Virtualization Software for Flash and NVRAM

Session D11

Presenter: Andy Mills, CEO/Co-founder

andy.mills@enmotus.com

Data Center Challanges

- Data centers dealing with huge amounts of data
  - 4 zettabytes[1] of world wide digital data being stored
  - As much as 40% data growth rates
  - 8 million data centers by 2018, 70% of which will be mega datacenters[2]
- Need to rapidly and non-disruptively scale compute and storage resources
- Focus on technologies and products that reduce system wide costs and increases automation and operational efficiency

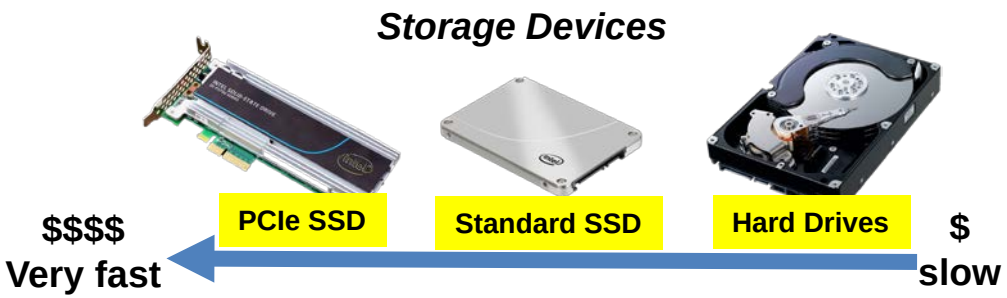

[1] 1 zettabyte = 1 billion terabytes or 1,000,000,000,000,000,000,000 bytes

Trends

- Dramatic shift in the way we are building large scale storage systems
  - Massively scalable data requires new architectures
  - Automation and adaptability has become a key requirement
  - Merger of compute and storage (hyper-convergence)

- Flash based storage gaining a firm foothold in data centers
  - All flash arrays now a ~$1bn market[1]
  - Hybrid Flash Arrays (HFAs) ~$10bn[1]

- Software defined data centers
  - Relying more on open/commodity hardware storage-servers, software and devices
  - Mostly still a flash caching or flash-primary storage world
  - Few “stack independent” software virtualization offerings addressing scale-out

[1] Source: IDC via web

Storage Virtualization

- Storage virtualization history
  - Designed into “big-iron” shared storage systems
  - Ultimate purpose is to facilitate 24/7 operation through storage move/change/adds i.e. increase operational efficiency of enterprise storage systems
  - Opened up new capabilities: thin provisioning, tiering, replication…
  - Tied to hardware and a vertical software stack
- Hyper-convergence and webscale
  - Based on commodity, commonly available building blocks
  - Software defined storage (SDS) gets us part of the way
  - Ceph, Swift Stack yet another approach
- Simplicity “it-just-works” is key for all virtualization

