



## Escaping VMware Lock-in and its Ever-Increasing Costs

*On-Premises and in the Cloud*

by, Marc Staimer

2023

### Premise – VMware Market Leadership in Decline

VMware – Broadcom to close acquisition of VMware 26 May 2023 – is well known for being the major force behind data center virtualization. They are the undisputed on-premises hypervisor market leader and the most prevalent hyperconverged infrastructure (HCI) hypervisor. VMware's well-known value proposition is the substantial reduction of IT infrastructure sprawl by enabling multiple operating systems and applications to run concurrently on the same server hardware. That value proposition was incredibly compelling and revolutionary when introduced more than two decades ago. It still is and made hypervisors must-have technology for IT organizations.

Much has changed since VMware first exploded onto the scene. Open source alternatives have become accepted and practical. Several other vendors have been providing leading edge innovation that leapfrogged the capabilities VMware had at the time.

VMware generally managed to keep pace with many of these new competitive threats by enhancing their core products, but each enhancement was aimed at catching up with the innovators so they could maintain their market dominance. However, their dominance slipped a bit every time they had to play catch-up.

The latest GigaOM research report – [Radar for Hyperconverged Infrastructure](#) – from 21 December 2022 clearly shows that VMware is not the technology leader it once was. That's not its only problem. Many of VMware's enhancements are not in the top tier of the market, often delivering insufficient performance. Some of the issues their customers face include, but are not limited to:

### Networking (NSX) Performance Issues

A significant VMware performance issue example is VMware NSX software defined networking. It is meant to replace top-of-rack switches. Sounds great in theory; however, there's a catch. Intel and AMD CPUs are not designed to be packet switches. Servers are notoriously poor at packet switching typically topping out at a miserly 10,000 packets per second (pps). Compared to actual top-of-rack switches, that's absurdly low. Most top-of-rack switches use custom switch ASICs that deliver at minimum several orders of magnitude greater pps. That's a huge difference and why there's a VMware NSX performance issue.

### Storage (vSAN) Performance Issues

Another VMware performance issue example is evident in VMware's vSAN software defined storage. The most commonly reported user performance issue is the excessive use of compute and network resources that otherwise could be used for applications. There is no free lunch. Something has to process the vSAN software stack. Inter-node communications and more importantly extensive storage data traffic has to flow over the same network consuming more compute resources. Application response times will degrade unless the number of clustered nodes is increased. This ultimately leads to diminishing marginal returns for each additional node. At some point, the next vSAN additional node can actually reduce overall performance. That vSAN performance is degraded when clustered by NSX contention for the same resources.

Customers have to figure out the best hardware configurations for their workloads. And because of the lack of engineered software and hardware synergy, they'll need more hardware infrastructure than they should to achieve their performance objectives.

### Additional VMware Storage (vSAN) Issues

Running vSAN in a cluster mesh is not trivial. Many of the vSAN users interviewed view cluster mesh vSAN as quisquous – that is, very hard to deal with. It takes significant expertise, skill, and time – not exactly for the infrastructure novice.

Further vSAN issues include inefficient node failure data protection. VMware vSAN does not enable access to data of a failed cluster node (server) unless that data has been replicated to a different node in the cluster. That replication effectively consumes 2x the storage. And that's just protecting against a single node failure. Making data available in the event of two concurrent cluster node failures requires replicating the data at least twice meaning 3x the storage. Most storage systems commonly provide continual data access even when nodes fail without requiring full copies of the data.

### VMware Reliability, Availability, and Serviceability (RAS) Issues

VMware is a software company. They do not make or sell hardware. VMware hypervisors run on server hardware. That means there are always multiple vendors involved in a VMware implementation. This is true regardless of whether it is do-it-yourself (DIY) infrastructure, converged infrastructure (CI), hyperconverged infrastructure (HCI), or cloud infrastructure.

Hypervisors are considered foundational software infrastructure similar to operating systems. Customers expect it to work all the time. When there is a problem, the first thing that must be determined is whose problem is it? Who supports it? This can and frequently does lead to finger pointing. The hardware vendors – and there can be several different hardware vendors for servers, storage, and networking – may take the first call, but they have no access to VMware source code. VMware may take the first call, but they have no access to the hardware microcode. All will do their best to resolve the problem. However, multiple support staffs can and often delay satisfactory problem resolution.

