Insight
64

The Incredible Shrinking Datacenter:

HOW HCI SIMPLIFIES IT AND LOWERS COSTS[1]

ABSTRACT

Early computers integrated computing, storage, and networking resources in a single system. As the need for capacity grew, these elements were disaggregated into separate fiefdoms within the IT infrastructure, making systems more capable but harder to manage. Over the past decade, a few pioneers have pursued the concept of re-integrating these resources into a single system that is more capable and easier to manage. This paper explores that concept, now known as Hyper-Converged Infrastructure (HCI), and demonstrates how it can make IT operations more agile while reducing overall expenses.

Table of Contents

Nathan Brookwood
Research Fellow, Insight 64

August 2020

The white paper was sponsored by AMD, but the contents and opinions expressed are solely those of Insight 64.

EXECUTIVE SUMMARY

IT organizations around the world struggle with the perpetual problem of trying to do more with less. For decades, practitioners could count on improved performance at constant or slowly falling prices to keep up with the workload without busting the budget. But recently, improved price/performance has ceased to work its traditional magic. Installed systems had become too complex and sclerotic. It would take a new approach to cut this Gordian knot. A few organizations have discovered that Hyper-Converged Infrastructures (HCI) can cut this knot, thus allowing them to continue their quest to deliver more for less.

HCI represents a complete rethinking of how organizations can acquire and deploy IT assets. For many organizations, it means replacing custom-configured systems that are expensive to buy and hard to maintain with scalable "IT appliances" that have been engineered and certified by their suppliers to perform as specified. If the organization needs more performance or higher availability, it just buys more appliances, and they all work together to deliver the services needed. Sounds too good to be true? According to IDC, HCI system sales accounted for about $7 billion in sales in 2020 and were growing at a rate of 8.3 per cent per year[2]. HCI is a real thing.

Until recently, most HCI deployments incorporated Intel Xeon processors. Not at all surprising, given that Intel's Xeon market share has been over 90 percent for most of the past decade. But AMD's 2017 reentry into the server segment with EPYC, and its 2019 refresh of the EPYC line with its 7nm 2nd Generation "Rome" offering offers users considering HCI new alternatives that cost less and perform better than the Intel alternatives. The AMD processors offer features like "Secure Encrypted Virtualization (SEV)" that Intel cannot currently match at any price. Most importantly, many HCI deployments need fewer cores and/or fewer CPUs to handle a given workload, and this can dramatically reduce the software license fees needed to implement an HCI system.

One of the best aspects of the competition between Intel and AMD in the server market is that almost all the software (operating systems, middleware, and application suites) that runs on Xeon also runs on EPYC. This makes it easy to benchmark the two environments and gather data that demonstrates the economic and performance advantages that AMD brings to this class of systems. HCI makes this even easier, since it is just a matter of moving an HCI stack from a vendor's Xeon-based system to that vendor's comparable EPYC based system. Don't you owe it to your organization to give this a try?

PROLOGUE: THE GENESIS OF HYPER-CONVERGED INFRASTRUCTURE

In the first decade the technology gods separated the data and the code and invented the mainframe. It was expensive to install, complicated to operate, and did not have much software. Then the technology gods created an IT department to feed and care for the mainframe. Corporate management thought it was too expensive, but the IT department solved more problems than it created, so they went along with the program.

In the second decade the technology gods invented the minicomputer. It was so affordable that any department could have their own system, and soon the corporate landscape was marked with many small computer islands, none of which could talk to each other or to the mainframe. Nevertheless, department managers could get solutions to problems without bothering the mainframe, and corporate management went along with the program.

In the third decade the technology gods invented the microprocessor, and the MPU begat the personal computer. Soon everyone in the company had a computer on their desk. Users developed their own databases which rarely agreed with one another. Chaos reigned and the IT gods were displeased, so they invented file servers to store and protect the company's data assets, and local area networks to connect personal computers to servers. IT spending increased, but employees were more productive, so corporate management went along with the program.

