



Building a 40 Gbps Next Generation Virtualized Security Processor

HOT CHIPS 23 – August 2011

Jeff Pangborn

Cavium, Inc

- Motivation for developing the NITROX[®] III Security Processor
- Major Components
 - Crypto Engines
 - Compression Engines
 - Interconnect
 - Virtualization
- Challenges
- Lessons Learned

- 1. Offer very high 2048-bit RSA Ops/Sec to meet next generation Data Center requirements**
 - NIST recommends using 2048-bit RSA keys as of January 1, 2011
 - Exponentially more complex than 1024-bit keys
- 2. Increase Gbps**
 - More bulk crypto throughput (up to 40 Gbps)
- 3. Expand feature set to meet additional Data Center & Enterprise requirements**
 - Compression/De-compression off-load support

4. Deliver the best Performance/Power

- Stringent Data Center/Enterprise power requirements
- < 25 W limitation for use as a PCI-Express adapter card

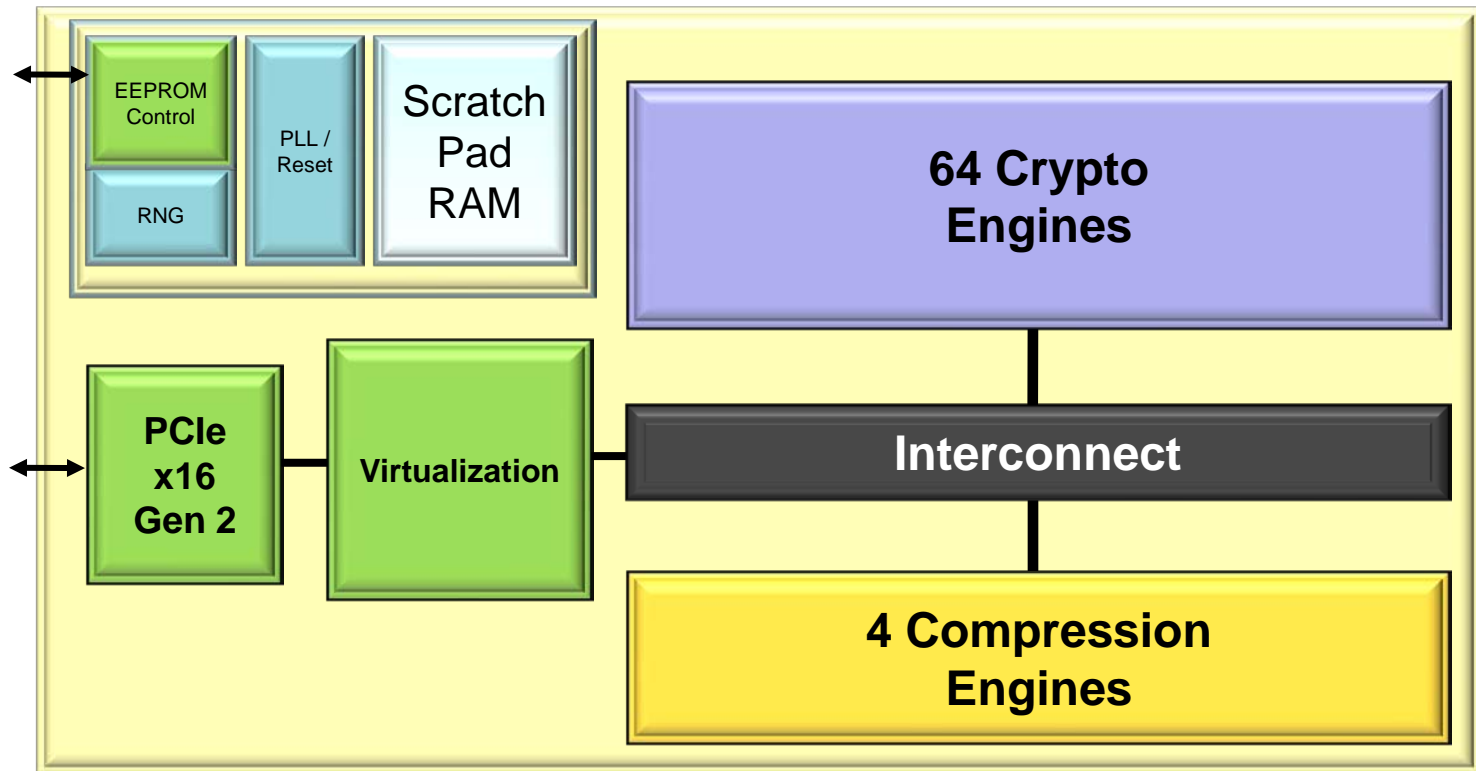
5. Support increasing Virtualization Requirements

