



SDN (Software Defined Network)

IBM System Networking 팀
김학열 차장 (mickykim@kr.ibm.com)



- IT Market Trend _Why SDN
- SDN/OpenFlow
- OpenFlow Research Time Frame
- OpenFlow의 시장성
- OpenFlow의 효율성
- IBM의 제공 가능 솔루션
- IBM의 SDN전략및 로드맵
- Appendix
- Q&A



IT Market Trend > Server Vendor + Network Vendor

IBM to Acquire BLADE Network Technologies
Technology to further optimize IBM servers for cloud computing, analytics and other new workloads

ARMONK, N.Y., 27 Sep 2011 - IBM (NYSE: IBM) today announced it has entered into a definitive agreement to acquire BLADE Network Technologies (BLADE), a privately held company based in Santa Clara, CA. BLADE specializes in software and devices that route data and transactions to and from servers. The acquisition is anticipated to close in the fourth quarter of 2011, subject to certain regulatory and other closing conditions and applicable regulatory reviews. Financial terms are not being disclosed.

BLADE is a leading and top drive provider of network servers to business and government customers in the top 500 list across 20 major industries including financial services, education, government, healthcare, defense and finance. IBM and BLADE have worked together since 2002, resulting in thousands of joint clients. In fact, over 50 percent of IBM System x Blade Centers currently attach to or use BLADE products (1).

The BLADE acquisition builds on the industry-leading capabilities and technologies IBM is applying to its systems, which are optimized to help clients manage a range of new, more demanding workloads. This year, IBM introduced a full line-up of new, workload-optimized systems that incorporate innovation at each level -- from microprocessors and firmware software to middleware and hardware. With BLADE, IBM can drive innovation at the systems networking level to enable clients to speed the delivery of key information from system to system -- for workloads such as analytics and cloud computing -- while also reducing data center costs.

Force10 Networks

Dell announces its intent to acquire Force10 Networks

Acquisition Overview
Dell intends to acquire Force10 Networks, a leading technology leader in high-performance data center networking. Force10's solutions provide enterprise performance and resilience while reducing overall total cost of ownership, simplifying network deployment and supporting open standards-based systems and management solutions. The revenue acquisition is Dell's latest investment as it broadens its enterprise portfolio to offer customers a complete range of data center products and solutions.

Customer Benefits
The acquisition will enable Dell to expand its portfolio and be able to offer customers a complete end-to-end infrastructure solution across storage, servers and networking. Dell remains committed to providing open, scalable and affordable solutions together with its Allston partners; current networking technology and solution partners will remain an important component of our global enterprise data center strategy.

Have Dell contact you by email or phone

Notes:
"Today's data center networks are too complex and require too much manual intervention. We've worked in the past is no longer viable in the virtual era. Dell's approach of offering customers open, capable and affordable solutions aligns with Force10 Networks' approach to offering customers new levels of flexibility, performance, scale and automation which is fundamental to changing the economics of data center networking."
Brad Anderson, Senior Vice President.

HP Completes Acquisition of 3Com Corporation, Accelerates Converged Infrastructure Strategy
PALO ALTO, Calif., April 12, 2011

HP today announced that it has completed its acquisition of 3Com Corporation in a price of \$7.90 per share. The acquisition adds 3Com's converged infrastructure capabilities to HP's existing HP BladeCenter portfolio, which is built on the open standards-based networking, management and security capabilities of HP's Converged Infrastructure. With this integration, customers will be able to simplify their networks, deploy on edge to core networks, fabric for the enterprise and improve IT service delivery capabilities.

The acquisition of 3Com expands HP's Ethernet switching offerings, adds routing solutions and significantly strengthens the company's position in China. 3Com also brings to HP network security capabilities through its TippingPoint portfolio. Together, HP is able to deliver one of the broadest network technology capabilities in the market to meet customer needs well into the future.

Further details on product integration will be announced at a later date.

HP creates new possibilities for technology to have a meaningful impact on people, businesses, governments and society. The world's largest technology company, HP brings together a portfolio that spans printing, personal computing, software, services and IT infrastructure to solve customer problems. More information about HP (NYSE: HPO) is available at <http://www.hp.com>.

Cisco Finally Releases UCS Market Share Numbers

May 24, 2011 - IDC came out with their 1Q 2011 worldwide server market revenue report today showing that Cisco has finally entered the market standings with a **3rd place standing at 9.4% factory revenue share**. IDC's findings also showed that both HP and IBM decreased their blade server market share from Q4 2010.

According to IDC, worldwide server sales (all servers, not just blade servers) for 1Q 2011 increased 12.1% year over year to \$11.5 billion in factory revenues. IDC also reported the blade server market accelerated and continued its strong growth in 1Q with revenue increasing 23.8% year over year with shipments increasing to 5.4% compared to 1Q 2010. Overall, blade servers represented 15.2% of the quarterly worldwide server revenues. Interestingly enough, 91% of all blade revenue is driven by x86 systems, a segment in which Dell now represent 20.2% of the market (factory revenue share).

Here's a summary of the 1Q 2011 worldwide server market share (factory revenue share):

- #1 market share: HP at 50.0% in Q1 2011
- #2 market share: IBM at 20.2% in Q1 2011
- #3 market share: Cisco at 9.4% in Q1 2011
- #4 market share: Dell at 8.2% in Q1 2011


IDC 1Q 2011 Worldwide Blade Server Market Share

Vendor	Market Share (%)
HP	50.00%
IBM	20.20%
Cisco	9.40%
Dell	8.20%
All Other Vendors	12.20%

IT Market Trend > Traffic의 변화


Gartner Group “Debunk the Myth of the single vendor network” 11/17/2010

Introducing a 2nd networking vendor will reduce TCO for most organization by at least 15-25% over 5 years
 Most organization that introduced a 2nd vendor report a lasting decrease in network complexity, compared with an all-Cisco network.
 Equipment cost premium that Cisco generally charges does not tend to off-set potential operating cost savings.



