

Best Practices for Power Management Configuration

best practices



Abstract.....	2
Introduction.....	2
Rack and Power Manager	2
Control and status.....	5
Summary	10
Call to action	11

Abstract

This document details the best practices for configuring the Rack and Power Manager software in a server environment to enhance reliable system operation. In particular, it addresses best practices for managing the system Uninterruptible Power Supply (UPS).

Introduction

Today's server systems commonly include an uninterruptible power supply (UPS) power backup system to guard against the possibility of data loss or system damage in the event of a power failure. HP systems allow actively managing the UPS by setting operating parameter limits, monitoring internal inputs, and receiving indication of issues that need attention through the Rack and Power Manager software¹. This software works together with the Console Management Controller (CMC) hardware to permit an administrator to program system behavior in response to specific events to achieve design goals.

The focus of this paper is the application of best practices to the management of UPS resources. In this context, the term *best practice* refers to an accepted method of performing a task, or finding the best way to do something for a particular customer situation. As it relates to the Rack and Power Manager software, best practice encompasses a range of topics.

At one level, best practice means taking the time to plan and design a deployment to minimize problems before they occur. This would include such things as facilities and power planning, as well as system installation and configuration. At another level, it includes functional tasks such as performance and system monitoring, system management, and backup. At the highest level, best practice means a set of consistent operating procedures that, when strictly followed, help to reduce errors by standardizing processes for upgrades, planned maintenance, and disaster recovery.

Rack and Power Manager

One of the most valuable aspects of the Rack and Power Manager is its ability to work in concert with a system UPS to gracefully shut down computer operating systems (OS) during a power failure. It is often desirable to prolong the runtime of critical computers (Web servers, domain controller, databases, and so on) while still allowing protection of essential data by closing down services completely before power is removed. During a power failure, it can be beneficial to shut down early any equipment that is not serving critical needs to prolong battery time for the more necessary systems. Power fail settings allow for prioritizing shutdowns of UPS load segments (banks of power outlets) and of the servers receiving power from the UPS (attached agents, for example). The following items should be considered when configuring the power failure settings:

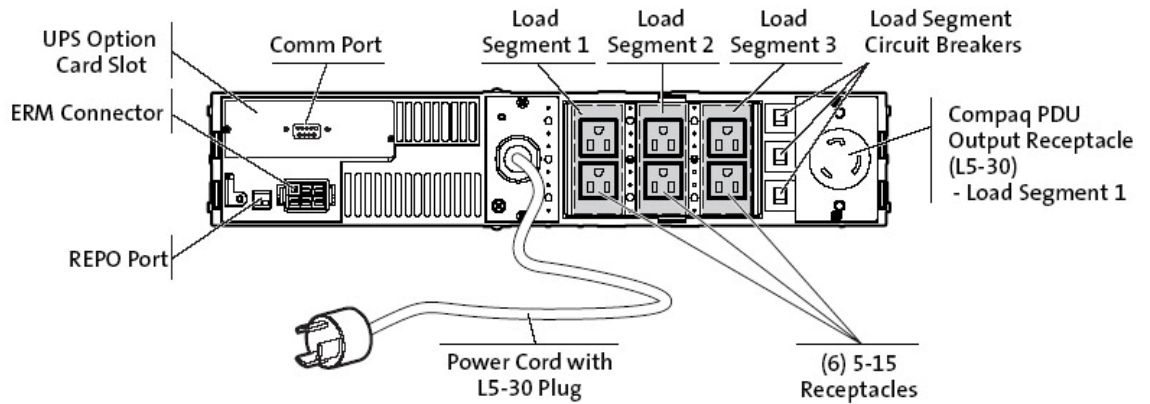
- The length of time for which battery protection is available increases as the load on the UPS decreases.
- After a power outage recovery, another outage could occur before the UPS batteries fully recharge. For more robust system operation, the time value should be set small enough to allow battery reserve for at least two shutdowns.

One basic aspect of UPS management is partitioning the outlets of the UPS into load segments for independent control. Each load segment may consist of several outlets (see Figure 1). As an example, one load segment may supply file and print servers that are non-critical and can be quickly shut down, while another segment supplies data base servers whose operation is considered critical or which take longer to shut down. A file-and-print server can shut down its OS rather quickly (on the order of a minute or two); but even after that, the machine is still on and consuming some power. In

¹ For reference, information on the installation and integration of the software can be found in [The HP Rack and Power Manager User Guide](#), which contains that information and more about configuration.

this case, the best practice is to turn off the corresponding UPS load segment altogether by switching off the outlets from the UPS. This ensures that the power will be shut off even when there is no one on site to go to the machine and switch off the power button.

