How VN2VN Will Help Accelerate Adoption Of FCoE

SNIA Ethernet Storage Forum

David L Fair / Intel
Joseph 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

How VN2VN Will Help Accelerate Adoption Of FCoE

- VN2VN is an enhancement to the ANSI T11 specification for Fibre Channel over Ethernet (FCoE) that promises to significantly reduce the cost of implementing an FCoE SAN. VN2VN allows end-to-end FCoE over an L2 network of DCB switches. VN2VN is part of the T11 FC-BB-6 draft standard that Completed T11 Letter Ballot in Jan/2013, and is now in T11 comment resolution process.
- Anticipating the value of VN2VN, some vendors have already started to release VN2VN capable products. Customers and bloggers are starting to discuss the impact of VN2VN on their environments.
- This presentation will begin with a brief overview of what VN2VN is and proceed to elucidate its real capabilities in illustrative usage models. How VN2VN can be used in typical customer deployments will be shown.

Don't forget the Hands-On LABs

- Virtual SAN
- Architecting and Implementing a Fibre Channel SAN

Why VN2VN ?

What Is VN2VN?

- **V**irtual **N**_Port **to V**irtual **N**_Port (VN2VN) FCoE allows endpoint ("ENodes") communication without a Fibre Channel Forwarder (FCF)
  - There are two modes of operation:
    - Point to Point which is similar to FC Point to Point
    - End point mesh which is logically similar to FC Private Loop
- VN2VN is specified in T11 FC-BB-6
  - FC-BB-6 adds a number of other unrelated enhancements as compared to FC-BB-5
  - Completed T11 Letter Ballot in Jan/2013, now in T11 comment resolution process
- VN2VN allows FCoE to operate over a pure Ethernet network
  - "Ethernet Only" – forwarding based on constructed MAC addresses
  - *In practice* requires switches to support Data Center Bridging (DCB)
- So clarifying some terminology to avoid confusion
  - **VN2VF** is FC-BB-5 ENode connecting to an FCF
  - **VE2VE** is FC-BB-5 FCF connecting to another FCF (aka FCoE ISL)
  - **VA2VA** is FC-BB-6 FCF/FDF distributed switch connectivity

Why VN2VN ?
- Simplicity and cost savings

- With FC-BB-5 one should see some level of cost saving
  - Eliminating dedicated Fibre Channel HBAs in the server
    - Only 10GbE CNAs with DCB are required
    - These increasingly come as standard on new servers
  - Sharing some LAN and SAN network cabling and switching costs
    - May be a partial overlap with converged access deployments
    - May be a full overlap with end to end converged networks
  - However server and network convergence savings are mitigated by
    - Added hardware cost of Fibre Channel Forwarding (FCF) switch hardware
    - Licensing fees for enabling necessary hardware and software functionality
- VN2VN increases savings…
  - By not needing Fibre Channel Forwarding (FCF) switches
    - DCB is becoming standard for 10GbE+ enabled servers and switches
    - Reduction in management overheads

Details Of VN2VN

Including the most recent enhancements to the draft standard

Fibre Channel Initiation Protocol (FIP) for VN2VF (FC-BB-5)

- FIP VLAN Discovery (optional)
  - Find which VLANS have FCoE
- FIP Discovery
  - Find the FCFs within the VLAN
- FCoE Virtual Link Instantiation
  - Login to the SAN and get your FCID and FCoE MAC
  - Setup the VN2VF connection
- FCoE Virtual Link Maintenance
  - Keep everything up and running
  - Ensure links are brought down correctly
- *NOTE: FCF is part of both control path and data path*

Fibre Channel Initiation Protocol (FIP) for VN2VN (FC-BB-6)

- FIP VLAN Discovery (optional)
  - Find which VLANs' have FCoE VN2VN
- *Discovery Phase (per L2 Ethernet Network/Domain)*
  - *Locally Unique End Point ID assignment*
    - *As no longer assigned from the FCF*
  - *End Point Learning / Discovery*
    - *As no name server in the FCF to query*
- FCoE Virtual Link Instantiation
  - Setup the VN2VN connections
- FCoE Virtual Link Maintenance
  - Keep everything up and running
  - Ensure links are brought down correctly

Comparing VN2VF and VN2VN

- For the Ethernet/IP network people
  - VN2VN is like IP within an Ethernet domain
    - End point driven automated address allocation and device discovery
  - VN2VF is like IP cross subnet / through routers
    - Along with the other services that make it work (DHCP, DNS, etc)
- For the long term FC people
  - VN2VN is like good old point to point and (logical) private loop
    - Though with the benefits of a switched infrastructure and lessons learned
  - VN2VF is good old classic FC with Ethernet as the transport
    - Leveraging Ethernet as a new physical layer with multipoint connectivity
- *VN2VN and VN2VF can be used separately or together*
  - ***Very*** *useful for the different deployment use cases*

VN2VN VLAN Discovery - Very Similar To BB-5 VN2VF

- As in BB-5, VLAN Discovery is an optional phase
  - Unified VN2VN/VN2VF VLAN discovery process
  - Devices discover the FCoE VLANs
  - Devices learn which FCoE VLAN is which type
  - Identification is based on the addresses used for discovery
- Discovery is through broadcast on the default VLAN
  - See who responds and on what VLAN they respond
  - Broadcast range can be controlled through Native VLAN settings and port membership of that VLAN

