

SoftNAS Architecture on AWS

April 2017

This paper has been archived

For the latest technical content about the AWS Cloud, see the AWS Whitepapers & Guides page:

<https://aws.amazon.com/whitepapers>



© 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

SoftNAS and the SoftNAS logo are trademarks or registered trademarks of SoftNAS, Inc. All rights reserved.

Archived

Contents

Introduction	1
About SoftNAS Cloud	1
Architecture Considerations	1
Application and Data Security	1
Performance	3
Using Amazon S3 with SoftNAS Cloud	9
Network Security	10
Data Protection Considerations	13
SoftNAS Cloud is Copy-On-Write (COW) File System	14
Automatic Error Detection and Correction	14
SoftNAS Cloud Snapshots	15
SoftNAS SnapClones™	16
Amazon EBS Snapshots	17
Deployment Scenarios	17
High-Availability Architecture	17
Single Controller Architecture	20
Hybrid Cloud Architecture	21
Automation Options	23
Conclusion	25
Contributors	25
Further Reading	26
SoftNAS References	26
Amazon Web Services References	26

Abstract

Network Attached Storage (NAS) software is commonly deployed to provide shared file services, data protection, and high availability to users and applications. SoftNAS Cloud, a popular NAS solution that can be deployed from the Amazon Web Services (AWS) Marketplace, is designed to support a variety of market verticals, use cases, and workload types. Increasingly, SoftNAS Cloud is deployed on the AWS platform to enable block and file storage services through Common Internet File System (CIFS), Network File System (NFS), Apple File Protocol (AFP), and Internet Small Computer System Interface (iSCSI). This paper addresses architectural considerations when deploying SoftNAS Cloud on AWS. It also provides best practice guidance for security, performance, high availability, and backup.

Archived

Introduction

Network Attached Storage (NAS) systems enable data and file sharing and are used for business-critical applications and data management. NAS systems are optimized to balance performance, interoperability, data reliability, and recoverability. Although widely deployed by IT in traditional data center environments, NAS software is increasingly used on AWS, a flexible, cost-effective, easy-to-use cloud-computing platform. Deploying NAS on [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) provides a solution for applications that require the benefits of NAS storage in a software form factor.¹

About SoftNAS Cloud

SoftNAS Cloud is a software-defined NAS filer delivered as a virtual appliance running within Amazon EC2. SoftNAS Cloud provides NAS capabilities suitable for the enterprise, including Multi-Availability Zone (Multi-AZ) high availability with automatic failover in the AWS Cloud. SoftNAS Cloud, which runs within the customer's AWS account, offers business-critical data protection required for nonstop operation of applications, websites, and IT infrastructure on AWS.

This paper doesn't cover all SoftNAS Cloud features. For more information, see www.softnas.com.²

Architecture Considerations

This section provides information critical to a successful SoftNAS Cloud installation. This information includes application and data security, performance, interaction with [Amazon Simple Storage Service \(Amazon S3\)](#),³ and network security.

Application and Data Security

Security and protection of customer data are the highest priorities when working with SoftNAS Cloud on AWS. When you use SoftNAS Cloud in conjunction with AWS security features, such as [Amazon Virtual Private Cloud \(Amazon VPC\)](#),⁴ Amazon VPC Security Groups, and [AWS Identity and Access Management \(IAM\)](#) roles, you can deploy a secure data storage solution.

SoftNAS Cloud uses the CentOS Linux distribution, which is managed, updated, and patched as part of a normal release cycle. You can use SoftNAS StorageCenter™, the web-accessible SoftNAS Cloud administration console, to check the current software revision and apply available updates. For security and compliance reasons, the SoftNAS technical support team should approve any custom package before it is installed on a SoftNAS Cloud instance.

Web-based administration through SoftNAS StorageCenter is SSL-encrypted and password-protected by default. Optional two-factor authentication is also available for use.

You can administer SoftNAS Cloud through SSH and a secure REST API. On AWS, all SSH sessions use public/private key access control. Logging in as root is prohibited. Administrative access through the API and command line interface (CLI) over SSH are SSL-encrypted and authenticated.

Iptables, a commonly used software firewall, is included with SoftNAS Cloud and can be customized to accommodate more restrictive and finer-grained security controls. Data access is performed across a private network by Network File System (NFS), Common Internet File System (CIFS), Apple File Protocol (AFP), and Internet Small Computer System Interface (iSCSI). You can also restrict the list or range of client addresses allowed to perform data access.

SoftNAS Cloud offers encryption options for data security – both in flight and at rest. If NFS is used, all Linux authentication schemes are available, including Network Information Service (NIS), Lightweight Directory Access Protocol (LDAP), Kerberos, and restrictions based on the user ID (UID) and group ID (GID). Using CIFS, you manage security through SoftNAS StorageCenter, facilitating basic Windows user and group permissions. Active Directory integration is supported for more advanced user and permissions management in Windows environments.

The SnapReplicate™ feature provides block-level replication between two SoftNAS Cloud instances. SnapReplicate between source and target SoftNAS Cloud instances sends all data through encrypted SSH tunnels and authenticates using RSA (public key infrastructure PKI). Data is encrypted in transit using industry-standard ciphers. The default cipher for encryption is Blowfish-CBC, selected for its balance of speed and security. However, you can use any cipher supported by SSH, including AES-256-bit-CBC.

SoftNAS Cloud uses the IAM service to control the SoftNAS Cloud appliance's access to other AWS services.⁵ IAM roles are designed to allow applications to securely make API calls from an instance without requiring the explicit management and storage of access keys. When an IAM role is applied to an EC2 instance, the role handles key management, rotating keys periodically and making them available to applications through Amazon EC2 metadata.

Performance

The performance of a NAS system on Amazon EC2 depends on many factors, including the Amazon EC2 instance type, the number and configuration of [Amazon Elastic Block Store \(Amazon EBS\)](#) volumes,⁶ the type of Amazon EBS volume, the use of Provisioned IOPS with Amazon EBS, and the application workload. Benchmark your application on several Amazon EC2 instance types and storage configurations to select the most appropriate configuration.

SoftNAS Cloud provides Amazon Machine Images (AMIs) that support both paravirtual (PV) and hardware virtual machine (HVM) virtualization. To take advantage of special hardware extensions (CPU, network, and storage) and for optimal performance, SoftNAS recommends that you use a current generation instance type and an HVM AMI with single root input/output virtualization (SR-IOV) support.

