Software Defined Networking
Disruptive Technologies

Renato Recio

IBM Fellow &
System Networking CTO

Agenda

- Software Defined Networking (SDN) Defined
- SDN Client Value
- SDN is a Discontinuous & Disruptive Technology
- SDN Adoption Curve Status
  - Current Products
  - Technology Investment Areas
- Summary

10,000 Foot Software Defined Networking Defined

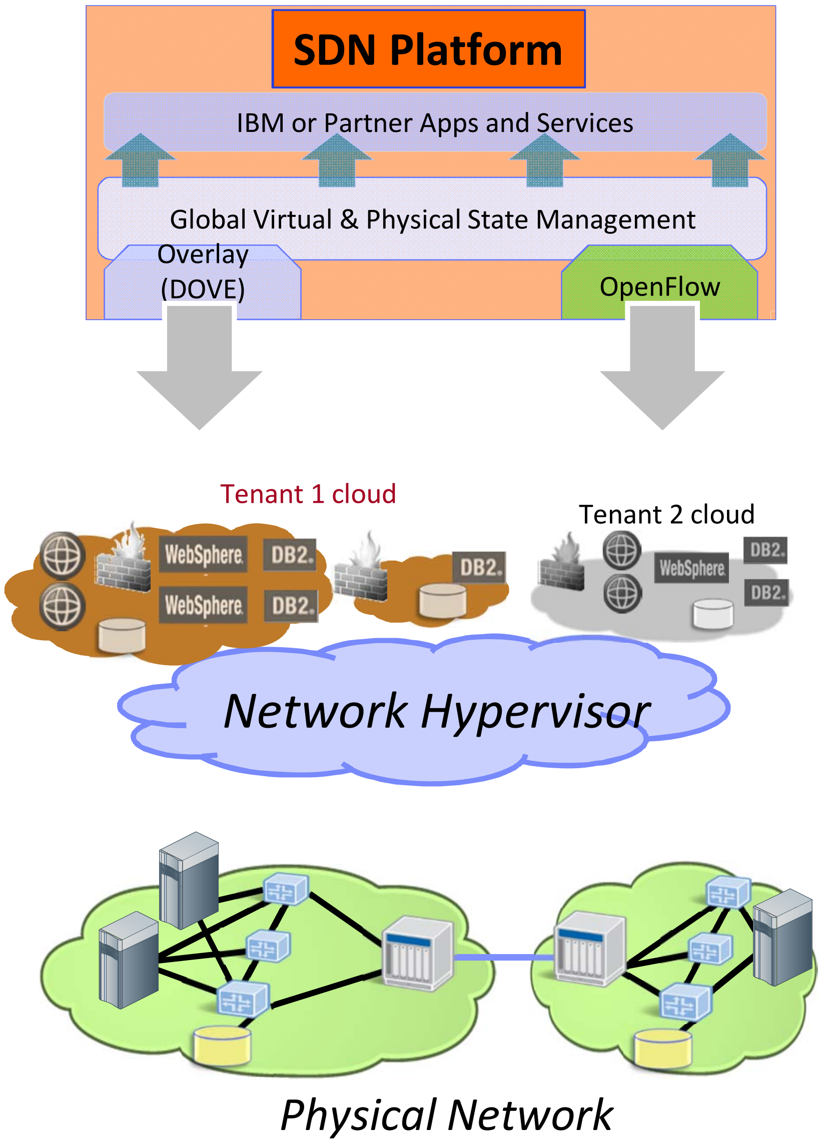

SDN Platform – *SDN Controller*

- Automates connectivity of network services (e.g. Firewall, IPS) used in multi-tier virtual systems, with multi-tenant network capability
- Optimizes traffic performance, availability and separation through fabric pathing services, with global network visibility and control
- Open APIs enable network applications

Network Hypervisor – *DOVE Network*

- Virtualizes the physical network thru a Network Hypervisor that enables a “wire once” physical network, analogous to Hypervisor for compute/IO

Optimized Fabric – *Ethernet & OpenFlow*

- Leverages OpenFlow to move network OS from physical switches to server based controller cluster, enabling: rapid protocol development time, workload aware network optimization, faster convergence times and global control

SDN value to IT Constituents

SDN has value across Data Center constituencies;
overall its disruptive potential is akin to server virtualization.

Constituency	Today's concerns	The benefit of SDN
CIO	Computing models such as cloud and compute virtualization limited by human "middleware" needed to instantiate	Network becomes a "virtualized" asset with automated linkages to computing and applications.
Line of Business Owner	Exploiting analytics within enterprise requires IT agility. Rapid connectivity to new sources of data and resources across-Departments, BU's, Enterprises.	Network can be rapidly reconfigured and modified. Virtualized DMZ concept can eliminate complicated security barriers and limitations
Application developer	Multi-tier compose-able applications require complex interaction between physically distributed systems and resources. Security, quality-of-service, etc need to be enforced.	New connectivity service provides abstracted connectivity model, without tight linkage to physical network configuration
Systems Manager	Provisioning, configuring, monitoring across server, storage, network is very complicated	Simplified, virtual network model makes it much easier to integrate with data center wide management systems
Networking Manager	Inability to evolve network rapidly enough to support changing workloads	Wire-once model limits need for physical network modification
CTO architect/strategist	Long cycle to deploy new, standard (IETF, IEEE) based networking functions	Rapid development cycle, leveraging OpenFlow's control plane separation

SDN Value to Data Centers

Example deployment models

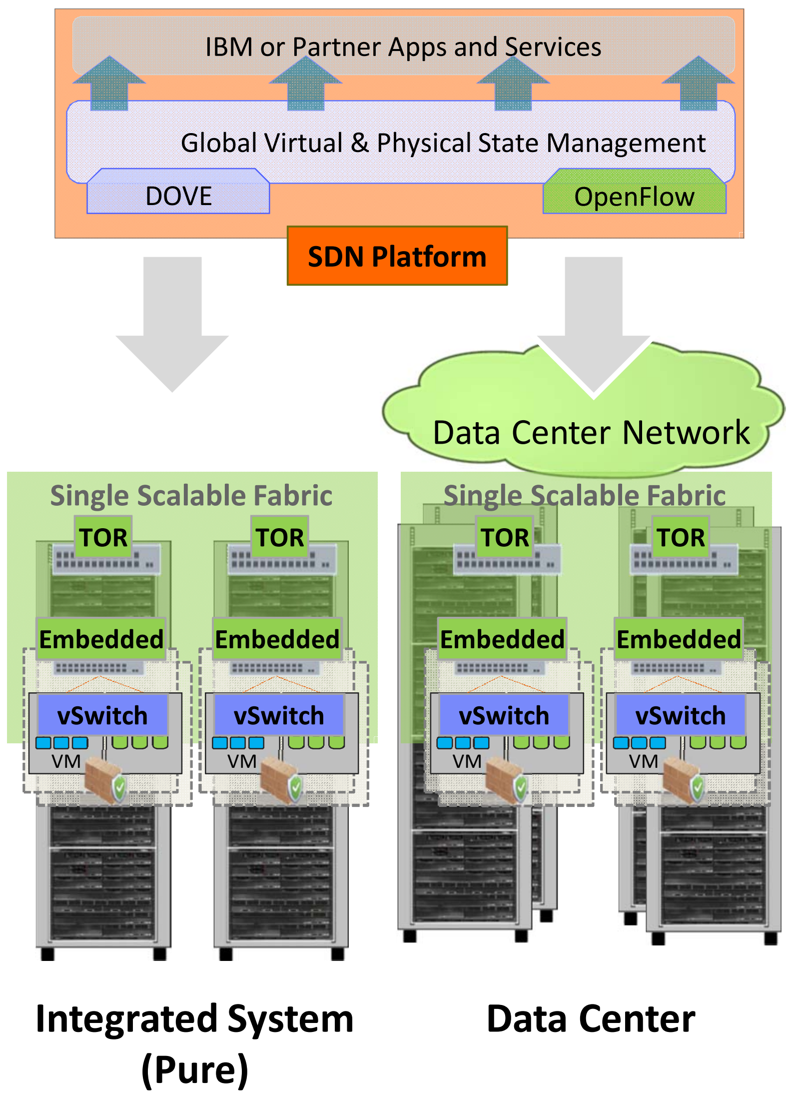

Integrated Network Software

Provide dramatic improvement in business efficiency by reducing application deployment times

Automated Network Virtualization

Provide business agility by making the infrastructure (network connectivity) completely dynamic

Optimized Fabric

Provide finer control of network traffic flow, enabling higher fabric utilization

Discontinuous Technologies

Discontinuity – a: the property of being not mathematically continuous; b: an instance of being not mathematically continuous; *especially a value of an independent variable at which a function is not continuous*

Current Discontinuous Technologies Examples

Discontinuous Technology → Impacted Technology

Flash & SSD → Magnetic Disk

CEE & FCoE → Fibre Channel

SDN* Overlays & OpenFlow → Traditional switching

* One can argue the InfiniBand Subnet Manager was an early SDN controller example.

Sustaining vs Disruptive Technologies

Sustaining – doesn't affect existing markets, evolves existing one's with better value .

