nexenta™

Global Leader in Software Defined Storage.

WHITE PAPER

Software Defined Storage Hydrates the Cloud

Table of Contents

Overview

As many companies have entered the storage virtualization market a new category has been created called Software-defined Storage (SDS). Although these companies use this term, true software-defined storage (SDS) is characterized by a separation of the storage hardware from the software that manages it. SDS is a key enabler in the trend of moving the Datacenter toward a more software-defined datacenter. As with any new technical term there are numerous interpretations in the market. Some vendors claim they offer SDS but instead are offering virtualized storage, characterized by a separation of capacity from specific storage hardware resources, meaning they are merely pooling storage devices. (EMC ViPR, NetApp Data ONTAP) Still other vendors who claim they have SDS solutions provide a hardware appliance which includes both hardware and software necessary to complete their solution. (Nutanix, HP StoreVirtual, Gridstore, Coraid)

The promise of SDS is that separating the software from the hardware enables enterprises to make storage hardware purchase decisions independent from concerns about over/underutilization or interoperability of storage resources. The key benefits of software-defined storage over traditional storage are increased flexibility, automated management, and cost efficiency. Nexenta's NexentaStor is the global leader in true SDS, offering a full featured NAS/SAN software platform that can be installed on standard commercial hardware with capabilities that meet and exceed the capabilities of traditional storage systems.

NexentaStor (Block & File Storage)

NexentaStor provides Security, Manageability, Availability, Reliability, lower TCO, and Scalability (SMARTS) across our storage infrastructure. Key attributes for this are:

- **Unified Storage** - NexentaStor concurrently supports block (FC, iSCSI) and file (NFS, SMB) access protocols across active/active controllers delivering a complete suite of data services: unlimited snapshots, clones, thin provisioning, inline deduplication, compression, and replication across all-HDD, all-SSD, and Hybrid configurations.
- **ZFS Foundation** - NexentaStor is built on the ZFS File System, a unique filesystem that provides simple administration, end-to-end data integrity, and immense scalability. Extended to 128-bit, the ZFS file system has the ability to handle virtually unlimited size of files. It also provides self-healing by using 256-bit checksums end-to-end to validate the stored data. ZFS offers software RAID, either single, dual or triple parity, through its RAID-Z and mirroring schemes which are not vulnerable to issues such as “Raid Write Hole”.
- **High Availability** - NexentaStor offers a proven Active/Active High Availability solution with transparent failover for NAS and SAN. Failover times are minimized to less than a minute.
- **Fault Management Architecture** - NexentaStor offers an improved Fault Management Architecture that allows I/O to continue in the face of hardware that has not yet failed but is not performing well with intermittent disk problems. Nexenta works with hardware vendors to comprehend their firmware specifications and embed intelligence to detect and handle drive issues thus minimizing their impact and keep IO flowing.
- **Resource Management at Scale –** NexentaStor is designed for large-scale growth with unlimited snapshots and copy-on-write clones. It also optimizes I/O performance with the ability to use 512GB memory cache per head. The total storage capacity of each cluster is designed to scale up to a petabyte or higher.
- **Replication** - Auto-sync is the NexentaStor periodic asynchronous data replication service that can provides remote data copy services for backup, disaster recovery, and archiving. Auto-sync improves management and control through compression and deduplication of transferred data and delivers recovery point objectives (RPOs) as low as a few seconds.
- **Automated Deployment** - In the interest of making it easy to create, configure, and manage storage pools, NexentaStor has a wizard in the GUI that enables an administrator to select preferred characteristics of the storage pool (e.g., number of cache devices and width of a RAID set) which then automatically determines an optimal pool layout and applies it. This can save significant time deploying large systems/environments.
- **Hardware Compatibility** - NexentaStor has been certified with a wide range of hardware platforms and components from vendors such as Dell, Cisco, HP, SuperMicro and Intel. Nexenta has developed a list of fully-certified, end-to-end reference architectures available through its partners. Please refer to nexenta.com for the latest information on certifications of these solution stacks.
- **Automation –** NexentaStor works with Cloud Automation tools such as Opscode Chef and Puppet, allowing storage administrators to rapidly react to performance and capacity issues. By integrating these automation tools companies are able to tune or modify their storage environments based on business, environmental or customer needs.

Software Defined Data Centers (SDDC)

Datacenters have gone through a major transition over the past ten years. The use of servers dedicated to a specific role or application has given way to server virtualization which is dramatically evolving the data center landscape for the compute portion of the datacenter, driving to higher levels of cost efficiency and flexibility. Until recently storage and networking have lagged behind in this transition but are now rapidly gaining ground with the use of software-based storage and networking running on a range of industry-standard hardware vs being tied to a specific vendor's hardware platform with the associated higher cost involved. The use of cloud automation suites enables higher use of software defined networking and storage. Datacenter infrastructure is built on three primary pillars: network, compute, and storage. These pillars are then augmented with a cloud automation suite before providing virtual machines with applications and data.

OpenStack

"OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control while empowering users with the ability to provision resources through a web interface." (OpenStack.org website)

OpenStack was originally developed with Hosting Companies/Service Providers as the primary target, but it is increasingly used in corporate private clouds. There is a large community of developers familiar with enterprise requirements that are working to enhance OpenStack and adding features required by enterprise customers.

