Tapping Solid State Storage to Tame the Power-Hungry Data Center

Marius Tudor – Fusion-io

Power and ***Cooling*** are "a

Pandemic in the world of the Data Center" – VP Research at Gartner Inc.,

The Data Center Power Dilemma

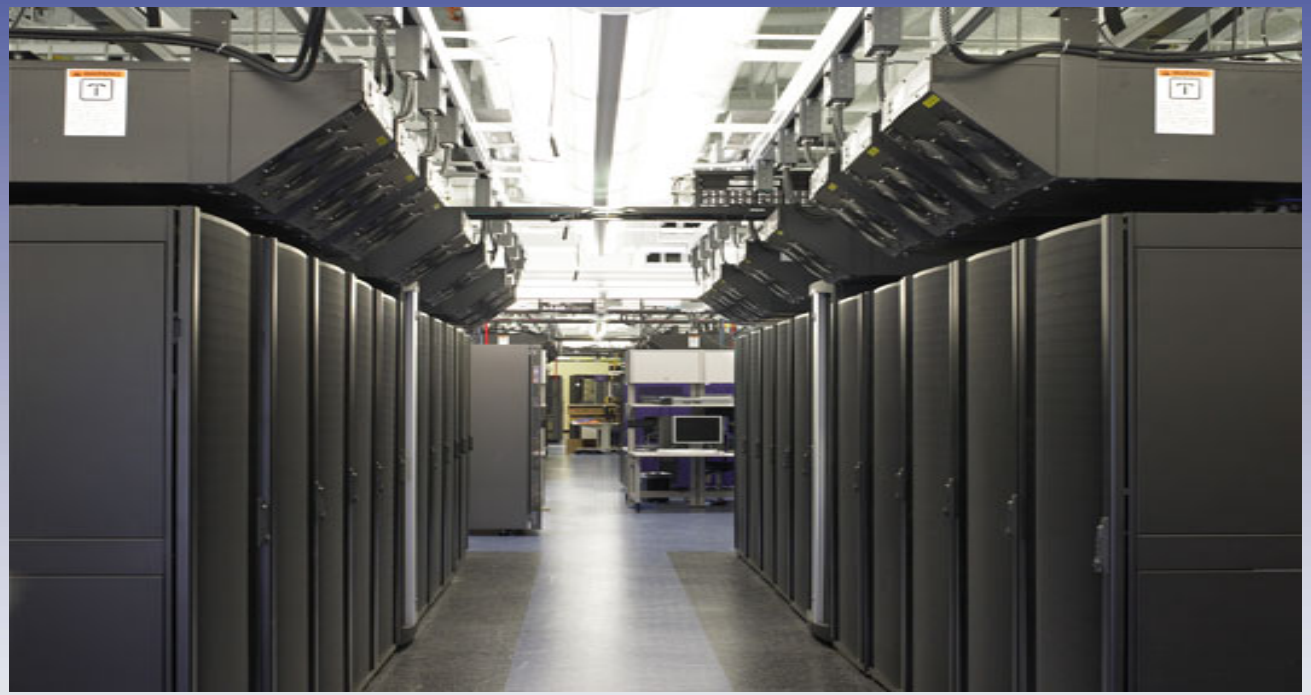

“70% of US CIOs believe that Data Center Power & Cooling is their No. 1 Challenge” – Gartner Group 2008

What is the magnitude of this 'Pandemic'?

- In 2006 Data Centers Consumed approx.

  61 billion kilowatt-hours (kWh)
  - By 2011 our power footprint may break the 100 billion kWh mark

OUCH!!!

- But wait… it gets Better… We will only need to build

10 additional power plants just to cool our Data Centers….Now that is Progress, right?!

What's the cost of our 1s & 0s staying cool??

Approx. **$7.4 billion**
with a 'B' Annually

*How much is **Cooling** alone out of that…?*

About half or…. A Way Cool $3.7 billion

So… How Did we Get Here???