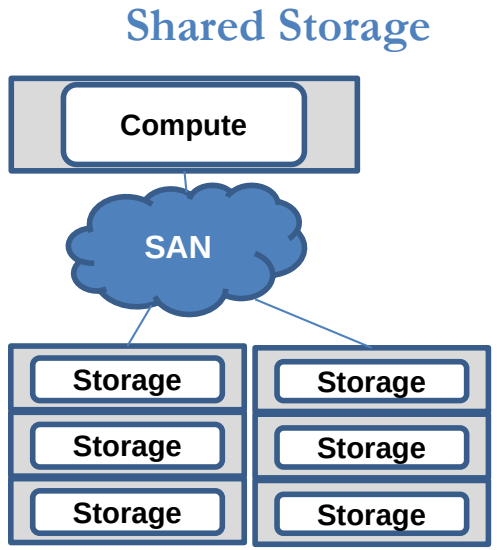

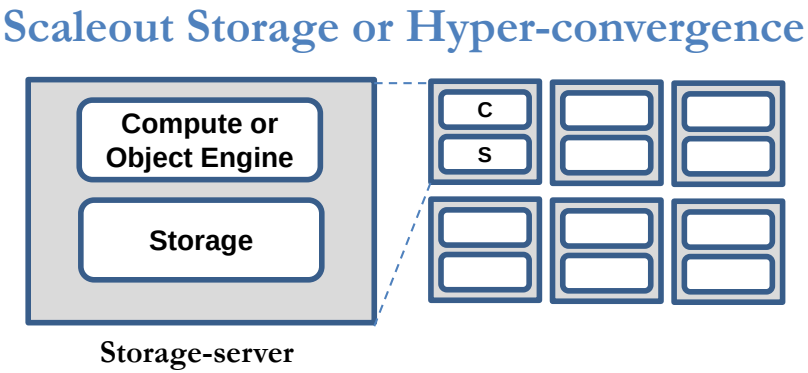

Options for Integrating Flash into Servers

$$$$$
- Put everything on flash or storage class memory
  - 600-700K random IOPs or 3GB/s+ streaming data
  - NVDIMM/NVRAM memory class with flash backup options promising higher performance
  - In-memory for extreme performance databases (several millions of transactions/sec)

$$
- Flash caching software
  - Most common and least risky approach for integrating server-side flash
  - Focus is on speeding up legacy hard disk systems
  - Sacrifice SSD capacity and performance for ease of integration
  - Limited so far to classic SAN architectures and some DAS applications
  - More complex versions appearing, but still essentially caching

$
- Full flash optimized virtualization software
  - Highest performance – flash and NVRAM primary storage in a hybrid storage pool
  - Greatest capacity efficiency – flash contributes to the user pool
  - Least amount of latency and overhead as non-CPU intensive
  - Most recent use real-time tiering and file pinning

All Flash or Hybrid?

- All-flash provides best performance but too costly for broad deployment
  - A large portion of data is not being accessed – so why pay for it to sit there
- Flash as a cache is limited to narrow set of legacy applications
  - Not suited to webscale/big data unstructured and streaming files
  - Capacity is “lost” as flash only holds copies of data – important for larger SSDs
  - Often demonstrates no visible application improvements
- Flash optimized virtualization provides broader options
  - Treats flash or memory class storage as direct attached primary storage
  - Typically 80% less than all-SSD costs
  - Handles moves/changes/adds more efficiently
  - Broader combinations of devices: memory-flash, flash-flash, flash-hard disk

Flash Optimized Virtualization and Tiering

Storage Devices

Fast
$2-4/GB

Flash (SSD)

Slow
4-8c/GB

Hard Disk Drives (HDD) or RAID

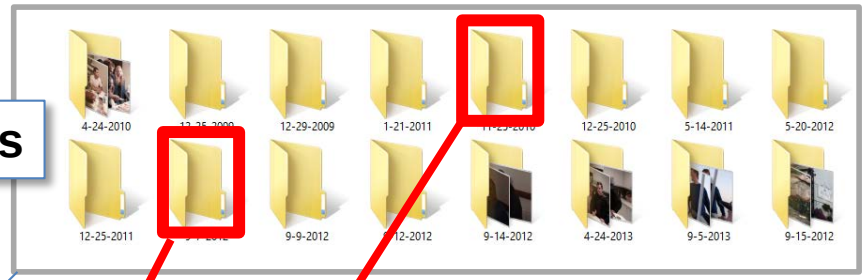

Folders • • • **Virtual Disk** **10-32c/GB**

Intelligently places active data on flash

Use File Pinning to tag work files for SSD performance or simply let it automatically migrate after first access

