



Education

Accelerating Applications and File Systems with Solid State Storage

Jacob Farmer, Cambridge Computer

- ◆ The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- ◆ Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- ◆ This presentation is a project of the SNIA Education Committee.
- ◆ Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- ◆ The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

➤ Accelerating Applications and File Systems with Solid State Storage

The data storage industry is experiencing a mini renaissance as a flurry of new solid state storage devices come to market. SSDs are now available in a broad spectrum of price, form factor, and performance. It is hard enough for a storage architect to sort out the myriad of choices, but that is nothing compared to the challenge of figuring out how to deploy SSD technology in an affordable and effective way. This lecture focuses on practical approaches for integrating SSD technology into your existing storage infrastructure. We explore different insertion points in the SNIA shared storage model to reveal a variety of methods for accelerating databases, email systems, virtualization platforms, content repositories, local file systems and network file systems (NAS).

- Overview of solid state storage devices and where they fit in the SNIA shared storage model
- Accelerating specific applications
- Accelerating virtual server environments
- NAS and file system acceleration
- SAN and block/volume-level acceleration
- Questions and answers

Random vs. Sequential I/O

- Random Disk I/O performance is not keeping pace with other advancements such as processor and network I/O
 - ◆ 15K drives are not that much faster than 10K!
 - ◆ Traditional RAID arrays make it difficult to aggregate disk spindles without host-based LVM software
 - ◆ As disk capacities grow, it becomes very difficult to balance capacity and performance
- Applications running in virtual servers are particularly constrained by random I/O
 - ◆ The virtualization environment randomizes traditionally sequential workloads
 - ◆ Booting and page file activity tend to share the same disk pools as applications and thus compete with apps for disk I/O

Why are SSDs Suddenly Hot?

- The price of DRAM has declined to the point where a meaningful capacity is affordable
 - ◆ Traditional DRAM SSDs were very pricey, niche technologies
 - 5 years ago – 20GB of SSD sold for about \$500K
- The emergence of enterprise-class flash technology
 - ◆ Recent innovations allow flash-based SSD to achieve meaningful performance with enterprise-class reliability.
 - ◆ Flash offers new tiers in between DRAM SSD and 15K disk
- We are nearing the tipping point where SSD goes viral.
 - ◆ Higher demand accelerates the decline in price
 - ◆ Commoditization inspires new software technologies that leverage commodity hardware

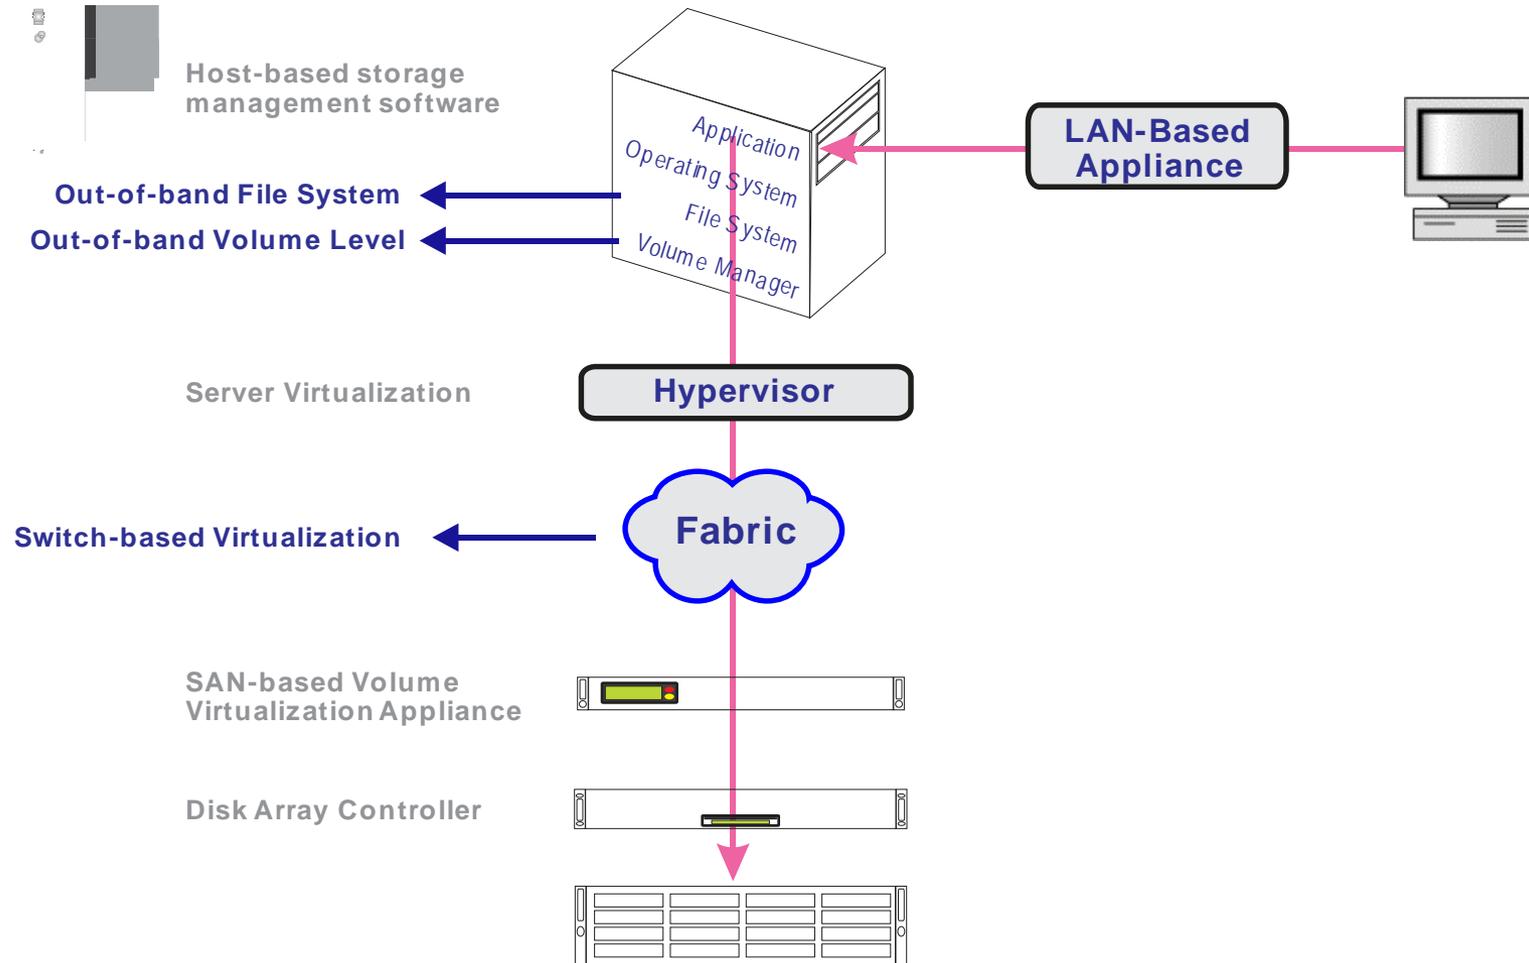
Attractive New Tiers Create Demand for Automating Tiered Storage