## Emergence of Kubernetes and Containers

Containers do not require a hypervisor. They provide application virtualization underneath the OS and don't need to include an OS within each container. They utilize much fewer hardware resources than hypervisor VMs. The same amount of server hardware can support up to 6x the number of containers compared to VMs.

Kubernetes has become the open source de facto standard for container orchestration. Many of the concerns about containers such as reliable backups and permanent storage for stateless containers, have been resolved. More application development on-premises and especially in the cloud are based on containers not VMs. Adoption among mid-tier to enterprise users is growing rapidly.

This puts pressure on VMware licensing revenues as the need for hypervisors declines. Although some applications need to be in a VM under a hypervisor many simply do not. Running containers under a hypervisor is a waste of resources.

## VMware Excessive Cost Issues

Candidly, all of those issues are acceptable to users when VMware competently gets the job done. These customers view it as reasonably cost effective. But cost effectiveness has been rapidly diminishing as VMware costs have continued to rise. While never having been viewed as inexpensive, every major addition such as NSX, vSAN, Tanzu, etc. has increased its cost. That was not the only cost increase. VMware used to license by socket. They determined they were leaving money on the table as CPUs increasingly increased their core count. This allowed customers to use less hardware and fewer VMware licenses, reducing VMware revenue. VMware flipped the script by changing their licensing to cores instead. More cores now mean more cost to the customer.

Higher cost may have been justified when VMware was the innovation leader—and practically a monopoly—offering unique problem solving services. Since they're no longer the innovation leader, it's no longer the case.

The market place has irreversibly evolved. There are better and lower cost alternatives to VMware today. Open source Linux by itself is a major VMware competitor. Linux includes many of the same VMware hypervisor, containers, and Kubernetes capabilities. In fact, all of the major public clouds—AWS, Azure, GCP, Oracle, and IBM have built their clouds on Linux or in the case of Microsoft, Windows. None are based on VMware. What about the VMware cloud?

VMware, as a proprietary offering, creates other cost issues. Moving VMs or apps to non-VMware alternatives requires conversions. Conversions require tools, time, skilled professionals, and cost. It's a way to de-incentivize a customer from leaving. It's a form of lock-in. Lock-ins limit options and always mean higher costs.

## One Other VMware Notable Concern for Customers

There have been a lot of ownership changes at VMware. In the latest, VMware is being acquired by Broadcom. VMware was initially a private company founded in 1998. It was acquired by EMC in 2004. EMC spun out a portion of VMware publicly in 2007. Dell acquired EMC in 2016. In 2021 Dell spun out the rest of VMware publicly. Each of these ownership changes creates uncertainty. There are cultural, support, workforce reductions, reorganizations, and policy changes that occur over time. Whether the Broadcom acquisition is good, bad, or indifferent to customers is an unknown. In the meantime, it can cause angst and disruption.

All of these VMware issues are fundamentally behind its noticeable market share decline. Clearly, users are looking for the exits.

## Obstacles to Moving Away From VMware

VMware customers researching exit plans have several major concerns:

- How much time, effort, training, and cost does it take to move from VMware?
- Do they have to become experts in software and hardware integration?
- What does it take to move VMs and applications from VMware to open source or other alternatives?
- What mission-critical capabilities will they lose?
- Will it weaken security?

- Will support be more difficult especially when troubleshooting issues?

All are valid concerns.

This research report examines how the Oracle Private Cloud Appliance X9-2, a.k.a. PCA X9-2, provides a dynamic alternative to VMware that's better, more complete, higher performance, simpler, and lower cost. It also details how the PCA X9-2 effectively deals with and solves each of these concerns.

## Oracle PCA X9-2

PCA X9-2 is a highly engineered, rack-based private cloud native infrastructure specifically architected for performance, reliability, and low cost. It's designed to efficiently consolidate a wide-range of middleware and application workloads with far fewer resources than VMware implementations. The PCA X9-2 combines and integrates high performance servers, high performance network switches, high performance storage, and intuitive cloud management. Another way of grasping the significance of PCA X9-2 is that it's a unique combination of converged and hyperconverged infrastructure, called Ultraconverged infrastructure, or UCI, that provides a level of synergy previously unobtainable. That synergy produces unmatched performance, simplicity, reliability, and manageability. It starts with performance.