To increase the performance of your system, you need to know which of the server's resources is the performance constraint. If CPU or memory limits your system performance, you can scale up the memory, compute, and network resources available to the software by choosing a larger Amazon EC2 instance type. Use StorageCenter dashboard performance charts and [Amazon CloudWatch](#) to monitor your performance and throughput metrics.⁷

Depending on the instance type and size chosen, EC2 instances are allocated varying amounts of CPU, memory, and network capabilities. Some instance families have higher ratios of CPU to memory, or higher ratios of memory to CPU. In general, to achieve the best performance from your SoftNAS Cloud virtual appliance, select an instance with large amounts of memory, up to 70 percent of which will be dedicated to high-speed dynamic random-access memory (DRAM) cache. If you require more than 120 MB/s NAS throughput for more demanding use cases, select an instance with advanced networking. AWS provides instances that support 10 and 20 Gbps network interfaces. If available,

choose an EBS-optimized instance, which uses a dedicated network path to EBS storage.

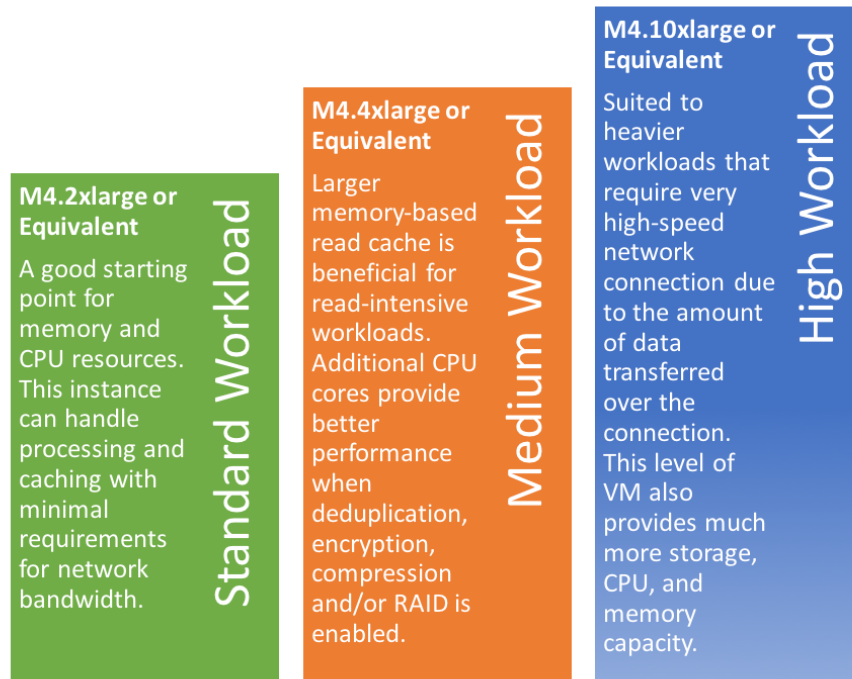
For production workloads, SoftNAS recommends starting with a larger EC2 instance size, coupled with monitoring of CloudWatch metrics as workloads are increased to their typical levels. This ensures applications have sufficient IOPS and throughput as they're brought online. Continue monitoring the application using SoftNAS StorageCenter and CloudWatch metrics, in particular CPU and network usage, to determine how well the chosen instance size is serving your unique workloads.

After a period of time (e.g., 30 days) with your workload in production, it will become apparent if the instance is well matched to the production workloads. As your load increases, if CPU or network usage reaches 75 percent or higher, you might need to increase instance size to achieve full throughput at low latencies. If CPU and network usage are below 40 to 50 percent, you can consider decreasing the instance size during a maintenance window to reduce operating costs.

SoftNAS does not recommend using T1 or T2 instances, as they are designed for burstable workloads and can run out of CPU credits during sustained heavy usage. At the time of this writing, SoftNAS recommends the m4.2xlarge as a minimum default AWS instance size, the m4.4xlarge for medium workloads, and the m4.10xlarge for heavier workloads, as seen in Figure 1 below. A SoftNAS representative can help with further sizing guidance.

About RAM Usage

SoftNAS Cloud allocates 50 percent of available RAM for use as Zettabyte File System (ZFS) file system cache. Remaining RAM is used by the Linux operating system, SoftNAS Cloud processes, and NAS services. It's typical to see 80 to 90 percent of RAM allocated and in use.



Later instance families also supported

Figure 1: AWS instance to workload

If your performance is limited by disk I/O, you can make configuration changes to improve the performance of your disk and caching resources.

Multilevel Cache

Read-intensive workloads benefit from additional RAM as level 1 cache (ZFS ARC), plus level 2 cache (ZFS L2ARC). Leverage the ephemeral SSD disks attached to certain EC2 instances to provide additional high-speed read cache.

Because data on ephemeral disks can be lost whenever an EC2 instance stops and restarts, or if underlying hardware changes or fails, use ephemeral disks only for read-cache purposes and never as a write log.

Amazon EBS Performance Optimizations

Because Amazon EBS is connected to an EC2 instance over the network, instances with higher network bandwidth can provide more Amazon EBS

performance. Some instance types support the Amazon EBS-optimized flag (ec2:EbsOptimized). This flag provides a dedicated network interface for Amazon EBS-bound traffic (storage I/O) and is designed to reduce variability in storage performance due to contention with network I/O. The chart [here](#) provides an outline of expected bandwidth, throughput, and Max IOPS per instance type and size.⁸

For SSD-based volume types, Amazon EBS measures an I/O operation as one that is 256 KB or smaller. I/O operations larger than 256 KB are counted in 256 KB increments. For example, a 1,024 KB I/O would count as four 256 KB IOPs. Amazon EBS also combines smaller I/O operations into a single operation where possible to achieve higher performance for all volume types.

Benefits of Each EBS Volume Type and Relevant Storage Application

Magnetic Backed

Magnetic-backed volume types support higher block sizes up to 1,024 KB.

Throughput Optimized HDD (st1) and Cold HDD (sc1) Amazon EBS volume types are based on magnetic storage technology.

