Transform Your SAN with Transparent Application Acceleration

Enterprise Caching Forum

Cameron T Brett

QLogic Corporation

Valued Acceleration Approaches

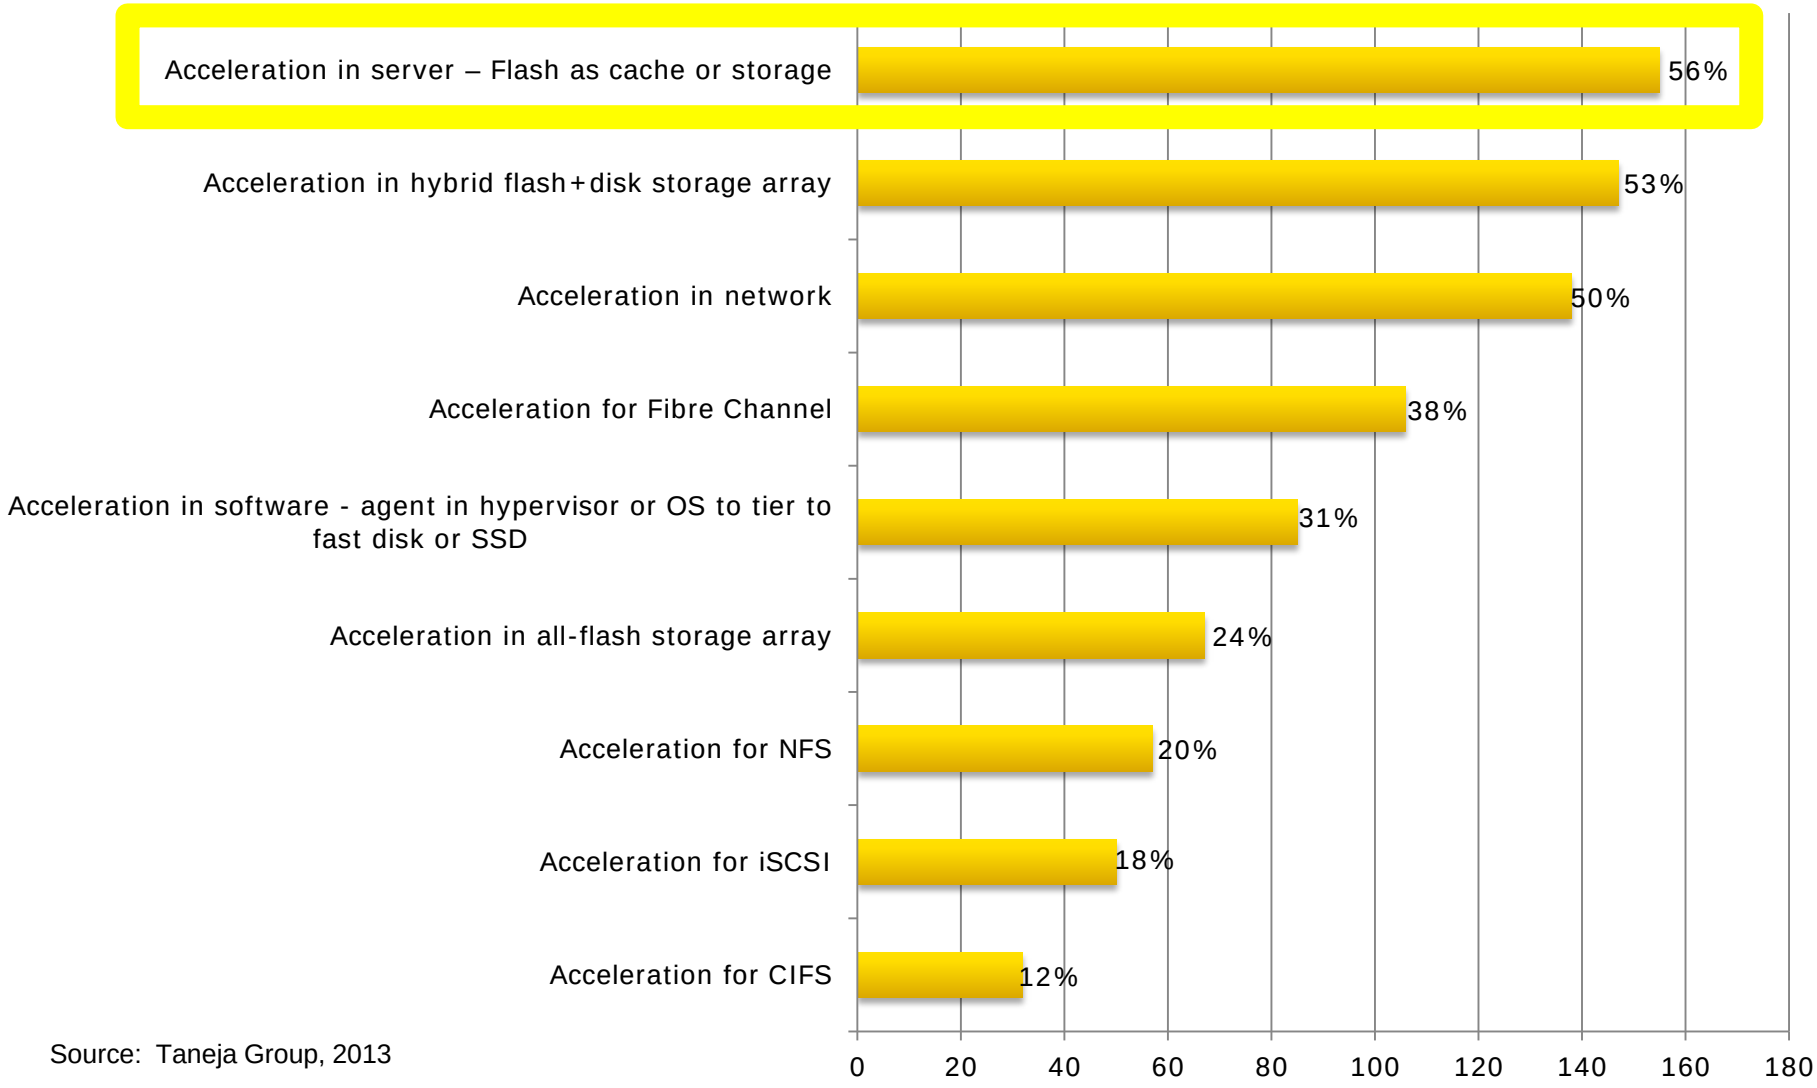

Source: Taneja Group, 2013

Flash Storage/Caching Solutions for the Data Center

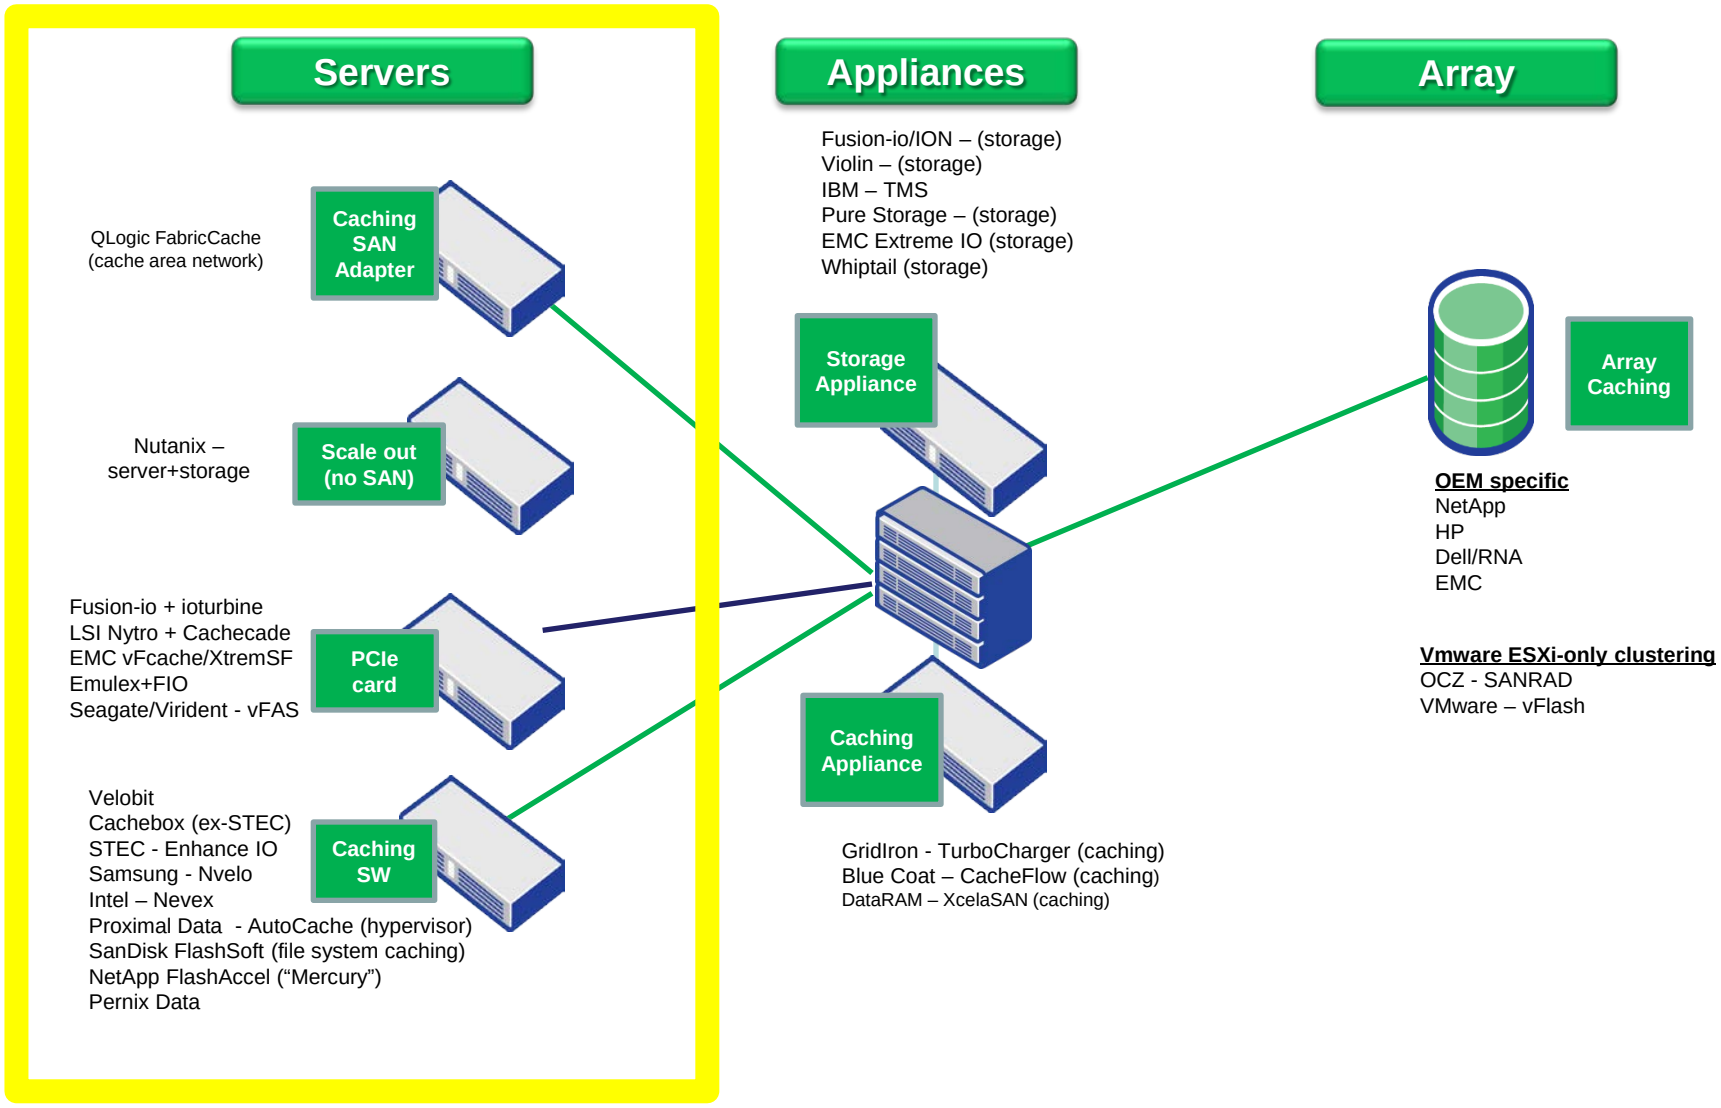

Where to Use Flash?