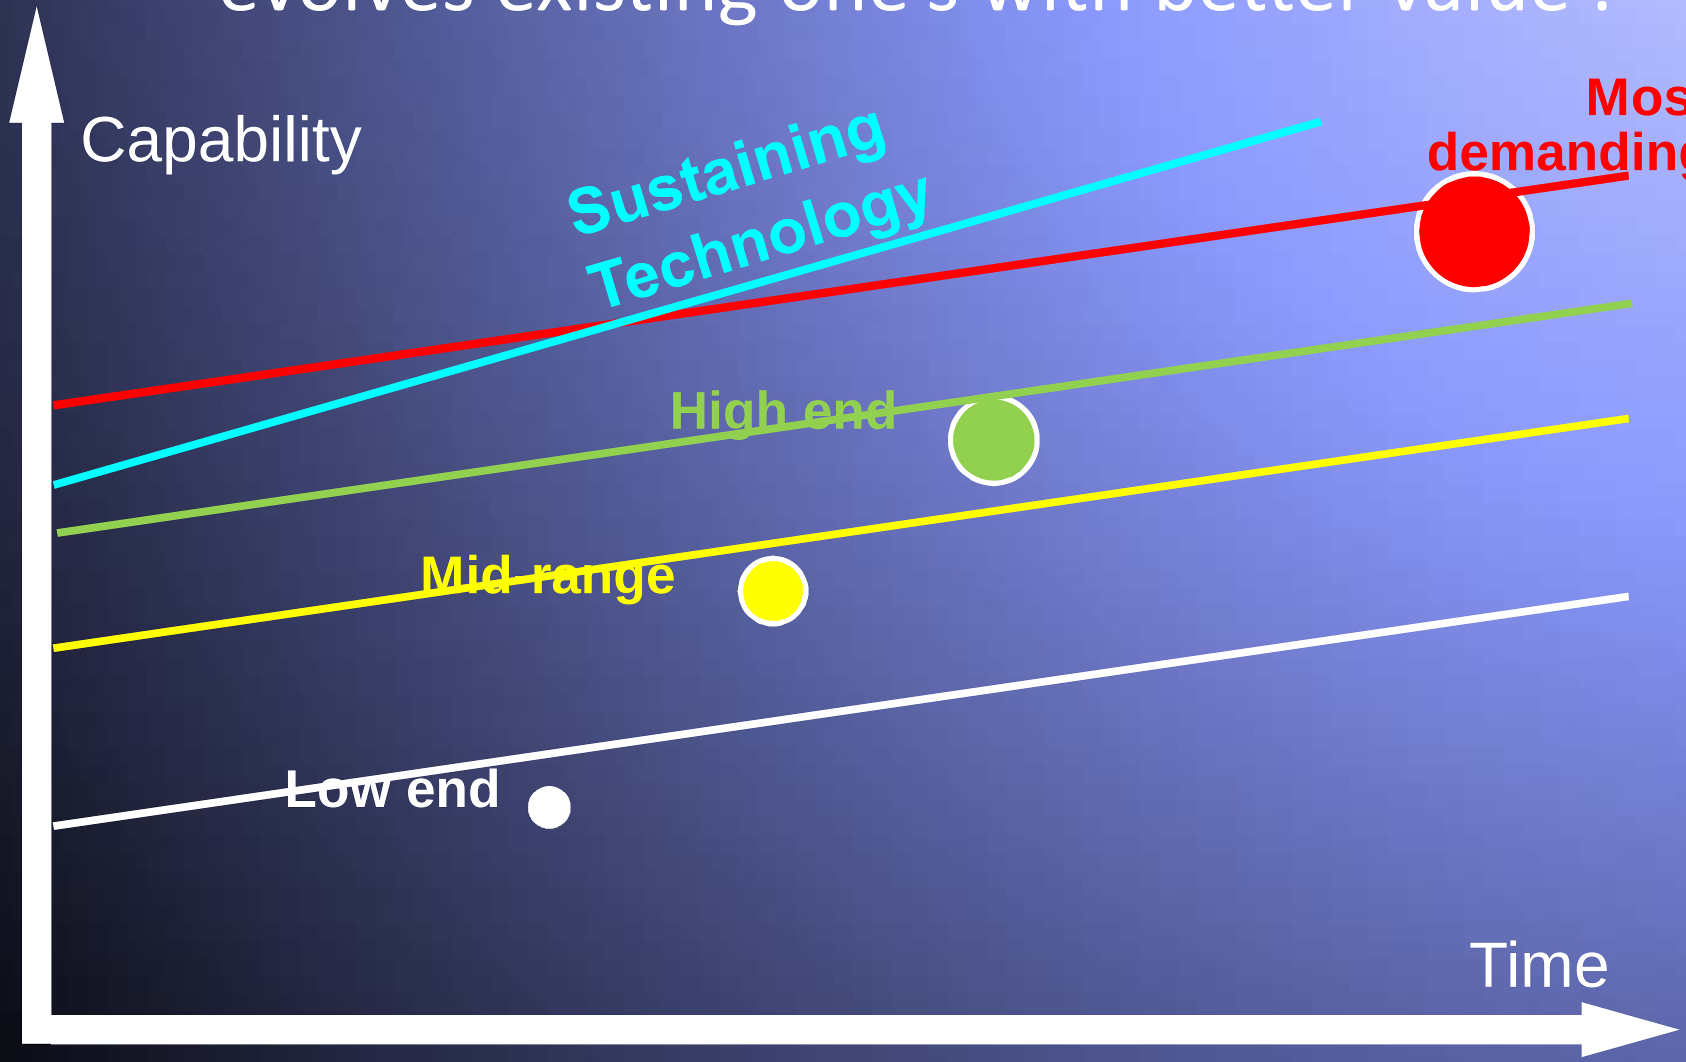

See: Clayton M. Christensen, The Innovator's Dilemma

Sustaining vs Disruptive Technologies

Disruptive – innovation that creates a new market or displaces existing technologies in a market.

