User Guide

PowerChute™ Network Shutdown v4.2

VMware®

990-4595E-001
07/2016

PowerChute™
— NETWORK SHUTDOWN

Schneider Electric

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

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

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

IPv6 support is available only for Network Management Card firmware 6.0.X or higher.

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.

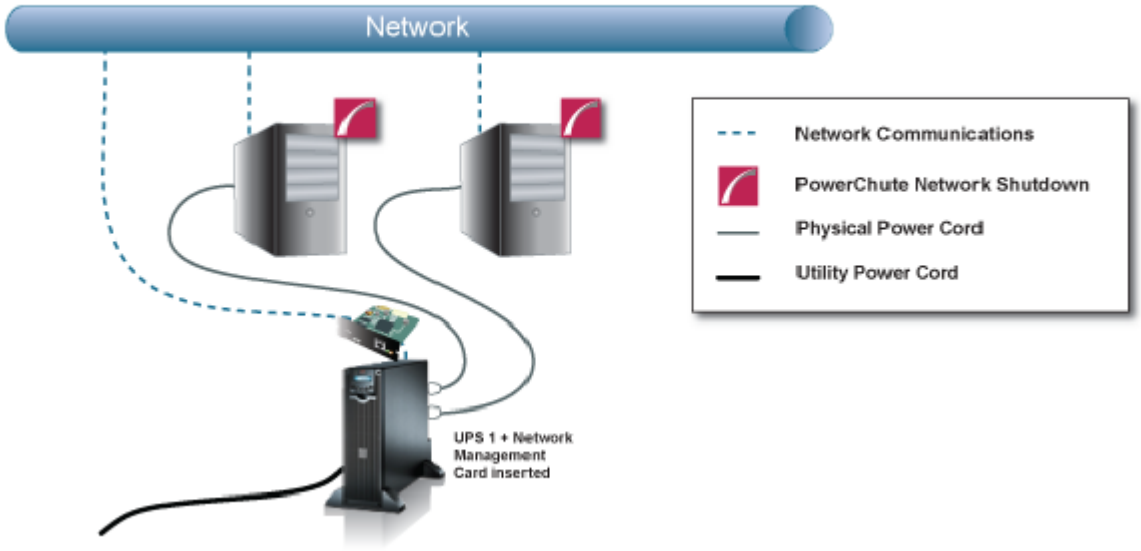

Redundant-UPS Configuration

Redundant-UPS Configuration: Two or more UPS's 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.

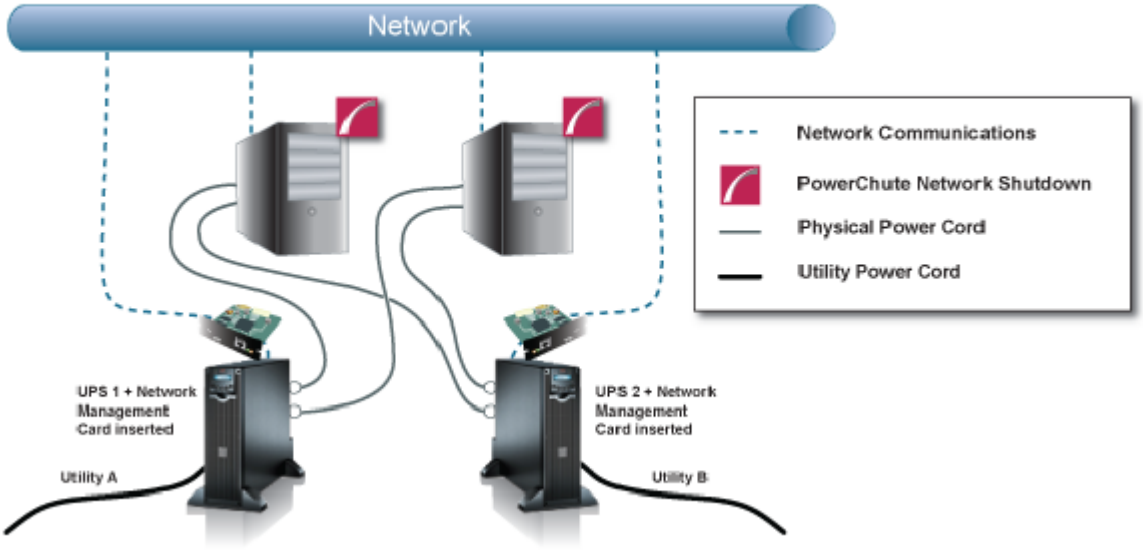

For detailed information, please view “Using PowerChute Network Shutdown in a Redundant-UPS Configuration” Application Note here.

Parallel-UPS Configuration

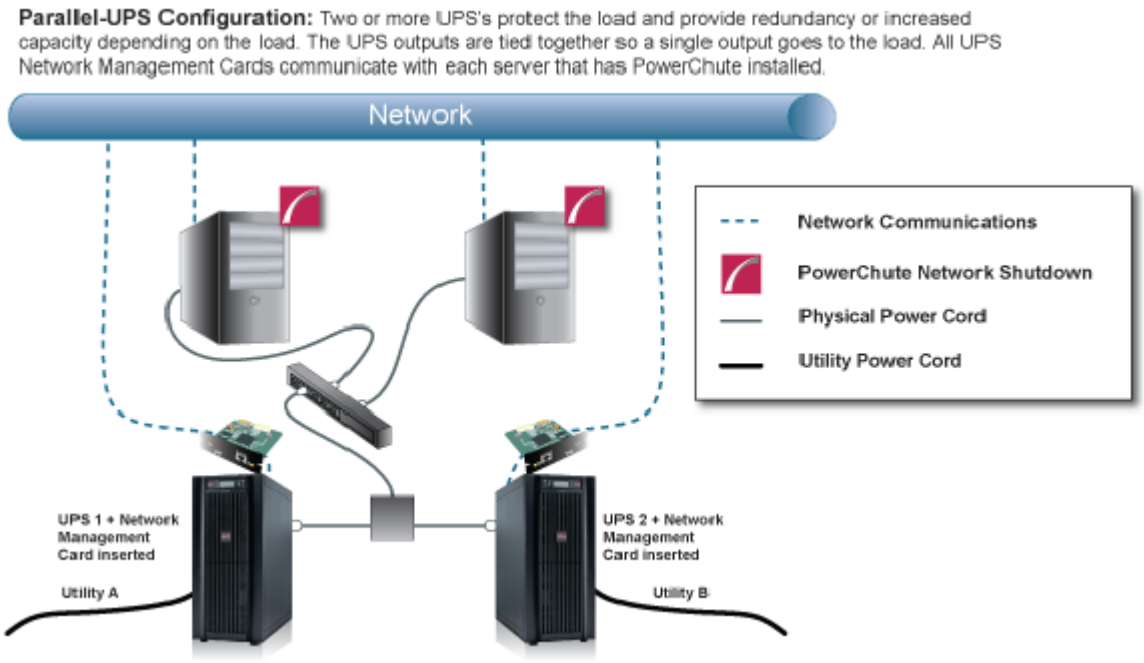

Note: To use the Parallel-UPS configuration, your UPS devices must already be configured to operate in parallel mode.

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.

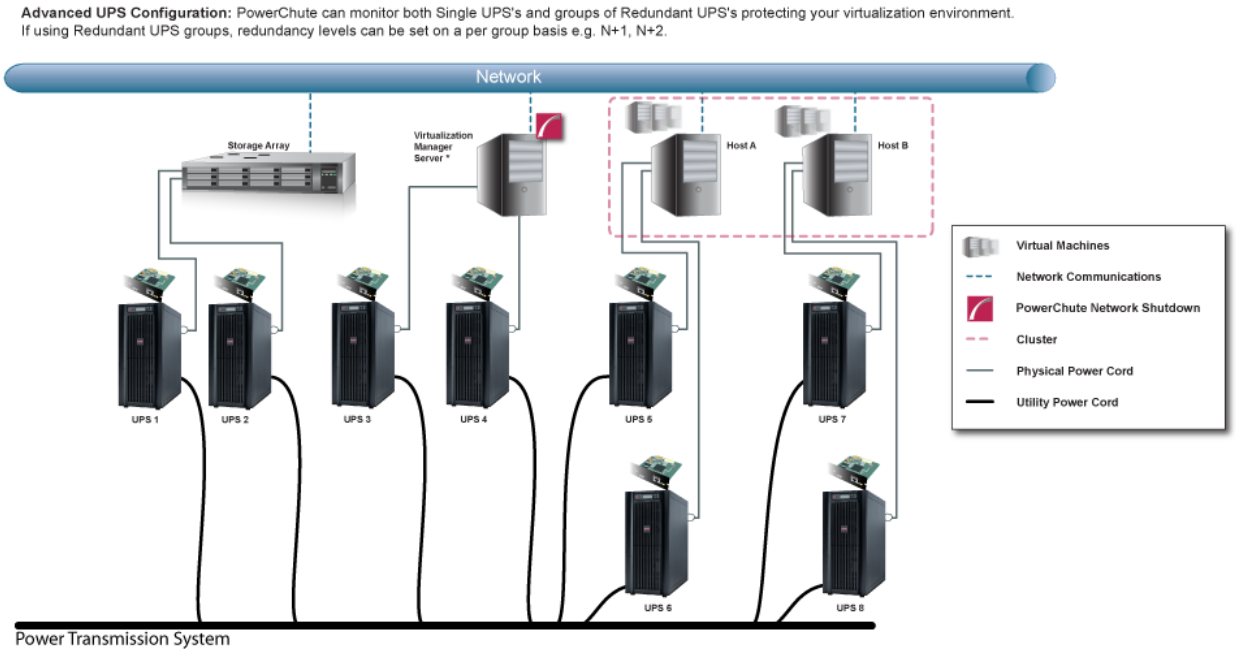

* Virtualization Manager Server is vCenter Server

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.

Network Management Card Connection

The Network Management Card uses the HTTP protocol by default. This can be changed to HTTPS through the NMC user interface. Based on the NMC protocol used, you can select either HTTP or HTTPS in PowerChute.

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.

The NMC uses a self-signed SSL certificate by default when HTTPS is enabled. You need to enable "Accept Untrusted SSL Certificates" to allow PowerChute to establish communication with the NMC if a self-signed cert is being used by the 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**.

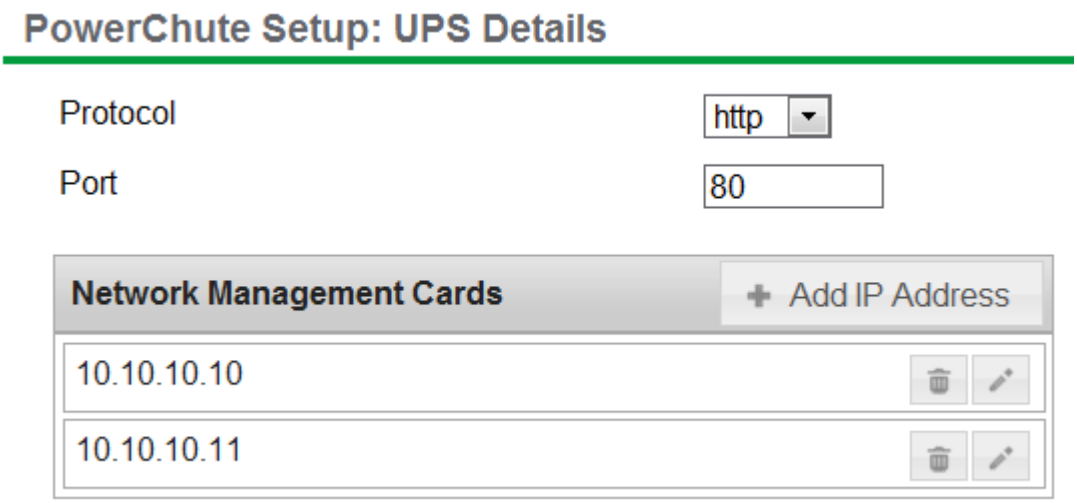

To edit an IP address, click the icon. To delete an IP address, click the icon.

Adding a Trusted Certificate to PowerChute for NMC communication

When using the HTTPS protocol to communicate with the NMC, you must select the **Accept Untrusted SSL Certificates?** check box. However, it is possible to create a Trusted Certificate file and add it to the PowerChute truststore.

Your NMC Security Handbook has details on the Security Wizard used to create the Trusted Certificate file with an extension .CRT. This file is then used to create components that can be uploaded to the NMC to replace the default self-signed certificate.

In order to facilitate the trusted SSL communication of PowerChute with the NMC, this Trusted Certificate file must then be added to the system Java **cacerts** keystore or to the **PowerChutekeystore** file. (You can do this using the Java keytool.exe; for details see the Java help documentation). Adding it to the cacerts keystore means it is available to all your applications as distinct from just PowerChute.

By default the **PowerChute-keystore** file is located in APC\PowerChute\group1. Its password is “password”. If you add the Trusted Certificate and you subsequently get a connection error with the NMC, then it could be because a) the certificate has expired, b) it is not yet valid, or c) it has been revoked. In any of these cases, you need to add a new Trusted Certificate to the PowerChute server or to upload a new valid SSL certificate to the NMC.

The **PowerChute-keystore** file only exists after the first attempt is made to communicate with the NMC using HTTPS (by using the configuration wizard for example). For this reason, for a silent installation you must add the Trusted Certificate to the Java cacerts keystore.

PowerChute only checks the keystore when its service starts. After you add the Trusted Certificate, you will need to re-start the PowerChute service if it’s already running.

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**:

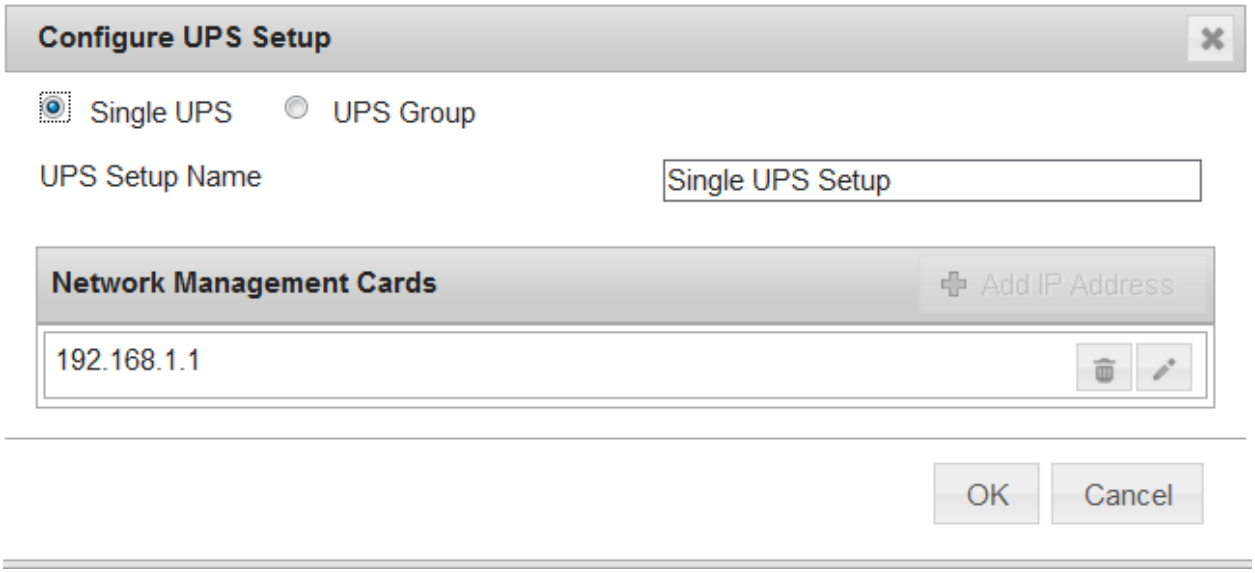

1. Enter a **UPS Setup Name** (with a maximum of 20 ASCII characters)
2. Click the **+ Add IP Address** button and enter the IP address of the Network Management card in the UPS. Click **OK**.
3. Click **OK** to complete Single UPS Setup.

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

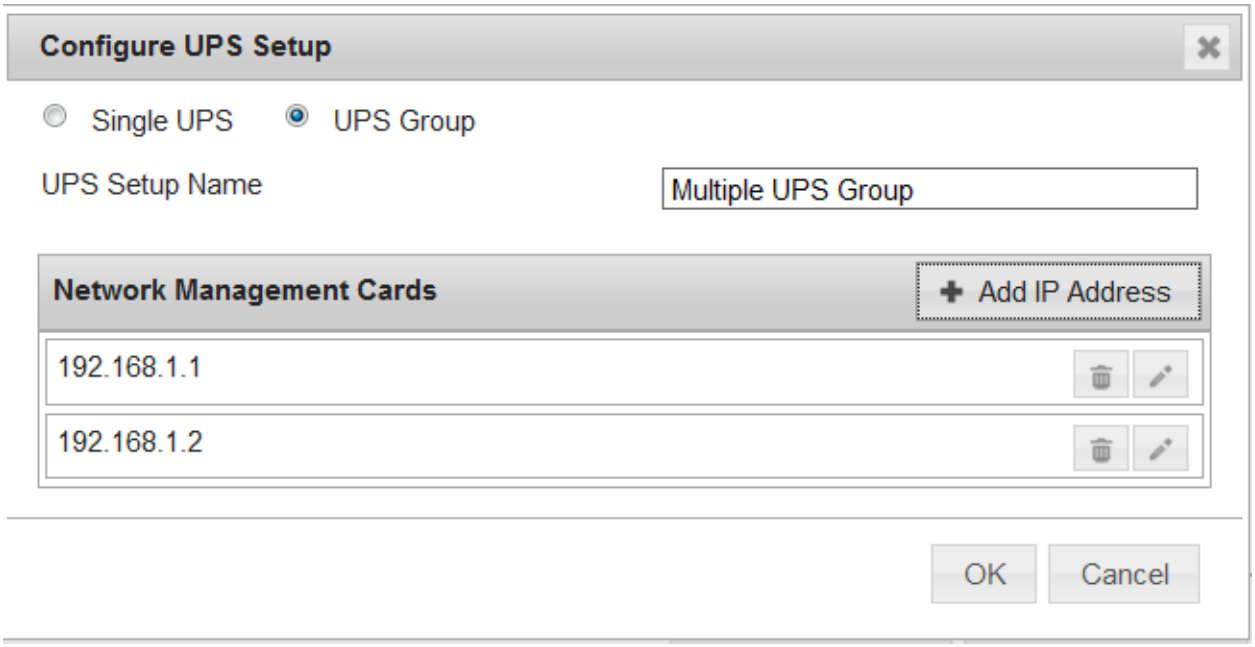

1. Enter a **UPS Setup Name** (with a maximum of 20 ASCII characters)
2. Click the **+ Add IP Address** button and enter the IP address of the Network Management card in the UPS. Click **OK**.
3. Repeat for each of the UPS devices to be added to the UPS group. A minimum of 2 IP addresses is required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

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.

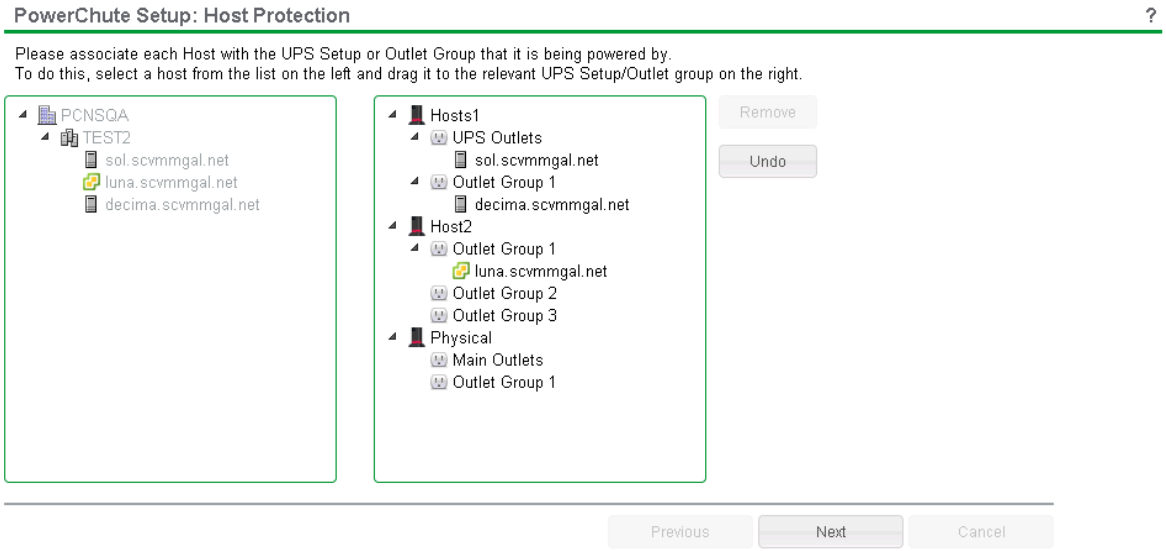

PowerChute Setup - VMware hosts associated with more than one outlet group on the same NMC

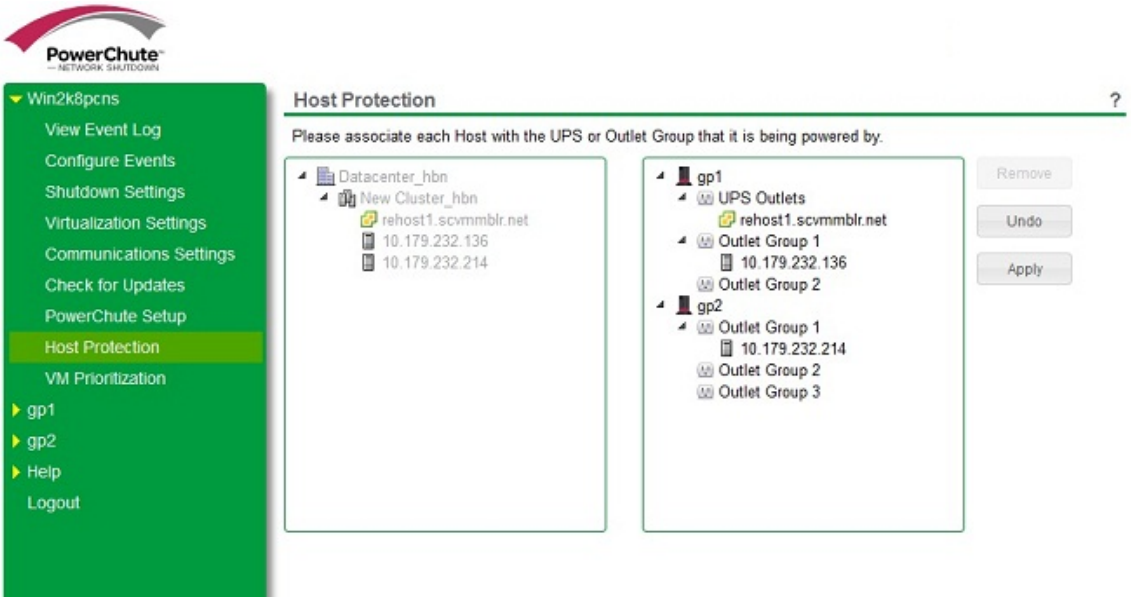

Host Protection - VMware hosts associated with more than one outlet group on the same NMC

Network Management Card Settings

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

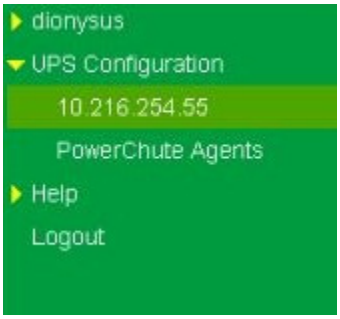

For Advanced UPS configuration, each UPS Setup is displayed as a menu item and the IP address of the NMC(s) with which PowerChute is communicating is displayed under each UPS setup.

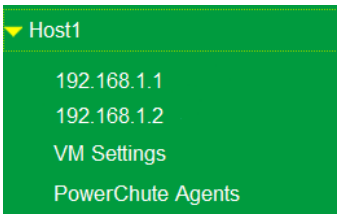

Click on the IP address to view the UPS information specific to that NMC. If the NMC has VMware Hosts associated with it, you can edit these settings for a specific NMC on this page. This overrides the global NMC settings configured via the initial PowerChute Setup or via the Shutdown Settings screen.

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:

- 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.

