



Education

Data Center Convergence

Ahmad Zamer, Brocade

- ◆ 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.

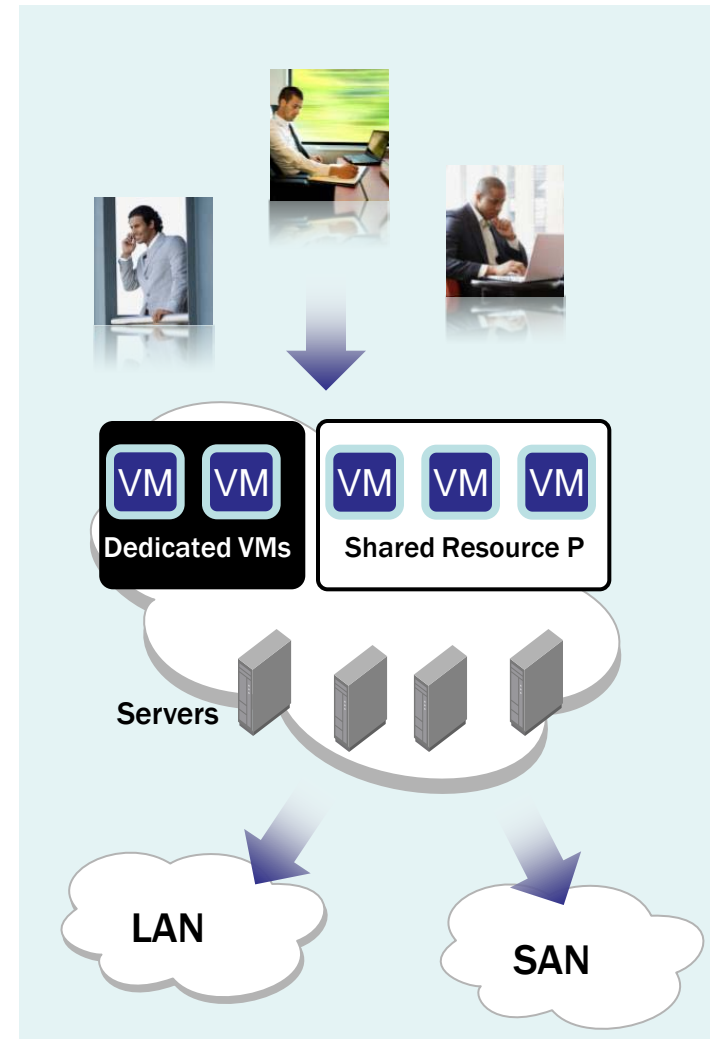
NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

➤ Data Center Convergence


- ◆ This session will appeal to Data Center Managers, and IT professionals seeking a fundamental understanding of the data center convergence trends. The session will define convergence as it applies to data centers, review upcoming changes and their impact on data centers. The audience will get acquainted with the emerging trend towards flatter data center networks and how to deploy them in existing environments or new deployments.

What if I Don't Converge?


- Your data center will not become obsolete
- Continue to manage and operate separate purpose built networks
- Take advantage of new technologies
 - ◆ 16G FC for FC SANs
 - ◆ 10GbE, 40G and 100G for enterprise networks and iSCSI SANs
- Deploy virtualization
- Upgrade management tools



