

Technology Research

Brian Marshall	212 653 9047	bmarshall@isigrp.com
Stephen Patel, CFA	310 473 0767	spatel@isigrp.com

1 of 7

**Tues
Nov 5
2013**

How the Largest Spenders of IT are Thinking about their Future Networks...

ISI's Thoughts from the Open Networking User Group (ONUG) Conference:

We were fortunate to attend the Open Networking User Group (ONUG) two-day conference held in NYC this past week. As a reminder, ONUG is an end-user comprised working group where senior IT representatives from virtually all the tier-1 financial service companies, leading global service providers and top hyperscale cloud companies gather every ~6 months to discuss the concept of "open networking" and how they are architecting these new technologies into their existing infrastructure. Importantly, no product vendors are allowed to attend the user discussions and we participated on a panel regarding the financial implications of software-defined networking (SDN) on the traditional data networking industry. While ~25% of the hundreds in attendance are conducting proof-of-concepts (POCs) for SDN and network function virtualization (NFV) currently, we were surprised to learn that not one of them are actually in production today with these new architectures. Obviously, we expect this to start to change as we exit 2013 and will likely enter 2014 with a running start as many of these trials are slated to migrate into production.

In our view, the most important question remains: "will the SDN market ramp in a similar fashion as server virtualization and emerge into a multi-billion market within a 5 year window?" While it is too early to answer this question with a high degree of conviction/accuracy, we do believe the following vendors are positioned to benefit as SDN trends start to meaningfully emerge: VMW (STRONG BUY, \$115 PT) with its NSX offering, Arista Networks with its open EOS software platform, PLUMgrid and Versa Networks driving layer 4-7 virtualization (e.g., NFV), and Cumulus Networks leading the whitebox switching landscape with its open network operating system (Cumulus Linux). Below are select highlights from the two-day event and please call/email for greater levels of granularity:

High Level Summary Points of ONUG October 2013:

- **Where is the Industry with SDN Today?** 1) Huge movement in open networking, 2) Need a software ecosystem (e.g., welcome Dev/Ops into this community), and 3) Need an architectural shift (e.g., using same format for 25 years now)
- **Why are Users Interested in SDN?** OPEX reduction mostly but now the reason is starting to become the concept of virtualized networks in the datacenter (e.g., ~50% are reviewing today, ~25% are in pilots and ~15% are in production today)
- **IT Needs Automation?** ~80% of engineer's time is spent on manual configuration on network (e.g., 1 engineer to 120 network switches vs. 1 to 10,000 servers in a hyperscale environment)
- **Open Networking Promise?** 1) Lower OPEX and CapEx, 2) Software ecosystem for rapid innovation, 3) Vendor independent network design flexibility, 4) Centralized network tools, and 5) Faster IT and efficient business process
- **Networking Trend Setters?** Vendors no longer setting the direction of networking industry...it's now set by service providers and hyperscale guys (e.g., users are more important today vs. the past)

