
THE FLASH

Flash memory talk
Felton Linux Group
27 August 2016

Jim Warner

Flash Memory Summit

Annual trade show at Santa Clara Convention Center

Where there is money, trade shows follow.

August 8 – 11, 2016

Borrowing liberally from . . .

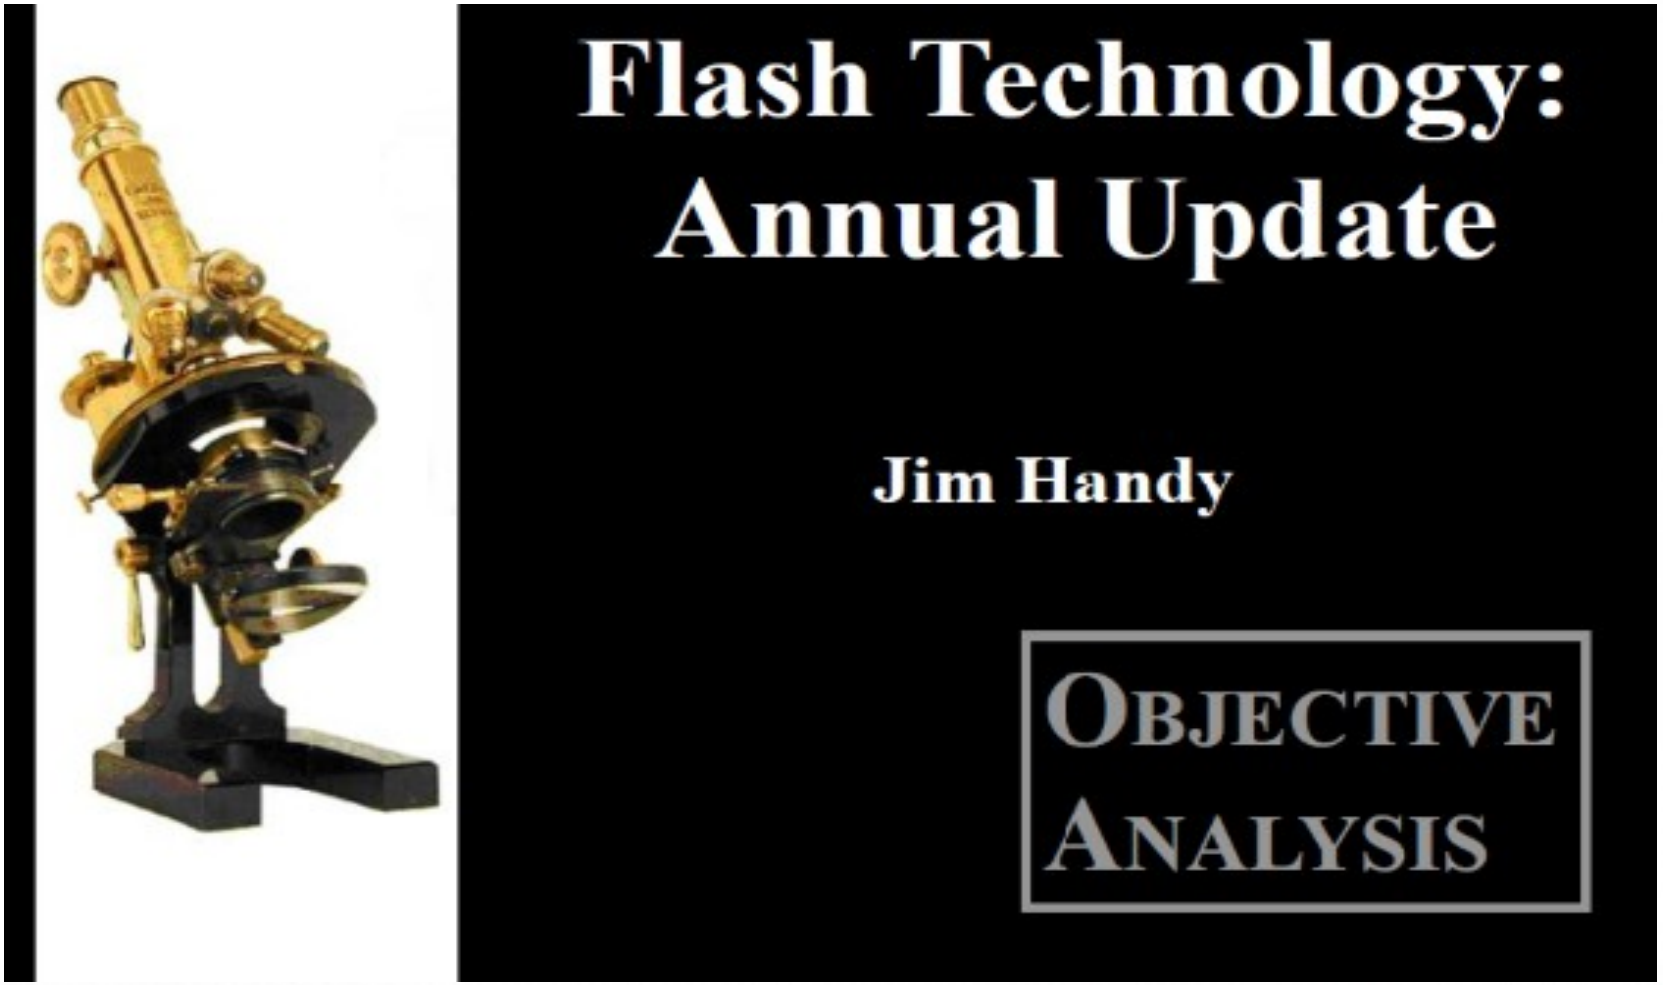

Industry analysts

SSD

- Flash memory is [now] component to make solid state disks
- Bits stored as charge in conductive regions surrounded by an insulator.
- Electric field from the stored charges is detected by a nearby field effect transistor

Remember ?

- UV Eraseable EPROMs ?
- These were the first devices to use isolated charge regions for computer storage.

How flash works . . .

"Each NAND flash memory cell is a floating gate transistor whose threshold voltage can be configured (or programmed) by injecting certain amount of charges into the floating gate."