➤ Today, data and storage networks are separate




**STORAGE
NETWORK**



**DATA
NETWORK**

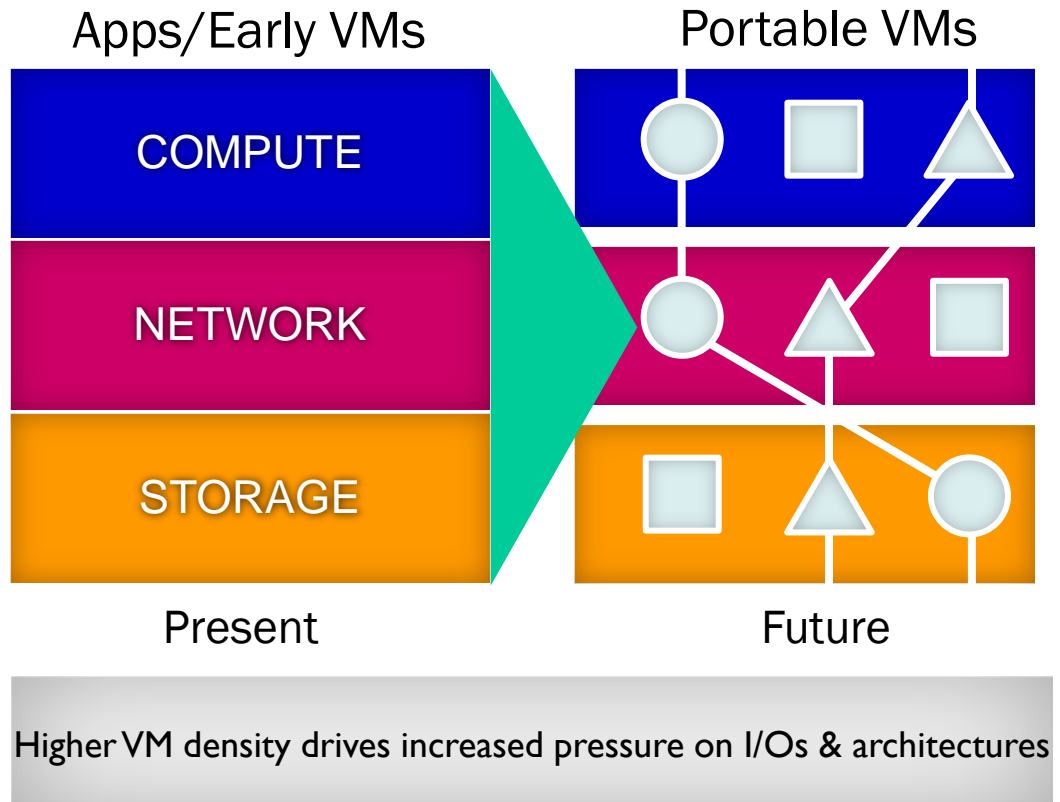
- The goal is to consolidate data and storage networks into one



**CONVERGED
NETWORK**


Virtualization is driving convergence

- Virtual data centers require highly portable VMs that move between compute nodes and distant data centers
- Data center resources need to be available on demand
- Simpler and scalable infrastructure is needed to support virtualization




Challenges of Today's Network

- Layer 2 performance, scalability, reliability
 - ◆ Limitations of Spanning Tree Protocol (STP)
- Scaling virtual server environment
- Virtual Machine mobility
- Infrastructure complexity
 - ◆ Lots of switches to manage
 - ◆ Layer 3 protocols to the edge
- Management overhead
 - ◆ High OpEx




Challenges of Scaling VM Environments




- ▶ **Layer 2: single active path**
 - ◊ STP disables other paths
 - ◊ Not Virtualization optimized
- ▶ **Add Virtual Machines**
 - ◊ Move to 10 GbE for simplicity and performance
 - ◊ Uplinks are stressed
- ▶ **Increase utilization using MSTP** (spanning tree per VLAN)
 - ◊ Increases complexity
 - ◊ Creates multiple single-path networks
- ▶ **Link failure – slow reconvergence**
 - ◊ STP reconfiguration may be too slow
 - ◊ Broadcast storms stress network
- ▶ **Layer 3 as an alternative**
 - ◊ Even more complexity; higher cost
 - ◊ VM mobility limited

What we need?

