THE OPEN COMPUTE PROJECT PROVIDES INSPIRATION, BUT DELIVERY IS UNCLEAR

COLLABORATION OF INDUSTRY LEADERS INSPIRES OPEN, EFFICIENT DESIGN METHODOLOGY BUT IS NOT YET DELIVERING ON PROMISE FOR MAINSTREAM ENTERPRISE IT

EXECUTIVE SUMMARY

The Open Compute Project (OCP) was developed by Facebook and other hyperscale computing leaders in 2011 to help the datacenter hardware community embrace open, efficient, and cost-optimized design principles. Since then, other industry leaders such as Microsoft, Apple, Google, financial services firms, and telecomm providers have joined OCP to take part in the collaboration and community efforts. One of OCP's key goals has been to bring design and supply chain efficiencies to mainstream enterprise IT by making hyperscale-inspired designs available to a wider array of end users.

OCP has done an excellent job creating awareness around the benefit of cost-optimized, efficient platforms and the overall value they can bring to datacenters around the world. However, after 5 years since the project's inception, OCP production-scale deployments are still limited to a small number of hyperscale players and, with a few exceptions, have not yet made their way into mainstream enterprise IT and smaller scale service providers. Customers looking to adopt OCP designs in their datacenters should go in with eyes wide open and be sure to evaluate hyperscale-inspired design choices from both OCP vendors and other leading global datacenter hardware providers.

MARKET DRIVERS FOR "OPEN" TECHNOLOGIES

Many in the industry use the term "hyperscale" to describe datacenter deployments at web giants (Amazon, Facebook, Google, Microsoft, Alibaba, Baidu, Tencent, and other large tier 1 service providers) and the buying patterns they use to support their IT operations. These top tier hyperscale buyers define leading-edge operational efficiency. They have the purchasing power to request custom and semi-custom hardware, and they do so when the up-front engineering investment will be offset by operational cost savings and performance benefits to help their businesses run more efficiently. Their scale means that saving a little power or slightly decreasing processing time often does justify the up-front resources to develop a custom design or build a custom datacenter from scratch. They claim efficiency targets that other organizations seek to emulate, and they leave a trail of new IT infrastructure products and open source code in their wake.

Enterprise IT organizations and cloud service providers outside the hyperscale category are under more pressure to improve efficiency and provide competitive differentiation for the business. Costs are increasing, and finding expertise is challenging, yet they are required to drive innovation and align with new business opportunities. As a result, minimizing capital expenses (CapEx) and operating expenses (OpEx) is a key priority for IT, and enterprise IT organizations are seeking disruptive innovations or new business models to achieve their goals. These organizations are paying close attention to hyperscale buyers' actions and are looking for ways to leverage hyperscale best practices in their own IT environments. To optimize visibility, improve efficiencies, and achieve collaborative design opportunities with the datacenter supply chain, many IT buyers are demanding technologies that are considered "open".

To meet the needs of both hyperscale and mainstream IT buyers, new IT supply chain business models have emerged over the last decade to provide "openness" and technology choice for both software and hardware solutions. On the software side, Linux has gained server market share versus Windows and other traditional licensed operating systems. There have been an increasing number of open source projects adopted and driven by increased interest in hyperscale purchasing. OpenStack has experienced significant momentum over the last few years to become the *de facto* standard open source cloud operating system. Initiatives like the Open Compute Project and Project Scorpio (now called the Open Data Center Committee or ODCC) have emerged as collaborative efforts to define open hardware platforms.

THE OPEN COMPUTE PROJECT HISTORY & VISION

The Open Compute Project (OCP) began in 2011 when Facebook, under the leadership of Frank Frankovsky, gathered a small group of industry leaders to create a collaborative movement for datacenter hardware development similar to the open source software model. Facebook, who shared its initial designs with the public, was joined by Intel, Rackspace, Goldman Sachs, and Andy Bechtolsheim to create the project and incorporate the Open Compute Project Foundation. The Foundation's stated project mission is as follows:

> *"The Open Compute Project Foundation is a rapidly growing, global community whose mission is to design, use, and enable mainstream delivery of the most efficient designs for scalable computing."*

OCP's mantra since the project's inception has been to improve hardware design efficiencies via open standards to help IT organizations lower costs, streamline the supply chain, and avoid proprietary technologies that create vendor lock-in. OCP technology innovation has been driven in large part by hyperscale computing focused providers. However, the OCP community's goal is to bring hyperscale best practices and efficient design principles beyond just the largest hyperscale players to mainstream enterprises and service providers.