Yangyang Pan, Guiqiang Dong, and Tong Zhang

Fowler-Nordheim tunnel

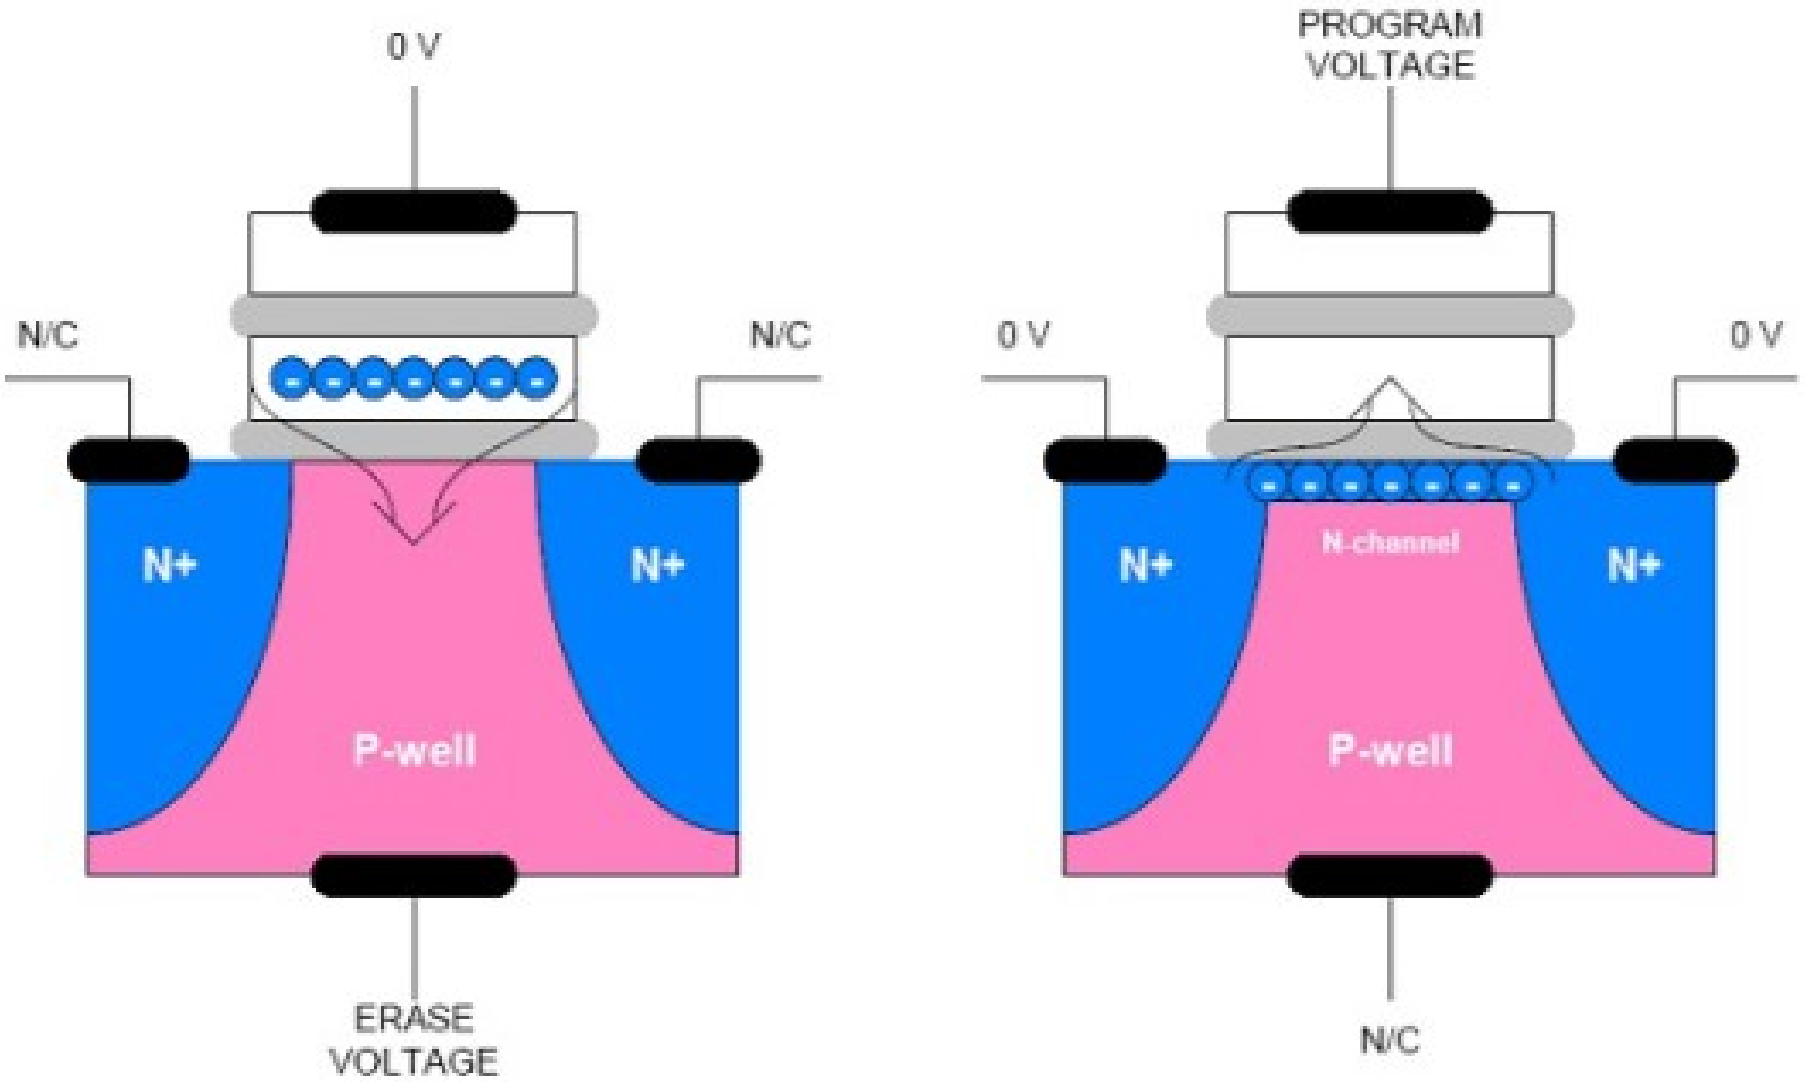

Two kinds . . .

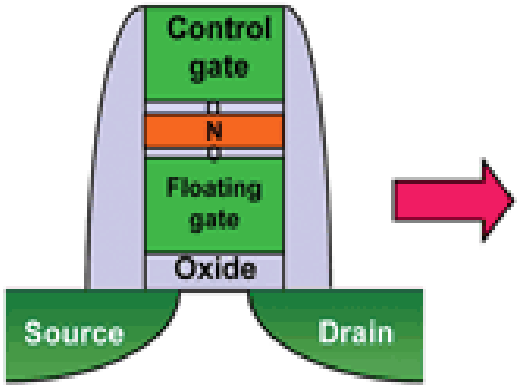

Single cell structure

	NOR	NAND
Cell Array	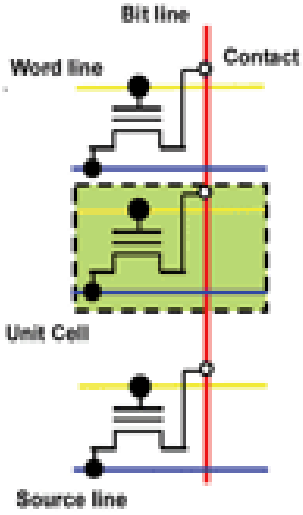	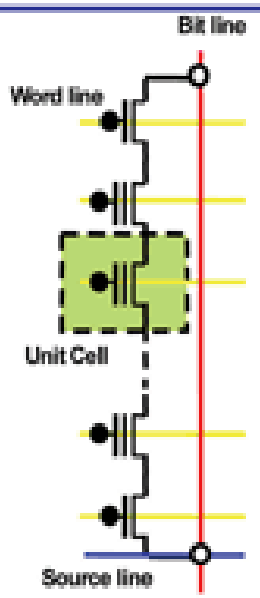
Layout	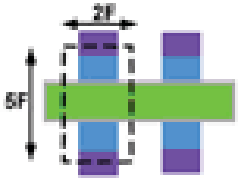	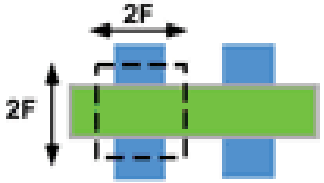
Cross section	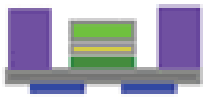 $10F^2$	 $4F^2$

Bits per cell

- **SLC** - Enterprise class flash stores a single bit in each cell
- **MLC** - Density can be increased by storing two bits in each cell. This is done by detecting four discrete charge levels.
- Apple music players made MLC popular.
- Devices are now available that store 3 bits per cell.

