

StorageCenter™	Web-based management console for managing SoftNAS Cloud® software
Update Proxy	Single Firewall Port Software Updates: Doesn't require multiple open firewall ports for connection to different software update depots
No Storage Downtime Guarantee™	SoftNAS Cloud® storage SLA when either DCHA™ or SNAP HA® is being used
Petabyte scale	Scales data storage from terabytes to 16 petabytes or greater
Multiple Storage Protocol Support	POSIX compliant file access to backend block and object storage via NFS, CIFS/SMB (with Active Directory), AFP (Apple File Protocol) and iSCSI block services
360-degree Encryption™ (In-Transit & At-Rest)	Military-grade Encryption for data in-transit and data at-rest (only customer controls keys)
Active Directory and LDAP Integration	Integration with Access Control systems to prevent un-authorized access
Compression	Amount of storage that data consumes is “shrunk” by removing extra “white space” between the blocks of data, thereby reducing the overall storage required
Inline Deduplication	Duplicated data is eliminated allowing for a more efficient and cost effective use of storage
Snapshots and Recovery	Scheduled or manual images of storage (snapshots) in order to rollback data (recovery) to a point-in-time
Dual Controller HA (DCHA™) – Object storage	Application resiliency across zones / regions and Data availability in a single zone / region
ObjFast™ / ObjectBacker™ (Patent pending)	An acceleration method that speeds up the reads, writes & deletion of data located object storage to near block level storage performance
Object storage support	Provides NAS file service support to native object storage so applications can use it without code changes
SSD read cache support	Use of a solid state drive (SSD), that provides an additional layer of cache, in addition to RAM memory cache
SSD write log support	Use of a SSD, preferably in a RAID 1 (mirror), to provide caching for incoming writes to be eventually written to lower-speed hard disk drive (HDD) storage
Block storage support	Provides NAS file service support to native block storage so applications can use it without code changes
SNAP HA® – Block and/or Object storage	Maximum uptime with both Application and Data resiliency across two zones / regions
SnapReplicate™	Replicate data from one storage pool to a duplicate storage pool. Used with either Block and/or Object storage
DeltaSync®	Reduce the Recovery Time Objective (RTO) to hours for cluster recovery from a high-availability failover event
SnapClone®	Create a new volume from a volume snapshot in order to recover from an event or for DevOps to test with
SmartTiers™ (Patent pending)	Automatically move data (Auto-Tiering) that is less frequently accessed to less expensive/performant storage - with savings up to 67% of public cloud storage costs
UltraFast™ (Patent pending)	Used to bulk transfer data from 1 or more locations to 1 or more other locations at speeds up to 20x faster than TCP/IP
FlexFiles™	A “drag and drop” dev environment that provides a seamless experience between design, control, feedback and monitoring for data integration and movement
Lift and Shift	An advanced cloud data management wizard that migrates on-premises file data to the public cloud
Apache NiFi	An open source, powerful and user friendly, “drag-and-drop” data ingestion and integration capability

**Pain Points**

**Benefits**

Regain control of user data and end server sprawl

Data is no longer stored centrally on file servers IT controls — it's everywhere and hard to manage and secure

Consolidate on-premises file servers

On-demand access to my tape archived data in the cloud

Tapes go bad over time and it takes days or even weeks to get tape data shipped back when we need it

Move archived data from tape to cloud

Store my backup files in the cloud

Data is growing so fast that backups are overrunning your local data storage and cloud backup is too slow

Secure NFS mount for on-premises Veeam plus highly durable cloud object storage

Predictable and reliable performance from my cloud workloads

Latency issues and unreliable performance in shared, multi-tenant public cloud NAS services

Balance compute and storage to create an optimal cost/performance cloud environment

Maximum uptime for my cloud workloads

Data loss and outages cause major harm to my business reputation and can cost us millions

Cross-zone high availability for maximum uptime

Improve data security and access control

A lack of granular data security and control is putting my customer data, business data and IP at risk

Data is encrypted in-transit and at-rest plus Active Directory integration

Highly-available NFS and CIFS file sharing in the cloud

Native cloud storage (block and object) doesn't support the file systems & storage protocols we use

Eliminate application reengineering or rewrites and cut cloud migration time by 90%

Move mountains of data to the cloud

Slow data movement and connection time outs caused by throttling and reliability issues impact our projects

Move data up to 20X faster than TCP/IP can over standard network connections

Move my data to the cloud fast and easily

Frequent HTTP 500 errors and forced restarts because of congestion, unreliable networks and latency issues

Move massive amounts of data with unmatched speed and reliability and keep it in sync

Get out of the data center and hardware business

Storage hardware is expensive to constantly purchase, deploy, and manage

Eliminate application reengineering or rewrites and cut cloud migration time by 90%

Move an existing SaaS business to the public cloud

Poor performance, limited scalability and high cloud file storage costs impact my business

Scale storage cost and performance up or down on-demand project by project

Move my legacy applications to the cloud

It will take too long and cost too much to re-architect or rewrite my existing applications

Get to the cloud in hours to days, not months or years by avoiding custom coding

Publish content from a central location to remote locations

I need a cost-effective, fast and reliable way to publish content in a variety of formats to a large number of locations

Integrate and transform data in a single process. Supports 25+ data file formats

Gather data from multiple locations and store it centrally in the cloud

Aggregating data content of various types and formats from many locations synchronizing it is impossible

Move data up to 20 times faster than TCP/IP can over standard network connections

Connect my remote offices, branch offices and factories with the cloud

I have "islands of data", making it difficult to gather all my data into one central cloud repository

Store offsite copies of my Veeam backups in the cloud

Meeting regulatory requirements for offsite backups with Veeam is too complex and expensive

Store my Veeam backups in the cloud because I'm out of space

Doing on-premises backups is impossible due to data sprawl and storage costs

Run cloud backups at near-block-level-performance at object storage pricing

Quickly restore VMware VMs from Veeam backups in the cloud

I store VMware backups in the cloud, but performance and timeout issues make it unreliable and slow