MOOR
INSIGHTS & STRATEGY

HP Apollo 6000 System: Density-Optimized Server Packs a Performance Punch

Rack-scale Throughput for HPC, Hosting, and Private Cloud

Executive Summary

Hyperscale and enterprise datacenters have become increasingly conscious of efficiency, ensuring the optimal amount of hardware resources are dialed-in to the specific needs of their workloads. For many workloads, density-optimized servers have evolved to deliver on the promise of OPEX savings (power/cooling, space) without sacrificing system performance.

HP is expanding their density-optimized portfolio with the release of the Apollo 6000 System. This new server is designed to address the needs of lightly-threaded HPC applications like Electronic Design Automation (EDA) and Monte Carlo Simulations (used for financial risk modeling and various engineering/scientific applications). The system offers high per-thread performance, robust network bandwidth, and rack-level shared infrastructure for efficiency.

The lightly-threaded HPC segment provides a nice entry point for the Apollo 6000 System, but it is relatively small. To increase market penetration for the Apollo 6000 System, HP's future trays should target other HPC and web scale workloads with differentiated capabilities like workload accelerators, storage options and other processor types.

Scale-Out Servers: One Size Doesn't Fit All

Less than a decade ago, x86 server form factors were relatively simple, falling into the basic categories of racks, towers, and blades. However, as various new form factors emerged to meet the needs of scale-out datacenters, the segmentation and descriptive terminology became increasingly complex. To complicate matters further, the advent of software-defined and workload-specific servers requires solutions to be optimized on various vectors such as compute node performance, storage capacity, energy efficiency, density, modularity, and network throughput. General-purpose standard servers are becoming an undifferentiated commodity and specialty workload-optimized solutions are winning, even in traditional enterprise. MI&S has created a new server segmentation model to address the evolving landscape.

Density and Throughput: Can They Coexist?

Modular and dense server form factors were originally introduced to address the datacenter floor space and power/cooling costs of hyperscale datacenters. These form factors were effective in creating TCO savings for a subset of web services and HPC workloads. However, for many workloads, the limitations of system performance and throughput from dense solutions just required more servers to be deployed—which eliminated the potential benefit of these more "efficient" solutions.

Server vendors continue to look at ways to broaden the reach of density-optimized solutions, and now, for a larger percentage of workloads, density no longer requires the sacrifice of system throughput. To expand the use of density-optimized servers even further, a broader set of configurations with various allocations of resources (network, storage, compute, *etc*.) that match the needs of specific workloads must be made available.

For compute, there are a variety of processor architectures used in dense and modular servers—all the way from "wimpy" ARM/Atom cores and standard Intel Xeon cores, up through workload-specific accelerators like GPUs (Graphics Processing Units), coprocessors, and DSPs (Digital Signal Processors). However, when evaluating the right solution to match a specific workload, it is more than just the processor choice and sockets-per-rack that count. As modern day scale-out workloads become increasingly dependent on rack-level throughput, it is just as important to ensure a server solution has the proper amount of network bandwidth to avoid communication bottlenecks.

HP's Scalable Server Portfolio

HP offers a series of density-optimized servers for hyperscale and high performance computing with a range of capabilities to meet the needs of different workloads and customers.

Table 1 provides an overview of these products.

Table 1: HP Scalable Server Portfolio Overview

	Key Attributes	Chassis Form Factor	Processor Architecture	Target Workloads
Moonshot	Energy efficiency, workload-optimized configurations	4.3U with up to 180 compute nodes	Intel Atom, AMD Opteron APUs, ARM (future)	Web serving, web hosting, hosted desktop
Proliant SL2500	Modularity, density, flexibility	2U with 4x2P compute nodes	Xeon E5-2600 v2	Web services, hosting, mainstream computing, HPC
Proliant SL4500	Storage density, serviceability	4.3U with 1-3 compute nodes & up to 60 drives	Xeon E5-2400	Big Data, object storage, exchange and parallel data processing
Proliant SL6500	Dense compute, workload acceleration	4U with up to 8x2P compute nodes, GPU/ coprocessor options	Xeon E5-2600 v2, NVIDIA Tesla/ GRID GPUs, Xeon Phi	HPC
Apollo 6000 NEW	Per-thread performance, throughput, density	5U with 20 compute nodes	Xeon E3-1200 v3	Lightly threaded HPC, hyper-virtualization, hosting
Apollo 8000 **NEW**	Innovative cooling (liquid) to improve efficiency/density	72 trays in a rack with up to 2x2P compute nodes per tray, GPU/ coprocessor options	Xeon E5-2600, NVIDIA Tesla GPUs, Xeon Phi	HPC