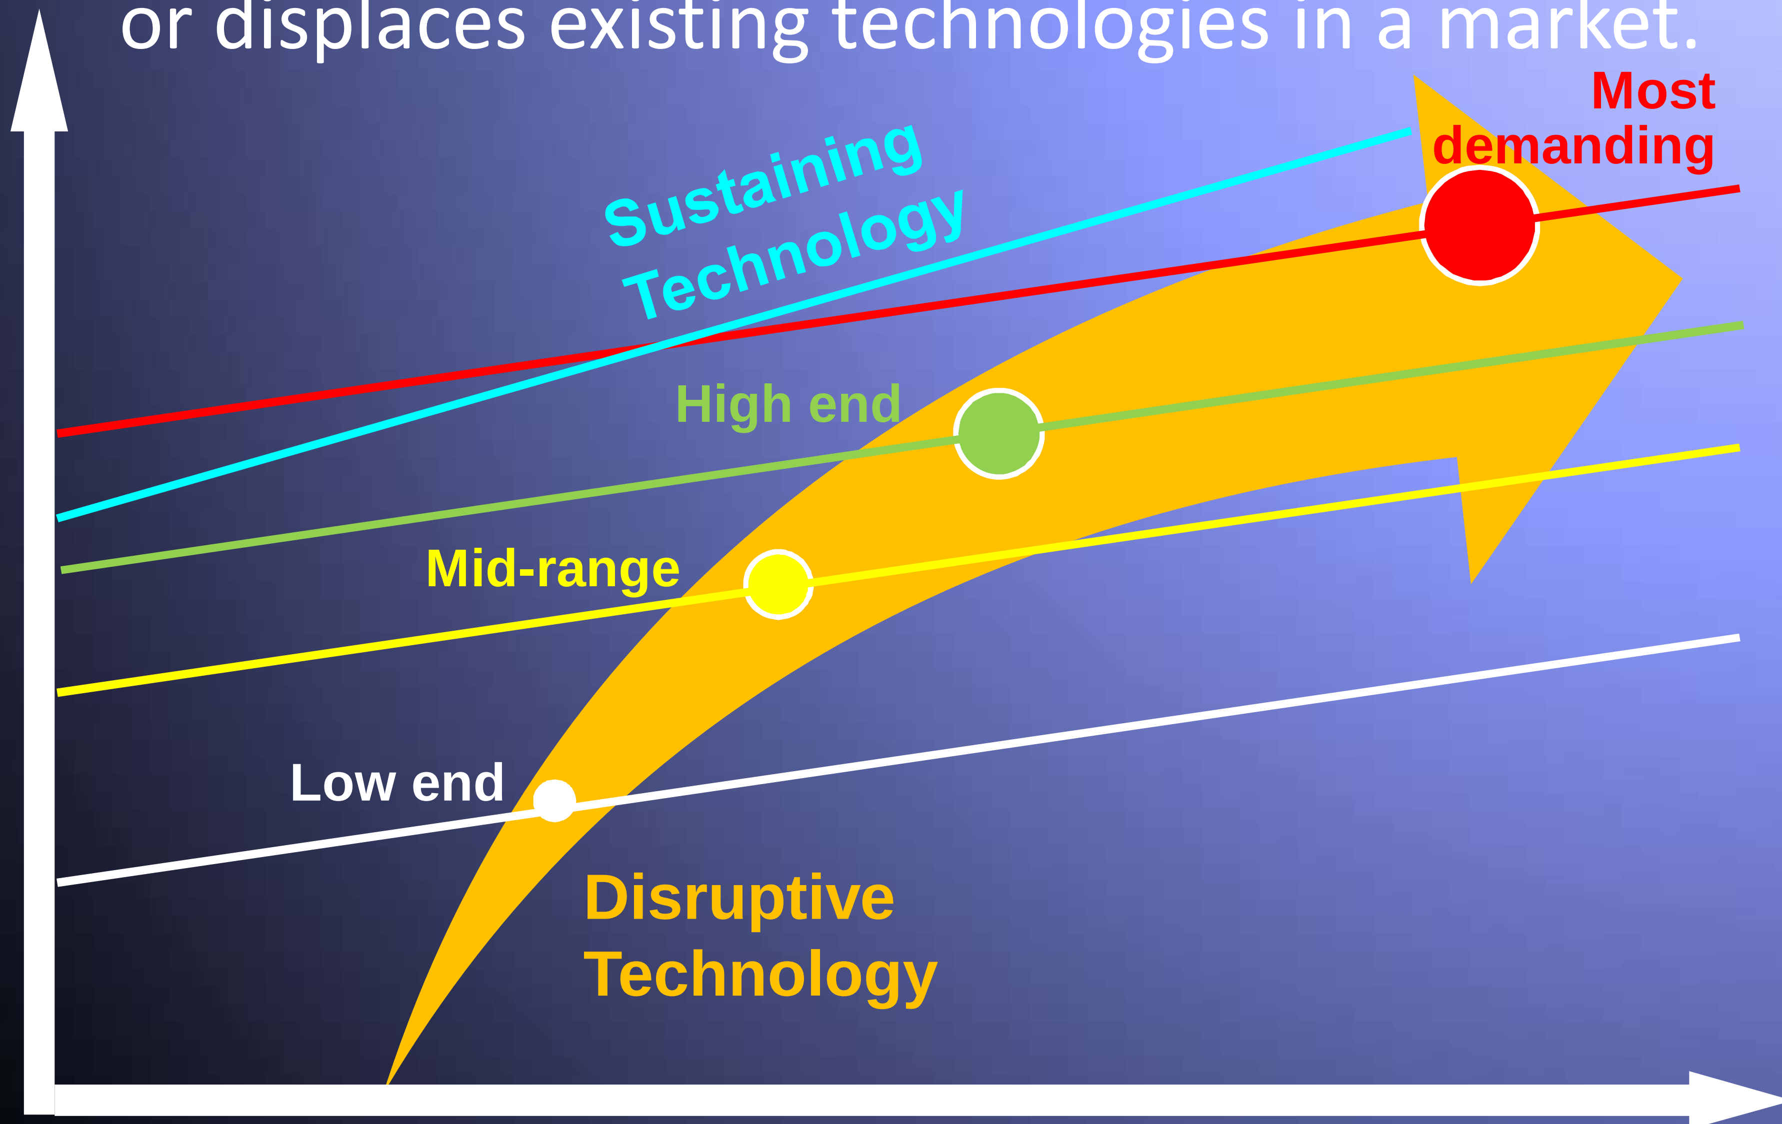

Current Disruptive Technologies Examples

PC & Microprocessor → Server markets

Handheld devices (e.g. iPAD) → Notebook PCs

Overlays & OpenFlow → Traditional switching

CEE & FCoE → Fibre Channel

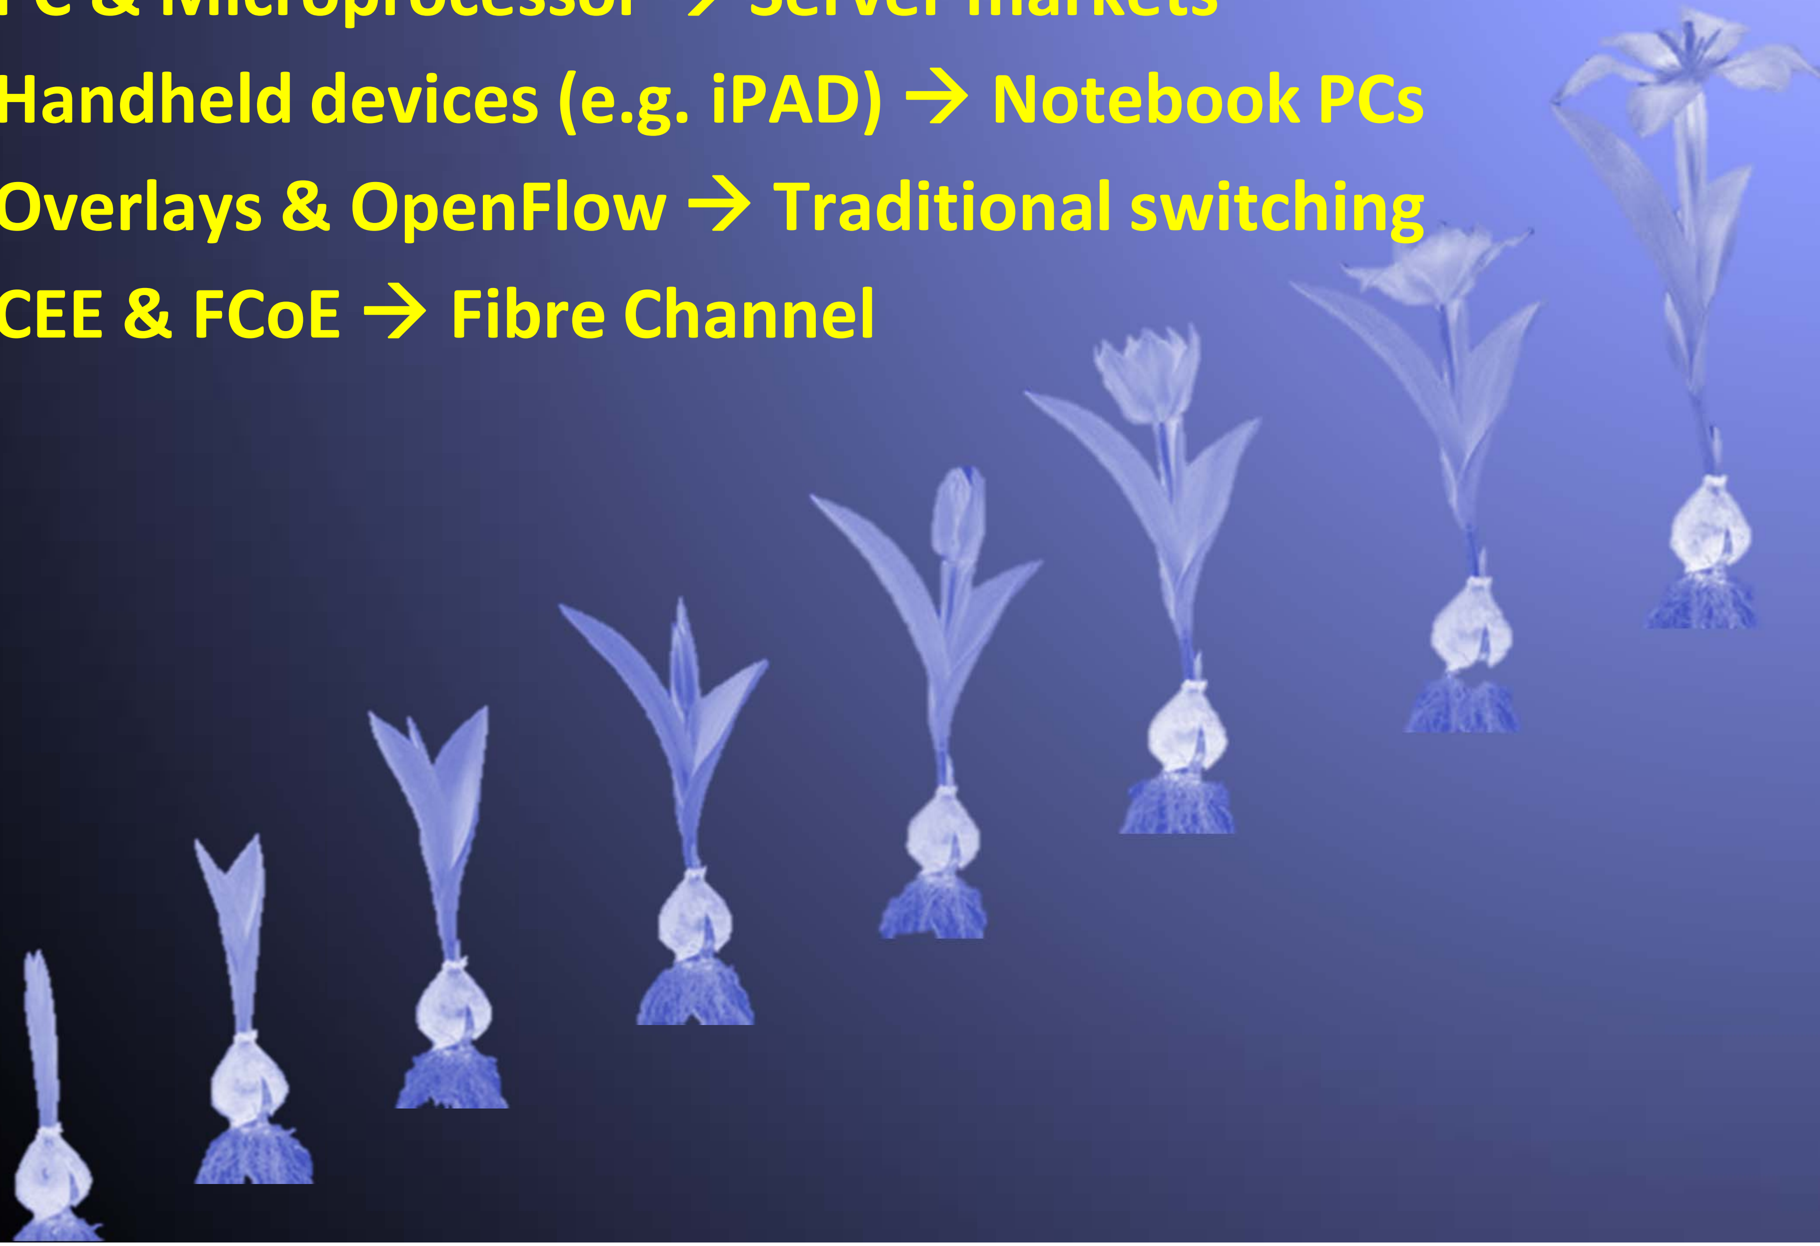

SDN (Overlay & OpenFlow) Client Examples

IBM

IBM SDN OpenFlow clients

For Global Network Services network between data centers
SDN Value: "Better network visibility & control"

SELERITY — Provider of ultra-low latency real-time financial information
SDN Value: "Policy driven content distribution & automated network configuration"

TERVELA *Data in Motion* — Provider of distributed data fabric for global trading & risk analysis
SDN Value: "Predictable network performance & rapid convergence"

IBM SDN Deployments in Enterprise Client Test & Development

2 service providers

SDN Value:
- "Programmable hosting network"
- "Global visibility and control"

