

How does hardware affect Oracle performance?

Don Burleson

http://searchoracle.techtarget.com/tip/0,289483,sid41_gci1034639,00.html?track=NL-93&ad=501858

The Oracle professional has always had to learn how improvements in hardware will affect their Oracle database performance. This article will explore the recent advances in hardware and show how Oracle pros can leverage these new hardware technologies to improve the performance of their databases.

At a high-level, we see these important trends for Oracle:

- All hardware gets cheaper every year
- CPU always gets faster ? RAM and disk does not
- RAM-Disk will soon replace platter disks

Let's take a closer look and understand how the hardware will affect Oracle databases.

Moore's Law and Oracle

The economics of server technology has changed radically over the past 40 years. In the 1960s, IBM dominated the server market with giant mainframe servers that cost millions of dollars. These behemoth mainframes were water-cooled and required a huge operations center and a large staff to support their operations.

Gordon Moore, Director of the Research and Development Laboratories at Fairchild Semiconductor, published a research paper titled "Cramming more components into integrated circuits" in 1965. Moore performed a linear regression on the rate of change in server processing speed and costs, and noted an exponential growth in processing power and an exponential reduction of processing costs. This landmark paper gave birth to "Moore's Law," which postulated that CPU power gets four times faster every three years

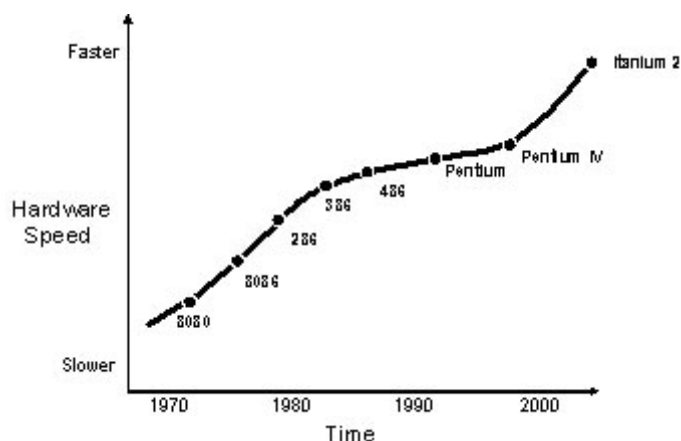


Figure 1. Moore's Law

Moore's law has held true over the three decades since it was first published, and Moore's law has been extended to cover RAM and disk storage costs as well.

However, the "real" Moore's Law cannot be boiled-down into a one-size fits-all statement that everything always gets faster and cheaper. Prices are always falling, but there are important exceptions to Moore's Law, especially with regard to disk and RAM technology (Figure 2).

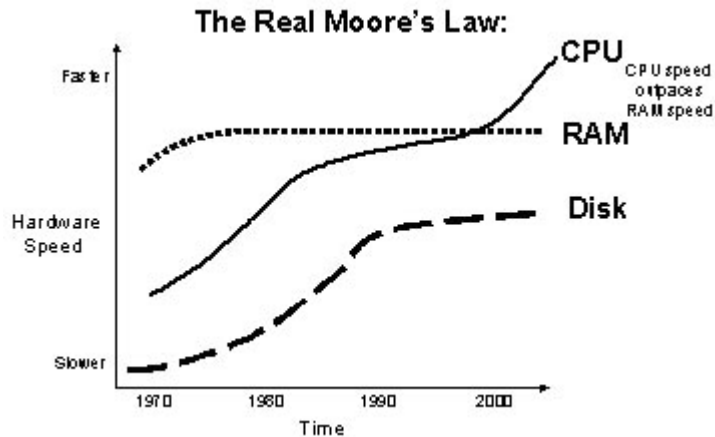


Figure 2. The "real" Moore's Law

As we can see, these curves are not linear and this trend has a profound impact on the performance of Oracle databases. Let's take a closer look.

Storage Costs

I'm old enough to remember when punched cards were the prominent data storage device. Every year I would get my income tax refund check on a punched card and we would make Christmas trees from punched cards in the "Data Processing" department.

College-aged kids have no idea what the term "Do not fold, spindle or mutilate" means and they missed-out on the fun of dropping your card deck on the floor and using the giant collating machines to re-sequence them.

- Punched Cards: (1960-1970)
- Drums: (1970-1980)
- Disks: (1980-2000)
- Solid-State Disk: (2000-2010)

In 1985, a 1.2 gigabyte (GB) disk (the IBM-3380 disk) sold for more than \$250,000. Today, you can buy 100 GB disks for \$200, and 100 GB of RAM Disk (solid-state disk) for \$100,000.

We see the following trends for Oracle disk in this decade:

- Disk storage costs fall 10x every year.
- Magnetic-coated spinning platters are obsolete and cannot go faster
- RAM-SAN replaces disk by 2008.

Note the change to Moore's Law for disks show the limitations of the spinning platter technology (figure 3).