The Apollo 6000 and Apollo 8000 Systems are the newest servers in HP's Scalable Product line. The Apollo 6000 System is the first of this portfolio to use Xeon E3 Processors V3 Family. It bridges the gap between the high performance HP ProLiant SL servers mainly built around Xeon E5 processors (larger cores, dual socket) and the HP Moonshot servers that deliver higher density but use Intel Atom, AMD Opteron APU (accelerated processing unit), and ARM processors (smaller cores, single socket). Apollo takes a similar approach to the Moonshot family with servers that are purpose-built and modular to tackle the specific resource requirements of various workloads. With Apollo's focus on larger cores and Moonshot designed for energy efficiency, these products target different workloads. They are well-suited to complement each other in scale-out datacenter deployments.

With Xeon E3's high performance-per-thread, HP is positioning the Apollo 6000 System for HPC workloads like Electronic Design Automation (EDA) and Monte Carlo Simulations (used for financial risk modeling and various engineering/scientific applications). These workloads tend to be lightly-threaded, requiring higher speed with fewer threads per server. The Apollo 6000 System also makes sense as a building block for highly virtualized IT environments looking to move to private cloud. For hosters and service providers, single socket servers are often the platform of choice, as they provide a better mechanism for managing customers' virtually hosted servers.

Key System Advantages

The Apollo 6000 System's balance of efficiency, density, and throughput make it a strong addition to the HP Scalable portfolio for compute-intensive, lightly-threaded applications, hosting, and Hyper-Virtualization. HP claims for an EDA workload, a rack scale deployment of the Apollo 6000 System provides a 4X performance/$/watt advantage over Dell M620 blades.

Below are some of the key advantages of the Apollo 6000 System.

- **Per-Thread Performance** Intel Xeon E3-1200 V3 processor based on the Haswell architecture offers the highest frequency per core/thread (up to 4GHz Max Turbo Frequency) of any Xeon-class solution. Intel leads their Xeon product roll-outs with the one-socket server family, which means the Xeon E3 Haswell-based product line is currently one generation ahead of Xeon E5 Ivybridge-based products. This newer architecture offers both power savings improvements and better compute efficiency. In addition to the benefits of Xeon E3 performance, one-socket servers can provide lower latency than two-socket, since there is no 2P cache coherency required.
- **Network Throughput** The HP Apollo 6000 System includes an I/O Innovation Zone which provides flexible configuration options for networking (1G, 10G, InfiniBand) depending on an application's needs. For workloads that need both the per-thread performance of Xeon-E3 and significant network bandwidth per rack, a system configuration with 10GbE could provide up to 1600 GbE Uplink bandwidth per rack.

Figure 1: Dual Rear ALOM (LAN on Motherboard)

- **Density Optimization** the HP Apollo 6000 System delivers a total rack density of 140 nodes in a standard 42U rack, or 160 nodes in a 47U rack. HP achieves this level of density by using shared infrastructure (enclosures, PCBs, fans, I/O, management module) and pooled power resources (Apollo 6000 Power Shelf). Apollo 6000 System's density is similar to Xeon E5-based SL2500 and SL6500 solutions. The system is front-serviceable and rear-cabled, providing servicing efficiency for standard datacenters**.**

Figure 2: HP Apollo 6000 System 47U Rack

- 8x 5U Enclosures
- 80x Compute Sleds (2 nodes per)
- 2x Power Shelves
- 4x 48 Port Network Switch

HP Proliant XL220a Compute Tray Overview

The HP Proliant XL220a is the first compute tray available for the Apollo 6000 System. Each vertically-mounted tray includes up to 2 independent server nodes based on Xeon E3-1200 v3 Series processors and the Intel C222 (Lynx Point) chipset with a maximum power of 169W per tray. Each Xeon E3 server node can support up to 4 DDR3 1600MHz un-buffered DIMMS with up to 32GB max memory per node (64GB max per tray). While this memory footprint is significantly lower than a standard Xeon E5-based system, 4 GB/core is sufficient for the Apollo 6000 System target workload classes.