Figure 1. Load segments on rear panel of an R3000 XR UPS



Because database servers may be running critical applications or simply require longer to shut down (an Exchange server database, for example, is quite large), it is desirable to shut them down only when necessary. Since most power outages last only 5 to 10 minutes, the best practice is to wait as long as possible to verify that the outage will be of sufficient duration to necessitate shutting down these more complicated applications. To be able to wait, though, the UPS must have sufficient battery capacity to keep supplying power to the servers. If the UPS can shed some of the other loads while waiting, that will decrease the overall power consumption and allow more time before requiring the shutdown of the database servers. The strategy, then, is to allocate limited resources where they are most needed and to postpone the shutdown of critical or complex operations as long as possible in hopes that the power outage will be of a short duration and the shutdown will prove unnecessary.

Figure 2. UPS Shutdown Settings screen

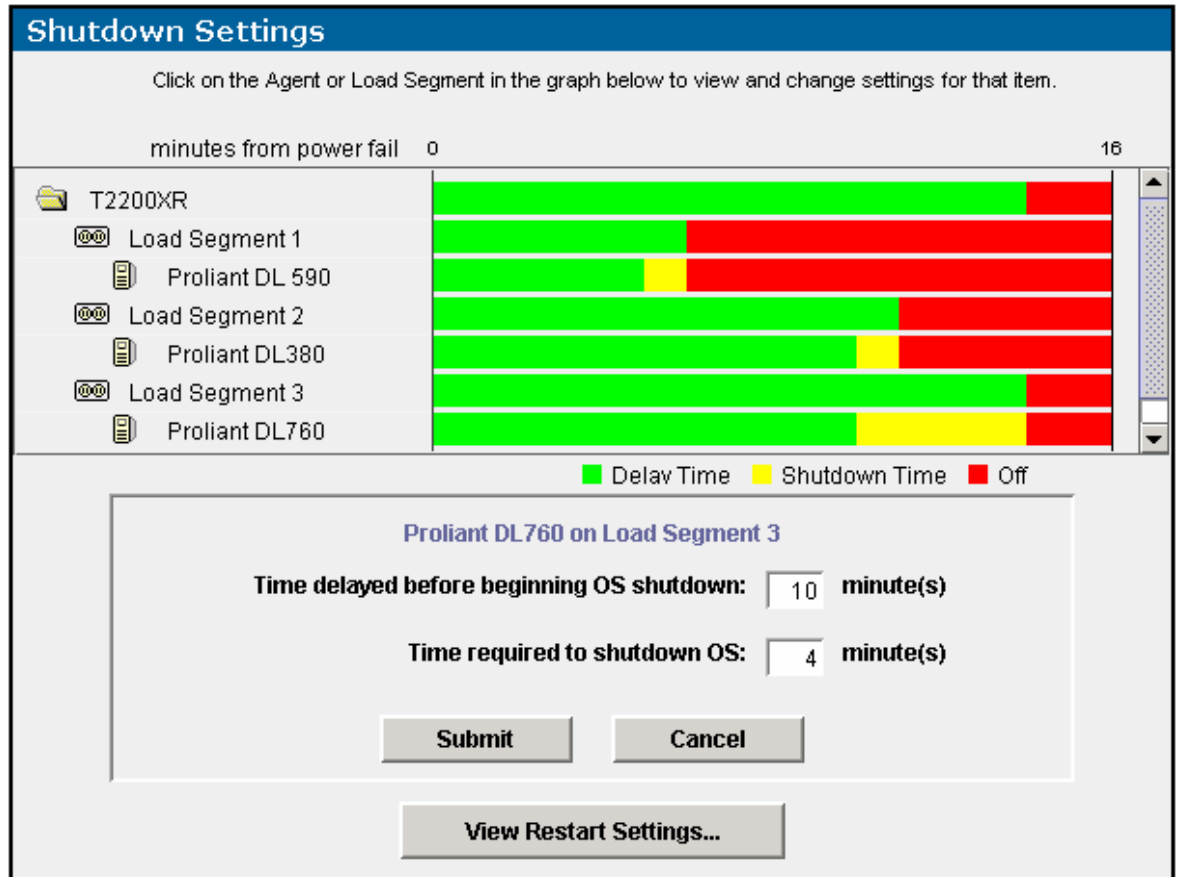


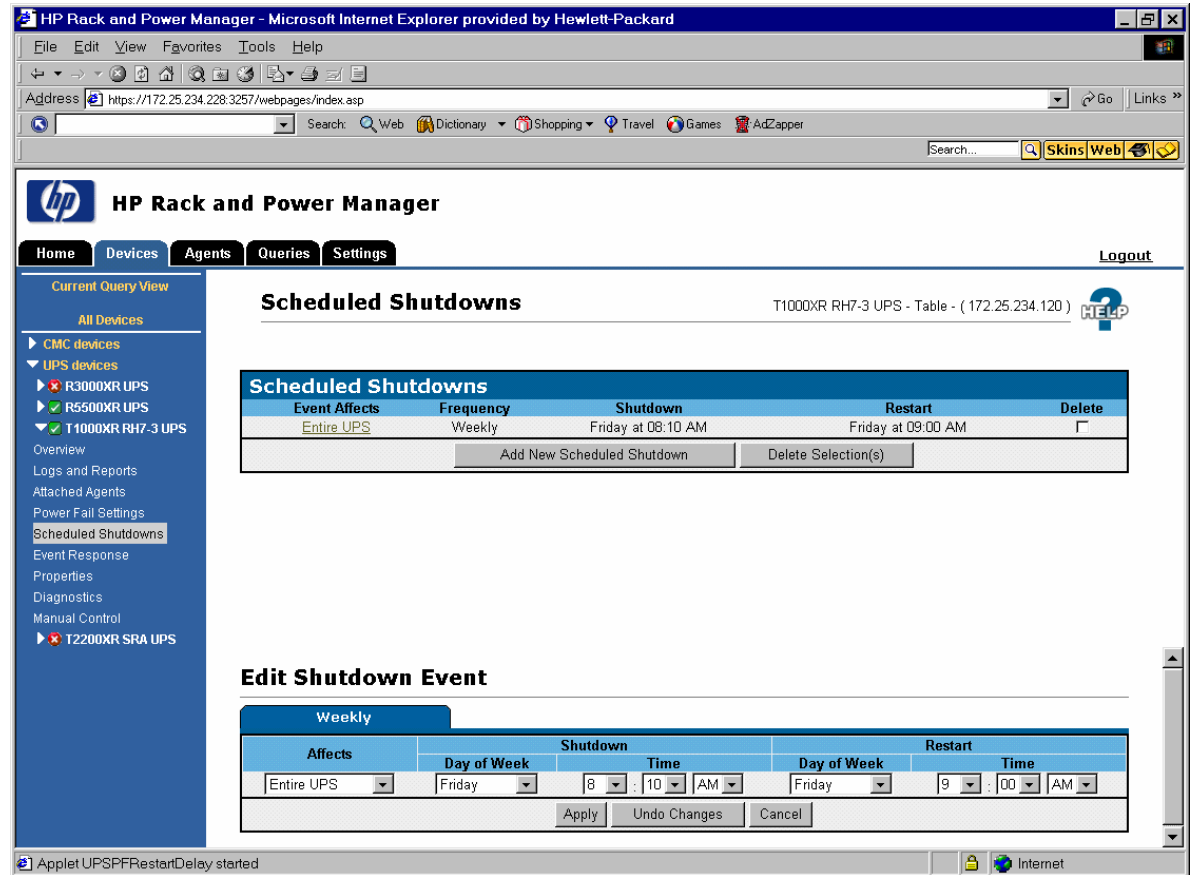
Figure 2, for example, shows a short shutdown time for the ProLiant DL590/64. This screen indicates the UPS will wait approximately 5 minutes before beginning the OS shutdown in this case. Once that time has expired, how long does it take to actually shut down the OS? Since the device will be unable to report when the OS has shutdown, the administrator will have to find out how long that will take and set the value appropriately. To learn this, someone must go through the process of shutting the OS down on the machine and recording how long it takes. This value must then be entered into the "Time required to shutdown OS" field on this screen, shown as approximately one minute for the DL590. By contrast, the bar graphs shown for the ProLiant DL760 on this page are much longer, and in this screen, the actual numbers are visible: 10 minutes before beginning OS shutdown and 4 minutes to accomplish OS shutdown. For the segment containing the DL760, the best practice is to wait as long as possible before going through this relatively lengthy shutdown. This example shows only one server attached to each load segment but that is not a requirement; there could be several. After shutting down the OS for all the machines on a given load segment, it is safe to turn off that segment. Turning off a load segment reduces the UPS load and extends the time available before shutting off other servers.