Power vs. Performance… Unnecessary Tradeoff

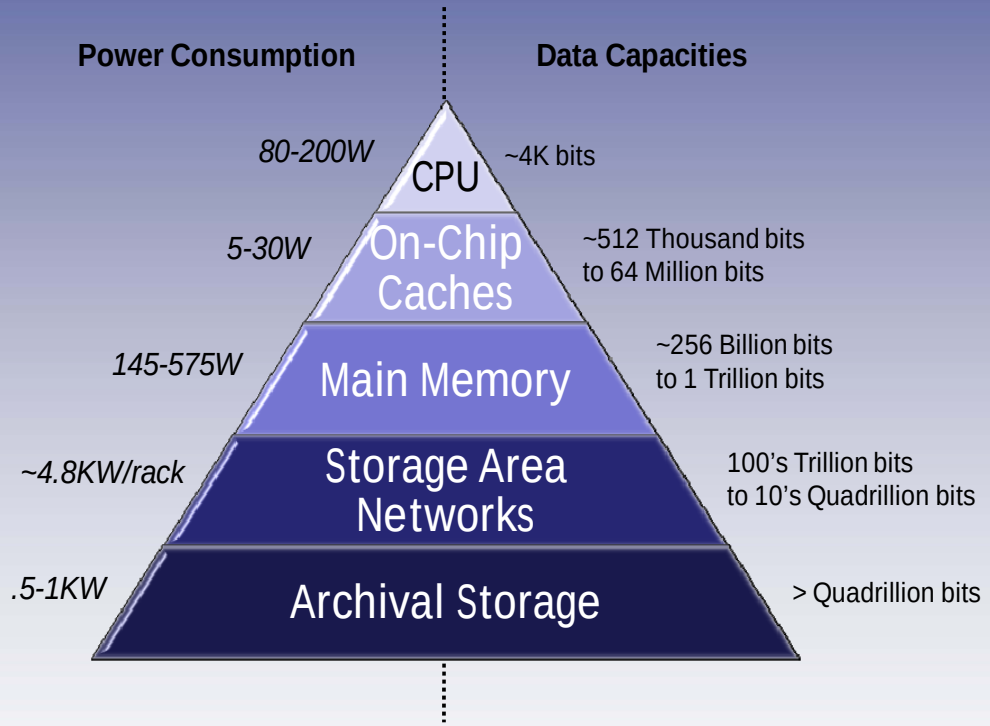

A long time ago in a not so distant universe of IT and smart people doing their best….
This type of performance disparity 'stalled' the CPU.
The need for faster monolithic processors has passed a while back…
The applications needed more performance from the storage, which threw everything but the kitchen sink at the CPU and still couldn't keep it busy…cycles where continually lost in space…
Going to dual, quad, octal core attempted to address some of the intrinsic challenges by bringing parallelism into play to solve the challenge, but it also exacerbated the weaknesses of the storage infrastructure
We added more & larger on-chip caches, more RAM, but this is costly, takes space and uses even MORE Power…Grrrr
Networked storage resolved some of the performance, sharing, failover protection, replication, but it threw even more tens of Millions of HDDs packaged neatly into 48U racks needing Power & Cooling in order to 'help' meet the demands of Millions of underutilized CPUs …

How do WE feel about some of these challenges???

Driving for Energy Reduction

2004 witnesses Intel's 4Ghz Pentium 4, fastest ever…even 5 yrs later. Why no progress since?

While the semiconductor universe dutifully followed Moore's Law and doubled processor performance every 18 months, the good ol' reliable spindle and the infrastructure built around it couldn't keep up. We'd have to have HDDs spinning faster than UFOs taking off in order to close the gap… *How fast is that?*

