

Magic Quadrant for Modular Servers

Published: 23 May 2016 **ID:** G00278242

Analyst(s): Andrew Butler, George J. Weiss, Kiyomi Yamada, Hiroko Aoyama

Summary

Any server that is hosted by a dedicated chassis or enclosure is termed a modular server by Gartner. The market chiefly comprises blade and multinode servers. I&O leaders should use this research to help choose both the right form factor and the right vendor for new projects.

Market Definition/Description

Modular Servers Overview

Gartner defines the modular server market as a class of server where the compute capabilities are delivered in a modular style and housed in a common, often proprietary, chassis. Storage and network components may also be housed within the same chassis housing. There is a broad range of modular servers that vary in configuration options, based on CPU type (reduced instruction set computer [RISC], x86 and system-on-chip [SoC]), memory and storage attachment, and switching technologies. Server design focuses on density of components and richness of infrastructure management software platforms. Modular servers have evolved from many years of blade, multinode and custom-made high-performance computing (HPC), supercomputing, and IT service provider designs. The term describes a style of server design that focuses on the ability to improve server management visibility, CPU architecture flexibility, server density and operational efficiency. Modular servers will be increasingly positioned to address many emerging web-scale IT and Mode 2 workload requirements.

The Modular Server Market

For Gartner to consider a server as "modular," the device must fit into a chassis or enclosure that allows for the easy and rapid addition or replacement of new hardware. We do not consider rack, tower and frame-based servers as modular servers; we also exclude "build your own" servers created from custom-made motherboard and component acquisitions. Blade servers and multinode servers comprise the great majority of this market (see "Market Definitions and Methodology: Servers" for more information).

The design of modular servers focuses on the density of hardware components and the ease by which technology can be extended and/or replaced. Blade servers additionally offer the benefits of a highly integrated switched backplane that benefits input/output (I/O) performance and infrastructure management. All modular servers boast CPU architecture flexibility and support the requirements of

many emerging, web-scale IT workloads. A wide range of modular server offerings is now available. These servers vary in configuration options, such as CPU type and memory, storage, and switching technologies.

Modular servers typically address smaller workloads and workloads that can scale out to achieve greater throughput. Blade servers are particularly well-suited for data center consolidation and modernization strategies, and most first-generation integrated system strategies have been based on blade servers. There are still some key differences between blade servers and multinode servers; blade servers can support scale-up as well as scale-out workloads (especially blades from vendors like Hitachi and Hewlett Packard Enterprise that support processor aggregation across multiple blades), whereas multinode servers only support scale-out workloads. The biggest difference is that a blade server will likely have a dedicated switch, while multinode servers are more than likely going to have some form of Ethernet-based switching. But the lines of separation between blades and multinode are blurring more and more with each generation.

The primary adopters of blade servers are large and midsize enterprises that desire maximum flexibility in their on-premises data centers. Most multinode demand is shared between large enterprises building out infrastructure for big data analytics, and external IT service providers (including those using hyperscale) that invest in massive scale-out architectures for public and hybrid cloud deployment.

So, the technology preferences for modular servers vary. Large enterprises turn to the technology for its sophisticated management capabilities, hardware density and ease of installation. External IT service providers, on the other hand, need only basic management capabilities and use the technology to reduce costs and increase agility, delivering IT services through lean, repeatable and dense system configurations. But all users have to ensure that the rack density of modular servers does not become a burden. A rack full of modular servers can exceed the heat dissipation capability of air-cooled racks, and power consumption has been a challenge for some power-constrained data centers as well. Data center leaders must always check that intense deployment of modular servers – especially very large-scale clusters – will not exceed the capability of their data center facilities.

The different use cases fit well with the concept of bimodal investment. In Mode 1 of bimodal IT, traditional IT teams need servers that are suited to business-critical workloads – such as blade servers – that justify investment in highly available hardware and services. However, in the Mode 2 scenario, a development team within infrastructure and operations will need servers where agility at the lowest cost is a leading characteristic – a situation better-suited to multinode servers. As a result, data center managers will have to decide whether to bifurcate their infrastructure sourcing strategy and create separate silos, or find a single solution that fits all their needs. From a workload perspective, a blade server design could address both sets of requirements. But it is likely to be overengineered and overpriced for the needs of most Mode 2 workloads.

The bifurcation of the industry creates new challenges – and opportunities – for vendors. In reality, IT infrastructure has always been bifurcated, where HPC investments have usually been driven by different teams within the enterprise – often reporting to the line of business rather than the IT organization. From an architectural perspective, the characteristics of an HPC workload are very similar to those of many new-generation cloud or analytics workloads. So Mode 2 is providing an opportunity for established HPC vendors like Atos, Cray and SGI to expand their coverage to include

new market needs, and potentially create bridgeheads into the legacy data center that have not been possible before now. Meanwhile, vendors with proven Mode 1 track records are benefiting from the trends in HPC that increasingly allow multinode, high-volume servers to address HPC workloads that were once the reserve of more specialized hardware.

Magic Quadrant

Figure 1. Magic Quadrant for Modular Servers



Source: Gartner (May 2016)

Vendor Strengths and Cautions

Atos

Since its 2014 acquisition of Bull, Atos has been working to integrate Bull's various technologies into relevant Atos go-to-market strategies, and to create a broader global awareness of the product portfolio. Today, over 90% of Atos server revenue comes from Western Europe (of which, around

two-thirds comes from France). So, while the Atos brand has wider global recognition than Bull ever had, it will still take some years to create wider market awareness of Atos as a server brand. Bull's past track record as a leading vendor in the high-performance computing market gives Atos a strong legacy upon which to build. In the November 2015 publication of the TOP500 list, Atos had five projects in the top 100. Its highest-ranked project was No. 44.

In November 2015, Atos launched the Bull sequana X1000 – a new supercomputer range that aims to address exascale technological challenges. The French Alternative Energies and Atomic Energy Commission (CEA) is an early public reference. Bull sequana supports both general-purpose Intel x86 CPUs, plus specialized processors such as general-purpose computing on graphics processing units (GPGPU) and field-programmable gate array (FPGA). Atos has developed a software ecosystem that is able to reconcile traditional HPC workloads and big data (nonstructured workloads).

The bullion S is a more established modular server platform targeted at both HPC requirements and enterprise workloads like SAP Hana. Bullion S scales from two to 16 Intel Xeon E7 v3 CPUs (in a symmetric multiprocessing [SMP] or scale-out configuration) and allows for up to 24TB of main memory. Such a system scored a new record benchmark result for the SPECInt_rate2006 benchmark in January 2016. The bullion S technology includes a number of features developed by Atos in addition to Intel's intellectual property, such as hardware partitioning, processor and memory aggregation, hot-swap memory, and predictive failure tools.

