

Veritas Access 3340 Appliance for Long Term Retention of Pure Storage® FlashArray™ Snapshots

This white paper provides a technical overview of Veritas Access Appliance as a secondary storage solution for Pure Storage FlashArray volume snapshots. It highlights the overall solution architecture components, integration flow, best practices, sizing guidance, and deployment of Access Appliance with FlashArray solution.



TABLE OF CONTENTS

| | |
|--|----|
| INTRODUCTION | 4 |
| EXECUTIVE SUMMARY | 4 |
| SCOPE | 4 |
| TARGET AUDIENCE | 4 |
| SOLUTION VALUE | 4 |
| SOLUTION ARCHITECTURE OVERVIEW | 5 |
| FLASHARRAY | 5 |
| ACCESS APPLIANCE | 6 |
| SOLUTION INTEGRATION | 8 |
| Snap to NFS | 9 |
| Snapshot Retrieval | 10 |
| SOLUTION SECURITY | 11 |
| DISASTER RECOVERY | 11 |
| BEST PRACTICES AND RECOMMENDATIONS | 12 |
| DATA LAYOUT ON ACCESS APPLIANCE | 12 |
| NFS PROTOCOL | 12 |
| COMPRESSION | 13 |
| NETWORK CONNECTIVITY | 13 |
| MONITORING | 13 |
| SIZING GUIDANCE | 13 |
| EXAMPLE CAPACITY SIZING OF ACCESS APPLIANCE FOR FLASHARRAY SNAPSHOTS | 14 |
| PERFORMANCE OF ACCESS APPLIANCE FOR FLASHARRAY | 14 |
| CONCLUSION | 15 |
| REFERENCES | 16 |
| APPENDIX | 17 |
| SOLUTION DEPLOYMENT | 17 |
| Access Appliance Storage Configuration and Provisioning | 17 |
| Configure Storage Pools | 18 |
| Enabling NFS | 20 |
| Creation of File System to Export as NFS Share | 21 |
| Modification of NFS Share Permissions | 26 |
| Configuration of FlashArray to Utilize Access as an NFS Target | 27 |
| Validation of Configuration | 29 |
| Snap | 29 |
| PERFORMANCE STUDY | 33 |
| Testing Strategy | 33 |
| Results | 35 |

Revision History

| Version | Date | Changes |
|---------|---------|-----------------|
| 1.00 | 06/2019 | Initial Version |

INTRODUCTION

EXECUTIVE SUMMARY

Pure Storage® FlashArray™ is an all-flash block storage designed for high performance and varied workloads such as databases, enterprise applications, and virtual environments. One of the features of FlashArray is the ability to copy volume snapshots to secondary storage for data protection, migration, cloning, and test and development operations. Access Appliance is a complementary cost-optimized storage option for offloading data volume snapshots residing on Pure Storage all-flash storage products. The Access 3340 Appliance is a low-cost, disk-based solution that is easy to manage and designed for long term retention, tape replacement, backup, and archival workloads.

SCOPE

The purpose of this document is to provide technical details to assist in understanding Access Appliance as a solution for the preservation of FlashArray volume snapshots for long-term retention. It describes the components of this solution, its value, sizing guidance, and some best practices. It is advised to refer to Veritas and Pure Storage product documentation for installation, configuration and administration of each of the products discussed in this whitepaper. **NOTE:** This document is updated periodically and is also available [here](#).

TARGET AUDIENCE

This document is targeted for joint customers, partners, and field personnel interested in learning more about the use of Veritas Access Appliance as a secondary storage for Pure Storage FlashArray volume snapshots. It provides a technical overview of this solution, guidance in sizing, and highlights some best practices.

SOLUTION VALUE

Pure FlashArray provides primary storage for organizations' critical digital assets. Therefore, having a strategy to protect primary data in case of a failure, disaster, or crisis is imperative for business continuity. Inherent to Pure FlashArray is ability to create a point-in-time snapshot for data protection, migration, and testing. To save space for primary workloads, FlashArray has the capability to offload the snapshots to secondary storage for long-term retention. There are several challenges that come to mind when talking about a long-term solution which include cost, complexity, control, visibility, and security. To address all these challenges, Veritas has designed the Access Appliance as a purpose-built, on-premises disk-based storage appliance for long-term retention use cases. The Access Appliance provides a resilient and cost-effective solution for the preservation of FlashArray data snapshots that companies want to retain and have readily available for further use.

Key benefits for utilizing Access Appliance for preservation of FlashArray snapshots are:

- **Minimize costs** – Access Appliance provides a low-cost, disk-based solution that is easy to manage. FlashArray creates efficient and portable snapshots which are compressed, further reducing the storage footprint on the Access Appliance and overall costs.
- **Increase visibility and control** - Having the data on-premises under the company's control and visibility allows for quicker restores.
- **Security** – The Access Appliance encrypts FlashArray snapshots utilizing an external key management system thereby providing a secure secondary storage solution.

- **Replication** – FlashArray snapshots stored on an Access Appliance can be replicated to another Access Appliance at a remote site for disaster recovery purposes.

SOLUTION ARCHITECTURE OVERVIEW

As shown in Figure 1, FlashArray is primary storage for enterprise applications, databases, files, virtual images, etc. and has capability to take point in time snapshots of data volumes. In some instances, customers may want an option to copy snapshots on their high-performance storage array and tier to low-cost, on-premises storage for secondary or long-term retention such as the Access Appliance. FlashArray has a feature called “Snap to NFS” that sends volume snapshots to the Access Appliance via NFS protocol. The following sections discuss each of these components and then the flow of the solution integration.

Figure 1 - FlashArray with the Access Appliance



FLASHARRAY

Pure Storage FlashArray is an all-flash array (100% NVMe) block storage platform optimized for mission critical (Tier 1) workloads requiring ultra-low latency and high performance. Built-in features include inline data reduction (on average, 5:1), snapshots, replication, encryption, and business continuity. In addition, it has the capability to run virtual machines and containers on the FlashArray. The FlashArray current model is the //X series which ranges from 55 TB to 3 PB of effective capacity, where effective capacity assumes high availability, RAID, and metadata overhead, GB to GiB conversion, data reduction, compression and pattern removal. For more details, refer to the FlashArray//X [product web page](#) and [datasheet](#).

Integrated and native to FlashArray is the snapshot capability for efficient data protection of its volumes. FlashArray snapshots are immutable, point-in-time images of the contents of one or more volumes. For simplicity and ease of management, FlashArray offers protection groups. FlashArray protection groups provide automatic snapshot and replication scheduling as well as retention options for volume(s).

Although the amount of storage space consumed by snapshots on FlashArray is minimized by Redirect-On-Write (redirection of all overwrites to new blocks to avoid additional writes or data copy), compression and array-wide deduplication technologies, FlashArray may not be optimal for the long-term retention of its volume snapshots. Snap to NFS and CloudSnap have been introduced with Purity version 5.2 to provide the means of protecting FlashArray volumes and their snapshots by offloading them to another media. Snap to NFS and CloudSnap allow customers to replicate protection group snapshots from FlashArray to any NFS target or AWS S3 for long term retention. This offload snapshot feature requires no additional license to use.

Snap to NFS is also highly available. During the FlashArray controller upgrade procedure or in the event of the controller failure, the offload application which runs in Purity Run environment fails over to the surviving controller and resumes operation from the point when it was interrupted.


ACCESS APPLIANCE

The Veritas Access Appliance is a turn-key storage platform designed for cost optimization and high capacity, making it well suited for long-term retention of FlashArray volume snapshots. Key features of the Access Appliance include:

- Disk-based solution
- Integrated high availability – 2 nodes configured as active/active cluster, hot spares, and redundant power modules, RAID controllers, disk power paths, and disk SAS signals.
- Multi-protocol support – NFS, CIFS/SMB, FTP and S3.
- Scales up to 2.8 PB (2.5 PiB) of usable storage capacity
- Veritas AutoSupport to provide proactive monitoring and alerting 24x7 on the health of the appliance to reduce risk and quicker resolution.
- Symantec Data Center Security (SDCS) intrusion detection system. SDCS is a real-time monitoring and auditing software. It performs host intrusion detection, file integrity monitoring, configuration monitoring, user access tracking and monitoring, and produces logs and event reports. For more information on the Access Appliance intrusion detection system, refer to the [Access Appliance Initial Configuration and Administration Guide](#).
- Encryption with an external Key Management System (KMS) to create the keys for encryption.
- Replication of data on an Access Appliance to another offsite Access Appliance for disaster recovery purposes.

The Access Appliance model 3340 is comprised of two clustered nodes and one primary storage shelf and up to three additional expansion storage shelves. The appliance can scale up to 2,800 TB of usable space as can be seen in Figure 2. Refer to the [Access Appliance datasheet](#) for more detailed specifications.

Figure 2 - Access Appliance Rack Units

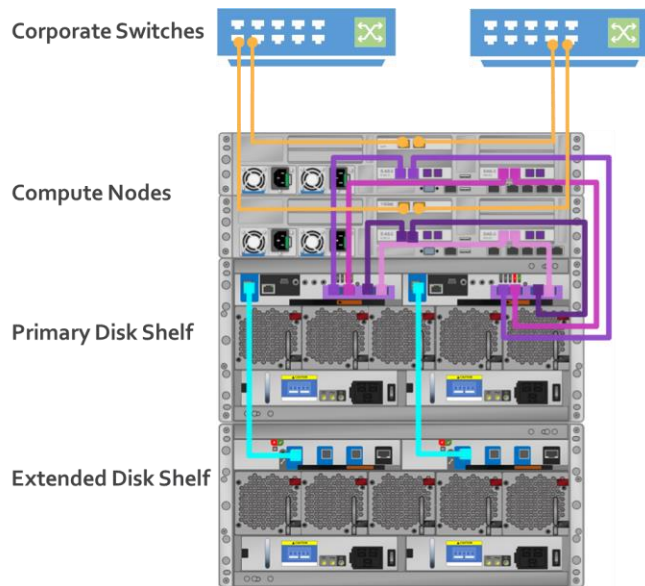


Note: TB - Capacity values are calculated using Base 10; TiB - Capacity values are calculated using base 2.

The two nodes are clustered in active/active configuration such that each node can handle I/O requests. Storage shelves are connected to each node and configured with dynamic multipathing, so I/O can be sent to either node for performance and

availability. There are four 10 GbE uplinks (2 per node) available for client connections. Figure 3 provides a view of these connections with 2 shelves.

Figure 3 - Access Appliance Connections Example



Access 3340 Appliance Connections




10GbE NIC – 4x 10GbE Small Form-factor Pluggable (SFP+) ports (2 per server node). 4 available for data transfer.



12Gb SAS-3 HBA – 16x SAS-3 ports (8 per server node, but only 4 per node are used) connect server nodes to primary storage shelf.

The redundant hardware RAID controller in the primary storage shelf configures and presents the shelves' physical disks into disk groups (volumes) protected by a RAID 6 storage layout. With a RAID 6 configuration, data with dual parity is striped across the configured volumes. There are 5 volumes per storage shelf with each volume containing 16 disks as pictured in Figure 4. Each data volume can remain operational despite two concurrent disk failures.

Figure 4 – Access Appliance Storage Shelves Disk Layout



Storage Shelf Drawers

Configuration

- 2 drawers per storage shelf
- 5 Data Volumes (16 drives per Volume configured as RAID 6).
- 41 drives (where 1 drive is hot spare) per drawer, 82 drives per shelf
- 10 TB, 7,200 rpm, SAS-3 drives
- SAS-3 host interface and internal connection

The nodes run RHEL 7.4 or later as the operating system platform and Access software version 7.3.2 or later. The Access Appliance supports multiple protocols, including NFS, CIFS/SMB, FTP, and S3. With FlashArray, the Access Appliance is seen as an NFS target.

When using Access as an NFS target, a "cluster file system" type is created where shares on that file system can be exported and then mounted on the FlashArray. The shares should be exported at the minimum with "rw" and sync option set on Access. Prior to configuring the share as a target in FlashArray, permissions of the FlashArray user which has a user id and group id of 1000 would need to be set appropriately for the share exported.

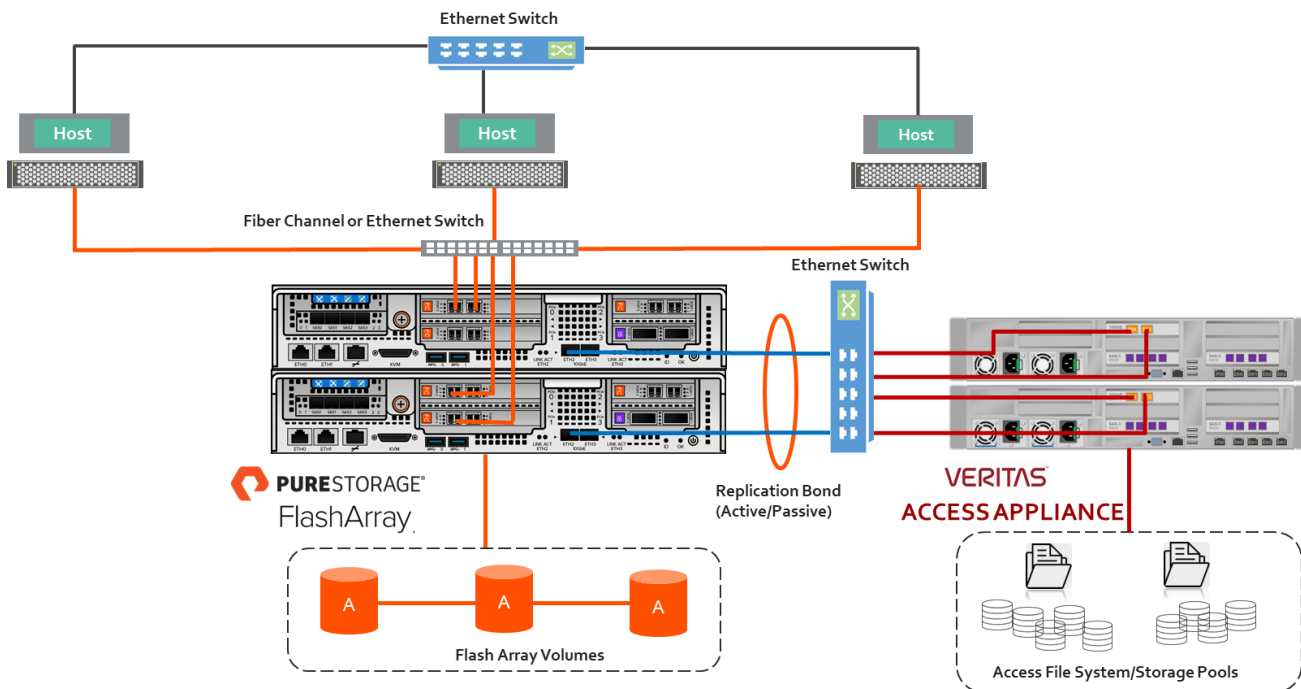
For management, the appliance can be managed by the command-line shell referred to as the CLISH and/or a web-based graphical user interface (GUI) where one can provision storage pools, create file systems and provision Access as an S3, NFS, and/or SMB targets.