- **Issue Exists in Software?** Software has become big, fat and bloated today while hardware is fine
- **Proper Approach?** Some players today like Arista are taking a networking device and making it look like a server that can be programmed (i.e., software is the important part here)
- **New Industry Demands Emerging?** New workloads possess new traffic patterns and require scale-out architecture
- **Challenges?** Include 1) scale, 2) business resiliency, 3) IT agility, and 4) asset utilization (e.g., active-active requires a ton of bandwidth and lots of storage = expensive)
- **New Underlay Network Architecture?** Datacenter core to spine to leaf (ToR). New 2-tier layer 3 network optimized for next-gen IP traffic architecture and east-west traffic
- **Old Networking Underlay Architecture?** WAN edge to core to distribution to access (3-tier layer 2 network designed for north-sound traffic)
- **Roadblocks for New Architecture?** Network segmentation, workload mobility, workload migration, physical device provisioning, traffic monitoring, networking software maintenance
- **Top Use Cases for SDN?** 1) SDN use case for Branch Wide Area SDN (WA-SDN), 2) Automated Network Segmentation on L2/3, 3) Virtual Networks and Overlay, 4) Traffic Monitoring and Visibility, 5) Layer 4-7 Network Services in SDN, 6) Whitebox Switches in Access Closets, and 7) Declarative Language for SDN Apps
- **Why Minimal Traction of SDN Today?** Not in production today because vendors aren't driving down lower OPEX yet...need this to occur. Vendors proposing proprietary solutions and customers need interoperability (i.e., "can't be one jail cell to the other jail cell")
- **Need for "Carbon-based Life Forms?"** Top U.S. Carrier offers several SDNs already but still need "carbon based life forms" and wrote their own controller and leveraging the AMD SeaMicro servers with orchestration layer
- **OCP in Hyperscale Environments?** Top social network has OCP boxes in the lab today and chose to start at ToR box because it is the easiest to build today
- **NFV?** Top cable company thinks it's important to do layer 4 in layer 3. Top carrier thinks the real value to be created in SDN is in NFV because it translates into raw dollars to customers (e.g., firewall, load balancers, VPN, etc.)
- **What do Social Networks Want from SDN?** "Agility is what we want and programmability is how you get it" (i.e., means to an end)
- **Incumbent Vendors?** Top U.S. carrier has made its bets but always looking for new better approaches (e.g., silos were wrong). Top social network thinks opportunity for vendor base due to customer desire for abstraction layer (These companies are software engineering shops and want access down to the silicon layer but many other companies don't want this level of visibility)
- **OCP switching?** Top social network looked at "cost plus model" and determined what margin they were willing to pay for not building switches themselves. In the datacenter, there isn't much differentiation in hardware but it's in the software today
- **Importance of Linux?** Need more Linux people in the networking space as the silos collapse and need a common OS and tools to manage. This will be Linux standardization. Don't need Linux down to the ASIC level however
- **Bare Metal Networking Economics?** Compelling considering a fully loaded x86 servers cost ~\$5k vs. an OEM Ethernet switch at ~\$25k
- **Open vs. Closed...Dramatically Different Growth Rates?** Look at x86 server market (e.g., closed OS like Windows and Solaris or open OS like RedHat and Ubuntu...apps closed with APIs like SAP or open with APIs like Hadoop)
- **Blob vs. Ecosystem?** User needs to decide whether they want to continue down the path of a "single vendor blob" (i.e., integration of app, OS, HW, etc.) or migrate to a more flexible "multi-vendor ecosystem" (e.g., decouple the apps, OS, HW, etc.)
- **Complexity in Ethernet Switch?** Resides in the switching silicon and PHYs but merchant silicon vendors have emerged to make whitebox switch market feasible (e.g., BRCM Trident2 chipset, Metcalf.)
- **Switch Architectures?** Started with a bus (e.g., CSCO Catalyst 5000) and moved to a switch fabric (e.g., DELL's Force10) and moved to a cell-based switching fabric (e.g., Arista's 7500E powered by BRCM's Dune Networks)
- **Open Compute Networking?** Initiated by FB and BRCM and put forth switch proposal (e.g., OCP Open Switch) for Open Rack 1U form factor (e.g., 21 inches vs. standard 19"). Based on BRCM Trident2 switch and AMD x86 control processor (e.g., quad core, 1.6GHz). Multiple network OS options (e.g., ICOS, Big Switch's Switch Light, Cumulus Linux, etc.). Standard Linux distribution/tools. Optional data plane processor from BRCM's NTEL acquisition. INTC also has an Open Networking platform based on their Fulcrum switch and IA processors