- L2 networks with no STP
- All paths in the networks are utilized with traffic automatically balanced
- Link failures not result in temporary outage and shortest paths are always used
- Networks with lossless transmission & low latency
- Networks that are built for convergence of data and storage networks



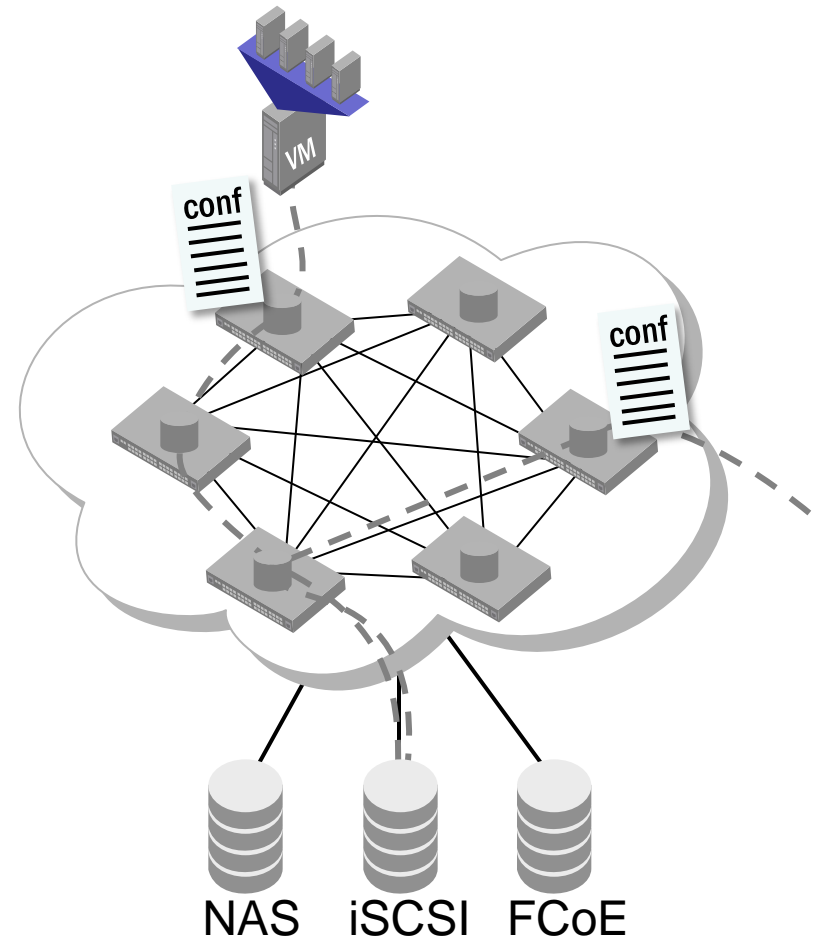
Challenges of VM Mobility



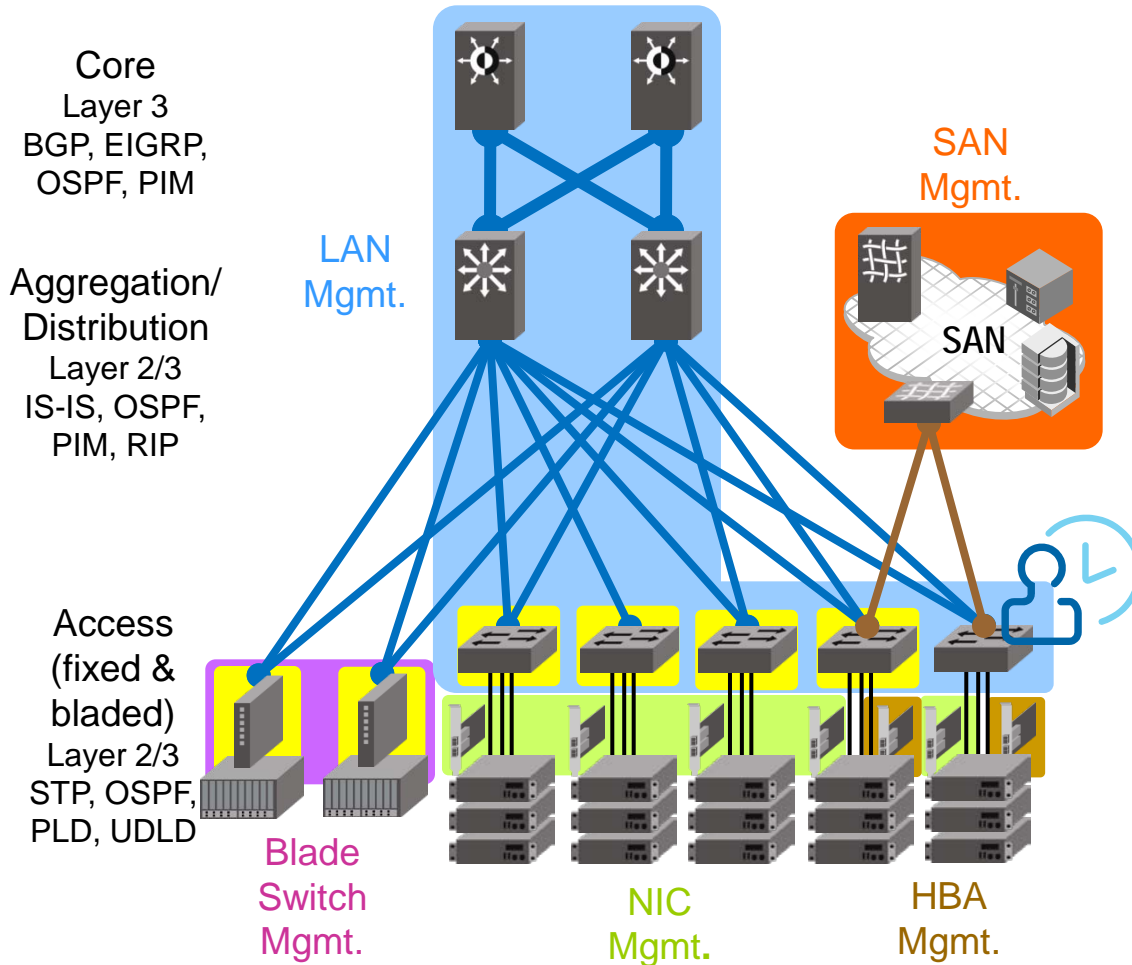
- ▶ **Limited sphere of mobility**
 - ◆ STP limits flexibility to a minimized, defined tree of switches
 - ◆ L3 limits mobility to a single rack
- ▶ **VM migration may break network access**
 - ◆ Manual adjustment of destination hosts and switches required
- ▶ **Services (VLANs, QoS, security) mapped to all physical ports**
 - ◆ Eases mobility, but breaks network and security best practices
- ▶ **Distributed Virtual Switch**
 - ◆ Service configuration consistency
 - ◆ But high overhead
- ▶ **Limited insight into where VMs are running**
 - ◆ Automation results in VMs existing anywhere in the HA cluster

What do we need?

- No physical barriers of VMs migration
- Networks that are aware of where VMs are running at all times
- Data centers with automated ability to balance utilization and increase application availability
- Fully distributed control plane
- Arbitrary topology with auto configuration or zero configuration



