White Paper

Evaluating Enterprise Hypervisors for Existing Workloads and Future Modernization

Sponsored by: VMware by Broadcom

Gary Chen
September 2025

THE CHANGING VIRTUALIZATION LANDSCAPE

Evaluating hypervisors and server virtualization is a complex task today and oddly enough, something new to many organizations even though they have been virtualizing for decades. Today's environment and complex needs are far different than when virtualization first exploded onto the scene many years ago. Many enterprises have simply built and expanded virtualization over time and upgraded the infrastructure year over year without doing a full reevaluation of their platform strategy and cost considerations. But today, in light of recent vendor/market changes, rapidly shifting modernization demands, and the coming AI revolution, enterprises are revisiting this technology in a way they have not considered before, with a fresh, bottoms-up look. Competitive market dynamics have also spun up rapidly in recent years, as several vendors have increased investments into virtualization capabilities and products, presenting users with more options than ever to consider. This paper will explore the key aspects when evaluating hypervisors that come up most often in customer engagements with IDC, including technical merit, cost, vendor viability, and modernization.

VIRTUALIZATION EVALUATION CRITERIAS

Technical Merit – A high bar today to meet today

Hypervisors run at the lowest layer of the stack and in a very privileged position as the first layer above the hardware. As the foundation, enterprise hypervisors must be rock solid. Additionally, the hypervisor market has existed for several decades and enterprise hypervisors have evolved considerably over the years. The bar is now very

high in the enterprise for hypervisor requirements, one that is feature rich and extremely stable. Key technical considerations:

- **Scale** - Physical and VM server sizes are constantly growing, as well as total compute footprints. This will become a larger factor as we enter the AI era, with more compute and data intensive workloads. Scaling issues need to be considered on several levels:
  - Within a node – What the hypervisor supports on the physical host in terms of CPU sockets, CPU cores, and RAM. Likewise, for an individual VM in terms of vCPUs and RAM. There is also a maximum number of VMs supported per host, though this limit is rarely reached by customers as physical resources usually run out before reaching this.
  - Across nodes – These factors would be nodes per cluster, clusters per management zone, and management across multiple datacenters. As customers move from basic virtualization to private cloud, scale and management for large footprints becomes a limiting factor. This is also highly relevant given the changing nature of modern workloads such as AI, which often requires large scale distributed infrastructure.
  - Theoretical vs tested limits - While these published maximums are a useful initial figure, it is important to keep in mind that some of these numbers are theoretical limits, not what has been tested and proven. While many customers will not exceed these limits, it is important to know when performance degradations can possibly happen and scaling issues as the footprint grows.
- **Performance** - Virtualization performance is a complicated topic with many aspects. Certainly, much of the performance discussion is in regard to hypervisor overhead with CPU, memory, and I/O performance. But consider these aspects of performance as well:
  - Support for the latest accelerators – GPUs and DPUs have become much more popular and there may be more types in the future, especially for AI workloads. Support for these accelerators are essential for many modern workloads as they cannot function without them and application performance is very dependent on how efficiently they are allocated. Additionally, consider if the platform has additional efficiency or sharing benefits based on virtualization technologies for these resources. This can greatly increase utilization and expand availability to more workloads, a key cost consideration as many of the newest accelerators are expensive and scarce.
  - Performance at load – Performance between various hypervisors can be more differentiated when the utilization is high. This is a key metric to

determine workload behaviors when the load increases. Ideally, enterprises want to achieve high utilization while maintaining SLAs. A key aspect of confidently achieving higher utilization is the robustness of the monitoring, remediation, migration, and orchestration capabilities when nodes become overloaded.

    - VM mobility underpins many key resiliency and scheduling features of modern datacenters and is key for overall cloud performance. When it comes to the speed of live migrations, it is important not only to consider the speed of transfer but also performance with multiple simultaneous transfers. There may also be additional features that can expand migration options such as migration between different CPU types. Not to be forgotten is also live storage migration, which can be just as important as the VM itself. While live migrations are a standard feature of virtualization today, technical details of the robustness of these features is still key to compare.

- **Security** – In an increasingly dangerous digital environment, security is one area that most enterprises do not underinvest in anymore, instead being one of the top IT requirements. Below are several key considerations, but the most important one is simply the vendor's commitment to security. This can sometimes be difficult to determine, but some key factors to weigh would be:
    - The development culture around continuously improving and prioritizing secure architecture design and producing secure code
    - The response infrastructure in place to respond to newly discovered flaws and ongoing incidents
    - Internal efforts to harden and hack their own code as well as external efforts to work with the larger security research community, for example with bug bounty programs.