Too fast even for the wishful thinking 20k rpm HDDs…

Two Decades of Processor-Storage Widening Gap brings us to…Today

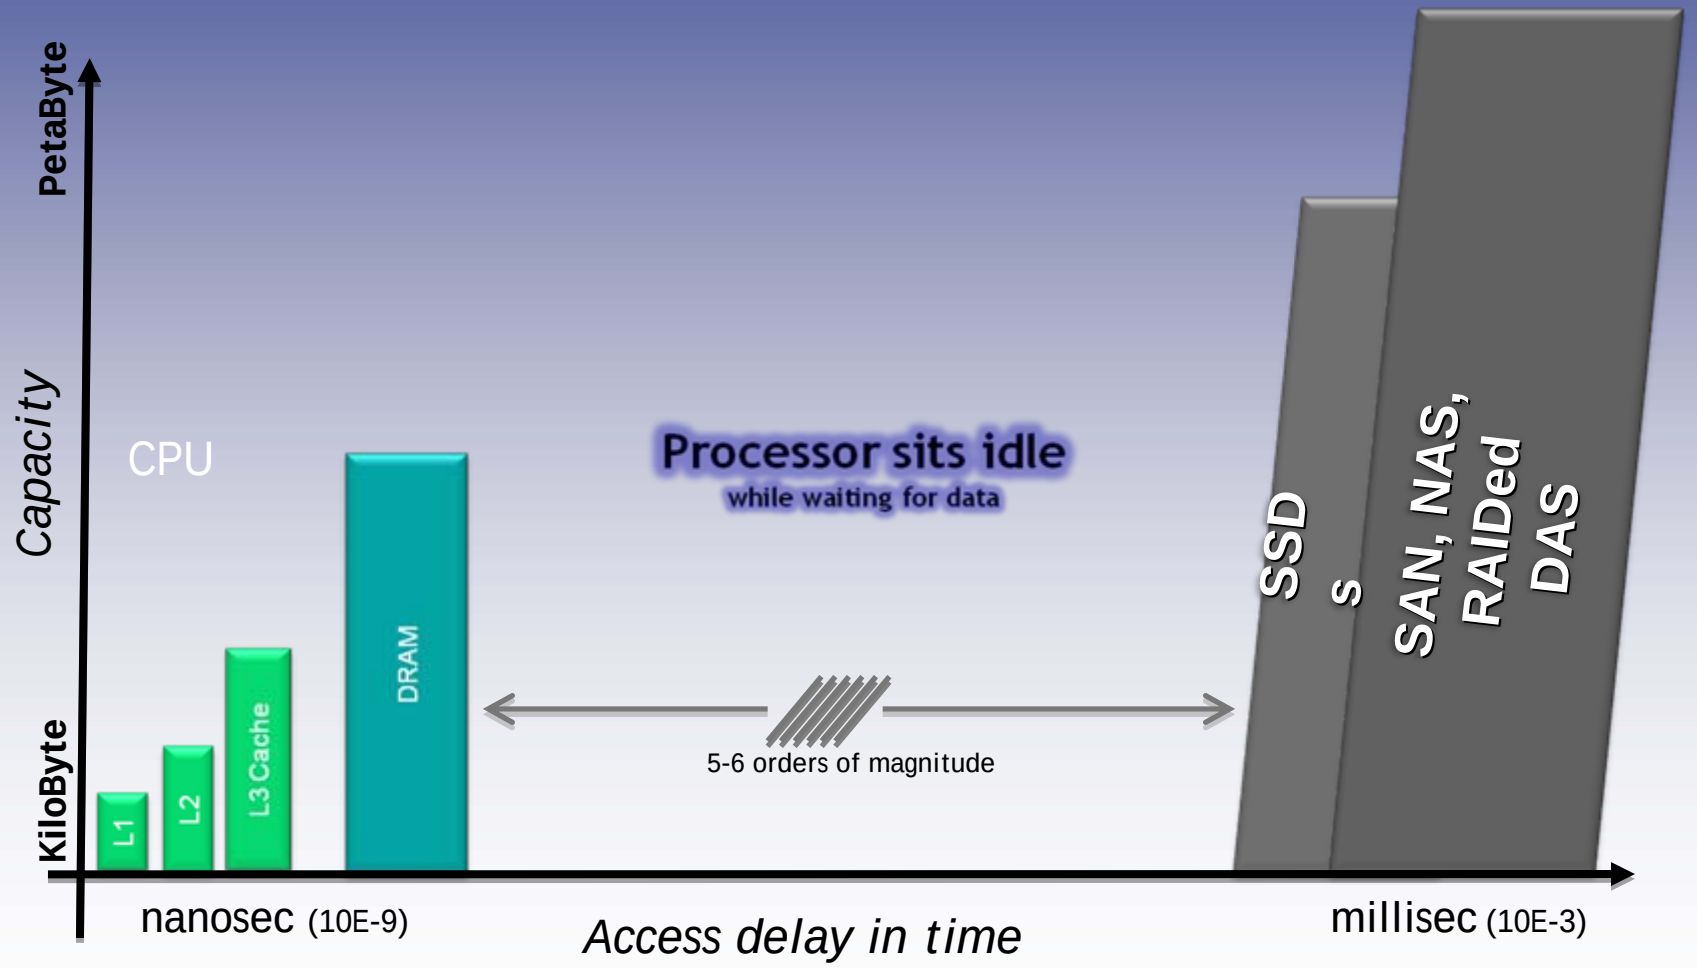

Classic Approach to Overcoming the Gap

- Aggregate more hard drives
  - Costly, wastes capacity, forces shared performance
  - Increases access time but adds performance
- Add DRAM memory
  - Costly, wastes power, but avoids having to go to disk
- Deploy additional Servers
  - Cheaper way to get lots of memory or lots of disks
  - Wastes floor space and power, and has high failure rates
  - Processors sit idle, under-utilized
- Hire (expensive) Engineering talent
  - Spend years to build & optimize the application to cope
  - Applications become inflexible, unable to be adapted