NOTE: For an example of how to deploy and configure the Access Appliance as an NFS target with FlashArray, refer to the Appendix section of this whitepaper.

SOLUTION INTEGRATION

FlashArray connects to Access Appliance as depicted in Figure 5 and as previously mentioned sends snapshots to the Access Appliance via the NFS protocol. The hosts can connect to FlashArray using fiber channel (16 or 32 Gb), iSCSI or NVMe/RoCE (RDMA over converged Ethernet). For Snap to NFS and CloudSnap, the FlashArray replication ports and bond are utilized, which depending on the FlashArray model are 10 or 25 GbE in active/passive mode. In this mode, only one network interface is active, and the other interface will only become active when the first network interface becomes unavailable. FlashArray can communicate with any of the four 10 GbE network interfaces on the Access Appliance side (two 10 GbE per node). For more bandwidth, the Access Appliance network interface can be bonded if desired.

Figure 5 - Host, FlashArray and Access Appliance Solution Connection Example

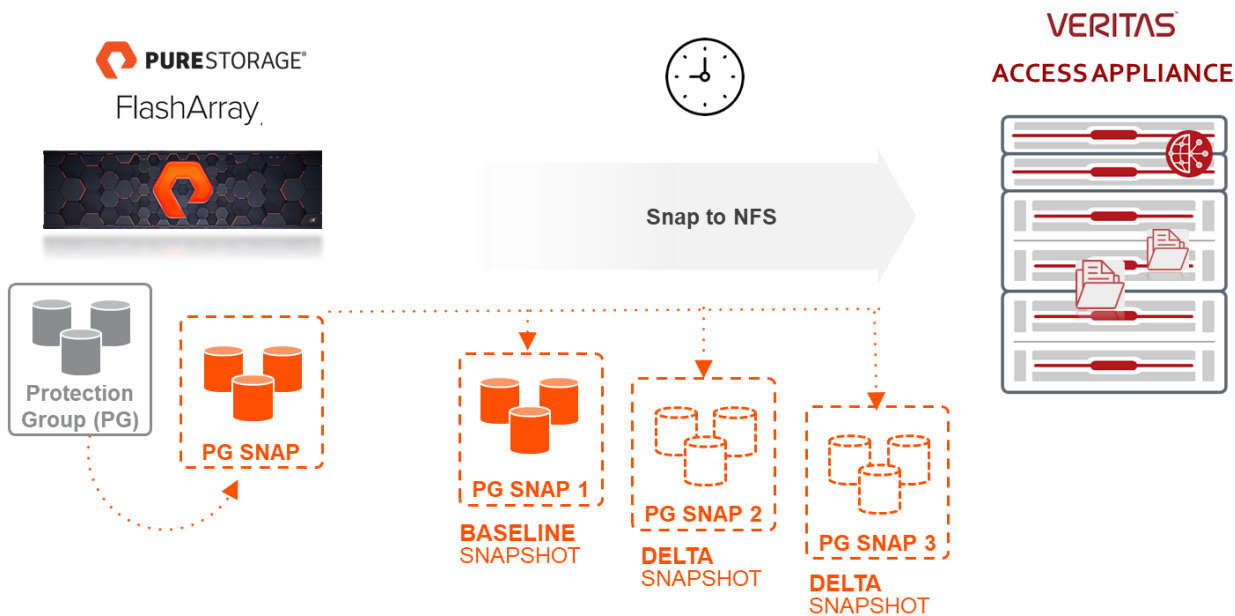


SNAP TO NFS

Currently, only one NFS offload target can be configured per FlashArray by specifying the NFS export share, the IP address or hostname of the NFS server and mount options. Once the connection is successfully established, the offload target can be defined within the protection group where volume snapshots can be offloaded or replicated. As previously stated, a protection group is where volume(s) can be specified to be protected using snapshot technology and then offloaded to a replication target. Point-in-time protection group snapshots of these volumes can be created manually or via a schedule on the FlashArray. These snapshots can then be sent or replicated to the Access Appliance either manually or by scheduling for long term retention.

As illustrated in Figure 6, there is a set of volumes defined within a protection group. Snapshots of the volumes belonging to the protection group are taken and stored locally on the FlashArray. The initial snapshot (baseline) sent to the offload target is the full, compressed volume data. For subsequent snapshots, only delta changes are sent to the offload target. This offloading technique results in lower network utilization, reduced space consumption on NFS target and less time to send the snapshot(s).

Figure 6 - Snap to NFS Example Flow



The offloaded FlashArray snapshots are fully portable, self-describing and allow for complete restoration of volume(s) to the source or another FlashArray. The data portions of the snapshots are sharded into files of maximum sizes of 70 MB and several meta-data files of less than 500 bytes. On the Access Appliance, the data and meta-data files of the snapshots are embedded within several directory structures as shown in Figure 7.

Figure 7 - Sample View of Files/Directories Representing the FlashArray Snapshot on Access

```

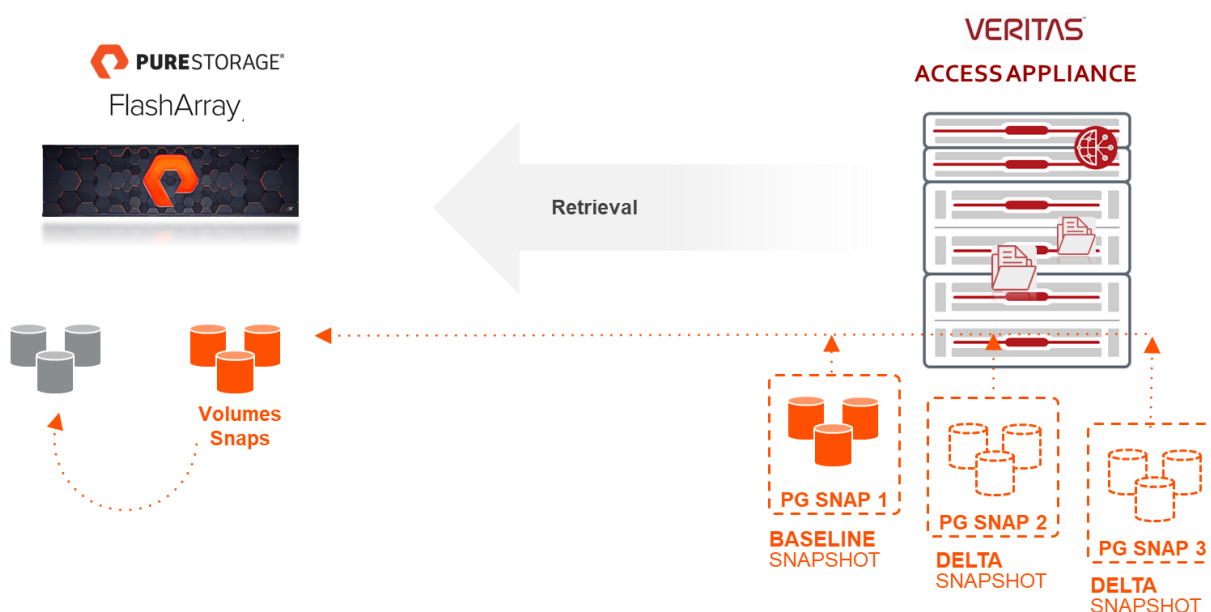
vaccess-02:/vx/va-nfs #
vaccess-02:/vx/va-nfs # ls -l
total 24
drwx----- 4 tomcat admin 96 May 7 17:10 10028266641579747074_14231117463407512331_125
drwx----- 7 tomcat admin 8192 May 7 17:09 11718279279661602163_9355406818422377885_125
-rw----- 1 tomcat admin 293 May 8 00:40 array_desc
drwxr-xr-x 2 root root 96 Apr 2 11:00 lost+found
-rw----- 1 tomcat admin 135 Apr 16 11:41 mount_desc
vaccess-02:/vx/va-nfs # ls -lh ./10028266641579747074_14231117463407512331_125/
0/
8/
vaccess-02:/vx/va-nfs # ls -lh ./10028266641579747074_14231117463407512331_125/
0/
8/
vaccess-02:/vx/va-nfs # ls 11718279279661602163_9355406818422377885_125
0 22 25 26 28 pgroup_desc
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_125/28/
data/
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/
11718279279661602163_9355406818422377885_8267/ 11718279279661602163_9355406818422377885_8267/ 11718279279661602163_9355406818422377885_8267/
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_8267/
11718279279661602163_9355406818422377885_8267/ 11718279279661602163_9355406818422377885_8267/ 11718279279661602163_9355406818422377885_8267/
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/
Display all 675 possibilities? (y or n)
10000/ 10026/ 10052/ 10133/ 10159/ 10261/ 10372/ 10398/ 10510/ 10536/ 10629/ 10655/ 10767/ 10793/ 10882/ 10908/ 11020/ 11046/ 11158/ 11269/ 11295/ 11404/ 11430/ 11520/ 11546/ 11572/
10001/ 10027/ 10053/ 10134/ 10160/ 10262/ 10373/ 10399/ 10511/ 10537/ 10630/ 10656/ 10768/ 10794/ 10883/ 10909/ 11021/ 11047/ 11159/ 11270/ 11296/ 11405/ 11431/ 11521/ 11547/ 11573/
10002/ 10028/ 10054/ 10135/ 10161/ 10263/ 10374/ 10400/ 10512/ 10538/ 10631/ 10657/ 10769/ 10795/ 10884/ 10910/ 11022/ 11048/ 11160/ 11271/ 11297/ 11406/ 11432/ 11522/ 11548/ 11574/
10003/ 10029/ 10055/ 10136/ 10162/ 10264/ 10375/ 10401/ 10513/ 10539/ 10632/ 10658/ 10770/ 10796/ 10885/ 10911/ 11023/ 11049/ 11161/ 11272/ 11298/ 11407/ 11433/ 11523/ 11549/ 11575/
10004/ 10030/ 10056/ 10137/ 10163/ 10265/ 10376/ 10402/ 10514/ 10540/ 10633/ 10659/ 10771/ 10797/ 10886/ 10912/ 11024/ 11136/ 11162/ 11273/ 11299/ 11408/ 11434/ 11524/ 11550/ 11576/
10005/ 10031/ 10112/ 10138/ 10240/ 10266/ 10377/ 10403/ 10515/ 10541/ 10634/ 10660/ 10772/ 10798/ 10887/ 10913/ 11025/ 11137/ 11163/ 11274/ 11300/ 11409/ 11435/ 11525/ 11551/ 11577/
10006/ 10032/ 10113/ 10139/ 10241/ 10267/ 10378/ 10404/ 10516/ 10542/ 10635/ 10661/ 10773/ 10799/ 10888/ 10914/ 11026/ 11138/ 11164/ 11275/ 11301/ 11410/ 11436/ 11526/ 11552/ 11578/
10007/ 10033/ 10114/ 10140/ 10242/ 10268/ 10379/ 10405/ 10517/ 10543/ 10636/ 10662/ 10774/ 10800/ 10889/ 10915/ 11027/ 11139/ 11165/ 11276/ 11302/ 11411/ 11437/ 11527/ 11553/ 11579/
10008/ 10034/ 10115/ 10141/ 10243/ 10269/ 10380/ 10406/ 10518/ 10544/ 10637/ 10663/ 10775/ 10801/ 10890/ 10916/ 11028/ 11140/ 11166/ 11277/ 11303/ 11412/ 11438/ 11528/ 11554/ 11580/
10009/ 10035/ 10116/ 10142/ 10244/ 10270/ 10381/ 10407/ 10519/ 10545/ 10638/ 10664/ 10776/ 10802/ 10891/ 10917/ 11029/ 11141/ 11167/ 11278/ 11304/ 11413/ 11439/ 11529/ 11555/ 11581/
10010/ 10036/ 10117/ 10143/ 10245/ 10271/ 10382/ 10408/ 10520/ 10546/ 10639/ 10665/ 10777/ 10803/ 10892/ 10918/ 11030/ 11142/ 11168/ 11279/ 11305/ 11414/ 11440/ 11530/ 11556/ 11582/
10011/ 10037/ 10118/ 10144/ 10246/ 10272/ 10383/ 10409/ 10521/ 10547/ 10640/ 10752/ 10778/ 10804/ 10893/ 10919/ 11031/ 11143/ 11169/ 11280/ 11306/ 11415/ 11441/ 11531/ 11557/ 11583/
10012/ 10038/ 10119/ 10145/ 10247/ 10273/ 10384/ 10496/ 10522/ 10548/ 10641/ 10753/ 10779/ 10805/ 10894/ 10920/ 11032/ 11144/ 11170/ 11281/ 11307/ 11416/ 11442/ 11532/ 11558/ 11584/
10013/ 10039/ 10120/ 10146/ 10248/ 10274/ 10385/ 10497/ 10523/ 10549/ 10642/ 10754/ 10780/ 10806/ 10895/ 10921/ 11033/ 11145/ 11171/ 11282/ 11308/ 11417/ 11443/ 11533/ 11559/ 11585/
10014/ 10040/ 10121/ 10147/ 10249/ 10275/ 10386/ 10498/ 10524/ 10550/ 10643/ 10755/ 10781/ 10807/ 10896/ 11008/ 11034/ 11146/ 11172/ 11283/ 11392/ 11418/ 11444/ 11534/ 11560/ 11586/
10015/ 10041/ 10122/ 10148/ 10250/ 10276/ 10387/ 10499/ 10525/ 10551/ 10644/ 10756/ 10782/ 10808/ 10897/ 11009/ 11035/ 11147/ 11173/ 11284/ 11393/ 11419/ 11445/ 11535/ 11561/ 11587/
10016/ 10042/ 10123/ 10149/ 10251/ 10277/ 10388/ 10500/ 10526/ 10552/ 10645/ 10757/ 10783/ 10809/ 10898/ 11010/ 11036/ 11148/ 11174/ 11285/ 11394/ 11420/ 11446/ 11536/ 11562/ 11588/
10017/ 10043/ 10124/ 10150/ 10252/ 10278/ 10389/ 10501/ 10527/ 10553/ 10646/ 10758/ 10784/ 10810/ 10899/ 11011/ 11037/ 11149/ 11175/ 11286/ 11395/ 11421/ 11447/ 11537/ 11563/ 11589/
10018/ 10044/ 10125/ 10151/ 10253/ 10279/ 10390/ 10502/ 10528/ 10554/ 10647/ 10759/ 10785/ 10811/ 10900/ 11012/ 11038/ 11150/ 11176/ 11287/ 11396/ 11422/ 11448/ 11538/ 11564/ 11590/
10019/ 10045/ 10126/ 10152/ 10254/ 10280/ 10391/ 10503/ 10529/ 10555/ 10648/ 10760/ 10786/ 10812/ 10901/ 11013/ 11039/ 11151/ 11177/ 11288/ 11397/ 11423/ 11449/ 11539/ 11565/ 11591/
10020/ 10046/ 10127/ 10153/ 10255/ 10281/ 10392/ 10504/ 10530/ 10556/ 10649/ 10761/ 10787/ 10813/ 10902/ 11014/ 11040/ 11152/ 11178/ 11289/ 11398/ 11424/ 11450/ 11540/ 11566/ 11592/
10021/ 10047/ 10128/ 10154/ 10256/ 10282/ 10393/ 10505/ 10531/ 10557/ 10650/ 10762/ 10788/ 10814/ 10903/ 11015/ 11041/ 11153/ 11179/ 11290/ 11399/ 11425/ 11451/ 11541/ 11567/ 11593/
10022/ 10048/ 10129/ 10155/ 10257/ 10283/ 10394/ 10506/ 10532/ 10558/ 10651/ 10763/ 10789/ 10815/ 10904/ 11016/ 11042/ 11154/ 11180/ 11291/ 11400/ 11426/ 11452/ 11542/ 11568/ 11594/
10023/ 10049/ 10130/ 10156/ 10258/ 10284/ 10395/ 10507/ 10533/ 10559/ 10652/ 10764/ 10790/ 10816/ 10905/ 11017/ 11043/ 11155/ 11181/ 11292/ 11401/ 11427/ 11453/ 11543/ 11569/ 11595/
10024/ 10050/ 10131/ 10157/ 10259/ 10285/ 10396/ 10508/ 10534/ 10560/ 10653/ 10765/ 10791/ 10800/ 10906/ 11018/ 11044/ 11156/ 11182/ 11293/ 11402/ 11428/ 11454/ 11544/ 11570/ id
10025/ 10051/ 10132/ 10158/ 10260/ 10286/ 10371/ 10397/ 10509/ 10535/ 10561/ 10654/ 10766/ 10792/ 10881/ 10907/ 11019/ 11045/ 11157/ 11183/ 11294/ 11403/ 11429/ 11455/ 11545/ 11571/
vaccess-02:/vx/va-nfs #
vaccess-02:/vx/va-nfs #
data_1_desc
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/10000/data_1
-rw----- 1 tomcat admin 60M May 7 23:53 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/10000/data_1
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/10001/data_1
-rw----- 1 tomcat admin 60M May 7 23:54 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/10001/data_1
vaccess-02:/vx/va-nfs # ls -lh 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/10002/data_1
-rw----- 1 tomcat admin 60M May 7 23:54 11718279279661602163_9355406818422377885_125/28/data/11718279279661602163_9355406818422377885_8267/10002/data_1
vaccess-02:/vx/va-nfs #