Figure 3: XL220a Server

As mentioned earlier in this paper, one of the unique advantages of the HP Apollo 6000 System is the robust networking capability available with the I/O Innovation Zone—a modular unit that plugs into the back of each server tray. Each tray has one Dual Rear ALOM (FlexLOM) with up to 8 ports per module. Options include 1Gb, 10Gb, and InfiniBand, allowing users to tradeoff between performance and cost depending on the networking requirements of their workload. Each independent I/O module can be configured differently and modified as workload needs change.

For storage, the Apollo 6000 System includes internal storage of up to 4 SFF Hot Plug HDD or SSD (2 per node) and an Optional Smart Array (HBA) via the IO slot. System Management capabilities are based on HP iLO4 with dedicated iLO for each tray and HP APM 2.0 (Advance Power Management)—making it straightforward for existing HP customers to manage Apollo 6000 servers in their existing management frameworks.

While the Proliant XL220a is the only officially released tray for the Apollo 6000 System at this time, HP indicates there could be other configuration options in the future that include compute accelerators or other capabilities to address the needs of additional workloads. This will help expand the market opportunity for Apollo 6000 as the lightly-threaded HPC segment is relatively small. We expect HP to use a phased approach for introducing new Apollo 6000 trays similar to what they have done with Moonshot: new workload-optimized cartridges for the existing Moonshot 1500 Chassis are launched every few months.

HP Apollo a6000 Chassis and Power Shelf Overview

The Apollo a6000 chassis is a 5U enclosure that holds up to 10 single compute trays (20 compute nodes) vertically. The chassis fits in a standard 19" rack and requires 1.0m depth per rack. NICs are rear-cabled and 5 x80mm redundant fans per chassis are shared across the trays for maximum cooling efficiency. The chassis includes a management module with redundant 1Gb NICs.

Figure 4: Apollo a6000 Chassis Connected to Apollo 6000 Power Shelf (rear view)

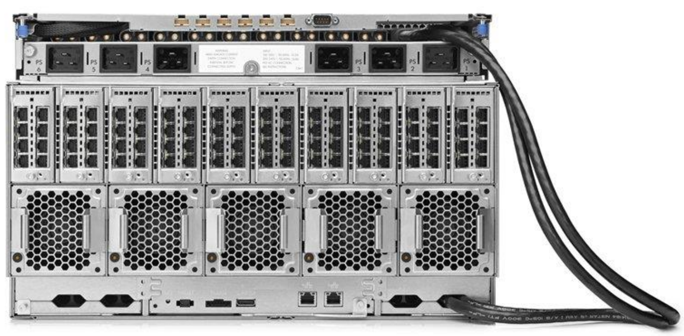

The Apollo a6000 chassis does not have an internal power source but rather relies on connections from the chassis to the HP Apollo 6000 Power Shelf which provides shared power across enclosures for efficiency and redundancy. The Apollo 6000 Power Shelf is 1.5U tall, and supports up to 6 power supplies (14.4kW or 15.9kW non-redundant or configurations that include N+1, N+N redundancy). One power shelf can support 3-4 fully loaded enclosures or up to 6 moderately provisioned chassis.

HP's Advanced Power Manager allows an IT administration to see and manage shared infrastructure, server, chassis, and rack-level power from a single console. For those datacenters with sophisticated power management capabilities, HP APM provides the ability to use dynamic power allocation and capping based on usage requirements.

Summary

Density-optimized servers are addressing an increasing number of web and HPC workloads for customers that require a balance of OPEX (power, space) with overall system throughput. HP continues to expand their density-optimized platform portfolio with workload-specific solutions to meet the requirements of a broader set of customers.

As the vast majority of the workloads in the datacenter are running on dual CPU platforms, the HP Apollo 6000 System is not necessarily a replacement for all of the servers in a datacenter. But it is an important additive element to help drive a better blended efficiency across workloads in the datacenter. The system's high per-thread performance, robust network bandwidth, and rack-level shared infrastructure for efficiency make it a strong potential fit for lightly-threaded HPC applications, hosting, and hyper-virtualized environments. HP's planned introduction of future server trays with other capabilities should help expand the market opportunity for the Apollo 6000 System in the future.

For more information: http://www.hp.com/go/Apollo

Important Information About This Paper

Author
Paul Teich, CTO & Senior Analyst at Moor Insights & Strategy
Gina Longoria, Analyst at Moor Insights & Strategy

Review and Publisher
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

Disclosures
HP is a research client of Moor Insights & Strategy and this paper was commissioned by HP. No employees at the firm hold any equity positions with any companies cited in this documented.

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.