



Enabling SSD and Flash for Virtualized
Environments

Vikram Joshi, Co-founder and CTO



Technical Vision

“Eliminate memory – I/O contention in virtualized environments.”



Technology Trends

- Number of CPU cores is going up
 - Need more VMs to use cores
 - I/O demands of a VM are modest, but in aggregate big
- Gap widening between CPU and disk speeds
- Flash – new point in memory hierarchy



The Trends Say...

*We must think of
storage I/O
differently.*



Problem

- Virtualization exacerbates the memory & I/O issues
 - More VMs means more I/O
 - Higher application latencies, lower throughput
- Sub-optimal CPU utilization due to I/O wait times
- Virtualization drives a new requirement – shared storage



Current Solutions

- Costly and Complex
 - Add more high-end storage (\$\$\$\$\$)
 - More HBAs, FC networks/switches (\$\$\$)
 - Add RAM (\$\$ & watts)
 - Front primary storage with a cache (\$\$\$)



Technology for a Better Solution

- Separate IOPS from disk-based storage without decreasing reliability
 - Continue to use primary storage for *data management*: Capacity & Reliability
 - Use Flash on compute server for *IOPS*: Performance



Technology for a Better Solution *(cont.)*

- Use Flash cache on compute server
 - Increase application performance by reducing I/O latency
 - Increase CPU utilization; run more VMs
 - Cache OS and application paging data
 - Cache I/O to persistent storage
 - Prevent IOPS surge caused by VMs
 - Transparent to applications and existing infrastructure
 - Dynamic allocation of Flash capacity as VMs come and go
 - Support hypervisor functionality such as vMotion
 - Extensible to virtualized and bare-metal environments



Why Now

- Historical evolution of Flash usage
 - Caching appliances
 - Storage arrays fronted with Flash
 - Flash/SSD as replacement to disks
- Recent trends driving this technology
 - Systems vendors are including Flash
 - Lower \$/GB and watts/GB than memory
 - Lower \$/IOPS than primary storage
 - Multiple Flash vendors have enterprise grade products
 - PCI card and SSD form factor options



Opportunity

- Virtualization is here to stay, but there are challenges
 - Many applications not ready to embrace virtualization given the current trade-offs
 - More CPUs, more data, but storage IOPS lag
- Incredible opportunity to:
 - Virtualize more applications
 - Accelerate VDI adoption
- 2010 is the inflection point with more virtualized servers deployed than physical servers. This trend will only continue.

–IDC, Server Virtualization Shipment Forecast



IOT Technology - Unfair Advantage

- Ability to direct the I/O performance to specific VMs, files, volumes, or disks
 - IOT is in the memory – I/O path
 - IOT is close to the application's data
 - Intercept all reasons to go to storage
 - OS paging, application paging, and persistence
- Semantic/contextual knowledge of data
 - Cache or bypass cache selectively for files, volumes, and disks
 - Shared file-level deduped cache
 - Compression on specific VMs, files, volumes, or disks
 - Enables application-specific performance profiles
- Unified solution for both server & desktop virtualization



IOT Technology - Unfair Advantage (cont.)

- Downstream caches are inclusive, including:
 - Hypervisor, caching appliances, storage
- IOT works directly on application data
 - Intelligent decisions made in the guest OS
 - For example:*
 - Cache `pagefile.sys` for Windows
 - Cache only one copy of a mirrored volume
 - Cache only one `kernel32.dll` in the compute server
 - Do not cache database redo
- Agnostic
 - SSD / Flash
 - Storage: SAN, NAS, & DAS
 - Hypervisor
 - Transparent to Guest OS / Applications



Summary

Virtualization is here to stay

IO Turbine

- Enables SSD / Flash to participate in virtualization market
- Enables directing I/O performance to specific VMs, files, volumes, disk
- Semantic / contextual knowledge enables application specific performance profiles



About our Speaker

- Systems software engineer spanning multiple CS disciplines: OS, parallel and distributed systems, databases, storage, media, and computer graphics
- Oracle: Doubled database TPCC perf. on 64-way SMPs; formed Exadata group & led initial effort
- SGI: Developed video on demand server; led Shared Worlds video game server effort at CosmoSoft
- Sun: Parallelized the Solaris virtual memory system to demonstrate linear scalability from 1-12 processors; worked on machine bringup; Spring microkernel (SunLabs)
- Founded PixBlitz Studios in 2006: virtual advertising technology
- MS (Hons.) Physics 1984, BE Engineering in 1984, Birla Institute of Technology and Science, Pilani, India

For More Information

www.ioturbine.com

info@ioturbine.com

408-246-9922