Clients Evaluating World-wide Enterprise & Service Provider Data Centers

IBM Systems Networking SDN products and technology investments

Controller Platforms

GA 2/2012
DVS 5000V Controller

GA 10/2012
IBM Programmable Network Ctrl

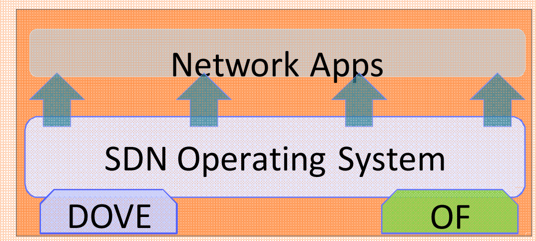

Open Source Based SDN Platform Investment

Network Virtualization

GA 2/2012
standards-compliant layer-2 virtual switch

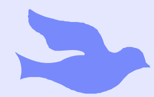

DOVE Technology Investment

OpenFlow Physical Switches

GA 11/2011

OpenFlow 1.0 10GB switch

OpenFlow Technology Investment

Optimized OpenFlow Network

IBM

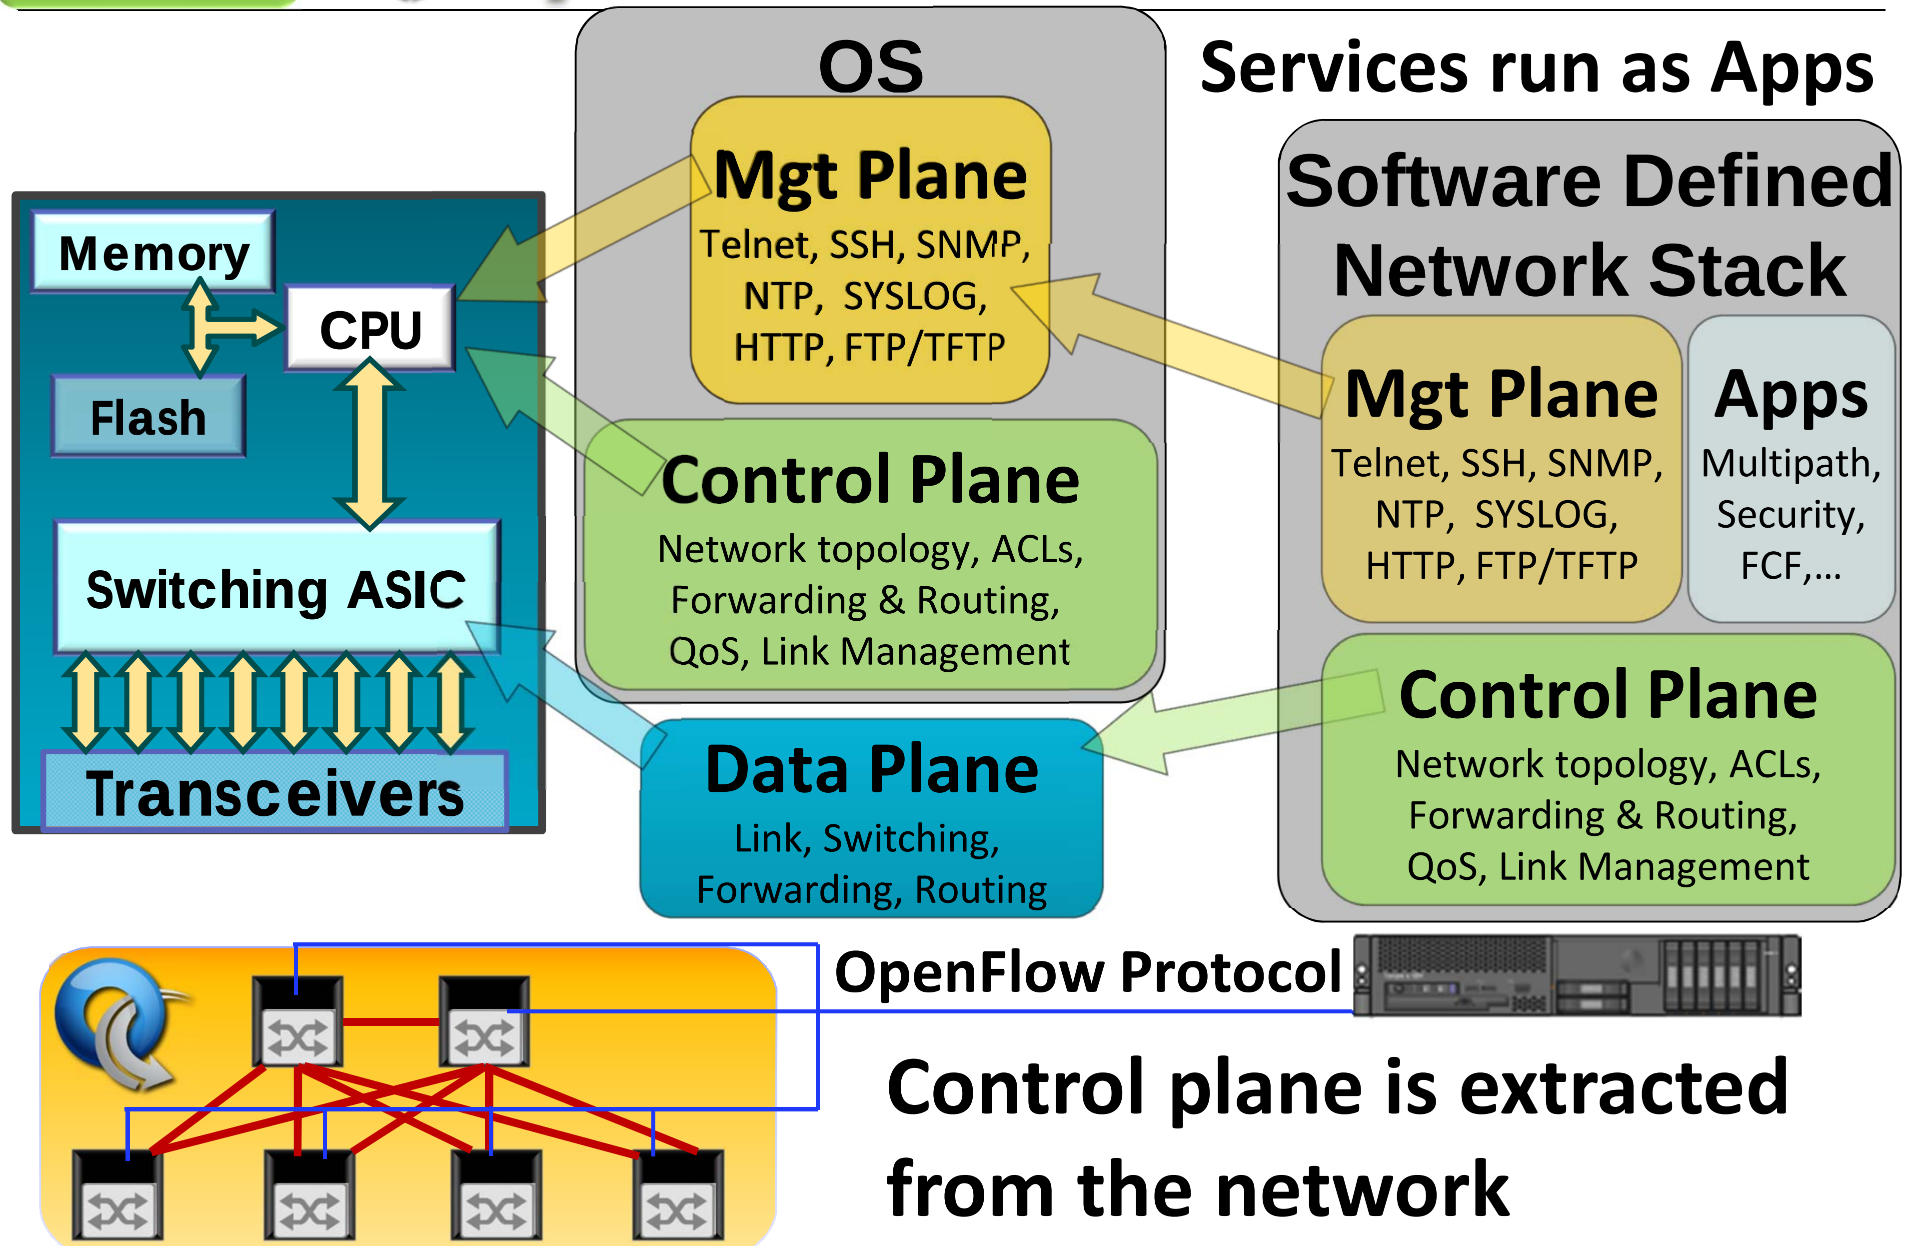

Control plane is extracted from the network

Optimized

IBM RackSwitch G8264 OF Switch

IBM

1st OpenFlow single chip switch to pass the 1 Terabit per second barrier!