Automated address allocation and device discovery

- Through this process an ENode configures itself and learns about each of the L2 Ethernet domains it can see
  - Determine per Domain locally unique FC_ID
  - Discover other devices and their capabilities in those domains
- Each ENode generates a candidate FC_ID
  - FC_IDs are local to a given L2 Domain (or VLAN)
    - So this is repeated for each domain in which an ENode is present
  - In the range 000001h to 00FFFEh
    - 00xxxxh is not used in FC-SW which uses domain IDs 1 through to 239 but not 00
  - ENodes run a protocol to ensure that there is no address collision
    - In the Probing State, tentative FC_IDs are tested for potential collisions
    - Passing that, in the Claiming State, a notification is broadcast and responses monitored
    - Passing that, in the Operational State, FIP FLOGI and FCoE PLOGI begin
    - Periodic beacons continue to insure the FC_ID remains locally unique

VN2VN Link creation

- Each VN port can connect to multiple other VN ports
  - Just as was the case with private loop
- Devices can create multiple VN ports in the same VLAN
  - For instance you may want one for each virtual machine
  - Another way for one ENode to connect to multiple other ENodes
- Devices can create VN ports in multiple VLANs
  - Each VLAN is a separate VN2VN Domain
  - May have different reach across the network
  - *Allows Logical Dual Rail for HA, Multi-tenancy, Private Storage*

Applying Target Driven Zoning *concept* from FC-SW-6 to VN2VN in FC-BB-6

- Traditional FC (FC-SW-6) provides for Target Driven Zoning (TDZ)
  - To allow a storage array to control SAN visibility to itself
- FC-BB-6 providing a comparable feature for VN2VN
  - Allows an ENode (*typically target*) to indicate that a discovering ENode (*typically initiator*) should not connect – please don't talk to me !
    - Would likely align with the arrays' port ACLs and LUN masking
    - If ignored then the storage array would likely reject connection anyway
- Improves scale attributes
  - Mitigates endpoint n^2 problems (everyone talks to everyone)
  - Sessions only formed where server allowed to access storage
- Simplifies Storage Management Security and Provisioning
  - Concentrates fabric management in the storage array, which is already managed by storage provisioning software that configures LUNs and establishes server access rights.

What About Security in a converged network ?

- FIP Snooping applies to VN2VN as well as VN2VF
  - Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) Snooping is an optional feature but strongly recommended
  - Switch establishes firewall filters to prevent unauthorized network access by unknown or unexpected virtual N_Ports transmitting FCoE traffic – both for setup and removal of VN2VN connections
- FIP snooping (both VN2VN and VN2VF) also makes the network more resilient to events like a cable or VM move
  - Assisted MAC Learning / Unlearning as well as ACL protection
- Enhanced monitoring and manageability
  - Device and connection learning for monitoring/support
  - Keep-alive monitoring
  - May monitor and act on FCoE primitives as well as FIP

Ethernet Rules clarified in FC-BB-6 (VN2VF and VN2VN)

- FC-BB-6 has an informational annex
  - Clarifies many points of discussion from FC-BB-5
  - Applies to both FC-BB-5 and FC-BB-6
  - Describes the implications of various implementation choices
- FIP Traffic guidelines
  - The FIP Ethertype is distinct from the FCoE Ethertype
  - In-order delivery is not required
  - Lossless (ie PFC) not required
  - FIP and FCoE can be on the same or different priorities
- FCoE (non FIP) Traffic guidelines
  - In-order delivery (within an exchange only) is required
  - Lossless (ie PFC) is required

Usage Models For VN2VN

Some interesting examples

Direct (and pseudo direct) Attach - Non Shared Storage

- FCoE VN2VN can be used to DAS an individual single disk
  - Limited use case
- FCoE VN2VN can be used to DAS a disk subsystem
  - Large database system
- FCoE VN2VN can be used to connect an entire JBOD to a single server
  - Logically like a private loop but packet switched data path

