
TOMORROW
starts here.

Cisco UCS Network Performance Optimisation and Best Practices for VMware

BRKCOM-2015

Chris Dunk

Technical Marketing Engineer, Cisco UCS

#clmel

Cisco live!

Agenda

- Server to Server East West Traffic Flow Architecture
  - Why it is important
  - UCS Flat Network Topology Comparing to Legacy
- Overview of UCS B series Connectivity Flexibility
- Latency Test Configuration and Results
- VM Migration Configurations Test Results for East-West Traffic Flow
- Best Practices and Observations to Achieve the Greatest Impact on East-West Traffic

East-West Traffic Architecture

Cisco live!

East-West Traffic Flow, Why Do You Care?

- East-West Data Traffic flow represents increasing amount of data centre LAN traffic, up to 80%
  - Growth of virtualisation
  - New requirements for back end communication networks, Storage, Big Data, Cloud
- Traffic between servers inside the data centre
- Application Architectures are changing – distributed API
- Traffic between servers and storage inside the data centre
  - Storage synchronisation, cluster file systems
  - Data centre storage virtualisation
- Most common use cases are virtualisation live migration
  - VMware vMotion
  - MSFT Hyper-V Live Migration

East-West Traffic Flow Requirements

- Provide low latency server to server communication
  - Need for a low latency predictable fabric
  - Reduce the number of hops between servers
- Highly available active/active data paths
- Large Layer 2 flat network design
- Host bandwidth flexibility and control

Legacy Single Chassis East-West Traffic

- Single Chassis IS the best case scenario
- Traffic between servers has to go through a L2 Switch inside the chassis
- Traffic does not traverse the midplane

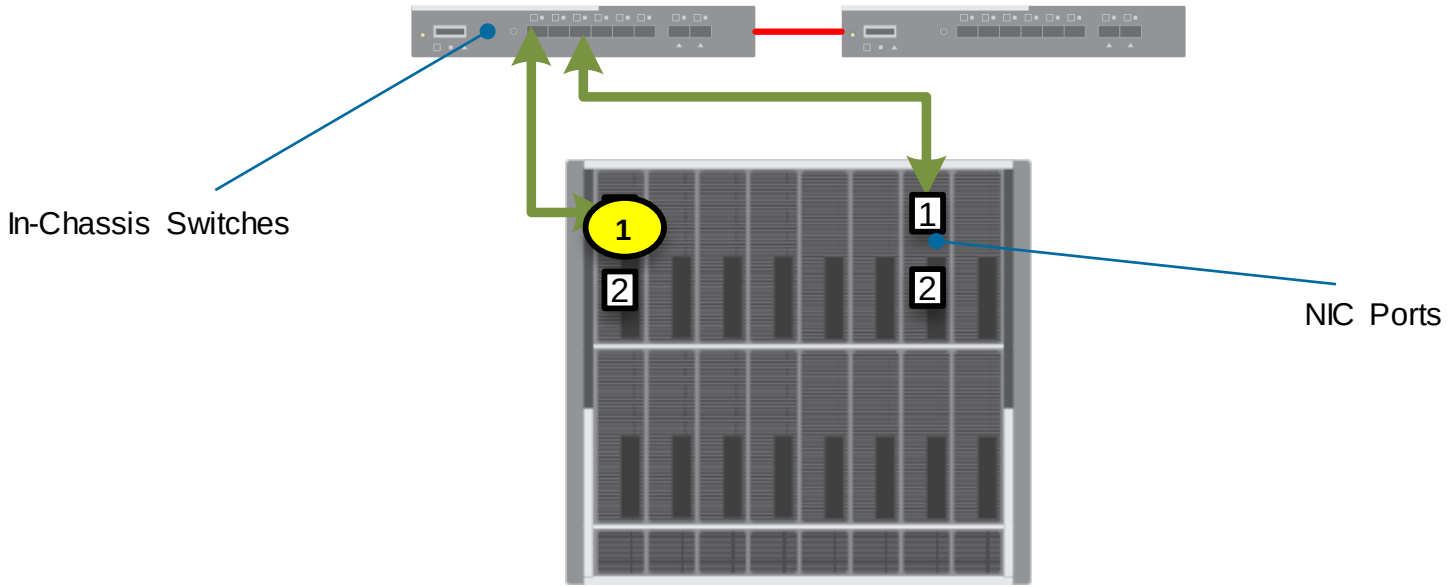

Legacy Multi Chassis East-West Traffic

- Common Data Centre Example with multiple Chassis connected to ToR Switch
- Traffic blade to blade will go through multiple switches

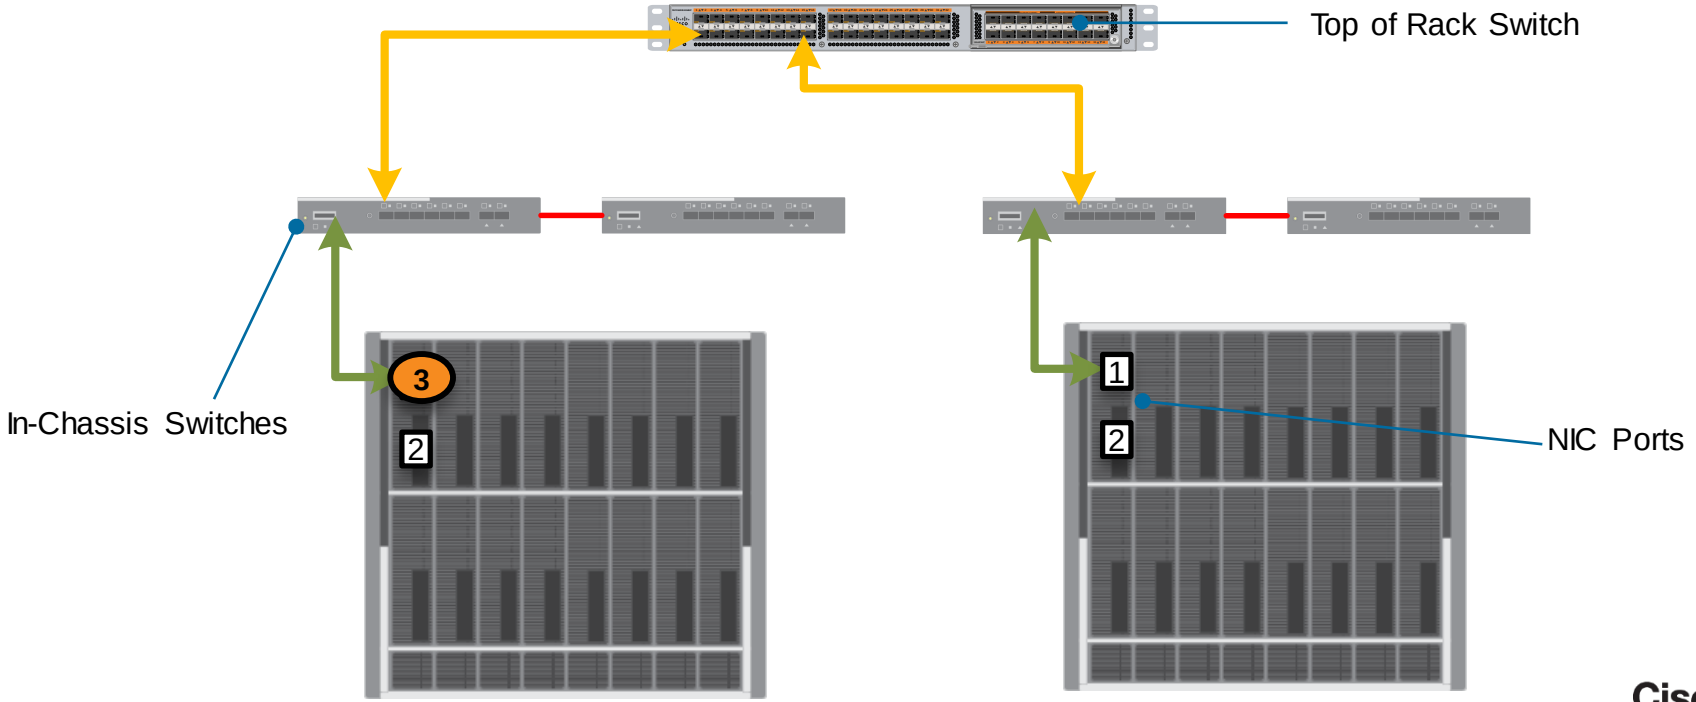

Legacy Multi Chassis Stacking East-West Traffic

- Hits multiple switches even with stacking
- As more stacking links are added latency and hops increase
- Hops become non-deterministic and Spanning Tree is required

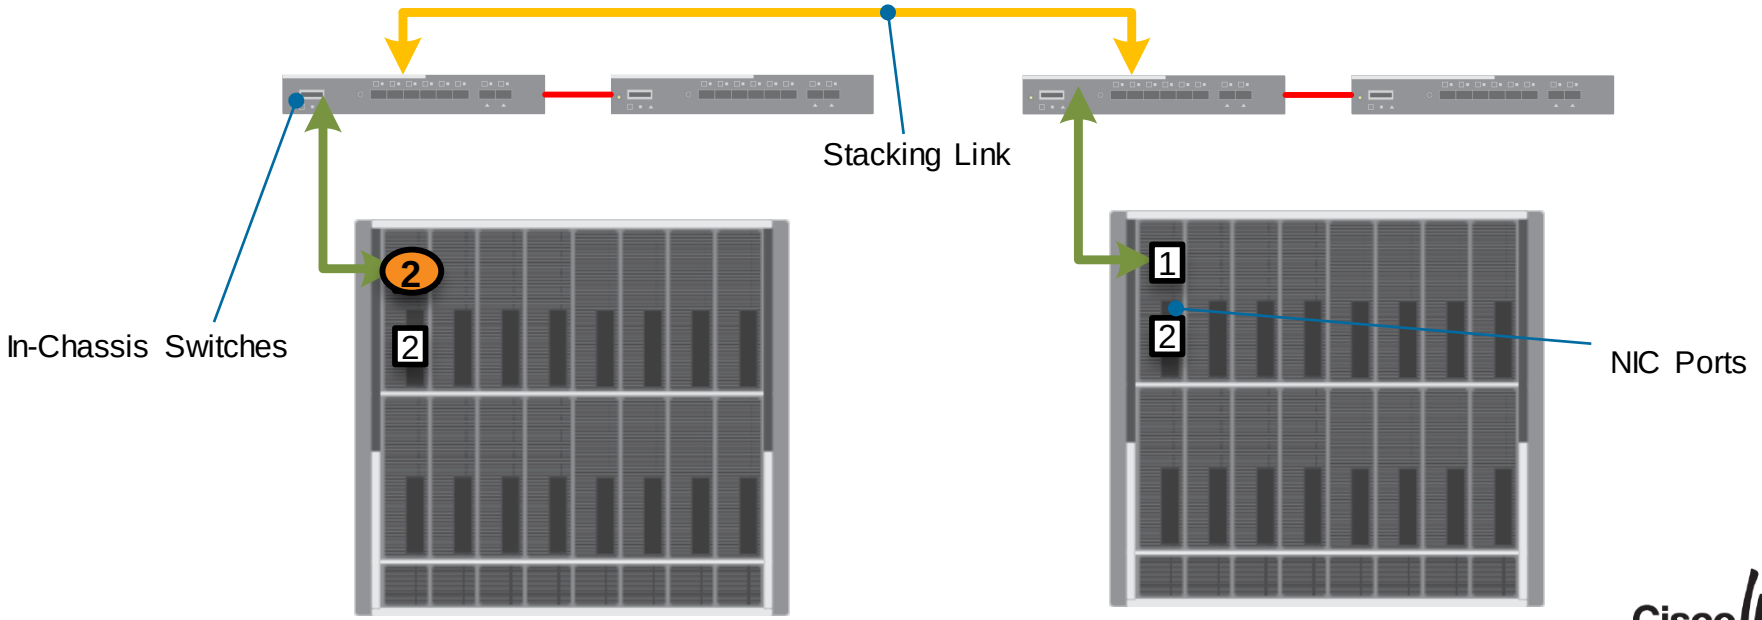

Legacy East-West Traffic Truths

- Server to Server communication optimised for only a SINGLE chassis of blades
- Chassis IO modules ARE switches (except for B22 and pass-thru modules)
- Requires many chassis switches that must be managed
- Blade Chassis to Chassis communication requires at least 2 or more hops
  - Including switch stacking
  - Increased latency
  - Non deterministic
- Customer data centres are larger than 1 chassis; must deal with this issue
- Intra chassis communication represents a small percentage of actual data flow

UCS Advantage Summary

- Cisco UCS offers a single flat L2 compute network for systems up to 160 servers
- Cisco UCS offers ‘right sizing’ the networking providing bandwidth where it’s needed, when it’s needed
- Simple and flexible, logical configuration of blade I/O connectivity managed with UCS Manager