Specifications	
Forwarding	• Delay less than 1us 1.28Tbps; 960Mpps
Number of ports	• 48 x 1 Gb/10 Gb SFP+ ports; • 4 x 40 Gb QSFP+ ports • Up to 64 x 1 Gb/10 Gb SFP+ ports with optional breakout cables
Model	• Airflow-type rear to front • Airflow-type front to rear
Dimensions	• 17.3” wide; 19.0” deep; 1U high
Protocol version	• OpenFlow 1.0.0
Number of instances	• 1
Protocols	• No legacy protocols running in OpenFlow switch mode
Management	• Telnet, SSH, SNMP, sFlow
Redundancy	• Power/fan

4x 40 GE uplinks or 4x*10G with QSFP to SFP+ cable

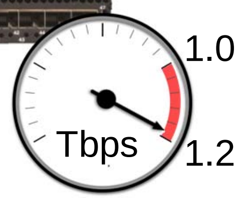

- OpenFlow-based flow handling in hardware at line rate (1.28Tbps)
- Support Layer 2 (MAC) forwarding table manipulated thru OF:
  – Layer 2 (destination MAC & destination MAC_+_VLAN) table: Max 128K flow entries
  – Layer 3 or anything other than layer-2 MAC/VLAN: Max 750 flow entries
- IBM and partner OF Controllers (PNC, BigSwitch Floodlight,…)

Optimized IBM Programmable Network Controller

- High performance OpenFlow based controller, provides:
  - Highly reliable end-to-end fabric that works over any OpenFlow 1.0 compliant switches
    - Automatically discovers OpenFlow network topology
    - Intelligent and dynamic multipath routing based on business policy
  - Virtualizes an OpenFlow network using Virtual Tenant Network (VTN) Application, where a VTN provides a:
    - customized layer-2 or layer-3 virtual network isolated from other virtual networks
    - Secure slice of the underlying physical network, with policy based networking
  - Automated, global end-to-end view & control of the network
    - Automated network topology discovery
    - Point and Click virtual network design
    - Network segments can be configured centrally
    - One touch point (IBM PNC) versus hundreds of touch points (Network Elements)
    - Policies are enforced throughout the network vs individual switch configuration
    - APIs to create, edit, and delete VTNs, as well as to add and remove policies

Optimized OpenFlow Optimization

IBM

Today's Network	SDN with OpenFlow

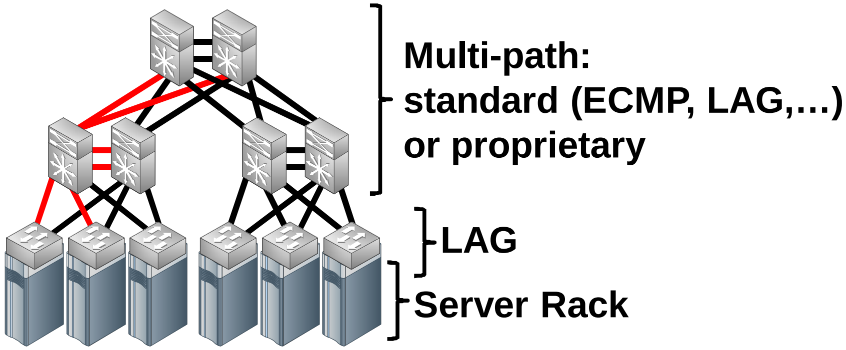

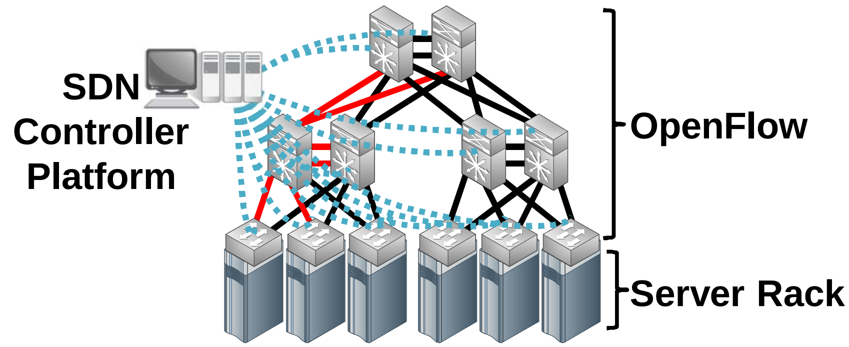

Multi-path fabric with LAG to server

\+ Scale wells

\+ Standard based multi-pathing protocols

\- Requires proprietary fabric to reduce managed switches

\- Long convergence times for Distance / Path vector algorithms (e.g. RIP / BGP), which iteratively send routing protocol packets to calculate routes.

\- Long time to add new functions to standard (e.g. OSPF, TRILL)

OpenFlow based multi-pathing

\+ Scales well with arbitrary topologies

\+ Optimized multi-pathing with short convergence times*

\+ New functions (e.g. disjoint multi-pathing can be added to SDN controller, without long-standard lead time)

\- Gated by SDN market adoption S-Curve

\- Established networking vendors will create FUD against this model

* Controller discovers switches & creates topology, when a device fails, neighbors report the event and controller loads new routing tables around the failure.

Automated

Distributed Overlay Virtual Ethernet (DOVE) Network Technologies

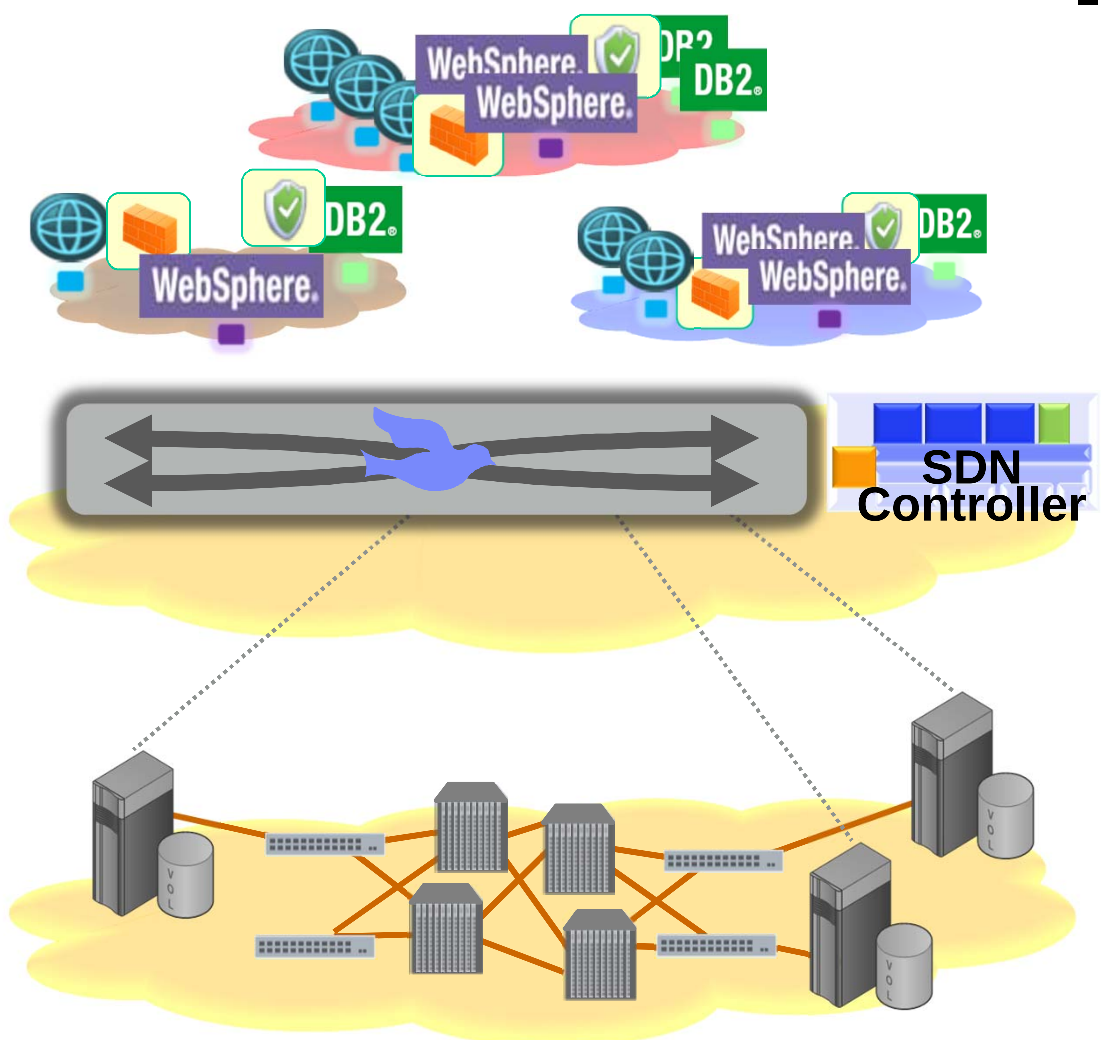

- Automated connectivity of services (e.g. Firewall, IPS) used in virtual system patterns, with multi-tenant network capabiliy
- Layer-3 DOVE switch decouples virtual networks from physical network
- Simple "configure once" physical network (vs configured per VM)

Automated Laying Foundation for DOVE Standards

IBM contributions to
Distributed Overlay Virtual Ethernet (DOVE) networking

- In early 2009, IBM Haifa research began work on overlay networking.
  - https://www.research.ibm.com/haifa/projects/systech/reservoir/research.shtml
  - Virtual Application Networks (VANs) are a host side solution that allows hosts to construct a fully virtualized network service overlay on top of a standard IP physical network.
  - Filed patents on VAN related material
- At the Ethernet Summit Conference in February 2011, IBM introduced DOVE concept:
  - Page 16 of www.ethernetsummit.com/English/Collaterals/Proceedings/2011/20110223_PlenaryEthernet_Recio.pdf
- At the September 2011 Data Center Converged And Virtual Ethernet Switching workshop IBM published a paper on DOVE networking:
  - http://www.itc23.com/workshops/dc-caves-program/
- At the November 2011 IETF, IBM:
  - Proposed creation of an IETF workgroup to work on DOVE networking: http://tools.ietf.org/agenda/82/slides/l2vpn-9.pdf
  - Provided a Problem Statement to begin the IETF standardization process: http://tools.ietf.org/html/draft-narten-nvo3-overlay-problem-statement-01
- At the March 2012 IETF, IBM:
  - Presented an updated problem statement (http://www.ietf.org/proceedings/83/slides/slides-83-nvo3-5.pdf) & (overlay) network virtualization working group charter (http://www.ietf.org/proceedings/83/slides/slides-83-nvo3-6.pdf)
- At the July 2012 IETF approved IBM's problem statement as an internet draft
  - http://tools.ietf.org/html/draft-ietf-nvo3-overlay-problem-statement-00