As a proven system integrator able to work with many other hardware vendors, Atos must juggle the business potential for high-end workloads like HPC and in-memory compute with the need to collaborate with other vendors who have broader modular market appeal. To achieve maximum success for sequana and bullion S, Atos must focus business development on the more complex workloads that can benefit from the high-touch, service-led approach that the company can deliver.

STRENGTHS

The combination of Atos' track record as a global system integrator with strong independent software vendor (ISV) relationships, and Bull's pedigree in HPC and mission-critical systems, positions the company well for market success.

Bull sequana positions Atos to address high-end digital simulation and big data challenges in demanding markets such as defense, manufacturing, energy, life science and finance where the company has established relationships.

With bullion S, Atos can leverage the fast-growing market for SAP Hana and target other high-end x86 server requirements beyond the scope of most other vendors.

CAUTIONS

While Atos can claim financial stability and a globally recognized service brand, few enterprises yet associate the Atos name with anything related to hardware. Similarly, the Atos brand does not bring obvious value for HPC buyers who will assess Atos versus the resurgent businesses of Cray and SGI.

The bullion S modular server range is best-suited to customers who have specialized, high-end projects, preferably in a region where Atos already has proven server sales capability. The platform will not help Atos gain recognition as a vendor able to address more generic workload requirements.

While Atos is making good progress in regions like EMEA and Latin America, the company faces fierce competition to penetrate the U.S. market, with its multiple entrenched established players.

Cisco

Within seven years of entering the market, Cisco has risen to become a leading modular server vendor; Cisco is the clear No. 2 blade server vendor by revenue, with a No. 1 position in the U.S. and a number of other countries. However, its overall modular server impact has lagged behind leaders such as Dell and HPE, who have far broader portfolios and proven success in the multinode server market. Much of Cisco's modular server revenue in 2014 and 2015 was generated by the Unified Computing System (UCS) Mini, a UCS-managed, distributed and remote office/branch office (ROBO) server. Market success has been limited so far, and Cisco has now extended the Mini's scale with a second chassis. In 2015, Cisco extended its overall modular portfolio with higher scale-up and scale-out configurations. M-Series and C-Series expanded configurations address higher compute and storage densities to enable Cisco to reach markets in web-scale, HPC and hyperconverged integrated system (HCIS) use cases.

UCS positioned Cisco well to ride the early wave of converged infrastructure demand that was based on blade server and shared storage technology (a market that is still buoyant today). Modular infrastructure markets employing direct-attached storage (DAS) – such as HCIS – are now demanding leadership, not only in servers, but also in components of finer granularity rather than monolithic technology. Cisco will be driving innovations and go-to-market strategies in the following areas to gain a greater footprint in the modular server market: (1) extending policy-based resource management to smaller building blocks in UCS Manager 3.1; (2) including a layer to deliver user and administrator portals for orchestrated service catalogs, design templates, hybrid and multicloud integration in UCS Director; (3) offering a common control plane unifying orchestration; and (4) continuing System Link Technology that enables shared microresources such as CPU, memory, accelerators and networking as a composable infrastructure of disaggregated resources. These innovations will help to further reinforce Cisco's Mode 1 presence, but will not help to create more demand in Mode 2 situations.

We expect partnerships (for example, Springpath) will provide Cisco functionally broader software-defined capabilities in DAS-based clusters during 2016. In addition, System Link presents a promising substrate for composable infrastructure and software-defined storage solutions with Application Centric Infrastructure (ACI) integration. On a go-to-market basis, Cisco will be faced with market challenges as a unified EMC and Dell drive market opportunity for Dell's modular portfolio.

STRENGTHS

Cisco is a major success story in the blade server market, with a more than \$3 billion dollar business since entering the market in 2009.

Cisco blade technology holds a market-leading position for converged infrastructure, with proven implementations that leverage the storage assets of EMC, NetApp, IBM, Pure Storage and Hitachi.

Cisco delivers strong global service and support for its servers, backed by a respected and responsive server engineering team. This is an impressive achievement for a vendor that has only been a market participant for seven years.

CAUTIONS

Cisco's integrated system success is not a prelude to modular and multinode success until it demonstrates successful large-account adoption and scale in a variety of workloads.

Cisco's UCS Manager and Director must advance and integrate as a seamless composable infrastructure to reinforce and expand its UCS market acceptance.

Cisco lacks the portfolio breadth, market coverage and pricing leverage that its competitors with greater market share have accumulated over a much longer period in the market.

Cray

Cray is arguably the most proven and respected name in high-performance computing. The latest TOP500 list of global supercomputers (published in November 2015) lists 33 Cray systems in the top 100, and no less than five Cray systems in the top 10. As a result, Cray is a recognized brand name across the industry – with clear leadership in the TOP500 – and a proven leader in the delivery of graphics processing unit (GPU)/accelerator systems. However, until recently, Cray has not been a brand that most enterprises would consider relevant for their non-HPC business needs.

Meanwhile, the HPC market is evolving as multinode high-volume servers displace the use of highly differentiated and high-margin offerings for many workloads. This is forcing HPC vendors to seek new market opportunities beyond their traditional space, and Cray is investing in product and partner development that will position the company to benefit from the Mode 2 drive toward analytics solutions.

Cray targets the HPC space with two powerful and proven solutions – the CS series and XC series product lines. The very highest multipetascale demands are addressed by the XC40; this system is designed to deliver the maximum level of scaling and upgrade potential that Cray can offer. The CS family product line is focused more on delivering maximum price/performance through the use of industry-standard technologies. To expand the market appeal of the CS series, Cray works closely with hardware partners (like Intel and Nvidia), allowing early testing and shipment programs, and enabling innovative designs such as the eight Tesla GPU nodes per chassis in the CS-Storm product. Cray also works closely with software vendors like Altair (a leader in HPC workload management), SUSE and Red Hat.

Cray's advance on the market for big data analytics is spearheaded by the Urika-XA system. This is an InfiniBand-based system that can support up to 1,500 cores and 512TB RAM, and is aimed at markets like Hadoop and Spark. By targeting the growth in analytics investment – which is happening across the span of enterprise IT – Cray aims to create a new revenue stream that will augment and complement its continued strong HPC performance.

STRENGTHS

Cray is the strongest of the recognized HPC brand leaders in terms of annual revenue and proven high-end implementations.

The company has a strong roadmap that aims to converge, over time, the technologies for analytics and HPC simulation with modeling.

With a large and diverse installed base that covers multiple segments, the company is well-positioned to exploit its strong customer partnerships and create new opportunities for non-HPC business; during 2015, Cray was able to more than double its non-HPC revenue to over 15% of total revenue.

