




Lies, Darn Lies and a **Zillion** IOPs

Swapna Yasarapu

*Director, SSD Product Marketing
STEC, Inc.*

IOPS are important but...

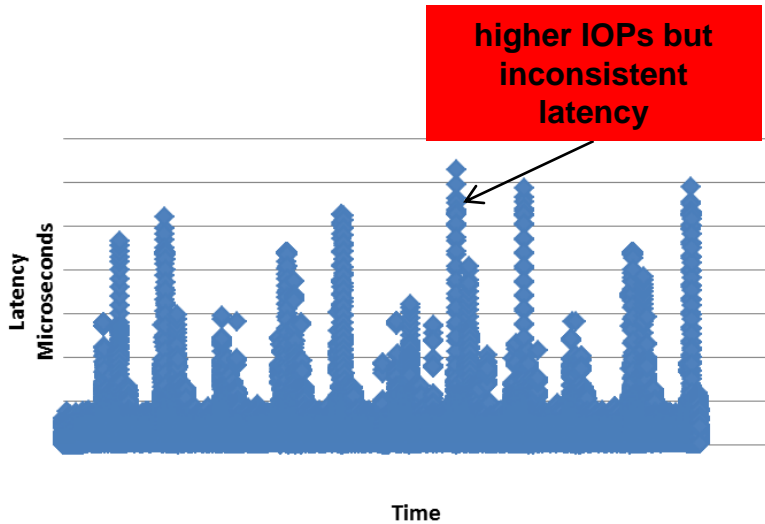
Speeding up applications



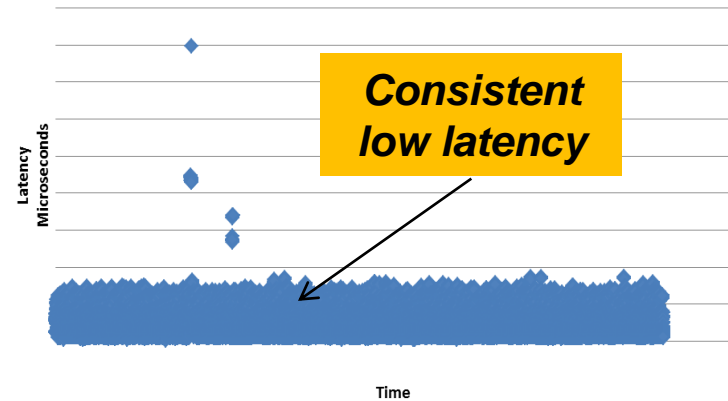
It's all about LATENCY!!

IOPs vs. Latency

Applications prefer consistent latency vs. IOPs



Slows down application performance



Preferred PCIe Solution

Why?

- 1 Dependent IO
- 2 Operating Range
- 3 Consistent Latency over service life


The case of “Dependent IO”

Benchmarks have **Independent IO**

Applications have **Dependent IO**

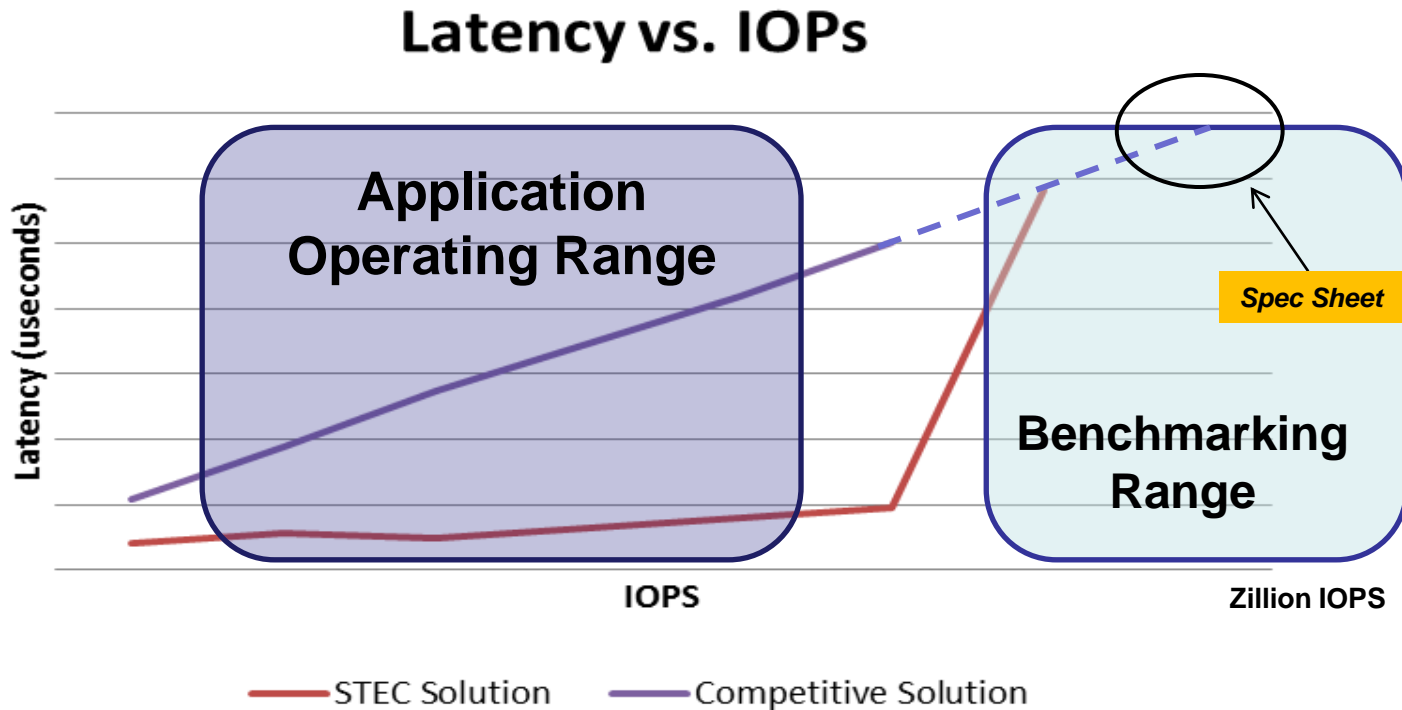
An Example of a RAID system

```
For (i=1..32)
{
  write (i-1) ←
  ddkjdf;
  ddjldfj;
  read (i); ←
}
```



