Clustering In A SAN For High Availability

Steve Dalton, President and CEO
Gadzoox Networks

September 2002

HP WORLD 2002
Conference & Expo

Agenda

- What is High Availability?
  - The differences between High Availability, System Availability and Serviceability
- Clustering in a SAN environment for High Availability
- Summary

High Availability Equals

- Data is never lost!!
  - Ports are never down
- System Level of Five 9's, where you have;
  - Redundancy of all active components (devices) within the SAN
  - A communication mechanism reporting the fitness of all components (devices)
  - A way to shut down a perceived non functioning component (device) (if possible)
- System (Data) Availability, where you have;
  - A graceful bringing up of the redundant replacement, with minimum time lapse and minimum affect on other system components (devices) that may exhibit unpredictable behavior

Levels of High Availability

- System Level of High Availability
  - Requires duplicating all elements and paths in the system, and having failover mechanisms for switches, paths, servers and storage.
  - In a high availability system there is more than just hardware redundancy. Applications, protocols, configuration and management within the redundant hardware environment complete the system level for high availability.
- Expensive and not always necessary
  - Especially in the SAN entry-level

System Availability

- Achieve high system availability in a clustered environment with redundant switches and storage
  - Allows access to applications and data regardless of type of network, storage or server failure
  - Examples of highly available systems regardless of where and type of failure

