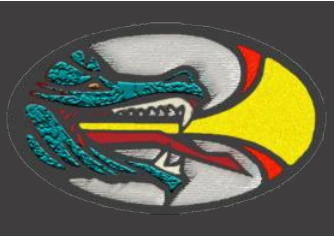

Author: Marc Staimer

W H I T E P A P E R

Oracle Private Cloud Appliance (PCA) X8

Obsoletes Nutanix, Dell, Cisco, HPE, NetApp HCI & CI

Oracle PCA X8 Obsoletes Nutanix, Dell, Cisco, HPE, NetApp HCI & CI

Introduction

Converged Infrastructure (CI) and Hyperconverged Infrastructure (HCI) have been data center rock stars for several years now. And yet at this point in time, there is no integration with applications and databases; they require extensive expertise, skills, and training; there's minimal automation; patching is manual and disruptive; there are significant security holes; and costs have tended much higher than expected. Oracle Private Cloud Appliance (PCA) X8 is the first system to address these problems that others have not. A deeper dive shows how.

Converged Infrastructure

CI integrates physical servers, hypervisors, networking, and storage into a single system with a common management interface. Systems are ordered as a single SKU. They're then staged, integrated, configured, and tested before shipment. The CI system is then managed, operated, and troubleshot as a single system. A key CI value proposition comes from having a single throat to choke for maintenance or root cause analysis, all of which simplifies the life of the administrator(s). Admins only have to turn it on, configure the network, configure VMs, install applications, provision the storage, and they're off to the races. What's the catch?

There are several. Cost is a major issue. CI is primarily a layer of integration, UI, management software, packaging and testing—all of which adds considerable labor cost for the vendor. There is a healthy premium for this CI convenience. Then there is the issue of rack space real estate consumption. The servers, switches, and storage systems are not rack optimized, consuming excessive floor tiles, power, cooling, cables, transceivers, conduit, admin time and allocated overhead adding more cost. The next issue is the failure to simplify expertise, training, and experience requirements in multiple disciplines including physical servers, virtual servers, hypervisors, containers, Kubernetes, networking, and storage. More expertise, training, and experience add yet more cost. Finally, there are few to no applications or databases integrated with the CI system. That increases time and expense to implement, integrate, test, and provision each application and database. These shortcomings led to the development of Hyperconverged Infrastructure (HCI).

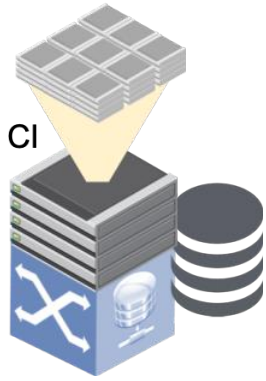

Hyperconverged Infrastructure

HCI combines several disciplines by virtualizing compute, storage, top of rack or leaf switching into a single managed system. It places everything in the x86 server. HCI is essentially data center infrastructure in a box that simplifies implementation, configuration, provisioning, management, orchestration, operations, and troubleshooting (root cause analysis) of the hypervisor, switching, and storage. It doesn't just integrate systems at the management layer; it integrates them as features of a single clustered system. This deep integration solves the multi-discipline skills and expertise issue of CI systems with a simple intuitive UI. This reduces costs. By putting everything under the hood of an x86 server, cost is further reduced. Utilizing commodity-off-the-shelf (COTS) x86 servers and storage media reduces hardware costs. Virtualizing the top of rack or leaf switch reduces rack space, cabling, power, cooling, etc. and more cost. Implementing software defined storage (SDS) and utilizing internal server storage media, eliminates storage controllers and separate arrays. Once again, less rack space, cabling, power and cooling, etc. resulting in yet more cost reduction. HCI is positioned as the private cloud alternative to public clouds because it offers many of the same dynamics, elasticity, and value at a lower cost than do-it-yourself x86 servers, storage and switches. It appears to have solved most of the CI problems.