```

SNAPSHOT RETRIEVAL

When restoring, FlashArray can retrieve one or several volumes in the protection group from the Access Appliance. In Figure 8, the compressed data that make up the requested snapshot is requested and then re-assembled onto the FlashArray as a volume snapshot. To reduce network bandwidth, only missing data blocks not already present on the FlashArray are transferred to rebuild the entire snapshot. The volume snapshot can further be restored to a volume with an option to overwrite or make a copy if the volume exists.

Figure 8 - Snapshot Retrieval Flow



SOLUTION SECURITY

At a transport level, FlashArray sends data over dedicated NFS network ports (2049, 111, 4001, and 4045) to Access. The ownership of the NFS exported share belongs to a FlashArray user with a user id and group id of 1000, hence, limiting the access of the contents to this specific user. For enhanced security, the Access Appliance also has encryption capabilities in conjunction with an external Key Management System (KMS). The appliance encrypts the volume that the “file system” resides on. An external KMS such as IBM KMS is required to create the keys for the encryption.

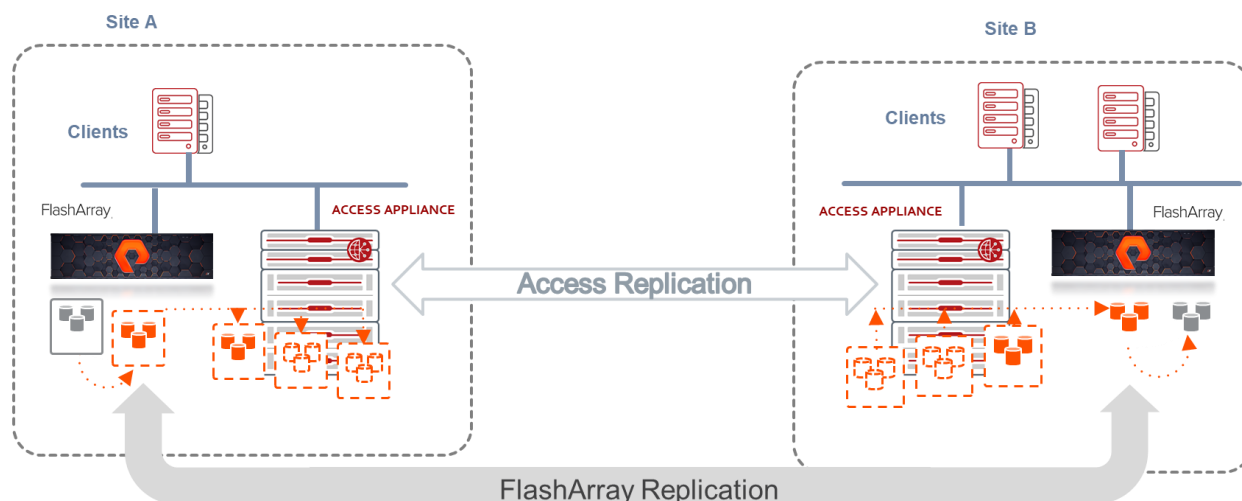
DISASTER RECOVERY

Having a disaster protection plan is imperative for business continuity. FlashArray has a capability to replicate its snapshots to another FlashArray at a remote site. For details on FlashArray replication, refer to [FlashRecover Replication Configuration and Best Practices Guide](#). The Access Appliance also has a feature to provide replication of FlashArray snapshots to another Access Appliance at a remote site for disaster recovery as pictured in Figure 9. There are two types of replication offered by the Access Appliance:

- **Episodic Replication** – file system replication asynchronously.
- **Continuous Replication** – block level replication of volumes in synchronous or asynchronous modes.

For more details on the Access Appliance replication, please refer to the [Access Appliance Administrators Guide](#).

Figure 9 - Access Appliance Replication of FlashArray Snapshots



BEST PRACTICES AND RECOMMENDATIONS

Following best practices is important in creating an optimally deployed solution. This section covers some best practices relating to the Access Appliance as a long-term retention storage solution for FlashArray volume snapshots.

DATA LAYOUT ON ACCESS APPLIANCE

The appliance contains hardware RAID 6 controllers and inherently does striping with dual parity across disks on the storage shelves for high performance and data durability. Selecting other layouts such as mirrored or erasure coding layout for data protection is not necessary. As a best practice, for backup, archival and long-term retention use cases, it is recommended to configure the Access Appliance using defaults such as clustered file system, simple layout, and block size of 8 KB. For additional performance, the layout can be configured to be stripe. **NOTE:** To maintain the stripe performance, when growing the storage pool, the volumes must be added in multiples of the stripe columns, and thus, it is advisable to plan or size the system appropriately.

NFS PROTOCOL

When using Access Appliance as secondary storage for FlashArray volume snapshots, it is recommended to mount the NFS share with the following options for better performance and minimize data loss.

- **sync** – data is flushed to disks on server instead of keeping data in the client or server caches. This option should be specified when exporting on Access to minimize data loss in case of failure.
- **nfsvers=3** – it was observed that the performance when utilizing NFS version 3 is better when compared to the NFS version 4 protocol.
- **rsize=1048576** and **wsize=1048576** – the best performance was observed when setting the read and write size values to 1 MB. This value specifies the number of bytes that NFS utilizes to read and write files on the NFS server.

Since multiple FlashArray can store snapshots on a single Access Appliance, it is recommended as a best practice to use a separate NFS share as the offload target for each FlashArray for segregation and ease of management. Snapshots from different FlashArray

can be seen in each FlashArray utilizing the same share. For instance, snapshots taken from two FlashArray can be seen in both FlashArray. Additionally, those snapshots can only be removed by the FlashArray where the protection group resides or where the snapshot was created. Furthermore, it is recommended to use the NFS share solely for FlashArray snapshots and not for other file sharing purposes. Better performance was also observed when utilizing different file systems for each share.

The optimum frequency to conduct "Snap to NFS" is generally dependent on size of the data, the network bandwidth, and system resources on the FlashArray. However, per [Snap to NFS Overview and Administration](#) documentation, it is recommended to offload data not more than once or twice a day.

COMPRESSION

Compression is a good feature for better storage utilization. The Access Appliance has the capability to do compression, however, it is not recommended for use with FlashArray snapshots. The snapshots received from FlashArray are already compressed by default.

NETWORK CONNECTIVITY

The Access Appliance has two 10 GbE uplinks per node. Each physical port maps to a virtual IP. Thus, there are four virtual IP addresses. Always present the virtual IP to clients or in this scenario FlashArray so it will automatically transition to the other node if one node fails, the physical links on one node fails, or the node becomes unreachable.

Bonding is an option on Access Appliance. Joining or bonding multiple network interfaces on the Access appliances into a single interface improves the bandwidth and network throughput through the combined interface. Bonding is only configurable via the Access command-line interface. As a best practice, the switch that the uplinks of the Access Appliance are connected to should be configured appropriately for the link aggregation.

For load balancing, virtual IP addresses of the nodes can be manually assigned to applications in a distributed manner. Balancing the load across the nodes and network interfaces improves overall performance especially when directing multiple FlashArray to a single Access Appliance.

MONITORING

It is important to monitor or be aware of the alerts, especially storage utilization warnings and hardware critical alerts. The AutoSupport features assists in this manner, but as a best practice, it is advisable to be pro-active instead of re-active. For instance, once the capacity reaches 70%, it might be a good time to revisit the storage utilization or plan for growth.

SIZING GUIDANCE

When planning or sizing the Access Appliance as a long-term retention solution for FlashArray snapshots there are two factors:

- Capacity - how many snapshots can be stored
- Performance - how much workload can the storage platform handle.

The Veritas and Pure Storage account team will assist in the sizing of the appliance based on your requirements using these factors. Some parameters that might enter in the equation when estimating long-term storage requirements include:

1. Volume of source data.
2. Daily data change ratio
3. Annual storage growth

4. Data retention.
5. Compression Ratio (how well your data compresses)
6. Performance and/or service level requirements.

