (such as **Low-Battery Condition: Occurred** OR **UPS turn off has been initiated**) occurring in succession on 2 UPS devices cause a shutdown. The shutdown is immediate and no configured delay is counted down.
- 2 identical user-configured critical events such as **UPS: On Battery** occurring in succession on 2 UPS devices will cause a shutdown. Any configured delay is counted down first.
- 2 different critical events (such as **Low-Battery Condition: Occurred** and **PowerChute cannot communicate with the Management Card**) occurring in succession on 2 UPS devices cause the event called **Multiple Critical Events occurred** which always leads to a shutdown. Prior to the PowerChute shutdown process starting, a 10-second delay is counted. No configured delay time is counted down.

Scenario 2: Three 10kVA UPS devices supporting a 23kVA Server Load (Parallel Capacity)

In this Parallel Capacity configuration, one critical event triggers a graceful shutdown of the server(s).

However, if 2 occur, they have these delays.

- If 2 *identical* critical events occur in a parallel capacity configuration, then the event is only reported once and any configured delay is counted down.
- If 2 *different* critical events occur, then both events are reported separately and the shortest shutdown delay of the two is counted down.

Scenario 3: Two 10kVA UPS devices supporting an 8kVA Server Load (Parallel Capacity)

In this Parallel Capacity configuration, one critical event occurs and graceful shutdown of the server(s) is not triggered as the single UPS device can still support the connected load.

Scenario 4: Two 10kVA UPS devices supporting an 8kVA Server Load (Parallel Capacity)

In this Parallel Capacity configuration, two *identical* critical events occurring causes PowerChute to trigger a graceful shutdown of the server(s). The shutdown is immediate and no configured delay is counted down.

Logging

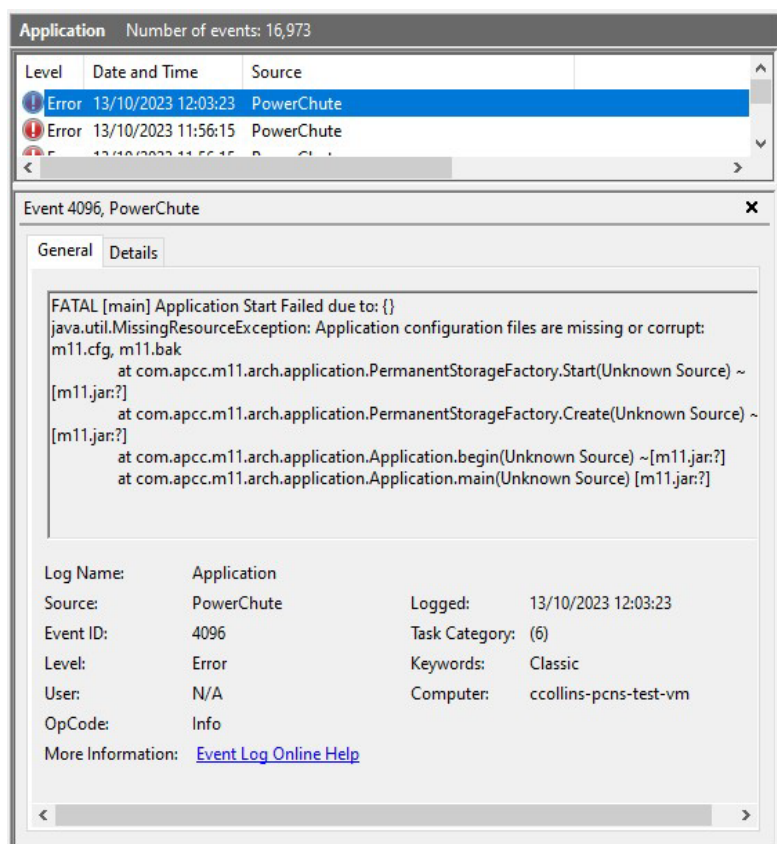
Fatal Errors

Microsoft Windows

On Windows operating systems, use Microsoft Event Viewer to view the messages written to the Operating System Event Log. PowerChute Logs are displayed under Windows Logs - Application. In the **Source** field, the “PowerChute” value is provided for PowerChute events.

The PowerChute service will not start if the configuration files are missing, or the service is prevented from accessing the files due to permissions.

The following image shows an example of a service start error:



Linux

On Linux operating systems, service logs are written to the system journal. To review all the service events related to PowerChute, use the following command: `journalctl -u PowerChute.service`

To review the most recent service action, use the following command as an alternative: `systemctl status PowerChute`

The service indicates “inactive (dead)” status, when the service has stopped, and the context provides additional information about why the service has ended. Additional logs in error.log file will display information on unexpected service halts.

```

root@linuxmint: /opt/APC/PowerChute/group1
File Edit View Search Terminal Help
root@linuxmint: /opt/APC/PowerChute/group1# systemctl status PowerChute
○ PowerChute.service - PowerChute Network Shutdown
  Loaded: loaded (/etc/systemd/system/PowerChute.service; enabled; vendor preset: enabled)
  Active: inactive (dead) since Thu 2024-01-11 16:20:44 GMT; 7s ago
  Process: 2237 ExecStart=/usr/bin/PowerChute start (code=exited, status=0/SUCCESS)
  Process: 2292 ExecStop=/usr/bin/PowerChute stop (code=exited, status=0/SUCCESS)
  CPU: 2.475s

Jan 11 16:20:43 linuxmint systemd[1]: Starting PowerChute Network Shutdown...
Jan 11 16:20:43 linuxmint PowerChute[2237]: PowerChute Network Shutdown, v5.1.0
Jan 11 16:20:43 linuxmint PowerChute[2237]: Copyright (c) 1999-2023, Schneider Electric. All Rights Reserved.
Jan 11 16:20:43 linuxmint PowerChute[2237]: Startup completed.
Jan 11 16:20:43 linuxmint systemd[1]: Started PowerChute Network Shutdown.
Jan 11 16:20:44 linuxmint PowerChute[2292]: Killing PowerChute Network Shutdown Group #1. PID= 2247
Jan 11 16:20:44 linuxmint PowerChute[2292]: /usr/bin/PowerChute: 111: kill: No such process
Jan 11 16:20:44 linuxmint PowerChute[2292]: Done.
Jan 11 16:20:44 linuxmint systemd[1]: PowerChute.service: Deactivated successfully.
Jan 11 16:20:44 linuxmint systemd[1]: PowerChute.service: Consumed 2.475s CPU time.
root@linuxmint: /opt/APC/PowerChute/group1#

```

Non-Fatal Errors

Error conditions that are significant to operators but do not cause the service to stop are logged in the error.log. Examples include network conditions preventing connection to VCenter, controller VMs, cluster hosts, and other collaborating systems. Operators are advised to review logs periodically to ensure PowerChute configuration and network conditions enable PowerChute to react to UPS events.

Logging to Third Party Systems

PowerChute manages logging through [Log4J](#). This is configured by the configuration file `log4j2.xml` in the installation directory.

Windows Event Log

On Windows installations, the Event Log records the below events. You can view these events with Windows Event Viewer in the Application Log.

1. UPS Events

- All events enabled for logging are reported to the Windows Event Log.

2. Service Fatalities

- Any event causing PowerChute to stop running.

3. Unsuccessful Command File execution 4. Unsuccessful SSH Action execution

PowerChute issues the Command File and SSH Action events at pre-defined stages of the startup and shutdown procedure and it considers them successful once they begin running. Warning or error events are not logged for command files or SSH Actions which run for extended periods, or for scripts which have errors during execution.

Syslog (Insecure)

PowerChute uses [Log4J SyslogAppender](#) to direct logs to an appropriately configured syslog server.

The following configuration is provided in the `log4j2.xml` file, in the installation directory.

```
<Syslog name="syslog" format="RFC5424" host="localhost" port="514" protocol="UDP"
appName="PowerChute" facility="ALERT" enterpriseNumber="1.3.6.1.4.1.318"
messageId="PowerChute"/>
```

This configuration will forward log events to a syslog server over UDP to port 514. This transport mechanism is considered insecure because data on the network is not encrypted, and the network endpoints do not mutually authenticate each other. This configuration is not recommended for production deployments.

Configure Insecure Syslog

1. Modify the **host** setting and specify the syslog server host details, with the IP address for Fully Qualified Domain Name.
2. Modify the **port** setting to specify the syslog server input port.
3. Modify the **protocol** setting to specify TCP or UDP transport.

Syslog (Secure)

PowerChute uses [Log4J SyslogAppender](#) to direct logs to an appropriately configured syslog server.

The following configuration is provided in the `log4j2.xml` file, in the installation directory.

```
<Syslog name="syslog_sec" format="RFC5424" host="localhost" port="6541" protocol="TLS"
appName="PowerChute" facility="ALERT" enterpriseNumber="1.3.6.1.4.1.318"
messageId="PowerChute">
  <SSL>
    <KeyStore location="/path/to/syslog_keystore.jks" type="PKCS12"
      password="apassword" />
    <TrustStore location="/path/to/syslog_truststore.jks" type="PKCS12"
      password="apassword" /> </SSL>
</Syslog>
```

This transport mechanism is considered secure because the data transported over the network is protected by TLS, and network endpoints can mutually authenticate each other.