Data Center의 10GbE 수용 증가

Data Type의 변화 및 Data 증가로 인한 요구 Bandwidth 증가
 대형 서버 벤더의 10G Ethernet 기본의 Server 발표로 인해 수요량 증가.
 40GbE Ethernet 장비 발표 및 100GbE Ethernet Switch의 Road Map 발표

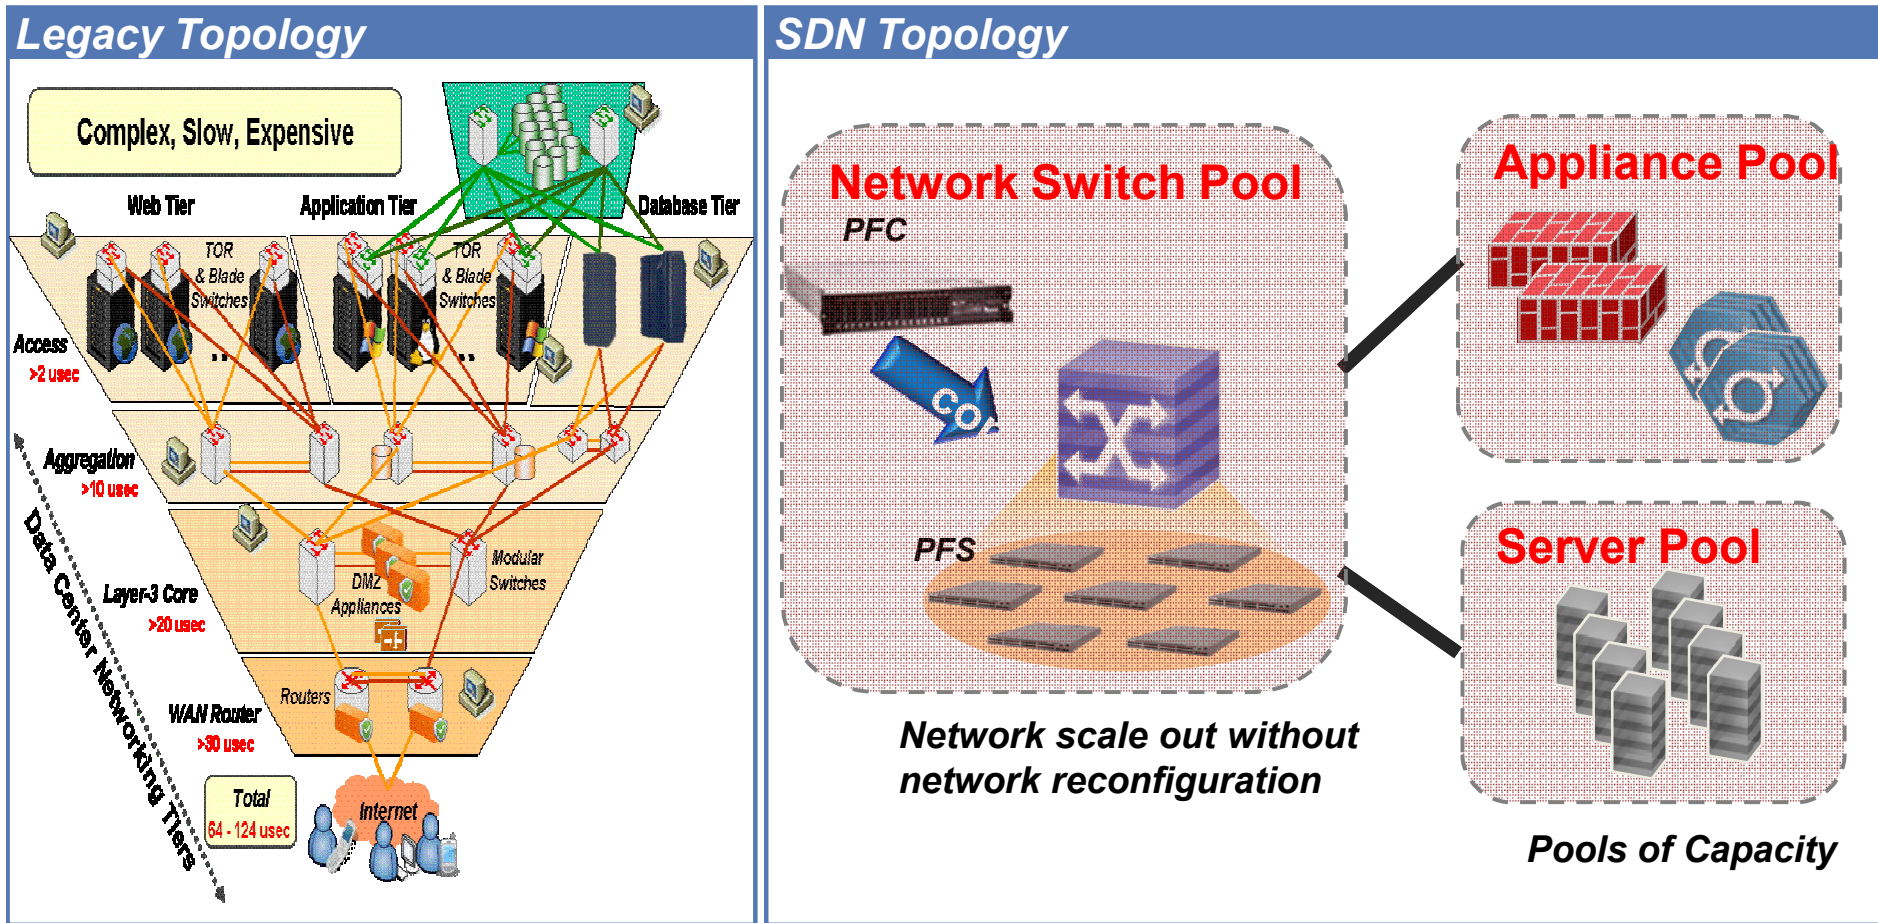


Source Dell'Oro Group 2012

SDN/OpenFlow > Why SDN?