Increasing Application Availability

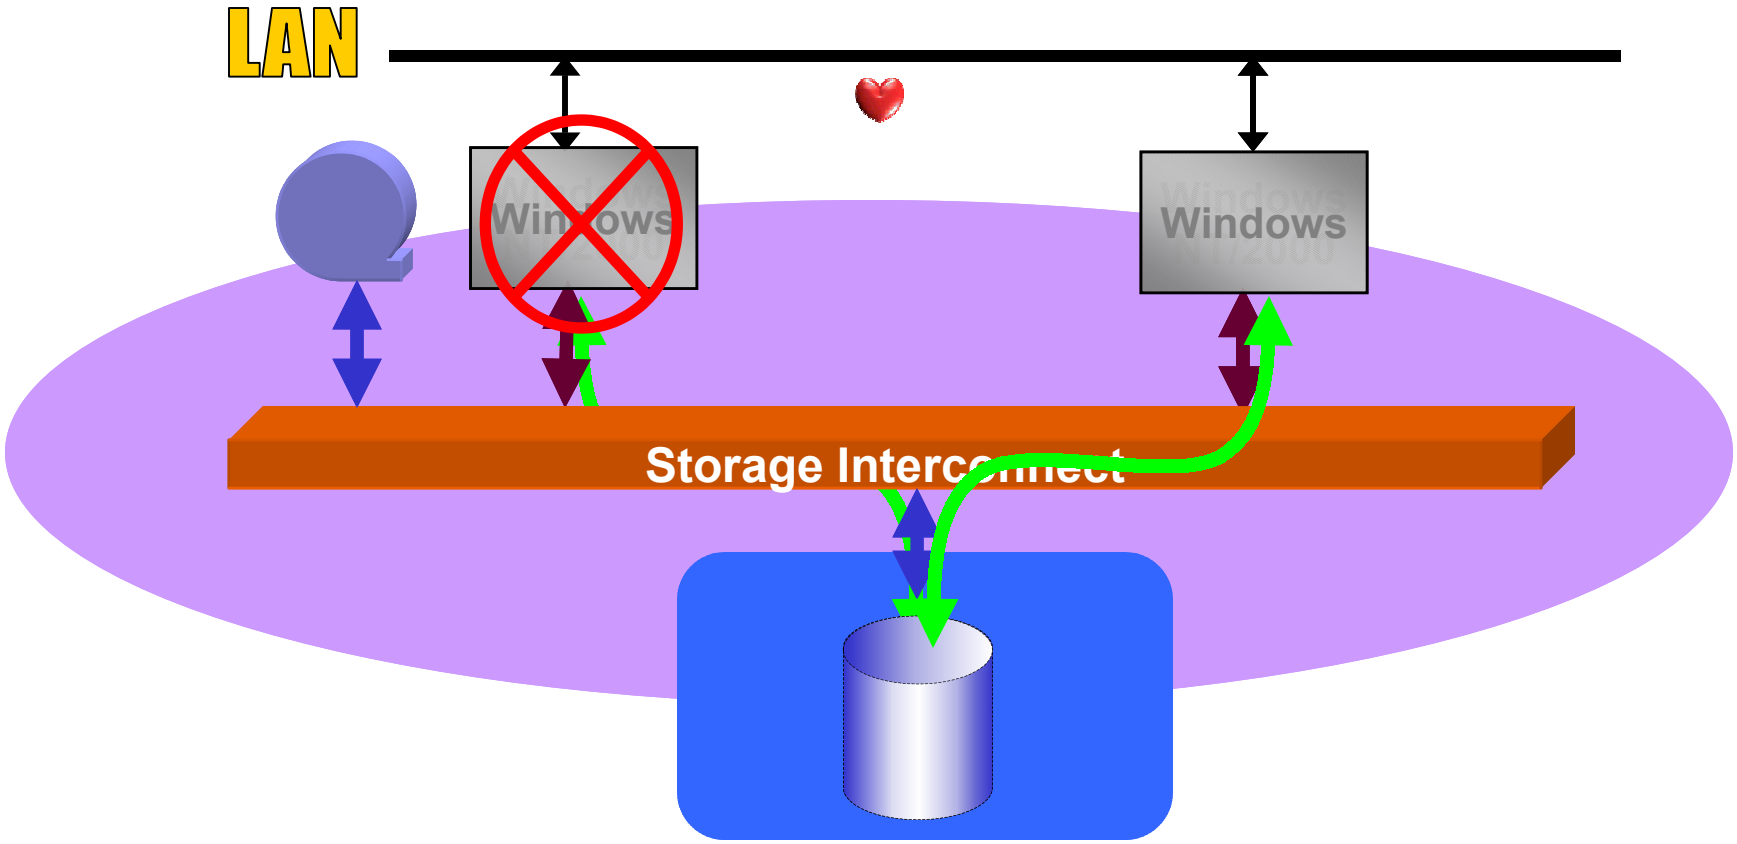

- Business operations protected from application failure
- Direct connections limit scalability