The Throughput Optimized HDD (st1) volume type is designed for sequential read/write workloads (e.g., big data). It can achieve very high throughput (500 MB/s) for sequential read/write workloads (compared to 160 MB/s and 320 MB/s for SSD-backed gp2 and io1, respectively). Generally, big data workloads operate on very large sequential datasets and generate data for storage in a similar way. The st1 volume type has a baseline performance of 40 MB/s per terabyte (TB) of allocated storage and, like gp2, can burst beyond the baseline performance for a short period of time.

The Cold HDD (sc1) volume type is designed for high density and infrequent access workloads. This volume type is suitable for cold storage (infrequent access) applications where low cost is important. Unlike st1, the baseline performance of an sc1 volume is 12 MB/s per TB of allocated storage. It's important to note that Amazon S3 achieves high availability (HA) by default within a single region, whereas sc1 volumes have to be mirrored across Availability Zones to achieve parity with Amazon S3 in durability and availability of the data. (This doubles and triples the cost of sc1 when compared to Amazon S3.) Nevertheless, depending on certain access patterns (e.g., cold

versus warm) of the data, the cost of sc1 volumes can be cheaper for certain workloads.

SSD Backed

General Purpose (gp2) and Provisioned IOPS (io1) SSD volumes can achieve faster IOPS performance and very high throughput on random read/write workloads when compared to magnetic disks, but at a higher price point. However, gp2 and io1 volume types are limited to a throughput of ≤ 320 MB/s (160 MB/s for gp2, 320 MB/s for io1).

General Purpose (gp2) volumes provide a fixed 1:3 ratio between gigabytes and IOPS provisioned, so a 100 GB General Purpose volume provides a baseline of 300 IOPS. Gp2 volumes less than 1 TB in size can also burst for short periods, up to 3,000 IOPS. You can provision General Purpose volumes up to 16 TB and 10,000 IOPS.

Provisioned IOPS (io1) volumes are intended for workloads that demand consistent performance, such as databases. You can create Provisioned IOPS volumes up to 16 TB and 20,000 IOPS. Over a year, Amazon EBS Provisioned IOPS volumes are designed to deliver within 10 percent of the Provisioned IOPS performance 99.9 percent of the time.

There are differences in total throughput capabilities between Provisioned IOPS (io1) and General Purpose SSD (gp2) volumes. Io1 volumes are designed to provide up to 320 MB/second of throughput while gp2 volumes are designed to provide up to 160 MB/second.

RAID

If you need more I/O capabilities than a single volume can provide, you can create an array of volumes with redundant array of independent disks (RAID) software to aggregate the performance capabilities of each volume in the array.

For example, a stripe of two 4,000 IOPS volumes allows for a theoretical maximum of 8,000 IOPS. RAID 0 and RAID 10 are the two RAID levels recommended for use with Amazon EBS.

RAID 0, or striping, has the advantage of providing a linear performance increase with every volume added to the array (up to the maximum capabilities of the host instance). Two 4,000 IOPS volumes provide 8,000 IOPS, three

4,000 IOPS volumes provide 12,000 IOPS, and so on. However, because RAID 0 does not provide redundancy, it has less durability than a single volume. It also aggregates the failure rate of each volume in the array.

RAID 10 is a good compromise because it provides increased redundancy, aggregates the read performance of all volumes in the array, and provides a mirror of stripes in the array. However, RAID 10 isn't without drawbacks. There is a 50 percent penalty to write performance and a 50 percent reduction in available storage capacity. This penalty is due to half of the disks in the array being reserved for a mirror. RAID 10 has the same write penalty as RAID 1.

RAID 5 and 6 are not recommended because parity calculations incur significant overhead without dramatically increasing the durability of the volume set. Such a large write penalty makes these RAID levels very expensive to run in terms of both dollars and I/O.

In general, RAID using mirroring or parity for increased durability adds extra steps and reduces performance, while not necessarily increasing the data's durability. Amazon EBS has its own durability mechanisms. It can be supplemented with Amazon S3-backed snapshots and SoftNAS replication to more than one Availability Zone.

DRAM cache can dramatically increase read IOPS performance. Choose instances with more memory for the best read IOPS and throughput. For an even larger read cache, choose instance types with ephemeral SSD locally attached disks and attach an SSD cache device to each storage pool. To ensure their availability, attach local SSD ephemeral disks to the SoftNAS instance when you create the instance.

Many instance types provide instance-store or "ephemeral" volumes. Although SoftNAS doesn't support the use of these volumes for dataset storage, you can use them as a read cache for storage pools. These volumes are located physically inside the underlying host of the instance and are not affected by performance variability from network overhead. Although this varies by instance type, most instance-store volumes (especially on newer instance types) are SSD volumes. However, stopping and starting an instance can move it to another underlying host, which causes all data on these volumes to be lost. This isn't an issue for caching, but is detrimental for dataset storage.

If you require additional write caching or IOPS, you can attach SSD-backed Amazon EBS volumes to a storage pool. The use of locally attached ephemeral disks for write cache isn't recommended.

Consider your workload requirements and priorities. If the amount of storage and cost take priority over speed, magnetic EBS volumes might be the right choice. General Purpose SSD or Provisioned IOPS volumes offer the best mix of price, performance, and total storage space. With AWS and SoftNAS Cloud, you can add more storage or configure a different type of storage on the fly to address a variety of price or performance needs.

Using Amazon S3 with SoftNAS Cloud

SoftNAS Cloud provides support for a feature known as SoftNAS S3 Cloud Disks. These are abstractions of Amazon S3 storage presented as block devices. By leveraging Amazon S3 storage, SoftNAS Cloud can scale cloud storage to practically unlimited capacity. You can provision each cloud disk to hold up to four petabytes (PB) of data. If a larger data store is required, you can use RAID to aggregate multiple cloud disks.

Each SoftNAS S3 Cloud Disk occupies a single Amazon S3 bucket in AWS. The administrator chooses the AWS Region in which to create the S3 bucket and cloud disk. For best performance, choose the same region for both the SoftNAS Cloud EC2 instance and its S3 buckets.

SoftNAS Cloud storage pools and volumes using cloud disks have the full, enterprise-grade NAS features (for example, deduplication, compression, caching, storage snapshots, and so on) available, and can be readily published for shared access through NFS, CIFS, AFP, and iSCSI.

When you use a cloud disk, use a block device local to the SoftNAS Cloud virtual appliance as a read cache to reduce Amazon S3 I/O charges and improve IOPS and performance for read-intensive workloads.