Wear out

- **Not** holes in the tunnel insulator
- High voltages during P/E cycles cause electrons to be trapped in insulating regions near the floating gate.
- These charges affect the transistor threshold adding noise to the read thresholds
- Rogue charges cannot be erased

NAND flash is doing well . . .

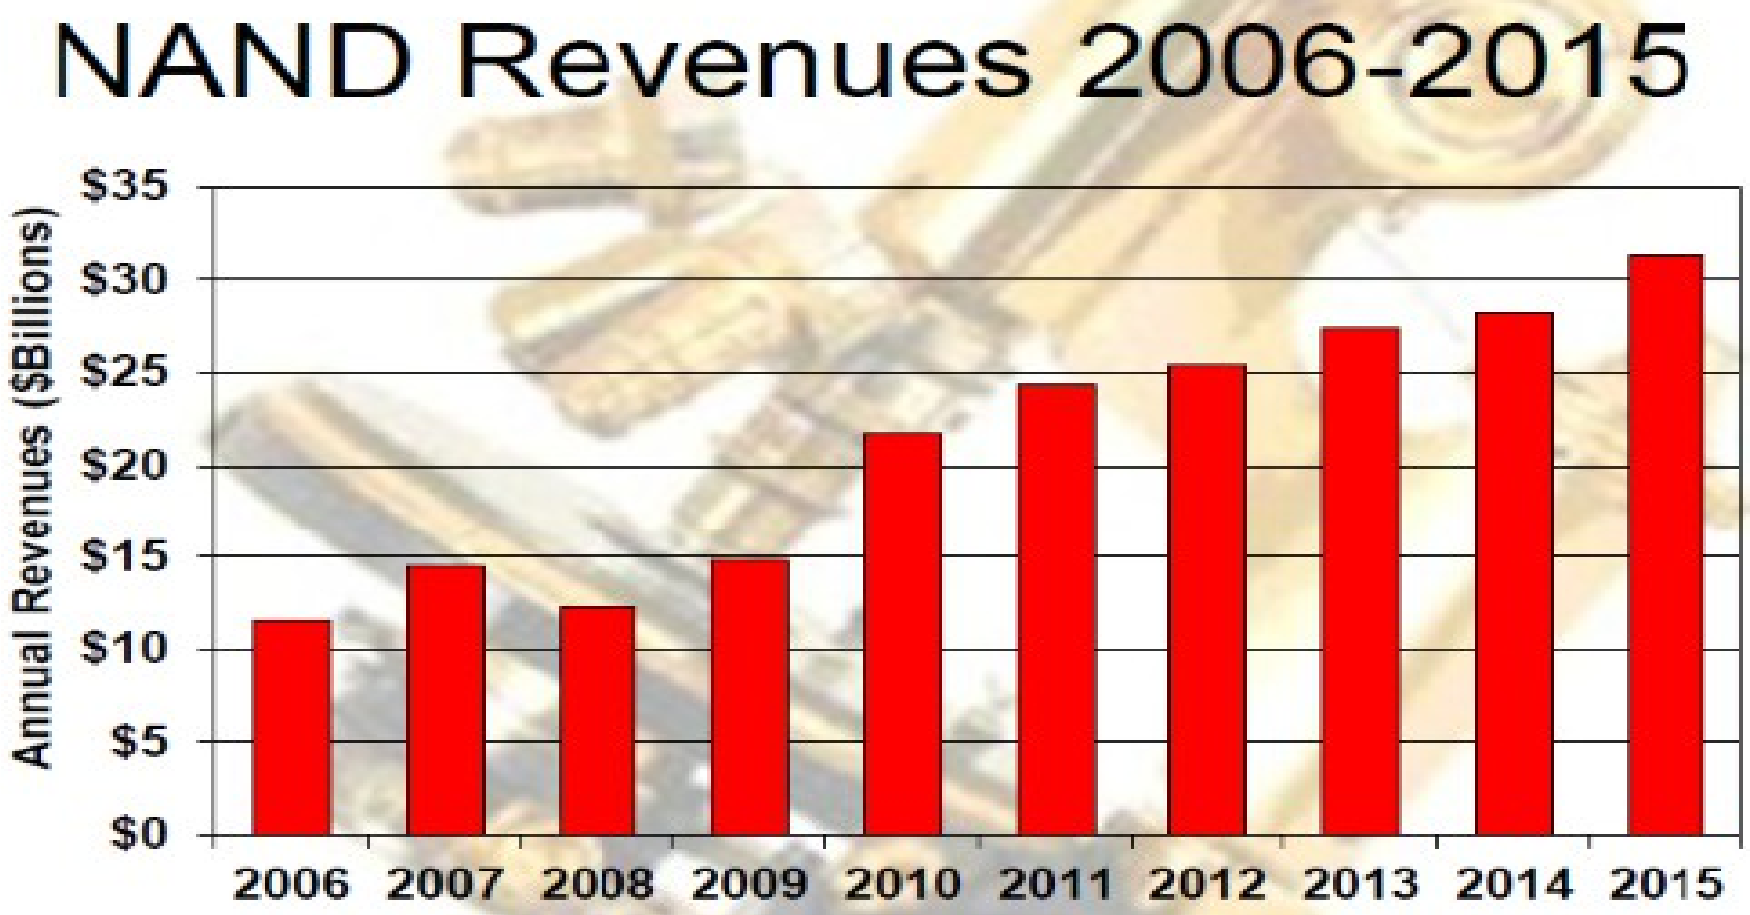

NOR flash, not so much

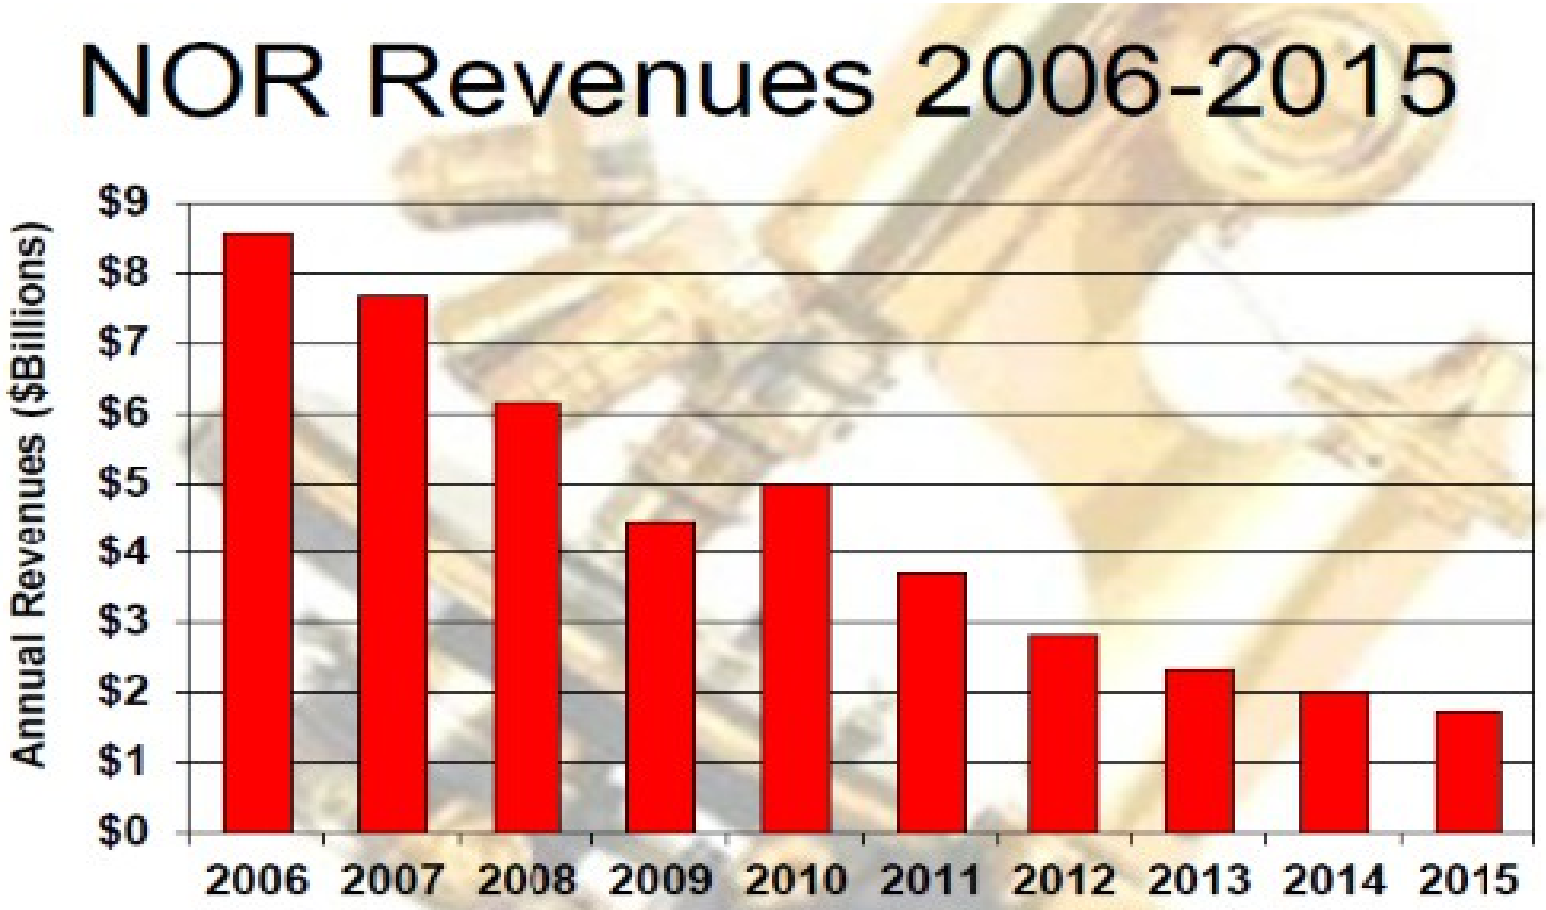

Memory prices . .