CAUTIONS

Cray, like other HPC vendors, is vulnerable to extended sales cycles and budgetary constraints caused by heavy exposure to the public sector. Business results come from a relatively low volume of fairly high-value deals, which causes more quarter-to-quarter variability than most server vendors.

The Cray brand has lacked enterprise recognition; however, that is changing as the company's commercial revenue is becoming an increasingly bigger part of its business.

Cray faces threats from mainstream server vendors targeting evolving supercomputing workloads that can be addressed by high-volume, x86-based computing clusters. This forces Cray (and other HPC vendors) to seek new Mode 2 market opportunities to amortize the risk.

Dell

During 2015, Dell enhanced its portfolio of modular servers and expanded its Blueprints program to target customer workloads across more specific vertical solutions by introducing new templates and capabilities to its Active System Manager for provisioning and automating a wide variety of infrastructure and workloads. Dell's modular server strategy is shifting from predominantly traditional rack-optimized and tower-based systems to greater investment and go-to-market emphasis on converged and extreme-scale infrastructures. Dell's modular server portfolio is composed of several PowerEdge systems: VRTX, FX series, M series and C series, and the Extreme Scale Infrastructure (ESI) comprising the Datacenter Scalable Solutions (DSS) and Data Center Solutions (DCS) families. The market and technology improvements were primarily aimed at high-density compute, scale-out and scale-up, high-performance workload-optimized and purpose-built configurations, ease of configuration and automated configuration replication. The FX2 and the C6000 are purpose-built for web-scale, high-density scale-out DAS storage to bolster extreme scale, capital expenditure (capex) and operating expenditure (opex) savings, efficient power management, footprint reduction, and strong scale-out performance. In the 2015 version of this Magic Quadrant, we alluded to FX2 as being the first of a series of modular configurations, and the 2016 portfolio has indeed expanded to several models with new options adaptable to virtualization, HPC and online transaction processing (OLTP), including VMware VSAN-ready nodes.

The C series targets hyperconvergence and web-scale workloads, and serves as the foundation for Nutanix and Evo:Rail (which is likely to include future VxRail nodes from the combined Dell-EMC entity). Dell also addresses hyperconvergence in a partnership with Microsoft Cloud Platform System (CPS). Dell continues to drive sales in small or midsize businesses (SMBs) and ROBOs with PowerEdge VRTX as an integrated system form factor. Dell broadened the VRTX configuration flexibility to include options for a second chassis and greater I/O and storage to accommodate more workloads and scale. These capabilities are helping to simplify management at the system level, but

Dell needs to create better market awareness of the enhancements being made to Active System Manager as it drives toward an open, extensible architecture that allows automation across heterogeneous infrastructure and customer-preferred applications via software development kits (SDKs). Introducing new management functionality, while attempting to drive innovation into the market for new sales opportunities requires a delicate and balanced portfolio approach. We expect that Dell's hardware portfolio will be the go-to strategy of the Dell-EMC merger, with no short-term attempt to displace Cisco servers and networking from the original VCE product portfolio.

STRENGTHS

Dell is the second-largest provider of x86-based servers in the world, and has a large global reach through an extensive partner network.

Partnerships and acquisitions have strengthened Dell's portfolio and market reach, and the company has strong alliances as a hyperconvergence platform, plus original design manufacturer (ODM) competitive opportunities.

Dell continues to strengthen its web-scale and hyperconvergence platforms via its expanded high-density PowerEdge FX2 systems, while providing support and migration options for existing M series users.

CAUTIONS

Dell's efforts to create a convincing management software story for integrated systems (above the standard system hardware management level) are well-suited to its broadening modular server, but remain limited in market awareness and penetration.

The acquisition of EMC will complicate Dell's efforts to solidify enterprise credibility as the new company coalesces go-to-market programs, branding, identity and messaging.

Dell's acquisition of EMC should make the new company even more attractive to enterprises, but will require a period of at least a year to seamlessly and efficiently manage the large mix of accounts, and to integrate IT portfolios with a strategic and coherent product delivery.

Fujitsu

Over the past two years, Fujitsu has steadily created clear branding and positioning for its modular server range, which complements its broad server portfolio and brand recognition in the global market. Under the Primergy brand that is applied to the majority of its x86-based servers, Fujitsu offers the Primergy BX blade server family and the Primergy CX multinode server family. Both can support cloud infrastructures, HPC and scalable enterprise infrastructures. Primergy BX and CX servers are also the compute components of the Primeflex Integrated System family, which has also undergone rebranding to help create better synergy and expand sales opportunities.

Fujitsu is a strong investor in modular server innovation, offering Cool-safe liquid cooling technology, which is unique to the CX400 multinode server range. Fujitsu is also investing in the next-generation bootable Many Integrated Core (MIC) Architecture from Intel, which will mainly be targeted at CX600 users running HPC workloads. Fujitsu has also developed a "software-defined server platform" (SDSP), which is its terminology for a form of integrated system that enables automation

management to support a fully virtualized enterprise environment, including private cloud. These technology investments are evidence that Fujitsu is using technology strength to better compete in the busy modular server market.

But while Fujitsu is able to demonstrate technology innovation, proven experience and a consistent modular server strategy, its market execution has not demonstrated the outcome we would have expected given the breadth of offerings. Fujitsu's challenge (albeit one that is shared by many other vendors in this research) is to show how it can promote a complex portfolio of products to existing customers and prospects in a way in which true synergy can be achieved between multiple product lines. The company also needs to address geographic gaps in market credibility and awareness – particularly in North America, where local recognition remains patchy.

STRENGTHS

Fujitsu's range of blade and multinode servers has proven recognition among enterprises in a variety of vertical segments, which enables cross-selling opportunities for other Fujitsu products.

Due to good market understanding, Fujitsu invests heavily in modular server innovation, especially liquid cooling, bootable MIC and tailor-made chassis options that meet customers' flexibility requirements.

Fujitsu has articulated a clear vision of its modular server portfolio, with the scope to address growing software-defined, cloud and HPC opportunities that help to gain prospective customers.

CAUTIONS

While Fujitsu is a highly proven server brand in Japan and German-speaking countries, and is well-recognized in many other global geographies, its modular server market share remains low and lags behind that of leading vendors.

Partly because Fujitsu has not sufficiently proved its global capabilities, the company has struggled to align its integrated system and modular sales strategies, leading to limited sales expansion.

Fujitsu is still in the process of developing a broad global set of channel partnerships, and this contributes to its frequent lack of presence when competing against global leaders.