Increasing Connection Availability

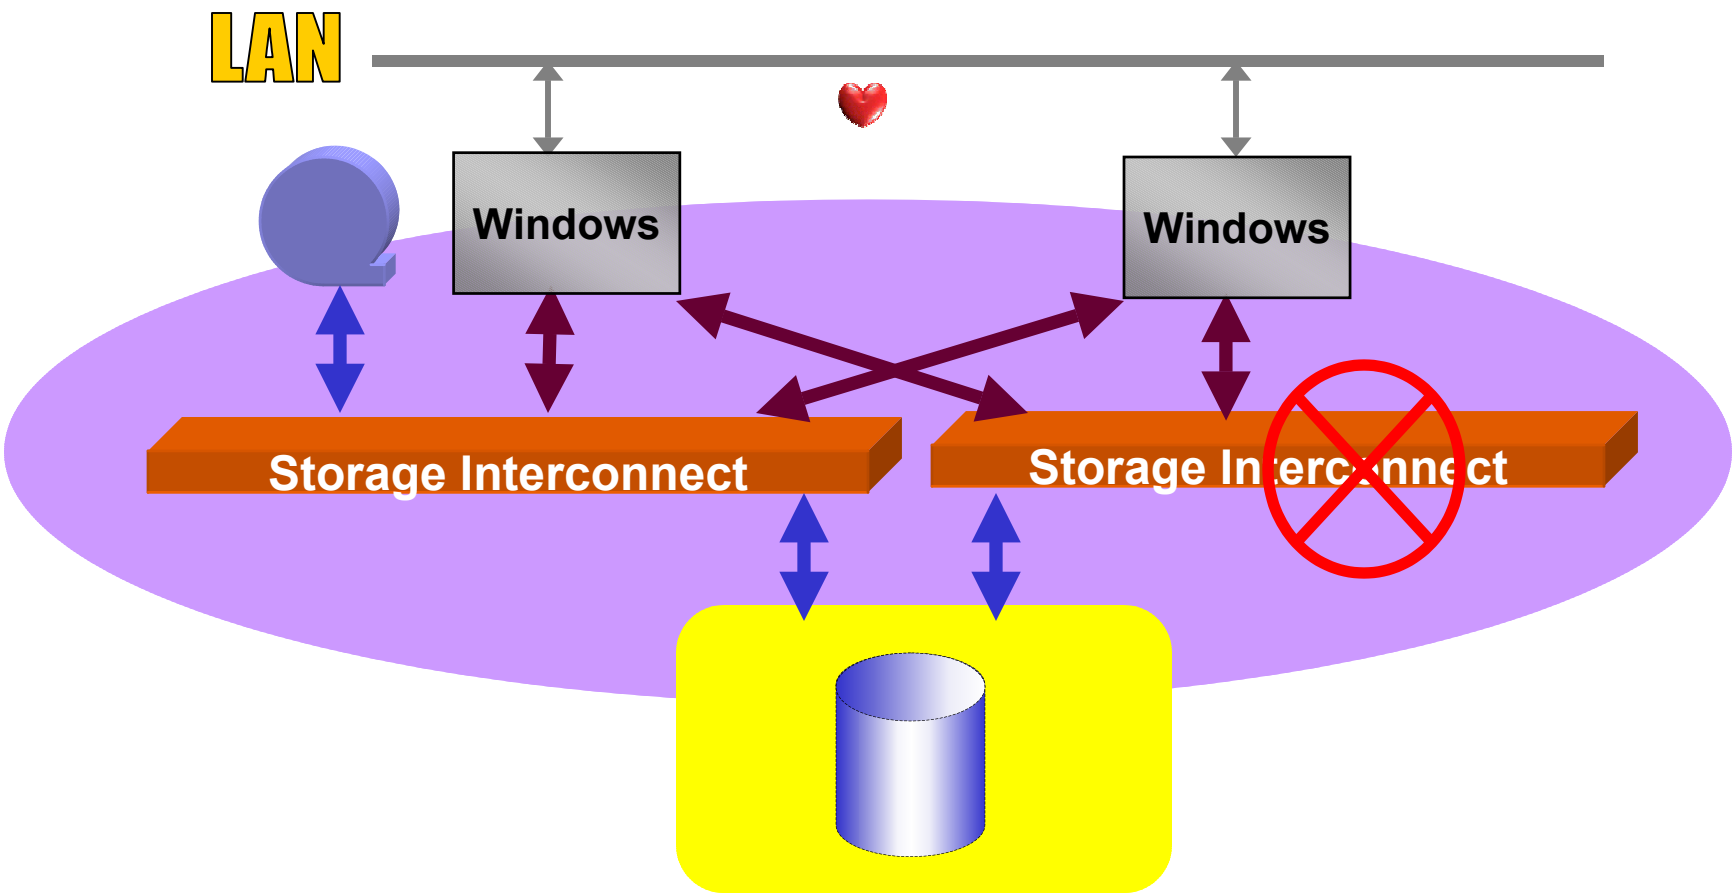

- Operations protected from connection failure
- Leverages switched SAN infrastructure