Configure Secure Syslog

1. Modify the **host** setting and specify the syslog server host details, with the IP address or Fully Qualified Domain Name.
2. Modify the **port** setting to specify the syslog server input port.
3. Modify the KeyStore and TrustStore file settings.

KeyStore and TrustStore file for Secure Syslog

The PowerChute KeyStore contains a certificate which is presented to the server to establish trust in the connecting client. The client certificate should be countersigned by your trusted Certificate Authority. The TrustStore file contains the syslog server certificate from the Certificate Authority and enables PowerChute to authenticate the server's certificate. For information on the passwords required for the KeyStore and the TrustStore, see [Log4J SSL configuration manual](#).

Create a KeyStore

1. To create a KeyStore, issue the following commands:

```
keytool -keystore syslog_keystore.jks -genkey -keysize 4096 keyalg RSA -alias client
```

2. Enter the KeyStore password when prompted and follow the prompts to set the certificate details for your location and organization. Ensure that the certificate **Common Name** contains the **Fully Qualified Domain Name** for the PowerChute host.

3. Generate a Certificate Signing Request from your new KeyStore:

```
keytool -certreq -alias client -file client.csr -keystore syslog_keystore.jks
```

4. Enter the KeyStore password when prompted. This will generate a certificate signing request file to `client.csr`. Use the certificate signing request to sign the client certificate with your trusted Root Certificate Authority certificate.

5. Import the public root Certificate Authority certificate into the KeyStore:

```
keytool -keystore syslog_keystore.jks -import -file rootCA.crt alias root-ca
```

6. Import the signed client certificate into the KeyStore:

```
keytool -keystore syslog_keystore.jks -import -file client.crt alias client
```

7. Copy the `syslog_keystore.jks` file into the PowerChute installation directory.
8. Modify the `log4j2.xml` configuration file to specify the full path to the file in the KeyStore location setting.
9. Modify the password setting to contain the KeyStore password.

Create a TrustStore file

1. To create a TrustStore file, issue the following commands:

```
keytool -keystore syslog_truststore.jks -genkey -alias deleteme
```

2. When prompted, provide the values for a new certificate in the file and create a password for the TrustStore.

3. Delete the certificate to create an empty KeyStore:

```
keytool -keystore syslog_truststore.jks -delete -alias deleteme
```

4. Import the Certificate Authority public certificate into the KeyStore:

```
Keytool -keystore syslog_truststore.jks -import -file  
/path/to/rootCA.crt -alias syslog-ca
```

5. Copy the `syslog_truststore.jks` file to the PowerChute installation directory
6. Modify the `log4j2.xml` configuration file to specify the full path to the file in the TrustStore file location setting.
7. Modify the password setting to contain the TrustStore file password.

General

This section contains information on the topics below:

- [Communications Settings](#)
- [PowerChute Agents](#)
- [PowerChute Configuration \(INI\) File](#)
- [Resetting your PowerChute username or password](#)
- [Java Update](#)
- [User Interface Session Timeout](#)
- [Check for Updates](#)
- [Customer Support](#)
- [Customer Experience Improvement Program \(CEIP\)](#)
- [License](#)

Communications Settings

PowerChute Access

HTTPS is enabled by default and provides secure access to the PowerChute user interface. You may change the Protocol to HTTP (unencrypted) and this will come into effect after you restart the PowerChute service. For more information, please see the PowerChute Security Handbook [here](#).

PowerChute Security

The Username and Authentication Phrase are used to authenticate communications between PowerChute and the NMC. Therefore, you must set these values to be the same in both PowerChute and the NMC.

- The maximum number of characters for the user name is 10.
- The authentication phrase must be 15 to 32 ASCII characters.
- The password specified here is unique to PowerChute. The password requires:
 - Minimum 8 and maximum 128 characters in length
 - One upper and lower case letter
 - One number and special character
 - The username cannot be part of the password.
 - Passwords with leading and trailing white spaces are removed.
 - Passwords containing a white space are kept.

Changes to the Username, Password and Authentication Phrase are logged to the event log. For more information see [Non-Configurable Events](#).



If you forget your password, see [Resetting your PowerChute username or password](#).

If PowerChute is registered with more than one NMC, they should all use the same administrator user name and authentication phrase.

Click the **Check Details** button on the Communications Settings page to validate that the PowerChute settings are the same as the NMC(s).

PowerChute Agents

The PowerChute Agents page, under the **UPS Configuration** menu option, lists all PowerChute Agents registered with the same NMC(s).

Loading this screen may be slow as PowerChute tries to resolve the host name for each Agent. If the host name cannot be resolved, just the IP address will be displayed. You can click on an IP address to launch the PowerChute user interface for that Agent.

A maximum of 50 PowerChute Agents can be registered with a single NMC. For more information, please view Application Note "PowerChute Network Shutdown with more than 50 computers" [here](#).



If you uninstall a PowerChute Agent, its IP address remains registered with the NMC and must be removed manually using the NMC UI.

If **System Problem Report** is displayed when accessing the screen, this is because PowerChute has not received the information it requires from the NMC(s). During normal operation, this can happen due to network traffic. Try the menu selection again in a few minutes.

This may also occur if PowerChute cannot establish communication with the NMC.

See [Network Management Card Troubleshooting](#).

PowerChute Configuration File

PowerChute stores all its settings in a configuration file called **pcnsconfig.ini**, located in the **group1** folder where PowerChute is installed.

This file is updated when running the PowerChute Setup and when you make configuration changes through the user interface, e.g. enabling shutdown actions for events.

After you have configured one installation of PowerChute with your required settings you can use the pcnsconfig.ini file to apply the same settings to another copy of PowerChute on a different machine. Certain settings such as the **localHostAddress** or **UnicastAddress** values in the **[Networking]** section will need to be edited manually for the target machine.

To apply the settings on the target machine:

Stop the PowerChute service. For more information, see Knowledge Base article [FA290624](#).

Replace the existing copy of pcnsconfig.ini in the group1 folder.

Start the PowerChute service.

Resetting your PowerChute username or password

If you forget your username or password, you can re-set them by editing the pcnsconfig.ini file.

In the **[NetworkManagementCard]** section of the INI file, set the following lines with your new values:

```
username= new user name
```

```
password= new password
```

Save the file and re-start the PowerChute service

Java Update

The Java Update feature enables you to change the Java Runtime Environment (JRE) used by PowerChute to any other JRE already installed on your system. Follow the steps below to update the Java version used by PowerChute.

1. Download a valid JRE on your system. JREs can be downloaded from the [AdoptOpenJDK website](#).



PowerChute supports AdoptOpenJDK 11. You can only update the Java version used with PowerChute to a 64-bit JRE.

The Java versions supported by PowerChute are posted on the APC website [here](#).

2. Navigate to the PowerChute installation directory, and create a new folder called “Updates”. If the default installation directory was chosen during installation, this location will be:
 - C:\Program Files\APC\PowerChute\Updates for Windows systems
 - /opt/APC/PowerChute/Updates/ for Linux systems
3. Copy the Java file downloaded in Step 1 above to the Updates directory.
4. Navigate to the **About** screen in the PowerChute UI (**Help > About**). In the **Java Update Available** field, the downloaded Java file will be listed in a drop-down box.
5. Select the Java version you want to update PowerChute to use from the drop-down box, and click **Upgrade**.
6. A confirmation dialog will appear. Click **OK**.
7. Another dialog will appear and PowerChute restarts. Wait 2-3 minutes for the Java version to successfully update.
8. When the PowerChute service restarts, refresh your browser and navigate to the **About** screen. The **Java Version** field will be updated to show the new Java used by PowerChute.



It is not necessary to keep the downloaded Java files in the Updates directory after the Java update is complete

User Interface Session Timeout

The PowerChute user interface has a ten minute session timeout by default. Following ten minutes of inactivity, the session will be terminated and the login screen will display to enter the username and password. It is possible to increase or decrease the duration of the session timeout by editing the **Web.xml** file.

To change the session timeout duration:

1. Stop the PowerChute service. For more information, see Knowledge Base article [FA290624](#).
2. Open the folder where the installed files are located, and locate the file at the following location:

```
\group1\comp\http\html\WEB-INF\web.xml
```

Open the **Web.xml** file with a text editor.

3. Locate the the **<session-config>** element, e.g:

```
<session-config>  
  
    <session-timeout>10</session-timeout>  
  
</session-config>
```

4. The duration value in the **<session-timeout>** element can be changed. For example to change the timeout to 15 minutes, change the **<session-timeout>** value to 15, e.g.:

```
<session-config>  
  
    <session-timeout>15</session-timeout>  
  
</session-config>
```

5. Save the Web.xml file.
6. Start the PowerChute service.

The PowerChute UI will now timeout following a period of inactivity that corresponds to the new **<session-timeout>** value.

Check for Updates

The **Enable Automatic Updates** feature is selected by default and informs you when a new software update is available.

When enabled, PowerChute checks for available software updates when the service is started and every seven days after that. You can also check for updates immediately by clicking the **Check Now** button.

When a new software version is released, key details and a link to download the new version are displayed on the Check for Updates screen and logged in the Event Log and MIB browser via SNMP traps.

