Storage Best Practices for Microsoft Server Applications

Dennis Martin
President, Demartek

Agenda

- Brief Company Overview
- Recommendations and Best Practices for Microsoft Server-based Applications
- Deployments of Flash in Enterprise Storage
- References and Contact Information

Demartek Company Overview

- Industry analysis with on-site test lab
- Lab includes servers, networking and storage infrastructure
  - Fibre Channel: 4 & 8 Gbps
  - Ethernet: 1 & 10 Gbps (with FCoE and iSCSI)
  - Servers: 8 cores, up to 96 GB RAM
  - Virtualization: ESX, Hyper-V, Xen
- We prefer to run real-world applications to test servers and storage solutions
  - Currently testing various SSD implementations
- Web: www.demartek.com

Best Practices

General Recommendations

- Configure storage for application servers with performance and availability as design criteria
  - Many server applications must satisfy high transaction rates
- Use more disks and faster disks for best performance
  - If you choose “desktop” disk drives, you’re often emphasizing capacity above performance (this choice may also reflect your budget)

Windows Storage Formatting

- Disk Alignment
- Format Allocation (Cluster) Size
- Stripe Size
- Thin Provisioning Storage
- SSD-aware operating system versions

Disk Alignment

- Windows Server 2003 or older:
  Align the file system to the disk offset recommended by the storage hardware vendor. If unknown use an offset of 64K.
  - Diskpart command:
    `create partition primary align=64`
- Windows Server 2008 and above uses default alignment of 1MB
  - Also applies to Windows Vista and Windows 7

Format Allocation Size

- Exchange Server
  - Databases: **64K**
  - Logs: can use default size (typically 4K)
- SQL Server: use **64K** for volumes dedicated to SQL Server
  - The SQL Server page size is 8K
  - SQL Server allocates disk from the operating system in units known as “extents” of 8 pages

Stripe Size

- Since SQL Server accesses disk storage in 64K blocks, the optimum disk array stripe size Microsoft SQL Server volumes is **64K**
- Similar recommendations for Exchange Server
- Your mileage may vary, always test for best results

Thin Provision Storage

- Thin provisioning storage systems use pointers, linked-lists and other similar techniques to minimize the consumption of capacity
- Always use the “quick format” option in Windows
- Defragmentation is not necessary from the operating system

SSD-aware Operating Systems

- Operating systems need to detect the presence of NAND-flash SSDs
  - Windows 7
  - Windows Server 2008 R2
- No defragmenting
- Trim – notify the underlying device regarding data that is no longer needed

O.S. Behavior with Flash and Trim

- Operating systems need to behave differently with flash SSDs
  - Trim – notify the underlying device regarding data that is no longer needed
    - Trim is currently available for SATA interfaces only. The SAS developers are investigating this.
  - Windows 7 and Windows Server 2008 R2
    - Defragmenting is off by default for flash SSDs
  - RHEL 6 with EXT4, but Trim is not enabled by default
- Third-party utilities available

Flash in Enterprise Products

- Disk array vendors
  - Primary storage: SSDs in standard HDD slots
  - Cache: SSD technology used as cache
    - Set it and forget it
- Appliance vendors – “Accelerators”
- Server vendors
  - Add flash on a PCI-Express bus card
  - Add flash directly onto the motherboard
  - Blade server mezzanine cards
- Is enterprise flash storage or memory?

Storage Vendor Trends

- Automated data movement
  - Applies to primary storage
  - Moves hot data to SSD tier
  - Scheduled by minutes, hours, days, etc.
  - LUN level already available; beginning to see sub-LUN level this year
- SSDs together in cache and primary storage
- SSD-only arrays for file and block storage

Demartek Lab Recent Results

- PCI-Express bus:
  - Jetstress IOPS up to 40,000+ with multiple PCI-Express SSD cards in a server
- 6Gb/sec SAS RAID controller:
  - IOmeter random read IOPS more than 100,000 with four SSDs in RAID0 stripe
  - IOmeter random write IOPS more than 20,000 with four SSDs in RAID0 stripe

SSD Performance Comments

- Enterprise applications only need small amount of SSD relative to total HDD capacity for significant performance gains
- Demartek tests with caching solutions show huge gains (5x – 8x) with only one or two SSDs in one disk enclosure

Demartek SSD Resources

- Demartek SSD Zone
  - http://www.demartek.com/SSD.html
- Look for my article *Making the Case for Solid-State Storage* in June online edition of Storage Magazine
  - http://searchstorage.techtarget.com
- Demartek Storage Interface Comparison
  - http://www.demartek.com/Demartek_Interface_Comparison.html

Free Monthly Newsletter

- Demartek publishes a free monthly newsletter highlighting recent reports, articles and commentary. Look for the newsletter sign-up at www.demartek.com.

Contact Information

Dennis Martin, President
Demartek
(303) 940-7575
dennis@demartek.com

www.linkedin.com/in/dennismartin
http://twitter.com/demartek