Education

Data Center Evolution and Network Convergence

Joseph L White, Juniper Networks

SNIA Legal Notice

- 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.

Abstract

Data Center Evolution and Network Convergence

- FC, FCoE, NAS, iSCSI, DCB, traditional LAN, internet/WAN, HPC, clusters, clouds, server virtualization, storage virtualization, network virtualization, and more are all colliding in your data center. Redundancy, resiliency, security, I/O consolidation, network convergence, dynamic application distribution, and thin provisioning with high levels of service are desired at all layers and all data center sizes across a broad spectrum of use cases. You worry about operational separation, buying decisions, investment protection, cost and energy savings, and evolving standards while maintaining very high levels of service and security. Is the technology evolving to a dream come true or a nightmare? If that doesn't keep you up at night nothing will.
- This tutorial will untangle, define, and illustrate the main ideas and concepts behind Data Center Evolution and Network Convergence to give context and a solid foundation for discussions with your vendors as well as for your further reading and investigation. The point of view taken for this presentation is that of the network and transport characteristics in the face of the changes taking place.

Topics Discussed

- Data Center Evolution
  - History
  - Influences, Trends, Drivers
  - Data Center Map
  - Network Convergence
- Techniques and Technology
  - Complexity – Scaling – Tiers
  - I/O Consolidation
  - SAN-LAN Convergence

Some History

- First Model
  - isolated systems
  - direct attach storage
- Second Model
  - limited networking
  - direct attach storage
- Third Model
  - Networking Explosion
  - direct attach storage + local cluster server-server storage
- Third and a half Model
  - Flexible storage via NAS and network file systems
- Fourth Model
  - SANs for block storage attach plus fully entrenched NAS
- Fifth Model ('current one')
  - Server Virtualization drives first hop I/O consolidation, increased SAN attach
  - "Cloud" (pooled resources of all kinds with uniform distributed access)
- Evolving Model
  - Network Convergence
    - Protocols for SAN and LAN on same infrastructure
    - Network Scaling via virtualization and simplification (tier collapsing, distributed control planes)

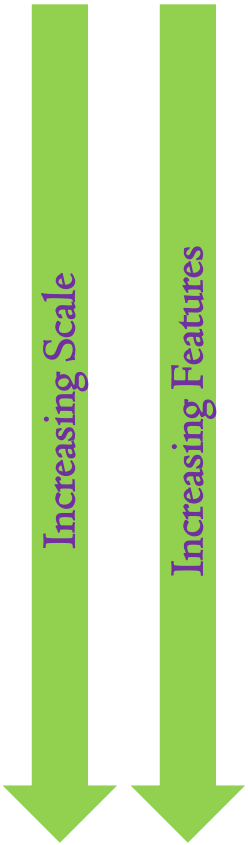

Recent Trends in the Data Center

Consolidation

- Mega DCs; 400K sq ft
- 4K racks, 200K servers

DC Scale

Server Trends

- Multi-core (8->16 >32,….128,…)
- Virtualization and VMs

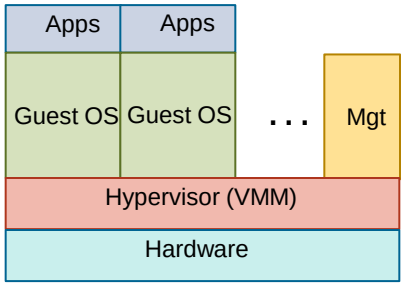

Want Low Oversubscription

Application Trends

- SOA, Web 2.0
- MapReduce, Hadoop, Grids

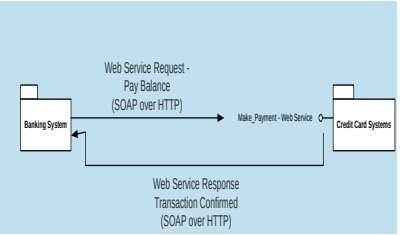

Increased East-West traffic

Interconnect Trends

- Convergence to 10 GE
- Enhancements to Ethernet

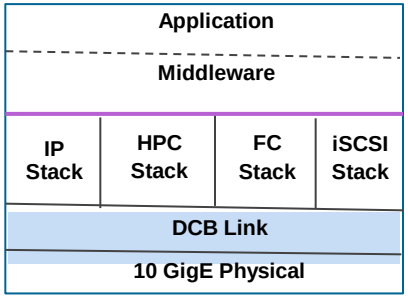

Large speed increases 10/40/100 GE

Explosion in infrastructure

- Function of new services, digital media, application design and deployment topologies
- Fueled by Moore's law and ever increasing price performance
- Driving issues of managing massive performance and scale

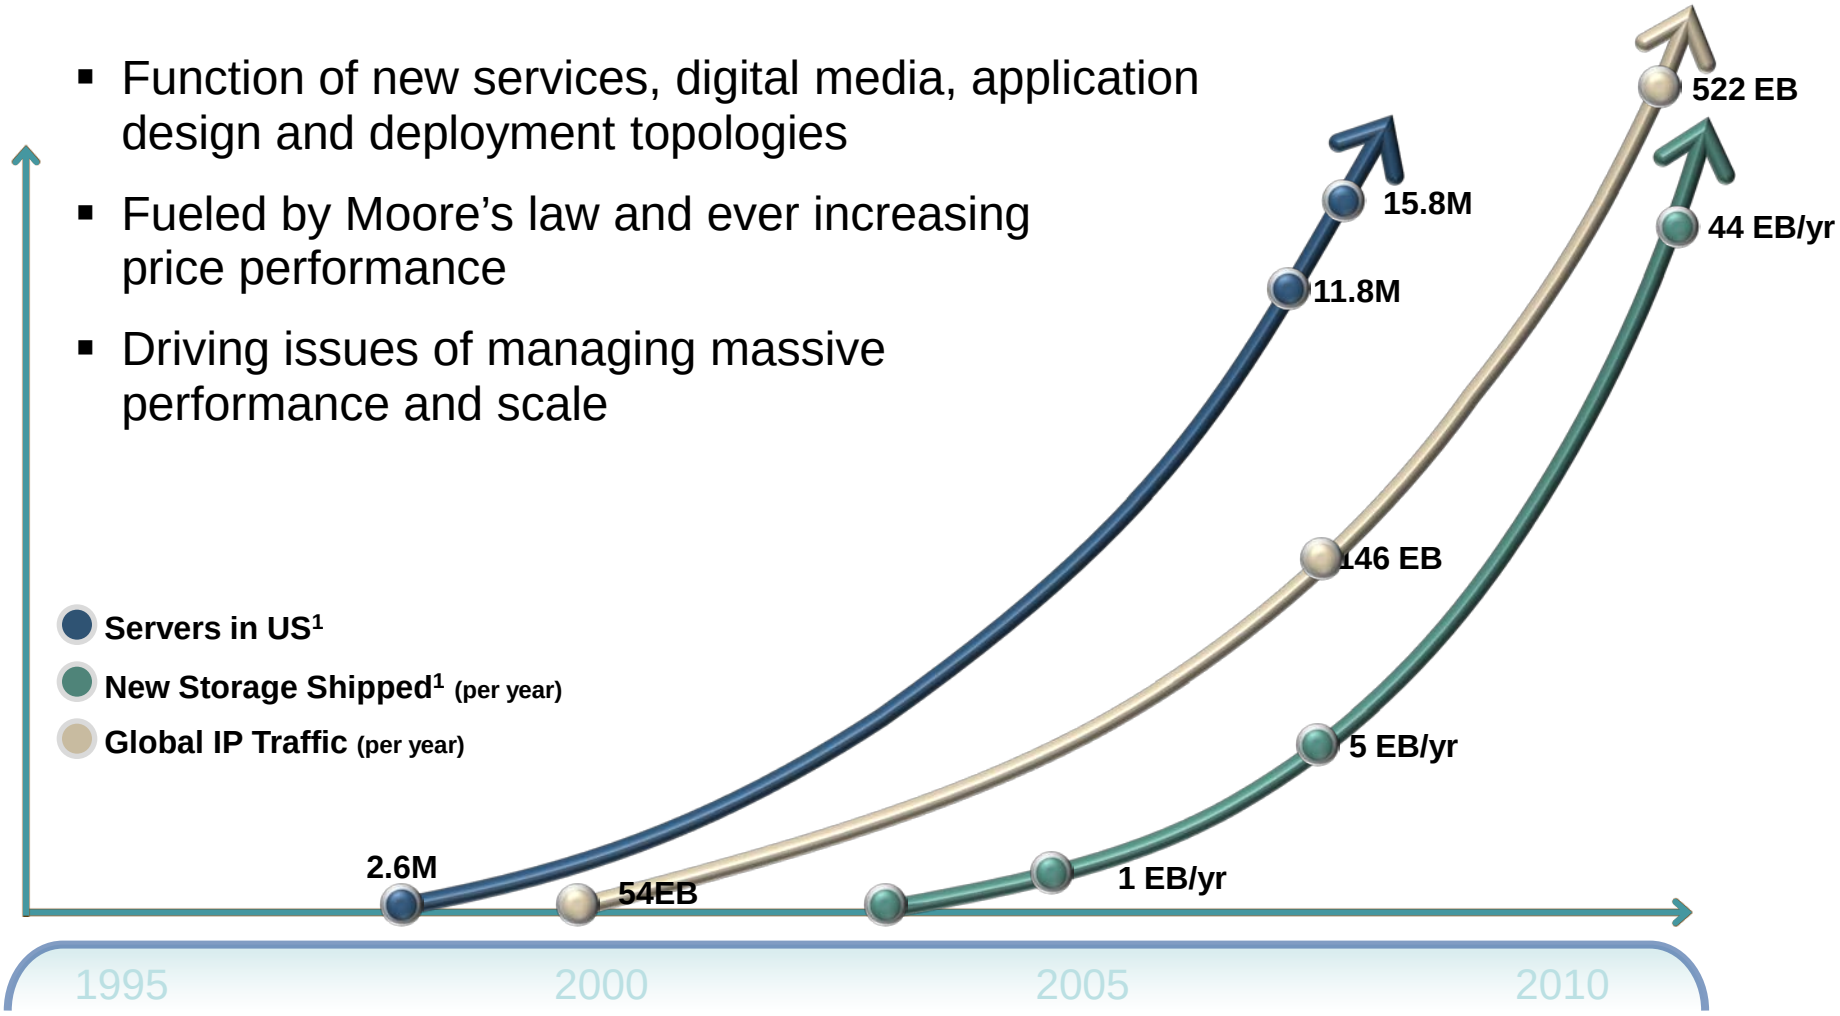

Complexity increases exponentially with scale [1]Source: IDC

Servers and Storage Attach

High End Servers with high IOPS & performance MF / Unix	**Storage is High End High IOPS**
Mid range servers Unix / Wintel	Storage with Reasonable balance of capability & performance
Rack Dense & blade servers Wintel / Linux	Storage with minimal load, no persistent data, just VM images

