



MLC Flash for Enterprise Storage Performance

Erik de la Iglesia

Chief Architect

GridIron Systems

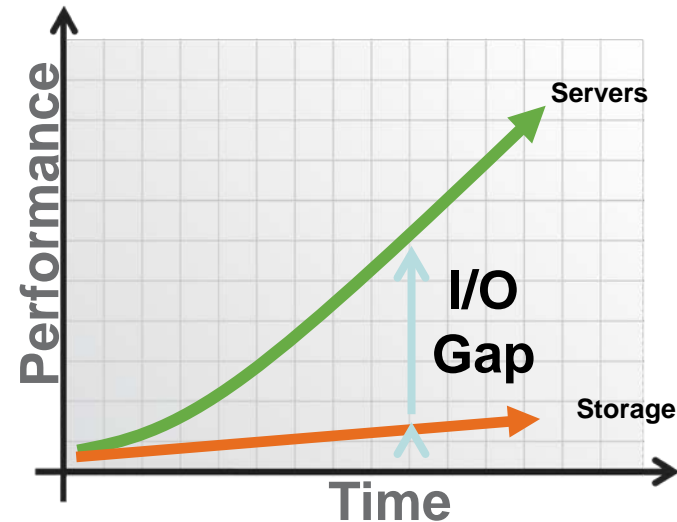
Flash Adoption in the Enterprise

■ I/O Gap Driving Adoption

- Accelerating business apps has high ROI
- Cost per IOP is compelling (vs disk)
- Multi-core servers rapidly outclassing storage