The Rack and Power Manager implements the ability to "calendar manage" the UPS. As an example, the building facilities team may decide to turn off the building power on a Sunday afternoon for 4 hours to do maintenance. Without a means for managing this case, the UPS protecting a system might inadvertently run itself down during the outage to provide unnecessary backup power. Since the goal is to make sure the UPS is always ready to service the load, that means the batteries need to be charged at all times. If a weekend shutdown discharges the UPS batteries, they might not have time to recharge before an accidental outage could occur, and that would prevent the desired coverage. Charging the batteries can take 3 to 4 hours, but could take longer if additional batteries

have been added to the UPS. Since there is only one charger for all the batteries in the UPS, extra batteries multiply the amount of time it takes to recharge them all until it begins to approach the better part of a day.

To avoid this potential problem, an administrator with advanced notice would enter the time and date of the anticipated shutdown, programming the UPS to shut down at the appropriate time and stay off until the programmed time and date to start it back up again (see **Figure 3**. Shutdown Scheduling). The UPS will then follow these settings from the rack and power manager automatically. The UPS has its own internal clock and is able to store and execute these instructions on its own. It does not need to have any software running on the machine to be able to perform this operation.

Figure 3. Shutdown Scheduling screen



Control and status

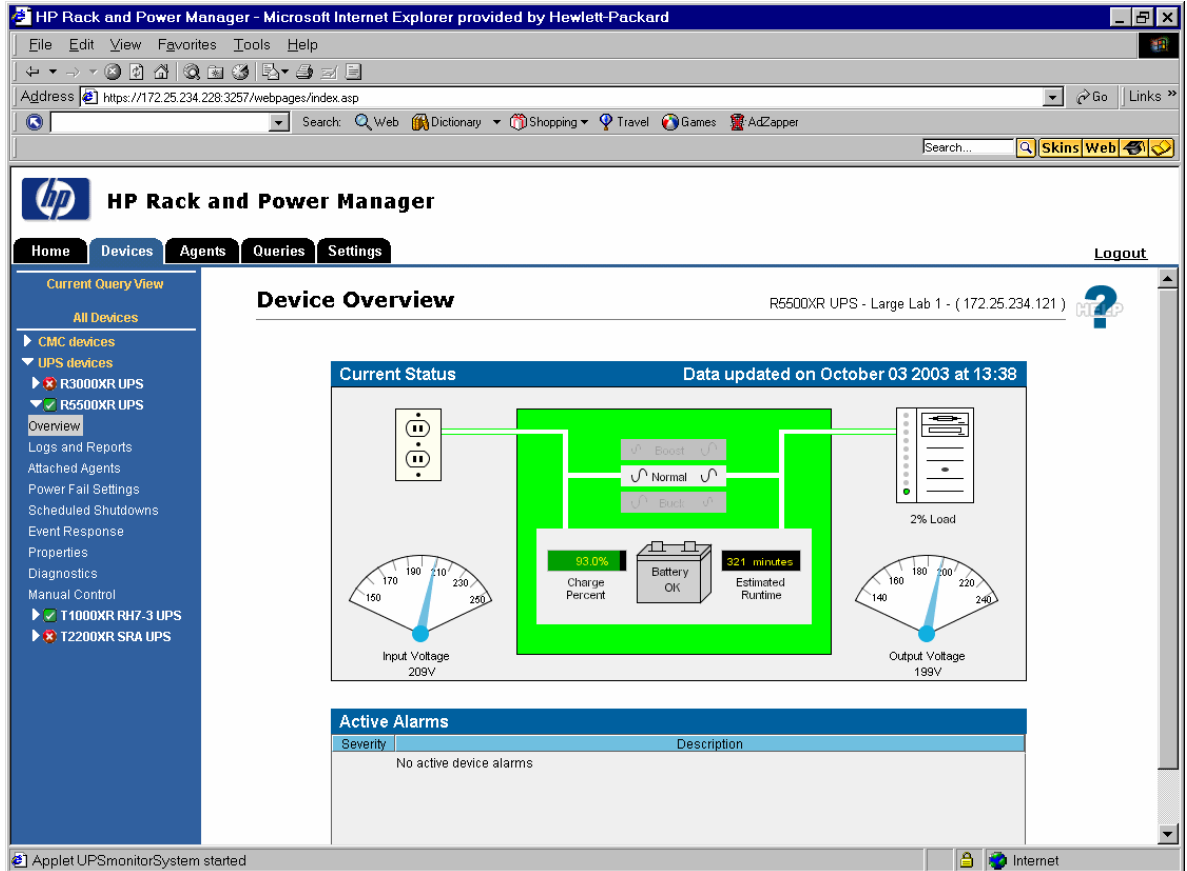
To better understand the capabilities of the software to manage the UPS functionality, it is helpful to examine what is available in terms of control and status for the unit. First, consider the screen shown in Figure 4, the UPS Current Status. The items shown here include:

- Input and output voltage levels need to stay within specified ranges for reliable system operation. If the input voltage is too low (on the order of 110 V to 112V), some of the protected equipment will not be able to operate properly. As a result, the UPS must increase or "boost" the voltage to an acceptable output level, and it handles this task automatically by means of internal circuitry. On the other hand, if the input voltage is too high, the UPS will need to reduce or "buck" it back down to

within the acceptable range. (The product QuickSpecs on the HP website give more detail on UPS operation.²)

- The estimated run time, based on the charge and the loading, serves as a basis for deriving the wait times in the shutdown settings shown in Figure 1. This screen shows the measured values for these parameters.

Figure 4. UPS Current Status screen



It is possible to create logs of power events, actions taken in response to them, or changes to the settings defined for the UPS (see Figure 5). Similarly, it is possible to create graphical logs of these parameters, as shown in Figure 6.

² See http://h18004.www1.hp.com/products/quickspecs/10671_na/10671_na.html for an example.

Figure 5. Event Log screenshot

The screenshot shows a web browser window titled "HP Rack and Power Manager - Microsoft Internet Explorer provided by Hewlett-Packard". The address bar shows the URL "https://172.25.234.219:3257/webpages/index.asp". The browser's search bar contains "Skins Web".

The main content area is titled "HP Rack and Power Manager" and includes a navigation menu with "Home", "Devices", "Agents", "Queries", and "Settings". A "Logout" link is visible in the top right corner. The "Devices" menu is expanded, showing "All Devices" and "Current Query View". Under "Current Query View", there are sub-menus for "CMC devices" (with "CMC 119" and "CMC 205" listed), "Overview", "Logs and Reports", "Sensor Setup", "Accessory Setup", "Event Response", "Properties", and "Manual Control".