But, and there is always a but, HCI has its own serious issues. The first issue is scalability. HCI is a scale-out, shared everything architecture. Most HCI clusters are limited by hypervisor constraints on clustering, typically topping out at 32 or 64 nodes. This is exacerbated by HCI's coarse grain scaling. It's generally architected to limit or prevent separate scaling of processing, networking, and storage. There are some HCI systems that have exceptions for external storage—but, not many. If the system needs more compute, it automatically adds more networking and storage even if it doesn't need it. If it needs more storage, it's going to get more compute and networking. That's coarse grain scalability at work. That nodal limitation becomes more apparent through over-utilization of the server resources.

The HCI architecture consumes the same x86 processors, memory, and IO for virtual machines, software defined networking (SDN), and software defined storage (SDS). All are CPU- and resource-intensive. Every node requires all of the software defined services. The most resource-intensive is SDS. It ties up a lot of resources making them unavailable to VMs, networking, containers, and applications leading to the deployment of yet more nodes, adding more cost. Another issue rarely discussed is how inefficient and ineffective x86 processors are at packet switching as compared to virtualized top of rack leaf switches when using SDN. The x86 processors switch multiple orders of magnitude fewer packets per second than packet switch chips. This is why every major networking switch in the market is based on packet switch chips from Broadcom, Intel (Barefoot Networks), or NVIDIA (Mellanox). Packet switching by x86 chips means once again more nodes and more cost. There are also server node real estate limitations for storage media. Each node in the HCI cluster has a hard limit for Flash SSDs and/or HDDs. That limitation ends up requiring yet even more nodes and, of course, more cost.

HCI data protection is another node inflation typically overlooked. All data must be mirrored on write to ensure that data is available to every node in the cluster in case a node fails. To protect the data against multiple concurrent nodal failures requires N+1 mirrored copies for the total number of failed node protections. Each copy of the data consumes 100% more storage, causing more node inflation.

Similar to CI, HCI only integrates at the infrastructure layer. Rarely, if ever, do HCI systems additionally integrate applications, databases, application management, performance tuning, multi-cloud management, devops, devsecops, monitoring and troubleshooting. By definition, it's a limited convenience. Similar to CI, deploying applications and databases on HCI requires admin time and expertise to integrate, test, tune and manage, resulting in unbudgeted costs and much longer time-to-value.

Disruptive Patching/Upgrades Plus Exposure To Cyberattacks

One of the biggest issues that both CI and HCI generally ignore is disruptive patching and upgrading. They attempt to mitigate the issue by using rolling patching/upgrades. Rolling patching/upgrades employ live migration or vMotion to patch and upgrade a hypervisor system or VM. It patches or upgrades a system, then moves the VMs from a non-upgraded or unpatched system to the upgraded or patched system. It then upgrades or patches the previous system. The problem with this approach is that the OS and application patching and upgrading are still disruptive and require scheduled downtime. The current HCI processes are manual, labor-intensive, repetitive, and subject to frequent human errors. More importantly, they are a massive security hole exposing the unpatched systems to cyberattacks.

Because OS patching is disruptive and requires scheduled downtime, it takes a while for a patch to be implemented. Ponemon Institute and ServiceNow IT cloud services interviewed more than 3,000 cybersecurity professionals worldwide[1] and found 48% of IT organizations experienced at least one and possibly more data breaches within a 2-year period. Of those experiencing a breach, 57% attributed it to a vulnerability where a patch had been available but not applied at the time of the breach. When an OS vulnerability patch is released all of the vulnerabilities and security holes that the patch fixes are documented. This starts the clock ticking. The release of the patch and documentation has just informed the cybercriminals where to strike. It becomes a race between the cybercriminals and the IT organization's ability to patch those vulnerabilities before they can be compromised.

The cybercriminals reverse-engineer those vulnerabilities in weeks or even days. They generally know that the vast majority of IT organizations take months, often many months to apply the patch. Even when they do apply a patch, it is often incomplete. The average IT organization patches fewer than 40% of their

[1] Today's State of Vulnerability Response: Patch Work Demands Attention

affected vulnerable systems within 30 days of disclosure of the vulnerability according to the 2019 Verizon Data Breach Investigations Report.