OpenStack operates as an Infrastructure as a Service (IaaS) framework consisting of several modules as follows:

- **Nova** – cloud compute component that provides computing resources to users through virtual machines.
- **Swift** – a massively scalable, redundant storage module that provides object-based storage to users. It is typically provisioned from the underlying file system on the cloud server nodes.
- **Quantum** – software defined networking (SDN) component. Recent improvements in Quantum have made OpenStack considerably more acceptable for enterprise use.
- **Cinder** – persistent block storage component. NexentaStor currently supports Cinder through drivers developed by Nexenta and included in the OpenStack code.
- **Horizon** – a common dashboard for OpenStack deployments allowing users to configure, provision, and automate OpenStack resources.
- **Keystone** – identity management component. This is relatively new, but critical for enterprise
- **Glance** – shared services component for VM image discovery, registration, and delivery services. It is also useful for storing and cataloging backups.

There are several other development projects underway at OpenStack including: Ceilometer for metering and billing (chargebacks), **Heat** for Orchestration, and **Tempest** for integration testing.

CloudStack

"Apache CloudStack is open source software designed to deploy and manage large networks of virtual machines, as a highly available, highly scalable Infrastructure as a Service (IaaS) cloud computing platform. CloudStack is used by a number of service providers to offer public cloud services, and by many companies to provide an on-premises (private) cloud offering, or as part of a hybrid cloud solution. CloudStack is a turnkey solution that includes the entire "stack" of features most organizations want with an IaaS cloud: compute orchestration, Network-as-a-Service, user and account management, a full and open native API, resource accounting, and a first-class User Interface (UI)." (cloudstack.apache.org website)

As it competes with OpenStack, analysts have credited CloudStack with being more "enterprise ready" and easier to install and configure, while not as flexible as OpenStack.

NexentaStor in the Cloud

OpenStack

NexentaStor is a superior choice for provisioning OpenStack clouds with rock-solid, high performance block storage. NexentaStor is currently in use for a number of very large and significant public and private cloud deployments, including systems that scale in the hundreds of Petabytes.

NexentaStor block storage and NFS FS are both supported for provisioning OpenStack Cinder. The choice depends upon the customer's preferences and individual feature sets of each. In either case, Cinder delivers Block Storage to other components of OpenStack and End Users. The advantages that NexentaStor has over other choices include HA, efficient snapshotting, diagnostics, and tiering.

Nexenta is a participating member of the OpenStack Foundation and supports the OpenStack initiative for cloud computing. NexentaStor drivers for Cinder are added into OpenStack as they are updated.

CloudStack

CloudStack primary storage is associated with a cluster, and it stores the disk volumes for all the VMs running on hosts in that cluster. This is ordinary NFS or iSCSI storage and NexentaStor is the superior choice for provisioning CloudStack primary storage.

CloudStack provisioning with NexentaStor can now be automated using **Enterprise Chef™** cookbooks, using the Knife API plug-ins, developed in a cooperative effort by Nexenta, Schuberg Philis of the Netherlands, and Opscode, the developers of Enterprise Chef.

NexentaEdge (Object Storage)

Object Storage, like **OpenStack Swift**, is not a traditional file system, but rather a distributed object storage system for static data such as virtual machine images, photos/images, emails, backups, and archives. Having no central "brain" or master point of control provides greater scalability, redundancy, and durability.

NexentaEdge, currently in development and scheduled for release in late 2014, will provide a unique, high performing, scalable object store service for Cloud users. NexentaEdge will integrate seamlessly with Cloud installations and operate under a common management interface with NexentaStor. So Object, File, and Block storage can all be managed using a common framework.

NexentaEdge is an enhancement for Cloud Software Providers offering object storage services to customers. Object storage is rapidly becoming the storage architecture of choice for very large document storage repositories. Distributed object storage allows the system to distribute workload and capacity requirements across a large number of physical systems, while at the same time reducing the time and number steps required to locate specific files. NexentaEdge will be an ideal solution for cloud environments that can create a stand-alone object store, a cloud storage repository to provision applications, or deliver storage-as-a-service.

Conclusion

Cloud software-based solutions provide companies with the ability to provision and use autonomous/commodity hardware to assume any of the three traditional roles (storage, compute, and networking) that are found in a Data Center. By building an infrastructure that can rapidly redefine the data center personality that hardware will assume, companies can quickly and efficiently scale and adjust to corporate needs. The Cloud provides us with a vision of a data center full of hardware that is indistinguishable from one component to the next. A traditional server would be able to provide compute resource, network services, and storage capacity.

Looking into the future, even the individual components will likely find their way on the network. Imagine having computer nodes, memory nodes, and disk nodes, thereby enabling software providers to logically build the resource pools needed to deliver on the next generation of demands.

Nexenta is at the foundation of this trend. Having a completely software-based solution we become the mechanism that the SDDC will use to build and define storage requirements. By working with manufacturers and driving innovation across the industry Nexenta is able to help accelerate the next technology wave.