For best S3 cloud disk performance and security, specify an S3 endpoint within the VPC in which you deploy SoftNAS Cloud. The S3 endpoint ensures S3 traffic is optimally routed through the VPC and not across the NAT gateway or Internet, which is slower and less secure.

You can also encrypt S3 cloud disks to protect all Amazon S3 I/O, should it need to travel over the Internet or outside a VPC (e.g., from on premises or across regions).

Network Security

Amazon VPC is a logically separated section of the AWS Cloud that provides you with complete control over the networking configuration. This includes the provisioning of an IP space, subnet size and scope, access control lists, and route tables. You can configure subnets inside an Amazon VPC as either public or private. The difference between public and private subnets is that a public subnet has a direct route to the Internet; a private one does not. When you configure an Amazon VPC for use with SoftNAS Cloud, consider the level of access that your use case requires. If the SoftNAS Cloud virtual appliance doesn't need to be accessed from the Internet, consider placing it in private Amazon VPC subnets.

To leverage SoftNAS S3 Cloud Disks, the SoftNAS Cloud virtual appliance must have a way to access the S3 bucket, either through the Internet or a configured VPC endpoint.

A VPC Security Group acts as a virtual firewall for your instance to control inbound and outbound traffic. For each Security Group, you add rules that control the inbound traffic to instances and a separate set of rules that control the outbound traffic. Open only those ports that are required for the operation of your application. Restrict access to specific remote subnets or hosts.

For a SoftNAS Cloud installation, determine which ports must be opened to allow access to required services. These ports can be divided into three categories: management, file services, and high availability.

Open the following ports to manage SoftNAS Cloud through the SoftNAS StorageCenter and SSH. As the following table indicates, you should limit the source to hosts and subnets where management clients are located.

Management

Type	Protocol	Port	Source
SSH	TCP	22	Management
HTTPS	TCP	443	Management

When providing file services, first determine which services you will provide. The following tables show which ports to open for security group configuration. As the tables indicate, the source should be limited to the clients and subnets that consume these services.

AFP

Type	Protocol	Port	Source
Custom TCP Rule	TCP	548	Clients
Custom TCP Rule	TCP	427	Clients

NFS

Type	Protocol	Port	Source
Custom TCP Rule	TCP	111	Clients
Custom TCP Rule	TCP	2010	Clients
Custom TCP Rule	TCP	2011	Clients
Custom TCP Rule	TCP	2013	Clients
Custom TCP Rule	TCP	2014	Clients
Custom TCP Rule	TCP	2049	Clients
Custom UDP Rule	UDP	111	Clients

Custom UDP Rule	UDP	2010	Clients
Custom UDP Rule	UDP	2011	Clients
Custom UDP Rule	UDP	2013	Clients
Custom UDP Rule	UDP	2014	Clients
Custom UDP Rule	UDP	2049	Clients

CIFS/SMB

Type	Protocol	Port	Source
Custom TCP Rule	TCP	137	Clients
Custom TCP Rule	TCP	138	Clients
Custom TCP Rule	TCP	139	Clients
Custom UDP Rule	UDP	137	Clients
Custom UDP Rule	UDP	138	Clients
Custom UDP Rule	UDP	139	Clients
Custom TCP Rule	TCP	445	Clients
Custom TCP Rule	TCP	135	Clients

Active Directory Integration

Type	Protocol	Port	Source
LDAP	TCP	389	Clients

iSCSI

Type	Protocol	Port	Source
Custom TCP Rule	TCP	3260	Client IPs

The following security group configuration is required when you deploy SoftNAS SNAP HA, which is discussed later in this whitepaper. As the table indicates, you should limit the source to the IP addresses of the SoftNAS Cloud virtual appliance.

High Availability with SNAP HA™

Type	Protocol	Port	Source
Custom ICMP Rule	Echo Reply	22	SoftNAS Cloud IPs or Security Group ID*
Custom ICMP Rule	Echo Request	443	SoftNAS Cloud IPs or Security Group ID*

* <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

Data Protection Considerations

Creating a comprehensive strategy for backing up and restoring data is complex. In some industries, you must consider regulatory requirements for data security, privacy, and records retention. SoftNAS Cloud provides multiple capabilities for data redundancy.

Always have one or more independent data backups, beyond the data redundancy provided by SoftNAS Cloud. You can back up data disks using EBS snapshots and third-party backup tools to create offsite or other backup copies to protect data.

SoftNAS Cloud provides multiple levels of data protection and redundancy, but it isn't intended to replace normal data backup processes. Instead, these levels of redundancy and data protection reduce risks associated with data loss or data

integrity, and provide features that enable rapid recovery, often without the need to restore from a backup copy.

SoftNAS Cloud is Copy-On-Write (COW) File System

SoftNAS Cloud leverages the reliable, mature ZFS. ZFS is a copy-on-write file System, which means that existing data is never directly overwritten. Instead, new data blocks are appended to each file, conceptually similar to a tape.

Figure 2 depicts how the file System inside SoftNAS Cloud maintains multiple versions, known as storage snapshots, without overwriting the existing data.




Figure 2: Copy-on-write file system

Automatic Error Detection and Correction

SoftNAS Cloud automatically detects and corrects unforeseeable data errors. These errors can occur over time for many different reasons including bad sectors, network, or other I/O errors. SoftNAS Cloud also provides protection against potential “bit rot”, disk media deterioration over time caused by magnetism decay, cosmic ray effects, and other sporadic issues that can cause data storage or retrieval errors.

Proven ZFS data integrity measures are implemented by SoftNAS Cloud to detect errors, repair them automatically, and ensure data integrity is maintained. Each read is validated against a 256-bit checksum code. When

errors are detected, the system automatically repairs the block with the corrected data transparently, so applications aren't affected and data integrity is maintained. Periodically, administrators can “scrub” storage pools to provide even higher levels of data integrity.

SoftNAS Cloud Snapshots

SoftNAS Cloud snapshots are volume-based, point-in-time copies of data. StorageCenter provides a rich set of snapshot scheduling and on-demand capabilities. As Figure 3 shows, snapshots provide file system versioning.