All reads and writes need to finish before application can move ahead


The Operating Range





Benchmarks test at Max 'Everything'
(duty cycle, IOPS, workload)

These measurement do not reflect the applications operating range

PCIe vs. SAS Latencies



← Extra data hops = Additional Latency →



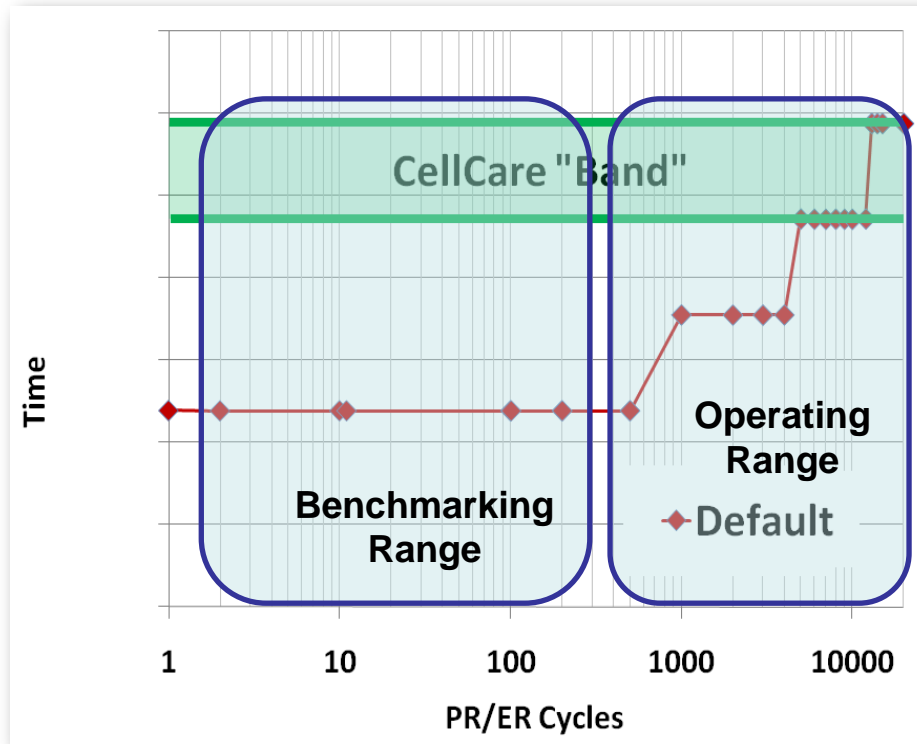
30-50us of latency improvement with PCIe by eliminating SAS/SATA protocols

Consistent Latency over service life

Just the Facts;

1. Flash reads are fast; Flash writes are slow
2. Background tasks to manage flash increase as device ages
3. Device latencies change as it ages, unless

Requires flash to be managed



STEC CellCare™ Manages flash ; maintains consistent performance over device life

Case in Point: Cloud Deployments

- Business Model:
 - Loading maximum number of users at a specific service level
 - More users per server node = saving in power, cooling.

Service level = Fast Response = Latency


Amazon loses 1% of sale for every 100ms it takes for the site to load

Data Center Saving* via 10:1 server consolidation ~ \$80,000/year

* Includes server cost + Power cost per year

What matters in the datacenter today?

- 1 Dependent IO
- 2 Operating Range
- 3 Consistent Latency over device life



***Solid State Devices
that deliver on these design parameters
are the winning solution***

It is all about the needs of the application

In other words...

- Optimized solutions for ***Dependent IO***
- ***Work in the applications Operating Range***
- ***with Consistent Latency over Device Life***

THANK
YOU

www.stec-inc.com