- The landscape of SSD offerings represents several possible tiers above traditional 15K disk
 - ◆ Media types: RAM, SLC Flash, MLC Flash
 - ◆ Interfaces: PCI-E, local storage channel, SAN-attach
 - ◆ Performance also varies as a function of product engineering
 - ◆ New devices are coming to market rapidly
- Automated Tiered Storage = The ability to scale performance and capacity independently of one another.
 - ◆ Need performance? Buy performance.
 - ◆ Need capacity? Buy capacity.

SSD Packaging: SSDs Come in Many Shapes and Sizes

- Hard Drive Form-Factor with SATA, SAS, or FC interface
 - ◆ Can be installed in server drive slots
 - ◆ Can be installed in enterprise SAN enclosures
 - › Add a new high performance tier (“Tier Zero”) to your SAN array
- PCIe Cards and PCIe External Appliances
 - ◆ Looks like a hard drive to the local OS
 - ◆ Lower latency than devices that use storage channels (FC, SATA, SAS)
 - ◆ Suitable for applications that can use direct-attached storage
 - › PCIe switches are slowly being adopted by enterprise applications
- SAN-Attached Appliances
 - ◆ Great for traditional enterprise apps, especially when high availability clustering is involved
 - › Failover Clusters, Server Virtualization Platforms, Oracle RAC
- DIMM Modules Installed on the Motherboard on the RAM Bus
 - ◆ Extremely low latency, especially on reads

The SNIA Shared Storage Model: Where Can You Insert SSD Technology?



© copyright Jacob Farmer, Cambridge Computer, 2010

- Any application or data set where the demand for random I/O performance is high relative to the demand for capacity
 - ◆ Databases: logs, temps, indexes, whole tables
 - › MS Exchange: logs and data
 - ◆ Operating systems: boot images, page files
 - › Next Generation Virtual Desktop Solutions (VDI)
 - › Multi-host Platforms: VMware, Citrix
 - ◆ Out-of-band file systems for metadata catalog
 - › Any file system where metadata can be isolated from file data
- Modern advances enable cache and SSD for general purpose computing
 - ◆ Cache and SSD can be inserted anywhere along the storage I/O path
 - › Dynamic tiered storage, life cycle management, caching

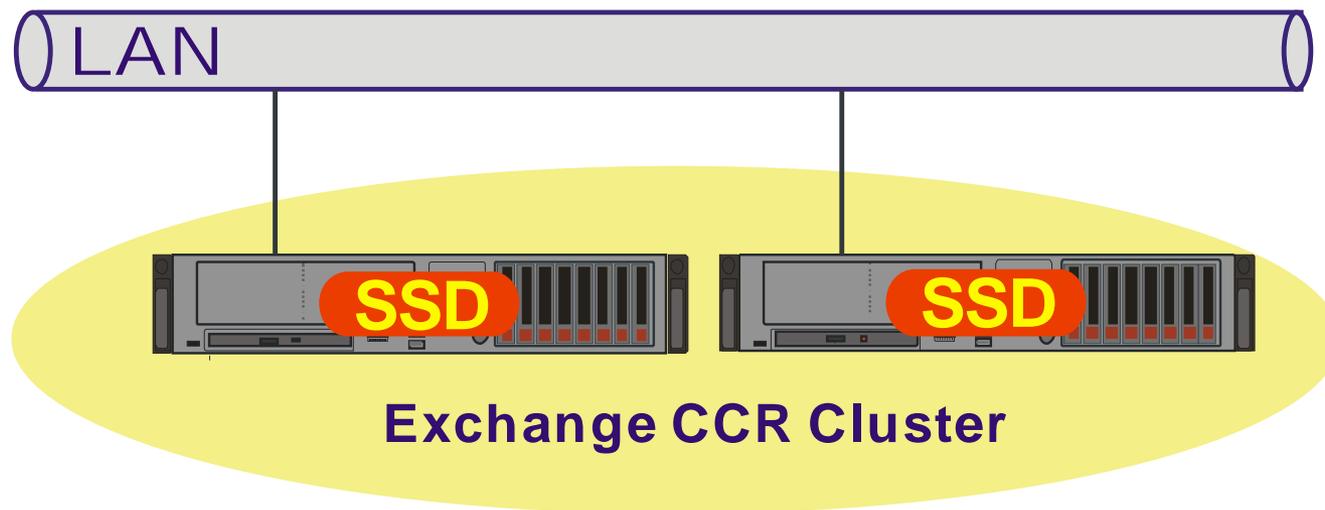


Education

Accelerating Individual Applications

- ◆ Sizing disk I/O for MS Exchange
 - ◆ 1/4 to 1/2 I/O per active user
 - ◆ Multiply by 4 for mobile users, multiple by 1.5 for any search, etc...
- ◆ What should you put on SSD
 - ◆ The entire data set if it will fit
 - ◆ EDBs holding mail for select power users
 - ◆ Log files
 - › Big win for write-intensive mail systems
- ◆ Exchange 2007 and 2010 offer HA clustering without the need for a shared storage SAN
 - ◆ CCR – Continuous Cluster Replication
 - ◆ DAG – Database Availability Groups
 - ◆ Leverage the low-cost of direct-attached SSDs