Some considerations when deploying and implementing FlashArray “Snap to NFS” include:

- Maximum number of offload targets per FlashArray: 1
- Maximum number of volume snapshots offloaded to NFS per FlashArray: 100,000
- Maximum number of volume snapshots on NFS per FlashArray: 200,000
- Maximum number of protection groups per FlashArray: 250
- Maximum number of FlashArray per NFS target (not NFS server): 4

EXAMPLE CAPACITY SIZING OF ACCESS APPLIANCE FOR FLASHARRAY SNAPSHOTS

The Access 3340 Appliance can scale up to 2.8 PB of usable capacity. Although the FlashArray volume snapshots are compressed, to create portable and efficient snapshot, FlashArray has about a 5% overhead. For this simplistic sizing exercise, capacity is evaluated with an initial source data of 100 TB in size, a known data compression ratio of 2:1, retention of 30 days, and daily change rate of 5%. The sample capacity required for 30 days would be as follows:

$$\begin{aligned}\text{Capacity Needed} &= ((\text{Data Size} + (\text{Retention} * \text{Change Rate} * \text{Data Size}) * \text{Compression Rate}) + (\text{Data Size} * 5\%)) \\ &= ((100 \text{ TB} + (30 \text{ days} * 0.05 \text{ per day} * 100 \text{ TB})) * 0.5 \text{ compression}) + (100 \text{ TB} * 0.05) \\ &= 125 \text{ TB} + 5 \text{ TB} \\ &= 130 \text{ TB}\end{aligned}$$

In addition, when old snapshots expire, the data blocks from these snapshots are not converted into free space immediately, so it conceivable that at some point in time the space consumed by the above snapshots will approach the following:

$$\begin{aligned}\text{Capacity Needed} &= 130 \text{ TB} / 70\% \\ &= 186 \text{ TB}\end{aligned}$$

Therefore, capacity on Access Appliance should be calculated to include room for the above overhead.

PERFORMANCE OF ACCESS APPLIANCE FOR FLASHARRAY

Pure and Veritas conducted a joint performance study to understand the performance of sending and retrieving snapshots from FlashArray to and from the Access Appliance via the NFS protocol. Table 1 exhibits the maximum throughput results when sending, retrieving and 50%/50% send and retrieve of the volume snapshots between one FlashArray and an Access 3340 Appliance with one shelf. If using a higher model of FlashArray, the throughput values can be higher. Based on the “sar” and “vxstat” outputs gathered during the testing, the Access Appliance system (CPU, memory, etc.), network and disk utilization were minimal. When another FlashArray was added, performance doubled as expected since the Access Appliance system resources was not overloaded. The overall system utilization on the Access Appliance did increase when two FlashArray were used, however, it did not fully saturate them. Thus, a single Access Appliance is capable of handling multiple FlashArray. Refer to the Appendix section for more details on this performance study.

Table 1 - Solution Performance

| Workload | Maximum Observed Throughput |
|-----------------------|---|
| Snap to NFS | 232 MB/s with one FlashArray - Access Appliance (1 node, 1 network interface, 1 file system) 457 MB/s with two FlashArray - Access Appliance (1 node, 1 network interface, 1 file system) 640 MB/s with two FlashArray – Access Appliance (1 node, 2 network interfaces, 2 file system) |
| Snapshot Retrieval | 180 MB/s - Access Appliance (1 node, 1 network interface, 1 file system) |
| 50% Send/50% Retrieve | 305 MB/s - Access Appliance (1 node, 1 network interface, 1 file system) |

Additional testing was conducted to determine the effects when modifying different parameters such as tunables, utilizing multiple FlashArray, multiple network interfaces, etc. A summary of observations includes:

- Best performance was observed when using a rsize=1048576 and wsize=1048576
- Best performance was observed when using NFS version 3 when compared to NFS version 4 protocol.
- When adding another FlashArray to send data to the Access Appliance, the throughput performance doubled. For instance, with 2 FlashArray the combined throughput observed was 457 MB/s when load was sent to 1 node and 1 network interface of the Access Appliance.
- Utilizing a higher model of FlashArray also produced higher results. For instance, utilizing a //FA10-R2, the maximum throughput observed for “Snap to NFS” was 300 MB/s.
- Separating the volumes in different protection groups did not affect the results.
- Performance is improved when distributing the workload across the two nodes, multiple network interfaces and different file systems of the Access Appliance. For example, with 2 FlashArray and when using 1 node, 1 network, and 1 file system on the Access Appliance the throughput was 457 MB/s whereas when using 1 node, 2 filesystems, and 2 network interfaces the combined throughput was 640 MB/s. This is an improvement of 40%.

CONCLUSION

The Access Appliance provides a competitive disk-based solution for long-term retention of FlashArray volume snapshots. It is also a complementary disk-based option with Pure Storage all-flash array products and offers a compelling solution in data protection, disaster planning and recovery as well as migration and test and development operations use cases. Implementing the Access Appliance as a long-term retention target for FlashArray volume snapshots provides security, minimizes costs, and improves control and visibility.

REFERENCES

- Access 3340 Appliance
 - Product Documentation
 - 7.3.2 - https://sort.veritas.com/documents/doc_details/AAPP/7.3.2/Appliance%203340/ProductGuides/
 - 7.4.2 - https://sort.veritas.com/documents/doc_details/AAPP/7.4.2/Appliance%203340/ProductGuides/
- FlashArray
 - Snap to NFS Overview and Administration
 - https://support.purestorage.com/FlashArray/PurityFA/Purity_RUN/Snap_to_NFS/Snap_to_NFS_Overview_and_Administration
 - FlashArray Concepts and Features
 - https://support.purestorage.com/FlashArray/PurityFA/FlashArray_User_Guide/Purity%2F%2F%2F%2FFA_Version_5.2.3/FlashArray_Overview/FlashArray_Concepts_and_Features

APPENDIX

This section provides an example deployment of the Access Appliance as an NFS target for FlashArray volume snapshots. It also provides the details on the performance study that was jointly conducted at Pure Storage Labs relating to this solution. It describes the testing strategy, environment, results and analysis associated with this study.

SOLUTION DEPLOYMENT

This section describes an example of the configuration and readers are expected to refer to the Veritas Access Appliance and Pure Storage FlashArray product documentation for definitive and specific installation, administration, and configuration details.

The deployment example walks through the configuration of the Access Appliance as an NFS offload target for FlashArray as shown in Figure 10.

Figure 10 - Configuration Used in this Example Deployments



Outline of the steps involved in this example include:

- Configuration and storage provisioning of the Access Appliances as an NFS target.
- Configuration of FlashArray to utilize the Access Appliance as an NFS target.
- Validation of configuration by conducting a send to NFS and retrieval of snapshot from the Access Appliance.

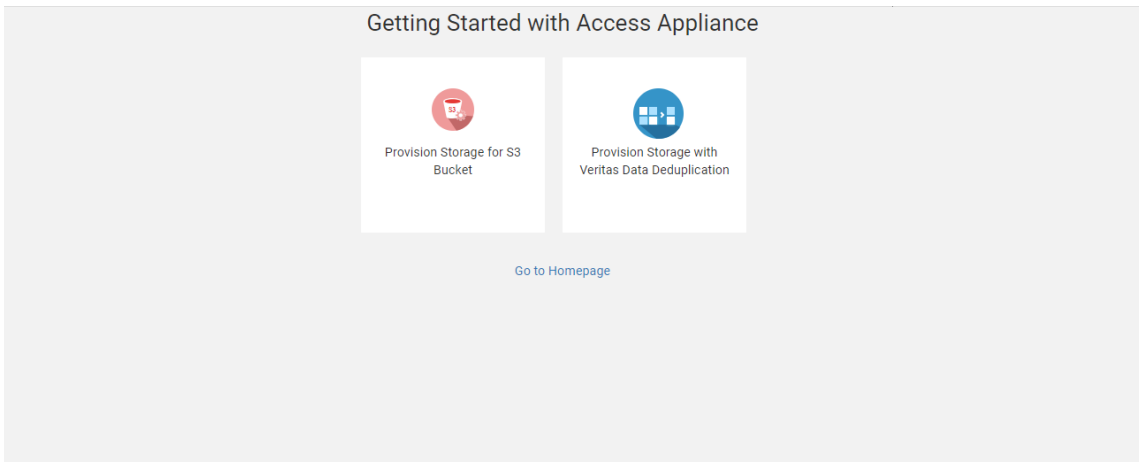
ACCESS APPLIANCE STORAGE CONFIGURATION AND PROVISIONING

In this example, the Access Appliance graphical user interface (GUI) and command-line (CLI) is utilized to provision the storage. It is assumed that the Access Appliance has already been installed and connected to the same network as the FlashArray. Summary of steps to configure and provision Access Appliance as an NFS target involves the following:

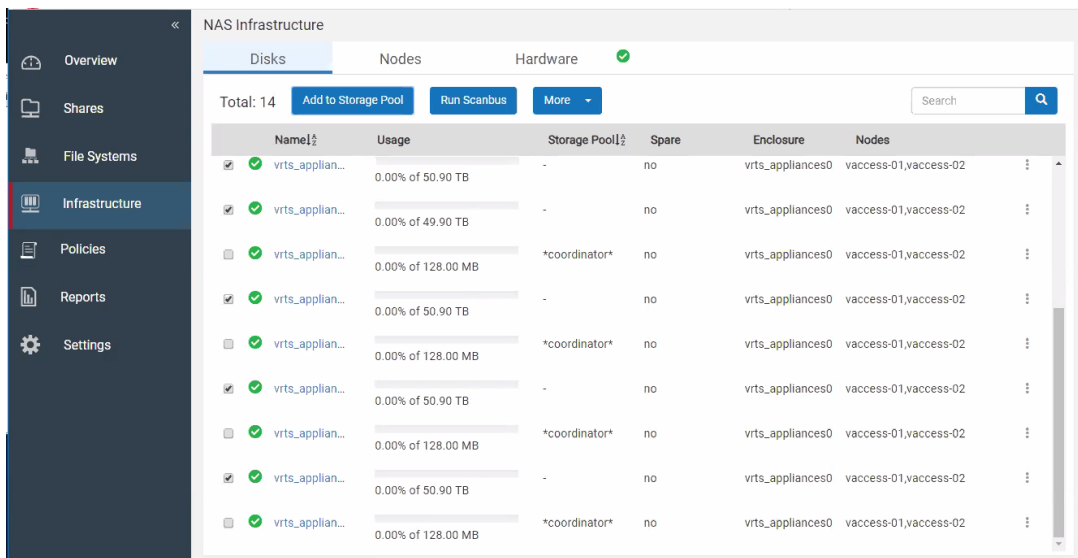
1. Configure the storage disk pools
2. Enable NFS services.
3. Create a file system and export the file system as an NFS share.
4. Modify the permissions of share to have a user id and group id of 1000.

CONFIGURE STORAGE POOLS

Step 1) Using a web browser, login to the Access Appliance graphical user interface: <https://<consoleIP>:14161>. Click on “Go to Homepage”.



Step 2) Click on **Infrastructure** on left pane. Select the number of disks or volumes to be added to storage pool and click “Add to Storage Pool”.



Access Appliance for FlashArray Snapshots

Step 3) Select **"Add to new storage pool"** and enter a **Storage Pool Name**, i.e. ppool and click **Next**. On next screen, click **Finish**.

Add to Storage Pool

Select Disks to Add into Existing or New Storage Pool

Selected: 5 Storage Pool Capacity: 0 byte

| Name | Usage | Storage Pool | Enclosure | Nodes |
|---------------------|-------------------|--------------|------------------|-----------------------|
| vrts_appliances0_10 | 0.00% of 50.90 TB | - | vrts_appliances0 | vaccess-01,vaccess-02 |
| vrts_appliances0_2 | 0.00% of 49.90 TB | - | vrts_appliances0 | vaccess-01,vaccess-02 |
| vrts_appliances0_4 | 0.00% of 50.90 TB | - | vrts_appliances0 | vaccess-01,vaccess-02 |
| vrts_appliances0_6 | 0.00% of 50.90 TB | - | vrts_appliances0 | vaccess-01,vaccess-02 |

Select Storage Pool : Add to new storage Pool

Storage Pool Name : ppool

Next Cancel

Step 4) Check the **Activity** icon, the **clock** on top and click on **"Show All Recent Activities"**. Wait until pool creation succeeds.

All Recent Activities

Click any task to view the details.

Total: 1 of 1

| Name | Status | Start Time | End Time |
|---------------------------|---------|---------------------|---------------------|
| Create Storage Pool ppool | Success | 2019-04-02 10:50:34 | 2019-04-02 10:51:07 |
| vxdisk_init | Success | 2019-04-02 10:50:40 | 2019-04-02 10:50:44 |
| vxdisk_init | Success | 2019-04-02 10:50:44 | 2019-04-02 10:50:49 |
| vxdisk_init | Success | 2019-04-02 10:50:49 | 2019-04-02 10:50:53 |
| vxdisk_init | Success | 2019-04-02 10:50:53 | 2019-04-02 10:50:57 |
| vxdisk_init | Success | 2019-04-02 10:50:57 | 2019-04-02 10:51:02 |
| Setting up DG | Success | 2019-04-02 10:51:02 | 2019-04-02 10:51:05 |

Output: ACCESS Pool SUCCESS V-493-10-2911 Created pool ppool successfully

Command executed: NAS_OUTPUT=json /opt/VRTSnas/clish/bin/clish -u master -c "storage pool create ppool vrts_appliances0_10,vrts_appliances0_2,vrts_appliances0_4,vrts_appliances0_6,vrts_appliances0_8"

Close

ENABLING NFS

Step 5) Click on **Settings** and in **Share Services Management** panes, click **slider** to right to enable NFS.

The screenshot shows the 'Settings - Services Management' interface. On the left is a navigation menu with options: Overview, Shares, File Systems, Infrastructure, Policies, Reports, and Settings (highlighted). The main content area is divided into several panels:

- Storage Services Management**: Displays discovery information for 'tmeaccess2-01', including last discovery status (Partial), start time (Tue, 11 Dec 2018 14:18:06), and end time (Tue, 11 Dec 2018 14:18:17). It includes a 'Run Full Discovery' button and a link to the REST API.
- Cluster Time Management**: Shows the current time (2018-12-11T14:20:47-08:00) and current time zone. It has 'Set Time' and 'Set Time Zone' buttons.
- NTP Server Management**: Shows the NTP service status as 'Off' and the NTP server as '127.127.1.0'. It includes a 'Sync Now' button.
- Share Services Management**: Shows the status of NFS and CIFS services, both currently 'Offline'. Each has a toggle switch to enable it.
- Key Management Service (KMS)**: Indicates that the KMS server is not configured and provides instructions for configuration, including a 'Provide Key & Certificates' button.
- CIFS Service Management**: Displays various CIFS settings such as NetBIOS Name (tmeaccess2), NTLM Auth (yes), Trust Domains (no), Home Dir File System, AIO Size (0), ID Map Backend (rid:10000-1000000), Work Group (WORKGROUP), Security (user), Domain, Domain User, and Domain Controller. It includes a 'Set' button.