In the fourth decade the technology gods invented the Internet and the World Wide Web. Organizations needed an infrastructure to support exponentially increasing e-commerce and employee productivity requirements. Massive datacenters sprung up overnight on empty fields. The servers in these datacenters could not physically accommodate all the storage capacity needed for these new users and applications, so the technology gods invented disaggregated storage systems – Network Attached Storage (NAS) and Storage Area Networks (SAN) – that allowed hundreds of storage devices to operate with dozens of servers. Then the technology gods created specialized IT departments to manage these new storage networks. The datacenter kept ahead of end-user computing requirements, so management went along with the program.

In the fifth decade the technology gods invented the smartphone and third generation (3G) wireless networks that allowed billions of users to access the internet. As applications migrated from desktop and laptop computers to smartphones, the demands on network bandwidth and computing resources grew. Those demands would wax and wane during the day, complicating the tasks of network administrators often constrained by service level (SLA) agreements. The technology gods saw this problem and created incremental resources in "the cloud" that allowed datacenters to dynamically adjust capacity to meet demand. Management saw that this approach could save money and went along with the program.

In the sixth decade the technology gods invented faster (4G) wireless networks and streaming video services that once again taxed the existing infrastructure. But when datacenter operators tried to expand their operations, the complexity of the systems they had put in place got in their way. Their siloed implementations made their systems less responsive and harder to manage. Even simple changes often required the approval of specialists in storage, networking, and server departments. They had to spend more of their time tweaking their existing hardware and software, which left less time to address new features and applications. Nobody was happy.

At the start of the seventh decade the technology gods saw that the system architectures they had evolved over decades had become sclerotic, unresponsive, and unmanageable. Although the calendar said they were due for a sabbatical, they deferred their rest period and used the time to design a simpler architecture that would be easier to deploy and operate. All major elements would be virtualized to simplify resource allocation and location. All physical elements – processors, networking, and storage – would reside in the same box. System capacity could be increased by using more and faster processors, more and larger storage devices, and faster networking connections. Multiple boxes, each containing processing, networking, and storage, could be combined to accommodate larger or more diverse workloads. Everything would be managed from a single command console. They summarized their conclusions in an email with the subject line "Hyper-Converged Infrastructure" that they sent to the CTOs of all the leading hardware and software suppliers.

And only then did the technology gods take a well-deserved rest.

Chapter 1. What is HCI and How Does It Improve the Datacenter?

As we discussed in the Prologue, Hyper-Converged Infrastructure (HCI) simplifies the tasks of acquiring and operating the equipment that goes into a modern datacenter. Part of its "secret sauce" entails standardizing on a few configurations of industry-standard server hardware, but much of the benefit comes from virtualizing all aspects of the system. Applications run on virtual machines, which makes them easier to install and move as needed in response to changing conditions. Storage is directly attached to physical servers, and managed virtually, which eliminates the need for storage area networks with their attendant complexity. In theory, even system network operations can be virtualized, although most deployments these days have yet to include network virtualization in their implementations. One step at a time.

The key elements of HCI systems are industry-standard (x86) servers with directly attached storage systems and network interfaces. Major server suppliers like Dell, Hewlett-Packard (HPE), and Lenovo offer fully-configured systems optimized for specific HCI workloads, including Virtual Desktop Infrastructure (VDI), high-performance computing (HPC), business-critical applications, collaboration applications like Exchange and SharePoint, data management packages like SQL Server, and big-data analytics. Recently, AI and machine learning applications are also being added to the mix. Several hardware suppliers offer "HCI appliances" that include all the hardware and system software needed to implement an operational datacenter but restrict the customer to a limited number of options that the supplier has completely checked out and certified. Many suppliers also offer "al la carte" HCI hardware and software with a wider range of configuration options. Until recently, customers seeking to gain the benefits of HCI were constrained to use Intel Xeon® processors, but last year when AMD rolled out its second-generation AMD EPYC processor (formerly code named "Rome"), most of these suppliers added EPYC to their HCI systems. We will examine later how these new alternatives can lower acquisition (CapEx) and operating costs (OpEx), and even more importantly, software licensing costs.

Software is the key enabling technology for HCI systems. The success of the industry-standard server market is largely an outgrowth of the broad range of operating systems, middleware and applications that have migrated to the x86-64 environment over the last two decades[3]. These systems now account for 93 percent of the $88.6 billion worldwide server market[4]. Customers have invested millions in software based on Windows and/or Linux, and in virtual machine environments including VMware, Hyper-V, Openstack, KVM and proprietary hypervisors. The advanced virtualization capabilities included in leading HCI environments allow organizations to preserve software investments, while moving applications to environments that are easier to manage and provide higher levels of availability.

