



# AppliedMicro X-Gene<sup>®</sup> ARM Processors

*Optimized Scale-Out Solutions for Supercomputing*

# AppliedMicro X-Gene<sup>®</sup> Processor Philosophy


- **Few workloads are compute bound**

- Most are limited by memory capacity, bandwidth, or I/O
- HPC workloads are better served by GPGPU



- **Scale-out versus scale-up**

- High density
- Performance per Watt
- Performance per \$




- **Balance**

- Strong CPU with an optimized ARMv8 core
- Large memory capacity / bandwidth – **adequate memory is not an upsell**
- Low power – **power efficiency is not an upsell**



- **Open Source**

- Open Source Software
- Open Source Hardware




A bit about X-Gene<sup>®</sup> ARM technology being deployed in the  
enterprise, today...

# Emergence of the Optimized ARM Scale-Out Data Center



## Real Solutions

*Validated Results  
Lower TCO*



### Processor Architecture

- *Strong compute*
- *Large memory*
- *Low power*
- *Cost-effective*



ubuntu



redhat



mongoDB



Apache



KVM




Xen



CentOS

### Software Ecosystem

- *Mature, optimized toolchain*
- *Broad Linux support*
- *Open-source workloads*



### System Architecture

- *High Density*
- *Power-efficient*
- *Top-Tier Suppliers*

# X-Gene<sup>®</sup> Technology in the Enterprise

*Demonstrating Value in Leading IT organizations - Today*

- **Node Density / Rack**

*900% Higher*


- **Power Consumption**

*85% Lower*

- **Acquisition Cost**

*45% Lower*

## ARM 64-bit compute Platforms



- Comparing Node to Node ARM delivers game changing \$/Watt/Cu FT benefit over traditional infrastructure
- Delivering traditional data center services like Firewall,VPN, LDAP, Kerberos,Virtualized etc.


Source: PayPal

# Real Workload Performance

## Web Server (WRK Benchmark)

AppliedMicro  
X-Gene<sup>®</sup> 2

Performance  
(higher is better)



Latency  
(lower is Better)




Bandwidth  
(higher is better)



- X-Gene 2 (8c @ 2.4 GHz)**
- 4 node 1U / ½ width sled
  - 64GB DDR3-1600
  - 4 x 10GbE (integrated)
  - Wall power: ~190 Watts

Intel Xeon<sup>®</sup>  
E5-2630v3



- Xeon e5-2630v3 (8c/16t @ 2.4 GHz)**
- 2P 1U / ½ width sled
  - 64GB DDR4-2133
  - 4 x 10GbE (NIC)
  - Wall power: ~180 Watts