You can disable the software update notifications feature on the last screen of the PowerChute Setup wizard or on the Check for Updates page.

Customer Support

For customer support options, please visit www.apc.com/support as a starting point.

The [Knowledge Base](#) there contains detailed troubleshooting information for product issues.

You can also browse discussion forums or submit a query using e-mail.

For country-specific support centers' contact details, go to www.apc.com/support/contact-us and select your country from the drop-down box list. This lists the contact details for support services you may require.

Customer Experience Improvement Program (CEIP)

PowerChute's Customer Experience Improvement Program (CEIP) provides us with the information that enables us to improve our product and services, and helps us to advise you on how best to deploy and configure PowerChute.

As part of the CEIP, we will collect certain information about how you configure and use PowerChute Network Shutdown in your environment. This information is completely anonymous, and cannot be used to personally identify any individual. For more information, please refer to the [CEIP Frequently Asked Questions](#) on the APC website.

By default, you are participating in the PowerChute CEIP. If you prefer not to participate, unselect the **Join PowerChute Customer Experience Improvement Program ("CEIP")** checkbox in the PowerChute CEIP page. You can join or leave the CEIP at any time.



localhost
Logout

- Home
 - Configuration
 - Virtualization
 - UPS Configuration
 - Help
- Overview
 - View Event Log
 - Customer Experience Improvement Program

PowerChute's Customer Experience Improvement Program ("CEIP") provides us with the information that enables us to improve our product and services, and helps us to advise you on how best to deploy and configure PowerChute.

As part of the CEIP, we will collect certain information about how you configure and use PowerChute Network Shutdown in your environment. This information is completely anonymous, and cannot be used to personally identify any individual. For more information, please refer to our [CEIP Frequently Asked Questions](#)

If you prefer not to participate in the PowerChute CEIP please uncheck the box below. You can join or leave the CEIP at any time.

Join PowerChute Customer Experience Improvement Program ("CEIP")

Apply Undo

License

PowerChute v4.5+ is a licensed product. A license can be purchased through your Schneider Electric IT Partner and is required for PowerChute to function and to perform graceful system shutdown.



NOTE: A PowerChute license is required for all configurations except local OS shutdown, configured by selecting **Do not enable Virtualization Support** in the installer or selecting the **Standalone VMware Host** option in the **VMware Configuration** page in the PowerChute Setup wizard.



Cloned PowerChute VMs are considered unlicensed.

PowerChute Setup: License ?

Contact your Schneider Electric IT Partner to purchase a [PowerChute Network Shutdown license](#).

Activate Online ?
 Activate Offline ?
 Local Licensing Server ?
 30-day Trial ?

Activation ID ?
This field is required.

Quantity ?

Proxy IP Address ?

Proxy Port ?

Proxy Username ?

Proxy Password ?

License screen in the PowerChute Setup wizard

Purchase a License

A PowerChute license can be purchased through your Schneider Electric IT Partner.

Activate License

You can activate your purchased license [online](#) using the Schneider Electric Cloud Licensing Server [offline](#). **NOTE:** The validation process requires an Internet connection to connect to the Schneider Electric Cloud Licensing Server.



You can also use the Schneider Electric Local Licensing Server to activate your license. See the [Local Licensing Server Guide](#) on the [APC website](#) for more information.



If you purchased more than one license, you can enter the same Activation ID for each PowerChute agent that you install. You can return a license before uninstalling PowerChute to reuse the license for a new agent. See [Return License](#).

Activate a Trial License

PowerChute v5.1.0 allows customers to avail of a trial license which can be activated during the Setup Wizard. A trial license enables all PowerChute functionalities for 30 days from the date of license activation. The purpose of the trial is to test PowerChute in your virtualized or HCI environment before purchasing a license.



The trial license can only be activated once. When the trial period has expired, PowerChute becomes unlicensed. Functionalities will be restricted and your environment will no longer be protected.

To enable full functionality, you must purchase a license which can be deployed at any time before, during, or after a trial license.

Online Activation

If the machine running PowerChute has an Internet connection, you can use the Schneider Electric Cloud Licensing Server to validate your license:

1. Select the **Activate Online** radio button.
2. Enter the license's **Activation ID**. This is in the format ACT-XXXX-XXXX-XXXX-XXXX.
3. Enter your **Proxy IP Address and Port Number**, **Proxy Username** and **Proxy Password** if required.
4. Enter the **Quantity** of nodes in your cluster protected by the PowerChute instance. One license is required for each node protected by PowerChute.
5. Click the **Activate** button.
6. PowerChute will contact the Schneider Electric Cloud Licensing Server to validate the provided Activation ID. If the license is successfully validated, PowerChute will display the license details.

Offline Activation

If the machine running PowerChute does not have an Internet connection, you can validate your license offline:

1. Select the **Activate Offline** radio button.
2. Enter the license's **Activation ID**. This is in the format ACT-XXXX-XXXX-XXXX-XXXX.
3. Enter the **Quantity** of nodes in your cluster protected by the PowerChute instance. One license is required for each node protected by PowerChute.
4. Click the **Generate Request** button. This will generate and download a request file called `request_[GUID].bin` in your Downloads folder. This file can also be found in the `user_files` directory. If the default location was chosen during installation, this location will be: `/opt/APC/PowerChute/user_files/`
 - `C:\Program Files\APC\PowerChute\user_files` for Windows systems
 - `/opt/APC/PowerChute/user_files/` for Linux systems
5. Click the **Open Licensing Portal** button to open the Schneider Electric Licensing Portal in a new browser tab.
 - a. Log in via your **Activation ID**.
 - b. Navigate to **Devices > Offline Device Management** and upload the `request_[GUID].bin` file generated in step 4.
 - c. Download the generated license (`capabilityResponse.bin` file).

6. Return to PowerChute and click the **Upload Response** button and select the downloaded `capabilityResponse.bin` file using the file selector.



If the `capabilityResponse.bin` file is modified or the download was corrupt, the file will be invalid, and the license cannot be activated. If an error message is displayed in the UI related to the capability response, follow the above steps again to generate and download a new `capabilityResponse.bin` file.

7. If the license is successfully validated, PowerChute will display the license details.

Renew License

The displayed **Renewal Date** is when your PowerChute license will expire. You can renew your license **before the Renewal Date** is reached. You will be notified 60 days in advance of your license expiring. If you do not renew your license before it expires, you must purchase and activate a new license.

Expired License

PowerChute provides a 60-day grace period after your license expires to allow you to purchase a new license. See [Purchase a License](#). Once you have purchased a license, update your license details in the **Renew License** screen for continued PowerChute functionality. **NOTE:** PowerChute remains functional during the grace period.

If you do not update your license details during the grace period, you will no longer have access to configuration options in the PowerChute UI and the PowerChute Setup wizard – except the **License** screen – and shutdowns and startups will be canceled.



PowerChute cannot perform graceful system shutdown if the software is unlicensed. It is highly recommended that you renew your license to prevent potential data loss.

Return License

You can return a license before uninstalling a PowerChute agent to be reused for a new agent. Returning a license will increase the **Quantity** associated with your **Activation ID** by the quantity that was returned. For example, if you returned 5 PowerChute licenses, your **Quantity** will increase by 5.

You can view how many licenses you have remaining in the [Schneider Electric Licensing Portal](#):

1. Log in via your Activation ID.
2. Navigate to **Activation & Entitlements > List Entitlements**.
3. You can view the total quantity, available quantity, and expiration date of the licenses associated with your Activation ID.


NOTE: When you return a license, the PowerChute agent will be unlicensed and no longer function.




If you want to deploy a new instance of the PowerChute virtual appliance, you must return the license on the old appliance before activating the same license on the new appliance instance. This also applies to Windows installations, you must return the license before uninstalling PowerChute so the same license can be activated when you re-install PowerChute.

This does not affect the license term or upgrades.

If you activated your license via the [online](#) method, click the **Return License** button in the **License** screen.

 It is not supported to return a license if you activated your license via the [offline](#) method. To reuse licenses on offline PowerChute machines, contact [Customer Support](#).

 When you have successfully returned a license, you can uninstall the PowerChute agent.

Free/Open Source License Details

PowerChute stores the details of the free and open source licenses used in PowerChute in the notice.txt file located in the PowerChute installation directory. If the default installation directory was chosen during installation, this location will be:

- C:\Program Files\APC\PowerChute for Windows systems
- /opt/APC/PowerChute for Linux systems

Reset Virtual Appliance Root Password

Follow the steps below to reset the PowerChute virtual appliance (AlmaLinux) root password:

1. Open the VMware console:

```

Welcome to PowerChute Network Shutdown 4.5 for VMware

Please complete the PowerChute Setup wizard to ensure that your VMware Hosts and Virtual Machines
are protected.

To configure PowerChute Network Shutdown using the Setup Wizard, browse to:

https://18.216.249.125:6547/
https://2001:db8:c18:4:258:56ff:fe87:5abd:6547/
localhost login: _

```

2. Reboot the virtual appliance:

```

AlmaLinux (4.18.0-348.23.1.el8_5.x86_64) 8.5 (Arctic Sphynx)
AlmaLinux (0-rescue-fd9665cbb0cc46318f608a986f0d9953) 8.5 (Arctic Sphynx)

Use the ↑ and ↓ keys to change the selection.
Press 'e' to edit the selected item, or 'c' for a command prompt.