UCS Advantage Fabric Topology

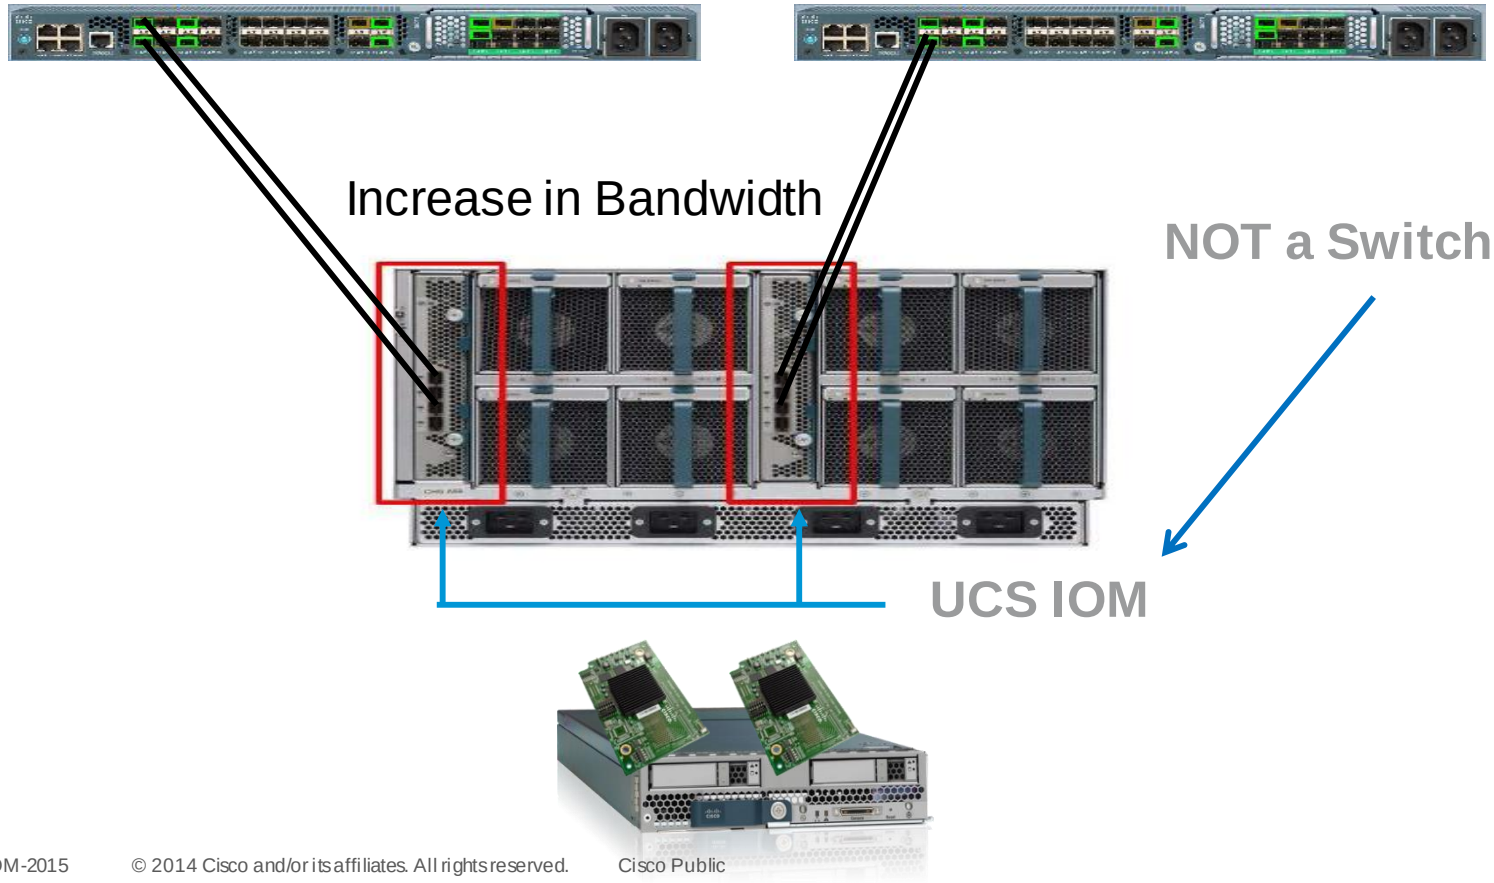

UCS Optimised Flat East-West Traffic

- Cisco UCS changes the legacy constraints
- Industries flattest L2 blade to blade network – 160 servers
- Single hop and predictable latency between ANY server in domain

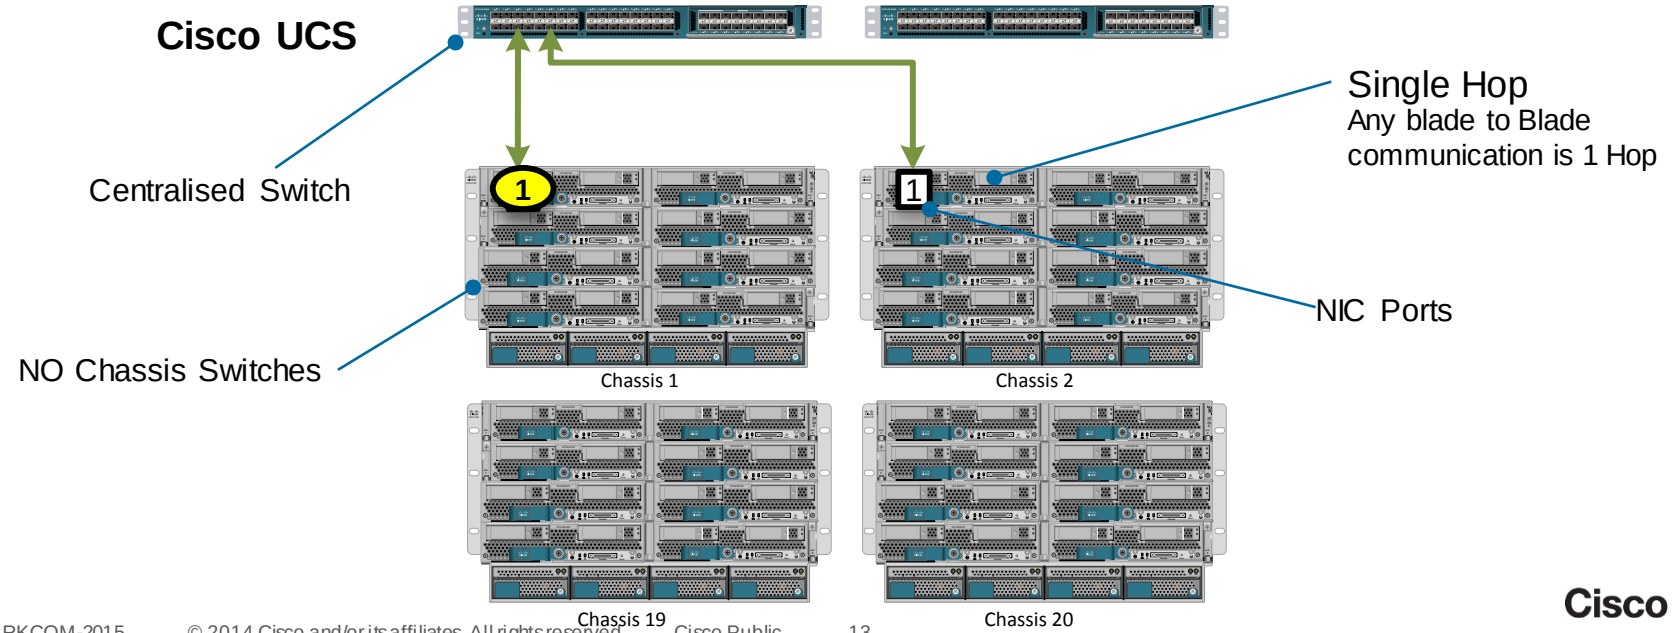

UCS Optimised Flat East-West Traffic

- Cisco UCS changes the legacy constraints
- Industries flattest L2 blade to blade network – 160 servers
- Single hop and predictable latency between ANY server in domain

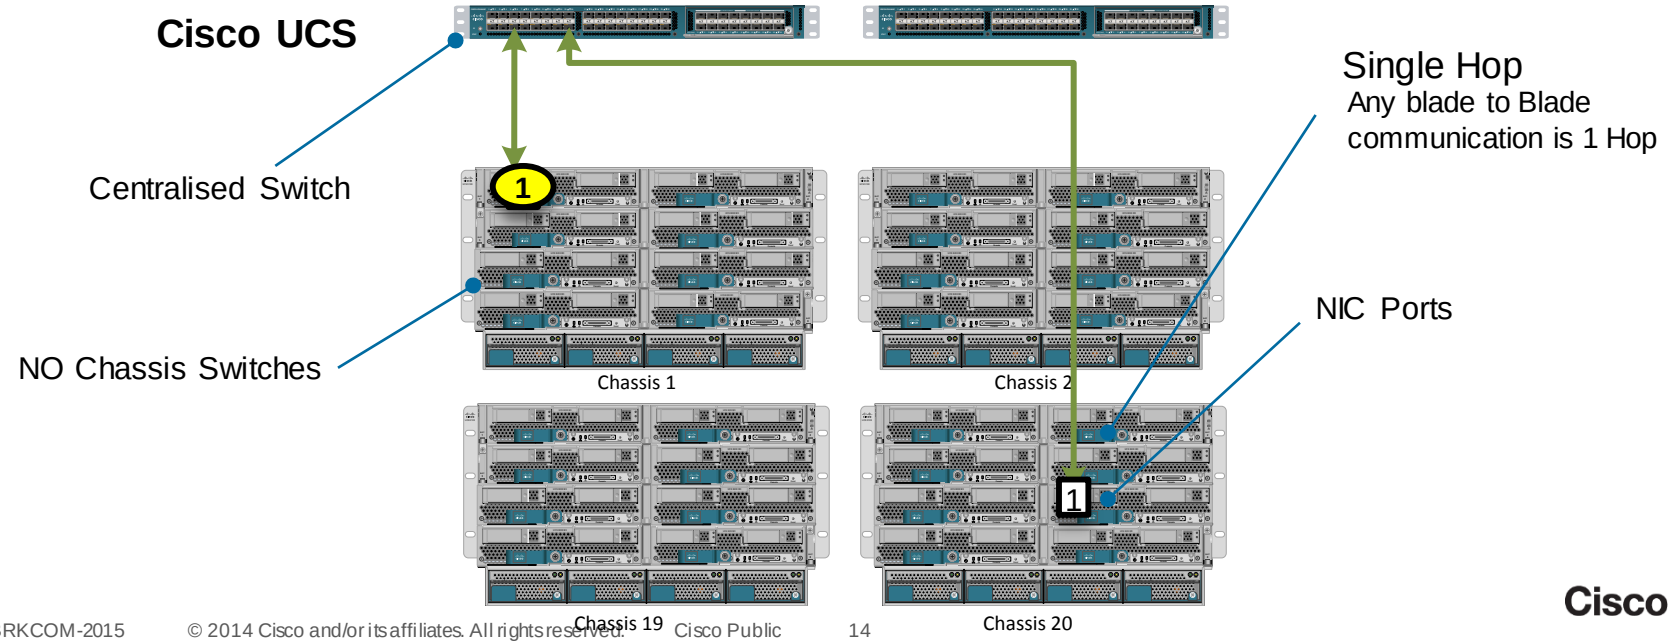

UCS East-West Traffic Truths

- Server to Server communication optimised for 160 blades across 20 chassis
- Chassis IO Modules are NOT switches
  - Fabric Extenders have no switching ASICs
  - Extension of the Fabric Interconnect

- Blade Chassis to Chassis communication requires only 1 hop
  - Lower system latency
  - Latency and performance predictable across entire domain

- NO Spanning Tree
  - Legacy stacking technologies require spanning-tree
  - Most legacy switches are non configurable

Cisco B-Series Networking Flexibility

Cisco live!

UCS B-Series Network Connectivity

- UCS B200 M3/M4 Half Width blades
  - Aggregate bandwidth up to **80Gb**
  - (1) mLOM connector
  - (1) Mezzanine slot
- UCS B260/420/440 Full Width blades
  - Aggregate bandwidth up to **160Gb**
  - (1) mLOM connector
  - (2) Mezzanine slot
- UCS B460 Full Width Double Slot Blade
  - Whopping **320Gb** of aggregate bandwidth
  - (6) Mezz slots, 2 for VIC 1240 and 4 additional slots

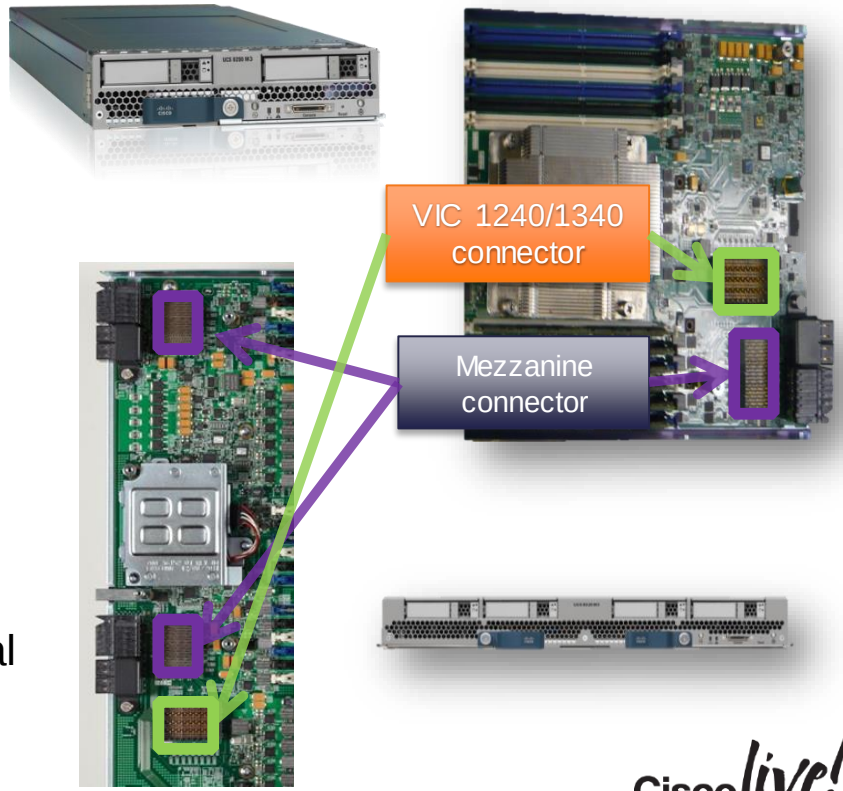

UCS Half-Height Blade Networking Options

- Cisco UCS Fabric Extender (FEX)
  - 2204 IOM 10Gb
    - 4 Uplinks / 16 Downlinks
  - 2208 IOM 10Gb
    - 8 Uplinks / 32 Downlinks

- Cisco Virtual Interface Card (VIC)
  - VIC1240/1340 10Gb
    - Up to 4 10Gb Interfaces
    - Offers 40Gb of connectivity PER Blade
  - Port Expander adds 4 additional 10Gb ports
  - VIC1280/1380 10Gb
    - Up to 8 10Gb Interfaces
    - Offers 80Gb of connectivity PER Blade

- All without adding expensive and complicated chassis switches

UCS B Series Bandwidth Scaling

- UCS B200 M3/M4 connectivity with Virtual Interface Mezz Card
  - VIC1240/1340 ships standard with B200 on Mezz Slot
  - Provides 2 10GbE links when paired to IOM 2204
  - Provides 4 10GbE links when paired to IOM 2208
- Optional Port Expander for VIC1240/1340
  - Enables 4 additional 10GbE links
  - Passive Connector
  - Fits into Mezz slot B200
- Simple to scale bandwidth
  - NO switches in chassis
  - Just increase chassis links
  - Add expander or mezz in blades