Increasing Data Availability

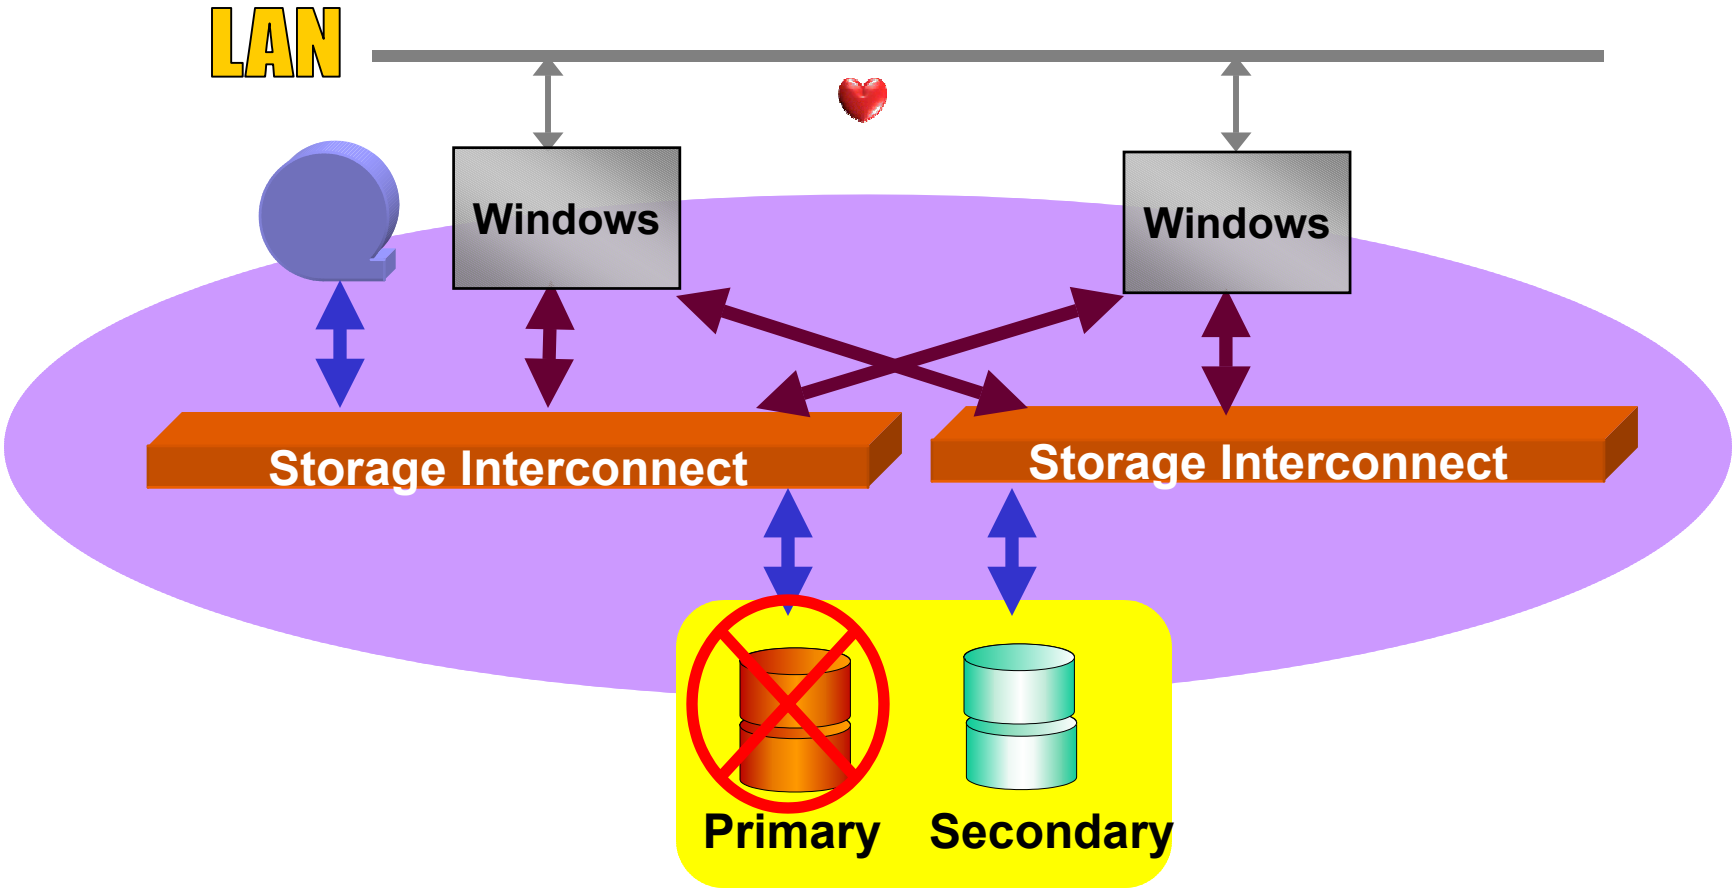

- Business operations protected from storage failure
- Leverages switched SAN infrastructure

Do You Need This?