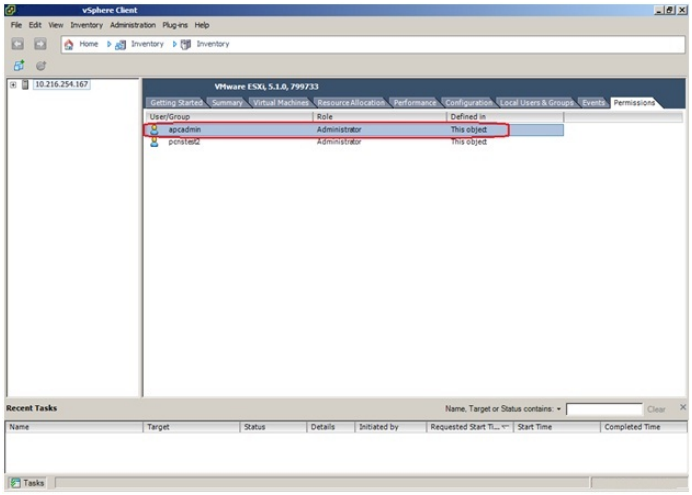

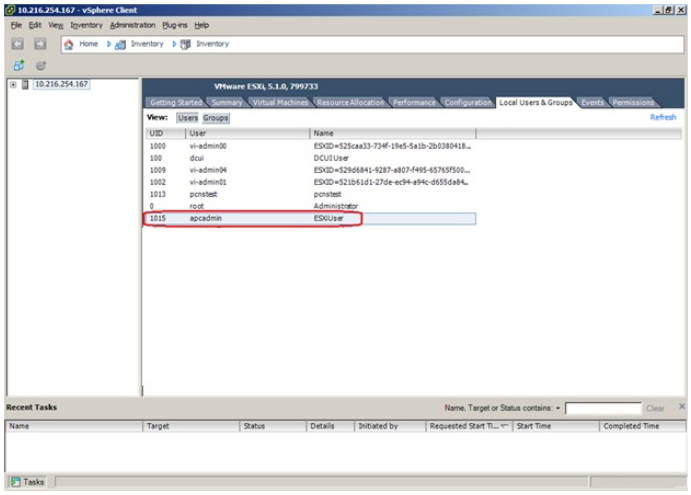

PowerChute does not directly shut down virtual machines with this configuration. In order to gracefully shut down VMs, you must use the **Virtual Machine Startup/Shutdown** settings for the VMware host in the vSphere Client. This is accessible on the **Configuration** tab under **Software** in the vSphere Client.

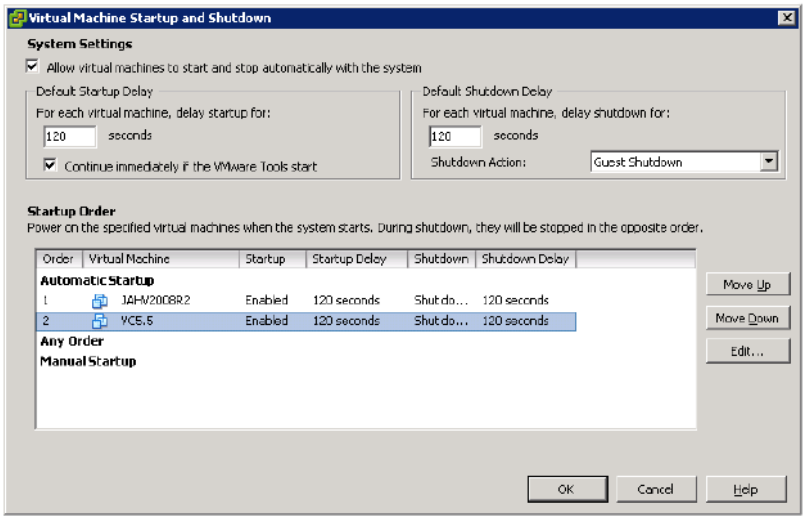

The shutdown action must be changed to "Guest Shutdown".

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 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.

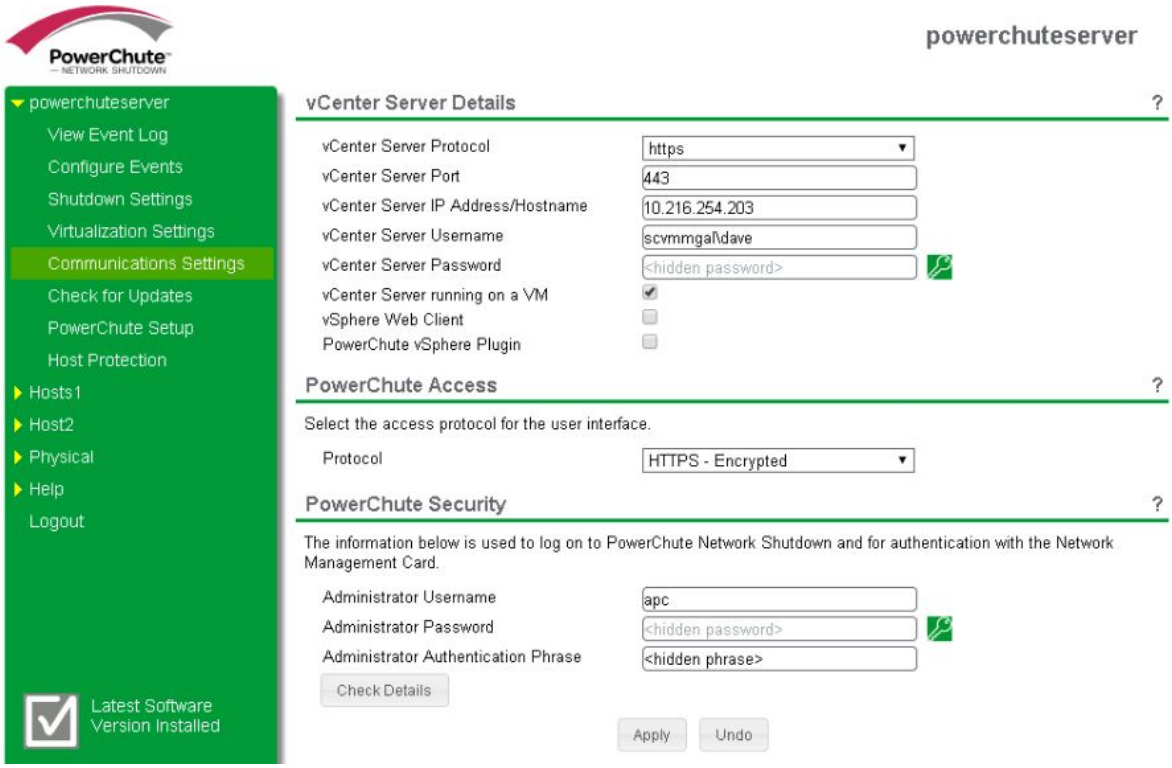

Communications Settings - Main UI - vCenter Server running on a VM option

If vCenter Server is running on a VM managed by vCenter Server, the option **vCenter Server Running on a VM** should 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.

The checkbox is only used to aid troubleshooting – it does not control whether PowerChute will perform vCenter Server VM shutdown. If unchecked and PowerChute locates vCenter Server running on a VM on an affected Host it will still shut down the Host.

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, 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.

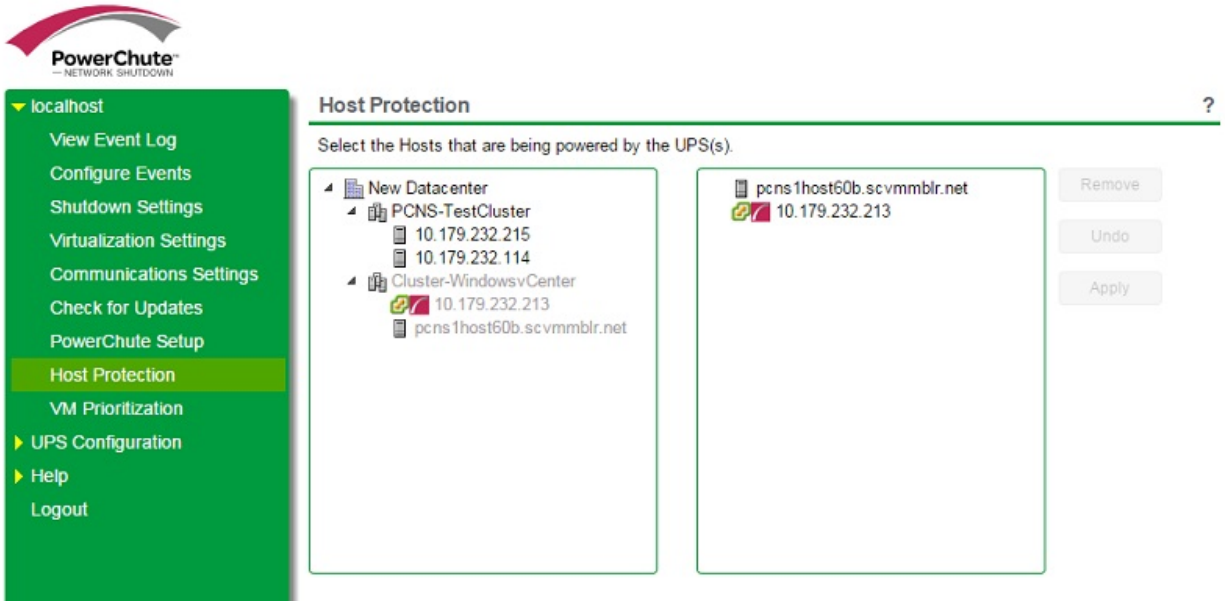

When a critical UPS event occurs, PowerChute will shut down VMs and the Hosts in the order that they appear in the right-hand panel. You can change this order by clicking on a host in the right-hand panel and dragging it up or down. PowerChute will re-start VMs on each host in the reverse order that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

If vCenter server or PowerChute is running on a VM they are shutdown 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

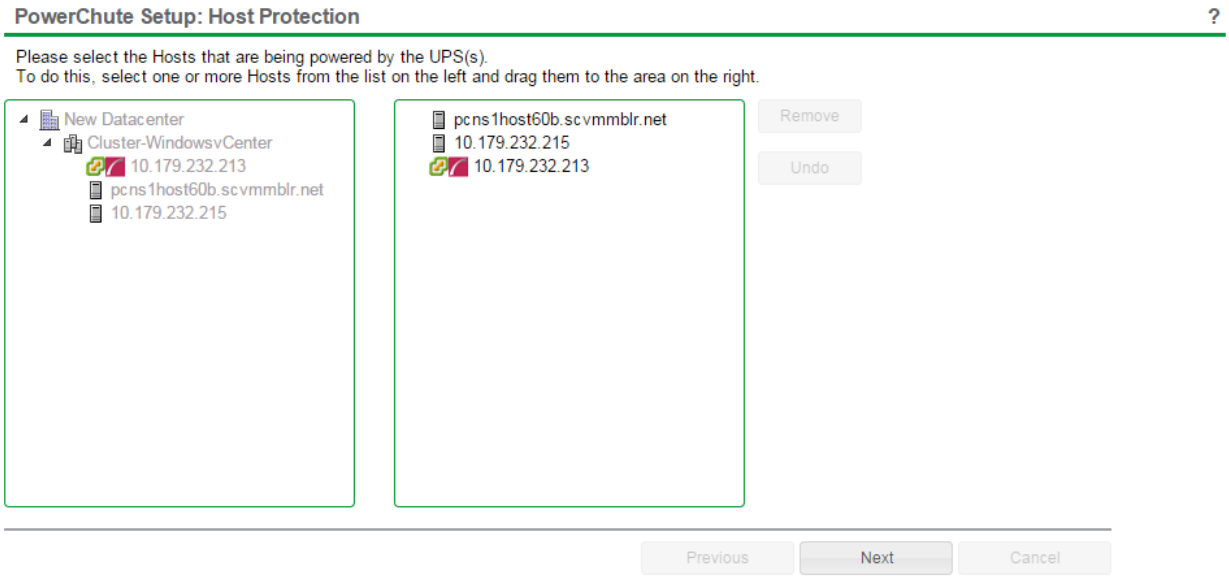

If vCenter is running on a VM on one of the ESXi hosts and none of the Hosts are highlighted with the logo, or the wrong Host is highlighted, this indicates there is a configuration issue in vSphere setup that will prevent PowerChute from shutting down the vCenter Server VM correctly. For more information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

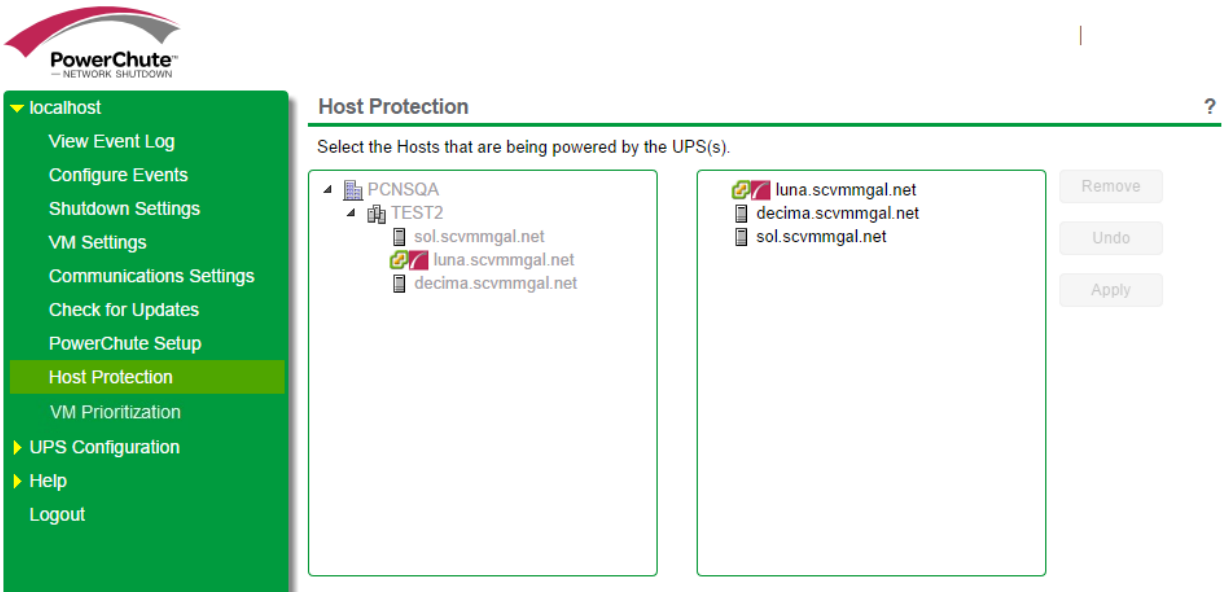

Advanced UPS configuration

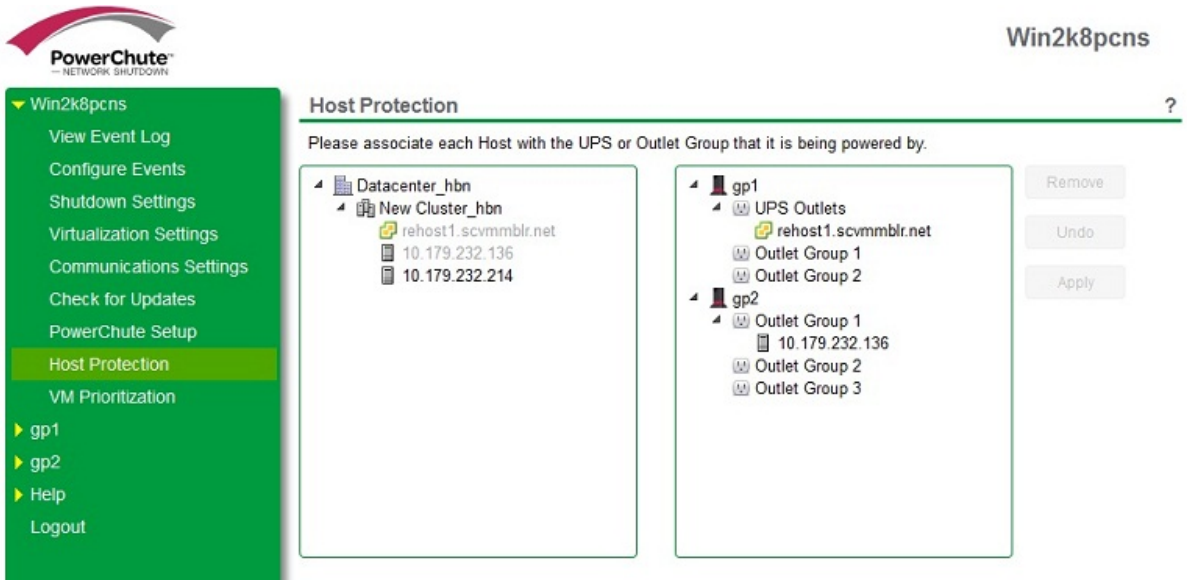

The VMware hosts are powered separately by one or more UPS(s). The right-hand panel shows the UPS(s) that PowerChute is registered with. You need to associate each VMware host in the left-hand panel with the UPS by which it is powered. To do this, drag each host to the UPS/UPS group in the right-hand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

For this configuration, PowerChute must be installed on a physical Windows machine outside the cluster. If PowerChute is installed on a Virtual Machine, a dialog displays to note that Advanced UPS Configuration is not supported.

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 Sequence -** 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.

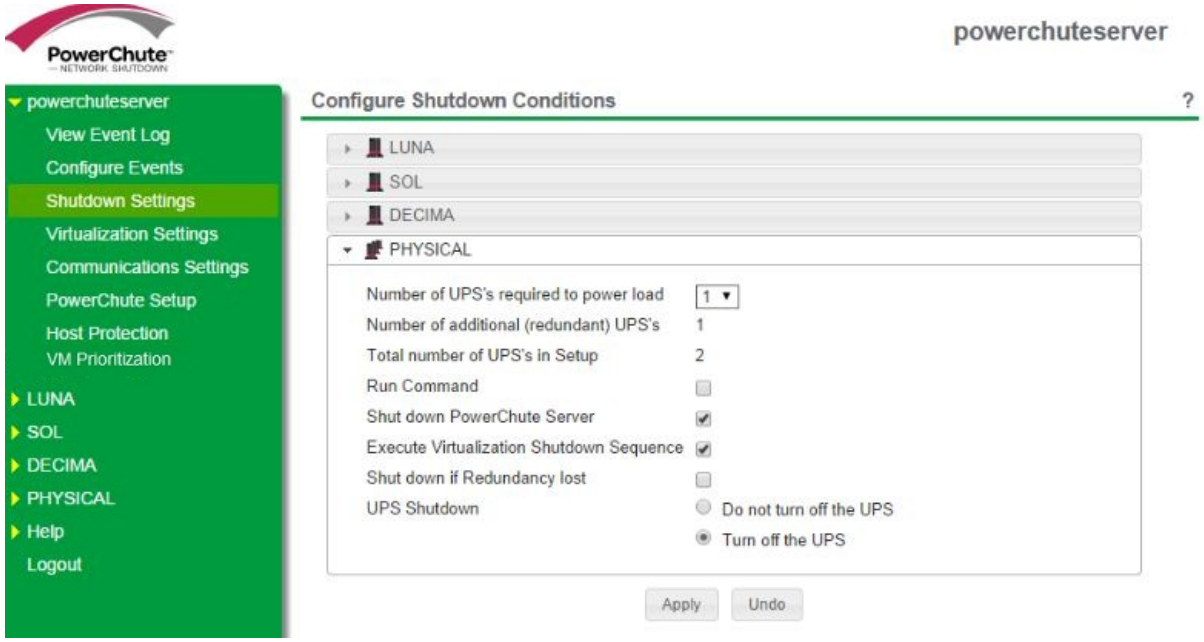

PowerChute User Interface - Configure Shutdown Conditions

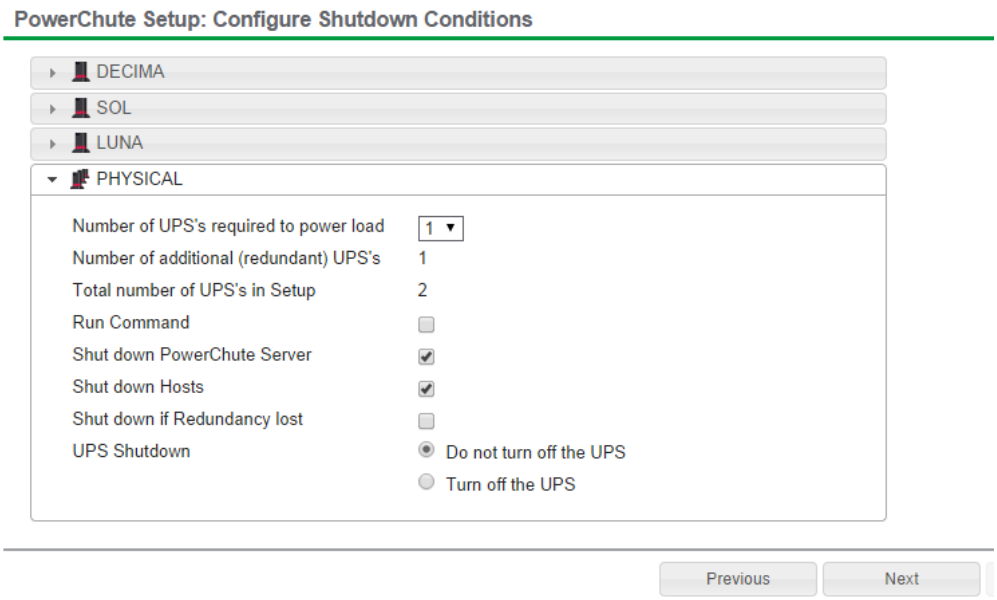

PowerChute Setup - Configure Shutdown Conditions

Virtualization Settings

Virtual machine (VM) settings like VM migration and VM shutdown can be configured in Virtualization settings.

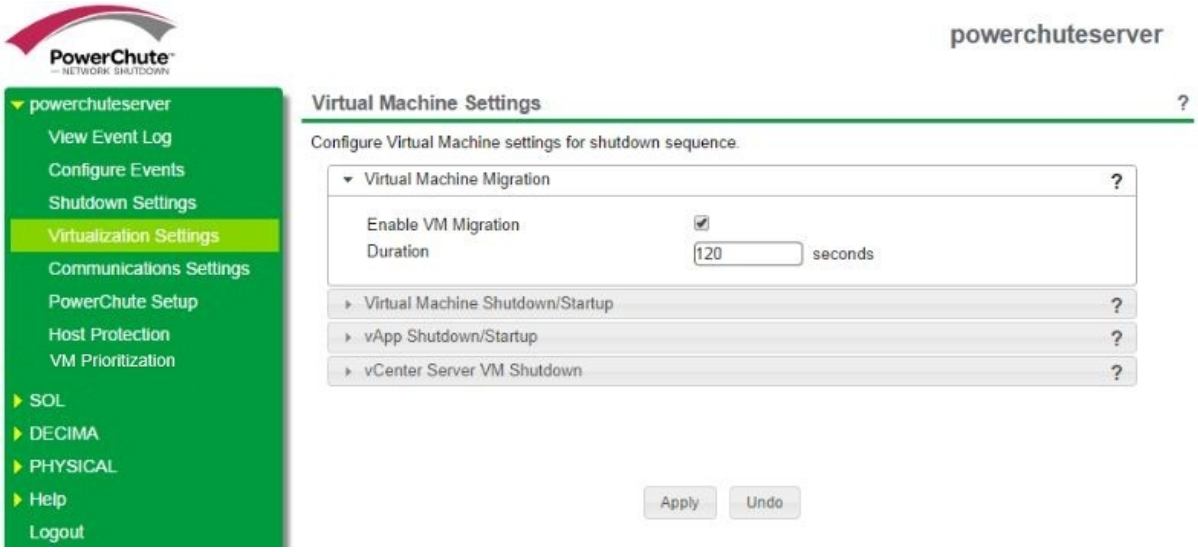

For more information see:

- Virtual Machine Migration
- Virtual Machine Shutdown/Startup
- vApp Shutdown/Startup

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.

PowerChute does not support installation on a host that is part of a vSAN enabled cluster. In a configuration that includes a vSAN enabled cluster, PowerChute must be installed on a physical machine, or on a host that is not part of the vSAN enabled cluster.

When PowerChute starts a maintenance mode task via vSphere API it does not pass any arguments for vSAN data migration option. Because of this the default option **Ensure accessibility** is used.

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 check box **Apply VM settings to all UPS Setups** is selected. This checkbox is enabled by default.

