

OCP 40GE Network Interface Cards

OCP Summit | Jan 2014


End to End Provider of Smart, Fast, Open Interconnects



Software Defined Networks



Open Platform




Overlay Network Tunnels



End-to-End InfiniBand and Ethernet Portfolio

ICs



Adapter Cards



Switches/Gateways




Host/Fabric Software



Metro / WAN




Cables/Modules



Open Compute Project - NIC

- First 40GbE adapter for OCP
- Open platform drives high volumes
- High performance OCP adapters
- Single and dual port 40GbE options
- Servers from multiple ODMs





OPEN
Compute Project



ConnectX-3 Pro Adapters for Open Compute Project (OCP)

- 40/56GbE, dual port, QSFP cages

- 40/56GbE, single port, QSFP cage



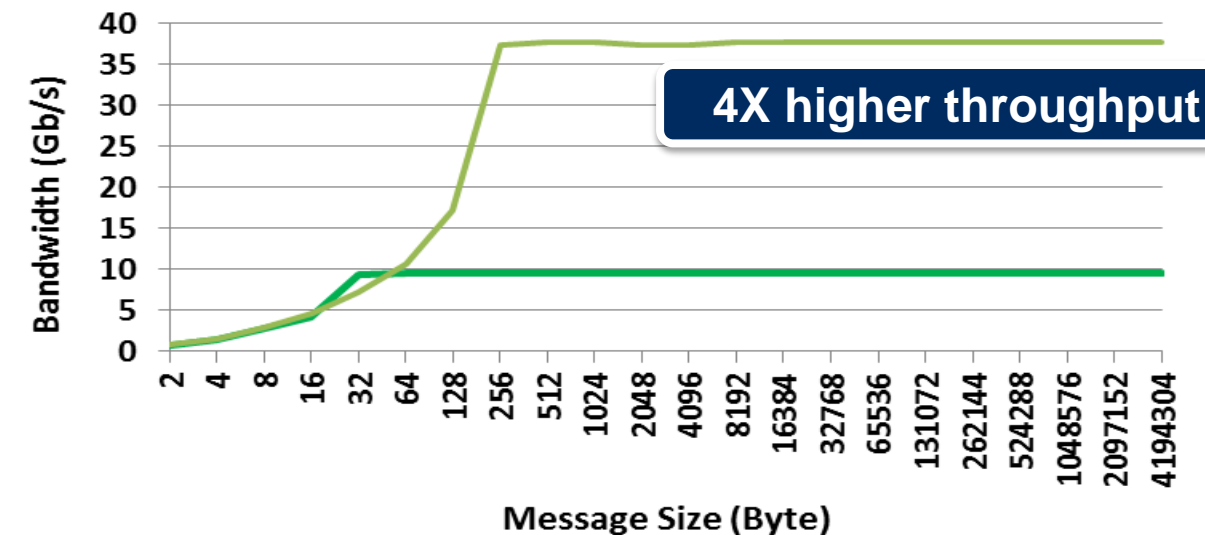
OPEN
Compute Project

40GbE OCP – Overview



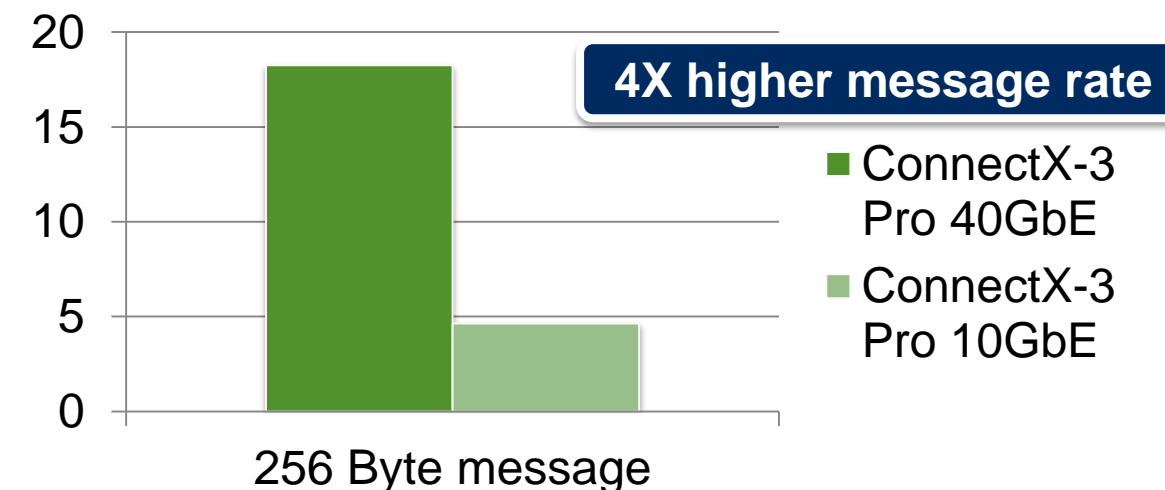
- **Single / Dual Port OCP Adapter**
 - PCIe Gen3 x8
 - 10Gb/s / 40Gb/s / 56Gb/s data rates support
 - <1usec latency
- **Unmatchable Performance in the Market**
 - Up to 4 times higher throughput than comparable 10GbE
 - Up to 4 times message rate
- **Advanced Features**
 - RoCE (RDMA over Converged Ethernet)
 - SR-IOV support with embedded switch
 - NVGRE / VXLAN Stateless Offloads
- **Host Management**
 - SMB / NC-SI Baseband Management Controller Interface
 - PXE and UEFI
 - IPv6 Support