- Dual network connections, dual host bus adapters and clustering with failover
  - Techniques that improve availability
  - Managed via software that detects actual failures, separates them from false failures, and executes a failover of the path, storage, HBA or server as needed, with minimum disruption.
- Sophisticated applications, supporting SAN redundancy clustering and failover
  - Effective in improving availability, and minimizing the redundancy needed in the SAN to achieve high availability.
- Or Serviceability?

Serviceability Enables High Availability

- System Availability at 100%
  - Ports are never down
- Very inexpensive
  - Reduce time to repair
  - No-cost FRU's
    - Allow a greater granularity for FRU replacement
- Key serviceability considerations
  - Cost vs. service time
  - Hot-swappable parts
    - GBICs
    - power supplies
    - fans
  - Ease of access
    - On multiple sides
    - Replace without disrupting other items

Clustering In A SAN Environment For High Availability

- Servers/Hosts
  - Clustering software from OS vendor or ISV
  - Can be active-active, or active-passive
    - Depends supported application
- Redundant controllers (ie HBA's/RAID)
  - Needed for path failover
  - Hot standby or load balancing
    - Load-balancing software in Open SAN OS
- Switches
  - Needed for path failover
  - Self-healing trunks
    - Software in Open SAN OS
- Storage
  - JBOD, RAID, Tape, Distance, Mirroring ,etc…

Start with Clustered Servers

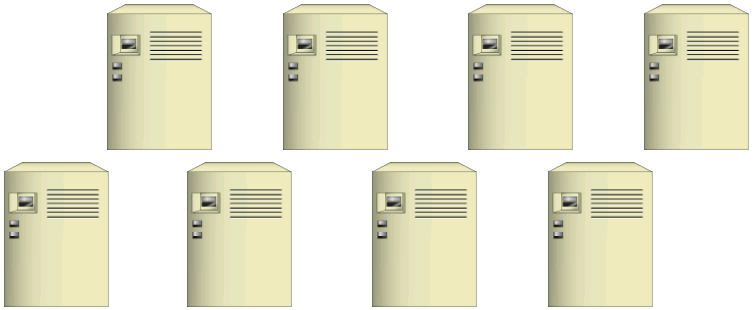

- Use clustering software from OS vendor or ISV
- Can be active-active, or active-passive
  - Typically depends on what the application can support

Install Two HBAs per Server

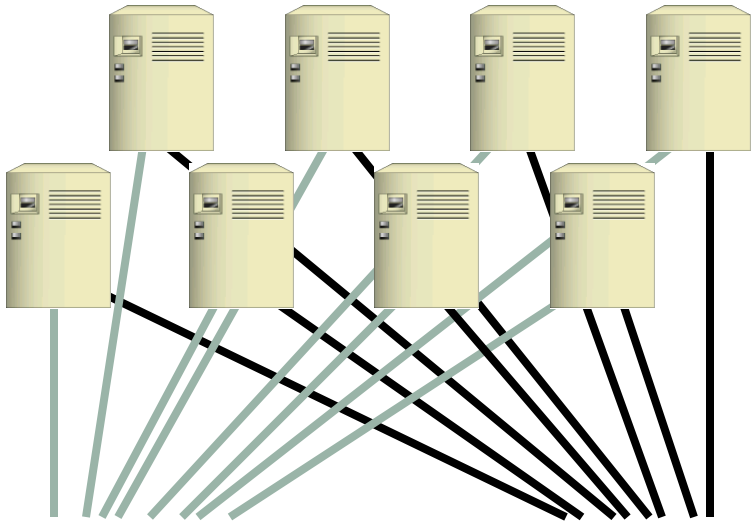

- Needed for path failover
- Hot standby or load balancing
  - Use load-balancing software if available