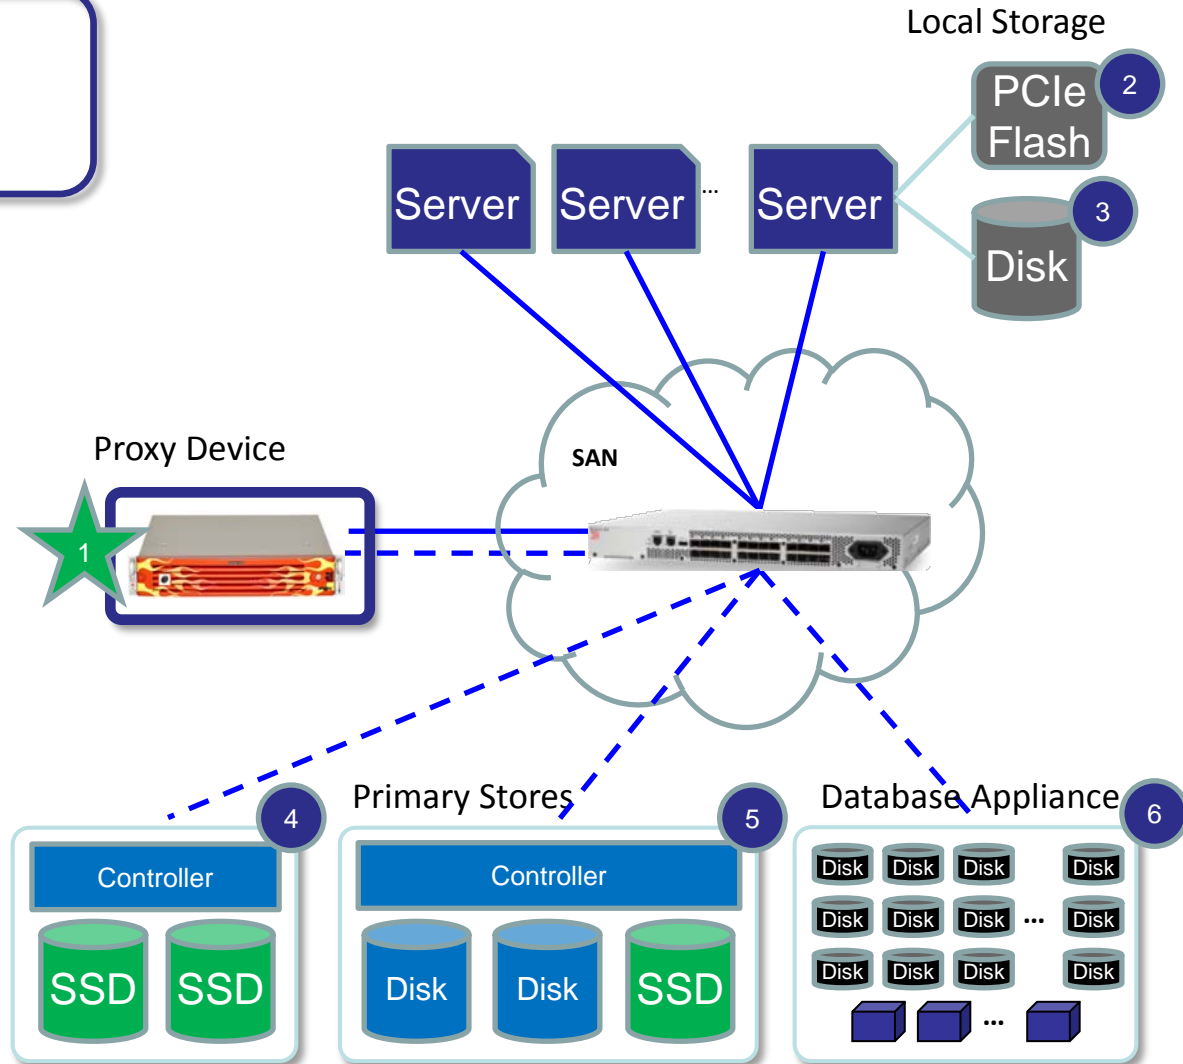
■ Challenges abound

- Integration or forklift
- Technology readiness
- Cost effectiveness

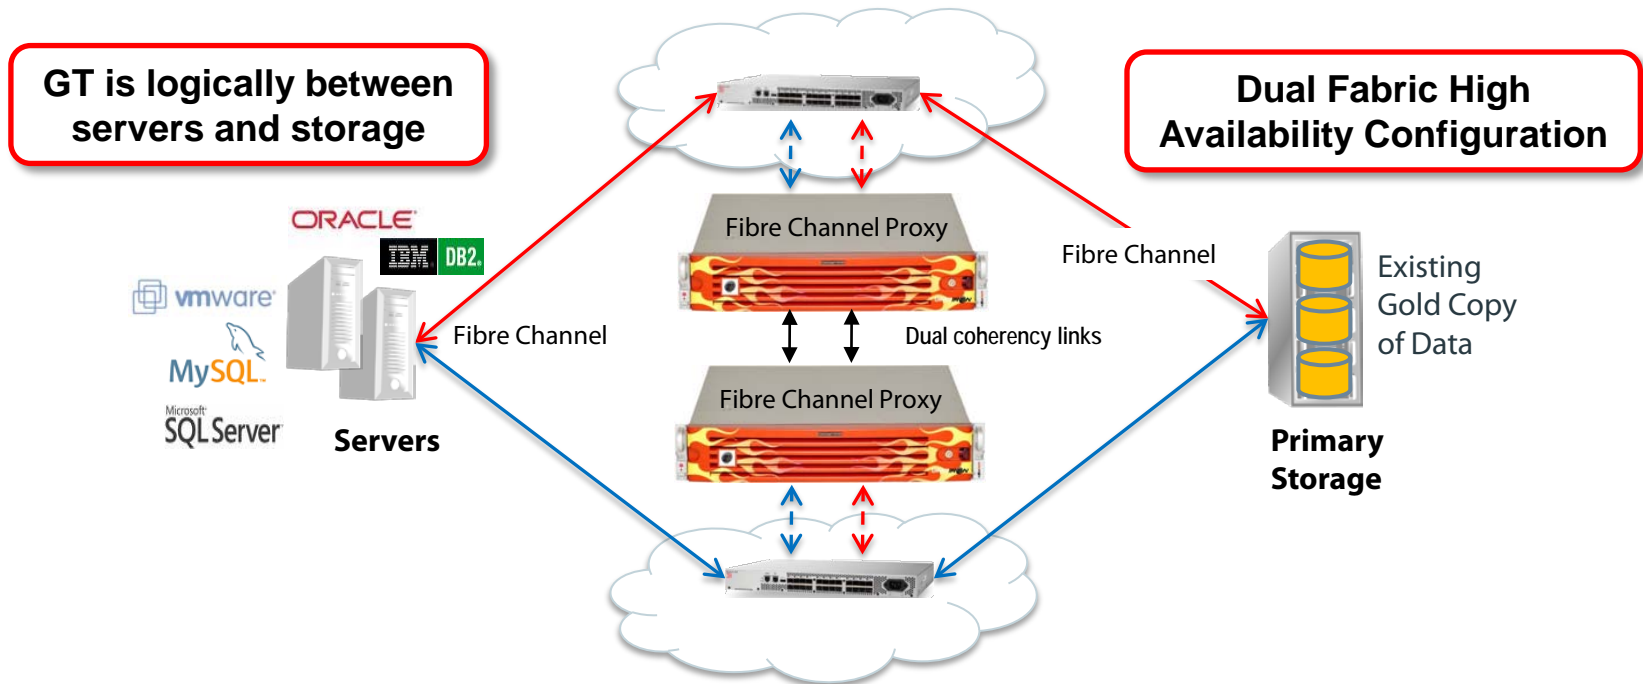


Many Deployment Options

1. **SAN Proxy Device**
Benefits all servers, storage and applications
2. **Flash PCIe Cards**
Benefits only a single server, no virtualization
3. **Local Flash SSD**
Data security, Degraded write performance over time
4. **Flash Primary Storage**
Applies only to that storage, Wear issue for heavy writes
5. **Flash Disk Replacements**
Controllers constrain SSD performance
6. **Specialized Processors**



Deploy Transparently in the SAN



Transparently accelerate data access in the SAN

Solid State Performance with no change to:
Software - Databases - Servers - Storage - Processes

High Bandwidth Real-time Tiering

GridIron provides

Higher concurrent IO bandwidth
Higher IOPS
Low latency multi-level cache


The GridIron Process

Learn data access graphs in real time
Use patterns to manage caching
Use feedback to continuously refine performance

*Cache only
performance
enhancing data*

GridIron IP






HW Analytics
FC Proxy
Billion entry MMU
Flash Mgmt.



Results

- Applications 2-10x faster
- Read latency 10-100x shorter
- Flash life a non-issue

Real-World Deployments: High-Impact and Cost-Effective

Customer	Application / Business Problem	FastData	GridIron Performance Improvement	CapEx Savings With GridIron
	Oracle Data Warehouse <ul style="list-style-type: none"> • “Real time” reports taking 6 hours • Lost revenue from delays • Over-provisioning storage for performance 	<ul style="list-style-type: none"> • Bandwidth: >10GB/s • Concurrency: 25+ users • Data set size: 40TB DWH • Data turnover: continuous ETL 	<ul style="list-style-type: none"> • Critical reports 6 hrs -> 30 mins. 	<ul style="list-style-type: none"> • \$2M from storage and server consolidation
	Large eDiscovery - MS SQL under VMware <ul style="list-style-type: none"> • Multi-hour query times affecting productivity • Need to support concurrent users • Serialized system impacting business 	<ul style="list-style-type: none"> • Bandwidth: >2 GB/s • Concurrency: 4 users 	<ul style="list-style-type: none"> • Query Times reduced by >50% • Increased query capacity by 6x 	<ul style="list-style-type: none"> • Saved \$775k on a storage upgrade (only 2x)