• Why SDN (Software Defined Network)?

- 네트워크 영역의 복잡성과 Vendor의 구속에서 벗어나기 위한 사용자들의 욕구
- SDN을 통해 네트워크 영역 또한 하나의 Device Pool로 만들어 간편함 및 확장성의 원활함 제공
- 네트워크 구성개선을 통한 TCO 절감 및 운영효율제공



SDN/OpenFlow > Who lead SDN?

- **Who lead SDN (Software Defined Network)?**

- Mail Border Member가 고객 으로 부터 시작
- 각 Vendor들은 고객의 요건에 따라가는 그림
- 기존 Network Vendor들과 None Network Vendor들의 접근법의 차이

ONF board members are the largest network operators (not networking vendors)




ONF Board



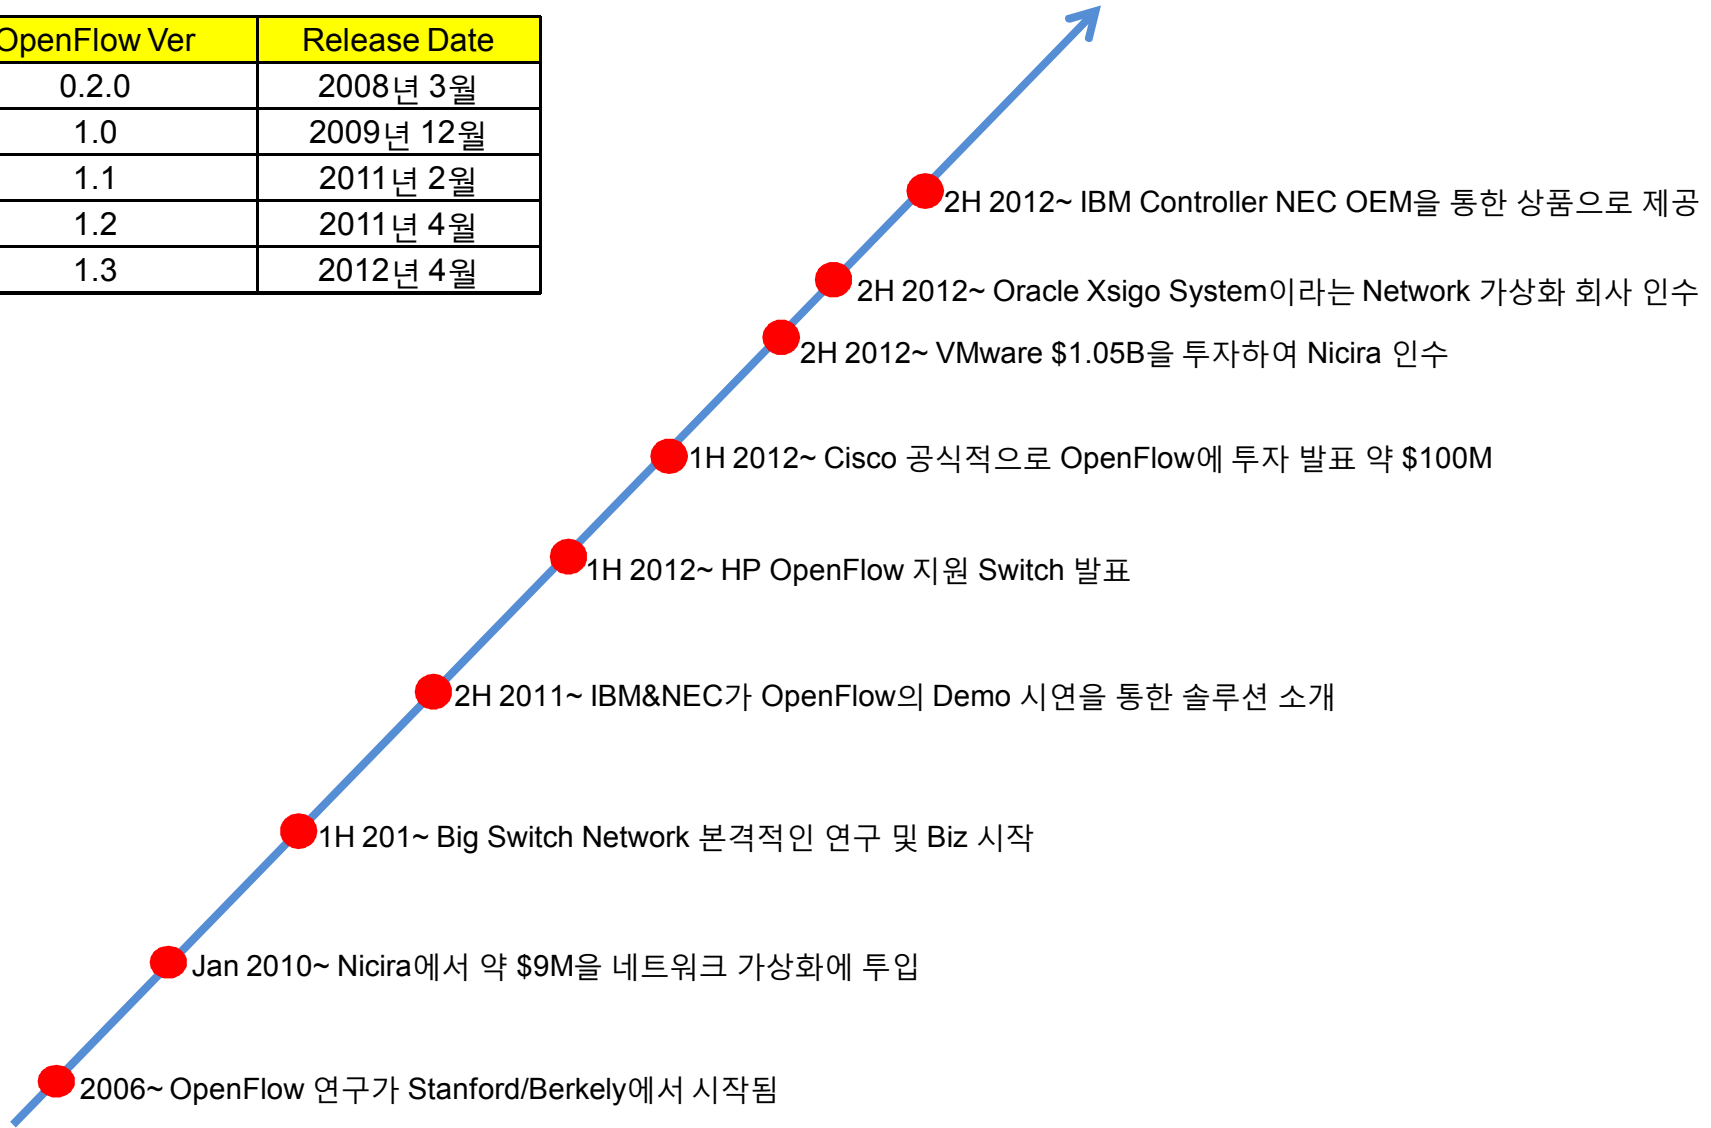