Why do IT organizations wait to apply patches to known security vulnerabilities? First, it's disruptive, requiring coordination between multiple departments and users to schedule downtime to apply the patch update. Second, patching is generally a lengthy manual labor-intensive task. According to the previously cited Ponemon Institute and ServiceNow report:

- 55% reported spending more time with manual processes than rapidly responding to vulnerabilities
- 61% feel frustrated about their reliance on manual processes when patching vulnerabilities
- 12 days on average were lost manually coordinating across teams for every vulnerability patched
- 65% recounted difficulties triaging which vulnerabilities to patch first and which can wait

NopSec has reported that the average IT organization takes more than 103 days (more than 3 months) to apply vulnerability patches. That number rises to 176 days for financial institutions.

Oracle recognized all of these CI and HCI issues and engineered the Oracle Private Cloud Appliance (PCA) X8 to solve them. It is the next generation of CI and HCI, surpassing both in a new category—Ultra-converged Infrastructure and making all other current CI and HCI systems…obsolete. Keep reading to see why.

Table of Contents

Known CI and HCI Issues Review

CI Issues

As previously stated, legacy CI has numerous issues including:

- Limited system integration typically only up to the management of the infrastructure layer
  - No integration at the database or application layer
- Too much required multi-discipline operational knowledge, training, expertise, and skill
  - Servers, networking, storage networking, and storage
- Extensive disruptive patching and upgrading downtime requires delayed scheduling
  - Decreased security with increased cyberattack vulnerabilities
- High cost:
  - Integration costs
  - Professional services costs
  - Power & cooling costs
  - Excessive data center rack space consumption
  - Admin labor costs

HCI Issues

HCI issues are somewhat different but just as numerous:

- Coarse grain scalability
  - Scales at the server level
  - Scaling processing, networking, or storage requires adding the other two as well, regardless of need
- Software defined networking limitations.
  - x86 processor virtual packet switching in the low thousands vs. millions for packet switching chips
- Very rarely works with external shared storage
- No integration at the database or application layer
  - Requires database and application expertise to test, tune, deploy, and manage
- Limited scalability:
  - Typically, no more than 64 nodes
- Extensive disruptive patching and upgrading downtime requires delayed scheduling
  - Decreased security with increased cyberattack vulnerabilities
- Costs higher than projected

The Oracle Private Cloud Appliance X8 (PCA X8) Engineered System

Oracle PCA X8 is a new breed of converged system. Indeed, it's a new category referred to as "Ultra-converged Infrastructure" or UCI. UCI goes the extra mile beyond CI and HCI. It does this by optimizing server, network, and storage hardware with virtualization, containers, databases, and applications, integrations with performance, automation/ease-of-use, security, and reliability. As such, PCA X8 surpasses CI and HCI by solving the issues that neither CI nor HCI currently address at a comparable or lower TCO than HCI. A more in-depth examination shows how.