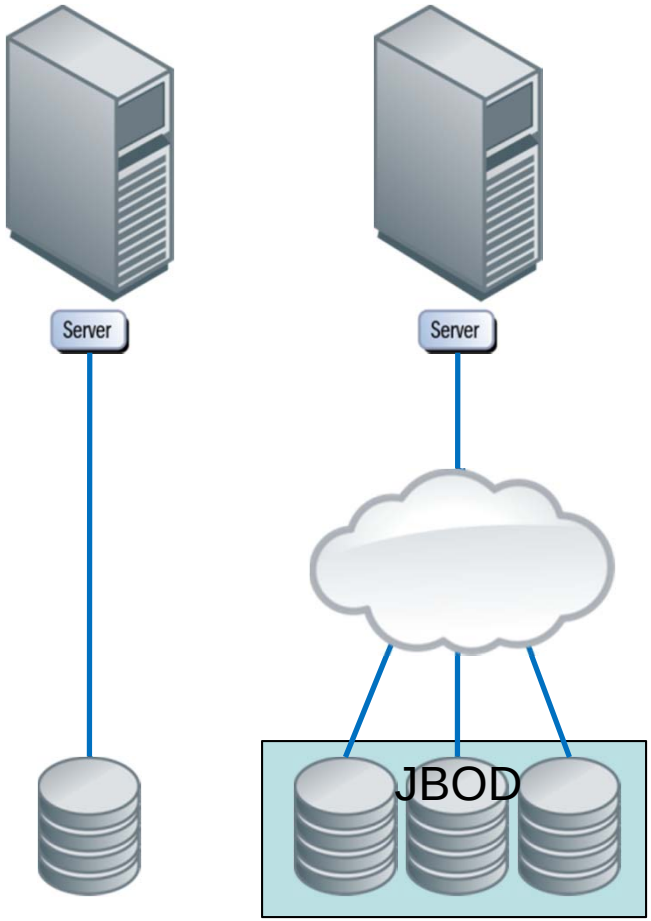

Small Business Deployment - simplicity is the key factor

- Simple DCB Switches
  - PFC, ETS, DCBX
  - FIP Snooping
- Simple Logical/Physical Dual Rail
  - Dual FCoE VLANs/Switches
- Unified Management
  - Simple plug and play
  - Configure the storage and off you go
- Managed by single team
  - Small business, small team

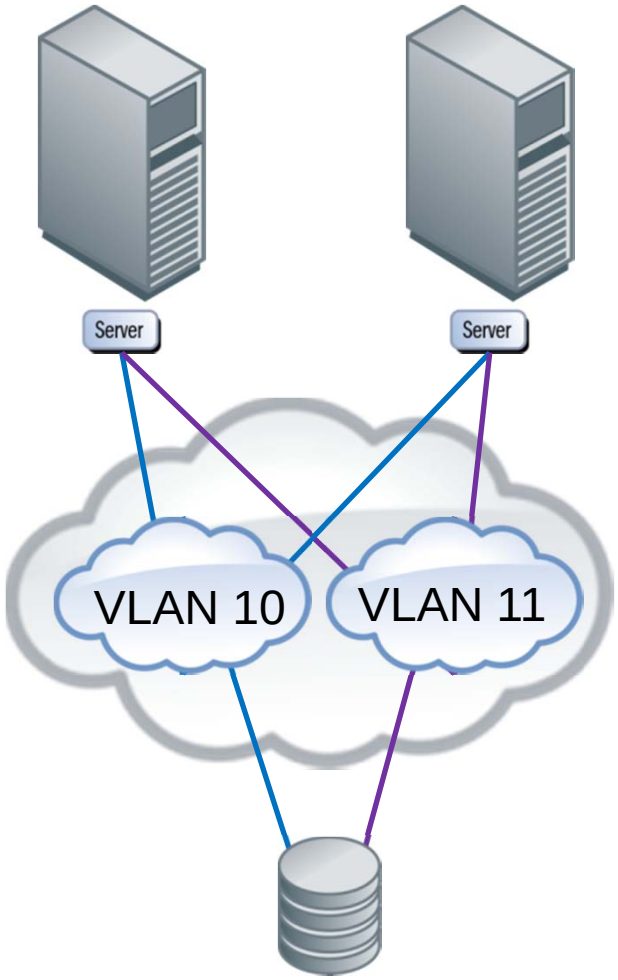

HPC Deployment
- High performance compute

- VN2VN allows direct use of low latency Ethernet switching infrastructure
  - Low latency Server to Server
  - Low latency Server to Storage
- Storage access via simple efficient end points
  - All the benefits of FC/FCoE
  - No TCP/IP Stack required
- Great fit with flash storage
  - Low latency disk access
  - May use cluster file system