Solid State Storage Bridges the Gap

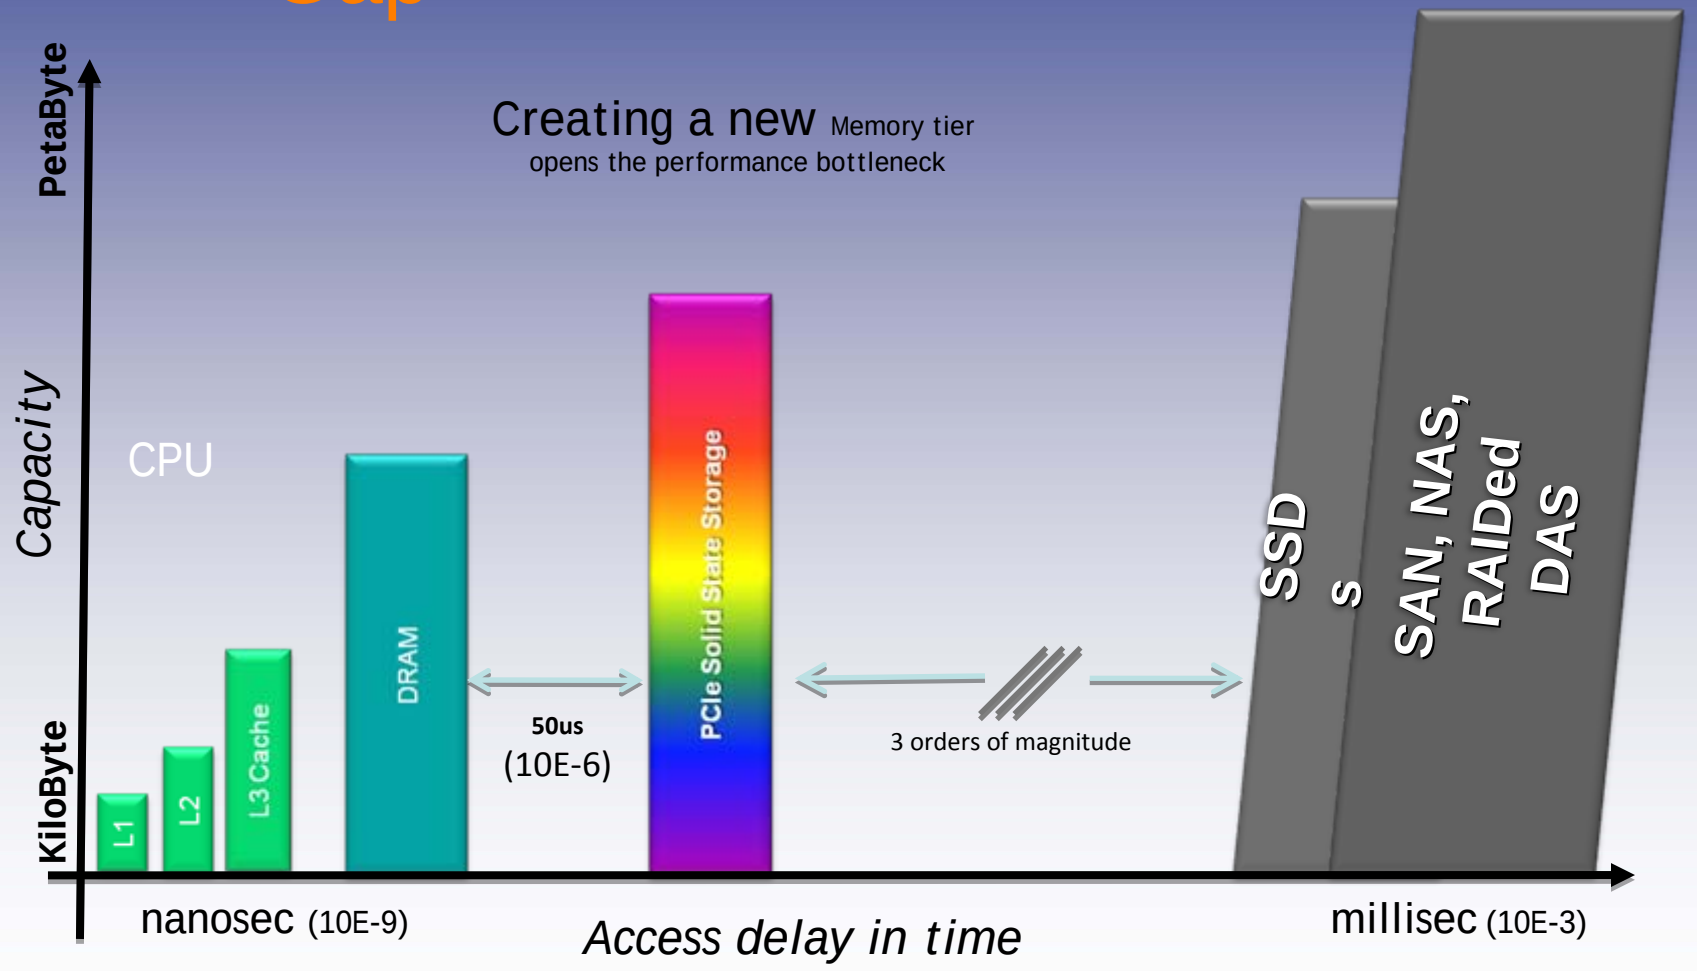

Case-in-Point: MySQL

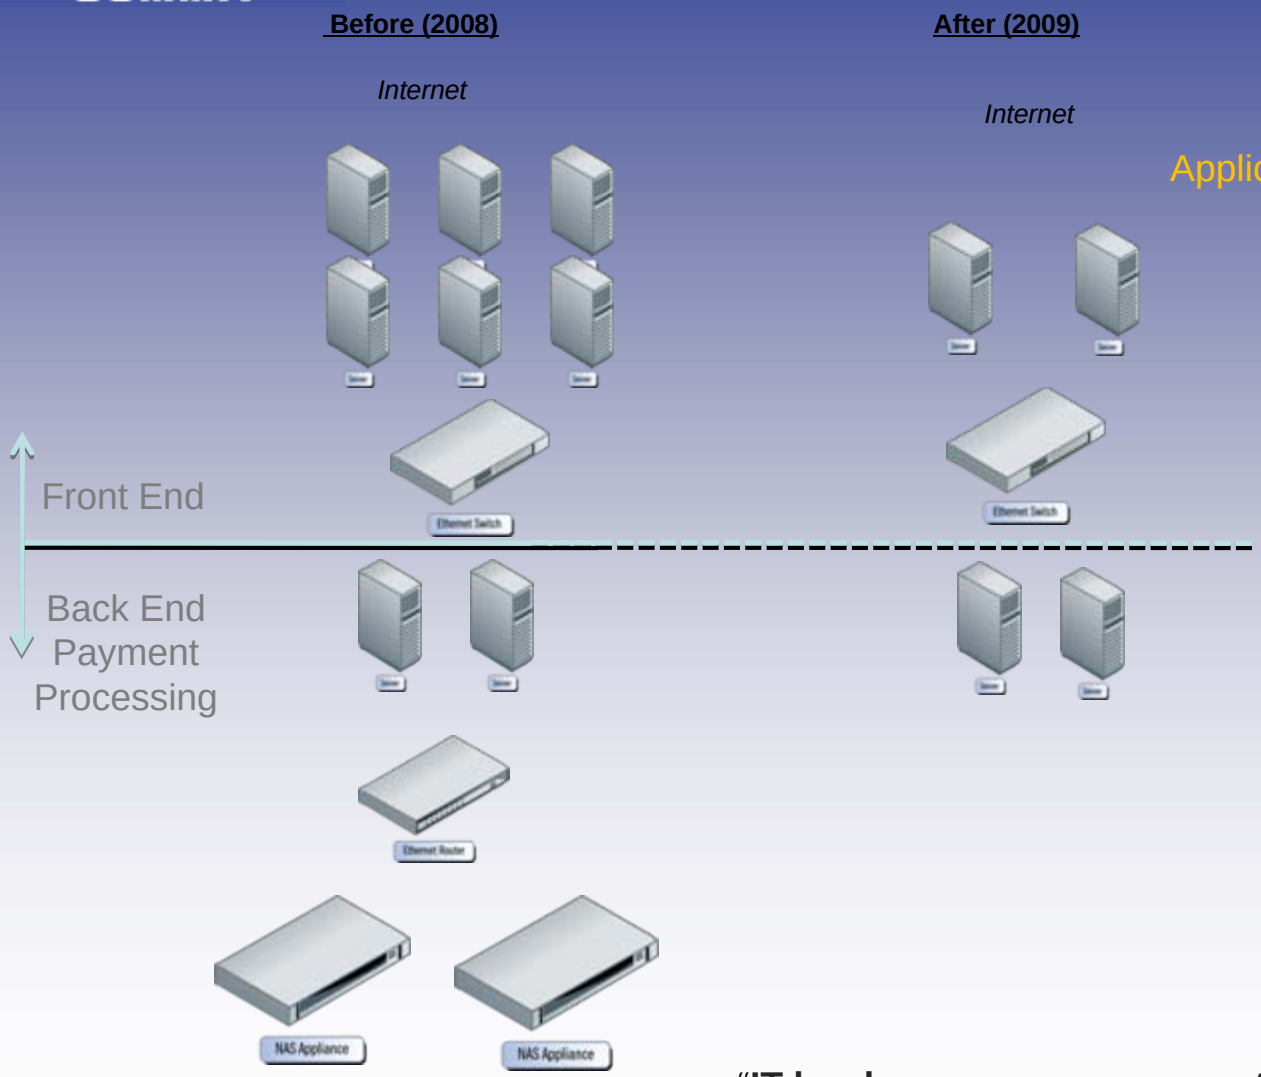

Application level replication for redundancy

Power & Cost Reductions

50% OpEx reduction
- 100 HDDs eliminated
- 3 CPUs turned off

62% reduction in Main Memory
33% Footprint reduction

Performance 'After' with PCIe SSS

12x improvement on write
- Latency down from 4 to <1ms

14x improvement on read
- Latency down from 12 to <1ms

"IT hardware power consumption was reduced by more than 90% for equivalent application performance."