OPEN COMPUTE TODAY

The OCP community has done an excellent job educating the market about the benefits of open, non-proprietary, cost-optimized hardware design. Their battle cry encouraging the supply chain to steer clear from "gratuitous differentiation" (anything that doesn't add value should be removed) has been heard and resulted in a movement toward no-frills, cost-optimized platforms for the hyperscale market. OCP has provided an avenue for like-minded customers to collaborate and establish metrics of innovation based on their own needs versus relying on manufacturers to set the pace of innovation. Vendors have also become more open minded about providing engineering services for large-scale customers and working with the supply chain on custom and semi-custom solutions that align with customer needs.

STRONG LEADERSHIP, BUT A LONG WAY TO MAINSTREAM ADOPTION

In 5 years, the OCP community has grown and diversified with a global presence and public support from industry leaders. The 2016 OCP Summit (the community's annual global event) included 2,400 people from 600 companies. What began as an effort largely driven by Facebook has evolved into something much broader.

- **Facebook**, the original OCP content contributor, is still a leading participant.
- **Rackspace**, a founding OCP member, has contributed many innovative designs to OCP. About half the systems used for Rackspace's cloud services are based on OCP designs, with a move toward 80-85% in the coming years.
- **Microsoft** joined OCP in 2014 and contributed the Open Cloud Server (OCS) design, which is the basis for its deployments today. Microsoft claims publicly that over 90% of the servers it installs today are based on the OCS spec.
- **Apple** joined OCP in 2015, but specific plans are not yet public.
- **Google**, who historically has been secretive about its hardware infrastructure IP, joined OCP in 2016. Google plans to donate a 48V rack spec and is expected to collaborate with Rackspace on an IBM POWER9 based server in the future.

- **Large financial services firms** such as Goldman Sachs, Fidelity Investments, JPMorgan Chase, Capital One, and Bank of America have collaborated on OCP designs and created versions compatible with standard 19" racks. A number of these organizations have been part of OCP since the project's inception. While many of these users remain quiet about their specific OCP plans, Goldman Sachs has committed that more than 80% of servers it deploys will be based on OCP. Fidelity claims OCP hardware reduced its datacenter energy bill by 20%.
- **Major telco companies** such as AT&T, Deutsche Telekom, EE, SK Telecom, and Verizon formed the OCP Telco Project in 2016. A number of these companies are looking to drive hardware and software standards around Network Function Virtualization Infrastructure (NFVi) and software-defined networking systems in their next generation datacenters.
- **The high-performance computing (HPC) community** has recognized OCP as a potential avenue for cost-effective, standardized systems. For example, the US Department of Energy's National Nuclear Security Administration selected an OCP design from Penguin Computing for a large installation to serve Lawrence Livermore, Los Alamos, and Sandia National Laboratories.
- **More points-of-entry for global customers** have been created with OCP enablement around the world. For example, ITOCHU Techno-Solutions Corporation, Murata Manufacturing, and NTT Data Intellilink are jointly developing an OCP compatible rack system targeted for Japanese datacenters.
- **Storage & networking** focus areas have been added to OCP over the last couple of years, which was originally focused primarily on servers.

Participation by these industry leaders may indicate many large-scale datacenter customers view OCP as a beneficial way to collaborate to create more efficient architectures. However, despite this momentum, the size of the actual OCP footprint in datacenters around the world remains relatively small, with a limited number of organizations embracing OCP today in their production scale deployments.

Is There an OCP Standard?

One of OCP's goals is to improve design and supply chain efficiency by leveraging common specifications and standards across the community. For example, the OCP Server Project emphasizes the importance of standards.