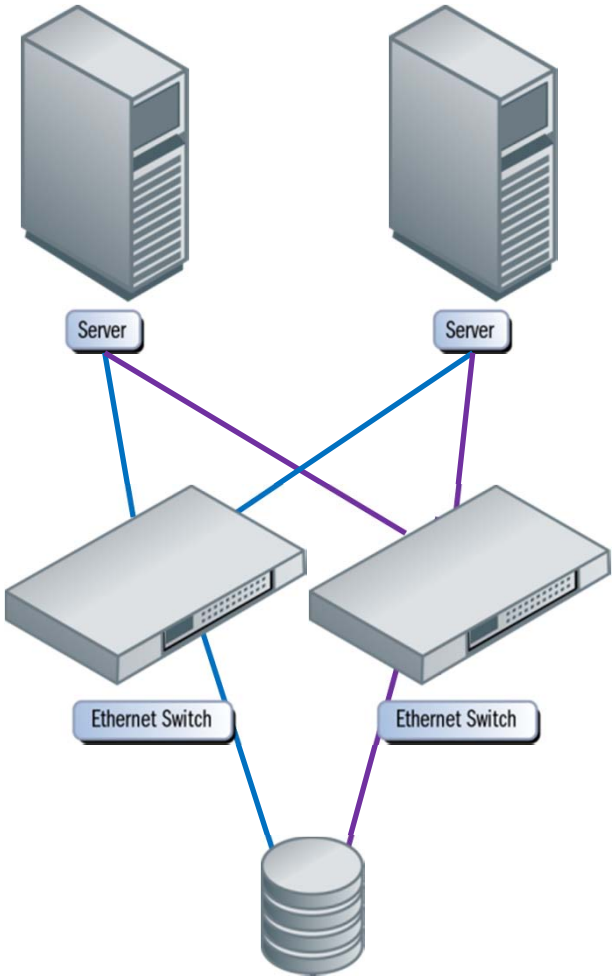

IAAS / Virtualized Datacenter - common multi-tenancy approach

- Simultaneous local/private and central/shared storage
- Locale/Private via VN2VN
  - Locally switched
  - Server Boot, VM Storage
  - Private / segregated client
- Central/Shared via VN2VF
  - Traditional High end disk
  - Existing Tape, archive, etc
- Unified Management
  - Leverage VLAN, QinQ etc. for all multi-tenancy needs

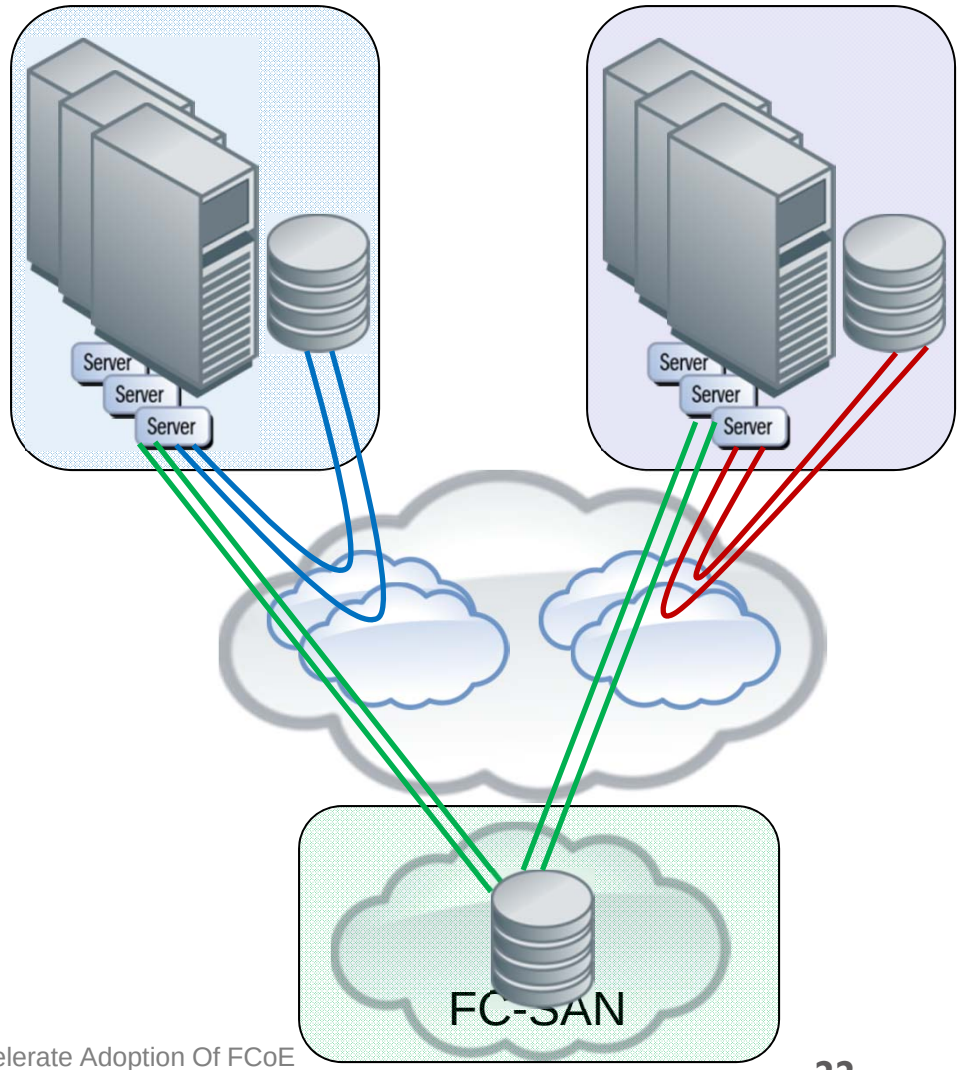

Private SANs
- building a storage subsystem

- Internal to storage array
  - Needs lots of simple connectivity
- Disaggregated storage array
  - Need lots of simple connectivity
- Storage virtualization
  - The SAN between the virtualizer and the storage
- Software defined storage
  - A conversation for another day

