

THE FIRST ENTERPRISE CLASS 64-BIT ARMv8 SERVER: HP MOONSHOT SYSTEM'S HP PROLIANT M400 SERVER CARTRIDGE

APPLIEDMICRO X-GENE-BASED SERVER DEMONSTRATES 35% LOWER TCO FOR SCALE-OUT WEB TIER/CACHING ENVIRONMENTS

EXECUTIVE SUMMARY

Ubiquitous cloud-enabled smart devices are a driving force behind a major shift in IT infrastructure. Service providers deploying context-rich services to these devices are building massive new datacenter capacity and looking to their vendors to optimize infrastructure for their specific workloads. But given the rapid rate of workload and application evolution, infrastructure optimization will be a continuous process for at least the next few years; optimization demands flexible hardware and software infrastructure.

System designers and silicon vendors have recognized that this shift in customer demands cannot be served effectively with traditional datacenter and server architectures, and they have responded with new approaches to system design. HP launched the HP Moonshot System hardware platform in 2013—a modular and ecosystem-driven approach targeted at specific workloads. In 2011, ARM Holdings and its partners began promoting the vision of extremely power efficient servers. Earlier this year, AppliedMicro's X-Gene, which combines server-class performance with mobile-pedigree power efficiency, was the first 64-bit ARMv8-based server system-on-chip (SoC) to begin shipping in production.

HP and AppliedMicro partnered to create HP's ProLiant m400 server cartridge for HP's Moonshot 1500 chassis—the world's first enterprise-class 64-bit ARMv8 server in the market with production shipments beginning in calendar Q4 2014. The ProLiant m400 provides choice in the marketplace versus traditional x86 servers. It is initially positioned as a web tier/caching solution for service providers, commercial internet providers, early adopters of ARM servers, and the ARM server/mobile software development community. This paper describes a TCO analysis that illustrates a 35% potential savings when replacing standard x86 1U servers with ProLiant m400 servers for web and application/caching tiers in a commercial internet environment.

Datacenter customers looking to evaluate the ProLiant m400 should determine which workloads are good candidates for this class of server, assess the 64-bit ARMv8 software ecosystem to determine readiness of their solution stacks, and work with HP to gain access to the ProLiant m400 for evaluation via the HP Discovery Lab or beta/evaluation programs.

MEGATRENDS AND MARKET DRIVERS

The datacenter industry is experiencing a profound shift that will be more or less complete by the end of this decade. The dynamic driving this shift is the proliferation of smart devices with increasingly ubiquitous broadband access that enables people and things to interact with remote cloud-based software. Cloud-based analytics are evolving to learn from the aggregate of people using it—referred to as “Big Data”—to provide much better context than a purely local application. Every one of these devices, one way or another, will be connected to a datacenter for control, management, and analysis. Service providers are creating new workloads to deploy their new context-rich services, and they are building massive new datacenter capacity to host these new workloads. Datacenters operating at this scale are often referred to as “hyperscale” datacenters. The required datacenter capacity to drive this shift cannot be served effectively with current datacenter and server architectures.

Unlike highly virtualized enterprise IT runtime environments, hyperscale services run individual, specialized workloads at such scale that they do not share infrastructure with other workloads at runtime. Instead of optimizing infrastructure to run any workload at a “least common denominator” of service, hyperscale customers are asking their suppliers for infrastructure that they can optimize for high value and specialized workload classes. There is a solid return for investing in an optimal balance of density, costs, and expenses for each workload class and even for specific workloads and applications when deployed at scale. Given the rapid rate of workload and application evolution, optimizing performance will be a continuous process for at least the next few years; optimization demands flexible hardware and software infrastructure.

HP MOONSHOT: AN ACCELERATOR FOR HYPERSCALE WORKLOADS

In 2010, HP chartered its Moonshot team to break out of HP’s mainstream enterprise value propositions and to build a solution that directly services the unique needs of its hyperscale datacenter customers. This team lived up to the challenge with the official launch of the [HP Moonshot System](#) hardware platform in April 2013. Thus began a

multiyear, multi-phased journey to enable a portfolio of solutions specifically tailored for workload flexibility and optimization. HP's Moonshot System provides a standard chassis that hosts customizable cartridges. In this multi-phased rollout, HP is developing an ecosystem of specialized cartridges that can use vastly different kinds of compute engines to deliver a wide range of specialized acceleration. Current and future compute engines include CPUs, GPUs, APUs, DSPs, and FPGAs from vendors such as AMD, AppliedMicro, Intel, Texas Instruments, and others.