Hitachi

Hitachi has made steady progress during the past year, both with its blade server strategy and the related strategy for Hitachi Unified Compute Platform (UCP) integrated system solutions. Helped by a global marketing initiative to create better broad market recognition, Hitachi is now realizing 46.9% of modular server revenue from Japan (up from 40% at the end of 2014). This is a major achievement. Even so, the company's modular server business track record is still patchy when compared at a regional level.

Hitachi Data Systems (HDS) – which is better-known as a major, global storage vendor – is the Hitachi subsidiary responsible for promoting the modular server (and the UCP solution) business globally. But Hitachi Data Systems has chosen to promote the Hitachi UCP integrated system message most aggressively, rather than focus on selling modular servers in its own right. This has resulted in good global recognition for Hitachi's presence in the integrated systems market, but

recognition for Hitachi's known modular server strengths remains muted. The implication is that Hitachi's Japanese parent company and Hitachi Data Systems have ongoing differences in sales strategy. So for global customers, Hitachi Data Systems primarily markets the blade server line only as a part of the broader UCP strategy, while Hitachi actively markets both the BladeSymphony blade server line and UCP to domestic Japanese customers. This situation can cause confusion for global customers who seek to buy the same product portfolio from Hitachi across multiple geographies and may be unaware that the blade offerings are technically available to all customers.

The Hitachi/Hitachi Data Systems modular server strategy is consistent, with the addition of a new entry-level platform to expand the opportunity for SMBs and midmarket enterprises plus new scale-out solutions that better address the needs of cloud service providers. Hitachi is also working to improve the number and quality of channel partnerships. Hitachi offers many blade server capabilities that are highly differentiated (and sometimes unique). For instance, BladeSymphony is able to process and aggregate processors and memory across multiple blades in a single chassis. This capability is very unusual in standard blade chassis architectures and makes Hitachi blades particularly suitable for mission-critical workloads that have vertical-scaling demand beyond the scope of a regular blade, or workloads that will benefit from the ability to implement logical partitions (LPARs). An LPAR is a composable function that allows dynamic provisioning of memory, CPU, disk and I/O functions. Such functionality is usually limited to expensive, high-end RISC and similar servers. Features like LPAR support and advanced management orchestration through the Hitachi UCP Director set Hitachi's blade server technology innovation apart from that of many other vendors.

STRENGTHS

To better meet regional needs, Hitachi has implemented variations in its product portfolios by actively marketing blade servers and the UCP range in Japan where it has greater market recognition as a broad server vendor, while focusing on growing awareness of its integrated systems in the rest of the world.

Hitachi is a huge global conglomerate with proven storage expertise and the technical strength to implement highly differentiated functionality – such as LPAR support and processor/memory aggregation – to strengthen its modular server strategy.

Hitachi has steadily expanded international sales beyond the domestic Japanese market, and it has achieved recognition from global users.

CAUTIONS

While Hitachi's modular server sales are improving outside Japan, the revenue contribution remains low, which impacts the company's ability to make its technical superiority noticed.

The strategy of limiting the marketing of blade servers to be part of a broader integrated system strategy outside of Japan is known to have caused the loss of sales opportunities for Hitachi. This can impact Hitachi's ability to reach Japanese customers wanting to deploy a consistent global server strategy, as well as other international customers who are aware of the technology differentiation Hitachi is able to achieve.

The transition to a new global Hitachi business strategy and the reorganization will create flux, as the ongoing corporate evolution is likely to create changes in technology roadmaps, the solutions portfolio and partnerships.

HPE

Hewlett Packard Enterprise (HPE) has arguably the largest portfolio of modular systems in the market. Its breadth of coverage includes: BladeSystem, Apollo, Moonshot, Integrity Superdome X, Edgeline, HC 380 hyperconverged and most ConvergedSystem variants. We can now add the new Synergy line recently announced for delivery in 2H16. These systems are positioned to serve specific market needs optimized for specific workloads in varying architectural configurations. HPE's approach to a software-defined infrastructure, labeled "composable" infrastructure, is the company's most innovative new strategy direction. The foundation of composable is based on modular principles of disaggregated building blocks that can be assembled to meet most workload demands through a common management interface: OneView. With the OneView solution, resources become fluid, templated and deployed for particular workloads in minutes through an autodiscovery and integration process delivered through an integrated management appliance. Because Synergy represents a divergence from the other architectures, HPE has created go-to-market strategies that position it first and foremost for leading-edge I&O leaders driven by infrastructure innovation. As it plants the seeds of the new Synergy product strategy, HPE intends to demonstrate and provide "bridges" to existing legacy environments such as its BladeSystems and new HC380 HCIS.

However, until it evolves a user portfolio transfer strategy, each of the systems stands on its own for the market segments most suitable to the technologies. Some products made an initial market splash, such as Moonshot, and then receded to within their specializations, as other products (for example, Apollo) took center stage. Knowing the complexities of go-to-market with such an expansive portfolio, HPE will seek to enable sales and channel partners by market segment, including installed-base, best-of-breed workload-optimized solutions, mission-critical business, and competitive plays through system conversions and integration. HPE's prior experiences in convergence were fraught with sales inefficiencies, lack of clarity and weak execution. The best approach to an acquisition of HPE technology depends on whether the decision is strategic or tactical. The decision process must judge into which bucket(s) the acquisitions should fall: legacy upgrade, legacy bridge (composable), best-of-breed mix and match, and bridges into other vendor environments. As mature products like BladeSystem approach an end-of-life scenario (HPE expects to sell BladeSystem products through 2019, with a five-year support window beyond that), HPE will focus investment to build market demand for innovation such as nonvolatile RAM (NVRAM), silicon photonics, and the unified management and automation that OneView is designed to implement. HPE is betting that in the emerging period of intense innovation, mergers and acquisitions, and new competitors it represents the one known and reliable factor – I&O leaders can depend on continued R&D, financial stability and rapid innovations.

STRENGTHS

HPE is a financially stable and highly accepted x86 server vendor, with consistent market share leadership.

HPE's strength in modular servers has broadened to a variety of configurations and workloads that exploit scale-out, scale-up, compute accelerators, direct-attached and outboard storage under an enhanced version of the unified user and administration interface of OneView.

HPE has just begun to roll out an advanced version of software-defined infrastructure (called composable) with the Synergy line, which will enable users to more granularly and quickly aggregate components to create optimized, best-fit-for-purpose workload solutions.

CAUTIONS

In the face of innovations in technologies and systems, HPE must restructure and organize sales, channels and marketing support to upgrade in-the-field training, consultation and integration services to propel market reception and create momentum.

HPE will be challenged to go deep in its own blue-chip accounts, as well as accounts where competitors have reinforced positions through convergence, networking and storage (for example, Cisco UCS, VCE Vblocks, EMC storage and NetApp FlexPod).