```

3. Press "e" during the boot process on the grub menu to enter the grub editor:

```

load_video
set gfx_payload=keep
insmod gzio
linux ($root)/vmlinuz-4.18.0-348.23.1.el8_5.x86_64 root=/dev/mapper/almalinux-
root ro crashkernel=auto resume=/dev/mapper/almalinux-swap rd.lvm.lv=almalinux\
/root rd.lvm.lv=almalinux/swap rhgb quiet
initrd ($root)/initramfs-4.18.0-348.23.1.el8_5.x86_64.img $tuned_initrd

Press Ctrl-x to start, Ctrl-c for a command prompt or Escape to
discard edits and return to the menu. Pressing Tab lists
possible completions.

```

```
load_video
set gfx_payload=keep
insmod gzio
linux ($root)/vmlinuz-4.18.0-348.23.1.el8_5.x86_64 root=/dev/mapper/almalinux-\
root ro crashkernel=auto resume=/dev/mapper/almalinux-swap rd.lvm.lv=almalinux\
/root rd.lvm.lv=almalinux/swap rhgb quiet rd.break enforcing=0
initrd ($root)/initramfs-4.18.0-348.23.1.el8_5.x86_64.img $tuned_initrd
```

4. Use the arrow keys to move to the end of the line that starts with “linux” and add the following at the end of the line: “rd.break enforcing=0”:
5. Press **Ctrl + X** to boot to emergency mode.
6. Mount the `/sysroot` directory with read and write permissions with the command “`mount -o rw,remount /sysroot`”:

```
Generating "/run/initramfs/rdsosreport.txt"

Entering emergency mode. Exit the shell to continue.
Type "journalctl" to view system logs.
You might want to save "/run/initramfs/rdsosreport.txt" to a USB stick or /boot
after mounting them and attach it to a bug report.

switch_root:/# mount -o rw,remount /sysroot_
```

7. Change the directory environment to `/sysroot`:

```
Generating "/run/initramfs/rdsosreport.txt"

Entering emergency mode. Exit the shell to continue.
Type "journalctl" to view system logs.
You might want to save "/run/initramfs/rdsosreport.txt" to a USB stick or /boot
after mounting them and attach it to a bug report.

switch_root:/# mount -o rw,remount /sysroot
switch_root:/#
switch_root:/# chroot /sysroot_
```

8. Use the command “`passwd root`” to set a new password for the root user:

```
Generating "/run/initramfs/rdsosreport.txt"

Entering emergency mode. Exit the shell to continue.
Type "journalctl" to view system logs.
You might want to save "/run/initramfs/rdsosreport.txt" to a USB stick or /boot
after mounting them and attach it to a bug report.

switch_root:/# mount -o rw,remount /sysroot
switch_root:/#
switch_root:/# chroot /sysroot
sh-4.4#
sh-4.4#
sh-4.4# passwd root
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
sh-4.4#
```

9. Exit `sysroot` by typing the command “`exit`” and make the file system read-only again by typing the command “`mount`”

```
-o ro,remount /sysroot”:
```

```
Generating "/run/initramfs/rdsosreport.txt"

Entering emergency mode. Exit the shell to continue.
Type "journalctl" to view system logs.
You might want to save "/run/initramfs/rdsosreport.txt" to a USB stick or /boot
after mounting them and attach it to a bug report.

switch_root:/# mount -o rw,remount /sysroot
switch_root:/#
switch_root:/# chroot /sysroot
sh-4.4#
sh-4.4#
sh-4.4# passwd root
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
sh-4.4# exit
exit
switch_root:/# mount -o ro,remount /sysroot
switch_root:/#
```

10. Exit emergency mode and reboot the appliance.

11. Login using the new root password.

Troubleshooting

This section contains information on the topics below:

- [Network Management Card Troubleshooting](#)
- [VMware Troubleshooting](#)
- [Nutanix Troubleshooting](#)
- [HPE SimpliVity Troubleshooting](#)
- [HyperFlex Troubleshooting](#)
- [Dell VxRail Troubleshooting](#)
- [Browser Troubleshooting](#)
- [SSH Actions Troubleshooting](#)
- [SNMP Troubleshooting](#)
- [General Troubleshooting](#)

Network Management Card Troubleshooting

PowerChute does not register with the Network Management Card(s) or PowerChute reports communications lost with the Network Management Card(s):

1. Verify that the Administrator Username and Authentication Phrase are the same for PowerChute and the NMC.
2. Verify that UDP port 3052 is not being blocked by a firewall.
3. Check the IP settings on the PowerChute machine and on the NMC user interface to verify that the default gateway and subnet mask are correct.
4. Check that the IP address of the Network Management Card has not changed after PowerChute was registered with the NMC. To check this, click on the Communications Settings menu item in PowerChute and check that the IP address shown in the Network Management Card section is correct.
5. Verify that there is a network connection between the PowerChute client computer and the Network Management Card. Attempt to access the Network Management Card from the PowerChute client computer, or use the ping utility from the Network Management Card.
6. A PowerChute Network Shutdown client that acquires its IP address through DHCP will lose communications with the Network Management Card when the client renews its DHCP address lease and acquires a different IP address. To resolve this issue, each system using PowerChute Network Shutdown must have a permanent IP address. Reserve IP addresses in the DHCP server by using the MAC address of the clients, so that they never change for the specified machines. The NMC should also have a static IP address.
7. Verify that the PowerChute service is started. If it is already started, stop the service and then restart it again.
8. Update the firmware on the NMC to the latest version which can be downloaded from the APC Web site, or contact ["Customer Support"](#).

VMware Troubleshooting

VM Migration

The issues below will prevent VMs from being migrated to other available Hosts. PowerChute does not log migration errors for individual VMs in the Event Log but the error.log file contains additional information.

vSphere Errors

When performing manual VM Migration PowerChute uses the MigrateVM_Task API function. When there are errors migrating VMs the following appears in the error.log file:

VM Migration Error Message -> [Error message]

For detailed information:

<https://www.vmware.com/support/developer/vc-sdk/> - Click on vSphere API Reference, Expand vSphere API Reference and go to All Methods->MigrateVM_Task. The faults section provides a table outlining the various errors that can occur when attempting to migrate VMs.

To troubleshoot individual VMs the logging level can be set to “warn” instead of “error”, see section on debug logging.

The events below are logged to the Event Log if some or all VMs do not successfully migrate when there are target hosts available:

“Unable to migrate all VMs from Host [hostname]”

“Unable to migrate any VMs from Host [hostname]”

The events below are logged to the error.log file:

“[VM] failed to migrate to Host [Host]”

“VMware Error message [Fault message]”

Examples of errors that can occur include:

- **InvalidState:** Thrown if the operation cannot be performed because of the virtual machine's current state or the target host's current state. For example, if the virtual machine configuration information is not available or if the target host is disconnected or in maintenance mode.
- **InsufficientResourcesFault:** Thrown if this operation would violate a resource usage policy.
- **MigrationFault:** Thrown if it is not possible to migrate the virtual machine to the destination host. This is typically due to hosts being incompatible, such as mismatch in network policies or access to networks and datastores.
- **VmConfigFault:** Thrown if the virtual machine is not compatible with the destination host. Typically, a specific subclass of this exception is thrown, such as IDEDiskNotSupported.

Note: The errors above and descriptions are taken from the vSphere API Reference Guide :

<https://www.vmware.com/support/developer/vc-sdk/>

There were no available Hosts to migrate VMs to.

This will occur if the other Hosts are in Maintenance mode, incompatible or have a UPS critical event active.

The event below is written to the PowerChute Event Log:

"Unable to find a suitable host to migrate VMs from Host: [Host]"

There is no communication with vCenter Server and/or there was an Active Directory authentication error.

The event below is written to the PowerChute Event Log:

"Cannot connect to vCenter Server. PowerChute will not be able to perform VM migration."

The migration time set in the Duration field was not long enough to allow all VMs to get migrated to another host in the cluster.

The error below is written to the PowerChute Event Log:

"Insufficient time to migrate all VMs using the VM Migration Duration configured."

The error below is written to the PowerChute error.log:

"Insufficient time to migrate all VMs using VM Migration Duration"

VM Shutdown**There is no communication with vCenter Server and/or there was an Active Directory authentication error.**

The errors below are written to the PowerChute Event Log:

"Cannot connect to vCenter Server. PowerChute may not be able to issue commands to Virtual Machines or Hosts."

"vCenter Server authentication error. PowerChute may not be able to issue commands to Virtual Machines or Hosts."

When vCenter Server Connection is not available PowerChute will attempt to connect directly to each ESXi host to perform VM shutdown. This requires a shared user account (see section on creating shared user accounts) to be present on each of the ESXi hosts. If PowerChute cannot connect to the ESXi hosts the following exceptions may appear in the error.log:

com.vmware.vim25.NoPermission

com.vmware.vim25.InvalidLogin

Solution:

Verify that the vCenter Server User Account configured in PowerChute can access each of the ESXi hosts by connecting directly to each host using the vSphere client. If the connection is lost, add the user account to each ESXi host. If the connection is successful verify that the user account has "Administrator" permissions.