Need at Least Two Switches

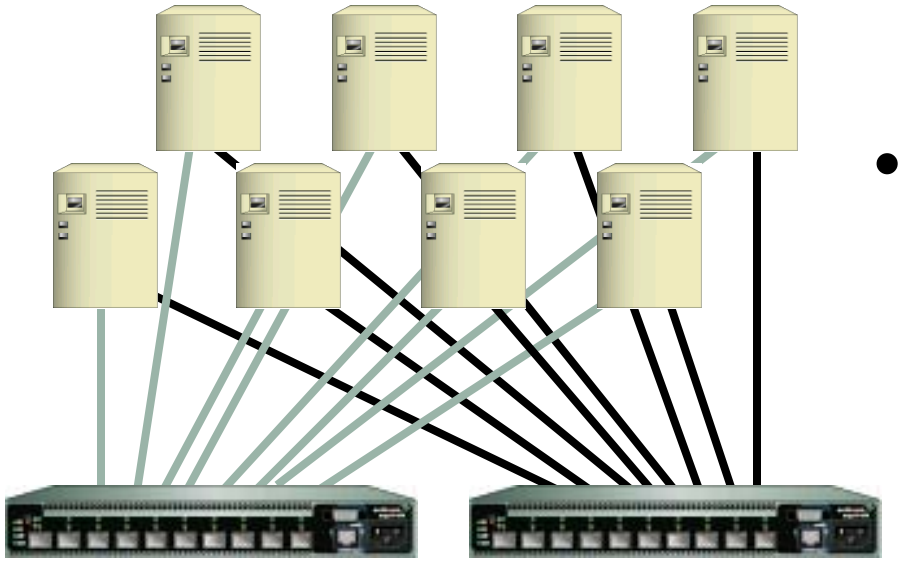

- Separate, redundant paths

Need at Least Two Switches

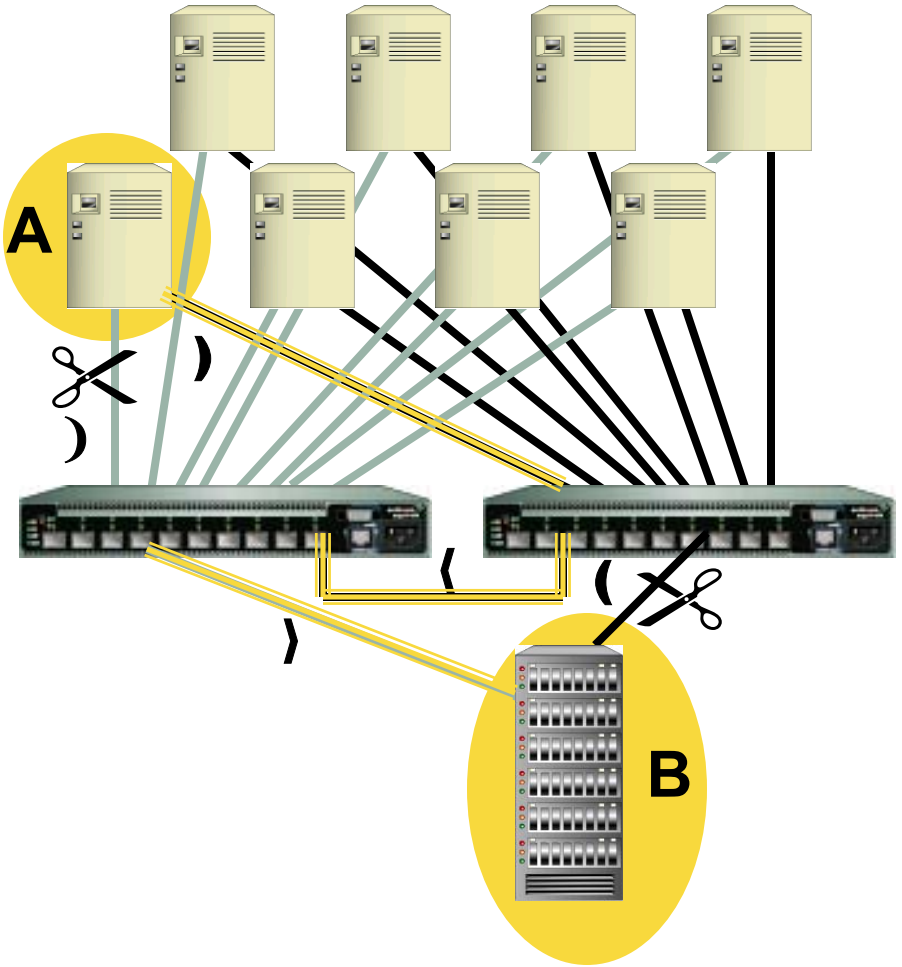

- With two switches, there is an advantage to having switches connected when they support self-healing trunking

Add the Storage Systems

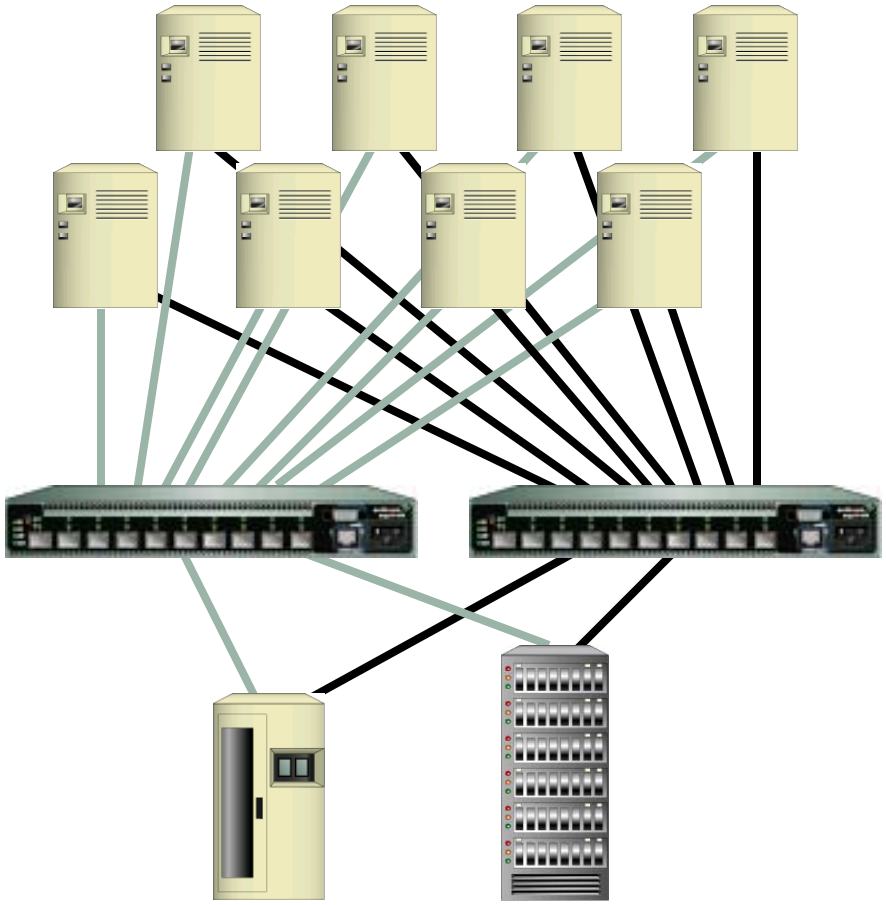

- Dual-ported RAID
- Tape system
  - LAN-free backup
- Exploit Fibre Channel distance capabilities
  - 10Km plus
  - Remote mirror or backup
- Direct error messages

