

......

The future of SDN: Transforming the REN in support of Big Data

Oct 4th, 2013

Eric Choi Senior Product Manager, Data Centre Business Unit APAC and Japan





Transforming REN Across Data Centres



C © 2013 Brocade Communications Systems, Inc. Proprietary Information

Data Intensive Science

Main Driver for High Capacity Research Education Networks



Large Hadron Collider- ATLAS

- ~10s of terabytes of data transfers
- Data distribution across continents
- Distribution rates at 40+ Gbps
- Requirements increasing over time



Genomics

- Data rates increasing exponentially
- Instrument data rate up 10x in 5 yrs
- Sequencing cost down 10x in 5 yrs
- New applications for genomics data as science improves



Square Kilometer Array

- 100Gbps rates needed to distribute data to international collaborators
- Large radio telescope
- ~1 sq km of combined signal collection area
- ~2800 telescopes in array, ~2 petabytes/sec at central correlator

BROCA

Uniqueness of Research Networks

	Commercial networks	Research/collaboration network
#flows	Many (mouse flows)	Few (elephant flows)
Flow sizes	Typically follows Internet mix (IMIX)	Very large- up to 50% of link bandwidth in a single flow
TCP effectiveness	Very good	Weak: result in loss sensitivity, increased CPU load



Typical 100GbE Forwarding Engine Architecture Suitable for internet traffic but not elephant flows in REN



COMPROMISED NETWORK PERFORMANCE DUE TO UNEQUAL LOAD BALANCING

DIVOCADE

Packet Buffer for 10Ge/100Ge port

Uniqueness of Research Networks



SUFFICIENT PACKET BUFFER IS REQUIRED TO TCP SESSION BACKOFF

© 2012 Brocade Communications Systems, Inc. CONFIDENTIAL—For Internal Use Only

Internet2

Requirement of On-Demand L2 Connectivity





Internet 2

- 49 Custom Location Facilities
- 15,500 miles of dark Fiber
- 8.8 Tbps of Optical Capacity
- Hybrid Mode with protected

SOME OF THESE PROBLEMS CAN BE ADDRESSED BY PROGRAMATIC NETWORK CONTROL

© 2013 Brocade Communications Systems, Inc. Company Proprietary Information

DIVOCADE

OpenFlow-Enabled Device Operation

How it works?

- The Flow Table contains Flow Entries Ma
 Each Flow Entry represents a Flow, e.g., packets with a given destination IP address
- Incoming packets are matched against the Flow Entries in priority order
- Matching stops when a match occurs, and the <u>actions</u> for that Flow Entry are performed
- Unmatched packets are typically dropped





SDN WAN Use Case 1

"Protected" Production IP Network with OpenFlow Overlay



- Hybrid port mode
 - OpenFlow overlay runs concurrently with traditional MPLS/IP routing
 - OpenFlow enabled on existing production network
 - Protected
 - OpenFlow experimentation by researchers does not affect production traffic
 - OpenFlow/Production traffic isolation in hardware



Customer 1

Customer 2

100 GE and 10GE WAN

DC 2

WAN Virtualization

DC1

SDN WAN Use Case 2

"Unprotected" Production IP Network with OpenFlow Overlay



- Hybrid port mode
 - OpenFlow overlay runs concurrently with traditional MPLS/IP routing
 - OpenFlow enabled on existing production network
- Unprotected
 - OpenFlow allowed as an "exception rule" to MPLS/IP forwarding
 - Service Provider can add new revenue-generating features on top of existing production network

BROC

Hybrid Mode Operation Flow Packet flow

Priority of lookup for hybrid mode ports are based on below order (high to low).

- Packets with Protected vlans
- Packets matching Openflow rules
- Packets with Configured unprotected vlans
- All other packets



Feature Reference Diagram

Openflow Hybrid port





MAP NI 5.5

© 2013 Brocade Communications Systems, Inc. CONFIDENTIAL—For Internal Use On



Transforming REN within Data Centre



10/7/2013

1 © 2013 Brocade Communications Systems, Inc. Proprietary Information

Software Defined Networking: Overlay Networking



- East-west tunnels between software switches (vSwitch) enable physical network abstraction (logical networks)
- Proposed tunnel technologies
 - VxLAN, STT: VMware
 - NVGRE: Microsoft
- Technology benefits
 - Supports millions of logical networks (no 4K limit as VLANs)
 - Support VMs with overlapping IP/MAC addresses



External Storage Capacity by Protocol (PB) IDC: By 2017, FC remains on top



Worldwide External Enterprise Storage Capacity

IDC Worldwide Enterprise Storage Systems Forecast Update, May 2013



How can we connect them together?

- Storage is still physical and non-VxLAN aware
- Some applications running on physical servers like Oracle DB etc.
- Existing appliances like firewalls and server loadbalancers.
- Access to existing L3 networks via existing routers.



BROC

Network (Overlay) Virtualization Introducing VxLAN Gateway

OVERLAY CONTROLLERS

vmware^{*}



Network Virtualization with VxLAN GW & VMware NSX

Solution Overview and Components



- 1. VXLAN Gateway Bridges virtual and physical infrastructure
- 2. VMWare NSX: Network Provisioning and Automation
- 3. Ethernet Fabric: Efficient, Resilient foundation for the Network Infrastructure

BROC

Summary

- Two Key Technologies in SDN: Device Based(Openflow) and Overlay Tunnel based.
- Network programmability, openness, control/data plane separation. Near-term use cases on WAN/DC network virtualization.
- Hybrid port mode simplifies transition to SDN
- Network Virtualization using Overlay Tunnels