When PowerChute attempts to connect directly to the ESXi hosts it must be able to connect to them using their Fully Qualified Domain Name or IP Address when vCenter Server is no longer available. If there are DNS issues or

Hostname resolution issues the following errors appear in the error.log file:

VI SDK invoke exception:java.net.ConnectException: Connection timed out: connect

Solution:

Verify that DNS lookups of the ESXi Host FQDN are working from the PowerChute machine using the nslookup command.

If nslookup is not successful, the ESXi hosts can be added to the /etc/hosts file on Linux (PowerChute Appliance/vMA) or **C:\Windows\system32\drivers\etc\hosts** file on Windows machines. If the only DNS server available is running as a VM on the ESXi hosts being protected then it is necessary to use the hosts file. Alternatively ESXi hosts can be added to vCenter Server using their IP address instead of FQDN.

The time set in the Duration field was insufficient to allow VMs to shut down gracefully.

The event below is written to the PowerChute Event Log:

"Insufficient time to shut down all VMs using the VM Shutdown Duration configured."

The event below is written to the PowerChute Error.log:

"Insufficient time to migrate all VMs using VM Migration Duration"

VMware tools have not been installed. This will cause a hard shutdown of the VM.

VMware tools must be installed to perform graceful Guest OS shutdown. VMware Tools status is shown in the summary tab for VMs.

DRS is set to fully automated for the cluster. Enable VM Migration and set the duration.

If DRS is enabled and set to fully automated for the cluster, VM Migration is enabled by default. If you disable VM Migration while DRS remains fully automated and enabled, when a maintenance mode task begins on the host, DRS will start migrating VMs to other available hosts. If PowerChute begins VM shutdown on the host at the same time as the DRS migration occurs, VMs that are in migration will not successfully shut down.

If DRS is enabled and set to fully automated, VM Migration must be enabled in PowerChute with a VM migration duration set, in order to allow Virtual Machines to migrate successfully.

VM Startup

The Host had insufficient resources to start the VM, see section on [HA Admission Control](#).

The event below is written to the PowerChute Event Log:

"Attempting to power on VMs on Host [hostname] that did not start."

The event below is written to the PowerChute Error.log:

"VM [VM] Power on failed."

The Host is reporting a HA Configuration error

When entering and exiting maintenance mode the Host's HA configuration might experience an error and end up in an invalid state. Invalid state will be shown on the Summary tab for the Host in the vSphere client. This will prevent PowerChute from powering on VMs and there will be repeated event log entries indicating that PowerChute is

attempting to re-start VMs that did not start.

Solution:

Right click on the Host and select "Reconfigure for vSphere HA"

OR

Right click on the Host, select "Disconnect" and then "Reconnect"

vApp Shutdown

There is no communication with vCenter Server and/or there was an Active Directory authentication error.

The following event appears in the PowerChute Event Log:

"Cannot connect to vCenter Server. PowerChute will not be able to perform vApp Shutdown."

vApp shutdown is not supported if the vCenter Server is unavailable during the shutdown sequence.

VMware Tools have not been installed on each VM in the vApp.

This is required to perform Guest OS shutdown. If Guest OS shutdown is enabled for a VM in the vApp and VMware Tools are installed this will cause the vApp shutdown task to be unsuccessful.

The Shutdown action in vApp settings is not set to Guest OS shutdown.

If the shutdown action is set to Power Off instead of Guest OS Shutdown, VMs in the vApp will not be shut down gracefully.

The time you set in the Duration field was not long enough to allow the vApp to shut down gracefully.

The event below is written to the PowerChute Event Log:

"Insufficient time to shut down vApp using the vApp Shutdown Duration configured."

The event below is written to the PowerChute Error Log:

"Insufficient time to shut down vApp using the vApp Shutdown Duration"

The vApp contains either the PowerChute VM or the vCenter Server VM.

vApp shutdown is not supported if either the PowerChute VM or vCenter Server VM are part of a vApp.

The events below are written to the PowerChute Event Log:

"vApp [vApp] will not be shut down as it contains the Virtual Machine running PowerChute. Please remove PowerChute from the vApp."

"vApp [vApp] will not be shut down as it contains the vCenter Server VM. Please remove vCenter Server VM from the vApp."

PowerChute was unable to shut down a vApp.

The event below is written to the PowerChute error.log:

[vApp] vApp could not be Powered Off

vCenter Server VM Shutdown

VMware Tools are not installed and running on the vCenter Server VM

This will prevent the VM from being shut down gracefully and it will also prevent PowerChute from identifying the VM as the one running vCenter Server.

PowerChute cannot identify the Virtual Machine running vCenter Server using the IP address or hostname configured under vCenter Server Details on the Communications Settings page

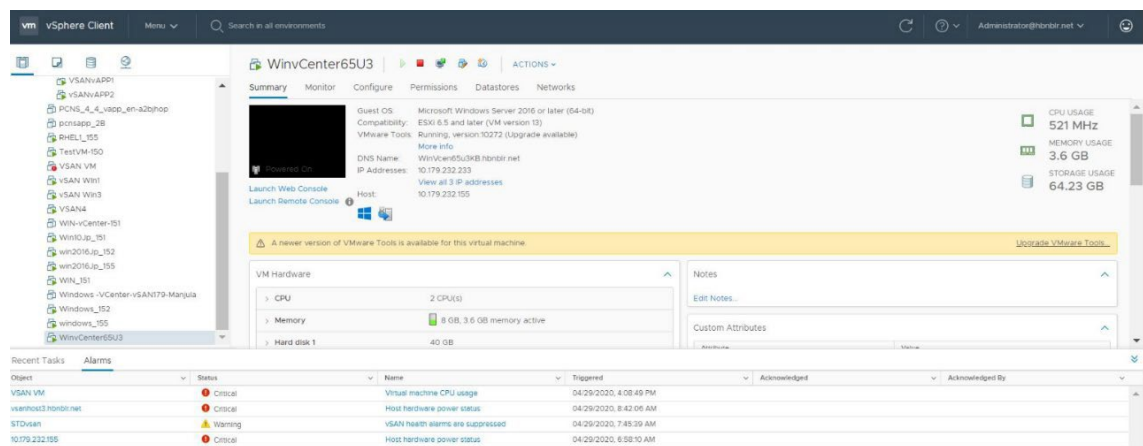
This can occur if there is a DNS configuration issue i.e. the IP address/Hostname cannot be resolved.

The event below is written to the PowerChute Event Log:

“PowerChute cannot locate the vCenter Server VM in the Inventory.”

Solution:

- Check that VMware tools are installed and running on the vCenter Server VM
- Check that the IP address or Hostname/FQDN for vCenter Server Connection in PowerChute matches the IP/Hostname shown on the Summary page for the VM.



The vCenter Server VM IP address is being used by another system or there is an old VM with the same static IP address in the vCenter Server inventory – even if that VM is powered off.

On the Host Protection page ensure that the Host running the vCenter Server VM is marked with the vSphere icon. Using vSphere client vMotion the vCenter Server VM to another host and verify that the Host protection page is updated to reflect this change.

The event below is written to the PowerChute Event Log:

“The vCenter Server VM found in the Inventory is powered off.”

Solution:

Delete any old copies of vCenter Server VM from the Inventory.

The shutdown time set in the VM Shutdown Duration field was not long enough to allow the

vCenter Server VM to shut down.

The event below is written to the PowerChute Event Log:

vCenter Server VM [VM] cannot be gracefully shut down. Please check vCenter Server VM Shutdown duration.

If the vCenter Server VM cannot be located.

The event below is written to the PowerChute Event Log:

PowerChute cannot locate the vCenter Server VM in the Inventory. See the troubleshooting section in the Online Help.

If the vCenter Server VM is powered off.

The event below is written to the PowerChute Event Log:

The vCenter Server VM found in the Inventory is powered off. See the troubleshooting section in the online help.

If the Domain name for the ESXi host running the vCenter Server VM is different for the vCenter Server Inventory compared to the Standalone ESXi host, this will cause vCenter Server VM shutdown and unsuccessful Host shutdown.

e.g. Domain is set to apcc.com in vCenter, but apcc on the ESXi host. This can be seen using the vSphere Client for each Host under Configuration > DNS and Routing > Host Identification > Domain.

To avoid this ensure that the same domain name is set in vCenter Server and on the ESXi host.

The following message appears in the error.log:

```
"checkForVCSAVMAndHostInCriticalHosts - cannot obtain HostSystem using findByIP and findByDnsName for critical host [hostname]"
```

Host Shutdown**If PowerChute cannot shut down a Host.**

The following message appears in the EventLog:

```
"Shutdown Host unsuccessful for Host [Host]."
```

PowerChute only supports the licensed version of ESXi (Essentials, Standard, Enterprise, Enterprise plus) and requires the vSphere API licensed feature.

The following exception appears in the error.log when attempting to perform VM actions or Host actions using the vSphere API using the unlicensed version of vSphere ESXi:

com.vmware.vim25.NoPermission

The account being used to connect to the Host does not have sufficient privileges.

The following is logged in the Error Log:

com.vmware.vim25.NoPermission

The account credentials being used to connect to the Host are incorrect or have expired.