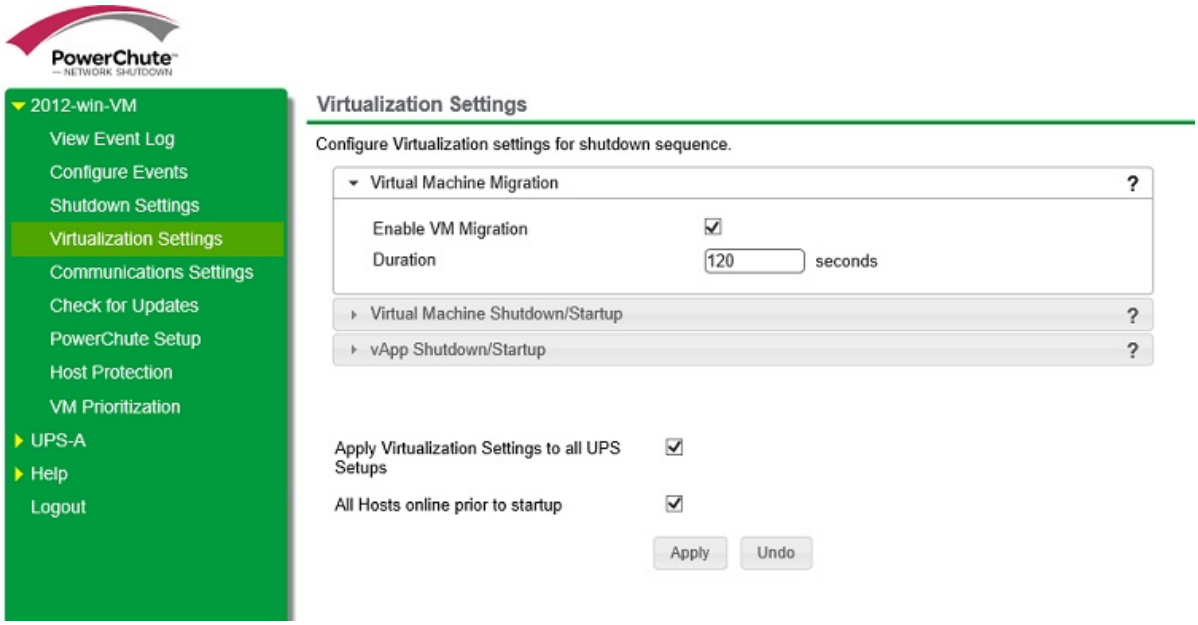

If you have applied settings to individual UPS Setups you should uncheck this option to prevent them from being overwritten.

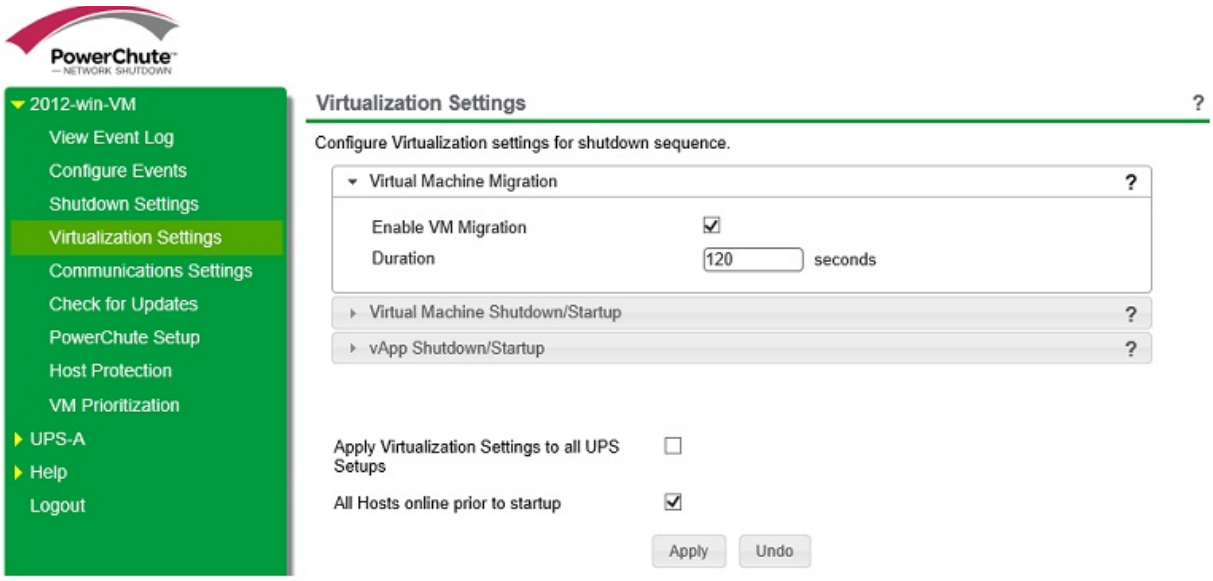

Apply VM Settings unchecked

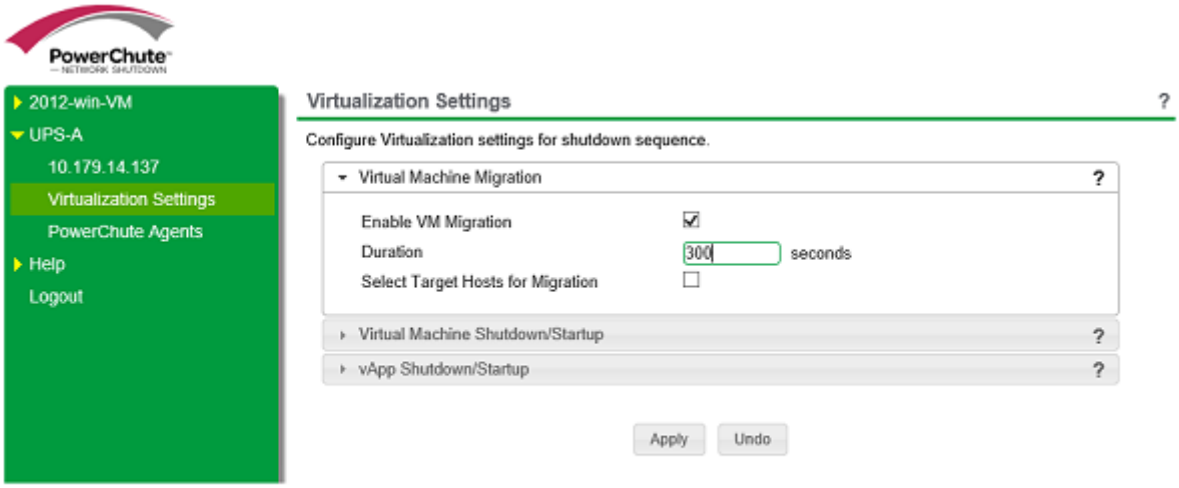

Individual UPS Setup Settings applied

Virtual Machine Migration

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.

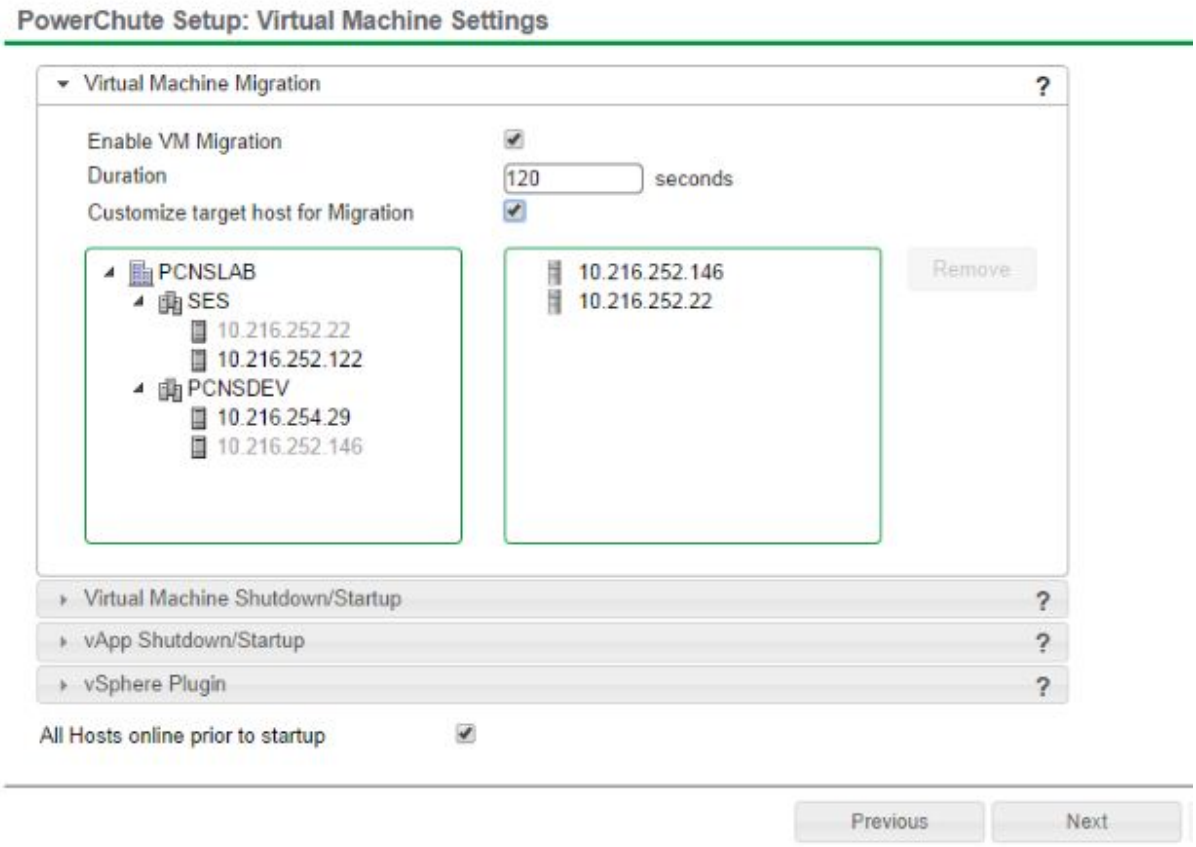

Select target host for Migration - PowerChute Setup Single/Redundant/Parallel UPS Configuration

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.

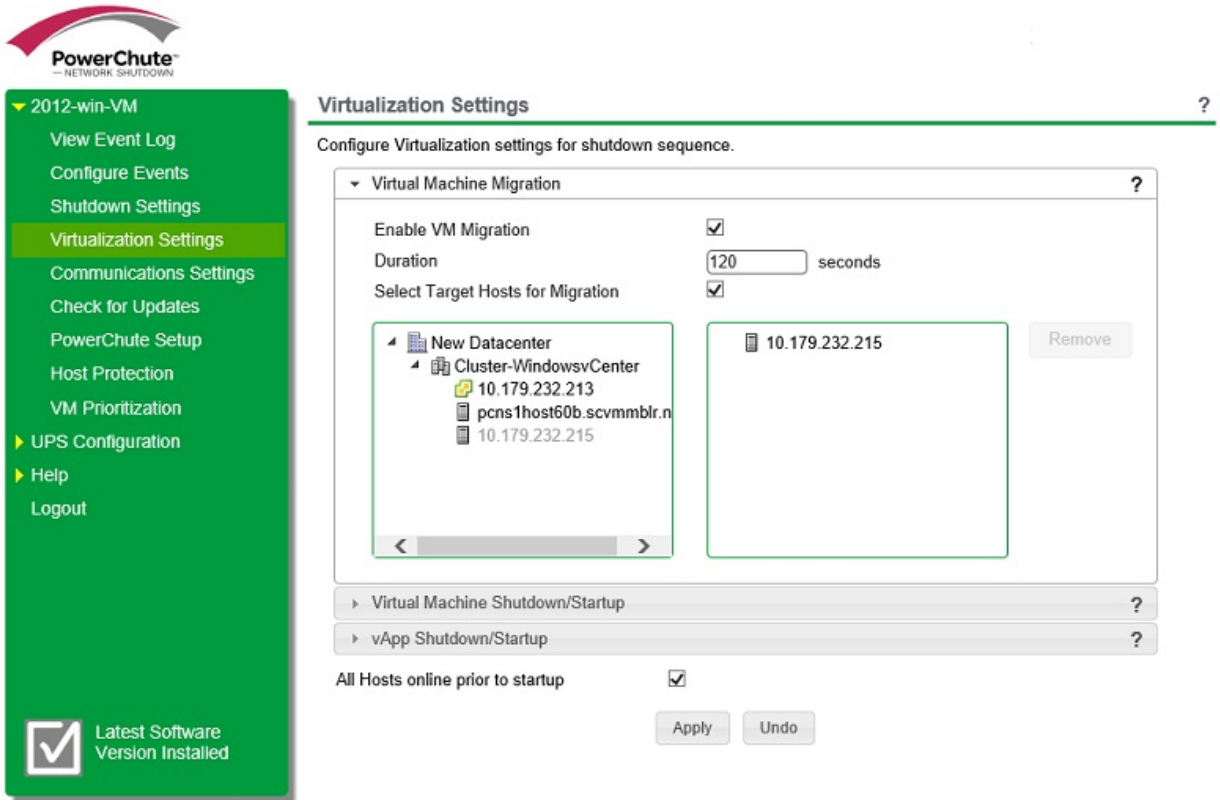

Select target host for Migration - Single/Redundant/Parallel Configuration

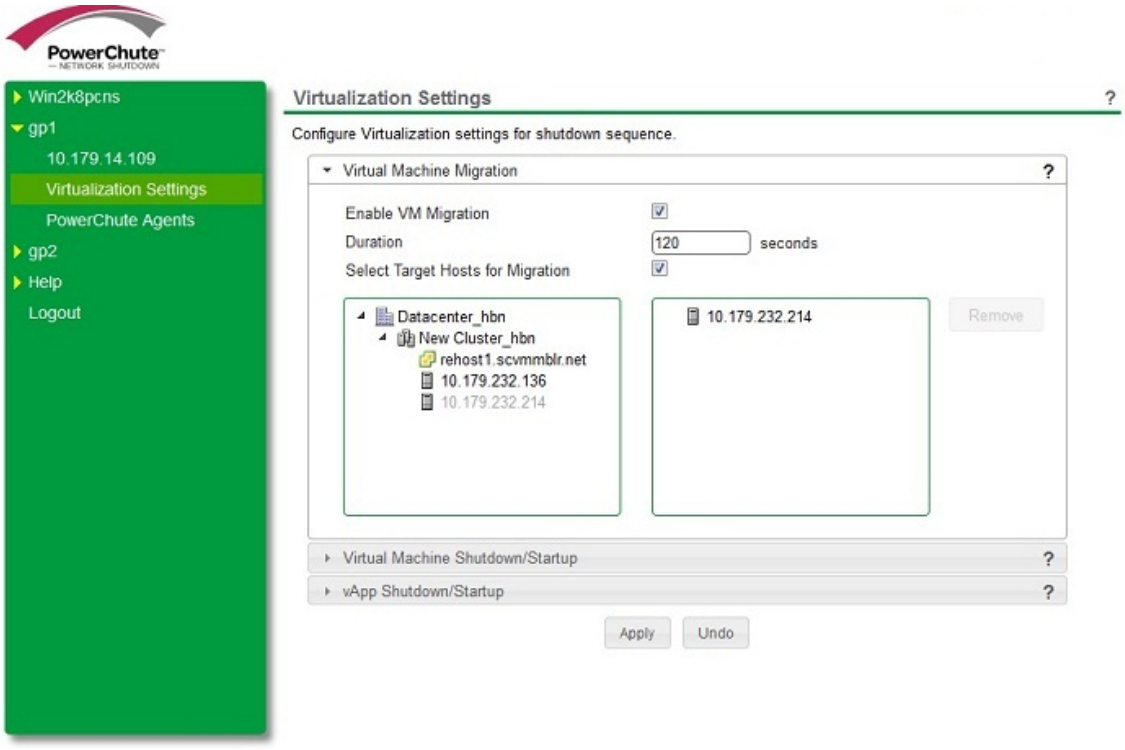

Select target host for Migration - Advanced Configuration

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, and allow DRS to migrate the VMs to other hosts in the cluster.

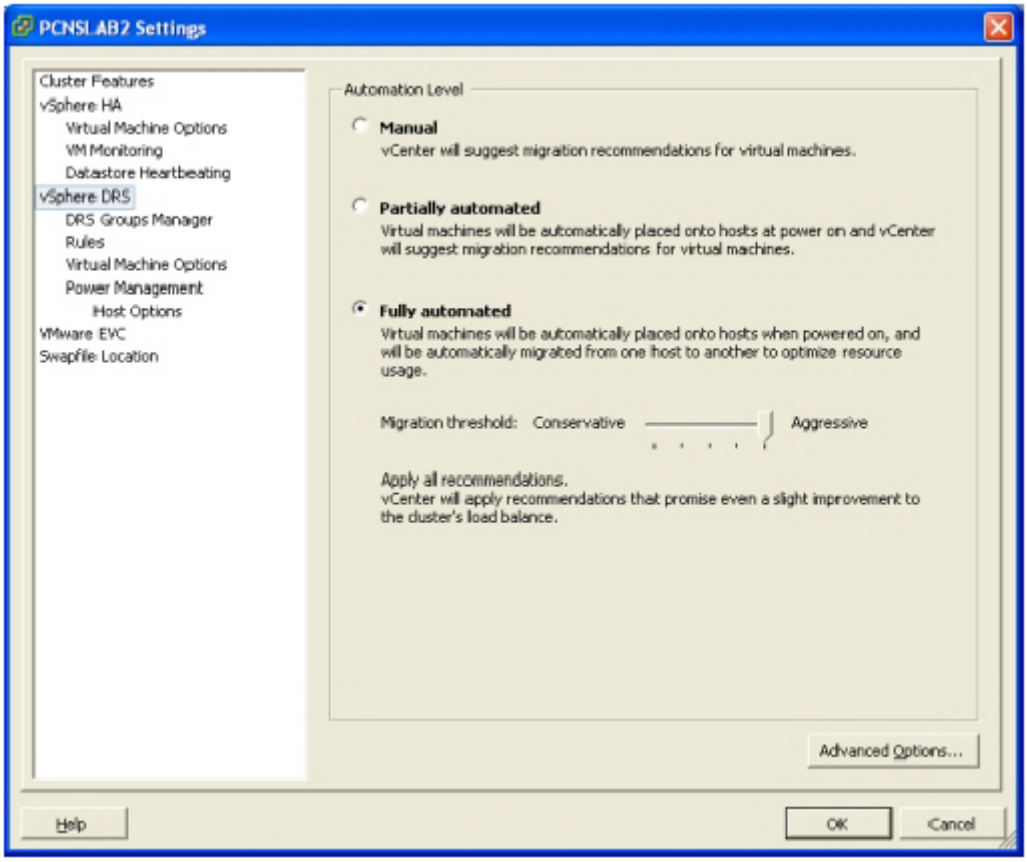

DRS is enabled in the Cluster Settings dialog of vSphere. To access Cluster settings, right-click on a cluster, choose Settings.

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

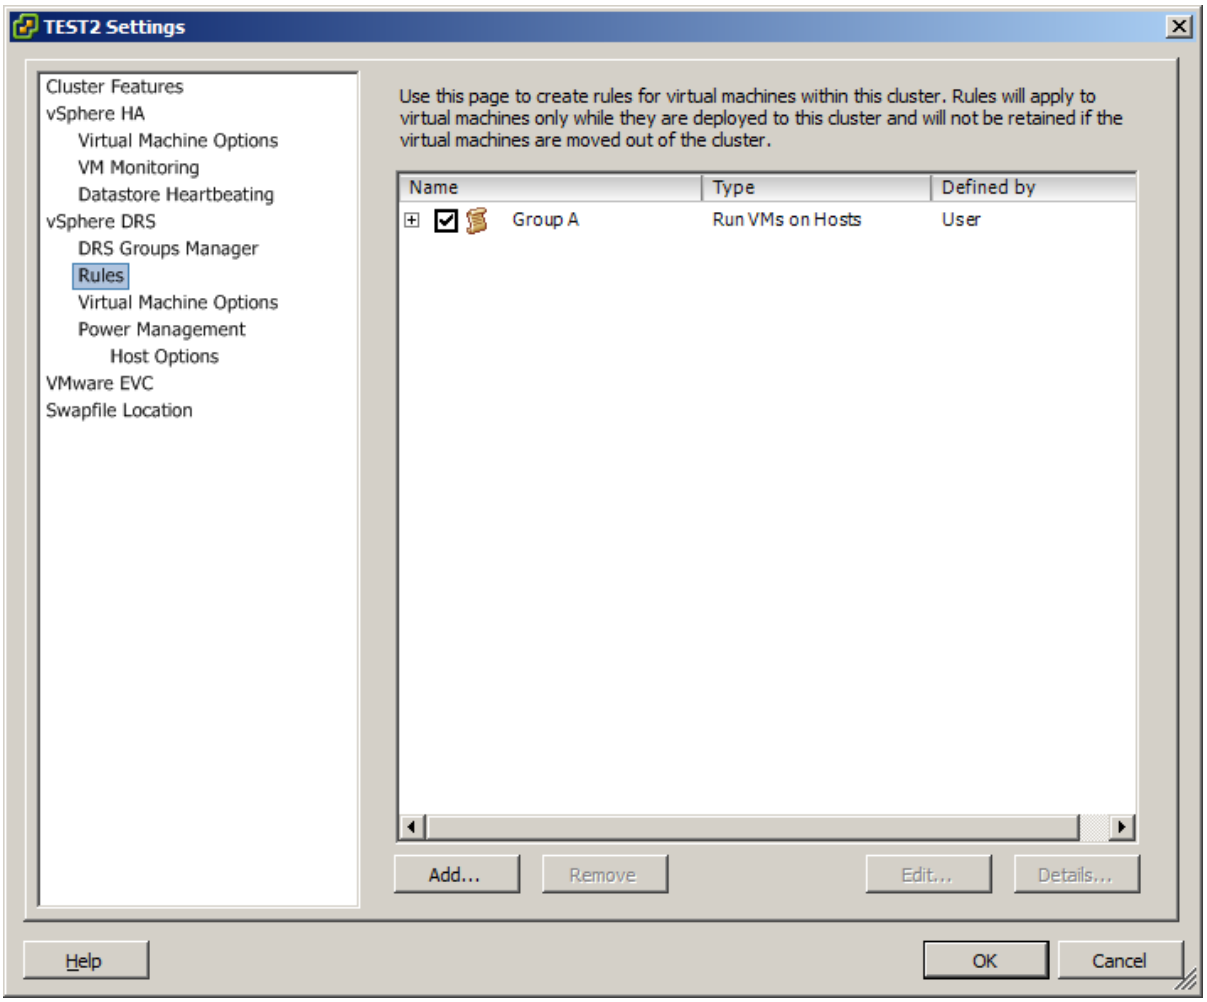

1. In the Cluster Settings dialog choose Rules under vSphere DRS. Click the Add button.

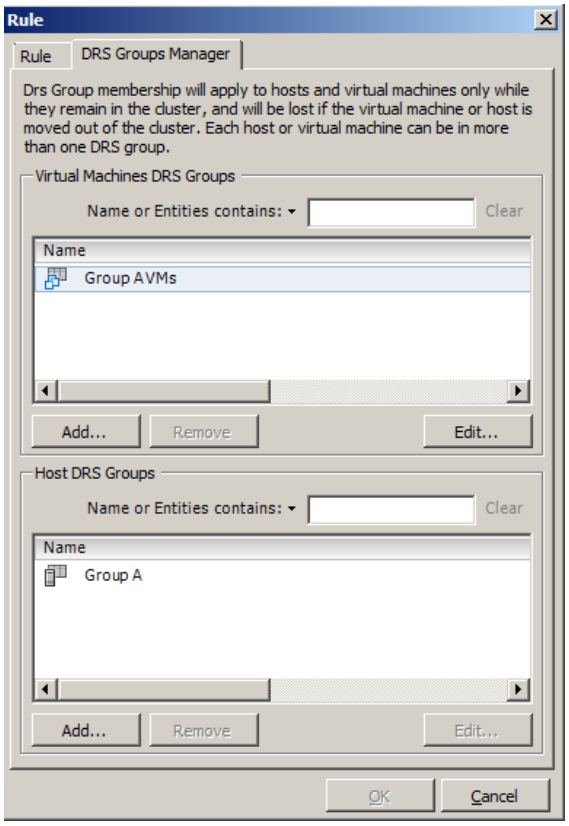

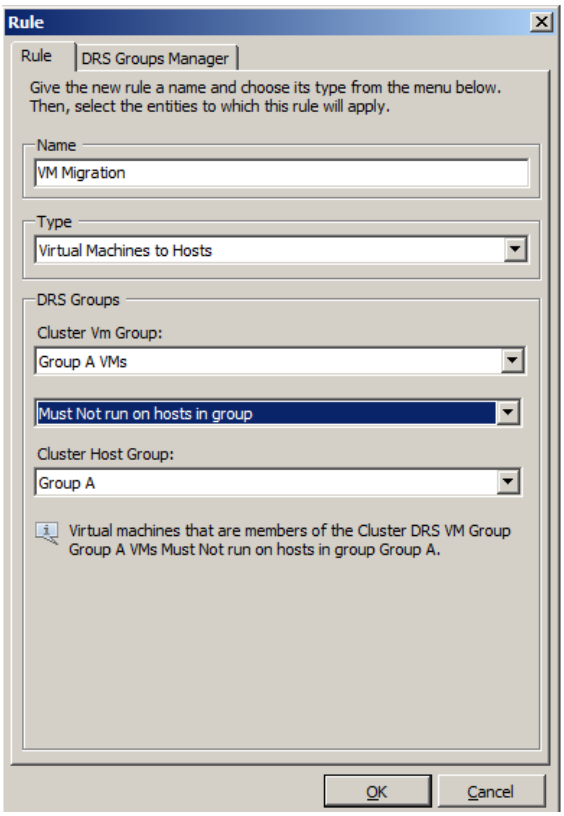

2. Choose the DRS Groups Manager tab:
3. Add a VM DRS Group - this contains a list of VMs to which the rule applies.
4. Add a Host DRS Group - this contains a list of Hosts to which the rule applies.
5. On the Rules tab, specify a rule for the VM and Host DRS groups. For example, to prevent VMs from being migrated to Hosts in the Host DRS Group created in step 4, select **Must Not run on hosts in group**.