 	Video game software builds under VMware <ul style="list-style-type: none"> • Builds taking 70 minutes to complete • Game quality impacted by long build time • Virtualized architecture not scalable 	<ul style="list-style-type: none"> • IOPS: 40,000 Random • Concurrency: 24 users with parallel builds 	<ul style="list-style-type: none"> • Build time reduced from 70 -> 8 mins. 	<ul style="list-style-type: none"> • \$800K vs. alternatives
	Post production video editing <ul style="list-style-type: none"> • Concurrent 2K streams needed • Support more users at lower cost • Complete system replacement is prohibitive 	<ul style="list-style-type: none"> • Bandwidth: >3 GB/s • Concurrency: Five 2K video editing streams 	<ul style="list-style-type: none"> • 5x improvement in concurrent 2K video streams 	<ul style="list-style-type: none"> • Avoided a complete system replacement

Handling Trends in SSDs

- **SLC / eMLC / heMLC / MLC / 3bMLC / ???**
 - Competing goals of endurance and economics
 - Lithography complicates choices
 - SSDs not the volume drivers for flash silicon

- **SSD flavors and toppings**
 - Interfaces, channels, controller horsepower
 - Smart wear reduction
 - Specified lifespan

Considerations for Evaluating SSDs

- **Initial performance will not (and should not) last**
 - Evaluate your workload over time
 - Must utilize intended capacity, alignments, operation sizes, etc.
 - More than 2X or continued degradation is suspect

- **Being “too smart” is not always good**
 - Compression and de-duplication are great – **for IOMeter**
 - Compression can cause high and uneven churn for indirection tables
 - Any background process in device can take away from IO processing

- **Consistency and predictability are key**
 - Architecture can plan around known performance corners
 - Characterize device over specific data rates and thread depth
 - Allows smart monitoring to detect deviations early

- **Who makes the drive, controller and flash chips**
 - Manufacturing, testing capability and history
 - DPM levels and FA process