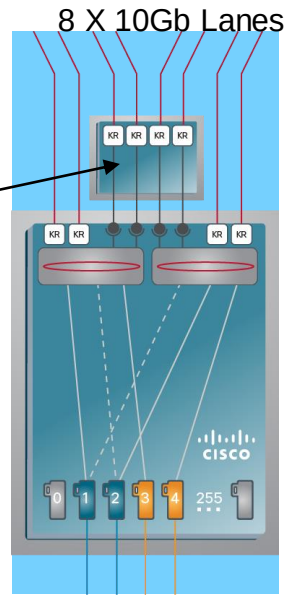

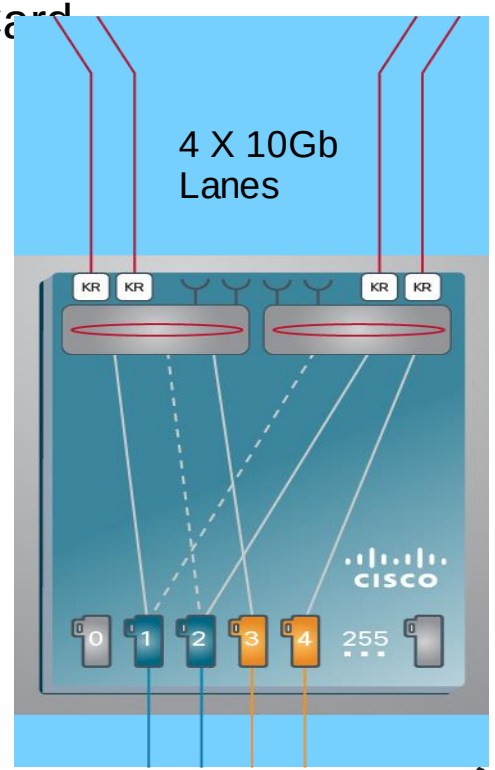

UCS B Series VIC Configuration

- VIC host ports are programmable by the Service Profile
  - OS Independent
  - No complicated BIOS configuration required
  - Place the vNICs on ANY fabric desired – not limited by Hardware placement
- Up to 256 vNIC/vHBA interfaces can be created
  - 8 reserved for internal use
  - OS limits apply
- VIC Fabric Failover (Cisco only Feature)
  - HW based NIC Teaming fail on fault
  - Enabled on each vNIC
  - Does not apply for vHBA

vNICs

Filter | Export | Print

Name	MAC Address	Desired Order	Actual Order	Fabric ID
vNIC Eth1	00:25:B5:00:A1:1E	1	1	A
vNIC Eth2	00:25:B5:00:A1:0E	2	2	A
vNIC Eth3	Derived	5	Unspecified	B A
vNIC Eth4	Derived	6	Unspecified	A B
vNIC Eth5	Derived	7	Unspecified	B A
vNIC Eth6	Derived	8	Unspecified	B

Block Diagram: Blade to FI Connection Test Setup

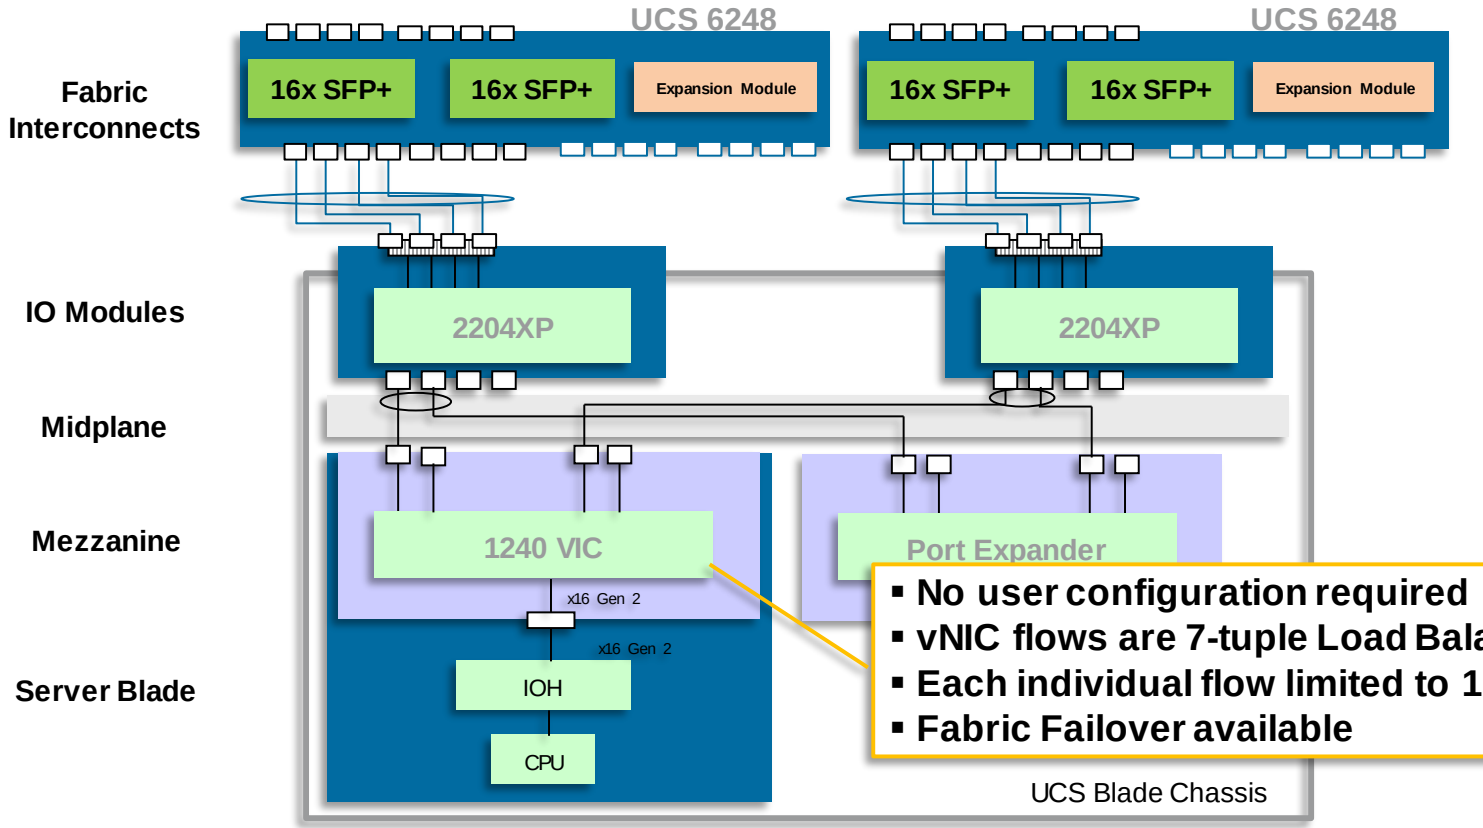

- No user configuration required
- vNIC flows are 7-tuple Load Balanced across links
- Each individual flow limited to 10Gb
- Fabric Failover available

Virtualised I/O

- Improved Performance and Visibility with Cisco FEX Architecture

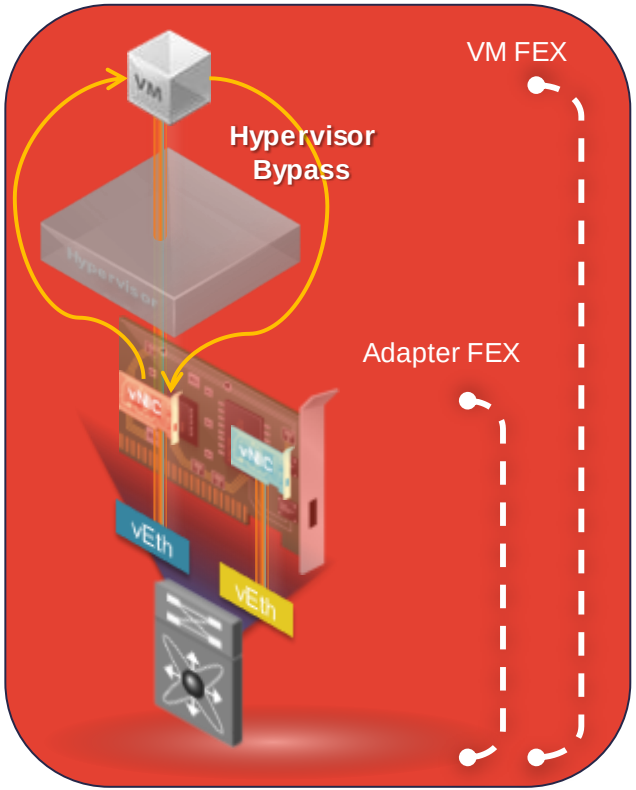

Features:

- Adapter FEX split a physical NIC into multiple logical NICs
- VM-FEX extends Adapter FEX technology to virtual machine
- Based on open Standards

Benefits:

- Same functionality in Physical and Virtual environments
- Increased bandwidth utilisation by sharing physical interface to multiple Applications and/or VMs
- Improved Performance and dynamic network & security policy mobility during VM migration with VM-FEX

Network Latency Test Setup and Results

Cisco live!

Latency Test Setup

- 2 X UCS B200 M3 with VIC1240
- 2 X 5108 Blade Chassis with 2204 Fabric Extenders
  - 1 link IOM to FI
- 1 10Gb VIC port configured on each host connected to fabric A
- Chassis to Chassis Local network communication
  - Single Fabric (A) used on all tests unless noted

- Purpose of the test
  - Determine if UCS design has more host based latency than other fabrics
  - Answer claims that Cisco UCS has higher latency for applications

Latency Test Setup

- Server BIOS optimised for Performance
  - Max Performance memory and CPU set
- OS NIC parameters set for performance
  - RX/TX Coalescing Timer –OFF
    - Forces HW to be used to measure latency
  - Interrupts localised to CPU Socket
- NetPerf used for host to host packet generation
- Ran 3 Latency measurement tests
  - UDP, TCP, TCP RTT
- 9K MTU shown for all test results
  - 1500 MTU was also tested but not shown

Test Architecture Overview

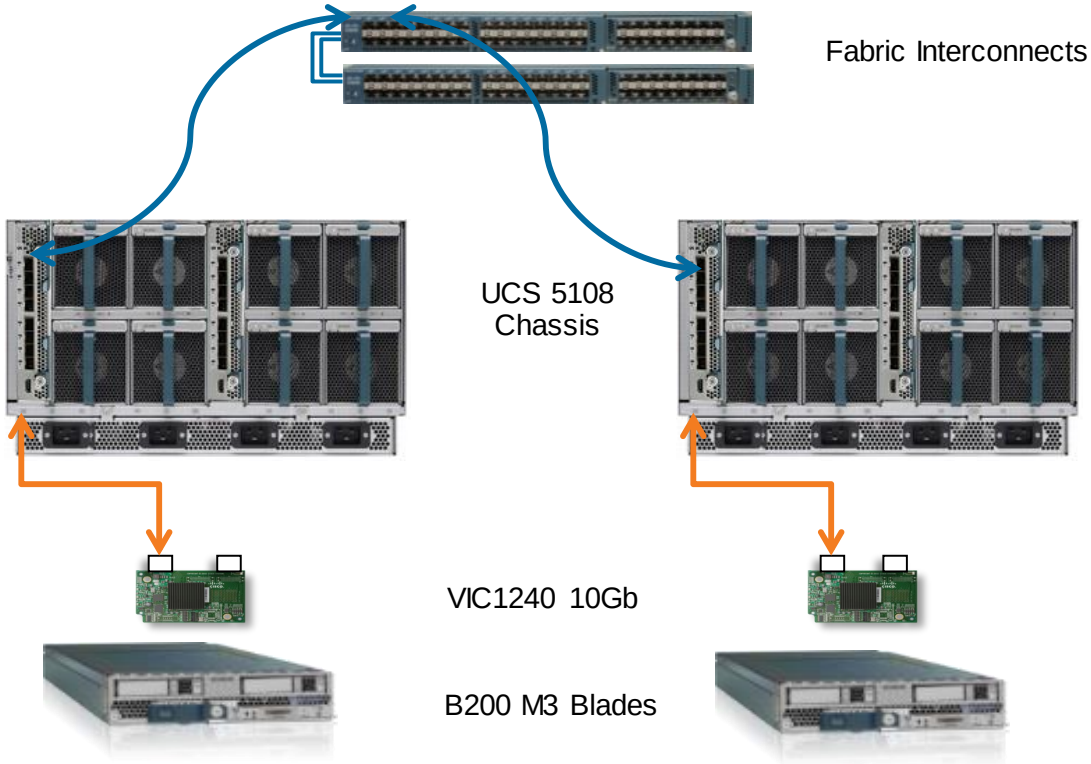

UDP Latency Test

- UDP Latency
  - UDP connectionless, reliable delivery no required
  - Fire and Forget
- Application Examples
  - Streaming audio and video
  - VoIP
  - Heartbeats
  - NFS
- Incrementally increased payload size from 1 byte to 1024 bytes
- Cisco UCS Chassis to Chassis results
  - Legacy both single chassis and chassis to chassis represented