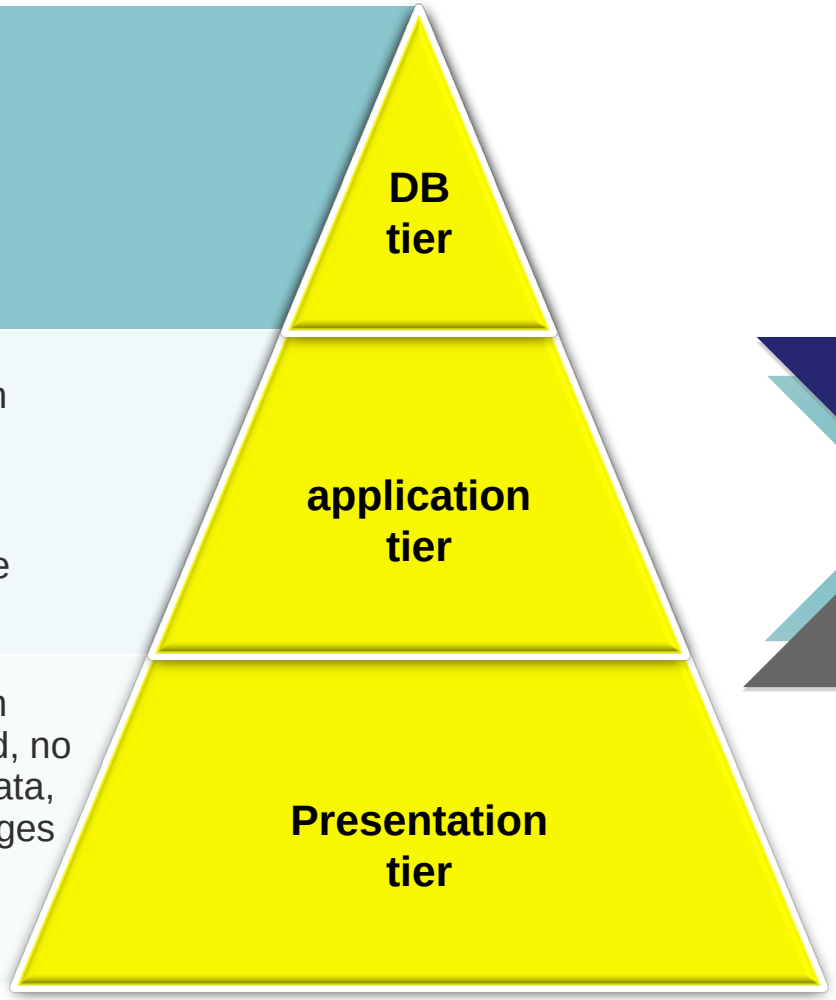

SAN Attach started at the high end

SAN Attach Penetrating lower

Server virtualization driving towards 100% attach

Server virtualization started at the low end

Steadily penetrating higher

Needs shared storage for VMs

Virtualization OF EVERYTHING

- Aggregate up and Virtualize down
  - many examples such as storage arrays, servers, ...
  - avoid Accidental partitioning
  - embrace Deliberate partitioning
- Aggregation
  - Physical and Software
  - Bring together and pool capacity with flexible connectivity
- Virtualization
  - logical partitions of the aggregated systems to match actual need
  - flexibility → fungible resources everywhere
  - Utility Infrastructure with just in time & thin provisioning

THIS IS HAPPENING TO NETWORKS AS WELL

Virtual to Physical Server Trend

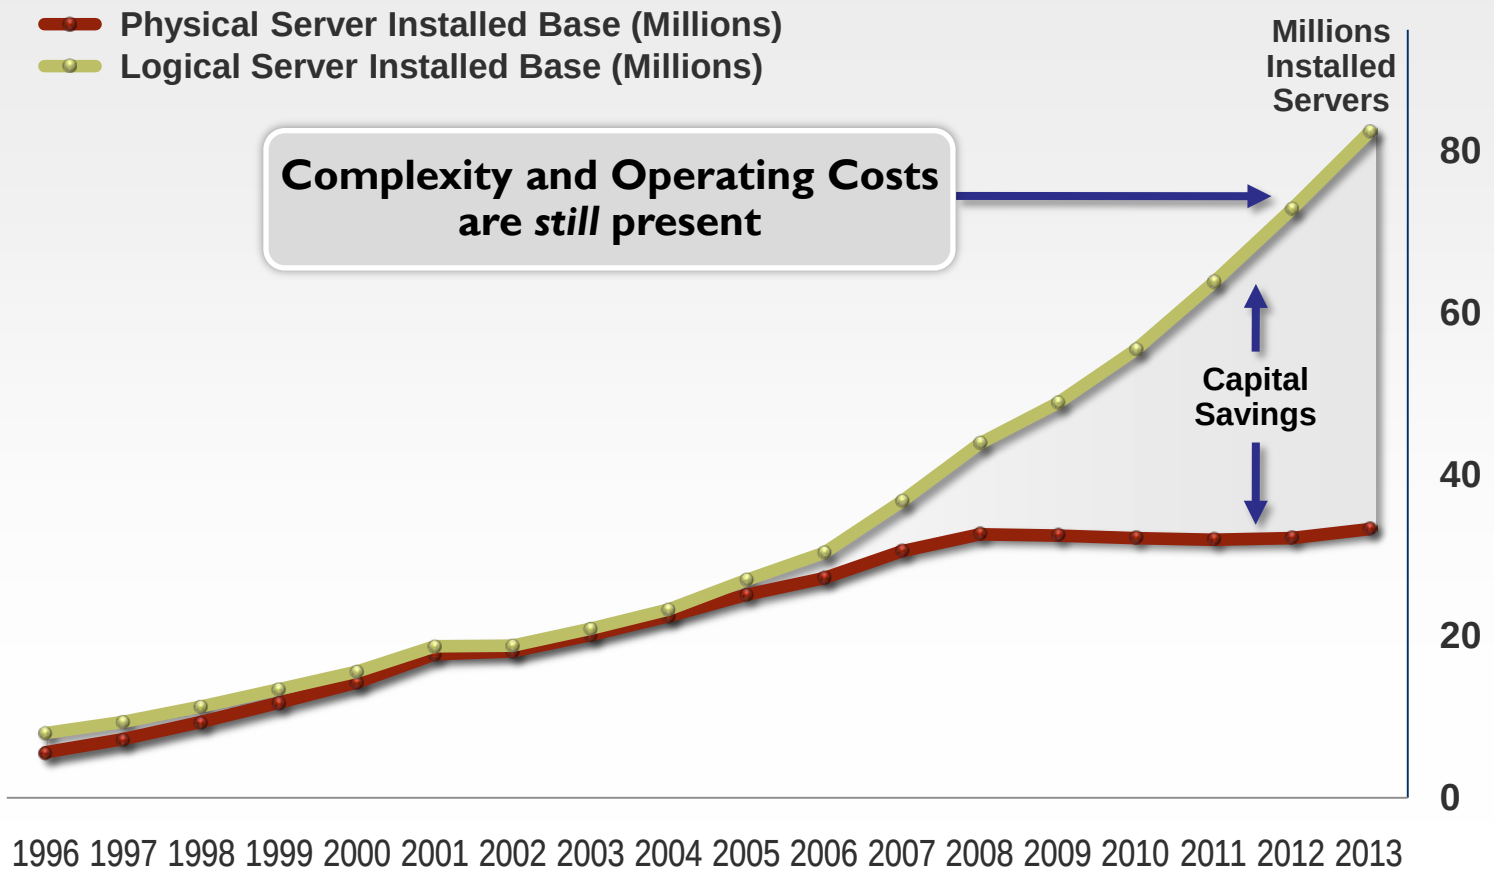

Source: IDC

Virtualization Drives Storage Connectivity

... because Data Centers are always in flux

Application life cycle

services introduced, updated, retired

Load on servers and networks constantly changing

can be unpredictable

Resource management challenge

- Minimize the need for excess capacity
  - Reconfigure
  - Reclaim/Reuse
- Adding resources is last resort

Dynamic shared resource pools address these issues

Enabled by Virtualization + Full Connectivity Networks

Any servers potentially needs access to any storage
Drives SAN attach from 20% to near 100%

If you don't converge you will end up connecting everything to everything anyway but across additional parallel networks.

Has Convergence already happened?

