Developing a Cloud-Like Experience for Mission-Critical Data

Storage Switzerland, LLC

kaminario

Cloud-First is a common mantra, popular among organizations of all sizes. The challenge is that many enterprises are discovering that it may not be the right mantra, at least not for mission critical workloads. While certain software as a service (SaaS), data protection, and file collaboration workloads seem like a good fit for the cloud use case, mission-critical workloads continue to need dedicated, deterministic infrastructures.

The challenge is that once organizations experience the cloud's incrementally scalable infrastructure and its "pay as you go" business model, they want to create that same experience for their on-premises mission-critical applications while maintaining the consistent high-performance they require.

1 The Requirements of a Cloud-Like Experience for On-Premises Mission Critical Workloads

The cloud provides computing, networking, and storage resources on a pay-as-you-go basis. IT resources become an operational expense. Most on-premises datacenter resources need to be purchased upfront. IT resources are a capital expense. The challenge with purchasing hardware and software upfront is that in most cases IT needs to buy with the end result in mind. They need to predict to what extent the workloads using these resources will grow and make sure that the systems purchased support that growth.

Scale-out and later scale-right architectures alleviated some of the concerns over a CAPEX business model. In a scale-out architecture, organizations can add nodes as the workload demands for performance and as capacity expands. Each expansion unit is a server, called a node, with storage capacity, storage performance, networking, and computing capabilities. The challenge with scale-out is that not all of the resources of a node are used at the same pace and eventually the scale-out system becomes out-of-balance. The out-of-balance resource problem was solved by a scale-right architecture, where each node could have additional resources added to it, capacity for example, so that resources were consumed more completely before new nodes were added.

The next step in the evolution from scale-up, scale-out, and scale-right architectures is composable architectures where virtual arrays are created from discrete components available in the cluster, enabling a fine-grained use of resources for maximum efficiency. A composable architecture enables IT to respond dynamically to the needs of the organization's workloads and deliver the exact performance and capacity required by each.

The other appealing component of the cloud is its business model. While scale-out, scale-right, and composable architectures provide the technological flexibility of the cloud, they lack its consumption model. Cloud resources can be purchased at the scale desired, expanded in lock step with

the workloads and then returned if the need passes. Most on-premises workloads, regardless of the architecture, require an upfront purchase for the full amount of the system. Also, each of these requires another full purchase of additional nodes as the environment scales.

These business models can be created on-premises but need to be more than a technology lease. Organizations looking to create a cloud-like experience for their mission-critical workloads need to be able to consume resources incrementally as needed. Additionally, they should have the option to have software subscriptions separate from hardware purchases – to deliver further flexibility, both from a technical and economic perspective. This enables them to change hardware configurations as required but not have to re-buy software. In the classic CAPEX model, each expansion requires the purchase of new software as most licenses are not transferable from one version of hardware to another.

If done correctly the vendors providing on-premises solutions have a distinct advantage. Cloud providers essentially have only one option. On-premises vendors that adapt their solution and partner correctly can provide many more options.

They can continue to provide a more classic CAPEX business approach to organizations that want it. Ideally, this appliance approach still disaggregates hardware and software. Vendors can also offer a subscription model for the software so that hardware and software purchases can be further disaggregated. Third, the vendor can offer a utility type of solution where hardware and software are bundled and charged for based on monthly consumption. Lastly, the vendor can also provide an “as a service” model which the customer can leverage with an existing model for workloads such as disaster recovery, or disaster recovery as a service (DRaaS), and dev/test.

Conclusion

Two big attractions of cloud computing are flexibility and its OPEX business model but they lack the deterministic, performance, and high security of on-premises solutions. On-premises solutions can benefit from a cloud-like model but vendors have to do more than offer a cloud model that is really a lease in disguise. They need to make technological advances in how their solutions allocate resources and how that solution can be consumed.

