

October 17, 2012  
 Jeremy Li

Credit: [INTERFACE 2012 – Knowledge to Weather the Storm](http://www.f2fevents.com)  
[www.f2fevents.com](http://www.f2fevents.com)

Interface 2012 was held on Oct. 17, 2012 from 9:00 a.m. to 4:30 p.m. at the Anaheim Marriott, Orange County. Thanks to MR2 Solutions, Inc.'s invitation for this technology event and its MR2 THEATER, where the Dell Converged Infrastructure: Datacenter in a Box; Dell Force 10: Virtual Network Architecture, were presented by knowledgeable IT professionals. The audiences will be able to learn the latest Dell Technology in its virtualization and cloud.

### Dell M1000e Chassis vs. Leading Competitors'

It is obviously that Dell holds the leadership by offering the highest computer density at this time, as illustrated in the picture below:

## Dell M1000e vs Competitors

M1000e density leading modular blade chassis

	  M1000e	  BladeSystem c7000	  UCS 5100
<b>Metrics</b>			
Physical Servers Per Blade Enclosure	32	16	8
Enclosure Rack Unit Space Required	10u	10u	6u
Maximum I/O Modules	6	8	2
Chassis Required to match Dell 32 servers	1	2	4

### Pervasive Reliability:

The industry's only dual-media Failsafe embedded hypervisor, Hot-swappable, fault-tolerant SSDs with hardware RAID.

EqualLogic PS-M4110, latest EqualLogic Blade Array, is only compared with both HP's and IBM's due to no this kind of products from EMC, NetApp and Cisco for comparisons. It gave me the impression that PS-M4110 provides Simplicity, Flexibility and Automatic Peer Scaling while both HP (e.g., P4800) and IBM (e.g., BladeCenerS) either requires complex servers and storage nodes to achieve the features mention here or no such features offered.

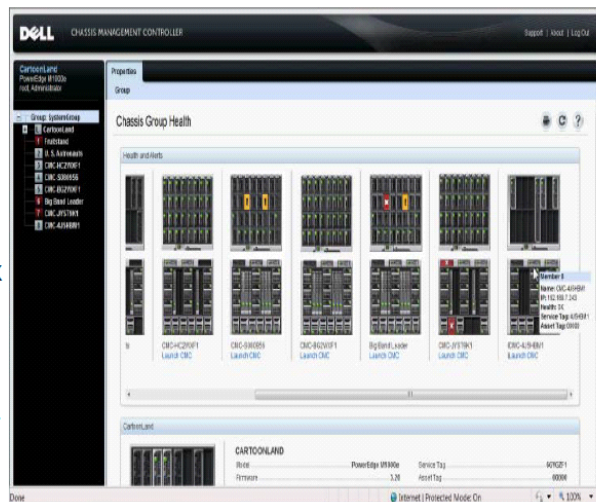
## Simplified and Scalable Management

Dell now offers many management tools without agents. It can challenge, in some cases, has surpassed HP Systems Insight Manager (HP SIM), a foundation for the HP unified server-storage management, to manage server and storage environments from a single console, with the feature rich, extensible and secure management tool.

Read below for Dell's Simple and Scalable Management Tools for details:

## PowerEdge M1000e CMC Simple and Scalable Management

- Simplified management
  - It's embedded, nothing additional to install
  - No agents, all out of band
  - Discovery is built-in
  - Easy to use, intuitive interfaces
- Scalable management
  - Interfaces and operations work across all chassis components
  - Automated 1:many operations for greater productivity
  - Update, deploy, configure, fix...
  - Ability to span multiple chassis (9)
- For monitoring, inventory, update, configuration and remediation