Contemporary virtualization solutions like VMware's vSphere and Microsoft's Azure Stack HCI abstract storage resources as well as processor and main memory. This means that physical storage attached to any node on an HCI cluster can be accessed by any application running on any VM in the cluster, as long as the application has permission to do so. As the storage needs of an app grow, it is relatively easy to add new storage capacity to a node that automagically goes into the pool of resources all VMs in the cluster can utilize. This means the IT staff can focus its efforts on new software features that add value, instead of overhead tasks like capacity planning that might be necessary in less automated environments.

The Cloud plays an increasingly large role in many organizations' IT operations, but the hybrid-cloud issue continues to challenge those organizations. HCI can play a role in managing these hybrid environments, since most HCI implementations are in essence miniature clouds. Microsoft was the first software supplier to address this issue, when it based the APIs it uses in the cloud on the ones it uses in Windows Server. This means applications written for Azure Stack HCI can move effortlessly to the Azure Cloud. A good idea like this is easy to copy, so a few years after Microsoft offered this capability, Nutanix came out with "Xi Cloud Services" that link on-prem Nutanix resources with cloud-based Nutanix resources running on AWS, Azure and Google Cloud platforms. Soon thereafter, VMware introduced VMware Cloud Foundation (VCF), a hybrid cloud platform that manages VMs and orchestrates containers that reside on-prem or in most popular clouds. Amazon forced its way into the hybrid segment with "AWS Outposts" which places an AWS-compatible rack of equipment on the customer's site. Google jumped in with Google Anthos, a software platform that supports existing and new application deployments across on- and off-premises systems. Hybrid clouds will be big, and HCI makes them easier.

Finally, today's sophisticated HCI environments simplify the task of managing hyper-converged environments with the use of dashboards that give administrators insight into all aspects system operation. VMware's vCenter, Microsoft's Windows Admin Center, and Nutanix's Prism all allow the datacenter manager to manage all aspects of the HCI cluster, including virtual machines, virtual storage, and networking, from a single control console with a contemporary graphic user interface. System administration has finally moved into the 21st century.

Chapter 2. Who's Who in the HCI Supplier Universe

The Software Players

Before marketers came up with the catchy "HCI" category, its underlying concepts were referred to as "software-defined infrastructure" (SDI) or "software-defined datacenter" (SDDC). Those labels made it abundantly clear that software plays a key role in building hyper-converged environments. There are many niche software players in the HCI software universe, but three companies – VMware, Nutanix and Microsoft -- account for the lion's share of the market.

VMware was started in 1998 by a small group of engineers who figured out how to efficiently virtualize the notoriously arcane x86 instruction set, thus allowing a single Pentium III processor

to host several virtual machines. Over the next two decades, the company added storage and network device virtualization to its bag of tricks, and those expanded capabilities drove its annual revenues from a less than $1M dollars to almost $11B in their most recent fiscal year. They must have been doing something right.

VMware's vSphere product line includes vSAN, a flash-optimized storage management system, vCenter Server, a centralized platform to manage all virtual resources, NSX Data Center, a network and security platform, and most recently the VMware Cloud Foundation (VCF) that provides management services for public and private clouds. Virtually all x86-based server suppliers resell some or all the vSphere product offerings, although Dell accounts for more than half of that volume.

Nutanix was started in 2009 by a few ex-Oracle managers who had the vision of bringing web-scale engineering – distributed systems running on commodity servers – to the masses. The technology companies like Facebook and Google use to deliver web-based services is scalable and highly cost-effective, but not something an ordinary IT department can deploy or manage. Nutanix set out to change that equation, and in the process became one of Silicon Valley's early Unicorns. It initially pursued Virtual Desktop Infrastructure (VDI) applications and then expanded its HCI software offering, now known as Acropolis, to support most of the IT applications ordinary companies need to function. Nutanix markets its products as "NX" turnkey appliances directly to end-users, but also sells its software through OEMs and channel partners who combine its software with their own hardware. The company's list of partners includes Dell, HPE, Inspur, Fujitsu, and Lenovo. Many, including Dell and HPE offer configurations that have been defined specifically for the Nutanix environment.

