Nexenta Storage Appliance

User Guide – Tutorial – Reference

Version 1.0.7

Table of Contents

1. Introduction

Nexenta Storage Appliance (NexentaStor) is a software based network attached storage (NAS [1]) appliance that meets the current feature sets of the best of breed NAS, including unlimited snapshots, snapshot mirroring (replication), NFS v3/v4, CIFS, and easy management of extremely large storage pools. NexentaStor delivers richly featured software in the form of a software appliance that is trivial to install and easy to manage.

The product is designed and built to operate as 2nd tier storage alongside pre-existing commercial storage, providing online continuation of data for months and years, with tapes relegated to archival purposes only. The appliance is targeted for 2nd-tier NAS and iSCSI applications requiring extremely low cost storage as well as dramatically simplified provisioning, expansion, backup, replication and archiving. NexentaStor can also be used as a primary NAS in businesses that wish to expand at closer to commodity pricing.

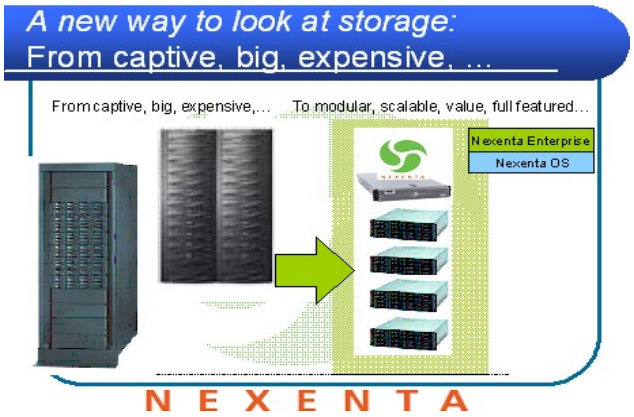

Login

NexentaStor is pre-configured with administrative user accounts: **root** and **admin**. After registering the appliance[2], you will receive the default password in auto-generated email, along with the product Registration Key. For more information, please see Sections "Getting Started with NMC" and "NMV Login".

Document Conventions

Command line example
Cross-reference
A notice, warning, conclusion, important remark.

1. NAS introduction, overview, history and usage is available on the web, see for instance http://en.wikipedia.org/wiki/Network-attached_storage
2. See NexentaStor Quick Start Guide document at http://www.nexenta.com/docs

Terminology

Term	Comment
NexentaStor	Nexenta Storage Appliance.
SA-API	Storage Appliance API. NMS (see next) is a sole provider of SA-API. The API provides access to the appliance's management objects and services. All client management applications use the same API (namely, SA-API) to monitor and administer the appliance. This ensures consistent view of the appliance from all clients, transactional behavior of all management administrative and monitoring operations, and easy third-party integrations.
NMS	Nexenta Management Server. There is only one server instance per appliance. The server provides public and documented Storage Appliance API (SA-API) available to all appliance management and monitoring clients, remote and local, including (but not limited to) NMC.
NMC	Nexenta Management Console. NMC can be used universally to view and configure every single aspect of the appliance: volumes and folders, storage and network services, fault triggers and statistic collectors. NMC communicates with the local NMS (see previous) and remote management consoles and management servers to execute user requests. Multiple NMC instances can be running on a given appliance. NMC is a single-login management client with a capability to manage multiple appliances and groups of appliances.
NMV	Nexenta Management View. Web client uses the same SA-API (above) to communicate with the NMS. NMV shows status of all appliances on the network, displays graphical statistics collected by "statistic collectors" (see below), and more. NexentaStor management software is further illustrated in Section "Functional Block Diagram" below.
Volume	NexentaStor volume is a ZFS[3] pool (a. k. a. zpool), with certain additional attributes. There is a one-to-one relationship between a volume and the underlying ZFS pool.
Folder	NexentaStor folder is a ZFS filesystem.
Auto-Snap	A type of appliance's storage service. The **`auto-snap`** service enables easy management of snapshots, providing regular multiple period scheduling on a per-folder or per-volume basis (with or without recursion into nested folders/filesystems). In addition, auto-snap allows to define a certain snapshot-retention policy. Snapshots can be kept for years, and/or generated frequently throughout the day.

3. Introduction to ZFS is available on OpenSolaris website at http://www.opensolaris.org/os/community/zfs/whatis/

Term	Comment
Auto-Tier	A type of the appliance's storage services. The **auto-tier** (or simply, "tiering") service can regularly and incrementally copy data from one host (local or remote, appliance or non-appliance) to a destination, local or remote, again of any type. NexentaStor **auto-tier** service runs on a variety of transports, and can use snapshots as its replication sources. This solution fits the more common backup scenarios found in disk-to-disk backup solutions. However, unlike regular backup solutions with only the latest copy available on the backup destination, this solution provides the advantage of both "the latest copy" and a configurable number of previous copies.
Auto-Sync	A type of the appliance's storage services. The **auto-sync** (or simply, "syncing") service will maintain a fully synchronized copy of a given volume or folder on another Nexenta Storage Appliance. Where tiering provides a copy, NexentaStor **auto-sync** service provides a true mirror, inclusive of all snapshots. The major difference between **auto-tier** (see previous) and **auto-sync** services that the latter transfers both data and filesystem metadata from its source to its (syncing) destination. This allows for standby hosts, as well as image-perfect recovery sources for reverse mirroring in case of a failure in the primary storage.
Auto-CDP	Automatic Continuous Data Protection (CDP) service. NexentaStor **auto-cdp** service provides remote mirroring capability. The service allows to replicate disks between two different appliances in real time, at a block level. Conceptually, the service performs a function similar to local disk mirroring scheme of RAID 1 except that in the case of **auto-cdp** this is done over IP network. Auto-CDP is distributed as Plugin (see below).
Trigger	Fault Triggers, or simply "triggers", are the appliance's primary means of fault management and reporting. Each fault trigger is a separate (runtime-pluggable) module that typically runs periodically at a scheduled interval and performs a single function, or a few related functions. Triggers actively monitor appliance's health, state of all its services and facilities, including hardware. See also "NexentaStor Runners" below.
Collector	Statistic Collectors, or simply "collectors" are, as the name implies, the appliance's means to collect network and storage statistics. A large number of network and storage IO counters is collected on a regular basis and recorded into SQL database. The data is then used to generate daily and weekly reports, and (via NMV - see above) various performance/utilization graphs and charts. The available collectors include **'volume-collector', 'nfs-collector', 'network-collector'**. See also "NexentaStor Runners" below.
Reporter	Yet another type of pluggable module tasked to generate periodic reports. The available reporters include **'network-reporter', 'nfs-reporter',**

Term	Comment
	'volume-reporter', 'services-reporter'. See also "Runners" below.
Indexer	Indexer is a special runner that exists for a single purpose: to index a specified folder, or folders. Once a folder is indexed, it can be searched for keywords, and the search itself takes almost no time. In a way, Indexers provide functionality similar to Internet search engines (think "Google"). However, in addition to searching the most recent raw and structured data, Indexer will allow you to search back in history - as long as there are snapshots available (that is, retained according to the **auto-sync/tier/snap** policies) to keep this history.
Runner	Triggers, Collectors, Reporters, and Indexers - also commonly called "Runners" - are pluggable modules that perform specific Fault Management, Performance Monitoring, Reporting, and archive Indexing tasks. All appliance's runners use the same SA-API (see above) provided by NMS (see above). The runners can be easily added – they are the source of future customizations in the product.
LUN	Physical and logical drives, attached to the appliance directly or via iSCSI or FC SAN, are commonly called LUNs. The terms "LUN", "hard drive" and "disk" are used interchangeably. See also http://en.wikipedia.org/wiki/Logical_Unit_Number
Zvol[4]	Emulated (virtual) block device based on a given appliance's volume. Can be used as additional swap partition but the primary usage: easy iSCSI integration. Zvol is a powerful and flexible tool also because of its tight integration with the appliance's storage services.
Plugin	NexentaStor extension, a pluggable module can be easily added (installed) and removed. Plugin uses the same SA-API (see above) as all the rest software components, and implements a certain well-defined (extended) functionality. At installation time, plugin integrates itself with the appliance's core software. The currently available plugins include: Auto-CDP, NMV based API browser, I/O performance benchmark, network and storage utilization. From the user's perspective, adding a plugin often[5] means that new menus, toolbars and new commands, man pages will show up in the GUI and CLI, respectively.. Each major component of the NexentaStor management software is, including the Management Server itself, is designed from ground up to be extensible with plugins.

4. In NMC, the capabilities and manual pages associated with the majority of the listed terms can be easily located using '**help keyword**' command (Section "Finding NMC commands")
5. Often but not always, since strictly NMS based non-UI extensions are supported as well

Functional Block Diagram

The following block diagram presents the major components of the NexentaStor management architecture. The management software includes Nexenta Management Server and its clients: NMC, NMV, NexentaStor *runners* (see Section "Terminology" above), NexentaStor plugins, 2nd tier storage services.