Engineered Multi-Level Integration

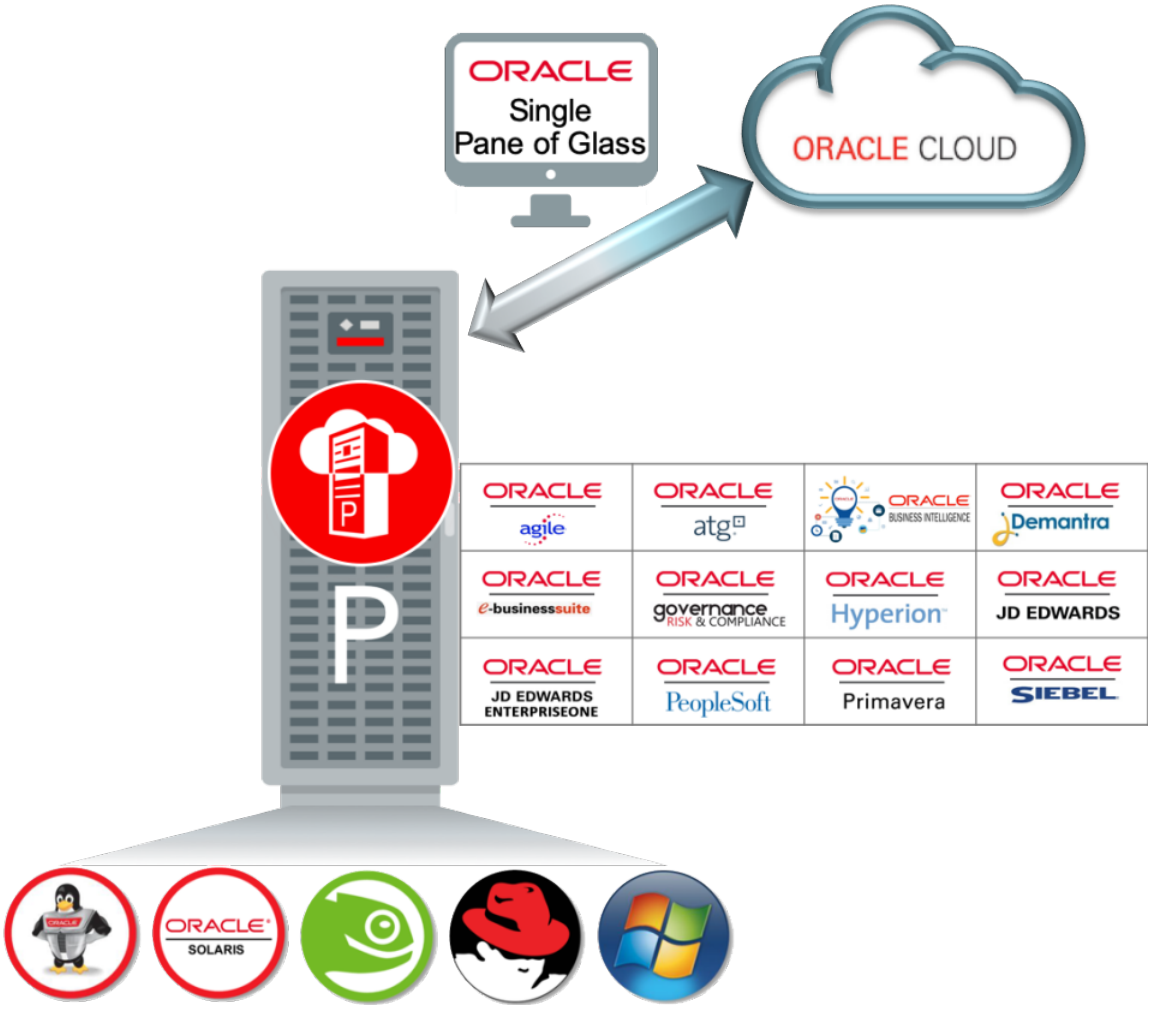

Oracle integrates the best Intel-based Oracle x86 servers (X8-2), Cisco high performance 100Gbps network switches with Oracle Networking software, and Oracle ZFS storage systems – it can utilize external shared file or block storage as well – in a small standard rack. It next integrates Oracle Virtual Machines (OVM), Oracle Linux, Kubernetes, Oracle Applications, Oracle Fusion Middleware, application containers, Docker containers, even Kata secure containers supporting Red Hat, Suse, Ubuntu, and Oracle Linux, Microsoft Windows, Oracle Solaris, and containerized cloud native applications. PCA comes complete with OVM and choice of OS.

Oracle next provides downloadable prebuilt Oracle VM templates that help deploy Oracle applications, Oracle Database, Oracle RAC, and middleware in minutes. In addition, PCA supports Oracle Linux Cloud Native Environment (OLCNE) for automating deployment, scaling, management, and migration of containerized applications, and multi-cloud implementations. Oracle Container Registry is a trusted source that provides simple access to Oracle products for use in Docker containers.

No CI or HCI offering delivers this level of unprecedented application, middleware, cloud native, and database engineered integration with built-in automation and optimization.