The following exceptions appear in the Error Log:

com.vmware.vim25.InvalidLogin

com.vmware.vim25.AuthenticationError

DNS Configuration issues may prevent PowerChute from connecting to the host e.g. a stale DNS record containing an invalid hostname/FQDN or IP address.

The following exception appears in the Error Log:

VI SDK invoke exception:java.net.UnkownHostException

The account being used to connect to vCenter Server does not exist on each of the VMware Hosts being protected by PowerChute. See section on configuring Shared Active Directory/Local User account.

The following exception appears in the Error Log:

com.vmware.vim25.InvalidLogin

If PowerChute cannot identify the VM on which it is running in the vCenter Server inventory, its VM may get shut down too early which results in the vCenter Server VM and the Hosts not being shutdown as expected.

The following event appears in the Event Log when this occurs:

“Cannot locate the PowerChute VM in the Inventory.”

Solution:

Check that VMware tools are installed and running on the PowerChute VM.

“The PowerChute VM found in the Inventory is powered off.”

Solution:

Delete any old copies of PowerChute VM from the Inventory.

Maintenance Mode

The following message appears in the EventLog if it takes too long for PowerChute to put a Host into Maintenance Mode:

"Maintenance mode task cancelled on Host [Host] as there are still powered on VMs. Please verify that sufficient time has been configured for VM/vApp/VCSA VM shutdown duration."

The following event appears in the error.log when this occurs:

"Maintenance Mode not entered for Host [Host]"

There is no communication with vCenter Server.

If PowerChute is configured to start a Maintenance Mode on hosts at the start of the shutdown sequence and the connection between PowerChute and vCenter Server is lost during the VM/vApp Shutdown step, this will result in the hosts being placed into Maintenance Mode later in the shutdown sequence during the Host Shutdown step via a direct host connection.

The following messages appear in the EventLog when this occurs:

"Starting Maintenance Mode Task on Host [Host]"

"Cannot connect to vCenter Server. PowerChute may not be able to issue commands to Virtual Machines or Hosts."

The following event appears in the error.log when this occurs:

"Error checking maintenance mode for Host [Host]"

General error.log messages

If the host is powered off PowerChute will be unable to perform Virtualization Shutdown (VM Migration, VM Shutdown, vApp Shutdown).

The following event appears in the error.log when this occurs:

"Host is powered off. No need for Virtualization Shutdown"

If the connection to vCenter Server (Physical or vCenter Server VM) was lost e.g. when vCenter Server VM is shut down during a Shutdown Sequence or if there was a network issue connecting to the vCenter Server.

The following events may appear in the EventLog when this occurs:

"Cannot connect to vCenter Server. PowerChute may not be able to issue commands to Virtual Machines or Hosts."

"Cannot connect to vCenter Server. PowerChute will not be able to perform VM Migration."

"Cannot connect to vCenter Server. PowerChute may not be able to perform vApp Shutdown."

The following exceptions may appear in the error.log when this occurs:

VI SDK invoke exception:java.net.ConnectException: Connection timed out: connect

VI SDK invoke exception:java.net.ConnectException: Connection refused: connect

If you experience issues during the shutdown/startup sequence, e.g. the host does not exit maintenance mode or does not shutdown as expected, the following events appear in the error.log:

```
java.rmi.RemoteException: VI SDK invoke exception:java.net.SocketTimeoutException: Read timed out
```

```
java.rmi.RemoteException: VI SDK invoke exception:java.net.SocketTimeoutException: connect timed out
```

These issues can occur due to network latency issues.

Solution:

Increase the values for the following configurable timeout settings in the Configuration INI file:

```
[HostConfigSettings]
```

```
VMware_connect_timeout = 10
```

```
VMware_read_timeout = 15
```

For example, to increase the timeout settings to 30 seconds each, change the values to the following:

```
VMware_connect_timeout = 30
```

```
VMware_read_timeout = 30
```

In a HA Cluster environment, VMs may not get powered on during startup when the ESXi host is taken out of maintenance mode. This can occur if the HA Cluster election process has not completed before PowerChute attempts to start the VMs.

The following event will appear in the EventLog when this occurs:

```
"Attempting to power on VMs on Host that did not start"
```

Solution:

Increase the values for the following configurable settings in the Configuration INI file:

```
delay_after_exit_maintenance_mode = 30
```

```
delay_after_vcsa_powered_on_and_connected = 30
```

For example, to increase the delay to 60 seconds each, change the values to the following:

```
delay_after_exit_maintenance_mode = 60
```

```
delay_after_vcsa_powered_on_and_connected = 60
```


No event log or warning message displayed when the VCenter Root Certificate on Host Protection page is not accepted

On the Host Protection page, ensure that the VCenter root certificate is accepted. If the certificate is not accepted, no error message will be displayed and the shutdown will not be completed as it will be unable to communicate with individual ESXi host.

Solution:

Ensure you always accept VCenter Root Certificate.

Nutanix Troubleshooting

 **NOTE:** The troubleshooting items below are only applicable when Nutanix support is enabled, and are to be used along with [VMware Troubleshooting](#).

Cluster Connection

If the connection to your Nutanix Cluster or Nutanix Controller Virtual Machines is lost, the following log messages and exceptions may appear in the error.log:

"Failed to start connection; nested exception is:"

"net.schmizz.sshj.transport.TransportException: Broken transport; encountered EOF"

"java.net.ConnectException: Connection times out: connect"

Nutanix Cluster not available for any cluster service stop operations

This error may occur if the Cluster is down when PowerChute attempts to stop the Cluster. This could also happen if your Nutanix Cluster credentials have changed after configuring PowerChute. Ensure the correct Cluster/Controller VM credentials are provided via the PowerChute Setup wizard, or the [Nutanix Settings](#) screen.

The error below is written to the PowerChute Event Log:

"Could not connect to the Nutanix Cluster. Cannot stop Nutanix Cluster services."

Nutanix Cluster cannot be stopped

This can be caused by an unsuccessful attempt to stop AFS or an unsuccessful attempt to power off all User VMs. To resolve this issue, ensure you configure a sufficient duration to successfully stop AFS and power off all the User VMs.

The error below is written to the PowerChute Event Log:

"Cluster cannot be gracefully shut down."

Unable to start Nutanix Cluster

Cluster may not start if the time drift is incorrectly set. To check if the time between the individual Controller VMs is correctly synchronized, execute the following command:

```
allssh date
```

For information on configuring time synchronization for your Nutanix Cluster, see the following [recommendations from Nutanix](#).

This could also happen if your Nutanix Cluster credentials have changed after configuring PowerChute. Ensure the correct Cluster/Controller VM credentials are provided via the PowerChute Setup wizard, or the [Nutanix Settings](#) screen.

The error below is written to the PowerChute Event Log:

"Cluster cannot be started."

Acropolis File Services (AFS)

Unable to stop AFS

If AFS cannot be successfully stopped, the Cluster cannot be stopped and Controller VMs cannot be gracefully shut down. To resolve this issue, increase the AFS stop duration. It is recommended that you manually test the time needed to stop AFS on your Cluster and specify that as the duration in the PowerChute UI.

To do this, connect to any Controller VM while AFS is running, and use the following command:

```
afs infra.stop <fileservname>
```

Note the duration it takes for the AFS service to fully stop and use this as the **AFS Shutdown Duration** in the Virtualization Settings screen in the PowerChute UI.

The error below is written to the PowerChute Event Log:

```
"Nutanix AFS cannot be gracefully shut down."
```

The error below is written to the PowerChute error.log:

```
"Failed to stop Nutanix AFS: [error received]"
```

Protection Domain

Unable to abort Protection Domain Replications

For information on why this did not abort, refer to the error message generated in the error.log.

The error below is written to the PowerChute Event Log:

```
"Nutanix Protection Domain replications cannot be gracefully aborted."
```

The errors below is written to the PowerChute error.log:

```
"Failed to retrieve Protection Domains replication status: [error received]"
```

```
"Failed to abort ongoing Protection Domain replications: [error received]"
```

Unable to disable Metro Availability

Metro Availability may not be disabled if Metro Availability was not correctly set up on the Cluster.

For more information, refer to the error message generated in the error.log.

The error below is written to the PowerChute Event Log:

```
"Nutanix Metro Availability cannot be disabled."
```

The error below is written to the PowerChute error.log:

```
"Failed to disable Protection Domains Metro Availability: [error received]"
```

Unable to re-enable Metro Availability

Metro Availability may not be re-enabled if Metro Availability was not correctly set up on the Cluster.

For more information, refer to the error message generated in the error.log.


The error below is written to the PowerChute Event Log:

"Nutanix Metro Availability cannot be enabled."

The error below is written to the PowerChute error.log:

"Failed to enable Protection Domains Metro Availability: [error received]"

HPE SimpliVity Troubleshooting

 **NOTE:** The troubleshooting items below are only applicable when HPE SimpliVity support is enabled, and are to be used along with [VMware Troubleshooting](#).

Cluster Connection

Cannot connect to Cluster

If the connection to your HPE SimpliVity OmniStack Virtual Controller (OVC) VMs is lost, the following log messages and exceptions may appear in the error.log:

"One or more OmniStack Virtual Controllers are offline. We cannot apply configuration for offline controllers."