- For some aspects of Convergence: YES
  - NAS Allows access to file based storage across the network
  - iSCSI Allows access to block based storage across the network
  - SANs have been bridged across metro and wide area networks for 10 years (FCIP, iFCP, & proprietary)
  - FCoE provides an accepted protocol for FC across an Ethernet
- These are good but not sufficient
  - Data center LANs have issues at scale
  - WAN IP SAN connections do not solve the Local Data Center problem
  - Operational Characteristics of FC based SANs desirable and entrenched
- **SO, we will explore the new wave of convergence**

Network Convergence

- Convergence occurring along 2 major themes
  - These are happening at the same time
  - We will discuss them separately
  - We will discuss how they merge together

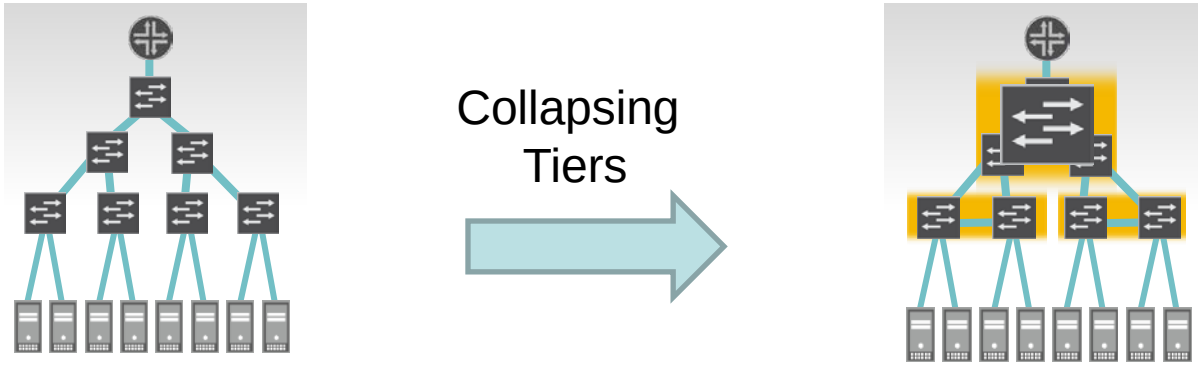

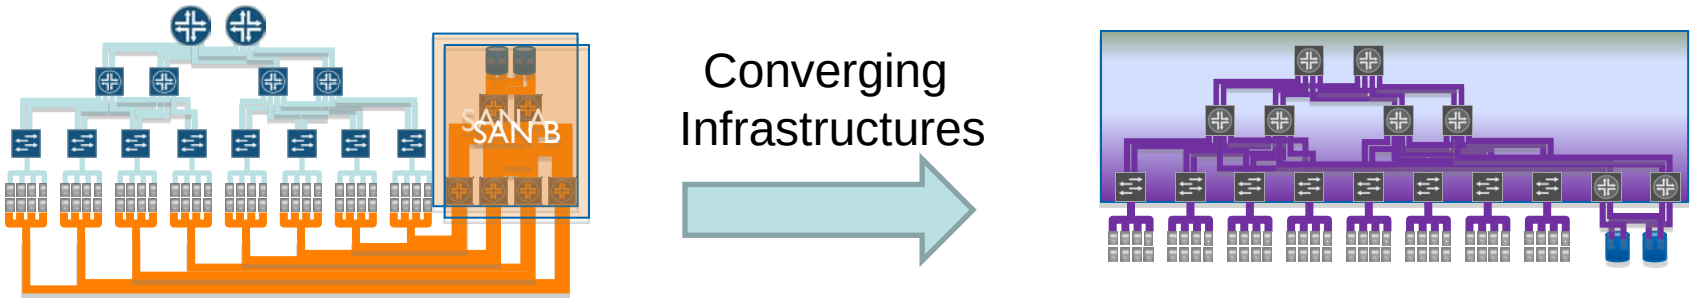

Data Center LAN (today)

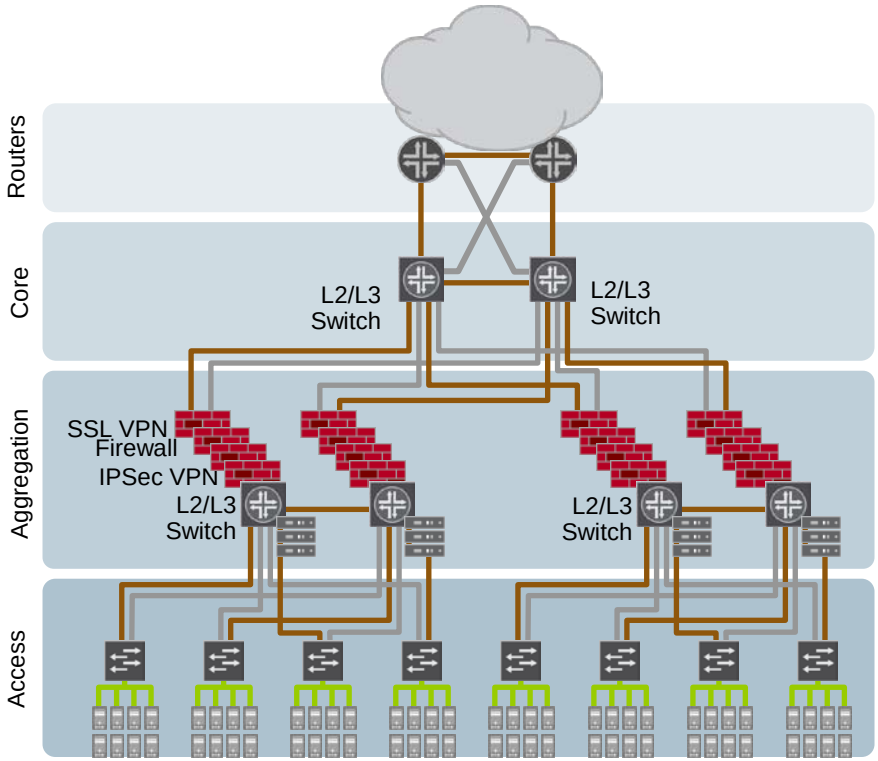

- Servers → Servers, NAS, Campus/MAN/WAN across switched network
- Multi-Tier
- 100's to many1000's of ports
- multi-link redundancy
- 100s of meters max diameter
  - oversubscribed
  - East-West Latency can be a problem
- Ethernet carrying predominantly IP traffic
- Firewalls and security in aggregation layer
  - have to be distribute in the data path due to efficiency forced by oversubscription

Evolution of the Data Center LAN

Education SNIA

Local Area

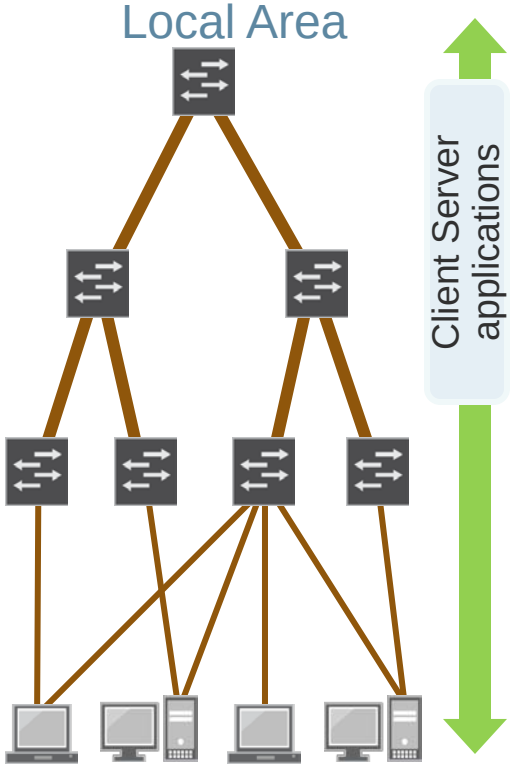

20 years ago the Ethernet switch was introduced to solve the LAN problem

And it became the basic building block of the network

Data Center

And so we wired the data center the same way

Need more redundancy: STP and disabled links OR lots of VLANs

SOA Applications etc

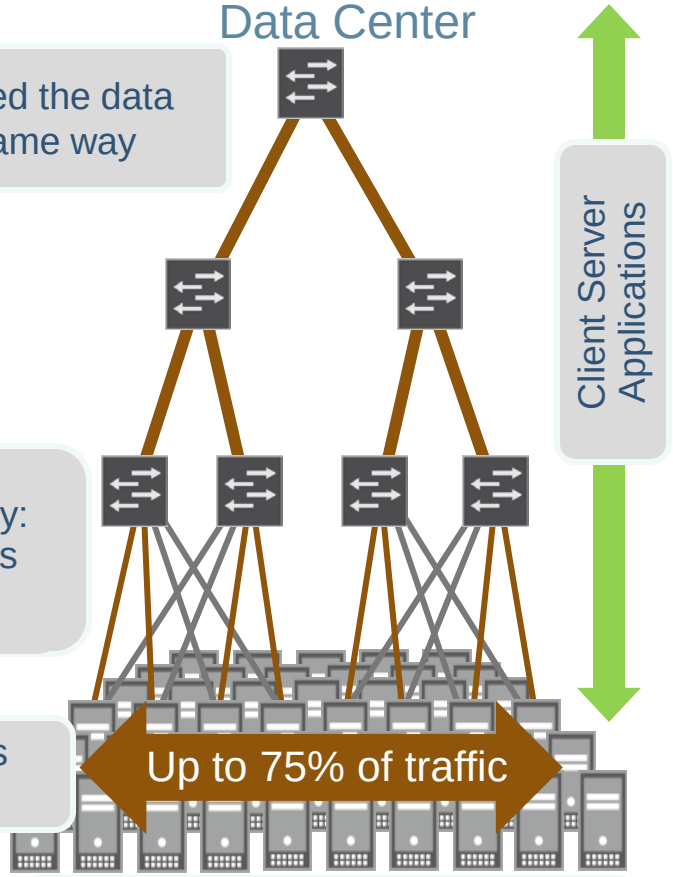

Adding in the complexity of storage networking as well, these pressures force the Ethernet network to change.