PCA's extensive engineered system integration means that everything is managed as a single system. Servers, networking, and storage are FRU[2] subsystems of a single highly tuned system. Not separate systems loosely connected at the management layer.

The result is PCA simplicity that favorably compares to HCI software defined datacenter offerings with the general application performance that matches or exceeds the best of breed CI subsystems. That simplicity

[2]Field Replaceable Unit

and automation enables power-on to production in minutes to hours at most, instead of the days or weeks required for HCI and CI.

Performance Optimization

PCA is extensively performance optimized for the Oracle Database for input / output operations per second (IOPS), throughput, RAC, and data protection. This type of engineered performance optimization is simply not available from any CI or HCI implementation. It makes PCA the second-best option for IOPS intensive workloads such as On-Line Transaction Processing (OLTP) databases and bandwidth intensive workloads such as data warehouses, business intelligence analytics, AI-machine learning, AI-deep machine learning, video processing, and animation rendering. The best option is still the Oracle Exadata X8M. Nothing currently comes close to it for performance, availability, automation, ease-of-use, scalability, and security.

In addition, PCA is performance-optimized for Oracle Fusion Middleware, Oracle Applications, OVM, Kubernetes, containers, and more. And, even though PCA utilizes standard Intel X86 servers, Oracle and Intel worked together to engineer in greater performance optimization for Oracle Linux, Oracle Database, and Oracle applications that are not available with any other processor. PCA then leverages the high-performance standard 100Gbps Ethernet to provide the extremely low latency high bandwidth connections required to get optimal performance between servers and storage.

Scalability

PCA scales non-disruptively from 2 server nodes to 25 in a single rack one or more nodes at a time. Each node includes 384GB, 768GB, or 1.5 TB of memory. PCA also supports up to 15 additional capacity and flash storage trays. The trays can be all flash, all HDD, or a mix of both. PCA similarly works well with external shared file and block storage arrays. Oracle VMs support up to 128 vCPUs and a variety of guest operating systems, including Linux, Microsoft Windows, and Solaris. What gives PCA effectively unlimited scalability is Oracle Enterprise Manager (OEM).

Oracle Enterprise Manager makes private and hybrid cloud deployments, operations, troubleshooting, and root cause analysis seamless. They become easily usable resources with different performance and cost parameters. OEM transforms PCAs into on-premises private clouds, empowering IT organizations to provide managed services internally such as Infrastructure as Service (IaaS), Platform as a Service (PaaS), and Oracle Database as a Service (DBaaS). It facilitates business users and developers getting rapid and self-service cloud services access with admin centralized governance, including ALL Oracle PCAs and attached Oracle ZFS Storage Appliances in a single pane of glass. Nothing smooths the management of private and public clouds combinations as seamlessly as Oracle Enterprise Manager.

CI and HCI hybrid cloud deployments, on the other hand, tend to be wasteful. They generally use the public cloud as a backup repository and/or standby DR site. Some HCI software can be deployed in the public cloud on bare metal to provide the same ecosystem as the HCI on-premises. However, doing so is costly and has limited scalability and performance. PCA and OEM do not have that problem. OEM utilizes the Oracle Linux Cloud Native Environment with the elasticity and low cost that implies.

Reduced Expertise, Skills, Knowledge, and Experience Requirements

PCA shields administrators from having to be an expert in its subsystems. The modern UI, automation, virtualization, and OEM make the PCA very easy for admins to deploy, configure, provision, manage, troubleshoot, and perform root cause analysis. Oracle ILOM provides advanced online fault diagnostics and fault isolation, combined with special hardware instrumentation to improve security and increase uptime. Operations are intuitive with unique automation. No field integration or expertise is required.

Non-disruptive Online Software Patching and Upgrades with Enhanced Security

Oracle has made a point of eliminating disruptive software patching for Oracle Database, Oracle Applications, Oracle Fusion Middleware, Oracle Linux, Oracle Virtual Machines, and the PCA system. Oracle assumes all Oracle applications, databases, and operating systems are mission critical and IT cannot afford

to take them out of service for patches or upgrades. Vulnerabilities need to be eliminated as quickly as possible to keep the organization in compliance and reduce risk.