- **SDN (Software Defined Network)**

- Control Plane과 Data Plane이 완전히 물리적으로 분리된 형태의 Architecture
- SDN은 네트워크 전체 영역에 초점을 두고 있다
- SDN의 일부분인 OpenFlow에 많은 관심과 연구가 진행중



OpenFlow Research Time Frame

OpenFlow Ver	Release Date
0.2.0	2008년 3월
1.0	2009년 12월
1.1	2011년 2월
1.2	2011년 4월
1.3	2012년 4월




Worldwide OpenFlow Revenue, 2012-2016 (\$M)

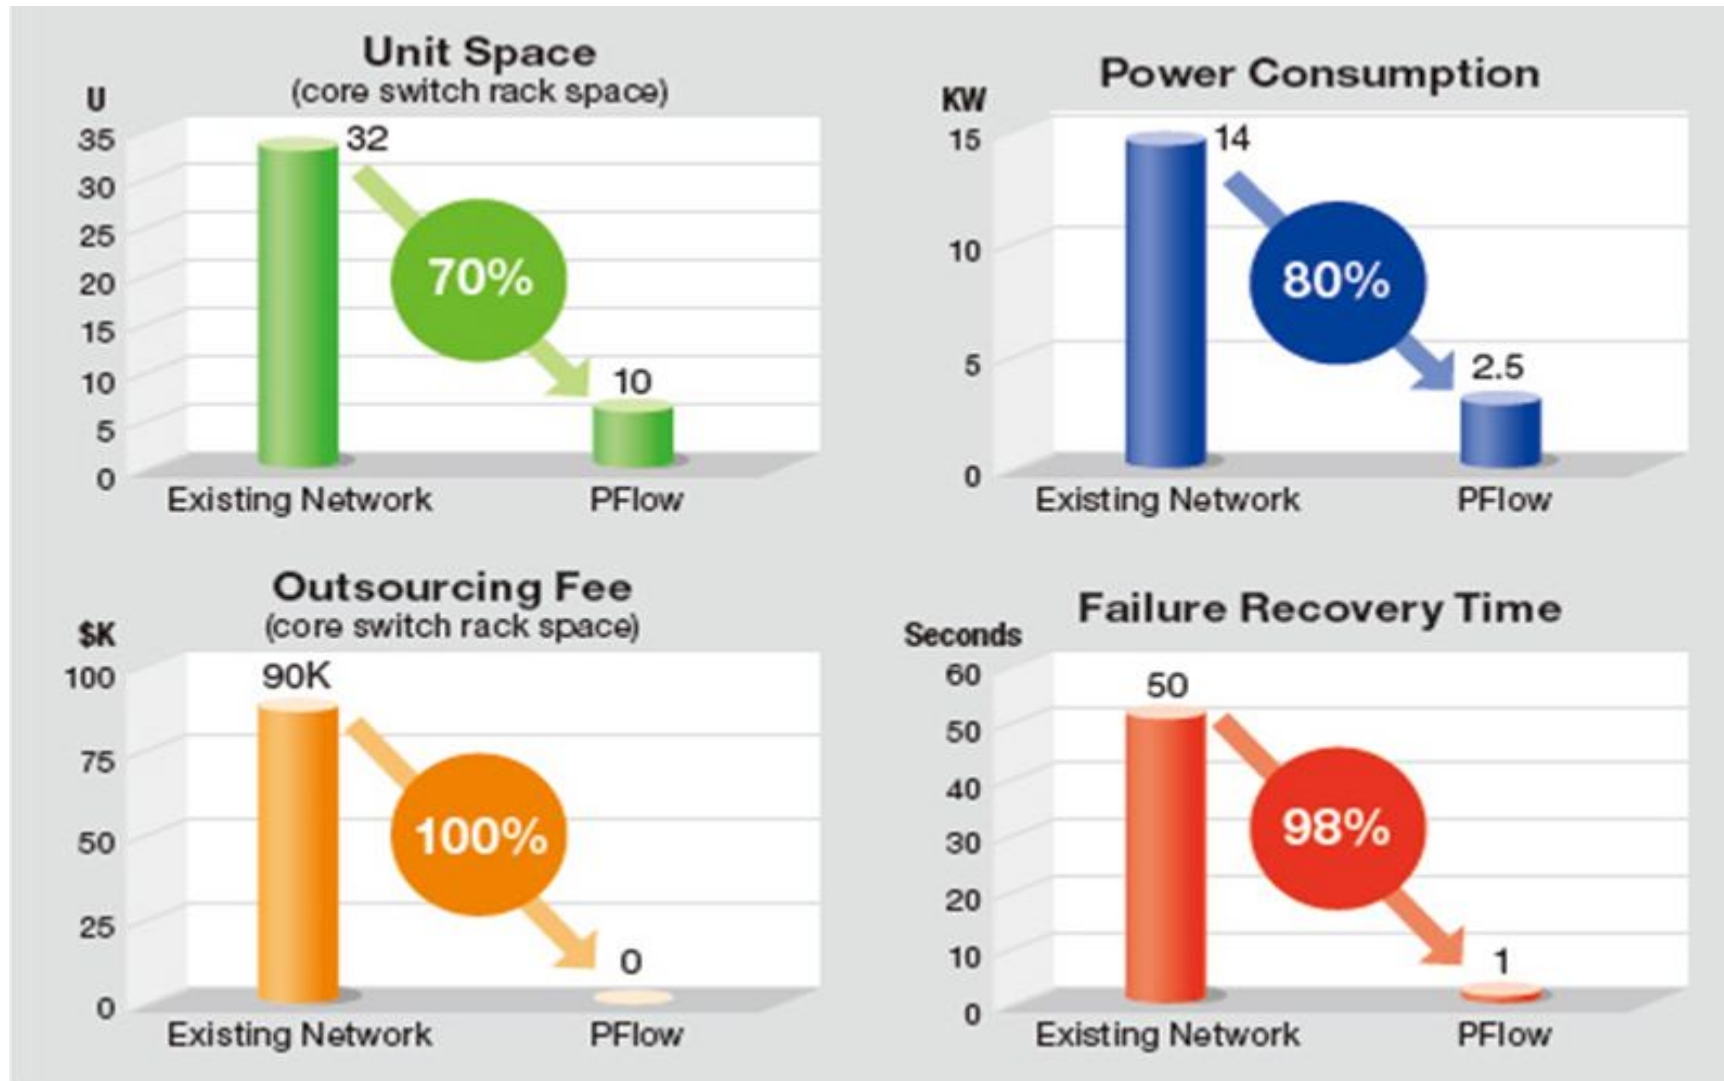
	2012	2013	2014	2015	2016	2012-2016 CAGR (%)
Worldwide	53.8	168.1	380.6	897.0	1,954.5	145.5
Growth (%)	NA	212.5	126.4	135.7	117.9	