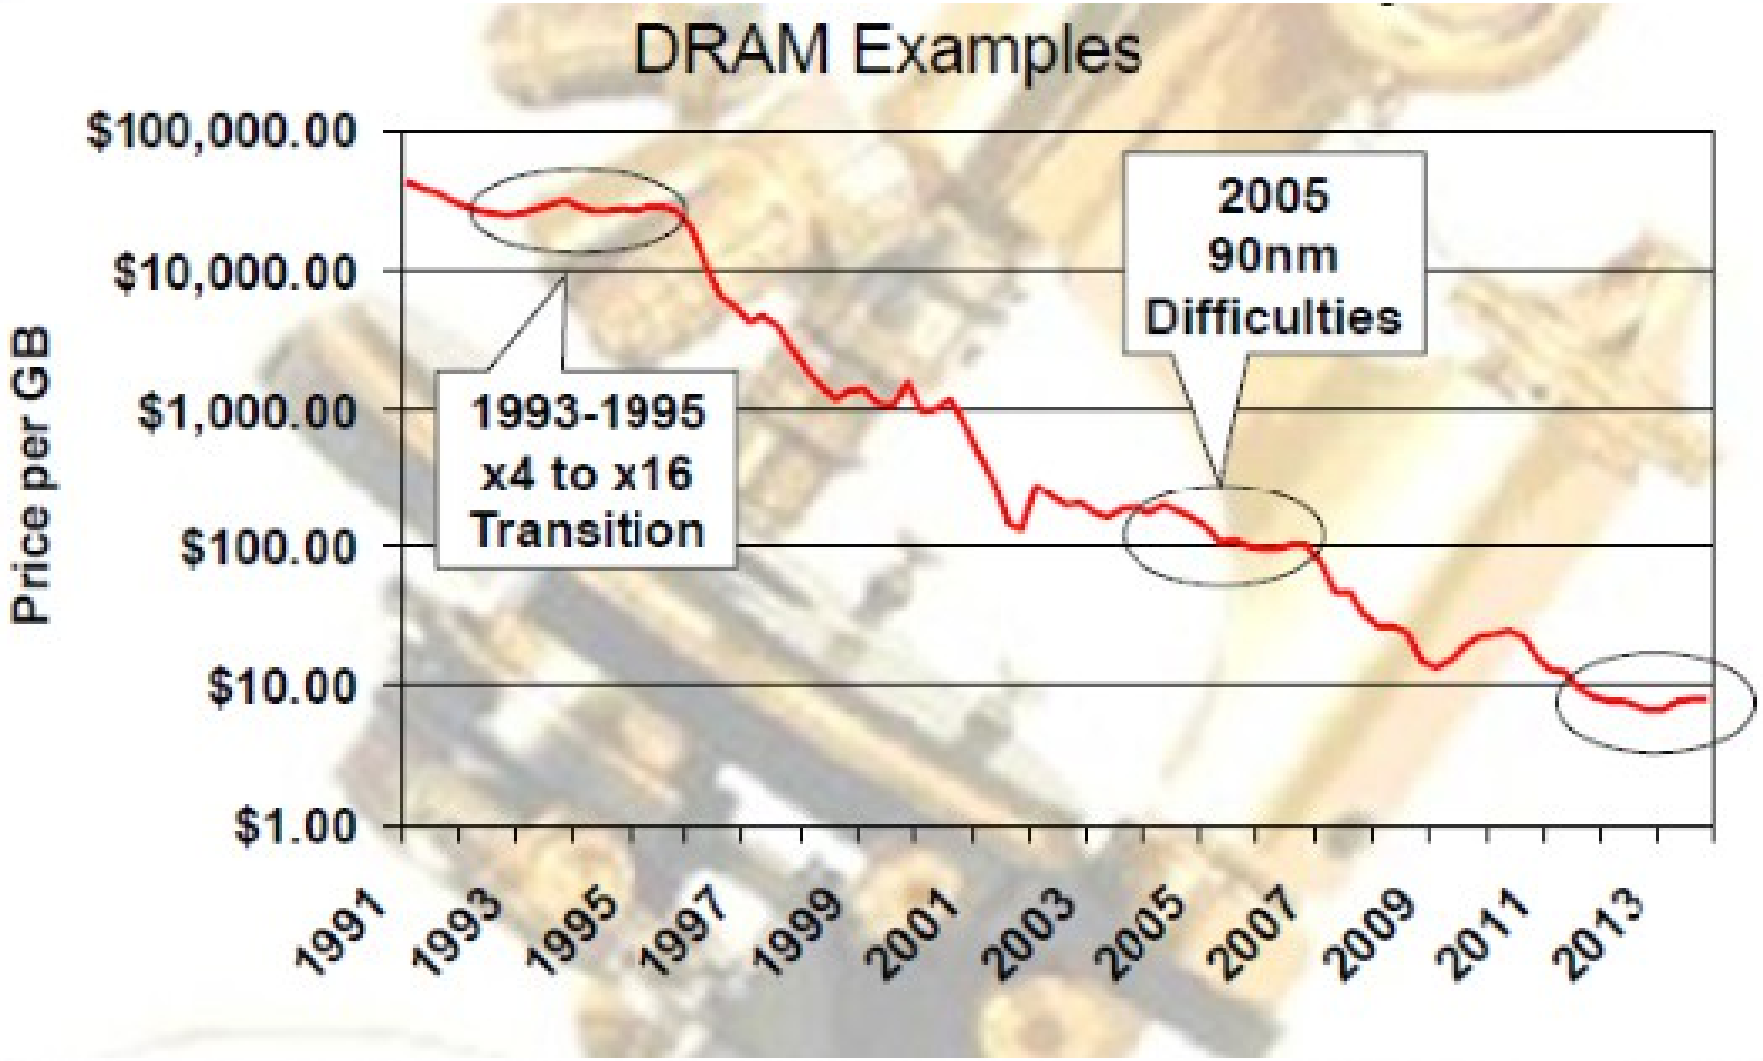

That is a factor of 1000 in twenty years. The flat spots are interesting . . .

Long term trends

- *Cost per Gigabyte* drops by 30 percent a year
- Now at about $0.33 per GB
- *Gigabytes produced* increase by 50 percent a year
- Q: When will flash pass magnetic disk ?
- A: Don't hold your breath.

Comparison w Disk

- Magnetic disk sectors can be re-used by writing over the contents. No preparation for re-use is required.
- Flash memory needs to be cleared to an all 1's state. Then new content can be written to it.
- The *clear* operation takes longer than either reading or writing.
- As a background task, released sectors are erased.

Comparison [more]

- Putting released blocks through an erase cycle is a form of garbage collection.
- The electronics that manages the flash is unaware of file system semantics that make blocks claimed or free.
- The *Fix* is for the OS to provide a list of blocks to be erased. The SATA TRIM command implements this. Rick Moen says kernels 3.8 and beyond do this right.

Comparison [more]

- Disk defragmentation was a good thing to do with your spinning disk. It improved speed by increasing the size of chunks of files that were stored contiguously.
- *Defrag* is unnecessary for SSDs because latency does not depend on location.
- Defrag would increase the rate of erase operations and shorten the life of the flash.

Formats

- SSDs can be packaged with electronics that makes them appear as SATA drives – but much faster because there is no rotational and seek latency.
- SSDs can also have SAS [Serial Attached SCSI] interfaces to get higher data rates.
- Neither of these match the speed of the devices. High performance SSDs connect directly to PCIe slots.

NVMe

- = **N**on **V**olitile **M**emory **e**xpress
- PCIe standard is not enough
- Need API with common command set to manage flash functions

Why are SSDs a good idea?