"Could not connect to the Cluster over the network."

"Connection error due to invalid login credentials."

Ensure the correct credentials are provided via the PowerChute Setup wizard, or the Communication Settings screen. Additionally, check that all OVCs in your HPE SimpliVity Cluster are powered on and have the correct IP addresses assigned, and provide a sufficient **Shutdown Duration** value to allow the Cluster Services to start successfully,

OmniStack Virtual Controller

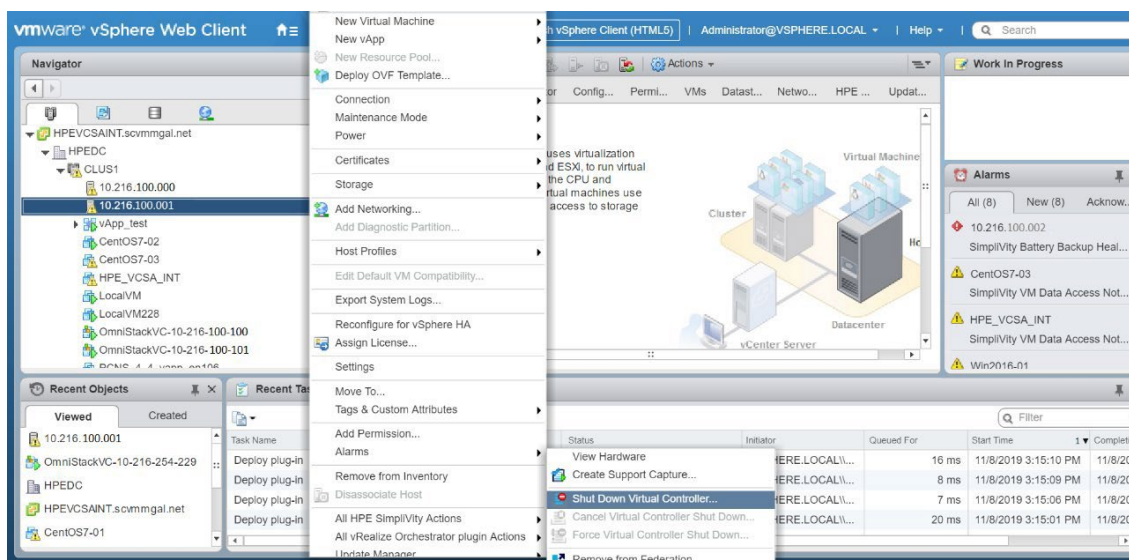
vCenter Server displays the same IP address for all OmniStack Virtual Controller VMs

After vCenter Server is restarted, VMware Tools installed on the VM incorrectly displays the same IP address for all OmniStack Virtual Controller VMs in the HPE SimpliVity Cluster. When this issue occurs, OmniStack Virtual Controller VMs cannot be shut down and no error message is logged to the Event Log. PowerChute may also report that it cannot communicate with the OmniStack Virtual Controller VMs as PowerChute cannot connect to the IP addresses reported by VMware Tools.

To resolve the issue, restart the OmniStack Virtual Controller VMs one at a time using the HPE vSphere plugin options:

1. Right-click on the Host with the required OmniStack Virtual Controller VM.
2. Click "All HPE SimpliVity actions" and click "Shut Down Virtual Controller..."

The correct IP addresses will be displayed after the OVC VMs have restarted.



Unable to shut down OmniStack Virtual Controller VMs

To successfully shut down HPE SimpliVity OmniStack Virtual Controller (OVC) VMs, HA Compliance must be reached and the first OVC VM must have completed its shut down process before the subsequent OVC VMs in the Cluster can be powered off.

Also ensure that you configure a sufficient duration to successfully shut down all OmniStack Virtual Controller VMs in the Cluster.

The errors below are written to the PowerChute Event Log:

"OmniStack Virtual Controller [OVC] did not shut down gracefully."

"Insufficient time configured for OVC shutdown. Please increase OVC Shutdown Duration."


Unable to start OmniStack Virtual Controller VMs

If the OmniStack Virtual Controller VMs in the Cluster cannot be started, ensure that you configure a sufficient duration to successfully start up all OmniStack Virtual Controller VMs in the Cluster.

The error below is written to the PowerChute Event Log:

"OmniStack Virtual Controller [OVC] has not started."

HyperFlex Troubleshooting

 **NOTE:** The troubleshooting items below are only applicable when HyperFlex support is enabled, and are to be used along with [VMware Troubleshooting](#).

Cluster Connection

Cannot connect to Cluster

If your HyperFlex credentials are incorrect and PowerChute cannot connect to HyperFlex, the following log messages may appear in the error.log:

```
"Credentials invalid, returning UNAUTHORIZED"
```

If your HyperFlex credentials were updated on HyperFlex and not PowerChute, the following log messages may appear in the error.log:

```
"Credentials invalid, returning UNAUTHORIZED"
```

Ensure the correct credentials are provided via the PowerChute Setup wizard, or the Communication Settings screen.

If the Controller VMs in your HyperFlex cluster are not available (for example, they are not fully powered on), the following log messages may appear in the error.log:

```
"Unable to connect to HyperFlex service, returning UNAVAILABLE, error: [error received]"
```

```
"Attempt to retrieve token failed, error: [error received]"
```

Ensure that all Controller VMs in your HyperFlex Cluster are powered on and have the correct IP addresses assigned.

Cannot shut down Cluster

If your HyperFlex Cluster cannot be shut down, ensure that you configure a sufficient duration to successfully stop the cluster service and shut down the cluster. For more information, consult the error.log.

The errors below are written to the PowerChute Event Log:

```
"Cluster cannot be gracefully shut down."
```

The errors below are written to the PowerChute error.log:

```
"Maximum number of retries [retry value] reached, cluster stop failed."
```

Unable to start Cluster

If your HyperFlex Cluster cannot be started, ensure that you configure a sufficient duration to successfully start the cluster service and start up the cluster.

The errors below are written to the PowerChute Event Log:

"Cluster cannot be started."

The error below is written to the PowerChute error.log:

"Failed to start HyperFlex cluster."

Controller VM**Unable to shut down Controller VMs**

To successfully shut down Controller VMs, ensure that you configure a sufficient duration to successfully shut down all Controller VMs in the Cluster.

The error below is written to the PowerChute Event Log:

"Controller VM [CVM] did not shut down gracefully."

The errors below are written to the PowerChute error.log:

"Error shutting down CVM: [CVM] - [error received]"

"CVM: [CVM] failed to power off with taskResult: [error received]"

If "Failed to retrieve VM entity for Controller VM" is written to the PowerChute error.log, increase the value of the "VMware_read_timeout" setting in the PowerChute configuration file (`pcnsconfig.ini`). The default value is 15 seconds, and it is recommended you increase this value to 30 seconds:

```
[HostConfigSettings]
```

```
VMware_connect_timeout = 10
```

```
VMware_read_timeout = 30
```

Unable to start Controller VMs

If the Controller VMs in the Cluster cannot be started, ensure that you configure a sufficient duration to successfully start up all Controller VMs in the Cluster.

The errors below are written to the PowerChute error.log:

"Error starting Controller VM: [CVM] - [error received]"

"Controller VM: [CVM] failed to power on with taskResult: [error received]"

Dell VxRail Troubleshooting



NOTE: The troubleshooting items below are only applicable when Dell VxRail support is enabled, and are to be used along with [VMware Troubleshooting](#).

Cluster Connection

Cannot connect to Cluster

If your Dell VxRail credentials are incorrect and PowerChute cannot connect to VxRail, the following error message may appear in the UI:

"VxRail Cluster is inaccessible. Please verify that the vCenter login credentials are correct and that the VxRail Manager IP or FQDN is accessible over the network."

If your vCenter Server account credentials were updated but not on PowerChute, the following log messages may appear in the error.log:

"Credentials invalid, returning UNAUTHORIZED"

Ensure the correct credentials are provided via the PowerChute Setup wizard, or the Communication Settings screen.

Cannot shut down Cluster

If your Dell VxRail Cluster cannot be shut down, ensure that you configure a sufficient duration to successfully stop the cluster service and shut down the cluster. For more information, consult the NMC Event Logs and the PowerChute error.log.

Check the NMC Event Logs

1. Log in to the NMC interface: https://<NMC_IP>.
2. Navigate to **Logs > Events > Log**.
3. Search for VxRail-related events. For example: VxRail cluster shutdown response from <IP address>: `{"request_id":"be0d7734-67d3-4787-ac49-b529f85ec099"}`

For VxRail events with a request ID, you can use the Swagger API interface to investigate the status of the VxRail cluster shutdown.

1. Log in to Swagger API: https://<vxrail_manager_ip>/rest/xvm/api-doc.html
2. Select **Requests** from the definition drop-down list:

The screenshot shows the Swagger API interface for VxRail Request Management. At the top, there is a header with the Swagger logo and the text "Select a definition" followed by a dropdown menu showing "requests". Below the header, the main content area displays "VxRail Request Management" with version "7.0.130" and "OAS3" tags. The URL is "/rest/vxrm/swagger-resources/api-specs/fgroup=requests". Below this, there is a "Servers" section with a dropdown menu showing "/rest/vxrm - VxRail Manager Server" and an "Authorize" button. The main content area is titled "Status of requests" and contains two API endpoints:

Method	Path	Description
GET	/v1/requests/{id}	Retrieves the operation status and progress report of the specified request.
GET	/v1/requests	Queries all of the requests.