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. 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.

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 Shutdown/Startup

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

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

If vCenter Server is running on a VM, it is shut down once all other VMs have been shut down. There is a separate duration for shutting down the vCenter Server VM, which is configurable.

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 .

Running PowerChute on a VM in an Advanced UPS configuration is not supported. The recommended setup is to install PowerChute on a physical Windows machine. This is recommended because if PowerChute is running on a VM, it will not be able to monitor UPS's until it is restarted by HA.

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.

Re-starting after a shutdown

Selecting **Enable VM Startup** re-starts any VMs 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.

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

In **Single**, **Redundant**, and **Parallel** UPS configurations, the option Wait for all Hosts Online is enabled by default. When enabled, PowerChute waits until all hosts in the cluster are back online before starting the VMs. The VMs are started on each host in the reverse order to how the hosts were shut down.

The VMs on the first host are started and PowerChute waits until the VM Startup duration has elapsed before starting VMs on the next host. The VM Startup duration is the delay between starting VMs on each host.

If vCenter Server is running on a VM and it was shut down by PowerChute, 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.

Disable option **Wait for all Hosts Online** to allow PowerChute to attempt to start VMs 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 on other hosts.

If vCenter Server was offline or unavailable when VMs were shut down, PowerChute will start the VMs 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 enabled 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 on other hosts. If the **Wait for all Hosts Online** option is disabled, PowerChute will attempt to start VMs on VMware hosts as they become available.

If PowerChute is running on a VM, hosts will remain in maintenance mode when they re-start 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.

HA Admission Control

HA Admission Control is enabled by default in vSphere. To change Admission Control settings, right click on a Cluster in the Inventory and select Settings.

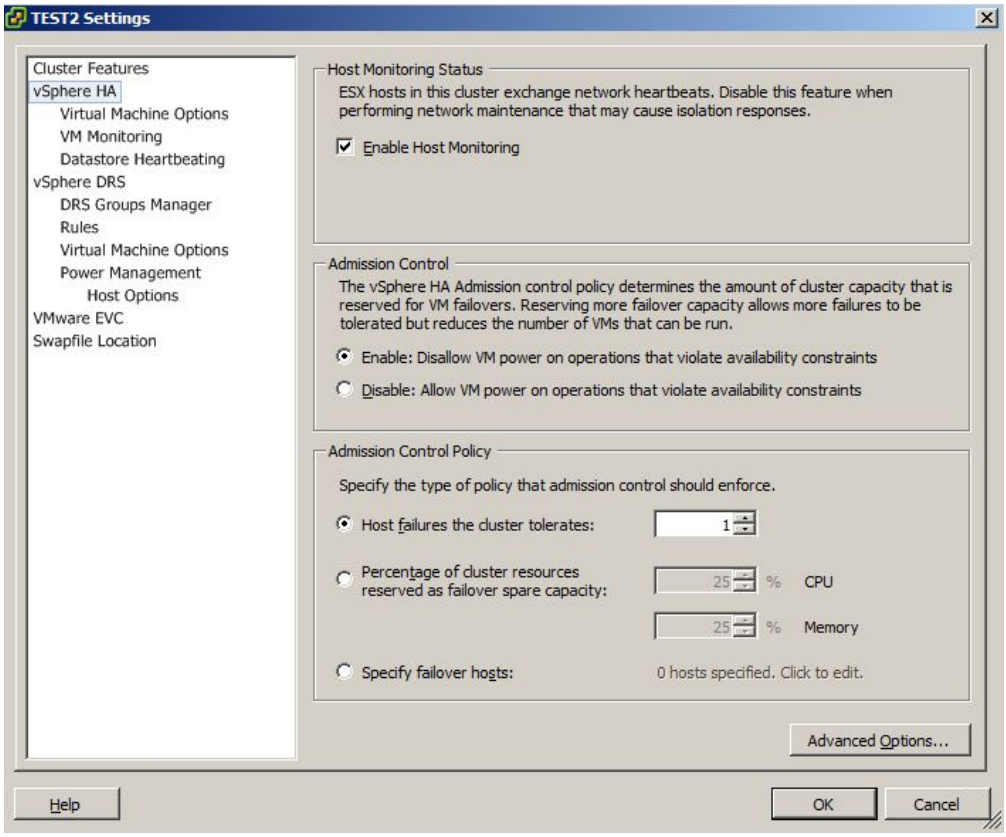

If PowerChute is installed on a Virtual Machine, it may not be restarted automatically when its Host is powered on after a critical event has been resolved. This can occur if HA Admission Control is enabled or if the Admission Control policy being enforced prevents it.

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.

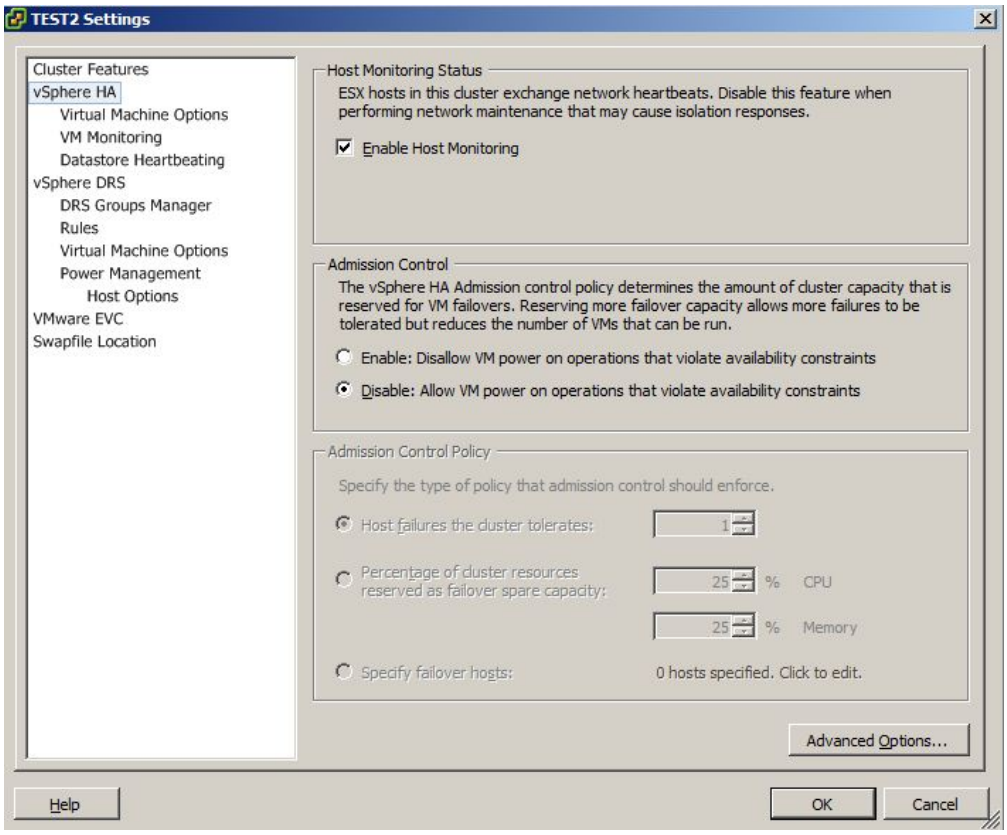

When Admission Control is disabled, HA will attempt to automatically start the VM on which PowerChute is running and PowerChute can begin to monitor associated UPS devices and automatically restart the VMs that it shut down.

To troubleshoot VM Startup issues see VMware Troubleshooting.

vApp Shutdown/Startup

With vApp shutdown, the order of VM shutdown and startup can be configured if required. To do this, set the order in which the VMs are shut down through vCenter Server.

Right-click the vApp in the left-hand panel of vCenter Server, click on **Edit Settings** and then on the **Start Order** tab.

Shutdown Action - **Operation** must be changed from **Power Off** to **Guest Shutdown** to ensure that the VMs in the vApp are shut down gracefully.

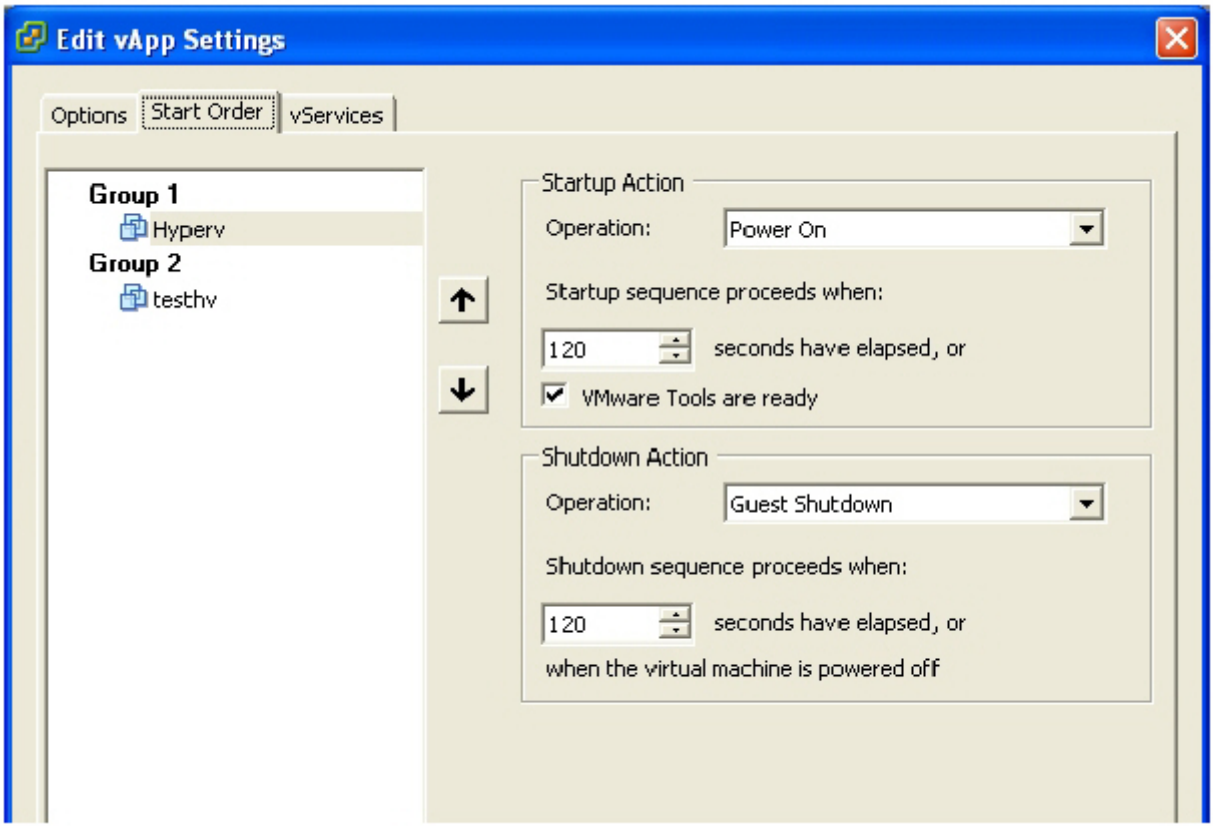

By default, there is a 120 second delay between shutting down each VM in a vApp. The next VM in the vApp will not be shut down until this delay has elapsed or the current VM is powered off. This should be factored into the PowerChute vApp Shutdown duration.

For example: If a vApp contains 3 VMs each with a delay of 180 seconds, the PowerChute vApp shutdown duration should be set to 540 seconds.

Each VM in the vApp must have **VMware Tools** installed.

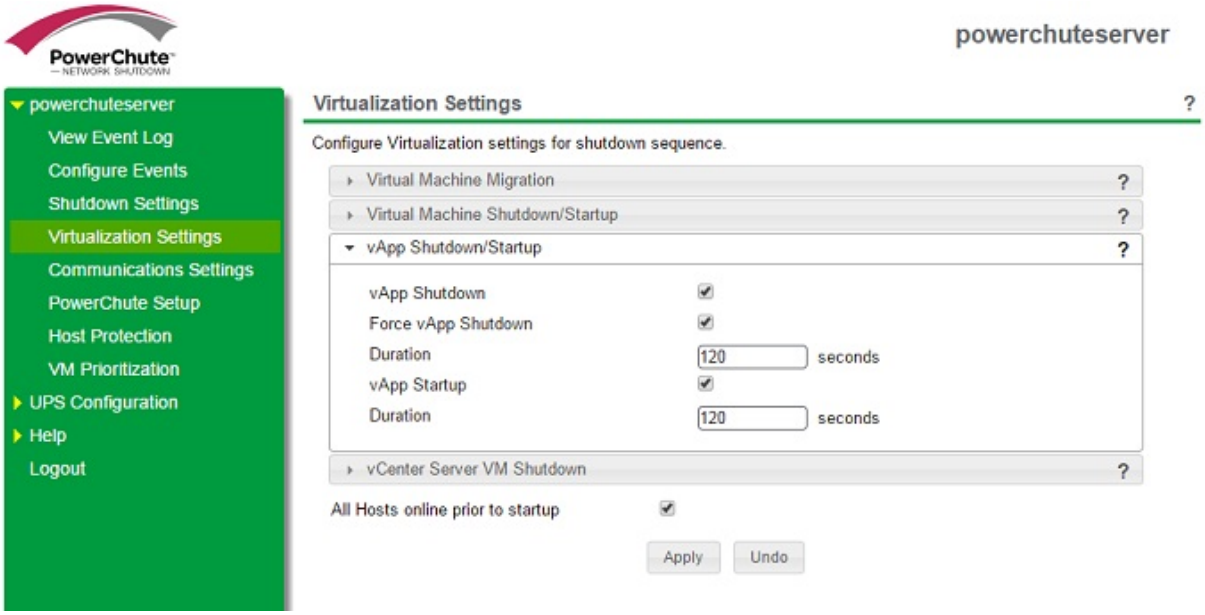

vApp Shutdown/Startup - Duration field

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

You need to ensure that you allow sufficient time using the Duration field to allow your vApp to be gracefully shut down before the hosts are commanded to shut down.

The **Enable Force vApp Shutdown** check box is enabled 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 check box is disabled, vApp shutdown will be skipped if some hosts are still available. In addition, the VMs running on the impacted hosts will be powered off.

Multiple vApps shutdown

The following points apply when there is more than one vApp present:

- PowerChute shuts them down in alphabetical order.
- PowerChute waits for the vApp shutdown command to complete before proceeding to the next vApp.
- The vApp shutdown timer must take into account the time required for the vApp shutdown command to complete for all vApps.

Re-starting after a shutdown

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

If there are multiple vApps in a cluster, vApps will be re-started in reverse order to how they were shut down. PowerChute will wait until the vCenter Server is accessible before starting vApps or VMs.

If vCenter Server is not available when a critical UPS event occurs, PowerChute will shut down any VMs that are part of a vApp (as well as the rest of the VMs that are not part of a vApp).

If PowerChute is running on a VM, hosts will remain in maintenance mode when they re-start 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 vApp.

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. It is available for all UPS configurations – Single, Redundant, Parallel and Advanced.

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

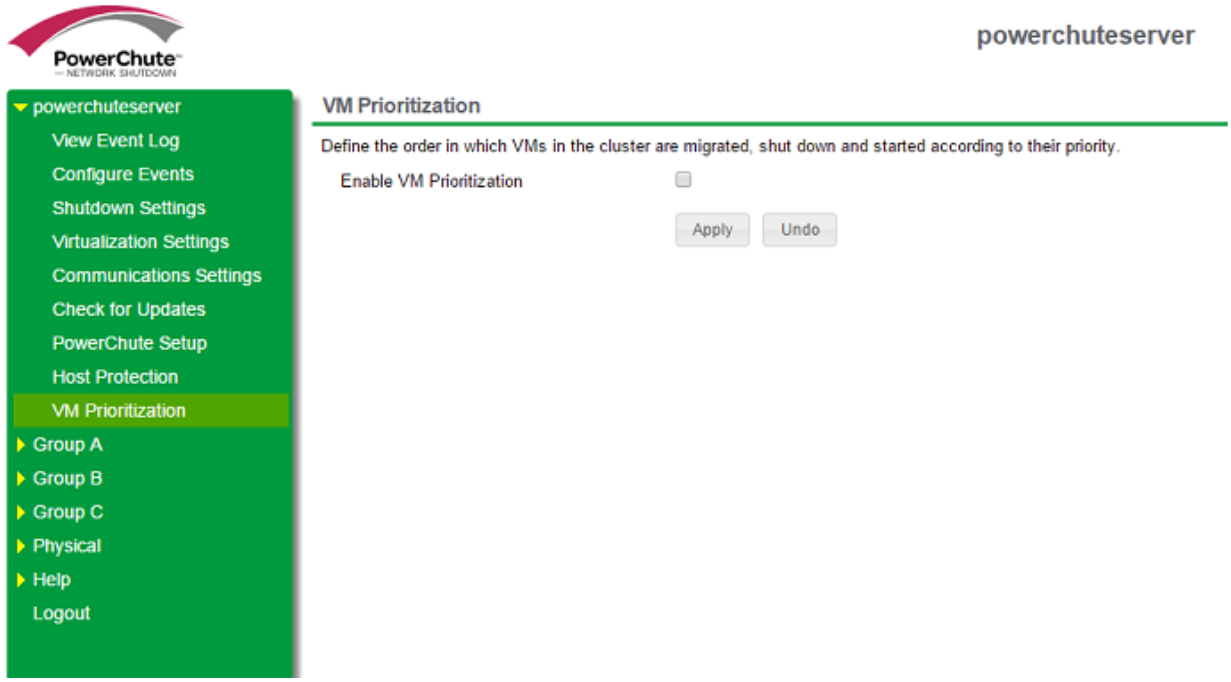

Enable VM Prioritization screen

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 three priority groups – High, Medium and Low. When VM Prioritization is enabled, an **inventory view** of the datacenter, clusters and VMs appears on the left. On the right, the High, Medium and Low priority groups are listed.

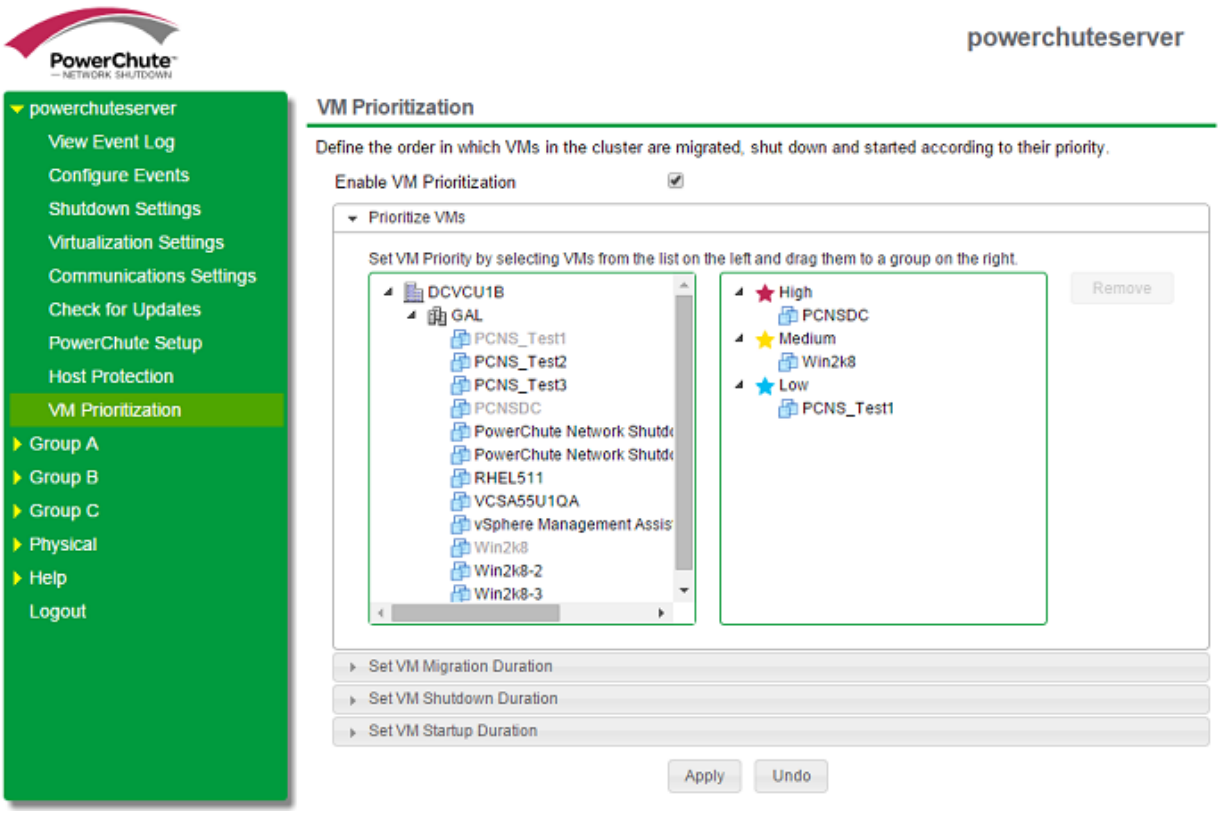

Datacenter inventory view and VM Priority Groups

You can assign a VM to a priority group by clicking on a VM on the left hand side and dragging it to a priority group on the right. You can select multiple VMs by pressing the CTRL key on the keyboard while clicking on the VMs you want to move. You can also click on the cluster icon to select all VMs within that cluster, or click on the datacenter icon to select all clusters of VMs within that datacenter.

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

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.

Note:

- 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 PowerChute and vCenter Server VMs are not displayed.
- 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.
- vApps are not displayed in the list and cannot be added to priority groups.

For PowerChute to correctly identify clusters, VMs 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 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 and Clusters each have a unique name.

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.
- **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.

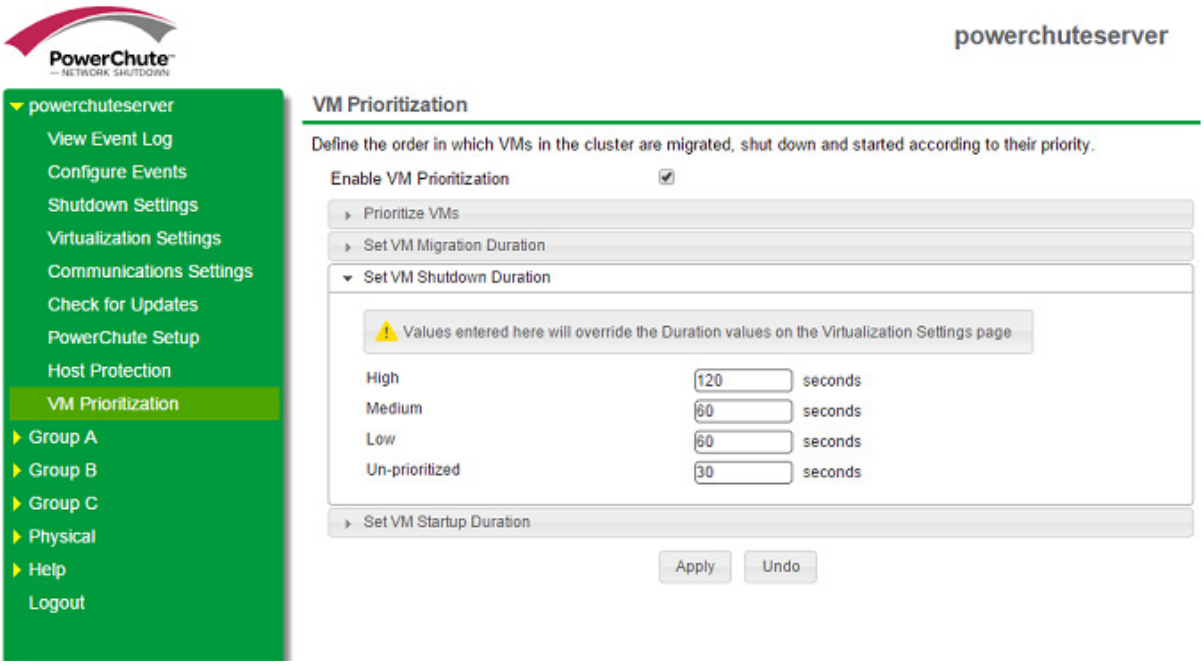

Set VM Operation durations - VM Migration, Shutdown and Startup

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

Durations for un-prioritized VMs are automatically set to the global duration values for VM Migration, VM Shutdown and VM Startup, as configured on the Virtualization Settings page.