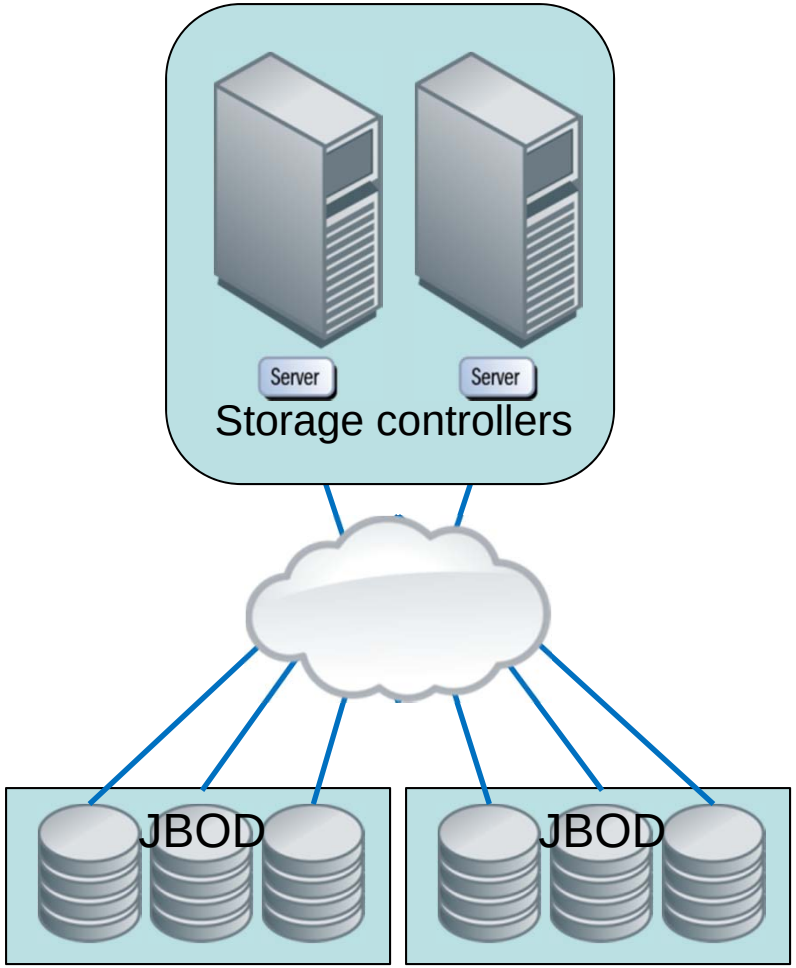

Private SANs
- Tape SAN

- Tape Library
  - Large collection of independent tape drives
  - Library Controller & Tape Load mechanisms
- Connectivity requirement
  - Often dozens of devices
  - Accessed from a few media servers

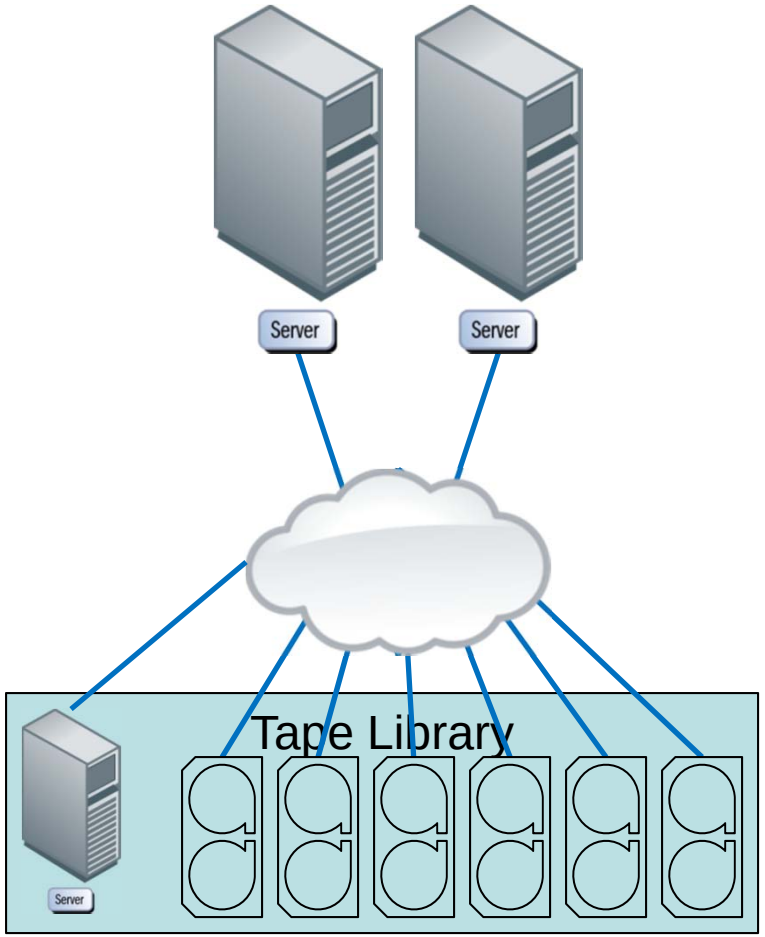

An interesting private SAN use case - pNFS using VN2VN

- pNFS allows flexibility in its choice of block storage protocol
  - Aka the data plane protocol
  - Already uses iSCSI and FC
- VN2VN would work well
  - Simple visibility rules
  - Access is mediated