3. Click **Authorize** and enter your vCenter Server credentials.
4. Click **Try it out** for GET /v1/requests/{id}
5. Enter the request ID and click **Execute**.

If PowerChute cannot successfully issue the REST API request to the NMC, the error below is written to the PowerChute Event Log:

"Request to Network Management Card {0} to initiate {1} Cluster shutdown did not succeed. Please refer to troubleshooting section in the User Guide."

The errors below are written to the PowerChute error.log:

"NMC {0} Rest service for {1} Cluster Shutdown was unable to initiate a cluster shutdown"

PowerChute relies on the Network Management Card (NMC) to shut down the Dell VxRail Cluster via a REST API. If the connection between PowerChute and the NMC is lost, the VxRail Cluster cannot be shut down. The error below is written to the PowerChute error.log if the NMC is unavailable:

"Cannot connect to the NMC {0} Rest service for {1} Cluster Shutdown."

If you have a Dell VxRail stretched cluster, ensure that the cluster hosts are synchronized. This is required for successful API calls between the NMC and cluster. For more information, consult Dell VxRail Knowledge Base article 000180885 - Dell VxRail: Failed to do cluster shutdown if the time of witness host is not synchronized with other hosts.

If you are using a FQDN for VxRail Manager, ensure that the hostname can be resolved on the NMC by adding the hostname to the **DNS Configuration** screen in the NMC Web UI.

Cannot start Cluster

When Dell VxRail and Management hosts are powered on at the same time in a stretched cluster configuration, the datastore may become inaccessible and cluster startup may not work as expected. To prevent this issue, ensure that the Management host is started first followed by the VxRail hosts so the Witness host and vCenter Server are available when the VxRail hosts start up. This can be achieved using one of the following methods:

- If hosts are manually started, the Management host must be started first followed by the VxRail hosts. Ensure you provide a sufficient delay to allow for the Management host, Witness host, and vCenter Server startup before manually starting the VxRail hosts.
- Log in to the Integrated Dell Remote Access Controller (iDRAC): https://<idrac_ip>/restgui and navigate to **Configuration > BIOS Settings**. Expand the System Security section, enable **AC Power Recovery** and set a delay of 600 seconds for **User Defined Delay**. Click **Apply And Reboot**. In PowerChute, enable the option to **Turn off the UPS** in the **Shutdown Settings** page

iDRAC9 | Enterprise

Dashboard System Storage Configuration Maintenance iDRAC Settings

TPM Firmware 1.3.2.0

TPM Hierarchy Enabled

> TPM Advanced Settings

Intel(R) TXT Off

SGX Launch Control Policy Unlocked

Power Button Enabled

AC Power Recovery On

AC Power Recovery Delay User Defined

User Defined Delay (60s to 600s) 600

UEFI Variable Access Standard

In-Band Manageability Interface Enabled

Secure Boot Disabled

Secure Boot Policy Standard

Secure Boot Mode Deployed Mode

Authorize Device Firmware Disabled


Apply Discard

> Redundant OS Control

> Miscellaneous Settings

Apply And Reboot At Next Reboot Discard All Pending

For more information, consult Dell VxRail Knowledge Base article 000190707 - Dell VxRail: 2-node ROBO cluster shutdown best practice.

 For additional troubleshooting items, visit <https://www.dell.com/support/home/en-ie/product-support/product/vxrail-appliance-series/docs>

Browser Troubleshooting

The PowerChute Web UI is accessed using a browser. For a list of supported browsers please view the [Operating System Compatibility chart](#).

PowerChute requires cookies and JavaScript to be enabled in the browser in order to function correctly. If cookies are being blocked this will prevent logging into the PowerChute UI. To avoid this, allow cookies for the PowerChute URL.

SSH Actions Troubleshooting

Error shown on screen when session times out

- When the user session times out or the PowerChute service restarts, clicking on the edit or delete icons on the SSH List View screen will result in an error message being shown on screen. Click on any of the menu items to return to the log in screen.

SNMP Troubleshooting

The Network Management System (NMS) cannot connect to PowerChute via SNMPv1:

1. Verify that there is a network connection between the NMS and PowerChute.
2. Verify that the SNMP Port specified during installation (161 by default) is not blocked for inbound communications by a firewall.
3. Verify that SNMPv1 is enabled in the PowerChute.
4. Verify that the Community Name specified in PowerChute matches the Community Name used by the NMS. The Community Name is case sensitive.
5. Verify that the NMS IP or Hostname specified in PowerChute matches the IP/Hostname of the NMS.
6. Verify that the Access Type specified in PowerChute is set to Read for SNMP Get requests, or Read/Write for SNMP Set requests.

The Network Management System (NMS) cannot connect to PowerChute via SNMPv3:

1. Verify that there is a network connection between the NMS and PowerChute.
2. Verify that the SNMP Port specified during installation (161 by default) is not blocked for inbound communications by a firewall.
3. Verify that SNMPv3 is enabled in the PowerChute.
4. Verify that the User Name specified in PowerChute matches the User Name used by the NMS. The user name is case sensitive.
5. Verify that the Authentication Protocol, Authentication Passphrase, Privacy Protocol and Privacy Passphrase used by the NMS match those specified in PowerChute.
6. Verify that the Access Type specified in PowerChute is set to Read for SNMP Get requests, or Read/Write for SNMP Set requests.

SNMP Traps sent by PowerChute are not received by the NMS:

1. Verify that there is a network connection between PowerChute and the NMS.
2. Verify that a Trap receiver has been added in PowerChute:
 - Verify that the UDP Port specified (162 by default) is not blocked for outbound communications by a firewall.
 - Verify that the SNMPv1 Community Name, or SNMPv3 User Profile used to send the Trap is configured in the NMS.
 - Verify that the Privacy Protocol selected is compatible with the Java JRE used by PowerChute:
 - Verify that the SNMP Trap Receiver Test was successful.
3. Verify that Traps for UPS Critical events are enabled.
4. Verify that Traps for Lost Communication events are enabled.

General Troubleshooting

The PowerChute service does not start if m11.cfg file is missing

If the m11.cfg file is moved, renamed, or deleted, the PowerChute service cannot start. If the m11.bak file is present, PowerChute will restore the m11.cfg file using this backup file when the PowerChute service is started. However, you will need to start the PowerChute service again after the m11.cfg file is restored.

If both the m11.cfg and m11.bak files are not present, you must uninstall and reinstall PowerChute.



Administrator access is required on all operating systems to open and edit the m11.cfg file. Ensure that the m11.cfg file is not modified in any way.

Parallel UPS Configuration is not possible with Easy UPS

Several SRV UPS devices with capacities of 10 kVA and above can be configured into a parallel configuration through the use of external hardware, such as an SRV Parallel Board.

Network Management Card(s) deployed within this configuration cannot detect this parallel configuration. PowerChute, which relies on Network Management Card data, is also not capable of recognizing this configuration as a parallel UPS setup.

The PowerChute service does not start if invalid PowerChute credentials are provided before deploying the virtual appliance

When deploying the PowerChute virtual appliance, the **Customize template** page allows you to specify some settings for the virtual appliance before deployment. The **PowerChute Username** and **PowerChute Password** are both required fields.

The OVF Deployment Wizard only validates the length of the PowerChute Web UI username and password. If the password provided does not meet the [password requirements](#), an error message is displayed in the virtual appliance **Console** stating that PowerChute services are disabled due to invalid credentials. To resolve this issue:

1. Open the pcnsconfig.ini file using an ASCII text editor.
2. In the [**NetworkManagementCard**] section of the INI file, set the following lines with your new valid values:

```
username= new user name
```

```
password= new password
```

3. Save the pcnsconfig.ini file.
4. Enable and start the PowerChute service using the following commands:

```
systemctl enable PowerChute
```

```
systemctl start PowerChute
```



Enabling the PowerChute service without providing a username or password via the pcnsconfig.ini file will start the service in an unconfigured state. This may result in denial of service if an untrusted user can access the PowerChute configuration wizard before credentials are set.

APC by Schneider Electric Worldwide Customer Support

Customer support for this or any other APC by Schneider Electric product is available at no charge in any of the following ways:

- Visit the APC by Schneider Electric web site, www.apc.com to access documents in the APC Knowledge Base and to submit customer support requests.
 - **www.apc.com** (Corporate Headquarters)
Connect to localized APC by Schneider Electric web site for specific countries, each of which provides customer support information.
 - **www.apc.com/support/**
Global support searching APC Knowledge Base and using e-support.
- Contact the APC by Schneider Electric Customer Support Center by telephone or e-mail.
 - Local, country specific centers: go to **www.apc.com/support/contact** for contact information.
 - For information on how to obtain local customer support, contact the APC by Schneider Electric representative or other distributor from whom you purchased your APC by Schneider Electric product.

© 2024 Schneider Electric. All Rights Reserved. Schneider Electric, APC and PowerChute are trademarks and the property of Schneider Electric SE, its subsidiaries and affiliated companies. All other trademarks are property of their respective owners.