enmotus

Example Virtual NVMe Flash-HDD Hybrid Performance

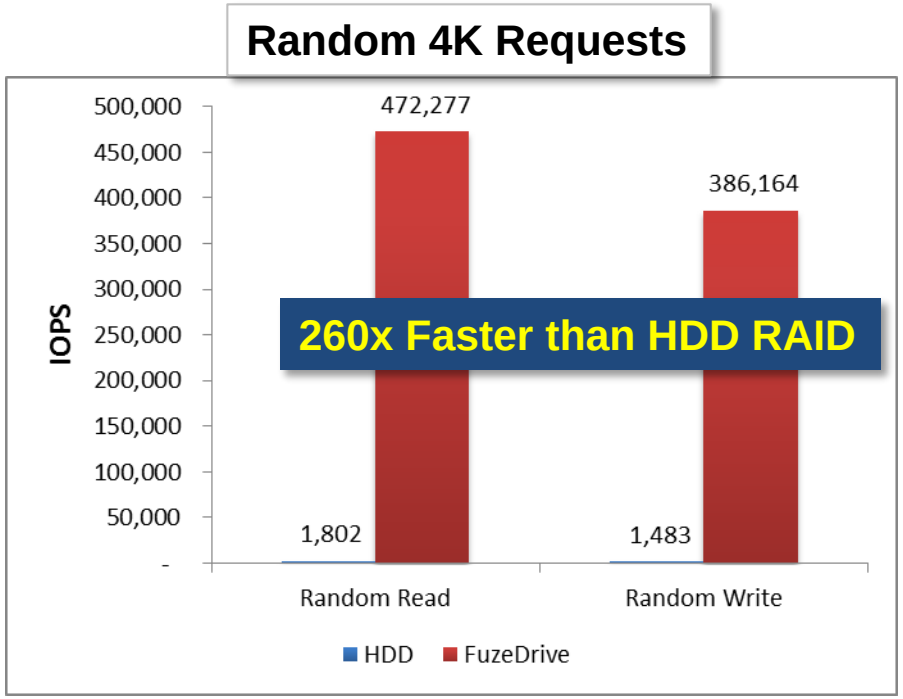

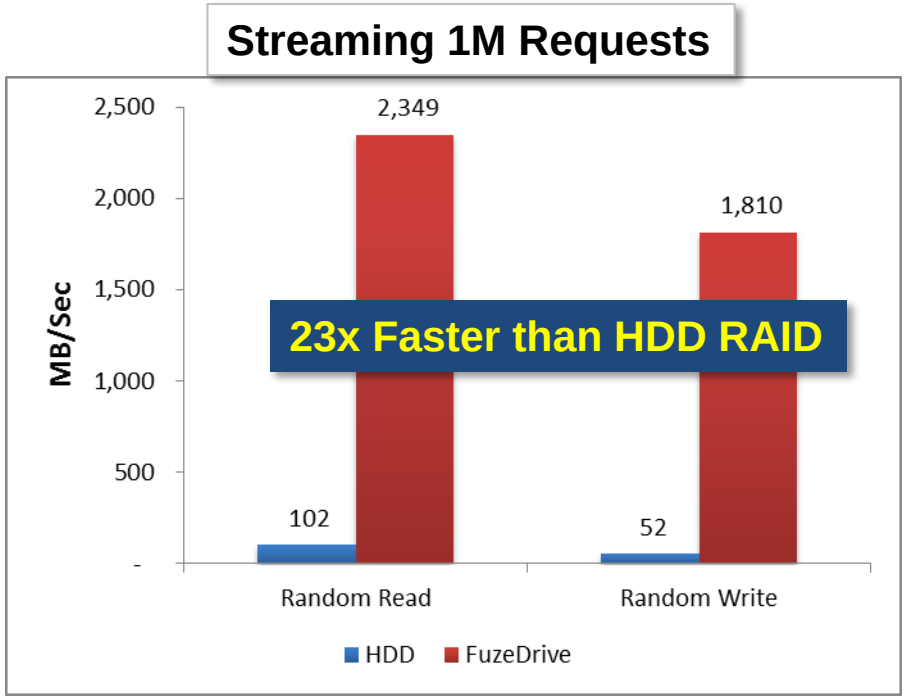

- Up to 260x faster in raw performance than RAID 6 for same capacity (NVMe drives)
- Proprietary PCIe SSDs (e.g. MicronP3/4xx) demonstrate as high as 3.8GB/s off a single SSD