**Up to 35% Higher Performance | Lower TCO**

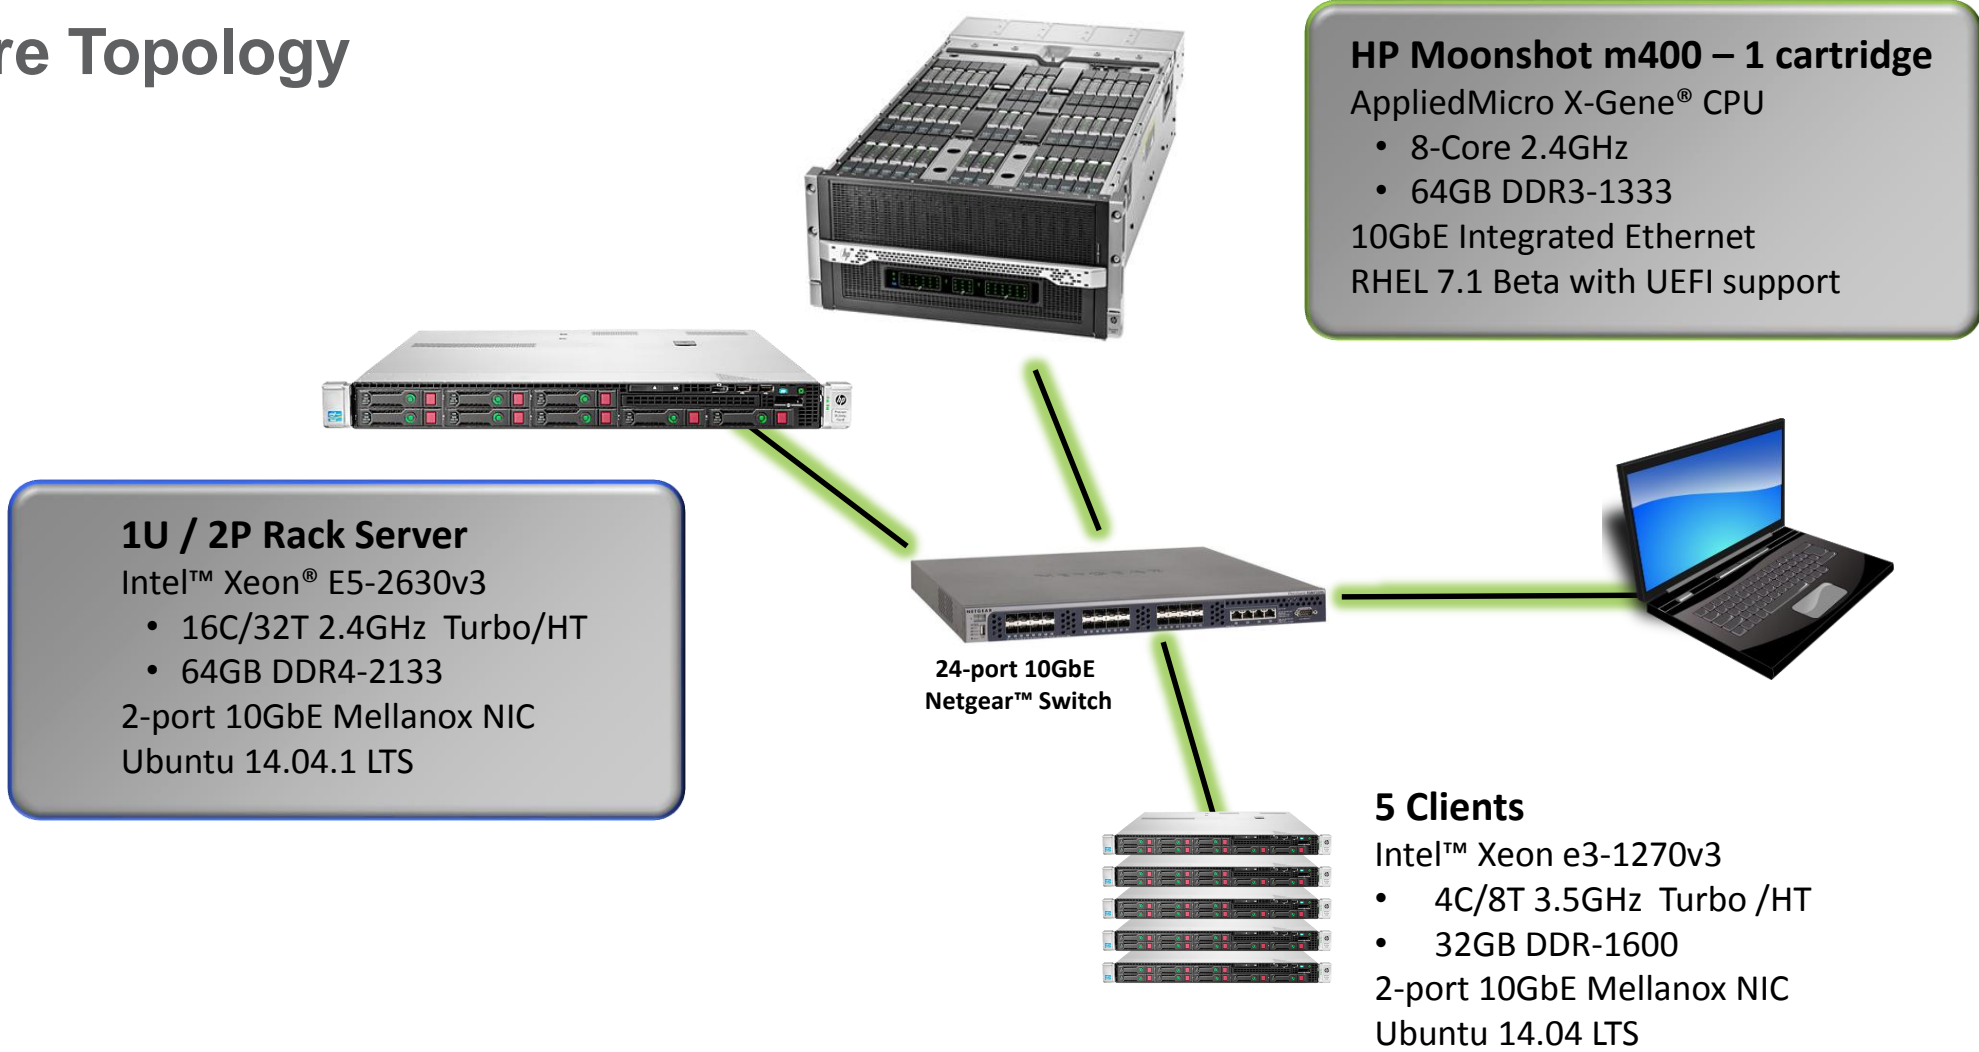
*Standard CPU benchmarks do not always translate to delivered workload performance*