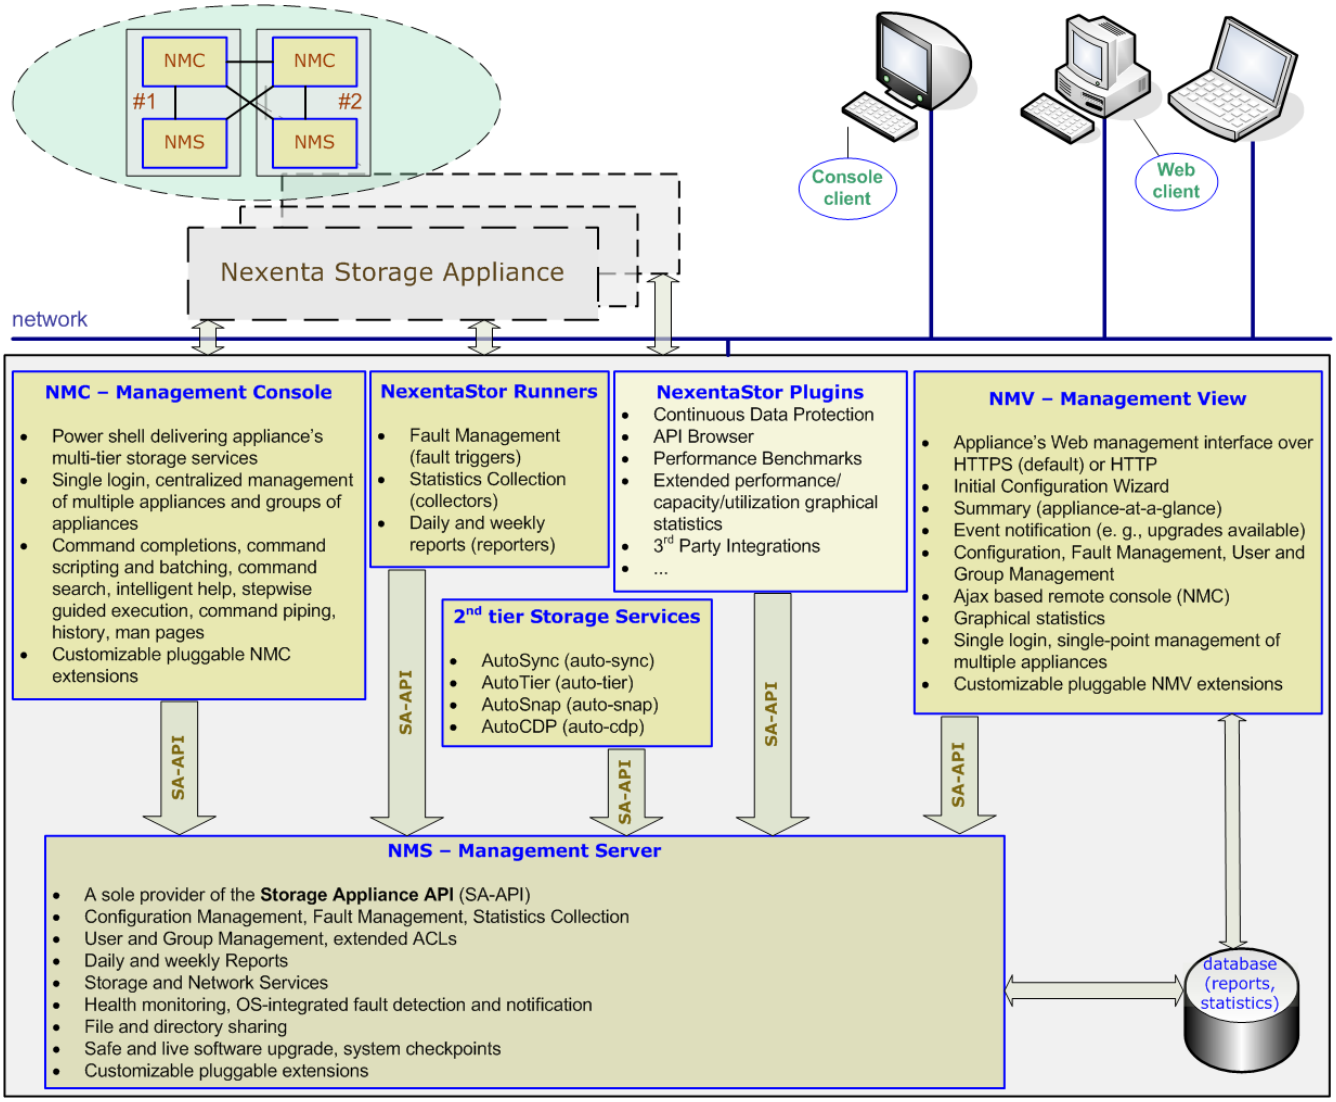

Getting Started

For step by step installation instruction, for initial configuration steps, and generally for guidance on how to get started with the appliance, please see NexentaStor Quick Start Guide document at http://www.nexenta.com/docs and/or Section "5. Initial Configuration Wizard" in this document. To get started with the Management Console, see Section "Getting Started with NMC", and for a quick introduction into Web GUI capabilities, see "6. Management GUI – Quick Introduction".

Finally, make sure to review NexentaStor F.A.Q. pages, which are constantly reviewed/updated.

2. NexentaStor Features – Overview

Overview

The Nexenta Storage Appliance product targets your end to end storage needs. Whether as your primary NAS for NFS or CIFS, or as secondary storage of the tiered or replicated variety, this ZFS-based product targets the full life cycle of your data, with flexible snapshots, retention policies, tiering and synchronization between various storage systems, and optional compression per storage container.

The product enables a wide variety of disk technologies, such as local SCSI, SAS, SATA and similar to be directly employed, but also adds the flexibility of iSCSI, Fibre Channel, and newer interconnected storage such as Infiniband-based SAS solutions.

NexentaStor also enables various network layer support, whether at the file system protocol layer, or the underlying hardware. First, for hardware, 10/100/1G BaseT and many 10G Ethernet solutions are directly supported, as well as aggregation (802.3ad) and multi-path I/O. For most installations, we recommend at least aggregated 100Mbps Ethernet at a minimum.

Currently the native CIFS server and Samba 3 variant of CIFS, and NFS v4/v3 are supported. For easy access from Windows, the provided WebDAV would be yet another option. The product also makes use of rsync, ssh, and zfs send/receive, CIFS and NFS transports for tiering and replication. In addition, block level replication (remote mirroring) is provided (Section "Setting up remote mirroring – CDP").

Associated with the above, LDAP have support for UID mapping, netgroups, X.509 certificate based client authentication, and other directory service based options one would expect in a NAS. Active Directory support is also provided. A full feature list can be found in the product data sheets provided with each major release, and is included below (Section "Data Sheet").

This document covers the major features of the product. We encourage you to examine both the online help (Section "Documentation – Registration – Technical Support") and usage comments of the product (Section "Tutorial"), as well as read further into the capabilities linked to off the web, at NexentaStor: an Overview.

Data Sheet

FEATURES	NexentaStor
2nd tier storage services	Included
Continuous data protection	Included[6]
RAID	RAID 0,1, 0+1, 5[7], 6[8]

6. Is available as a separate plugin
7. More exactly, RAID-Z – which is similar to RAID-5 in terms of data/parity scheme. RAID-Z eliminates a so called "RAID-5 write hole", and provides for better than RAID-5 performance.
8. More exactly, RAID-Z2 which is similar to RAID-6 in terms of parity/data protection and (similar to RAID-6) can sustain multiple simultaneous drive failures.

FEATURES	NexentaStor
USB storage device support	Included
Superior data protection	Included
Block size	Runtime variable, from 512B to 128KB
Limits	16 exabyte files and filesystems, 2^{64} devices in a pool
Compression	Included
Encryption	Included in near future release[9]
End-to-end checksum	Included
Quotas	Included
Volume support	Included[10]
CIFS	Included
NFS	Included
File sharing	Available via CIFS, NFS, RSYNC, FTP, WebDAV[11]
Snapshots	Included (up to 2^{48} snapshots per filesystem)[12]
Read/Write Copies (clones)	Included
Snapshots accessible over NFS	Included
iSCSI	Included[13]
Expansion	Simple[14]

9. Nexenta Systems, Inc may decide to provide this and possibly other new features as a value-adds available only for commercial deployment.
10. Practically unlimited number of devices, filesystems, and snapshots per volume; unlimited growth; integrated management
11. Web-based Distributed Authoring and Versioning (WebDAV) allows users to collaboratively edit and manage files on remote servers. WebDAV is yet another file/directory sharing option, in addition to provided (and somewhat more conventional) NFS, CIFS, RSYNC, and FTP.
12. Immediate (online) snapshots; rollback (to a snapshot) and cloning (of a snapshot) capabilities included as well
13. Software Initiator and Software Target, integrated console and web GUI management
14. May be TB-limited, as per NexentaStor product license. Please see the license text for details.

FEATURES	NexentaStor
Management Console	Included
Initial Configuration Wizard	Included
Management web GUI	Included
Appliance Upgrade	Automatic, Transactional, Unlimited Checkpoints
Windows backup and restore	Included[15]
Built-in Search Engine	Included
LDAP integration	Included
Active Directory integration	Included
Access Control List (ACL) integration	Included[16]
I/O benchmark	Included[17]
WORM	Write Once, Read Many (WORM) included[18]

Data Replication

The continuing growth of disk based storage has had two primary affects. The amount of data to backup is increasingly difficult to fit onto tape or within a backup window, and the costs of capacity of storage makes it feasible to build online backups out of disk subsystems themselves. One of NexentaStor's primary uses is in this new digital archiving role. Where as tapes will always find use, the development of disk based backup systems regulates tape to the final tier of archiving, where offline preservation is the requirement. You will find this product fits many roles, including primary storage, secondary storage to any primary storage array, and even remote site replication and archival.

What makes multi-tier storage possible in Nexenta's solution is the "**`auto-tier`**" service, which can regularly copy data from one source, local or remote of any nature, to a destination target again of any type. The only limitation is that at least one of either the source or destination must be local. In large arrays where the appliance encompasses both first tier and second tier storage, you'll even see local-to-local tiering. Tiering is accomplished by taking a given filesystem or share, breaking into smaller manageable chunks, and replicating that data at that point in time to another volume. Using snapshots at the target end, one can maintain a full efficient backup of the primary storage at unique intervals typical of backups. Where as you may have hourly and daily snapshots on your primary NAS, auto-tiering with snapshots will generally have daily, monthly,

15. Section "7. Windows Backup and Restore".
16. The ACLs are CIFS and NFSv4 compatible. The included user and group management provides an easy way to manage ZFS extended ACLs.
17. Included as a separate plugin
18. Included as a separate plugin

and even yearly snapshot points, with the same policies for retention of any given periodicity.