This is why the PCA and all Oracle Engineered Systems practice proactive non-disruptive patch management. When Oracle discovers vulnerabilities, it packages the patches in a complete stack. That patch stack is implemented non-disruptively online for Oracle Linux, Oracle Virtual machines, Oracle supported containers, Oracle Database, and Oracle Applications.

No other CI or HCI system does that today. Oracle PCA is the first and only.

PCA additionally provides secure multitenancy with up to eight easy to set up and operate, trusted isolated tenant groups, eliminating noisy and nosey neighbors.

Reducing the Cost of Performance and Convenience

PCA combines the best of CI, HCI, and unique Oracle co-engineering. Fine grain scaling of servers, networking, and storage reduces wasted cost. Shared high performance fully redundant and highly available storage reduces storage consumption.

Unique Oracle Hybrid Columnar Compression (HCC), an Oracle Database feature, radically reduces storage consumption for data warehouses by 10-15X. That's approximately 4 to 6X better data reduction than the best deduplication and compression algorithms on the market. HCC is only available on Oracle engineered systems and storage.

PCA specifically engineered integration of Oracle Databases, Oracle Fusion Middleware, Oracle Applications, Oracle Linux, OVM, Kubernetes, and others is not available from any other vendor and delivers much greater performance with significantly less hardware. Less hardware means less cost.

Trusted Partitions reduce Oracle Database, Oracle application, and other application license costs by providing the flexibility to license based on the cores used and not on the system's total capacity.

That's just below the line cost savings. PCA also increases above the line returns. It does this by accelerating time-to-market and enabling greater market share, revenues, and profits.

Conclusion

Converged Infrastructure (CI) and Hyperconverged Infrastructure (HCI) have been great boons to IT pros. However, they are far from perfect: no integration with applications and databases, high expertise requirements, minimal automation, disruptive manual patching, security holes that invite cyberattacks, and higher than expected costs are persistent issues for customers.

PCA X8 is the superior choice over Nutanix, Dell EMC VX Rail, or any other converged or hyperconverged infrastructure for the following:

- General mixed-use workloads.
- Extensive file or object-based workloads.
- Exchange™, SharePoint™, web services.
- Dev ops workloads.
- Design and simulation workloads.
- Heavy deep machine learning projects.

Oracle PCA X8 is the first system to address all of these issues and more. It creates a new category—Ultra-converged Infrastructure (UCI)—that obsoletes offerings from Nutanix, Dell, Cisco, HPE, and NetApp.

For More Information on Oracle PCA X8

Go to: Oracle PCA X8

Paper sponsored by Oracle. **About Dragon Slayer Consulting:** Marc Staimer, as President and CDS of the 22-year-old Dragon Slayer Consulting in Beaverton, OR, is well known for his in-depth and keen understanding of user problems, especially with storage, networking, applications, cloud services, data protection, and virtualization. Marc has published thousands of technology articles and tips from the user perspective for internationally renowned online trades including many of TechTarget's Searchxxx.com websites and Network Computing and GigaOM. Marc has additionally delivered hundreds of white papers, webinars, and seminars to many well-known industry giants such as: Brocade, Cisco, DELL, EMC, Emulex (Avago), HDS, HPE, LSI (Avago), Mellanox, NEC, NetApp, Oracle, QLogic, SanDisk, and Western Digital. He has additionally provided similar services to smaller, less well-known vendors/startups including: Asigra, Cloudtenna, Clustrix, Condusiv, DH2i, Diablo, FalconStor, Gridstore, ioFABRIC, Nexenta, Neuxpower, NetEx, NoviFlow, Pavilion Data, Permabit, Qumulo, SBDS, StorONE, Tegile, and many more. His speaking engagements are always well attended, often standing room only, because of the pragmatic, immediately useful information provided. Marc can be reached at marcstaimer@me.com, (503)-312-2167, in Beaverton OR, 97007.