Array

PROs

- Easy to manage; shared
- Simple to implement
- Improved I/O performance

CONs

- Distance from server processor (latency)
- Limitations of storage array controllers
- Potential bottleneck

Storage cache / tier

Networks

PROs

- Centralized cache
- Shared cache by all servers
- Can cache network storage

CONs

- Single point of failure
- Appliances can be expensive
- Potential bottleneck

Caching appliance

Servers

PROs

- Critical data close to application and workload
- Improved storage performance and utilization

CONs

- Complexity & dependencies
- Cached data is captive in server “DAS SSD”
- Clusters not supported

Server cache

Performance Bottleneck on Storage Array SAN Ports

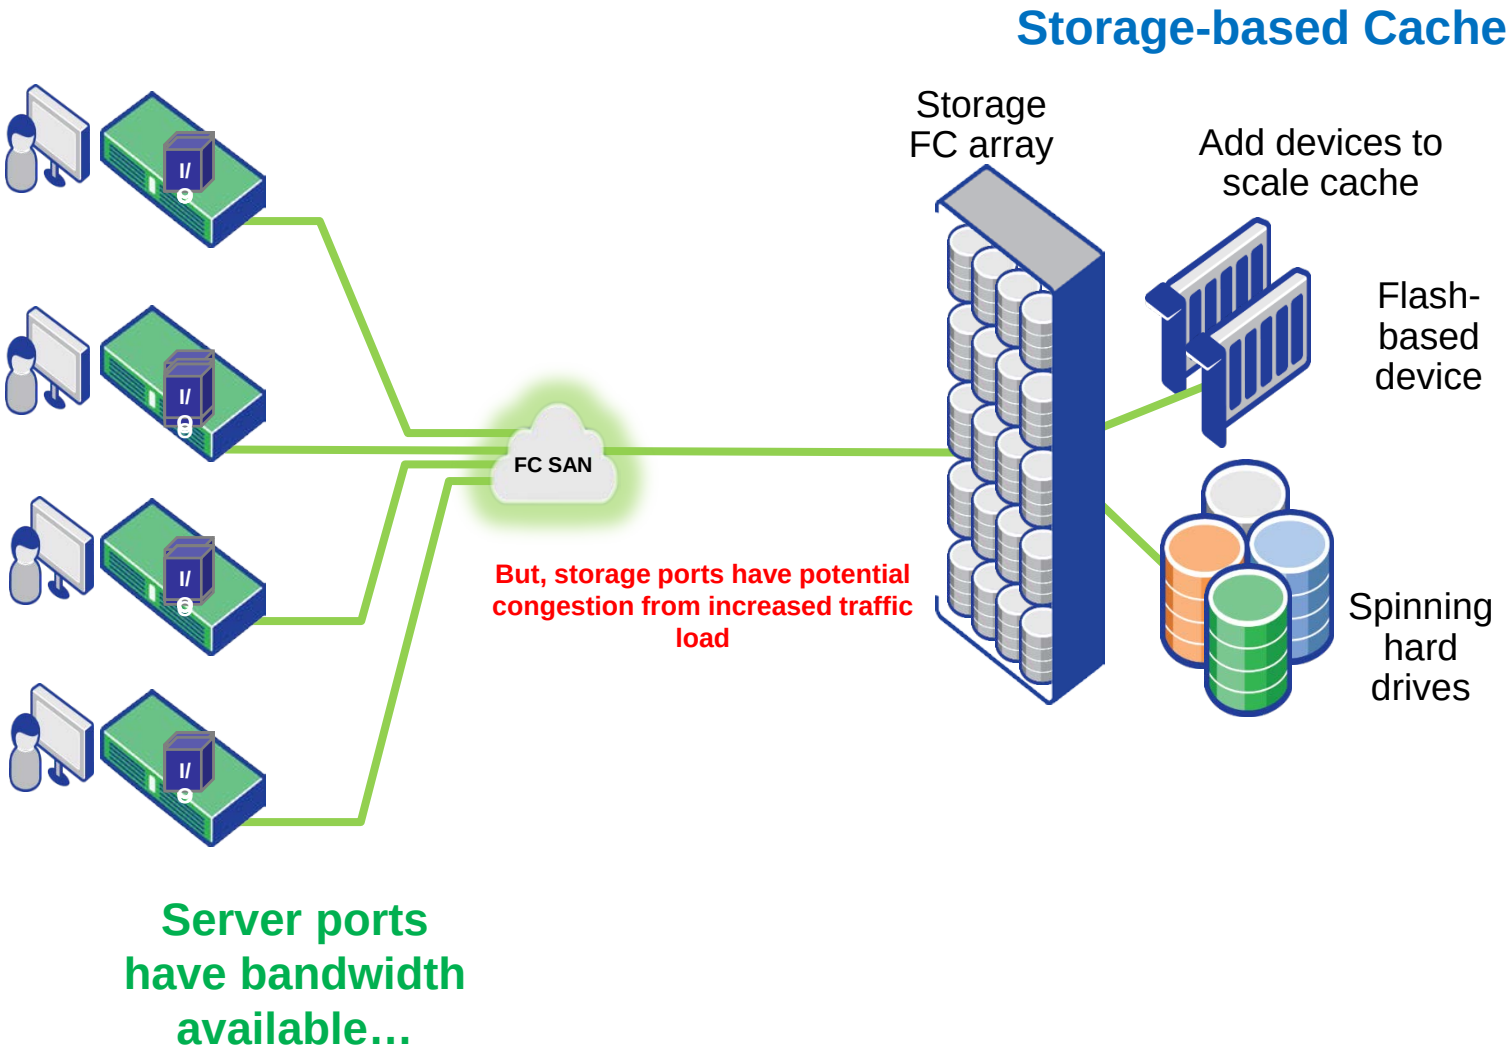

Flash-based Cache in Network Caching Appliance

Server ports have bandwidth available…

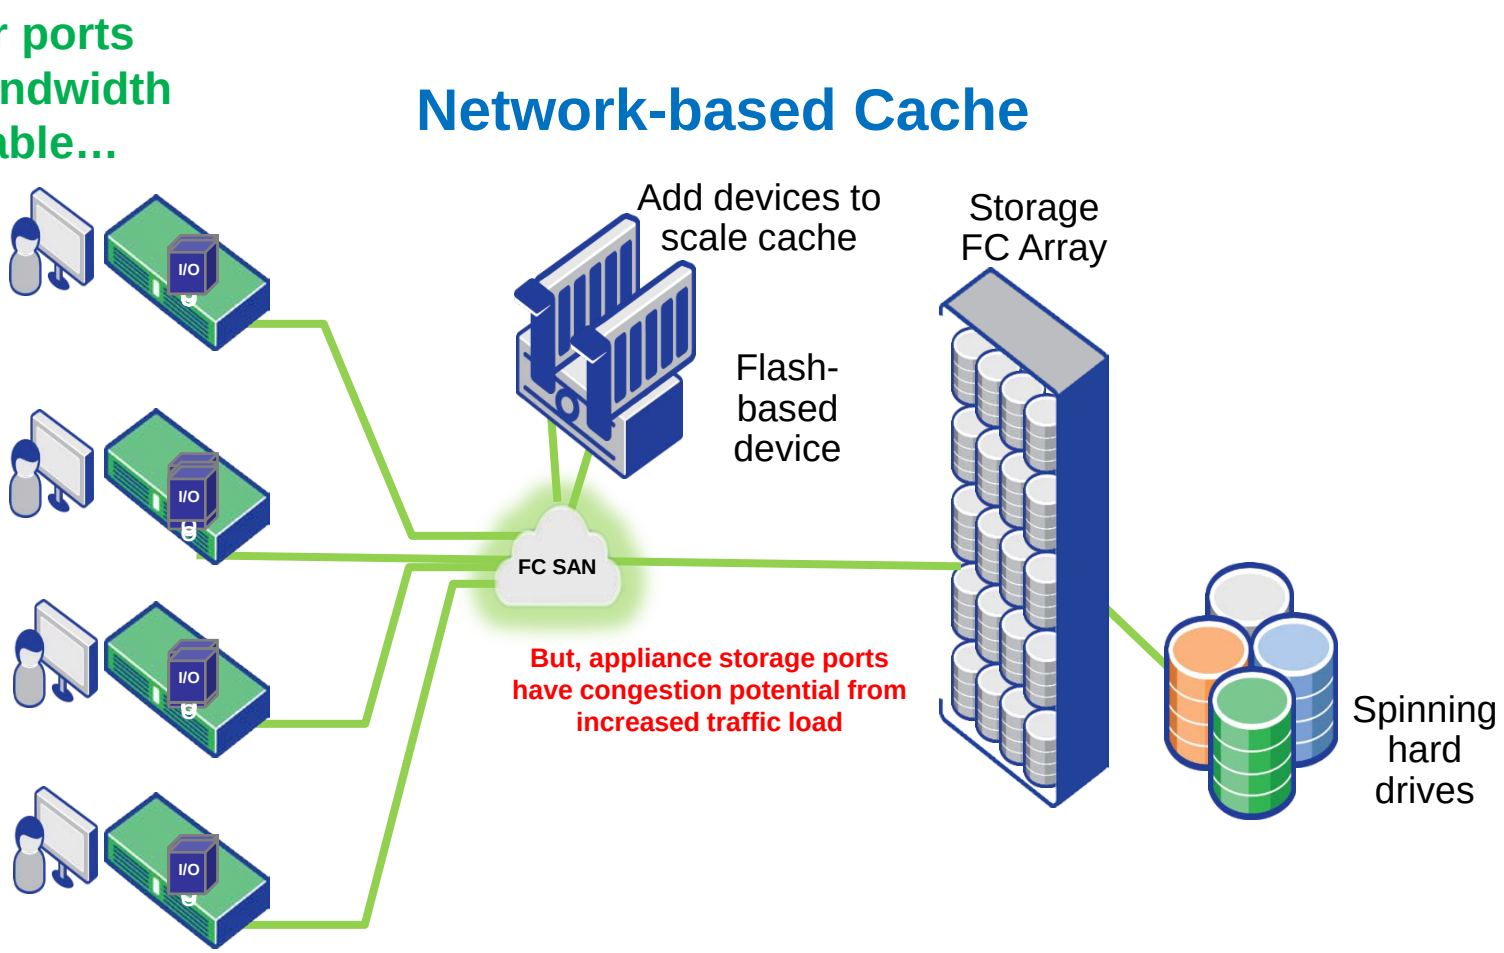

Legacy compute architecture

Compute to Disk Connection (physical)

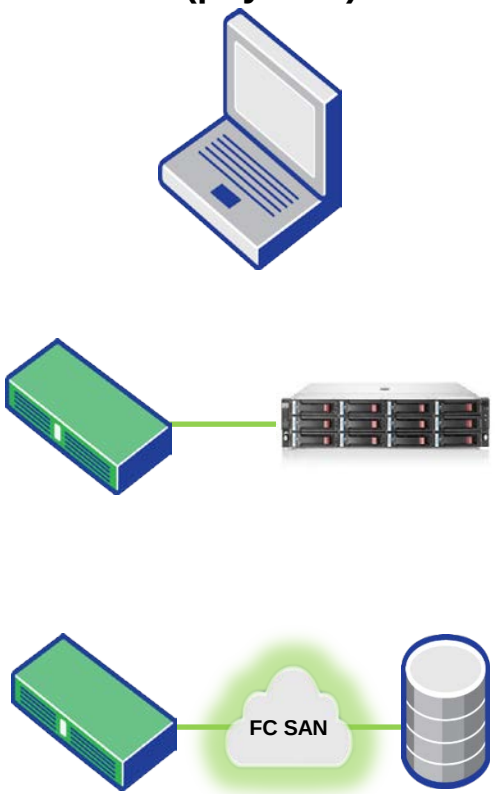

Workload to Storage Binding (logical)

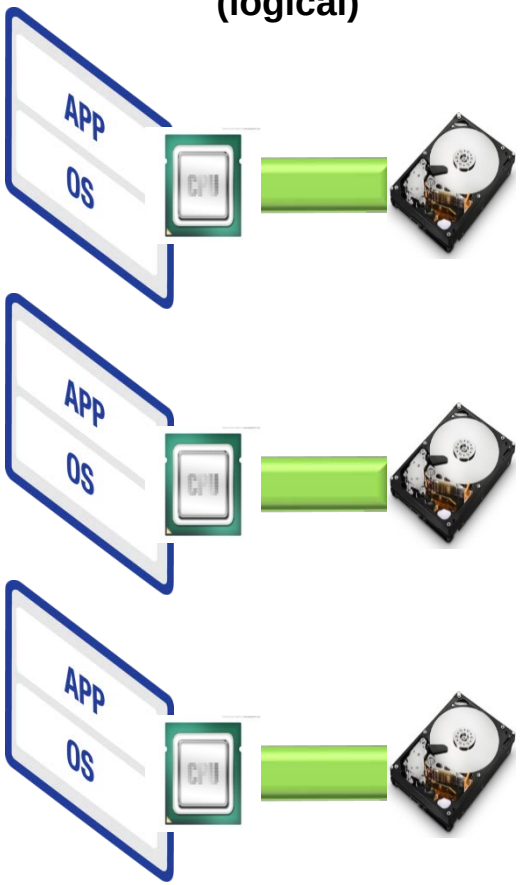

Legacy hardware environment
- static physical hardware
- static workload assignment

Legacy application environment
- workload tied to specific storage (LUNs)

Simple: Server to SAN with Standard SSDs/Flash

Compute to Disk Connection (physical)

Workload to Storage Binding

Server-based Flash Acceleration?