Looking at past responses can be a good indicator of a vendor's security culture. Beyond culture, specific functionality to consider for virtualization would include:

- **Compliance certifications** – A broad list of compliance certifications is a good sign that the vendor has a healthy business across diverse industry segments. While different customers will only be impacted by a small set of specific requirements, it can be a broader indication of the vendor’s overall security focus
- **Ecosystem** – Complete security requires that virtualization vendors work with a broad array of specialist security vendors. Built-in security integration capabilities and a strong partner program is essential to building out a large interoperable ecosystem. This also extends to server hardware and the silicon as many virtualization security enhancements are built in at that level to protect against CPU and memory vulnerabilities. It requires tight coordination between hardware and software vendors to enable these features in the hypervisor.

- **Securing the management plane** – The management plane is an attractive target today because it is a target with high reward and access to larger parts of the infrastructure. There does not have to be a software flaw for the management plane to be at risk, most issues arise from accidental or unknowingly insecure misconfigurations. The age-old problems of configuration checks, drift, and remediation are still difficult operational issues today and platforms that develop more robust monitoring and automation technologies are key to closing these gaps.

Cost – Not as simple as it appears

Cost is an important aspect of any infrastructure, but cost can be a complex model that goes well beyond hypervisor licensing cost. Most enterprises are quite familiar by now with the differences between perpetual and subscription as well as the myriad of metrics used in enterprise software licensing (sockets, cores, RAM, etc.). OS and application licensing differences in virtual environments are also well understood by now. However, true cost and value calculations for virtualization must extend beyond raw licensing cost comparisons. Consider the following in cost models:

- **Hardware** - Since hypervisors are the closest software to the bare metal, hardware refresh cycles and hardware configurations tie closely into total cost figures. Hardware can have a massive impact on VM performance and density, so it is important to factor in the alignment of refresh cycles and the impact of next generation configurations. One key factor is the scope of the virtualization provider's certified hardware list, which will dictate the hardware options to choose from. Hypervisors also can vary in how well and how soon they support the newest silicon features, which can impact performance, density, and security. Co-engineering and technical partnerships with hardware vendors are a key factor in enabling better and faster support. These factors also apply to hardware accelerators, which are becoming common, especially to support workloads like AI. DPU's for example can affect VM density since they offload storage and networking tasks to the accelerator, freeing the main system for more VMs. GPUs are evolving extremely fast and demand outstrips current supply, making them expensive and scarce resources that must be fully utilized to achieve payback. Exposing and sharing GPUs is still an area ripe for further optimization.
- **Efficiency** - Hypervisors can vary in how densely they can pack servers and how they maintain performance at high loads. Being software-defined, there are many clever hypervisor "tricks" around memory optimizations, leveraging flash to extend memory, storage performance optimizations, and compute optimizations that can pack more VMs per set of hardware. These cannot simply be evaluated as check box items, as the robustness of these features require

deep technical knowledge to test and evaluate true capabilities and limitations. The other factor that affects efficiency is the robustness of the monitoring and remediation systems (leveraging orchestration and automation capabilities), which allows customers to run confidently at consistently higher utilization levels. These systems are essential to maintaining performance when workloads change and even more essential in cloud deployments that span across large computing estates.

- **Operational costs** – As enterprises grow their compute footprints to meet demand and increase the complexity of their infrastructure as they move toward cloud operating models, the cost of operations is just as important, if not more, as VM density and performance. Moving to a cloud model ties up large swaths of hardware and software into a coordinated system, but this also has the implication that one problem in one area could cascade into the rest of the cloud. Operating at scale creates many challenges:
  - The lifecycle management of nodes and the control plane of modern clouds is challenging problem at scale. One key challenge is modern infrastructures are highly complex and software defined, spanning compute, storage, networking, monitoring, observability, application performance management, and many more areas. Often these systems are interdependent and controlling the constant patching and upgrade versions creates high operational overhead. Impeding the speed of upgrades due to complexity can also affect the speed of innovation and security. Beyond lifecycle challenges, the day-to-day operations of highly virtualized infrastructures and private clouds are complex as operations span multiple infrastructure disciplines and management planes. Applications are becoming more complex as well, very different from the traditional monolithic architectures. Customer need to have a vision of the operating model and the platform services they aspire to offer, where hypervisors and server virtualization are just one small piece of that platform and its overall cost.
- **Silos are a fact of life in IT and there are still many silos of varying degrees across the infrastructure.** These silo walls inject resistance when trying to cross them; some are minor, while some can be very hard boundaries. Breaking down technical and organizational silos are a key aspect of moving to a modern cloud operating model and platform engineering approaches. Skills gaps in the organization are also a challenge as modern cloud operating models and newer computing technologies like containers are becoming the focus for new applications like AI.
- **Downtime and security breaches** – Being a mature technology that is widely deployed as the foundation for datacenter infrastructure comes with a different set of requirements than an emerging technology that is experimental, cutting