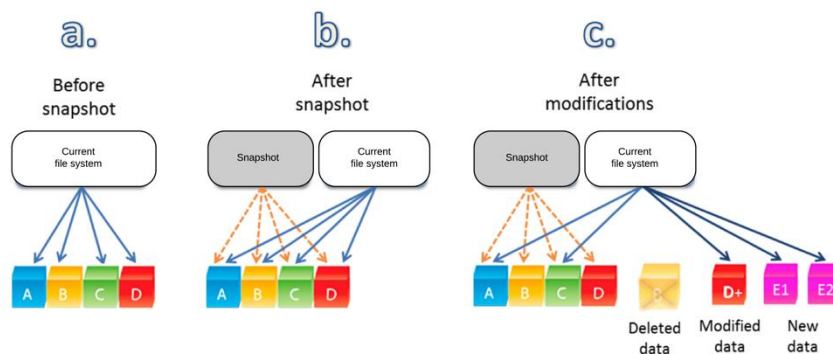


Figure 3: SoftNAS Cloud volume-based snapshots

SoftNAS Cloud snapshots are integrated with Windows Previous Versions, which is provided through the Microsoft Volume Shadow Copy Service (VSS) API. This feature is accessible to Windows operating system users through the **Previous Versions** tab, so IT administrators don't need to assist in file recovery. Microsoft server and desktop operating system users can use scheduled snapshots to recover deleted files, view or restore a version of a file that was overwritten, and compare file versions side by side. Operating systems that are supported include Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012.

Snapshots consume storage pool capacity, so you must monitor usage for over-consumption. Storage snapshots grow incrementally as file system data is modified over a period of time. SoftNAS Cloud automatically manages snapshots based on snapshot policies to prevent snapshots from consuming all available space. Several snapshot policies are provided as a starting point, and you can also create custom snapshot policies. Snapshot policies are independent

of each volume, so when a snapshot policy is changed it's applied across all volumes that reference that policy.

When allocating storage pool space and choosing snapshot policies, be sure to plan for enough additional storage to hold the snapshot data for the retention period you require.

SoftNAS SnapClones™

SnapClones provide read/write clones of SoftNAS Cloud snapshots. They're created instantaneously because of the space-efficient, copy-on-write model. Initially, SnapClones take up no capacity and grow only when writes are made to the SnapClone, as shown in Figure 4. You can create any number of SnapClones from a storage snapshot.

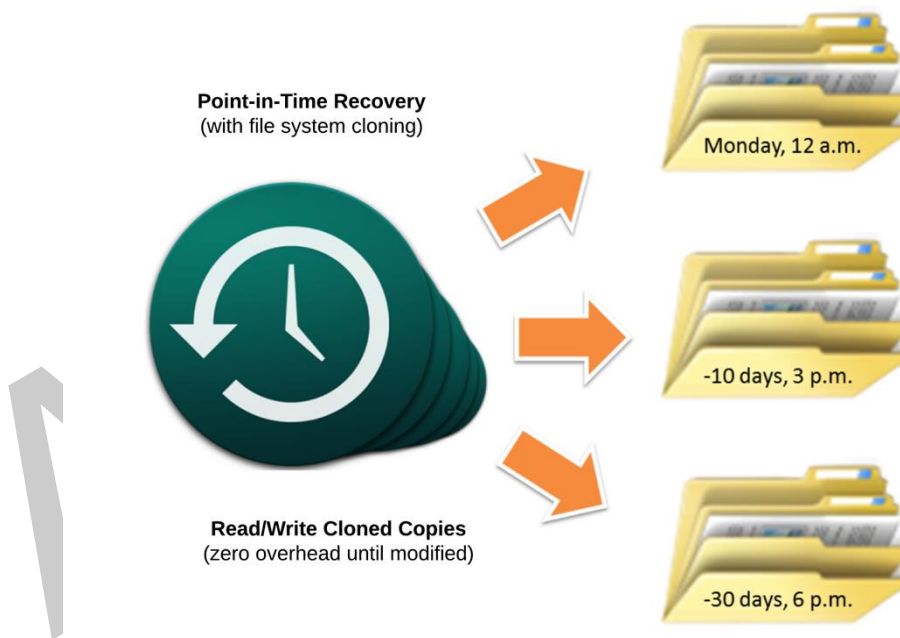


Figure 4: SoftNAS SnapClones

You can mount SnapClones as external NFS or CIFS shares. They're good for manipulating copies of data that are too large or complex to be practically copied. For example, testing new application versions against real data and selective recovery of files and folders using the native file browsers of the client operating system. You can create a SnapClone instantly, even for very large datasets in the tens to hundreds of TBs.

Amazon EBS Snapshots

SoftNAS Cloud has a built-in capability to leverage Amazon EBS point-in-time snapshots to back up EBS-based storage pools. The Amazon EBS snapshot copies the entire SoftNAS Cloud storage pool, for backup and recovery purposes. Advantages include the ability to use the AWS Management Console to manage the snapshots. Capacity for the Amazon EBS snapshots isn't counted against the storage pool capacity. You can use Amazon EBS snapshots for longer-term data retention.

Deployment Scenarios

The design of your SoftNAS Cloud installation on Amazon EC2 depends on the amount of usable storage and your requirements for IOPS and availability.

High-Availability Architecture

To realize high availability for storage infrastructure on AWS, SoftNAS strongly recommends implementing SNAP HA in a high-availability configuration. The SNAP HA functionality in SoftNAS Cloud provides high availability, automatic, and seamless failover across Availability Zones.

SNAP HA leverages secure block-level replication provided by SoftNAS SnapReplicate to provide a secondary copy of data to a controller in another Availability Zone. SNAP HA also provides both automatic and manual failover.

High availability and cross-zone replication eliminates or minimizes downtime. It is not, however, intended to replace regular data backups, which are always required to fully protect important data, especially in disaster recovery scenarios.

There are two methods for achieving high availability across zones: Elastic IP (EIP) addresses and SoftNAS Cloud Private Virtual IP-based HA.

The use of Private Virtual IP-based HA is recommended for best security, performance, and lowest cost. All NAS traffic remains inside the VPC.

Support for EIP is available for situations that require a “routable” IP address, or the rare cases where data shares must be made available over the Internet. Of course, access via EIP addresses can be locked down using Security Groups.