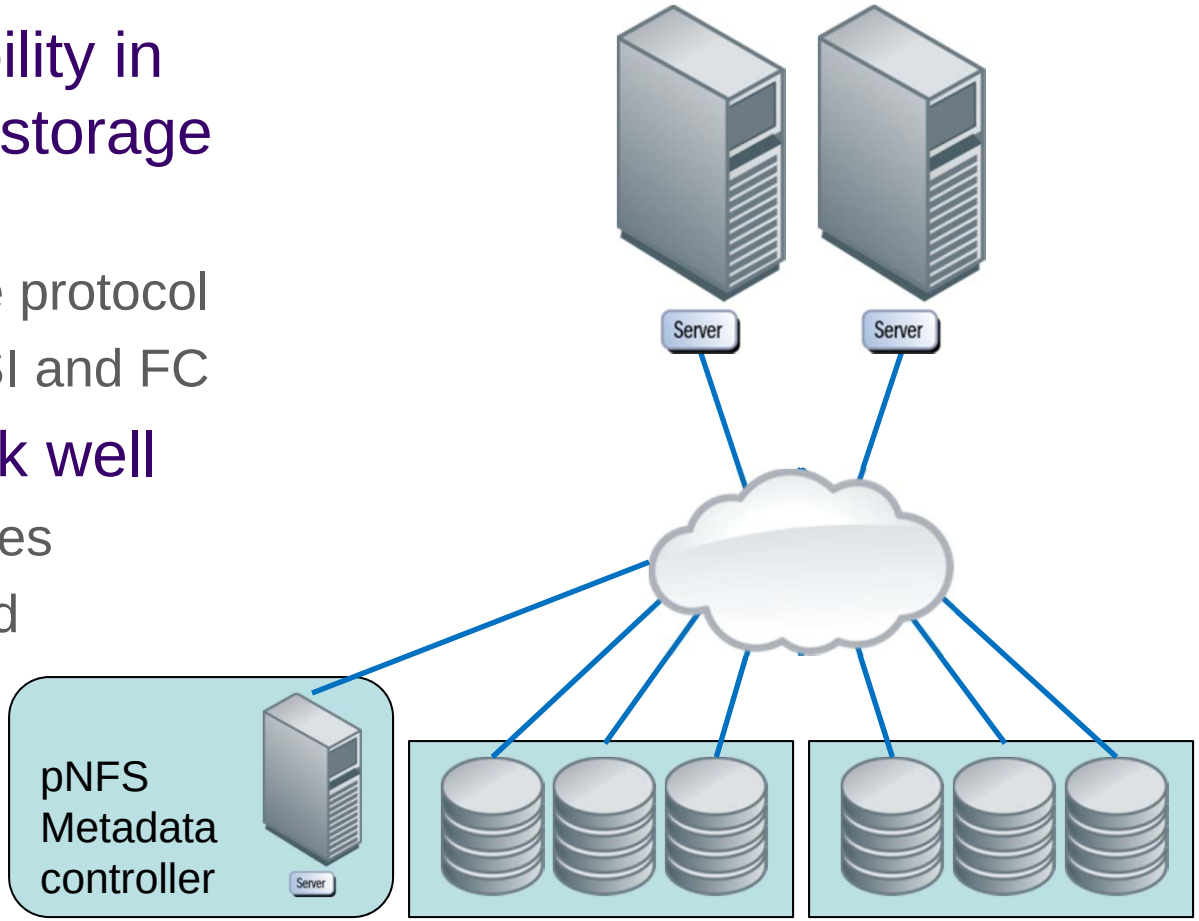

Open-FCoE Linux with VN2VN

VN2VN has been incorporated in to the Open FCoE Linux code (www.open-fcoe.org)

