

The Blade Form Factor May Not be the Best Choice for Data Centers Going Forward

Article

Blade servers were popular 10 years ago offering a reduced data center footprint, central management and less cabling. As processors increased from 150W to 400W+ and memory and networking power increased, it became difficult for blades to meet the required power and cooling requirements. At the same time rack servers and multi-node servers increased their capabilities and flexibility to better support the increased workload requirements.



In this article we list 10 reasons why rack servers or multi-node servers are better positioned than blade servers in 2024 and beyond.

1. Density & scalability

Blades do not provide the density advantages vs. rack servers as they once did:

- Cisco 7U enclosure holds 8 nodes
- HPE 10U enclosure holds 12 nodes
- Dell 7U enclosure holds 8 nodes

Blades average 0.86U per node, slightly better than a 1U rack server, but multi-node servers provide 0.5U per server (that is, 4 nodes in 2U of rack space).

Another density issue with blade density occurs when a blade enclosure is not 100% full. If a customer has four nodes installed in the blade enclosure, it still takes up the full blade enclosure (i.e., 10U) height in the rack. Four 1U rack servers will only take up 4U in height.

With blades scalability a customer is required to purchase and install a new blade enclosure as soon as they exceed the capacity of the first enclosure. If a customer needs 10 servers and their blade enclosure supports 8 nodes, they need to purchase two 7U blade enclosures with the second enclosure only containing 2 nodes. With rack servers the customer purchases and installs 10 servers.

2. Power requirements

As processor, DDR5 DIMM and networking power requirements increase, the combined power requirements will exceed the chassis power supply capabilities.

Less granular power leads to more expensive power interconnects between the PSUs and the blades. There is less physical space for the bulk power connectors. That reduction in space leads to the use of more expensive power connectors that have higher power density.

Blade power is less distributed, compared to single node rack servers. In blades, all power flows through the midplane. Higher resistive power losses are incurred compared to individual rack servers which distribute the power more evenly. These losses result in the creation of wasted heat which must be removed from the rack and data center.

The limited motherboard real estate on a blade restricts the space for voltage regulators which usually leads to limitations of the maximum processor power that can be supported. Blade motherboards generally have less space to place components, compared to the equivalent rack systems. This makes it more of a challenge to place all the VRDs that should be right next to the CPU socket.

3. Cooling challenges

As processor, memory and networking power requirements have gone up and up, the blades form factor has struggled to cool the systems resulting in limiting or not supporting many configurations. When nodes/blades are only partially utilized in a blade infrastructure, most, if not all, of the fans would consume power and air flow where as in a rack or multi node server the fans only associated with the node will consume power and air flow. The cooling solution for each node in a multi-node server can be customizable whereas in the blade chassis for each compute node an impedance matched filler has to be designed in the unpopulated bays/slots to balance air flow.

4. Local storage & HCI incompatibility

Blades are very limited to the amount of local storage each node/server provides. Most blade servers only support 2-6x 2.5-inch storage drives. 1U rack servers typically support 10-12x 2.5-inch storage drives

When blades were first introduced, they were a natural fit for converged infrastructure with servers and networking in the chassis along with fiber channel adapters to connect to external storage. With the increasing popularity of hyperconverged infrastructure (HCI), the blade's lack of internal storage rendered blades suboptimal. 1U rack servers with support for 10-12x 2.5-inch drive bays and multiple networking slots is a much more natural fit for hyperconverged infrastructure.

The key distinguisher between hyperconverged infrastructure and a blade server architecture is that in hyperconverged systems, the storage is networked and then pooled to create a huge virtual SAN. New innovations such as software-defined infrastructure take this further, to the point that the storage pool and the networks connecting the appliances are virtualized and controlled automatically by orchestration software. This allows tenants of an HCI-based cloud to add and subtract to their configurations using scripts and policies, without central IT intervention. In short, HCI is far easier to manage than a typical blade infrastructure because your compute, storage and network can all be managed from one central control plane eliminating silos and bottlenecks.

5. Limited I/O & GPU support

Blades all have custom I/O designs. These custom designs limit the choice of vendors, slows time to market of new technology, and reduces solution agility as software requirements change. The Blades unique form factor limits the number of slots for Ethernet and Fiber Channel adapters. Typically, these adapters are only available in unique form factors from the blade vendor.

