

The Alternative Hypervisor Handbook

Exploring your options, from Nutanix and Proxmox to Red Hat and beyond



Table of contents

Taking a new path forward

Many organizations have made the decision to reduce or eliminate VMware from their IT infrastructure and incorporate one or more other hypervisor platforms. There are many choices available, from long-standing commercial offerings to open source software. But which best fits your organization’s needs?

This guide explains some of the more popular alternatives, the advantages and potential drawbacks, and how Everpure™ can make your transition easier and more successful. Note that this guide focuses on data center–based solutions and does not address public cloud offerings.

Hypervisor considerations	3
The hypervisor landscape	4
Nutanix Cloud Platform	5
Microsoft Azure Local/Hyper-V	7
Red Hat OpenShift Virtualization (KubeVirt)	9
OpenStack and KVM	11
Proxmox VE	13
Summing it up	15

Hypervisor considerations

VMware is regarded as a best-in-class hypervisor with one of the richest feature sets, especially for large enterprises. As organizations plan to reduce or eliminate VMware, it is important to ensure any new hypervisor meets organizational requirements.

Here are some specific areas to consider:

Concern area	Details
System scalability	Horizontal vs. vertical growth: Evaluate the maximum number of virtual machines (VMs) per host and hosts per cluster. Consider the stability of the management plane—how does the interface perform when managing hundreds vs. 1,000+ instances across multiple geographic sites?
High availability	Automated recovery: Assess the logic used for automated restarts during hardware failure. Does the system support live migration (moving a running VM between hosts without downtime) and dynamic load balancing (moving VMs automatically to optimize host resources)?
Disaster recovery	Orchestration and recovery point objective/recovery time objective (RPO/RTO): Look for the ability to automate the failover of entire data centers. Important factors include native replication capabilities, support for immutable snapshots to combat ransomware, and the ability to test disaster recovery plans without disrupting production.
Storage integration	Performance and protocols: Verify support for standard storage architectures such as storage area network (SAN) and network-attached storage. The hypervisor must handle high-throughput I/O requirements for modern workloads such as databases, AI, and large container deployments.
Migration path	Ingress and egress complexity: Review the availability of virtual-to-virtual tools. Consider the risk of cold migrations (requiring downtime) versus warm migrations. The complexity of converting disk formats and reconfiguring networking/drivers is a primary cause of project labor costs.
Ecosystem and tooling	Third-party compatibility: The hypervisor does not exist in a vacuum. It must provide robust APIs for existing backup suites, security agents (antivirus/endpoint detection and response), and monitoring tools. A lack of native support from major vendors can force an organization to replace its entire secondary software stack.
Licensing and TCO	Cost structure: Beyond the initial purchase price, analyze the ongoing total cost of ownership (TCO). This includes per-core vs. per-socket licensing, mandatory support contracts, and the hidden tax of required management software or gateway appliances.
Operational skillset	Talent availability: Evaluate the learning curve for the existing infrastructure team. A shift in the underlying hypervisor technology often requires a transition from one management philosophy to another (for example, GUI-driven to a command line–driven or infrastructure-as-code-driven philosophy).

The hypervisor landscape

In 2026, the market for VMware alternatives has matured into distinct categories based on organizational needs, from turnkey enterprise clouds to lightweight open source platforms. While this list is not inclusive, here are some of the most popular alternatives.

1 Nutanix Cloud Platform

The most enterprise-ready direct replacement for those running full-stack VMware, Nutanix Cloud Platform supports both a hyperconverged infrastructure (HCI) model and external storage. The Nutanix hypervisor (AHV) is bundled for free with its storage and management software. It is favored for its one-click simplicity and mature migration tools (Nutanix Move).

2 Microsoft Azure Local/Hyper-V

The primary choice for Windows-centric organizations, Microsoft Azure Local/Hyper-V leverages the familiar Hyper-V hypervisor but connects natively to the Azure portal for management. It is cost-effective for teams already paying for Windows Server Datacenter licenses, as the hypervisor rights are essentially included. Azure Local includes external storage support.