Figure 3. Moore's Law for disk

Platters can only spin so fast without becoming aerodynamic and the disk vendors were hard-pressed to keep their technology improving in speed. Their solution was to add a RAM front-end to their disk arrays and sophisticated asynchronous read-write software to provide the illusion of faster hardware performance.

RAM Costs

Today, you can buy 100 GB of RAM for only \$100,000, with access times 6,000 greater than traditional disk devices. By 2013, a gigabyte of RAM will cost the same as a gigabyte of disk today (about \$200). RAM I/O bandwidth capacity doubles every ten years, making the 128-bit architecture due in about 2010:

- 8 bit 1970s
- 16 bit 1980s
- 32 bit 1990s
- 64 bit 2000s
- 128 bit 2010s

Note that Moore's Law does not apply at all to RAM and the speed has been fixed for nearly 30 years.




Figure 4. Moore's Law for RAM

Because CPU speed continues to outpace memory speed, RAM subsystems must be localized to keep the CPUs running at full capacity.

Processor Costs

The same trend also exists for processor costs and speed. In the 1970s, a 4-way SMP processor costs over \$3,000,000. Today, the same CPU can be purchased for under \$3,000. CPUs continue to increase speed by four times as much every three years and cut cost in half.

So, with this ongoing improvement in speed and rapidly declining costs, where does the Oracle shop decide to adapt to a new technology? This is a very important question to the Oracle professional, and one that warrants careful investigation. Here is how the hardware is going to change our lives with Oracle in the next few years:

- CPU speed continues to outpace memory speed. This means that RAM sub-systems must be localized to keep the CPU's running at full capacity and that RAM will move out of the data buffer caches and in to the back-end disks.
- Platter disks are being replaced by [solid-state RAM disk](#). Within the next decade all Oracle databases will be 100% solid-state.
- Oracle databases are shifting from being I/O-bound to CPU-bound as a result of improved data caching. This is reflected in the new `cpu_cost` feature of the Oracle10g SQL optimizer.
- Using hardware to correct Oracle performance problems has become increasingly attractive to management.

Let's take a closer look at how the super cheap, super fast hardware is changing Oracle tuning activities:

Hardware hits the Oracle bottlenecks

Remember, every Oracle database has a bottleneck, no matter how well it is tuned, and a quick look at the top five timed events will reveal the bottleneck. Remember, even a properly-tuned Oracle database system has a bottleneck (top-5 timed events), and the ONLY way to speed-up a well-tuned Oracle database is to hit the bottleneck.

- Faster CPU will speed-up ANY CPU-Bound Oracle database.
- High-speed Solid-state disk will speed up ANY I/O-bound Oracle database.
- Faster network will speed-up ANY "SQL*Net" bound Oracle database.

For example, here is a CPU-bound Oracle10g database:

Top 5 Timed Events

Event	Waits	Time (s)	% Total Ela Time
CPU time	4,851	4,042	55.76
db file sequential read	1,968	1,997	27.55
log file sync	299,097	369	5.08
db file scattered read	53,031	330	4.55
log file parallel write	302,680	190	2.62

Assuming that the library cache and the SQL are already optimized, more CPU's or faster CPUs will improve the performance of this database.

The super-cheap Intel Itanium2 processors are making huge inroads for these types of Oracle databases. Last month I upgraded a CPU-bound database (already SQL tuned) to these faster 64-bit processors and their batch jobs went from 2 hours to less than 15 minutes.

But what about an Oracle database that needs tuning? Is it "legitimate" to throw hardware at a poorly-tuned database? Let's take a closer look.

Using Hardware as an Alternative to Oracle tuning?

Throwing hardware at an Oracle problem is a legitimate management decision, and in some cases a cost-effective one, especially when the database required a huge amount of manual effort.

For example, when the schema design is messed-up, even materialized views may fail to correct the performance issues. In my experience, IT management DOES NOT want to hear about an expensive re-design of the database for several reasons:

- Fear of Blame -- No manager wants to admit that a bad database design was implemented under their watch.
- High Downtime -- The time to re-implement a table re-design can mean DAYS of downtime
- High Cost -- If the crappy implementation costs \$200k, what do you think the chances are that management will spend the money all over again to re-design the system properly? Zero.

These are the shops that want a fast, save-face approach to improving performance, and like it or not, it happens every day.

Just last week I had a client who was having a huge CPU bottleneck, and the root cause was excessive parsing and really sub-optimal SQL execution plans. They chose to spend \$50k for faster processors (15 minutes to fix) rather than spend \$100k to tune 2,000 SQL statements (6 weeks to fix).

Now does that make sense? Well, it depends on how you look at it. From a management perspective it made sense as they saved \$50k and got a fast, unobtrusive solution to a complex problem.

Was it elegant, heck no! Was it the "right" thing to do? You be the judge...

Oracle professionals cannot ignore the economic reality of database tuning. Time and time again, it's too costly (in both time and money) for a shop to tune their SQL.