- Support many virtual functions (VFs)
- Virtual Machine (VM) guests must be given a fair share of resources
- Similar software drivers for guest and non-virtualized host

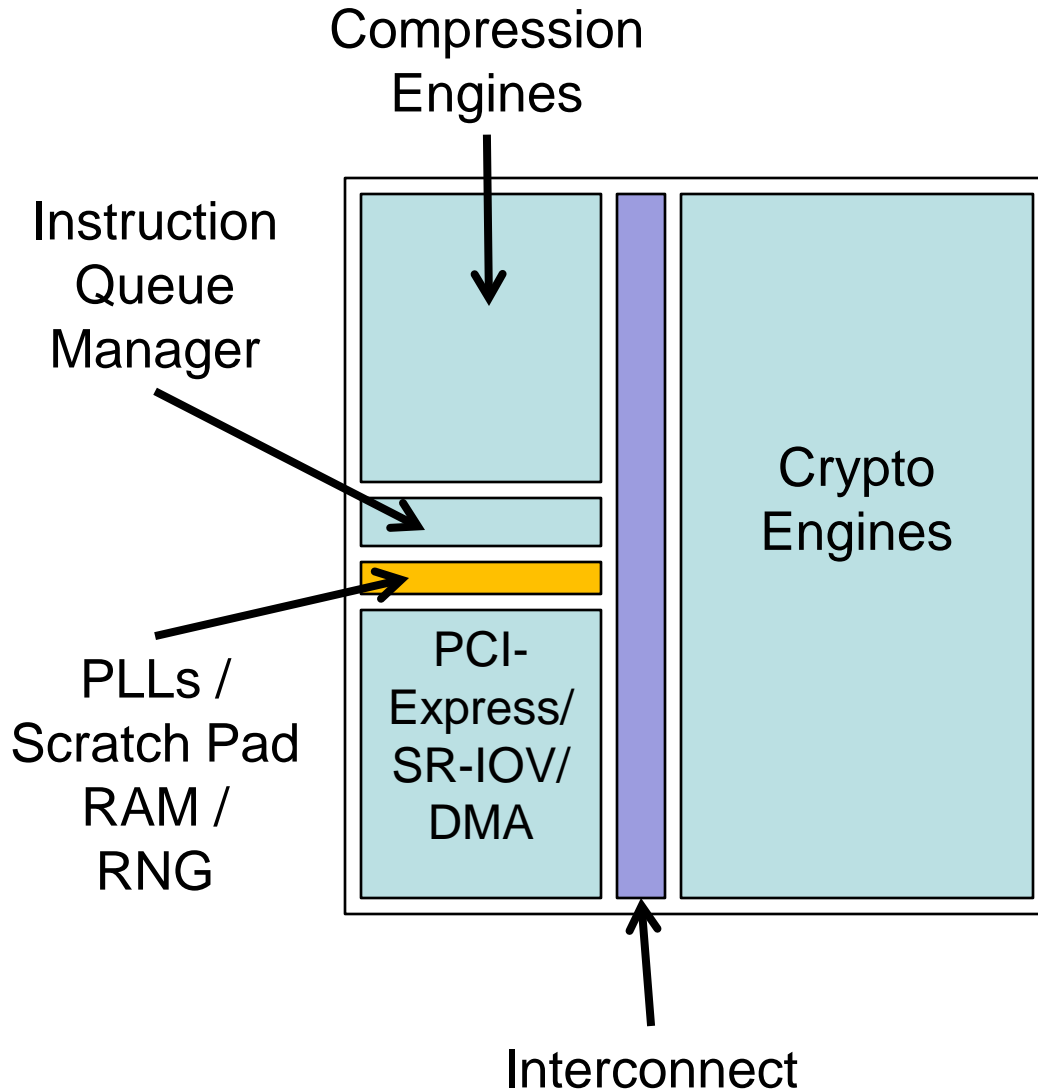
6. Software compatibility with previous generation