Yes! Workload and Storage bound to single server

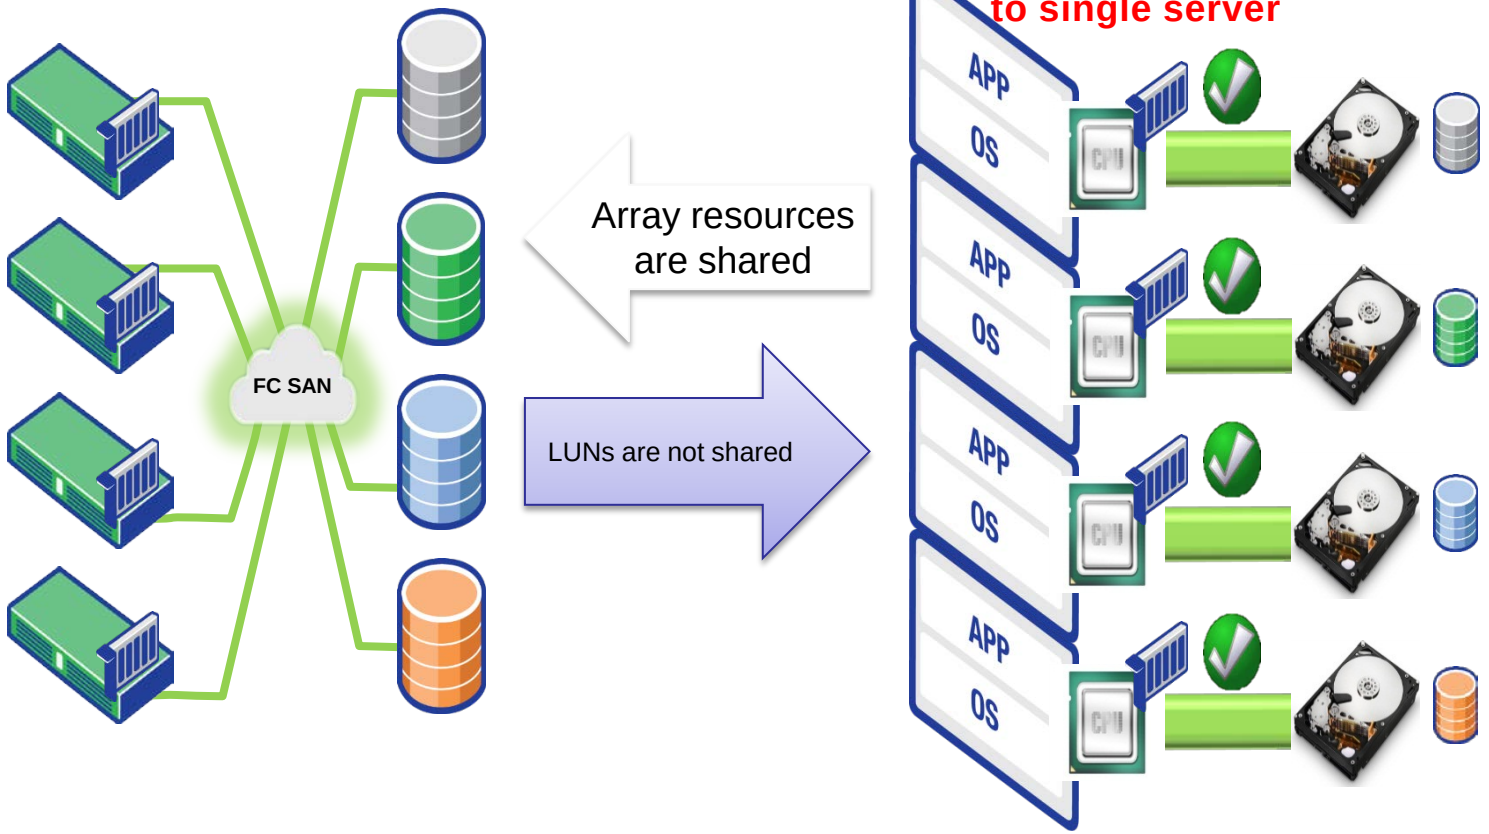

SAN enables sharing of storage resources

- Centralized provisioning and management

Legacy application environment

- ***LUNs not shared***

Multi-Server Virtualized: Server Cluster with SSDs/flash

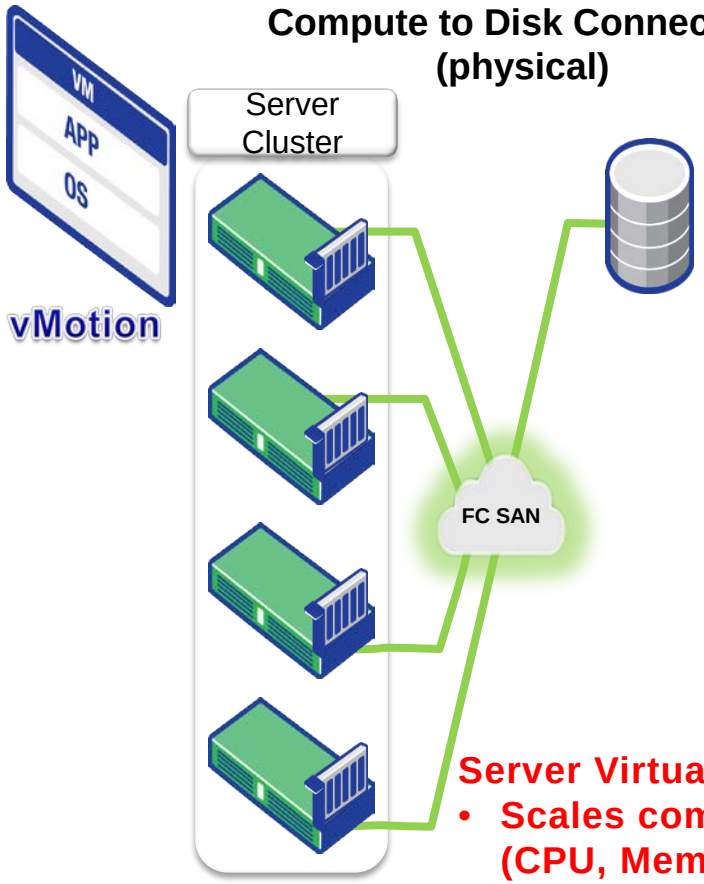

Server Virtualization:
- **Scales compute resources (CPU, Memory)**
- **Workload is now mobile**

Virtualized hardware environment
- elastic, scalable physical hardware
- dynamic workload assignment