FIGURE 1: HP MOONSHOT 1500 CHASSIS



Traditional servers rely on dedicated components, including management, networking, storage, power cords, and cooling fans all in a single chassis. In contrast, the Moonshot 1500 chassis shares power, cooling, management, and fabric for 45 individually-serviceable hot-pluggable server cartridges in a 4.3U chassis—enabling them to use less energy, cabling, space, and cost.

Each server cartridge has access to three independent network fabrics: an Ethernet switch fabric, a storage fabric, and a cluster fabric.¹ In the HP Moonshot 1500 chassis, network access for its server cartridges is implemented as two removable Ethernet switch modules that can be configured for redundancy or for maximized bandwidth. This solution uses 10GbE downlink switches with 40GbE uplink switches. The cluster fabric is an independent local interconnect topology in the shape of a 2D torus: groups of three server cartridges are connected north-south in independent rings and groups of

¹ Support for specific fabric capabilities and features vary for each ProLiant server cartridge model.

15 server cartridges are connected east-west in independent rings. This cluster fabric approach provides customers with the flexibility to configure the compute clusters within the chassis in an optimal way for their specific application(s).

Each server cartridge also has access to four SAS or SATA storage lanes and is designed to support external iSCSI storage. HP designed-in modular disk sharing and the ability to share slices of drives across this independent storage fabric, thus providing a flexible solution that can be customized dependent on the storage needs for each workload. The Moonshot 1500 Chassis has a built-in chassis management (HP iLO) module, and system power is delivered through a pooled-power backplane to make the full capacity of the hot-plug power supplies available to all cartridges.

IS THE MARKET READY FOR ARM-BASED SERVERS?

ARM-BASED SERVER HISTORY AND APPLIEDMICRO MARKET LEADERSHIP

In 2011, ARM Holdings and its partners began promoting the vision of building servers with extremely power efficient cores, SoC design, and mobility-derived power management to help reduce both capital and operational expenses for workloads that were not fully using the compute horsepower of traditional x86 server processors. In addition to the power efficiency advantages, ARM Holdings licenses their IP to multiple silicon partners, thus providing broader choice and the potential for customers to have a greater influence on future designs based on their specific workload requirements.

The early ARM-based server experiments were based on 32-bit cores from Marvell and Calxeda. The market was interested in the possibility of ARM-based in servers but decided against supporting 32-bit operating systems based on anticipated delivery dates for the first generation of 64-bit ARMv8 server SoCs and production servers. The first 64-bit ARMv8 licensee to ship functional processors was AppliedMicro who began shipping development kits based on their 64-bit ARMv8-based X-Gene product to key customers in late 2013. AppliedMicro is now shipping production silicon to key customers and partners, and 64-bit ARMv8-based solutions are now on track to ramp into production in the next 12 months.

FIGURE 2: APPLIEDMICRO X-GENE



ARM Holdings licenses their chip designs and the instruction set architectures to third parties who design their own products that implement one of those architectures—including systems-on-chips (SoC). ARM Holdings also offers an "architectural license" for their instruction sets, which allows the licensees to design their own cores that implement one of those instruction sets. AppliedMicro holds an architectural license for the 64-bit ARMv8 architecture. AppliedMicro's X-GENE is a SoC solution that combines 10/40Gbps mixed signal I/O with eight (8) 64-bit ARMv8 cores running at up to 2.4GHz with an enterprise-class memory subsystem. AppliedMicro's goal with X-GENE is to marry server-class thread level performance and feature set with mobile-pedigree power savings capability.

64-BIT ARMV8 SERVER SOFTWARE ECOSYSTEM

As is true with every new microarchitecture introduction, one of the primary drivers of market readiness is software ecosystem support. The Linaro Consortium was established in 2010 to help drive development of the open source ARM software ecosystem. The Linaro Enterprise Group (LEG) was formed in 2012 specifically to focus on server OS enablement across 64-bit ARMv8 processors with supporting members from ARM Holdings, silicon partners, and system providers. AppliedMicro was a founding member of LEG and has been on the forefront of 64-bit ARMv8 server ecosystem development. Their time-to-market advantage has resulted in the majority of 64-bit ARMv8 server development and optimization work thus far to be accomplished using X-GENE-based hardware.