### PCA X9-2 Superior Performance

#### PCA X9-2 Entire Software Stack

What makes the PCA X9-2 unique is that it delivers a cloud-like experience with its own control plane and administrative services. The PCA X9-2 is built on Oracle Linux – Red Hat binary compatible – a hybrid cloud native system that includes KVM, Oracle Cloud Native Environment (OCNE), Kubernetes, Docker containers, and Verrazzano that looks, feels, and operates faithfully like the IaaS in Oracle Cloud Infrastructure (OCI) with OCI compatible APIs. Each PCA X9-2 is an Availability Domain (AD) containing three Fault Domains (FDs). These act as fault isolation boundaries consisting of one or more compute nodes.

A big chunk of PCA X9-2's superior performance is directly attributable to the software stack. While KVM is a type 1 hypervisor similar to VMware, it delivers considerably greater performance at much lower overhead. KVM consumes far less server resources than VMware. This frees up those resources for application performance. KVM is much faster at creating and starting a virtual machine than VMware ESXi and has much faster VM performance. It runs applications at near-native speeds. The same cannot be said for VMware hypervisors. The [SPECvirt\\_sc2013 benchmark](#) results proves this point. The PCA X9-2's full-stack performance optimizations accelerate crucial and mission-critical customer workloads by up to 40% compared to traditional DIY architectures<sup>1</sup>.

That's simply the beginning on how PCA X9-2 delivers superior performance. Equally important is how Oracle engineered and integrated PCA X9-2's hardware with the software stack.

#### Internal PCA X9-2 Networking

PCA X9-2's internal network is based on the Oracle Cloud Infrastructure (OCI) network. It supports L3 networks, virtual cloud networks (VCNs), subnets that rely on high-speed physical infrastructure, all with a separate data plane, and management networks. PCA X9-2's networking leverages isolation and physical redundancy for performance and security.

This internal networking alone is a massive performance improvement over VMware on DIY or HCI infrastructures. PCA X9-2 uses a very high performance private internal network that's not exposed to the customer's datacenter network. Each PCA X9-2 has a pair of high-performance leaf switches, a second pair of high performance spine switches, and a separate management switch.

Each leaf and spine port delivers 100Gbps maximum throughput. The spine switches use 5 interlinks (500Gbps total). The leaf switches use 2 interlinks (200Gbps) and 2x2 cross links to each spine. The interconnect is a two-layer design. Leaf switches connect all of the PCA X9-2 rack components. Each leaf switch is connected to each spine switch. Spine switches are the internal interconnect backbone handling routing. Each of the switches are rated at billions of packets per second (Bpps), not thousands as with

---

<sup>1</sup> Per CSP of Brazil.

VMware NSX. This private internal network supports VLANs, VCNs, subnets, gateways, security lists, routing tables, and more. The management switch is a 48-port 1Gbps Ethernet switch.

Put simply, PCA-X9-2 has an extremely fast internal network that completely outclasses any other VMware ESXi HCI or DIY implementation. It eliminates outside networked contention, hot spots, or blocking that enables it to deliver much lower end-to-end latencies, a lot more IOPS, and orders of magnitude more throughput.

#### PCA X9-2 Compute Nodes

PCA X9-2's compute nodes are architected to provide the optimal balance of cores, memory, and I/O throughput for enterprise applications. It is built specifically for the demands of enterprise and virtualization workloads and uniquely run Oracle Database.

Each of PCA X9-2's compute nodes is a dual-socket 1 RU server. It comes with 32 DIMM slots and up to 2 TB of DDR4 ECC DIMM memory using 32 GB or 64 GB at DDR4-3200 dual rank. They also come with up to 32 cores per socket, one Platinum, two Gold, or one Silver Intel® Xeon 3<sup>rd</sup> generation processors, two boot flash drives of M.2 SATA 240GB (480GB total), and three PCIe 4.0 16-lane expansion slots. There are also four small form factor SFF/U.2 drive bays per server enabling configurations up to four 6.8 TB NVM Express SSDs for a total capacity of 27.2 TB of low-latency, high-bandwidth flash. That hot-pluggable, high-bandwidth flash is engineered to work synergistically with the Oracle Database Smart Flash Cache when it's running in this virtual environment.