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

PowerChute will not permit a VM Shutdown duration of zero; it must have a value greater than zero for High/Medium/Low priority Groups that contain VMs. This ensures that the VMs are shutdown gracefully.

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/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
- Un-prioritized: 60 seconds

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

90 + 60 + 60 + 60 = **270 seconds**

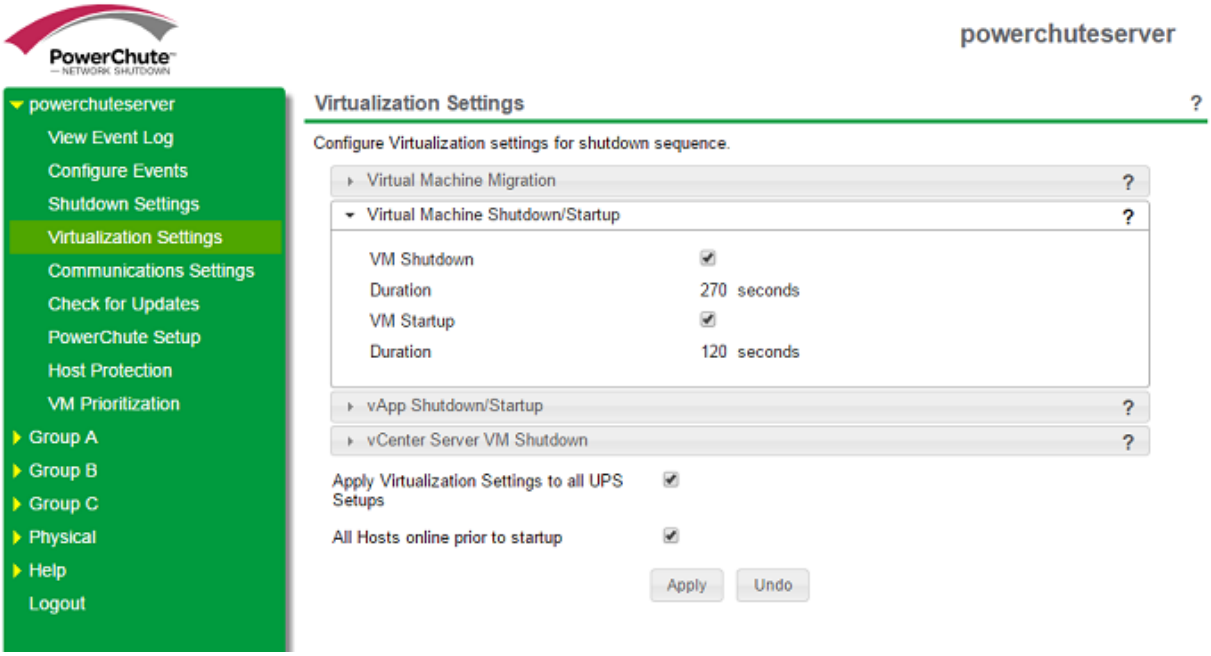

Virtualization Setting page - VM Shutdown duration

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

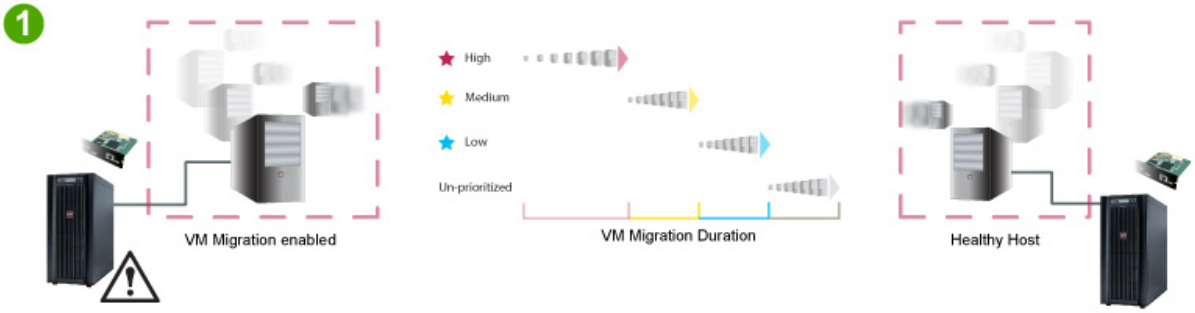

A UPS Critical event occurs and PowerChute starts a maintenance mode task on each protected host. First the High priority group VMs migrate, in parallel. When all High priority Group VMs have migrated, the Medium priority VMs migrate in parallel, followed by the Low priority group and finally the un-prioritized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

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

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

Priority Group VM Shutdown

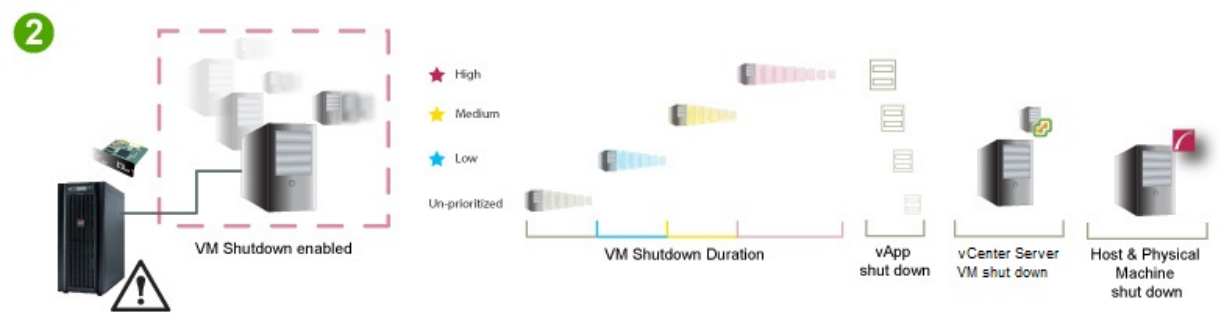

As the critical event continues, PowerChute begins sequenced VM Shutdown. First, the un-prioritized VMs are shut down sequentially. After the duration for un-prioritized VMs elapses, the Low priority VMs are shut down sequentially, followed by Medium priority VMs and finally the High priority VMs are shut down sequentially. When all priority group VMs have been shut down, PowerChute shuts down the vApps, followed by the vCenter Server VM if vCenter Server is running on a VM. PowerChute then shuts down the hosts and the physical PowerChute server.

Priority Group VM Startup

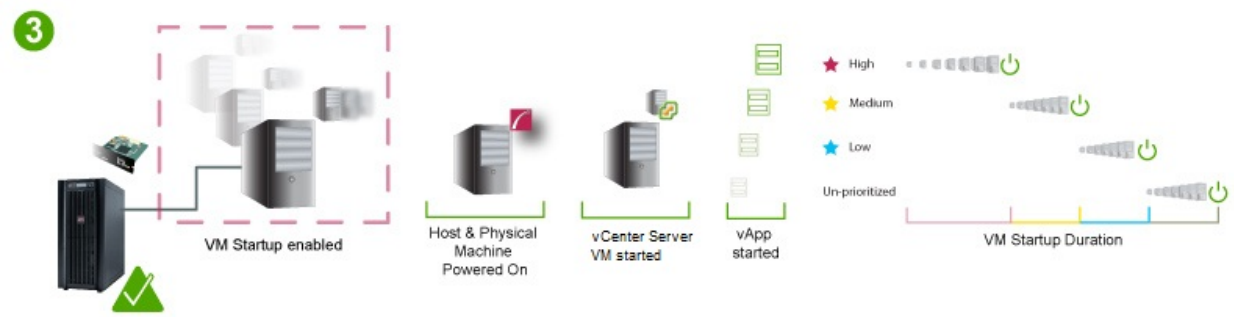

When the UPS critical event is resolved and the hosts and physical machine running PowerChute are powered back on, the hosts are taken out of maintenance mode and the vCenter Server VM is started. The vApps are started first and then PowerChute begins to start the High priority VMs with respect to their VM startup duration. When the High priority startup duration elapses, the Medium priority VMs are started, followed by the Low priority VMs and finally, the un-prioritized VMs 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.

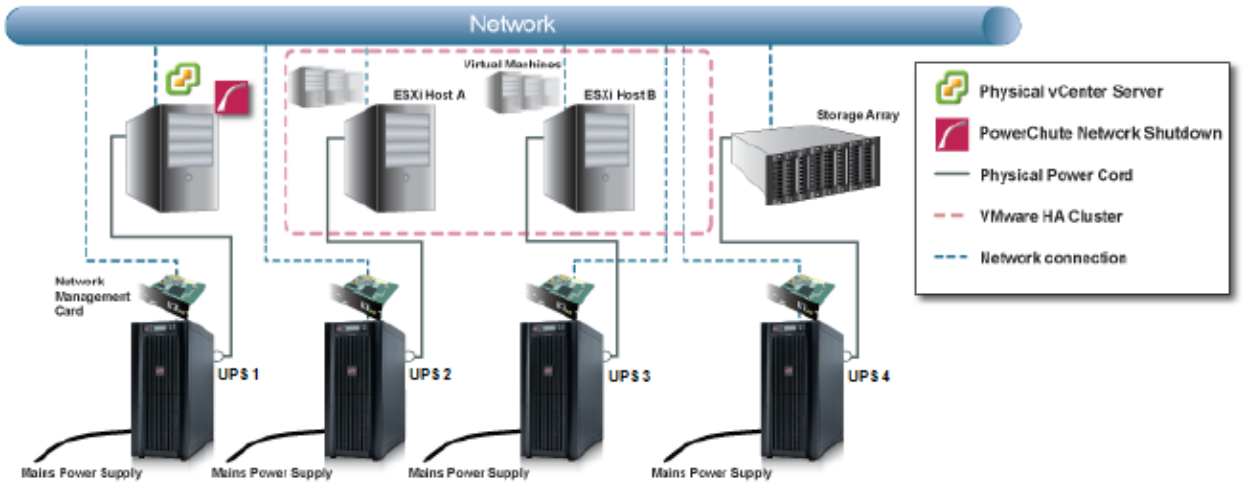

In this example, UPS 2 may experience a critical event at a different time to UPS 3. If the UPS Setups for Host A and B have VM Prioritization enabled, the VM operations that are performed on the priority groups of the two hosts will occur at different times. For example, if UPS 2 experiences a critical 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.

PowerChute vSphere Plugin

Enable the vSphere plug-in option to integrate PowerChute with vCenter Server.

The PowerChute vSphere Plugin is available as:

1. a vSphere Client Plugin (vSphere Desktop Client)
2. a vSphere Web Client plugin

This can be enabled on the Virtualization Settings page in the PowerChute Setup Wizard or on the Communications Settings page. You can enable only one of the plugin options.

To access the vSphere Client plugin:

Log into vCenter Server using the vSphere client and access the PowerChute UI using the views below:

Home - Management View

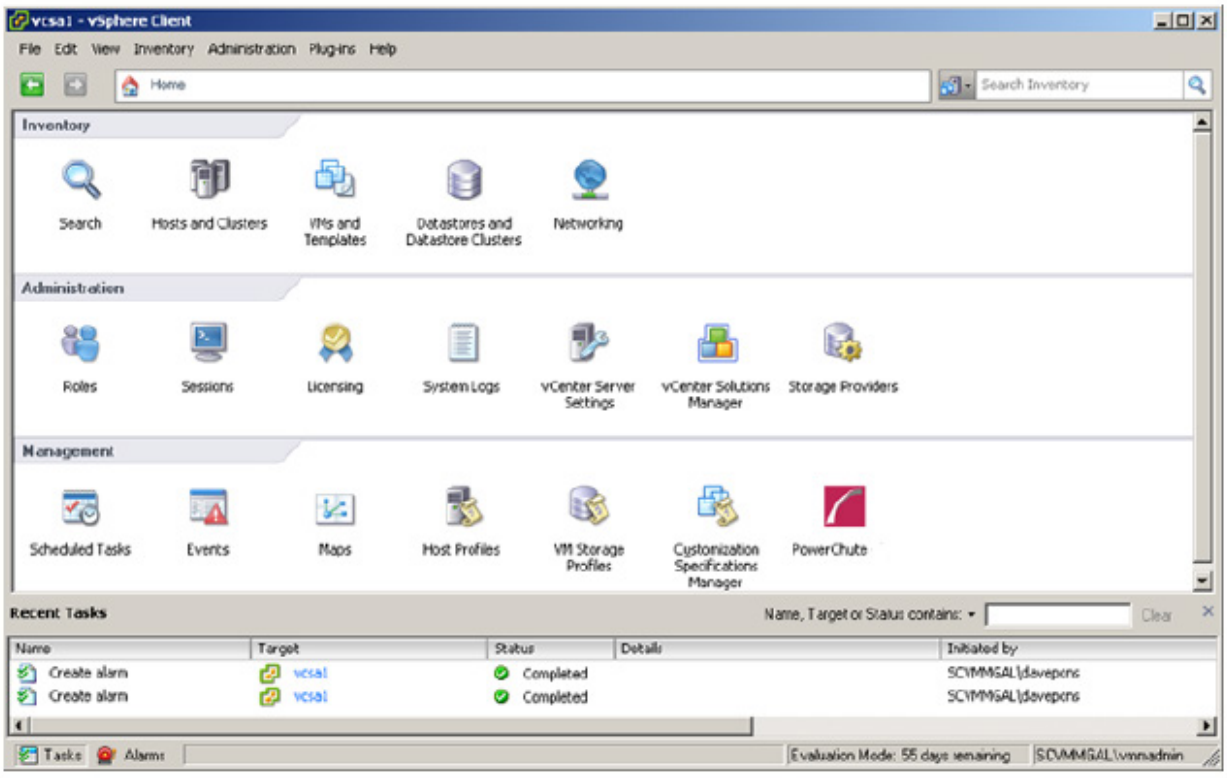

Home - Inventory - Hosts and Clusters view - select the root level in the left-hand pane and click the **PowerChute** tab in the right-hand pane.

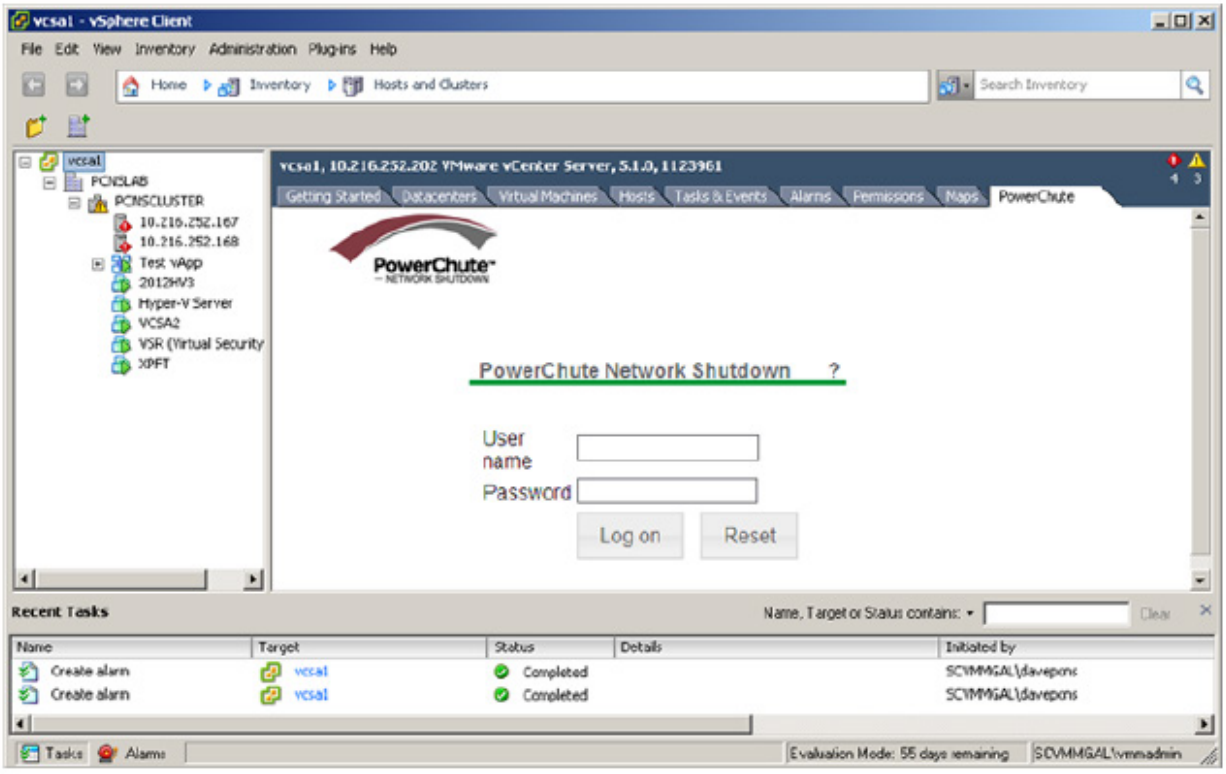

To access the vSphere Web Client plugin:

Log into vCenter Server using the vSphere Web client and access the PowerChute UI by clicking on the Home icon:

Home tab - **Monitoring**

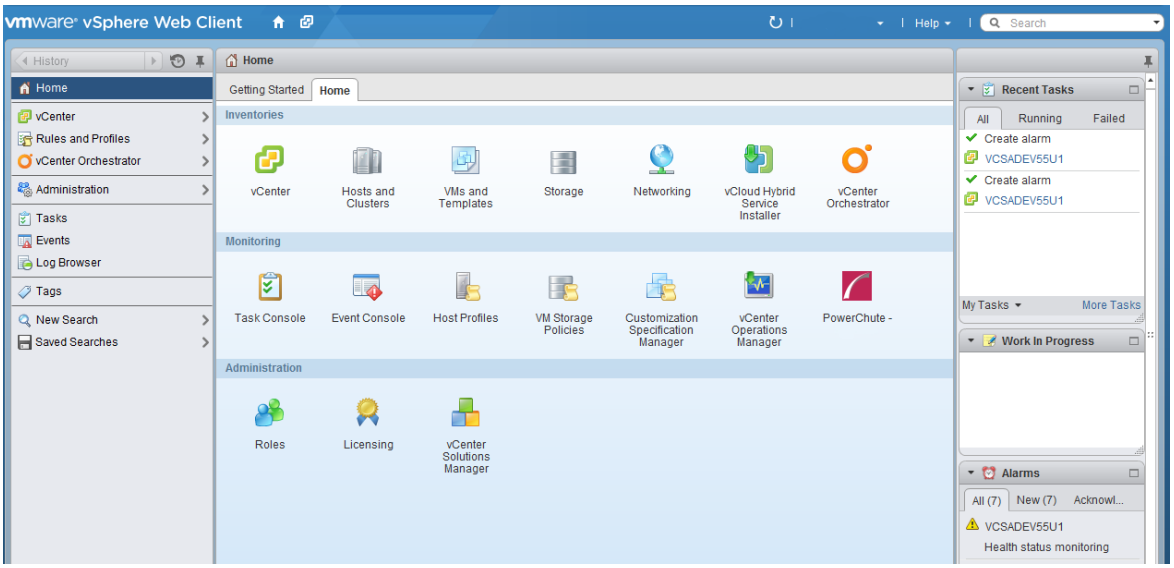

Click on the PowerChute icon to log in to PowerChute:

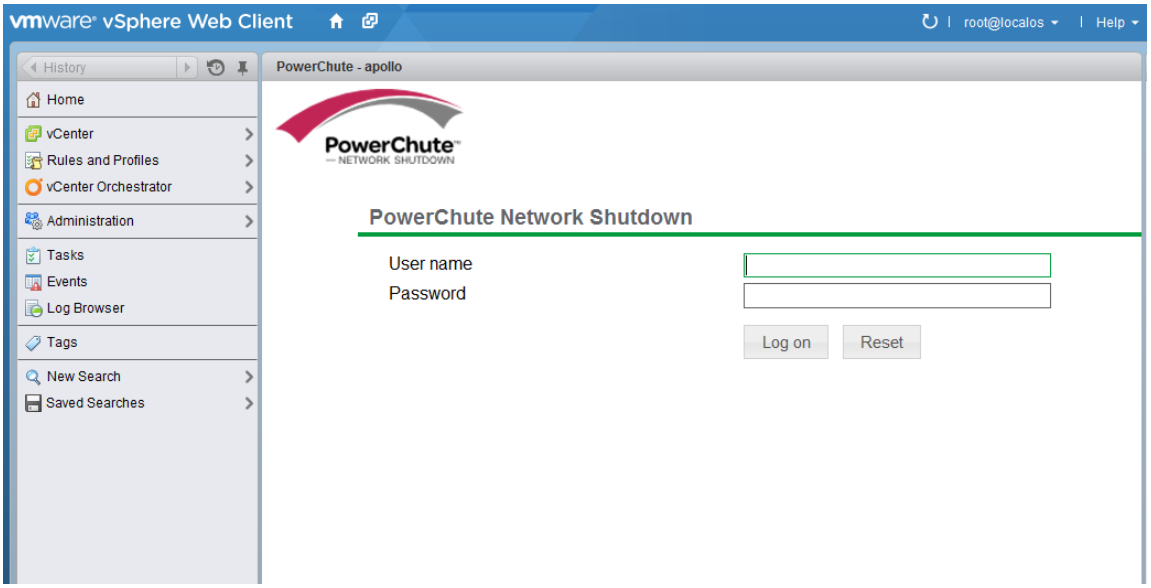

The PowerChute vSphere Web Client plugin is only available for vSphere Web Client v5.5 update 1 or later.

If Internet Explorer Enhanced Security Configuration is enabled you must add the URL for the machine or VM where PowerChute is installed to the Trusted Sites zone.To do this, select Tools-Internet Options in IE and click on the Security tab. Select Trusted Zone and Sites. Add https://<PowerChute hostname/IP address>:6547 to the list of Trusted Sites.

PowerChute vCenter Server Alarms

Enabling either of the vSphere plug-in options also creates a custom PowerChute vCenter Server Alarm.

In the vSphere desktop client plug-in, alarms can be configured to carry out actions using the **Actions** tab in the **Alarm Settings** dialog. For example you can configure an action to send a notification e-mail to an administrator when the alarm is triggered

In the vSphere Web Client plugin, alarms can be configured to carry out actions using the **Actions** tab - **Settings** - **Manage** - **Alarm Definitions**.