Nutanix differentiates itself from other HCI software suppliers by its flexibility regarding hypervisor environments. Originally the company used VMware's hypervisor, and most of its customers still use VMware to implement the virtual machine layer of their systems. Later it expanded its support to Hyper V since Microsoft bundled it with its Server operating system. Eventually the company developed a hypervisor of its own, dubbed the Acropolis Hypervisor (AHV), which it bundles at no additional charge with its other software. Nutanix gently pokes its HCI competitors by suggesting it has eliminated the "virtualization tax" that they use to boost their revenues.

Microsoft was founded in 1975 when Bill Gates dropped out of Harvard to write a Basic interpreter for the Altair 8800, the first personal computer. Today the company has 144,000 employees and is valued by the stock market at $1.5trillion. Although its early products focused on desktop operating systems and applications, it expanded its product line into server software, including operating systems, email servers, and enterprise databases, and now its Windows Server packages power more than 70 per cent of the enterprise server market. The company translated its expertise in building large scale data centers to support xBox gaming and Outlook email exchanges into a public cloud infrastructure now known as Azure. It was the first CSP to offer the same application interfaces (APIs) for software that could run in the cloud or on premises and dubbed these APIs and their supporting software "Azure Stack." The company added storage virtualization to the mix and branded the extension "Azure Stack HCI." This version has not been around long enough to make much of a dent with the folks like IDC and Gartner who assess market share data, but given Microsoft's presence in the datacenter, its likely to become a major contender in the HCI market.

The Hardware Players

The leading HCI hardware suppliers are (surprise!) also the top enterprise server suppliers – Dell and HPE. Both companies integrate key HCI software environments – Azure Stack from Microsoft, vSphere from VMware, and Acropolis from Nutanix – and market those HCI systems through their usual channels. HPE's HCI line also offers two proprietary products – SimpliVity and Nimble dHCI – that it acquired through acquisitions in 2017.

Dell owns 81 per cent of VMware, so it is not surprising the companies work closely in their technology and marketing programs. Dell goes to market with three different VMware-based HCI solutions. It brands its preconfigured HCI appliances as "Dell EMC VxRail." Customers plug the machines in, supply an IP address or two, and they are off and running. A simple and painless way to update the company's infrastructure to the latest technology. For customers for whom VxRail is not a perfect solution, Dell offers a line of VMware "ReadyNodes" that have been fully tested with VMware but are more configurable than the VxRail. Lastly, Dell sells vSphere components that can run on any server in Dell's catalog, or in mixed vendor environments, on an a la carte basis, but the customer needs to do some homework to ensure the chosen configurations make sense.

In addition to its VMware lines, Dell provides systems optimized for Microsoft's Azure Stack HCI. The company adds its Storage Spaces Direct (S2D) storage management software to Microsoft's Windows Server 2019 that includes the software-defined networking (SDN) facilities. Microsoft was ahead of the pack when it comes to on premises/cloud integration, so software written for Azure Stack HCI can move effortlessly to the cloud and vice versa, a feature that other cloud service providers (CSPs) have only recently adopted.

Dell's HCI portfolio also includes its PowerEdge XC family optimized to run the Nutanix HCI software suite. XC systems can be configured with a range of processors, main memory capacities and all-flash array capacities to suit a variety of workloads.

These systems position Dell as the leading hardware supplier in the HCI market. The company offers a range of systems optimized for specific HCI use cases, including VDI, big data analytics, and high-performance computing. Dell's product line includes five EPYC based systems optimized for HCI operation. A pair of single-socket systems come in 1U and 2U boxes (PowerEdge R6515 and R7515) that accommodate up to two terabytes of DRAM. A pair of dual-socket systems come in 1U and 2U boxes (PowerEdge R6525 and R7525) that accommodate up to four terabytes of

DRAM. A fifth system, the compute-optimized dual-socket PowerEdge C6525, accommodates up two terabytes of DRAM and 24 SAS/SATA/NVMe drives. EPYC's ability to support large DRAM configurations with 128 lanes of PCIe Gen 4 I/O enables single-socket EPYC systems in many cases to match and sometimes exceed the performance of dual-socket Xeon systems. This capability can have a substantial impact on software license fees and power consumption, as we will explore later in this paper.