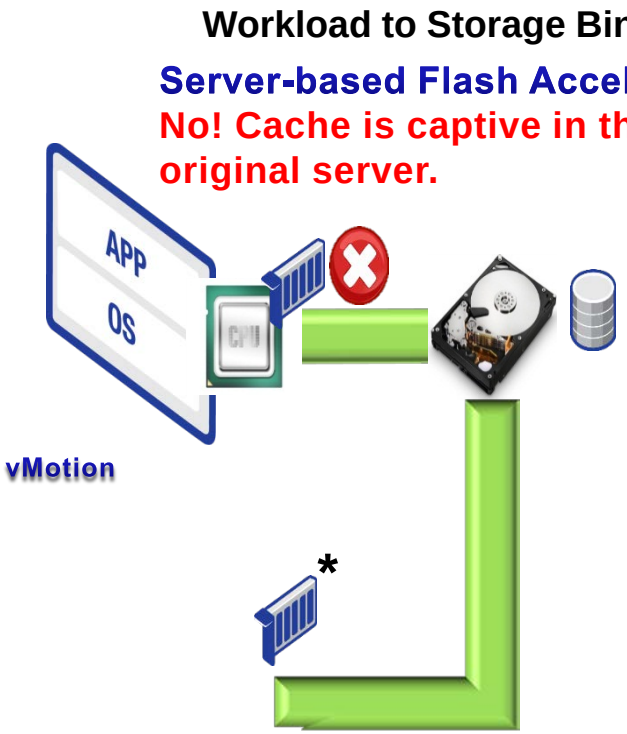

Workload to storage binding is still tied to the same specific LUN in the array

Server Cluster environment
- ***LUNs not shared***

* Empty cache at migration

Caching SAN Adapter

- Data is placed closest to the application
- Cache is transparent to hypervisor and guest applications
- Shared cache addressable by all servers
- No change to switches or storage arrays

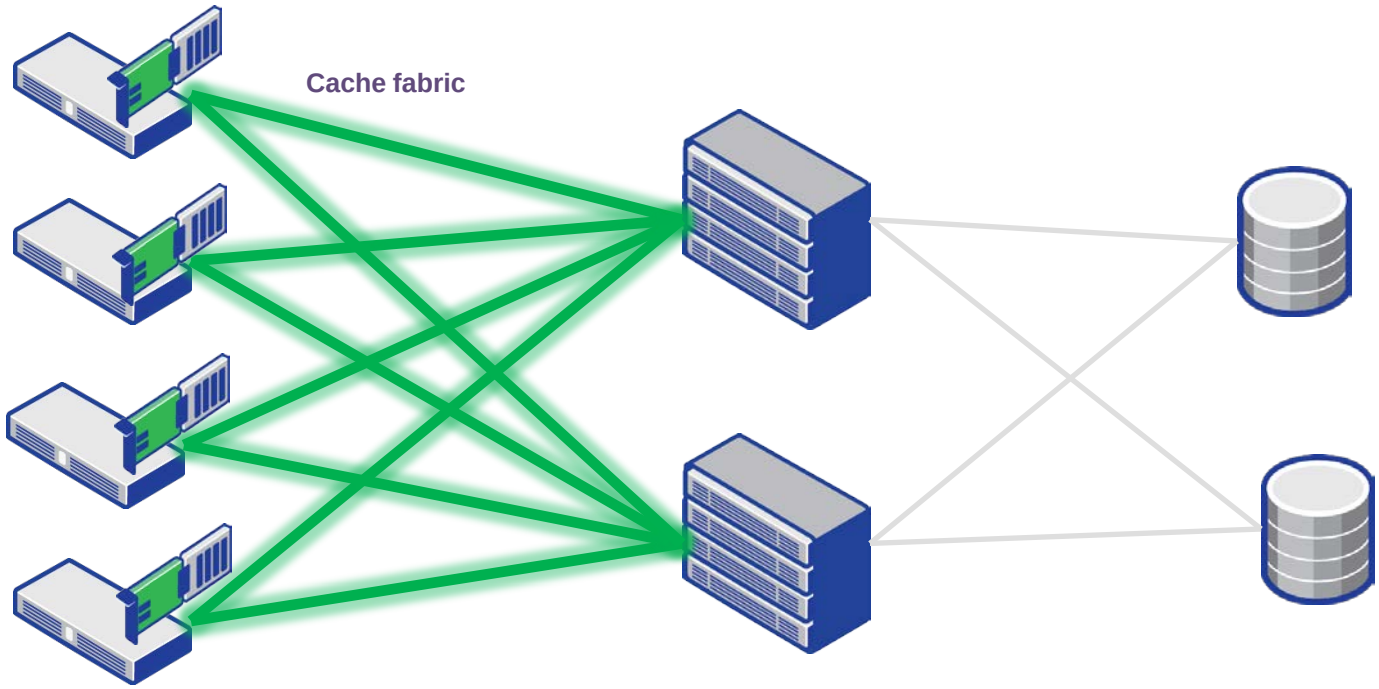

Multi-Server Virtualized: Caching SAN Adapter

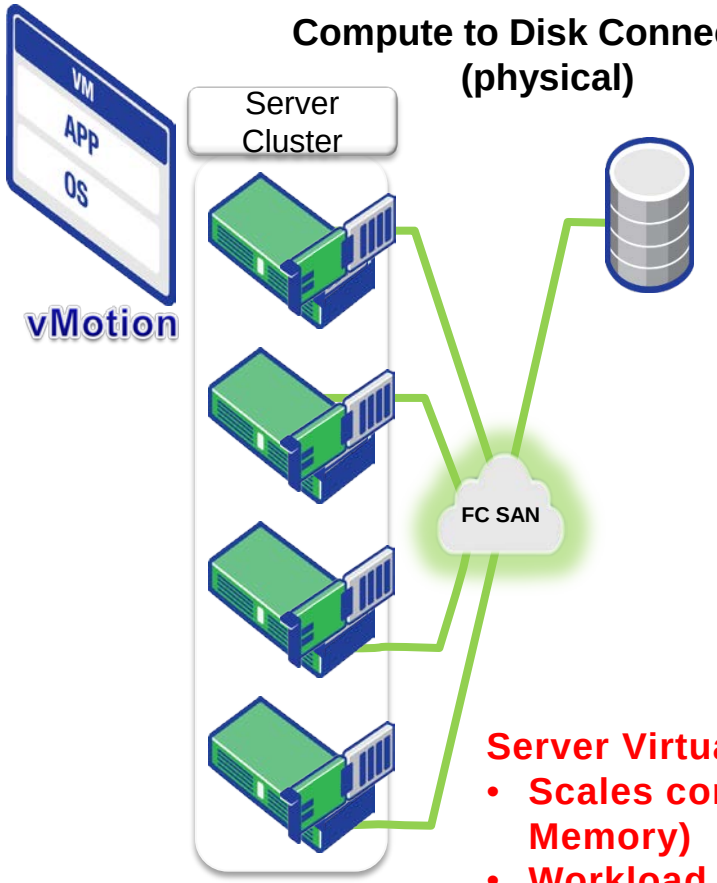

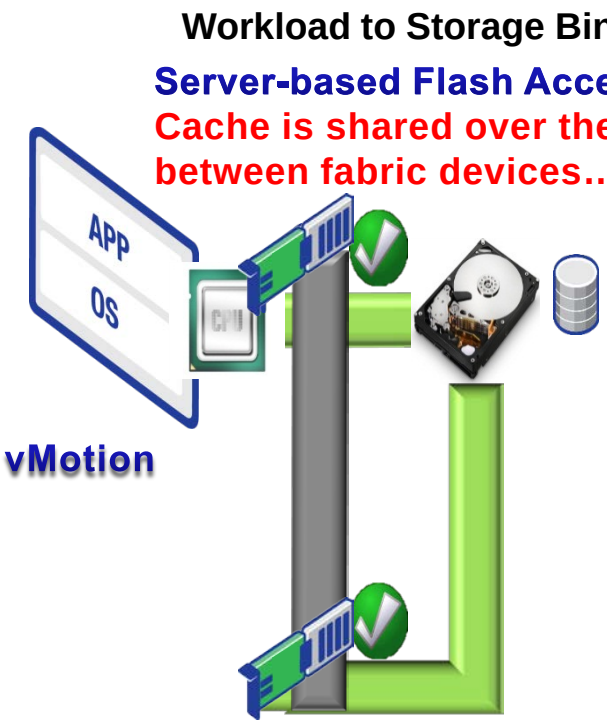

Server Virtualization:

- **Scales compute resources (CPU, Memory)**
- **Workload is now mobile**

Workload to storage binding is still tied to the same specific LUN in the array

Virtualized hardware environment

- elastic, scalable physical hardware
- dynamic workload assignment

Server Cluster environment

- ***LUNs not shared***

Application/Server Cluster: with Standard SSDs/flash

Compute to Disk Connection (physical)

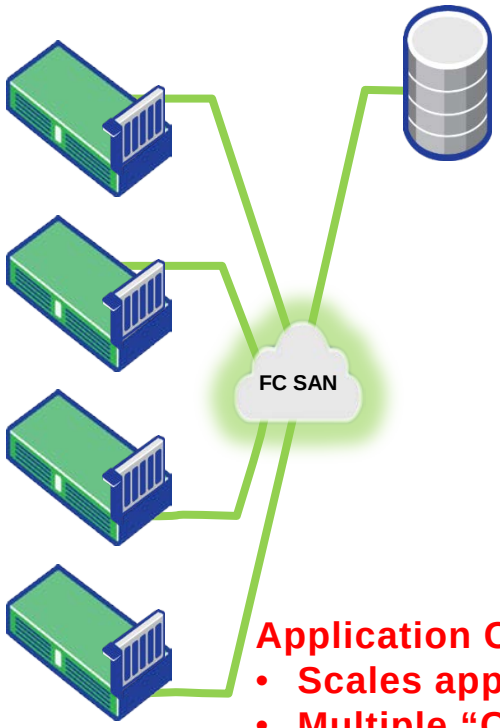

Application Clustering:

- **Scales application resources**
- **Multiple "Cloned" server nodes**
- **All nodes require access to the same LUN(s)**

Workload to Storage Binding

Server-based Flash Acceleration?

No. Cache is not coherent between devices…

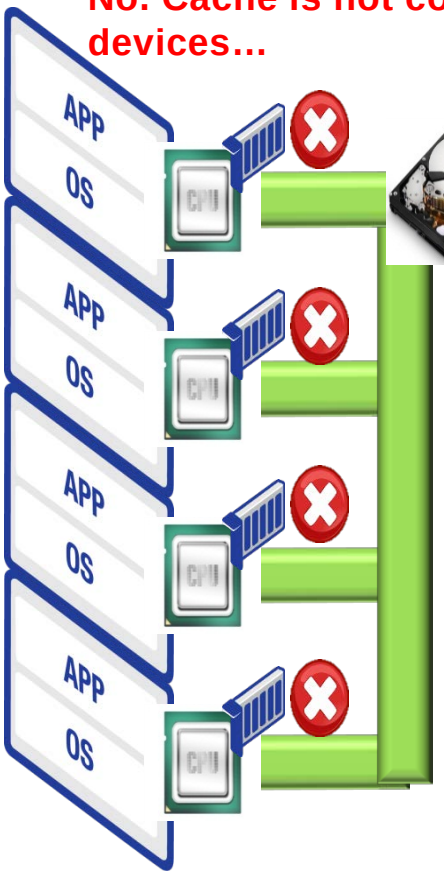

Workload to storage bindings for all nodes are tied to the same specific LUN(s) in the array

Application Cluster file system must manage the LUN sharing…

SAN enables sharing of storage resources

- Centralized provisioning and management

Clustered application environment

- ***LUNs must be shared***

Application/Cluster Server: Caching SAN Adapter

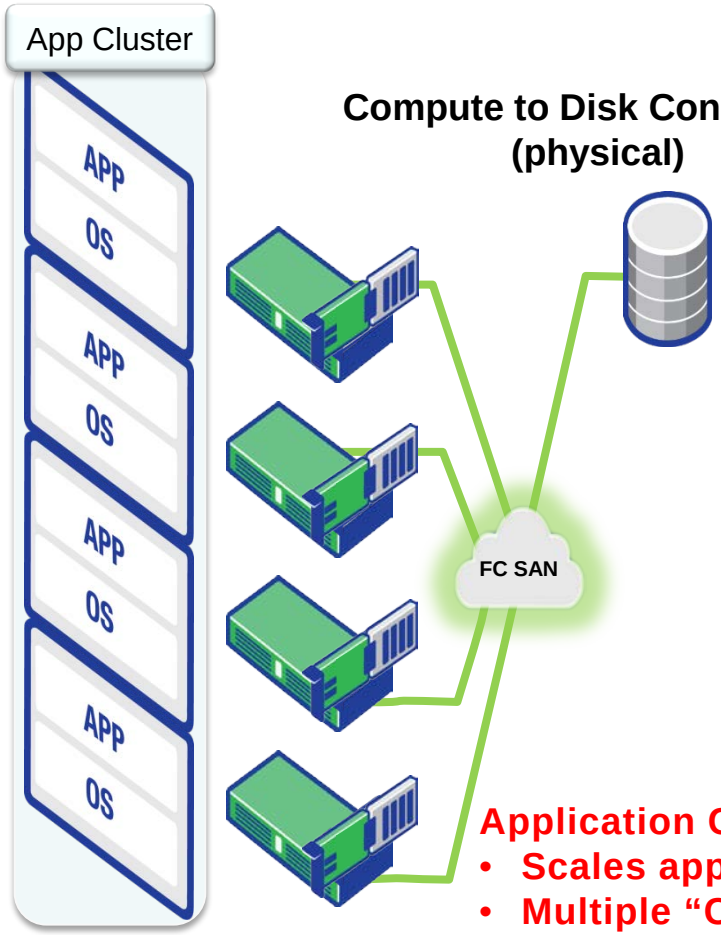

Application Clustering:

- **Scales application resources**
- **Multiple "Cloned" server nodes**
- **All nodes require access to same LUN(s)**