- **Open Network Install Environment (ONIE)?** Provide an install environment for networking equipment. ONIE is like PXE installer for servers but designed for switches (e.g., zero-touch install, etc.). Cumulus and Big Switch are working on this. Jail breaking an OEM switch is very difficult to do. Customer standpoint? Order a switch from manufacturer with ONIE pre-installed and then: 1) unpack and rack switch, 2) connect to mgmt network, 3) power up switch, 3) ONIE (small Linux kernel) discovers OS, and 4) OS first boot
- **How are Users Deploying Puppet?** Cumulus uses this mgmt/automation tool to rebuild servers. Some tier-1 financial service companies use it for auditing purposes (e.g., easily see what is on a server) and because it's free!
- **Target Market for Whitebox Switches/Routers?** Cumulus target market? Likely not going after CSCO CRS core routers (e.g., too complex as it peers with too many devices via many protocols, etc.). The datacenter is where this market will focus and likely target spine and aggregation switches in addition to the core market of top-of-rack (ToR) switches
- **Catalyst for SDN?** Nicira started the value of SDN via abstraction of apps/network and decoupling (e.g., \$1.3bil acquisition of Nicira by VMW). Dynamic, programmable and automated
- **Why Virtualize Networks?** 1) Easy to deploy and orchestrate workloads in a multi-tenant cloud environment and that is why network virtualization is gaining traction, 2) Hybrid cloud for bursting is another reason, 3) Scalability and efficiency limited by vLANs (e.g., Goal is 80% utilization rate but reality is 25% today. Network design limited by poor isolation as well. L2 doesn't scale because of need to track MAC addresses. L3 routing scales but traditional architecture doesn't support IP overlap between tenants), 4) Poor orchestration of virtualized L4-7 appliances, 5) VMs aren't treated as first class citizens (e.g., east-west traffic is poorly managed today as over 70% of servers are VMs). Lack of priority action and rate limiting at VM level, and 6) Dynamic workload over multiple clouds is tricky today
- **Current Internet?** Made up of expensive, complex and closed boxes
- **Future Internet?** SDN will require only specialized packet forwarding hardware for the data plane and the control plane and all services are abstracted outside the box and reside in network operating system, controller platform or via northbound interface APIs. SDN enables lower complexity/cost, granular traffic management that is dynamic and automated. For example, a top Japanese carrier is using Quagga open source routing software today
- **Modes of SDN Deployment?** 1) "in-network model" (e.g., existing/green-field network fabric is upgraded to support OpenFlow (OF) and OF-enabled port connects to OF controller cluster), and 2) "overlay model" (e.g., no change to existing network but intelligence is added to the edge such as an additional appliance or an enhanced server kernel bridge like open vSwitch in x86 hypervisors. Tunnels used (e.g., STT [stateless transport tunneling], VxLAN, MPLS over GRE [generic routing encapsulation], etc.)
- **Select Public SDN Deployments Today?** 1) GOOG, 2) T, EBAY, Fidelity Investments, NTT, RAX, and 3) Genesis Hosting
- **Business Potential of SDN?** 1) reduce time to revenue, 2) new revenue opportunities, 3) improve policy compliance, 4) OPEX saving, and 5) reduce OPEX during upgrades
- **SDN Deployment Issues?** Sometimes scalability (e.g., only ~250 VMs in 5 racks, etc.)
- **OpenFlow?** Offloads control intelligence to a remote server. v1.3.2 is most current version of OF today. Conducts various "match" and "actions" (e.g., encapsulate, header rewriting, etc.). Match: L1 (e.g., Tunnel ID, Switchport), L2 (e.g., MAC address, VLAN ID, Ether type), L3 (e.g., IPv4/6, ARP, etc.), and L4 (e.g., TCP, UDP, etc.)
- **Open vSwitch (OVS)?** Kernel module that replaces the standard Linux bridge to provide significant packet matching and processing flexibility
- **OVSDB?** API that is alternative to OF and includes configuration and control. Manages slow moving state. Lightweight, transactional, not SQL, persistent
- **Open-source OF controllers?** Ryu (NTT), Open Daylight (Linux Foundation) based on Java, Floodlight (Big Switch), Trema (NEC), Nox, etc.
- **Open Daylight?** Vendor driven consortium (e.g., CSCO, IBM, etc.) developing open-source SDN controller platform. SAL (service abstraction layer) and Open Daylight APIs (REST)
- **OpenStack?** Networking using Neutron today (e.g., create network, associate a subnet, boot a VM and attach it to the network, delete the VM, delete any ports, delete the network). Neutron is API used to program network (plugin to OpenStack)
- **Goal of Deploying Network Virtualization?** Abstract apps from network via SDN-enabled infrastructure. Requires traffic isolation across virtual networks, orchestrating virtual L4-7 service chaining, scalably identifying individual VM's traffic, and integration with legacy world.
- **Insertion Choices?** 1) bare metal (e.g., native OS with baked in containerization), 2) hypervisor (e.g., KVM, Xen or Hyper-V), or 3) appliance (e.g., VMW ESX, DVS or distributed virtual switch)

Our “Four-Legged Stool” Analogy of Enterprise Infrastructure

We created the genesis of the “Four-Legged Stool” (“FLS”) analogy of enterprise infrastructure back in the summer of 2009 (originally dubbed the “Three-Legged Stool”) to assist the investment community by providing a simple representation of the required components used in the creation of a next-generation datacenter. Since then, the FLS has evolved to include a separate “security leg” and is widely used as an investment and competitive framework by investors as well as corporations. Specifically, **the “Four-Legged Stool” is a tool used to map and organize the strengths and weaknesses of respective product silos of the large mature information technology (IT) companies** (e.g., IBM, Hewlett-Packard, Cisco, Oracle, DELL and Huawei). The following are the key components of our “Four-Legged Stool” analogy:

- **Leg 1 (Enterprise Storage):** enterprise storage product portfolio aimed at the retention and analysis of corporate data (e.g., SAN or storage area networks, NAS or network attached storage, flash memory arrays, disk arrays, tape libraries, de-dupe, etc.)
- **Leg 2 (Networking/Telecom Equipment):** connectivity within, to and from the datacenter is established with networking/telecom equipment (e.g., routers, switches) which basically provides the “plumbing” responsible for transporting data from one location to another
- **Leg 3 (Servers):** data in the form of binary code (e.g., 1s and 0s) requires “computational horsepower” in order to be processed; primarily occurring in the computing elements of the network (i.e., servers)
- **Leg 4 (Security):** as corporate networks become more virtual (or hybrid) in nature, the importance of enhanced security features has grown rapidly and is near the top of most IT spending priority lists (e.g., authentication, threat protection, network forensics, etc.)
- **Stool (Software & IT Services):** intelligence or “glue” (i.e., software and IT services) is required to seamlessly integrate all disparate legs and create one unified “stool” or network aimed at increasing efficiency and productivity (e.g., virtualization, operating systems, middleware, applications, etc.)


ISI Group’s “Four-Legged Stool” Analogy of Enterprise Infrastructure



Source: company reports and ISI Group estimates.

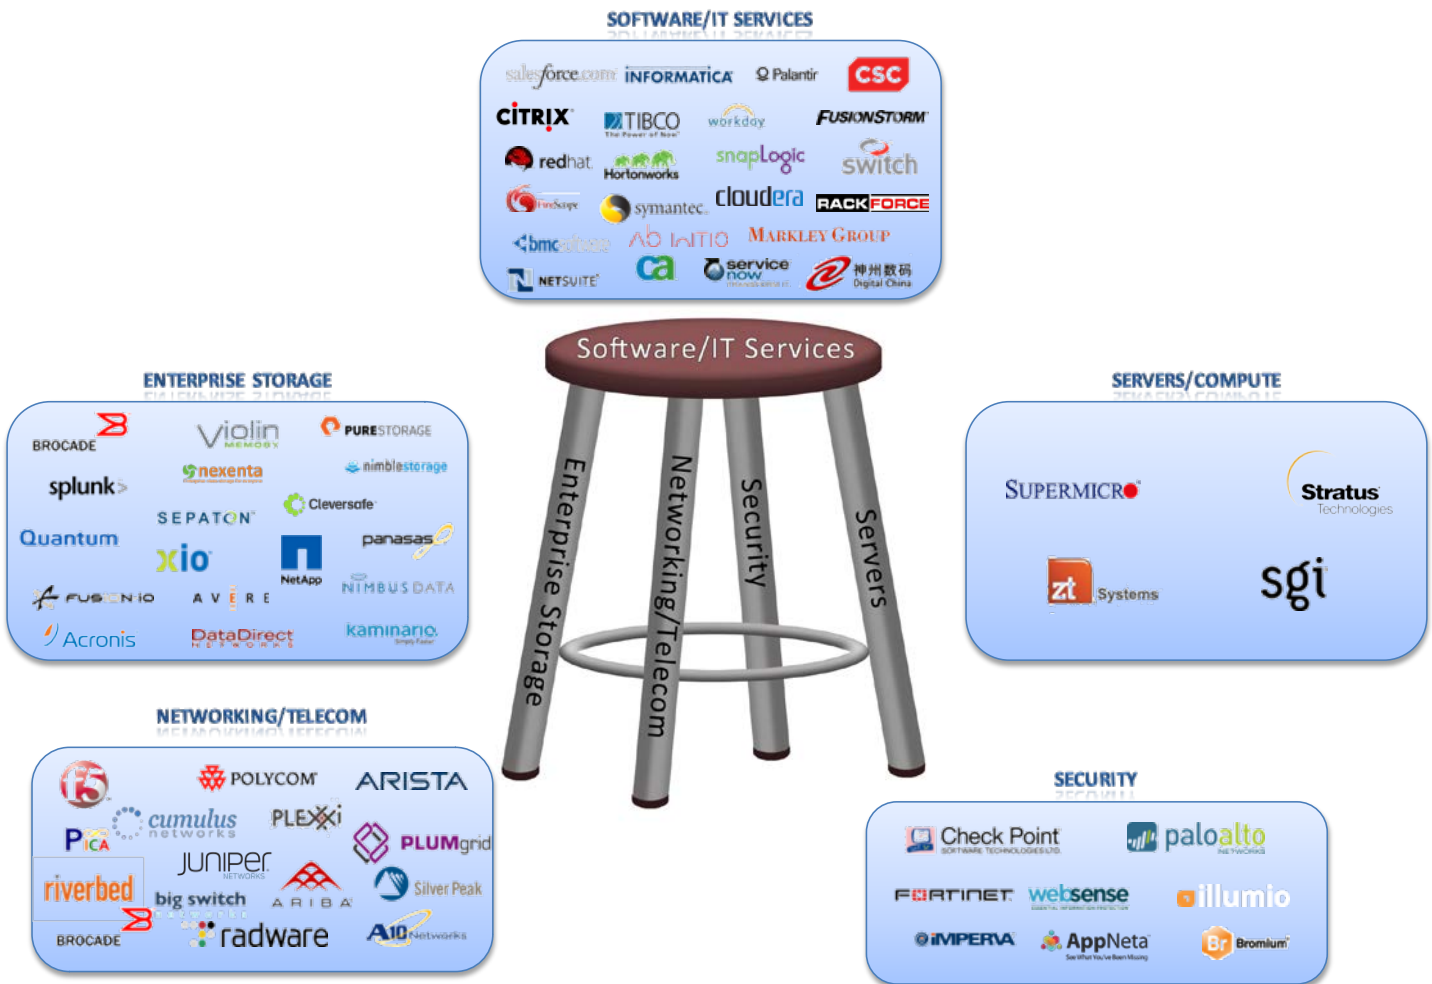
Total Addressable Market (TAM) of the “Four-Legged Stool”

