

The Forrester Wave™: AI Infrastructure, Q4 2021

The 13 Providers That Matter Most And How They Stack Up

December 13, 2021

By Mike Gualtieri, Tracy Woo with Srividya Sridharan, Caroline Provost

FORRESTER®

Summary

In our 25-criterion evaluation of AI infrastructure providers, we identified the 13 most significant ones — Amazon Web Services (AWS), Dell, Exxact, Google, Hewlett Packard Enterprise (HPE), IBM, Inspur Information, Lambda, Microsoft, NVIDIA, Run:AI, Spell, and Supermicro — and researched, analyzed, and scored them. This report shows how each provider measures up and helps infrastructure operations and technology architecture professionals select the right one for their needs.

Additional resources are available in the [online version](#) of this report.

AI Infrastructure Keeps AI Teams Working, Not Waiting

AI is rocking the world. It's rapidly gone from "if" to "when" to "now." [Seventy-four percent](#) of surveyed global data and analytics decision-makers whose firm is implementing or expanding its use of AI say that AI adoption has had a positive impact. AI platform software that many of these enterprises rely on is all well and good, but it needs hearty infrastructure — compute, storage, networking — to keep AI teams working, not waiting. IT ops professionals are the key decision-makers because they are already a company's primary infrastructure buying experts. This evaluation will help them make the best AI infrastructure decisions for their organization to handle AI-specific workloads.

As a result of this momentum in AI adoption, AI infrastructure buyers should look for providers that:

- **Design for diverse AI workloads.** AI is not one single workload. It is three workloads: data management, training, and inferencing. Each of these workloads has starkly different characteristics for throughput, latency, and enterprise qualities. Training workloads require infrastructure optimized for machine learning (ML) algorithms to analyze large and/or complex data sets. Inferencing, on the hand, must be optimized for models used in production applications, and those requirements can be any combination of requirements such as high-concurrency, low-power edge, and/or model type. Enterprise buyers should understand how vendor solutions map to each of these workloads specific to their current and future workloads. AI infrastructure solutions in this evaluation cover all workloads; however, it will make sense for enterprises to choose multiple vendors based on specific needs. For example, an enterprise might choose an on-premises solution for data management and training but choose a hyperscaler for inferencing, and vice versa.
- **Keep up with breakneck innovation.** Far from excelling on commodity hardware, AI training and inferencing workloads run best on purpose-built systems. That's because of how the underlying algorithms work to analyze data to create models — especially deep learning algorithms. This has initiated a renaissance in hardware innovation focused on AI infrastructure, including AI accelerator chips. NVIDIA is well-known for dominating the market for deep learning training with its graphics processing units (GPUs). To wit, every vendor in this evaluation supports NVIDIA GPUs. But it is early days, and a competitive market has emerged with Intel offering a full range of AI accelerators, including Habana Gaudi. And VC money is

pouring into startups like Cerebras, Graphcore, and others. IBM Research is developing “brain-inspired” analog AI chips. Even the hyperscalers have designed AI accelerators, such as Google’s tensor processing unit (TPU) and AWS Inferentia. Enterprise buyers must understand a vendor’s roadmap to incorporate this rapid pace of innovation in addition to what meets their organization’s needs today.

- **Manage complexity with software.** Managing the many moving parts of IT infrastructure is difficult enough. Now add in massive amounts of data, performance and latency requirements, and security measures for a highly complex environment, all of which makes AI infrastructure management a monumental task to take on manually. Luckily, most AI infrastructure systems come with a management layer that helps abstract away the complexity. Table stakes capabilities include role-based access control (RBAC), platform and resource scheduling, dashboard visibility for the entire system, integration with common third-party logging and monitoring solutions, and authentication protocols. Enterprise buyers must understand how a vendor’s AI infrastructure management layer can be incorporated with its existing infrastructure management tools, policies, and IT ops practices. If an enterprise has already standardized on a vendor’s non-AI infrastructure, then using that vendor’s AI infrastructure can be attractive.