TCP iperf 16 Streams

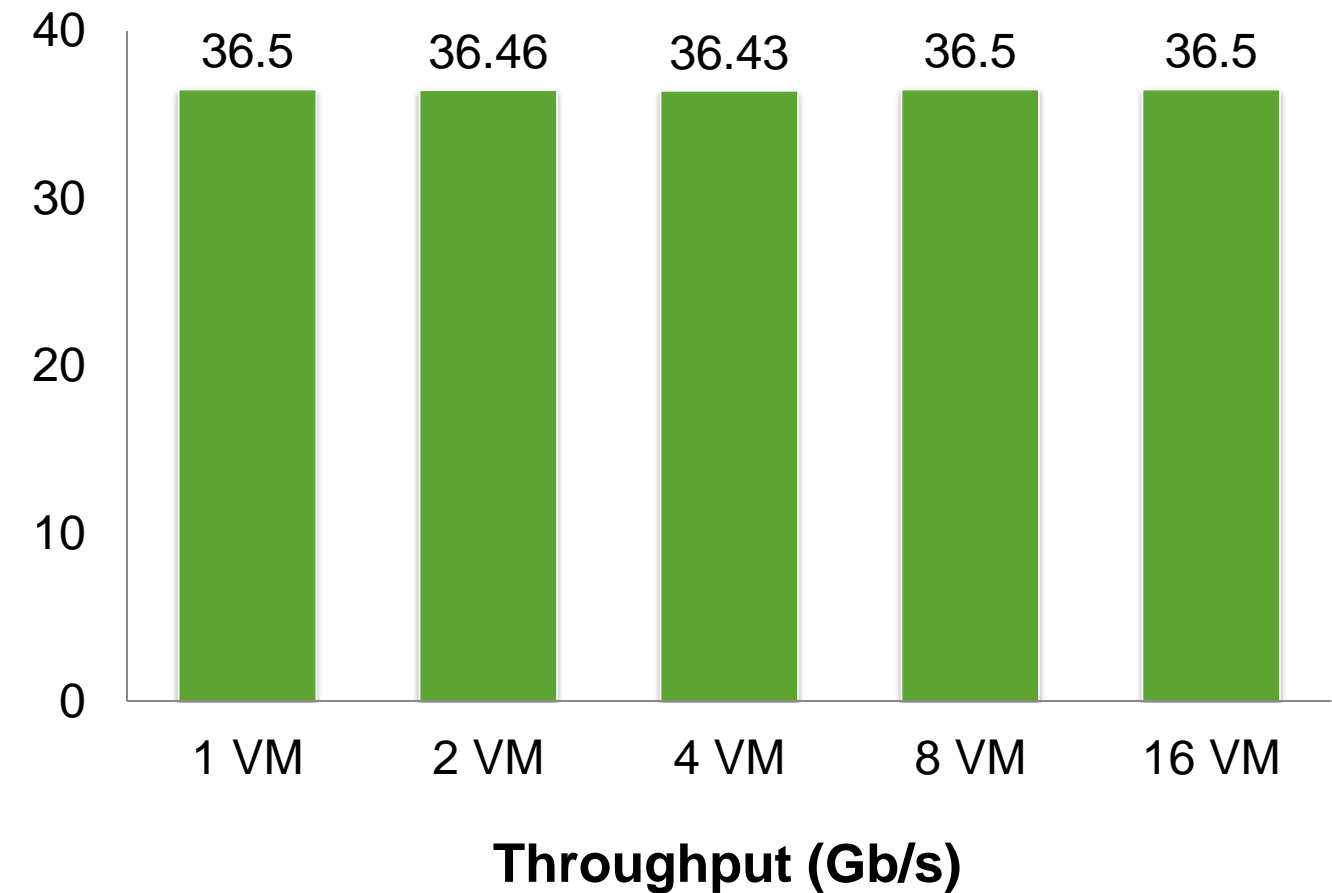


— ConnectX-3 Pro 10GbE — ConnectX-3 Pro 40GbE

Message Rate (million messages / sec)