Challenges of Network Management



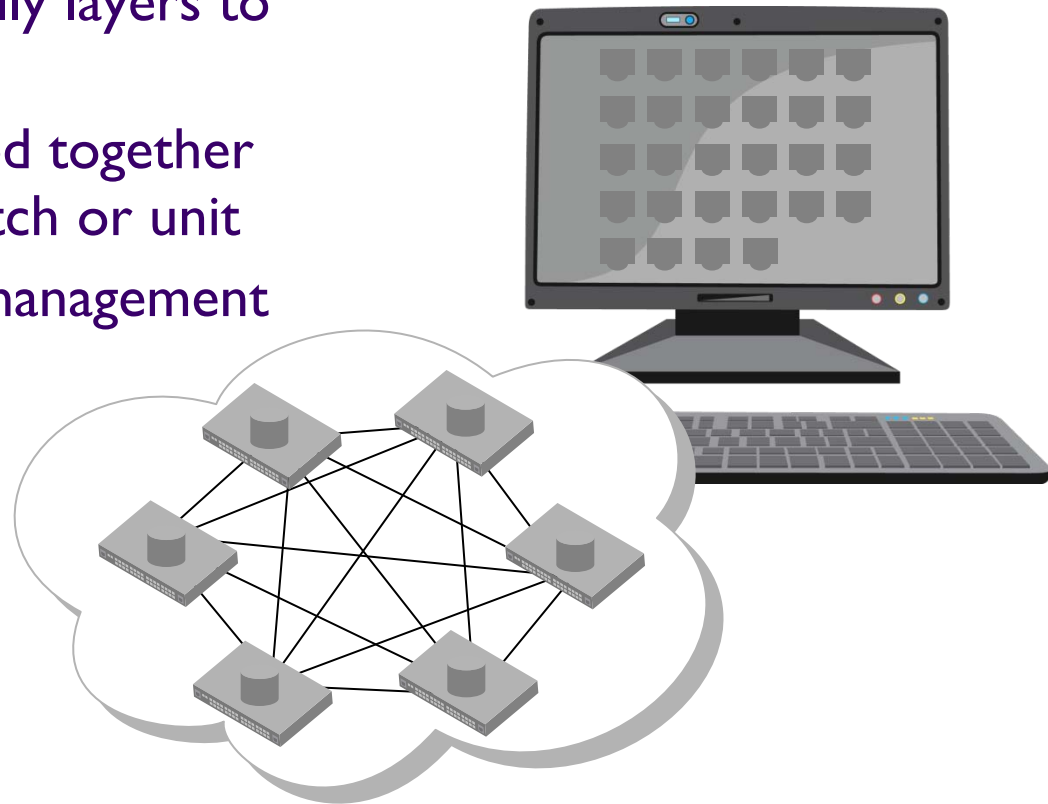
- ▶ Too many network layers
 - Utilize many L2/L3 protocols
- ▶ Lots of small-form-factor switches at the edge
 - Each switch has to be managed
 - Because of the number, they need to be aggregated
- ▶ Configuration time when deploying new switches
 - Switch has to be set up
 - Templates have to be loaded
- ▶ Separate management tools for LAN, SAN, NICs/HBAs
 - Management silos do not fit in a virtualized data center
 - Drives up OpEx

Elements of network layers are simplified for the purpose of this presentation

What do we need?


- Networks with fewer logically layers to deal with
- Switches that can be grouped together and managed as a single switch or unit
- Centralized or distributed management

- Universal or common tools to manage all converged network resources




Today's Common Architecture

Core




- Layer 3 only

Aggregation/ Distribution




- Layer 2/3

Access



- L2 edge switches
- STP

Devices



Servers with
10 Gb CNAs

Mix 1 Gb, 10 Gb
Servers



FCoE
Storage



iSCSI, NAS
Storage


10 Gb DCB

1 Gb Ethernet

10 Gb Ethernet


Next Architecture

Core




- Layer 3 only

Edge



- L2 switches
- Layer 3 to Core layer
- Managed as one, auto configure, No STP

Devices



10 Gb DCB

1 Gb Ethernet

10 Gb Ethernet

What Do I get With Convergence?

- Lower costs
 - ◆ Consolidate & optimize resources; simplify configurations
- Increased performance and reliability
 - ◆ Faster, consistent access; more with less
 - ◆ Minimize disruption, recover quickly with more resilient L2 infrastructures
- Agility & Scalability
 - Deploy/re-deploy resources quickly
 - Scale based on business needs
- Improved Virtualization
 - ◆ Applications deployment and mobility