HPE is putting a clear stake in the ground on composable infrastructure, but depending on industry reception of the concept and proven success stories, HPE will face an initial period (12 to 18 months) of potential uncertainty where any failures could cause collateral damage to its other portfolio members.

Huawei

Huawei is a global provider for networking, telecommunications and data center businesses with a broad IT product portfolio. The company has been focusing more on enterprise data center markets in recent years, and has made some progress in its modular server sales. It has expanded geographical coverage by targeting various verticals, including manufacturing, financial, utility and retail, as well as by leveraging its strong relationships with telecom carriers. With increasing demand for cloud platforms from enterprises, the company is taking advantage of its expertise to cater to the need for service providers. Huawei is ramping up efforts to better penetrate the U.S. market, but the sales and market acceptance is limited to date. However, Huawei has attained the No. 3 position in the 2015 worldwide modular server market in volume.

Huawei has a wide range of modular server products in its portfolio, targeting both scale-up and scale-out workloads. The E Series Blade Servers represent its converged infrastructures based on E9000 blade chassis and FusionServer CH series. Its multinode server line, X6800, targets scale-out applications such as big data analytics and HPC. The X6000 aims to cater to service provider workloads (that is, web hosting), while the X8000 is based on Open Compute Project (OCP) and Project Scorpio – a Chinese standards consortium led by major local public cloud vendors like Tencent, Baidu and Alibaba. Additionally, it plans to introduce a big data appliance in 2016. In addition, the X Series density-optimized rack-mounted servers and E9000 blades form the foundation of Huawei's FusionCube integrated systems portfolio. Huawei's co-marketing programs with SAP for Hana have helped the company to bring higher levels of brand recognition, as well as lead to new sales.

Huawei has strong engineering capabilities (with its main R&D center in China and multiple regional R&D centers worldwide). Huawei's engineering prowess is demonstrated by its ability to build and ship a 32-socket blade-based system to address specialized customer demand, and Huawei's labs have advanced projects underway to bring new technology (like silicon photonics and nonvolatile memory) to market. Huawei's broad portfolio of IT and communications technology, including storage, networking, telecom, wireless and cloud solutions, should help it further cultivate its place in the modular server market.

STRENGTHS

Huawei is targeting a number of global markets with its modular server products, and a wide range of workloads including scale out to scale up.

Huawei historically has a strong relationship with telecom companies and service providers from its networking device sales, enabling the company to cross-sell its modular servers and other data center technology.

Huawei is heavily investing in innovative technologies across full infrastructure stacks, including fabric-based disaggregation and open-source software technologies such as OpenStack.

CAUTIONS

Huawei's efforts to leverage its telecom market momentum to build data center brand recognition are at a nascent stage; still more work is needed for it to be recognized as a major enterprise solution provider.

Even though sales have been picking up in Europe, they have been limited in the U.S. due partly to geopolitical issues between the United States and China, but also through local lack of market recognition as a viable server brand.

Huawei's vision for software-defined infrastructure (aiming to provide a more flexible and agile solution) is relatively unknown outside of China. The company's FusionCube and investment in broader cloud computing solutions should help increase awareness (assuming success of those strategies).

Inspur

Inspur was one of the first companies to invest strongly in blade and other modular servers, and has become a successful vendor — albeit with limited global recognition — with strong and growing credibility in the Chinese domestic market. As a result, Inspur is the Chinese modular server market share leader with around 25% revenue share. The company is now investing in a global go-to-market strategy with an active focus on channel recruitment and local sales/marketing presence in multiple countries. This offensive is intended to create international awareness of the Inspur brand to help stimulate new business potential. The company has also opened an R&D center and a separate assembly plant in the U.S., and a service support center in India. Attaining global recognition is an exercise that will take many years to reach fruition, but Inspur is developing a track record of international successes, with most visible progress in Russia, Japan, some Middle East countries, and emerging nations in Africa, Asia and the Americas. Inspur's relationships with global ISVs like SAP, VMware, Oracle, Microsoft and Red Hat are nascent, which will limit Inspur's ability to achieve a foothold in the market for Mode 1 data center modernization.

Inspur categorizes its technology portfolio – and target audience – into two main communities: enterprises and service providers. Enterprise needs are addressed by two product lines: I9000 Blade System and InCloudRack system for enterprise. I9000 is a proven blade architecture that can scale up to eight sockets in a single node, so it is suited to most x86 enterprise workload requirements. Inspur also sells GPU and storage expansion blades that enable users to tailor their configuration needs. The InCloudRack system utilizes all the I8000 base technology in a single SKU solution that enables Inspur to claim a foothold in the integrated system market. Inspur targets cloud service providers primarily through its SmartRack technology, which is a dense, modular, multinode design that has evolved steadily over the past six years, with the ability to mix and match different generations. For extreme density needs, the i4008 design can support up to eight two-socket servers or four storage nodes in a single 4U footprint.

Inspur is a diverse provider with a keen investment in emerging data center technologies, which positions the company well for growing Mode 2 market opportunities. Inspur already has extensive success selling to Chinese customers in the hyperscale, HPC and financial sectors. But in order to gain international credibility, the company also needs to demonstrate its deep technical expertise to avoid being pigeonholed as a vendor that is only capable of addressing commoditized market needs. While most modular server revenue is based on x86, Inspur invests in alternative technologies to open up additional market opportunities. This includes active collaborative investment in ARM and IBM Power. Inspur also has a legacy product line of non-uniform memory access (NUMA)-based Itanium scale-up servers and a GPU computing solution, plus investment in solutions for Intel's Xeon Phi processor. The company is also an OpenPOWER Foundation member and a strong contributor to Project Scorpio that drives hardware compatibilities similar to those of the Open Compute Project. Inspur claims that over 50% of its server R&D resources are focused on modular server development, with commitments to improve infrastructure and resource management, and to offer complete hardware/software solutions for selected use cases.

STRENGTHS

Inspur is a long-standing vendor of modular servers, and has become the clear volume leader in China (outselling Lenovo, Huawei, and Western vendors like HPE and Cisco that have strong global presence).

In keeping with other Chinese vendors, Inspur has strong relationships with leading public cloud vendors like Alibaba, Qihoo 360 Technology, Tencent and Baidu, which helps to maintain an annuity revenue stream for its multinode server designs.

Inspur is a technology innovator, with exploratory investments in emerging communities like OpenPOWER and Project Scorpio, which will position the company well if market demand takes off.

CAUTIONS

Until very recently, Inspur has been inwardly focused and only present in the Chinese market, so its business is still very skewed toward a single national market.

Building an international channel and local sales and marketing infrastructure in multiple countries is a major exercise that will take Inspur many years to accomplish.