Bubbles of Optimal Performance

Typical tree configuration

Location matters in a tree architecture

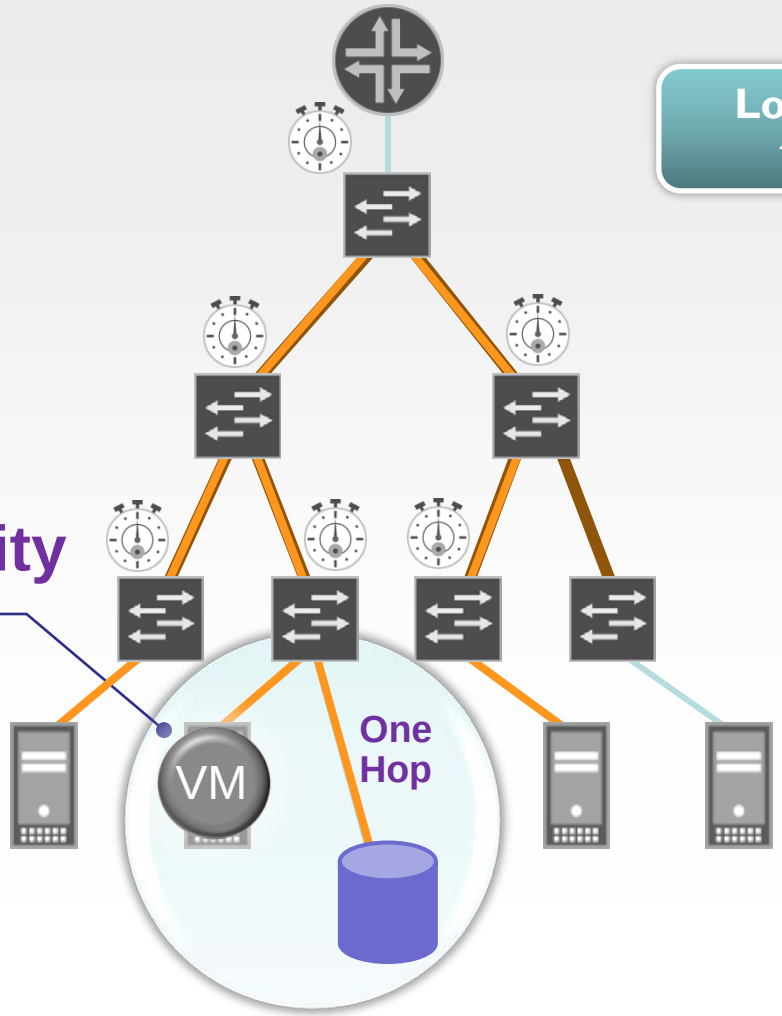

Local Connectivity Bubble
for Optimal Performance

Appliances Create Shadows

Location matters in a tree architecture

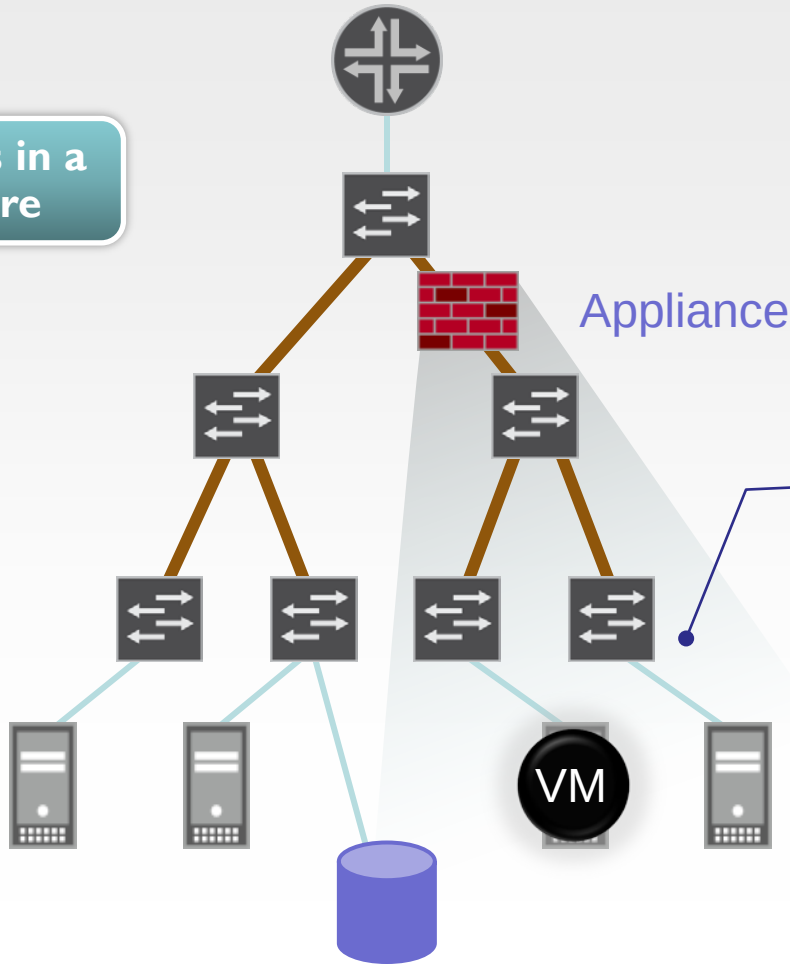

Appliances and VLANs

Shadow of accessibility

Simplifying the Data Center LAN

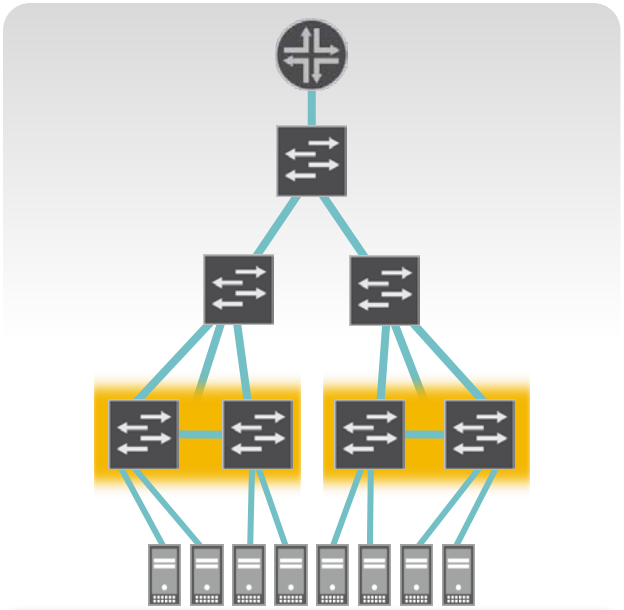

Aggregate Switches

Multiple physical switches that operate as a single logical device

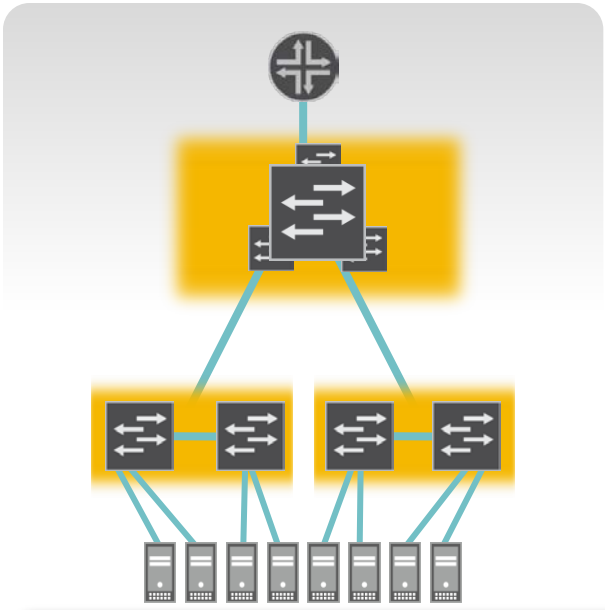

Collapse Tiers

Use an aggregated switch to do the work of multiple tiers of switches

Allows aggregated access to services

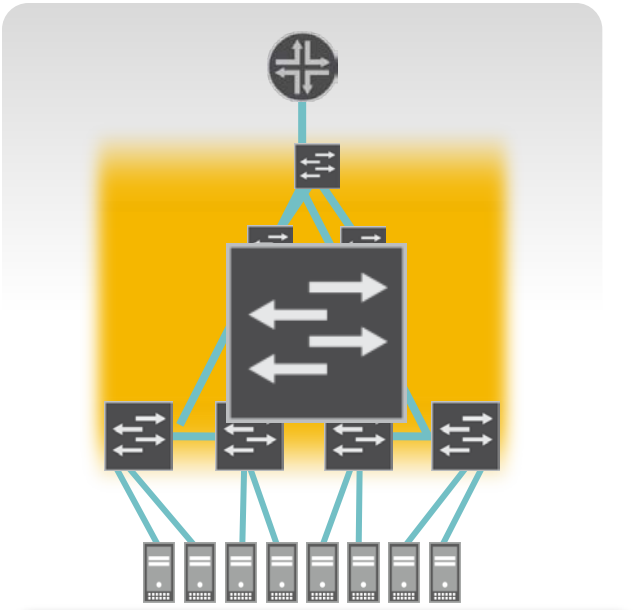

Create A Fabric

Use both techniques at the same time to build a fabric based infrastructure

WE WANT:

1. Operational Simplicity
2. Lower Latency
3. Lower cost

Data Center SAN (FC based)