Source: AppliedMicro

# Real Workload Performance

## *In-Memory Database (MongoDB - YCSB)*

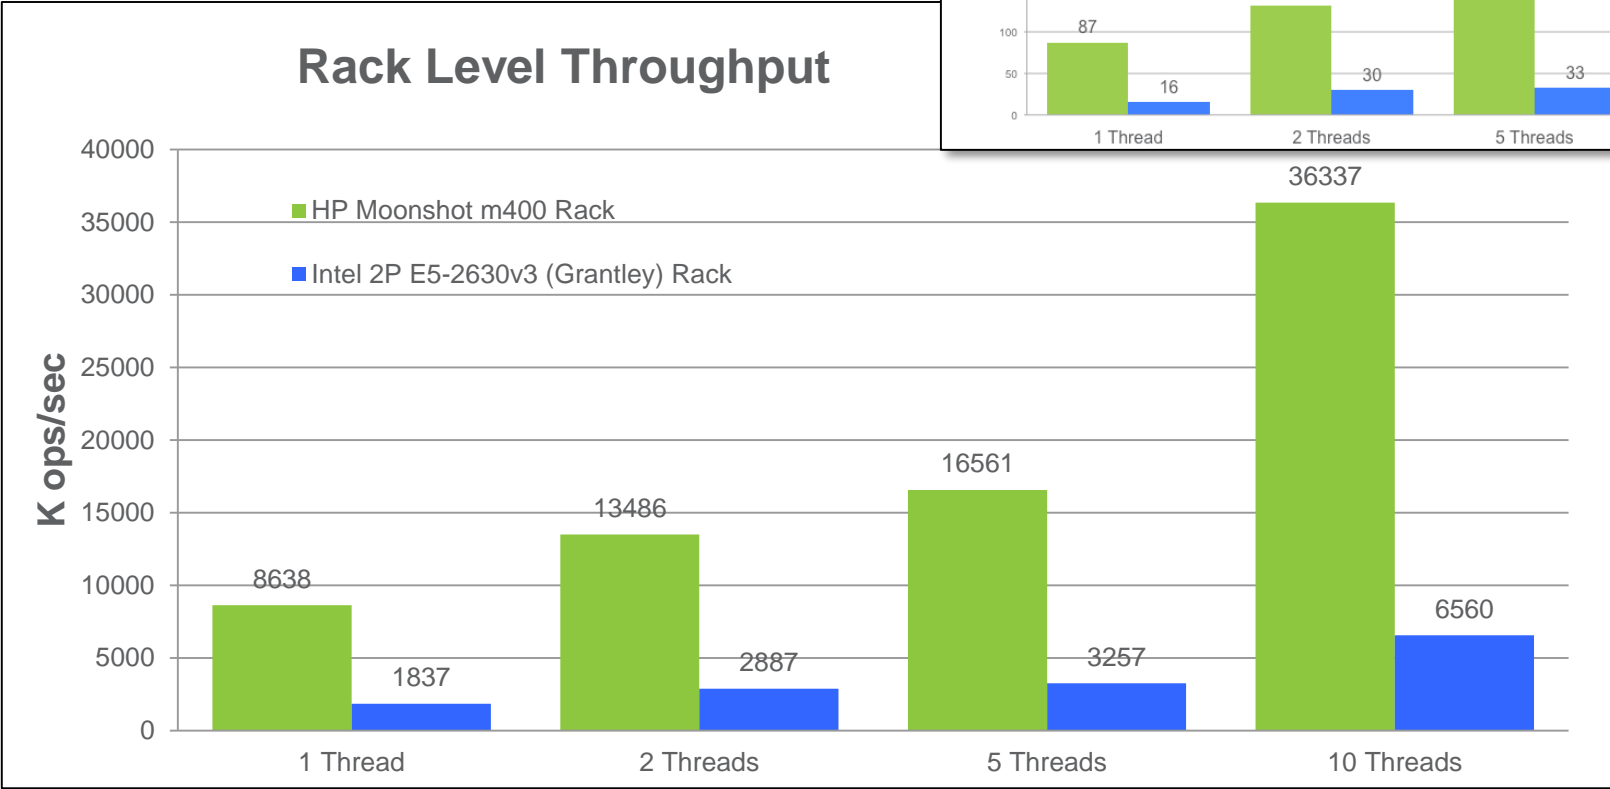
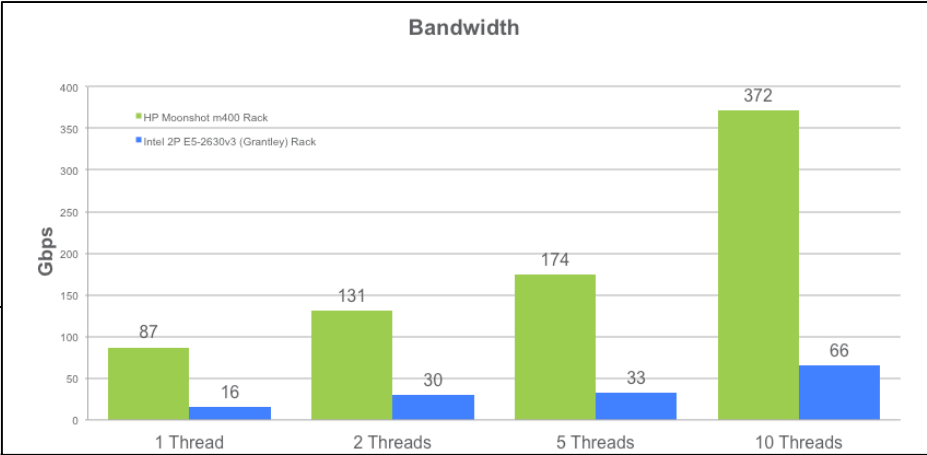
### Hardware Topology