Currently, LEG has enabled the creation of a base set of open source and commercial software for 64-bit ARMv8 servers including basic tools: compilers, libraries, debuggers, etc. The latest UEFI (Uniform Extension Firmware Interface) specifications now provide support for 64-bit ARMv8. UEFI boot support is a requirement for enterprise operating systems such as Red Hat and Microsoft Windows server today on x86. UEFI solutions

for 64-bit ARMv8 are expected to be commercially-available in the coming months, so server vendors will be able to provide the same capability on their ARM-based servers.

In terms of operating systems, Canonical announced official support for Ubuntu 14.04 LTS on 64-bit ARMv8 in April 2014. Red Hat began their support for 64-bit ARMv8 on Fedora over a year ago and recently announced their Red Hat ARM Partner Early Access Program (PEAP). Hypervisor support is now available for KVM and Xen. Java optimizations are underway for both OpenJDK and Oracle JRE with general availability of Oracle JDK 8 for 64-bit ARMv8 on track for the first half of calendar 2015. The investments and public support of Canonical, Oracle, and Red Hat provide significant validation in the potential of the 64-bit ARMv8 architecture for servers.

At the application layer, there are a number of open source and commercial application projects underway to support and optimize 64-bit ARMv8 servers in target workload areas. These include LAMP, web serving/caching (Nginx, Apache, Memcached), Apache Hadoop, OpenStack and especially its compute component (Nova), enterprise applications (SugarCRM, Elastic Search), and storage (GlusterFS, Ceph).

INTRODUCING THE HP PROLIANT M400 FOR MOONSHOT

HP ProLiant m400 server cartridge will be the first enterprise-class production 64-bit ARMv8 server in the market with shipments beginning in calendar Q4 2014. The cartridge has one AppliedMicro X-Gene SoC with 8 cores running at a frequency of 2.4GHz. Each of four core pairs share a 256KB L2 cache, and all 8 cores share an 8MB L3 cache.

Two key advantages of this product, compared to other ProLiant cartridges for the Moonshot System available today, are a doubling of the addressable memory to 64GB per cartridge and significantly higher memory bandwidth made possible by X-Gene's four memory channels (two DDR3L-1600 SO-DIMMs per channel for a total of eight x 8GB DIMMs per cartridge).

FIGURE 3: HP PROLIANT M400 SERVER CARTRIDGE: FRONT AND REAR VIEWS



Another advantage of the ProLiant m400 is improved throughput for I/O intensive workloads with two 10Gbps Ethernet channels using the Mellanox Connect-X3 Pro Dual 10GbE NIC. The cartridge includes low latency storage access with one M.2 flash storage module for local OS booting with configurations that support 120GB, 240GB, and 480GB capacities. iSCSI (software initiator) external storage capability is also available. The ProLiant m400 currently supports the Ubuntu 14.04 OS for 64-bit ARMv8. Typical power draw per is 55W per cartridge including processor, NICs, memory, and everything else on the cartridge.

An HP Moonshot 1500 chassis packed with ProLiant m400 server cartridges will offer significant compute density advantage over a traditional rack-mount server with on average 4X more web-serving compute in the same power, cooling, and space plus 2X the memory of today's typical mid-range servers.² For workloads that are I/O intensive, the ProLiant m400 provides a nice balance of I/O throughput, memory bandwidth and capacity, and compute capability.

PROLIANT M400 TARGET CUSTOMERS AND WORKLOADS

The HP ProLiant m400 server cartridge is built for those customers who are looking for server-class performance while maintaining focus on power efficiency. HP is positioning the ProLiant m400 to address the following target workloads and customers:

² Based on HP's internal testing of web workloads running HP Moonshot with the ProLiant m400 Server Cartridge in a 10GbE network environment compared to an x86 1U rack-mount server.