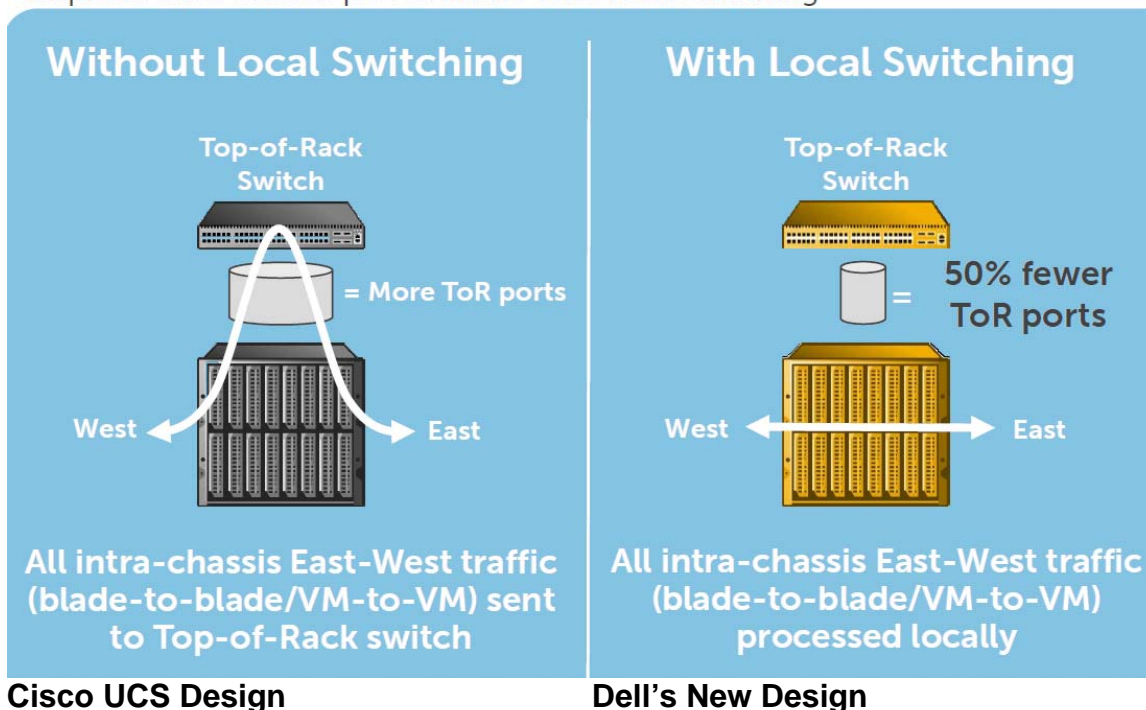
CMC V4.11 required

## North–South Traffic vs. East–West Traffic

It looks like HP, IBM and Dell have moved away from the traditional design within the blade chassis to a newer design – keep all intra-chassis East-West traffic processes locally with local switching technique, as illustrated in the picture below:

# Architected for Performance

Improve intra-chassis performance with local switching



The Gartner has predicted that 80% network traffic will be East-to-West by 2014.

Therefore, Dell's newer design with local switching (put switches inside its chassis) can reduce the latency and cables. In other words, all intra-chassis East-West traffic (e.g., blade-to-blade or VM-to-VM) processed within the chassis. This new design is better than Cisco's UCS design in the data center.

## Dell Force 10 Switches with Virtual Network Architecture

Dell™ Force10™ specialist presented an outstanding session to introduce its newly acquired switches, formerly known as Force 10 switches based on Dell™ Force10™ operating system architected for consistent high performance, resiliency and portability across switch platforms while the Force 10 switches are at 1/5 cost and 1/4 power consumption of the Cisco's 7000 series.

The presenter impressed the audience that its Force 10 switches have replaced the Cisco's 6509 switches in NASDAQ and have been used by Facebook for many years to prove its reliability. The Force 10 switches, generally regarded as analogous to Rolls Royce in a high-end switches market, now are able to be brought into an Enterprise environment thanks to its acquisition in July, 2011.


A detailed comparison chart was displaced on a large projector screen showing that the Force 10 switches use much less power than Cisco's. For example,

Force 10 switches consume only 30 watts per 10G port, while Cisco's at 130 watts per 10G port.

Note: The flagship Cisco Nexus 7000 Aggregation Layer switch uses its very old chips from the Cisco 6500 series, which consumes heavy power, especially on the 10GbE ports. The Nexus 5000 switch via Cisco's acquisition is much better in its power consumption by relying on up to date technology.