Typical Business Deployment

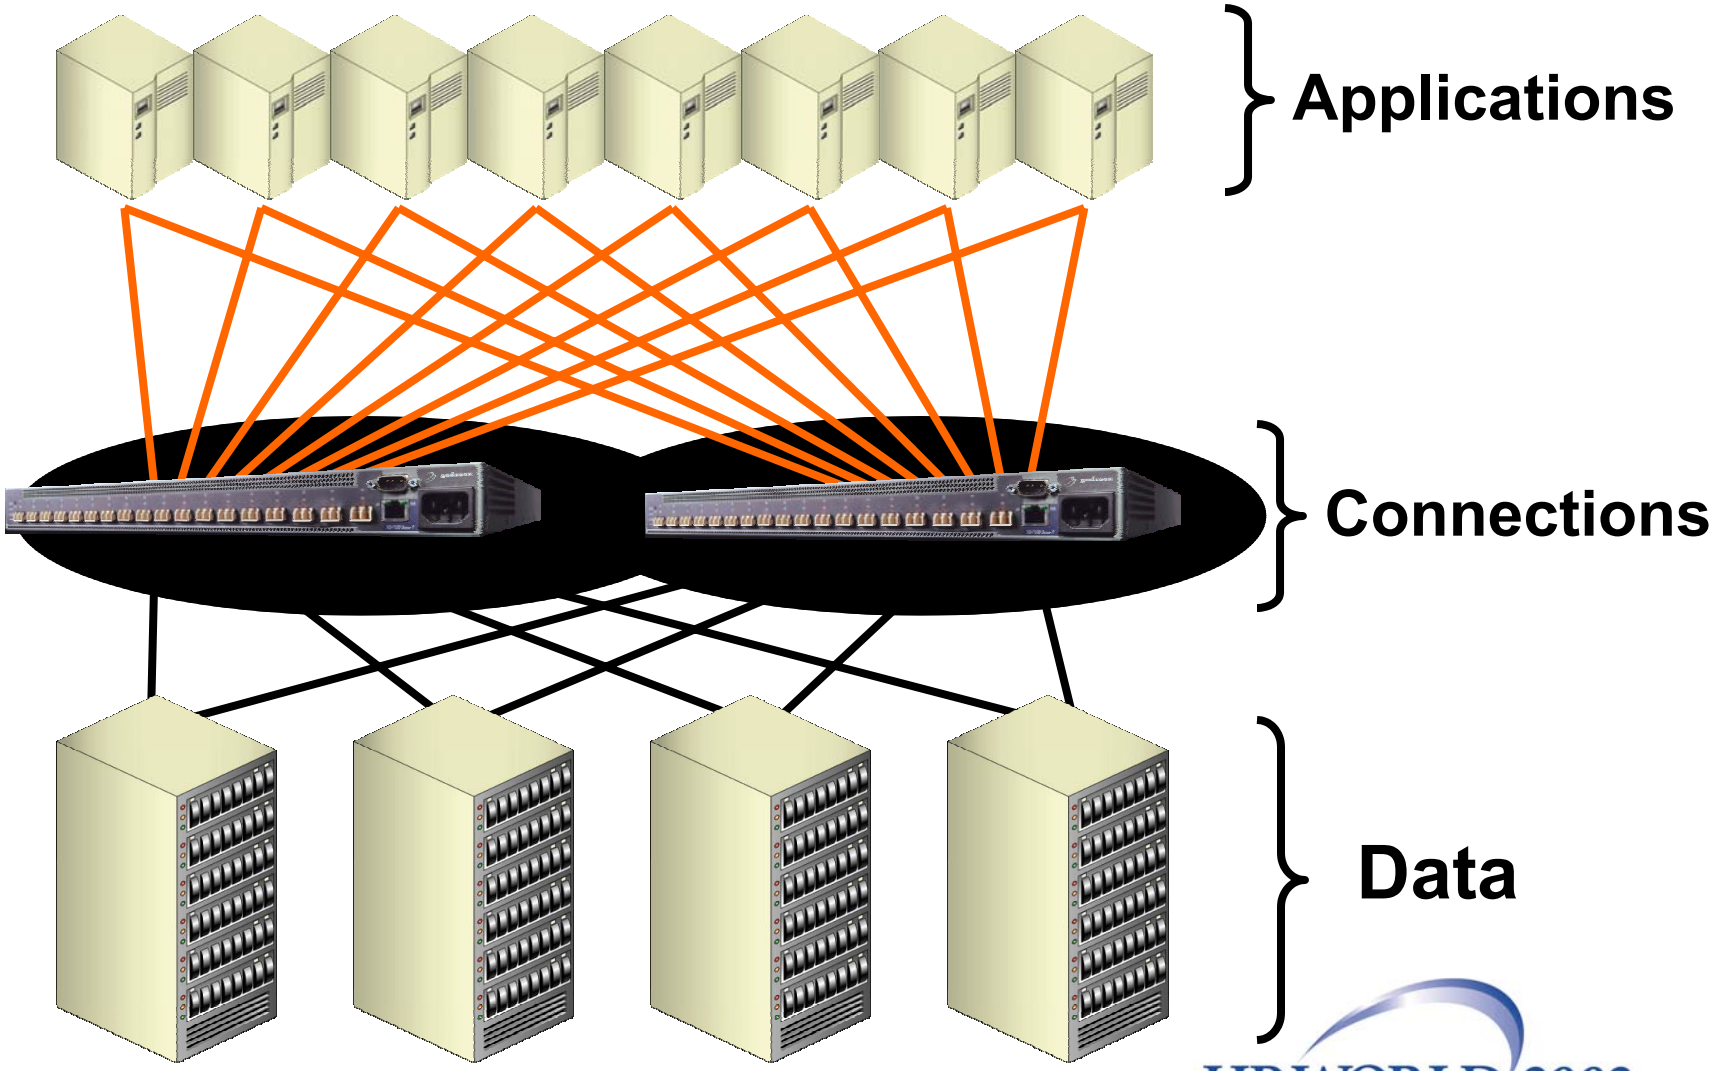

Scale Using the Same Principals

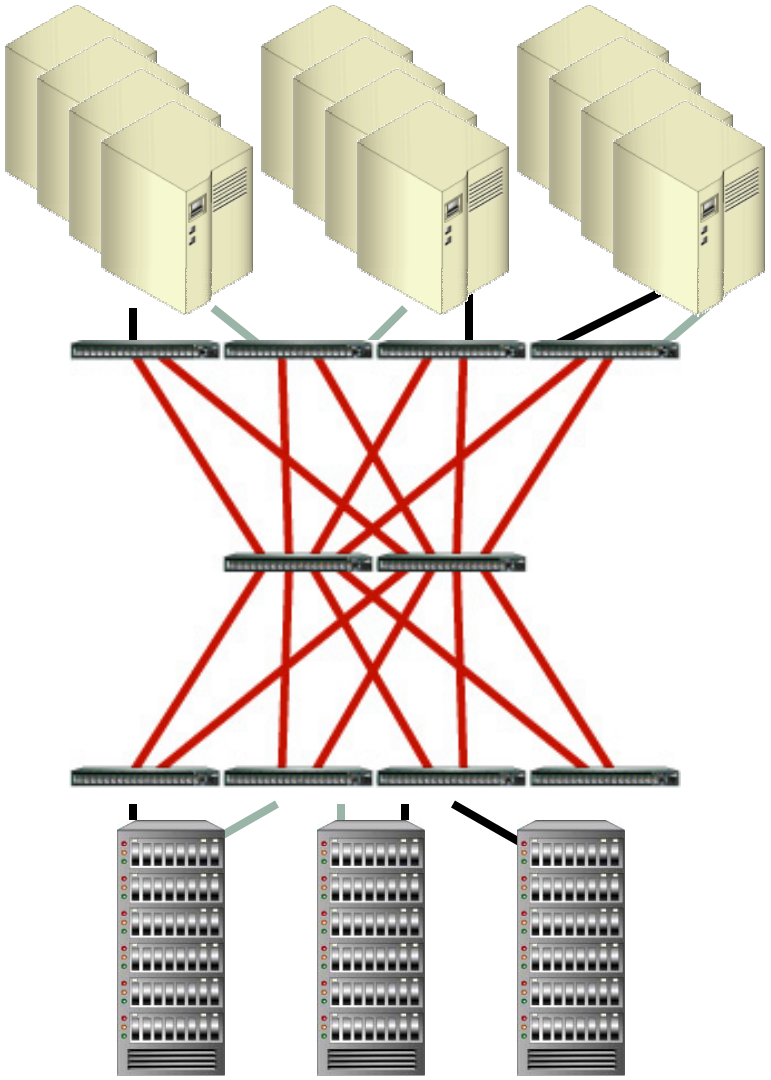

- Redundant components
- Separate, redundant paths
- Mesh, mesh-tree, backbone, or core & edge topologies

