

The Future of Cloud-Scale Application Infrastructure

Strategies for Building Cloud-Scale Datacenters Based on Lessons Learned from HPC

By Frederic Van Haren
CTO & Founder of HighFens, Inc.

Software as a Service (SaaS) and cloud-scale applications bring new requirements for datacenter infrastructure strategies. But where have we seen this before? The evolution of High-Performance Computing (HPC) offers lessons that provide a framework for projecting the fast-evolving needs of SaaS and cloud-scale application infrastructure.

This whitepaper compares the evolution of HPC with that of modern SaaS and other cloud-based software to better understand the technologies and architectures that will dominate the datacenter in the age of digital transformation. Everything old is new again.

- 2** | Introductions
- 3** | Learning from HPC Infrastructure Strategies
- 4** | Emerging Technologies: The Rise of NVMe and NVMe Over Fabrics
- 5** | Envisioning the Future of As-a-Service Infrastructure
- 6** | Conclusions
- 6** | About the Author

Introduction

Modern application architectures are intersecting with High-Performance Computing (HPC) paradigms, and the evolution of HPC can serve as a guidepost for envisioning the future of cloud-scale application infrastructure strategies.

The cloud is rapidly transforming the software industry. The growth of SaaS and cloud software is anticipated to dominate the growth of traditional software over the next several years. Virtually all independent software vendors are delivering—or are planning to deliver—cloud-based versions of their offerings. As the SaaS business model has gained momentum, new “as-a-service” datacenter infrastructure technologies and application delivery models are advancing rapidly. In addition to enterprise SaaS offerings such as CRM, ERP, HR, financials, etc., there are many other examples of cloud-scale applications that are delivered using the same datacenter strategies.

For example, customer-facing online services are advancing rapidly in many industries, including financial services (e.g. online trading, online banking, and financial data-as-a-service), healthcare (e.g. insurance claims, pharmacy, patient portals, practice management, and electronic health records), retail (e.g. eCommerce), and travel (e.g. booking services and travel recommendations). The explosive, global growth of SaaS and online services is leading to major changes in enterprise infrastructure, with new application development methodologies, new database solutions, new infrastructure hardware and software technologies, and new datacenter management paradigms.

These new approaches for delivering applications to end users stem from the need to support the specific needs of as-a-service infrastructures. And this growth will only accelerate as emerging Internet of Things (IoT)-enabled technologies like connected health, smart industry, and smart city solutions come online in the form of as-a-service businesses.

The Modern Cloud-Scale Datacenter

The modern cloud-scale application is characterized by high volumes of data that are streaming in from numerous sources and a significant dependency on sophisticated, real-time analytics. SaaS applications are architected to deliver functionality to a large number of users from a large number of organizations. This multi-tenant, as-a-service cloud delivery model is architected for cost efficiency, resource optimization, and ease of management. The nature of many cloud-scale applications mandates an emphasis on security, availability, and SLA-protected performance that supports high levels of user experiences.

Scale is a major challenge for applications running over multi-tenant infrastructure. Think about it: SaaS solutions have to support many users—accessing shared infrastructure—using multiple types of devices—from different geographical locations. This requirement for ever-increasing scalability demands new architectures that can meet increased demands for SaaS offerings. But multi-tenancy, security, and scale are not the only characteristics that are challenging SaaS providers; cloud-scale applications require the swift processing of increasing data volumes using high-performance analytics. For example, eCommerce applications require the ability to deliver personalized shopping experiences based on real-time customer behaviors.

What Can We Learn from HPC?

HPC paradigms were born from the need to apply sophisticated analytics to large volumes of data gathered from multiple sources. With the advent of cloud technologies, HPC applications have leveraged storage and processing delivered from shared, multi-tenant infrastructure. HPC technologies and datacenter architecture strategies have evolved substantially over the past few decades.

Many of the same challenges addressed by HPC practitioners are now facing modern SaaS applications. The future of SaaS infrastructure can, therefore, be effectively informed by comparing the evolution of HPC infrastructure to the future evolutionary requirements of SaaS technologies, applications, and solutions. In other words, as you try to envision the future requirements for SaaS infrastructure, it is helpful to draw the parallels to the historical evolution of HPC to gain a longer-term perspective.

Learning from HPC Infrastructure Strategies

HPC was started over 60 years ago, primarily to serve academic and industrial research. Organizations that needed to handle large and complex computational problems that couldn't be handled by general-purpose computers turned to HPC—if they could afford it. Only a handful of vendors supplied the market with specialized supercomputers, often delivered as an all-in-one solution.