2 Alternatives to the Public Cloud

As previously stated, cloud computing is commonly heralded as the future, in large part due to the ability to procure and to pay for IT infrastructure resources on demand via cloud. The problem is that the cloud does not provide the levels of control or performance that are required by mission-critical applications. In this chapter, we will explore in more detail the merits and drawbacks of typical on-premises alternatives to the public cloud.

Traditional Scale-Up Storage

Scale-up storage architectures are the longest standing. These arrays are configured with a fixed number of controllers (typically two). Additional resources, such as disk drives, may be added to the system as capacity requirements increase. Typically, scale-up architectures are ideal for serving point workloads with fairly predictable requirements. They are also typically capable of delivering the fastest levels of performance of per volume performance on fixed workloads. The problem is that workloads are becoming more dynamic and unpredictable, and that the performance of scale-up systems tends to degrade as more capacity is added, because the controller compute resources are not upgradeable and become overwhelmed with Input/Output (IO).

Scale-up storage systems are limited by the maximum levels of drives and performance that the storage controllers can handle. If a system's limits are exceeded, the controllers either need to be upgraded (if the vendor's architecture allows) or an entirely new system must be added for additional compute and storage. Silos of data are created, and multiple systems must be purchased, racked, stacked, and managed. Additionally, the balance of compute and storage that was once ideal for the specific workload can quickly become imbalanced. Resources can quickly become over, or under, provisioned.

Scale-Out Storage

Scale-out architectures emerged to address some of these pain points. In a scale-out approach, resources (including storage capacity) across multiple systems are aggregated into a centralized, shared pool. The pool of resources is no longer bound by the constraints of a singular system. A new system may be deployed to expand capacity or performance. This creates the opportunity for more flexibility, better system utilization, simplified management, and a stronger ability to run multiple workloads in parallel (a capability that is increasingly becoming a requirement for modern businesses).

The problem with the scale-out approach is that a minimum configuration of three nodes is typically required, which may be more than a particular project requires for a long time (if ever). In a scale-out approach, servers are clustered together so that compute, storage capacity, and networking may be pooled. These clusters are typically called *nodes*. Additionally, performance is often impacted by the heavy levels of communication between systems over the network that are required to keep systems in sync and to meet data protection service level agreements. Scale-out storage systems require parallelism to deliver the performance they promise: either through many workloads using the system or workloads that are innately parallel.

Composable Storage

The concept of composable storage is developing to further boost infrastructure agility and utilization through even more granular and fluid resource allocation. Composable storage virtualizes storage resources into a shared pool, from which virtual storage resources may be created on demand and automatically, according to application or workload requirements. Composable storage stands to bring the enterprise closer to the resource agility that is inherent in the public cloud.

Conclusion

The problem with scale-up, scale-out, and composable storage approaches alike when it comes to facilitating a cloud-like model on premises, is that infrastructure still needs to be procured upfront. Composable storage brings the enterprise closer to a highly utilized and agile infrastructure, but storage professionals should look further for the opportunity to pay on a consumption-based approach.