> *"The OCP Server Project provides standardized server system specifications for scale computing. Standardization is key to ensure that OCP specification pool does not get fragmented by point solutions that plague the industry today."*

However, OCP has had a difficult time adhering to a finite set of specifications due to the wide range of needs from contributors and vendors in the community. In reality, what works for one company does not always work for the next. Many industry leaders have contributed their own unique designs to OCP. To date, Google, Facebook, Microsoft, Rackspace, financial services firms, and many supply chain vendors have contributed or plan to contribute projects to meet the specific needs of their own workloads, business models, or datacenter environments. We expect as the OCP Telco Project evolves, additional specifications will be added to meet that segment's highly specialized needs. And because telcos are making a significant shift from proprietary hardware approaches to open, commodity-based hardware, we expect it will take much time and effort to gravitate to one common standard for this community.

Many US-centric businesses have different needs and requirements from those in other regions. For example, ITOCHU Techno-Solutions Corporation, Murata Manufacturing, and NTT Data Intellilink found the OCP rack specifications too large in size and not supportive of the conventional single-phase power mainly supported in Japan, making the OCP rack designs more expensive than traditional 19" racks. This group of Japanese companies is solving the power source, rack size, and cost issues, and adding earthquake-resistant design elements to their own version of the OCP rack.

Fragmentation of specifications means it will be more difficult for OCP to achieve the promised supply chain efficiency improvements the community desires. Fragmentation can also cause customer confusion about what OCP hardware really is and what choices to make for their specific environments. Some companies are even starting to intermix the terms "OCP" and "hyperscale", adding more confusion to what exactly OCP means when trying to reference a specific design.

In 2015, the OCP Foundation stepped in to help with a set of standards for certification and interoperability that attempt to distinguish what is and is not OCP hardware. Below is the current set of OCP labels available today.

- **OCP Inspired**: The product must meet one of 3 criteria: 1. it must fit into a 19" rack, OCP Rack, or Open Rack, or 2. the spec must be submitted to OCP IC and Foundation staff, or 3. the product must have a hardware license assigned from OCP or OWF (Open Web Foundation). Additionally, the product must be associated with the target segments of OCP which include IT compute, network, and / or storage functions, and must be manufactured (a physical product, even if it is just a prototype or engineered sample). OCP Inspired is the lowest requirement with a relatively loose set of qualification guidelines. A product that receives this certification can be resubmitted for OCP Accepted certification.

- **OCP Accepted**: This distinction has a much stricter set of criteria than OCP Inspired. The product must be manufactured, be available for purchase within 120 days, meet all criteria of an approved OCP specification, go through a self-assessment for each SKU, and be submitted to the Foundation with the proper documentation for re-assessment. After these criteria are met and the Foundation confirms the products meets the criteria via re-assessment, the OCP Foundation lists the SKUs on the OCP web portal and the vendor can use the OCP Accepted distinction for their product. No products are currently listed.
- **OCP Certified**: According to the OCP website, this distinction is “Temporarily closed; to be redefined”. An advanced level certification may be a challenge for OCP, as it requires substantial resources to implement and investments in time and money from member companies seeking the OCP Certified label.

Standards can help end users navigate the OCP landscape, but the current process appears to be in its infancy and not yet a useful metric for buying OCP-based hardware. With one distinction under construction and another with no products currently listed, an end user would be hard-pressed to use the OCP labels available today as a starting point with vendors to find qualified OCP hardware that meets their needs.

WELL-DEFINED SPECIFICATIONS LACKING IMPLEMENTATION DETAILS

It is important to note the significant differences between hardware and software and what each one takes to bring a new product to market. Delivery models for open source initiatives are fundamentally different for hardware and software. Below are some of the major differences to consider when looking at open initiatives in each area.

- **Economies of Scale**: Hardware is different from software with respect to open economies of scale. Manufacturing, inventory, shipment, and installation are all physical processes, and their efficiency models do not scale like software.
- **Manufacturing & Distribution Costs & Effort**: Software product investment is a developer time investment. Unlike hardware, software has low manufacturing and distribution costs; lead time for new software distribution and delivery is relatively straightforward and low-cost. Bringing new hardware to market requires procurement of inventory, manufacturing, logistics, and distribution management.

- **Dependence on Vendor Implementations & IP**: Hardware differs from software in that each atom of a hardware design cannot practically be sourced from open projects. At some point, there is a manufactured implementation different from other manufactured implementations, and the differences are not documented for competitive or other reasons. An open hardware design based on a specific SoC vendor's product does have a proprietary component even though that component may be readily accessible to buyers at a reasonably low price. Open hardware is also complicated, because differing implementations of components with practically equivalent functionality produce operational differences that impact the interchangeability of parts in an open design.
- **Design Costs**: Creating a specific implementation or change in software based on open source code is dependent on developer time as the primary cost. However implementing any hardware design modification is likely to cost significantly more, as each change impacts the entire hardware value chain.