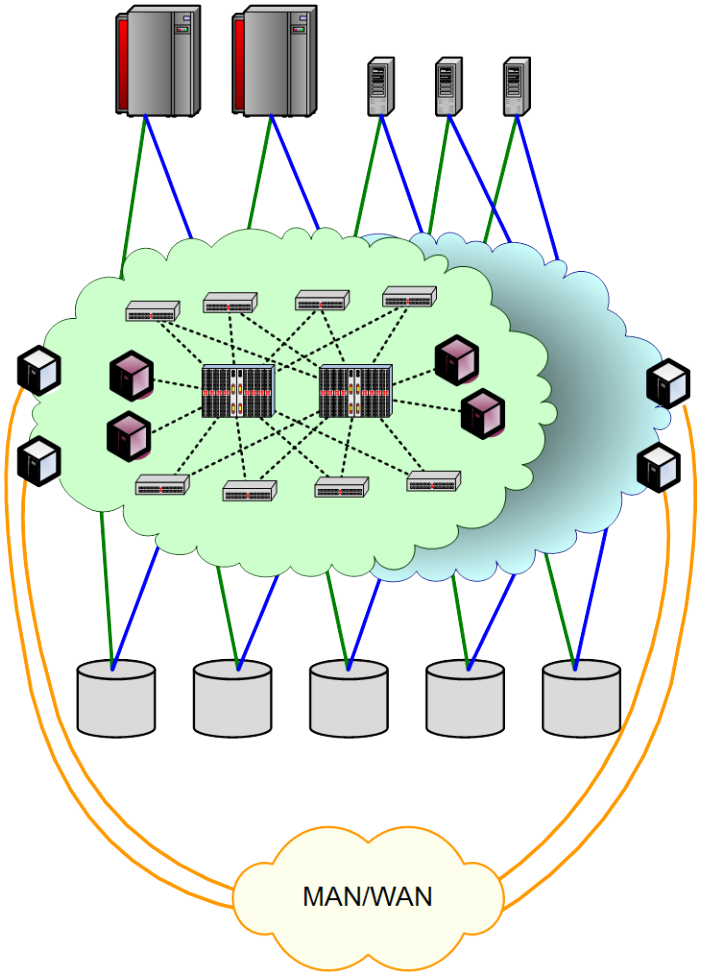

- Servers → Storage across switched network
- Core – Edge or Edge – Core – Edge
  - in effect one level of tier collapse is done
- 10's to 1000's of ports
- Full Dual Rail Redundancy
- 100s of meters max diameter
  - High Bandwidth, Low Latency
  - Lossless Links
- Fabric Services provide Discovery, Access Control, and Change Notification
- Gateways and specialized extension devices provide remote access for BC/DR
- Attached Appliances provide data services
  - Encryption, Block Virtualization

Put the SAN on the LAN

- What happens when you overlay
  - the SAN picture.... onto the LAN picture

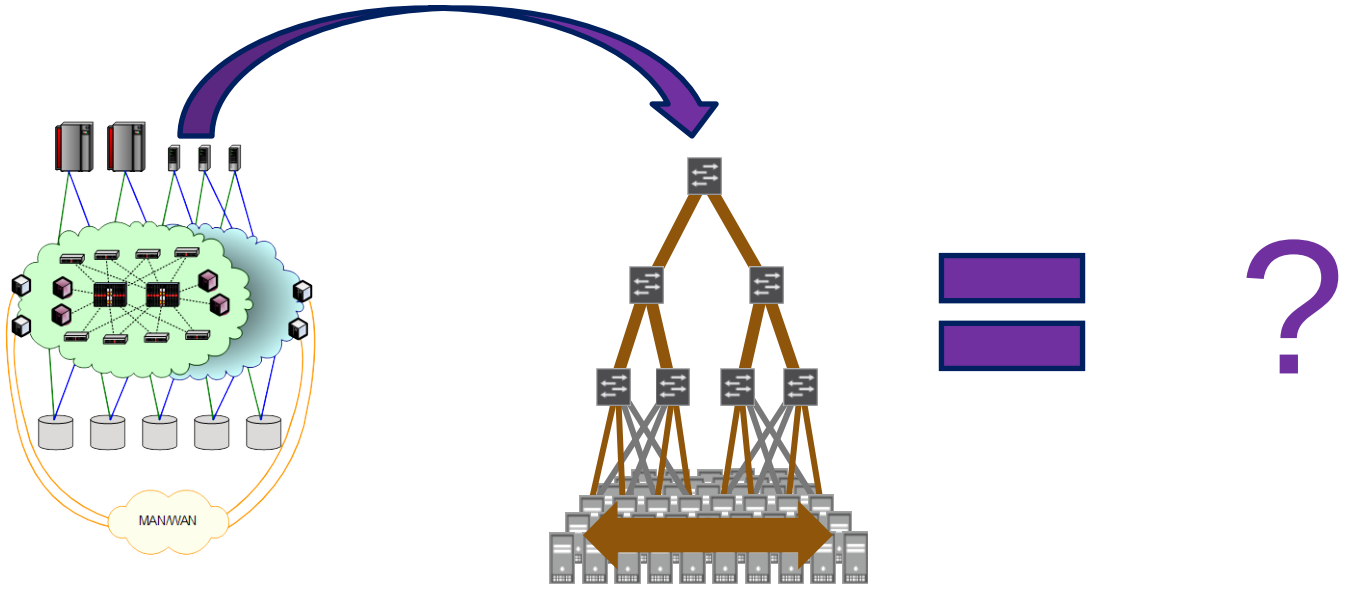

- SAN protocols drive the single biggest piece of the convergence story

At a high level you end up with...

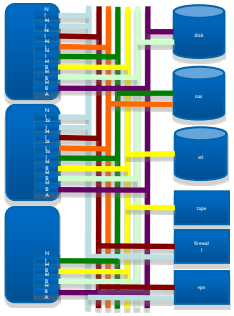

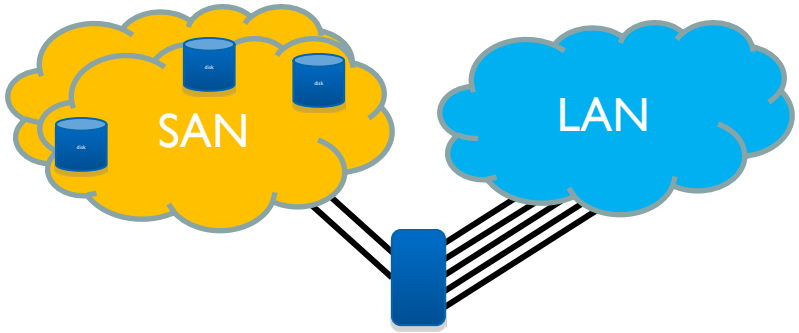

I/O Consolidation + Network Convergence

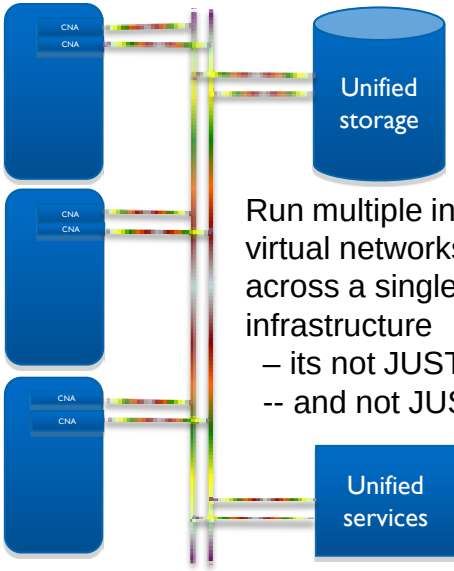

Run multiple independent virtual networks 'planes' across a single physical infrastructure
- – its not JUST an edge thing
- -- and not JUST a SAN thing

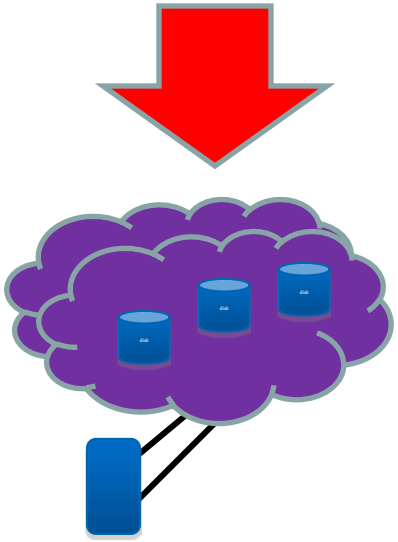

Evolution of Network Convergence

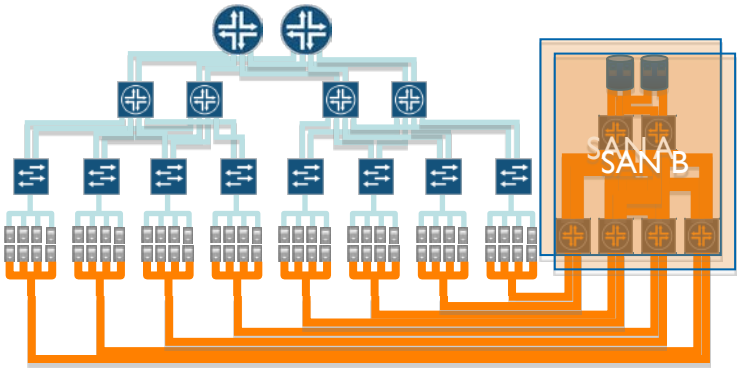

- Separate Networks
  - No overlap of operations & management
  - separate network characteristics
  - separate teams run them (typically)