Source: IDC, 2012

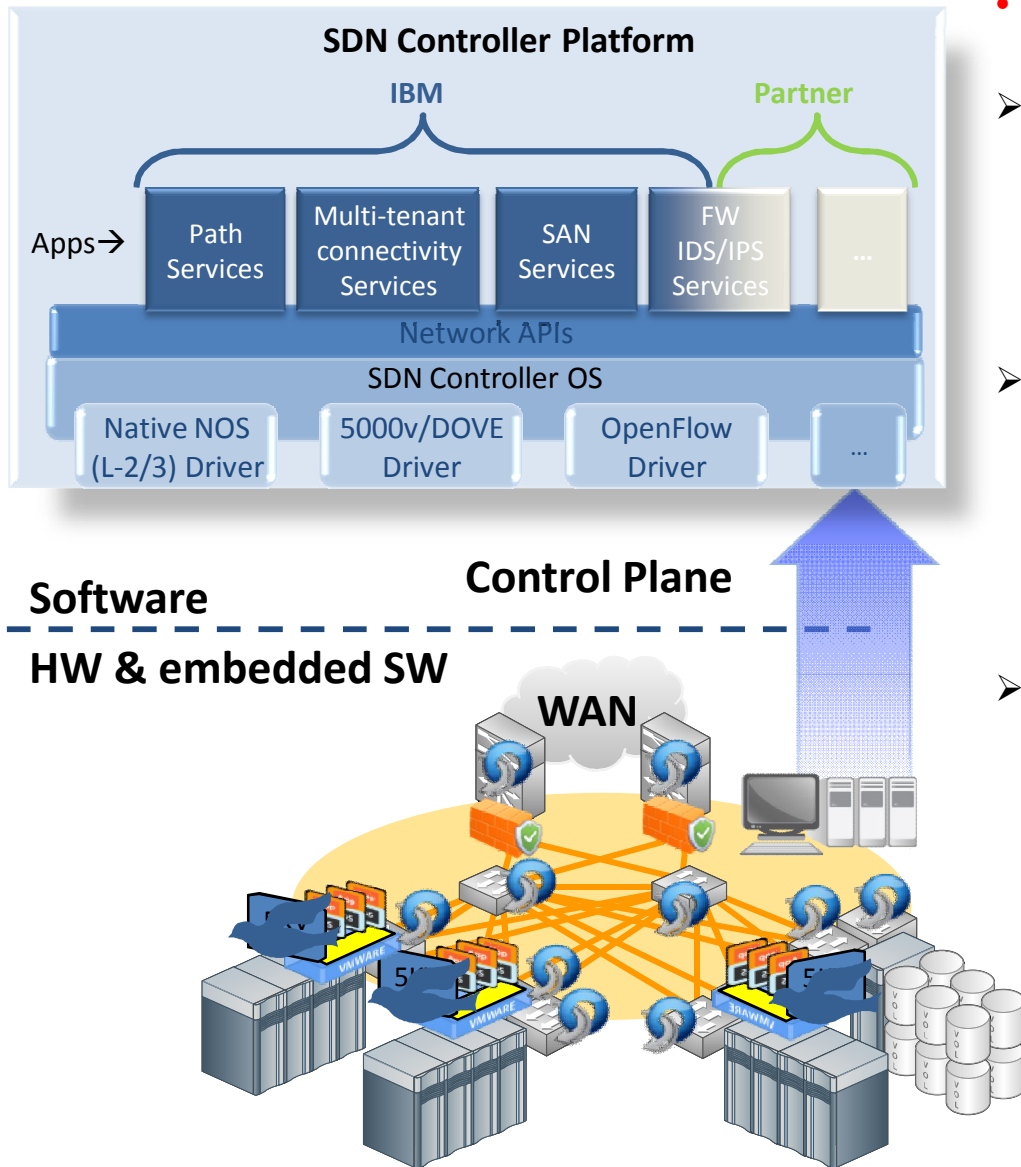
Worldwide OpenFlow Revenue, 2012-2016



Source: IDC, 2012



Source: NEC <http://www.nec.com/en/case/nittsu/>



What is IBM's value with SDN?

- Network services and Application Ecosystem
 - IBM can deliver the SDN key Integrated System elements
 - IBM can deliver new network services and applications

- Workload scaling simplicity: Cloud-scale, network Hypervisor
 - IBM can deliver the Distributed Overlay Virtual Ethernet (DOVE) network's "driver"


- Physical switch scaling simplicity
 - IBM can deliver the OpenFlow network's (OpenFlow Switch/OpenFlow Programmable Controller) "driver"

Controller	Platform	Language	Author
NOX	Linux	C++/Python	Vmware(Nicira)
POX	Any	Python	
BEACON	Window/Linux/MAC/Android	Java	Standford
Treama	Linux	Ruby/C	IBM/NEC
Floodlight	Any	Java	Big Switch
Maestro	Window/Linux/MAC	Java	Rice
RoutFloor	Linux	C++/Python	CPQD (Brazil)