The high performance CPUs, 320 GB/sec of bidirectional I/O bandwidth, very fast flash, combined with the high core and memory density makes each PCA X9-2 compute node a performance powerhouse. A PCA X9-2 rack clusters a minimum of three compute nodes and as many as 20 clustered compute nodes.

#### PCA X9-2 Storage

PCA X9-2's storage subsystem is based on the dual Oracle Storage Appliance ZS9-2 HA cluster. It supports block, file, and OCI object stores. It supports both high-performance flash SSDs and high-capacity HDDs. ZS9-2's storage controllers are specifically engineered for exceptional performance, low latencies, high IOPS, and high throughput. ZS9-2 systems in production in the field have proven this out. That outstanding performance comes from several unique ZS9-2 innovations:

- Application enhancing highly efficient large DRAM and flash caching that enables sustained high IOPS.
- Very low-latency DRAM focus – up to 4TB – supports IO intensive applications and thousands of VMs per ZS9-2.
- Massively parallel hardware and software designed to provide sustained high throughput at over 18GBps for data intensive applications.
- Automated prioritization of database IOs that eliminates those IOs from being delayed by non-mission critical workloads.
- Built-in data protection optimizations that recognize those workloads and reduce data movement while increasing its throughput, deliver up to 30TB/hr backups, and 37 TB/hr restores.
- Consistent and persistent flash high-performance in configurations with up to 8.8 PB of all flash.

All of these ZS9-2 capabilities provide much higher storage performance. So too does the hardware it runs on. The ZS9-2 hardware is tested, quality assured, production proven, and high performance. Each ZS9-2 storage controller node comes with dual 24 core 2.1 GHz Intel® Xeon® processors, 1 TB DRAM. ZS9-2 HA clusters are configurable with both high performance (supports up to 47 enclosures) and high capacity (supports up to 48 enclosures) storage media in 2U drive enclosures. Each high-performance drive enclosure configures up to 24 SAS-3 (7.68TB) 2.5 in drives. Each high-capacity drive enclosure configures up to 20 SAS-3 (18 TB) 3.5-inch 7200 RPM drives. The high-capacity drive enclosures also include 2x read SSD accelerators and 2x write SSD accelerators.

#### PCA X9-2 Oracle Database Performance Optimization

PCA X9-2 is specifically augmented for Oracle Database. The integration of Oracle Database HCC with ZS9-2 enables 10-50x compression for data warehouses and archives with customers achieving an average of 12x .HCC which greatly increases data warehouse and archive efficiency while substantially reducing Oracle

Database storage capacity consumption. It also speeds up Oracle Database processing of the compressed data because the database can perform many operations directly on the compressed data. The average data reduction for Oracle Database data warehouses on ZS9-2 is 3-6x better than the best competitive primary storage data reduction results<sup>2</sup>.

And for those database applications demanding the absolute lowest latency –  $\leq 19\mu\text{s}$ , highest SQL IOPS – up to 27.6 million 8K SQL read IOPS, and throughput – up to 1 TBps – need Oracle Exadata. There is no faster combination of any database type and the hardware it runs on than Exadata. Exadata plugs directly into the PCA X9-2 internal switches via multiple 25 Gbps Ethernet links. No other DIY, CI, or HCI delivers anywhere close to the same level of database performance.

It's this extensively engineered combination of software and hardware synergy of the software stack, interconnect, compute, and storage that delivers PCA X9-2's unprecedented performance in an ultraconverged infrastructure appliance.

### PCA X9-2 Simplicity

A significant PCA X9-2 advantage comes from extensive automation. That automation makes it much easier and faster to implement, operate, manage, and troubleshoot. It's aimed at both the non-expert and expert administrator alike. PCA X9-2 auto-configures compute servers, networking, and storage for optimal performance reducing time-to-value to [hours or days instead of weeks](#).

For storage, it automatically discovers and configures PCA X9-2 capacity when first powered up and when capacity is added. One of three Fault Domains (FD) is automatically provisioned to provide application resiliency. PCA X9-2 rack compute capacity is also automatically configured.