A traditional HPC environment consists of storage, network, and compute resources as well as software to logically glue these elements together. It is the combination of the components that makes up the overall system performance. Increasing the performance or the ability to process more workloads within a certain timeframe means adding more components until the required performance is obtained. It is also important that the components are properly sized and adjusted to efficiently leverage the other components. For example, it doesn't make sense to add more storage for performance if the network won't be able to handle it, because the network will become the bottleneck and hold back the other components.

The Move to the Cloud

Infrastructure for HPC has historically been located on premises, with all system components connected over a local network to ensure the best performance and simplify management of the HPC environment. But cloud computing is increasingly utilized as a viable option for running certain HPC applications on public or private cloud infrastructure. By offloading processing and storage to cloud infrastructure, organizations have been able to manage the exponentially greater volumes of data while serving users connecting remotely from a variety of devices.

A More Data-Centric Approach

While the HPC architectures were evolving and maturing at a steady state, the need for solving more complex problems led toward more data-centric system architectures able to move larger amounts of data. New tools and architectures had to be created to process unstructured data. The rise of Hadoop, MapReduce, and Spark helped achieve the more data-centric approach by bringing data as close as possible to where the processing happens. Hadoop, for example, relies on its own Hadoop Distributed File System (HDFS) to split the large blocks of data into smaller units so they can be distributed to the nodes in the cluster. Small portions of code are transferred to the nodes to process the data in parallel, thus creating data locality. The underlying concept is that processing data local to a node is more efficient and faster than it would be with a traditional HPC system.

The Shift to Open Source Software & Commodity Hardware

The new approach driven by the open source and research communities have driven the shift away from expensive hardware to lower-cost, commodity hardware. This trend opened up the HPC market to use “white box” infrastructure to control costs. Movement of HPC infrastructure to the cloud, the need to analyze unstructured data, and the shift to open source software running on commodity hardware platforms gave organizations a greater choice in HPC infrastructure adoption. This is allowing them greater selection opportunities for HPC resources based on a careful analysis of cost, performance, and the many other factors that influenced HPC investments.

Advances in Big Data

Meanwhile, Big Data tools were gaining in popularity and vendors started to provide professional services. As companies began to implement and deploy Big Data storage and server environments, they realized the need to extract better economic value from their data. Organizations began integrating discovery techniques such as machine learning and deep neural networks into their HPC workflows. This has led to the realization that implementing an emerging technology such as machine learning provides the benefit of merging both worlds. Organization could merge the performance advantages of HPC with the increased usability and flexibility of Big Data running on commodity hardware platforms. Instead of reinventing the wheel, the HPC and Big Data compute-intensive paradigms are now coming together to provide organizations with the best of both worlds.

Overcoming the Limitations of North-South Architectures

A traditional HPC system consists of three independent hardware layers, each of which has its own performance, latency, and capacity criteria:

- Storage
- Network
- Compute

The system design is designed in a North-South architecture, with each layer proportionally sized. For example, the compute layer needs a certain amount of storage, and the network has to be able to transport the data back and forth to the storage devices without becoming a bottleneck. As the system grows, each of the layers has to grow proportionally. When the growth is linear, this type of architecture works well. But if the growth is dominated by one layer, it can be costly because the other layers must also grow proportionally.

Emerging Technologies: The Rise of NVMe and NVMe Over Fabrics

Emerging technologies are allowing infrastructure architects to overcome the limitations of North-South architectures and scale datacenter resources more cost effectively.

Non-Volatile Memory Express (NVMe)

The performance advantages of Solid State Drives (SSDs) are leading organizations to replace Hard Disk Drives (HDDs) with SSDs at an astonishing rate. SSD-based storage arrays have leveraged the same interconnect technologies used for HDD systems, such as Serial ATA (SATA), Serial Attached SCSI (SAS), or Fibre Channel. However, SSDs are effectively memory chips bundled in the same form factor for backward compatibility. The reality is that those memory chips easily outperform the interconnects. So if SATA, SAS, and Fibre Channel are not the appropriate interconnects, then what is?

The industry has been integrating memory chips on a device that plugs right into a PCI Express (PCIe) bus that operates at a higher bandwidth and extends the memory chips to their full capabilities. As a logical device interface, NVMe has been designed from the ground up to capitalize on the low latency and internal parallelism of flash-based storage devices. It can provide high volumes of input/output operations per second (IOPS) per device with a reasonable amount of capacity, and a server can have multiple NVMe cards in a single chassis to deliver an unprecedented amount of IOPS. NVMe has increased the practicality of accessing fast storage in the datacenter by solving the problem of scaling non-volatile memory in a system.

NVMe Over Fabrics (NVMeF)

NVMe over Fabrics is a specification that extends the benefits of NVMe to larger fabrics beyond the reach and scalability of PCIe. It enables NVMe message-based commands to transfer data between a host computer and a target SSD over a network such as Ethernet, Fibre Channel, or InfiniBand. An end-to-end NVMe solution reduces access latency and improves performance, particularly when paired with a low-latency, high-efficiency transport such as Remote Direct Memory Access (RDMA).