One commonly sets up up tiering locally, over NFS or CIFS, or using rsync directly with or without SSH. A simple example of tiering data from an NFS file server to our example volume would be to first create a filesystem to tier to and then to setup an **`auto-tier`** from our source NFS server.

`Auto-tier` and **`auto-sync`** are not limited just to the first two tiers, as tertiary tiering for more critical data is also common. As legal and business drivers dictate, tiering will also include access policy enforcement, limiting data access to restricted personnel to over longer periods of time.

As either a primary or secondary storage server, NexentaStor can pull or push data regularly at arbitrary intervals, transferring only the periodic changes seen. This can be done as frequently or as sparingly as required, thus being ideal for both large tiering as well as replication needs or for providing WAN-base off site mirroring.

NexentaStor provides the complete range of data replication services:

1) **Auto-Tier** - In the case of "**`auto-tier`**" (or simply, **`tiering`**) service, NexentaStor makes use of snapshots and user definable source and destination points to regularly replicate a single copy of a file system to another storage pool, whether local or remote. Using snapshots on the target end, the tiered copy may have arbitrarily different retention and expiration policies and can be administered separately. NexentaStor tiering service runs on a variety of transports, and can use snapshots as its replication sources. This solution fits the more common backup scenarios found in disk-to-disk backup solutions. The auto-tier service is not limited just to the first two tiers, as tertiary tiering for more critical data is also common. As legal and business drivers dictate, tiering will also include access policy enforcement, limiting data access to restricted personnel to over longer periods of time.
2) **Auto-Sync** - Another option provided is the "**`auto-sync`**" (or simply, **`syncing`**) service, which will maintain a fully synchronized copy of a given volume, file system, or emulated block device (a. k. a. **`zvol`**, Section "Using ZVOLs") on another NAS. Where tiering provides a copy, **`auto-sync`** provides a true mirror, an exact replica of data, inclusive of all snapshots and file system properties. Auto-sync uses the built-in ZFS snapshot capability to efficiently identify and replicate only changed blocks. This allows central mapping of multiple snapshots of a file system onto remote storage, all the while maintaining control of the retention and expiration of that data at the replication source. This facility is the most ideal for full disaster recovery.

Both **`auto-sync`** and **`auto-tier`** are schedulable, fault-managed, tunable NexentaStor Data Replication services that can be used in a variety of backup, archiving, and DR scenarios. The primary difference between **`auto-sync`** and **`auto-tier`** is that the former will transfer not only data (files, directories) but filesystem meta-data as well, including snapshots. By default[19], **`auto-sync`** is based on '**`zfs send/receive`**' transport.

19. A given auto-sync instance can use a variety of transport mechanisms. In addition to the default 'zfs send/receive' transport, rsync is also provided, as an option.

There is one common mistake, in terms of using **auto-sync**, and this is related to changes at its destination. Note that for **auto-sync** over '**zfs send/receive**', its destination cannot and should not be mounted, because otherwise it would be exposed to changes, including changes in the ZFS meta information. Any change in the data or meta-data at the destination will break the subsequent **auto-sync** run.

To re-cover, use NMC '**setup auto-sync <name> reset**' command, or a similar GUI management functionality.

NexentaStor includes a variety of management capabilities, to quickly determine the services' state and status. This includes daily and weekly reports, ability to show the current data transfer rate, and more.

3) **Hybrid-tier/sync** - NexentaStor provides a hybrid tiering-syncing service which enables a history of changes on the tiering destination. Unlike regular backup solutions with only the latest copy available on the backup target, this solution would have the advantage of both "the latest copy" and a configurable number of previous copies - the latter in accordance with the retention policy.
4) **Auto-CDP** - Continuous Data Protection, or remote mirroring. The service allows to replicate NexentaStor volumes remotely, in real time, at a block level. Conceptually, the service performs a function similar to local disk mirroring scheme of RAID 1 except that in the case of Auto-CDP this is done over IP network. For details, please see next Section "Continuous Data Protection".

For more information, please see Sections "Setting up a tiering service", "Auto-sync: the service that replicates both data and metadata", "Setting up remote mirroring – CDP".

Continuous Data Protection

NexentaStor **auto-cdp** service is based on Sun StorageTek Availability Suite, with the following functionality[20]:

Feature	Comment
Link neutral	Can use any network adapter that supports the TCP/IP protocol
Reverse replication	The direction of replication can be reversed at any time. The operation is also known as reverse update. The typical scenario includes: 1) failure of a primary volume 2) importing and continued usage of the secondary volume 3) reverse synchronization secondary => primary
Active logging	Continue logging operations whenever the Remote Mirror software is disabled or interrupted.
Multihop sets	Replicate data from one primary volume to a secondary volume; the secondary volume then replicates the data again to another secondary volume, and so on, in daisy-chain fashion.
Mutual backup	Concurrently transmit data copies to and receive data copies from a remote volume. Also known as a *bilateral relationship*.
Optimized resynchronization	Resynchronize volumes following disk, link, system, and storage platform outages; you only replicate those blocks that were modified since the last synchronization

20. For complete set of Sun StorageTek Availability Suite features, see http://docs.sun.com/source/819-6148-10/chap1.html#97729

Feature	Comment
RAID support	Use RAID volumes as part of your Remote Mirror software strategy. Volumes can be any RAID level.
Well known port and firewall	Port 121 is the default TCP/IP port used by the Remote Mirror software. The firewall must be opened to allow RPC traffic to/from this well-known port address.

NexentaStor data replication services address Disaster Recovery, Continuous Data Protection, and other emerging technology solutions as every aspect of business goes digital, and must be protected, preserved, or potentially expired as required. NexentaStor appliance can also be used just as a tool or conduit, providing all the data replication services when both the source and targets are remote.

Designed from ground up for consistency and ease of use, NexentaStor provides the same general interfaces to manage all storage services. All NMC commands, NMV UI pages, and all general principles apply to all.

Yet, there are important differences between the storage services. One size does not in fact fits all – if it weren't so a single replication service would suffice.

For more information, please see Sections "Setting up a tiering service", "Auto-sync: the service that replicates both data and metadata", "Setting up remote mirroring – CDP".

Fault Management, Reporting, HA

The appliance comes with a number of standard "runners" – pluggable modules that perform essential Fault Management, Performance Monitoring and Reporting functions.

The runners include Fault Triggers[21], Statistics Collectors, daily and weekly Reporters, and Indexers. See Section "Terminology" for more background on each of those distinct types. Typically, each runner is a simple module that performs a single identifiable task. For instance:

- fault trigger 'memory-check' monitors free memory in the appliance,
- collector 'nfs-collector' collects and records in the database NFS statistics,
- reporters 'volume-reporter' and 'network-reporter' generate storage and network utilization reports at configurable (default, weekly) intervals, respectively.

NexentaStor is designed and built to provide a reliable 24/7/365 service. To that end, Nexenta appliance provides a rich set of fault management and HA features. Part of the fault management facility is realized through Fault Triggers. A fault trigger, or simply, a trigger, is a special kind of a pluggable runner module that performs a certain fault management and monitoring operation(s). Each trigger monitors one, or a few related conditions, including:

- memory and CPU utilization
- state and status of appliance's hardware, including physical disks
- state and status of appliance's services and facilities
- storage services

21. Or, simply, "triggers"

- state and status of other runners
- state and operational status of the Nexenta Management Server.

Default triggers actively monitor your appliance's health and ensure that critical services run as you have scheduled them to. Triggers are also a source of future customization in the product.

Reporting makes use of this service monitoring as well as periodic analysis to alert you regularly and as necessary to the condition of your appliances. Regular daily and weekly reports on the overall health of your volumes and completed tasks for the day are available once the appliance is informed of your mailing preferences.

One of the first things an administrator of the appliance should do is enable reporting. This can be accomplished in a variety of ways – via "Initial Configuration Wizard", NMV web interface (Section "Management GUI – Quick Introduction"), and management console (Section "NMC – Management Console").

In NMC, run:

```
nmc:/$ setup appliance mailer smtp_server
```

and

```
nmc:/$ setup appliance mailer smtp_addresses
```

to set the outbound mail server and the recipient email addresses as necessary. Optionally, you will need to setup also user name and password associated with the specified SMTP server.

For more information on setting up appliance's mailer via NMC, please see the corresponding NMC man page:

```
nmc:/$ setup appliance mailer -h
```

Alternatively, use one of the available GUI management interfaces, see:

- Section ""Initial Configuration Wizard

  and/or

- Section "Management GUI – Quick Introduction"

With this done, regular reporting and alert notifications will commence.

For more information on NexentaStor runners, please see Section "Pluggable Modules – Runners".

NexentaStor supports a number of high availability (HA) and fault correction functionality. This certainly includes ZFS self-healing capability, advanced data protection, ZFS hotplug[22], and other enterprise-grade features delivered along with the filesystem.

In addition, each appliance monitors other appliances on the network (including both ssh-bound hosts and dynamically discovered NexentaStor appliances).

The appliance's unique upgrade/checkpointing facility utilizes snapshots, provides system checkpoints, and is designed to support 24/7/365 availability (see Sections "Appliance Software

22. Assuming the appliance's data volumes are configured with hot spares, ZFS hotplug feature will make sure that faulted drives are automatically replaced by those spares. For more information, please see "Note On Redundant Configurations".

Upgrade" and "Tutorial").

Finally, a number of HA features is being tested right now, with more explicit HA features planned for a near-future release.