Step 6) Select **KNFS** to use Kernel NFS and click **Start**.

The screenshot shows a 'Start NFS Server' dialog box. It contains a section titled 'Select the NFS server type to start.' with two radio button options: 'KNFS' (which is selected) and 'GNFS'. At the bottom of the dialog, there are 'Start' and 'Cancel' buttons.

Step 7) Once complete, you will see that NFS is set to **ONLINE**.

Settings > Services Management

Storage Services Management
Discovery information
Management Console: tmeaccess2-01
Last discovery: Partial
Tue, 11 Dec 2018 14:18:06
Last discovery start time: Tue, 11 Dec 2018 14:18:06
Last discovery end time: Tue, 11 Dec 2018 14:18:17
[Run Full Discovery](#)
[Click here for REST API](#)

Cluster Time Management
Current Time: 2018-12-11T14:22:29-08:00
Current Time Zone: [Set Time](#) [Set Time Zone](#)

NTP Server Management
NTP Service Status: ☒
NTP Server: 127.127.1.0
[Sync Now](#)

Share Services Management
NFS: Online ☒ KNFS
CIFS: Offline ☐

Key Management Service (KMS)
KMS server is not configured
To configure KMS server:
STEP 1: Provide the client SSL key and certificate and KMS server SSL certificate.
[Provide Key & Certificates](#)
STEP 2: Configure KMS server.
[Configure KMS Server](#)

CIFS Service Management
NetBIOS Name: tmeaccess2
NTLM Auth: yes
Trust Domains: no
Home Dir File System: [Set](#)
AIO Size: 0
ID Map Backend: rid:10000-1000000
Work Group: WORKGROUP
Security: user
Domain:
Domain User:
Domain Controller:
Clustering Mode: normal

CREATION OF FILE SYSTEM TO EXPORT AS NFS SHARE

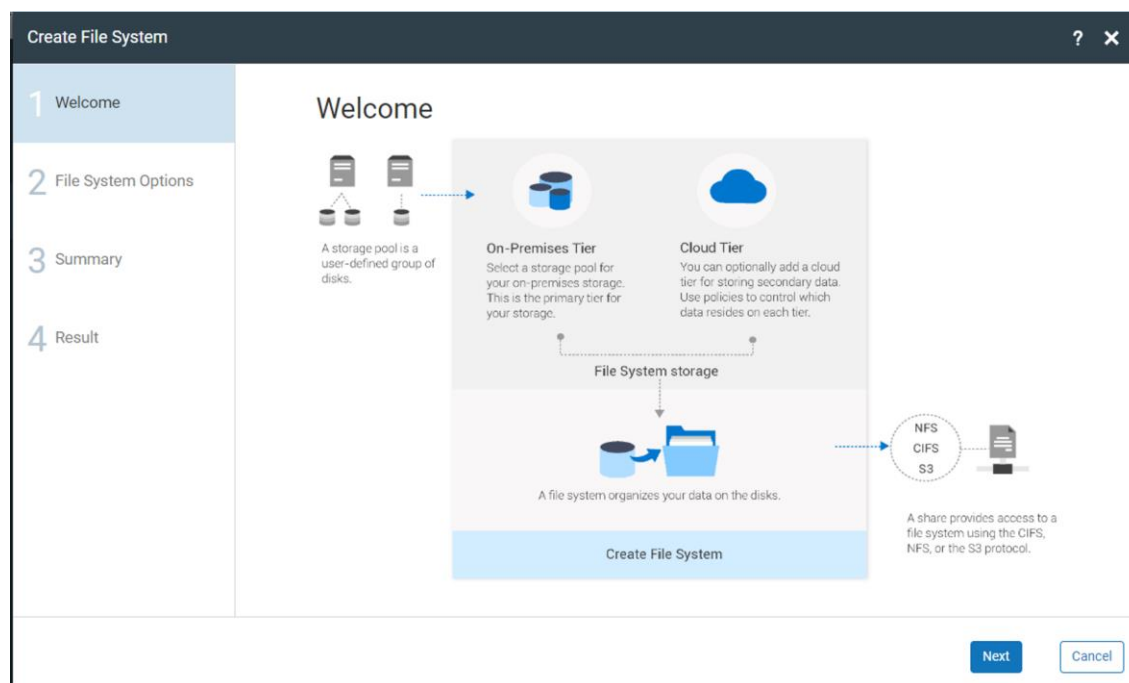
Step 8) On left pane, click on **File Systems** and select **Create File system** in the page view.

File Systems

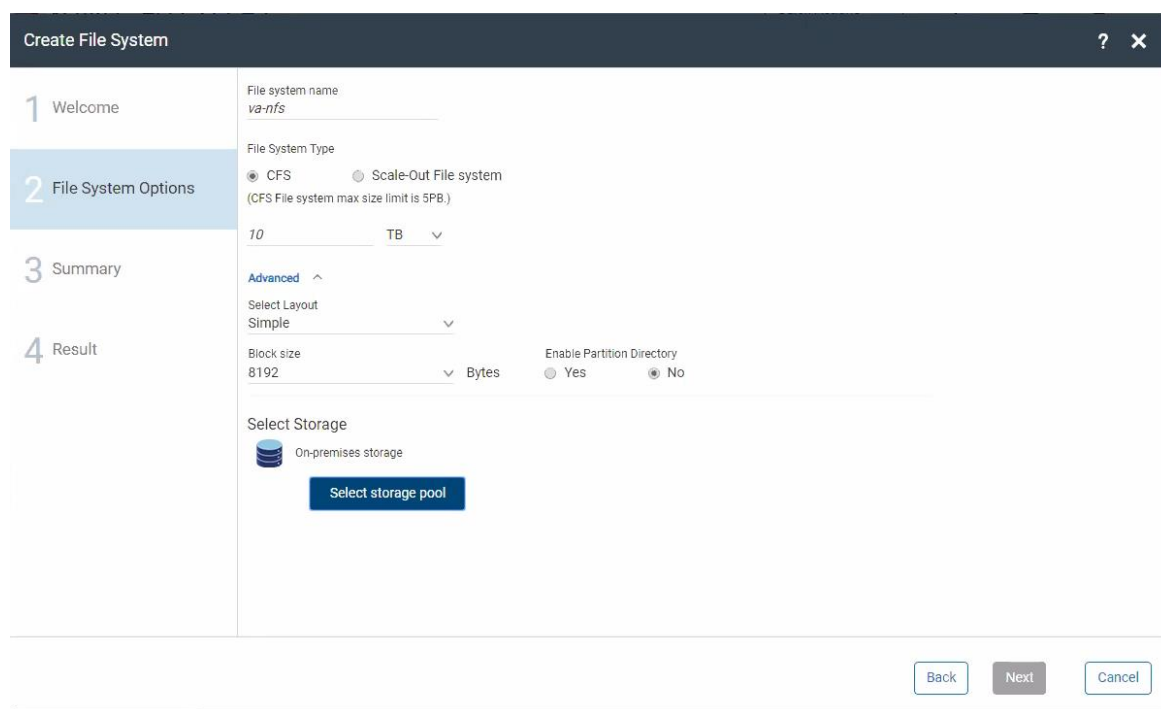
Total: 20 [Create File System](#)

| Name | Status | Usage | Storage Pool | Layout | S3 Bucket | Encryption |
|-------|--------|------------------|--------------|--------|-----------|------------|
| afs1 | Online | 0.02% of 5.00 TB | tsmpool | simple | No | No |
| afs10 | Online | 0.11% of 5.00 TB | tsmpool | simple | No | No |
| afs11 | Online | 0.10% of 5.00 TB | tsmpool | simple | No | No |
| afs12 | Online | 0.10% of 5.00 TB | tsmpool | simple | No | No |
| afs13 | Online | 0.11% of 5.00 TB | tsmpool | simple | No | No |
| afs14 | Online | 0.02% of 5.00 TB | tsmpool | simple | No | No |
| afs15 | Online | 0.02% of 5.00 TB | tsmpool | simple | No | No |
| afs2 | Online | 0.02% of 5.00 TB | tsmpool | simple | No | No |
| afs3 | Online | 0.02% of 5.00 TB | tsmpool | simple | No | No |
| afs4 | Online | 0.02% of 5.00 TB | tsmpool | simple | No | No |
| afs5 | Online | | tsmpool | simple | No | No |

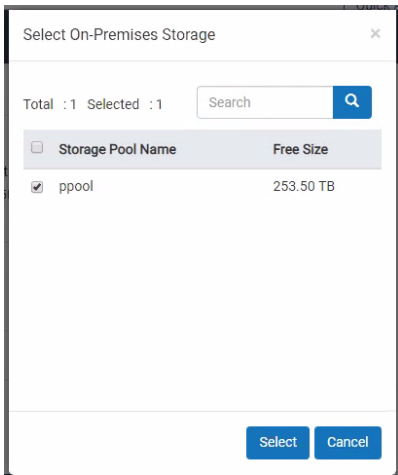
Step 9) A wizard pops up and then follow wizard. Click **Next**.



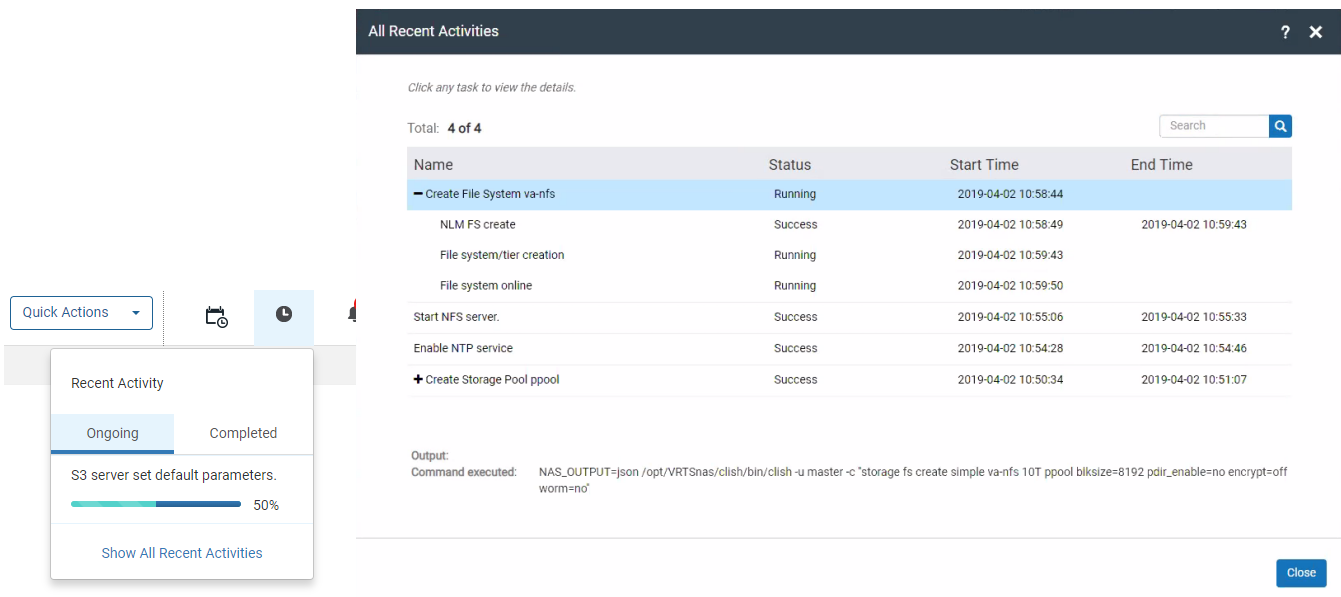
Step 10) Enter **name**, **file system type of CFS**, **size of the file system**, then just use the defaults of simple layout and block size of 8192. Click on **Select Storage Pool**.



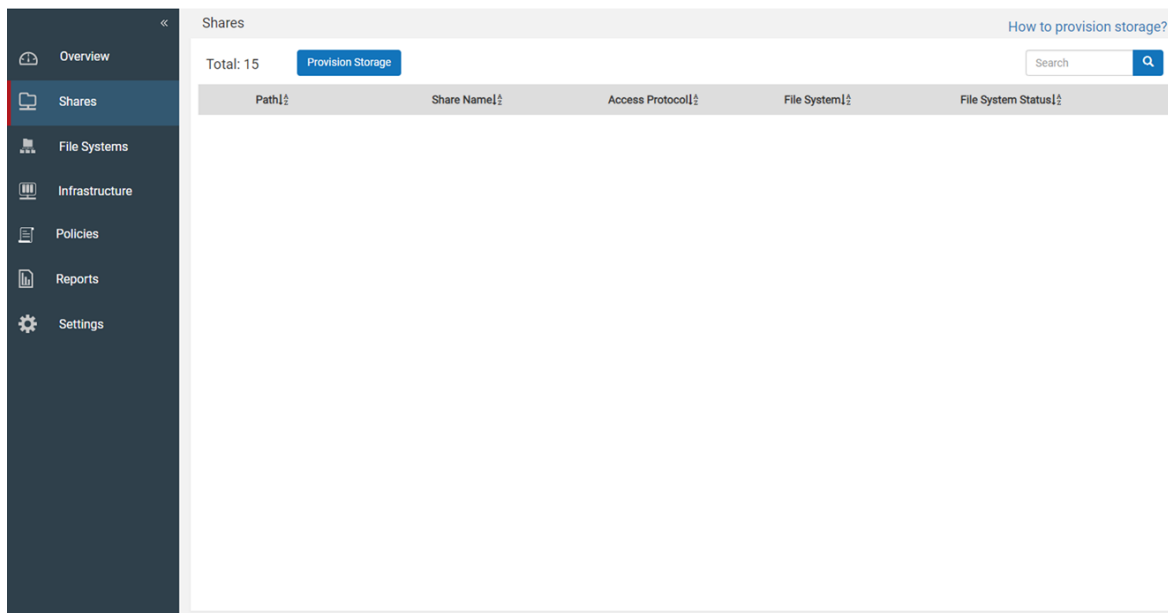
Step 11) Select the **storage pool** created in previous steps (i.e. ppool).