Evaluation Summary

The Forrester Wave™ evaluation highlights Leaders, Strong Performers, Contenders, and Challengers. It’s an assessment of the top vendors in the market and does not represent the entire vendor landscape. You’ll find more information about this market in our reports on [AI infrastructure](#).

We intend this evaluation to be a starting point only and encourage clients to view product evaluations and adapt criteria weightings using the Excel-based vendor comparison tool (see Figure 1 and see Figure 2). Click the link at the beginning of this report on Forrester.com to download the tool.

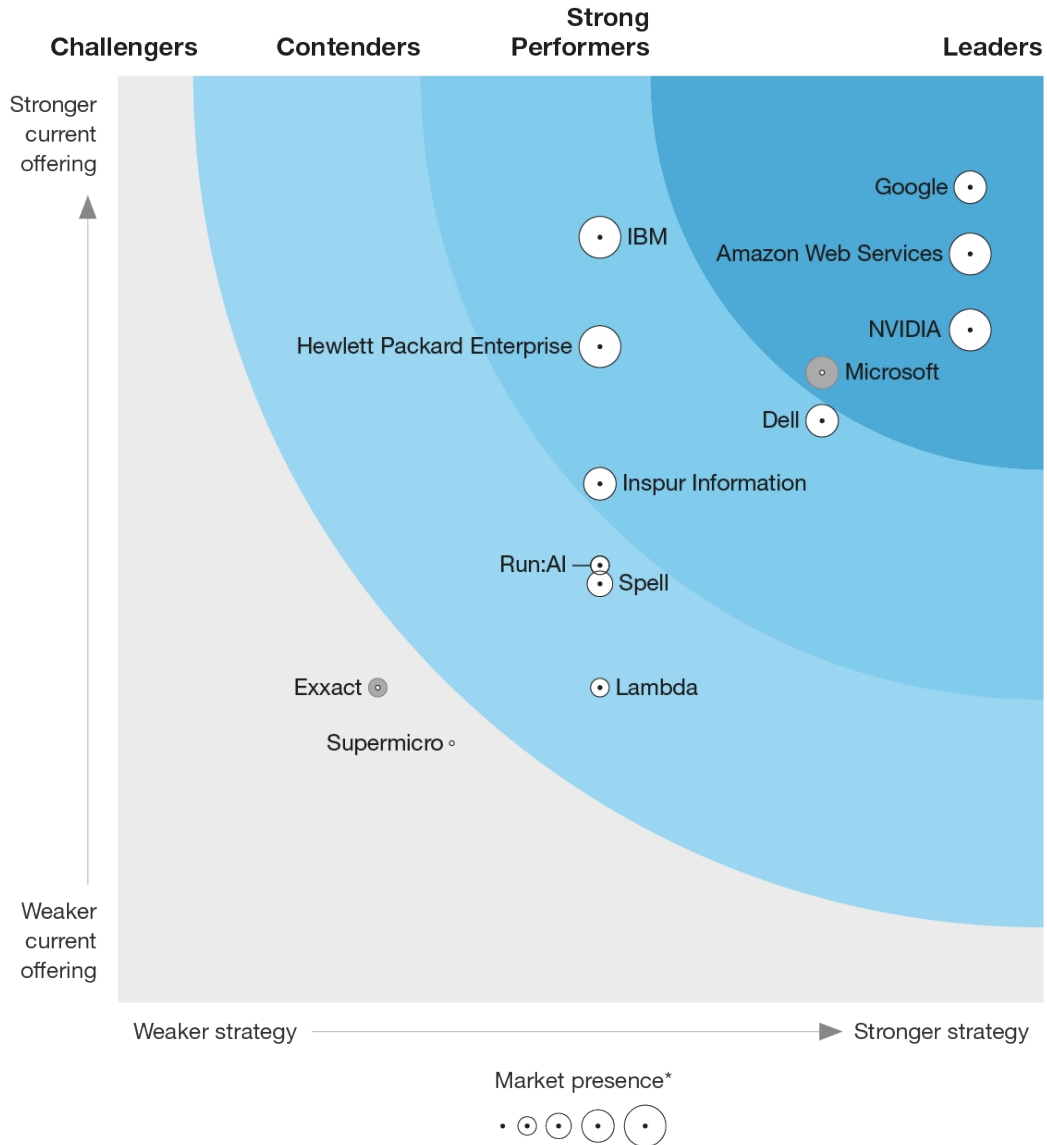
Figure 1

Forrester Wave™: AI Infrastructure, Q4 2021

THE FORRESTER WAVE™

AI Infrastructure

Q4 2021



*A gray bubble or open dot indicates a nonparticipating vendor.

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners. For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

Figure 2

Forrester Wave™: AI Infrastructure Scorecard, Q4 2021

	Forrester's weighting	Amazon Web Services	Dell	Exxact*	Google	Hewlett Packard Enterprise	IBM	Inspur Information
Current offering	50%	4.04	3.14	1.70	4.40	3.54	4.13	2.80
Architecture	20%	5.00	5.00	3.00	5.00	5.00	5.00	4.00
Training	20%	4.50	3.00	2.50	5.00	4.00	4.00	3.00
Inferencing	20%	4.34	3.00	1.00	4.32	3.68	4.32	2.34
Management	20%	4.00	3.00	1.00	4.00	2.00	3.00	3.00
Deployment	20%	2.34	1.68	1.00	3.66	3.00	4.34	1.68
Strategy	50%	4.60	3.80	1.40	4.60	2.60	2.60	2.60
Execution plan	20%	5.00	3.00	1.00	5.00	1.00	3.00	3.00
