BROOKE REDUCES RISK AND CUSTOMER CHURN LIKE A BOSS. BOTH LITERALLY AND METAPHORICALLY.

A SOLIDFIRE PAPER

A Service Provider's Perspective: Scale-Up vs. Scale-Out

NetApp®

For the service provider, one of the foundational tenets of the web-scale next generation data center (wNGDC) is that of using scale-out based architectures. This principle is particularly pertinent to storage arrays where the concept of an elastic, scale-out topology is positioned to become the predominant architecture used in service provider cloud infrastructure deployments. The ability to offer seamless, transparent resource expansion without the cost and complexity of traditional infrastructure migrations adds significant value to the service provider in terms of both top-line revenue generation and cash-flow optimization.

Market Trends have driven the need for scale-out storage architecture

Traditional scale-up architecture can be complex and risky

A number of years ago, the traditional, siloed data center relied upon dedicated, physical servers with their associated stranded resources. These servers, along with their storage footprint, dictated the maximum amount of data that could be managed or stored. After the drivers of standardization and consolidation had been implemented, traditional service providers gained significant efficiency savings and server utilization through virtualization.

Virtualization was built on the concept of sharing compute, memory and I/O resources between multiple workloads. Technologies such as KVM, ESXi, and HyperV took this further, enabling resource clustering that enabled the parallel processing of applications, load balancing, and fault tolerance (redundancy) in the cloud. The pooling of resources such as CPU, memory, and I/O, enabled service providers to perform quick, non-disruptive modifications using a singular orchestration software suite and thus quickly meeting their changing business needs.

The challenge came that as compute and network administrators began leveraging the efficiencies of scale-out architectures, storage technology remained mostly static. Service providers had little choice but to use separate storage networks and disparate storage pools to deliver the “tiers” of performance that the market was demanding. The storage administrators managing these traditional scale-up architectures were typically spending significant proportions of their time on planned and unplanned maintenance instead of enabling business innovation.

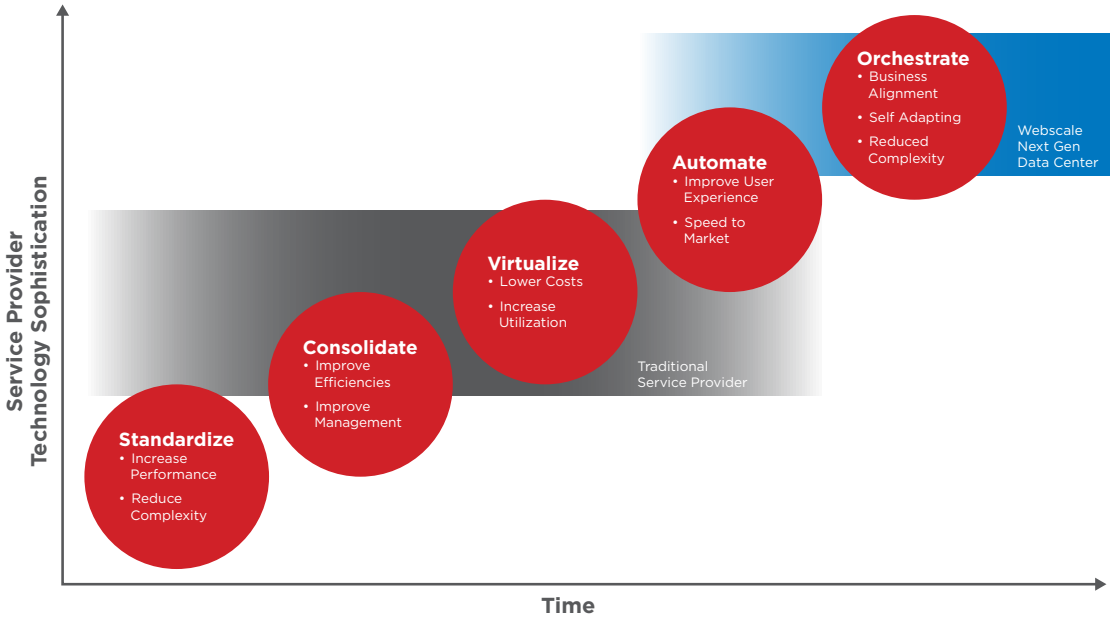

Storage administrators are typically spending time on planned and unplanned maintenance instead of enabling business innovation.

A classic example of storage complexity — that remains to this day — is the process of "scaling up" traditional dual-controller storage architecture where drive shelves and/or incremental compute is required. Not only are there computational and planning complexities to this task, there is also a financial risk in the form of an often large, upfront capital expenditure that is needed to pay for all the resources that are anticipated for the next three years of expansion. As a result, not only is this process complicated and time consuming, it is fraught with risk.

All of these challenges can be eliminated if a scale-out based next generation storage architecture is used.

Scale-up versus scale-out

With a scale-up architecture, the hardware you operate determines the limit of variables that can be expanded. The capability of each component limits the collective performance.