Cisco, traditionally, requires a customer to use 2N redundancy. For example, Cisco Hot Standby Routing Protocol (HSRP) must be used for 2N redundant Active/Passive network (the most expensive approach), which requires to purchase two or multiple switches [e.g., Catalyst 6509 switch (9 blades, each blade has 12, 24 or 48 ports) or Catalyst 6513 switch (13 blades, each blade has 12, 24 or 48 ports) with supervisor engine such as 720, which is the heart of the switch running an IOS), while Dell's Force 10 promotes N + 1 redundancy (one spare switch] for more economical deployments to achieve a similar result and delivers the lower total cost of ownership (TCO).

A traditional three-tier architecture and network (e.g., core, distribution, and access) is illustrated in the following diagram in a large production environment:



A client (e.g., users in VLAN 2) connects to an access (edge) switch (e.g., Cisco 2960 or 4006, but the picture is omitted here), which connects to a distribution switch (e.g., Cisco 3750, labeled as Dist2 and Dist 3 here), which connects to a core switch (e.g., Cisco 6509, labeled here as Core A and Core B.)

A recent new design has been moved towards a two-tier architecture, which deliver a lower-latency and high-performance network by removing a distribution layer.

## Interoperability

The reason, Cisco 6509 switches at NASDAQ were replaced by Dell's Force 10 switches, is simple. NASDAQ IT follows the Industry Standard Protocols such as LLDP (not Cisco's proprietary CDP), VTP (not Cisco's proprietary GVRP) or OSPF (not Cisco's proprietary EIGIP). As a result, replacing one vendor's product with Force10 switches, or co-existing two vendors' switches made easier

Dell Force 10 switches can work together with Cisco's without any issues as long as an environment follows an open standard.

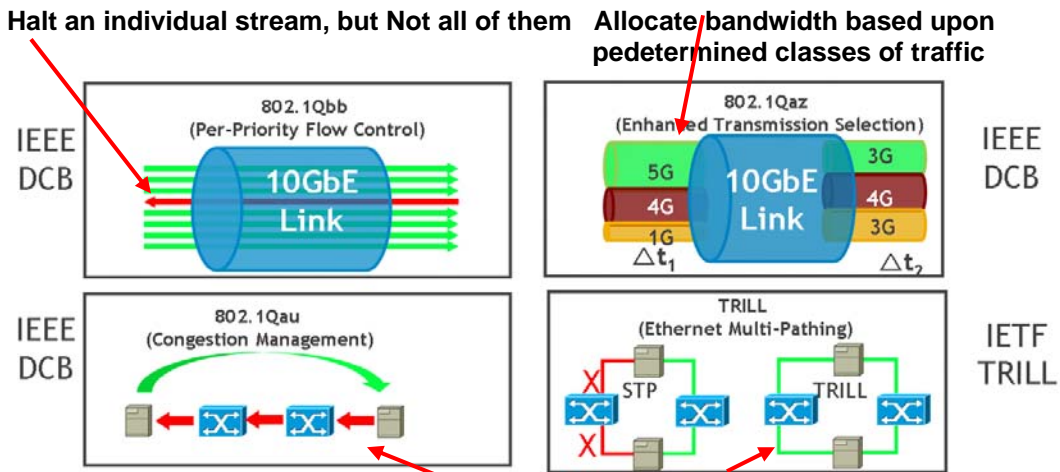
Note: Cisco Nexus 5000 Series Switches are the latest switches based on modular Cisco NX-OS operating system and has won FCoE Market Leader Award. It offers active/active redundant feature, instead of its traditional active/passive redundant feature. Click on the link <http://www.cisco.com/en/US/products/ps9670/index.html> for details.