# Real Workload Performance

## In-Memory Database (MongoDB - YCSB)

**Rack-Level Scalability**  
*5x the throughput of a Haswell Xeon® e5  
2P rack server implementation*



42U Rack, 9 Moonshot m400 chassis/rack



# Lower TCO with X-Gen<sup>®</sup> Technology

## Web / Application Tier @ 30kW/Rack


Traditional Intel Xeon<sup>®</sup> E5 2P/1U



**270 Web Servers + 45 Application Servers**

- Nine racks
  - 35 servers per rack
  - 2 TOR switches per rack
- 315 total nodes

HP Moonshot with m400 (X-Gen<sup>®</sup> CPU)



**Web servers (32GB)  
App servers (64GB)**

**270 Web Servers + 90 Application Servers\***

- One rack
  - 8 Moonshot Chassis
  - 2 TOR Switches
- 360 total 1P m400 nodes

**55%+ Hardware Acquisition Cost Savings**

*Additional TCO reduction via  
simplified management and lower power*

# X-Gene<sup>®</sup> ARM Processor Software Ecosystem



redhat.  
PEAP

ubuntu



debian



CentOS

Operating Systems



OpenJDK



Hypervisors & Java



Compilers



American  
Megatrends



Tools & BIOS

...but what about High Performance Computing?