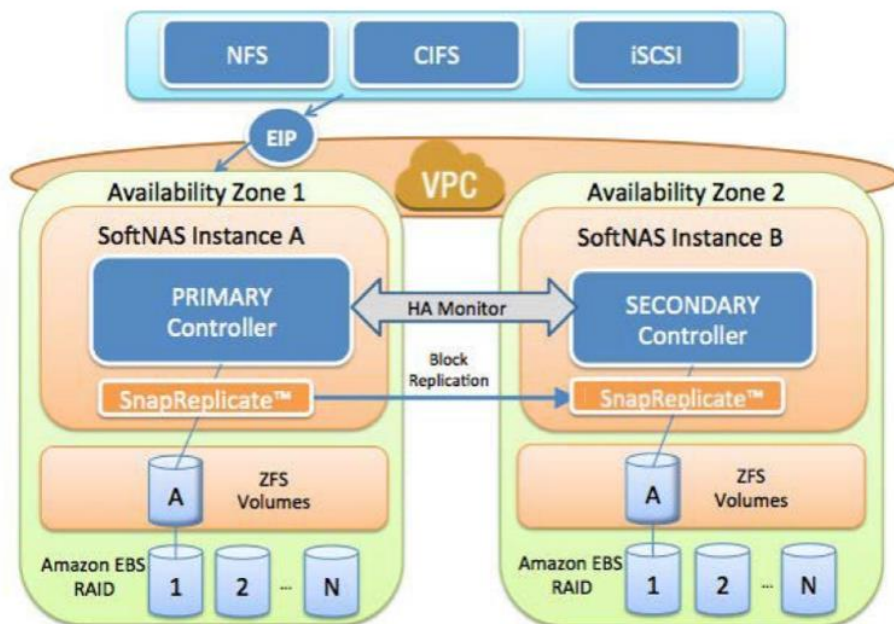


Figure 5: Task creation and result aggregation

Multi-AZ HA operates within a VPC. Optionally, you can route NAS traffic through a floating EIP combined with SoftNAS patented⁹ HA technology. That is, NFS, CIFS, AFP, and iSCSI traffic are routed to a primary SoftNAS controller in one Availability Zone, and a secondary controller operates in a different Availability Zone. NAS clients can be located in any Availability Zone.

SnapReplicate performs block replication from the primary controller A to the backup controller B, keeping the secondary updated with the latest changed data blocks once per minute. In the event of a failure in Availability Zone 1 (shown in Figure 5), the Elastic HA™ IP address automatically fails over to controller B in Availability Zone 2 in less than 30 seconds. Upon failover, all NFS, CIFS, AFP, and iSCSI sessions reconnect with no impact on NAS clients (that is, no stale file handles and no need to restart).

HA with Private Virtual IP Addresses

The patented⁹ Virtual IP-based HA technology in SoftNAS Cloud enables you to deploy two SoftNAS Cloud instances across different Availability Zones inside the private subnet of a VPC. Then you can configure the SoftNAS Cloud instances with private IP addresses, which are completely isolated from the Internet. This allows for more flexible deployment options and greater control over access to the appliance. In addition, using private IP addresses enables faster failover because waiting for an EIP to switch instances isn't required. Further, Virtual IP HA is less costly because there is no I/O flowing across an EIP. Instead, all traffic remains completely within the VPC.

For most use cases, Multi-AZ HA using private virtual IP addresses is the recommended method. Failover usually takes place in 15 to 20 seconds from the time a failure is detected. SoftNAS Cloud uses patented⁹ techniques that allow NAS clients to stay connected via NFS, CIFS, iSCSI, and AFP in case of a failover, ensuring that services are not interrupted and continue to operate without downtime.

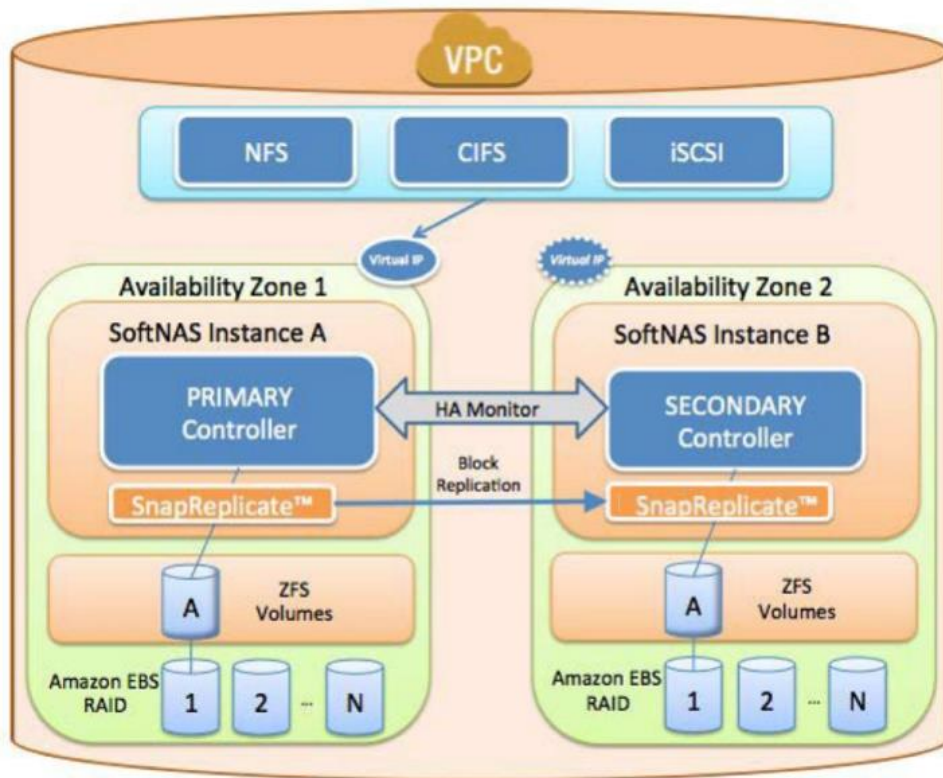


Figure 6: Cross-zone HA with virtual private IP addresses

For more information about implementation and HA design best practices, see the [SoftNAS High Availability Guide](#).¹⁰

Single Controller Architecture

In scenarios where you don't require high availability, you can deploy a single controller.

Figure 7 shows a basic SoftNAS Cloud instance running within a VPC.