With the additional complexities associated with open hardware, a detailed assessment of the hardware value chain must be considered to understand the necessary steps for bringing a new product to market. Figure 1 illustrates the technology hardware value chain and the steps delivered by OCP versus steps that are responsibilities of the end user customer or independent hardware solution provider.

FIGURE 1: THE HARDWARE VALUE CHAIN WITH OCP

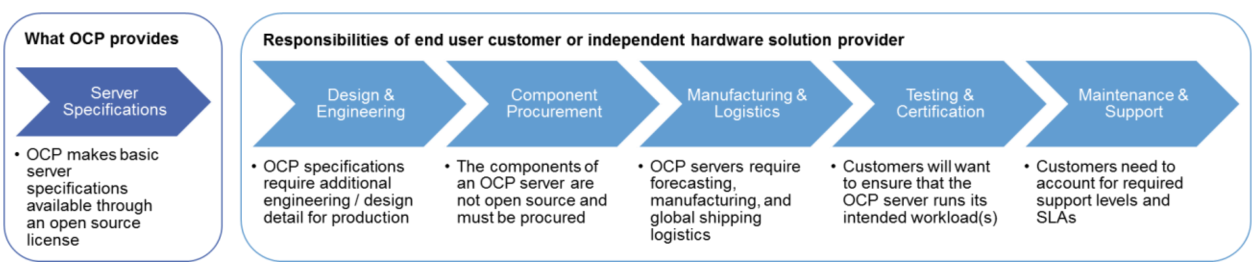

OCP is a set of specifications with mechanical CAD files but, in most cases, no motherboard schematics, PCB layout files, or BIOS code (OpenPOWER is the exception with a full design package including layout file and BIOS). This paper's Appendixes provide an overview of what is available for each of the technology contributions.

Design is a small piece of the hardware puzzle. Significant investments in time and money are required to complete the remaining value chain steps to create a datacenter-class hardware product. End users who have relied on the traditional datacenter hardware supply chain for design, manufacturing, and ongoing support are not well-staffed to take on these steps. Such customers will rely on their independent hardware solution provider to lead them through building and deploying OCP hardware.

The Supply Chain Has Grown, Yet Challenges Persist

Many hardware vendors have drawn design inspiration from OCP to create innovative and efficient hyperscale designs. OCP has likely had a hand in accelerating innovations in software-defined networking, software-defined storage, and hyperconverged systems. OCP has inspired new business models to help serve those interested in cost effective, no-frills designs more effectively. For example, HPE partnered with Foxconn to create the Cloudline brand of servers to compete directly with ODM whitebox vendors in the service provider market.

Almost every major vendor in the server space has claimed to develop an OCP-based design, but few are shipping significant volume of OCP hardware today, outside the Facebook and Microsoft deployments of their own donated designs. As of April 2016, OCP lists 8 official solution providers on their site. In reality, the vendors that are known to be shipping OCP in any significant volume include one or two ODMs (Quanta, *etc*.) with likely limited volumes by a few specialty vendors (Hyve, Penguin, Stack Velocity).

Many ODMs and specialty vendors involved in OCP have business models optimized for organizations like hyperscale and HPC users who have large internal engineering support staff of their own. Due to the nature of their historical business models, many OCP solution providers have limited sales, service, and support capabilities compared to tier 1 global OEM vendors. Enterprise IT and service providers accustomed to working with global OEM vendors may find that OCP solution providers have insufficient warranty support, field service capabilities, and staff to deal with technical issues. Moreover, the OCP supply chain caters to hyperscale customers making massive purchases, so smaller customers looking to buy OCP hardware may find their orders deprioritized. With these limitations in mind, using an ODM or smaller specialty vendor may result in a “hidden tax” for IT organizations by requiring extra employees to handle supply chain, service, and support of the hardware.