# AppliedMicro X-Gene<sup>®</sup> HPC Philosophy

- **Workloads are compute bound**
  - ...but 'general purpose' compute is not the path to exascale



- **Scale-out versus scale-up**

- High density
- Performance per Watt
- Performance per \$




- **Balance**

- Power-efficient CPU with an optimized, power-efficient ARMv8 core
- High performance GPUs for the 'heavy lifting'
- A better alternative to 'brute force' high performance computing



- **Open Source**

- Open Source Software
- Open Source Hardware



# X-Gen<sup>®</sup> Processor Platforms

*Multiple SKUs from Leading OEM and ODM Partners*



**Gigabyte MP30-AR0**



**Multiple New  
Platforms in  
Development**



**HP ProLiant m400**



**Cirrascale RM1905D**



**E4 ARKA RK003**



**Mitac Datun**

# The ARM Revolution has Expanded to Supercomputing

## *64-bit X-Gen<sup>®</sup> ARM Servers in production today*

- The “one size fits all” data center is no longer sufficient
- AppliedMicro is powering the transition
  - Proven: real customers in production today
  - Performance: balanced 64-bit ARM compute with large memory
  - Economics: TCO savings via both lowered CapEx and OpEx



“HP Moonshot is a first-of-a-kind system that’s enabling us to **extend the range of our calculations to solve really complex problems** in a highly efficient 64-bit architecture.”

**James Ang, Technical Manager,  
Sandia National Laboratories**




“HP Moonshot offers capabilities that will be critical to the future of cloud computing. It empowers researchers to **develop fundamental breakthroughs** that have the potential to change the performance, reliability, and security of future clouds.”

**Robert Ricci, Research Asst. Professor of Comp.  
Science, University of Utah**


University of Utah Cloudlab on 315 ARM nodes: <https://www.youtube.com/watch?v=yIA4FKibfXU&sns=em>

# X-Gene<sup>®</sup> Processors in HPC

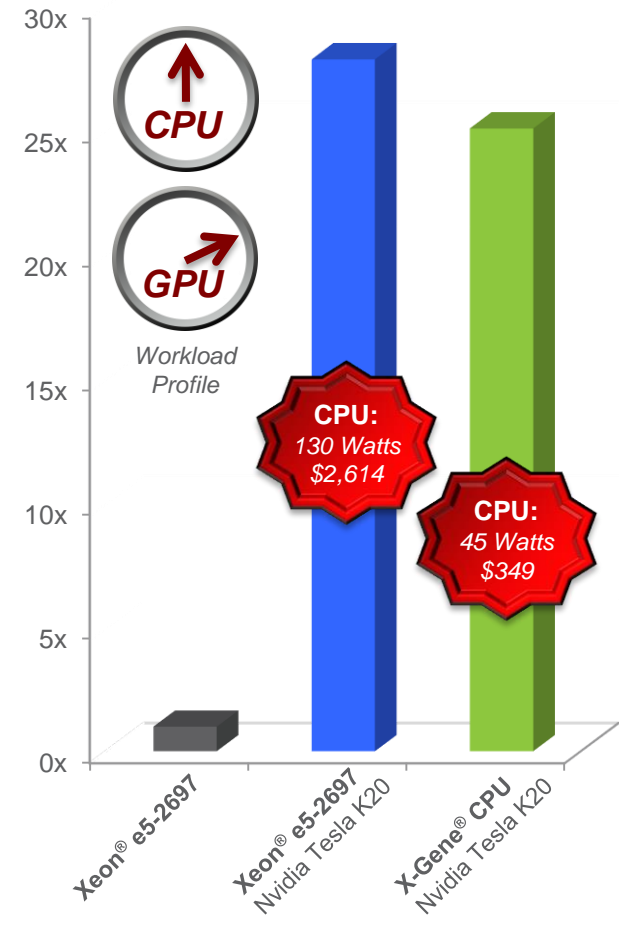
*The Efficiency of ARM & the Power of Tesla<sup>™</sup> GPUs*



NAMD



HOOMD



HPCG

All code recompiled to ARM64, no optimizations

Source: nVidia

# Delivering Performance that Matters.

---

There is a better answer to 'brute force' HPC: heterogeneous compute

Platforms with X-Gene® ARM technology and Nvidia GPUs is in production

The software ecosystem is established

The results are compelling