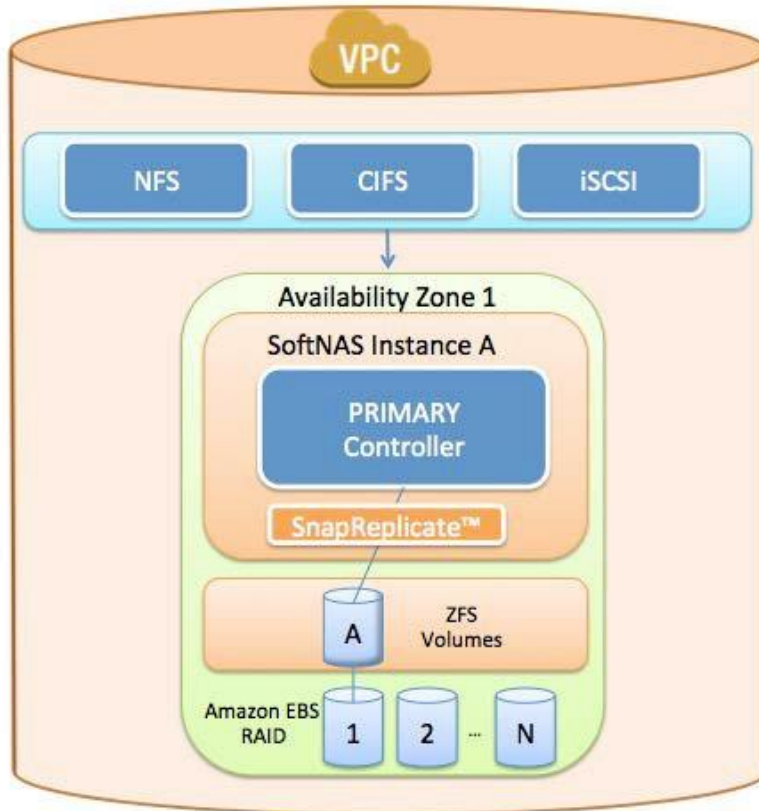


Figure 7: Basic SoftNAS Cloud instance running within a VPC

In these scenarios, you can combine EBS volumes into a RAID 10 array for the storage pool to provide usable storage space with no drive failure redundancy. You can also configure storage pools using a SoftNAS S3 Cloud Disk for RAID 0 (striping) for improved performance and IOPS. These examples are for illustration purposes only. Typically RAID 0 is sufficient, as the underlying EBS and S3 storage devices already provide redundancy.

Volumes are provisioned from the storage pools and then shared through NFS, CIFS/SMB, AFP, or iSCSI.

Hybrid Cloud Architecture

You can deploy SoftNAS Cloud in a Hybrid Cloud architecture in which a SoftNAS Cloud virtual appliance is installed both in Amazon EC2 and on premises. This architecture enables replication of data from on premises to Amazon EC2 and vice versa, providing synchronized data access to users and

applications. Hybrid Cloud architectures are also useful for backup and disaster recovery scenarios in which AWS can be used as an off-site backup location.

Replication

You can deploy SoftNAS Cloud in Amazon EC2 as a replication target using SnapReplicate. This enables scenarios such as data replicas, disaster recovery, and development environments by copying on-site production data into Amazon EC2, as shown in Figure 8.

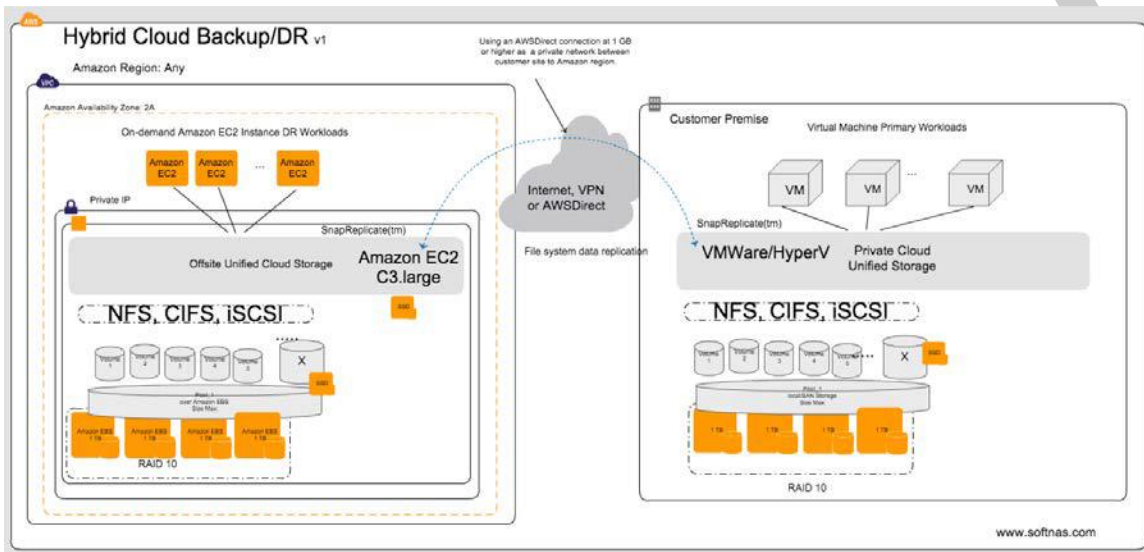


Figure 8: Hybrid Cloud backup and disaster recovery

File Gateway to Amazon S3

You can deploy SoftNAS Cloud in file gateway use cases, where SoftNAS Cloud operates on premises, deployed in local data centers on popular hypervisors, such as VMware vSphere. SoftNAS Cloud connects to Amazon S3 storage, treating Amazon S3 as a disk device. The Amazon S3 disk device is added to a storage pool where volumes can export CIFS, NFS, AFP, and iSCSI. Amazon S3 is cached with block disk devices for read and write I/O. Write I/O is cached at primary storage speeds and then flushed to Amazon S3 at the speed of the WAN. When using Amazon S3-based volumes with backup software, the write cache dramatically shortens the backup window.

