

Rack UPS Setup

Šarūnas Burdulis
Version 1.0, December 12, 2005

Server rack in 1U Bradley is powered from two rack-mounted UPS'es, each having 1500 VA power capacity (APC Smart-UPS 1500 RM). Rack also holds six IBM e325 1U server modules, 15" LCD monitor, Ethernet and KVM switches. Each UPS powers three servers. Switches and monitor are distributed among the two too.

In the event of a prolonged power failure servers have to be shutdown before UPS batteries are exhausted. In a "one-UPS-per-Linux-machine" setup this is easily achieved by using `apcupsd` daemon to monitor battery status and to send appropriate signals to the system. When more than one machine is powered from the same UPS, one of the machines (the one connected to UPS) has to operate in master mode, providing battery status to other machines (slaves) via network.

Table below gives a comparison of essential configuration settings in `/etc/apcupsd/apcupsd.conf` for standalone, master and slave machines. The first two use USB to connect to UPS.

Standalone	Master	Slave
<code>UPSCABLE usb</code>	<code>UPSCABLE usb</code>	<code>UPSCABLE ether</code>
<code>UPSTYPE usb</code>	<code>UPSTYPE usb</code>	<code>UPSTYPE smartups</code>
<code>DEVICE</code>	<code>DEVICE</code>	<code>#DEVICE</code>
<code>UPSCCLASS standalone</code>	<code>UPSCCLASS netmaster</code>	<code>UPSCCLASS netslave</code>
<code>UPSMODE disable</code>	<code>UPSMODE net</code>	<code>UPSMODE net</code>
	<code>NETPORT 6544</code>	<code>NETPORT 6544</code>
	<code>SLAVE 129.170.28.60</code>	<code>MASTER 129.170.28.59</code>
	<code>SLAVE 129.170.28.61</code>	

Master nodes will send messages to slaves, which should be listening on indicated TCP `NETPORT`. Therefore both TCP wrappers and `netfilter` on slave machines have to allow corresponding incoming traffic from their master. For example on slave to `129.170.28.59` master:

- `/etc/hosts.allow:`
`apcupsd: 129.170.28.59`
- `/etc/iptables.rules:`
`iptables -A INPUT -s 129.170.28.59 -p tcp -dport 6544 -j ACCEPT`

As of December 12, 2005 master-slave arrangement is:

UPS	Master	Slaves
#1	<code>math-01.grid (129.170.28.59)</code>	<code>math-02.grid (129.170.28.60)</code> <code>math-03.grid (129.170.28.61)</code>
#2	<code>math-04.grid (129.170.28.62)</code>	<code>math-05.grid (129.170.28.63)</code> <code>math-06.grid (129.170.28.64)</code>

apcupsd caveat for system reboot after poweroff

Problem: systems will randomly **not** power up after mains power to UPS is restored. This happens in spite of the correct BIOS setting for “AC Power Restore = ON”.

The problem seems to be caused by `apcupsd` not being able to send the signal to UPS to poweroff. After computer has been shutdown, there is no power consumption and the remaining UPS battery charge can last “forever”, i.e. computer will never see power restored, as it was never gone (UPS did not poweroff).

UPS is powered off by `ups-monitor poweroff`, which executes `apccontrol killpower`, which in turn executes `apcupsd --killpower`. For the latter to succeed, filesystems have to be mounted back (read only). This is done by `/etc/apcupsd/killpower`. However `remount` commands are commented out by default and therefore have to be uncommented!

Hopefully this fixes the problem. Tested on `hilbert` and it seems “to do the trick”.