Innovation roadmap	20%	5.00	3.00	1.00	5.00	1.00	3.00	3.00
Supporting products and services	20%	3.00	5.00	1.00	3.00	5.00	5.00	3.00
Partner ecosystem	20%	5.00	5.00	1.00	5.00	5.00	1.00	3.00
Commercial model	20%	5.00	3.00	3.00	5.00	1.00	1.00	1.00
Market presence	0%	4.67	3.67	1.67	4.00	4.67	4.33	3.67
Revenue	34%	5.00	4.00	2.00	4.00	5.00	4.00	4.00
Number of customers	33%	5.00	4.00	2.00	5.00	4.00	4.00	4.00
Breadth of offering	33%	4.00	3.00	1.00	3.00	5.00	5.00	3.00

All scores are based on a scale of 0 (weak) to 5 (strong).

*Indicates a nonparticipating vendor.

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

	Forrester's weighting	Lambda	Microsoft*	NVIDIA	Run:AI	Spell	Supermicro*
Current offering	50%	1.70	3.40	3.63	2.36	2.26	1.40
Architecture	20%	3.00	4.00	4.00	2.00	2.00	2.00
Training	20%	2.50	4.00	4.00	2.00	2.00	2.00
Inferencing	20%	1.00	3.66	3.66	3.00	3.00	1.00
Management	20%	1.00	3.00	3.50	2.50	2.00	1.00
Deployment	20%	1.00	2.34	3.00	2.32	2.32	1.00
Strategy	50%	2.60	3.80	4.60	2.60	2.60	1.80
Execution plan	20%	3.00	5.00	5.00	3.00	3.00	1.00
Innovation roadmap	20%	3.00	3.00	5.00	3.00	3.00	1.00
Supporting products and services	20%	1.00	3.00	5.00	1.00	3.00	1.00
Partner ecosystem	20%	1.00	3.00	5.00	3.00	1.00	1.00
Commercial model	20%	5.00	5.00	3.00	3.00	3.00	5.00
Market presence	0%	2.00	3.33	5.00	1.66	2.32	1.00
Revenue	34%	2.00	3.00	5.00	1.00	1.00	1.00
Number of customers	33%	3.00	4.00	5.00	1.00	2.00	1.00
Breadth of offering	33%	1.00	3.00	5.00	3.00	4.00	1.00

All scores are based on a scale of 0 (weak) to 5 (strong).

*Indicates a nonparticipating vendor.