Centralized Management

NexentaStor aims to provide both command line as well as web interfaces to configure and monitor all facets of your storage. In both interfaces, the concept of "**switch**"ing to another appliance is implemented, allowing you to connect to other authentication-bound appliances. When you switch, the focus of your actions are relative to that appliance, allowing single point administration of multiple appliances. Management of group of appliances is supported as well, so that with a single keystroke from NMC you will be able to run management operation(s) on all appliances in the group. The following NMV screenshot shows two appliances that can be managed individually, and/or as a group:

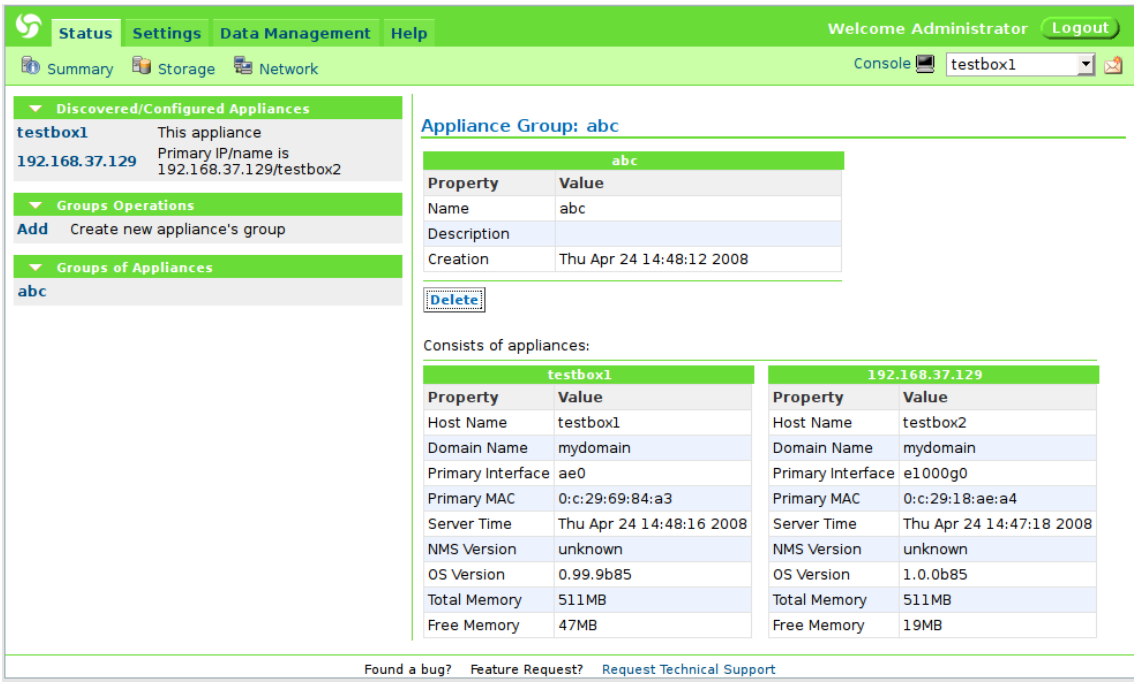

For more information, please refer to Section "Command Reference – switch". NexentaStor management suite is designed from ground up to provide administrators with global access management.

The various services, such as **auto-tier** and **auto-sync**, make use of binding of trusts between appliances when using certain protocols. The NexentaStor management layer makes use of this strong authentication, and its further used for data path discovery and the above "**switch**" command. Basic discovery of source or destination volumes and paths are available today in the command line interface. The appliance's web interface (Section "Management GUI – Quick Introduction") takes this further to make creation of a multi-tier architecture as easy as possible.

Storage Appliance API

All client management applications (including management console and management GUI, fault

triggers and storage services, etc.) use Storage Appliance API (SA-API) to communicate to the Nexenta Management Server. This ensures consistent view of the appliance from all clients, transactional behavior of all management administrative and monitoring operations, and also opens possibilities for third-party integrations.

SA-API is public, documented, and available. Please see [2] for details and downloads.

The following diagram illustrates the common-API-connectivity concept: it shows two management consoles running on two different appliance "boxes" #1 and #2 simultaneously connected to two management servers, local and remote[23]. The console clients, as well as all the rest NexentaStor management clients, use the same networked public API, with bindings provided for a variety of programming languages including C, C++, Perl, Python, and Ruby.

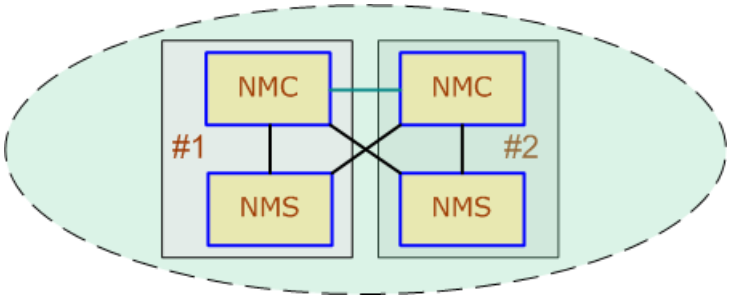

NexentaStor API reference documentation is currently provided along with NexentaStor Software Development Kit (SDK, see [2]). The NexentaStor Software Development Kit (SDK) is intended for software developers and users of the NexentaStor appliance. The SDK can be used to take a full advantage of the NexentaStor capabilities, which include unlimited including unlimited incremental backups or 'snapshots', snapshot mirroring (replication), block level mirroring ('CDP'), integrated search, and the inherent virtualization, performance, thin provisioning and ease of use of the ZFS file system. The entire suite of appliance's capabilities is available via Nexenta API (SA-API).

Access to the appliance via SA-API is secure (Section "Secure Access" below).

For the latest releases, updates, and documentation, please visit http://www.nexenta.com

Management Software

NexentaStor management suite provides intuitive graphical and management console interfaces for centralized management of Nexenta storage solutions. Designed from ground up for ease-of-use, NexentaStor management software places the cutting edge NAS functionality at your fingertips. With NexentaStor management software you can:

- create, grow, and destroy, export and import storage volumes
- share data folders and files via NFS, CIFS, RSYNC, FTP, and WebDAV[24]

23. The connectivity between two management consoles (shown in a different color) is a separate NexentaStor feature called "fast switch". See "Command Reference" for details.
24. Web-based Distributed Authoring and Versioning (WebDAV) allows users to collaboratively edit and manage files on remote servers. WebDAV is yet another file/directory sharing option, in addition to provided (and somewhat more conventional) NFS, CIFS, RSYNC, and FTP. For a background on WebDAV, please see http://en.wikipedia.org/wiki/WebDAV. WebDAV community and resources can be located at http://www.webdav.org/

- take snapshots, custom and automated, periodic and on-demand, and implement a user-defined snapshot retention policies
- automatically and periodically generate storage and service reports
- run a variety of storage services, periodically and on-demand
- discover, configure, and create iSCSI sessions, deploy iSCSI-attached LUNs[25]

  and much more

In the end, this reduces the overall storage management burden, maximizes storage utilization and optimizes IT resources and accessibility.

The management suite consists of the following components:

- Nexenta Management Server (NMS)
- Nexenta Management Console (NMC)
- Nexenta Management View (NMV)
- Storage Services
- NexentaStor Runners: Fault Triggers, Statistic Collectors, and Reporters.

For management software architecture and API, please see [3] (Section "References").

Appliance Software Upgrade

NexentaStor provides advanced software upgrade and system checkpointing functionality. The feature is designed from bottom up to allow uninterrupted, safe, and timely delivery of security fixes, new features and extensions. Once/if upgrades are available, users get notified via Nexenta management interfaces, for instance:

Notice the little "you've got mail" icon on the top right of the NMV screen. After the appliance get upgraded, this icon disappears (and will not show up until new upgrades become available).

Two types of upgrades are supported:

- ✔ safe upgrade, via cloning a system folder (see Section "Upgrading appliance software – Safe")
- ✔ in-place (live) upgrades (see Section "Upgrading appliance software – Live")

Unless the "safe" option is explicitly specified, the system will automatically detect whether the upgrade will require reboot, and if so, it will clone the active system folder and safely perform the software upgrade within this clone. A user then has two options: reboot into the new (upgraded) system folder or continue working (and possibly activate the upgrade and reboot into it later).

The live upgrading, as the name implies, happens in-place on the running system, and without reboot. NexentaStor upgrades are transactional – each upgrade has a start (atomic creation of a unique system checkpoint), and an end (activation of upgraded system folder). Being a complex transaction, an upgrade can be totally rolled back, thus resetting appliance to an exact atomically saved system state (checkpoint) in the past. The appliance supports multiple system checkpoints – the number of checkpoints is configurable. The upgrades are fast and easy to manage. And most importantly, the upgrades are safe.

25. See http://en.wikipedia.org/wiki/Logical_Unit_Number. Often, the terms "LUN", "hard drive" and "disk" are used

The philosophy behind this functionality is simple: provide 24/7/365 service. The appliance upgrade functionality preserves the running (active) system at all times. For more information and usage examples, please see Section "4. Tutorial".

Indexing NexentaStor Archives

NexentaStor built-in Search Engine (beta[26]) is realized via a a special kind of a "runner" called "indexer" (Section "Terminology"). To create an indexer simply means associating the latter with a storage folder – for the subsequent indexing of this folder snapshots and the most recent content. Symmetrically, to destroy an indexer means to disassociate it from the corresponding folder, and in addition, to remove from storage all the corresponding indexing data.

Once a folder and possibly some/all of its snapshots get indexed, the corresponding data can be searched. Indexing a large amount of storage takes time, which is why indexer works offline, according to the specified schedule. However, the search on an indexed folder is almost immediate. Search results integrated both into NMV GUI and management console (although Nexenta Management View is definitely a preferable way to view the results).

Features

- ✔ Supports Unicode (including code-points beyond the BMP), and stores indexed data in UTF-8.
- ✔ anked probabilistic search - important words get more weight than unimportant words, so the most relevant documents are more likely to come near the top of the results list.
- ✔ Relevance feedback - given one or more documents, Indexer can suggest the most relevant index terms to expand a query, suggest related documents, categorize documents, etc.
- ✔ Phrase and proximity searching - users can search for words occurring in an exact phrase or within a specified number of words, either in a specified order, or in any order.
- ✔ Full range of structured boolean search operators (for instance: "stock NOT market", "title AND virtualization", "NexentaOS OR NexentaStor", etc.). The results of the boolean search are ranked by the probabilistic weights. Boolean filters can also be applied to restrict a probabilistic search.