Hewlett Packard Enterprise (HPE) was the first major server supplier to offer AMD based servers in 2004, and the two companies have had a long and productive relationship ever since. Today HPE is a close second to Dell in overall server market share, but a distant second to Dell in the HCI segment. This year, it has taken several steps to improve its market position and differentiate its offerings by combining HPE's latest EPYC based servers – the 1U DL325 and the 2U DL385 – with the capabilities it obtained in two 2017 acquisitions – SimpliVity and Nimble Storage Systems.

SimpliVity started in 2009, about the same time as Nutanix, and the two are regarded as HCI pioneers. Like Nutanix, SimpliVity initially used VMware as its virtualization layer, but has since added support for Hyper-V and KVM. The company initially focused on the VDI segment, where much of the data in each virtual machine (Windows libraries, application code, etc.) is replicated. SimpliVity developed a hardware accelerator, known as the OmniStack Accelerator Card, to off-load the calculations needed to handle data compression and deduplication. With the availability of more and faster cores on the latest AMD EPYC processors, SimpliVity was able to move these calculations to the CPU, thus lowering the system cost and freeing-up a system PCIe slot for other uses. With or without hardware acceleration, HPE offers a "SimpliVity Hyperguarantee" that simply and boldly states "If you use HPE SimpliVity hyperconverged infrastructure and its built-in VM-centric backup capability, you will achieve 90 per cent capacity savings across storage and backup combined, relative to comparable traditional solutions."

The EPYC/SimpliVity marriage, which HPE markets as the "HPE SimpliVity 325 appliance," allows HPE to double the number of virtual desktop machines it can support on a single system[5], driving down the per-user cost by 50 per cent, since there are no additional software license fees required. VDI has become an increasingly important workload for HCI deployments during the 2020 Pandemic.

Nimble Storage was an early leader in all-Flash arrays that were usually attached to host systems via Fibre Channel or iSCSI arrangements. This year, HPE married the Nimble array to its DL325 and DL385 systems and refers to the combination as "HPE Nimble Storage dHCI," a disaggregated HCI platform. This sounds like an oxymoron until you realize that the "I" in "HCI" is for Infrastructure, not Integration. Nimble has integrated its dHCI user interface into VMware's vCenter Control panel that masks the underlying complexity of managing the storage arrays.

In addition to its proprietary offerings, HPE markets its DLxxx server line with VMware, Nutanix or Azure Stack HCI software environments. Many of these HCI offerings (SimpliVity, Nimble, VMware, Nutanix and Azure Stack HCI) can be acquired via HPE's innovative "Green Lake" marketing program, which is a really big deal, although HPE has not done nearly enough to educate the market about its capabilities[6]. Simply stated, Green Lake lets HPE own and operate the customer's datacenters and its customers pay for services rendered in those centers as they are incurred, in much the same way they pay for cloud services. It basically transforms the customer's existing fixed-cost IT infrastructure into a pay-as-you-go SaaS model.

Chapter 3. How AMD EPYC Processors Make HCI Even Better

Until recently, IT customers investing in an HCI infrastructure had only three processors from which to choose. Did they want to pay a steep premium for Xeon Platinum, a slightly smaller premium for Gold, or just go with the bargain-basement Silver edition. The introduction of AMD's EPYC server processor line not only gave these buyers a new alternative but has also encouraged Intel to lower its Xeon prices.

Although Intel can adjust its prices, its archaic single-chip CPU implementations are no match for the chiplet design AMD uses in its EPYC processors. With its 2nd Gen EPYC processors, AMD packs eight CPU cores onto a 7nm "Compute Core Die (CCD)" that measures only 74 mm2 and contains 3.9 billion transistors. It puts 2, 4, 6, or 8 of these CCDs, along with a single 14nm "Input/Output Die (IOD)" that that measures only 416 mm2 and contains 8.34 billion transistors. The IOD has Infinity Fabric™ links to all the CCDs and handles memory and I/O functions for the complex. AMD packs 24 billion transistors into a 32-core 2nd Gen EPYC processor that uses about 712 mm2 of silicon, but those CCD chips are easily manufactured with today's technology. Intel's latest 14nm "Cascade Lake" Xeons incorporate a monolithic die that measures 698 mm2 but contains only 8 billion transistors. Chips that big are difficult to manufacture, and you cannot fit as many on a typical silicon wafer, compounding the problem. 32-core 2nd Gen EPYC CPUs have 16 billion more transistors than Intel's 28 core Xeon, and most of those transistors are used in the massive caches on the 2nd Gen EPYC chip. The comparison is even more lopsided when one looks at the 64-core 2nd Gen EPYC chip which has 31 billion more transistors than the Intel Xeon[7].