## Date Center Bridging (DCB)

DCB (10GbE or higher) is a set of standards to enhance Ethernet for the data center. With DCB technology, the traffic can be shared among Layer 2 lanes. In other words, if one channel is preconfigured with network bandwidth of 5Gbps and this lane has traffic of 3Gbps with additional 2Gbps more capacity to be used, then, this lane can be shared by other lanes' traffic in order to fully utilize its existing bandwidth of 5GbE, as illustrated in a screenshot of top right pane below:

IEEE Standards:

- **PFC** – Priority-based Flow Control (*802.1Qbb*)
- **ETS** – Enhanced Transmission Selection (*802.1Qaz*)
- **CN** – Congestion Notification (*802.1Qau*)
- **DCBX** - Data Center Bridging capability eXchange (*802.1Qaz*)



**End-to-End Communication between end-points**      **Multi-pathing for Ethernet Replace Spanning Tree**

Windows Server 2008 R2 and Windows Server 2012 support DCB. If iSCSI with DCB enabled, it can achieve a balanced iSCSI throughput.

The benefits of using 10GbE with DCB:

- Convergence of applications and networks to use a single Infrastructure
- Different classes of traffic with different service guarantees
- Configuration of end-stations by switches through DCBX
- Mixed speed bridging with 10GbE to 40GbE/100GbE

## Spanning Tree Protocol vs. Trill

A newer technology to replace a traditional slow Spanning Tree Protocol (STP for short) network has emerged today in order to meet a new challenge in Cloud environment, such as Virtual Desktop Infrastructure (VDI), Business Intelligence (BI) with Big Data, Video Conferencing and etc. The STP is used to detect and prevent loops, which makes a network resilient and is able to deal with a sudden failure of a link between switches

Troubleshooting a STP related issue is not an easy task, often requiring a great deal of time to locate the root cause of the failure from a STP-based network.

Other network vendors (e.g., HP) have developed a newer technology to overcome network complexity with a more flat network design that will be faster to converge, easier to scale and simpler set up and manage. As a result, it can overcome the legacy multi-layer, slow and expensive STP-based network, as illustrated in the table below:

	<b>HP IRF-based designs</b>	<b>STP-based designs</b>	<b>Competitive solutions</b>
<b>Rapid failover</b>	<b>Yes</b> Sub 50 milliseconds	<b>No</b> Often measured in seconds	Varies by protocol
<b>Design simplification—common across data center/campus core/edge</b>	<b>Yes</b> Virtualizes up to 9 switches Common across layers/devices	<b>No</b> Switch by switch configuration	<b>No</b> Different protocols used at each layer and device type
<b>Support for Layer 2/3, MPLS, IPv6 protocols</b>	Yes	<b>L2 only</b>	Varies by protocol
<b>Performance</b>	<b>Very high</b>	<b>Low</b>	<b>High</b> Design/Device specific
<b>Geographic resiliency</b>	<b>Yes</b> Supported across 70 km	<b>No</b> L2 only, limited to site	Unclear Varies by protocol
<b>Overall administration effort and cost</b>	<b>Low</b> No extra cost	<b>High</b> No extra cost	<b>High</b> Often requires additional license and hardware

**Source: HP White paper – Reducing network complexity, boosting performance with HP IRF technology**

Simplicity of next generation network design is a key for a successful cloud computing now and in the future. Therefore, most potential or existing problems will be instantly disappeared and the troubleshooting time will be greatly reduced from often weeks, if not months down to a few hours or minutes.

## Next Generation Network Infrastructure

Next generation Network Infrastructure (LAN/WAN) is very important and must be addressed ahead of time prior to the Data Center virtualization taking place. Any organization will feel pain after a virtualization took place, then, has realized that the network performance is not good as expected. As a result, many users had to request to restore from a VDI solution to the original desktop solution.

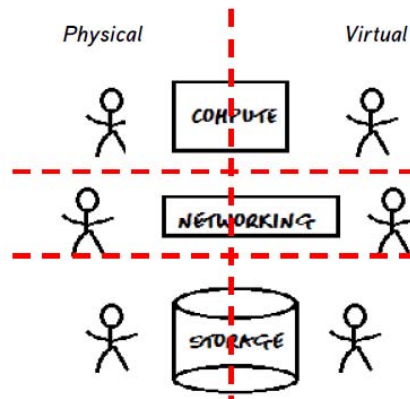
Therefore, any enterprise(s) should look for next generation network infrastructure (e.g., the “Networking” portion, as illustrated in a screenshot below) to prepare for cloud computing.