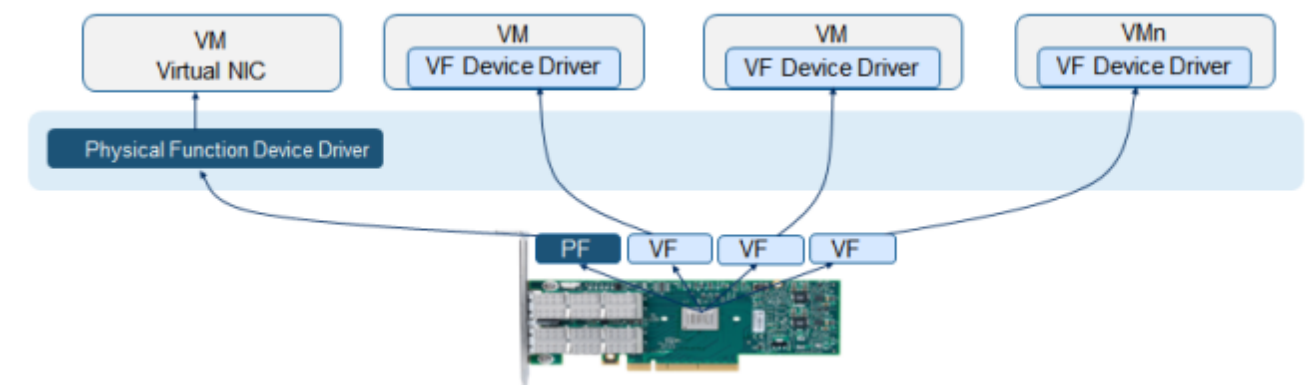



- **Highest performance in the industry**
 - Latency as low as 1us
 - Reliable/unreliable, connected/datagram
 - Unicast and multicast
- **Fast storage access**
 - RoCE and iSER support for storage access
- **RoCE for virtualized environments**
 - RoCE for SR-IOV connected VMs
 - Near 40GbE throughput in virtualized environments



Leading Solution for Virtualized Environments

- Single Root I/O Virtualization (SR-IOV) support
- Highest throughput available
 - Enabling more VMs on single machine
 - Highest traffic rate for each VM
- Overlay networks offloads
 - NVGRE and VXLAN
 - Breaking the 10GbE throughput barrier
- CPU offloads
 - SR-IOV enables application-direct access
 - Improving CPU utilization