UDP Latency Test Results Chassis to Chassis

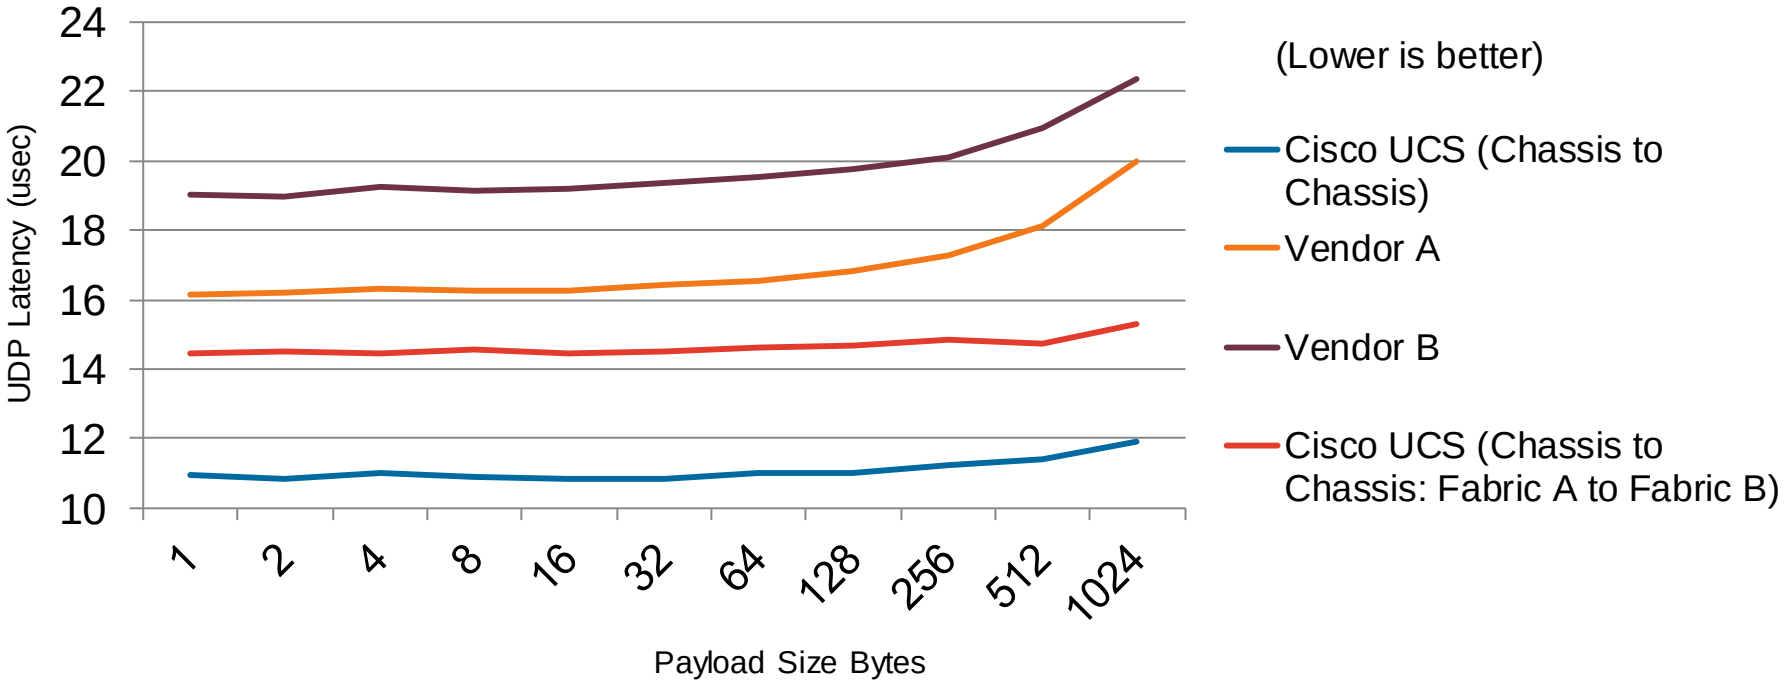

Source: http://www.cisco.com/en/US/prod/collateral/ps10265/le_40202_ibmlatencypb-130717.pdf; http://www.cisco.com/en/US/prod/collateral/ps10265/le_40201_pb_hplatency_130717.pdf; http://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/network_latency.pdf

Cisco live!

TCP Latency Test Results Chassis to Chassis

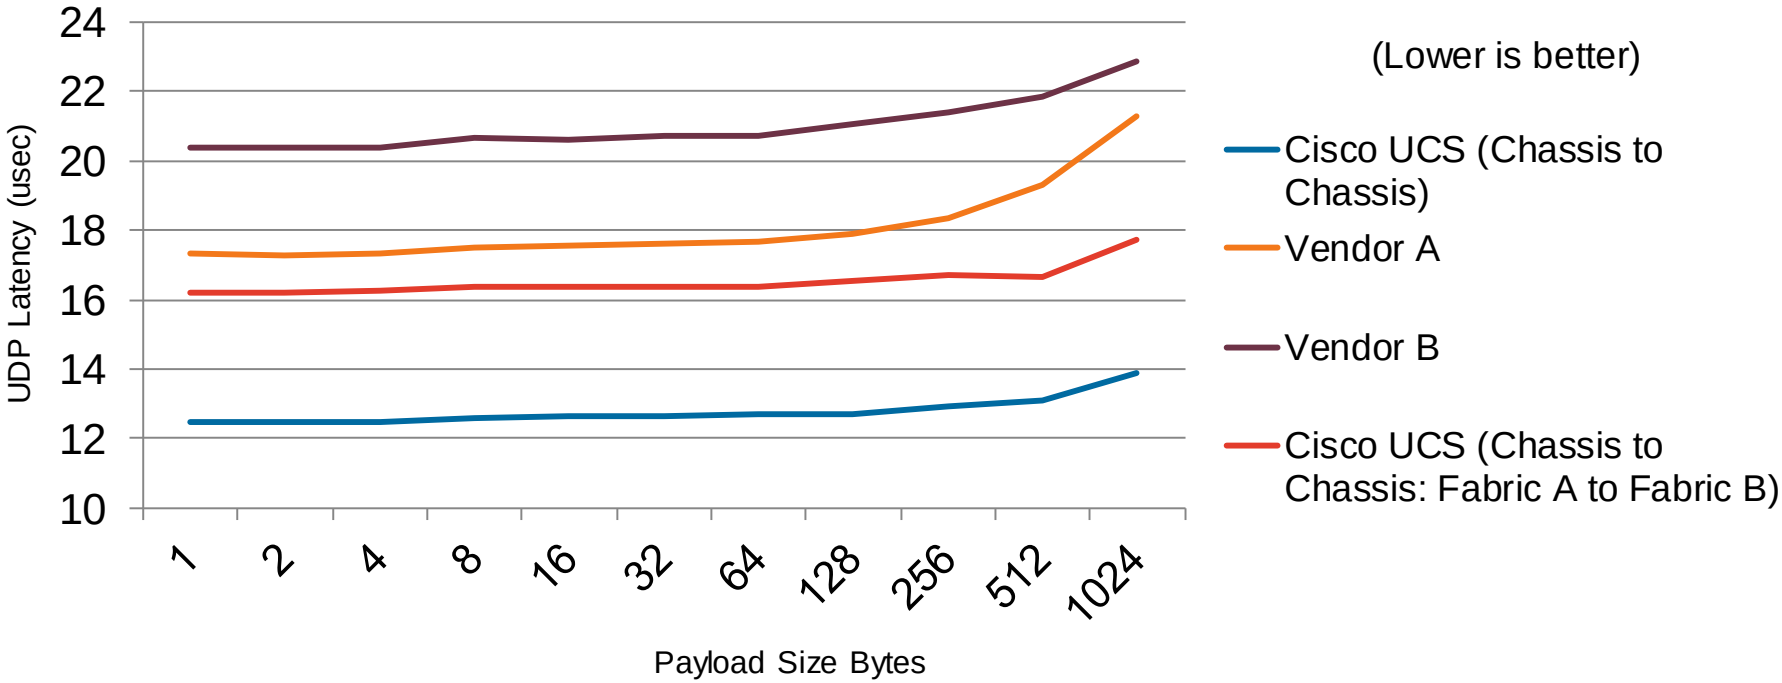

Source: http://www.cisco.com/en/US/prod/collateral/ps10265/le_40202_ibmlatencypb-130717.pdf; http://www.cisco.com/en/US/prod/collateral/ps10265/le_40201_pb_hplatency_130717.pdf; http://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/network_latency.pdf

TCP Transaction RTT Latency Test

- TCP Transaction RTT
  - Represents the TCP connection opening, sending, receiving, closing complete round trip time
  - Ping Pong example
  - Measuring Transactional Performance
- New connection for each request/response pair, transactional workload
- Mimics web server HTTP protocol
- Incrementally increased payload size from 1 byte to 16384 bytes
  - Large payload sizes will expose switch port buffer issues
- Cisco UCS Chassis to Chassis results
  - Legacy both single chassis and chassis to chassis represented

TCP RTT Latency Test Results Chassis to Chassis

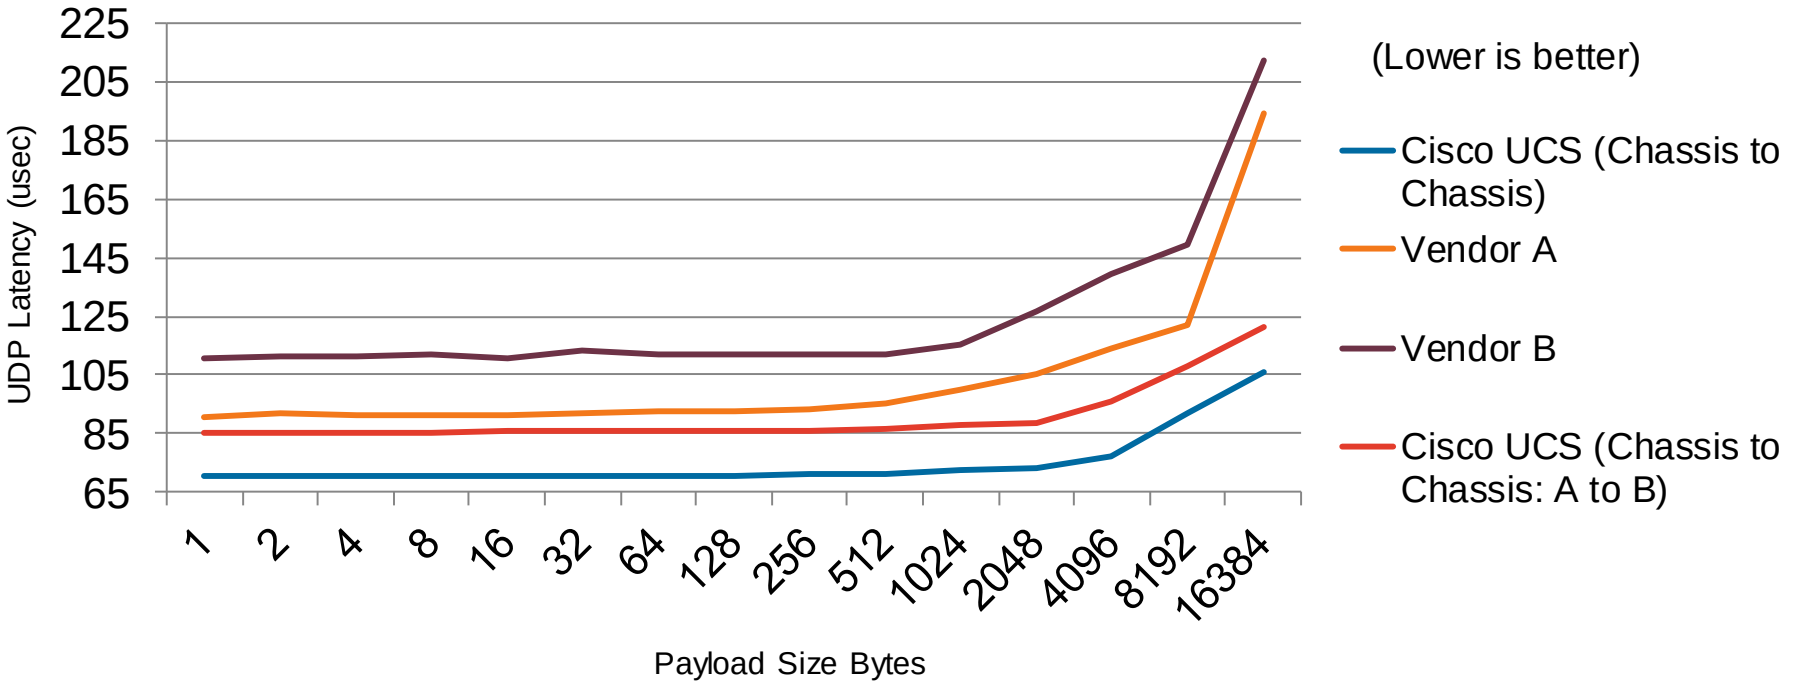

Source: http://www.cisco.com/en/US/prod/collateral/ps10265/le_40202_ibmlatencypb-130717.pdf; http://www.cisco.com/en/US/prod/collateral/ps10265/le_40201_pb_hplatency_130717.pdf; http://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/network_latency.pdf

Cisco live!

VM Migration Test Setup and Results

Cisco live!

VM Migration Test Setup

- B200 M3 with VIC1240 + Port Expander
  - 2 X 2680 Intel Processors
  - Host 64GB RAM
  - 1,2 or 4 Ethernet vNICs configured
- 2 X 5108 Blade Chassis with 2204 Fabric Extenders
  - 1, 2 and 4 links to FI tested
- Windows 2008R2 Guest VMs running Prime95 load tool
  - Each VM running 100% memory and processor capacity
  - ESXi Host memory fully utilised
- 2 different test stacks
  - 8 VMs with 8GB memory each
  - 4 VMs with 16GB memory each

Test Architecture Overview

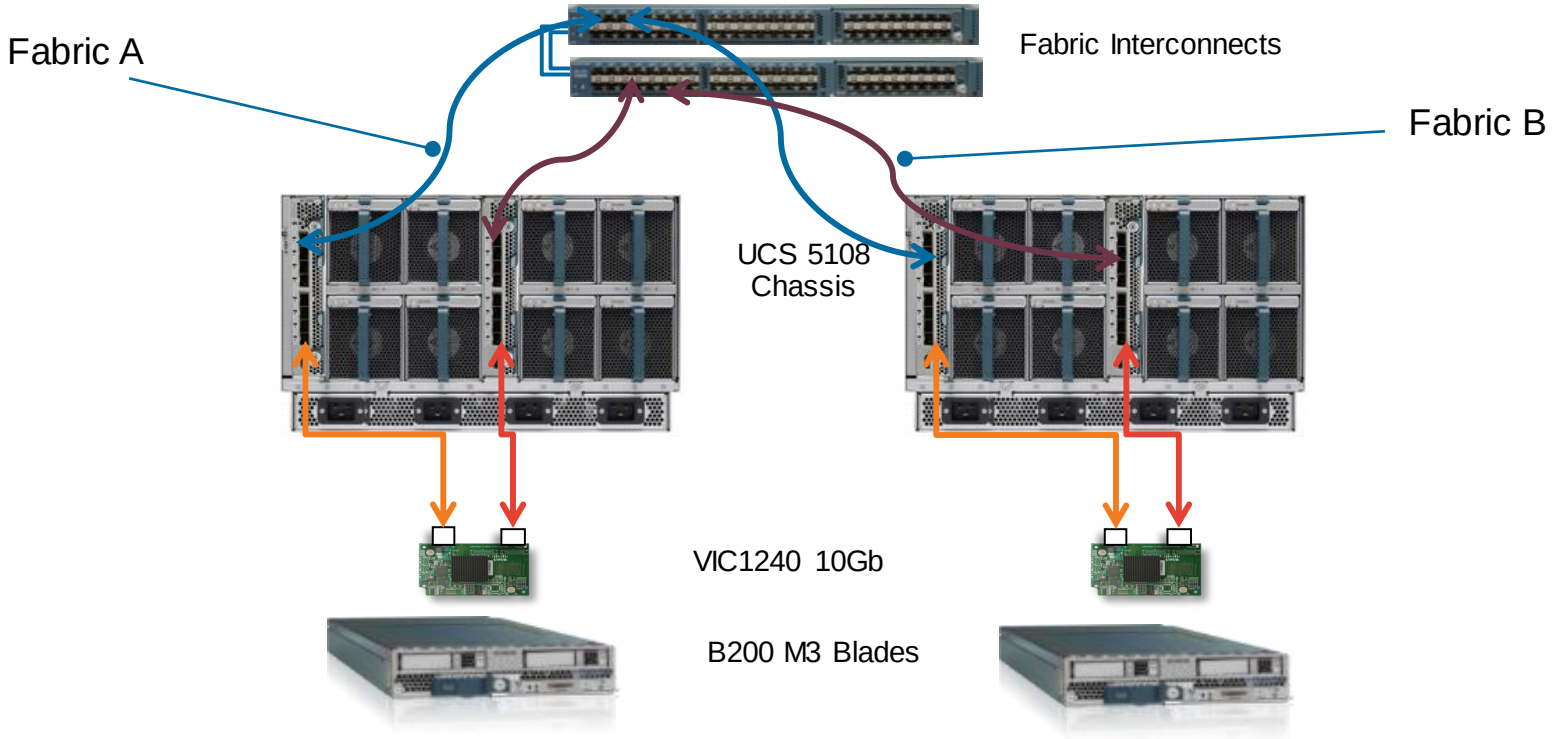

VM Migration Test Methodology

- Baseline Migration test overview
  - 2 IOM links to FI
  - 1vNIC configured and presented to VMware
  - Single vSwitch for vMotion
  - Traffic only on Fabric A for vMotion