Inspur has weak relationships with top-tier global ISVs, which limits the applicability of its products to Mode 1 data center modernization.

Lenovo

The x86 server business acquisition from IBM has brought Lenovo a global-class portfolio of x86 servers and support capabilities, covering 160 countries. Lenovo's modular servers target a wide range of workloads, including cloud, big data/analytics and HPC. Above all, its HPC business has considerable influence in its modular server and storage sales. In addition, the company tries to leverage its expertise into enterprise and service providers. Lenovo's modular server portfolio mainly consists of IBM-inherited legacy products, including blade servers such as Flex System and BladeCenter and an innovative multinode server named NeXtScale System. Additionally, Lenovo has another multinode server line originally designed for Chinese hyperscale companies such as Tencent. Its new ThinkServer SD350 targets enterprise workloads that require greater density. In-house resource management tool XClarity holds servers, storage and network together. Lenovo has global-class service/support capabilities, including 11 global research labs, 3,000 service locations and 42 contact centers with 10,000 IT support specialists.

The emergence of software-defined anything (SDx) is eliminating hardware silos and requires system providers to work with partners. As such, the company has been expanding its market by partnering with key system players, as well as major ISVs such as SAP, VMware or Microsoft. Recent partnership announcements include Nutanix, SimpliVity, Juniper Networks and Nexenta to augment its converged infrastructure and software-defined offerings. Lenovo is also able to leverage strong historical relationships with system vendors like EMC, as well as its continued alliance with IBM. However, the future of its EMC relationship will be influenced by how the Dell acquisition of EMC evolves.

The company plans to launch a new brand to unify all portfolios, including IBM-inherited products and its legacy lines. The new brand will launch in 2017 and all server products will be refreshed under it. The new brand launch will bring fresh looks to Lenovo's products and create a unified approach for its server business. However, it also means that the company may lose the strong brand values established from the legacy System x and ThinkServer eras.

STRENGTHS

Lenovo has a very broad portfolio of proven products with installed-base strength in most countries.

The Flex System and NeXtScale brands help to position Lenovo as a proven player in the modular server market.

By incorporating several thousand ex-IBM staff around the world, Lenovo has been able to achieve immediate localized expertise and the capability to execute complex projects.

CAUTIONS

Lenovo needs to increase its server revenue to demonstrate the validity of the acquisition. The company's 2015 modular server revenue (and its total x86 server revenue) was measurably lower than the combined revenue from IBM and Lenovo in 2014.

With many ex-IBM executives' departures, Lenovo needs to demonstrate its ability to succeed in the competitive landscape (especially for its high-end servers) without the legacy from IBM.

As in the cases for other Chinese providers, geopolitical issues between the United States and China will prevent server sales in particular verticals, such as government and telecom in the U.S.

SGI

Twenty-five percent of the participants in this Magic Quadrant are vendors that have achieved their greatest market success by focusing on the narrow – but enduring – market for high-performance computing. SGI is such a vendor, bringing deep supercomputing experience to the broader modular server market. So, while SGI is a recognized brand name and widely known in HPC, the brand remains unfamiliar to many enterprises who are planning Mode 1 data center transformations.

SGI, as we know it today, is the outcome of a 2009 acquisition of Silicon Graphics Inc. by Rackable Systems, a company with considerable previous experience in the market for multinode servers. In the HPC market, high-volume general technologies and products have been displacing highly differentiated and high-margin offerings for some years, and Rackable Systems saw the opportunity to achieve synergy through the acquisition of an established HPC vendor.

Seven years after the acquisition, SGI's modular server family has three main ranges. The SGI Rackable range is the evolution of the preacquisition Rackable Systems portfolio. Rackable servers are sold mainly for HPC and other compute-intensive scale-out clusters that can be addressed by small, stand-alone nodes. At the other extreme, SGI's Integrated Cluster Environment (ICE) X and XA systems are aimed at extreme scale-out HPC workloads, with the ability to support up to 288 sockets per rack and to extend to thousands of nodes. But the server platform that is helping SGI to gain the most new business attention is the UV range. UV servers are available in various configurations that can scale up to 256 sockets and 64TB of memory in a single image, if required. While most UV systems are deployed for HPC projects that require scale-up node sizes beyond the scope of regular cluster nodes, the technology is proving to be adaptable for other roles. SGI has certified a 20-socket/15TB memory configuration for SAP Hana use, with certification pending to scale up to 32 sockets and 24TB of memory. With so much industry momentum for SAP Hana, there is a new awareness and credibility for SGI as a potentially viable vendor for enterprise choice. SGI has also been able to sign up Dell and Cisco as partners, which creates further opportunities to leverage the channels of those vendors as well. In addition, SGI has certified a 32-socket/24TB memory configuration for Oracle Database 12c In-Memory DBMS deployment, and has recently signed an OEM agreement with HPE for HPE to sell an eight-socket variant of the system.

These efforts are helping SGI to become more recognizable to the general IT market as a viable vendor, and, in doing so, become less dependent on the variability of the HPC market that can create frequent spikes and troughs in revenue for more specialized vendors like SGI.

STRENGTHS

SGI is one of the most instantly recognizable HPC brands, with a strong product portfolio that provides an annuity revenue stream from the HPC segment.

The company is expanding its market reach with enterprise-oriented strategies that are gaining strong recognition for SAP Hana and other in-memory compute workloads.

SGI's technology innovation is creating partnering opportunities with top data center vendors, which provides SGI with additional market visibility and revenue.

CAUTIONS

Most SGI business is focused on high-value deals that entail longer sales cycles and project implementation times that extend the period by which SGI can recognize revenue.

Business results are most influenced by a relatively low volume of high-value deals, which creates quarter-to-quarter financial variability that is accentuated by SGI's modest size and specialized market focus.

As more and more HPC workload demand is met by predominantly standard high-volume x86-based compute clusters, vendors like SGI face the potential for growing competition from incumbent data center vendors like Dell, HPE and Lenovo, plus a new generation of low-cost rivals like Supermicro.

Supermicro

The Supermicro brand is less-known to large enterprises, but it is an established and recognized server logo for SMBs and midmarket enterprises, as well as external IT service providers and hyperscale organizations. The name is also becoming increasingly identifiable thanks to the market success of hyperconvergence vendors like Nutanix, whose appliances are frequently built by Supermicro under contract.

The company has grown steadily for more than 20 years by focusing on the provision of optimized solutions for specific customer groups. Enterprise buyers have not been a priority until quite recently, and the company has never needed to invest in brand awareness. But Supermicro is now targeting a broader market need with a portfolio of products that includes multiple families of modular server offerings.