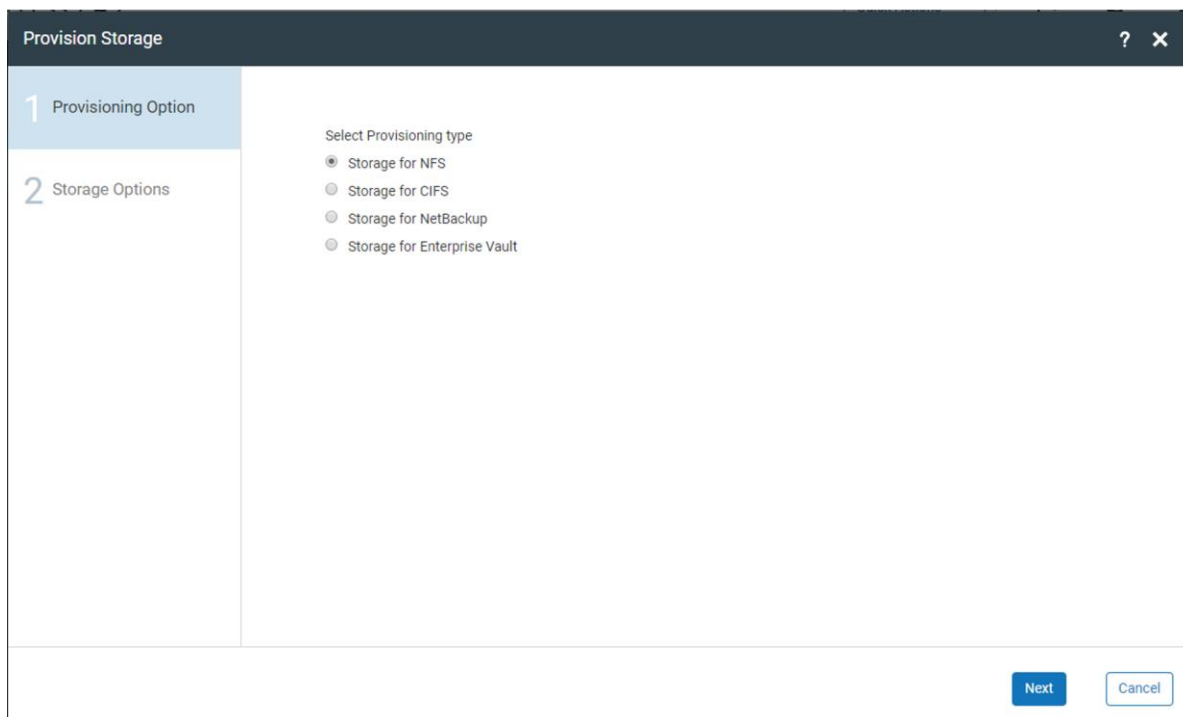
Step 12) Click **Next** and then click **Finish**. Click on **clock icon** on top of page and wait for the creation of file system to succeed.



Step 13) Click **Shares** on left hand side and then click **Provision Storage** as an NFS Share.



Step 14) Select "Storage for NFS".



Step 15) Select the **file system** created in previous steps, for instance, "va-nfs".

Provision Storage

1 Provisioning Option
2 **Storage Options**
3 Share Options
4 Summary
5 Result

Select Storage type

☐ Policy
☒ File System

Select a File System:

Total 1

| Name | Status | S3 Bucket | Layout | Usage |
|---|--------|-----------|--------|-------------------|
| <input checked="" type="radio"/> va-nfs | online | No | simple | 0.01% of 10.00 TB |

Back Next Cancel

Step 16) Set the pathname of the form "/vx/<file system name>". Expand Options and **select Read Write and Synchronous** and click on "**Set**". Click on **Next**.

Provision Storage

1 Provisioning Option
2 Storage Options
3 **Share Options**
4 Summary
5 Result

Add New Share

NFS Options Actions

Specify Client to Export
Any Client

Access Type
☐ Read Only ☒ Read Write

Advanced ^

NFS Export Options
☒ Synchronous ☐ Secure ☐ Secure Locks ☒ Root Squash ☐ Write Delay ☐ SubTree Check

Set

Back Next Cancel

Step 17) Review the summary details and click **Next** and **Finish**. Then, **click on clock icon** and wait for export of share to finish.

Provision Storage

1 Provisioning Option

2 Storage Options

3 Share Options

4 Summary

5 Result

Summary

Share Details

Selected File System: va-nfs

Directory Path: /vx/va-nfs

Protocol Option: NFS

Export Options: rw,sync,insecure,insecure_locks,no_root_squash,no_wdelay,no_subtree_check

Client: *

Back Next Cancel

MODIFICATION OF NFS SHARE PERMISSIONS

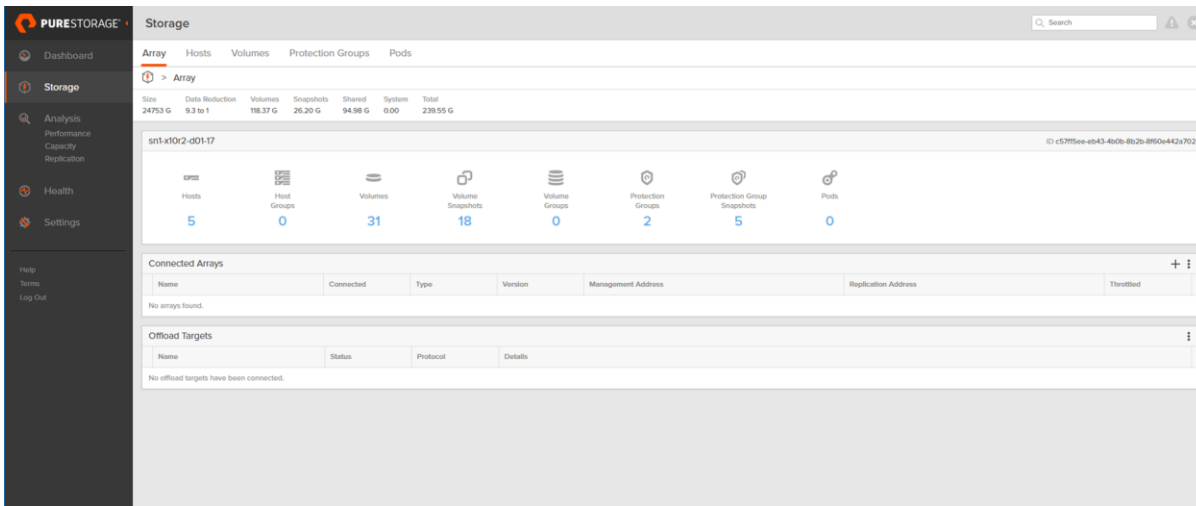
Step 18) FlashArray requires a specific user and group (1000, 1000) to be able to write into the share. Thus, the ownership would need to be modified on the file system that will be exported as an NFS share. To modify the ownership, ssh onto one of the nodes to enter to the Access Appliance command-line interface (CLI). (NOTE: This is different from the Access cluster CLI). Once you enter the CLI, get into the Support and then the Maintenance views. Enter password as instructed and elevate to get into the bash shell of the Access Appliance. From here do a **chown** of file system as shown below. In this example, /vx/va-nfs was modified with user id and group id of 1000 which maps to user tomcat and admin in the Access Appliance. The "tomcat" user is not used within the Access Appliance.

```
vaccess-02.Main_Menu> Support
Entering Appliance support view...
vaccess-02.Support> Maintenance
<!-- Maintenance Mode -->
maintenance's password:
maintenance-!> elevate
vaccess-02:/home/maintenance # chown 1000:1000 /vx/va-nfs
vaccess-02:/home/maintenance # ls -ld /vx/va-nfs
drwxr-xr-x 5 tomcat admin 8192 May 8 00:40 /vx/va-nfs
```

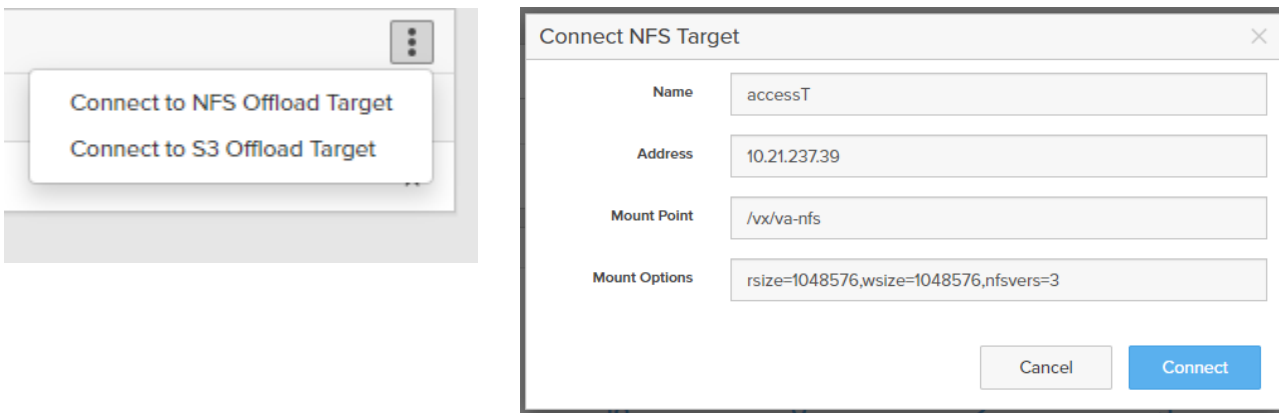

CONFIGURATION OF FLASHARRAY TO UTILIZE ACCESS AS AN NFS TARGET

Offload application needs to be installed to enable Snap to NFS and CloudSnap features. Thus, for offload application (Snap to NFS) installation contact Pure Storage support at support@purestorage.com


Step 1) Connect to the FlashArray GUI using a web browser, click on **Storage** on the left pane.



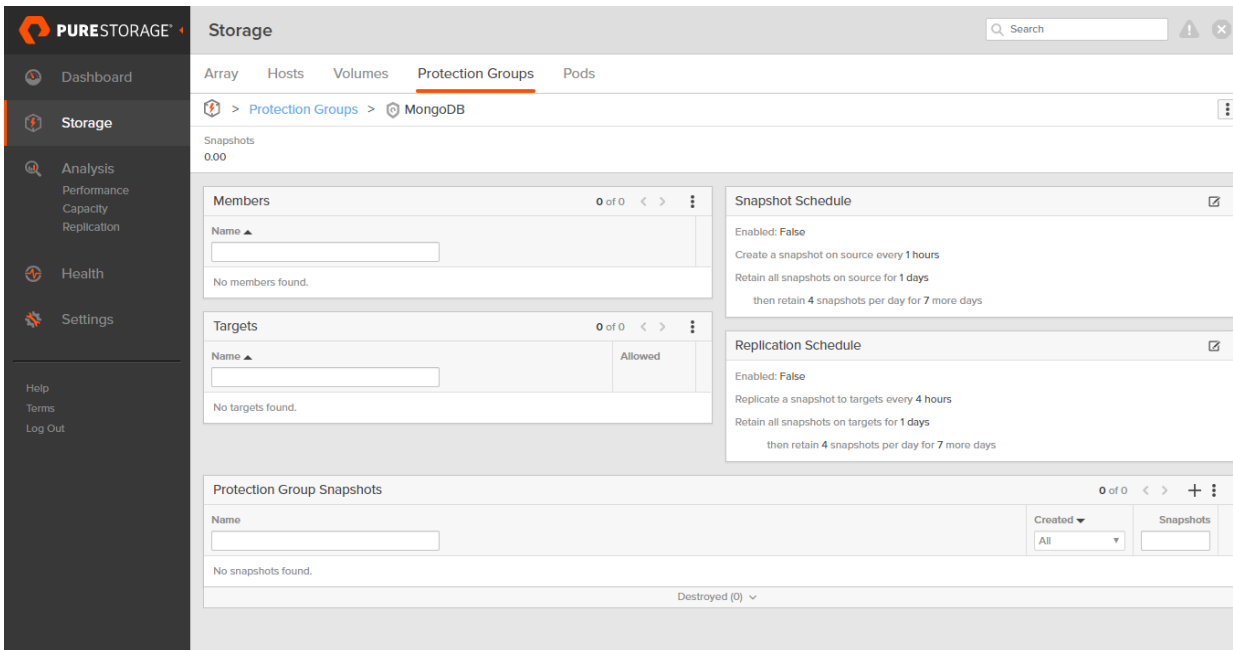
Step 2) Click the three **menu dots** (vertical ellipses) on the far right of the Offload Targets section and select **Connect to NFS Offload Target**. Then, enter the name, IP address or hostname, mount point and mount options as shown below.



