

ACHIEVING HIGHER RESOLUTION STORAGE FABRICS THROUGH VIRTUALIZATION

Nevoa Networks' storage virtualization suite, Nevoa Storage System (NSS), enables the deployment of high-resolution storage fabrics while simplifying resource provisioning and policy automation.

As storage virtualization solutions continue to gain popularity and take center stage, the implications these technologies carry are still being learned. Storage virtualization is not just one more technology to be added to an IT infrastructure, it is a whole new **information ecosystem** impacting purchasing, maintenance and data migration cycles. As the fundamental layer that offers a **generic storage stratum**, how can you ensure that the adaptive potential of your virtualization solution is an asset instead of a source of disruption? **Nevoa Networks'** Nevoa Storage System (NSS) is designed to support and manage such information environments using an unprecedented level of **storage fabric resolution**. The result is greater control over storage resources, fine-grained provisioning, highly automated management, and massively scalable storage fabrics.

Possibilities and Potential

Even a cursory look at the touted benefits of storage virtualization will provoke visions of unlimited possibilities or immediate skepticism in most IT managers. The ease with which one can organize and restructure one's resources and data within a virtualized environment that offers robust abstractions of the different layers has universal appeal. Achieving new levels of automation with precise policy descriptions that can maintain, migrate, and protect the data is now supposed to be simpler. There is also a significant economic appeal in that it **tips the balance** between vendor lock-in and client fidelity in favor of the customer and down plays the need for high-end hardware.

Enchantment of Consolidation

Often presented as one of the more tangible benefits of storage virtualization, storage infrastructure consolidation is well-understood and its immediate impacts are characterized as cost-saving ones. Much like a network operations center (NOC), storage consolidation solutions provide a global view of the storage fabric which makes quick identification of hotspots, utilization levels, under provisioned sub-systems, and failures possible. Such solutions also tend to include tools for remote management of the storage nodes, although vendor interoperability is often limited.

While all of these are desired characteristics and can save large amounts of valuable administrative time, there are others that determine the overall flexibility of the system and extent of interruption delays due to restructuring of the storage fabric. Unfortunately, these qualities of the storage fabric are usually only tested after consolidation has already occurred. If the chosen storage virtualization solution supports high-resolution fabrics, then negative impacts of restructuring are reduced and extensive automation is sustained. The **lower the resolution of the storage fabric**, the more cumbersome automation and resource reallocation become.

Understanding the Importance of Resolution

A simple definition of resolution is the number of like units within a finite space. Common measures of resolution are presented by digital cameras, scanners, high-definition TV sets and such. Higher-resolution means that the unit (i.e. a pixel) represents a smaller portion of the information present in the finite space. Resolution has a strong relation to density, a term often applied to storage, which also uses finite space as a constraint. When applied to storage fabrics, resolution refers to the number of individual resources that represent the storage capacity of the storage network.

For example, prior to consolidation by storage virtualization, a storage network contains four storage hosts. The aggregate capacity of all the hosts totals 5.5 terabytes (TB), which is the extent of the finite space. Once the hosts are consolidated into a single storage system we can determine the resolution, **gigabytes per resource** (gpr), of the storage fabric (see Figure1). In some cases a




Figure 1.
Each storage system is seen as a single storage resource by the storage fabric.

single storage host can actually represent multiple storage resources. Large capacity storage systems are often divided into logical volumes, creating smaller storage resources, but increasing overall storage fabric resolution (see Figure 2). Logical volumes usually represent a set of disks managed by the same host and must be defined and formatted before being made available to users and




Figure 2.
Two of the storage systems are each seen as two storage resources, increasing the storage fabric resolution.

applications. To overcome the increased probability of failures due to the usage of multiple disks, the fault tolerance strategies are built-in to the chosen format. Even when only the storage capacity of one disk is sufficient, it is paired with another in order to support mirroring. For both SAN and NAS storage solutions, the **effective unit** of storage is always a disk set containing more than one disk bound by host-applied formatting.