3 Red Hat OpenShift Virtualization (KubeVirt)

Red Hat OpenShift Virtualization (KubeVirt) is the leading choice for cloud-native modernization. Instead of a traditional hypervisor, it uses Kubernetes to manage VMs. This allows an organization to run legacy VMs and modern containers on the same platform and hardware, managed by the same team. SUSE is also an available option.

4 OpenStack

OpenStack is best understood not as a hypervisor itself, but as a cloud operating system that sits on top of hypervisors—typically Kernel-based Virtual Machine (KVM)—to turn them into an automated, self-service cloud. While a hypervisor like VMware ESXi or KVM manages individual VMs on a single physical host, OpenStack orchestrates thousands of those hosts into a unified pool of compute, storage, and networking.

5 Proxmox Virtual Environment (VE)

The top open source contender for mid-market and small and medium-sized businesses (SMBs), Proxmox VE combines KVM for VMs and Linux Containers (LXC) for lightweight containers in a single, clean web interface. It is widely praised for its built-in backup server and clustering features without mandatory per-core licensing fees.

The following sections cover these hypervisors in detail.



Additional hypervisor options

The full hypervisor ecosystem has even more available options. Here are some other products not covered in this guide that may be worth considering for specific use cases.

Oracle Linux KVM

Oracle Linux KVM is a safe harbor specifically for organizations with heavy Oracle Database or enterprise resource planning footprints. It is an enterprise-hardened version of KVM that is free to use (paid only for support). It is often chosen because it is the only non-VMware hypervisor that Oracle officially recognizes for hard partitioning license pinning.

Vates XCP-ng

Vates XCP-ng is an enterprise-grade, fully open source successor to Citrix XenServer. Managed via the Xen Orchestra console, it is popular with European enterprises and those seeking a sovereign cloud because it avoids vendor lock-in entirely while offering high-performance backup and migration features.

Scale Computing SC//Platform

The go-to solution for edge computing and distributed environments (like retail stores or branch offices), Scale Computing SC//Platform is designed to be self-healing and extremely lightweight, requiring significantly less RAM and CPU overhead to run the hypervisor than VMware or Nutanix.

1

Nutanix Cloud Platform

The enterprise-ready solution

The Nutanix Cloud Platform attempts to make the hypervisor “invisible.” The closest experience to VMware vSphere, it is the choice for organizations that want to move off VMware without retraining their entire staff or rearchitecting their applications.

Nutanix works in both HCI mode and with certified external storage partners. While not all features are available with external storage, this deployment model has its own advantages (see the following page).

Best for: Large enterprises that want to keep their existing VM-centric workflows but eliminate the VMware licensing uncertainty.



Concern area	Nutanix Cloud Platform
System scalability	Nutanix scales to thousands of VMs. In HCI mode, each node adds a fixed amount of CPU, RAM, and storage simultaneously. You cannot easily add just storage or just compute nodes. In external storage mode, you can scale compute (Nutanix nodes) and storage independently. This eliminates the “HCI tax,” where you were previously forced to buy more CPU/RAM just to get more storage capacity.
High availability	AHV does not just restart VMs; the underlying Acropolis OS (AOS) actively monitors storage health. If a drive fails, the system automatically redistributes data across the cluster to return to a healthy state without human intervention. In external storage mode, if a host fails, AHV restarts the VM on another node, which immediately reconnects to the same external volume. It mimics the traditional VMware/SAN high availability behavior but with NVMe speeds.
Disaster recovery	AHV includes asynchronous replication (RPO of 1 hour or more), near-sync replication (RPO <1 minute), and synchronous replication (RPO of 0). Options are limited in external storage mode but are continuing to develop.
Storage integration	In HCI mode, AHV keeps a VM’s data on the local disks of the host it is running on. This provides extremely low latency but requires a mindset shift for admins used to managing external storage arrays. With external storage, Nutanix Prism can speak directly to external arrays. You can provision a volume directly from the Nutanix console, treating the external SAN as a native Nutanix storage container.
Migration path	Nutanix Move is one of the industry’s most mature migration tools. It allows for warm migrations where data is seeded in the background, and the final cutover (VMware to AHV) typically takes less than two minutes of downtime.
Ecosystem and tooling	As of 2026, there is nearly universal support for Nutanix from major vendors like Veeam, Rubrik, Commvault, and HYCU. However, some legacy niche apps may still only be certified for VMware, requiring a careful audit of your software stack.
Licensing and TCO	AHV is license-free, meaning the cost is bundled into the Nutanix Cloud Platform. You stop paying per socket for the hypervisor, but you must factor in the subscription cost for the Nutanix Cloud Infrastructure (NCI) software layers.
Operational skillset	The management interface (Prism) is highly intuitive and HTML5-based. Most VMware admins can become proficient in Nutanix in less than a week, as the concepts of vMotion and DRS map directly to AHV live migration and ADS.



