Datasheet

Element Software Quality of Service

Performance predictability turns down noisy neighbors while increasing efficiencies

Key Features

Predictable

- Deliver consistent performance to thousands of applications
- Secure performance in unpredictable, mixed workload environments
- Scale performance independently of capacity on demand

Efficient

- Consolidate multiple applications to a single system
- Decrease the cost of operations with fewer storage platforms
- Reduce the number of tools needed to manage storage

Results

- Establish and meet customer SLAs
- Drive increased revenue with tiered offerings
- Increase operational efficiencies

The Challenge

Businesses today are tasked with working out how to build a flexible, scalable platform that can support multiple workloads while improving operational efficiency. To enable this business model, IT environments are transitioning from siloed IT environments to a single infrastructure with a broad set of capabilities.

At the heart of delivering a single infrastructure, on demand, and as a service is multitenancy, in which multiple applications and/or customers reside in the same storage infrastructure. At first glance, the opportunity to run a broad array of applications in a single system might sound easy, but the reality for today's IT managers is very difficult. When many performance-sensitive applications are consolidated onto a single platform, noisy neighbor applications arise and cause resource conflicts, unpredictable application performance, and frustrated customers. For administrators, managing and troubleshooting storage performance are a labor-intensive process that takes up valuable time. IT professionals ask, "How do I maintain consistent performance for a given application?" The answer is quality of service (QoS) for storage systems.

Storage systems with native QoS capabilities have become the key transitional element for companies making the shift to a cloud infrastructure that is home to applications and workloads that run smoothly and efficiently, rather than lagging and freezing.

What to Look for in Quality of Service

Solving for application resource contention requires balanced performance at the system level. From this starting point, a storage system can then deliver performance to serve the unique needs of different applications. This ability to finely allocate raw storage performance is important, but the predictable and consistent delivery of that performance are what ensure that every application has the resources required to run without variance or interruption. In servicing these workloads, IT must consider how well the underlying storage architecture solves for the following:

- Unpredictable I/O patterns
- Constantly changing workload and application performance requirements
- Scaling of performance and capacity resources on demand

The Solution: Native Quality of Service

QoS is a critical enabling technology for enterprises and service providers that want to deliver consistent primary storage performance to business-critical applications in a multitenant infrastructure.

When you consolidate these applications onto a single storage system, you risk having too much performance variability to ensure that any particular workload gets the performance it needs. What do you do?

NetApp® SolidFire® native QoS lets you allocate and manage performance independently from capacity, making certain performance is available to every application you run. Each storage volume in a SolidFire system can be allocated a

NetApp®

precise amount of capacity and performance, both of which can be changed on the fly without migrating data or interrupting I/O. This approach enables firm performance SLAs for applications, workloads, and tenants across your infrastructure, a key differentiator for service providers and enterprises alike.

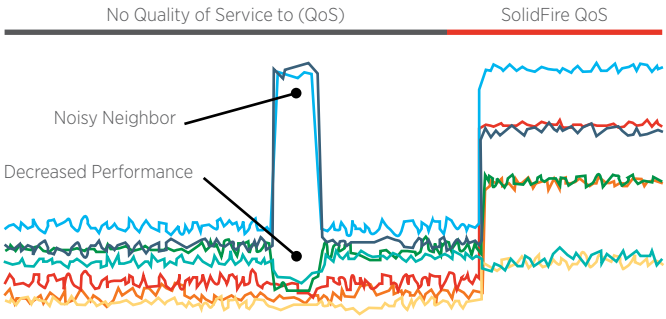

Figure 1) With QoS, applications no longer compete for performance

With SolidFire QoS, you can assign fine-grained levels of performance (IOPS and bandwidth) to thousands of volumes residing in a single storage system while meeting performance SLAs to internal or external customers.

With SolidFire QoS, you can:

- Ensure performance to every volume with fine-grained QoS settings
- Manage performance in real time without affecting other volumes
- Allocate storage performance independently of capacity

This approach proactively provides applications with the performance they require from day one throughout the life of their deployment. With SolidFire QoS, applications no longer compete for performance, and administrators no longer have to hassle with complex tiering systems or prioritization schemes.

This unique technology enables service providers and enterprises alike with two key operational accelerators:

- The ability to control performance and capacity independently from one another
- The ability to set fine-grained performance service levels on a per-volume basis

Volumes provisioned in a SolidFire system are assigned three performance values: max IOPS, burst IOPS, and min IOPS. Each value can be monitored, tracked for chargeback, and changed on the fly without affecting volume or system performance.

Max IOPS is the maximum number of sustained IOPS a volume is allowed to process over an extended period. Applications are not permitted to consistently exceed this level and affect other applications, ultimately eliminating the noisy neighbor.