For networking, PCA X9-2 automates configuration for best performance substantially decreasing deployment and operational workload complexity for IT network management teams. PCA X9-2's highly integrated software-defined networking (SDN) is very different from VMware in that it operates multiple internal virtual networks over a single very high performance physical one. It doesn't attempt to convert server resources ill-suited for networking to run the network.

A big PCA X9-2 simplicity advantage stems from the way it minimizes disruptions or outages. All compute, switch, and ZS9-2 node patches, upgrades, and repairs are performed non-disruptively via rolling implementations. Before a compute node is to be patched, upgraded, or repaired, its application instances are moved, and placed in maintenance mode, then returned upon completion.

Containers are the latest iteration of application virtualization. PCA X9-2 deeply integrates Oracle Cloud Native Environment (OCNE) making it quite easy to automate application container deployment, scaling, and management with Kubernetes.

An essential part of PCA X9-2's simplicity comes from manageability. PCA 9X-2 delivers this in spades with comprehensive monitoring and visualization with Prometheus and Grafana. It streamlines customer IT team workloads by enabling full-stack solution management, including Kubernetes container deployments. Combined with its advanced management and analytics engine, PCA X9-2 provides real-time actionable insights that accelerate troubleshooting. IT organizations can quickly understand why individual workloads are creating performance-limiting congestion and fix it.

Cloud portability is another area greatly simplified by PCA X9-2. Preparing for a move to or from the cloud has never been easier. PCA X9-2 services are compatible with OCI. The workloads, user experience, tool sets, and skills are completely portable between PCA X9-2 and OCI. The infrastructure is consistent between PCA X9-2 and OCI for compute, network, storage, and identity. The OCI designer and visualization toolKIT (OKIT) streamlines application deployment on either the PCA X9-2 or OCI providing incomparable private and public cloud portability.

Stress-free flexible scalability is greatly simplified by PCA X9-2. The base PCA X9-2 system starts with 180 CPU cores, 180 TB of usable storage, and fully redundant 100 Gbps Ethernet internal networking. These resources can be scaled up to 1080 CPU cores and nearly 9 PB of usable storage. PCA X9-2's flexibility means compute, network, and storage can either be scaled, upgraded, and innovated at completely different time

---

<sup>2</sup> Users of primary storage with deduplication and/or compression consistently report to DSC getting no more than 2 to 2.8x Oracle Database reductions.

frames or at the same time. That allows a new server, new network switch, or new storage enhancement to be implemented separately without having to replace or upgrade the entire infrastructure.

Oracle Database AWR integration with ZS9-2 storage analytics makes the database administrator's (DBA) life much simpler when it comes to performance troubleshooting. That unprecedented tight correlation provides DBAs unmatched actionable insights into IO storage performance on an IO type basis enabling them extraordinary control.

Another well-received exceptional automation capability comes with ZS9-2's auto-tuning. With other systems, storage, database, and application administrators spend enormous amounts of time tuning storage to get their best performance. PCA X9-2's ZS9-2 storage eliminates that burden. Oracle Intelligent Storage Protocol version 2 (OISP2) auto-sets up and auto tunes the ZS9-2, reducing steps and processes by as much as 70%. This enables administrators to focus on more strategic efforts versus mundane but previously necessary ones. Those important mundane efforts are automatically completed for them.

Whereas the PCA X9-2's value proposition is quite compelling, a very important concern for many VMware users is the time, effort, and cost it will take to move their VMs and applications to PCA X9-2. The answer is, not much. Oracle partnered with [RackWare](#) to make moving to PCA X9-2 a trivial process. RackWare makes it very easy and intuitive to migrate existing VMware workloads on-premises or in the cloud to PCA X9-2, OLCNE workloads to PCA X9-2, remote workloads to PCA X9-2, Nutanix workloads to PCA X9-2, and other Cloud (AWS, Azure, GCP) workloads to PCA X9-2.

One thing that reassures new PCA X9-2 customers is the OCI standalone open source OCI Designer Toolkit. It permits customers to rapidly design and prototype PCA X9-2's infrastructure. It's that rapid design prototyping when combined with the PCA X9-2's deeply integrated management that reassures them the vast majority of common manual IT tasks are easier than their prior experience. And they love that administrators can manage multiple PCA X9-2s from a single dashboard.