Annual Power for 100,000 IOPS

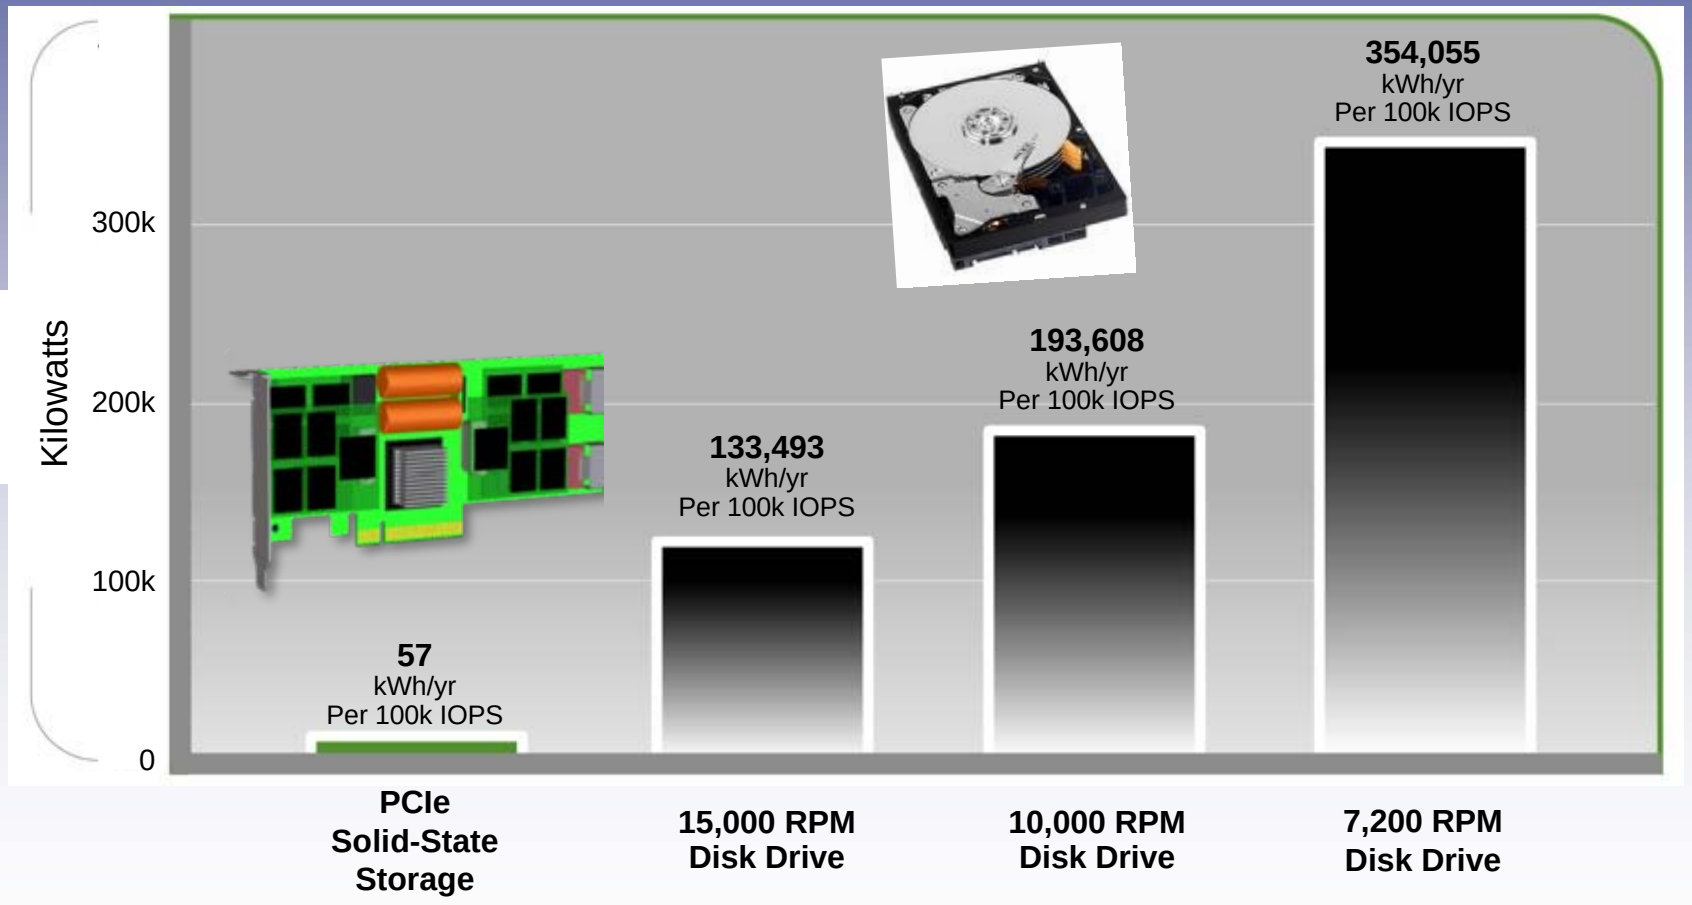

CapEx and OpEx Shrunk. Why would I integrate expensive SSS?

- Make the CIO smile: Think Green thru the prism of savings
  - The storage reqs still need met.

    They are a Necessity
  - Run TCO analysis to clearly see

    the SSS derived savings: you're avoiding Additional racks of HDDs, Power, Space, Licenses

- **Bottom-line**: $$$ Saved are $$$ Earned

End-Result: Power savings with enhanced application performance

- **No need for…** huge arrays of power-hungry paralleled disk farms in order to reduce the CPU wait time
- **No need for…** power-hungry 15k RPM HDDs since the speed of the disks would no longer be the limiting factor
- **No need for…** HDD short-stroking to improve system performance wile using 10% of the disk space, thus further increasing the number of power-hungry disks needed to handle the required capacity at the needed performance
- **No need for…** large, power-hungry main memory DRAM since the performance penalty of not finding the data in the system cache would no longer be a disaster, as we have it stored in SSS instead of HDDs
- **No need for…** large battery backup systems. Simply protect the CPU registers, caches and memory. All of these can be saved in Flash SSS and reloaded quickly when power is restored.

Thank you!