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

Vendor Offerings

Forrester included 13 vendors in this assessment: Amazon Web Services, Dell, Exxact, Google, Hewlett Packard Enterprise, IBM, Inspur Information, Lambda, Microsoft, NVIDIA, Run:AI, Spell, and Supermicro (see Figure 3).

Figure 3

Evaluated Vendors And Product Information

Vendor	Product evaluated	Product version evaluated
Amazon Web Services	AWS AI Infrastructure Services	
Dell	Dell Technologies Validated Design for AI	
Exxact	AMD Radeon PRO Platform	
Google	Vertex AI (ML platform); Deep Learning VM, Deep Learning Container, and TensorFlow Enterprise (ML runtime); Dataproc and Dataflow (GPU-powered data processing); Cloud GPU, Cloud TPU, and TPU Pods (ML accelerators)	
Hewlett Packard Enterprise	HPE Greenlake; Determined AI; HPE Ezmeral	
IBM	IBM Z; IBM Z15; IBM LinuxONE III; IBM Power; IBM Power AC922; IBM Power IC922; IBM Spectrum Scale; IBM Elastic Storage System; IBM Cloud; IBM Cloud Pak for Data; IBM Cloud Pak for Data Systems	
Inspur Information	Inspur Full-Stack AI Portfolio	
Lambda	Lambda Stack	
Microsoft	Microsoft Azure	
NVIDIA	NVIDIA DGX Systems	NVIDIA DGX A100
Run:AI	Run:AI Cloud-Native Compute Management Platform	Run:AI 2.0
Spell	Spell ML Platform	43.7
Supermicro	NVIDIA Ampere; Habana Gaudi	

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

Vendor Profiles

Our analysis uncovered the following strengths and weaknesses of individual vendors.

Leaders

- **Google offers tightly integrated AI software with elastic infrastructure.** Google delivers its platform and services via a vast portfolio of AI solutions, ML-specific

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.

For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

cloud-native instances (deep learning virtual machines and containers), and ML accelerators (GPU, TPU). To address the increasing diversity of AI workloads, Google created its own TPUs but also offers GPUs. It has a wide range of GPU and TPU accelerators available in preemptible machines and linearly scaling pods. It offers accessible AI solutions built on top of its Vertex AI platform. Google's strategic strengths include its advanced AI and ML capabilities. However, Google can do a better job explaining its strategy for using its own AI chips.

Google's strengths in its current offering are in architecture, training, data throughput, and latency. Its sweet spot is in its product offering, Vertex AI, which has core AI compute capabilities and machine learning operations (MLOps) services for end-to-end AI lifecycle management. Still, its AI solution portfolio has gaps. The company is bolstering its security capabilities and working to provide better compatibility with external tooling. Reference customers like its data warehousing capabilities but would like to see more data and model management. Google is a best fit for companies that are looking for top-of-the-line deep learning capabilities in the public cloud.

- **AWS offers a superior fully managed platform.** AWS's capabilities codify into three pillars: frameworks and infrastructure; Amazon SageMaker; and AI services, which includes its suite of off-the-shelf models, developer APIs, and business solutions. The company claims "tens of thousands" of customers are now standardizing on Amazon SageMaker, its fully managed platform to build, train, and deploy ML models. At the infrastructure level, AWS offers P4d instances that feature the latest A100 GPUs, 400 Gbps networking, and 8 terabytes of local NVMe storage. AI and ML have been a major strategic focus for the company as it works to tackle the biggest enterprise challenges, such as the ML skills gap, lack of advanced ML infrastructure, and missing optimal use cases. AWS continues to build out its AI/ML portfolio with a strong execution and innovation roadmap. And it has a massive ecosystem and partnership to pull on for gaining further traction in the market.

AWS offers AI services that range from managed services to industry-specific solutions meant for building, deploying, and managing ML applications. Its strengths include its architecture design, architecture components, and training data and software. However, it still needs to build out stronger management capabilities. Reference customers value the vendor's deep learning chops and pace of innovation but would like to see more model management capabilities and premade algorithms. AWS is a good solution for enterprises heavily invested in its public cloud or looking for leading-edge public cloud AI capabilities.