Huge caches are one of the keys to EPYC's extraordinary performance in virtual environments. In HCI systems almost all the code operates in virtual machines. 32-core EPYC CPUs typically contain 128MB of cache, and 64-core versions include 256MB. Double those numbers in two-socket configurations. That is a lot of cache. By way of contrast, Intel's top of the line Xeon Platinum 8280 processor only has 38.5MB of cache. Modern processors can zip through code at amazing speeds when their code and data reside in the cache. But when the cache misses and the CPU must go to main memory, things slow to a crawl until that data arrives in the cache. As CPU speed increases, cache misses exact a bigger and bigger toll, as Gene Amdahl predicted 50 years ago. AMD's architects and engineers clearly paid close attention to their professors at MIT, Cal Tech and UT Austin[8].

Although most programs love caches, some codes benefit less than others. EPYC architecture contains features that enhance these programs as well. Computer architects talk about the need "to feed the beast;" i.e., move data from I/O devices like Flash memory arrays and 10 gigabit Ethernet networks into the system's main memory, and from there, into the chip's caches. EPYC is one of the first commercial systems to support PCIe Gen 4, the latest version of the venerable PCIe I/O architecture that is twice as fast as it Gen 3 predecessor – 16 gigatransfers/second, or 32 GB/second for 16-bit interfaces. More than enough for the most demanding I/O loads. Those platinum, gold and silver systems are still stuck on Gen 3.

Virtual machines run in virtual memory, but at some point, virtual memory must be mapped into the system's physical memory. The more physical memory a system has, the more VMs it can support. Most single socket EPYC systems can address two terabytes of main memory, and most dual-socket systems can accommodate four terabytes. This means physical memory is rarely a constraint on the number of VMs in a system. Those fancy metallic-branded systems usually max out at one terabyte.

So, what limits the number of virtual machines a system can support? Sometimes it is just raw processing power. Too much work, too few CPU cycles. EPYC addresses this issue by packing more CPU cores onto a chip. Up to 64 cores and/or 128 threads on the high-end models. Those 64-core processors still have four-digit price tags, unlike the ones with shinier badges.

When all is said and done, the most expensive line on your HCI budget will likely be the software fees that make all this possible. The "hyper" in "Hyper-Converged" refers to the need for a hypervisor to create and manage virtual resources in the system. Which hypervisor you use may depend on the rest of your software environment, although it is often possible to mix-and-match hypervisors to align with your organization's preferences. A few hypervisors (like KVM) are offered in open-source form, but most are licensed software products from public companies like VMware, Nutanix and Microsoft. These companies need to generate a return for their shareholders and calibrate their price structure to maximize revenues. Buyers, in turn, seek to minimize acquisition costs, and must consider license fees as well as hardware in their economic analyses. It was a lot easier in the days when ISVs had a set "per socket" fee but now the industry is moving to "per core" pricing.

This change in ISV pricing strategies may impact the way your organization approaches its virtualization and HCI strategy. Instead of "the more cores the merrier," you might want to consider how to get more work done with fewer cores. Fewer cores, lower software fees. To execute this strategy, you will want to deploy on systems with the highest performance per core. AMD saw this change coming and adapted its EPYC offering to optimize your software license fees. In April, the company launched a new line of high frequency 2nd Gen EPYC Processors, adroitly named the "7FX2" series. These new versions trade off higher core counts for higher performance per core. In April, an HCI cluster consisting of four Dell PowerEdge R6525 servers, each with two EPYC 7F72 processors and 1GB of DRAM, scored 13.7 @ 14 tiles on the VMmark 3.1 benchmark. The EPYC result was 47.4 per cent higher than that achieved on comparably configured cluster with Intel Xeon Platinum 8276L processors[9].