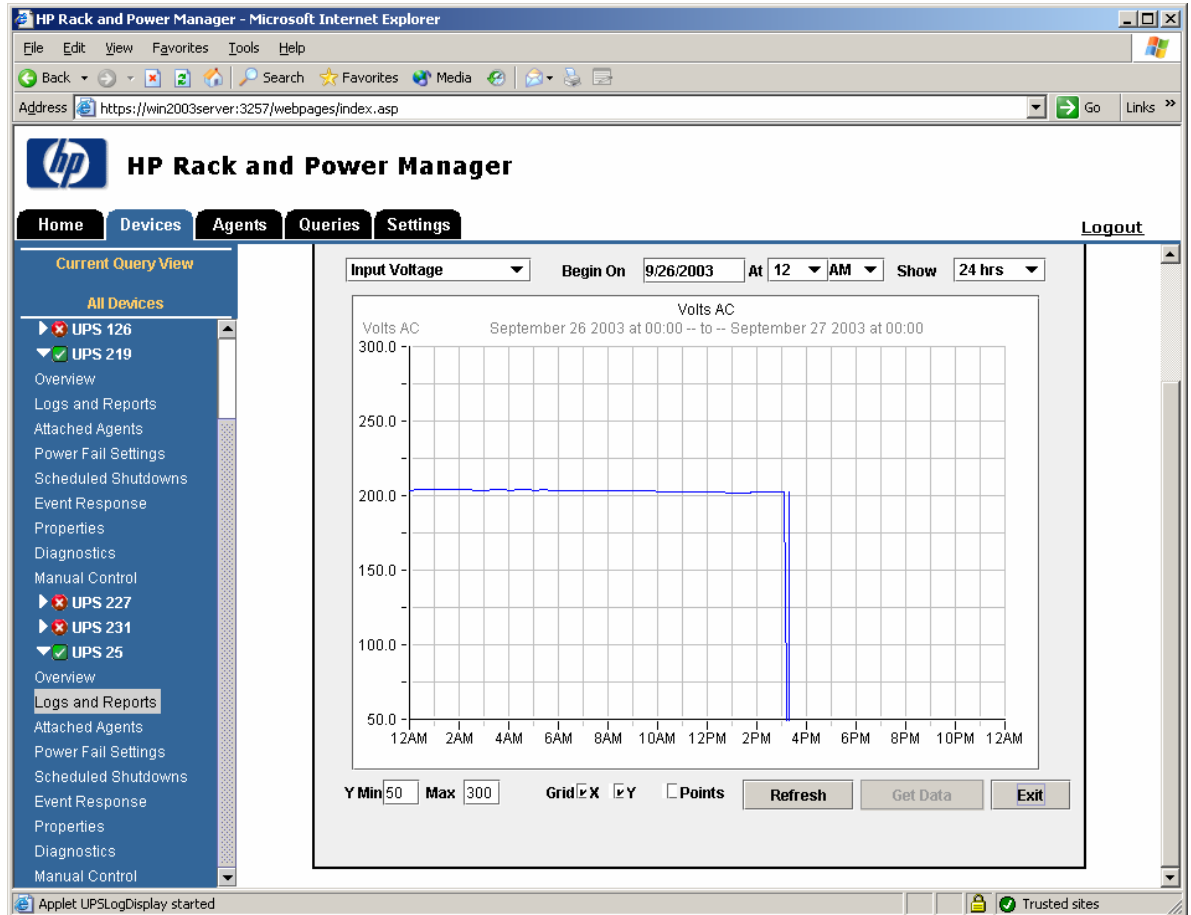
The "Logs & Reports" section is active, displaying a table of "All Logs" for "CMC 205 - Houston - (172.25.234.205)". The table has three columns: "Severity", "Description", and "Date&Time". The logs include:

Severity	Description	Date&Time
Warning	Added Event Action to Event: Temperature 1 Above Warning	Friday September 26 2003 at 10...
Warning	Added Event Action to Event: Temperature 1 Above Warning	Friday September 26 2003 at 10...
Info	Device connected	Friday September 26 2003 at 03...
Info	Shock Cleared	Thursday September 25 2003 at...
Warning	Shock Detected	Thursday September 25 2003 at...
Info	Device connected	Thursday September 25 2003 at...
Info	Device connected	Thursday September 25 2003 at...
Info	Device connected	Thursday September 25 2003 at...
Info	Device connected	Thursday September 25 2003 at...
Info	Device connected	Tuesday September 23 2003 at ...
Info	Device connected	Tuesday September 23 2003 at ...
Info	Device connected	Saturday September 20 2003 at ...
Info	Device connected	Saturday September 20 2003 at ...
Info	Device connected	Saturday September 20 2003 at ...
Info	Device connected	Saturday September 20 2003 at ...

Below the table, there are radio buttons for "All Logs" (selected), "Environmental Logs", "Event Logs", and "Change Logs". At the bottom of the log window are buttons for "Refresh List", "Graphical Logs Screen...", "Print Reports", and "Clear Logs".

The status bar at the bottom of the browser shows "Applet CMCEventLog started" and "Internet".