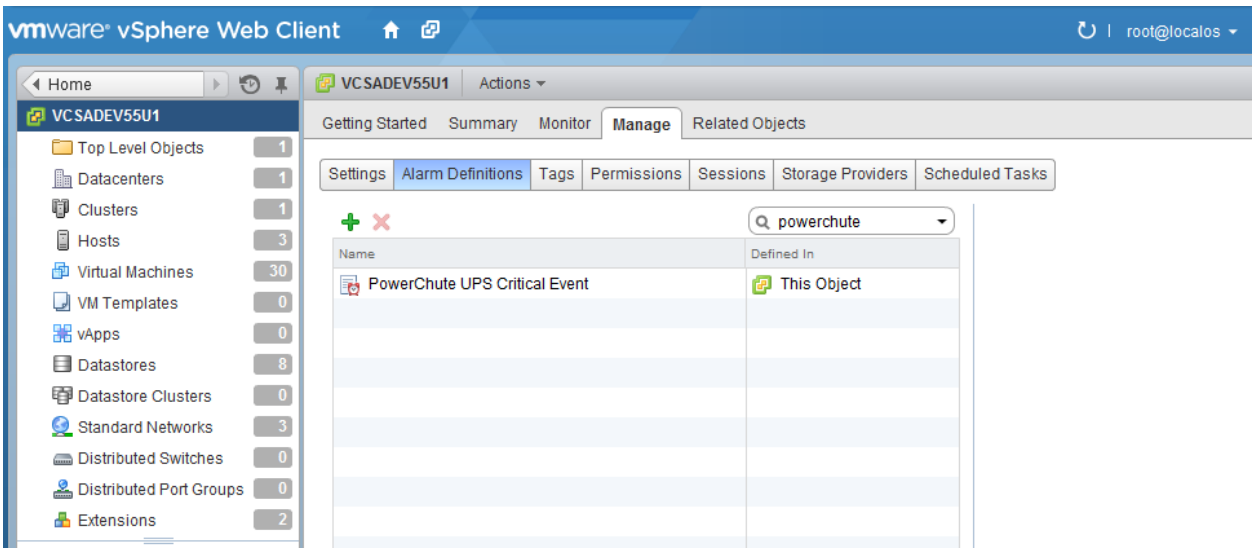

PowerChute UPS Critical Event

This alarm will be triggered with “Warning” status on the Triggered alarms view for the VMware hosts when a critical UPS event occurs and PowerChute starts the shutdown sequence.

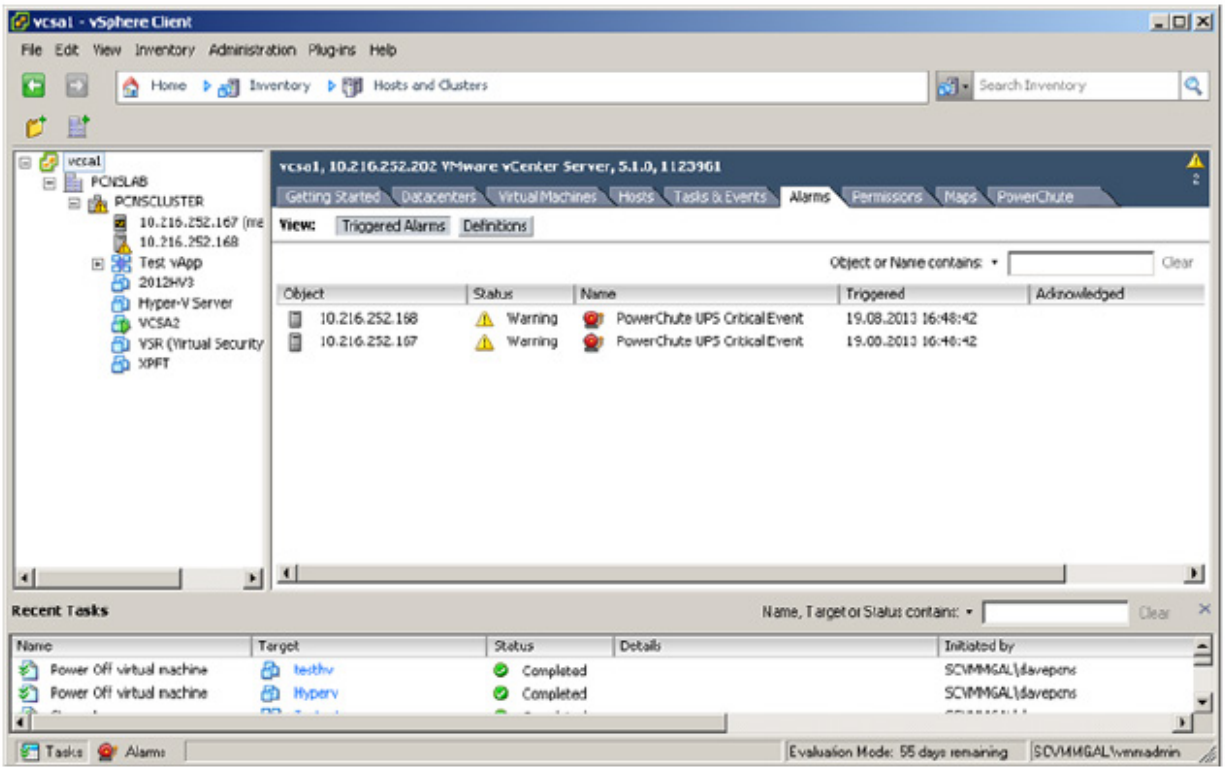

When PowerChute has finished shutting down VMs and vApps the Alarm status will change to “Alert” as the VMware hosts are commanded to shut down.

Removing Alarms

PowerChute vCenter server alarms are removed when:

- vCenter server plugin is uninstalled
- vCenter server plugin is disabled

The alarms can also be removed manually through vCenter Server.

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.

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>:9443** - using default vCenter Server administrator account – **administrator@vsphere.local**.

4. Click on Administration – Single Sign On – Configuration and then on the Identity Sources tab.

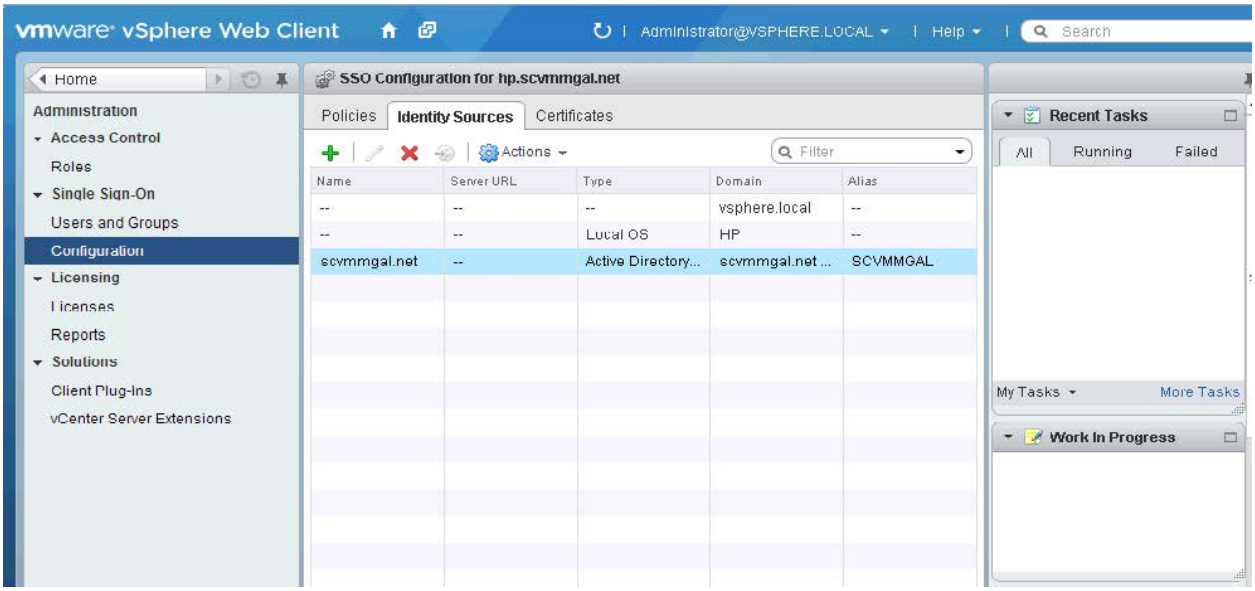

5. Click on the ➕ symbol to add a new identity source.

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

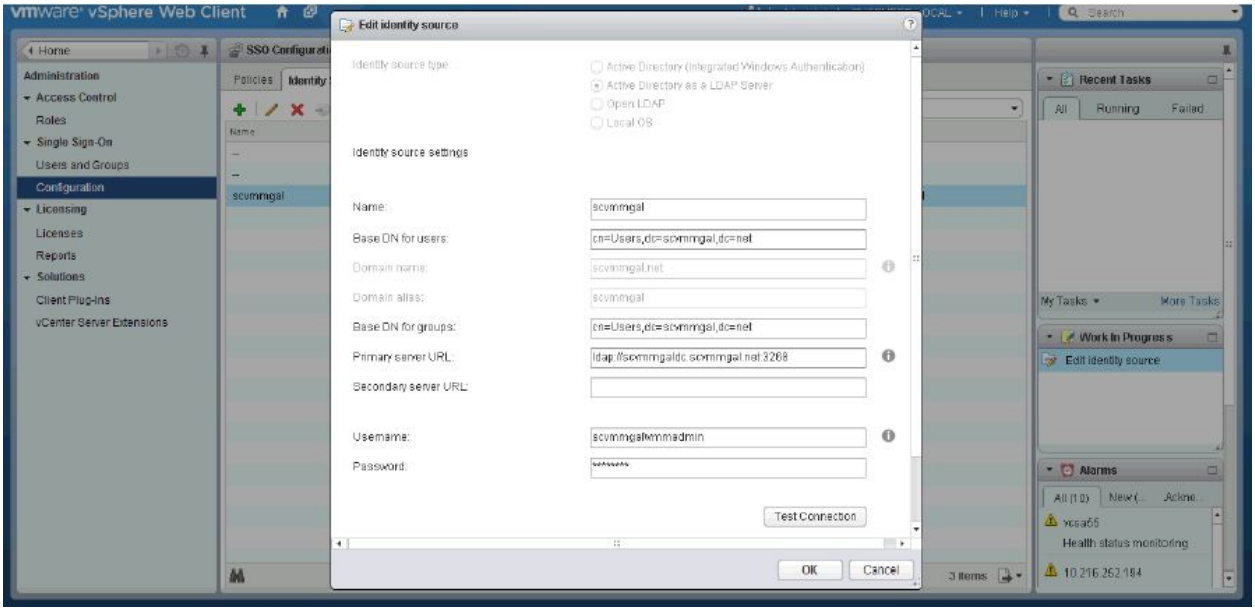

8. Click OK.

9. Click on **Set as default domain**.

10. If using vSphere 5.5 and have logged into vSphere Web Client from a Windows machine that is part of the Active Directory domain select **Active Directory (Integrated Windows Authentication)**.

11. Select **Use Machine Account**.

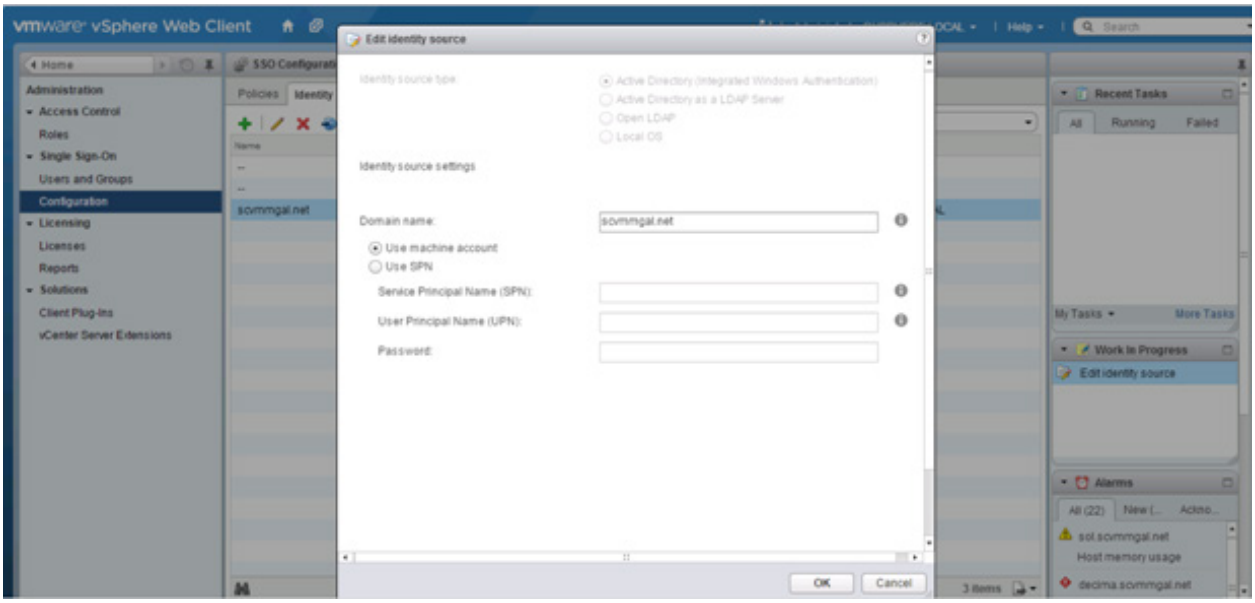

12. Click OK.

13. Log into vCenter using the vSphere client and select the root folder.

14. Next click on the Permissions tab, right click in the right hand pane and select **Add Permission**. Change the Assigned Role to Administrator. Click **Add** under Users and Groups. Select your Active Directory domain from the dropdown list. Select the group **ESX Admins** and click **Add**.

15. Click OK.

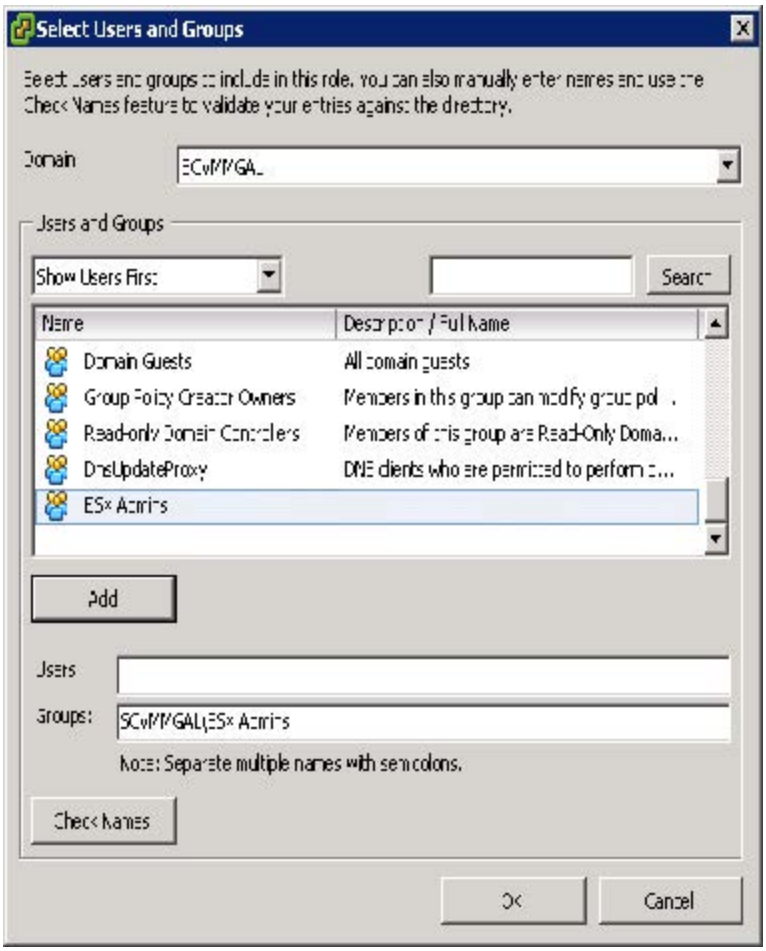

13. Confirm "Propagate to child objects" is selected.

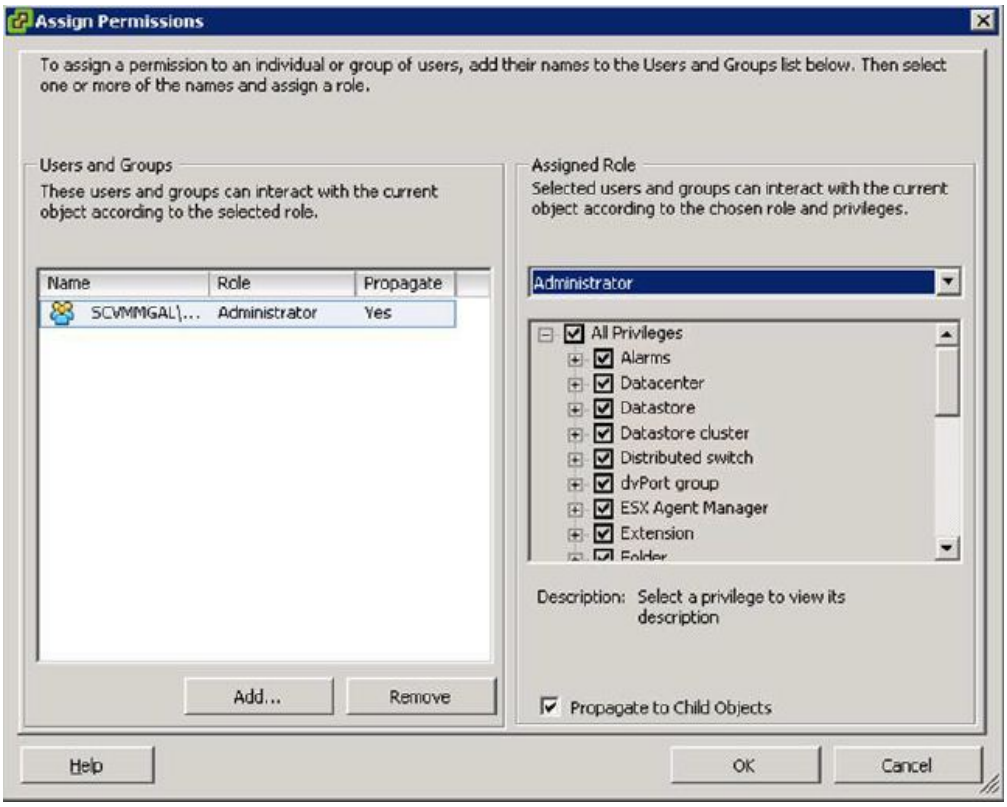

14. Click OK again.
15. Select each host in the Inventory and go to Configuration – Authentication Services under Software. Click on properties and join the host to your Active Directory Domain.
16. When entering the vCenter Server details in the PowerChute Setup Wizard enter a domain user account that is a member of the **ESX Admins** Active Directory User group.
17. More information on VMware Single Sign On Configuration is available from the VMware site at: http://www.vmware.com/support/pubs/
18. Log in to each VMware Host using the vSphere Client.
19. At the root level in the inventory click on the Permissions tab.
20. Verify that the **ESX Admins** group is present and has been assigned the Adminstrator role.

Shared Local Account for vCenter Server and VMware hosts

Create Shared Local Account on vCenter Server

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.
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.

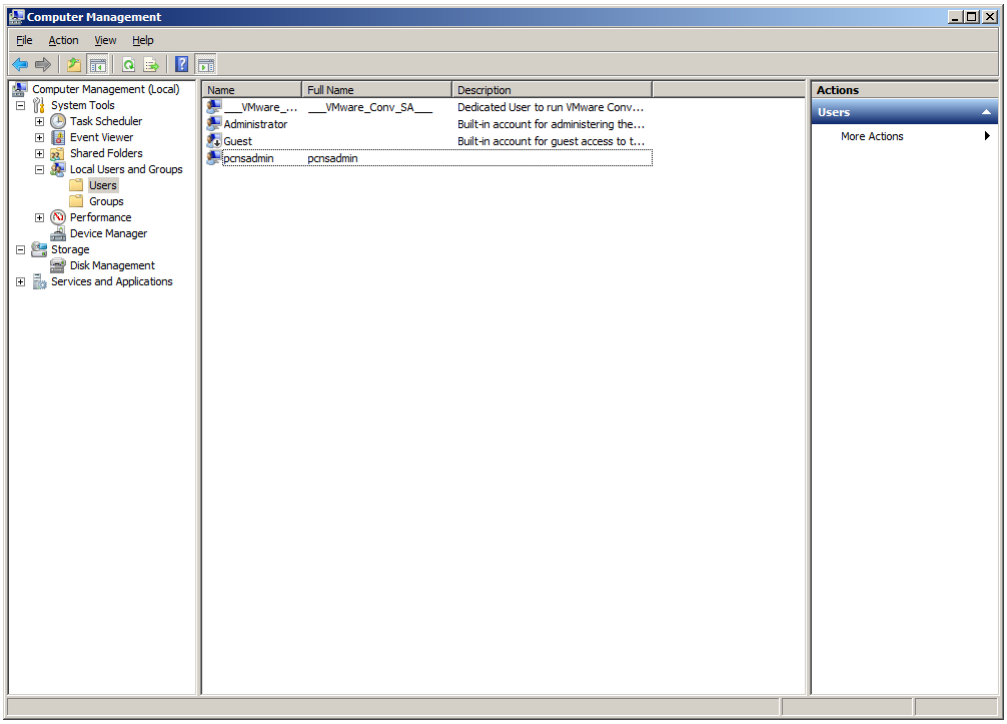

On Linux/vCenter Server Appliance use the terminal commands “useradd” and “passwd”.

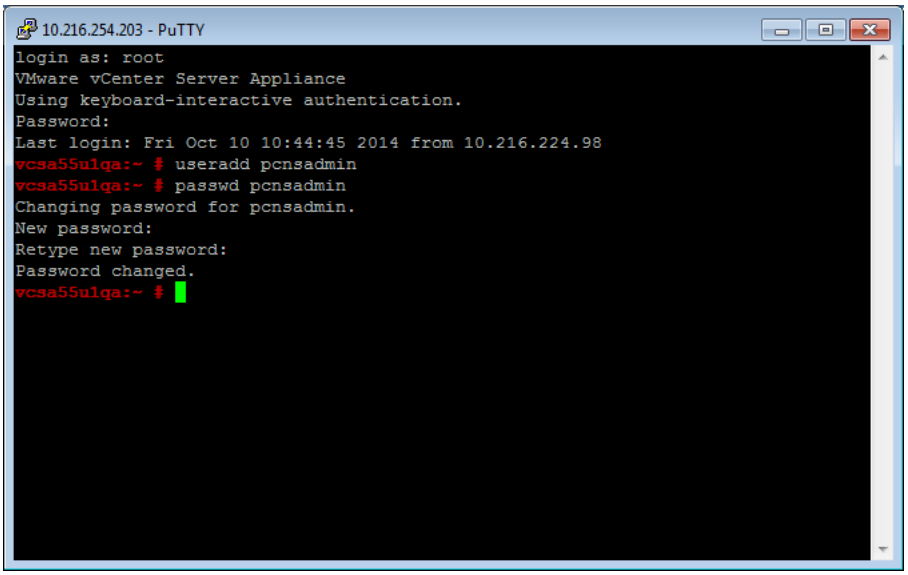

5. Log in to vCenter Server using the vSphere Client and click on the Permissions tab at the root inventory level.
6. Right click and select Add Permission.
7. In the Assign Permissions dialog click Add.
8. Select (server) under domain, select the User that was added in step 4 and click Add.
9. Click OK.
10. Change the Assigned role to Administrator.
11. Select “Propagate to Child Objects” and Click OK.