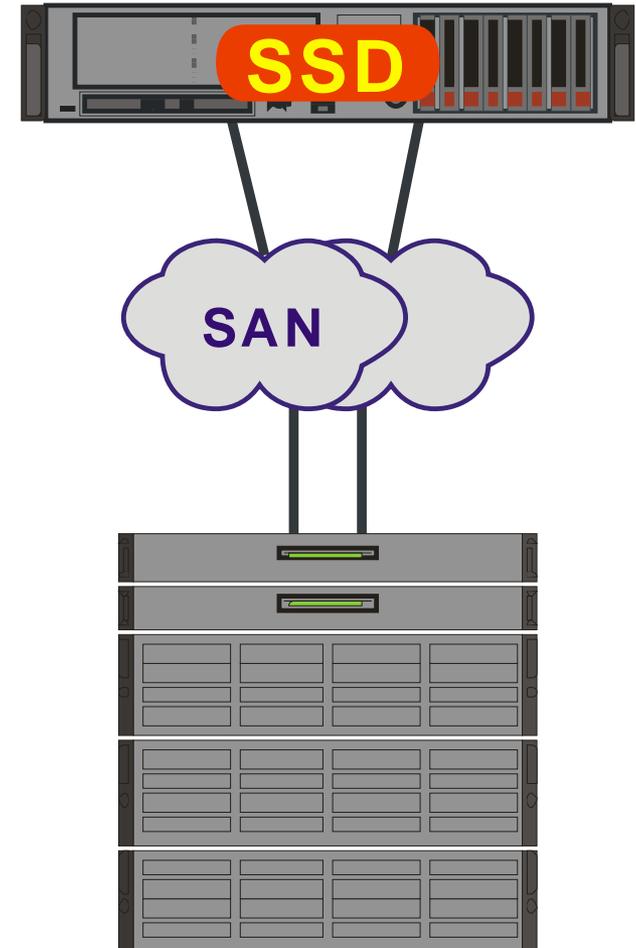
Direct-Attached SSDs with Cluster Continuous Replication



- CCR – Cluster Continuous Replication (2007)
- DAG – Database Availability Groups (2010)
 - ◆ MS Exchange replication and HA clustering without the need for shared SAN storage

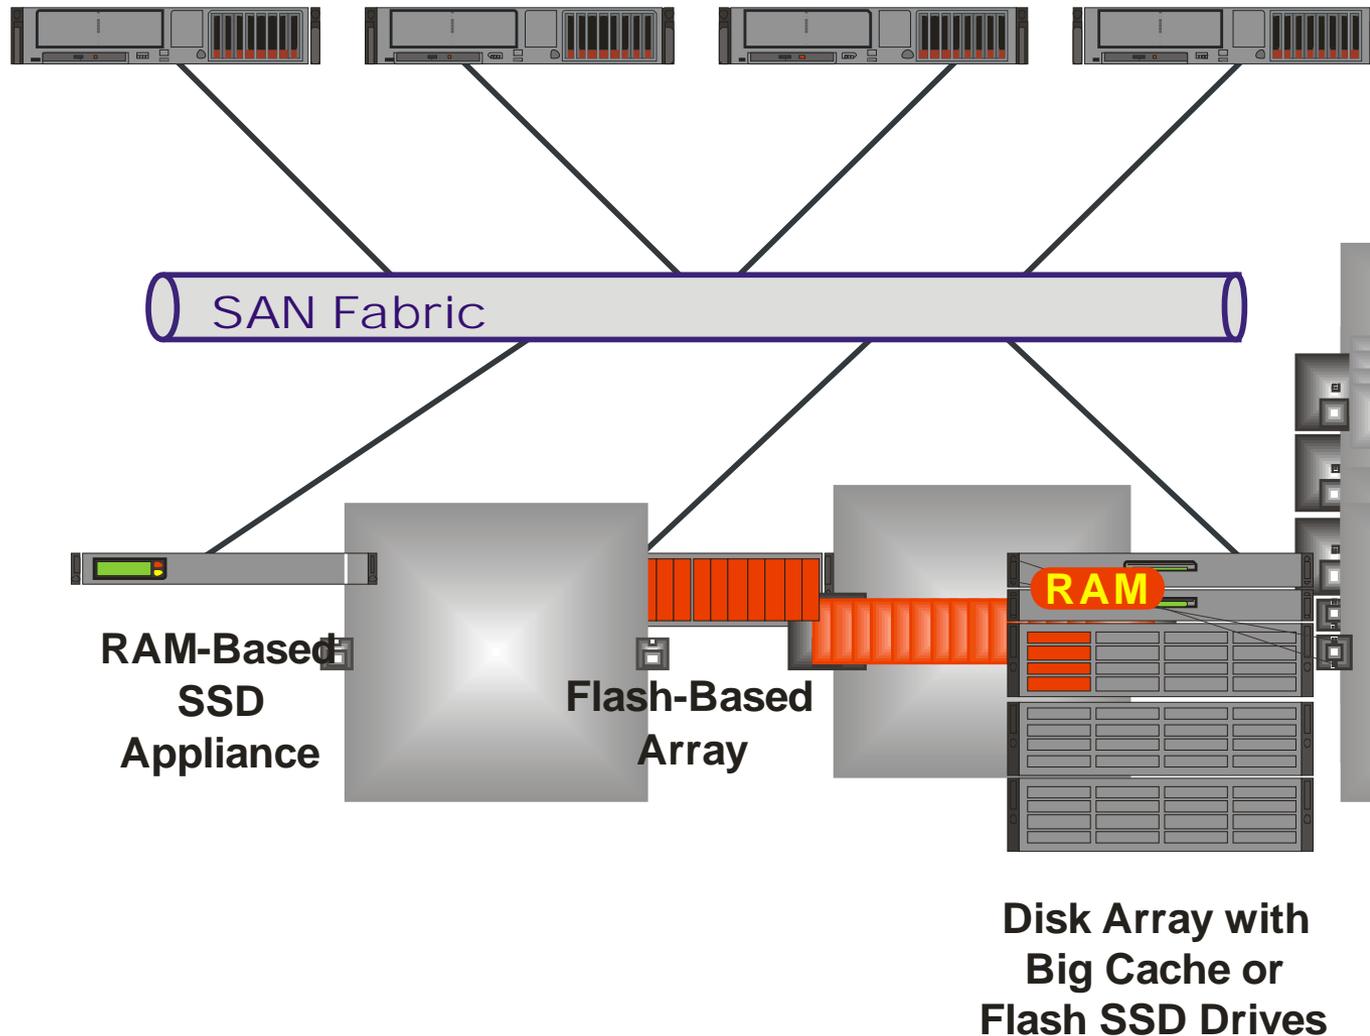
MS Exchange: Using Direct-Attached SSDs for Log Files Only