As more and more workloads require the use of GPUs for workloads such as AI inference or VDI these are typically not supported in blades due to their power draw or physical size that is unable to be supported in a blade.

6. Vendor lock-in & options

Blades have unique form factor network adapters, Fiber Channel and RAID cards that are required and are typically only available from the blade vendor. Once a customer chooses their blade enclosure, only that vendor's hardware can be supported in the blade chassis. With blades, a customer can only use the specific switches, network adapters, Fiber Channel adapters, and RAID adapters that are designed for that specific blade system.

Rack servers, on the other hand, offer more flexibility in vendor selection and interoperability. Rack server chassis and components are typically standardized and compatible with hardware from multiple vendors, allowing customers to choose the best-fit solutions for their specific requirements.

7. Blade cost & performance

Rack servers have a lower initial cost than blade servers due to their standalone nature and simpler infrastructure requirements. With rack servers, a customer needs to invest in individual server units and rack space without the additional expense of a blade enclosure. This makes rack servers more economical.

Blade servers typically involve higher upfront investment due to the need for a blade enclosure along with the blade servers themselves. Rack servers can be more efficiently cooled and therefore typically outperform blade servers in many performance benchmark tests.

8. Flexibility & customization

Rack servers offer better flexibility and customization with different configuration and models available to meet specific needs. Blade servers are fairly limited in models and configurations. Rack servers are also more portable allowing customers to relocate an independent server without needing to move an entire enclosure along with the installed blades.

9. Installation complexity & enclosure management

Initially, blade servers can be more complex to set up and deploy than rack servers. Installing a blade enclosure involves integrating multiple blade servers, power supplies, cooling modules, and networking components within the enclosure. Configuration and cabling tasks must be carefully executed to ensure proper connectivity and system functionality. Blades also require separate blade enclosure management tools and software that must be learned and maintained.

Rack servers, by contrast, are generally easier to install and deploy due to their standalone nature. Each rack server can be installed independently, requiring minimal integration and configuration.

10. Blade market and investment in blades

Vendors have reduced their investment in Blades. HPE, Dell, Cisco and Lenovo previously had new blade enclosures and 2-3 different blade nodes for each CPU generation. Now vendors have a single blade node for each generation running in an older blade enclosure. Lenovo has exited the Blade market.

Customers have come to realize all the facts stated above and therefore blade shipments from Cisco, Dell, HPE and Lenovo have declined every year since 2015.

- 88% decline in units from 2015 to 2023
- 43% decline in units from 2022 to 2023