#### PCA X9-2 RAS

While the definition of RAS is Reliability, Availability, and Serviceability, what it means to a user is that that the solution has redundancy and resiliency, limited to no planned outages or disruptions, and layered security. PCA X9-2 was designed specifically with those RAS requirements in mind. It leverages the same production proven technologies used in Oracle Engineered Systems, such as Exadata, and Oracle Cloud Infrastructure that maximize uptime for customer workloads.

The hardware and software are built for resiliency and redundancy to avoid any single points of failure (SPOF). PCA X9-2's high availability architecture has three fault domains for compute and no SPOF for storage and networking making sure crucial customer applications continue to operate when there are component failures.

System services resources are separated from customer workloads ensuring their workloads are not impacted by system resource consumption. This improves workload availability, reliability, consistency, and security. In addition, customer workload instances resources are isolated from one another to ensure deterministic, repeatable performance, which increases availability, reliability, and improves security. The PCA X9-2's cloud-like tenancy and granular storage encryption also enhances security. It enables companies to support up to 8 separate customer-managed tenancies per system.

There's more. As previously discussed, PCA X9-2's non-disruptive patching, maintenance, and upgrades are implemented without interrupting operations, applications, or users. It's a significant increase in serviceability and availability. This is especially important when vulnerability patches are released. DSC research has found the average time to implement vulnerability patches in mid to large enterprises is more than 100 days because they're disruptive, generally requiring scheduled downtime. That's a lengthy gap for the cyber crooks to take advantage of vulnerabilities.

PCA 9X-2 eliminates that gap by enabling those security vulnerability patches to be implemented as soon as they are released without operational disruption.

Disaster recovery (DR) is an essential part of every IT organization. PCA X9-2 includes built-in VM-level disaster recovery using Oracle Site Guard between multiple locations. This speeds up recovery time objectives (RTO) allowing IT teams to quickly restore normal customer operations from a major outage emergency, thus increasing availability.

It is often said the best ability is availability. The PCA X9-2 delivers best-in-class availability. Much better than VMware DIY or HCI alternatives.

### PCA X9-2 Lower Cost

The essential software of PCA X9-2 is open source. KVM virtualizations, Kubernetes, Verrazzano, OCNE, containers, and Linux are all open-source. There is no proprietary software lock-in, and the associated high pricing found in VMware. There are no license or subscription costs.

Every PCA X9-2 includes OCNE, KVM virtualization, and Oracle OS system software at no added cost. Even Oracle support costs cover all of the software that comes with the PCA X9-2. There are no hidden costs. As has become clear, the PCA X9-2 is a complete, pre-built, pre-tested, robust, engineered system, minimizing, or eliminating integration costs, automating much of the management, and accelerating time to value.

Make no mistake, PCA X9-2's one throat to choke support of the entire ultraconverged infrastructure including the entire software stack, reduces costs. Reduced costs from IT ops, patching, troubleshooting, outages, disruptions, labor, and time – the most precious resource.

The combination of all of these factors provides customers a highly performant, agile, intelligent infrastructure at a much lower cost than VMware DIY or HCI implementations. And, with Oracle's partner, RackWare customers should feel confident about moving their VMware applications to PCA X9-2.

### Conclusion

VMware's proprietary nature, loss of innovation leadership, and excessive cost combined with the emergence of lower cost, and open source alternatives with equivalent or better functionality has led customers to question their loyalty to VMware. Consequently, they are looking to lower costs while maintaining mission-critical capabilities, minimal specialized training or expertise, exceptional support, ongoing technological innovations, and a single throat to choke.

Oracle provides the ideal solution to all of these needs with PCA X9-2's ultraconverged infrastructure. As detailed above, it delivers much higher performance, noticeably better simplicity and ease-of-use, significantly better RAS, one throat to choke support, at much, much lower costs. VMware users looking to escape VMware lock-in and its ever-increasing costs, on-premises and in the cloud, should take a long hard look at the differentiated capabilities that Oracle PCA X9-2 delivers.

### For More Information on the Oracle PCA X9-2

Go to: [Introducing the Oracle PCA X9-2 Video](#)

[Oracle PCA X9-2](#)

[Oracle PCA X9-2 Datasheet](#)