DOVE Network Technology Components

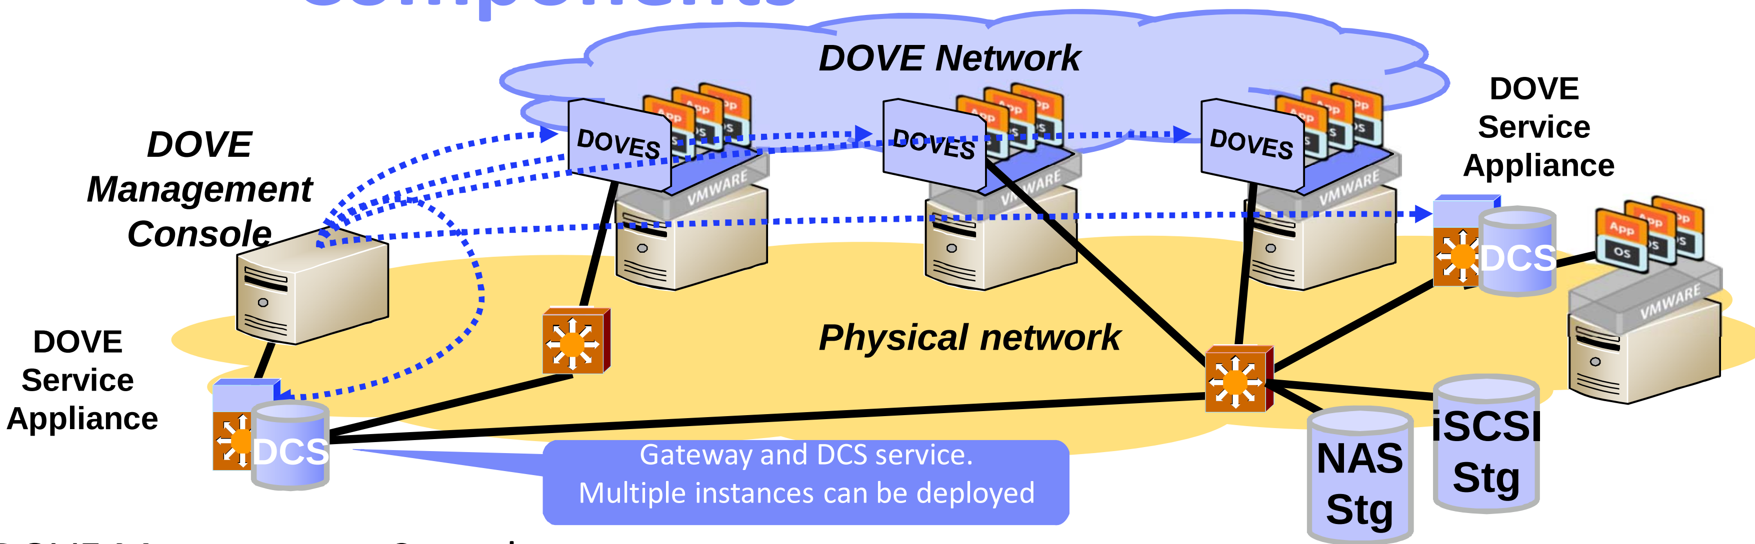

- DOVE Management Console
  - Provides GUI and APIs for management of DOVE network, groups and policies
- DOVE Switches (DOVES)
  - Provides layer-3 & layer-2 over UDP overlay (header format same as VXLAN)
  - Performs data and some control plane functions
- DOVE Service Appliance
  - Distributed Connectivity Service (DCS):
    - discovers & disseminates VM location (physical server)
    - maintains policy (e.g. allow, deny, insert service appliance) and works with DOVE switches to apply policy
  - Gateway Service: Connectivity to non-DOVE networks
- Virtual and Physical (network and storage) Appliances provide services for DOVE Network

Overlay Operation

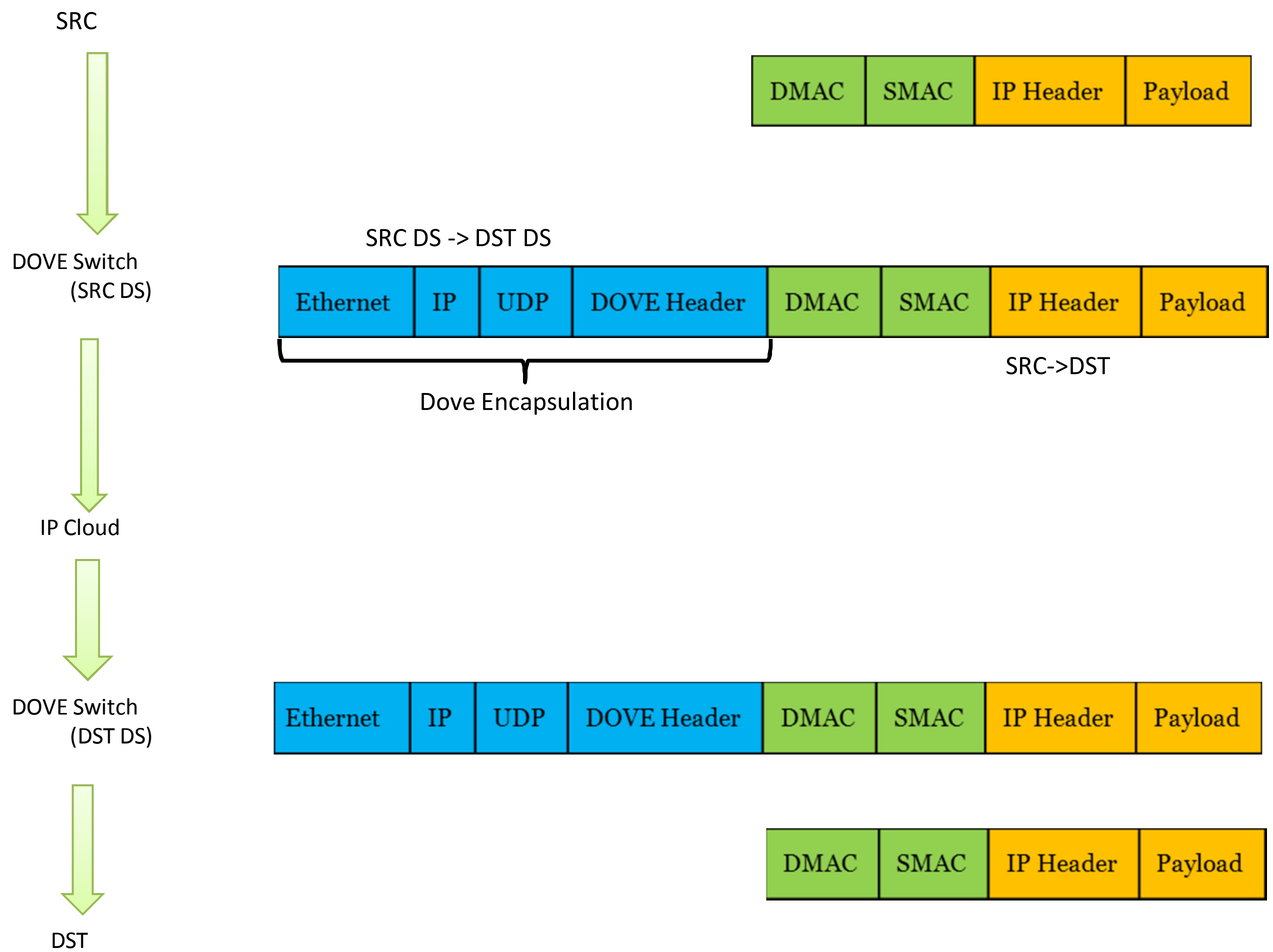

DOVE Header (VXLAN/OTV):

R	**R**	**R**	**R**	**I**	**R**	**R**	**R**	**Reserved (24-bits)**
DOVE Domain Group ID (24-bits)								**Reserved (8-bits)**

Automated DOVE Networks provide an SDN Connectivity Service that enables Virtualized Systems

Clients build DOVE Networks, which interconnect workload groups through a set of middle-boxes. A DOVE Network provides the network connectivity required to build virtualized multi-tier systems, that can be deployed multiple times with minimal manual configuration.