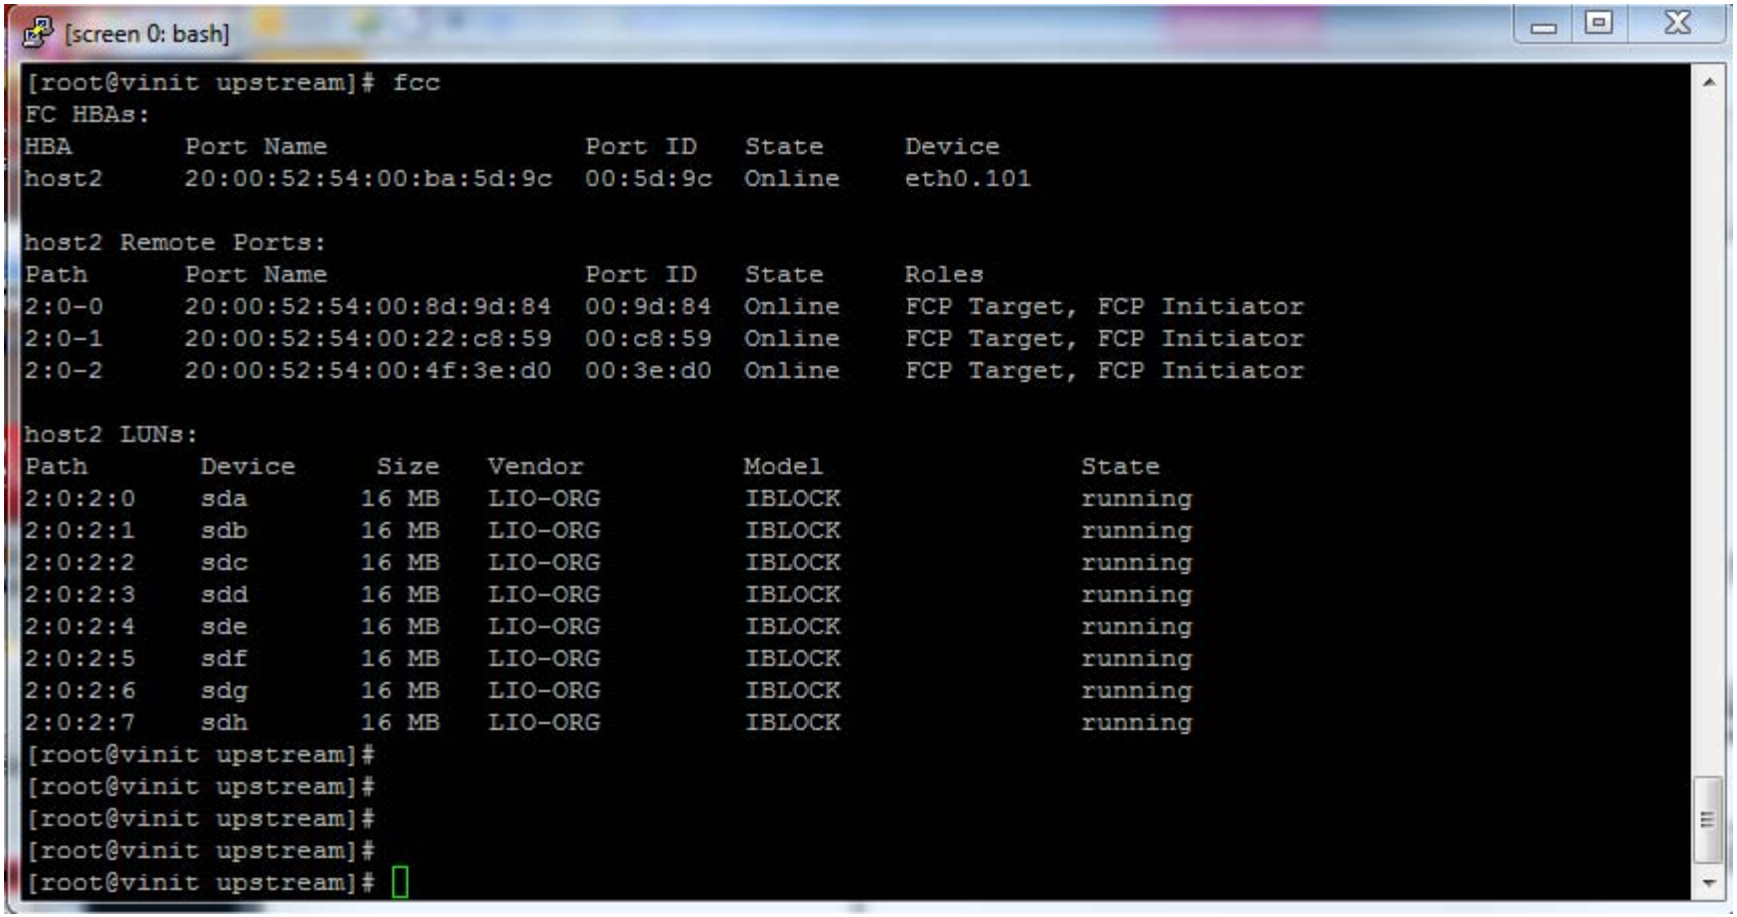

Windows FCoE stack with VN2VN

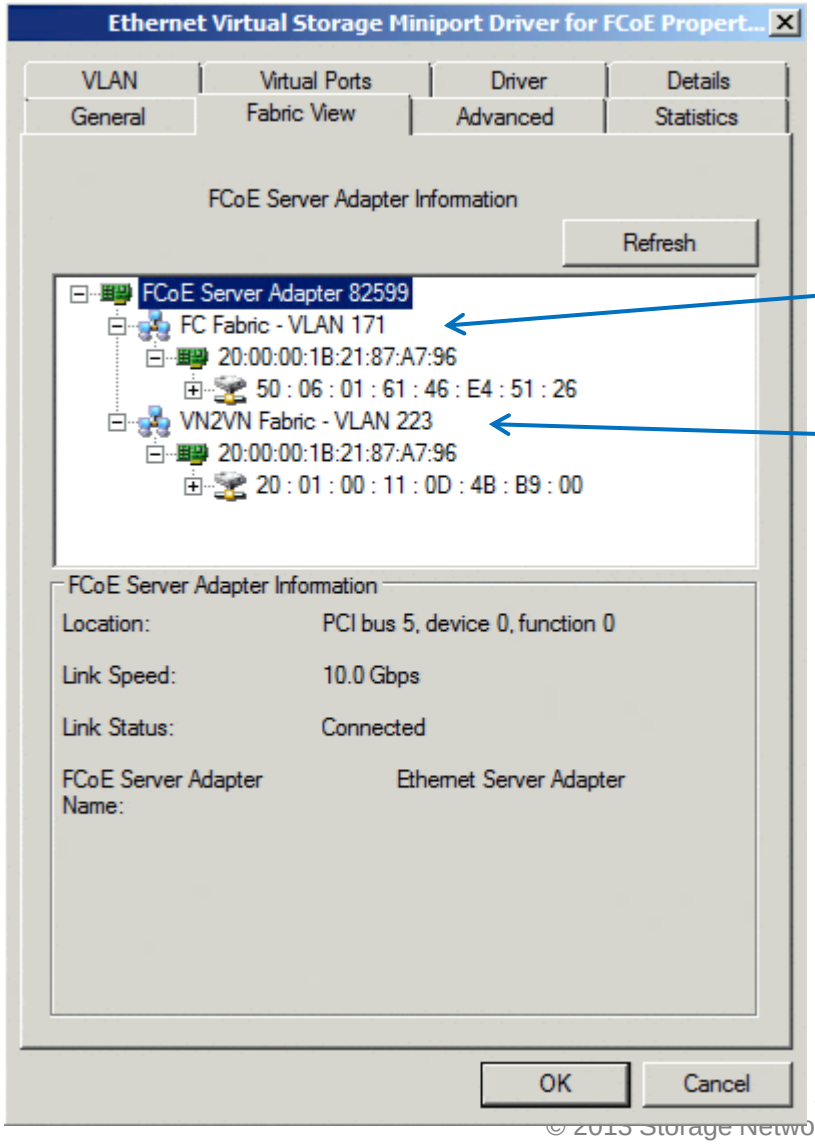

CNA logged on to an FCF

Same CNA logged on to an VN2VN fabric at the same time

Best Practices For VN2VN Deployment

Dual-Rail Deployments

- Leverage Ethernet Layer 2 Domains
  - Can be physical dual rail
    - Common in traditional FC
  - Can be logical using VLANs
    - Common in traditional LANs
- Control Plane VLANs'
  - VLAN Discovery starts on the default VLAN
    - Native VLAN (one per rail) can be used to control discovery of FCoE VLANs'
    - So provides both control plane & data plane logical dual rail
- Data Plane VLANs'
  - Dedicated to FCoE Traffic
    - Non FCoE Traffic can be blocked
    - FIP Snooping secures FCoE Traffic
  - Logical Dual Rail Enforcement
    - See Control Plane VLANs'

Managing FCoE VLAN Usage and security

- Storage array as the anchor point
  - Manually configure storage in to desired FCoE VLANs
  - Block FIP on non FCoE VLANs
  - Leverage VN2VN "Target Initiated Zoning" equivalent
    - Centralize all storage access control mechanisms with the storage array
- Allow the servers to discover the storage device VLANs
  - Server side auto configuration
  - Use of native VLAN to force logical dual rail
  - Use of native VLAN to create SAN islands if desired
  - Only useful VN2VN initiator/target connections formed

What About Security in a converged network ?

- Security can be applied at multiple levels and places
  - Ethernet based Security
    - 802.1X port based network access control
    - VLANs membership
    - Vendor specific switch features
  - End to End Fibre Channel Security
    - FC-SP from source to destination device
    - Enhanced to protect from man in the middle attacks
  - Storage Array based security
    - Port ACLs' and LUN Masking
    - Target initiated zoning
- All in addition to FIP Snooping
  - As recommended in the FC-BB-5 and 6 standards

Ethernet QoS

- With VN2VN just like VN2VF you have IO consolidation and network convergence so you have to configure your network and endpoints appropriately