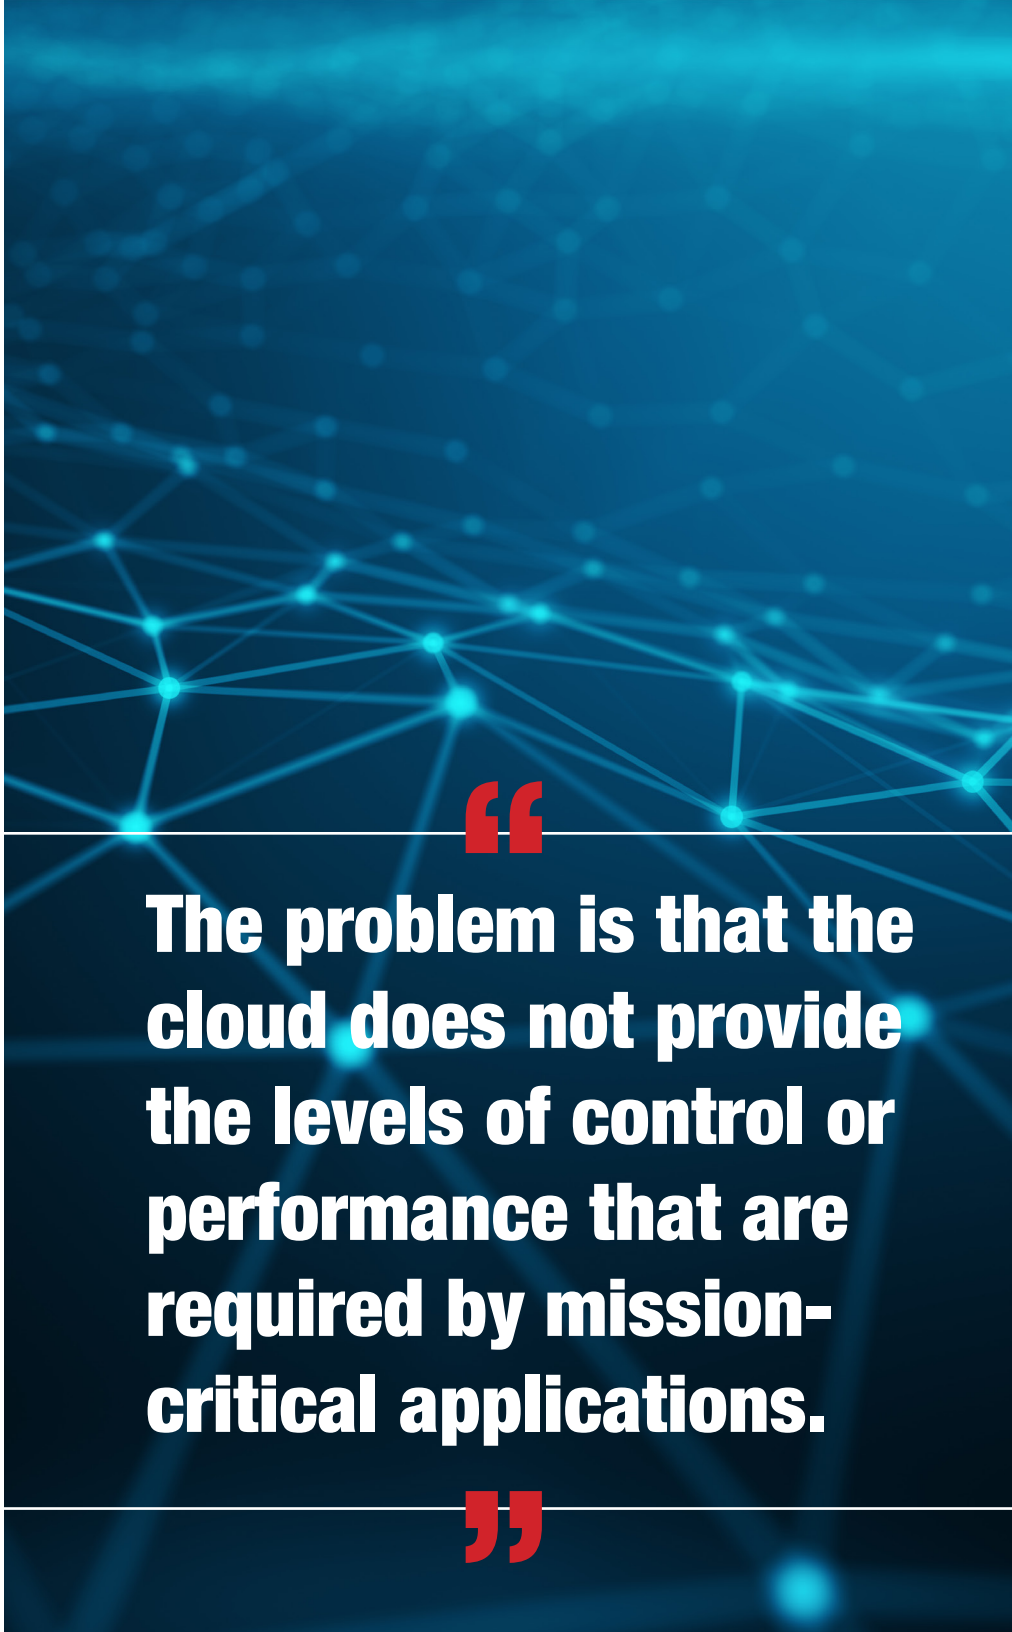

3 Multi-Cloud Needs Composable Infrastructure

Most organizations now live in a multi-cloud world. Workloads can run on-premises one day and then in the cloud the next, only to return back to on-premises soon after. A challenge for organizations looking for this type of workload mobility is the differences between the infrastructures. Current on-premises solutions, despite attempts to make them more cloud-like, haven't worked. Most vendors forget that the public cloud is more than just a business model -- it is an actual way of doing business. On-premise solutions should mirror the flexibility – both in technology and economics – of the public cloud. The best approach to achieve this is by leveraging a composable infrastructure.

What is Composable Infrastructure?

Most on-premises storage systems, even scale-out systems, are hard wired. They have a given amount of storage, compute, IO, and capacity resources that are shared across all the workloads attached to them. This sharing of all resources leads to storage silos which leads to data center inefficiency and budget overages.

Composable infrastructure enables organizations to specifically define what components of a storage system will be assigned to a specific workload or group of workloads. For example, a performance-sensitive, Oracle environment may need extra controller power in order to keep pace with a high transaction application as well as a high number of flash drives. A VMware environment, on the other hand, may only need two controllers, and -- thanks to deduplication -- a very finite number of drives. Lastly, an analytics process may need more bandwidth than random IO performance, so it may also only need two controllers but require additional network connections.

The needs of each of these workloads may also change from day-to-day or week-to-week. In order to achieve high levels of efficiency, IT may want to reallocate resources dynamically between these workloads as the need arises, instead of having expensive resources sitting idle.

All of the movement of these resources is controlled through the storage software. IT administration can rapidly define a new virtual storage system out of the discrete physical components.