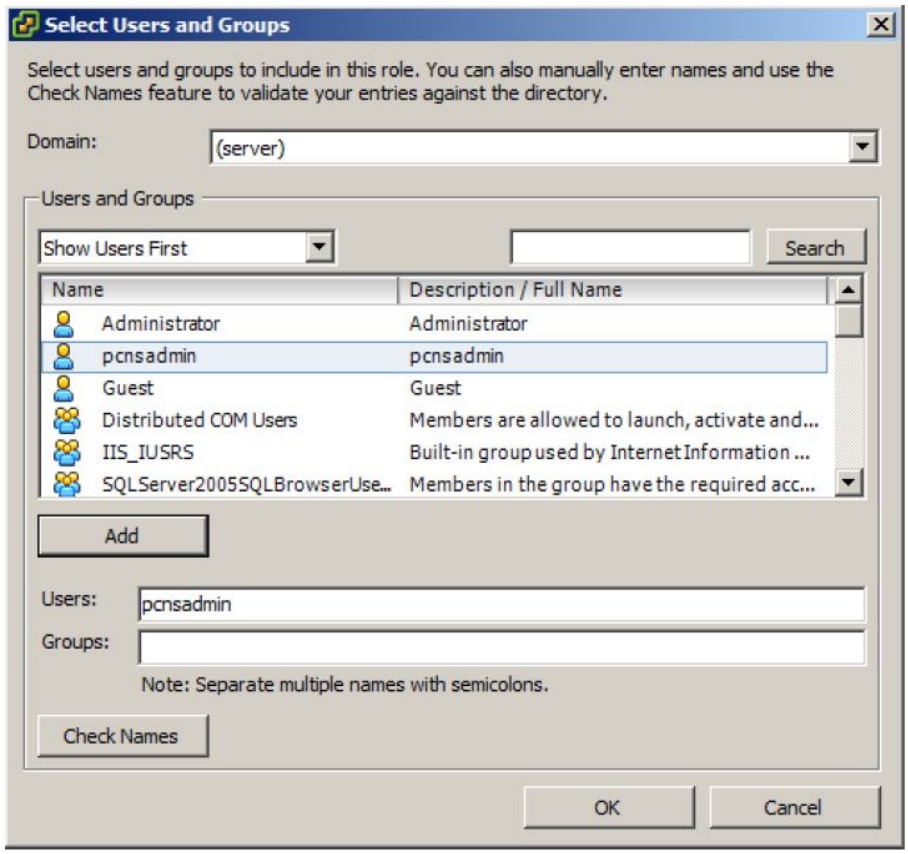
Select Users and Groups
Select users and groups to include in this role. You can also manually enter names and use the Check Names feature to validate your entries against the directory.
Domain:
(server)
Users and Groups
Show Users First
Search
Name
Description / Full Name
Administrator
Administrator
pcnsadmin
pcnsadmin
Guest
Guest
Distributed COM Users
Members are allowed to launch, activate and...
IIS_IUSRS
Built-in group used by Internet Information ...
SQLServer2005SQLBrowserUse...
Members in the group have the required acc...
Add
Users:
pcnsadmin
Groups:
Note: Separate multiple names with semicolons.
Check Names
OK
Cancel

Shared Local Account for vCenter Server and VMware hosts

Create Shared Local Account on each VMware host

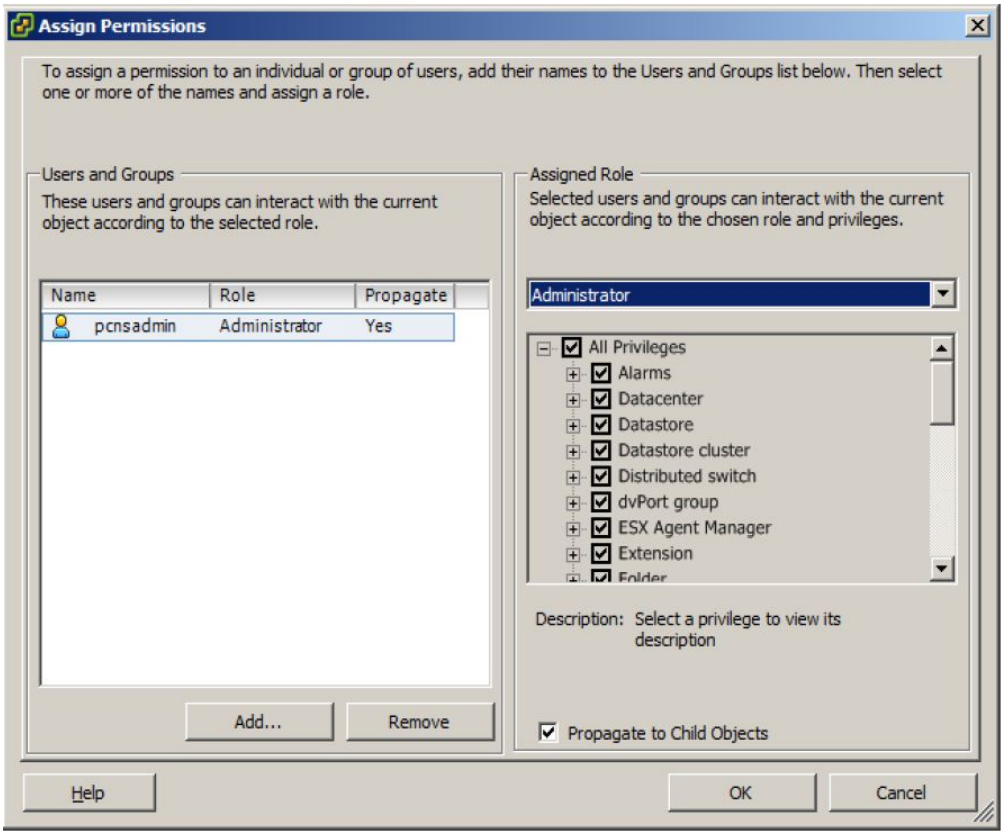

1. Log in to each ESXi host using the vSphere client and click on Local Users and Groups tab.
2. Right click and select Add...
3. Enter the same username and password that was used when adding the local user to vCenter Server.
4. Click on the Permissions tab.
5. Right click and select Add Permission.
6. In the Assign Permissions dialog click Add.
7. Select (server) under domain, select the User that was added in step 3 and click Add.
8. Click OK.
9. Change the Assigned role to Administrator.
10. Select “Propagate to Child Objects” and Click OK.

A shared local account should be used when vCenter Server is running on a VM and Active Directory is unavailable. A shared local account can also be used if the Active Directory Domain Controller is running on a VM and will be shut down.

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's.

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.

Full path to command file: You must specify the full path name of the command file, including the disk drive or volume name. On Windows®, the file should be a .cmd or .bat file. For Linux and Unix systems, it should be a .sh file with execute permissions of **chmod +x [command file name]**.

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:

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td>Number of UPS’s required to power load</td><td>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.

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.

In redundancy terminology, this is the N in N+x.

This setting is not displayed for UPS Setups with a Single UPS device.</td></tr>
<tr><td>Number of additional (redundant) UPS’s</td><td>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:

<table><tr><th>Redundancy level</th><th>No. of critical events that will trigger a shutdown sequence</th></tr><tr><td>N+1</td><td>2</td></tr><tr><td>N+2</td><td>3</td></tr><tr><td>N+3</td><td>4</td></tr></table>
Multiple critical events occurring on the same UPS does not impact the above table values.

In redundancy terminology, this is the x in N+x.

This setting is not displayed for UPS Setups with a Single UPS device.</td></tr>
<tr><td>Total number of UPS’s in Setup</td><td>This is the total of the above two rows and is calculated automatically.</td></tr>
<tr><td>Run Command</td><td>When a shutdown sequence is triggered you can configure PowerChute to execute a command file.

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.</td></tr>
</table>

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td></td><td>See Shutdown Command Files.</td></tr>
<tr><td>Shut down PowerChute Server</td><td>This is enabled by default and is used to gracefully shut down the physical machine running PowerChute.

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.

This option is not available if PowerChute is installed on vMA or deployed as a virtual appliance.</td></tr>
<tr><td>Shut down if Redundancy lost</td><td>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.

This option is associated with the number of UPS critical events required to trigger shutdown:
<table><tr><th>Redundancy level</th><th>No. of critical events that will trigger a shutdown sequence</th></tr><tr><td>N+1</td><td>1</td></tr><tr><td>N+2</td><td>2</td></tr><tr><td>N+3</td><td>3</td></tr></table>Multiple critical events occurring on the same UPS does not impact the above table values.

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.</td></tr>
<tr><td>UPS Shutdown</td><td>Use this option to set the required UPS behavior after connected equipment or servers are gracefully shut down. For more information see UPS Shutdown.</td></tr>
</table>

Field	Description
Execute Virtualization Shutdown Sequence	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. 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.

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. SNMP support is enabled during the Installation process; see the *PowerChute Network Shutdown Installation Guide* available on the APC website for more information. Once installed, 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.

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.
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:

   

   *Due to US Export restrictions, the Java JRE used by PowerChute may require an Unlimited Strength Jurisdiction Policy to be installed before long encryption keys (such as AES-192 or AES-256) can be successfully used. See http://www.oracle.com/technetwork/java/javase/downloads/index.html for details of the latest Java Cryptography Extension (JCE). For more information, see Knowledge base article FA290630 (Enter "FA290630" at http://www.schneider-electric.us/sites/us/en/support/faq/faq_main.page). For operating systems that require a custom JRE (such as HP-UX or AIX), see the manufacturer's website for JRE guidelines.

   † **Note:** Certain SNMP Network Management Systems use a non-standard AES key extension algorithm for 192 and 256 bit key lengths. This non-standard implementation or “Key extension algorithm” is specified by the IETF.
   If your NMS requires the use of the Key Extension algorithm, select Privacy Protocol options **AES-192 Ex** or **AES-256 Ex**.

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.

Select **Send Trap when condition is cleared** to be notified when the Event is cleared.
Note: If the PowerChute server is shutdown 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 `pcnsNmcTable`. 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.

* **Note**: These OIDs are not available for Advanced UPS Setups. See below for equivalent OIDs for Advanced UPS Setups.

Object Identifier Name	Access	Description
pcnsRuntimeRemainingThreshold	read-write	This defines a threshold for runtime remaining. When the UPS in 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.

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

Object Identifier Name	Access	Description
		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

Object Identifier Name	PowerChute UI Reference
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
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
pcnsEnviornmentCommunicationEstablished	Communication Established with EMC
pcnsEnviornmentCommunicationLost	Communication Lost with EMC
pcnsTempInRangeProbeX	Temperature Probe X In Range

Object Identifier Name	PowerChute UI Reference
pcnsTempOutOfRangeProbeX	Temperature Probe X Out Of Range
pcnsHumidityInRangeProbeX	Humidity Probe X In 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 pcnsRunTimeBelowThreshold. 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
pcnsRunTimeBelowThresholdCommandFileThreshold	read-write	If Runtime Remaining falls below this threshold, the command file is executed.
pcnsRunTimeBelowThresholdShutdownThreshold	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.
pcnsNMCComsLost	Severe	The NMC cannot communicate with the UPS.
pcnsNMCComsLostResolved	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 Windows, the message will be sent to all users who are on the same network. 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.

For Windows operating systems, PowerChute can only send notifications if the operating system supports the messenger service. If not supported, there is no option displayed in the UI. See KBase article FA169440 for more information. (If you have difficulty with this link, enter "FA169440" at http://www.schneider-electric.us/sites/us/en/support/faq/faq_main.page).

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. On Windows, the file should be a .cmd or .bat file. For Linux and Unix systems, it should be a .sh file.

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.

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.

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 minute 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.

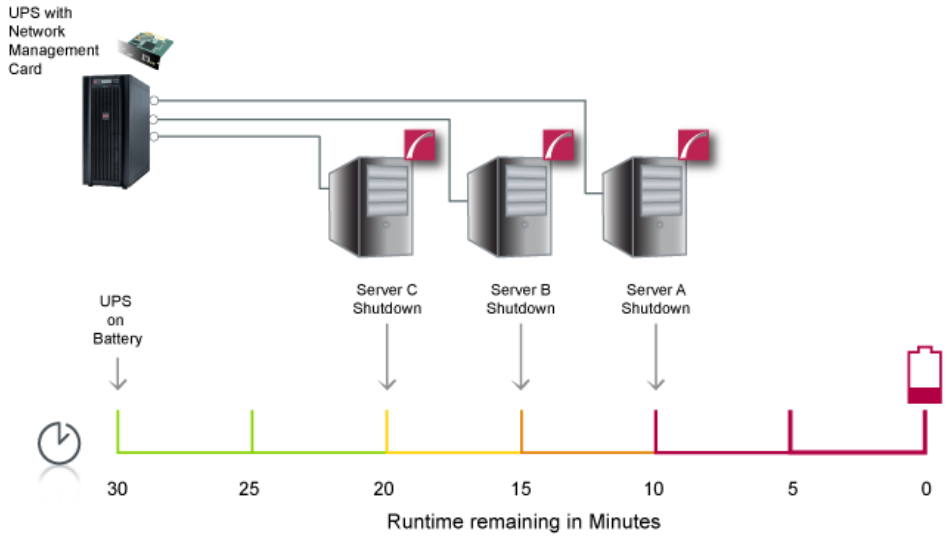

Sample Shutdown Scenarios

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. VM shutdown and vApp shutdown have been enabled with 120 second delay configured for each.

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

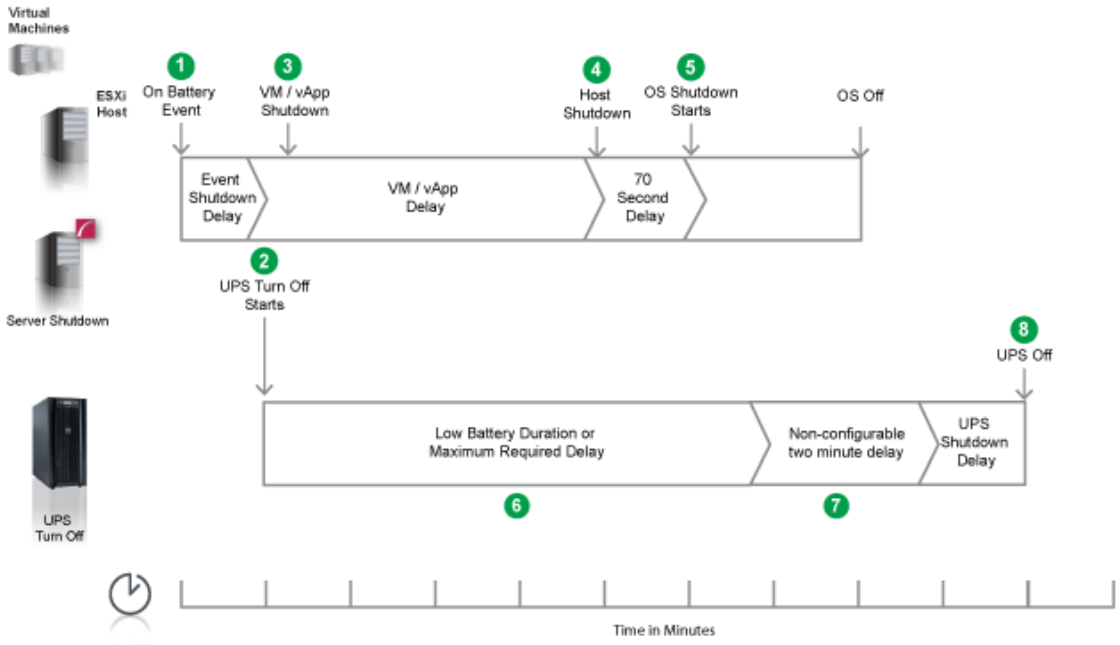

1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM shutdown delay = 120, vApp shutdown delay = 120), 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. VM shutdown and vApp shutdown have been enabled with 120 second delay configured for each.

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

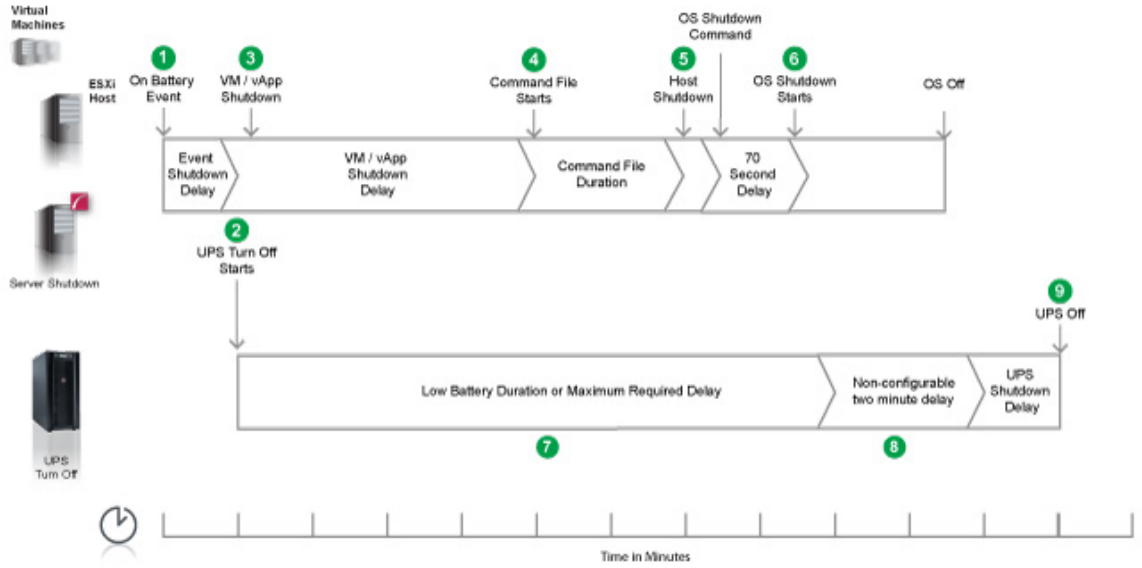

1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the UPS. UPS turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM shutdown delay = 120, vApp shutdown delay = 120), 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. VM shutdown and vApp shutdown have been enabled with 120 second delay configured for each.

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

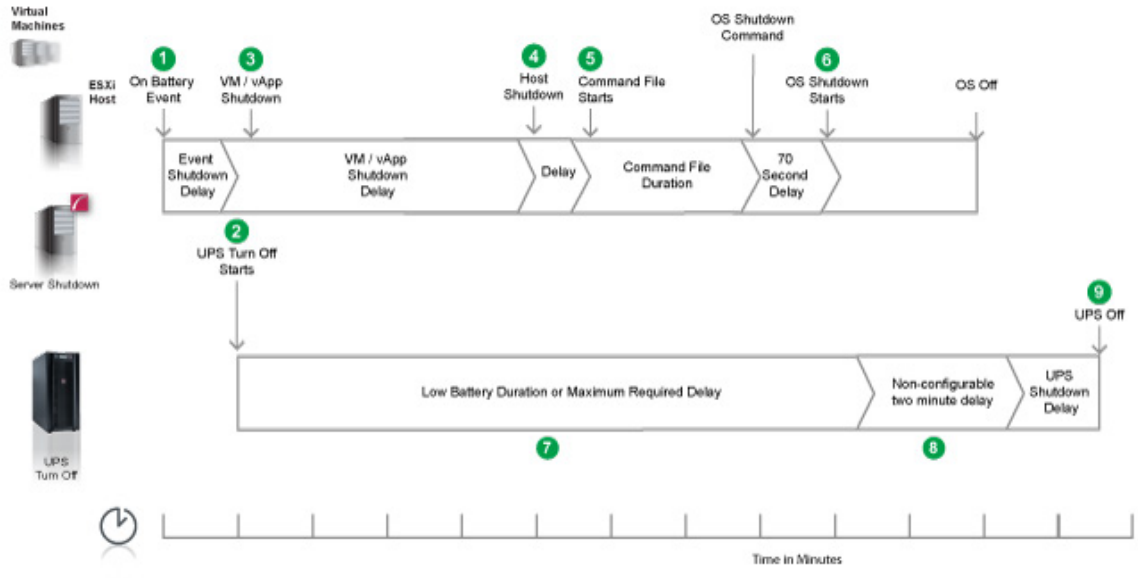

1. PowerChute reports that the UPS is on battery.

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 shutdown delay = 120, vApp shutdown delay = 120), 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.

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. VM shutdown and vApp shutdown have been enabled with 120 second delay configured for each.

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

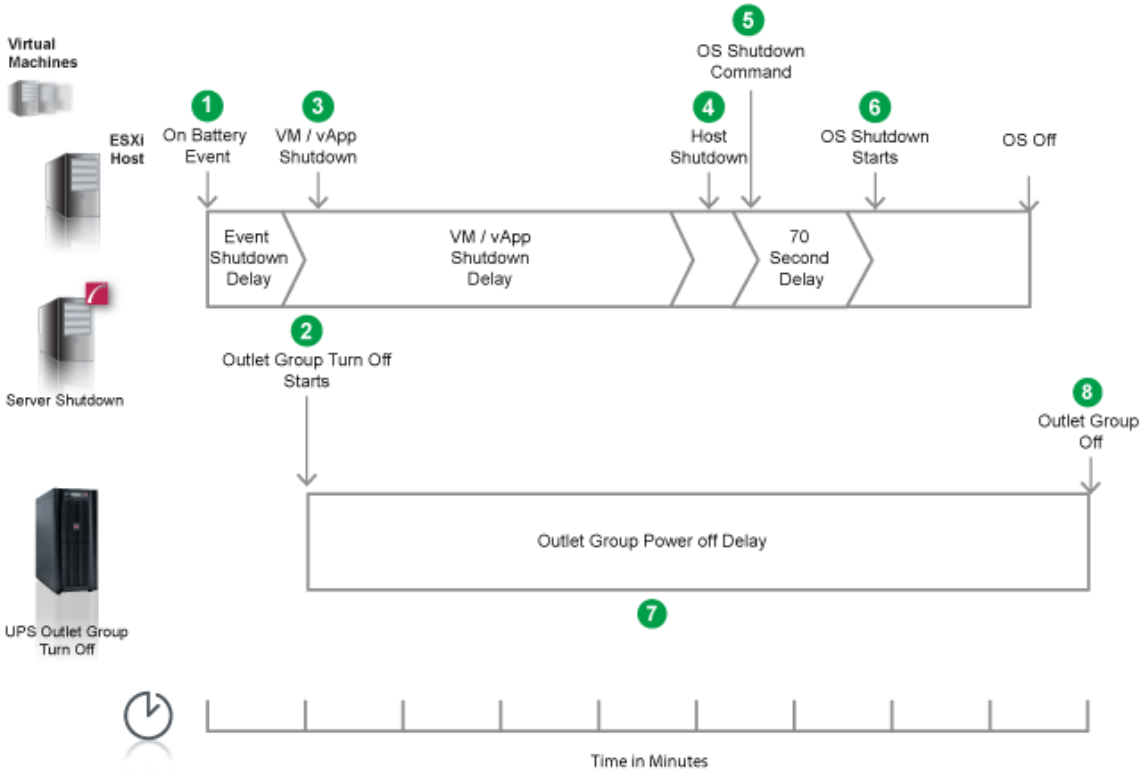

1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute issues a command to turn off the UPS outlet group and the outlet group turn off starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM shutdown delay = 120, vApp shutdown delay = 120), 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. VM shutdown and vApp shutdown have been enabled with 120 second delay configured for each.

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

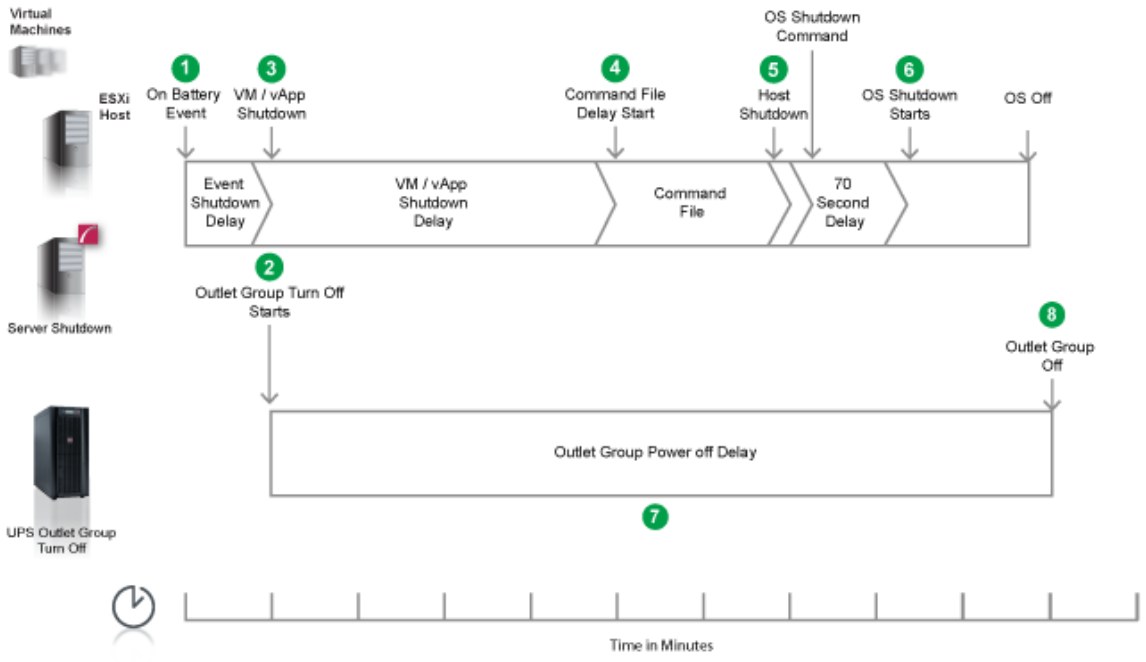

1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the outlet group and the outlet group turnoff starts.
3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM shutdown delay = 120, vApp shutdown delay = 120), 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. VM shutdown and vApp shutdown have been enabled with 120 second delay configured for each.

