NEXSAN A DIFFERENT KIND OF STORAGE EXPERIENCE.™

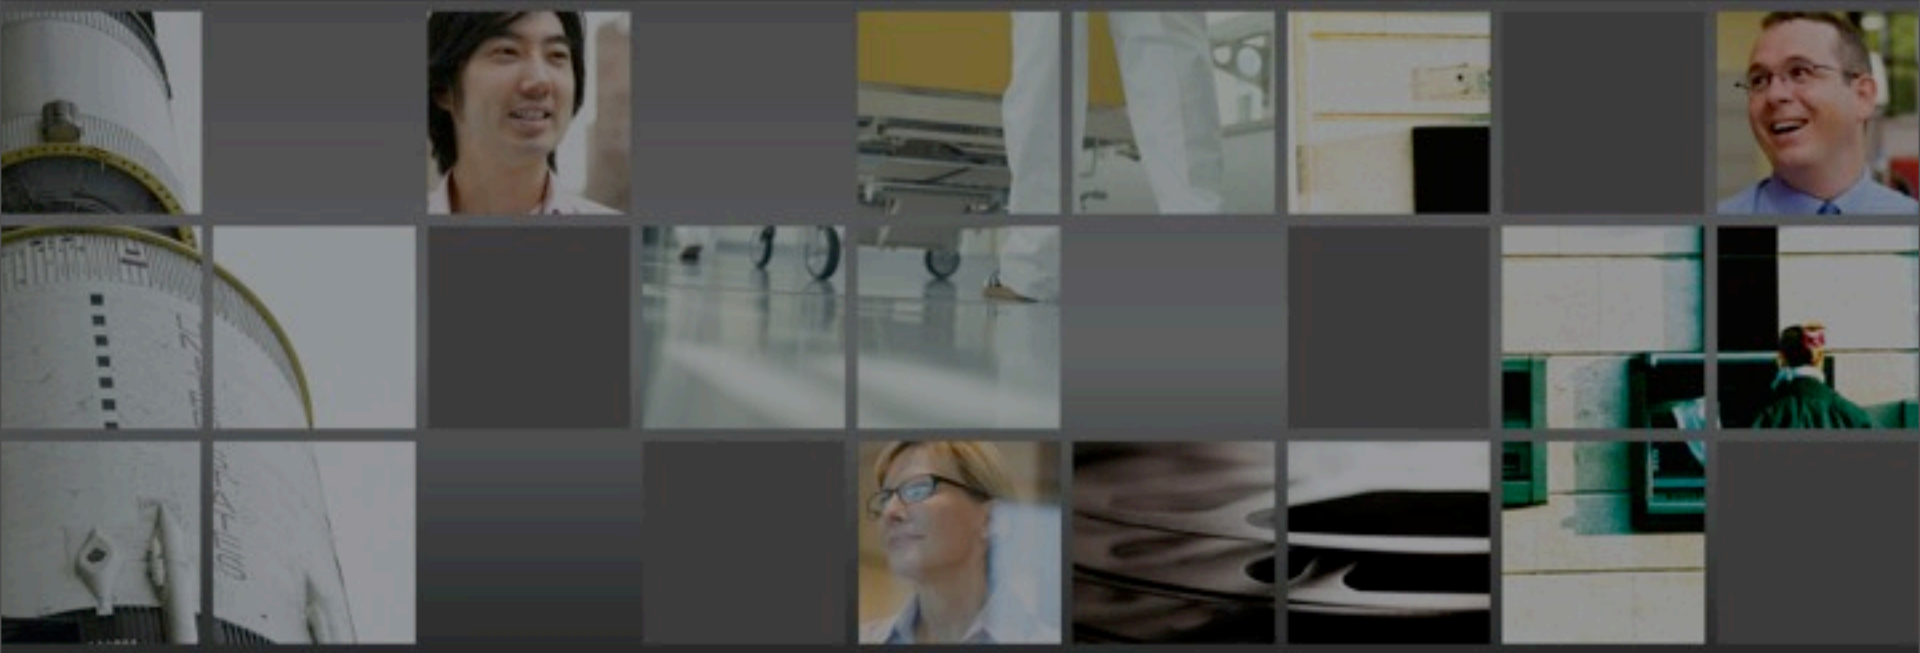

SSD Technologies for Hybrid Storage

Refined Performance

Gary Watson
Chief Technology Officer

August 2012

20 seconds about what we make…

	NEXSAN E18™ NEXSAN E18X™	NEXSAN E48™ NEXSAN E48X™	NEXSAN E60™ NEXSAN E60X™	NEXSAN NST5000™	NEXSAN Assureon®
					
	High Density Expandable FC/ iSCSI Disk Array	High Density Expandable FC/ iSCSI Disk Array	Ultimate Density Expandable FC/ iSCSI Disk Array	Hybrid Storage CIFS, NFS, iSCSI Leverages E-Series	Secure Archive Cloud Storage Highly Secure
	2U, 18 Drives	4U, 48 Drives	4U, 60 Drives	3U, 16 Drives	5U, 42 Drives
Drives	1 / 2 / 3 / 4 TB SATA			1 / 2 / 3 TB SATA	1 / 2 TB SATA
	450 / 600 GB 15K SAS				