- **NVIDIA's DNA is in every other AI infrastructure solution we evaluated.** It's an

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.

For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

understatement to say that NVIDIA GPUs are synonymous with AI infrastructure. Breakthroughs in deep learning around 2012 brought AI into focus, but only NVIDIA had the strategy, vision, and roadmap to invest in supporting these now mainstream AI workloads. Every other vendor in this evaluation supports NVIDIA GPUs and related components. NVIDIA also offers its own NVIDIA DGX systems engineered specifically for AI workloads.

NVIDIA has strengths in architectural components, throughput, latency, and overall product strategy. The vendor's sweet spot for its DGX systems are for customers that want a complete system that is engineered by NVIDIA to include its latest component technology for AI workloads. The vendor's primary focus has been on training workloads, and it recently added additional capabilities for inferencing workloads. It is in the catbird seat, but other chip giants such as Intel and several startups are on the verge of providing competitive alternatives. Reference customers appreciate the vendor's thought leadership in AI, its frameworks designed to run on NVIDIA GPUs, and having first access to chips coming out of NVIDIA R&D. The vendor's sweet spot is for customers that want a system engineered by NVIDIA.

- **Microsoft accommodates any size and type of AI workload.** Microsoft Azure offers a full range of AI capabilities, including an ML platform that AI developers can use to build ML models, pretrained "cognitive services" for out-of-the-box functionality, and AI infrastructure. The company's AI infrastructure includes the latest accelerators for both training and inferencing workloads and can be scaled from a single instance to a cluster of thousands of nodes. Microsoft is known for providing the AI infrastructure for the huge GPT-3 language model, which was trained on 45 TB of data using 285,000 CPU cores and 10,000 GPUs. Microsoft's strategy is to offer its Azure customers a cost-effective infrastructure for early experimentation to regular production AI workloads. This AI infrastructure strategy dovetails with Azure's larger portfolio of services to leverage that massive and loyal Microsoft developer community to build the next generation of AI applications.

Microsoft has strengths in architecture, components, software, training throughput, inferencing latency, and management tools. It is best fit for enterprises that want an elastic cloud-based solution that supports both training and production workloads. Microsoft declined to participate in the full Forrester Wave evaluation process; this assessment is based on publicly available information.

Strong Performers

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.

For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

- **Dell has a full-stack AI solution but shows weaknesses in its software**

capabilities. Dell offers a vast and broad infrastructure portfolio that spans compute, storage, networking, data protection, converged infrastructure, and hyperconverged infrastructure. The company has also grown its software side by recent expanding its CloudIQ AIOps platform as well as identifying the five levels of ML-based automation, which range from AI insights to full automated remediation. Dell continues to build on its AI portfolio, as AI is one of the company's growth pillars for its future strategy. While the company is strong in its performance, it lacks differentiation in its execution and innovation roadmap.

Dell's strengths are in its infrastructure, particularly its high-performance storage. Its sweet spot is in its architecture and software capabilities in operations management and monitoring. However, the vendor faces stiff competition from the other infrastructure providers that have stronger software capabilities. Dell is best suited for enterprises with a large on-premises footprint looking to bridge to public cloud.

- **IBM envisions and delivers transactional AI but lacks a unified message.**

IBM's AI portfolio includes the Spectrum computing family for cluster management and a hyperconverged infrastructure appliance for big data, warehouse, and AI applications, as well as IBM's z15 and Power10. IBM offers configurable cores via z15 and has accelerated platforms for training and inferencing class GPUs. The company bases its products on the philosophy of "scale-out and scale-up systems" for all AI workload types, which include high-performance computing, big data, and configurable cores. Its products are designed to serve hybrid cloud environments and support all workloads through bare metal and virtual servers. IBM has committed to AI as one of its major growth strategies and has made key acquisitions (Instana, myInvenio, Turbonomic) to build on that portfolio. Although IBM has impressive AI capabilities, particularly with its systemwide, silicon-level optimizations, its product offerings lack a central unifying theme.