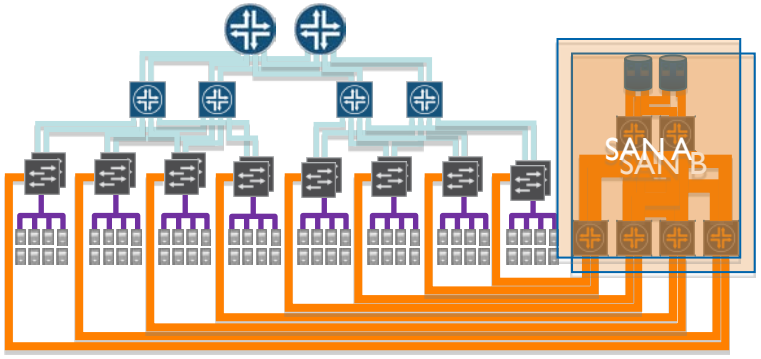

- I/O Consolidation
  - overlap confined to server and 1st hop
    - team consensus needed
  - several detailed ways to accomplish
    - transit switch L2 to access SAN
    - gateway

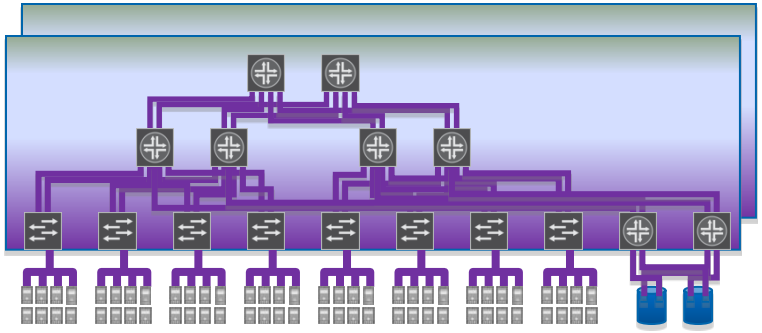

- SAN-LAN Convergence
  - unified management and operational model
    - run by single team
    - common redundancy
    - shared bandwidth and connectivity

Evolution of Server Connections

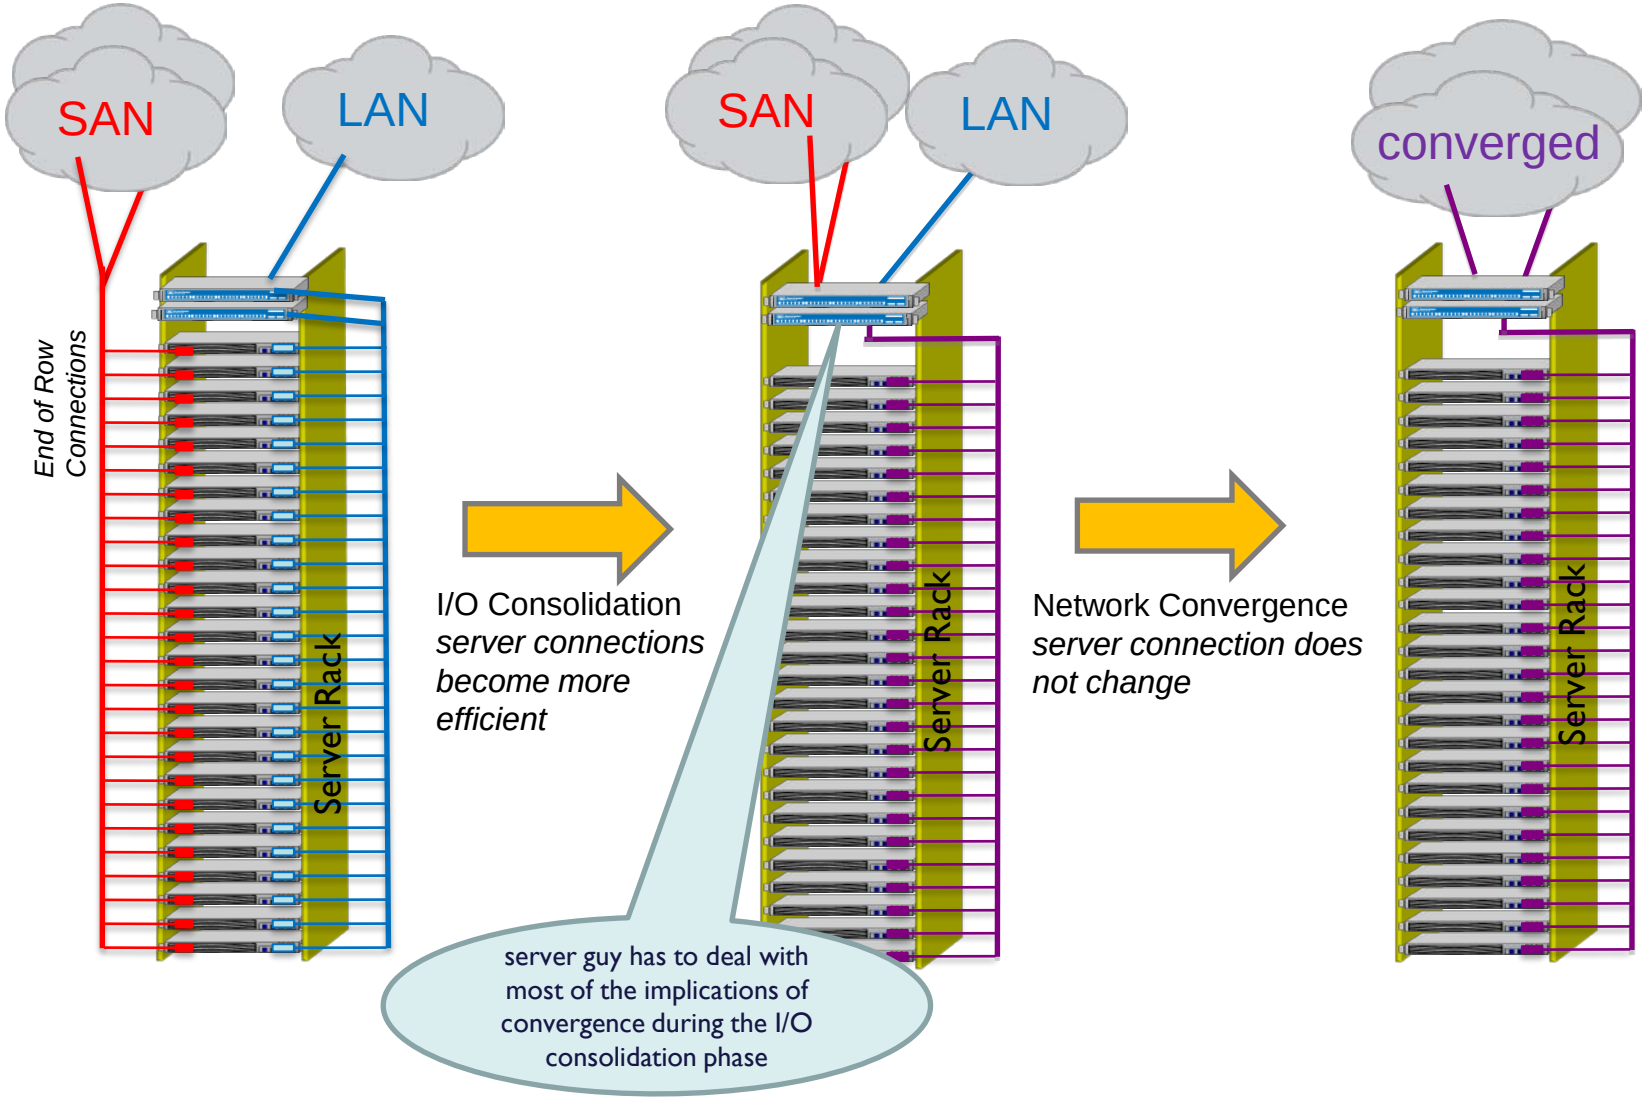

HOW MANY NIC'S IN A SERVER

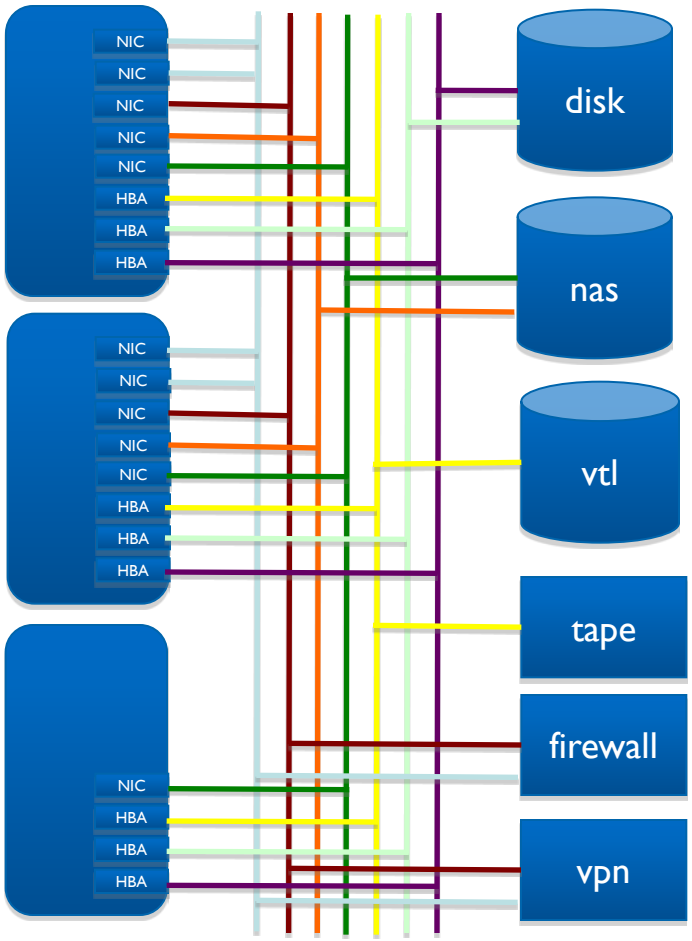

- Multiple network connections
  - Separate subnets
  - Separate VLANs
  - May be separate networks
  - May be separate switches in the same network
  - Each configured with different QoS settings
- Depends on the Server…
  - file servers and/or NAS points to campus
  - Presentation servers and intranet point to campus
  - Application servers
  - Database servers
  - Backup Servers, etc etc
- For Example Something like this is possible:
  - Campus Access (1+)
  - LAN Backup NIC (1)
  - Application Cluster NIC (1)
  - Vmotion NIC (1)
  - Presentation to application to database private NICs (1+)
  - Storage Access Cards (2)