	100 / 200 / 400 GB SSD				
In-Chassis	72 TB	192 TB	240 TB	45 TB	56TB per node
With Expansions	216 TB	576 TB	720 TB	1080TB	Scale-out Nodes
I/O	(4) 1Gb iSCSI & (0-4) 1Gb iSCSI or (0-4) 8Gb FC or (0-4) 10Gb iSCSI or (0-4) 24Gb SASx4	(4) 1Gb iSCSI & (0-4) 1Gb iSCSI or (0-8) 8Gb FC or (0-8) 10Gb iSCSI or (0-8) 24Gb SASx4		(4) 1GbE & (8) 1GBE or (4) 10GbE	(4) 1GbE (2) InfiniBand

Four Approaches to Leveraging Solid-state

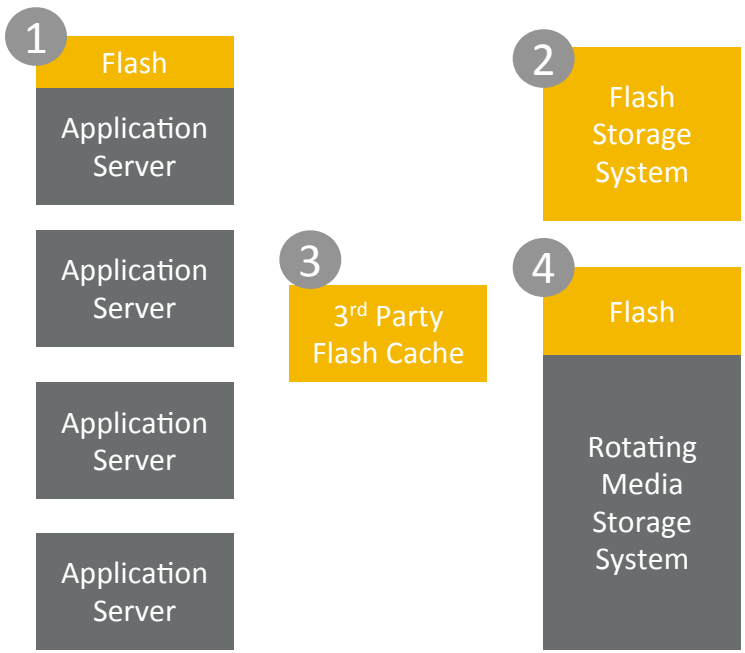

1 **In the Application Server**

- Fastest choice for reference data
- Challenges for H/A and serviceability
- Sharing of data involves high latency which may negate high IOPS of flash

2 **Solid-state Storage System**

- Relatively low capacity, expensive, but fast!