- Motivation for developing the NITROX[®] III Security Processor
- **Major Components**
 - Crypto Engines
 - Compression Engines
 - Interconnect
 - Virtualization
- Challenges
- Lessons Learned

NITROX III Block Diagram



NITROX III Floor plan



- 40nm process
- Multiple clock domains:
 - PCI-Express
 - Crypto engines
 - Compression engines
 - Interconnect


- Cavium-designed Crypto engines
 - Concurrent support of multiple protocols (e.g. IPSec, SSL, RSA)
 - Micro programmable
 - Off-load more than just crypto
 - Flexible – adaptable to new protocols
 - Backwards-compatible with previous generation engine
 - No external memory – lower power consumption
- IPSec/SSL off-load
 - 40 Gbps bulk crypto
- 200k 1024-bit RSA Ops/sec
- 35k 2048-bit RSA Ops/sec

“RSA Op” = RSA_Private_Key_Encrypt

Crypto Engines (cont.)



- Performance scales linearly with number of engines

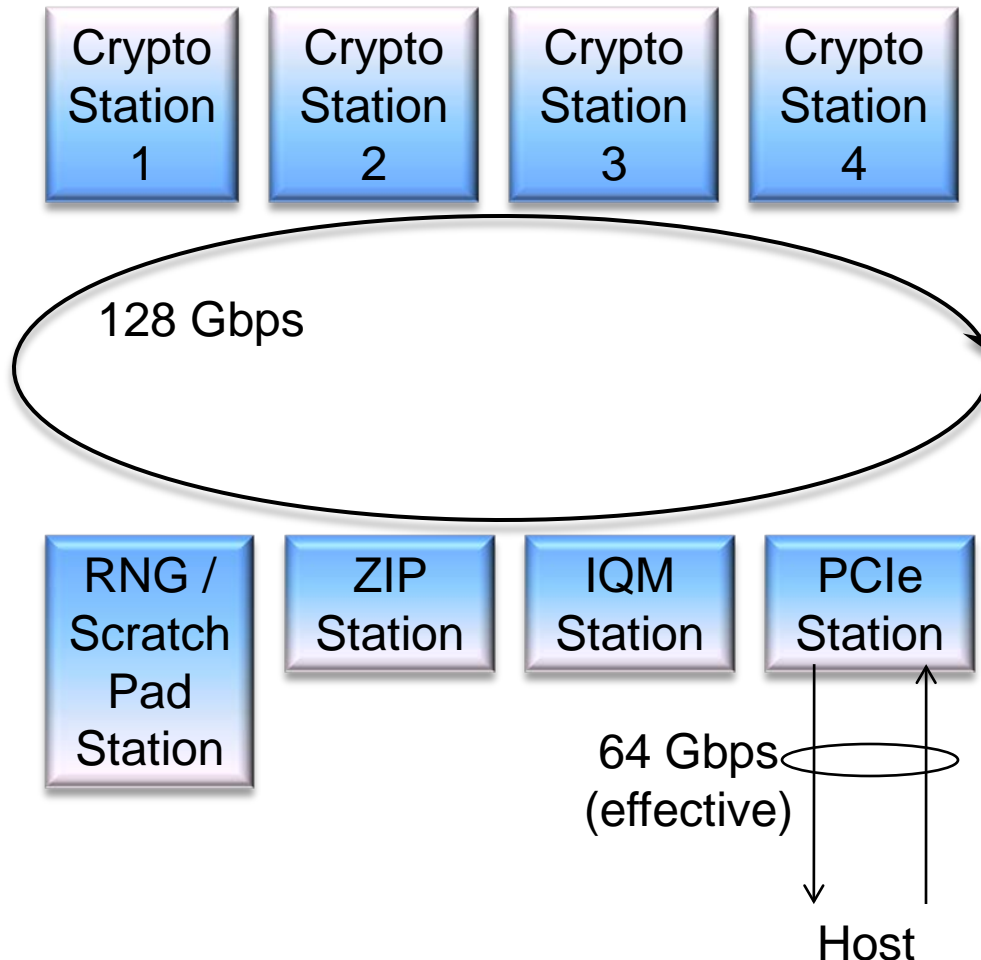


Example shown is for cores running at lower frequency bin
Higher frequency bins are available

- Compression/De-compression
 - Support common protocols
 - DEFLATE (RFC1951)
 - ZLIB (RFC1950)
 - GZIP (RFC1952)
 - LZS
 - Support up to 20 Gbps compression offload
 - Programmable scatter/gather DMA

- Multi-point internal data path and command interconnect
- Key Requirements
 - Compression requires high data bandwidth
 - Customize design to match data flow
 - Configurable bandwidth between clients
 - Debug capability

Interconnect (cont.)



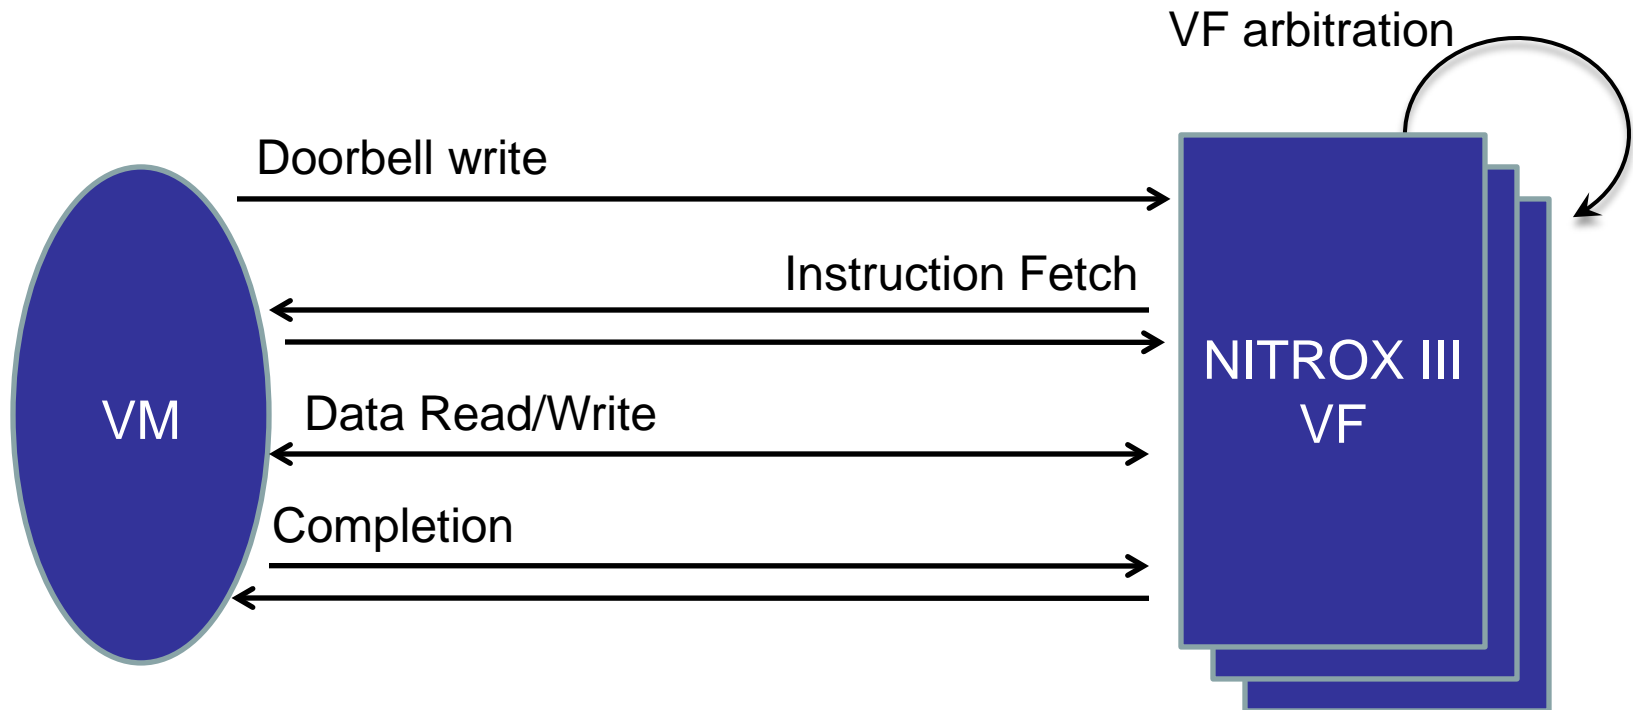
- In-house developed interconnect
- Ring Topology
 - Four Crypto Engine Stations
 - Compression Engine (ZIP) Station
 - Instruction Queue Manager (IQM) Station
 - PCI-Express Station
 - Random Number Generator / Scratch Pad Memory Station