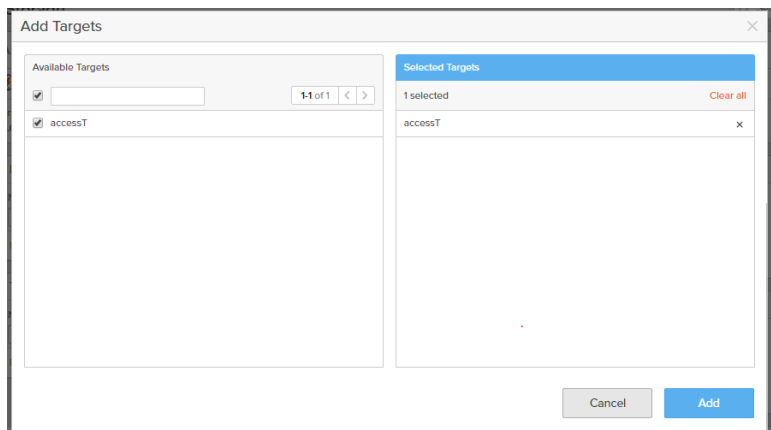

Step 3) Once the connection succeeds, the NFS target details entered will appear under the **Offload Target** with a green dot.



Step 4) Click on **Protection Groups** at the top and click on a **Protection Group** name (e.g. MongoDB) to get into page to specify the Target.



Step 5) In the **Targets** section, click on the 3 dots (vertical ellipses) on the far right and select **Add**. Select the NFS target (e.g. accessT) and click **Add**.



Step 6) The selected target will appear on the **Targets** section.

| Targets | | | 1-1 of 1 | < | > | ⋮ |
|---------|--|---------|----------|---|---|----|
| Name ▲ | | Allowed | | | | |
| accessT | | True | | | | 🗑️ |

VALIDATION OF CONFIGURATION

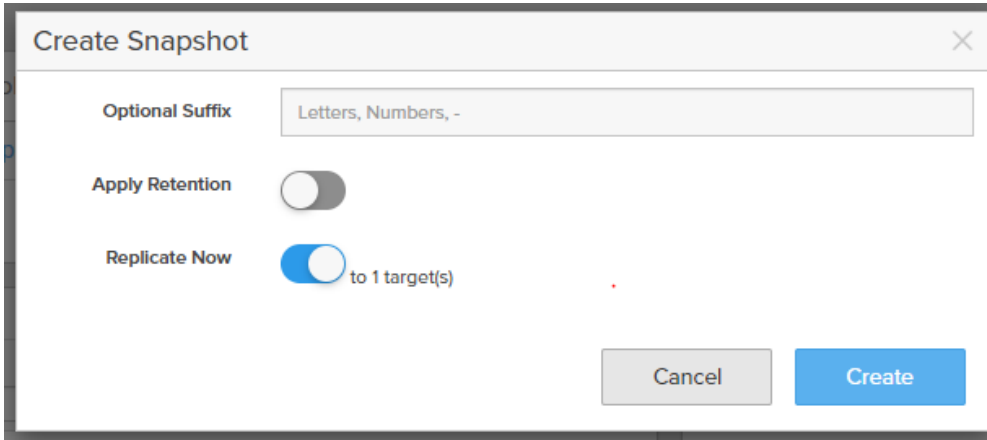
A snapshot and replication schedule can be defined as well, however, in this example, a manual snap and retrieval is described. Refer to the [Snap to NFS Overview and Administration](#) on how to setup a scheduled snapshot and/or replication.

Snap to NFS

Step 1) Click on "+" (plus) sign in **Protection Group Snapshots** section in the bottom of the Protection Group pane.

The screenshot displays the Pure Storage FlashArray web interface. The left sidebar contains navigation links: Dashboard, Storage (selected), Analysis, Health, and Settings. The main content area is titled 'Storage' and shows the 'Protection Groups' tab selected. The breadcrumb path is 'Protection Groups > MongoDB'. Below this, the 'Snapshots' count is 0.00. The 'Members' section shows a single member 'mongo01'. The 'Targets' section shows a single target 'accessT' with 'Allowed' status 'True'. The 'Snapshot Schedule' and 'Replication Schedule' sections are both 'Enabled: False'. The 'Protection Group Snapshots' section at the bottom shows '0 of 0' snapshots and a 'Destroyed (0)' count.

Step 2) Click on “Replicate Now” and then “Create” to initiate the send to NFS. This will create a snapshot locally of the volume(s) within the protection group and then replicate them to the Access Appliance.



Step 3) Click on the **Array** at the top and click on accessT in the Offload Target and view the snapshot from FlashArray to see the progress of the snap to NFS.

PURESTORAGE

Storage

Dashboard
Storage
Analysis
Performance
Capacity
Replication
Health
Settings
Help
Terms
Log Out

Array
Hosts
Volumes
Protection Groups
Pods

Array > Offload Targets > accessT

Status: connected
Protocol: nfs

Protection Groups

Name

Source

Remote

| | | |
|---------------------------|------------------|---------|
| sn1-m20-c08-20:MongoDB | sn1-m20-c08-20 | accessT |
| sn1-m20-c08-20:access | sn1-m20-c08-20 | accessT |
| sn1-m20-c08-20:access1 | sn1-m20-c08-20 | accessT |
| sn1-m20-c08-20:mongopg | sn1-m20-c08-20 | accessT |
| sn1-x10r2-d01-17:accesspg | sn1-x10r2-d01-17 | accessT |

Protection Group Snapshots

Name

Source

Remote

Created

Started

Completed

Transferred


Progress

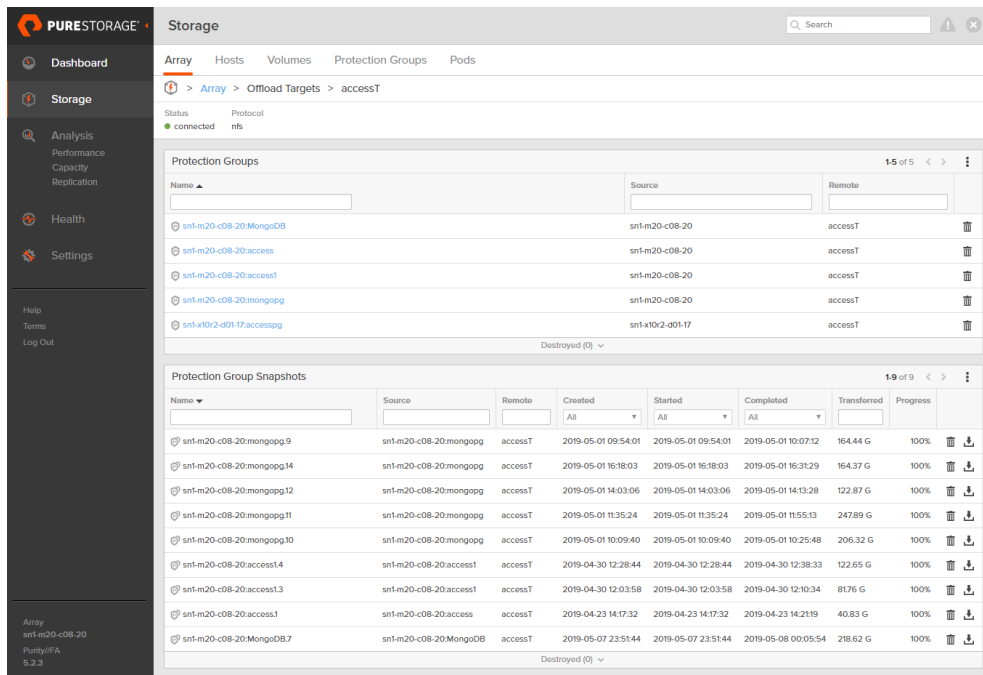
| | | | | | | | |
|---------------------------|------------------------|---------|---------------------|---------------------|---------------------|----------|---------|
| sn1-m20-c08-20:mongopg.9 | sn1-m20-c08-20:mongopg | accessT | 2019-05-01 09:54:01 | 2019-05-01 09:54:01 | 2019-05-01 10:07:12 | 164.44 G | 100% |
| sn1-m20-c08-20:mongopg.14 | sn1-m20-c08-20:mongopg | accessT | 2019-05-01 16:18:03 | 2019-05-01 16:18:03 | 2019-05-01 16:31:29 | 164.37 G | 100% |
| sn1-m20-c08-20:mongopg.12 | sn1-m20-c08-20:mongopg | accessT | 2019-05-01 14:03:06 | 2019-05-01 14:03:06 | 2019-05-01 14:13:28 | 122.87 G | 100% |
| sn1-m20-c08-20:mongopg.11 | sn1-m20-c08-20:mongopg | accessT | 2019-05-01 11:35:24 | 2019-05-01 11:35:24 | 2019-05-01 11:55:13 | 247.89 G | 100% |
| sn1-m20-c08-20:mongopg.10 | sn1-m20-c08-20:mongopg | accessT | 2019-05-01 10:09:40 | 2019-05-01 10:09:40 | 2019-05-01 10:25:48 | 206.32 G | 100% |
| sn1-m20-c08-20:access1.4 | sn1-m20-c08-20:access1 | accessT | 2019-04-30 12:28:44 | 2019-04-30 12:28:44 | 2019-04-30 12:38:33 | 122.65 G | 100% |
| sn1-m20-c08-20:access1.3 | sn1-m20-c08-20:access1 | accessT | 2019-04-30 12:03:58 | 2019-04-30 12:03:58 | 2019-04-30 12:10:34 | 81.76 G | 100% |
| sn1-m20-c08-20:access1 | sn1-m20-c08-20:access | accessT | 2019-04-23 14:17:32 | 2019-04-23 14:17:32 | 2019-04-23 14:21:19 | 40.83 G | 100% |
| sn1-m20-c08-20:MongoDB.8 | sn1-m20-c08-20:MongoDB | accessT | 2019-05-10 16:43:08 | 2019-05-10 16:43:08 | - | 1.77 G | 89.724% |
| sn1-m20-c08-20:MongoDB.7 | sn1-m20-c08-20:MongoDB | accessT | 2019-05-07 23:51:44 | 2019-05-07 23:51:44 | 2019-05-08 00:05:54 | 218.62 G | 100% |

30

© 2019 Veritas Technologies LLC. All rights reserved. Veritas, the Veritas Logo and NetBackup are trademarks or registered trademarks of Veritas Technologies or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

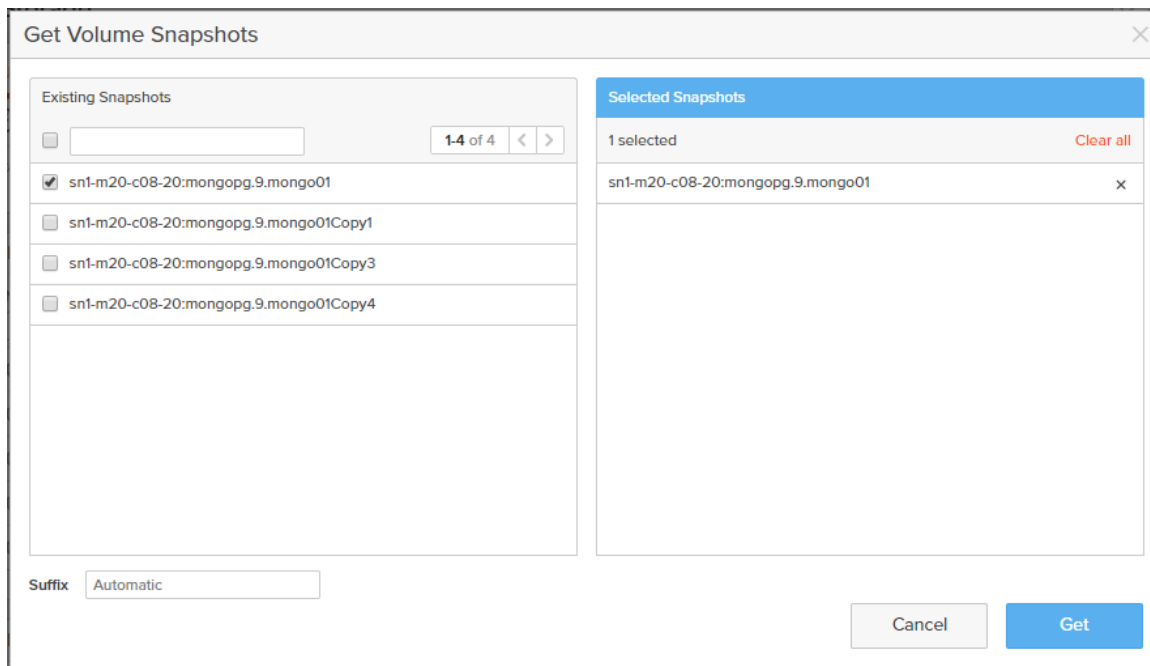
Snapshot Retrieval

Step 1) Click on the **Array** at the top and go to the **Offload** section and click on the **target** (e.g. accessT). From the list of **Protection Group Snapshots**, click on the **down arrow icon** .



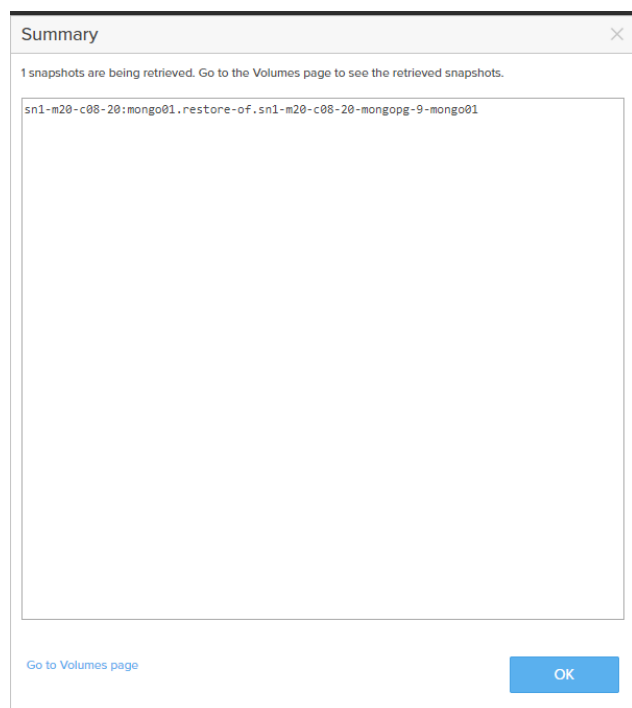
The screenshot shows the Pure Storage console interface. The left sidebar contains navigation options: Dashboard, Storage, Analysis, Health, and Settings. The main content area is titled 'Storage' and shows the 'Array' tab selected. Under 'Array', the 'Offload Targets' section is active, showing a list of 'Protection Groups'. Below this, the 'Protection Group Snapshots' section is displayed, showing a table of snapshots with columns for Name, Source, Remote, Created, Started, Completed, Transferred, and Progress. The table lists several snapshots, including 'sn1-m20-c08-20:mongopg.9' and 'sn1-m20-c08-20:mongopg.14'. The 'Completed' column shows the status of each snapshot, and the 'Progress' column shows the percentage of data transferred.