System Setup: Supermicro X9 server class motherboard, Linux CentOS 7 36-bay storage-server,
Single PCIe NVMe SSD fuzed with RAID6 8-drive 6TB drives

Source: http://www.colfax-intl.com/nd/solutions/enmotus-fuzedrive.aspx

Virtual Flash Benchmark Factory 7.1 Score: TPC-C

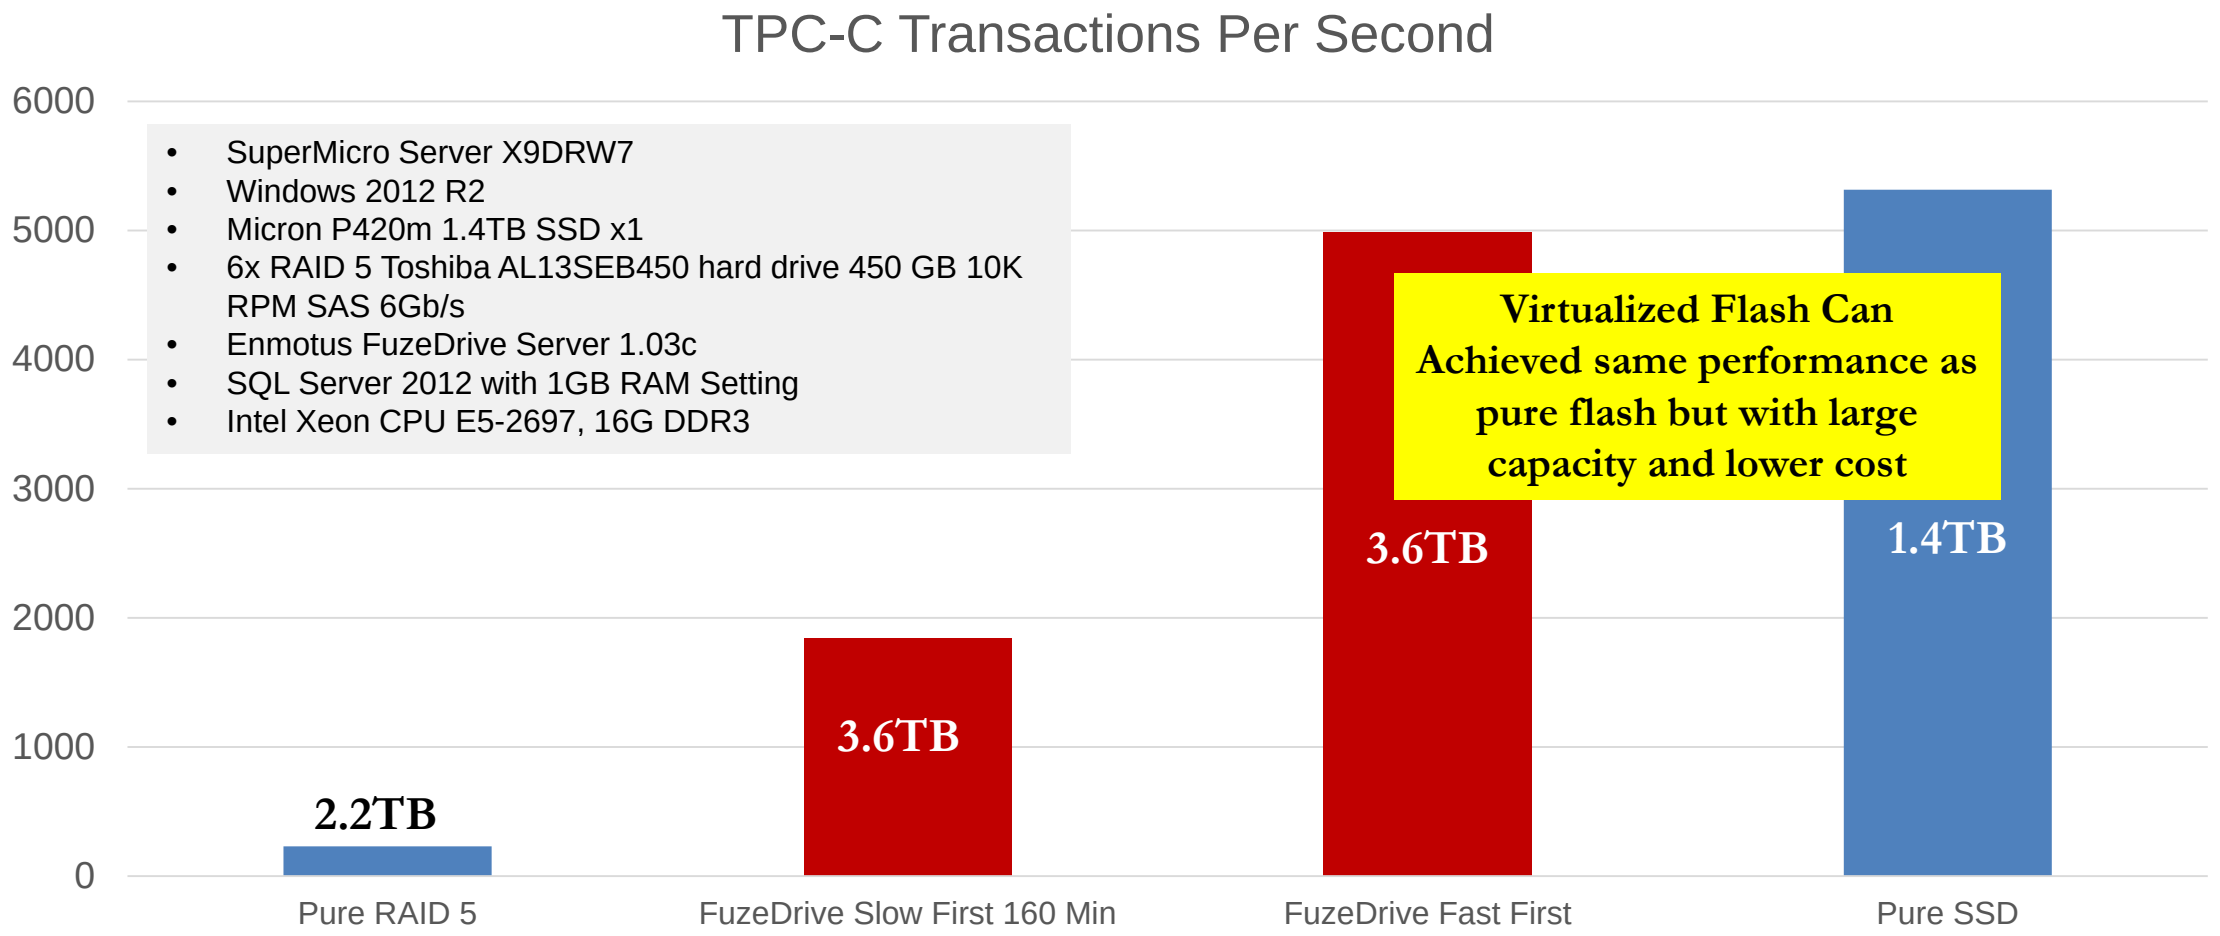

Virtualized Flash TCO Benefits Analysis

- Storage Setup
  - 532 TB total storage
  - 2 replicated sites
    - 266TB/site
  - 38 servers per site
    - 7TB/Server
  - Replica 2 Configuration yields 66.5 TB useable storage
- Server Configuration
  - Dual XEON 2.4GHz CPUs
  - 32 GB DRAM per server