- Why
  - Simple test setup that anyone could reproduce
    - No special benchmarking modifications or tuning
  - Mirror real world working scenarios
    - Customer inspired problem statement
    - Place busy hosts into maintenance mode and vacate VMs quickly
  - Find out the truth of claims being made against Cisco UCS network design

VM Migration Test Reliability

- How
  - VMs migrated between two blade nodes on two different chassis
  - PowerShell script to start migrations and capture times
  - Averaged times of all samples across multiple test runs
  - Pause delay between each move

- Quality
  - Obtain insight on system performance with multiple iterations of the tests
  - Each scripted test was run 5 times with 100 samples each (500 migrations minimum)
  - Nearly 30,000 samples captured across all tests
  - All samples used to calculate average migration time
    - Balances out vCenter variances
  - After each test run all hosts sat until idle and VMs rebooted

Multiple VMkernel Interfaces for vMotion

- VMware introduced multi-nic vMotion with ESX 5.0
  - using multiple 10Gb VMkernel ports designated for vMotion traffic
- Simple to configure
  - Add Physical NIC Ports to vSwitch
  - Add and configure 1st vMotion VMkernel portgroup
    - Configured NIC Teaming select 1st vmnic to active and 2nd vmnic to standby
  - Add and configure 2nd vMotion VMkernel portgroup
    - Configured NIC Teaming select 2nd vmnic to active and 1st vmnic to standby

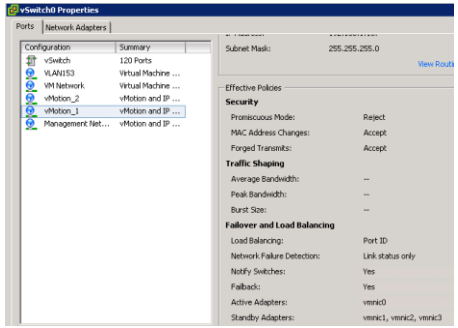

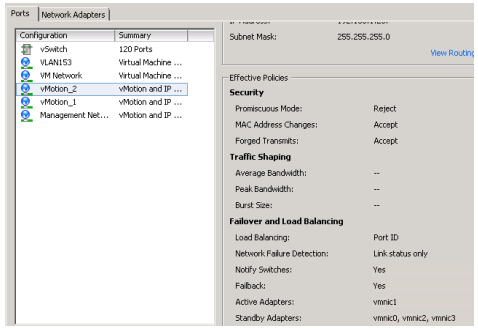

Multiple VMkernel Interfaces for vMotion

- Mark extra NICs as standby instead of unused
  - If you lose a physical NIC no network connection issue
  - Keeps the vmotion network stable
  - NIC failure traffic is routed to other connections

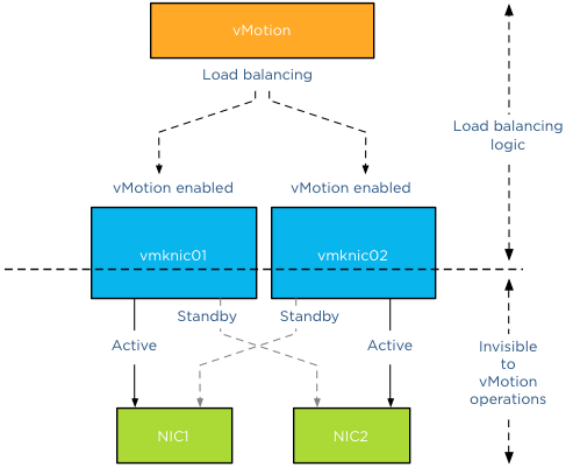

VMknic	Active NIC	Standby NIC
vmknic0	NIC1	NIC2, NIC3
vmknic1	NIC2	NIC1, NIC3
vmknic2	NIC3	NIC1, NIC2

vMotion Pairing Behaviour

- vMotion vmknic pairings are arbitrary setup by vCenter
  - No guarantee that any vmknic can communicate with any other specific vmknic
  - Placing on different subnets, likely to fail when pairings change
  - No preferred pairings of vmknic ports
  - Based on vmknic link speed and provision order (either of which may change depending on host behaviour)
  - Persistent only for a single vMotion
- ALL NICs in a cluster for vMotion must be able to communicate
  - Primary and secondary vmknics must be able to communicate on failure
- vMotion assigns a default interface on One vmknic for connection management
- VMware does not support dedicating vMotion on a specific fabric
  - No ability to prefer pairings of vmknics on a specific fabric

Multiple VMkernel Interfaces for vMotion

- VMware vMotion why is it not working the way I want?
- Do I put them on different subnets?
  - NO this is not supported by Vmware
  - This may work but not guaranteed, may fail at any time
- VMKernel NIC pairs are made first by
  - Highest speed
  - Order of creation on the ESXi host
- vMotion distributes across all available vmknics
  - even when a single VM is migrated multiple links are used

vMotion Behaviours

- Links speed matters
  - vMotion requires 10Gb links for 8 migrations
  - Mixing NIC speeds results in reduction of 4 migrations
  - Multi NIC vMotion still caps at 8 migrations
- vMotion operations
  - Work at the vmknic level not physical layer
  - No knowledge of pnic status
  - Trusts connection from vmknic between source and destination
- Ensure successful configuration
  - Build the vNICs in the same order on all hosts
  - Make all connections identical (speed, MTU, fabric connection)
  - **Leverage UCSM LAN connection policies**

Sample Setup Scenarios

- UCS upstream 5K configured VLANs
- Separate VLAN and IP ranges
- All vMotion vmknic ports are pingable from each host
- Successful migrations with no errors
- Would not work if remote host lost primary NIC and standby was used with different subnet/VLAN

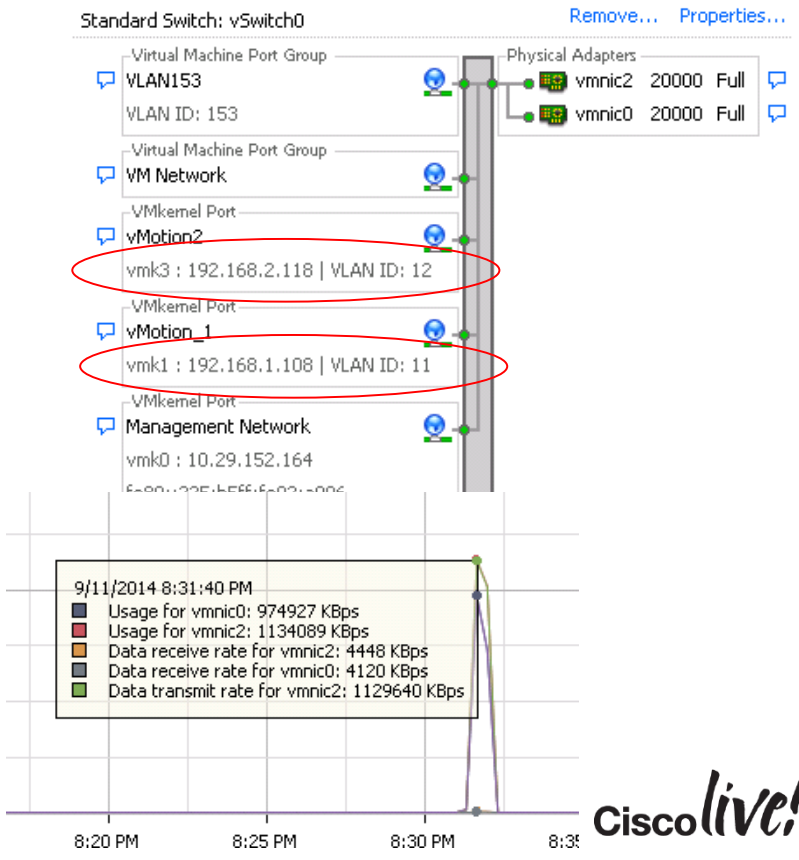

Sample Setup Scenarios

- UCS Private VLANs not configured on 5K
- Same VLAN different IP ranges
- All vMotion vmknic ports are pingable from each host
- Successful migrations with no errors
- Would not work if remote host lost primary NIC and standby was used with different subnet

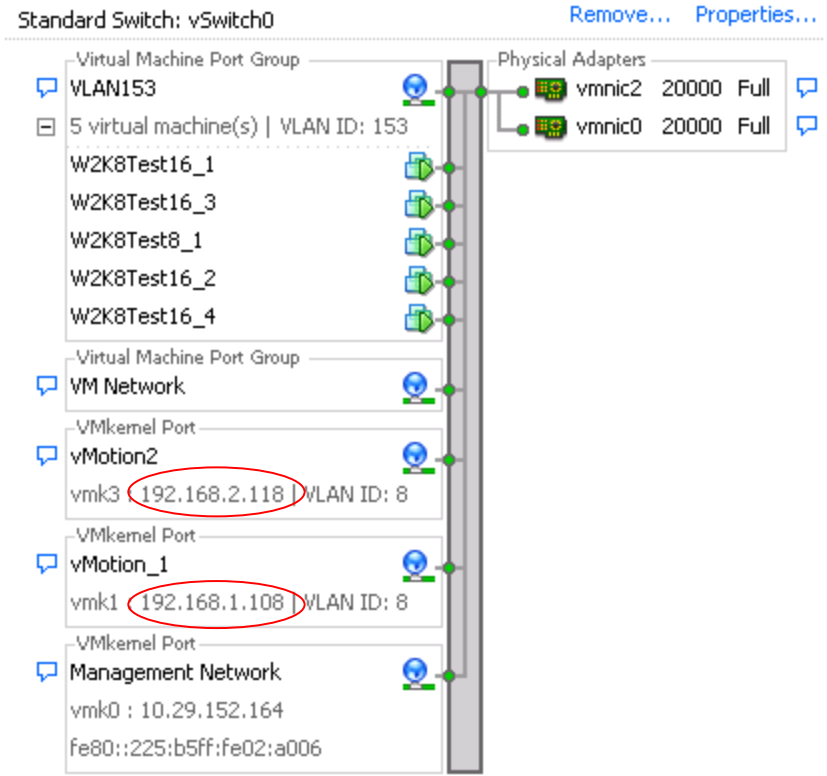

Sample Setup Scenarios

- UCS Private VLANs not configured on 5K
- Same Subnet, Same IP ranges
- Some vMotion vmknic ports are NOT pingable from each host
  - vMotion2 vmknic are not reachable on each host from each other
  - All other interfaces work (local and vMotion 1 vmknic)
- Vmknic pairings going over two different fabrics that are not connected
- Migrations do NOT Work

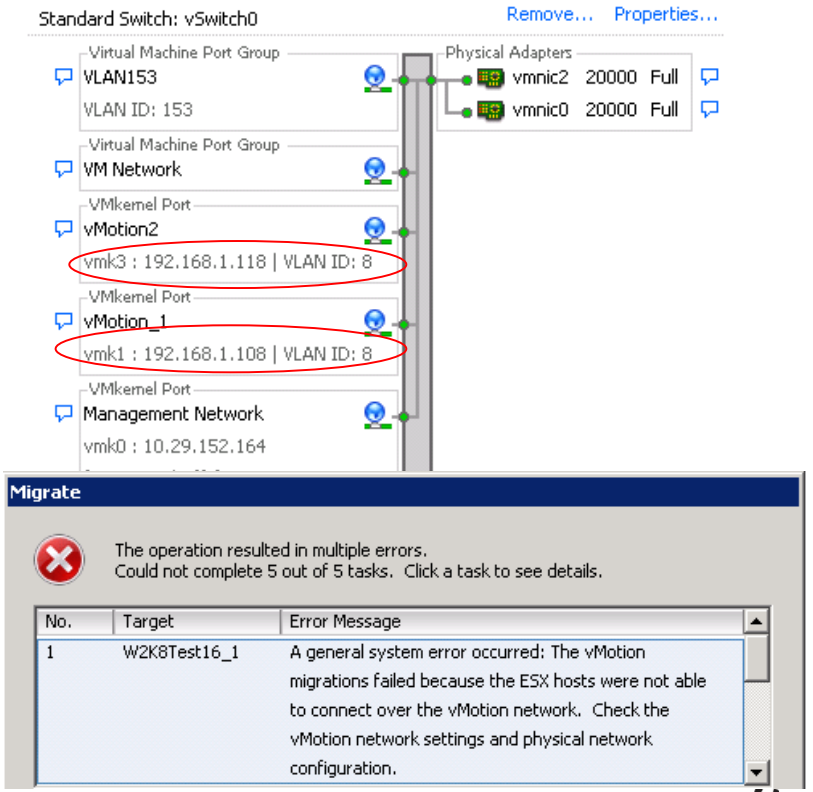

Demo Multi NIC Setup

East-West Traffic Test Objectives

- NOT a Performance benchmark of vMotion itself
- NOT a study to discover why vCenter behaves the way it does
- NOT a networking tuning benchmark effort
- Compare equally configured Legacy systems and Cisco UCS
- Performance impact on Cisco UCS with fabric tunings
  - UCS has many options and capabilities for traffic engineering
- Only observations made during testing outlined
  - Enable customers to choose configurations that provide the greatest impact to their environment