- Exchange logs are stored on internal SSD
- Exchange databases are stored on traditional enterprise SAN
 - ◆ Snapshots
 - ◆ Replication
 - ◆ Provisioning



- For small databases, put the entire database on SSD
- Typical best practice is to determine which elements of the database will benefit from SSD
 - ◆ Usual suspects: Logs, Indexes, Temps, select tables
 - ◆ Performance monitoring tools are available for all major databases
 - › Example: Oracle StatsPack
- Use volume manager, intelligent SAN, or database tools to migrate hotspots to SSD
 - ◆ Oracle ASM, for instance, allows seamless migration from one storage device to another.

Adding SSD Devices for Database Hotspots

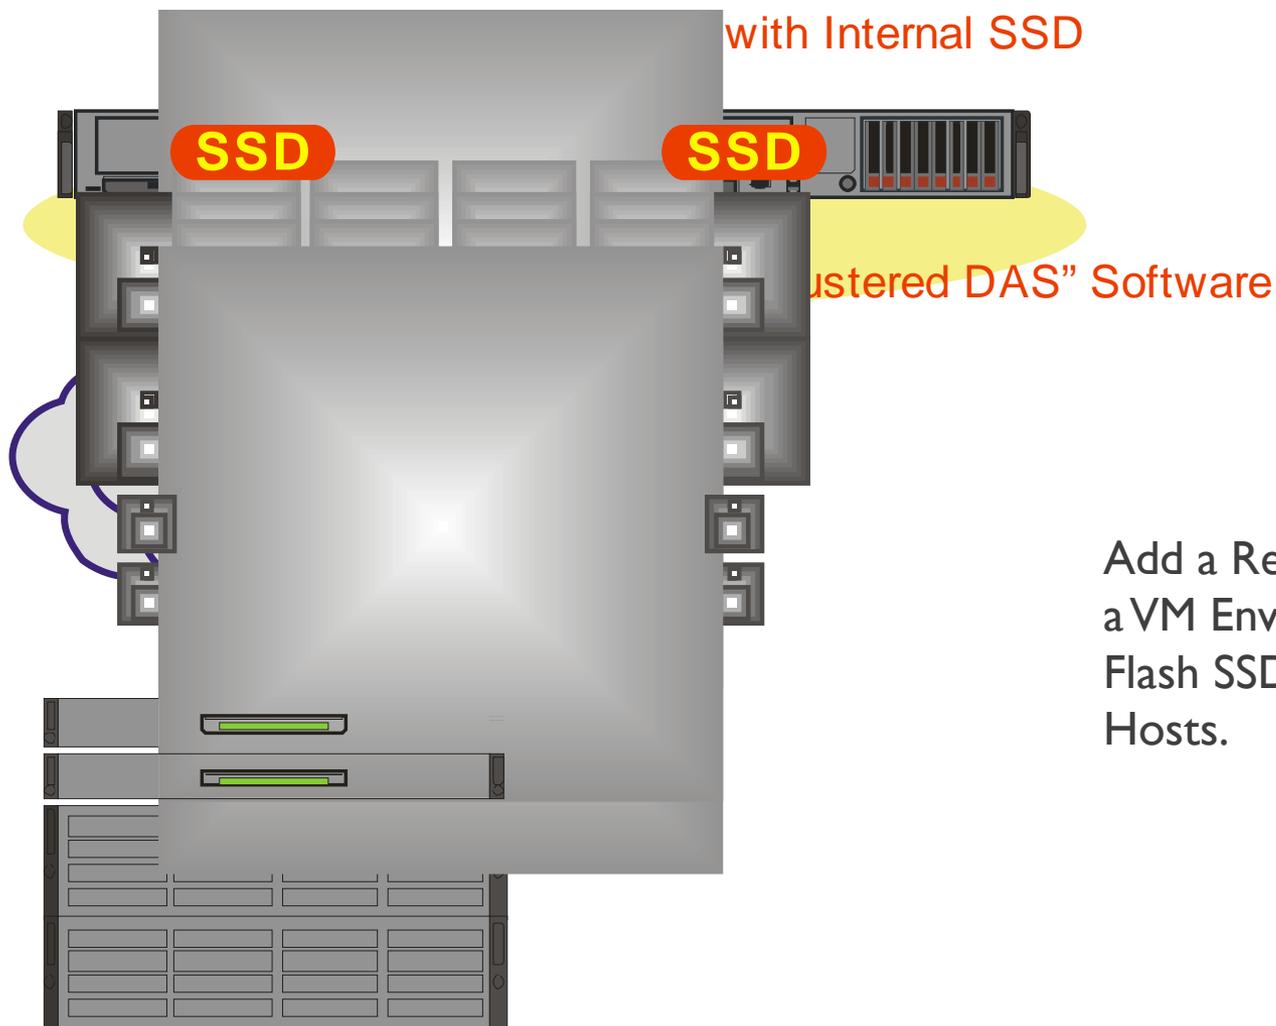




Education

SSDs with Server Virtualization

SSD with Clustered DAS Logical Volume Manager on Hyper-V/vSphere/XenServer



Add a Redundant Tier of SSD to a VM Environment by Installing Flash SSD Cards into the VM Hosts.

Shrinking the Cost of Citrix XenServer

XenServer performance is often bound by disk I/O. Adding internal SSDs reduces the number of hosts.



...therefore eliminating the need for additional servers, licenses, management, and support.