OCP with Offload Engines for Overlay Network Protocols



■ Overlay Networks Offload

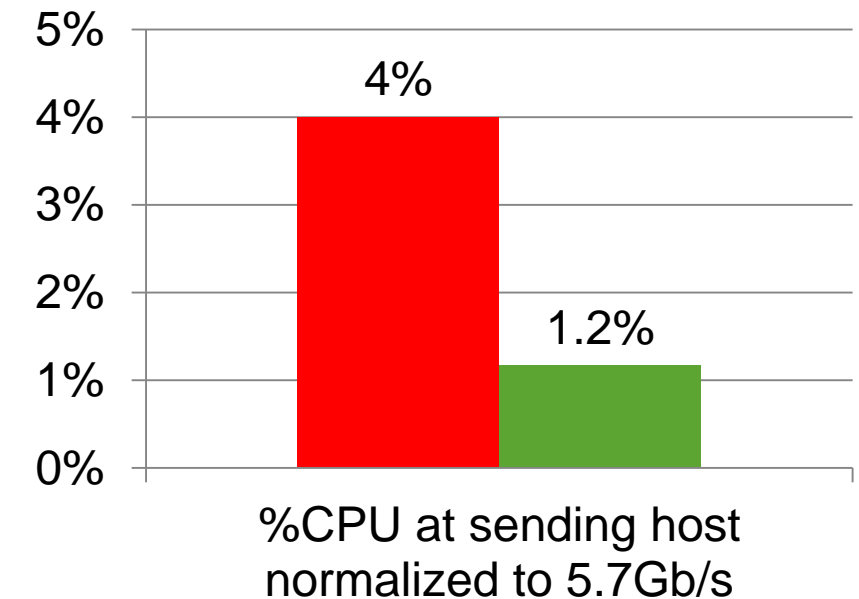
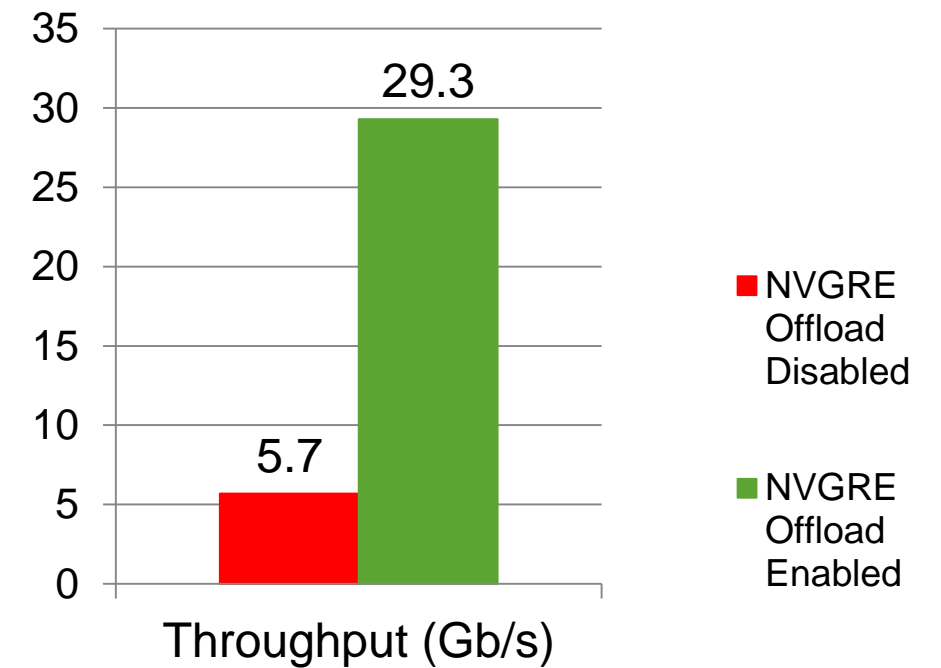
- Enabling unmatched throughput for virtualized environments
- VXLAN and NVGRE supported

■ Mellanox unique hardware offload of Overlay protocols

- Checksums, LSO, Flow ID calculation, VLAN Stripping / insertion
- Combined with steering mechanisms: RSS, VMQ

■ Mellanox offloads increase application performance

- Up to ~30Gb/s when NVGRE offload enabled
- Above 60% decrease in CPU usage when NVGRE offload enabled



OCP Cards – Product Family



OPEN
Compute Project

Ethernet					
	<table border="1"> <tr> <td>MCX341A-X (single port) MCX342A-X (dual port)</td> <td>MCX345A-B (single port) MCX346A-B (dual port)</td> </tr> <tr> <td>SFP+ 10GbE</td> <td>QSFP 40/56GbE</td> </tr> </table>	MCX341A-X (single port) MCX342A-X (dual port)	MCX345A-B (single port) MCX346A-B (dual port)	SFP+ 10GbE	QSFP 40/56GbE
MCX341A-X (single port) MCX342A-X (dual port)	MCX345A-B (single port) MCX346A-B (dual port)				
SFP+ 10GbE	QSFP 40/56GbE				
Connector and Port Speed					
PCIe 3.0 Speed	8.0GT/s (52Gb/s)				
Features	ConnectX-3 and ConnectX-3 Pro Flavors				
OS Support	RHEL, CentOS, SLES, OEL, Windows, ESX/vSphere, Ubuntu, Citrix, Fedora				


Thank You

- Low Power

- Typical 2-port 40GbE OCP - 8.68W

- Temperature (according to spec)

- Operational: 0°C to 35°C
- Non-operational: 0°C to 70°C



	Linux tests (slides 5-7)	Windows tests (NVGRE, slide 8)
Gen3 System	IBM X3650 M4	HP DL380 Gen 8
CPU	Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz	Intel(R) Xeon(R) CPU E5-2690 0 @ 2.90GHz
Num of cores	16	16
Distribution name	Red_Hat_Enterprise_Linux_Server_release_6.3_(Santiago)	Windows
Os version	2.6.32-279.el6.x86_64	Windows server 2012 R2
Guest Os	Red_Hat_Enterprise_Linux_Server_release_6.3_(Santiago)	Windows server 2012 R2
Driver version	MLNX_OFED_LINUX-2.1-1.0.0	WinOf 4.6
Firmware	2.30.8000	2.30.8000
MTU	1500 ETH	1500 ETH
PCIe	3	3
Width	x8	x8
40GbE Adapter	MCX354A-FCCT	MCX354A-FCCT