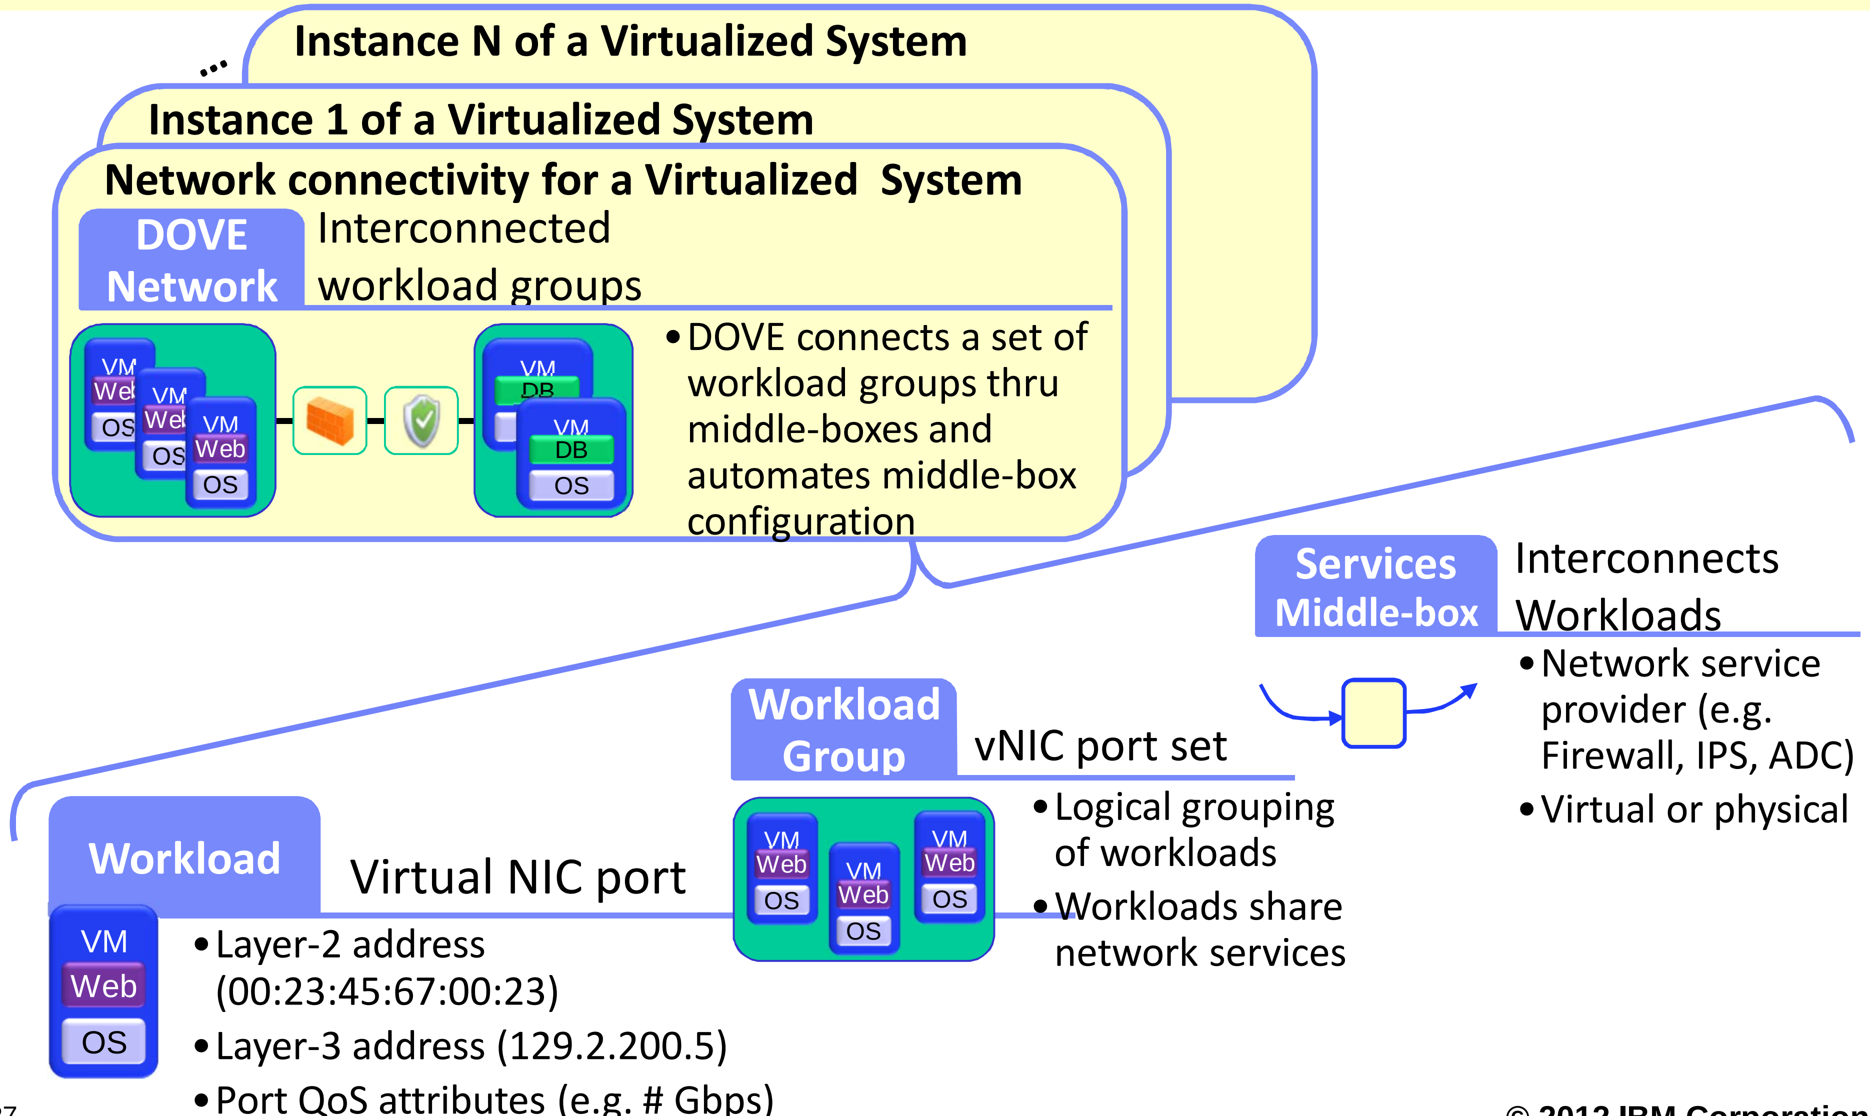

Automated

SDN Connectivity Services + Virtual App Patterns = Rapid Virtual App Pattern Deployment

IBM

When combined with Virtual App Patterns:
the network connectivity service interconnects and configure the middle-boxes; and the workload deployment manager (e.g. PSM) configures the pattern's workloads.

...

Instance N of a Virtual App Pattern Instance

Instance 1 of a Virtual App Pattern Instance

Network connectivity for a Virtual App Pattern

DOVE Network Interconnected **& configured** workload groups

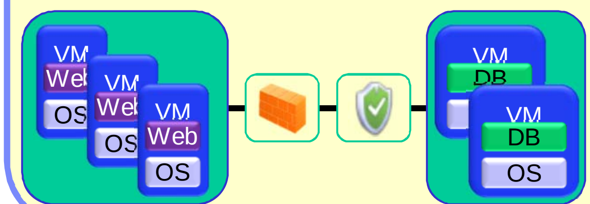

- DOVE connects & configures network
- Workload deployment manager configures workloads

Services Middle-box Interconnects Workloads

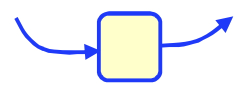

- Network service provider (e.g. Firewall, IPS, ADC)
- Virtual or physical

Workload Group vNIC port set

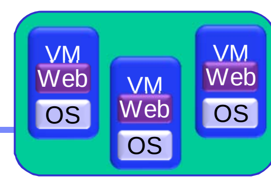

- Logical grouping of workloads
- Workloads share network services

Workload Virtual NIC port

- Layer-2 address (00:23:45:67:00:23)
- Layer-3 address (129.2.200.5)
- Port QoS attributes (e.g. # Gbps)

Automated Network Service Automation

IBM

Today's Network

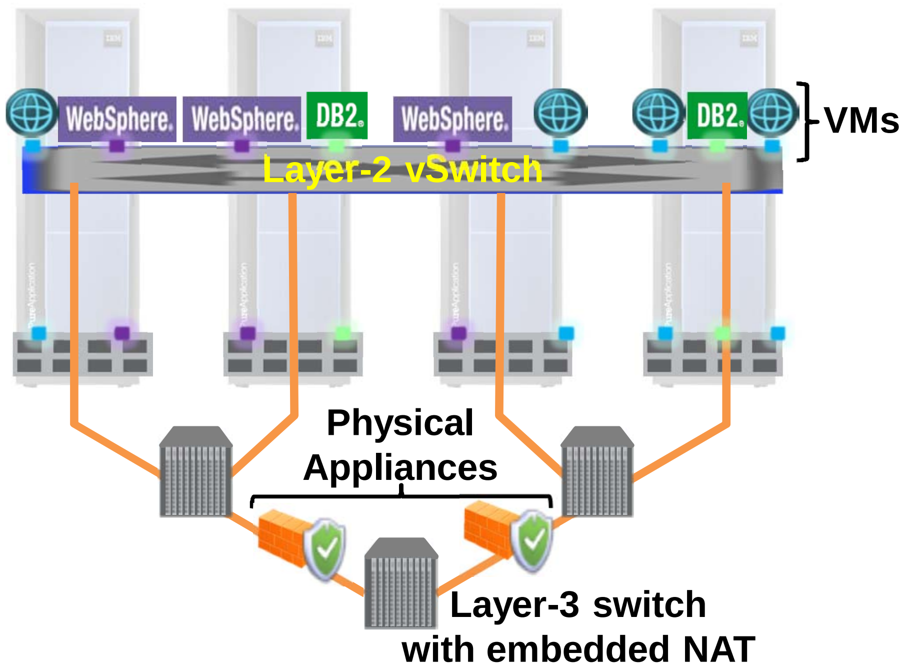

Layer-2 automated

+ Access switches are network aware,
  using IEEE 802.1Qbg standard.