For some customers, the switch from per socket to per core pricing may have little or no impact. The change obviously does not impact those who use open source software, which has no license fees. It also will have minimal impact on VMware customers, since that company has priced a 32-core license close to the previous charge for a single-socket license. This means your software fees will only increase if you are using EPYC processors with more than 32 cores; 48-core and 64-core systems will need two 32-core VMware licenses. Even with the need for an additional 32-core license, a single-socket 64-core system will likely still cost less to acquire and operate than a dual-socket system with two 32-core CPUs.

With all those VMs filling up terabytes of virtual memory, any responsible IT executive with just a small tinge of paranoia should be worried about how the organization can ensure that a malicious user could compromise system security and access data that should be off-limits. Memories of Spectre and Meltdown still linger[10]. Hardware that helps keep the bad guys out is always a good idea. AMD's security architects are as paranoid as any on the planet, and they conceived a way bad guys could create VMs that might initially contain data left over from an earlier VM instance. AMD invented "Secure Encrypted Virtualization (SEV)" to guard against such attacks. With SEV, code and data associated with a VM is encrypted for that VM and may only be decrypted from within the VM. Even the hypervisor cannot see the decrypted data. Hypervisor vendors like VMware need to adapt their software to support SEV. VMware has stated they will add this capability, but the timing of their release has yet to be determined. Take that, bad guys!

Chapter 4. Is HCI Right for Your Organization?

In earlier parts of this paper we discussed how HCI works and we explored a few of the more popular vendor offerings that might be relevant as you pursue your HCI journey. We focused primarily on the technical aspects, but you may also want to view HCI through financial or organizational lenses as well. In this chapter, we address those issues.

If your IT organization has no budgetary constraints, your end-users are happy with the services you deliver, and your company is insulated from competitors using technology to gain market share at your expense, then you might not need to investigate HCI as a means of delivering IT services. Unfortunately, very few companies find themselves in this strategic Utopia. On the other hand, if your IT budget is shrinking, and most of it goes to holding existing applications together with baling wire and duct tape, then it may be the time to start thinking about bringing your datacenter into the 21st century. There are many paths to datacenter modernization, but HCI should be on the list of items to explore.

HCI can help you eliminate technology silos, which is generally regarded as a good thing. But you may have groups within the IT department who have specialized in the care and feeding of SANs or NAS farms, and those roles may no longer be needed in a hyper-converged environment. The same applies to your networking gurus. Some of these employees may welcome the arrival of easier to deploy and administer systems, but some old-school workers may resist. In fact, most humans tend to resist change. You may need to allocate retraining expenses into your plan.

HCI has become a popular way to build Virtual Desktop environments, which make even more sense now than they did before the Covid-19 pandemic forced many employees to work from home. Some employees even like it! VDI allows you to keep your organization's data assets under lock and key while your employees can access virtually the same desktop/laptop environment they use in the office. HPE and Dell have armies of sales and support people who can guide you through this transition. The longer-term impact of WFH may ripple through commercial and residential real estate markets for years to come, so you better make sure your IT infrastructure can easily adapt to all these changes.

HCI may allow you to aggregate your processing requirements over a smaller number of physical servers. This can reduce the footprint of your datacenter, the electrical power needed to operate those systems, and the electrical power needed to cool all that equipment.

One of the first decisions you will face in your HCI journey is whether to fit your requirements into an HCI appliance from your preferred vendor or work with that vendor's a la carte HCI options. Most suppliers offer pre-configured standard nodes that include all the hardware and software elements you need. You may find that these appliance bundles also reduce your software license expenses. The appliance approach is the path of least resistance if you can live with the assumptions the appliance supplier made when defining that appliance.As we noted above, software licenses are likely to be one of the biggest line items on your purchase order, but there are ways you can minimize these costs. Even with VMware's recent switch from "per socket" to "per core" pricing, it costs less to run VMware on one 32-core server as opposed to two 16-core servers. Single-socket EPYC powered servers with 2 TB of DRAM and 128 lanes of PCIe Gen4 I/O can often outperform dual-socket 1 TB Xeon systems, and in this instance you can halve your software license expense by this approach. Those software savings can also apply to other expensive packages like SQL server, business intelligence and HR management. You will likely also save on the power used to run those systems.