IBM has strengths in architecture, training, inferencing latency, and deployment efficiency. Its sweet spot is in its end-to-end data pipelines and data stores, which are used for real-time inferencing and high-performance training. Reference customers like IBM for its pace of innovation and overall performance. The vendor is best fit for enterprises that have a massive IBM footprint and for those looking for a highly performing and innovative team.

- **HPE delivers AI at scale but requires bolstering in its management**

features. Hewlett Packard Enterprises' 2019 Cray acquisition bolstered its strength in the supercomputing and AI market. HPE retained the Cray name to show its focus on leading-edge, high-performance computing technology. The same

technology underlies its strengths in performance along with its large partner ecosystem. The company's AI portfolio includes its HPE Greenlake managed service offering, software solutions, and platforms (via Determined AI and HPE Ezmeral) that are delivered on its compute and storage solutions. However, HPE lacks a clearly defined execution or innovation roadmap, showing an uncertain path forward to growth.

HPE's strengths are its architecture, training data and throughput, and software inferencing. Its weaknesses in its current offering are its breadth of integrations with third-party platforms and its security capabilities. Still, reference customers like the pace of innovation and cost flexibility. HPE is best suited for enterprises looking for a strong infrastructure backbone combined with a fast pace of innovation.

- **Inspur Information is a major IT solution provider in China but lacks traction abroad.** Inspur Information is the leading server provider in China with its Inspur Full-Stack AI Portfolio. For software it provides a full-scale suite via AIStation, which enables model development and training model deployment and inferencing. It builds and delivers a full-stack AI solution that ranges from servers to software (e.g., training/inference platform, algorithm toolkit). Inspur Information also plays a big role in the open source community by supporting a large number of open source hardware projects. To shorten deployment times, its hardware and software can be deployed in an integrated appliance. Inspur Information's strategic sweet spot is its major presence in China, which will serve it well as one of the few integrated AI appliance makers in that region. Despite this, it still lacks significant depth in its partner ecosystem or traction outside of the Chinese market.

Inspur Information's strength is in the breadth of its product portfolio, especially its integrated appliance solution. Reference customers highlight the company's full-stack capabilities but would like to see better delivery mechanisms and more-specific industry packaging. Inspur Information is a good fit for enterprises that want AI capabilities and would like delivery in the Chinese market.

Contenders

- **Run:AI virtualizes AI infrastructure, but customers must BYOH.** Run:AI is a software-only AI infrastructure solution designed to orchestrate AI workloads on customer-supplied AI infrastructure hardware — i.e., bring your own hardware (BYOH). Run:AI's customers provide their own hardware, whether on-premises, cloud, or both, and then use Run:AI software to provide policy-based scheduling, elastic pools of virtualized GPUs, and infrastructure monitoring.

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.
For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

Run:AI has strengths in management tools, training, and flexibility in components. Reference customers like its ability to finely define how to prioritize and share resources among larger teams. Run:AI's software-only position is differentiated for AI infrastructure, but it must add even more software value that carefully straddles the line between AI infrastructure hardware and AI platforms. Run:AI's sweet spot is enterprises that want complete flexibility in the hardware they choose to use and where they choose to run it.

- **Spell unifies MLOps with AI infrastructure, but some features overlap with ML platforms.** Spell bills itself as “MLOps for deep learning.” To wit, the company is included in this evaluation as a software-only AI infrastructure solution because it operates AI workloads on customer-supplied AI infrastructure hardware. Spell's strategy is to eliminate the friction between ML engineering and the infrastructure it must run on.

Spell shines in providing software that manages the many iterative ML training experiments that are typical in AI solution development. The company also has strengths in management, efficiency, and flexibility in components. Reference customers appreciate its ability to define, launch, and manage experiments as a single workload. Spell's software-only position is differentiated for AI infrastructure, but it must forge closer partnerships with hardware vendors to fill gaps in the AI infrastructure market. Spell's sweet spot is for enterprises that want to unify how their AI teams use AI infrastructure with a common interface.