IBM Programmable Network Controller at a glance

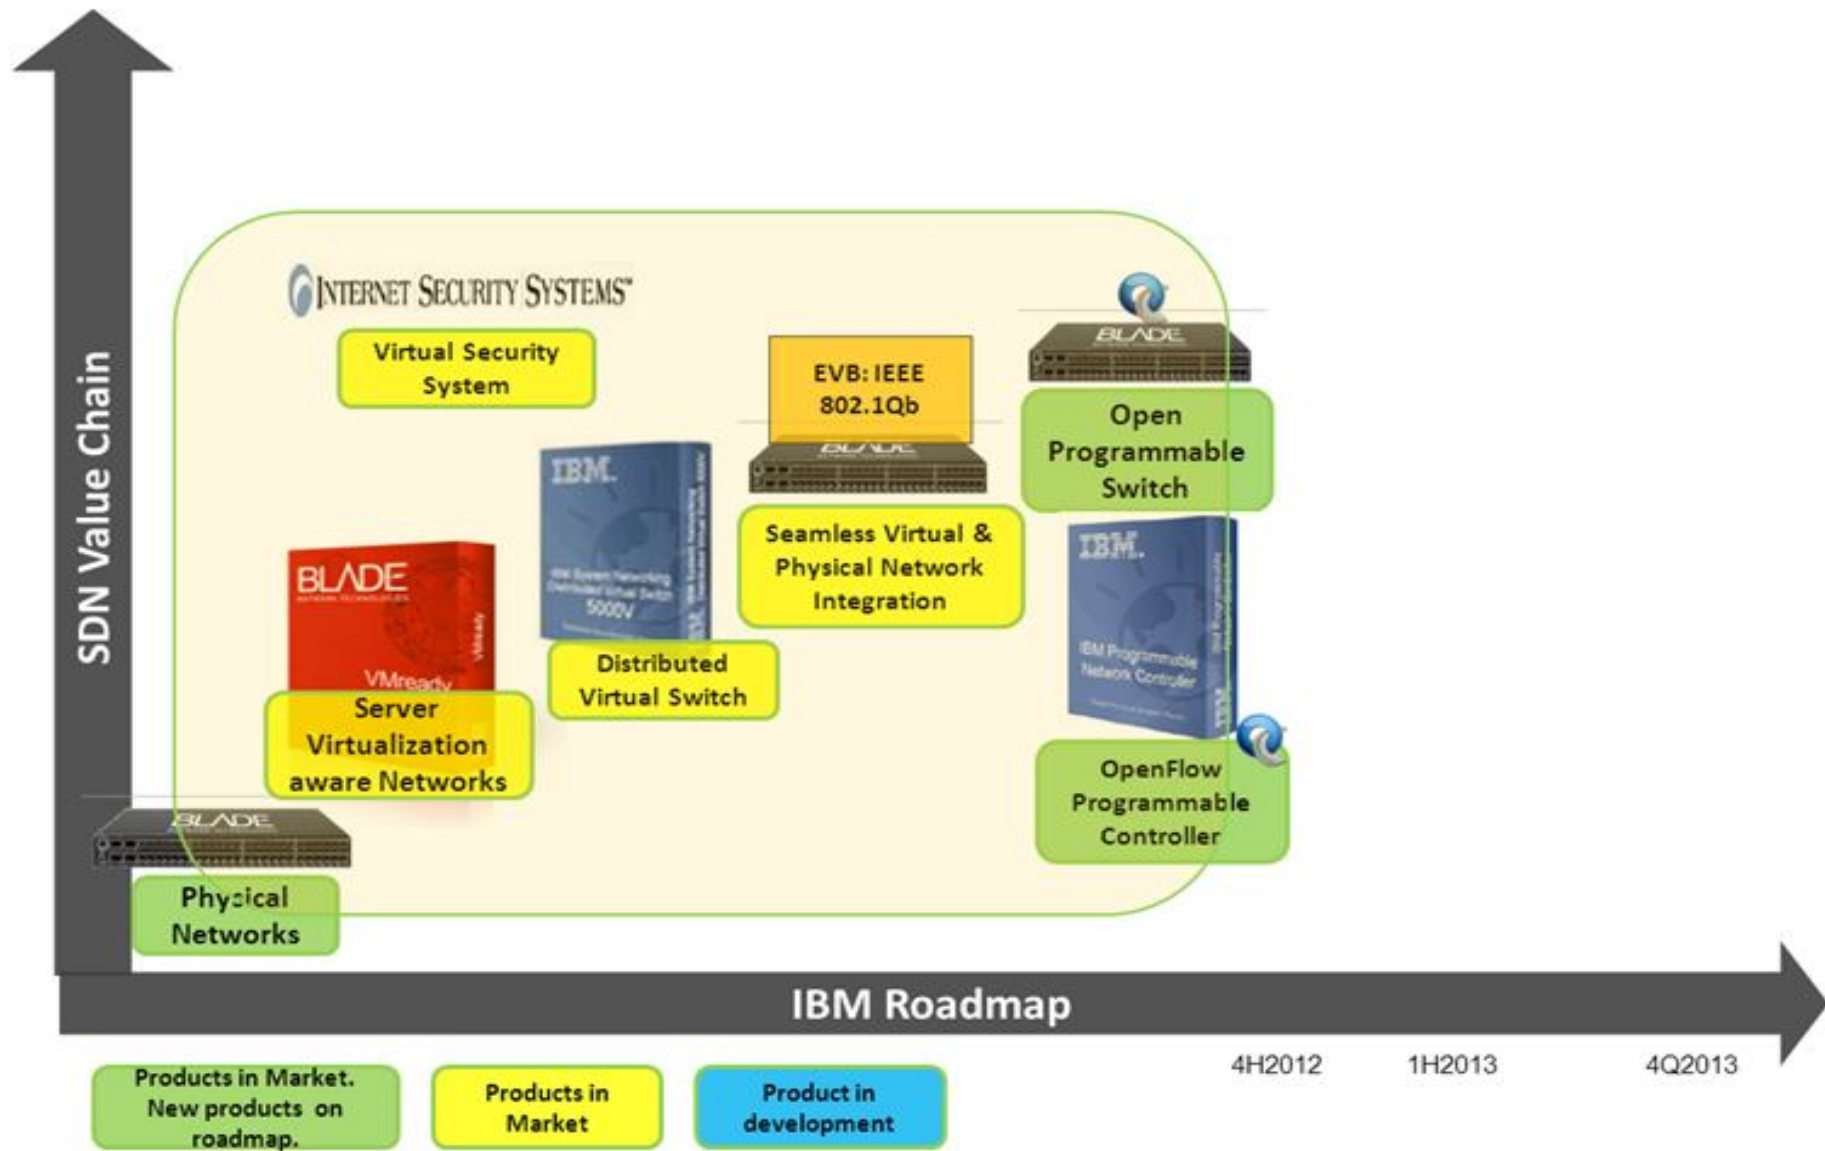
OpenFlow standard	<ul style="list-style-type: none"> OpenFlow 1.0 compliant
Verified OpenFlow Switches	<ul style="list-style-type: none"> Aggregation Switch: IBM RackSwitch G8264, NEC PF5240 Edge Switch (Edge of OpenFlow domain): IBM RackSwitch G8264, NEC PF5240
OpenFlow Vendor Extensions	<ul style="list-style-type: none"> Bit masking In band Broadcast/Multicast for wire speed forwarding
Virtual Tenant Network (VTN)	<ul style="list-style-type: none"> vRouter (L3) vBridge (L2) vFilter
North bound API	<ul style="list-style-type: none"> Web API
Number of VTN	<ul style="list-style-type: none"> 1,000 (Extended VLAN mode: 10,000 VLANs)
Number of Flows	<ul style="list-style-type: none"> 300,000
Redundancy Features	<ul style="list-style-type: none"> IBM PNC Active/Standby
Switch and Link Discovery	<ul style="list-style-type: none"> Topology discovery
IP	<ul style="list-style-type: none"> IPv4 IPv6 (L2 forwarding)
ARP	<ul style="list-style-type: none"> ProxyARP