edge, and small scale. Enterprises have a lot riding on this infrastructure and consequently more to lose when downtime or security breaches happen. This is about more than just product features for resiliency, uptime, or security, but rather a collective industry operating maturity and experience. Decades of operational history, knowledge, scripts/playbooks, and battle tested responses are an asset in itself with value. As anyone experienced with operationalizing new platforms in IT knows, that is a long journey that is paved with lessons often learned the hard way.

Maturity and long-term viability

Most IDC conversations with customers around compute and virtualization do not start with technical features, but rather with evaluating long term vendor and product viability. The bar is different for virtualization as it is a foundational technology and the convergence point for many other technologies and subsystems. Key discussion and proof points include:

- **Real world deployment proof points**: The age of the code base is only one aspect of maturity. Customers also want to know how many customers have run in similar environments in terms of scale, complexity, and workload types. Enterprises tend to be more conservative with foundational infrastructure and want proven success stories around operations, resilience, availability, and security.
- **Enterprise core platforms have a very long lifecycle and migrations and deprecations can be very painful and expensive.** Key vendor questions to determine are:
  - Financial stability and health of the vendor. Will the vendor still be in business in 10+ years?
  - The level of investment the vendor has in the product. Is it a side project or a main priority for the company? Will the product exist in 10 years, or will the vendor exit the business? Is the company investing significant amounts into R&D?
  - Beyond just the cost of support, can the vendor offer support appropriate for my size, scale, and service level? Does the vendor have the resources and investment necessary to provide enterprise class support at the level I need? Length of support cycle and product lifecycles are also key in many industries.
  - Does the vendor have a strong ecosystem? A strong ecosystem is a major factor in the viability and value of a platform. There are a huge variety of add-ons, extensions, and attached technologies to virtualization, so universal support gives customers choice and a wide range of targeted solutions.

Strong partnerships also bring co-engineering and co-innovation advances that are important as the platform constantly evolves and modernizes.

- Hypervisors no longer are a standalone solution. Today virtualization is one technology in a larger platform strategy that can extend into private cloud, hybrid cloud, public cloud, edge computing, and containers. Virtualization vendors need to have a vision and strategy for the platform of the future.

Modernization, AI, and future needs

While virtualization has powered datacenters for decades now, future workloads are just as important a consideration as existing ones. Modernization is a big investment area for enterprises today, across many axes. Some of the key ones related to infrastructure are:

- **The shift to cloud as an operating model and architecture for on-premises enterprise datacenters.** Private cloud has been deployed by some organizations with varying scopes and successes, but achieving a consistent, organization-wide, universal platform has been encumbered by complexity and cost at scale. To make cloud models more accessible and bring down costs, many platforms are aiming to provide higher integration across the various subcomponents of cloud such as compute, storage, networking, and multiple management areas. To make operations less complex, the orchestration and automation control plane is becoming more sophisticated and intelligent, increasingly leveraging more machine learning and AI. Cloud models go beyond building infrastructure of course. Enterprises also need to devote investment into making clouds more consumable, with innovative services that solve users' problems. It is not just about offering self service and exposing a bunch of APIs, it is a top to bottom remaking of both IT architecture and the IT service model that has to be driven by people and culture as much as technology.
- **Containers and Kubernetes are the default application platform for new cloud native, microservices workloads of which AI is a prime example of.** Container and VM technologies are highly interlinked today. Modern Kubernetes needs an agile cloud like infrastructure underneath it and cloud architectures built on virtualization are the way to provide that. Kubernetes does not absolve the user from building an agile cloud-like infrastructure for containers to consume. Most containers today run in VMs (like in the public cloud) and that will likely continue for the foreseeable future, though this is not to say that bare metal containers do not have a place for some deployment types. Coordination between containers and the underlying VMs is critical to making container deployment and management easier, both for deploying clusters/control planes and for managing scalable worker nodes. For applications, it is extremely common today to see applications made up of VM and container components.