Ideally, the storage fabric should be configured to its highest functional resolution prior to virtualization and hand over the responsibility of data reliability strategies, like erasure coding and replication, to the virtualization mechanisms. This is rarely done because it can involve heavy data migration and reconfiguration of each storage device or appliance. Consequently, storage fabrics built using virtualization usually inherit the resolution of the former storage network. This still provides the benefits of infrastructure consolidation, **but strongly undermines the benefits of virtualization** when defining automated resource reallocation and sub-system reconfiguration.

So what is a high-resolution storage fabric? To answer this question it is important to understand that your storage infrastructure represents a collection of storage devices. Here, a device refers to the individual units that will read and write the data such as a hard disk drive, optical drive, or tape drive. Breaking our example infrastructure down even further, we can see that a host can represent a set of storage resources and that each storage resource further represents a set of storage devices (see Figure 3). Therefore, the highest attainable resolution for this fabric is reached when each storage device represents a single storage resource, resulting in a 1:1 resource to device ratio.




Figure 3. Decomposition of the four storage systems into single storage devices. This shows the maximum resolution for non-NSS storage fabrics.

High-Resolution Virtualization Solution

Based on logistical networking technologies, the Nevoa Storage System (NSS) was designed from the beginning to support large numbers of virtualized storage resources regardless of size, performance, and underlying composition. This means that it inherently accommodates the highest attainable resolution of your storage fabric. In fact, because of its underlying storage virtualization technology, you can go beyond a 1:1 resource to device ratio by creating **virtual partitions** on storage devices (see Figure 4).

Even in non-virtualized storage networks, one of the strategies used is to create a pool of idle disks that can be drawn from in order to grow active storage resources as utilization increases. This is a favored strategy because growing a storage resource through joining is usually a simpler operation than splitting, or splintering, an active resource. By supporting high-resolution storage fabrics, **NSS can scale active resources using smaller units of storage**, or gpr values. In

NSS supports virtual partitioning of individual storage devices, **breaking** the former resource to device ratio barrier of 1:1





Figure 4.
Virtual partitions in NSS are presented as individual storage resources, further increasing storage fabric resolution.

other words, **active storage resources can grow using a disk set, a single disk, or a portion of a disk.** This is significant because cost/benefit analysis favors the acquisition of higher capacity disks. So while the storage system can support provisioning activities using a single disk, that disk continues to represent a larger percentage of the system's capacity. For example, a 5TB storage system using twenty 250GB disks inherently has a higher resolution than a similar system using ten 500GB disks. Adding 500GB instead of 250GB to an active storage resource may be excessive and can tie-up storage capacity soon needed by other resources. The disk could be partitioned, but without virtualization, the necessary data migration involved could be disruptive. High-resolution fabrics using NSS virtualization technologies give you the control needed to overcome these limitations and the tools to build advanced resource management policies:

- ❁ **Profile-based Storage Pools:** NSS virtualization technology allows you to create multiple storage pools based on resources with similar attributes. Each pool can have a different resolution, represent different types of storage devices, and be available only to certain active resources.
- ❁ **Format-free Virtual Volumes:** because the storage volumes are virtual, NSS works independently from the underlying formatting and file system of each storage device. Data reliability settings, such as RAID 0+1 or erasure encoding, for a storage volume can be changed without migrating the data off the volume, reformatting the physical resources, or experiencing downtime.
- ❁ **Active Provisioning:** NSS virtual volume monitoring intervals and acceptable utilization and integrity levels are all configurable. Once these are set, NSS will automatically take actions to maintain the volume within operational specifications when faced with device or communication failure and changes in utilization patterns.
- ❁ **Hybrid Deployments:** NSS can easily be deployed on existing storage infrastructures creating a hybrid data environment. Idle storage capacity on current hosts can be recovered and offered as new storage resources, or used for data staging during a physical-to-virtual storage migration.

- 
Flexibility: The NSS architecture uses a layered approach and is highly modular, making initial configuration decisions less critical and future reconfigurations trivial. Services at each layer can easily be duplicated, providing better fault-tolerance, hierarchical arrangements, and extending the NSS to all locations of a network.