➤ Storage

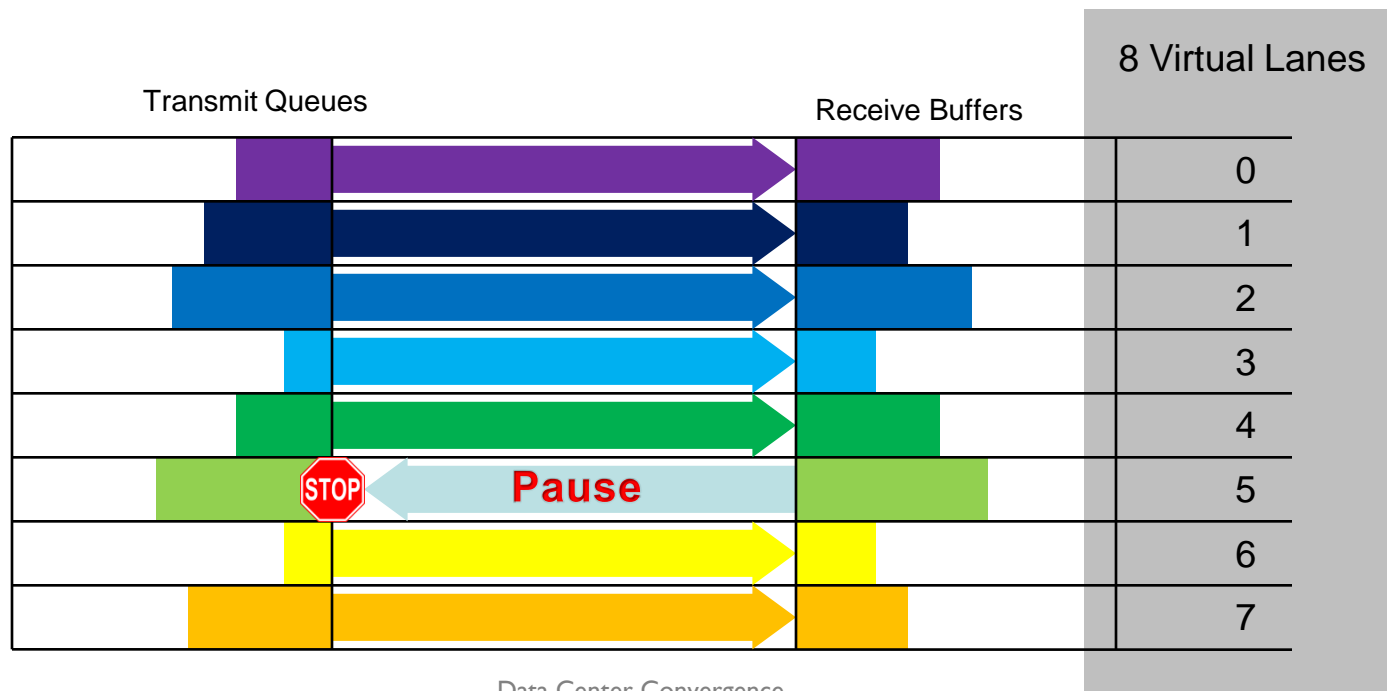
- ◆ FCoE: encapsulation of FC over Ethernet - see tutorial
- ◆ iSCSI: encapsulation of SCSI over TCP/IP - see tutorial

➤ Networking

- ◆ DCB: Lossless Ethernet – see tutorial
- ◆ TRILL: Layer 2 multi-path and multi-hop capabilities
- ◆ New L2 architectures: flatter networks


Priority-based Flow Control

- During periods of heavy congestion
 - Ensures delivery of critical data
 - Latency sensitive traffic continues normal operation



Enhanced Transmission Selection

- Capability to apply differentiated treatment to different traffic within the same traffic class enabled by ETS



Transparent Interconnection of Lots of Links



A proposed data center L2 protocol being developed by an Internet Engineering Task Force (IETF) workgroup

Mission


*“The TRILL WG will design a solution for **shortest-path frame routing** in multi-hop IEEE 802.1-compliant Ethernet networks with arbitrary topologies, using an existing link-state routing protocol technology.”* - source IETF

Scope

*“TRILL solutions are intended to address the problems of ..., inability to multipath, ... within a *single Ethernet link subnet*”* - source IETF


TRILL – no STP

- Multi-path Layer 2 switching
 - ◆ Multiple active paths
 - ◆ Reclaim network bandwidth and improve utilization
- Establishes shortest paths through Layer 2 networks
- Fast response to failures
- Backward-compatible and connects into existing infrastructures
- Deliver multiple hops for all traffic types (including FCoE)




Deployment: Server Edge Top of Rack

- Adding FCoE and DCB at the edge or top of rack switches
- Replace top of rack switches, but preserve the rest of LAN & SAN configurations
- Non-disruptive addition to existing environments