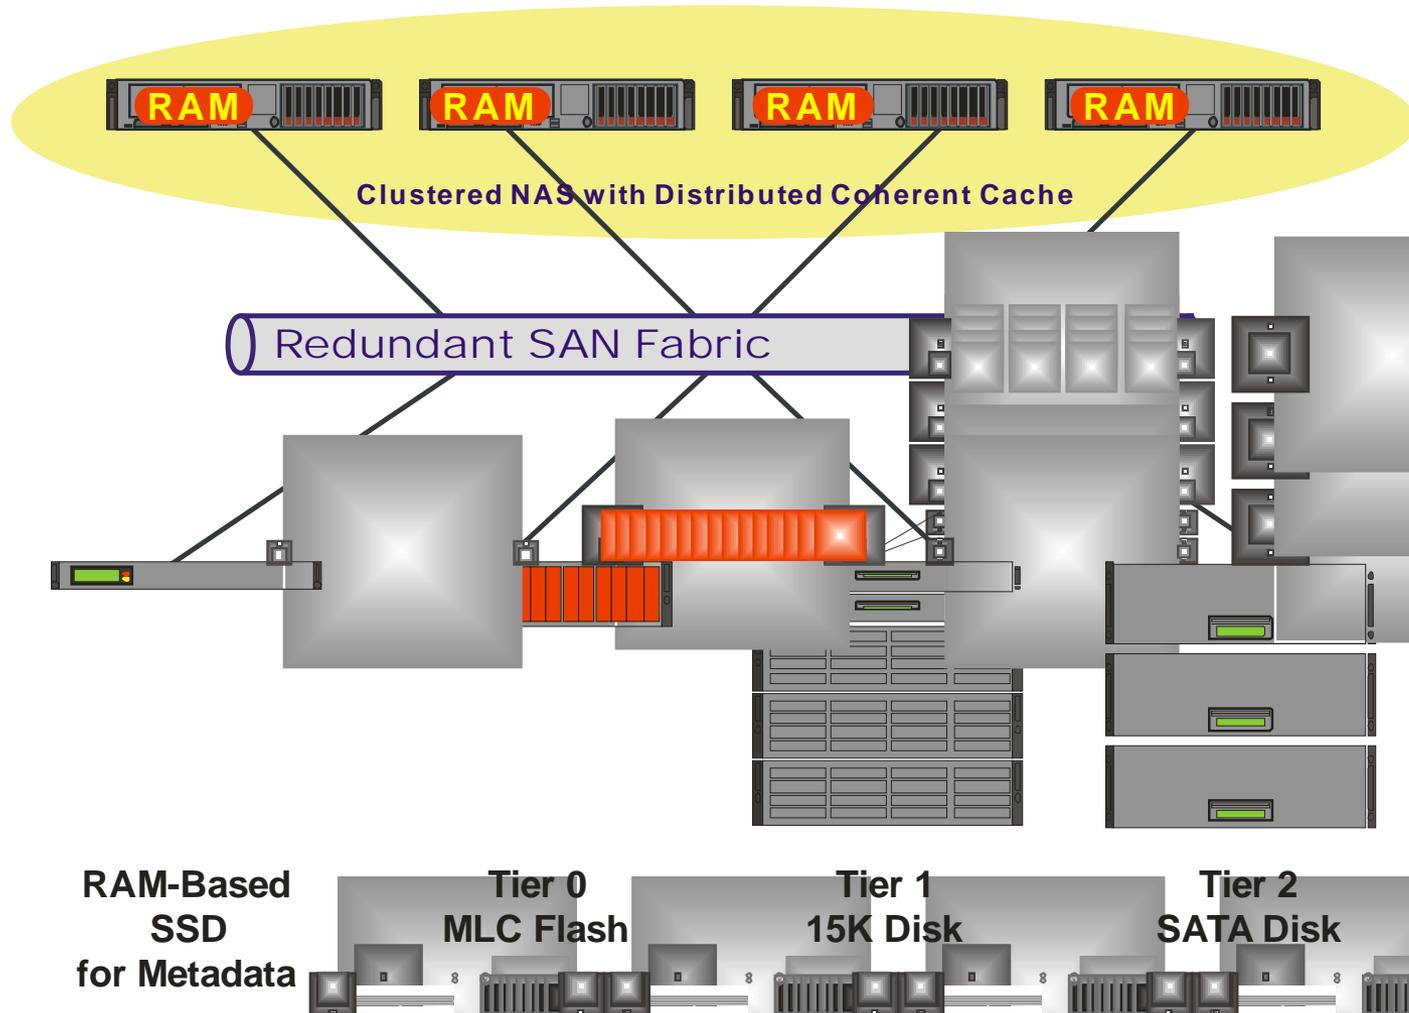


Education

File Systems and NAS Acceleration

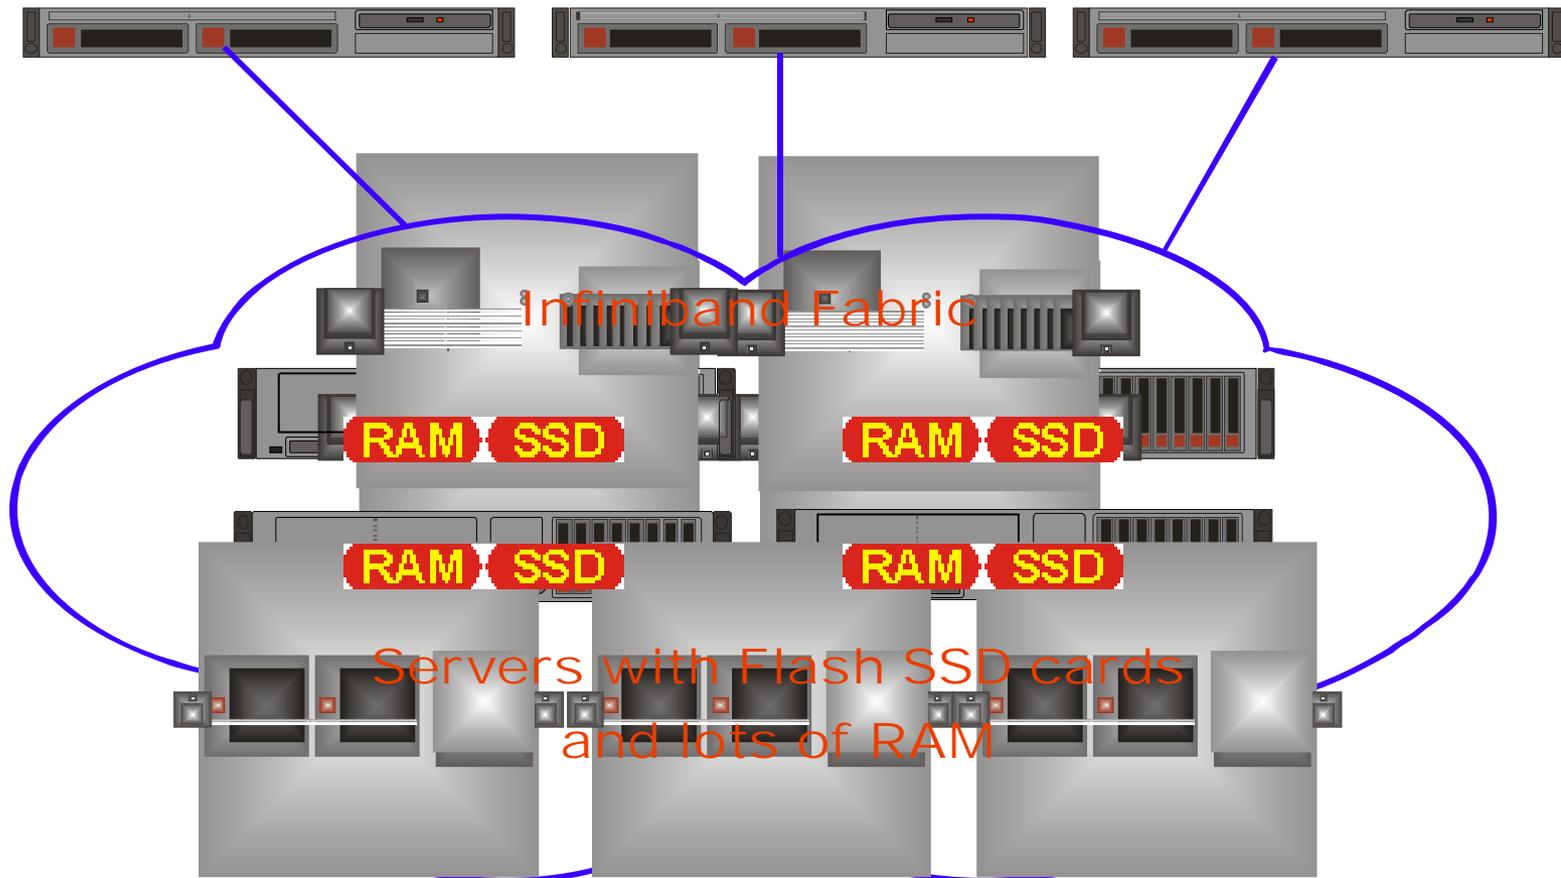
- ◆ Store file system metadata on SSD tier
 - ◆ Many modern file systems allow metadata to be stored separately from the file data.
 - › Object file systems
 - › SAN files systems with out-of-band metadata
 - ◆ Small I/Os and metadata operations (e.g. crawling the file system) can be greatly accelerated
- ◆ Many file systems offer dynamic file tiering in which case most recently accessed files and new files could be stored in an SSD tier.
- ◆ Some file systems allow the journal to be stored separately on an SSD tier
- ◆ Maybe your whole file system will fit on SSD
 - ◆ Use SSDs with hard drive form factor in place of a hard drive.

Clustered Scale-Out NAS: Living the Dream!