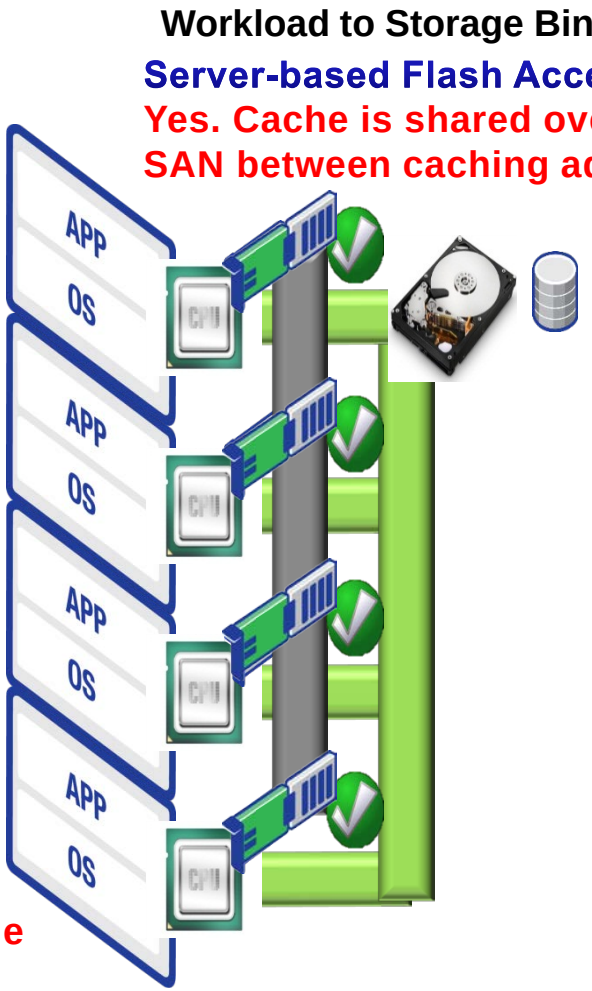

SAN enables sharing of storage resources

- Centralized provisioning and management

Clustered application environment

- ***LUNs must be shared***

Virtualized and Clustered: Caching SAN Adapter

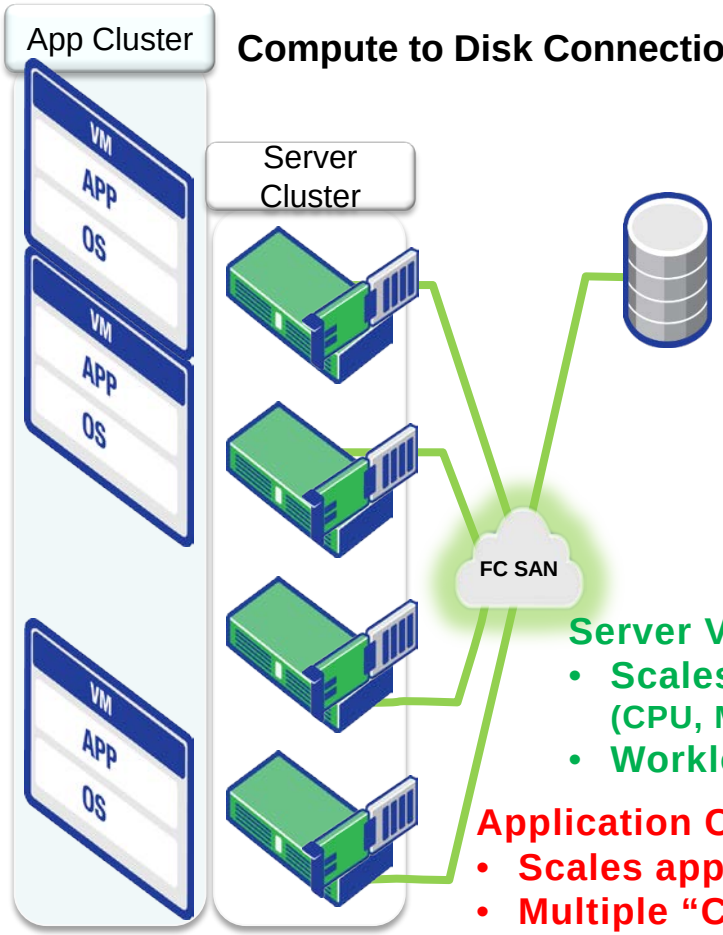

Server Virtualization:
- **Scales compute resources (CPU, Memory)**
- **Workload is now mobile**

Application Clustering:
- **Scales application resources**
- **Multiple "Cloned" server nodes**
- **All nodes require access to same LUN(s)**

Workload to Storage Binding

Server-based Flash Acceleration?
Yes. Cache is shared over the SAN between caching adapters

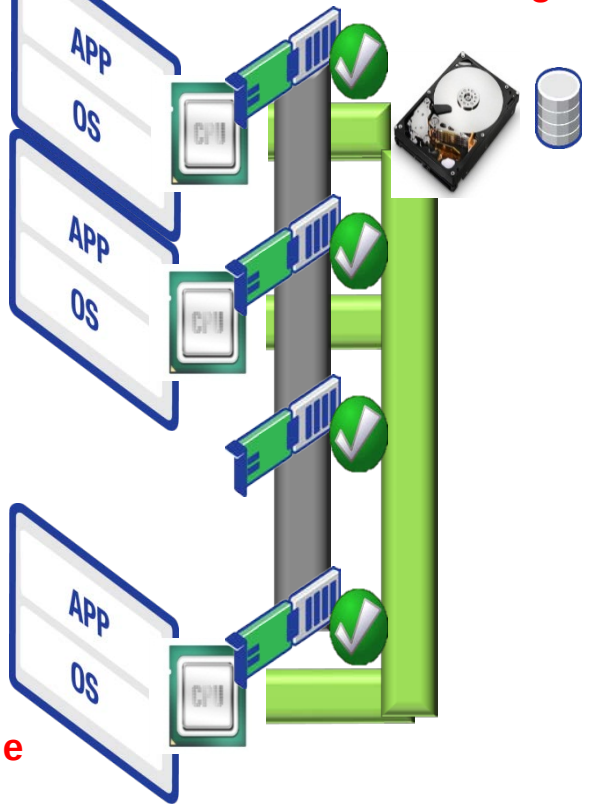

SAN enables sharing of storage resources
- Centralized provisioning and management

Clustered application environment
- ***LUNs must be shared***

Enterprise Acceleration Requirements

- Must be simple to deploy and manage – minimal SW footprint with drivers, caching SW in host
- Able to accelerate enterprise application software including cluster and virtualized
- Cache must be shared – coherent, no duplication
- Non-disruptive to existing network/SAN, policies, management tools
- Proven, stable technology – can't take chances with critical data

In Closing…

- Many implementations available to accelerate
- Move to server-based solutions
- Caching SAN adapter is a new type of solution for the enterprise
- Wide variety solutions available for simple applications
- More complex applications require more capable solution
- Keep it simple, non-disruptive, cost effective and proven