The challenge for IT is that of the few solutions that provide composable infrastructure, most are delivered as a turnkey hardware/software solution. This means the customer pays a premium for the hardware with the software licenses locked to the specific hardware. If the customer adds an additional storage node, they must also purchase the storage software again instead of transferring the license from the old node.

Composability 2.0 - Software Only

The next generation of composable storage is solutions where the hardware and software is disaggregated. While they may still come bundled with certified hardware solutions, software and hardware are independent of one-another. Software independence enables the transferring of licenses which enables organizations to quickly move to new hardware platforms when new technology justifies it. Software independence also enables a true consumption or subscription model. While the customer needs to purchase the hardware upfront, the software only needs to be activated as the organization uses that portion of the hardware. Most consumption models built on legacy storage solutions, or even composable storage 1.0 solutions, require a minimum purchase by the customer and upgrades that are in large chunks instead of small bites.

Conclusion

The second generation of composable infrastructure is more compatible with the cloud. Although composable storage 2.0 is actually more flexible than the cloud, it does share similar properties in terms of seamless scale, upgrades, and business models. Organizations looking to maintain mission-critical applications on-premises, while also having the option to move to and from various cloud providers, will find composable storage 2.0 an ideal foundation from which to build that strategy.

4 Kaminario Drives Composable Storage 2.0

Accelerated performance and usage-based storage for business-critical applications

A new generation of composable infrastructure that is truly software-only is required to deliver the levels of multi-cloud data mobility, automation, efficiency, and flexibility required by modern workloads. These solutions enable the independent transfer of software licenses across underlying hardware, as well as a true consumption-based or subscription model.