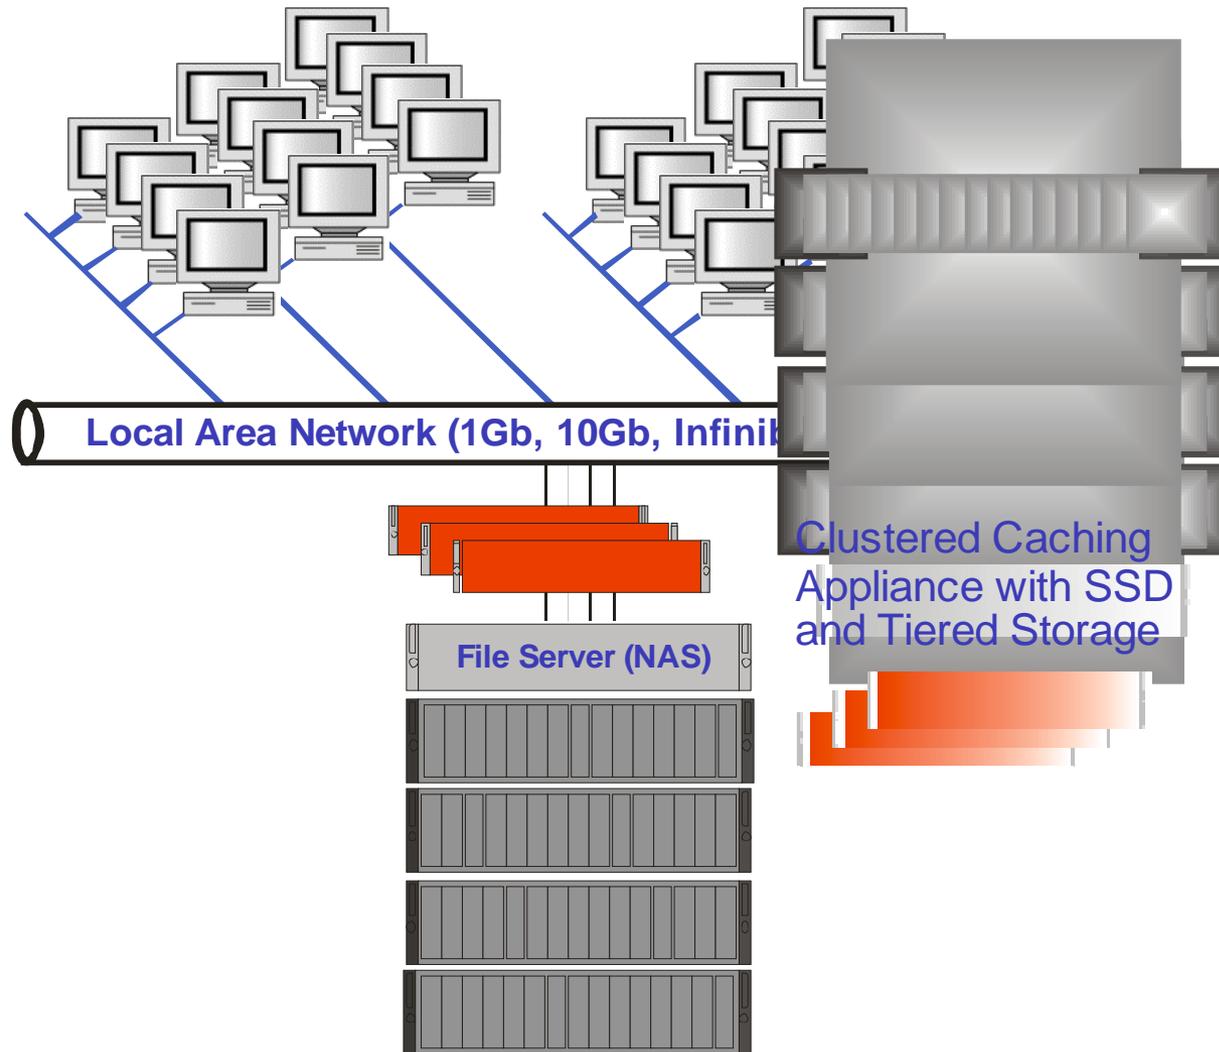


A Clustered File System on 10Gb Ethernet or Infiniband with Internal (DAS) SSD

Application Servers and/or Compute Nodes



Clustered In-Line NAS Cache with SSD and Tiered Storage

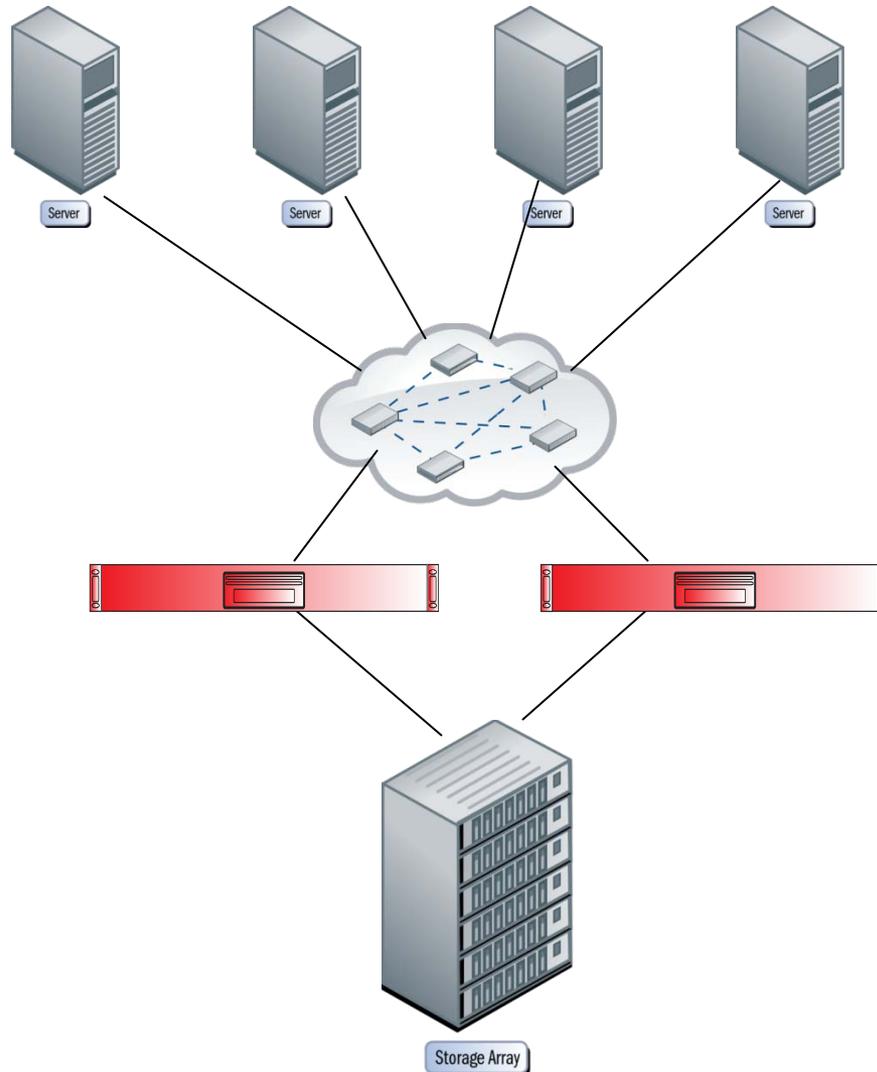




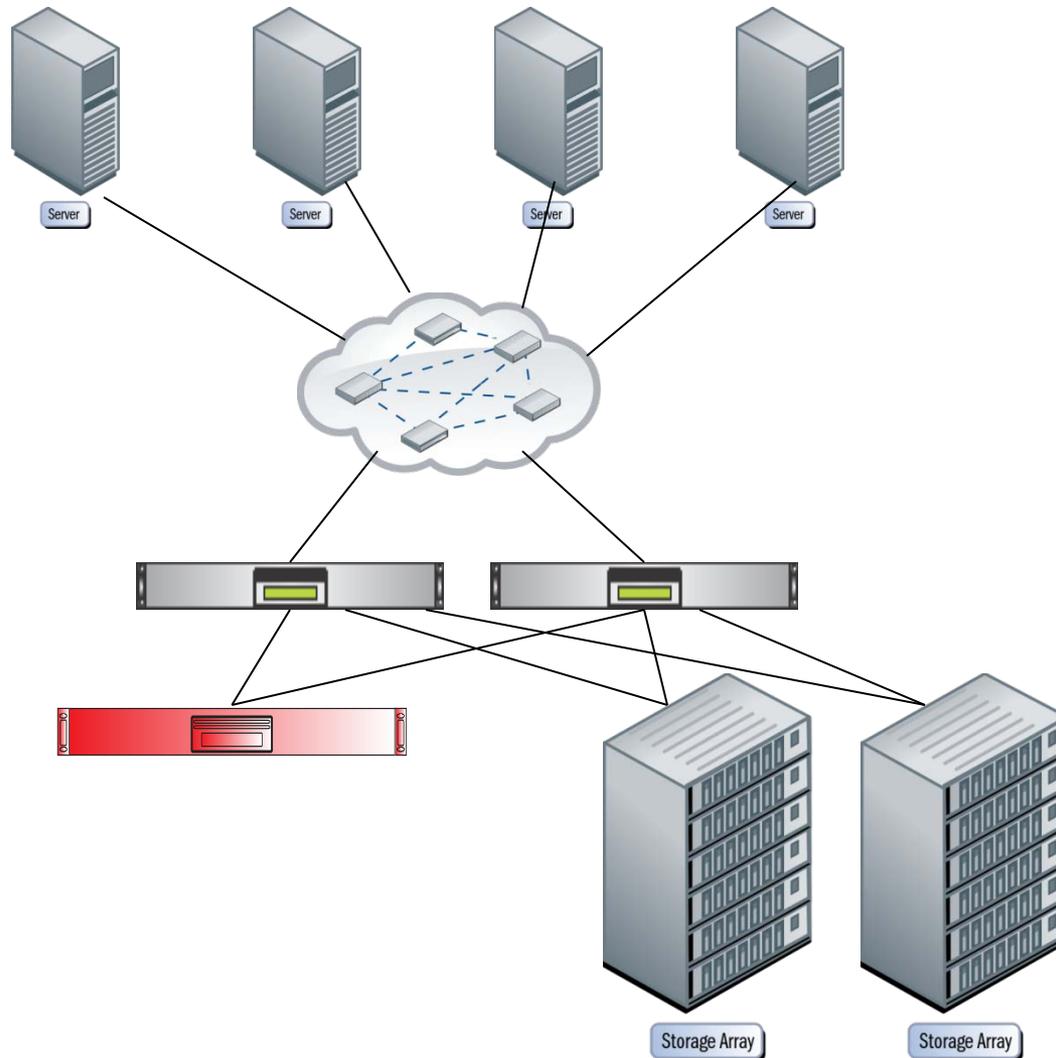
Education

Accelerating the SAN

In-band SAN Cache



Hot-Spotting or Caching with SSD Using SAN Volume Virtualization



- SSD can be leveraged across the entire array with dynamic block-level tiering or extended caching.
 - ◆ Lower tier for cache
 - ◆ Tier 0 for disk
- Hottest blocks or fresh new writes can be stored on SSD and migrate to lower tiers when not in active use.
- Makes it easy to deploy SSD without having yet another tier to manage.



Storage Array

- SSD solutions can be applied to specific applications or added to the general infrastructure
- The ultimate goal should be to scale performance and capacity independently of one another
 - ◆ SSDs allow you to buy performance in a wide variety of costs and performance capabilities
 - ◆ Dynamic tiering solutions allow you to get the most from each tier of device
- SSDs can fit just about anywhere in the SNIA shared storage model.
 - ◆ The trick is to figure out the optimal place for price, performance, reliability, and simplicity.

- Please send any questions or comments on this presentation to SNIA: tracksolidstate@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

Jacob Farmer