Because Supermicro is identified with low-end volume servers, it is easy to assume that the company adds little technology value to its products. This perception is wrong. Supermicro's strength lies in its engineering skill and custom design. Unlike many competitors, Supermicro's design and testing is all done in-house, which ensures strong IP ownership and control, plus the ability for the company to be early (if not first) to market with new technology offerings, such as Intel Xeon Phi-based servers and GPU-based blades that use the nVidia Tesla chipset.

Supermicro's modular products focus on density of computing and are available globally. But the company needs to expand the depth and breadth of channel programs to sell directly to enterprises, and it needs to develop stronger alliances with top-tier integrators and ISVs, whose allegiance will be essential for Supermicro to become a viable global threat to enterprise stalwarts like HPE, Dell and Cisco. Brand recognition also needs to be improved, although this is happening to a degree through the indirect visibility that Supermicro is gaining thanks to the market momentum of the hyperconvergence vendors.

STRENGTHS

Supermicro's product portfolio breadth is extensive and comprises modular server offerings including the Twin, FatTwin, TwinPro, MicroCloud, MicroBlade and SuperBlade products.

Supermicro is the primary hardware provider behind Nutanix's success in the hyperconverged market, and has become the prevalent choice for most other HCIS startups.

Supermicro has had a strong track record with hyperscale customers, plus enterprises and government agencies implementing HPC applications. These successes position Supermicro well to benefit from the anticipated wave of investment in Mode 2 workloads like big data analytics, public cloud infrastructure, mobility and digital enterprise.

CAUTIONS

Supermicro's relationships with top-tier enterprise solution ISVs and managed service providers/outsourcers are often immature and not well-understood by enterprise buyers. The company lacks the perceived track record that enterprises value, and needs to significantly increase investment in relationship building and the promotion of relevant success stories.

For enterprises, Supermicro's own management software tools are regarded as weak compared to those of other leading vendors. This creates complexity where operations managers are more dependent on third-party tools, which impacts Supermicro's ability to differentiate from other vendors.

Supermicro is investing in a separate focused product strategy for enterprise customers. But until this strategy is better-understood, the huge company product portfolio is more likely to be regarded as a drawback by enterprises, who value a clear and concise set of marketing messages and believe that a complex portfolio creates constraints with respect to inventory and product quality.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

Added

Inspur has been added to this year's Magic Quadrant. The company was cited last year for its broad portfolio and strong focus on the Chinese market. Inspur is now investing heavily in a more international focus, hence its inclusion.

Dropped

No vendors have been dropped for 2016.

Inclusion and Exclusion Criteria

The inclusion criteria have remained similar to those used in 2015. The inclusion criteria represent the specific attributes that analysts believe are necessary for inclusion in this research. Vendors must meet all criteria in order to be included:

Servers included in the assessment must satisfy the modular server definition described above. Some vendors will market a combination of blade and multinode designs; others will address a more focused market. But the role of some sort of chassis or enclosure for easy and rapid server addition or replacement is essential.

Servers included for assessment must be publicly available SKUs, therefore excluding "private" SKUs or custom-engineered or -designed systems created for a single company, or closed community group of companies.

Servers included must be marketed and sold under the vendor's brand name. Not included in the ratings are products that a provider chooses to resell, where products are designed and manufactured by another OEM. But vendors may choose to use contract manufacturers to build their products.

Vendors must achieve at least \$50 million in revenue for modular server sales during the period of 4Q14 through 3Q15, based on our published server market statistics estimates.

This research evaluates vendors on a global basis. Vendors must be actively marketing their modular servers in at least two or more country markets as defined by Gartner's standard geographies; we are looking for a guideline of around 5% of revenue coming from one or more countries outside the vendor's home country. Vendors focused on a single domestic market would be excluded, regardless of revenue size.

Products must be deployed in at least 10 live customer production environments – as of 1 January 2016 – and customer references must be available for potential follow-up.

There are many more vendors in the market than this research can present. Some vendors will be significant participants in a single domestic market, but lack global reach. But for some end-user clients, these vendors may deserve consideration. NEC is such a vendor.

NEC is a recognized global server vendor, but is still struggling to expand its modular server business outside of Japan. The company's modular server portfolio shows progress with both the established Express5800 line (Express5800/B and E series) and the new DX Series. The DX1000 adds to the current Intel Atom-based Micro Modular Server line by implementing 5,600 cores per rack. Meanwhile, the new Scalable Modular Server, DX2000, is based on Intel Xeon Processor D and focuses on Mode 2 workloads, such as real-time processing (big data, in-memory database), bare metal cloud and the Internet of Things (IoT). NEC has a strategic initiative with Hortonworks to target big data analytics, using the brand "Data Platform for Hadoop." This integrated system comprises Apache Hadoop and Apache Spark. NEC's global marketing strategies come across as inconsistent by region, which contributes to weak sales of modular servers due to the lack of partner channels. NEC needs to redefine its strategies to target global opportunities and growth markets that will help position it as a true innovator, rather than the current perception that regards NEC as largely irrelevant outside its domestic market.

Evaluation Criteria

Ability to Execute

Gartner analysts evaluate providers on the quality and efficacy of the processes, systems, methods or procedures that enable IT provider performance to be competitive, efficient and effective and to positively impact revenue, retention and reputation. Ultimately, providers are judged on their ability and success in capitalizing on their vision. The weightings for market responsiveness/record and operations have been increased this year, to reflect the growing importance of excellence in these categories.

Table 1. Ability to Execute Evaluation Criteria

Evaluation Criteria	Weighting
Product or Service	High
Overall Viability	High
Sales Execution/Pricing	Medium
Market Responsiveness/Record	High
Marketing Execution	Medium
Customer Experience	High
Operations	Medium

Source: Gartner (May 2016)

Completeness of Vision

Gartner analysts evaluate providers on their ability to convincingly articulate logical statements about current and future market direction, innovation, customer needs and competitive forces, and how well they map to the Gartner position. Ultimately, providers are rated on their understanding of how market forces can be exploited to create opportunity for the provider. The weighting for marketing strategy has been increased this year, to reflect the growing importance of excellence in this category.

Table 2. Completeness of Vision Evaluation Criteria

Evaluation Criteria	Weighting
Market Understanding	High
Marketing Strategy	High

Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	Low
Innovation	High
Geographic Strategy	Low

Source: Gartner (May 2016)

Quadrant Descriptions

Leaders

Leaders in the modular server market distinguish themselves by offering a service suitable for strategic adoption and by having an ambitious roadmap. They have a track record of successful delivery, significant market share and many referenceable customers. Because the market is loosely bifurcated between blade servers and multinode servers, the Leaders can serve a broad range of use cases. They do not have to excel in all areas, may not necessarily be the best providers for a specific need and may not serve some use cases at all. But as market demand switches toward multinode servers, Leaders need to have a viable strategy to address growing Mode 2 market opportunities.