- **Service providers, Co-Locators (CoLo), and mid-sized commercial internet providers** that need to save money on power and other infrastructure costs, while maintaining as small a physical footprint as possible.
- **Early adopter customers** including government, education, and scale-out service providers looking at ARM-based servers for **web front-end environments**, where advantages in performance-per-dollar and performance-per-watt are critical.
- **Software developers for ARM-based mobile devices** looking for enterprise-class reliability and platforms from a qualified, trusted server vendor. Currently, application developers for mobile devices are using bread-racks full of PCs with tablets/phones attached via USB to test their solutions at scale; they could use the ProLiant m400 to provide a “developer platform as a service” to improve efficiencies.
- **Software developers for ARM-based servers** looking to achieve fast turnaround of software that has been developed on a fully validated, production-class 64-bit ARMv8 system.
- **Enterprise customers deploying private/hybrid cloud models looking for an alternative to x86** which often requires tradeoffs between performance and price/power. One of the key promises of 64-bit ARMv8 is the opportunity to increase performance while keeping both prices and power use low.

The initial featured solution stack for the ProLiant m400 is a web-serving/caching solution based entirely on open source components:

- [Canonical Ubuntu 14.04 LTS OS](#)
- [Nginx web serving tier](#)
- Generic web application tier (Perl, Python, and Java apps)
- [Memcached in-memory caching](#)
- [MySQL database](#)

We expect that the software support and ecosystem will continue to evolve for ProLiant m400 as beta customers and software developers prepare for larger scale production use of 64-bit ARMv8 servers over the next 12 months.

WHERE THE RUBBER MEETS THE ROAD: EVALUATING TOTAL COST OF OWNERSHIP (TCO) FOR THE PROLIANT M400

The key factor that will drive adoption of ARM-based servers will be whether or not they can deliver measurable value for customers when compared to their existing server solutions. For the majority of large scale datacenters, this value is measured in total cost of ownership (TCO) which includes both the acquisition cost of the equipment and the operating cost of using this equipment during its lifetime.

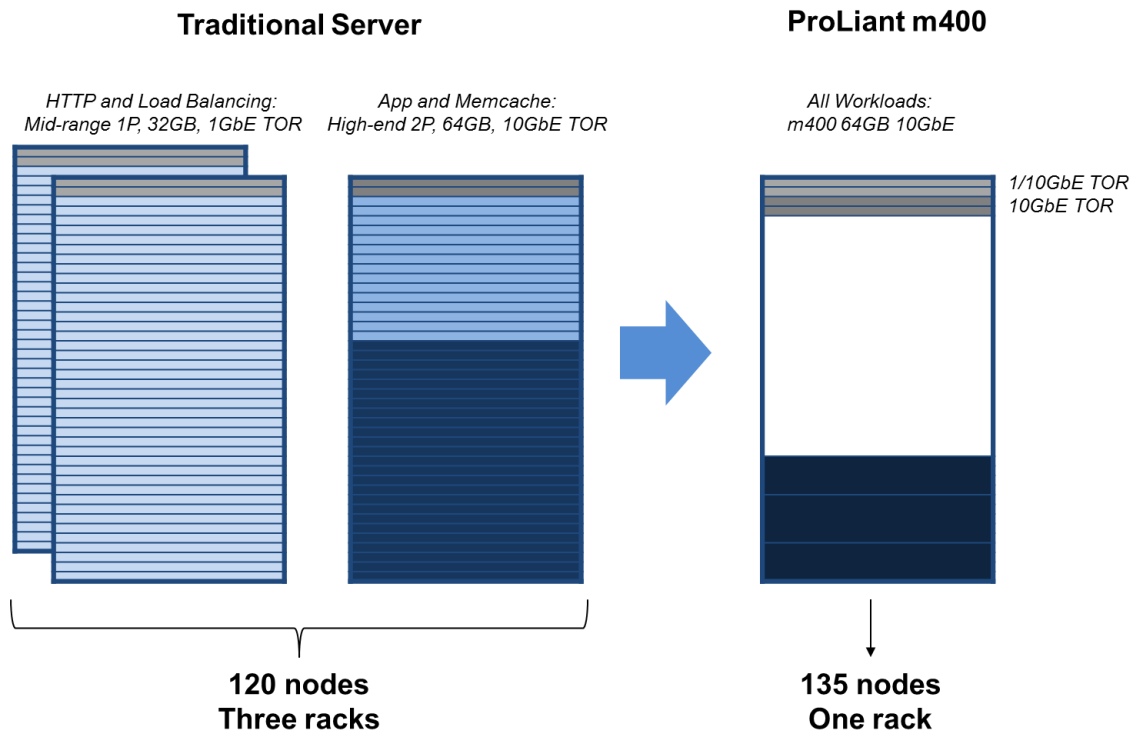
The TCO equation is different for every datacenter customer; there is no one-size-fits-all methodology that will tell the whole story. While the data in this analysis could be used as starting point to understand the potential value of the ProLiant m400, we expect that all customers looking at multi-rack scale deployments of 64-bit ARMv8-based servers will perform detailed evaluations that take into account their own specific workloads and other variables for their datacenter environments.