Conclusion:

In this document we have explored a new way to structure your IT operations that can simplify your computer operations, make your systems more responsive and your organization more agile, and potentially lower the amount you spend accomplishing these tasks.

If it has been more than four or five years since you last revisited your IT strategy, it is past time to review it now. A lot of water has flowed under the bridge. The cloud has become a force unto itself, and you need to find a way to harness it for your organization. The HCI technology we have discussed here provides a simple way to recreate your IT environment so that you can blend on premises and cloud resources in new and exciting ways. Whether you end up pursuing an HCI approach or not, you will be better equipped to manage your company's IT infrastructure when you are done.

Additional resources:

AMD offers an interactive Virtualization TCO Estimation tool that lets you compare the costs of using EPYC processors with comparable servers using your choice of 1st or 2nd generation Intel Xeon Scalable Processor. Surprise! AMD almost always saves you money.
www.amd.com/en/processors/epyc-tools

AMD keeps track of the world records its EPYC processors have won at
https://www.amd.com/en/processors/epyc-world-records

A simple demo of the way AMD's Secure Virtualization (SEV-ES) keeps prying eyes away from data inside a VM can be found at
https://www.amd.com/en/processors/epyc-world-records

The folks at IT Central Station (https://www.itcentralstation.com) post reviews and user comments from knowledgeable IT practitioners. They list many comments relative to VMware, Nutanix and Azure Stack HCI. They also publish useful buyers guides that cover many HCI product offerings.

Patrick Kennedy at ServeTheHome (STH) has an excellent video that explains the "frequency optimized" segment of the server market that AMD entered with the recent launch of is EPYC 7FX2 product line.
https://www.servethehome.com/amd-epyc-7f52-benchmarks-review-and-market-perspective/

Patrick also explores the changes VMware made to its licensing policies at
https://www.servethehome.com/licenseageddon-rages-as-vmware-overhauls-per-socket-licensing/

Hassan Mujtaba at wccftech.com has an excellent article on the physical implementation of AMD's Rome processor. I used the data he posted there as the basis for my analysis of die sizes and transistor counts for the various chiplets used in "Rome".
https://wccftech.com/amd-2nd-gen-epyc-rome-iod-ccd-chipshots-39-billion-transistors/

AMD HCI Landing page:
https://www.amd.com/en/processors/epyc-for-hyperconverged-infrastructure

---

1. This white paper was sponsored by AMD, but the contents and opinions expressed are solely those of Insight 64.
2. IDC Press Release, June 18, 2020
3. It is a little-known fact that operating systems, compilers, and other system software embed the code "AMD64" in their source code to identify specific developer options that target the x86-64 environment. AMD was responsible for the technical innovations that extended the 32-bit x86 architecture to 64 bits at a time when Intel was still pursuing its "Itanium is the best 64-bit architecture" strategy.
4. IDC Press Release, March 27, 2020
5. HPE Press Release, May 7, 2020. The claim is based on validation by LoginVSI benchmark testing
6. There is an apocryphal story about an HP marketer who failed when he described "sushi" as "cold raw fish."
7. To be completely fair, Intel does offer a 56 core Xeon Platinum 92X2 series that puts two monolithic 698 mm2 CPU dies in a multichip package that must be soldered onto the motherboard. The combined package measures 1400 mm2 and contains "only" 16 billion transistors. It also consumes 400 watts of power, which mandates a liquid cooling thermal solution. The company has not listed prices for these chips, but industry analysts speculate prices start at $25,000. Not surprisingly, there have not been too many takers.
8. It also helps that AMD's CEO Lisa Su has three degrees from MIT. I was lucky to get one.
9. www.amd.com/system/files/documents/amd-epyc-7Fx2-dell-vmmarkvsan-dualsocket-perfbrief.pdf
10. AMD's checks for speculative execution memory requests have made EPYC less vulnerable to-date than Xeon in certain "side channel" attacks. Most processors introduced since early 2019 have included changes that eliminate or at least minimize susceptibility to Spectre and Meltdown, but hackers are always looking for new ways to compromise systems. PID# 20613195-A