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

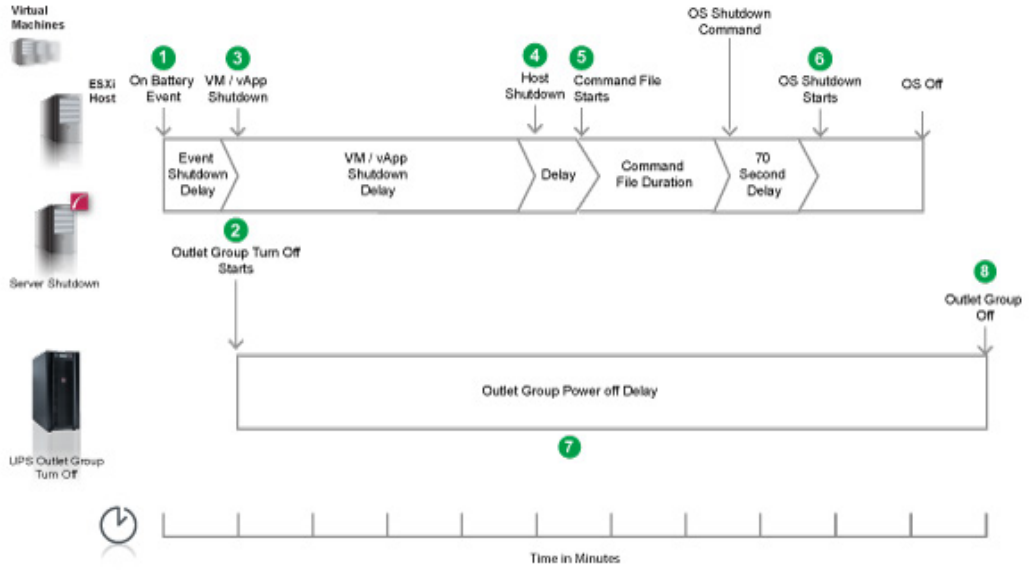

1. PowerChute reports that the UPS is on battery.
2. After the shutdown delay configured for the On Battery event has elapsed, PowerChute sends a command to turn off the outlet group and the outlet group turnoff starts.

3. PowerChute starts a Maintenance mode task on the VMware Hosts and starts to shut down VMs and vApps.
4. After 4 minutes (VM shutdown delay = 120, vApp shutdown delay = 120), 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.

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
- Execute Command File after Host Shutdown delay
- Shutdown Command File 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 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.

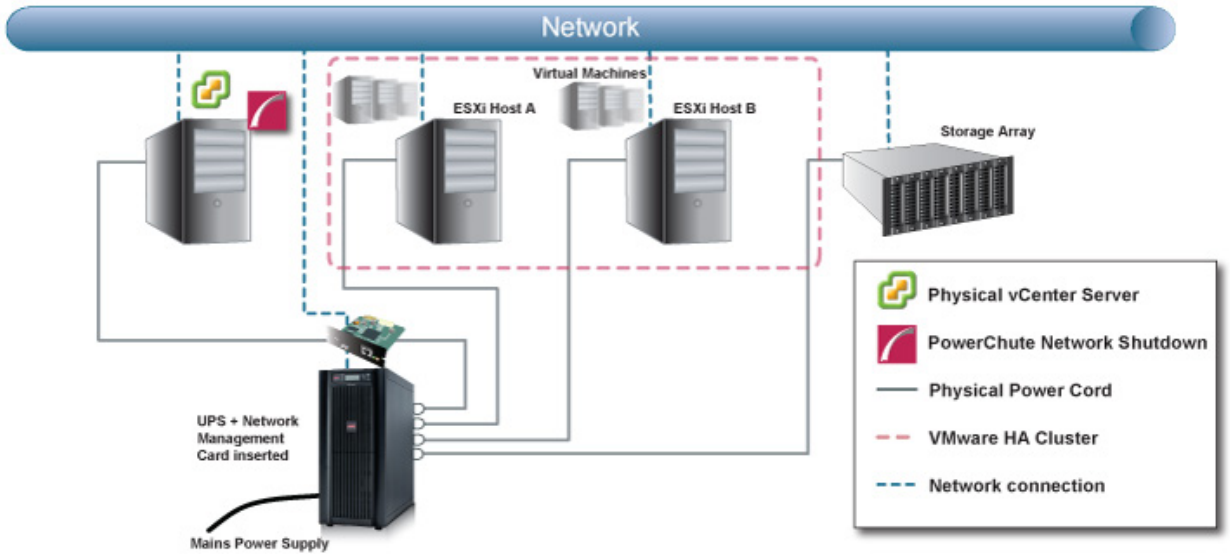

The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

1. The UPS has been running on Battery power for x number of seconds.
2. PowerChute and 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.

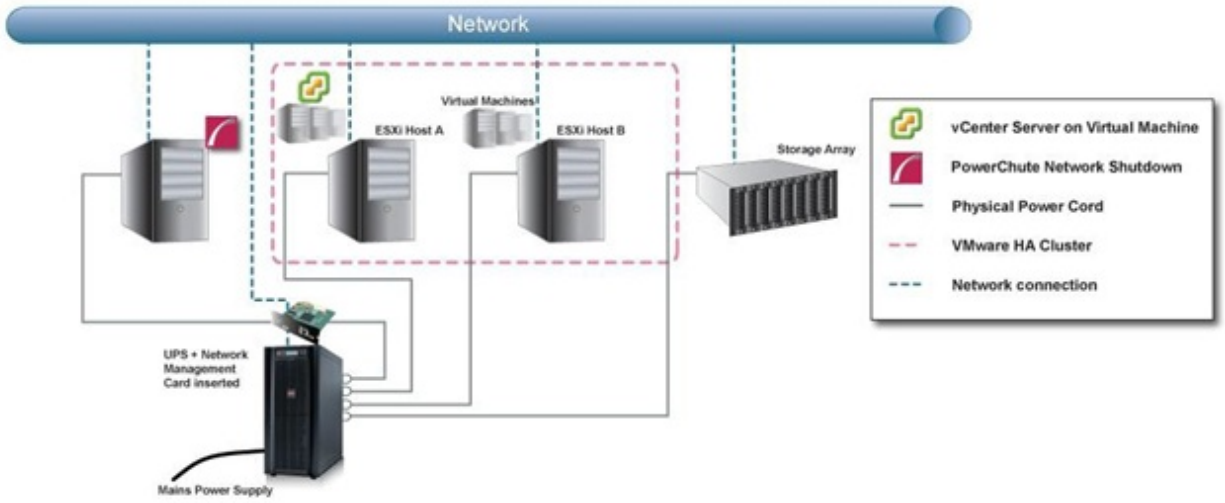

When a critical UPS event, such as UPS on Battery occurs the following sequence is triggered:

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 shutdown followed by 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

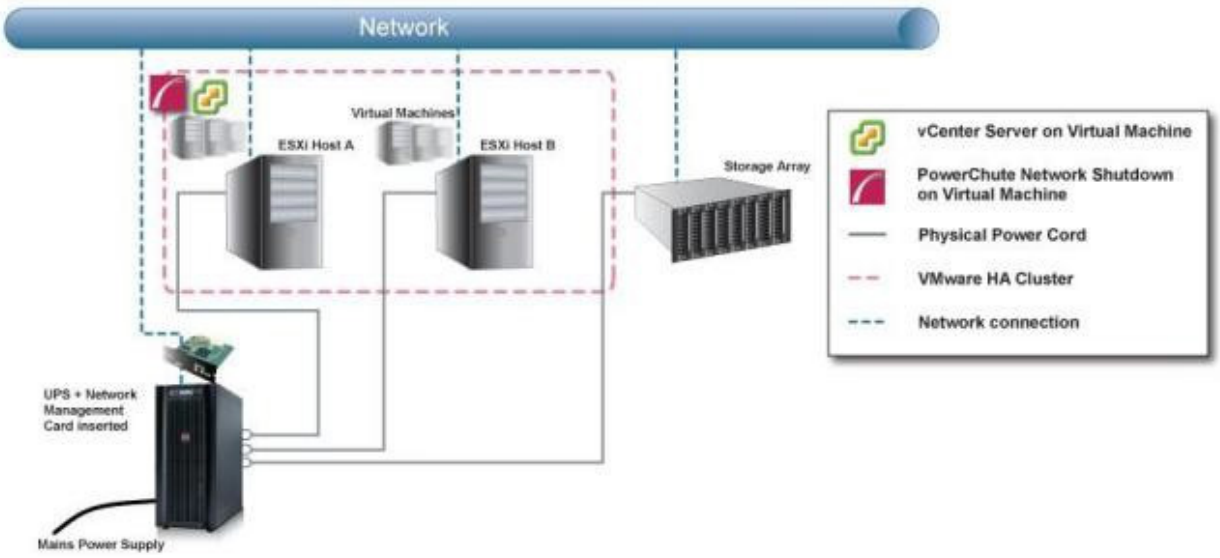

When a critical UPS event, such as UPS on Battery occurs the following sequence is triggered:

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 shutdown followed by 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 - Advanced UPS Configuration

Here, separate UPS devices are powering two VMware hosts: the vCenter Server and a storage array. PowerChute is installed on the vCenter Server machine and is monitoring all UPS's. A shutdown command file has been configured for UPS Setup #1 containing UPS #1.

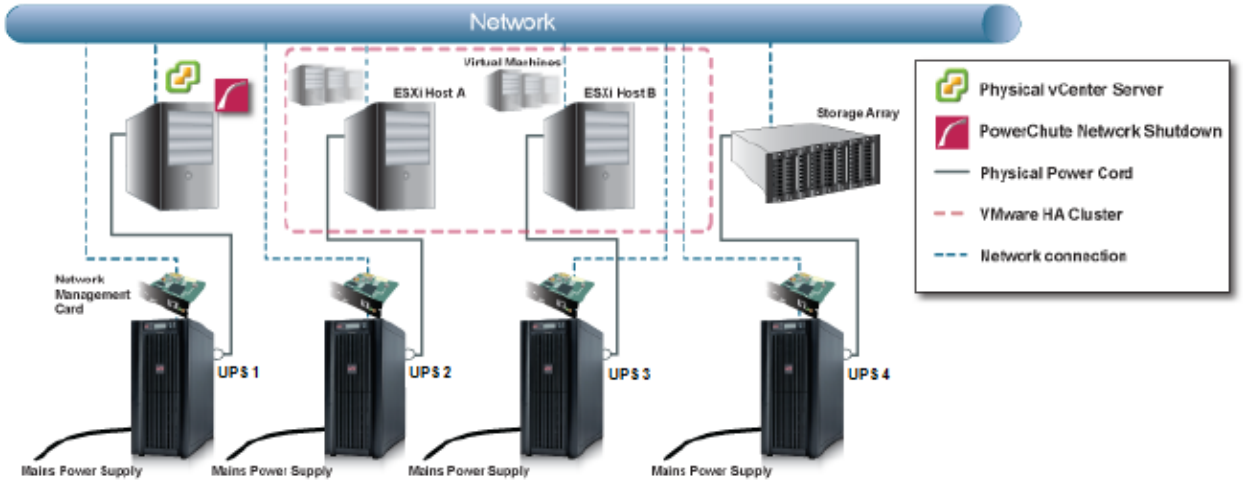

Critical event on UPS #4: Option to shut down virtual hosts is enabled for this UPS. VM Migration is not enabled.

1. UPS #4 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 two VMware hosts in the cluster.
4. PowerChute starts a maintenance mode task on the two VMware hosts, and shuts down the VMs on the two VMware hosts.
5. PowerChute shuts down any vApp that is powered on.
6. PowerChute shuts down the two VMware hosts.
7. As the physical machine is not affected PowerChute continues to run.

Critical event on UPS #1: Option to shut down virtual hosts and the physical machine are enabled for this UPS. VM migration is not enabled.

1. UPS #1 goes on battery.
2. PowerChute issues command to gracefully turn off UPS #1, if configured.
3. A UPS critical event is triggered for the two VMware hosts.

4. PowerChute starts a maintenance mode task on the two VMware hosts, and shuts down the VMs on the two VMware hosts.
5. PowerChute shuts down any vApp that is powered on.
6. The shutdown command file is executed.
7. After the shutdown command file duration has elapsed, if all VMs have shut down, the maintenance mode task completes and PowerChute shuts down the two VMware hosts.
8. PowerChute shuts down the vCenter Server machine.

For detailed information, please view “Using PowerChute Network Shutdown in a VMware HA Cluster” Application Note here.

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 are prioritized into High, Medium, Low priority groups.

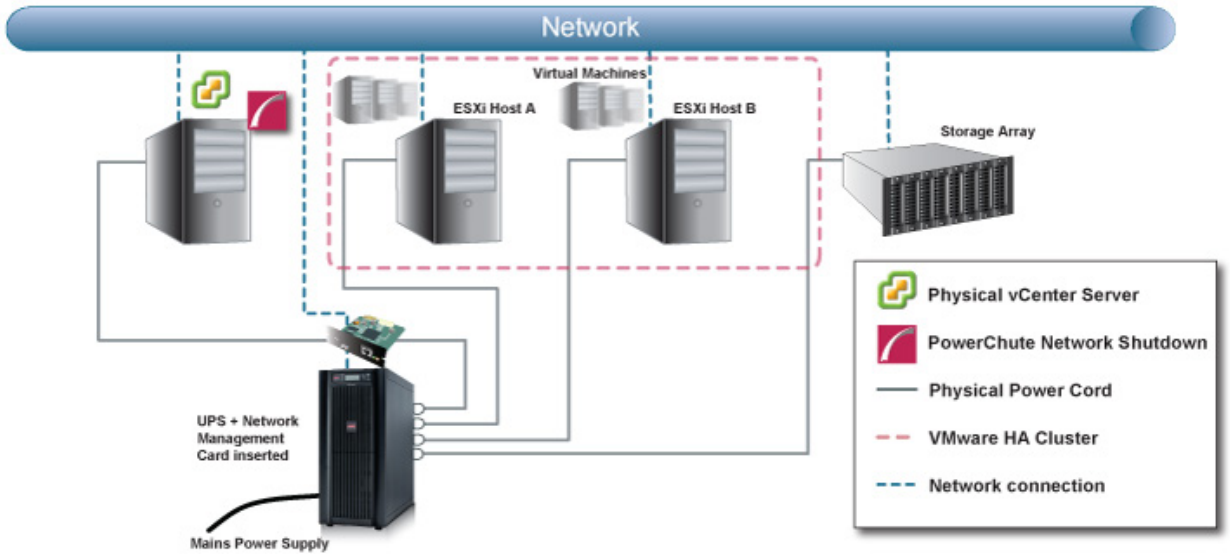

The following shutdown sequence occurs when the shutdown action is enabled for the On Battery event.

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 on VMware hosts A and B in the order in which they are prioritized:

   Un-prioritized > Low > Medium > High

   First, the un-prioritized VMs are shut down sequentially. As the duration for un-prioritized VMs elapses, the Low priority VMs are shut down sequentially, followed by Medium priority VMs and finally the High priority VMs are shut down sequentially.
   PowerChute sends a shutdown command to all VMs in each priority group at the same time. The VMs within each priority group are not shut down in a particular order.
3. As VM Shutdown duration elapses, PowerChute shuts down the vApp, if configured.
4. PowerChute runs the shutdown command file and issues a command to turn off the UPS, if configured.
5. After the shutdown command file duration has elapsed, PowerChute shuts down the VMware hosts.
6. PowerChute shuts down the vCenter Server Machine.

VMware Shutdown - VM Prioritization - Advanced UPS Configuration

In this example, separate UPS devices are powering two VMware hosts: the vCenter Server and a storage array. PowerChute is installed on the vCenter Server machine and is monitoring all UPS's. VM Prioritization is enabled and VMs are prioritized into High, Medium, Low priority groups.

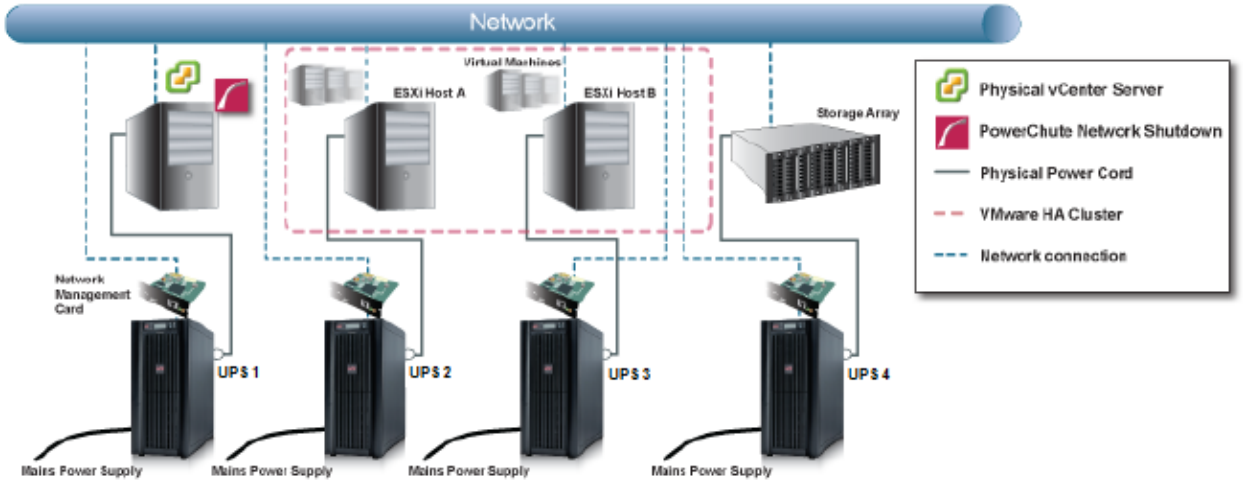

Critical event on UPS #2: Option to shut down virtual hosts is enabled for this UPS. VM Migration and VM Prioritization are enabled.

1. UPS #2 goes on battery.
2. A UPS critical event is triggered for Host A and PowerChute starts a maintenance mode task on Host A.
3. PowerChute migrates the VMs to healthy Host B, in order of prioritization:

   ★ High > ★ Medium > ★ Low > Un-prioritized

4. As sequenced VM Migration is enabled for priority groups, VM migration begins. 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 priority group and finally the un-prioritized VMs migrate.
5. As VM migration duration elapses, any VMs that have not been migrated will be shutdown in the VM Shutdown sequence, in the order in which they are prioritized:

   Un-prioritized > ★ Low > ★ Medium > ★ High

   First, the un-prioritized VMs are shut down sequentially. As the duration for un-prioritized VMs elapses, the Low priority VMs are shut down sequentially, followed by Medium priority VMs and finally the High priority VMs are shut down sequentially.
   PowerChute sends a shutdown command to all VMs in each priority group at the same time. The VMs within each priority group are not shut down in a particular order.
6. As VM Shutdown duration elapses, PowerChute shuts down the vApp, if configured.
7. After the shutdown command file duration has elapsed, PowerChute shuts down Host A.

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 (Enter "FA290624" at http://www.schneider-electric.us/sites/us/en/support/faq/faq_main.page).
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
  - Shutdown command file 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 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.

  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.

- 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 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 in 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 Communication 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 failed because PowerChute received an untrusted SSL certificate from the NMC https://[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.

  To accept the certificate, enable the option "Accept Untrusted SSL Certificates?" on the UPS Details page of the PowerChute Setup Wizard or add the certificate to the PowerChute-keystore.

- 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.

  If the option **Accept Untrusted SSL certificates** is enabled, PowerChute will automatically add self-signed and untrusted certs to its local keystore.

- 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.

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 they 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.

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 available host.

  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.

- VMs on Host [Host1] successfully migrated to Host [Host2]

  A critical event has occurred on a UPS associated with [Host1] and PowerChute has successfully migrated virtual machines to [Host 2]

- Start maintenance mode:[Host].

  PowerChute starts a maintenance mode task on the host when a critical UPS event occurs.

- 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.

- Exit maintenance mode: [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 available host.

  A critical event has occurred on a UPS associated with this host and PowerChute is trying to migrate virtual machines to another host.

- 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.

- Shutting down vApp [vApp].

  PowerChute is shutting down the specified vApp.

- Starting vApp shutdown process.

  PowerChute is starting to shut down any vApps.

- No vApp to shutdown 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].

  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 linked 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 physical 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 failed for Host [Host].

  A critical event has occurred on a UPS associated with [Host] and PowerChute was unable to shutdown [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] 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 VM 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 shutdown is triggered.

  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 in 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 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.

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** 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.

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 5kVA UPS devices supporting an 8kVA 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 5kVA UPS devices supporting a 13kVA 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.

General

This section contains information on the topics below:

- Communications Settings
- PowerChute Agents
- PowerChute Configuration (INI) File
- User Interface Session Timeout
- Check for Updates
- Customer Support

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 Application Note “PowerChute Network Shutdown Security Features Overview” 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.

In the NMC, the default administrator username and password are both **apc**. The default authentication phrase is **admin user phrase**. We recommend that you change the defaults for security reasons.

- 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 length must be 3 to 32 characters and special characters are allowed.

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:

1. Stop the PowerChute service. For more information, see Knowledge Base article FA290624 (Enter "FA290624" at http://www.schneider-electric.us/sites/us/en/support/faq/faq_main.page).
2. Replace the existing copy of pcnsconfig.ini in the group1 folder.
3. 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.

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 (Enter "FA290624" at http://www.schneider-electric.us/sites/us/en/support/faq/faq_main.page).
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.

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 and select your country from the drop-down box list. This lists the contact details for support services you may require.

Troubleshooting

This section contains information on the topics below:

- Network Management Card Troubleshooting
- VMware Troubleshooting
- Browser Troubleshooting
- SNMP 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:

“Failed to migrate all VMs from Host [hostname]”

“Failed 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 polices 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.

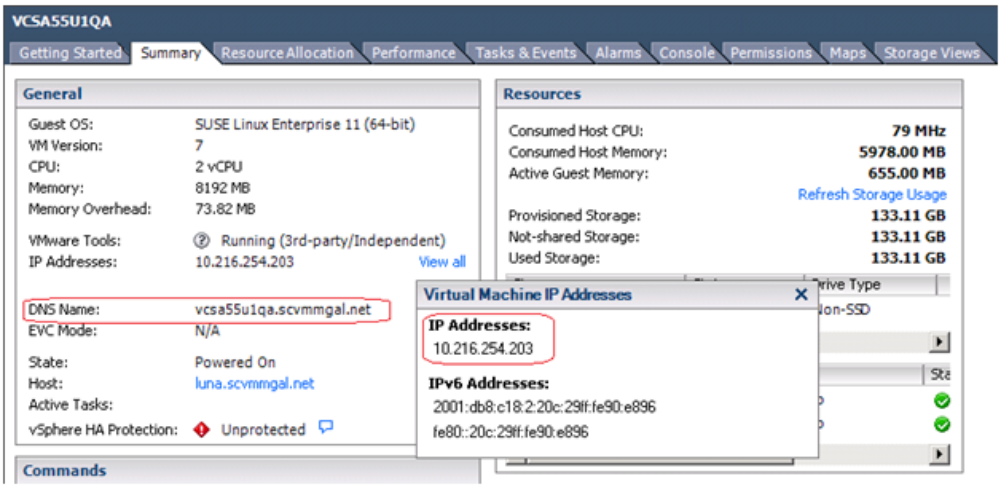

The vCenter Server VM IP address is being used by another system or there is an old VM with the same static IP address in the vCenter Server inventory – even if that VM is powered off.

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 Host shutdown to fail.

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 fails to shut down a Host.

The following message appears in the EventLog:

“Shutdown Host failed 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]

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
```

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.

Known Issues with Internet Explorer

- IE Enhanced Security is enabled by default for most Windows operating systems and this can block JavaScript for the PowerChute Application - to avoid this add the PowerChute URL to the Trusted Sites List under Tools - Internet Options - Security.

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:

Due to US Export restrictions, the Java JRE used by PowerChute may require an Unlimited Strength Jurisdiction Policy to be installed before long encryption keys (such as AES-192 or AES-256) can be successfully used. For most operating systems, see http://www.oracle.com/technetwork/java/javase/downloads/index.html for details of the latest Java Cryptography Extension (JCE). For more information, see Knowledge base article FA290630 (Enter "FA290630" at http://www.schneider-electric.us/sites/us/en/support/faq/faq_main.page). For operating systems that require a custom JRE (such as HP-UX or AIX), see the manufacturer's website for JRE guidelines.

Note: Certain SNMP Network Management Systems use a non-standard AES key extension algorithm for 192 and 256 bit key lengths. This non-standard implementation or “Key extension algorithm” is specified by the IETF.
If your NMS requires the use of the Key Extension algorithm, select Privacy Protocol options **AES-192 Ex** or **AES-256 Ex**.

- 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.

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.

© 2016 APC by Schneider Electric. APC, the APC logo, and PowerChute Network Shutdown are owned by Schneider Electric Industries S.A.S., or their affiliated companies. All other trademarks are property of their respective owners.

990-4595E-001 7/2016