Scale Using the Same Principals

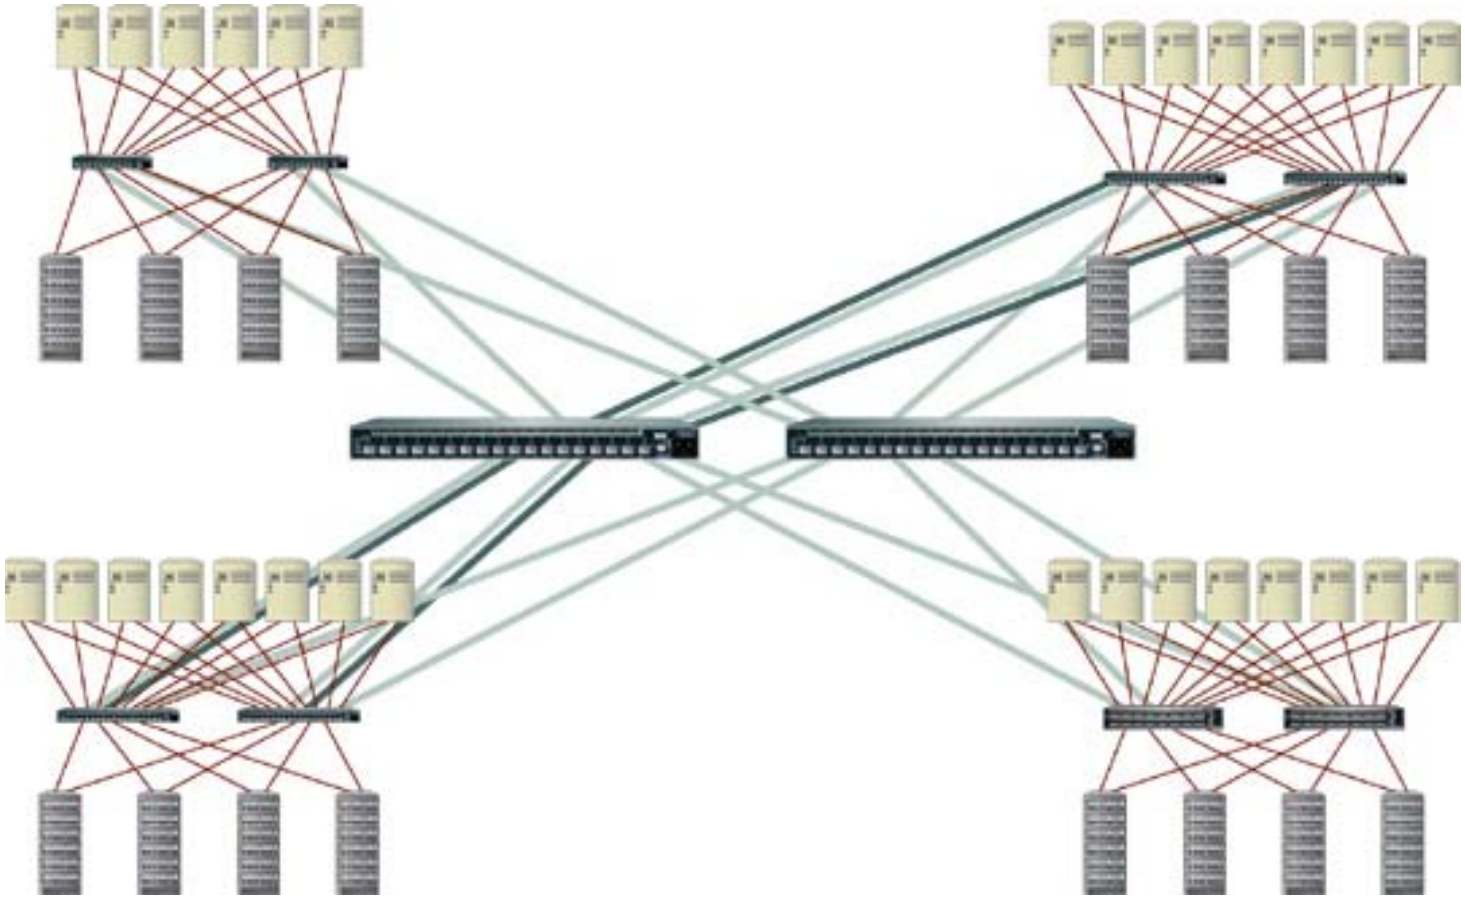

- Core and edge with flexible trunking

Entry-level Typical Deployment

- ***RAID technology protects from disk failure***
- ***Redundant controllers protect from card failure***
- ***Put the switch in the Array to reduce complexity***

MSA 1000

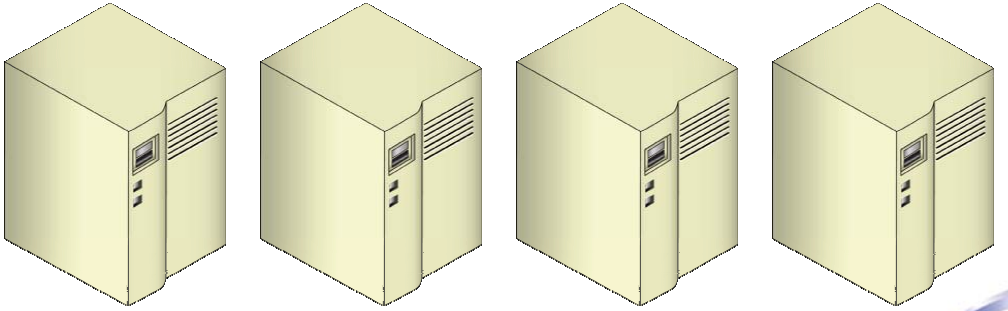

High Availability Considerations

- A system solution
  - Little benefit to placing an “HA” product in a non-HA topology
- Redundancy is best
  - More than n+1
  - 2 components - you can lose one and still have the application access the data
- SANs provide the redundant connections between servers and reliable external storage devices

Key Switch Considerations for High Availability

- Cost
  - Five 9's High Availability vs. Continuous Data Availability
- Five 9's High Availability→No Service Time Needed
  - Dual Active-Active internal data paths
  - Redundant internal modules and connectors
    - Power Supplies, Fans, etc…
  - Redundant components on the data-path
  - Dependencies between ports
- Continuous Data Availability→Scheduled Service Time (FRU's)
  - Dual Active-Active internal/external data paths
  - Minimum internal modules and connectors
  - Minimum dependencies between ports
  - Minimum components on the data-path
  - Distribution fabric architecture
  - Power supply redundancy
- Remote download F/W with dual image buffers for download error recovery
- Configuration download and up-load for quick switch mirroring
- Trunking for links' load sharing and seamless failover traffic

Thank You!