How Everpure can help

Everpure FlashArray™ is deeply integrated with the Nutanix Cloud Platform via a joint development effort between both companies, with a rich feature roadmap. The integration between Everpure and Nutanix products is engineered to minimize workflow disruption by bridging internal and external storage resources via the Nutanix compute cluster, which is connected to the FlashArray via NVMe over TCP. Once deployed, system administrators working through the Nutanix Prism management console can deploy new AHV-based VM vDisks to the datastore on the FlashArray system.

Management is simplified and familiar, as it’s very similar to the external storage model developed by VMware. Current FlashArray users can make use of their existing arrays to run Nutanix, even if they have existing VMware workloads.

The extreme performance and density of the FlashArray//XL™ and FlashArray//X™ family can deliver massive VM consolidation. Users moving off VMware on legacy storage can see significant savings on power, space, and cooling, as well as management overhead. Deploying external storage also greatly reduces the amount of compute nodes needed, lowering server costs in a time of rapid price increases and limited availability. A FlashArray can also support both VMware and Nutanix at the same time, simplifying migrations.

Learn More About Everpure and Nutanix

2

Microsoft Azure Local/Hyper-V

Familiar and easy

In 2026, Microsoft unified its hybrid infrastructure under a single banner: Azure Local. This brand now encompasses what was formerly known as Azure Stack HCI, Azure Kubernetes Service (AKS) for on-premises environments, and Azure IoT Edge. The shift to Azure Local means you are no longer just buying a hypervisor (Hyper-V); you are buying a slice of the Azure cloud that happens to live in your data center.

One of the biggest drivers for this platform is the ability to run Azure Virtual Desktop (AVD) in your own data center. This provides the Azure desktop experience but with the low latency of being on your local area network (LAN).

Best for: Organizations already committed to the Microsoft/Azure ecosystem who want integrated billing and hybrid-cloud management.



Concern area	Azure Local/Hyper-V
System scalability	Azure Local instances can now scale to 64 nodes (up from 16 in previous years). While individual clusters are smaller, they are managed as a single fleet via the Azure portal. However, you must account for the Arc resource bridge overhead required for cloud-based management.
High availability	A major 2026 update allows for local availability zones. You can split a single cluster across different racks or even different floors. If a top-of-rack switch or power feed fails, the system automatically shifts workloads to the surviving rack.
Disaster recovery	With Azure Site Recovery (ASR) integration, disaster recovery is a native toggle in the portal. You can fail over from your local hardware to a native Azure VM in the cloud. Site affinity ensures low-latency replication between your specific data center and the nearest Azure region.
Storage integration	While Storage Spaces Direct (S2D) is the standard HCI choice, Azure Local now officially supports external SANs from multiple storage vendors. This allows you to scale storage independently from your Hyper-V compute nodes, a critical need for high-density AI and database workloads.
Migration path	With Azure Migrate, Microsoft has optimized the migration path from VMware. The tool now supports agentless discovery, allowing you to move hundreds of VMware VMs to Azure Local with minimal manual configuration of virtual drivers.
Ecosystem and tooling	To get the best experience, you must use the Azure portal for day-to-day tasks. While Windows Admin Center (WAC) still exists for disconnected emergencies, the primary orchestration, patching, and security policy (Azure Policy) all run through the cloud.
Licensing and TCO	Azure Local is a subscription service. You pay a monthly per-core fee through your Azure bill. While there is no upfront software cost, you must factor in the requirement that the cluster connect to Azure at least once every 30 days to remain licensed.
Operational skillset	Your team needs to move from being Windows admins to Azure admins (a cloud ops mindset). It will manage VMs using ARM templates, Bicep, or Terraform, which is a significant departure from the traditional Hyper-V Manager GUI.