Routing options	<ul style="list-style-type: none"> Shortest Hop ECMP (L2/L3) Avoid switch routing




- OpenFlow-based flow handling in hardware at line rate (1.28Tbps)
- 10GbE(SFP/SFP+) x 48 ports + 40GbE(QSFP) x 4 ports
- Support Layer 2 (MAC) forwarding table manipulated through OpenFlow
 - Layer 2 (MAC) table—Max 97K flow entries
 - 12 tuple flow table—Max 750 flow entries

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
Weight	• 9.98 kg
Power/consumption	• 50–60 Hz; 100–240 V; 330 watts
Temperature	• 0–40 degrees C
Humidity	• 10–90% non-condensing
Altitude	• 3,050 meters (10,000 feet)
MTBF	• 165,990 hours at 40 degrees C
Transceivers/cables	• QSFP+ 40GBASE-SR, 1M/3M/5M • QSFP+ DAC Breakout Cable, 1M • QSFP+ to QSFP+ cable • SFP+, 0.5M/1M/3M/7M DAC • SFP+, 1000Base-T (RJ-45) • SFP+, 1000Base-SX SFP
Protocol version	• OpenFlow 1.0.0
Number of OpenFlow table entries	• 80–128K (Layer 2 MAC table for OpenFlow) • 1,000 (12 tuple table)
Number of instances	• 1
Protocols	• No legacy protocols running in OpenFlow switch mode
Management	• Telnet, SSH, SNMP, sFlow
Redundancy	• Power/fan

