WHY POWER OUTAGES ARE BAD FOR YOUR DATA

1. Introduction

A lot of people assume that when power is cut from their computer, it should be able to handle that gracefully and no data should be lost. Unfortunately, it is not that simple. For a lot of people the risk is low, but when you're using your computer even slightly more seriously and your data is equally important, you might want to consider using a UPS.

This article explains what happens to the hardware and software when the power fails and the possible consequences it has. It should give you enough insight to decide whether you need a UPS or not.

---

2. What happens hardware-wise

2.1. Direct result

When the power fails, no individual component gets a clean shutdown command; power is just removed. When this happens, some parts of the machine may last longer than other parts. One of the first things that will happen, is that the memory DIMMs will no longer be refreshed properly (DRAM needs to be refreshed constantly otherwise it will lose its data) and very rapidly, the memory will contain only garbage. The hard drives and DMA controller however, will run a bit longer; so if data is being written to disk, the DMA controller will keep reading data from

memory, but it has no idea that this data is corrupted. Some file systems are more sensitive to this kind of failure, because of the different kinds of journaling they do.

There are certain machines which are protected against this type of data corruption, by having the power supply send an interrupt to the operating system when power fails, but ordinary class PC hardware does not.

There is another side to this story, however. Researchers have shown that encryption keys can be retreived from memory minutes after the computer has been shut down. This would suggest that the memory doesn't corrupt so quickly at all when power is removed. However, because machines are still unstable when you decrease the refresh cycle of your RAM, it is apparent that some corruption still occurs.

2.2. Indirect result

Not directly related to a power loss as much as any kind of shutdown, is the fact that hard drives which have reached a certain age tend to die when the machine is powered up again. This doesn't necessarily happen immediately, but can take a few days. Additionally, you can imagine what will happen in this scenario when you have an array of identical disks of equal age; always put spare disks in your array, in power safe mode so they don't wear down.

---

3. What happens software-wise

3.1. Disk cache

Disk write cache is used to collect and delay transfers to the disks in favor of speed, because memory is faster than disks. When you shut down a machine

when there is uncommitted data in the cache, you will lose this data, or corrupt it because only part of the cache is written. This can be illustrated very nicely by booting your Linux machine with the kernel parameter "init=/bin/bash". This will start a shell instead of the initialization procedure. You can then edit files, like /etc/shadow, should you want to reset your password. If you then press ctrl-alt-del without running the "sync" command to commit the disk cache first, your changes will not be committed to disk.

There are different kind of cache systems in existence. Two important ones for write cache are write-through and write-back. The former is safe, because it reports the data as written when the data is committed to disk. The latter is unsafe, because it reports the data as written when it has been written to cache, while it hasn't been written to disk yet. Even when you have a UPS it's unsafe, because there are several other reaons when power can suddenly disappear.

Write-back cache is used in a lot of disks these days. If possible, you may want to consider turning it off.

3.2. (Encrypted) file systems

Most people will think that because of journaling, file systems are protected against power failures. It's true that filesystems with journaling are more robust than those without, but it should be clear by now there are some things the file system cannot protect against.

Then of course, there are different ways that journaling can be done. Ext3 is more resillient against power failures then XFS and ReiserFS, because ext3 does physical block journaling. Ext4, however, by default does something called "delayed allocation," which means that meta-data is saved more often than the

data itself. In the event of a power failure, this can data corruption, for obvious reasons. Linux Torvalds has an outspoken opinion against delayed allocation.

As the Gentoo Wiki states, you are even more susceptible to data loss in the event of a power failure when using an encrypted file system. The reason for this, is that hard disks are block devices. Normally, if a few bits are flipped as you write it to disk, you can take advantage of properties inherent in the data (depending on what it is) to recover it. At the very least, you only lose a bit of the block. Should the block be encrypted and trashed as it is written, the flipped bit will cause decryption of that block to fail, so it is all lost. Entropy of most computer data on disk is quite low, while encrypted data is essentially indistinguishable from true random; guessing gets a lot harder. Additionally, it is possible for multiple subsequent blocks to also be lost due to the initialization vectors used in IV chaining being unrecoverable. This depends on whether or not you're using them, how the cipher is configured, and other factors, but is a consideration.