Migration Test Results Cisco and Legacy

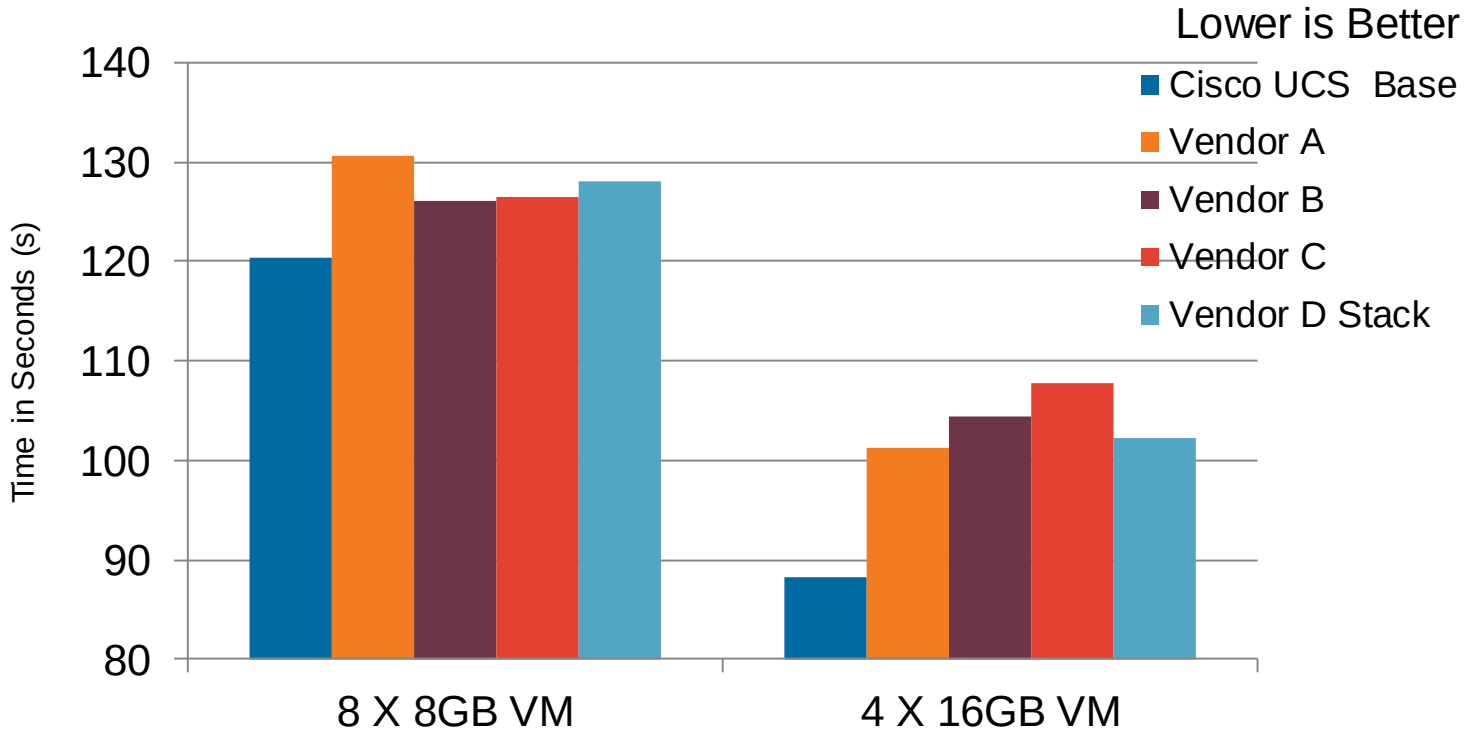

Source: http://www.cisco.com/en/US/prod/collateral/ps10265/le_40202_ibmlatencypb-130717.pdf; http://www.cisco.com/en/US/prod/collateral/ps10265/le_40201_pb_hplatency_130717.pdf; http://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-7000-series-switches/network_latency.pdf

Communication Performance Observations

- Viewing link performance during the migration tests observed
  - Real time performance between blades and uplink ports to the other side
  - Used built in measurement tools for other vendors switches to view real time data

- UCS with 1 Link between IOM and FI with 1 vNIC
  - **Averaged 9.5Gb/s**

- Other vendor intra-chassis switch performance
  - **Averaged 5.5 and 6 Gb/s**

Migration Base Configuration

- Base system using 2 IOM uplinks
- 1 vNIC configured with a single vSwitch dedicated for vMotion
- Traffic running over a single dedicated fabric

- **What if** we create two VMkernel ports on the same vSwitch?
  - All traffic on Fabric A
  - 2 VMkernel portgroups configured

- **What if** we configure Jumbo frames?

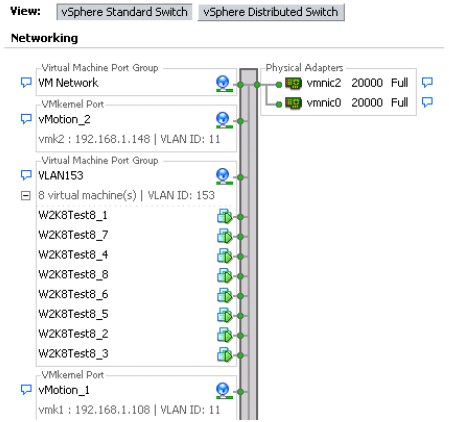

Migration Cisco UCS Test Base vs 2 vNIC

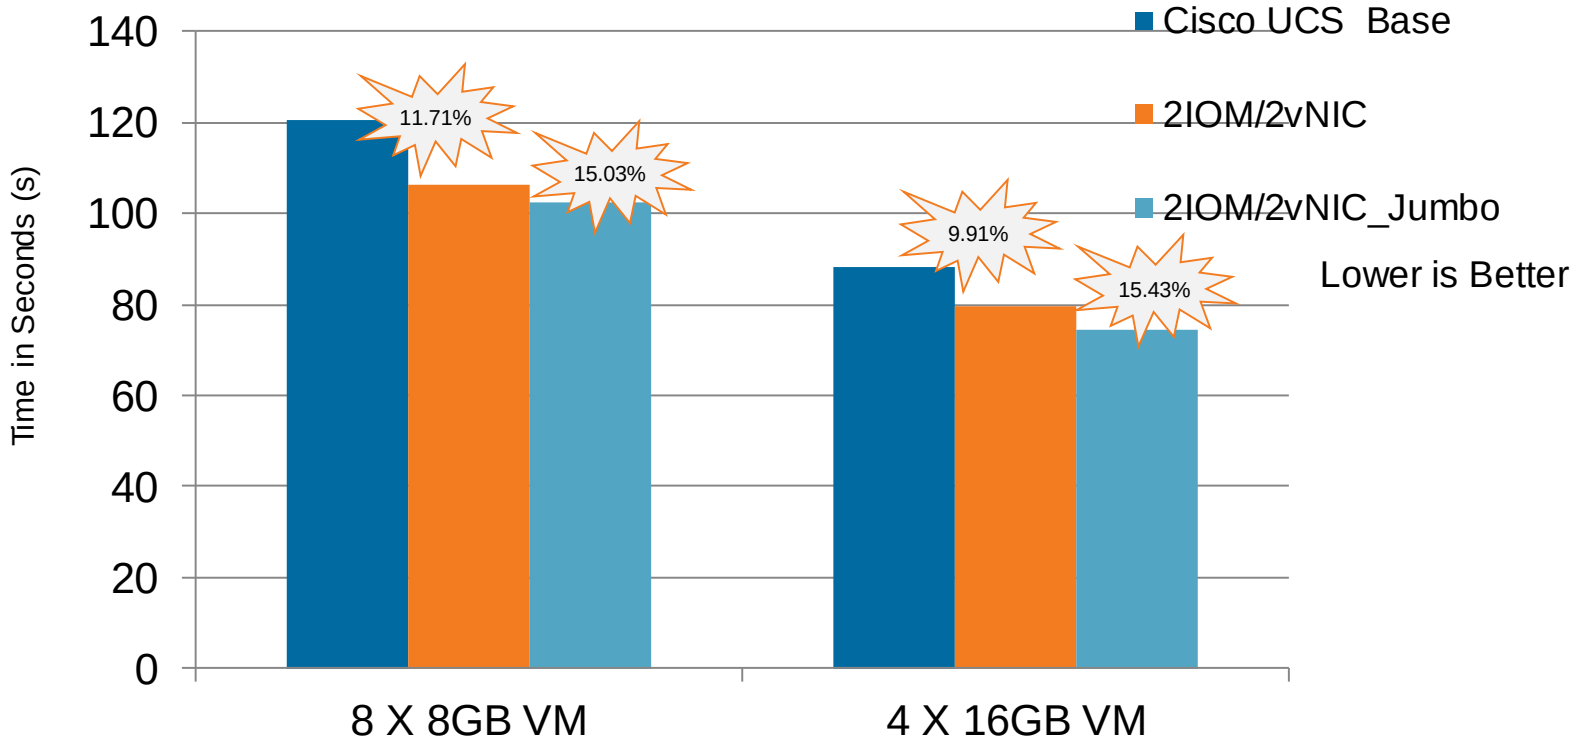

Source: Cisco Internal Testing

IOM Traffic Example

- Traffic is load balanced across all the available ports
  - VIC to IOM
  - IOM to FI

Single Fabric A

Average Bandwidth 13-14Gb/s

```
fex-1# show  platform software woodside rate
+--------++------------+-----------+------------++------------+-----------+------------+-------+-------+---+
| Port   || Tx Packets |  Tx Rate  |   Tx Bit   || Rx Packets |  Rx Rate  |   Rx Bit   |Avg Pkt|Avg Pkt|   |
|        ||            | (pkts/s)  |    Rate    ||            | (pkts/s)  |    Rate    | (Tx)  | (Rx)  |Err|
+--------++------------+-----------+------------++------------+-----------+------------+-------+-------+---+
| 0-BI   ||         25 |         5 |   4.52Kbps ||         21 |         4 |   4.18Kbps |    93 |   104 |   |
| 0-CI   ||         10 |         2 |   5.66Kbps ||         12 |         2 |  20.43Kbps |   334 |  1044 |   |
| 0-NI1  ||    2927298 |    585459 |   7.21Gbps ||     340801 |     68160 | 250.64Mbps |  1519 |   439 |   |
| 0-NI0  ||    2505339 |    501067 |   6.16Gbps ||     306456 |     61291 | 201.70Mbps |  1517 |   391 |   |
| 0-HI27 ||         17 |         3 |   5.52Kbps ||          0 |         0 |   0.00 bps |   183 |     0 |   |
| 0-HI26 ||          0 |         0 |   0.00 bps ||          1 |         0 | 288.00 bps |     0 |   164 |   |
| 0-HI19 ||     273030 |     54606 | 213.19Mbps ||    2116818 |    423363 |   5.20Gbps |   468 |  1515 |   |
| 0-HI18 ||     520805 |    104161 | 258.97Mbps ||    3316024 |    663204 |   8.17Gbps |   290 |  1520 |   |
+--------++------------+-----------+------------++------------+-----------+------------+-------+-------+---+
fex-1# show  platform software woodside rate
```


Cisco live!

Migration

- **What if** we configure more VMkernel ports (4) for vMotion?
  - All traffic on Fabric A
  - 4 VMkernel portgroups configured

- **What are** the results of Jumbo and non-Jumbo frames being enabled?