This allows applications to achieve fast storage response times whether the SSDs are attached locally or accessed remotely across enterprise or datacenter networks. For example, Intel and Micron have been working jointly on a disruptive non-volatile technology called 3D XPoint. They claim performance with read latencies at less than 10 microseconds and writes at less than 20 microseconds, and Intel has hinted that its 3D XPoint products could be offered at prices comparable to current flash memory products.

Implementing an East-West Architecture

New interconnect technologies make it possible to implement “East-West Architectures” that bring the data closer to the compute resources and as much as possible avoid having it go through the other layers. For example, you can connect the compute nodes with each other through NVMe Fabrics over Ethernet to eliminate traditional bottlenecks. The compute nodes are then configured in an East-West architecture, with the nodes exchanging data directly with each other.

An East-West architecture does not replace but instead complements a North-South architecture. The concept is that the primary copy of the data is still in the above-mentioned North-South storage layer, with a secondary copy at the compute node level in an East-West architecture. Keeping a working subset of the data as close as possible to the compute nodes enables faster access and lower latency while also solving the scalability problem. By moving to an East-West architecture and connecting storage and server resources over extremely fast connections, you can achieve the same performance as traditional North-West architectures through a far more flexible infrastructure layout.

Envisioning the Future of As-a-Service Infrastructure

The evolution of HPC infrastructure over time has largely been driven by increasing the power of analytics, supporting a more scalable infrastructure and decreasing infrastructure costs. Sound familiar? These are the same issues facing cloud-scale applications like SaaS today. As modern as-a-service infrastructures continue to grow in scale while delivering increasingly sophisticated analytics, we will see a move toward new architectures that look a lot like the East-West architectures employed by modern HPC applications.

Datacenter architects will move toward the efficiency of independently scaling compute and storage resources to deliver the required performance. The continued advancement of commodity hardware platforms will lead to continued price-performance improvements for storage, processor, and networking technologies. The East-West architecture framework, supported by flexible software-defined datacenter technologies, will create new levels of agility for modern datacenters.

We increasingly see new and existing datacenter technology vendors talking about highly scalable platforms leveraging underlying technologies like NVMe and NVMeF. All are aiming to solve limitations of current enterprise datacenter architecture to address the needs of the as-a-service world. In addition, leading analyst firms are beginning to cover new datacenter infrastructure technologies that will play a role in transforming the modern datacenter.

Shared Accelerated Storage

Gartner recently added a new technology category to their storage taxonomy termed “Shared Accelerated Storage.” This category encompasses approaches that disaggregate compute and storage architectures that implement flash memory centrally so it can be accessed by multiple servers connected via high-speed networks. According to Gartner:

“*Shared accelerated storage is an architecture that has been designed to tackle new data-intensive workloads by bringing high-performance and high-density next-generation solid-state-shared storage to the compute servers over a low-latency network. This technology delivers benefits of shared storage with performance of server-side flash by leveraging standardized NVMe low-latency technology.*”

Gartner

Composable/Disaggregated Infrastructure and Rack Scale Architectures

According to IDC:

“ *Composable/disaggregated infrastructure systems are an emerging category of infrastructure systems that seek to aggregate compute, storage, and networking fabric resources into shared resource pools that can be available for on-demand allocation.* ”



This is an evolution of converged and hyper-converged infrastructure, with hardware becoming disaggregated and software leveraging composability via a provisioning, orchestration and automation layer.

Conclusions

This is an exciting and challenging time for SaaS and other as-a-service providers. Technology is changing rapidly, and architectural shifts and the increasing use of commodity hardware are offering tremendous opportunities. But the path to managing increasing volumes of data while improving analytics can be a difficult one to walk because it forces architectural decisions that will significantly impact the cost model of your fast-growing infrastructure.

As you proceed on this journey, keep in mind the many similarities SaaS providers now have with organizations who have embraced HPC and have overcome these challenges over time. Learn from their experiences so you can avoid making strategic errors and accelerate your abilities to control costs while increasing your analytics, creating new revenue opportunities and the ability to build cloud-scale datacenters based on lessons learned from the evolution of HPC infrastructure.

About the Author



Frederic Van Haren is CTO of HighFens, Inc., a consulting firm he founded in 2016. With more than twenty years of experience in high tech with a focus on highly scalable solutions and data management, Frederic offers clients expertise and knowledge in the areas of storage, networking and computing. For over fifteen years he served as Senior Director of R&D Labs for Nuance Communications and previously served as a Director with Lernout & Hauspie and as a software developer for Alcatel-Lucent (now Nokia).

Frederic Van Haren