In the service provider space, data and applications usually only proliferate and, ultimately, the maximum capacity of compute/storage will be reached for that architecture. As more and more applications are added, the infrastructure resources become spread out, management overhead increases, and performance degradation creeps in.

In contrast, an elastic scale-out architecture expands the once siloed, static storage ecosystem into a clustered set of pooled storage resources (GB, IOPS). It is by its very nature a distributed architecture and thus is not limited to the resource capacity of any single machine. It delivers linear expansion of all variables such as CPU, memory, IOPS, and storage, allowing the service provider to incrementally scale on an as-needed, on-demand basis. The result is that service providers are able to more effectively align storage array purchases to user consumption and better manage their cash flow.

In addition to this, web-scale next generation data center (wNGDC) scale-out storage provides the flexibility to independently and non-disruptively scale both capacity and performance in a predictable linear pattern over time. This means service providers can strategically scale in or out and distribute data and traffic over any number of nodes while increasing the scope of data services.

Successful scale-out storage isn't automatic...

While the benefits of scale-out storage are clear, the successful implementation of this type of infrastructure can only occur when a number of key architectural capabilities are in place. These required capabilities can be summarized in the following five keys that address both the usability of the array as well as the end user experience:

Key capabilities of a scale-out storage array

- **Simple, intuitive process of adding/removing resources (nodes) from the cluster**
- **Automatic resource balancing (compute, memory, IOPS, storage etc.) after scaling out/in**
- **Flexibility to non-disruptively scale independent resource pools of capacity (GB) and performance (IOPS)**
- **The addition/removal of resources (scale-out/in) must be transparent to the customer and not affect workload performance; no-downtime**
- **Minimum performance levels must be protected and controlled on an individual workload by workload basis using Quality of Service techniques.**

The first two keys have significant impact on the management and maintenance overhead of the scale-out cluster and directly affect the resources that the service provider has to apportion to the array. The addition of new resources must be easy and intuitive, including not only a simple GUI but also a minimum of cabling requirements. Current market leading technology uses a 1U form factor node, requires only four Ethernet cable connections, and can be added to a cluster with four mouse clicks in an easy-to-follow user interface. Typical installation times from initial switch on to the provisioning of volumes is now down to a mere 10 minutes.

The other keys address possibly one of the most critical business attributes of the service provider: ensuring user experience and minimizing customer churn. The need for individual workload performance control (often known as Quality of Service) that ensures any application has a guaranteed minimum level of IOPs is critical to ensure good customer experience and minimize customer churn.

As a result, in a successful scale-out storage architecture, minimum, maximum, and burst IOPS levels are implemented on a workload by workload basis, creating predictable, controlled application performance that is independent of any other activity on the cluster (such as from the “noisy neighbor”). This allows providers to differentiate their services through software-defined performance tiers and aggressive SLAs that encourage customer retention.

One of the most critical business attributes of the service provider is ensuring user experience and minimizing customer churn.

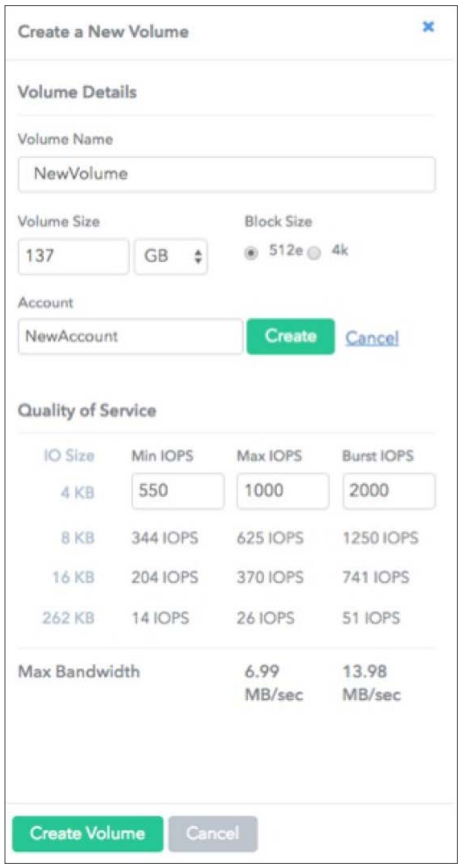

The future of service provider storage architecture

As the service provider market has matured, the need to provide next generation web-scale storage infrastructure has grown as a result of the limitations of dual controller based scale-up architectures. The expense of implementation, complexity of maintenance, and lack of scalability has driven the design and adoption of clustered, scale-out storage architectures that can transparently and non-disruptively expand and contract on an as-needed basis.

The design concept of NetApp SolidFire's next generation storage architecture meets the needs of the service provider through an elastic, scale-out storage array that incorporates full application control at a granular level. With comprehensive API, orchestration integration, and data protection capabilities, SolidFire continues to be globally adopted as the standard for service providers looking to grow their business and maximize profit margins.

For a more detailed explanation of this topic and how to differentiate your business in the service provider market, download the white paper: A Service Provider's Perspective: Designing the Next Generation Data Center at www.solidfire.com/sp-designing-the-ngdc.