I/O Consolidation reducing this to 2 NICs can save a bundle here

I/O Consolidation Deployments

- server direct to gateway

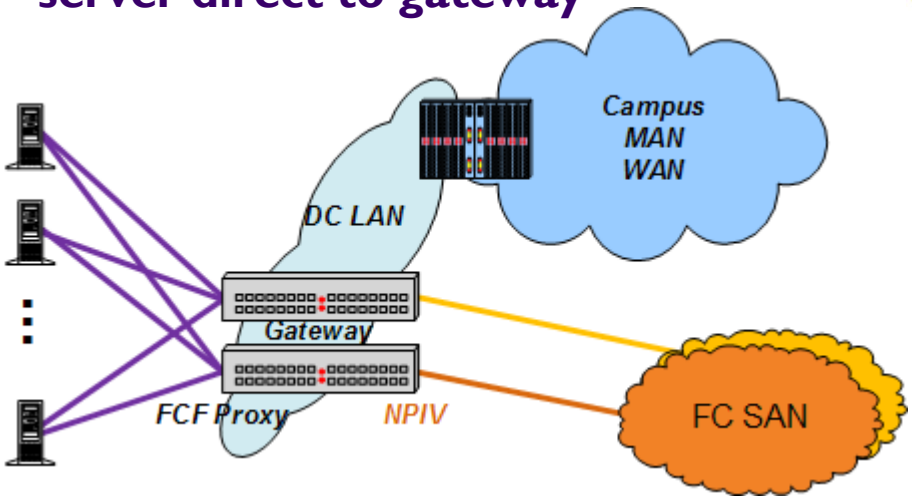

- server direct to FCF

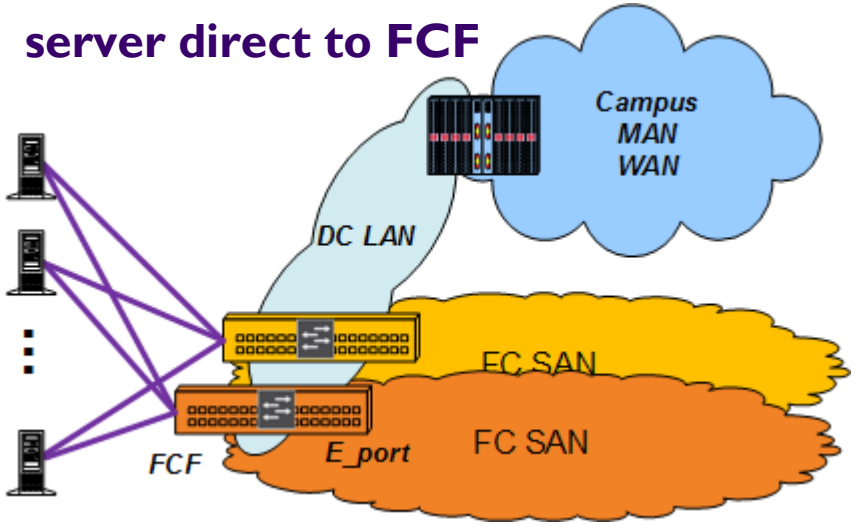

- server to L2 to gateway

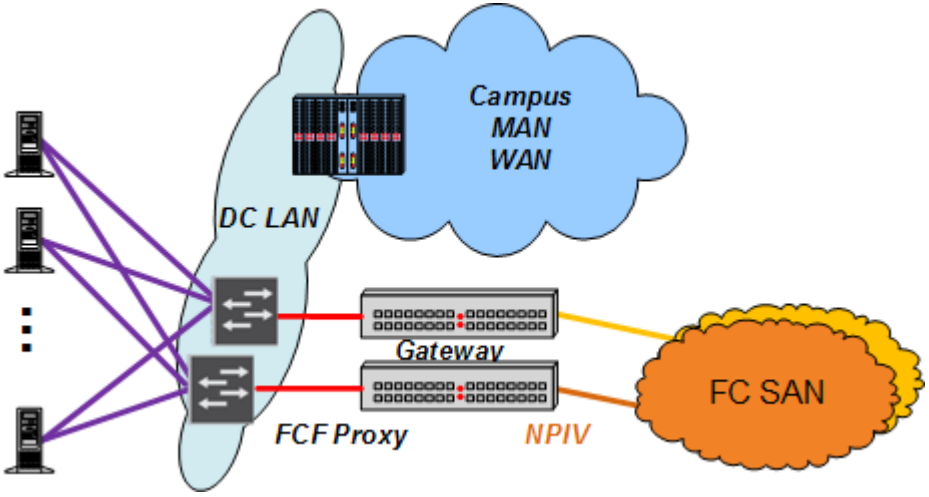

- server L2 to FCF

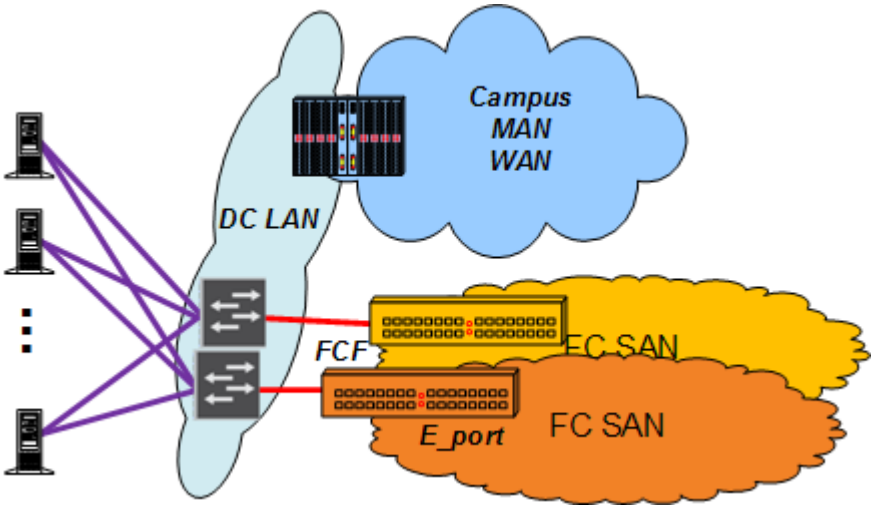

Making Convergence Happen

- Convergence needs the following
  - Transport Convergence
    - DCB for Ethernet is 'complete'. Deployments happening.
    - Switch aggregation is happening
    - Tier Collapse is happening
  - Protocol Evolution
    - This has happened with iSCSI refined by DCB
    - This is happened with FCoE (FC-BB-5) refined by FC-BB-6
  - Operational
    - This has started with I/O Consolidation out of server
    - Better Connected L2 networks
    - Collections of physical switches acting as single logical switch

IF we use FCoE for Network Convergence

- Scaling
  - FCF at the TOR hits Domain Scaling problems
  - FC-BB-6 intends to address this with FCF-FDF distributed domains
- Configuration of L2/L3 separation
  - VLAN → Virtual Fabric mapping
  - L2 network as access between server and FCF
- Full FC topology overlay onto data center network
  - Multi-hop FCoE across multiple L2 networks and several FCFs
  - Multiple VLAN/Virtual Fabric Configuration and Management

IF we use iSCSI for Network Convergence

- Services
  - services means discovery, access control, notifications, security
  - iSNS exists for these but not commonly deployed
- Management
  - Tools have limited support for iSCSI attach
- Performance
  - to match local data center SAN requirements need DCB
  - this is true even though TCP/IP is the transport

Protocol Independent Convergence Considerations

End to End performance

- bandwidth & latency & frame processing rate
- congestion handling
- logical overlays onto physical infrastructure

MAN and WAN extension for storage

- Under convergence the local network is Ethernet. The devices that make up the local Ethernet are not typically suited to directly support lossless distance extension.