THE SCENARIO

The scenario for this TCO analysis is a mid-sized commercial internet provider servicing their web requests using a typical three-tier web infrastructure (web tier, application/caching tier, and database tier). In this scenario, the model focuses on using the ProLiant m400 to replace the existing x86-based 1U web tier and application/caching tier while continuing use of a standard x86-based configuration for the database tier. To simplify the comparison, the specifics of the database tier were removed from the analysis. This scenario also assumes a 12kW per rack power budget.

Figure 4 provides a visual representation of the two comparable configurations for the web tier and application/caching tier used in this analysis.

FIGURE 4: TCO ANALYSIS: TRADITIONAL AND ARM-BASED CONFIGURATIONS



TRADITIONAL SERVERS

On the traditional server side, different servers were used in the client-facing tier (Mid-range) versus the application/caching tier (High-end). Many datacenter customers often standardize on configurations within a given tier. In this case, all mid-range servers for the front end are identical, and all high-end servers for the application/caching tier are identical.

The “Mid-range” server configuration for the web tier included:

- 40 1U rack-mount servers per rack (2 racks)
 - 1P Xeon E3 1200 v3
 - 32GB memory
 - 1 SSD or 1 HDD
 - On-board 1GbE LOM
 - Single power supply
 - 1 year Linux
- 2 1GbE top-of-rack switches per rack (4 total)

The “High-end” server configuration for the application/caching tier included:

- 40 1U rack-mount servers per rack (1 rack)
 - 2P Xeon E5 2600 v3
 - 64GB memory
 - 1 SSD or 1 HDD
 - One 10GbE NIC
 - Dual hot-plug redundant power supplies
 - 1 year Linux
- 2 10GbE top-of-rack switches per rack (2 total)

ARM-BASED CONFIGURATION

On the ARM-based server side, identical nodes were used across all applications which provides a potential advantage for easier load balancing across servers. The integrated switches in the Moonshot System also allow for a reduced number of TOR ports required when compared to the traditional 1U server solution.

The ARM-based configuration included:

- 45 ProLiant m400 Cartridges per chassis (3 chassis)
 - 1P 64-bit ARMv8-based X-Gene
 - 64GB of memory per cartridge
 - 1 120Gb M.2 solid state boot device
 - 1 dual-port 10GbE NIC
 - 1 year Linux
- Chassis includes two Moonshot 10GbE switches and 4 hot-plug redundant power supplies
- 4 top-of-rack switches in the rack cabinet (two 10GbE and two 1GbE)

PERFORMANCE ASSUMPTIONS

One of the key underlying assumptions of this analysis is that the three (3) racks of traditional 1U servers (120 nodes) can be replaced with the three (3) HP Moonshot 1500 Chassis with ProLiant m400 Server Cartridges (135 nodes) and still meet the same Service Level Agreements (SLAs) for these workloads. This assumption is based on the following set of criteria:

- **Web Tier** Based on guidance from HP's internal performance analysis for static web workloads, one ProLiant m400 server cartridge was used to replace one mid-range Xeon E3 1U server. In addition, the ProLiant m400 provides twice the memory capacity for those web workloads that require more memory.
- **Application Servers** Depending on the workload, more than one ProLiant m400 server may be needed to perform the same work as a heavily-used application server. Based on guidance from HP's internal performance analysis for dynamic web workloads, two ProLiant m400 cartridges were used to replace one Xeon E5-based application server for this illustration.
- **Caching Servers** HP has received feedback from customers that real world SLAs for Memcached are driven by cache hit versus cache miss rates rather than metrics similar to synthetic benchmark performance. Since both servers are configured with the same memory capacity (64GB), one ProLiant m400 cartridge was used to replace one Xeon E5-based Memcached server.