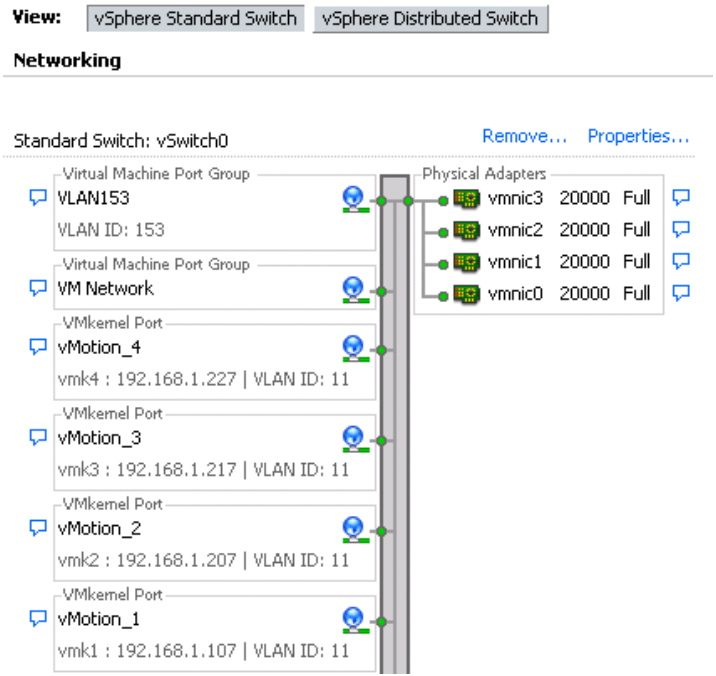

Migration Cisco UCS Test Base vs 4 vNIC

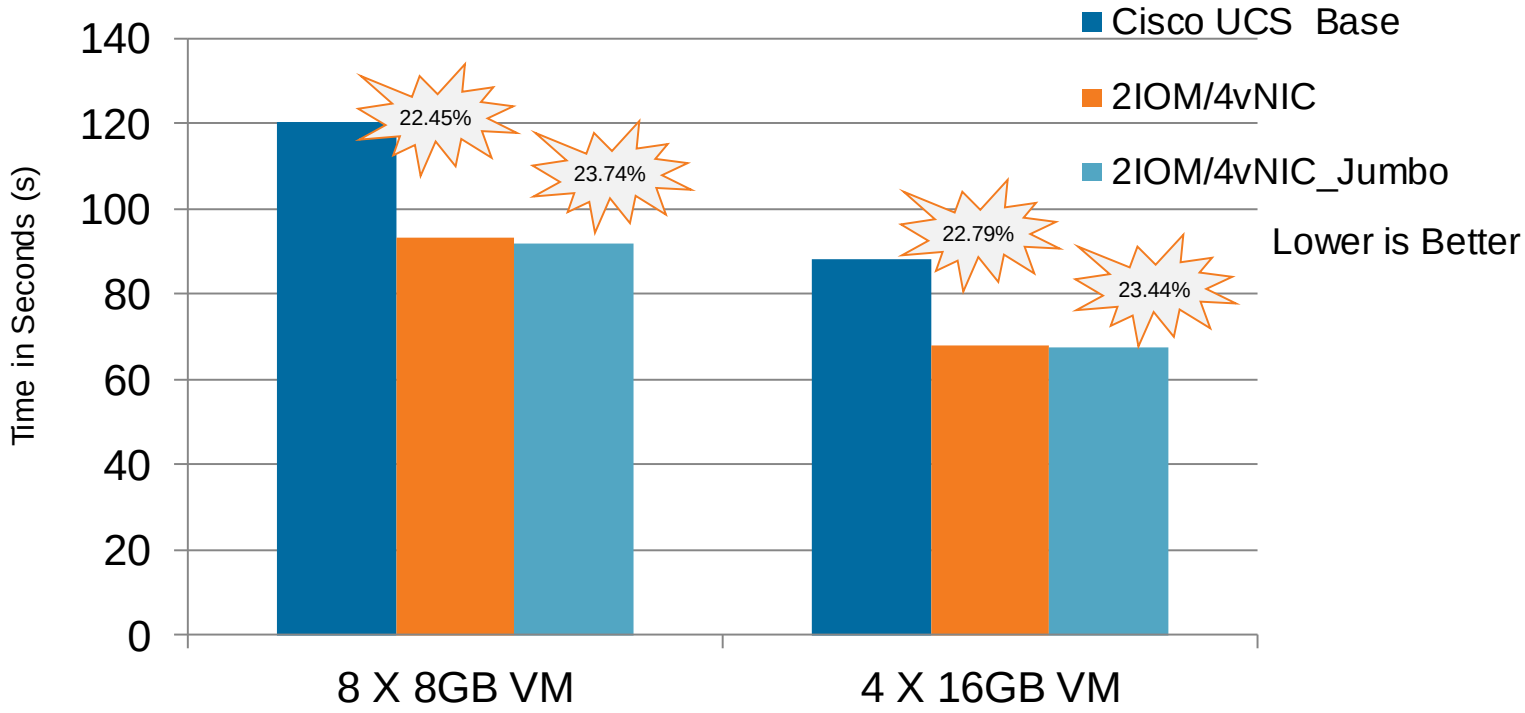

Source: Cisco Internal Testing

IOM Traffic Example 4vNIC

- Traffic is load balanced across all the available ports - note no change in host port connections
- vCenter can run 4 independent TCP Flows

Single Fabric A

Average Bandwidth up to 15Gb/s

```
fex-1# show platform software woodside rate
```

Port	Tx Packets	Tx Rate (pkts/s)	Tx Bit Rate	Rx Packets	Rx Rate (pkts/s)	Rx Bit Rate	Avg Pkt (Tx)	Avg Pkt (Rx)	Err
0-BI	13	2	2.75Kbps	9	1	1.74Kbps	112	101	
0-CI	14	2	3.96Kbps	15	3	11.13Kbps	157	444	
0-NI1	238325	47665	38.28Mbps	2740896	548179	6.78Gbps	80	1527	
0-NI0	339500	67900	54.54Mbps	3507941	701588	8.68Gbps	80	1527	
0-HI27	3	0	576.00 bps	0	0	0.00 bps	100	0	
0-HI19	2971312	594262	7.35Gbps	193220	38644	31.05Mbps	1527	80	
0-HI18	3275421	655084	8.11Gbps	384595	76919	61.76Mbps	1527	80	

Cisco live!

Migration

- **What if** we increase the IOM links from 2 to 4 and increase the number of VMkernel ports to 4?
- **What if** we configure jumbo frames?

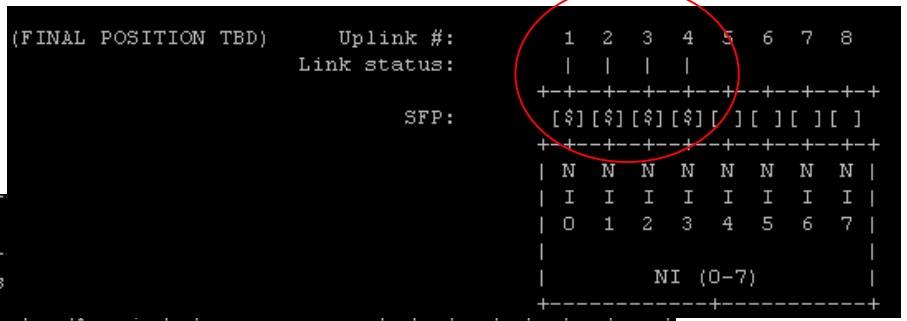


```
fex-1# show platform software woodside rate
```

Port	Tx Packets	Tx Rate (pkts/s)	Tx Bit Rate	Rx Packets					
O-BI	14	2	2.90Kbps	10	2	1.92Kbps	109	100	
O-CI	13	2	9.48Kbps	20	4	34.43Kbps	436	1056	
O-NI3	99257	19851	16.05Mbps	1020626	204125	2.53Gbps	81	1532	
O-NI2	320113	64022	345.97Mbps	1939933	387986	4.81Gbps	655	1530	
O-NI1	228915	45783	227.40Mbps	1744105	348821	4.33Gbps	600	1533	
O-NI0	118380	23676	19.15Mbps	1289526	257905	3.19Gbps	81	1529	
O-HI27	12	2	3.36Kbps	0	0	0.00 bps	155	0	
O-HI26	1	0	480.00 bps	0	0	0.00 bps	283	0	
O-HI19	3060155	612031	7.59Gbps	367310	73462	313.20Mbps	1530	512	
O-HI18	2915766	583153	7.24Gbps	399383	79876	295.35Mbps	1533	442	

Migration Cisco UCS Test Base vs 4 vNIC_4IOM

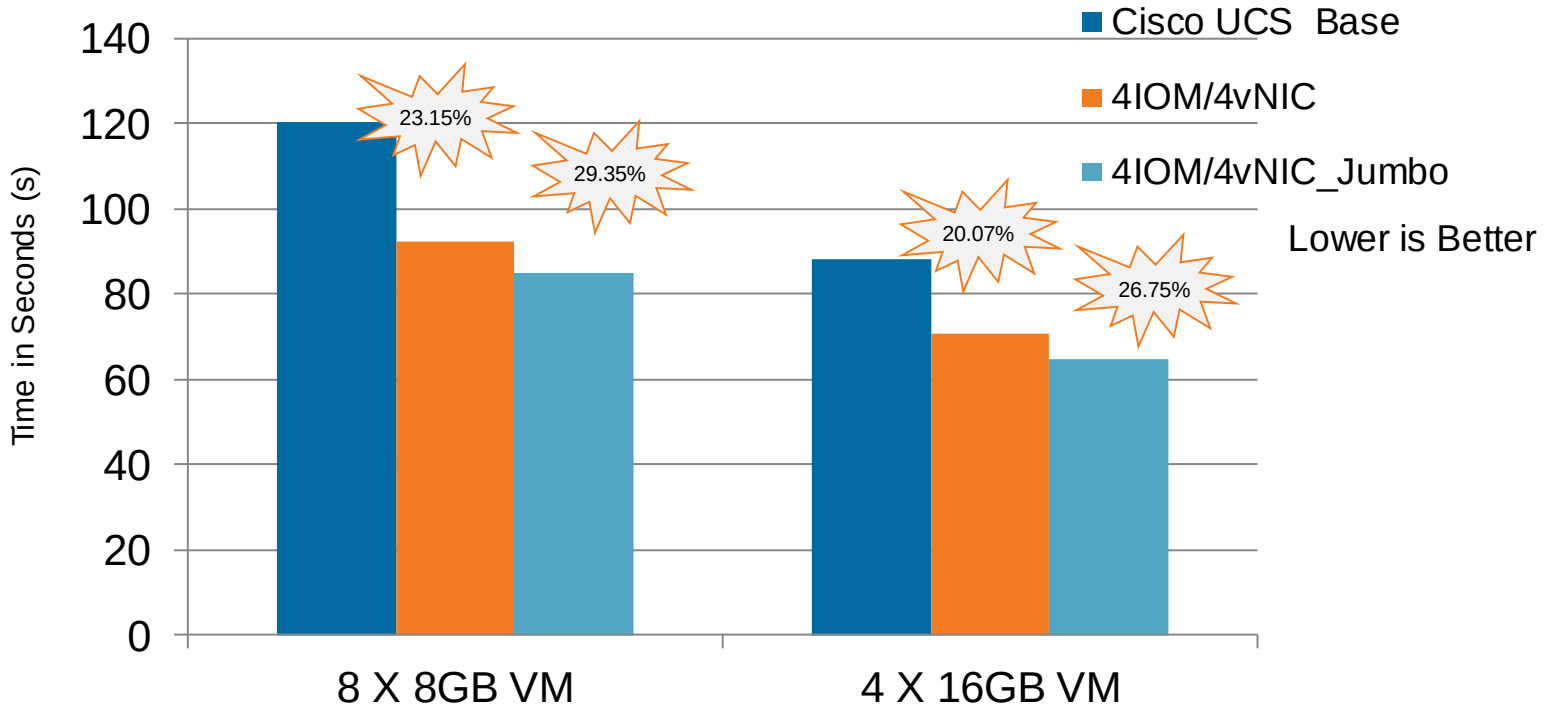

Source: Cisco Internal Testing

Migration

- What happens now if we run the same tests but spread the VMkernel ports over both UCS Fabrics?
  - Configured both 2 and 4 VMkernel portgroups
  - Spread the Physical NIC ports over UCS Fabric A and Fabric B

- Measure using 2 IOM to FI links then 4 IOM to FI links
- Enable Jumbo frames, what happens?

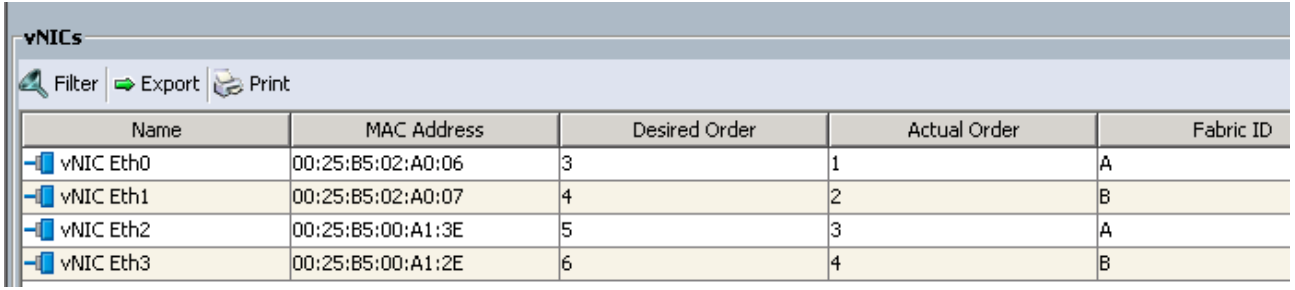

Name	MAC Address	Desired Order	Actual Order	Fabric ID
vNIC Eth0	00:25:B5:02:A0:06	3	1	A
vNIC Eth1	00:25:B5:02:A0:07	4	2	B
vNIC Eth2	00:25:B5:00:A1:3E	5	3	A
vNIC Eth3	00:25:B5:00:A1:2E	6	4	B

IOM Dual Fabric Traffic Example

- Traffic is load balanced across all the available ports
- Multiple VMKernel vNICs are configured assigned to Fabric A and Fabric B

10.29.152.97 - PuTTY

Port	Tx Packets	Tx Rate (pkts/s)	Tx Bit Rate	Rx Packets	Rx Rate (pkts/s)	Rx Bit Rate
0-BI	65	13	11.12Kbps	57	11	19.60Kbps
0-CI	19	3	7.95Kbps	17	3	12.36Kbps
0-NI1	13	2	8.54Kbps	1949040	389808	4.82Gbps
0-NI0	244	48	198.25Kbps	2406090	481218	5.95Gbps
0-HI27	0	0	0.00 bps	1	0	288.00 bps
0-HI26	1	0	152.00 bps	1	0	352.00 bps
0-HI19	1941444	388288	4.80Gbps	3125	625	565.10Kbps
0-HI18	2412413	482482	5.97Gbps	2949	589	401.62Kbps

Fabric A

Bandwidth Average combined fabrics 18Gb/s

Fabric B

10.29.152.96 - PuTTY

		(pkts/s)	Rate		(pkts/s)	Rate
0-BI	19	3	3.97Kbps	12	2	4.24Kbps
0-CI	12	2	3.76Kbps	17	3	13.43Kbps
0-NI1	563691	112738	90.95Mbps	2492069	498413	6.17Gbps
0-NI0	292840	58568	47.32Mbps	1782000	356400	4.41Gbps
0-HI27	8	1	2.89Kbps	0	0	0.00 bps
0-HI26	1	0	632.00 bps	0	0	0.00 bps
0-HI19	2996389	599277	7.42Gbps	417501	83500	67.32Mbps
0-HI18	1276838	255367	3.16Gbps	443908	88781	71.60Mbps

Migration Cisco UCS Test Base vs Multiple Fabrics

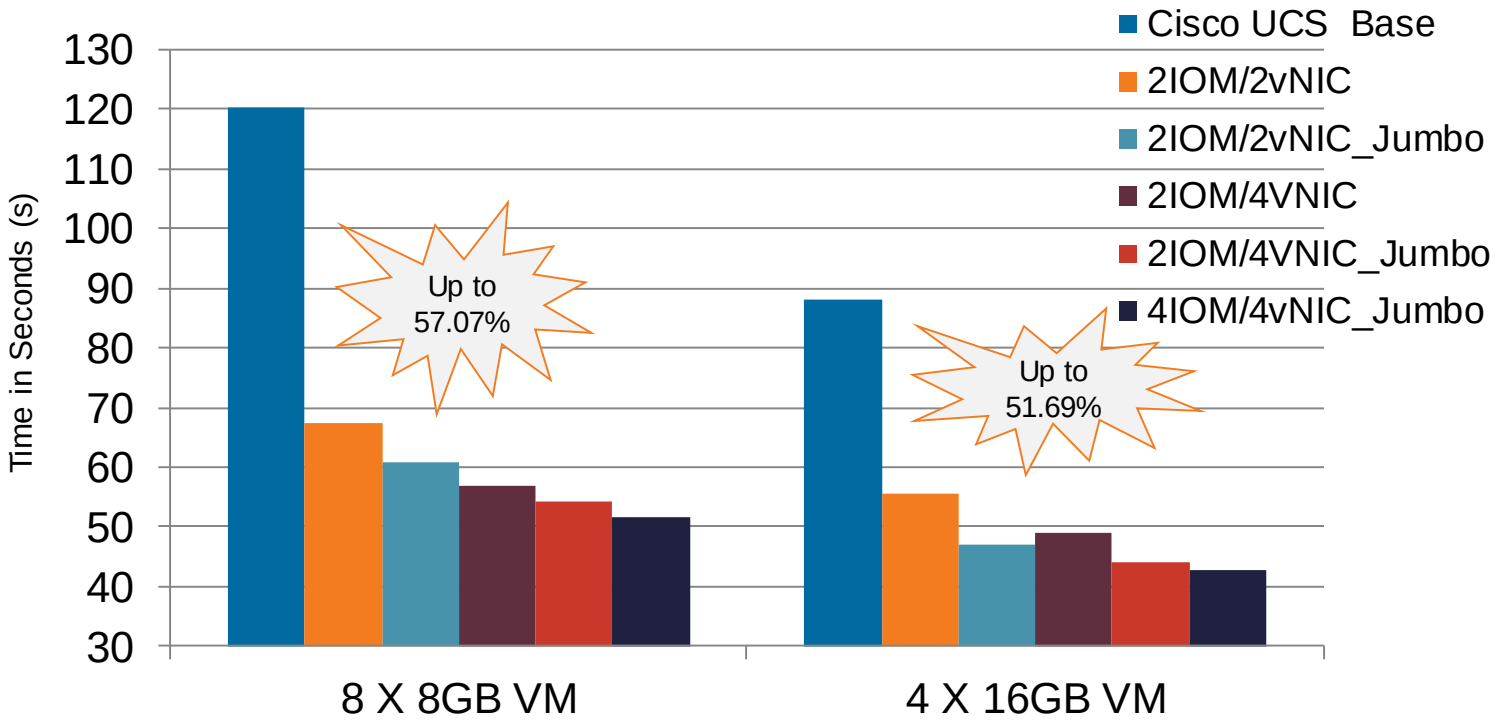

Source: Cisco Internal Testing

Cisco live!