  The following screenshot illustrates expression power of the appliance's built-in search engine. This performs a query "CIFS AND MAYBE snapshot AND NOT upgrade" on yesterday's content (note Searchable Snapshot field below) of a given folder.

26. This development is in progress and will be finalized in the upcoming near-future releases.

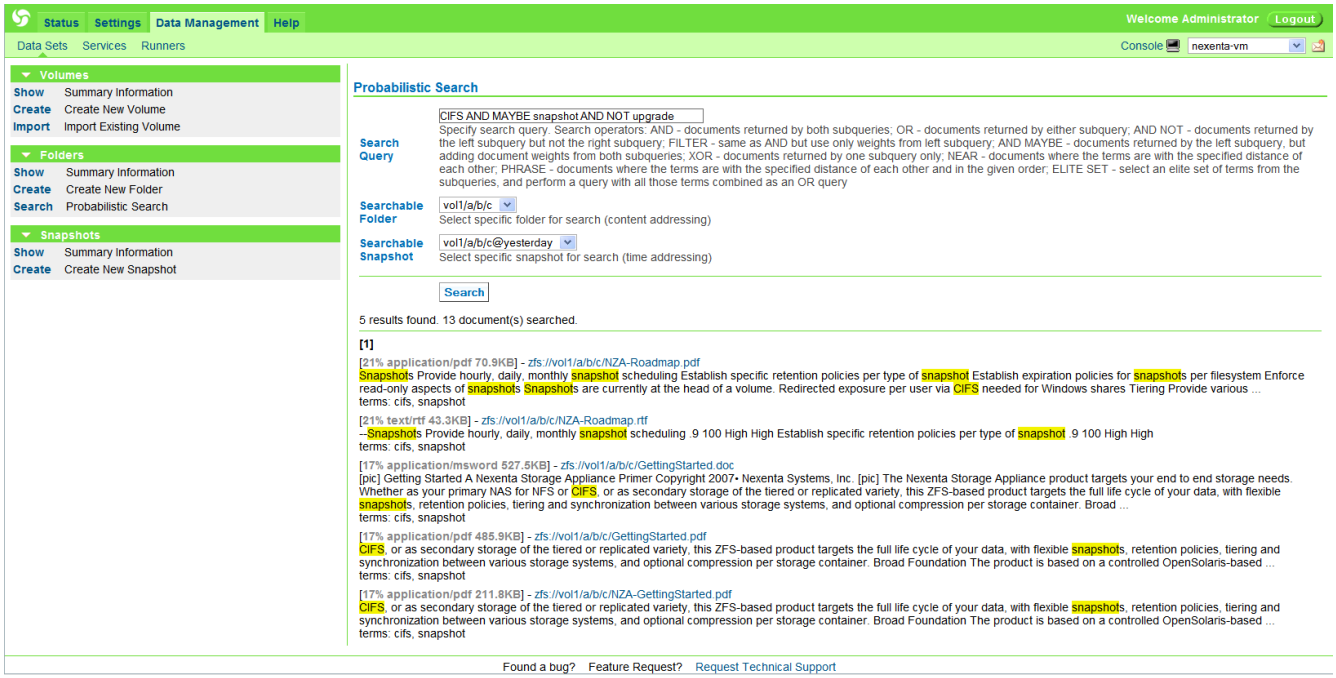

- ✔ Stemming of search terms (e.g., a search for "football" would match documents which mention "footballs" or "footballer"). This helps to find relevant documents which might otherwise be missed. Stemmers are currently included for Danish, Dutch, English, Finnish, French, German, Hungarian, Italian, Norwegian, Portuguese, Romanian, Russian, Spanish, Swedish, and Turkish.
- ✔ Wildcard search is supported (e.g. "nexen*").
- ✔ Synonyms are supported, both explicitly (e.g. "~cash") and as an automatic form of query expansion.
- ✔ Suggest spelling corrections for user supplied queries. This is based on words which occur in the data being indexed, and will work even if a word cannot be be found in a dictionary (e.g., "nexenta" would be suggested as a correct for "nexnenta").
- ✔ Supports database files > 2GB - essential for scaling to large document collections
- ✔ Allows simultaneous update and searching. New documents become searchable right away.

Formats

- ✔ Supports the following "raw" data formats:
  HTML, PHP, PDF, OpenOffice/StarOffice, Microsoft Word/Excel/Powerpoint/Works Word Perfect, AbiWord, RTF, Perl POD documentation, and certainly, plain text.
- ✔ Supports the following "structured" data formats:
  MySQL, PostgreSQL, SQLite, Oracle, DB2, MS SQL, ODBC, LDAP

Operations

Here is how indexer works. The indexer runs in a background, i.e., offline, at a scheduled intervals. Similar to other periodic services and runners (see Section "Terminology"), the intervals can be specified in multiple units of minutes, hours, days, weeks, or months. Each time the indexer runs, it "notices" new snapshots – that is, snapshots created since its last run on a given folder. It is expected that those new snapshots are created in accordance with the **auto-sync**, **auto-tier**, **auto-snap** policies provisioned for the appliance – for more information on the appliance's 2nd tier services see Sections "Data Replication" and "Multi-Tier Storage".

The indexer then proceeds to indexed those new snapshots into space-optimized index database. Snapshots are indexed only once. Every scheduled pass may index up to a configurable number snapshots. If an indexed folder gets destroyed, the corresponding index get destroyed as well. If a snapshot get destroyed, the corresponding part of the index database is cleaned up too.

Note that a folder without any snapshots can be indexed as well. The feature is called "continuous indexing" and is enabled by default.

The rest of this section is provided in a form of questions and answers.

Question: How much space will index take?

Answer: This is an estimate: it usually takes from 0.001% to 0.1% of the size of the original source. Note that the indexing database is stored in a NexentaStor folder with enabled compression, which should provide for even better space utilization.

Question: Where the indexing database is located?

Answer: Indexing database root is global and configurable through NMS option 'indexroot'. The default value is: **syspool/.index**

Question: How can I manage (enable, disable, tuneup) an individual index in NMC?

Answer: In the management console, run:

```
nmc:/$ setup indexer <name> [enable|disable|property] ...
```

In NMV, simply point your Internet browser to /data/runners?selected_runner=<name>

Question: Any guidelines to make the indexer run faster?

Answer: Use option "recursion_limit". It will limit the depth of source scanning. Default is 0 (unlimited). For instance, set it to 3 and see if it helps. In addition, limit a number of snapshots to index during a single indexer run - the corresponding option is called "snaps_per_schedule". The default is value 3, that is, by default the indexer, when scheduled, will work on up to 3 new snapshots.

Question: How to perform a search?

Answer: The built-in search engine supports full range of structured boolean search operators. The following NMV screenshot illustrates a search for documents that contain the word 'CIFS' and possibly 'snapshot' but not 'upgrade'.

Notice the discovered document URLs – each of those can be opened in your Internet browser for further research. For detailed usage guide with more illustrated examples, please see the corresponding sections under "NMC – Management Console" and "Management GUI – Quick Introduction".

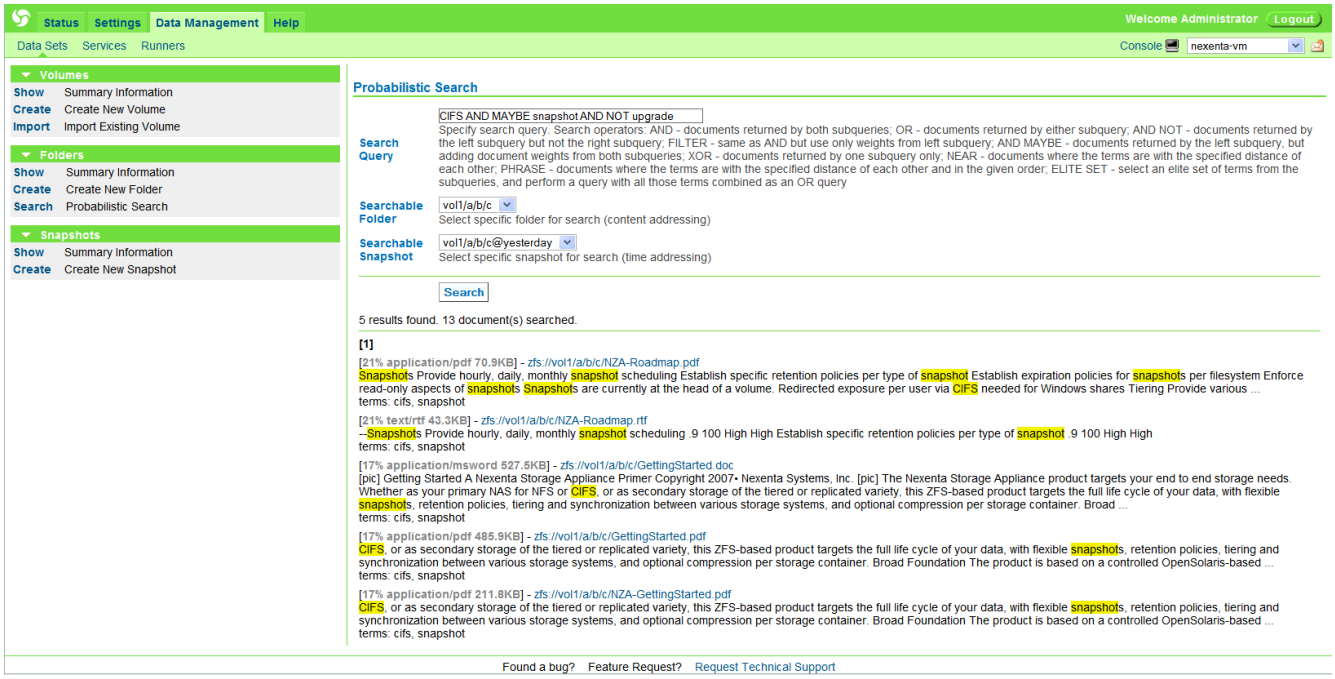

LDAP integration

The Lightweight Directory Access Protocol (LDAP) is a common protocol interface to Network Directory Services. Widely deployed directory services are Domain Name Service (DNS), NIS (Network Information Service), etc. They provide the clients with information such as host IP addresses, usernames, passwords and home directories. LDAP is a widely-deployed, simple and efficient network protocol for accessing information directories. LDAP typically runs over TCP; it has the potential to consolidate existing Network Directory Services into a single global directory.