HP encourages customers to use their own metrics to assess the specific performance comparisons that may be applicable for their workloads. In order to help facilitate this, HP has a number of Discovery Labs around the world to allow customers, partners, and the software development community to test and analyze application performance on various HP Moonshot configurations as well as on traditional servers. The ProLiant m400 has been available in the Discovery Labs for a number of months, and beta systems also began shipping to key customers and partners in the summer of 2014.

The typical synthetic benchmarks traditionally used to measure server performance are not representative of the scale-out workloads that customers will deploy on high-density architecture like HP Moonshot. With this in mind, HP and others in the industry are looking to find more representative ways to evaluate performance for these new classes of workloads that accurately represent overall system throughput capability versus being artificially constrained by one element (compute, I/O, memory) within the server.

COST COMPARISONS

With the assumption that these two configurations offer similar performance for this set of workloads, acquisition and operating costs were calculated for each configuration to compare TCO.

Acquisition cost comparisons included:

- Comparable list prices for the traditional 1U servers (average actual) and the ProLiant m400 servers (at launch). All server prices included a 1 year Linux subscription.
- Comparable list prices for 1GbE and/or 10GbE top-of-rack switches required for each configuration
- Typical server rack costs

Today's total acquisition cost of the traditional 1U server solution is \$709,057. Launch pricing of the ProLiant m400-based solution is \$506,800. The difference is a 29% savings in acquisition costs driven by the reduced number of top-of-rack switches with the ProLiant m400 solution, lower average price-per-server, and fewer number of racks required.

Operating costs included:

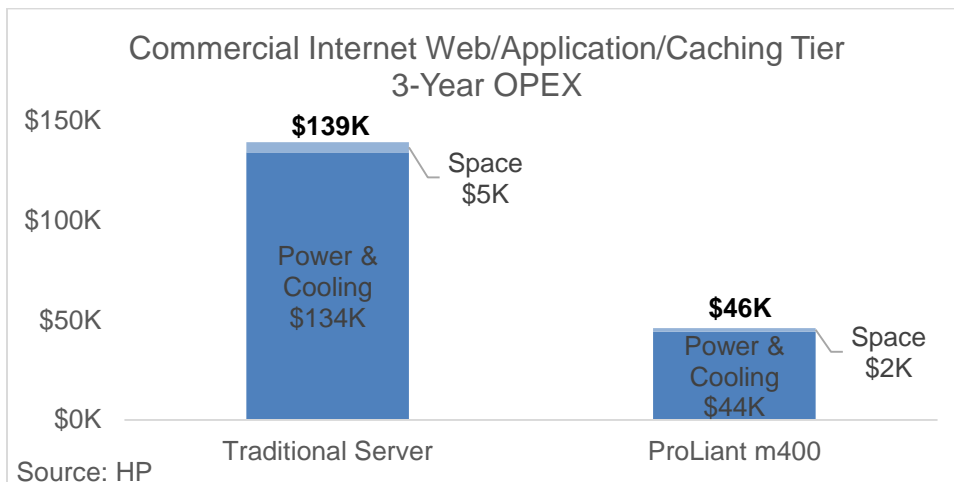
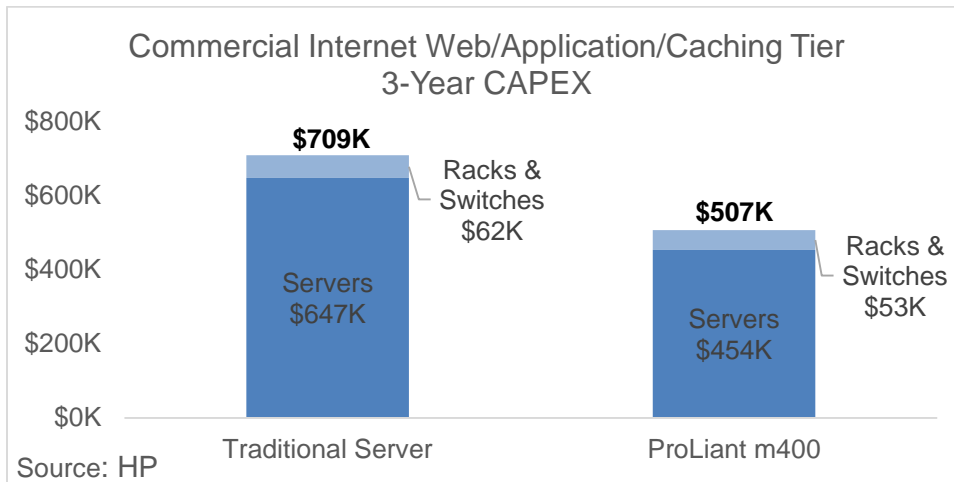

- 3 year power and cooling costs based on total rack power consumed using \$0.10 per kWh and 1.7 PUE.
- Cost per square foot of rack space.