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



- **Server specification for Production Network**
 - Two servers in a redundant configuration with
 - CPU – 1 Intel Xeon E5-2600 series processor with at least 8 cores
 - Memory – 64GB DDR3
 - HDD – 2 300GB, 10000RPM, RAID-1
 - NICs – 6 1000Base-T Ethernet ports
 - Redundant Power supplies
 - DVD-ROM optical drive
- **Server specification for a PoC or Demo or Test Network**
 - One server with
 - CPU – 1 Intel Xeon E5-2600 series processor with at least 8 cores
 - Memory – 4GB DDR3
 - HDD – 1 40GB
 - NICs – 2 1000Base-T Ethernet ports
 - Redundant Power supplies
 - DVD-ROM optical drive
- **Operating System**
 - Red Hat Enterprise Linux - RHEL 6.1(x86_64)
- **Recommended IBM servers**
 - System x3550 M4
 - System X3650 M4



Certified Interoperability into Cisco Networks**

- Auto-Negotiation
- 802.1p/Q VLAN Tag Propagation
- Jumbo Frame Support
- Link Aggregation Control Protocol (EtherChannel)
- 10GbE LAN PHY Support
- Rapid Spanning Tree Protocol (RSTP or PVRST+)
- Multiple Spanning Tree Protocol (MSTP)
- FCoE/DCB***



* MTM specific to AAS Ordering system
 ** <http://tolly.com/DocDetail.aspx?DocNumber=209116>
 *** <http://www.tolly.com/DocDetail.aspx?DocNumber=210140>

© Copyright IBM Corporation 2012

IBM
 Route 100
 Somers, NY 10589
 U.S.A.

Produced in the United States of America
 February 2012
 All Rights Reserved

IBM, the IBM logo, ibm.com BNT, BLADEHarmony, NMotion, Rackonomics, Server Mobility, and VMready and logos are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: ibm.com/legal/copytrade.shtml

Other product, company or service names may be trademarks or service marks of others.


References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

IBM RackSwitch for 1/10GbE Top of Rack Deployments

Products			
Products	G8000 1/10GbE Aggregation	G8052 1/10GbE Aggregation	G8124 / G8124-E 10GbE SFP+ Switch
IBM part numbers	G8000R:0446013, G8000F: 7309CFC, G8000DC: 7309CD8	G8052R: 7309G52, 1455-48E* G8052F: 730952F	G8124R: 0446017; G8124ER: 7309BR6 /1455-24E*, G8124F: 7309BF9; G8124EF: 7309BF7; G8124DC: 7309BD5
Form factor	1U Rack	1U Rack	1U Rack
Ports	48 1GbE (44 RJ-45 & 4 SFP) up to 4 x 10GbE CX4 or SFP+	48 x 1GbE RJ-45 4 x 10GbE SFP+	24 x 10GbE SFP+
Max thru-put (bidirectional) Gbps	176 Gbps	176 Gbps	480 Gbps
Average port to port latency	4.6 microseconds	1.8 microseconds	570 nanoseconds
Virtualization	VMready	VMready	VMready and Virtual Fabric
Availability	Hotlink, VRRP, Uplink Failure Detection	vLAG, Hotlink, VRRP, Uplink Failure Detection	vLAG, Layer 2 failover, Hotlink, VRRP
Other	Stacking	Hot-swappable power /fans	DCB/CEE/FCoE
Power	120W	130W	115W to 168W
Recommended for:	Stacking, zero over-subscription, low power consumption, optimized airflow	Low latency 1/10G, no over-subscription, low power consumption, optimized airflow	Lossless 10G, low latency, SFP+ low power consumption, optimized airflow

IBM RackSwitch for 10/40GbE Top of Rack Deployments

Products			
Products	G8264 10/40GbE SFP+/QSFP+	G8264T 10/40GbE 10GBase-T/QSFP+	G8316 10/40GbE SFP+/QSFP+
IBM part numbers	G8264R: 7309G64, 1455-64C* G8264F: 730964F	G8264R: 7309CR9 G8264F: 7309CF9	G8316R: 8036ARX G8316F: 8036AFX
Form factor	1U Rack	1U Rack	1U Rack
Ports	48 to 64 x 10GbE SFP+ up to 4 x 40GbE QSFP+	48 10GBase-T 10GbE RJ-45 up to 4 x 40GbE QSFP+	16 x 40GbE QSFP+ up to 64 x 10GbE SFP+
Max thru-put (bidirectional) Gbps	1,280 Gbps	1,280 Gbps	1,280 Gbps
Average port to port latency	880 nanoseconds	3.2 microseconds	800 nanoseconds
Virtualization	VMready and Virtual Fabric	VMready	VMready and Virtual Fabric
Availability	vLAG, Layer 2 failover, Hotlink, VRRP	vLAG, Layer 2 failover, Hotlink, VRRP	vLAG, Layer 2 failover, Hotlink, VRRP
Other	DCB/CEE/FCoE, hot-Swappable power/fans	Hot-Swappable power/fans	DCB/CEE/FCoE, hot-Swappable power/fans
Power	275W	385W	330W
Recommended for:	Massive scalability, 10/40G, low latency, low power consumption, optimized airflow. Also supports OpenFlow	Massive scalability, 10/40G, flexible, low cost CAT6 10GBase-T connectivity, optimized airflow	Cost-efficient 40G alternative to expensive core switches. Aggregates multiple racks of servers. Low latency, low power consumption, optimized airflow



Thank You