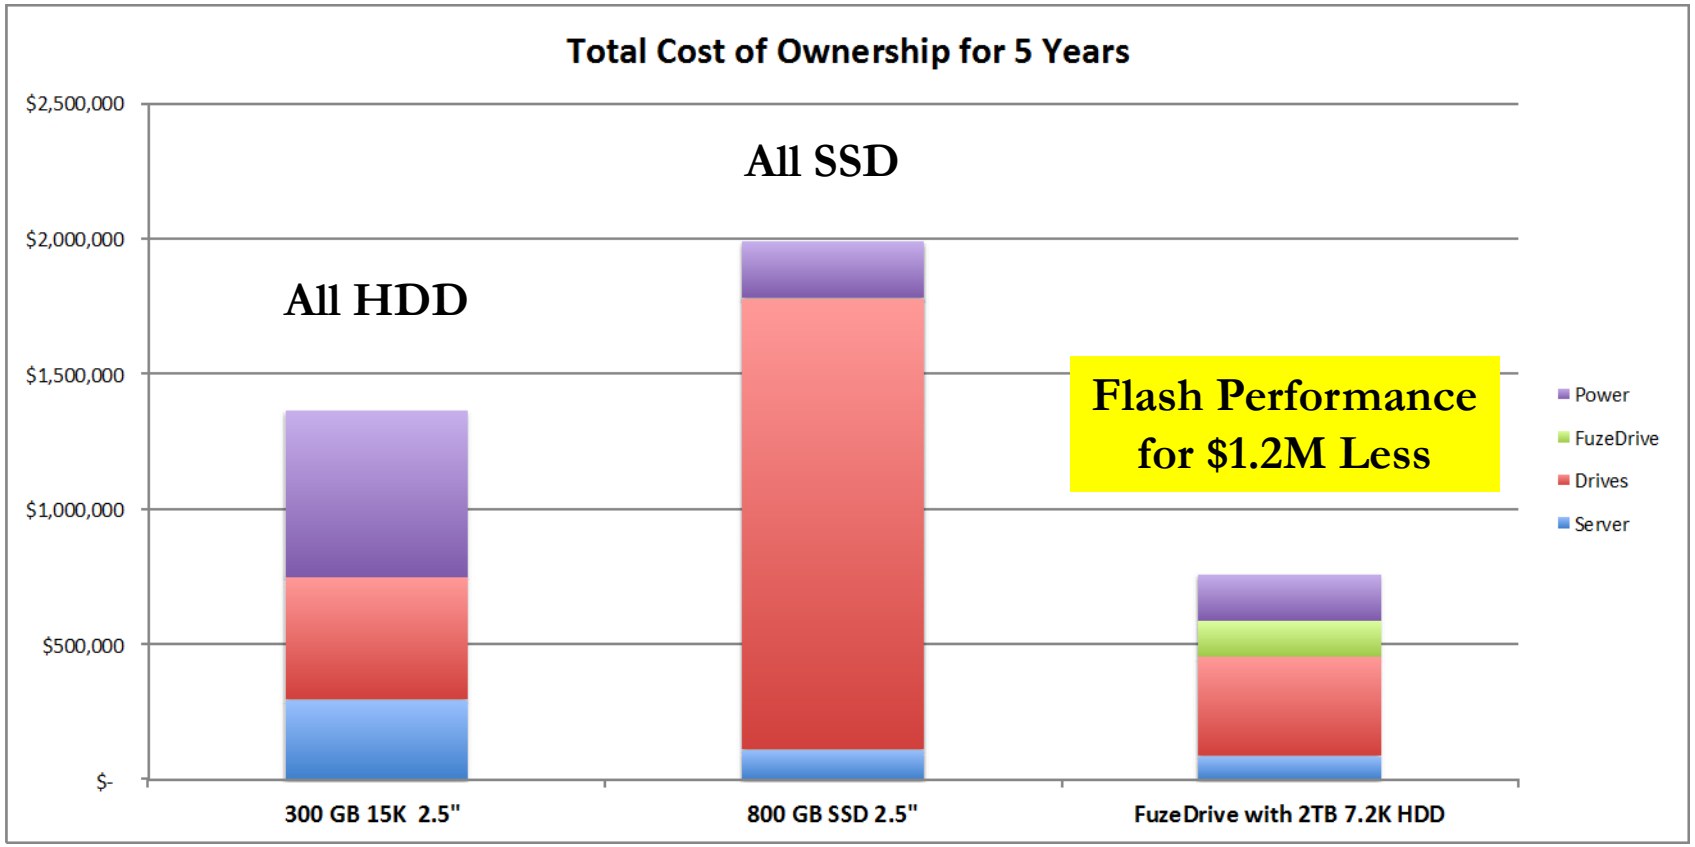

TCO does not include Application SW costs
Enmotus FuzeDrive fast tier = 20% of volume

Other Benefits of Flash Virtualization

- File pinning
  - Ability to pre-pin specific files to flash when needed
  - Example – accounting needs specific records on flash by 4PM on Friday for report generation
- Non disruptive add/move/changes
  - Ability to add more flash and increase capacity