3 **Between the Application Servers and the Storage System**

- Lacks integration, but is shared by servers; tends to be short-term solution.

4 **Hybrid Storage System**

- Cost-effective and shared by all servers
- Single System to Manage
- Performance + Capacity

NEXSAN

What is Hybrid Storage?

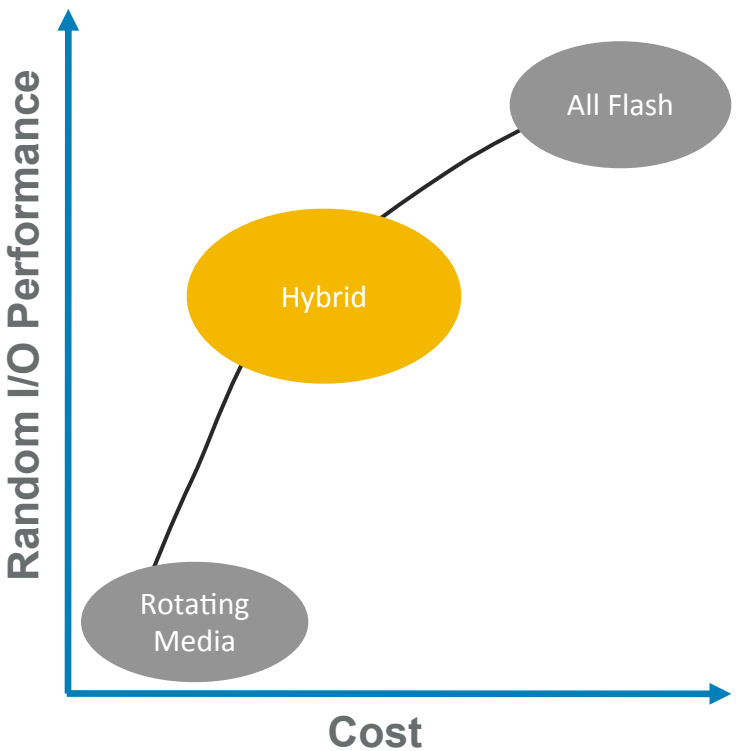

- **Use a modest amount of SSD to accelerate spinning disk performance.**
- **Can be cache or tier or both.**
- **Can be block, file, or unified.**
- **Ideal for transactional, file server, virtual machine, VDI, or mixed applications.**
- **Hot data migrates to SSD cache**

- **Reduces latency**
- **Increases IOPS**
- **More consistent performance**

NEXSAN

Our Implementation of Hybrid

- **We call it "FASTier"**
- **"Tier" in our case refers to different tiers of SSD types**
- **We use NV-DRAM SSD for write cache functions: coalesce, journal, metadata, etc.**
  - Infinite write durability – also much faster than best SLC flash
  - Dual ported pluggable device – all cluster processors can see it
- **We use enterprise SLC flash for read cache**
  - Write durability in the PB's per device is important when used as cache
- **Also optional is eMLC for main data storage**
  - But today's talk assumes spinning 7200 or 15,000 RPM drives