NexentaStor provides easy to use LDAP integration, specifically for usage in the mixed CIFS/NFS environments. In addition, LDAP user and group management can be deployed with NFSv4 – the default NFS version provided by the appliance. In general, LDAP based user and group management is required to consistently utilize ZFS extended Access Control List (ACLs) across heterogeneous file services instead of POSIX permissions and attributes.

To use NexentaStor with LDAP server, make sure the server is available.. You will need your base DN, with either anonymous or authenticated SASL binds (the latter requiring account DN and password), and netgroup, user, and group subtree DNs if known. Netgroup (a group of hosts) is only necessary if currently supported by the LDAP server and is of interest.

For more information, please see Section "Using LDAP" of the chapter "4. Tutorial".

Note that in addition to Unix based LDAP, NexentaStor provides Active Directory integration - an implementation of LDAP directory services by Microsoft for use primarily in Windows environments (Section "Using Active Directory").

Finally, NexentaStor LDAP client provides integrated ability to authenticate itself using X.509 certificates. Management console and management UI both provide the corresponding interfaces[27].

27. For the management console, please check out '**`setup network ssl-bind`**' and associated '**`ssl-unbind`**' and '**`show`**'

User, Group and ACL Management

NexentaStor provides native extended Access Control Lists (ACLs), capable of handling CIFS ACLs, as well as NFSv4 ACLs, as well as POSIX permissions natively in the same filesystem.

The appliance supports full management of per-user, per-group, per-folder ACLs in its user interface, while also populating the system with accounts and groups that you may have already defined in Active Directory or other LDAP-based directory service.

NexentaStor User and Access Control management has the following characteristics:

- ✔ Support both local and LDAP (or AD) managed users and groups. In LDAP or Active Directory configurations, the local users and groups can be used to override centralized settings.
- ✔ Native extended Access Control Lists (ACLs), that are both CIFS and NFSv4 compliant.

Following are two screenshots that show, first, appliance users (most of which are retrieved from LDAP in this case[28]), and the management GUI capability to administer access control to a given folder (and its subfolders – all operations on ACLs are recursive, to reduce the amount of administration).

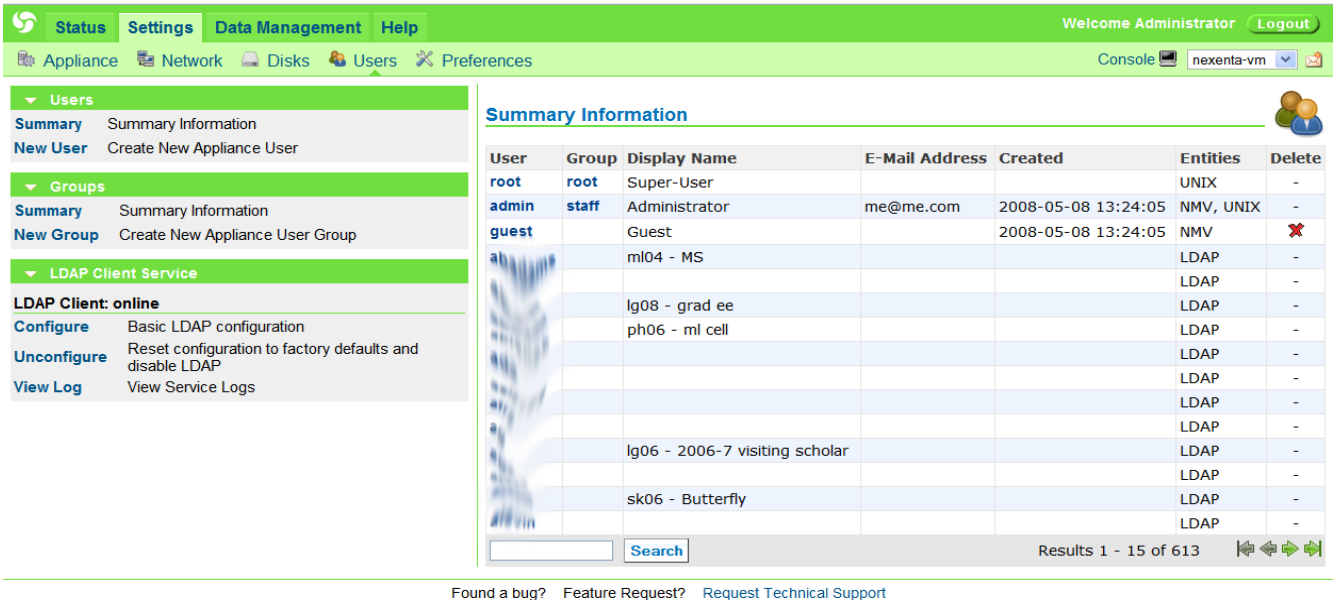

Notice that in the case below a local 'test-user' and LDAP-defined 'rfgroup' are granted a special set of permissions:

interfaces.

28. User names on the picture are intentionally skewed to protect users' identities.

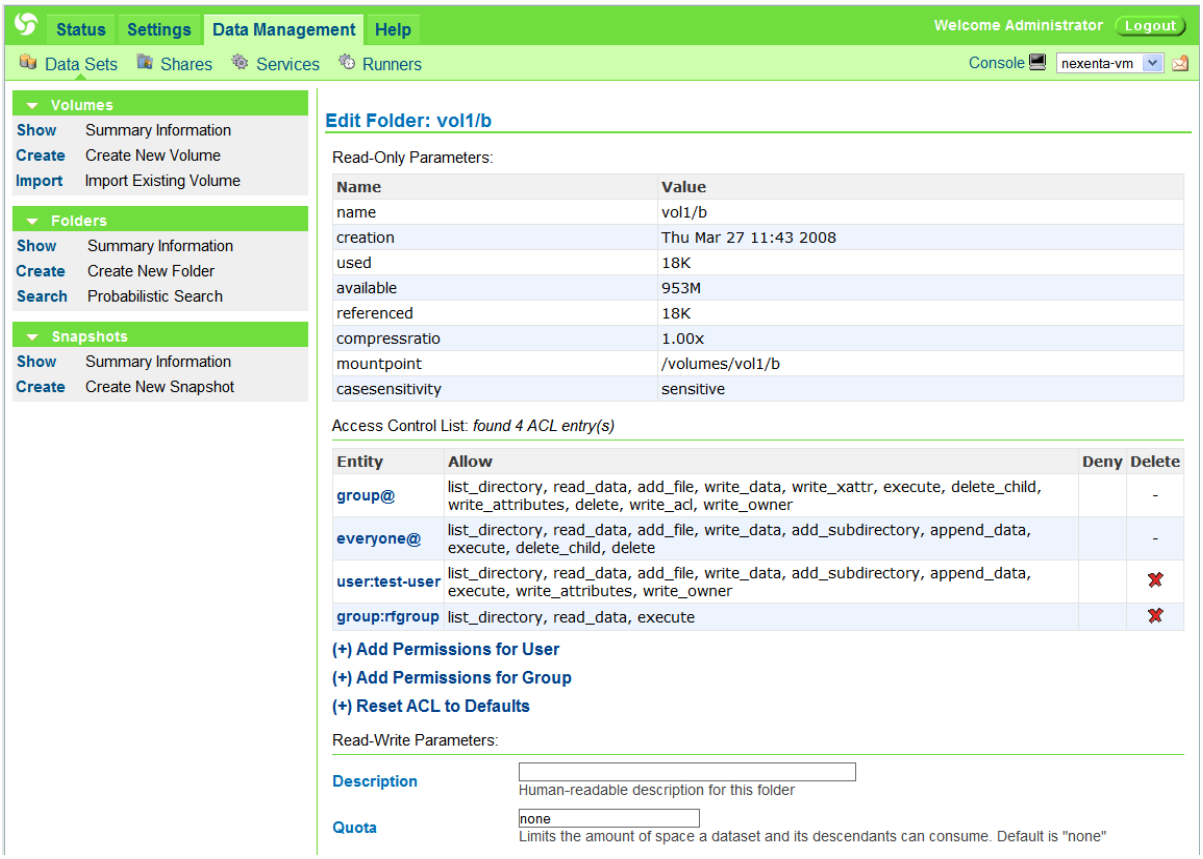

As always, NexentaStor CLI management client provides the same capabilities – via command line. The users and groups can be retrieved (that is, 'shown'), created and deleted, extended permissions can be modified and all the rest related management operations can be executed using either NMV or/and NMC.

NexentaStor ACLs are native across ZFS, CIFS, and NFSv4, and as such have no conflict in how they are operated on. Generally, one accomplishes ACL management via the following tasks:

1) local user or LDAP configuration
2) definition of per-user or per-group capabilities per volume or folder
3) overall management of ACLs and ACEs system wide, allowing overriding of end user activity via CIFS/NFS

A note on NFSv3 vs. ACL

NFSv3 relies on POSIX permissions, which are a subset of ZFS extended ACLs. Thus, NFSv3 clients will only check with the POSIX level permissions.