Today’s OCP Servers are Not for Every Workload

Many hyperscale datacenters have architected their applications with resiliency in the software and designed them to survive dozens or hundreds of hardware failures at a time. These organizations consider hardware to be a commodity that can easily be replaced in the instance of a failure without impacting their applications' ability to support the business. OCP platforms have been designed with software resilience in mind); therefore, applications written for high availability hardware may experience risks moving to today's OCP platforms. While enterprise IT is moving toward software-defined, fault tolerant workloads for newer applications, many legacy applications in the datacenter today would need to be re-architected before moving to OCP hardware or the hyperscale IT blueprint.

CALL TO ACTION

The Open Compute Project has influenced the industry to embrace the benefits of open, efficient design principles to help hyperscale and hyperscale-inspired IT buyers lower costs, streamline the supply chain, and avoid proprietary technologies that create vendor lock-in. However, enterprise IT organizations and service providers must go into the OCP evaluation process with eyes wide open to understand the potential benefits and shortcomings today.

- **What is right for Facebook and other hyperscalers is not necessarily right for a smaller scale datacenter**. Understand your pain points and technology requirements to make sure OCP can meet the needs of your applications.
- **There are a limited set of vendors who are committed to bringing OCP hardware to the market today**. Make sure these vendors can meet your needs for the entire hardware delivery value chain including design, manufacturing, supply chain, solution validation, service, support, and even decommission.
- **OCP hardware is not right for every application**. Ensure the applications you are looking to move to OCP hardware are designed with fault tolerance and resilience built into software.

There are other avenues to consider for those looking to move to open, cost-optimized design principles that may address some of the shortcomings seen with OCP today. Some of the specialty vendors who sell OCP hardware have broader portfolios that may be a better fit for organizations outside the hyperscale category. Large global vendors like Dell, HPE, and Lenovo have business lines that provide hyperscale-inspired designs that come with the service and support that many enterprise IT organizations and service providers require.

MOOR INSIGHTS & STRATEGY

APPENDIX A: DESIGN GUIDANCE AVAILABLE FOR OPEN RACK SUBMISSIONS

		Title	Version	Latest Update	Contributor	Whitepaper	Specification	Mechanical Drawings	Mechanical Files (CAD)	Schematics	Board Layout Files (Gerber)	BIOS
Open Rack and Power	General	Whitepaper on Deploying Freedom OCP Gear in a Co-location Facility	n/a	Jan-13	Facebook	✓				Not Applicable	Not Applicable	Not Applicable
		Open Rack Design Guide for IT Gear Builders	1.1	Aug-13	Facebook	✓						
		Open Rack Standard	**1.1**	**Mar-14**	**Open Rack Project**		✓					
	Freedom	Facebook Freedom Rack Hardware Specification	0.6	Sep-12	Facebook		✓					
		Facebook Freedom Triplet Rack CAD Files	0.6	Sep-12	Facebook				✓			
		Methode V1 DC clip (2D pdf & CAD)	A	Oct-12	Methode			✓	✓			
	Version 1	Facebook V1 Open Rack Hardware Specification	1	Jan-13	Facebook		✓					
		Facebook Open Rack V1 Specification	1.8	Mar-13	Facebook		✓					
		Facebook Open Rack V1 CAD files	11	Dec-13	Facebook			✓	✓			
		Facebook V1 Power Shelf Specification	0.3	Jan-13	Facebook		✓					
		Facebook V1 AC/DC PDU assembly CAD	n/a	Oct-12	Facebook			✓	✓			
		Facebook V1 gPDU Specification	0.3	Jan-13	Facebook		✓					
	Version 2	Facebook V2 Implementation of Open Rack	1	Jul-15	Facebook		✓	✓				
		Facebook Open Rack V2 Power Shelf Specification	1	Feb-15	Facebook		✓					
		Facebook Open Rack V2 Cubby Sub-Chassis Specification	0.41	Jun-15	Facebook		✓					
		Open Rack V2 DC PDU Specification	1	Feb-15	Facebook		✓					
		Open Rack V2 EU 230V AC PDU Specification	1	Feb-15	Facebook		✓					
		Open Rack V2 277V AC PDU Specification	1	Feb-15	Facebook		✓					
		Open Rack V2 Battery Backup Module 3600W Specification	2	Jul-15	Facebook							
Open Rack Server	V1	Server Chassis and Triplet Hardware	1.0	Apr-11	Facebook		✓			Vendor Responsibility	Vendor Responsibility	Vendor Responsibility
		Facebook server Fan Speed Control Interface	0.1	Jan-14	Facebook		✓					
	Version 2	Server sled and shelf specification for Open Rack V1	0.3	Jan-13	Facebook		✓					
		Server Shelf mechanical model for Open Rack v1	1.0	Sep-13	Facebook			✓	✓			
		Facebook server Intel Xeon motherboard v2	2.0	Apr-12	Facebook		✓					
		Facebook server Intel Xeon motherboard v2 - DXF	2.0	Apr-12	Facebook			✓	✓			
		Facebook server Intel Xeon motherboard v2 - 3D model (4 files)	2.0	Apr-12	Facebook			✓	✓			
		Facebook server AMD Opteron motherboard v2	2.0	Apr-12	Facebook		✓					
	Version 3	Facebook server Intel Next Generation Xeon motherboard v3.1	3.1	Jan-16	Facebook		✓					
		Facebook server Intel Next Generation Xeon motherboard v3	3.0	Jan-14	Facebook		✓					
		Leopard_ddr4_dfx_v1_20150210.pdf (157 KB) simple PDF of DFX	3.1					✓				
		Leopard_ddr4_dfx_v1_20150210.dfx (2MB) DFX file							✓			
		Facebook addin Card Thermal Interface Spec for intel Motherboard V3.0	3.0	Jun-15	Facebook		✓					
	Open	Overall Spec v0.8.2	0.8.2 DVT	Jan-16	Rackspace		✓					
		EE Components -- BoM, Board Files, Gerbers, Schematics, CPLD Source	0.8.2 DVT	Jan-16	Rackspace	✓		✓		✓	✓	✓
		ME Components -- 3D Files and Spec	0.8.2 DVT	Jan-16	Rackspace		✓	✓	✓			
1U/2U 19"	Monolithic	Intel Decathlete Board Standard v2	2.1	Mar-15	Intel		✓					
		QuantaGrid D51B-1U (Decathlete) System	1.0	Nov-15	Quanta		✓		✓	✓	✓	
		Open Compute Project AMD Open 3.0 Modular Server Specification (PDF)	0.5	May-13	AMD		✓					