  - Lossless (PFC) Configuration
    - Select one or more priorities to be lossless
  - Hierarchical Bandwidth Allocation
    - Port based traffic shaping
    - ETS (group of priorities) bandwidth allocation
    - 802.1p based (individual priority) bandwidth allocation
    - VLAN based bandwidth allocation
  - Buffer / Queue Configuration
    - Based on distance and burst absorption requirements

FCoE SAN Design/Operation

- Follow many of the usual FC SAN best practices
  - Balance the traditional FC and Ethernet best practices
- In other words – *Check out other SNIA Tutorials*

  - Sensible topologies
    - High enough bandwidth, Low enough latency
    - Not too many hops (in this case layer 2 Ethernet hops)
    - ***"The Confusion Arising from Converged Multi Hop FCoE Topologies"***
  - Well thought out high availability
    - Physical and Logical Dual Rail, FC or Ethernet Practices
    - *"a new tutorial is needed here"*
  - Sensible operational deployment
    - Configuration control procedures
    - Appropriate management & monitoring tools
    - ***"The Unintended Consequences of Converged Data Center Deployment Models"***

In summary

- VN2VN allows you to build practical SANs for a wide variety of use cases at lower complexity than traditional FC/FCoE SANs…
- In deploying VN2VN you have considerable flexibility in adopting traditional FC or traditional Ethernet best practices depending on the use case IT organisation
- Could be net-new SAN deployments, or cap & manage of existing FC/FCoE SANs, or part of a cap & migrate strategy

Attribution & Feedback

The SNIA Education Committee thanks the following individuals for their contributions to this Tutorial.

Authorship History

Original Authors:

David L Fair (ESF) / April 2013
Simon Gordon (ESF) / April 2013

Updates:

Additional Contributors

Prafulla Deuskar (Intel)
Joseph White (Juniper)
Mark Wunderlich (Intel)

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***