

Re: Hard Disk Performance

Source: <http://unix.derkeiler.com/Newsgroups/comp.unix.sco.misc/2004-05/0281.html>

From: Bela Lubkin (*belal_at_sco.com*)

Date: 05/19/04

Date: Wed, 19 May 2004 11:54:27 GMT

To: scomsc@xenitec.ca

IanC wrote:

- > *We have about 250 SCO Openserver installations around the UK. All of the*
- > *installations run Openserver 5.0.2 on a variety of 'low spec' single CPU*
- > *PC's with IDE disks. The specs range from a 1998 266MHz Pentium with 2Gbyte*
- > *disk up to 1009MHz Pentium with 20Gbyte disk.*
- >
- > *Some of the installations are taking much longer than others to run disk*
- > *intensive jobs.*
- >
- > *Every night each system does a backup to a single disk file (Using Backup*
- > *Edge from Microlite). The single disk file is then copied to tape. We have*
- > *used the time taken to perform the backup to disk as a measure of the*
- > *systems disk performance. A typical backup size is 800Mbytes of data which*
- > *is compressed into a single backup file of 250Mbytes.*
- >
- > *Our problem is that some sites take as little as 5 minutes to perform the*
- > *disk backup whilst others take over an hour. It is NOT the speed of the*
- > *processor or the spec of the disk drive which determines how long the*
- > *process take. Some older machines perform very well and some modern ones*
- > *perform badly.*
- >
- > *We have taken two systems with a similar spec. One takes 10 minutes to*
- > *perform the backup. One take 30 minutes. We find that they have different*
- > *types of hard disk drive. (One is a 40Gbyte Western Digital and the other a*
- > *40GByte Samsung) If we swap the hard disk drive then the backup times move*
- > *with the hard disk. We have tried installing Linux on these two systems and*
- > *the backup times under Linux are fast on both machines – (6 minutes)*
- >
- > *We have deduced that this may be some sort of driver problem which would be*
- > *solved by installing a new SCO disk driver (rather than the old one which is*
- > *in 5.0.2)*
- >
- > *Has anybody got any thoughts on this and can anybody tell me if it is*
- > *possible to install the latest disk drivers onto 5.0.2. On the SCO web site*
- > *driver packages are only available for 5.0.6 and 5.0.7*

The "wd" IDE driver in OSR502 has no DMA support at all. All IDE I/O is done with programmed I/O (PIO). The driver does not attempt to control the PIO mode; it relies on the BIOS to have set the drive to a sensible mode.

Take one of the systems where you noticed a big speed difference between two similar-spec drives. With one drive in the system, go into BIOS setup and see if it shows you anything about the chosen PIO mode. Many (but not all) BIOSes offer choices like "auto" and "PIO0-PIO4". The BIOS will likely be set to "auto". Hopefully it will then show what PIO mode it has automatically chosen. Assuming the BIOS is cooperative (by showing what it's done) — you will probably be able to see that it's chosen different modes for the two drives.

Next, read the drives' manuals and check what PIO modes they actually support. 40GB drives will be fairly new and probably support PIO mode 4, but it's best to check.

Assuming the BIOS allows you to manually override, tell it to use the highest mode that the drive and BIOS both support. (The BIOS may top out at PIO mode 3, for instance — either because it was written before PIO mode 4 was invented, or because it knows the hardware can't handle that.)

The modes may be presented differently by the BIOS. Some use mode numbers, others use speeds. The modes are:

PIO mode 0 3.33 MB/sec
PIO mode 1 5.22 MB/sec
PIO mode 2 8.33 MB/sec
PIO mode 3 11.11 MB/sec
PIO mode 4 16.67 MB/sec

=====
The current (OSR507+) version of the "wd" driver supports Ultra DMA modes up to 133 MB/sec. If you could use it, it would probably make a substantial performance difference compared to the PIO-only OSR502 driver.

The current "wd" will not link into an OSR502 kernel, but the incompatibility is very small, easily worked around. I have tested linking an OSR502 kernel with the current driver and the following workaround. I have not booted the machine (it isn't mine to boot). I think it will work, but if you attempt this you should first make sure you have good backups and crash recovery disks, be prepared for a potential crash.

If it passes the "smoke test" and boots successfully, it's probably fine — I would expect a pretty immediate failure if there's a problem.

comp.unix.sco.misc: Re: Hard Disk Performance

The workaround: well, first install the "Wd Driver Supplement" from <ftp://ftp.sco.com/pub/openserver5/drivers/OSR507/btld/wdsupp/>. Do not use the BTLT from ../wd; it won't work on OSR502.

Once the supplement is installed, if you try to relink the kernel, it will fail with a missing symbol "swab". To correct that, edit /etc/conf/patch.d/wd/space.c. Find "#if !defined(MODESELECT10_CMD)". A few lines after that, after the "#endif" line, add the following (everything between but not including the ===== lines):

```
=====
/*
 * swab() was new in OSR506. Provide a version of it here, if needed.
 */
#if !defined(WCE_BIT) /* WCE_BIT first appears in <sys/scsi.h> in OSR506 */
void
swab(src, dst, n)
char *src;
char *dst;
int n;
{
    while (n > 1) {
        n -= 2;
        dst[0] = src[1];
        dst[1] = src[0];
        src += 2;
        dst += 2;
    }
}
#endif
=====
```

Again, I think this will work but I haven't tested it. If you test it, please let me know (1) if you got it to link in, (2) if it booted, (3) if you saw any performance change.

With the new driver, you can add a bootstring "wd.debug=iug". The letters are 'i' for IDE controller identification, 'u' for UDMA configuration and 'g' for geometry determination. It might be interesting to boot with "wd.debug=u" once with each hard drive, see if the configurations differ.

>Bela<