However, even though POSIX permissions may otherwise grant a permission to a user, that will be nullified if the extended ACL on the server is defined and otherwise denies that access.

Guidelines on using (or not using) extended ACLs in a mixed (NFSv4, NFSv3, CIFS) environments are discussed elsewhere.

For more information, please see Section "Notes on User Management and Access Control".

Jumbo Frames

The screenshot (below) illustrates network interface configuration, and in particular, an easy way of using Jumbo frames. The latter boils down to setting MTU size, with a considerable complexity of vendor-specific operations on the respective network interface drivers completely hidden from the user. The appliance removes this complexity and recovers from any error resulting from a misconfiguration.

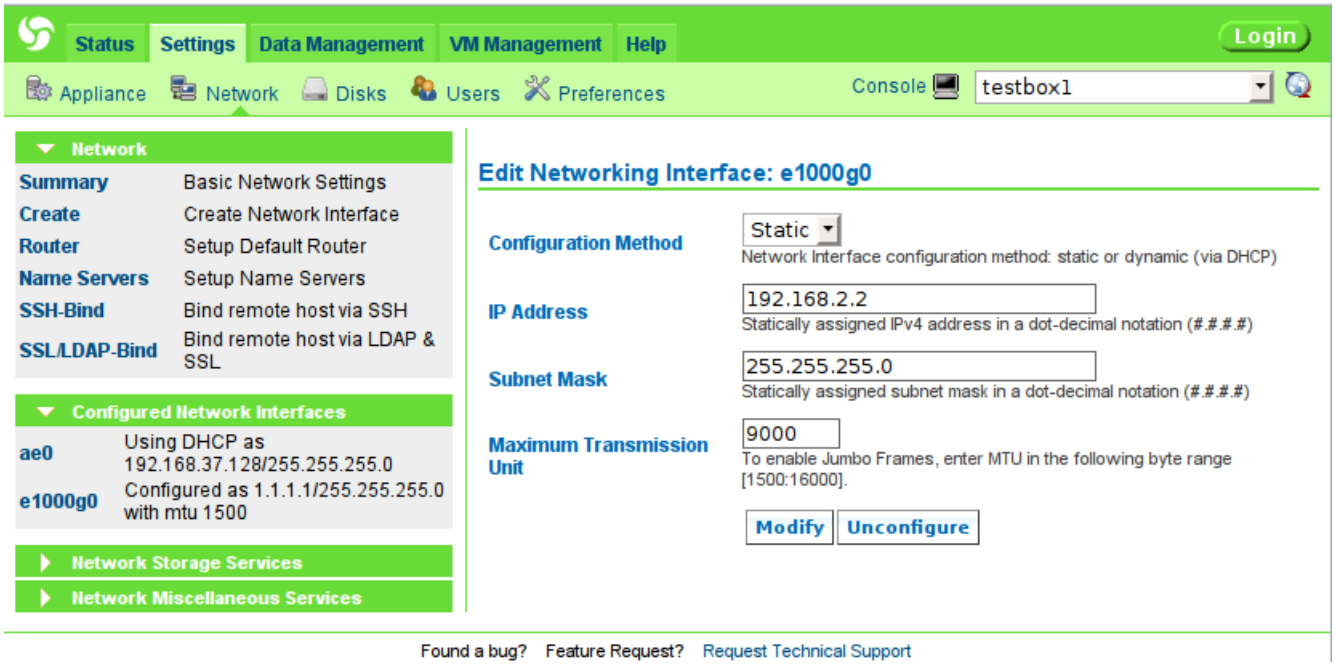

As always, the same functionality is available via NMC, via '**`setup network interface`**' operation.

WORM

WORM (Write Once, Read Many) is attributed to data storage which is physically incapable of being rewritten. WORM storage media includes CD-R and DVD-R.

NexentaStor provides WORM capability via pluggable module. Any data folder can be made Write Once, Read Many, which will then effectively prevent the folder's content to be modifiable. The latter includes the entire content, including files, directories, and existing sub-folders.

Any data folder without exception can be WORM activated; the opposite operation, however, is not permitted: WORM cannot be deactivated. Hence, an important note:

> Once activated, WORM cannot be undone. WORM folder stays WORM for its lifetime.

WORM protects the contents and to this end, certain functionality is limited. Specifically:

- sub-folders of the WORM folder cannot be created[29]

29. The corresponding functionality may be added in the future

- rollback of the WORM folder is not permitted (and will fail)
- WORM folder can be destroyed only after its entire content expires (see option **keep-days** below)

In NMC, run:

```
nmc:/$ setup folder <name> worm
```

to review available WORM sub-commands,

```
nmc:/$ setup folder <name> worm activate
```

to WORM the specified folder, or:

```
nmc:/$ setup folder <name> worm activate -h
```

for details and usage examples.

At WORM activation time you can specify **keep-days** – the number of days to ensure read-only access to each file in the WORM folder. The option allows to set WORM time limit on a per file basis. If omitted, write access is never allowed (which is equivalent to saying "keep forever").

Secure Access

NexentaStor appliance provides secure access to other NexentaStor appliances as well as administrative management client applications on the network. The picture below illustrates an appliance (with its main functional blocks) being accessed from/by another appliance and two management clients. The inter-appliance access is executed either via SSH, or via SA-API (Section "Storage Appliance API"), or both. All management client applications, whether developed internally by Nexenta Systems, Inc and/or by 3rd parties, access appliance via SA-API.

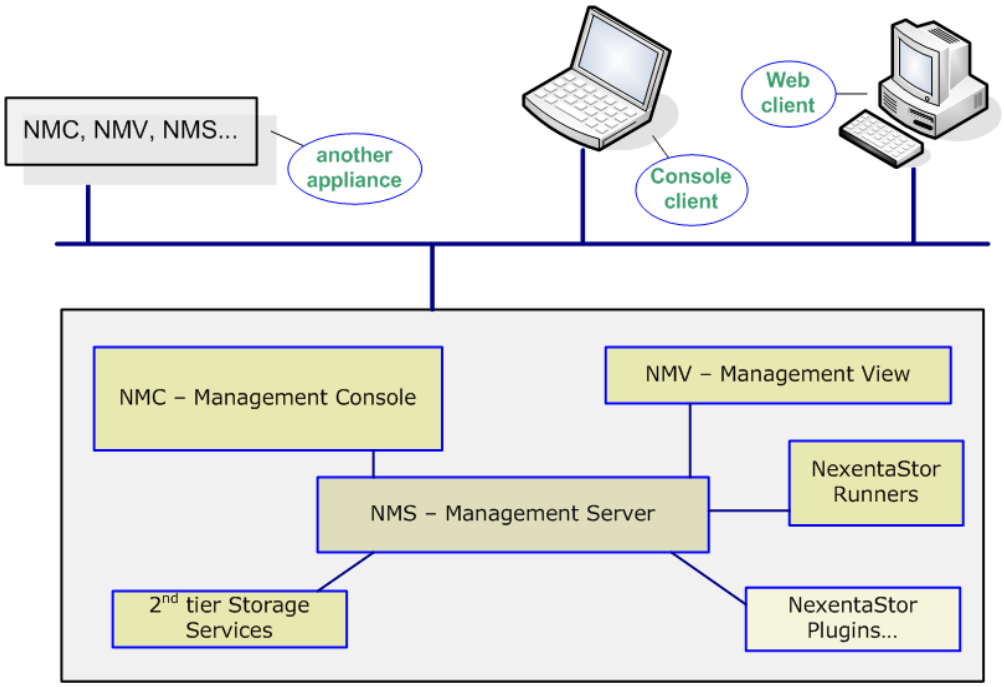

In all cases, access to appliance requires client authentication. NexentaStor supports two authentication mechanisms:

1. via IP address of the client machine
2. via ssh-keygen generated authentication keys.

The 2nd, ssh-keygen based, mechanism is the preferred one. This is the mechanism used by NexentaStor appliances to communicate between themselves. The latter is required to run storage replication services, to execute in a group mode, to **`switch`** between appliances for the purposes of centralized management. To enable inter-appliance communication, simply use NMC '**`ssh-bind`**' command (see "Note on SSH Binding"). Once the appliances are ssh-bound, all the capabilities mentioned above are enabled automatically, and executed in a secure way. Please see Section "Centralized Management" for more information.

To use IPv4 address based authentication, simply make sure that IP address of your management client machine is recorded on the appliance, via NMC '**`setup appliance authentication`**' command, and select '**`iptable`**' option. Administrative access to the appliance is required to perform this command. Alternatively, to use ssh-keygen generated authentication keys with your management application running on Windows, Linux or any other platform, use the same NMC command '**`setup appliance authentication`**' command, and select option '**`keys`**'.

3. NMC – Management Console

NMC can be used universally to view and configure every single aspect of the appliance: volumes and folders, storage and network services, fault triggers and statistic collectors.
Nexenta Management Console (NMC) is:

- Power shell that delivers appliance's multi-tier storage services, data replication, fault management, and reporting
- Single-point management client - all Nexenta appliances in the network can be administered without re-login. In fact, "switching" between appliances is instantaneous...

With NMC you can:

- ✔ create, grow, and destroy, export and import storage volumes
- ✔ share folders and files via NFS/CIFS/RSYNC/FTP/WebDAV ([24])
- ✔ take snapshots, custom and automated, periodic and on-demand, and implement a user-defined snapshot retention policies
- ✔ automatically and periodically generate storage and service reports
- ✔ run storage services, periodically and on-demand
- ✔ record and non-interactively replay NMC (batch) sessions
- ✔ discover, configure, and create iSCSI sessions, deploy iSCSI-attached LUNs,

  and more.

NMC provides:

- ✔ group execution mode to manage a group of appliances as a single entity.
- ✔ enhanced batch execution mode via record/replay mechanism.
- ✔ high-level context-defined command completion mechanism
- ✔ numerous manual pages (a.k.a. usage guides), and summary helps
- ✔ context-defined usage "cues" and hints
- ✔ interactive dialogs with step-wise guided command execution,

  and more.