How Everpure can help

Everpure and Microsoft have partnered to integrate Everpure FlashArray with Azure Local, creating a high-performance, hybrid-cloud solution that transcends traditional HCI limitations. By supporting external block storage, this architecture enables the independent scaling of compute and capacity while delivering the submillisecond latency required for mission-critical applications.

This integration enables the use of existing FlashArray SANs with Azure Local via Fibre Channel connectivity, allowing for side-by-side deployment with S2D, where SAN volumes are used exclusively for user workloads like VMs and Kubernetes.

The exceptional storage density of FlashArray allows you to consolidate many VMs onto a single storage platform while also delivering savings on energy and space. The solution roadmap includes support for iSCSI and NVMe connectivity.

[Learn More About Everpure and Azure Local](#)

3

Red Hat OpenShift Virtualization (KubeVirt)

The future is now

Red Hat OpenShift Virtualization (based on the KubeVirt project) is the leading choice for organizations pursuing cloud-native modernization. Instead of running a traditional hypervisor alongside a separate container platform, Red Hat OpenShift Virtualization brings VMs directly into Kubernetes, allowing organizations to run legacy VMs and modern containers on the same platform managed by the same team with the same tooling.

Choosing Red Hat OpenShift Virtualization is a decision to move away from siloed infrastructure, where VMs and containers are managed by different teams with different tools and different operational models, and toward unified platform engineering. VMs become first-class citizens inside a Kubernetes cluster. Teams use the same CI/CD pipelines, role-based access control security policies, and observability stack for both.

Best for: Organizations that view the VMware exit as an opportunity for long-term modernization, not just short-term migration. If your strategy includes moving workloads to containers within the next three to five years, consolidating VM and container operations onto one platform, or building a foundation for AI, Red Hat OpenShift Virtualization addresses all three on a single platform.



Concern area	Red Hat OpenShift Virtualization
System scalability	Red Hat OpenShift scales by nodes, not by clusters. A single Red Hat OpenShift cluster can manage thousands of nodes and tens of thousands of VMs. Scalability is governed by the Kubernetes control plane, meaning you need to right-size etcd and API server capacity as you grow. The architectural ceiling is significantly higher than traditional hypervisors that hit cluster-size walls.
High availability	High availability is handled natively by the Kubernetes scheduler. If a node fails, the virt-launcher pod is automatically rescheduled on a healthy node. Restart times depend on the storage layer. With the right storage orchestration, pod recovery can complete in 60 to 120 seconds, comparable to or faster than traditional hypervisor high availability for most workloads. Live migration (the vMotion equivalent) is fully supported for planned maintenance.
Disaster recovery	Disaster recovery in Red Hat OpenShift is managed via Red Hat Advanced Cluster Management (ACM). Unlike traditional hypervisors that fail over individual VMs, Red Hat OpenShift can redeploy an entire application stack (VMs, containers, config, and data) to a secondary cluster using a GitOps workflow (Argo CD) and a data management platform like Portworx®. This guarantees environment-identical recovery rather than just VM-level restart. Synchronous and asynchronous replication options are available depending on your storage layer and RPO requirements.
Storage integration	Red Hat OpenShift uses the Container Storage Interface (CSI) as the standard API to connect to storage back ends. CSI enables broad compatibility, and almost any modern storage array is supported. ReadWriteMany (RWX) block volumes are required for live migration. It is worth understanding what CSI does and does not do: CSI was designed to attach and mount volumes. It was not designed to optimize performance, coordinate resiliency, or manage data at scale. For production VM workloads, most enterprises layer a Kubernetes-native data management platform on top of CSI to fill those gaps.
Migration path	Red Hat provides the Migration Toolkit for Virtualization (MTV)/Forklift, a native tool that automates the discovery, planning, and migration of VMs from VMware and other platforms into Red Hat OpenShift. It handles the conversion of virtual disks to persistent volumes and supports warm migrations to reduce cutover downtime. Migration speeds can be further accelerated with purpose-built storage migration tools from ecosystem partners like Everpure.
Ecosystem and tooling	You gain access to the entire Cloud Native Computing Foundation ecosystem. Network security is handled by Service Mesh and network policy. Runtime security is handled by Advanced Cluster Security (StackRox). Observability comes from the built-in monitoring stack (Prometheus/Grafana).
Licensing and TCO	Red Hat OpenShift Virtualization is an operator included in the Red Hat OpenShift subscription, with no additional hypervisor license. You pay for the Red Hat OpenShift nodes. For teams already running Red Hat OpenShift for containers, adding VMs to the same cluster means the hypervisor cost is effectively zero. Organizations migrating from VMware consistently report 30% to 50% reductions in total platform costs when consolidating onto Red Hat OpenShift. ¹
Operational skillset	The skills gap is the biggest hurdle. Your infrastructure team must shift from a GUI-driven, point-and-click management model to declarative infrastructure: YAML manifests, kubectl, and infrastructure-as-code workflows. Admins need to understand Kubernetes concepts (pods, nodes, namespaces, operators) to troubleshoot a VM effectively. The learning curve is real, but Red Hat OpenShift Virtualization now includes a ClickOps-style console for VM administrators that bridges the gap, and Red Hat offers extensive training and certification paths to accelerate the transition.