Step 2) Select the desired volume snapshot to retrieve and click **Get**.

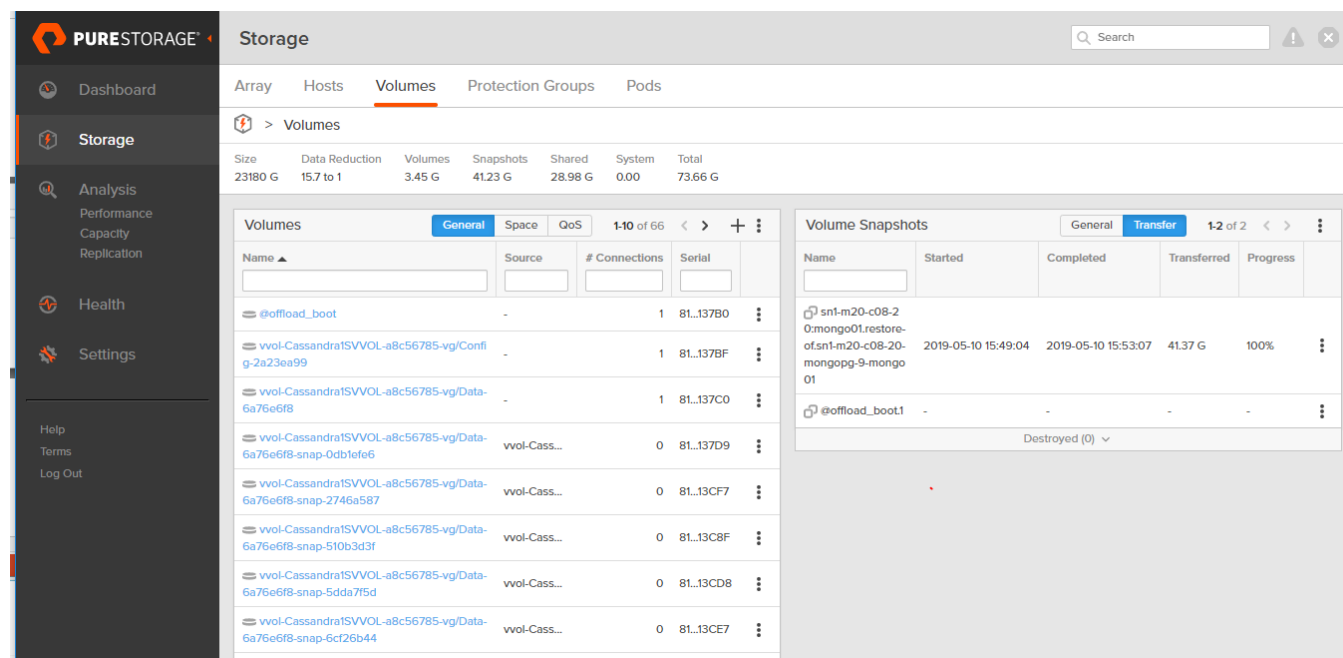


The screenshot shows the 'Get Volume Snapshots' dialog box. It has two main sections: 'Existing Snapshots' and 'Selected Snapshots'. The 'Existing Snapshots' section lists several snapshots, including 'sn1-m20-c08-20:mongopg.9.mongo01', 'sn1-m20-c08-20:mongopg.9.mongo01Copy1', 'sn1-m20-c08-20:mongopg.9.mongo01Copy3', and 'sn1-m20-c08-20:mongopg.9.mongo01Copy4'. The 'Selected Snapshots' section shows that one snapshot, 'sn1-m20-c08-20:mongopg.9.mongo01', has been selected. At the bottom of the dialog, there is a 'Suffix' field set to 'Automatic' and two buttons: 'Cancel' and 'Get'.

Step 3) View the Summary and either click **OK** or **"Go to the Volumes Page"**.



Step 4) From the Volumes page in the Volume Snapshots section, click on **Transfer** and the list of volume snapshots transferred is listed.



PERFORMANCE STUDY

This section describes the performance study details of when using Access 3340 Appliance as an NFS target for FlashArray volume snapshots. It includes information relating to the testing strategy, test environment, and results.

TESTING STRATEGY

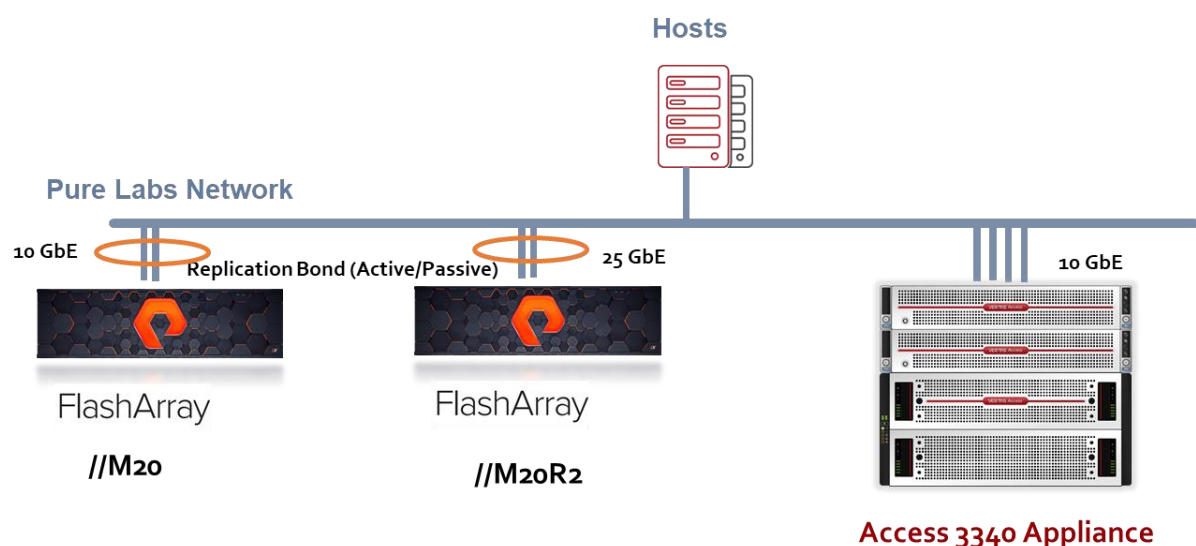
The performance study conducted involved the following two scenarios:

- Snap to NFS – sending volume(s) snapshots to the Access Appliance
- Snapshot retrieval – retrieving volumes(s) from the Access Appliance

Topology

The network topology of the test environment is shown in Figure 10. There were two FlashArray (//M20 and //M20R2) connected to the network utilizing the 10 GbE or 25 GbE replication bonded (active/passive) ports. The four 10 GbE ports of the Access Appliance were also connected to the network. The bulk of the tests used the FlashArray //M20. The other FlashArray //M20R2 was utilized only to observe how well the Access 3340 Appliance performs when multiple FlashArray are utilized.

Figure 11 - Network Topology



Hardware and Software

Table 2 provides the specifics relating to the hardware and software used in this study. As stated, the FlashArray //M20 was used for the bulk of the testing.

Table 2 - Hardware and Software Specifications

| Resource | Specification | Software |
|-------------------------|---------------|--------------|
| Pure Storage FlashArray | //M20 | Purity 5.2.1 |

| | | |
|-------------------------|--|--|
| | //M2oR2 | Share mounted with rsize=1048576, wsize=1048576, nfsvers=3 |
| Access Appliance | 3340 with one shelf containing 4TB disks | Access v7.4.2 <ul style="list-style-type: none"> • Storage Pool (entire disk shelf) • CFS File system type with defaults (8K block size, Simple layout, 10 TB in size) • File system exported as rw, sync • No bonding of network interfaces |

Snap to NFS Tests

The "Snap to NFS" tests involved sending 1 to "n" number of snapshot volumes to the Access Appliance. The source of the data was a MongoDB database. Multiple volume copies of the MongoDB database were created on the FlashArray. These volumes were then added into a protection group where the Access Appliance was specified as the NFS offload target. The size of each volume snapshot of the MongoDB was 40 GB. Via the FlashArray GUI or command-line interface, a create snapshot with "Replicate Now" to 1 target is initiated. Thus, the volume(s) snapshots are first taken, stored locally and then sent to the Access Appliance. Prior to each run, the local snapshots on the FlashArray and the Access Appliance were removed, the caches on the Access Appliance were cleared, and the tests repeated for results consistency. The FlashArray GUI indicates the start time, completed time and total amount of data transferred to the Access Appliance. The data was gathered, and the throughput was calculated using these values.

Snapshot Retrieval Tests

The volume snapshots sent in "Snap to NFS" tests were used for the retrieval tests. Prior to retrieval of the volume(s), the local snapshots on FlashArray were removed as well as the volumes that were used to create the snapshots. Also, the caches on the Access Appliance were cleared prior to each run. Then snapshots from 1 to "n" volumes were retrieved. The start time, end time and the amount of data transferred was captured from the GUI and the throughput was calculated using these values.

Mixed Workload Tests

The mixed workload consisted of sending and retrieving the same number of volume snapshots. The volume snapshots retrieved were different from the volume snapshots being sent. The caches were cleared on the Access Appliance prior to running the test. The test was run multiple times for consistency of the results. As in the other tests, the start and end times and the amount of data transferred or sent were extracted and throughput was calculated.

Monitoring Tools

Several monitoring tools such as sar and vxstat were utilized on the node of the Access Appliance. These tools were used to observe the system utilization of the Access Appliance such as CPU, memory, network and disks and assist in identifying any bottlenecks.

RESULTS

Performance Metrics


In this study, throughput in MB/s is the metric used to indicate the performance of sending and retrieving snapshots between the FlashArray and the Access Appliance. The average throughput is defined as the combined throughput for all the volumes sent and/or retrieved:

$$\text{Average Throughput} = \frac{\sum_1^{\text{\# Volume Snapshots}} \text{Quantity of Data}}{\text{Elapsed time for each snapshot to complete transfer}}$$

Snap to NFS

The throughput results for “Snap to NFS” are shown in Figure 11. The maximum throughput achieved was 232 MB/s where 4 volumes of size 40 GB each were sent from a FlashArray //M20 to the Access 3340 Appliance. The NFS target was configured with mount options of rsize=1MB, wsize=1MB and nfsvers=3 and was mapped to a single share mapped to a single file system, 1 virtual IP (1 network interface) and 1 node. Access exported the share with rw and sync options. The Access Appliance was minimally loaded in the sense of system resources and disk utilization. Thus, when sending volume snapshots from multiple FlashArray, the performance doubled to 457 MB/s. As previously discussed, when sending snapshots to two FlashArray and an Access Appliance configured to utilize 1 node, 2 file systems, and 2 network interfaces the combined throughput was 640 MB/s. The Access Appliance system resources utilization did go up however, they did not fully saturate.


Figure 12 - Snap to NFS Throughput Results



Snapshot Retrieval

Results of snapshot(s) retrieval are illustrated in Figure 12. As can be seen from the graph, the maximum observed throughput of 180 MB/s was when retrieving 5 volumes of 40 GB in size. The FlashArray //M20 and a single Access 3340 Appliance were utilized. Just like in the “Snap to NFS” scenario, the NFS target was configured with mount options of rsize=1048576, wsize=1048576 and nfsvers=3 and was mapped to a single share mapped to a single file system, 1 virtual IP (1 network interface), and 1 node. Access exported the share with rw and sync options.


Figure 13 - Snapshot Retrieval Throughput Results



Mixed Workload

For the mixed workload where 50% send and 50% retrieves of FlashArray snapshots were performed to and from the Access Appliance, the maximum combined throughput observed was 305 MB/s (send was 202 MB/s and retrieve was 103 MB/s) where 3 volume snapshots sent and 3 different volumes snapshots were retrieved as shown in Figure 14. The FlashArray //M20 and single Access 3340 Appliance with the same configuration as in the previous tests was used for the mixed workload tests.

Figure 14- Mixed Workload



DISCLAIMER

THIS PUBLICATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS PUBLICATION. THE INFORMATION CONTAINED HEREIN IS SUBJECT TO CHANGE WITHOUT NOTICE.

No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

ABOUT VERITAS TECHNOLOGIES LLC

Veritas Technologies empowers businesses of all sizes to discover the truth in information—their most important digital asset. Using the Veritas platform, customers can accelerate their digital transformation and solve pressing IT and business challenges including multi-cloud data management, data protection, storage optimization, compliance readiness and workload portability—with no cloud vendor lock-in. Eighty-six percent of Fortune 500 companies rely on Veritas today to reveal data insights that drive competitive advantage. Learn more at <http://www.veritas.com/> or follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

Veritas World Headquarters
500 East Middlefield Road
Mountain View, CA 94043
+1 (650) 933 1000
www.veritas.com

For specific country offices
and contact numbers,
please visit our website.

VERITAS™
The truth in information.