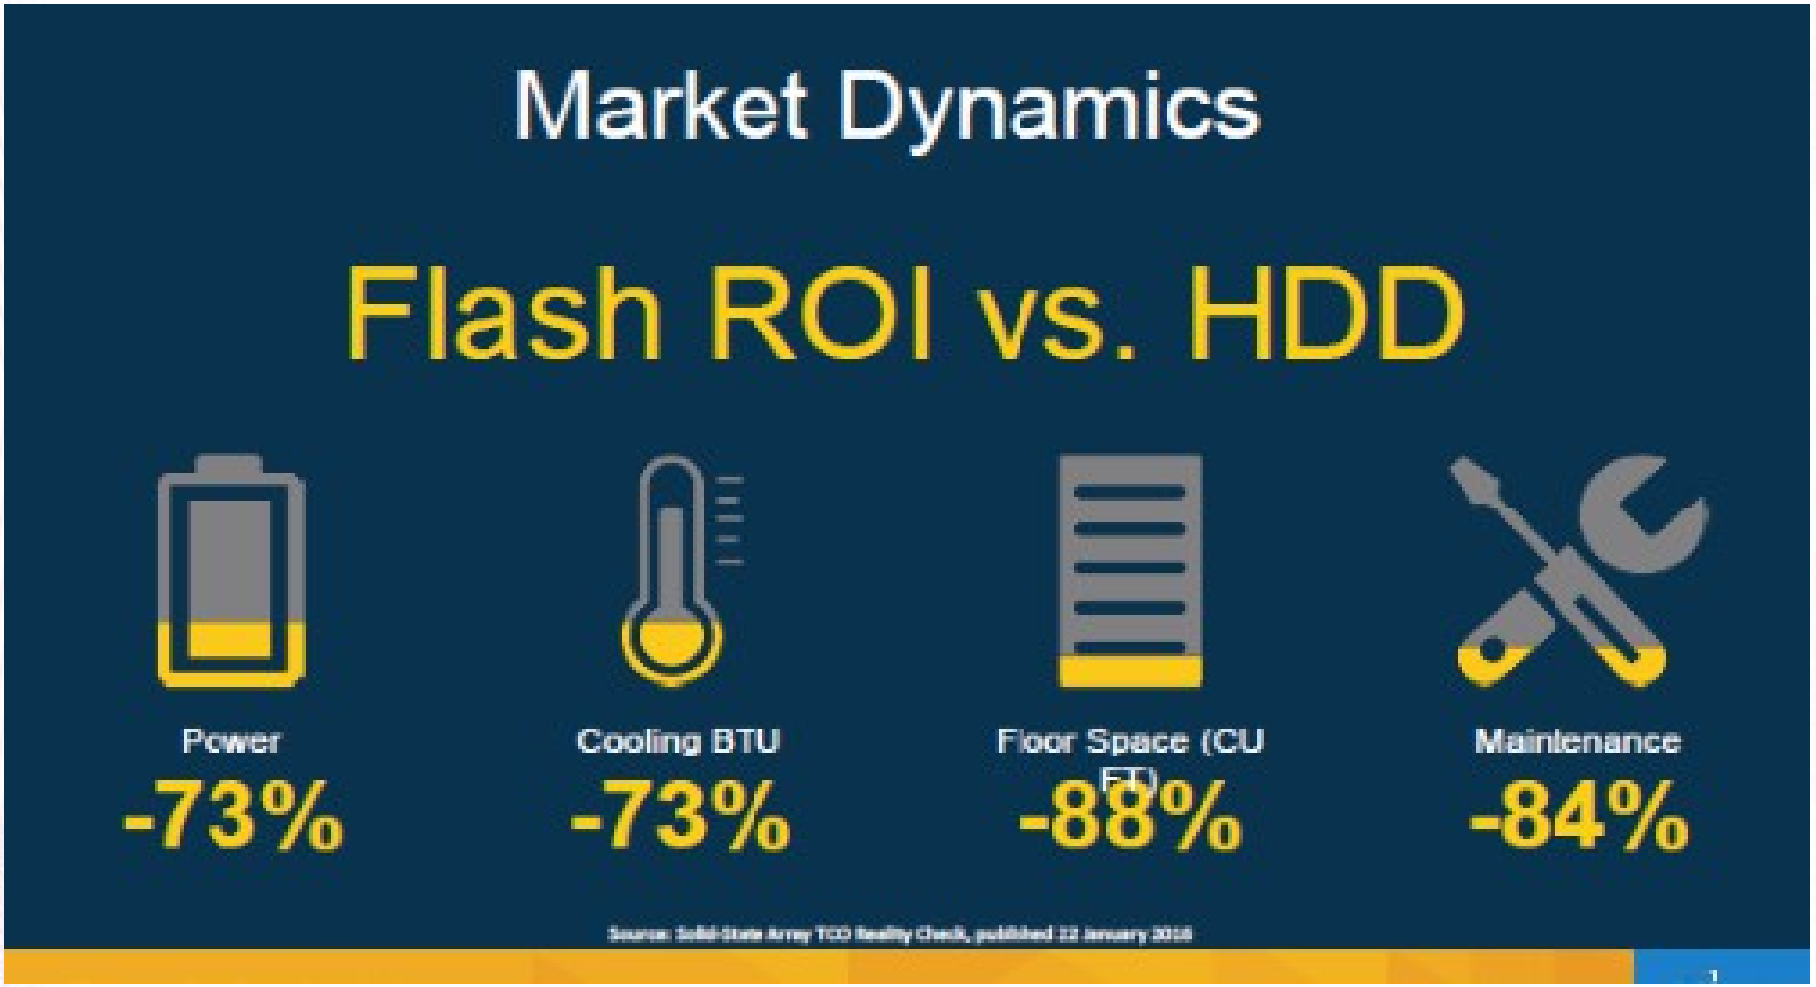

SSD is also 5x faster

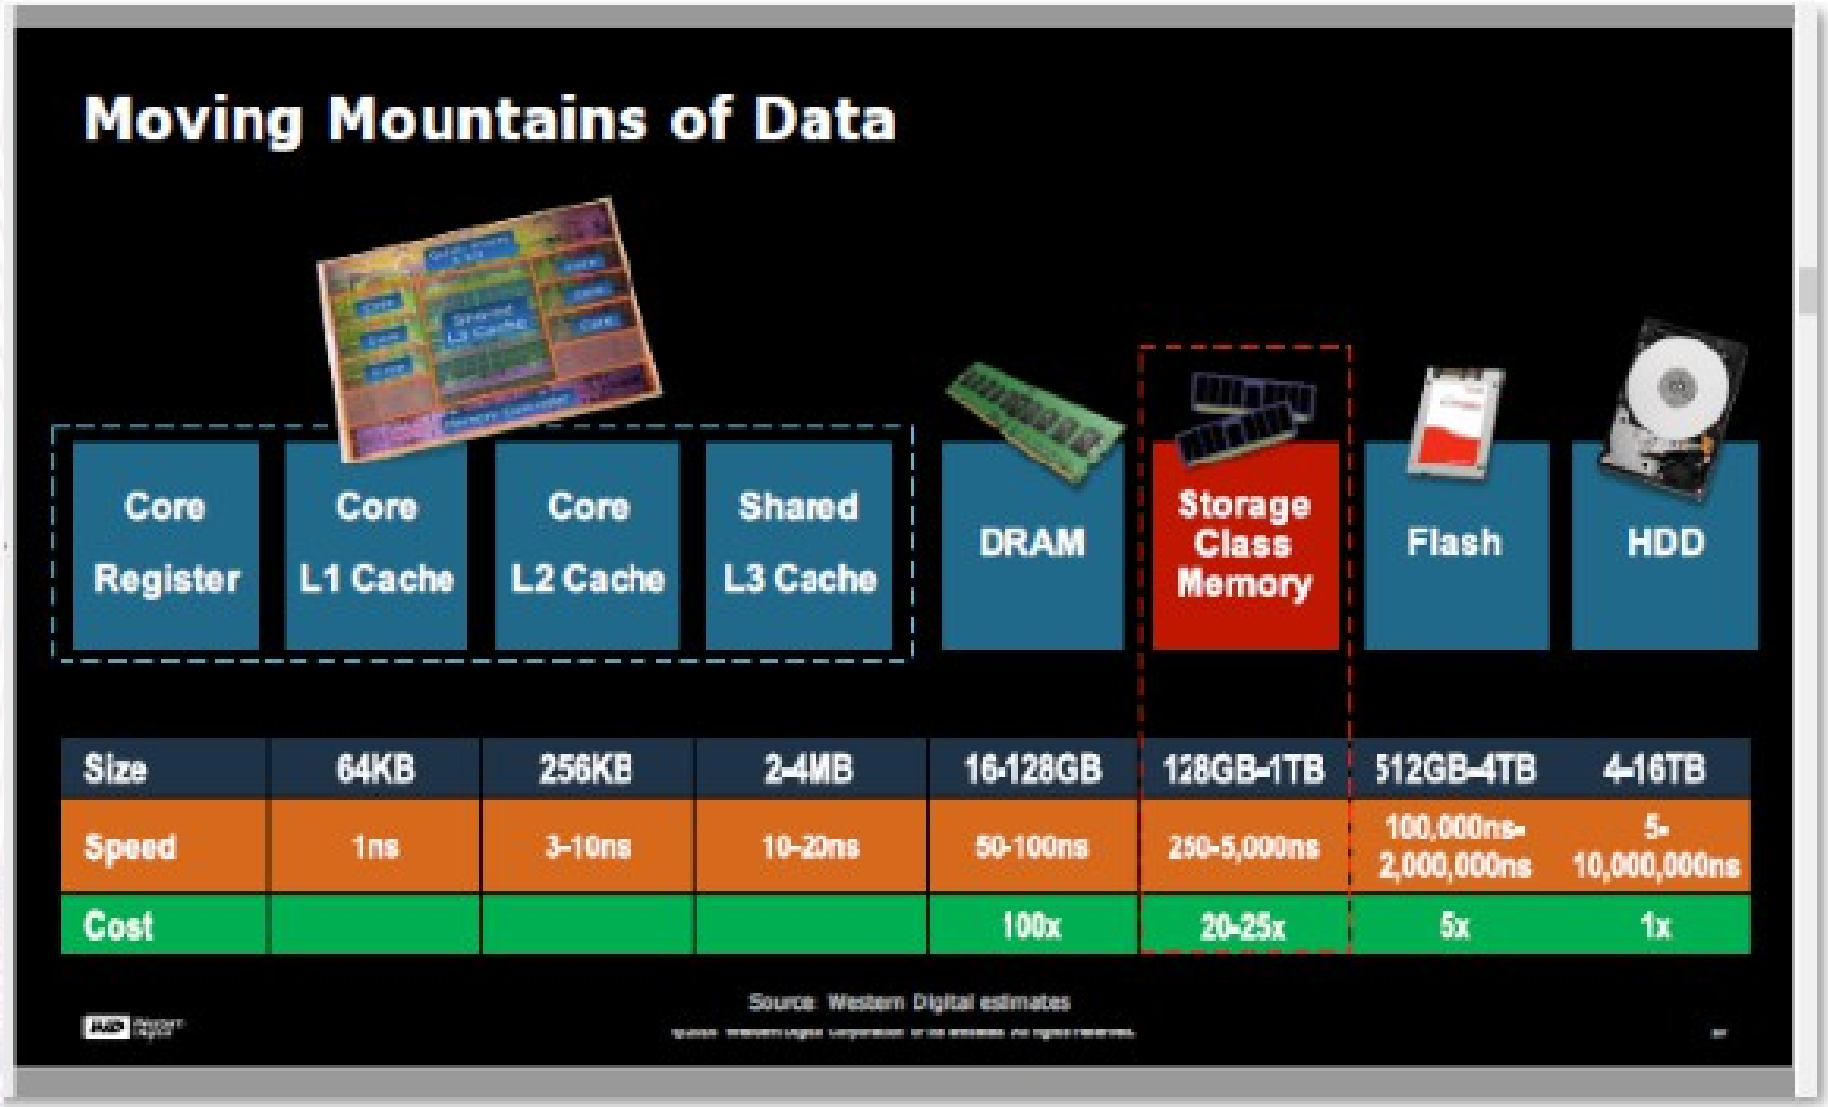

Rotational latency

- Spinning disk latency determined by rotational speed.
- The high performance end of the disk drive market – 10,000 RPM and 15,000 RPM
- These are dead products; this market now dominated by flash.
- Spinning disk has a cost-per-byte advantage using 7200 RPM drives.
- Laptops use 5400 RPM drives.

3D Flash

Next big thing will be 3D Flash. This is what it sounds like – layers of flash on the silicon substrate.

3D flash is available now, but it costs more to make than 2D flash.

This is [we hope] a temporary problem.