How Everpure can help

Portworx® by Everpure delivers the storage and data management foundation for building a modern virtualization stack on Red Hat OpenShift, providing the benefits of Kubernetes without rearchitecting existing VMs.

With Portworx, organizations gain a unified, Kubernetes-native data management platform that supports both VMs and containers running on Red Hat OpenShift. Portworx replaces fragmented CSI tooling with a single platform for live migration, synchronous and asynchronous disaster recovery, VM backup and restore, automated capacity management, and I/O control. The Portworx Red Hat OpenShift dynamic plug-in brings these workflows directly into the Red Hat OpenShift console, giving VM administrators a familiar, GUI-driven experience for storage operations without leaving the platform.

For organizations migrating from VMware, Portworx provides enhanced storage migration that accelerates migration speeds significantly over the default migration toolkit approach, plus VMware-familiar storage workflows via Kube Datastore that smooth the operational transition. Portworx runs on any block storage back end (across edge, core, and cloud) with zero platform lock-in and offers a dedicated SKU for modern virtualization workloads to optimize cost efficiency.

More than 45 enterprises across healthcare, manufacturing, government, financial services, and media run VMs on Kubernetes with Portworx and Red Hat OpenShift Virtualization today.

[Learn More About Portworx and Red Hat](#)

4

OpenStack and KVM

Build it yourself, or buy it supported

OpenStack is one of the most compelling alternatives to VMware, but it's worth clearing up a common misconception: OpenStack is not purely a do-it-yourself proposition. It is open source software, yes, but a mature ecosystem of commercial vendors—Red Hat (via Red Hat OpenShift and Red Hat OpenStack Services on OpenShift), Canonical, Mirantis, Platform9, IBM, Breqwatr, and Virtuozzo, among others—ship fully supported, enterprise-grade OpenStack distributions. You can get the economics of open source without betting your infrastructure on a team of two who learned it from a blog post.

That said, OpenStack is still not a shrink-wrapped appliance. It is a composition of loosely coupled services—Nova for compute, Neutron for networking, Cinder for block storage, and Glance for images—that talk to each other through a message bus and a shared database. How you assemble, tune, and operate those services is where the complexity lives. With a commercial distribution, the vendor absorbs a large portion of that burden. Without one, you own all of it.

Best for: OpenStack is the leading choice for government, defense, and European organizations with strict digital sovereignty requirements. Because the code is fully open, it can be audited for backdoors and telemetry—something no proprietary hypervisor can offer. OpenStack has also matured significantly for GPU workloads: KVM's vGPU and PCI-passthrough support allows VMs to consume physical NVIDIA or AMD GPUs at near-bare-metal performance, making it a natural fit for on-premises AI training infrastructure.