Adopting NSS virtualization technology has an **immediate impact** on hardware purchasing cycles and trends by bringing the raw cost-per-gigabyte back to the center of the decision making process. The storage system itself does not need to be endowed with expensive RAID controllers, SAN adapters, nor specialized storage cluster components like head-nodes. In fact, because of the advantages of high-resolution fabrics, low-level storage device grouping is not recommended. Virtualization also lets failed storage systems or devices be replaced with non-identical solutions, eliminating possible supply bottlenecks, extensive downtimes, and vendor lock-in.

Real-World Examples

The following examples demonstrate the impact of applying NSS virtualization to different types of storage systems common to corporate storage infrastructure. It is important to remember that this is largely possible due to the fact that data reliability policies are handled by NSS, making the low-level hardware configurations less critical.

Table 1 – Fabric resolution and economic scalability comparison

System Characteristics	Host-attached storage	NAS Device	SAN Fabric	Storage Cluster*
Number of enclosures	1	1	6	7
Number of drives	12	12	60	28
Drive types	SAS	SATA	SATA	SATA
Number of volumes	2	4	20	14
Total capacity	1800GB	6000GB	45000GB	21000GB
Internal bandwidth**	36Gbps	36Gbps	180Gbps	84Gbps
External bandwidth	2Gbps	2Gbps	8Gbps	14Gbps
Cost per terabyte	US\$5.4k	US\$2.5k	US\$2.2k	US\$1.9k
Effective resolution (capacity/number of volumes)	900 gpr	1500 gpr	2250 gpr	1500 gpr
gpr = Total Capacity/Number of Volumes		*deployed on eight 1U commodity servers **aggregate device throughput		

The *gpr* metric shows the average size of a storage resource that can be reassigned without undoing the data reliability strategy (RAID levels, replication) applied to the resource. Reassignment operations in NSS can draw on idle or active resources, buying time while more storage capacity is being supplied. On

As a 100% software-based solution, NSS deployment does not require any new hardware purchases.

Table 2 – Post-virtualization resolution using NSS

System Characteristics	Host-attached storage	NAS Device	SAN Fabric	Storage Cluster
Number of enclosures	1	1	6	7
Number of drives	12	12	60	28
Drive types	SAS	SATA	SATA	SATA
Number of volumes	12	12	60	28
Total capacity	1800GB	6000GB	45000GB	21000GB
Internal bandwidth	36Gbps	36Gbps	180Gbps	84Gbps
External bandwidth	2Gbps	2Gbps	8Gbps	14Gbps
Effective resolution* (capacity/number of volumes)	150 gpr	500 gpr	750 gpr	750 gpr
<i>*higher resolutions can be achieved using NSS virtual partitioning technology</i>				

average, it takes less than a minute to add an NSS-ready storage node and less than five minutes to configure an existing host to participate in an NSS storage fabric.

The shift towards higher-resolution as reflected by the *gpr* scores can also be interpreted in terms of costs. These represent how the investment in storage infrastructure is allocated amongst user groups. The result is that an increase in storage fabric resolution reduces the denomination of the currency, thereby increasing the liquidity of the underlying storage assets.

Unprecedented Control and Accessibility

High-resolution fabrics give administrators the level of granularity they need in order to create automated policies that can maximize resource usage and optimize provisioning tasks. NSS does that and more allowing storage administrators to draw from any, active or idle, resource pool in order to adjust to changes in storage demands without sacrificing precious uptime or impacting resource availability. It also makes the power of storage virtualization accessible to businesses of any size by giving them complete freedom of choice over future hardware purchases.

There is great potential and significant economic gain when adopting storage virtualization, but it must be harnessed and adapted to the evolving entity that is the corporation. Not all information has the same value and as long as it can be recovered, it is always in transit. Nevoa Storage System is aware of these realities and offers the technology, tools, and learning needed to get the highest possible return from crucial corporate resources.

For more information, contact a Nevoa Networks sales representative or a Nevoa Resource Partner, or visit www.nevoanetworks.com.

Corporate Headquarters

João Pessoa, Paraíba, Brazil

Tel: +55 (83) 3226-3544

info@nevoanetworks.com



© 2007 Nevoa Networks Ltda. All Rights Reserved.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Nevoa Networks. Nevoa Networks reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Nevoa Networks sales representative for information on feature and product availability.