  - Replenish flash that is close to wear out
  - Migrate data in real time
- Low CPU tax and dependency
  - Mapping takes far less CPU cycles than caching
  - Run more VMs

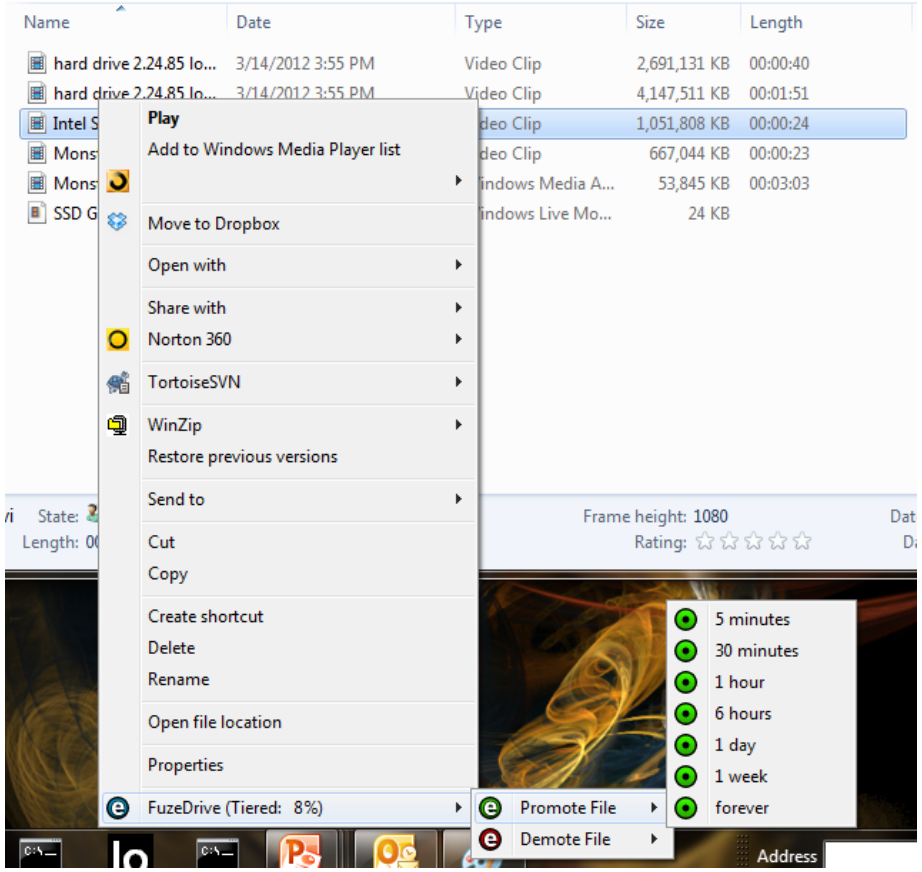

Flash Memory SUMMIT

Enmotus:

Delivering a new class of flash and memory optimized virtualization for massively scaleout servers

Come see Flash Optimized Virtualization in action at Booth 834!

enmotus