Concern area	OpenStack + KVM
System scalability	OpenStack is designed for hyperscale. Cells v2 architecture supports thousands of nodes across multiple data centers. The practical ceiling is usually determined by the stability of your RabbitMQ message bus and MariaDB cluster, not the compute plane itself.
High availability	API services (Nova, Neutron, and so on) run as multiple instances behind load balancers. VM-level high availability is handled by Nova compute’s evacuation mechanism—with a mandatory STONITH fencing layer to prevent split-brain data corruption when a host goes dark.
Disaster recovery	Native multi-region support lets you synchronize VM images (Glance) and block volumes (Cinder) across sites. Failover to a second on-premises cluster, or to a public OpenStack provider, is a well-trodden path.
Storage integration	Ceph is the traditional default: a unified pool for block, object, and file. But it is operationally heavy. NVMe-oF support in Cinder’s volume driver layer now gives you a credible alternative—external all-flash arrays accessible at near-local latency without the Ceph management overhead.
Migration path	No single migration tool matches what Nutanix Move does for Nutanix. The standard approach is a combination of OpenStack APIs and Ansible playbooks to convert VMware VMDK images to QCOW2 format. This is labor-intensive and requires methodical testing. One notable exception: as of OpenStack Epoxy (2025.1), the Everpure Cinder driver supports direct ingestion of vVol-backed VMware volumes, which meaningfully shortens the migration path for organizations already running Everpure on the VMware side.
Ecosystem and tooling	OpenStack’s API surface is one of the broadest in the private cloud world. Native integrations with Terraform, Kubernetes (via Magnum), and bare metal (via Ironic) are mature. Verify that your specific network interface card and switching hardware have a supported Neutron ML2 driver before committing.
Licensing and TCO	With zero software license cost to the OpenInfra Foundation, the TCO equation shifts entirely to people. The savings over annual Broadcom/VMware fees are real—but they need to be reinvested in engineering headcount.
Operational skillset	This is the highest operational bar of any option in this guide. Your team needs to be Linux system engineers fluent in software-defined networking, distributed databases, and message bus troubleshooting—not just hypervisor administrators who’ve learned some CLI. Commercial distributions reduce this bar substantially, but they don’t eliminate it.



How Everpure can help

Ceph is OpenStack’s default storage answer, but it carries real operational weight. Tuning placement groups, managing object storage daemon failures, and debugging network saturation under write-heavy workloads are not trivial tasks. Everpure FlashArray and FlashBlade® integration via the OpenStack Cinder and Manila drivers is the most widely deployed external storage path in the OpenStack ecosystem, and it sidesteps Ceph’s complexity almost entirely.

The driver handles all of the configuration. There is no array preconfiguration requirement—no LUN masking templates to build or volume groups to prestage. Provisioning, snapshots, replication, and host attachment are all surfaced through the standard OpenStack GUI and API. Block storage supported protocols span iSCSI, Fibre Channel, NVMe-RoCE, and NVMe-TCP.

For organizations migrating from VMware, the Epoxy-era Cinder driver adds a material capability: direct ingestion of vVol-backed VMware volumes into the OpenStack environment. For customers already running Everpure on the VMware side, this significantly compresses the migration timeline compared to a full VMDK-to-QCOW2 conversion cycle.

Everpure is qualified across the major commercial OpenStack distributions: Red Hat, Canonical, Mirantis, IBM, Platform9, Breqwatr, and Virtuoizzo.

These are just some of the ways Everpure can help you with OpenStack.

[Read the OpenStack Solution Brief](#)

5

Proxmox VE

The low-cost alternative

Proxmox VE has moved from being a niche open source tool to a primary contender for corporate VMware exits. It is unique because it is a stand-alone Debian-based distribution that combines two types of virtualization—VMs (KVM) and lightweight containers (LXC)—in a single, integrated interface.

Proxmox is the economic pragmatist's choice; it offers 90% of VMware's functionality at roughly half of the overall cost (depending on deployment size). While the software is free, the paid subscription gives you access to the Enterprise Repository, where updates are heavily tested and vetted before release. This provides the peace of mind that IT departments require for production workloads.

Best for: Secure or air-gapped environments, as Proxmox has no mandatory cloud connection. This is often the number one reason it is chosen. It is also a top choice for organizations seeking to lower costs as much as possible, such as SMBs.