Reminder: The Protocols

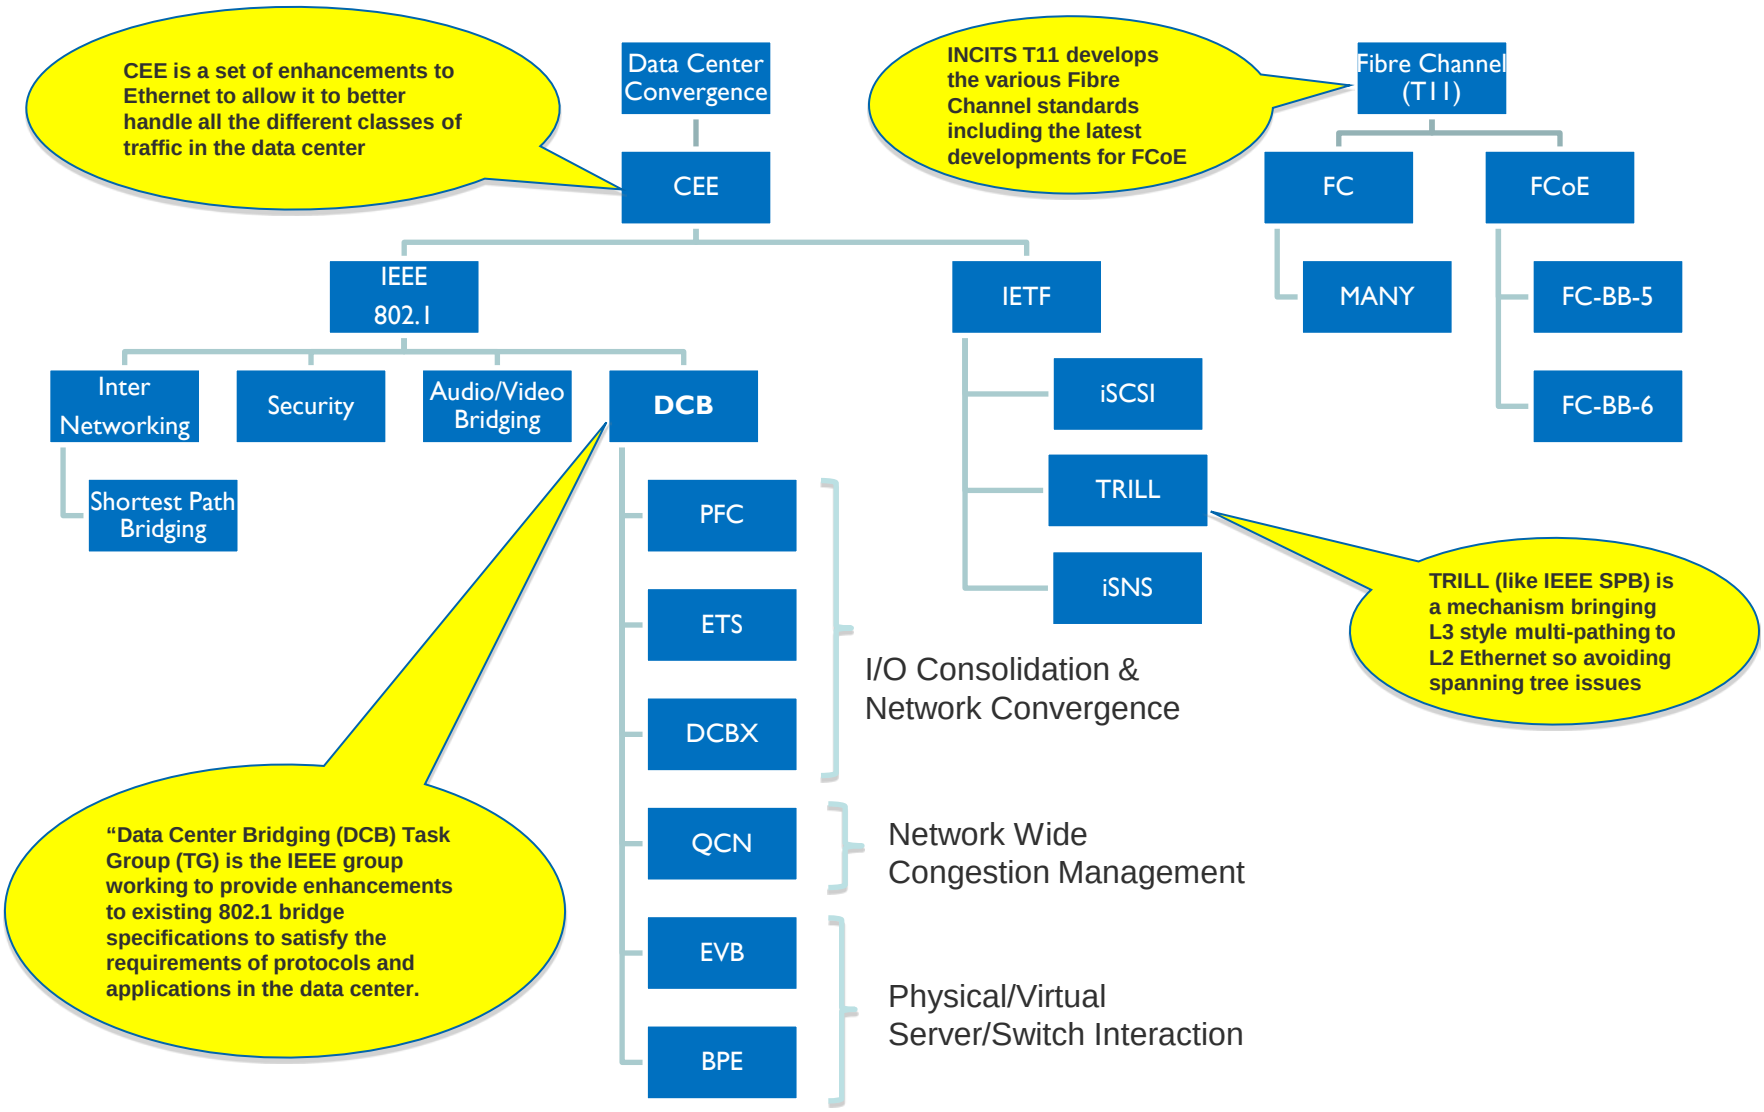

Full Network Convergence Benefits

- results in logical overlays for forwarding on single, shared HW infrastructure

- Benefits
  - stocking of spare FRUs
  - combined operations
  - fewer stranded resources
  - better utilization
  - lower latency
  - better flexibility

Visit the Hands-On Lab:

Unified Storage Infrastructure,
Unified Storage IP Solutions

Inhibitors and Accelerators

If convergence is so great why hasn't it already happened?

Actually it is in progress,

but there is a balance between various factors influencing the adoption rate

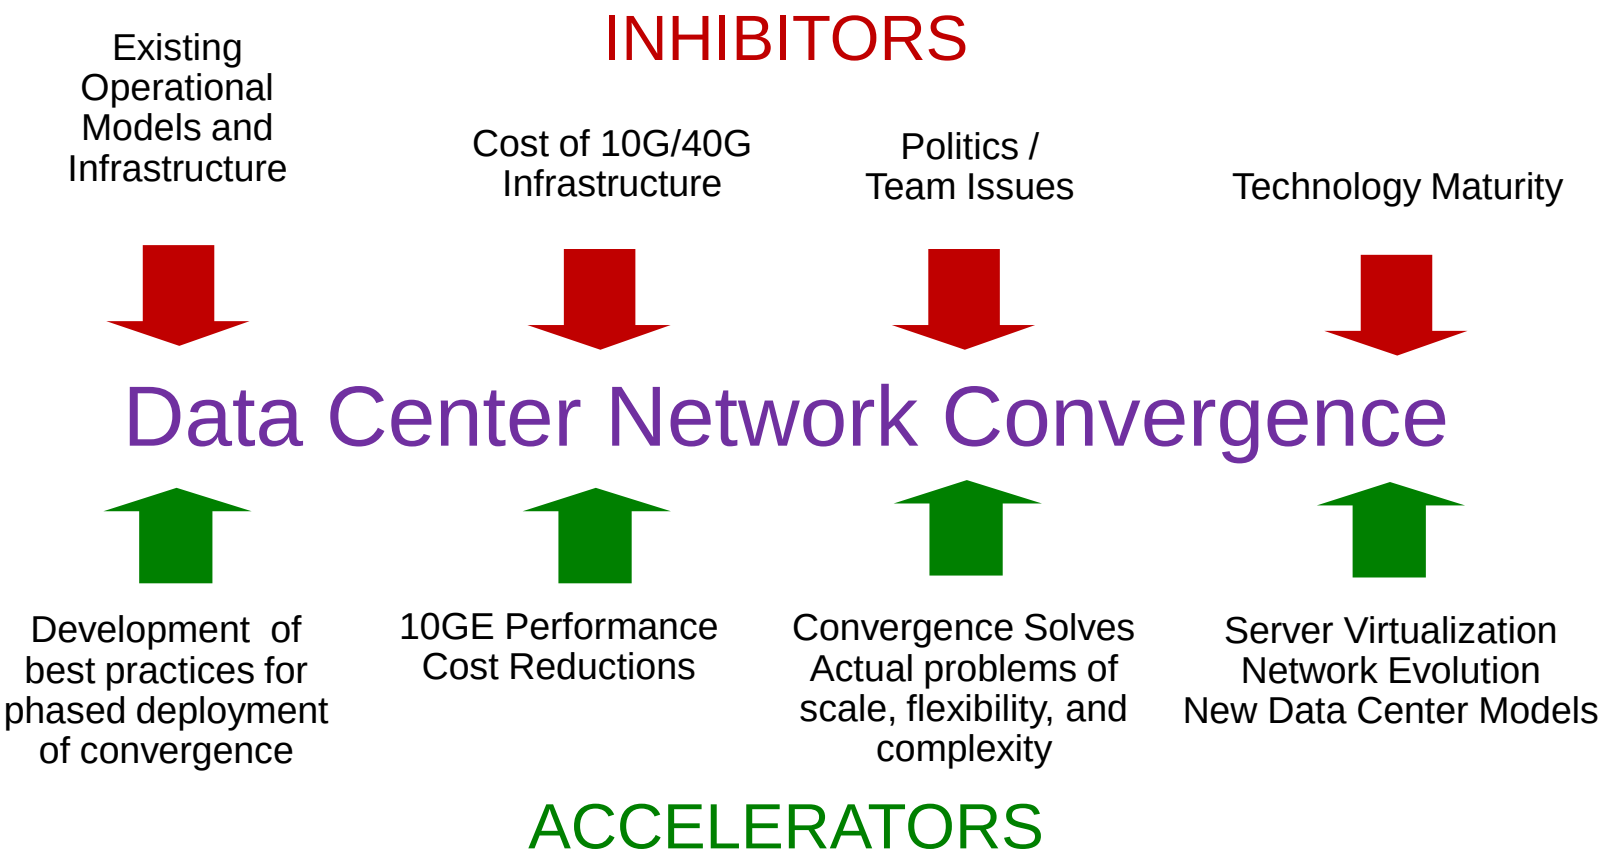

Q&A / Feedback

- 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

Joseph L White
Simon Gordon
Gunes Aybay
Charles Waters
Andy Ingram