- **Lambda delivers desktop-ready deep learning but faces cloud competition.** Lambda builds and delivers servers and workstations, and its GPU Cloud offers virtual machines of GPUs with a few clicks. Lambda servers and workstations provide AI researchers and developers with the power of AI infrastructure — quite literally, in many cases, under their desk. Every Lambda system includes the Lambda Stack, which includes preloaded software: PyTorch, TensorFlow, CUDA, cuDNN, and NVIDIA drivers. The company's strategy to add a GPU cloud is wise, since AI teams will now consider the merits of using hardware and/or cloud services.

Lambda has strengths in architecture and on-premises systems. It continues to experience tremendous growth that tracks the popularity of AI and in particular GPUs. However, the company will start to face competition from large enterprise infrastructure companies as AI moves from experimentation to production. Reference customers appreciate Lambda's ability to deliver fast systems that are preloaded with the Lambda Stack. Lambda is best fit for enterprises that want to immediately empower their AI researchers and/or AI developers with blazing-fast workstations and/or servers.

Challengers

- **Supermicro offers highly customizable solutions but lacks value-added software.** Supermicro is a US-based, global provider of high-end computing solutions for data centers and cloud computing. Among those solutions, the company offers customizable, rack reference architectures for AI, centered on NVIDIA GPUs. However, the strategic direction does not translate to how Supermicro communicates its approach, as the company's website lists out-of-date ML frameworks such as Caffe and refers to Habana Gaudi even after Intel acquired Habana in December 2019.

Supermicro has strengths in component customization for training workloads. To become more competitive in AI infrastructure, the company must add more value than the ability to customize servers. With that said, existing or new Supermicro customers that know exactly what they want in AI infrastructure can engage Supermicro to build and deliver it. Supermicro declined to participate in the full Forrester Wave evaluation process; this assessment is based on publicly available information.

- **Exxact delivers customized performance but is geared toward smaller teams of experts.** Exxact is an infrastructure systems and component supplier that has found success with high-end use cases, including scientific computing and AI. The company offers low-cost starter AI workstations, a full range of servers, and fully blown "AI clusters." Exxact has leveraged its well-known position in scientific computing to offer those same research teams with a full range of AI infrastructure.

Exxact has strengths in components and systems for training workloads but lacks management tools and software layers that many enterprises demand for larger teams as well as efficient sharing of resources. The vendor is best fit for academic and enterprise research teams that can specify configurations for AI infrastructure. Exxact declined to participate in the full Forrester Wave evaluation process; this assessment is based on publicly available information.

Evaluation Overview

We evaluated vendors against 25 criteria, which we grouped into three high-level categories:

- **Current offering.** Each vendor's position on the vertical axis of the Forrester Wave graphic indicates the strength of its current offering. Key criteria for these solutions are architecture, training, inferencing, management, and deployment.

- **Strategy.** Placement on the horizontal axis indicates the strength of the vendors' strategies. We evaluated execution and innovation roadmaps, supporting products and services, partner ecosystem, and commercial model.
- **Market presence.** Represented by the size of the markers on the graphic, our market presence scores reflect each vendor's revenue, number of customers, and breadth of offering.

Vendor Inclusion Criteria

Forrester included 13 vendors in the assessment: Amazon Web Services, Dell, Exxact, Google, Hewlett Packard Enterprise, IBM, Inspur Information, Lambda, Microsoft, NVIDIA, Run:AI, Spell, and Supermicro. Each of these vendors has:

- **A comprehensive, differentiated AI infrastructure solution.** Evaluated vendors must offer a system and/or cloud service that provides capabilities that minimize the latency of data preparation, model training, and/or model inferencing.
- **Marketed its solutions as AI infrastructure.** Evaluated vendors must offer solutions that are specifically marketed to target enterprise customers specifically shopping for AI infrastructure, as opposed to customers shopping for more generic infrastructure. For example, AI workloads can run on any generic infrastructure such as services or virtual machines, but the vendor has not productized, designed, and/or marketed the infrastructure specifically targeting AI workloads.
- **Install base and revenue requirements.** The vendor must have at least 10 paying, named enterprise customers using the version of the enterprise AI infrastructure solution that we evaluated. The vendor must also have provided Forrester with three customer references that were willing to fill out a confidential survey. Included vendors must also have proven revenue generated by customer adoption of the vendor's AI infrastructure solution.
- **Sparked client inquiries and/or has technologies that put it on Forrester's radar.** Forrester clients often discuss the vendors and products through inquiries; alternatively, the vendor may, in Forrester's judgment, warrant inclusion or exclusion in this evaluation because of technology trends, market presence, or lack of client interest.

Supplemental Material

Online Resource

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.

For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

We publish all our Forrester Wave scores and weightings in an Excel file that provides detailed product evaluations and customizable rankings; download this tool by clicking the link at the beginning of this report on Forrester.com. We intend these scores and default weightings to serve only as a starting point and encourage readers to adapt the weightings to fit their individual needs.

The Forrester Wave Methodology

A Forrester Wave is a guide for buyers considering their purchasing options in a technology marketplace. To offer an equitable process for all participants, Forrester follows [The Forrester Wave™ Methodology Guide](#) to evaluate participating vendors.

In our review, we conduct primary research to develop a list of vendors to consider for the evaluation. From that initial pool of vendors, we narrow our final list based on the inclusion criteria. We then gather details of product and strategy through a detailed questionnaire, demos/briefings, and customer reference surveys/interviews. We use those inputs, along with the analyst's experience and expertise in the marketplace, to score vendors, using a relative rating system that compares each vendor against the others in the evaluation.

We include the Forrester Wave publishing date (quarter and year) clearly in the title of each Forrester Wave report. We evaluated the vendors participating in this Forrester Wave using materials they provided to us by September 9, 2021, and did not allow additional information after that point. We encourage readers to evaluate how the market and vendor offerings change over time.

In accordance with [The Forrester Wave™ And New Wave™ Vendor Review Policy](#), Forrester asks vendors to review our findings prior to publishing to check for accuracy. Vendors marked as nonparticipating vendors in the Forrester Wave graphic met our defined inclusion criteria but declined to participate in or contributed only partially to the evaluation. We score these vendors in accordance with [The Forrester Wave™ And The Forrester New Wave™ Nonparticipating And Incomplete Participation Vendor Policy](#) and publish their positioning along with those of the participating vendors.

Integrity Policy

We conduct all our research, including Forrester Wave evaluations, in accordance with the [Integrity Policy](#) posted on our website.

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.
For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.

FORRESTER

We help business and technology leaders use customer obsession to accelerate growth.

FORRESTER.COM

Obsessed With Customer Obsession

At Forrester, customer obsession is at the core of everything we do. We're on your side and by your side to help you become more customer obsessed.

Research

Accelerate your impact on the market with a proven path to growth.

- Customer and market dynamics
- Curated tools and frameworks
- Objective advice
- Hands-on guidance

[Learn more.](#)

Consulting

Implement modern strategies that align and empower teams.

- In-depth strategic projects
- Webinars, speeches, and workshops
- Custom content

[Learn more.](#)

Events

Develop fresh perspectives, draw inspiration from leaders, and network with peers.

- Thought leadership, frameworks, and models
- One-on-ones with peers and analysts
- In-person and virtual experiences

[Learn more.](#)

FOLLOW FORRESTER



Contact Us

Contact Forrester at www.forrester.com/contactus. For information on hard-copy or electronic reprints, please contact your Account Team or reprints@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research, Inc., 60 Acorn Park Drive, Cambridge, MA 02140 USA
Tel: +1 617-613-6000 | Fax: +1 617-613-5000 | forrester.com

Not Licensed For Distribution.

© 2022 Forrester Research, Inc. All trademarks are property of their respective owners.
For more information, see the [Citation Policy](#), contact citations@forrester.com, or call +1 866-367-7378.