Source: 2023Q4 IDC Server Tracker – March 7, 2024

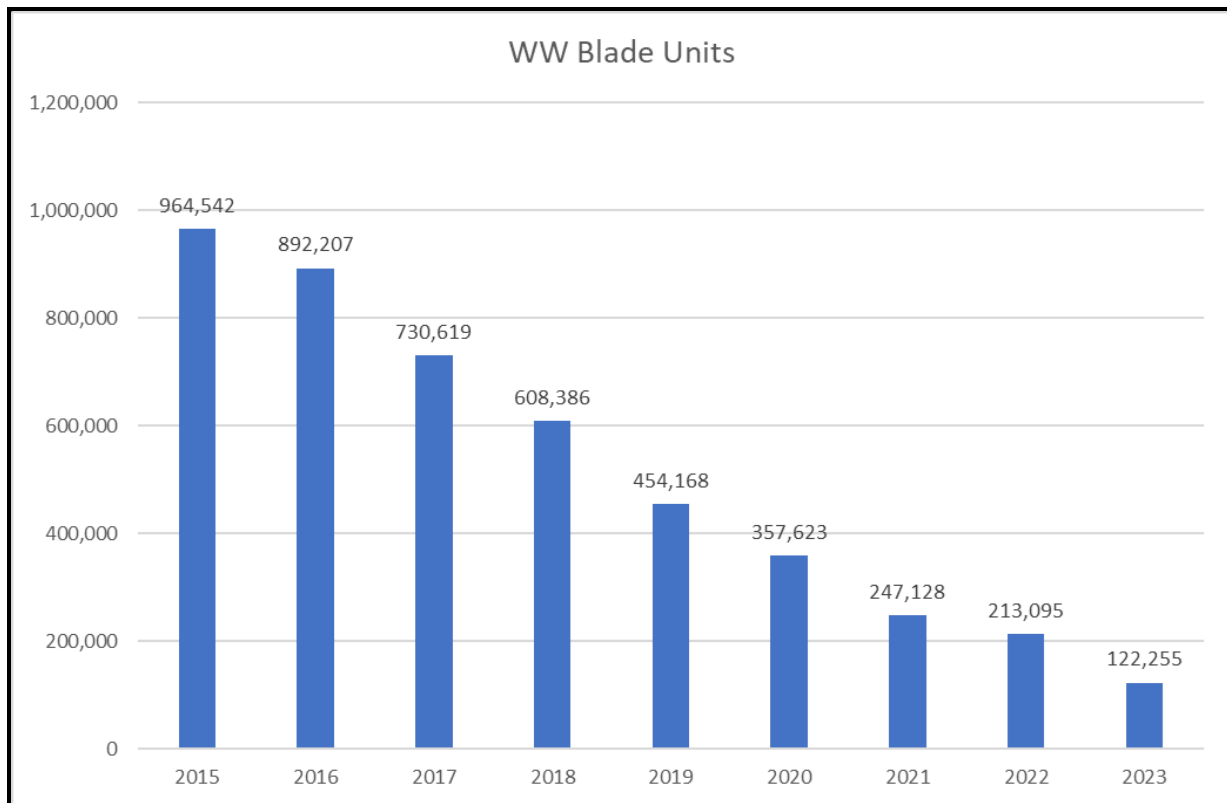


Figure 2. Blade shipments 2015 to 2023

The alternatives to blades

Now that you understand that blades may not be the best form factor, what Lenovo servers should you consider? Lenovo has several 1U rack servers, ThinkAgile and multi-node servers that can be added in addition or a replacement to blades.

Click the links to view the datasheet for each server:

- ThinkSystem 1U rack servers:
 - [ThinkSystem SR630 V3](#) – 1U2S Intel rack server
 - [ThinkSystem SR635 V3](#) – 1U1S AMD rack server
 - [ThinkSystem SR645 V3](#) – 1U2S AMD rack server
- ThinkSystem multi-node servers:
 - [ThinkSystem SD530 V3](#) – 1U2S half-width Intel multi-node server
 - [ThinkSystem SD535 V3](#) – 1U1S half-width AMD multi-node server
- ThinkAgile MX with Microsoft Azure Stack HCI
 - [ThinkAgile MX630 V3](#) - 1U2S Intel rack server
- ThinkAgile HX with Nutanix Cloud Platform
 - [ThinkAgile HX630 V3](#) (1U2S Intel)
 - [ThinkAgile HX645 V3](#) (1U2S AMD)
- ThinkAgile VX Integrated Systems for VMware HCI environments
 - [ThinkAgile VX630 V3](#) (1U2S Intel)
 - [ThinkAgile VX635 V3](#) (1U1S AMD)
 - [ThinkAgile VX645 V3](#) (1U2S AMD)

Blades were once a strong choice for reducing data center footprint, central management and less cabling but as servers grew and became more powerful, it became difficult for blades to meet the required power and cooling requirements.

We recommend you include Lenovo rack servers, multi-node servers and ThinkAgile servers in your reviews, request for proposals, and bids for your next data center expansion project.

Author

Randall Lundin is a Senior Product Manager in the Lenovo Infrastructure Solution Group. He is responsible for planning and managing ThinkSystem servers. Randall has also authored and contributed to numerous Lenovo Press publications on ThinkSystem products.

Related product families

Product families related to this document are the following:

- [1-Socket Rack Servers](#)
- [2-Socket Rack Servers](#)
- [Multi-Node Servers](#)
- [ThinkAgile HX Series for Nutanix](#)
- [ThinkAgile MX Series for Microsoft Azure Stack HCI](#)
- [ThinkAgile VX Series for VMware](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1982, was created or updated on June 28, 2024.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1982>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1982>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkAgile®

ThinkSystem®

The following terms are trademarks of other companies:

AMD is a trademark of Advanced Micro Devices, Inc.

Intel® is a trademark of Intel Corporation or its subsidiaries.

Microsoft® and Azure® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.