FASTier™ Acceleration Technology

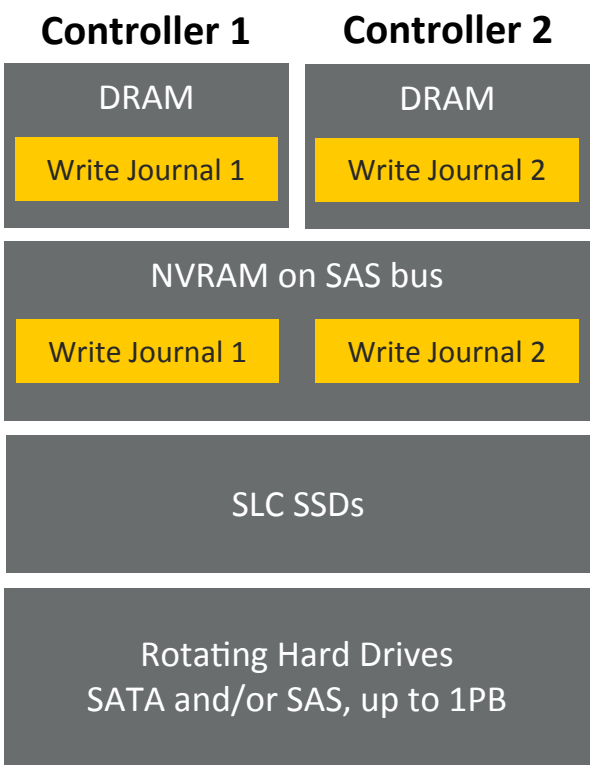

- **Solid-state Devices in the NST5000**
  - Up to 96GB of DRAM per Controller
  - NV-DRAM devices on the SAS bus (8GB each)
  - 100GB and 200GB SLC SSDs
  - FASTier is located before the RAID stack and external boxes
- **FASTier™ Automated Caching**
  - Write Cache
    - Log writes to DRAM (extremely fast)
    - Shadow copy of the write journal on NVRAM for fault-tolerance
    - Used to coalesce data
    - Used for metadata
  - Read Cache
    - From 100GB to 2.8TB of SLC SSDs
    - On demand and opportunistic read-aheads
    - Can feed from write cache or from spinning media