Note: System on Chip Servers and the OCP Mazzanine Card available at http://www.opencompute.org/wiki/Motherboard/SpecsAndDesigns

MOOR INSIGHTS & STRATEGY

APPENDIX B: DESIGN GUIDANCE AVAILABLE FOR OPEN CLOUD SERVER (OCS) SUBMISSIONS

		Title	Version	Latest Update	Contributor	Whitepaper	Specification	Mechanical Drawings	Mechanical Files (CAD)	Schematics	Board Layout Files (Gerber)	BIOS
Open Cloud Server	V2.1	Chassis Manager Update spec+collateral (aka CMv2)	2.1	Feb-16	Microsoft		✓					
		Blade spec update	2.1	Feb-16	Microsoft		✓					
	V2.0	Chassis	2.0	Oct-14	Microsoft		✓					
		Chassis Management	2.0	Oct-14	Microsoft		✓					
		Blade	2.0	Oct-14	Microsoft		✓					
		Blade NIC Mezzanine	2.0	Oct-14	Microsoft		✓					
		Tray Mezzanine	2.0	Oct-14	Microsoft		✓					
		Chassis Manager Service Open Source Location				https://github.com/MSOpenTech/ChassisManager						
		OCS Operations Toolkit Open Source Location				https://github.com/MSOpenTech/OCSOperationsToolKit						
		OCS_Mechanical_120314.zip, 66MB	2.0	Oct-14	Microsoft			✓	✓			
		OCS Tray backplane and power distribution board	2.0	Oct-14	Microsoft		✓					
		OCSv2_TBP_PDB.zip, 31MB	2.0	Oct-14	Microsoft					✓	✓	
		OCS Open CloudServer Power Supply v2.0	2.0	Oct-14	Microsoft		✓					
		OCS Open CloudServer Power Supply Mechanical v2.0	2.0	Oct-14	Microsoft		✓					
		OCS Open CloudServer Solid State Drive v2.0	2.0	Feb-15	Microsoft		✓					
		OCS Open CloudServer Solid State Drive v2.1	2.1	Aug-15	Microsoft		✓					
	V1.0	Blade	1.0	Jan-14	Microsoft		✓					
		Chassis	1.0	Jan-14	Microsoft		✓					
		Chassis Management	1.0	Jan-14	Microsoft		✓					
		JBOD Blade	1.0	Jan-14	Microsoft		✓					
		NIC Mezzanine	1.0	Jan-14	Microsoft		✓					
		SAS Mezzanine	1.0	Jan-14	Microsoft		✓					
		Chassis Manager Schematics, Board File Collateral (15.9MB)	1.0	Jan-14	Microsoft					✓	✓	
		Power Distribution Board Schematics, Board File Collateral (25.3MB)	1.0	Jan-14	Microsoft					✓	✓	
		Tray Backplane Schematics, Board File Collateral (26.3MB)	1.0	Jan-14	Microsoft					✓	✓	
		Chassis, Tray, Blade Mechanical Collateral (128.7MB)	1.0	Jan-14	Microsoft			✓	✓			
		3-D CAD Model of chassis with blade (4.9MB PDF)	1.0	Jan-14	Microsoft			✓				
		Chassis management source code, located in GitHub.				https://github.com/MSOpenTech/ChassisManager						
	Other	Netronome OCS NIC MEZZ adapter	1.0	Mar-15	Netronome Systems			✓	✓	✓	✓	
		Bom: BOM_AMDA0079-001_03.xlsx BOM with vendor P/N				✓	✓	✓				