## Why haven't we figured this out yet?

*Convergence should be customized to drive specific Customer business outcomes*

### challenges:

- *Complex*
- *Distributed*
- *disconnected*
- *Slow*
- *Wasteful*
- *OPEX > budget*
- *Missed SLA's*



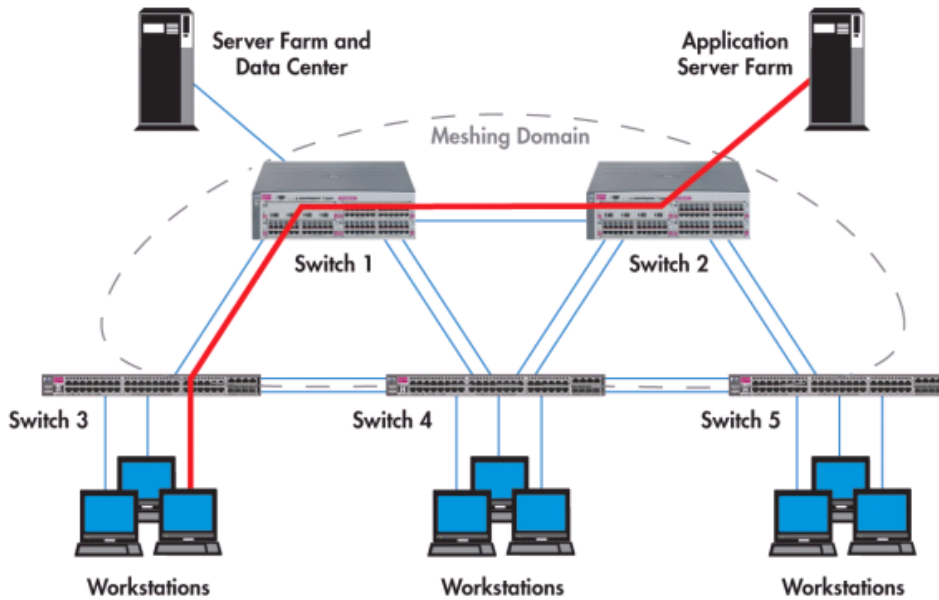
### Outcomes:

- *Consolidated*
- *Connected processes*
- *Agile*
- *collaboration*
- *Easier to manage*
- *Control*

Enterprise applications demand more and more network bandwidth besides the need for switches with high-density, high-speed ports. Therefore, simplicity of next generation network design is a key for a successful cloud computing now and in the future, in addition, most potential or existing problems will be instantly disappeared and the troubleshooting time will be greatly reduced from weeks, if not months down to a few hours or minutes.

The point-to-point connection is an ideal situation in the next generation network environment, while most companies are seeking a 2-tire network solution, as illustrated in a screenshot below, since a point-to-point solution is expensive, often, it is impossible.





**A 2-Tier Network Infrastructure (Access-to-Core Layer)**

## iSCSI vs. Fibre Channel


Implementing a 10GbE Fabric solution will reduce a cost by 48% and 50% in term of the Cable Infrastructure and LAN Infrastructure per port, respectively, as illustrated in a table below:

Value of 10G	
LAN Component	10 GbE Fabric Effect on Cost
Cable Infrastructure	48% Cost Reduction
LAN Infrastructure (per Port)	50% Cost Reduction
Server	12% Cost Increase
Storage Infrastructure	0% (Same for Both Fabrics)
<b>Total Savings (per Server)</b>	<b>18.5% Overall Cost Reduction</b>


For example, implementing an Intel X540 NIC not only enables 10GbE traffic using 10GBase-T over standard Cat6 copper cables with an RJ-45 connector, but also reduces the cable infrastructure by up to 48%, as illustrated in two pictures below:



The advantage of the 10GB NIC is cost per port, port management, cable management and virtualization. In some enterprises, a FC network, as illustrated in a picture below,



has been replaced by a 10GbE network, as illustrated in a picture below:



It is much economical to use iSCSI instead of Fiber Channel in most small and medium businesses (SMBs) market unless an environment has an absolute reason to select Fiber Channel for a new implementation. Here are some of the reasons listed below:

Why is Fibre Channel So Expensive?	iSCSI is Much Economical
<ul style="list-style-type: none"> <li>▪ FCAs provide full hardware acceleration                             <ul style="list-style-type: none"> <li>– 1-2% CPU load moving 800MBps</li> </ul> </li> <li>▪ FC Switches have advanced features                             <ul style="list-style-type: none"> <li>– Full duplex, non-blocking designs</li> <li>– FSPF Routing</li> <li>– Hardware-enforced flow control</li> <li>– Switch provides default services</li> </ul> </li> <li>▪ FC SANs are deployed as 2N Redundant networks</li> </ul>	<p>A TCP/IP Offload NIC can provide comparable performance with 1-3% CPU load moving 800MBps</p> <p><b>Note:</b> Every server shipped comes with 2 or 4 TCP/IP Offload NICs now, while dual HBAs (FC) must be purchased for every host.</p> <p><b>iSCSI Switches Features</b></p> <p>Full duplex, non-blocking designs</p> <p><b>No – FSPF Routing</b></p> <p>Hardware-enforced flow control</p> <p>Switch provides default services</p> <p>Optional: iSCSI SANs can be deployed either 2N or N+1 Redundant networks by using multi-path via iSCSI initiator</p>
<ul style="list-style-type: none"> <li>• DCB (Not Available??)</li> </ul>	<p>Yes, the DCB can share the bandwidth among many lanes.</p>
<p>8Gbps is in common, while 16Gbps is just available at 2% market share</p>	<p>10Gbps is in common and 40Gbps is available, while 100Gbps is in development</p>

Again, selecting iSCSI or FC deployment really depends on an environment.

## **CAT 5e vs. CAT 6 Cable**

Why the CAT 6 cable is more expensive than the CAT 5? It is because it contains four twisted wire pairs and more shields and run up to 250 MHz.

The CAT5e cables will be fine for a Gigabit transmission speed in most scenarios. If the distance between the client and switch is over 50 meters and the Gigabit speed is required, it will be best to ensure that a CAT6 cable is in use. Both CAT 5e and CAT6 cables are designed for 100 meters. As a general rule of thumb, any new cable installation should be either CAT6 or even CAT6e if there is a 10Gbps potential in the near future. Click on the link below for details.

[http://www.ehow.com/about\\_6521486\\_specifications-cat5\\_-cat6-cat6e-cables.html](http://www.ehow.com/about_6521486_specifications-cat5_-cat6-cat6e-cables.html)

## **Conclusion**

Interface 2012 gives IT professionals an opportunity for face-to-face interaction among local IT professionals, industry-leading manufacturers, and regional value-added resellers. This one-day event is designed for IT professionals to learn and get business-to-business matchmaking based on the current trends in Information/Network Security, BCP/DR Planning, Data Storage, and Enterprise Communications.

## **Additional Readings:**

1. How New York City is going to Consolidate 50 Data Centers from 40 City Agencies into One Location:

<http://www.informationweek.com/news/government/state-local/229219575>

2. NASA uses Amazon's cloud computing in Mars landing mission

<http://www.latimes.com/business/technology/la-fi-tn-amazon-nasa-mars-20120808,0,3551686.story>

3. The New York Public Library is Powered by Google Cloud

<http://www.nypl.org/collections/articles-databases/google-book-search>

4. Debunking the Myth of the [Single-Vendor Network](#)

<http://www.dell.com/downloads/global/products/pwcnt/en/Gartner-Debunking-the-Myth-of-the-Single-Vendor-Network-20101117-published.pdf>

## **Acknowledgement**

Thanks for MR2 Solutions, Inc.'s invitation to this event

(<http://www.mr2solutions.com>)

Thanks Dell for the courtesy of diagrams, which are used in my notes.