Complex, limited & costly appliances

- Configuration complexity (requires per workload instance configuration of physical: NAT controllers, security appliances, workload balancers, …)
- Only within the boundary of security appliance are protected (migration across boundaries requires physical configuration changes)
- Expensive appliances; no open eco-system

Software Defined Networking

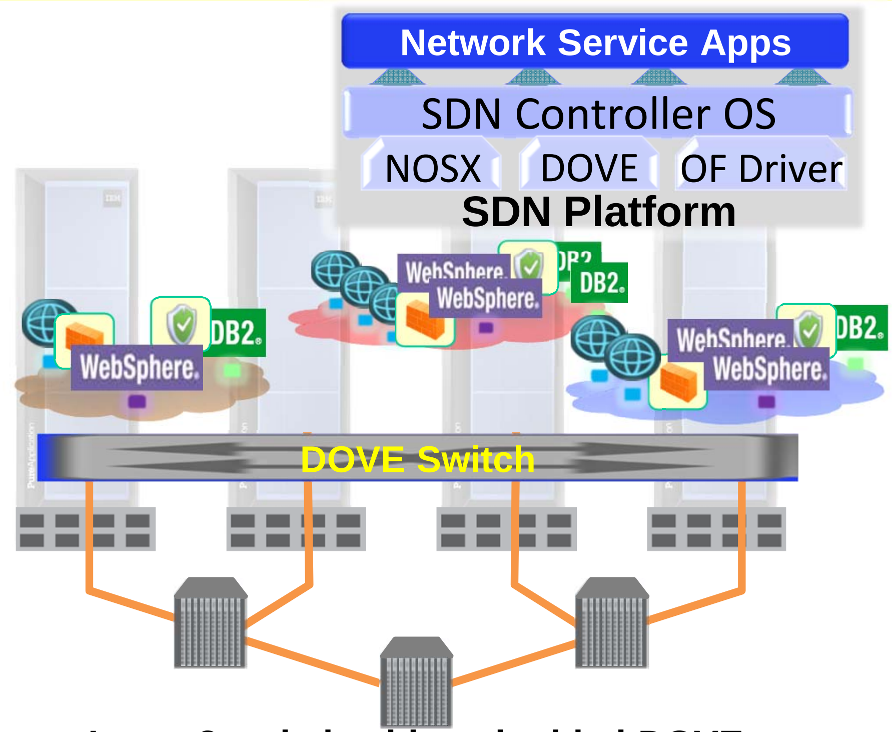

Layer-3 switch with embedded DOVE gateway

Full Virtual Network Automation

+ Simple configuration based on DOVE network pattern
+ VMs are free to move around DC, VM's network service attributes are not tied to physical location
+ Low cost appliances running on open SDN eco-system
- Gated by SDN market adoption S-Curve

Automated Multi-tenant, Optimized Virtual Network

Today's Network

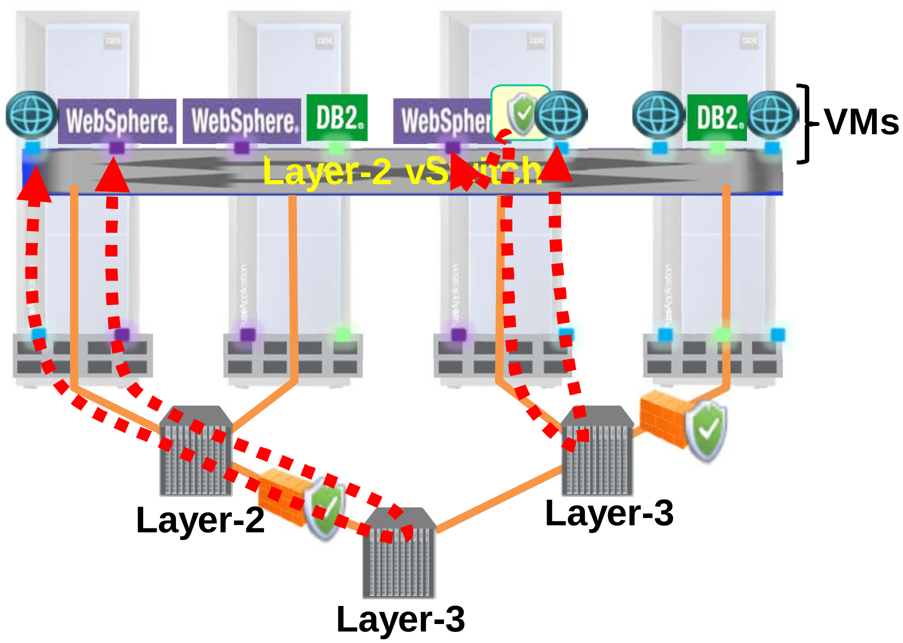

Sub-optimal traffic flow, limited scale

- Traffic is not optimized across groups (cross-subnet VM-VM traffic must go North-South), even for virtual network services
- Limited scaling (lacks multi-tenancy)

Software Defined Networking

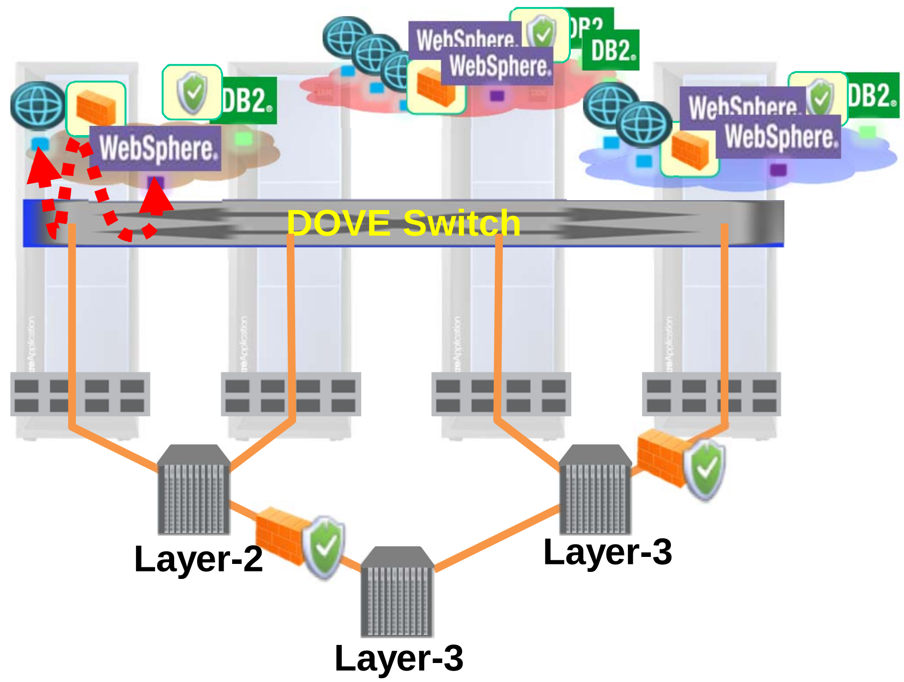

Optimized traffic flow, multi-tenant scale

+ Optimizes traffic within **and across** groups (cross-subnet VM-VM traffic stays in server)
+ Multi-tenant scaling for: cloud service providers; clients that consolidate infrastructure after a merger; …

Integrated

Software Defined Networking Technology Value

IBM

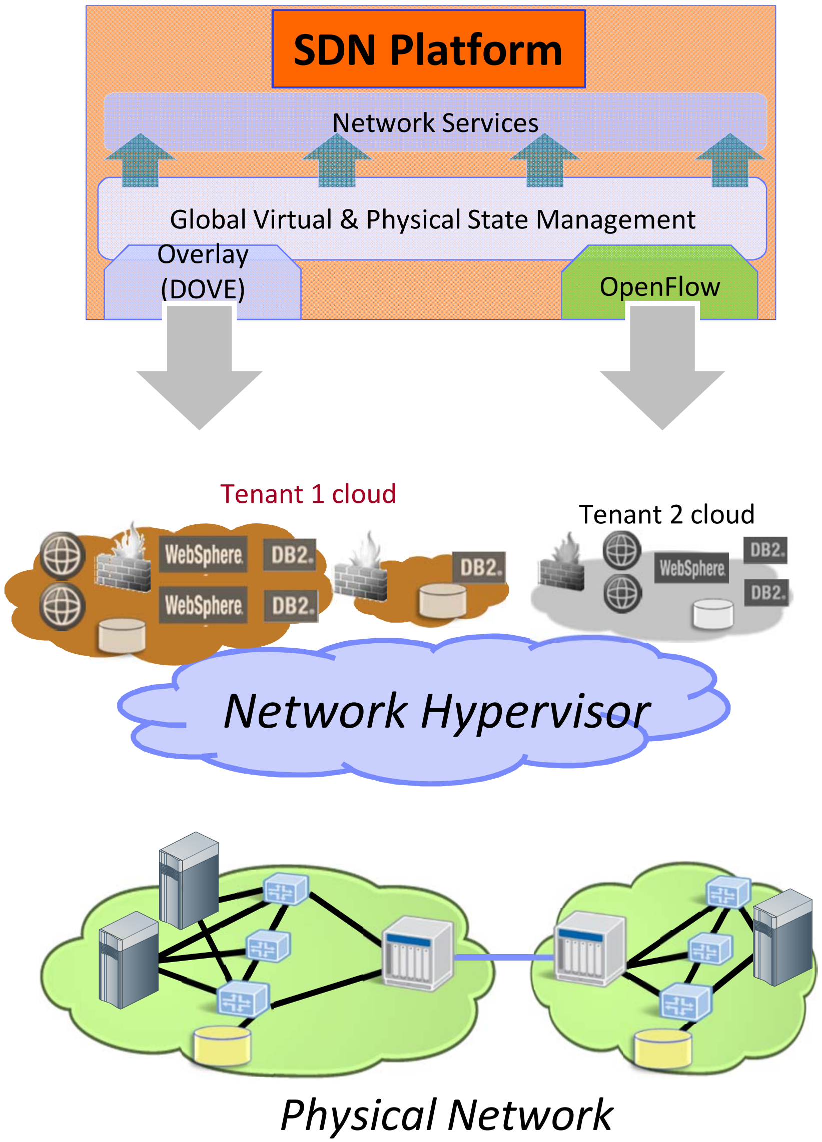

- **Network Services value:**
  - Eco-system for network Apps vs today's closed switch model
  - ***Connectivity Service for DOVE automates layer-3 and above***
  - ***Multi-pathing Service for OpenFlow lowers convergence time and time to new function***
- **DOVE Network value:**
  - ***Automated network resource (layer-2 and above) provisioning***
  - De-couples virtual network from physical network
  - Simple "configure once" network (physical network doesn't have to be configured per VM).
  - Cloud scale (e.g. multi-tenant)
- **OpenFlow value:**
  - ***Global physical network control & visibility***
  - De-couples control plane from data plane

Thank You


IBM
Renato J Recio
IBM Fellow & Systems
Networking CTO
11400 Burnett Road
Austin, TX 78758
512 973 2217
recio us ibm com


QUESTIONS