East-West Traffic Impact Summary

Cisco live!

Testing Summary

- Cisco UCS outperformed legacy environments on all Latency and VM migration tests
- Cisco UCS is specifically designed for compute East-West communication
- Cisco UCS offers the most flexible networking configuration options for architects
  - Traffic Pinning
  - Manage vNIC Fabric Placement – not stuck with legacy hardware rules
  - Add more bandwidth easily with chassis to FI uplinks – NOT adding expensive switches
- Quality and Reproduction of Testing
  - Thousands of migration samples gathered then averaged
  - Removes vCenter time variances
  - Not a benchmark, anyone can replicate test setup

Observations and Recommendations

- Biggest impact is to spread load across both UCS Fabrics
  - Example 1 vNIC on Fabric A and 1 vNIC on Fabric B
  - vCenter seems to keep traffic on each respective fabric – not guaranteed
  - Traffic does not appear to cross from Fabric A to B unless there is a failure event
- Adding multiple VMkernel interfaces yields excellent results
  - Multi-NIC vMotion capability
  - Increases the number of TCP flows not subject to the 1 NIC 10Gb limit
- Adding more IOM Links do not get much performance increases without increasing the host ports as well.
- Jumbo frames have little impact on performance (2vnic/4IOM) but do make a **bigger** difference on 2vnic/2iom and 4vnic/4iom configurations.
  - Impact greater when vNIC and IOM uplinks are evenly matched

Migration Cisco UCS Test Summary of All Tests

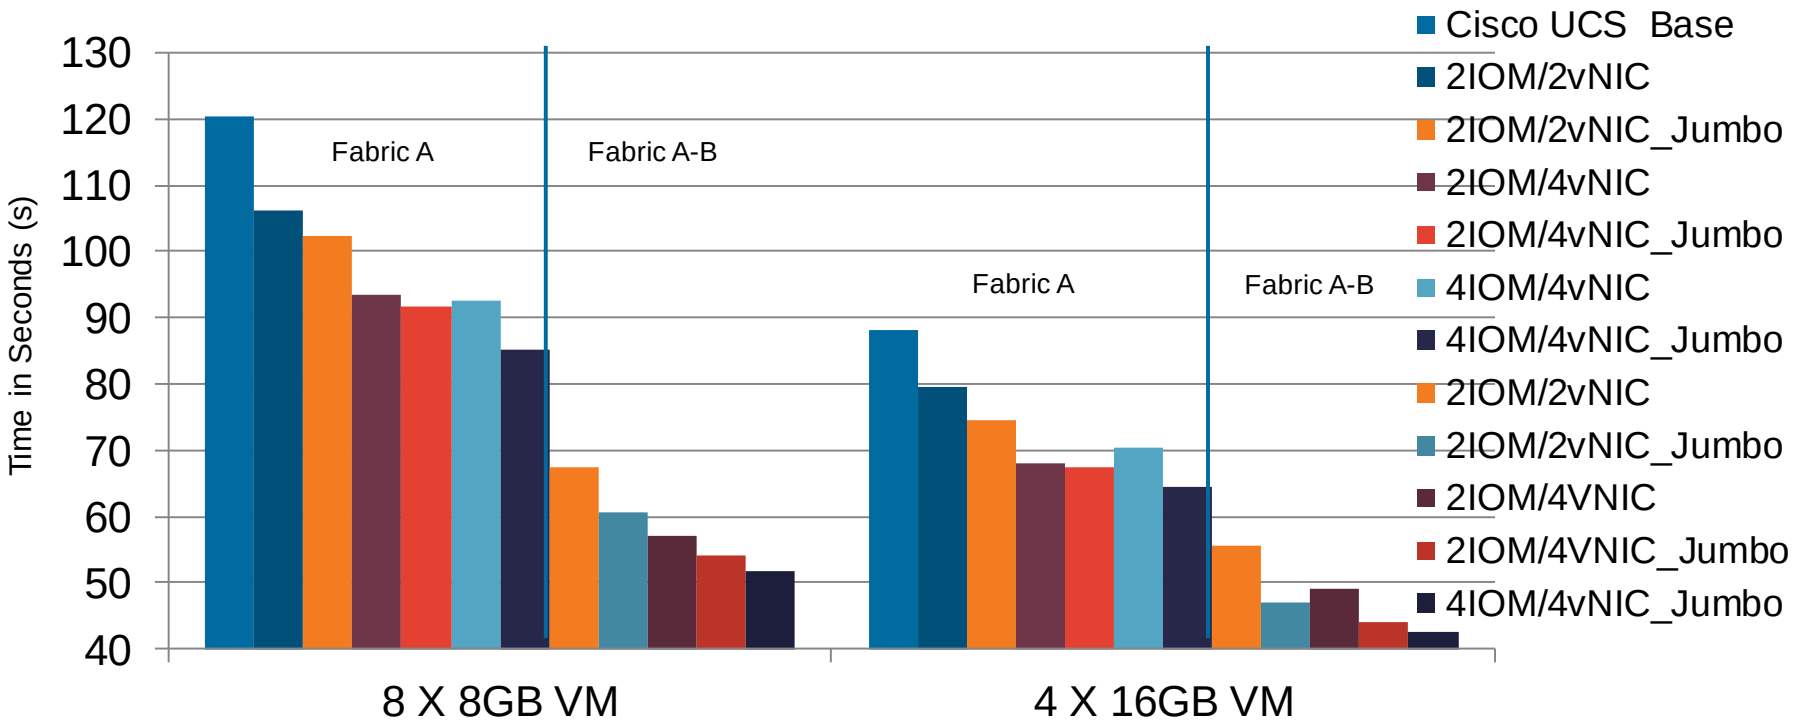

Source: Cisco Internal Testing


Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2015 T-Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site http://showcase.genie-connect.com/clmelbourne2015
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm

Learn online with Cisco Live!
Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

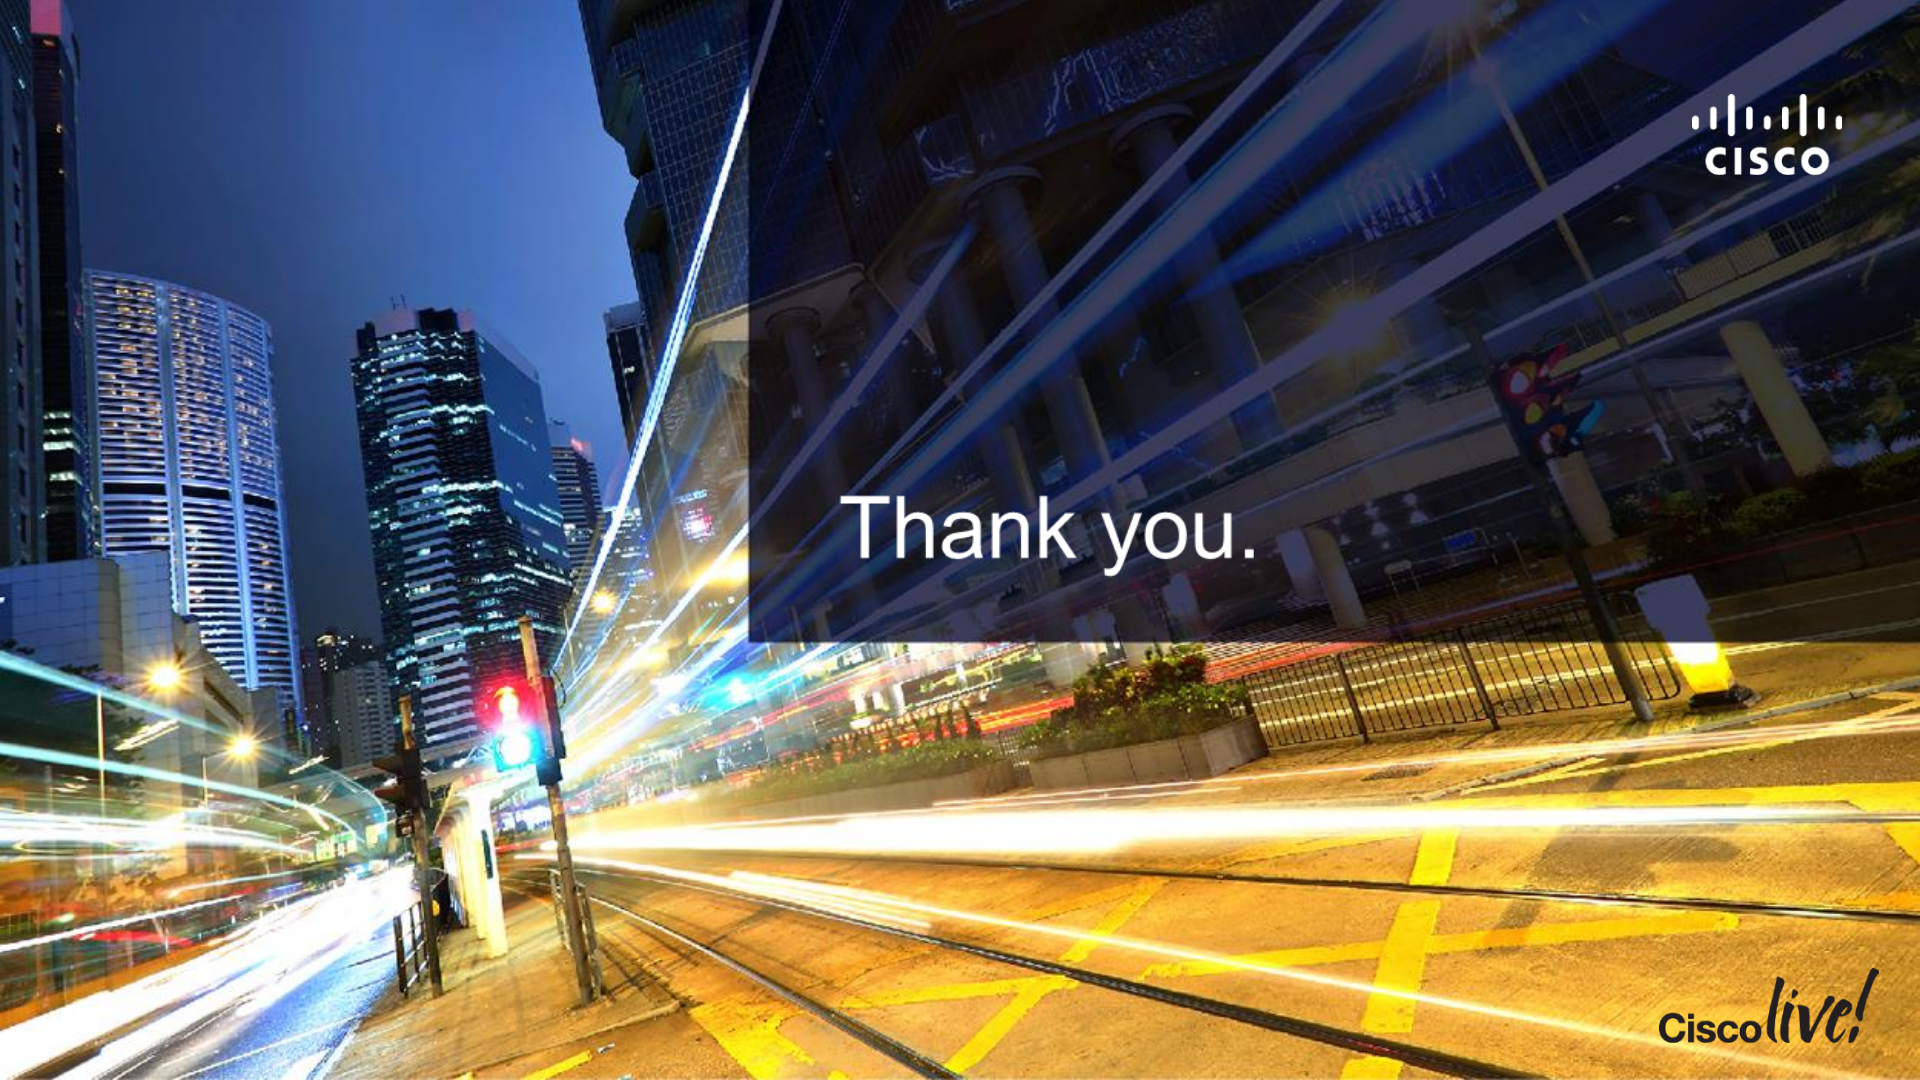
Thank you.


CISCO