Concern area	Proxmox VE
System scalability	Proxmox uses a unique Proxmox cluster stack (pvecm) that supports 32–64 nodes per cluster comfortably. While it doesn't reach the hyperscale of OpenStack, it is more than sufficient for 95% of enterprise data centers. A new 2026 feature, Proxmox Datacenter Manager now allows for the centralized management of multiple global clusters.
High availability	High availability is built in and relies on Corosync. If a node fails, VMs are automatically fenced and restarted on surviving nodes. It requires a three-node minimum to maintain a quorum (or a specialized two-node and QDevice setup) to avoid split-brain scenarios where two nodes think they are both the master.
Disaster recovery	Proxmox offers the highly efficient Proxmox Backup Server (PBS), which handles incremental, deduplicated, and encrypted backups. For disaster recovery, it supports storage replication, which uses ZFS snapshots to send VM data to a remote site as often as every minute, providing a very low RPO.
Storage integration	Proxmox is the only platform that provides a full GUI for managing Ceph (distributed storage) and ZFS (high-integrity local storage). It also supports NVMe over TCP for disaggregated storage, allowing you to connect to external storage arrays.
Migration path	The Proxmox standout feature is a native ESXi import wizard. You can point Proxmox directly at your existing VMware ESXi hosts and import VMs live over the network. It handles the disk conversion (VMDK to QCOW2) and basic driver injections automatically, significantly lowering the barrier to entry.
Ecosystem and tooling	While Proxmox has a massive community, some Tier 1 enterprise backup and security vendors (like Veeam) only recently added official support. You must verify that your specific Day 2 tools natively support the Proxmox API (PVE API).
Licensing and TCO	The software is 100% free (AGPL, v3). You only pay for a support subscription, which is priced per CPU socket (not per core). This typically results in a 20–50% reduction in annual costs compared to VMware vSphere Foundation. ²
Operational skillset	While the GUI is excellent, Proxmox is essentially Debian Linux. To troubleshoot at an advanced level, your team needs to be comfortable with the Linux command line, SSH, and basic bash scripting. It is an admin-friendly platform, but not a black box appliance.



How Everpure can help

Everpure FlashArray is an ideal storage foundation for Proxmox VE, pairing an open, Linux KVM–based virtualization stack with enterprise-class, shared storage. FlashArray delivers consistently low latency and high throughput, so Proxmox clusters can host dense, mixed workloads—VMs, databases, and VDI—without noisy neighbor issues or complex tuning. Always-on global compression and deduplication increase effective capacity and help lower TCO, while rich data services such as instantaneous snapshots, fast clones, and policy-driven quality of service make it easy to protect and right-size VMs at scale.

With nondisruptive upgrades, high availability, and built-in cyber resilience features, Everpure FlashArray keeps Proxmox environments online and recoverable even during maintenance, giving you a modern, flexible alternative to hyperconverged storage architectures.

[Learn More About Everpure and Proxmox](#)

Summing it up

The hypervisor market has moved beyond finding a cheaper hypervisor to selecting a platform engineering destination. Here is a short summary of the key decision points discussed in this guide. Every organization has different requirements, so take these as recommendations, not prescriptive guidance.

- **Choose Nutanix** if your priority is minimizing disruption. It is the most enterprise-proven replacement that mirrors VMware’s maturity without the licensing volatility.
- **Choose Azure Local** if your strategy is cloud-first. If your team already spends 40 hours a week in the Azure portal, this is the path of least resistance.
- **Choose Red Hat OpenShift Virtualization (KubeVirt)** if you want to eliminate the VM silo. This is the only choice that turns your infrastructure into a single Kubernetes-native fleet where VMs and containers are the same.
- **Choose OpenStack** if you are building a factory. If you have 500+ servers and the engineering talent to maintain them, OpenStack offers the only path to 100% vendor independence.
- **Choose Proxmox** if you need to slash your budget by 90% immediately. Proxmox has evolved into a viable production contender for mid-market firms tired of the vendor tax.

Next steps

- Consult with your Everpure account team for a complete discussion on how Everpure can help you in your virtualization journey.
- Request a virtualization workshop.

Explore Virtualization Solutions



1 | <https://www.everpuredata.com/company/newsroom/press-releases/portworx-collaborates-with-red-hot-cost-savings.html>

2 | <https://redresscompliance.com/proxmox-vs-vmware>

Visit Our Website

800.379.PURE