APPENDIX C: DESIGN GUIDANCE AVAILABLE FOR OTHER SUBMISSIONS

Storage

There are accepted projects but no approved specification
http://www.opencompute.org/wiki/Storage

Networking

Accepted Hardware

Alpha Networks - SNX-60x0-486F - 48-port 10G SFP+ & 6-port 40G QSFP+ - Leaf Switch
Alpha Networks - SNQ-60x0-320F - 32-port 40Gb QSFP+ Leaf/Spine Switch
Broadcom/Interface Masters Open Leaf and Spine Switch specification
Edgecore Networks AS5712-54X - 48-port 10G SFP+ & 6-port 40G QSFP+ - Leaf Switch
Edgecore Networks AS6712-32X - 32-port 40G QSFP+ - Leaf/Spine Switch
Edgecore Networks AS7712-32X - 32-port 100G QSFP28 - Leaf/Spine Switch
Edgecore Networks ORSA-1RU - Open Rack Switch Adapter
Inventec DCS6072QS - 48x10GB SFP+ & 6x40GB QSFP+ - Leaf Switch
Inventec DCS7032Q28 - 32x100GB QSFP28 - Leaf/Spine Switch
Mellanox MSX1410OCP - SwitchX-2 48x10GB SFP+ & 12x40GB QSFP+ - Leaf Switch
Mellanox MSX1710OCP - SwitchX-2 36x40GB QSFP+ - Leaf/Spine Switch
Facebook Wedge - 16x40GB QSFP+ - Leaf/Spine Switch

Accepted Software

Open Network Install Environment
Open Network Linux
Switch Abstraction Interface

http://www.opencompute.org/wiki/Networking/SpecsAndDesigns

Hardware Management

LAVA LMP

https://github.com/opencomputeproject/lava-lmp-firmware/subscription
https://github.com/opencomputeproject/lava-lmp-lmpd/subscription
Cloud Server Multi Node System Specification version V0.7.5, Date August 19,2015 PDF

http://www.opencompute.org/wiki/Hardware_Management/SpecsAndDesigns

Data Center

Data Center - Mechanical Specifications v1.0

http://files.opencompute.org/oc/public.php?service=files&t=2db7e886bcef44bb0421dcab84bb6bfc

http://www.opencompute.org/wiki/Data_Center/SpecsAndDesigns

Important Information About This Paper

Author

Gina Longoria, Senior Analyst at Moor Insights & Strategy
Jimmy Pike, Technologist at Moor Insights & Strategy

Review

Patrick Moorhead, President & Principal Analyst at Moor Insights & Strategy

Editor / Design

Scott McCutcheon, Director of Research at Moor Insights & Strategy

Inquiries

Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

Citations

This 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

Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. 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.