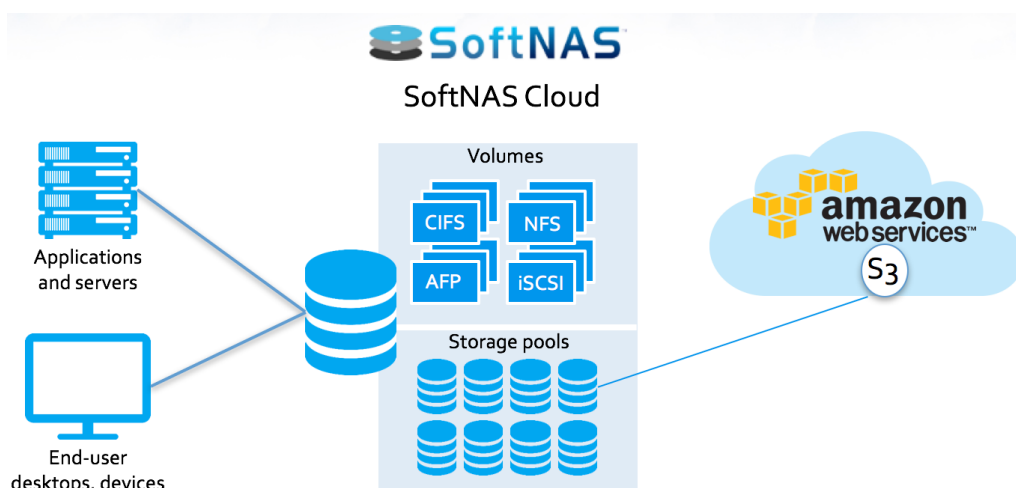


Figure 9: SoftNAS Cloud

Automation Options

This section describes how the SoftNAS Cloud REST API, CLI, and [AWS CloudFormation](#) can be used for automation.

API and CLI

SoftNAS Cloud provides a secure REST API and CLI. The REST API provides access to the same storage administration capabilities from any programming language using HTTPS and REST verb commands, returning JSON-formatted response strings. The CLI provides command line access to the API set for quick and easy storage administration. Both methods are available for programmatic storage administration by DevOps teams who want to design storage into automated processes. For more information, see the [SoftNAS API and CLI Guide](#).¹¹

AWS CloudFormation

The AWS CloudFormation service enables developers and businesses to create a collection of related AWS resources and provision them in an orderly and predictable way.¹²

SoftNAS Cloud provides sample CloudFormation templates that you can use for automation. You can find these templates [here](#) and in the [Further Reading](#) section of this paper. When you work with CloudFormation templates, pay

careful attention to the **Instance Type**, **Mappings**, and **User Data** sections, which are shown in the following examples.

List all the instance types that you want to appear. Edit this section with the latest instance types available.

```
InstanceType:
  Type: String
  Default: m4.2xlarge
  Description: Softnas HVM EC2 instance type
  AllowedValues:
    - m4.2xlarge
    - m4.4xlarge
    - m4.10xlarge
    - m3.2xlarge
    - c4.large
    - c4.xlarge
    ...
```

Map to the appropriate AMIs here (SoftNAS regularly updates AMIs, so this section must be updated accordingly).

```
Mappings:
  RegionMap:
    ap-northeast-1:
      AMI: ami-xxx
    ap-southeast-1:
      AMI: ami-xxx
    ap-southeast-2:
      AMI: ami-xxx
    eu-central-1:
      AMI: ami-xxx
    eu-west-1:
      AMI: ami-xxx
    us-east-1:
      AMI: ami-xxx
    us-west-1:
      AMI: ami-xxx
    us-west-2:
      AMI: ami-xxx
    ...
```

This section is used to pass variables to the SoftNAS Cloud CLI for additional configuration.

```
UserData:
"Fn::Base64":
  !Sub |
    #!/bin/bash -v
    # Configure NFS / CIFS Shares
    /var/www/softnas/scripts/initproc.sh 2>&1
    wget http://www.softnas.com/docs/softnas/v2/api/softnas-
cmd.zip
    unzip softnas-cmd.zip
    mv softnas-cmd /usr/local/bin/ \n", "chmod 755
/usr/local/bin/softnas-cmd
    INSID=`curl http://169.254.169.254/latest/meta-
data/instance-id`
    /usr/local/bin/softnas-cmd login softnas $INSID --
base_url https://localhost/softnas --pretty_print >>
/tmp/cf.tmp 2>&1
    /usr/local/bin/softnas-cmd createpool
/dev/xvdj:/dev/xvdk pool1 1 on -t >> /tmp/cf.tmp 2>&1
    /usr/local/bin/softnas-cmd createvolume volume1 pool1
filesystem thin exportNFS=on shareCIFS=on dedup=on
enable_snapshot=on schedule_name=Default hourlysnaps=5
daily snaps=10 weeklysnaps=0 -t >> /tmp/cf.tmp 2>&1
```

Conclusion

SoftNAS Cloud is a popular NAS option on the AWS Cloud computing platform. By following the implementation considerations and best practices highlighted in this paper, you will maximize the performance, durability, and security of your SoftNAS Cloud implementation on AWS.

For more information about SoftNAS Cloud, see www.softnas.com.

Get a [free 30-day trial](#) of SoftNAS Cloud now.¹³

Contributors

The following individuals and organizations contributed to this document:

- Eric Olson, VP Development, SoftNAS
- Kevin Brown, Solutions Architect, SoftNAS

- Brandon Chavis, Solutions Architect, Amazon Web Services
- Juan Villa, Solutions Architect, Amazon Web Services
- Ian Scofield, Solutions Architect, Amazon Web Services

Further Reading

SoftNAS References

[SoftNAS Cloud Installation Guide](#)

[SoftNAS Reference Guide](#)

[SoftNAS Cloud High Availability Guide](#)

[SoftNAS Cloud API and Cloud Guide](#)

AWS CloudFormation Templates for [HVM](#)

Amazon Web Services References

[Amazon Elastic Block Store](#)

[Amazon EC2 Instances](#)

[AWS Security Best Practices](#)

[Amazon Virtual Private Cloud Documentation](#)

[Amazon EC2 SLA](#)

Notes

¹ <http://aws.amazon.com/ec2/>

² <http://www.softnas.com/>

3 <http://aws.amazon.com/s3/>

4 <http://aws.amazon.com/vpc/>

5 <http://aws.amazon.com/iam/>

6 <http://aws.amazon.com/ebs/>

7 <http://aws.amazon.com/cloudwatch/>

8

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html#ebs-optimization-support>

9 U.S. Pat. Nos. 9,378,262; 9,584,363. Other patents pending.

10 <https://www.softnas.com/docs/softnas/v3/snapha-html/index.htm>

11 <https://www.softnas.com/docs/softnas/v3/api-html/>

12 <http://aws.amazon.com/cloudformation/>

13 http://softnas.com/trynow?utm_source=aws&utm_medium=white-paper&utm_campaign=aws-wp-2017