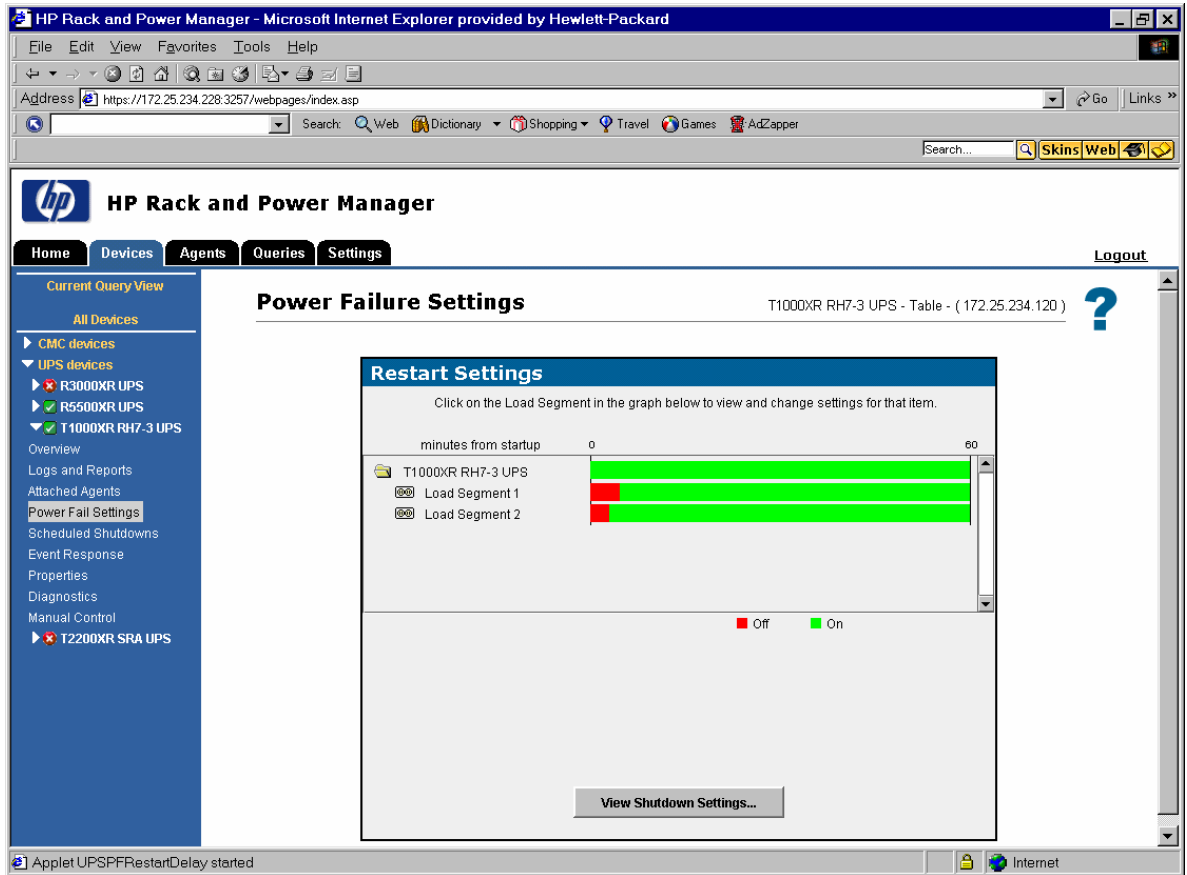
Figure 6. Graphical Log screen



The Power Fail settings (see Figure 2) allow defining the shutdown time on a load-segment basis. An example might be to program a wait time of 10 minutes, and then allow 5 minutes for OS shut down. The load segment must have loads attached to it in order to understand the assignment of time values for shutting it down, and that is a function of its attached agents. Choosing into which load segment to connect this load allows the administrator to define the shutdown operation and to decide how the segments should come back up when restoring power.

When power is restored to the room, what needs to be done? First, all the UPS devices need to be powered on, and then their loads need to be powered up as quickly as possible. However, it is usually not best to simply turn them all back on at once. Instead, best practice is to sequence the order of bringing them back up, much as they were sequenced going down (see Figure 7). The UPS itself is already awake, but staggering the restart of the devices it controls will ease the load burden on the UPS and on the electrical infrastructure.