The total operating cost over 3 years for the traditional 1U server-based solution would be \$139,068 while the ProLiant m400-based solution would be \$45,909. The resulting 67% savings in operating costs is driven by lower total power of the ProLiant m400 solution and fewer number of racks required.

THE BOTTOM LINE

Figure 5 provides a summary of the total cost of ownership comparisons for the traditional 1U server solution versus the ProLiant m400 solution.

FIGURE 5: TOTAL COST OF OWNERSHIP SUMMARY



In summary, this scenario illustrates a 35% potential savings when replacing traditional 1U servers with ProLiant m400 servers for web and applications tier in a commercial internet environment. In addition to these savings, the ProLiant m400 for Moonshot offers 66% fewer racks for the same performance and identical infrastructure across tiers to help ease load balancing.

EARLY CUSTOMER FEEDBACK: SANDIA NATIONAL LABS

One of the customers who had early access to the ProLiant m400 were Sandia National Labs researchers. They used the solution to develop a test bed for the Mantevo Project. The Mantevo Project is focused on developing tools to accelerate and improve the design of high performance computers and applications by providing application and library proxies to the high performance computing community. As a part of this project, a set of mini-applications (small self-contained proxies for real applications) have been developed to evaluate application performance tradeoffs in hardware platforms, runtime environments, compilers, languages, algorithms, and more.

Sandia researchers became interested in ARM-based servers for the Mantevo Project due to the potential performance benefits for their mini-applications. The lab researchers also believe there could be a significant benefit to the multi-vendor silicon development model for ARM with a lower barrier to influence future hardware concepts. According to limited tests by Sandia researchers, the AppliedMicro X-Gene-based ProLiant m400 supports more memory bandwidth per unit of processing capability than an equivalent, conventional x86 processor.

Sandia researchers will look to scale their testing to more nodes in the coming months with the vision of rolling out thousands of ARM-based server nodes in future years, as interconnect and memory technologies scale to meet the needs of their HPC workloads. For those who are interested in replicating the Sandia analysis, or learning more about the specific application areas tested, the Mantevo mini-applications are open source and [available for download](#).

CALL TO ACTION

The HP ProLiant m400 server cartridge will be the first enterprise-class production 64-bit ARMv8 server in the market with shipments beginning in calendar Q4 2014. Initial TCO studies and early customer feedback indicate this solution could be promising for customers who are looking to deploy ARM-based servers in web environments and also for server/mobile software developers who are interested in evaluating the performance

and TCO potential of 64-bit ARMv8 processors for their applications. Customers who fit this profile should consider taking the following next steps:

- Determine which workloads in your datacenter are a good candidate for initial evaluation of the ProLiant m400. Ideal targets would be applications that are I/O or memory bandwidth intensive versus CPU intensive.
- Evaluate the 64-bit ARMv8-based server software ecosystem readiness for your specific target solution stack to ensure key components are in place to begin evaluation.
- Work with HP to get access to the ProLiant m400 for testing and evaluation of the solution in your environment via their Discovery Labs or beta/evaluation programs.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHORS

[Gina Longoria](#), Analyst at [Moor Insights & Strategy](#)
[Paul Teich](#), CTO & Senior Analyst at [Moor Insights & Strategy](#)

REVIEW

[Patrick Moorhead](#), President & Principal Analyst at [Moor Insights & Strategy](#)

EDITOR

[Scott McCutcheon](#), Director of Research at [Moor Insights & Strategy](#)

INQUIRIES

Please contact us here if you would like to discuss this report and Moor Insights & Strategy will promptly respond.

CITATIONS

This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

HP and Applied Micro are research clients of Moor Insights & Strategy and this paper was commissioned by HP and Applied Micro. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2014 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.