- Benefits
 - Distributed arbitration
 - Easy timing closure
 - Scales well with large number of engines
 - Avoids unnecessary point-to-point connections
 - Supports transfer of virtualization information



- Protection and isolation of guest systems
- PCI-Express SR-IOV 1.1
- Instruction Queue Manager
 - Programmable QoS between VFs
 - Arbitration among and within VFs
- Interconnect must be VF-aware
- Replication/Sharing of resources
 - CSRs
 - Scratch Pad Memory

Virtualization (cont.)

- Similar mechanism used for both crypto and compression operations



Feature Comparison

	Current Generation NITROX PX	NITROX III
Crypto Engines	8	64
Bulk Data Crypto Bandwidth	2.5 Gbps	40 Gbps
1024-bit RSA Performance	17k Ops/Sec	200k Ops/Sec
Performance Increase		12 - 16X
Power Consumption	4 Watts	20 Watts
Crypto Performance per Watt	0.625 Gbps / W	2 Gbps / W
Compression Support		
Virtualization Support		

- Motivation for developing the NITROX[®] III Security Processor
- Major Components
 - Crypto Engines
 - Compression Engines
 - Interconnect
 - Virtualization
- **Challenges**
- Lessons Learned

- Several new components
- New technology node
- Completely new verification infrastructure due to virtualization support
 - Very large configuration space
- New environment
 - Use two simulators to catch more issues

- Single-pass design
 - Programmable performance statistics gathering
 - Software-observable internal state
 - Configurable interconnect
 - Micro programmable crypto engines

- Motivation for developing the NITROX[®] III
- Major Components
 - Crypto Engines
 - Compression Engines
 - Interconnect
 - Virtualization
- Challenges
- **Lessons Learned**

- Using many virtual functions presents several challenges
 - How does one verify VF isolation when using shared resources?
 - Multiple functions allows for a large configuration space
- Use RTL tricks to allow faster simulation
 - Usage of `define flags in RTL can allow simulation shortcuts
- Using multiple verification simulators can be both a blessing and a curse