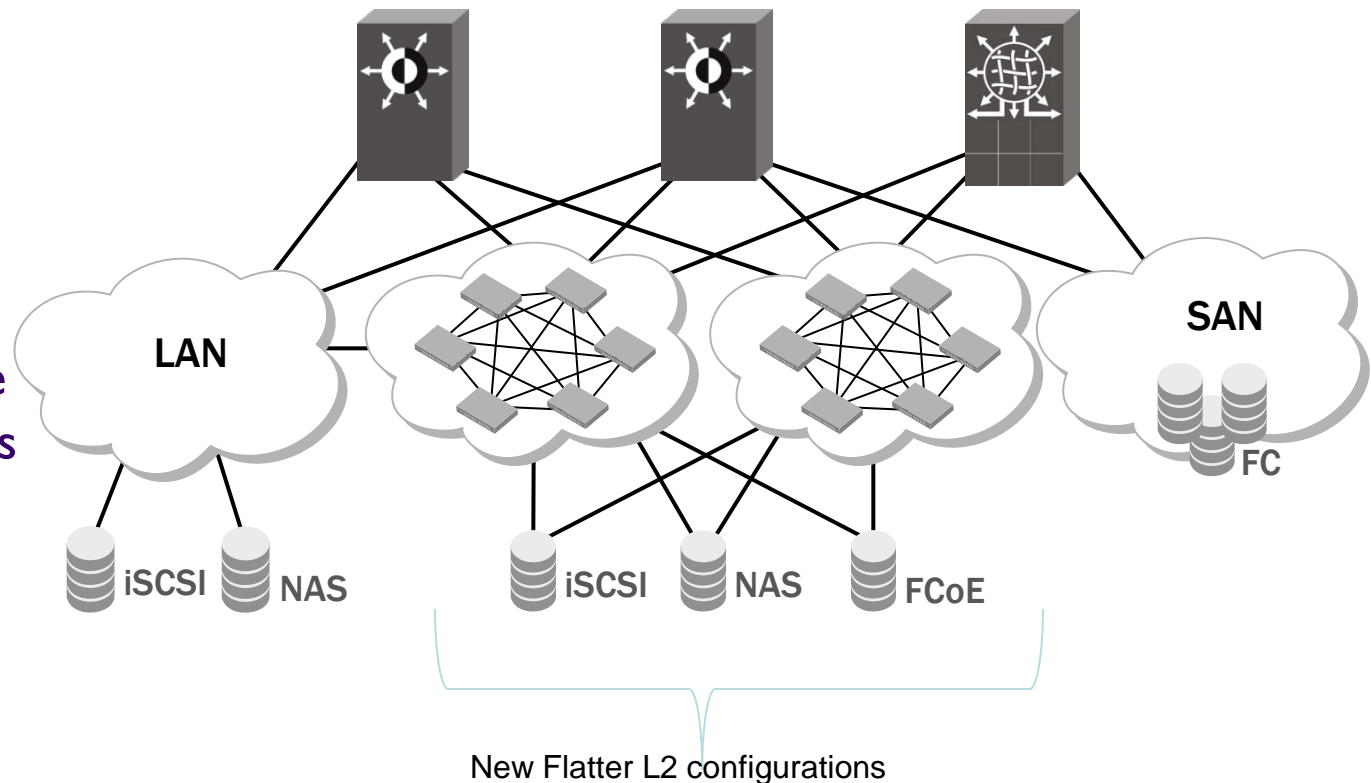
Deployment: End-to-end FCoE

- End-to-end FCoE, from edge to storage. Utilize converged switches throughout the network
- Add native FCoE storage, which connects to converged switches
- DCB and FCoE added to existing infrastructure. In this environment, FC, iSCSI, NAS and FCoE storage devices may coexist



Converged Data Center

- Expanded L2 layer deployments will be added to existing infrastructure
- New L2 deployments may not accommodate STP configurations
- New bridging devices will eventually enable integration of disparate data center networking devices



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

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

- SNIA Education Committee

TBD