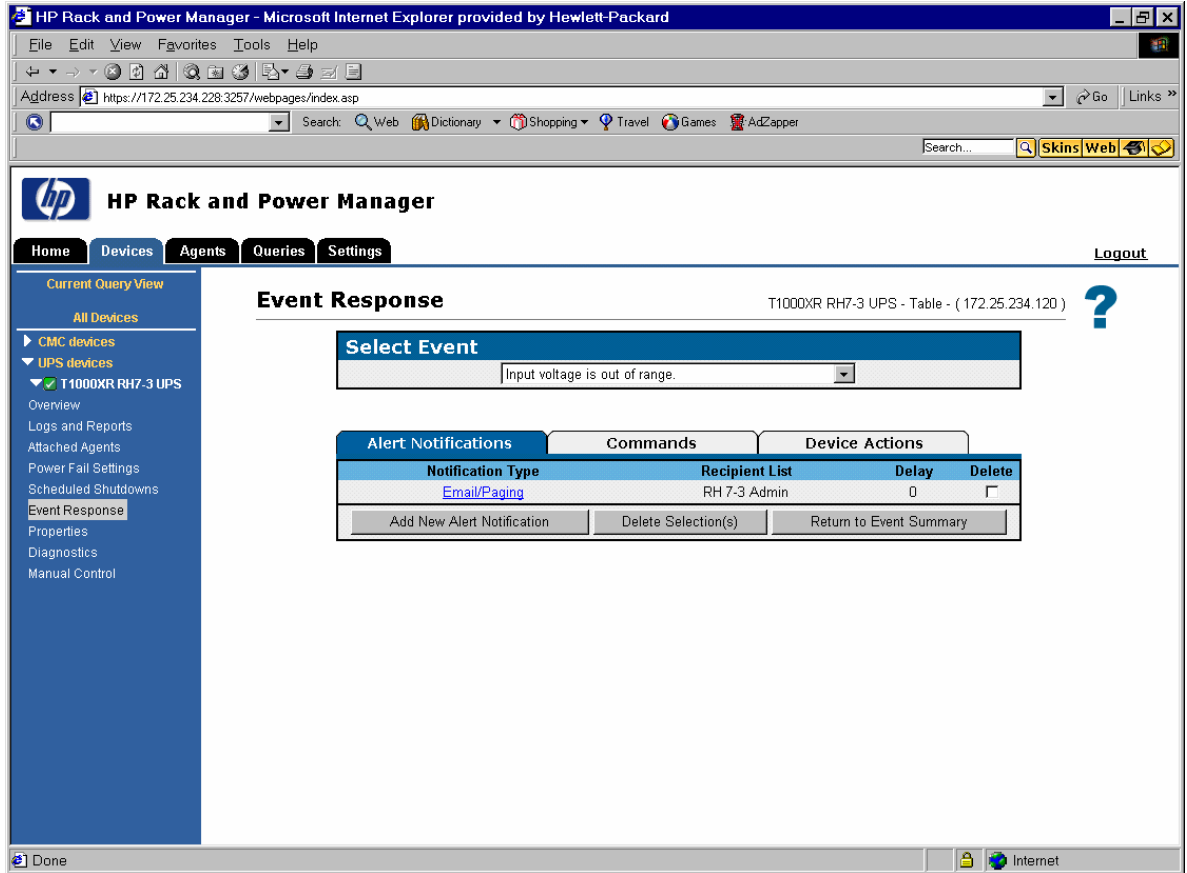
Figure 7. Restart Settings window



On scheduled shutdowns, it is much the same (see Figure 3). The operator chooses the date and time for shutting the UPS down and then a date and time for restoring the power. This can be on a one-time, daily, or weekly basis. For example, it might make sense to shut down the file and print server every week on Friday at 5 p.m. because there is no one in the office to use it, and then bring it back online on Monday at 7 a.m.

For event response, different screens are available (see Figure 8). An administrator can perform an "Add alert notification", specifying that the alert be accomplished with email, broadcast alerts that everyone sees, SNMP traps, etc. When an event occurs, such as a low battery warning, the system can report it and possibly launch an application in response. In the example of a low battery, simply turning off some load segments may be all that is needed.

Figure 8. Event Response screens



Summary

This paper provides a brief overview of the parts of the Rack and Power Manager software that are relevant to UPS control, along with an explanation of some best practices for using it. The focus is on two items: managing the intelligent shutdown of load segments to conserve battery energy, and automating the shutdown and restart of a UPS so it is not inadvertently allowed to run the batteries down. The Rack and Power Manager tool helps ensure that the UPS will maximize both run time and the availability of customer hardware.

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

TC040601, 06/2004