Challengers

Challengers are well-positioned to serve most current market needs, and should be regarded as "Leaders in waiting" if they can resolve market challenges presented by the Leaders. They may also be traditional segment leaders who are facing short-term challenges. They will usually be proven vendors with the ability to address a particular set of use cases, and they will have a track record of successful delivery. However, they are not adapting to market challenges quickly enough or do not have a broad enough scope of ambition.

Visionaries

Visionaries have an ambitious vision of the future and are making significant investments in the development of unique technologies. Visionaries may be new market entrants or they may be existing providers who are reinventing their business. Vendors who only address emerging Mode 2 market demands would typically be Visionaries because the market itself is only at a nascent stage. Their services are still emerging, and they have many capabilities in development that are not yet generally available. While they may have many customers, they might not yet serve a broad range of use cases well.

Niche Players

The Niche Players in this Magic Quadrant are all viable providers for the use cases in which they specialize, but may not serve a broad range of use cases well or have a broadly ambitious roadmap. They may be relatively new entrants to this market or may not yet have gained significant market share. The more highly targeted and specialized your needs, the more likely it is that there will be a Niche Player that may be more appropriate for your situation than a vendor in another Magic Quadrant position.

Context

Data center managers should determine the suitability of modular servers for their needs because not all workloads are appropriate. Modular servers generally address scale-out workloads and deliver many facilities-related benefits like optimum floor space utilization, excellent power density and cooling efficiency, and very granular component replacement and addition. Having established the business case for modular servers, data center managers must then determine whether blade servers or multinode servers are the best fit. Blade servers are best-suited for highly virtualized workloads and data center modernization strategies. Following Gartner's bimodal IT approach, this makes blade servers a more natural fit for Mode 1 scenarios. Multinode servers are more naturally oriented toward Mode 2 workloads, with most early demand coming from hyperscale companies and enterprises investing in public cloud, analytics and HPC workloads.

As hyperconvergence becomes an ever greater portion of the integrated systems market, this also puts pressure on the longer-term viability of blades. Most early generation integrated systems are based on blades and SANs, and a viable market for such technology will continue through this decade. But with hyperconvergence growing rapidly, the need to invest in integrated systems will no longer automatically be an opportunity for blade vendors.

Blade server buyers should place particular emphasis on the age of the chassis (also known as enclosure); that technology has a typical nine- to 10-year life before vendors have to develop a new design. At this point, it becomes counterproductive to keep adding innovations to a dated chassis that will not be ideal for the degree to which the market has evolved. In addition, data center managers should not automatically assume forward or backward compatibility between chassis designs. When a chassis is due for replacement, they can expect a period of disruption to follow. Data center managers who have invested in a superseded chassis rarely have a need to migrate when a new chassis is launched. Expect that all responsible vendors will offer at least five years of support, and most vendors now offer at least a year or two where both the old and the new can coexist on their price lists. But it is inevitable that users of an older chassis need to start planning for migration once a new design is launched. So users should always verify the age of chassis technology and the resulting impact on roadmaps. Multinode servers are not immune from this effect, as they are also based on chassis-based technology. So buyers also have to consider the impact of a chassis replacement in their vendor assessments — particularly regarding forward or backward compatibility between designs.

Market Overview

Modular servers are a class of server with modular compute capabilities housed in a common chassis or enclosure. The enclosure is usually a proprietary design owned by the vendor concerned. For Gartner to consider a server product as "modular," the product must have a chassis or enclosure that allows for the easy and rapid addition or replacement of servers (and potentially other components, like storage).

In 2015, the modular server market accounted for 27.2% of server sales, putting the modular server market at \$15.2 billion, with growth of 14.1% over 2014. Gartner estimates the modular server market will be worth \$18.0 billion by 2019, accounting for 30% of server spending and growing at a compound annual growth rate (CAGR) of 4% (2015 to 2019).

The primary adopters of modular servers are large enterprises and, most recently, external IT service providers. Among these two groups, external IT service providers, including hyperscale providers, have witnessed the fastest growth in modular server investment.

All of the 12 vendors assessed in this research have viable strategies and proven track records in this sector.

Evidence

"Market Share: Servers, Worldwide, 4Q15 Update"

Evaluation Criteria Definitions

Ability to Execute

Product/Service: Core goods and services offered by the vendor for the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether offered natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability: Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will advance the state of the art within the organization's portfolio of products.

Sales Execution/Pricing: The vendor's capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support, and the overall effectiveness of the sales channel.

Market Responsiveness/Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization's message to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This "mind share" can be driven by a combination of publicity, promotional initiatives, thought leadership, word of mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure, including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' wants and needs and to translate those into products and services. Vendors that show the highest degree of vision listen to and understand buyers' wants and needs, and can shape or enhance those with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling products that uses the appropriate network of direct and indirect sales, marketing, service, and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including vertical markets.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes.

Geographic Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the "home" or native geography, either directly or through partners, channels and subsidiaries as appropriate for that geography and market.



(<http://gtnr.it/1KsfgQX>)

© 2016 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. or its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. If you are authorized to access this publication, your use of it is subject to the Usage Guidelines for Gartner Services (/technology/about/policies/usage_guidelines.jsp) posted on gartner.com. The information contained in this publication has been obtained from sources believed to be reliable. Gartner 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 publication consists of the opinions of Gartner's research organization and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice. Gartner provides information technology research and advisory services to a wide range of technology consumers, manufacturers and sellers, and may have client relationships with, and derive revenues from, companies discussed herein. Although Gartner research may include a discussion of related legal issues, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner is a public company, and its shareholders may include firms and funds that have financial interests in entities covered in Gartner research. Gartner's Board of Directors may include senior managers of these firms or funds. Gartner research is produced independently by its research organization without input or influence from these firms, funds or their managers. For further information on the independence and integrity of Gartner research, see "Guiding Principles on Independence and Objectivity." (/technology/about/ombudsman/omb_guide2.jsp)"

About (<http://www.gartner.com/technology/about.jsp>)

Careers (<http://www.gartner.com/technology/careers/>)

Newsroom (<http://www.gartner.com/newsroom/>)

Policies (http://www.gartner.com/technology/about/policies/guidelines_ov.jsp)

Privacy (<http://www.gartner.com/privacy>)

Site Index (<http://www.gartner.com/technology/site-index.jsp>)

IT Glossary (<http://www.gartner.com/it-glossary/>)

Contact Gartner (http://www.gartner.com/technology/contact/contact_gartner.jsp)