(\$mil)	CY12
IT Services	\$840,000
Software	\$240,000
Storage	\$50,000
Networking	\$35,000
Servers	\$35,000
Security	\$10,000
Total	\$1,210,000



Source: company reports and ISI Group estimates.

Top Acquisition Candidates within the “Four-Legged Stool”



Source: company reports and ISI Group estimates.

Secular Growth Pedestal (“SGP”) Roster

EMC²
where information lives!

vmware

JUNIPER
NETWORKS

Members of the “SGP” represents “best-in-class” investment vehicles over a multiple year horizon.

The two key drivers of outsized returns for technology stocks include:

- Secular revenue growth
- Attractive operating margin structures

Common Characteristics of the “SGP”:

- Secular revenue growth targeting 15%+
- Attractive operating margin structures targeting 20%+

ISI Disclaimer

ANALYST CERTIFICATION: The views expressed in this Report accurately reflect the personal views of those preparing the Report about any and all of the subjects or issuers referenced in this Report. No part of the compensation of any person involved in the preparation of this Report was, is, or will be directly or indirectly related to the specific recommendations or views expressed by research analysts in this Report.

DISCLOSURE: Neither ISI nor its affiliates beneficially own 1% or more of any class of common equity securities of the subject companies referenced in the Report. No person(s) responsible for preparing this Report or a member of his/her household serve as an officer, director or advisory board member of any of the subject companies. No person(s) preparing this report or a member of his/her household have a financial interest in the subject companies of this Report. At various times, the employees and owners of ISI, other than those preparing this Report, may transact in the securities discussed in this Report. Neither ISI nor its affiliates have any investment banking or market making operations. No person(s) preparing this research Report has received non-investment banking compensation from the subject company in the past 12 months. ISI does and seeks to do business with companies covered in this research Report may have or has received non-investment banking compensation in the past 12 months.

DISCLAIMER: This material is based upon information that we consider to be reliable, but neither ISI nor its affiliates guarantee its completeness or accuracy. Assumptions, opinions and recommendations contained herein are subject to change without notice, and ISI is not obligated to update the information contained herein. Past performance is not necessarily indicative of future performance. This material is not intended as an offer or solicitation for the purchase or sale of any security.

ISI RATING SYSTEM: Based on stock's 12-month risk adjusted total return.

STRONG BUY	Return > 20%
BUY	Return 10% to 20%
NEUTRAL	Return 0% to 10%
CAUTIOUS	Return -10% to 0%
SELL	Return < -10%

ISI has assigned a rating of STRONG BUY to 16% of the securities rated as of 9/30/13.

ISI has assigned a rating of BUY to 32% of the securities rated as of 9/30/13.

ISI has assigned a rating of NEUTRAL to 48% of the securities rated as of 9/30/13.

ISI has assigned a rating of CAUTIOUS to 2% of the securities rated as of 9/30/13.

ISI has assigned a rating of SELL to 1% of the securities rated as of 9/30/13.

(Due to rounding, the above number may add up to more/less than 100%).

Price and Ratings Charts: Please refer to <http://www.isigrp.com/isia/arc.jsp> for price and ratings charts of individual companies

Risks and Valuation: Please find a discussion of risks and valuation in individual company reports available at <http://www.isigrp.com/isia/arc.jsp>