3D architecture

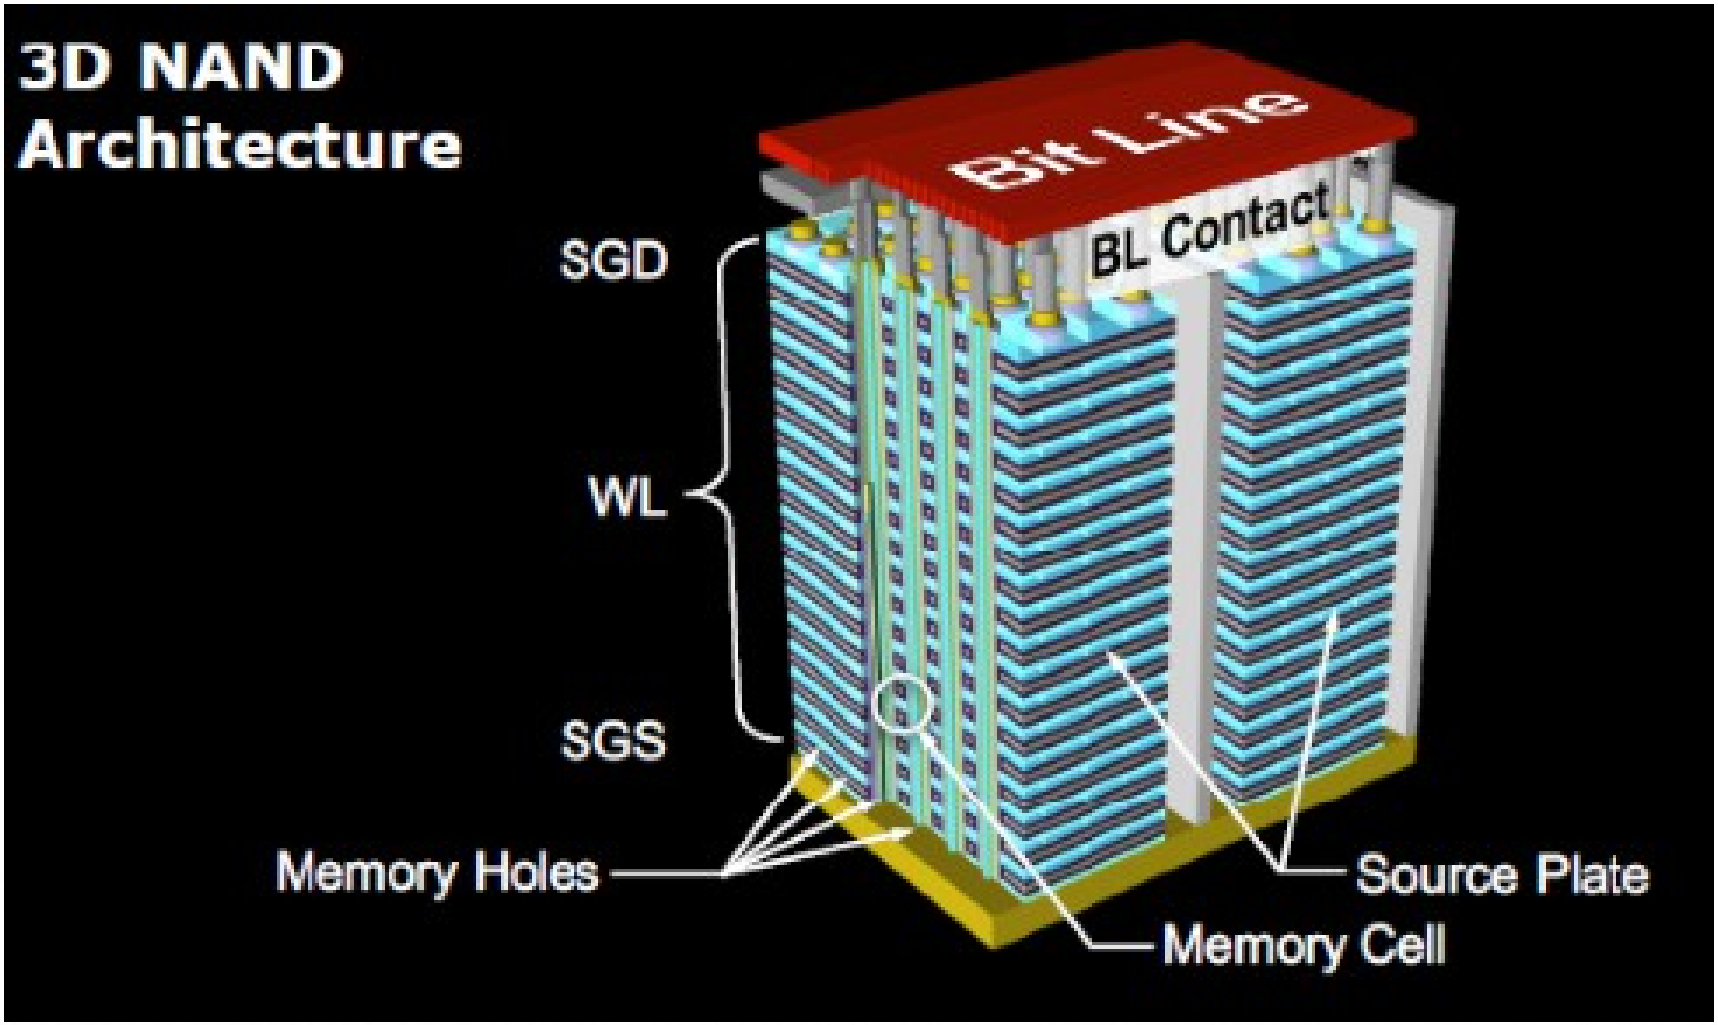

3D Trouble spots

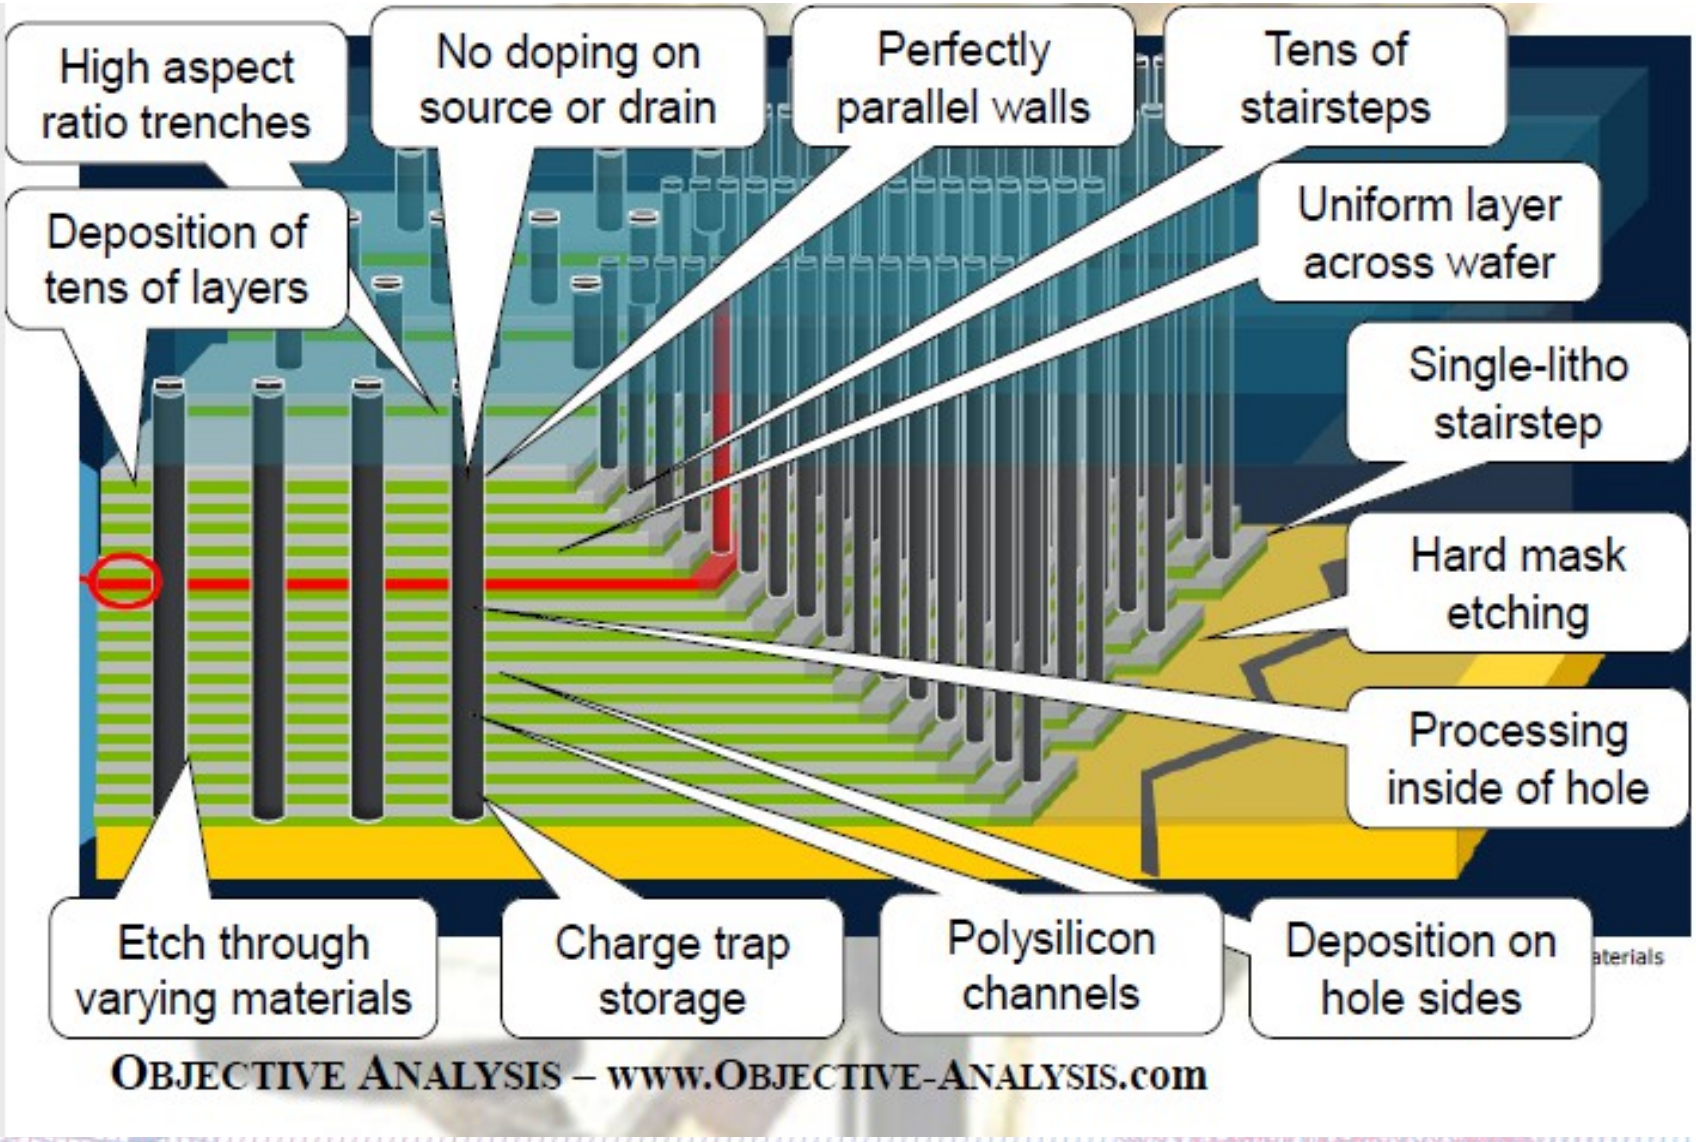

3D Predictions

- The industry always says *Next Year*
- Manufacturers are either producing product [Samsung] or making samples, so good signs.
- Jim Handy says 2018
- 3D will create a capacity glut
- Expect prices to collapse til demand catches up

What's next ??

- Storing bits as insulated charge is about 30 years old
- Bits can be remembered in phase state changes.
- Intel/Micron are sampling Xpoint memory in the 1-10 uS speed region. Not clear what *it* is.
- Intel needs this and may sell at a loss to get it.
- Otherwise, needs to be cheaper than DRAM

THANKS !


THE FLASH