Burst IOPS is the maximum number of IOPS a volume is allowed to process over a short period. When a volume uses less than its max IOPS, it accumulates credits, which can be used to burst to a volume's burst IOPS limit for a short period. Burst IOPS is particularly effective for virtual machine reboots, migrations, large file transfers, and other heavy loads that need to be completed in a short period. This functionality is only allowed when system performance resources are available, preventing any impact on other applications.

Min IOPS is the minimum number of IOPS that an administrator grants to a volume; it sets the floor. This IOPS level is effectively the focus of most conservative service-level agreement (SLA) provisions. Min IOPS values come into play only if the system becomes bound by I/O capacity, at which point the system scales all volumes back from their max IOPS levels proportionally toward their min IOPS values. Setting SolidFire min IOPS ensures fair resource allocation when the system is heavily loaded and offers a prioritization mechanism to give more important volumes priority at times of heavy load, while others are scaled back more dramatically. In all cases, the min IOPS setting ensures a predictable level of performance rather than the random performance degradation typically seen in unpredictable I/O patterns and/or performance constrained situations.

SolidFire native QoS is purpose-built to enable IT organizations of enterprises and service providers to allocate, manage, and assure storage performance, making it faster and easier to respond to changing demands of applications and the business than ever before.

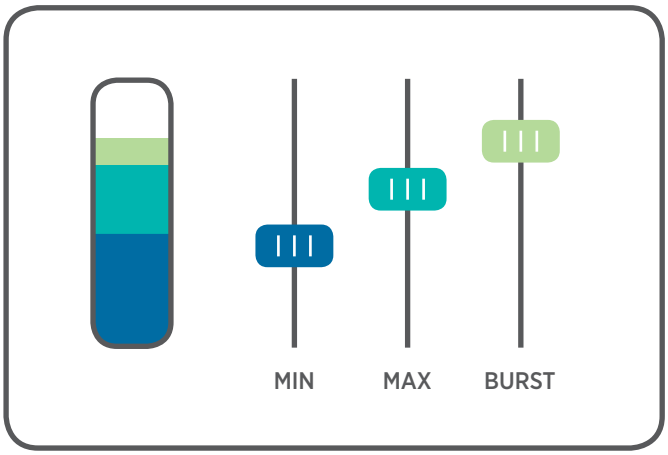

Figure 2) Native QoS ensures fair resource allocation

Who Benefits

Service Providers

The unique design of multitenant infrastructures leads to more inherent challenges than those experienced by a traditional enterprise. Without QoS, applications can run rampant and quickly become noisy neighbors. For the service provider, performance variance caused by a noisy neighbor can have huge consequences if those neighboring applications are those of another company or customer. Because hosting business-critical applications in the cloud represents a large revenue opportunity for service providers, the ability to deliver predictable hosting

services free of noisy neighbors is a critical capability. 451 Research states that only 32% of enterprise applications are running in a hosted infrastructure, and until storage performance is predictable and guaranteed, service providers won't be able to programmatically deliver services that attract the other 68% of those enterprise workloads. Is there a solution? Yes, and the answer is storage QoS architected from the ground up with predictable performance in mind. The storage system's ability to establish a minimum level of performance makes writing SLAs a snap. Regardless of system condition or an application's activity, performance is certain and has become a sure way to attract new enterprise revenue enabled by tiered services.

Enterprises

Enterprises today are tasked with consolidating siloed storage appliances, each corresponding to a separate workload and building a hybrid cloud infrastructure that can support multiple applications while reducing operations.

Without QoS, storage administrators have had to spend the bulk of their time tuning, tweaking, planning, and troubleshooting storage performance for individual workloads on disparate systems.

The primary advantage QoS offers to enterprises is the ability to consolidate multiple workloads, particularly those that have been isolated in separate storage silos onto a single platform. Being able to deploy many applications onto a single storage system with QoS makes it easier to address performance-related challenges quickly. QoS provides enterprises with the ability to deliver consistent performance to thousands of applications with far less hardware, improved operational efficiency, and simplified management so they can focus on moving the business forward.

Business Advantages

Today, large enterprises and IT service providers face the biggest challenges around growth and data management complexity. That is why predictability and automation are key to improving IT service delivery and staying on top of business demands and future transformation.

With SolidFire QoS, enterprises and service providers now have more flexibility and predictability from a solution that allows them to provision capacity and performance independently and uniquely for every application while improving operational efficiency. SolidFire QOS:

- Reduces the number of storage platforms, so the cost of operations goes down, and the number of tools needed to manage storage is decreased.
- Provisions capacity and performance independently out of a single pool, so less overprovisioning is required to meet the needs of the workloads.
- Provides a scalable platform that can grow or shrink based on the collective needs of the business, so enterprises can make more efficient use of capital, space, power, and management resources.

Enterprises and service providers expect applications to get the exact experience their customers are paying for: no more, no less.

About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit www.netapp.com. #DataDriven