For example, the application database is deployed in a VM while the front end is in containers. Running separate VM and container silos makes deploying and managing these composite applications more difficult, as well as operations like monitoring and observability. A recent trend in IT is the convergence of these platforms to create a more unified experience, as enterprises will be managing both compute types for the foreseeable future and there will be many interdependencies between them.

- **AI is a special type of workload that is impacting all of IT today as it grows into an application class of its own and gets grafted onto nearly everything in current existence.** From an infrastructure view, it is creating increased demands not just for compute, but challenges around big data and high volumes of internal east-west traffic. The role of compute accelerators for AI cannot be underestimated, and the virtualization, management, and efficiency of these devices will grow into a major infrastructure function as AI grows. From a service model view, AI will place increased demands for more agile and scalable infrastructure as well as a modern container layer for the application platform. AI workloads are generally built with highly modern container architectures, but the needs go beyond just container management. Container platforms are extending to become highly AI centric, offering platform capabilities such as model serving, templated AI frameworks and patterns, and inference infrastructure. The technology for AI is evolving so fast, users need to seek vendors and products that are committed to R&D, the capability to adapt quickly to the needs of the market, and the market influence to work with the larger AI ecosystem of partners.

As can be seen from the above discussion, there are many factors to consider when choosing a hypervisor. IDC has seen that it is generally a very long process and rightly so given the critical role hypervisors play in today's environments. Evaluations and testing will be lengthy, especially for complex and large scale environments. Keep in mind that the virtualization market is undergoing rapid change, so vendors, products, and workloads such as AI will be evolving over the next several years. What is true today may be different a couple years from now. Cost models have many subtleties that are key to accurate total cost figures and returns on investment, and cost assumptions can also change rapidly due to market conditions and vendor competition. The workloads and world of compute are in a very different state than when virtualization was sweeping through datacenters and this paper only touches on the many factors that must be considered as enterprises build for the future. What is certain is that the virtualization market is becoming much more dynamic than in recent years, which lends a level of uncertainty in whatever decision an organization makes today.

VMWARE PROFILE

VMware revolutionized the datacenter many decades ago when it developed a hypervisor that enabled server consolidation and introduced modern software-defined operations. Today, virtualization is the standard server technology for servers and VMware has grown into a broad and robust infrastructure platform. vSphere is the core virtualization technology for VMware and it underpins all of VMware's offerings, from essential virtualization packages to the flagship VMware Cloud Foundation (VCF). VCF is a comprehensive and cohesive enterprise platform that integrates vSphere along with multiple other technologies (containers, compute, storage, networking, security, and management) into a modern cloud that is easier and less costly to deploy and operate.

While the hypervisor has evolved greatly over the years, VMware shows that innovation is still happening in the hypervisor. vSphere 9 is the latest version of VMware's core hypervisor and brings several key improvements:

Modern workloads

- vSphere 9 enhances the Supervisor cluster with a more flexible configuration that allows users to deploy a minimal initial configuration and add services as desired later. Previously the Supervisor had to be specified during deployment and was largely immutable afterwards. The Supervisor can now also be independently upgraded from vSphere. The Supervisor cluster is a special privileged Kubernetes cluster that hosts vSphere's modern services, one of which is the vSphere Kubernetes Service (VKS). Other key services include the VM service, a modern way to run VMs with Kubernetes and VCF Automation.
- VKS is an embedded Kubernetes runtime within vSphere that enables unified management of VMs and containers. As it is an integral part of vSphere, it also inherits and shares vSphere's subsystems, such as storage, networking and security. As part of the Supervisor, VKS can be upgraded independently from vSphere/vCenter, enabling faster access to newer Kubernetes versions. VKS also allows multiple Kubernetes versions to run simultaneously, allowing for more flexible upgrades. This enables much easier and efficient deployment and simplifies the Kubernetes infrastructure.
- VCF also now includes key container components such as multi-cluster management (formerly Tanzu Mission Control) and the Istio Service Mesh (formerly Tanzu Service Mesh) that help build out a fuller container platform around VKS.