About Kaminario

Kaminario got its start in 2008 as a provider of flash storage area network (SAN) arrays. The company remains focused on the all-flash array segment in the market, but in 2018 pivoted away from hardware, in favor of providing a software composable storage platform. Kaminario's architecture, which it calls its Composable Storage Platform, is comprised of three core software platforms that it sells as a combined bundle:

- VisionOS, which turns industry-standard hardware into software-based resources and services for data control and optimization. VisionOS includes data reduction and protection capabilities including in-line, global, and selective deduplication; zero detect; data-at-rest encryption; native, array-based snapshots and replication; no single point of failure; and a robust RAID scheme. It offers the flexibility to be managed through a simplified graphical user interface (GUI), scriptable command-line interface (CLI), or third-party tools via a RESTful application programming interface (API).

- Clarity, a management and monitoring software that applies machine learning for more intelligent and predictive insights. For example, it can support better resource optimization and capacity planning, as well as streamlined, faster detection and resolution of hardware and software outages and performance issues.

- Flex, Kaminario's storage orchestration platform that provides its capabilities around composability. Flex builds upon VisionOS by taking the disaggregated resources and allocating them on a policy-driven basis according to application and business needs.

The combination of Kaminario's VisionOS and Flex software enables customers to build racks of commodity, off-the-shelf compute and solid-state drives (SSDs), and then compose those resources into logical storage resources that are tailored to specific requirements. For example, to serve high performance workloads, the configuration can be CPU-heavy to drive controller functionality with a smaller amount of SSDs. More capacity-oriented workloads may require just two controllers but larger amounts of high-capacity flash. The most basic configuration of these storage resources include two active-active storage controllers and storage capacity in the form of SSDs and Fibre Channel (FC), or iSCSI connectivity. These storage resources can be scaled as needed, adding one or more storage controllers for additional performance, or one or more JBOF shelves for additional capacity. This enables creation of scale-out clusters which complement the ability of each node to scale-up with additional flexibility.

Kaminario delivers flexibility in payment and procurement of its Composable Storage Platform. Customers may purchase appliances based on SuperMicro servers, Western Digital and Samsung SSDs, and Mellanox or Broadcom networking, through the vendor's partnership with distributor TechData. Kaminario's software can be procured "as a service" through a subscription model. Customers may purchase a global license-based subscription, or they may opt for a monthly subscription that is based on the high watermark of a customers' usage per month. Kaminario also has utility-based pricing for hardware and software, and it is working with a number of managed service providers for full "as a service" procurement of hardware and software.

StorageSwiss Take

Kaminario's background and pure focus on accelerating application performance is the company's leading differentiator as it shifts to composable storage. According to Kaminario, its ability to scale out and add performance enables it to achieve performance levels of up to 1.6 million input/output operations per second (IOPS) in a 4 K-Block node. These levels of performance, coupled with the increased agility and infrastructure resource utilization, will be draws for Tier One enterprise applications, such as online transactional processing (OLTP), and for large enterprises and service providers looking to consolidate multiple workloads for cost efficiency without performance degradation.

Conclusion

IT professionals face a multitude of options, in the form of on-premises solutions and cloud services, when it comes to meeting their workload consolidation and application acceleration needs. What makes this decision even more challenging is that the business model, in terms of how the technology is procured, must be evaluated in addition to the technology speeds and feeds themselves.

Storage Switzerland, LLC

THE FIRM

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, whitepapers, and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances, and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.

kaminario.

OUR PARTNER

Kaminario delivers a Storage as a Service (STaaS) platform for business-critical applications hosted on and off premise. Its software composable architecture combined with advanced analytics and automation allows customers to realize the agility and simplicity of public cloud, while gaining the performance and control of dedicated infrastructure. With disruptive economics, Kaminario delivers flexible, consumption-based licensing that aligns with customers' business priorities. Headquartered outside of Boston, Kaminario works with a network of strategic partnerships globally. For more information, please visit www.kaminario.com.