NEXSAN

Entry Level: 14x 7200RPM, single controller

- **NST5110 NAS**

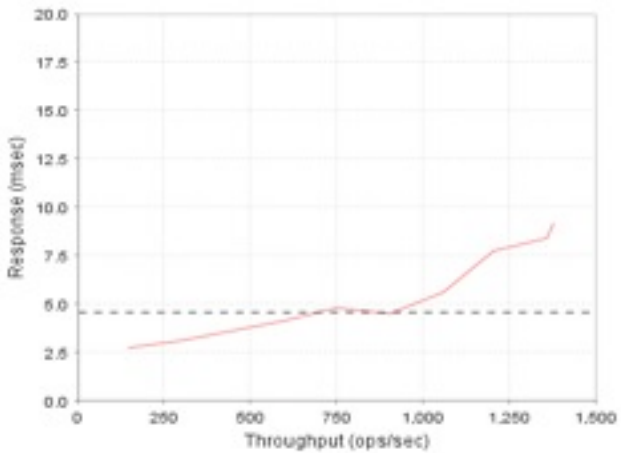

No FASTier

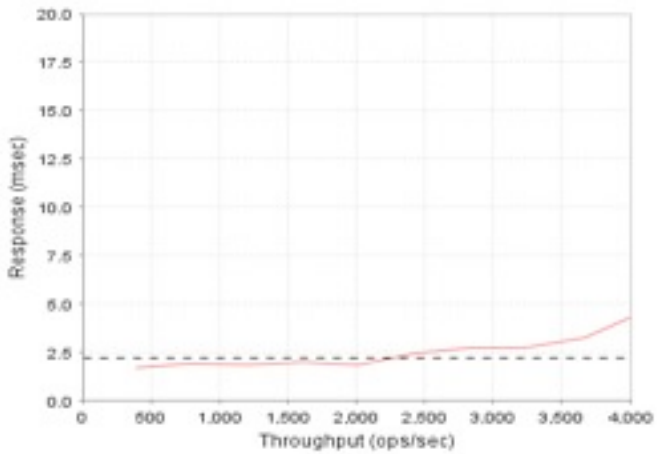

100GB / 8GB FASTier

NEXSAN

Medium: 120x 7200RPM, 16GB / 200GB cache

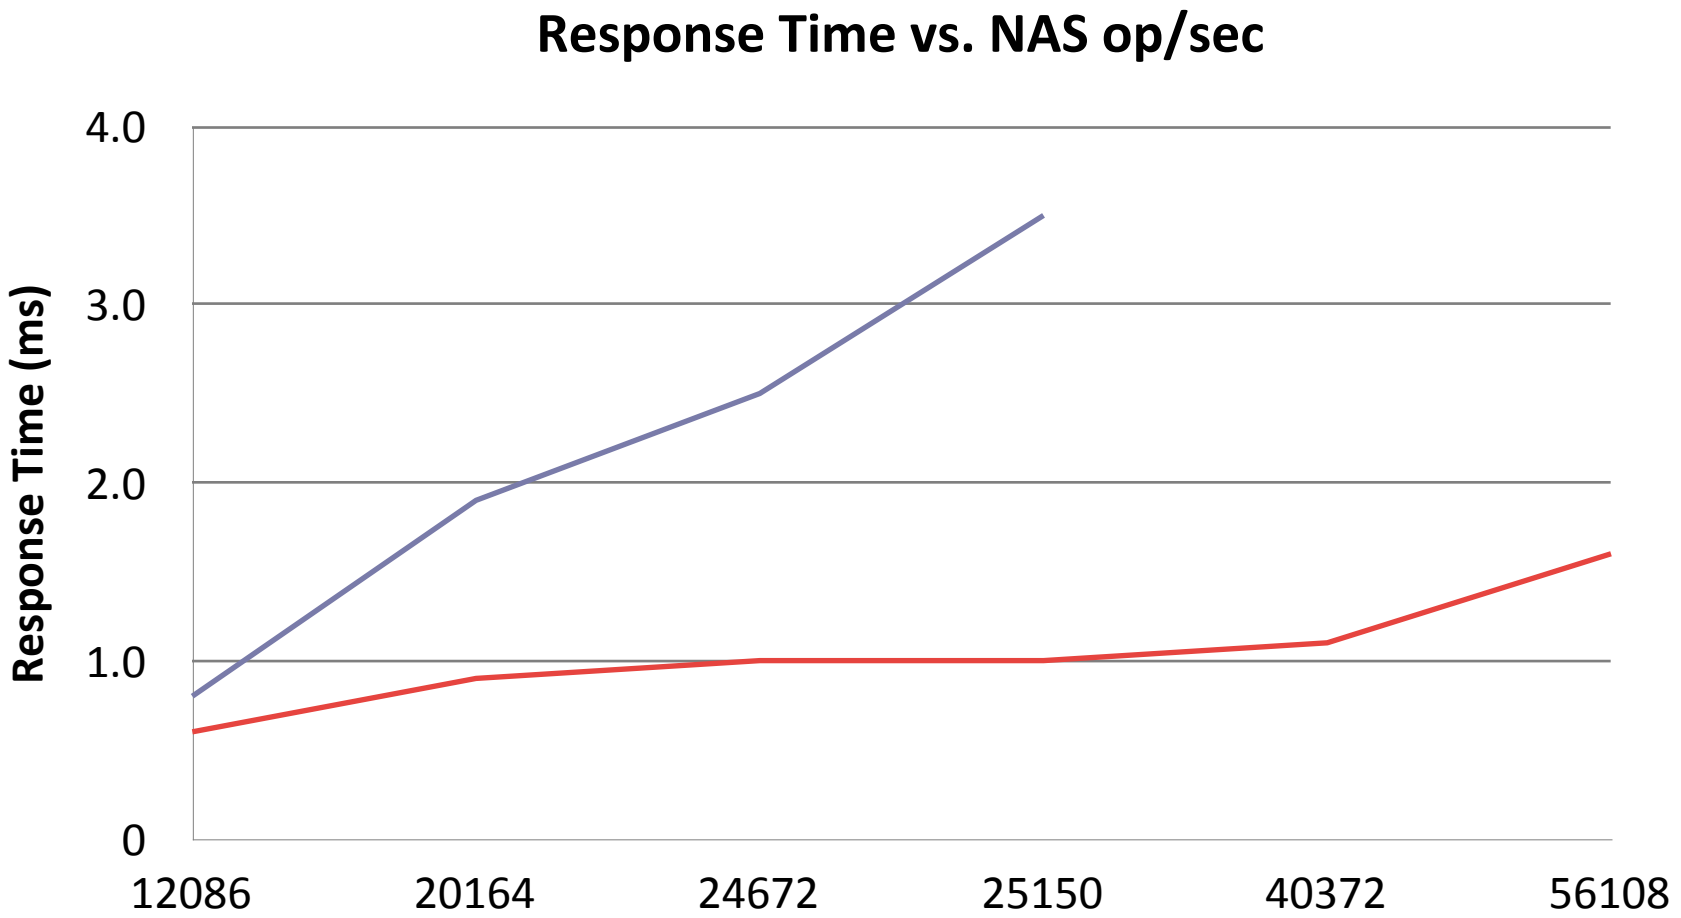

NEXSAN

Large: 96x 15K drives, 64GB / 1600GB cache

Throughput (ops/sec)	Response (msec)
4009	0.9
8023	1.0
12045	1.1
16069	1.1
20099	1.2
24137	1.2
28184	1.1
32224	1.2
36233	1.3
40299	1.3
44365	1.2
48488	1.2
52549	1.3
56990	1.3
61540	1.2
65676	1.1
70048	1.2
73872	1.2
78009	1.2
80805	1.5
85245	1.5
89135	1.6
93188	1.7
96817	1.8
98132	1.9

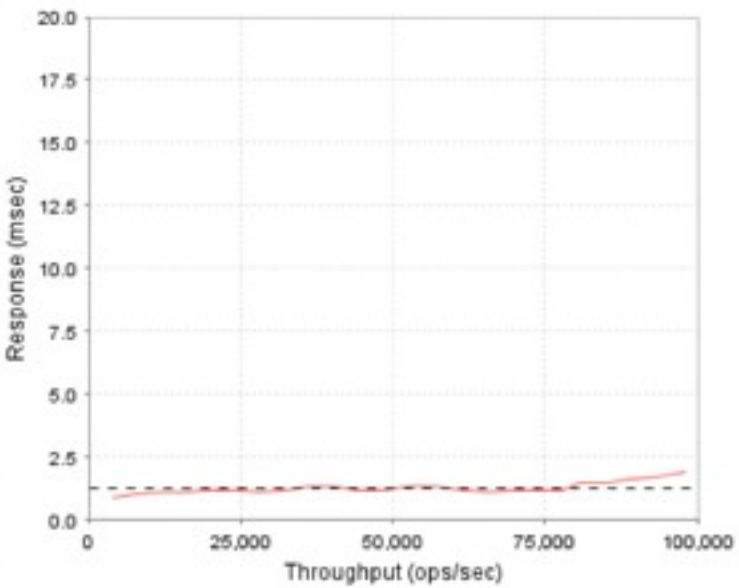

NEXSAN

Selecting Hybrid:

Selecting Hybrid:

- **Consider NV-DRAM for fastest changing data**

Selecting Hybrid:

- **Consider NV-DRAM for fastest changing data**
- **New SLC and eMLC devices are also very good**

Selecting Hybrid:

- **Consider NV-DRAM for fastest changing data**
- **New SLC and eMLC devices are also very good**
- **Relatively small SSD cache can get you 2x to 10x performance**

Selecting Hybrid:

- **Consider NV-DRAM for fastest changing data**
- **New SLC and eMLC devices are also very good**
- **Relatively small SSD cache can get you 2x to 10x performance**
- **Technology changes every quarter – keep informed!**

THANK YOU

Gary Watson
CTO

NEXSAN