Operations

- NVMe memory tiering allows expanding RAM with cheaper NVMe flash storage. This creates a larger memory pool that can provide higher VM density. It can be very efficient for VDI workloads and offer general workloads higher density in exchange for an acceptable performance hit. This can also enable software licensing savings in certain scenarios, for customers who have configurations that are memory constrained but have available cores. Customers can configure the management of memory tiering for workloads and key functions like DRS and vMotion are memory tiering aware.
- Larger scale limits - Up to 960 cores per socket on the host and 960 vCPUs and 16TB of memory for VMs, which can be needed for workloads such as large in memory databases.
- 6x faster vMotion for GPU workloads, using parallel TCP connections and a 100Gbps network. It also leverages offload accelerators that can securely transmit the data without overhead. The speed of vMotion is critical for AI workloads which have special vMotion challenges due to the use of GPU accelerators.
- Provision and manage VPC networks (with NSX) in vCenter. This reduces the complexity of network configuration and enables cloud users easier access to advanced networking and security. Previously, VPCs were only manageable with NSX tools with a complex integration process with vCenter.
- Live patching for ESXi, allowing security patches to be applied with no downtime. Hosts do not need to be taken offline or evacuated.
- Faster vCenter upgrades in less than 5 minutes, enabling higher uptimes and faster scheduling of important security patches. There is also a rollback capability in case any updates cause problems.
- A unified VCF SDK to deliver a consistent automation experience across all VCF components.

Security

- A more secure platform with TLS 1.3 that is supported in vSphere 9 and across the entire VCF stack. Setup time is reduced by half with faster and more responsive HTTPS connections.
- FIPS 140-2 compliance by default.
- Easier SSL certificate management with automatic, fleet-wide renewal across all VCF components.

Challenges/Opportunities

VMware is the market share leader for virtualization, owning a broad install base and technology leadership. It is the standard that many customers still measure competitors by. However, the landscape is rapidly changing with public cloud, containers, and AI, all of which are changing the nature of infrastructure and applications. While virtualization technology is still a key part of modern architectures, the evolution of the hypervisor's role is something that can bring disruption to incumbents like VMware. However, VMware has much to gain if it can deliver its vision of easier and cheaper private cloud and the convergence of VMs and containers.

There are many factors that can push customers away from ideal technical architectures and away from logical, data driven decision making. Often, enterprises are twisted by cost constraints, soured vendor relationships, internal organizational structures, and inertial resistance to change. Success in the market is as much about product and technology-led success as well as the human element of the enterprise IT business.

CONCLUSION

Evaluating enterprise hypervisors is increasingly complex due to evolving workloads, modernization demands, and the rise of AI. Many organizations are revisiting their virtualization strategies, considering technical merit, cost, vendor viability, and future needs. Hypervisors must meet high technical standards as users today expect a wide range of features and ultra high reliability. Security is also a much bigger focus than when virtualization started, with everything today a target, including hypervisors and its management plane.

Cost evaluation extends beyond licensing and needs to factor in hardware refresh cycles, VM density, operational efficiency, and downtime risks. Evolving beyond basic virtualization and into modern private cloud infrastructures demands robust lifecycle management, the breaking down of silos (technical and organizational), service culture change, and addressing skills gaps. Enterprises must also consider vendor stability, product maturity, and ecosystem strength, as hypervisors are foundational technologies with long lifecycles and are highly integrated into larger infrastructure and application platforms.

Container integration is also a key factor as modern and AI-centric workloads are already being built on Kubernetes and containers as the standard. Containers are not a standalone system; they have a dependence on the virtual and physical infrastructure underneath, so the coordination, integration, and convergence between containers and VMs are critical factors as organizations consider a more unified platform.

The virtualization market is undergoing rapid change, both technically in terms of modernization and workloads, as well as the competitive market dynamics between vendors. Many enterprises are conducting more thorough evaluations of virtualizations than they have in a long time, as they seek to align hypervisor and platform choices with current and future needs

MESSAGE FROM SPONSOR

Virtualization remains the backbone of modern IT — delivering unmatched performance, security, and efficiency across enterprise workloads. VMware vSphere stands as the industry's leading virtualization platform, trusted by organizations worldwide to power mission-critical applications, simplify infrastructure management, and maximize resource utilization. As the foundation of VMware Cloud Foundation, vSphere brings together decades of innovation to support everything from traditional applications to emerging AI workloads. Modernize your infrastructure with the proven power of VMware vSphere — the virtualization platform built for what's next. For more information about vSphere, please visit this webpage.

ABOUT IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology, IT benchmarking and sourcing, and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives. Founded in 1964, IDC is a wholly owned subsidiary of International Data Group (IDG, Inc.).

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com