3.3. (Linux software) RAID

Linux software RAID, and any RAID basically, needs to know if the disks of the array are still properly matched to eachother when the array is initialized. When power fails, or when you press reset, they will be in a "dirty" state, and the system may need to recreate the array. That is, if it can. I've never tried it, but I can imagine that a RAID0 can be completely destroyed by a power failure. But, don't take my word for that...

Modern Linux kernels (2.6.16 and newer) and raid tools (mdadm 2.4.1 and newer) luckily have a precaution against that, namely a write intent bitmap. When using Linux software RAID, I'd advise you to enable this. There are enough resources on the internet where you can find how, like the Gentoo Wiki. Unfortunately, write

intent bitmaps are very slow. When you have the protection of a UPS, not using a write intent bitmap becomes more acceptable.

3.4. Databases

When you use a database system with (good) transaction support, data corruption will not happen when the power is removed. That is, when you're using disks that don't lie about their cache status (see above). You will, of course, lose all the uncommitted transactions, which can be annoying in itself.

When you're running a database without transaction support, like MySQL MyISAM, data corruption is of course likely.

Applications like LDAP directories, source control management repositories, etc, are also potentially susceptible to the same kind of failures, depending on if and how well they use transactions. Like I describe in my backup article, it's also important to make scheduled dumps of such applications, to make sure have a backup in a robust, self-contained archive.

---

4. Surge protection

UPSes also protect against surges on the mains power. However, only an online UPS (the expensive sort) does it properly, by always running the load from the battery. The offline variety merely uses MOVs, which is exactly the same thing as those ordinary power strips with surge protection. The effectiveness of those things can be questioned because of delay time, impedance (resistance) of the safety earth, longetivy of the MOVs because of frequent surges, etc. The internet is filled with information about the fallibility of MOV surge protection.

The surge protection on UPSes also often includes protection for ethernet and/or telephone networks. I really advice against using those. When there is a surge, the MOVs temporarily short the line containing the surge with the safety earth, but it will also connect the data networks to it. This safety earth, however, does not have infinitely low impedance, and therefore it's possible that some of the excess current will travel up the network, as opposed to down the safety earth. The exact details of this are more complex than this, but as always, the internet is your tool should you want to find out more.

---

5. Recommendation

It should come as no suprise that I would advise a UPS if your data is important to you, especially when the machine in question is heavily used data-wise (with a lot of writing to the disks), like an office file server, or when it uses a database of some sort. And even more so if you use XFS or ReiserFS. It's also convenient to have your external USB disk, router, cable modem, telephone switch, or similar devices on the UPS. The router and network switches is particularly neat, because then the machine can notify you (by e-mail or SMS) of the power failure, and possible connections from the internet (like your SSH shell...) will be maintained.

5.1. Selecting a UPS

In my experience, servers often use a lot less power than you'd think (at work, our dual core 1.8 GHz Intel with three disks only uses about 100W), so you don't really need a big UPS. I'd advice getting/borrowing a power/VA meter to gauge

the power use of the machine in question, and size the UPS accordingly. When doing so, remember that it will use more power when the CPU is highly utilized.

Our 500 VA APC CS 500 can supply one server, one pentium-100 (internet router), phone switch and network switch for about 15 minutes. It's loaded at 40% under normal use.

I would not connect a CRT monitor to a UPS. When the degaussing coil of the picture tube is triggered (when the monitor is turned on for example), an enormous surge of current is drawn. When you're running on battery power, this is enough to make computers connected to it reboot. Should the connected devices ever start on battery power, each connected device will suffer a momentary power outage because of the high current drawn by the degaussing coil of the monitor. This kind of outage is not very good for your hardware. Besides, what use does it have to run the monitor off a UPS? The primary function of a UPS is to avoid data corruption because of an ugly shutdown, not because you are too lazy to save your work all the time...

Source : http://www.halfgaar.net/why-power-failures-are-bad-for-your-data