This section is a user guide to NMC. A brief introduction with navigation guidelines and the most basic conventions is followed by top-level commands and their subcommands, Tutorial pages[30], numerous usage examples, collection of tips and guidelines, and more.

30. With additional playback demonstrations available on the nexenta.com website, at http://www.nexenta.com/demos

Getting Started with NMC

Nexenta Management Console is pre-configured with administrative super-user account: **root**. The default password is sent to you in email, along with the product Registration Key. Don't forget to change this password. For more information, please see NexentaStor Quick Start Guide at http://www.nexenta.com/docs.

To login, you could use either directly attached console, or ssh as **root** to the appliance's IP address[31].

The fastest way to start with NMC is to press TAB-TAB, type the first a few letters of a selection (a.k.a. command completion - note that the available choices are displayed just below the cursor), and then repeat this routine until the command is complete. At which point you press Enter.

This is similar in a way to completing the directory and executable names in a regular Unix shell.

By definition, a command is complete if either:

- TAB-TAB does not produce any new output

  or

- command is executable[32]

Completing NMC commands via TAB-TAB is one alternative. There is another: type a word or two and press Enter. (For instance, type '**show**' or '**setup**' and press Enter). The selection, as demonstrated by a screenshot below[33], is then done in a menu-style: using arrows and Enter.

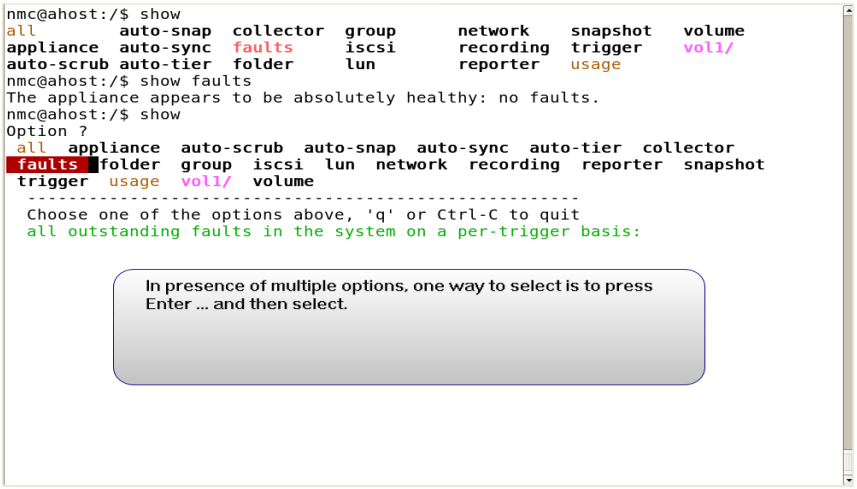

This menu-based way of *completing* NMC commands can be disabled via the 'ask_incomplete' NMC option. See Section "Preferences" for details.

31. or hostname, if propagated via DNS
32. Which is indicated by presence of **<Enter>** – see Section "Navigation" and screenshots below
33. Some of the screenshots throughout this document may be outdated – captured from older versions of NexentaStor. This may be reflected in minor differences with actual appliance screens you may notice..

You can interchangeably use the “TAB-TAB” approach for command completion, type command names or partial command actions to enter a menu driven mode, or add “-h” as necessary to most secondary commands for full usage statements and examples.

Whichever way you use to enter commands, NMC will present a number of (completion) choices. To quickly find out the meaning of all those multiple options, type '?' and press Enter. For instance, type 'show appliance', and press TAB-TAB or Enter:

```
nmc@myhost:/$ show appliance
```

In response NMC will show a number of options - in this particular case appliance's services and facilities that can be "shown". Note that **<?>** is part of the '**show appliance**' completion set - its presence indicates availability of brief per-option summary descriptions.

Next:

- type '?'
- observe brief descriptions
- decide which (completion) option to use
- repeat the sequence, if needed

See also:

Section "Navigation" for further general-usage discussion on how to navigate NMC.

Section "Tutorial" for usage cases and commentary.

Section “Command Reference” for NMC command reference.

Navigation

To navigate NMC, you can interchangeably use the “TAB-TAB” approach for command completion, type command names or partial command actions to enter a menu driven mode, or added “-h” as necessary to most secondary commands for full usage statements and examples.

NMC is a management shell, and as such, it provides all the features typically expected of command shells.

Management console maintains a log (a history) of executed command. The command history is accessible via:

- '**history**' (type '**history**' and press Enter)
- Up and Down arrows at the prompt
- Ctrl-R to perform reverse search

NMC provides I/O redirection, command completion, piping and background execution, and so on. See Section "Tips and Guidelines" for more usage tips and examples.

A Note On NMC Viewer

NMC employs the popular **less** utility as its internal Viewer. To allow convenient paging, scrolling and searching, NMC runs the viewer automatically - if and when when the operation produces more than a single screen of an output.

When in '**less**', you will see a short instruction at the bottom of the screen:

```
Up/Down/PgUp/PgDn/w/z - scroll, /- search, q - quit
```

In addition to 'q' (quit), you could always hit Ctrl-C to terminate the viewing mode (as well as any other NMC operation and dialog). The viewer supports a multitude of options, allows to search NMC operations' result set, save it in a text file, and more. Run 'help less' for more information.

All NMC operations are accompanied with usage information and examples. While entering a command, look for **<-h>** - its presence would mean that the usage text for this command is available. For instance, type 'create volume' and press TAB-TAB.

With NMC you'll never need to do much typing. For instance, in the example with 'create volume', you could first type 'cr' and press TAB-TAB. NMC will complete it to 'create'. Then type 'v' or 'vol', and press TAB-TAB – here's what will show up:

```
nmc@myhost:/$ create volume
show  <Enter> <help|-h> <?>
```

This effectively presents the following completion choices:

show	show existing volumes
<Enter>	enter a dialog, fill out the required parameters and subsequently execute the command (by pressing Enter)
<help|-h>	get the command's manual page, with synopsis, full description of all command-line options that would allow to execute the command in a batch (non-interactive) mode, and usage examples.
<?>	get a few-lines short summary description.

Top-Level Commands

NMC supports 11 top-level commands[34]; the exact list can be obtained at any time by pressing TAB-TAB at a prompt. These are:

show setup

query switch help

destroy create run share unshare record

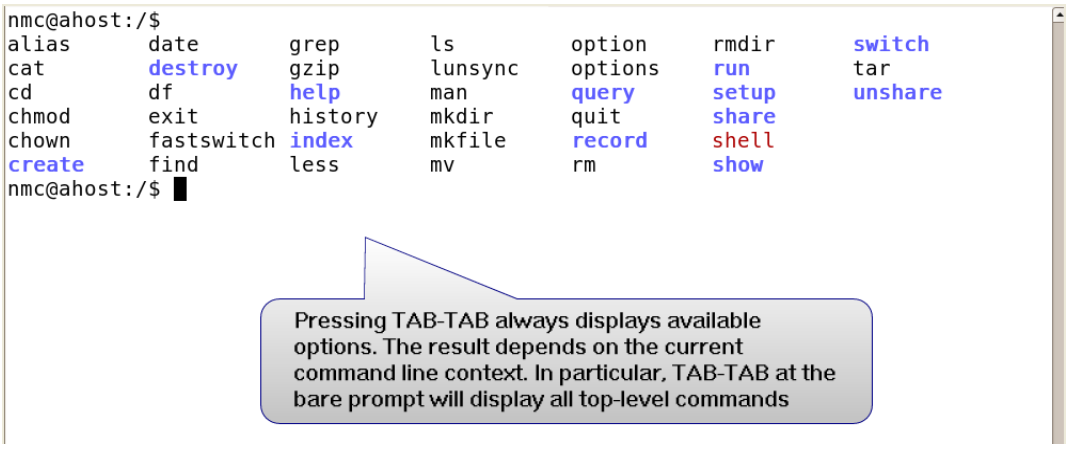

The above example shows a "TAB-TAB" key action at NMC prompt resulting in a NMC commands[35]. Selecting, for instance, **setup** and another "TAB-TAB" sequence will list the available sub-commands. At any point pressing the return key will begin a menu driven command mode that will allow any further options or necessary arguments to be answered interactively. One can always exit an interactive command by pressing the "Ctrl-C" key combination.

For each NMC command, to get a manual page with a complete set of its use cases, simply run **help <command>.** For instance:

```
nmc@myhost:/$ help setup
```

Super Commands

Out of total 11, the first two top-level commands are somewhat notable: the majority of administrative management actions can be performed using **setup** and **show**.

The appliance is sub-divided into components – managed services and objects – which in turn consist of sub-components, and so on. The taxonomy, that is, the entire management hierarchy reveals itself via **setup** and **show** "super commands".

As an example, type **setup** and press TAB-TAB or Enter. The menu of choices will include **'appliance', 'network', 'group',** etc. From this point on you can traverse the tree-like appliance's management hierarchy. For a context-sensitive summary, type or select '?' at any step of the way. For instance, '**setup ?**' will display:

34. At the time of this writing..
35. Some of the screenshots throughout this document may be outdated, which may be reflected in minor differences you may notice. For the latest updated documentation please refer to Nexenta Systems, Inc. website

```
* Configure Nexenta Storage Appliance, deploy storage services,
  create, configure and grow storage volumes, manage groups of appliances,
  iSCSI, network, statistic collectors, etc. - all aspects of the system.

  appliance       Appliance configuration: general settings
  auto-scrub      Configure and deploy auto-scrub storage services
  auto-snap       Configure and deploy auto-snap storage services
  auto-sync       Configure and deploy auto-sync storage services
  auto-tier       Configure and deploy auto-tier storage services
  collector       Enable, disable, and run appliance collectors
  folder          Create, destroy, share, unshare, snapshot, and
                  configure folders
  group           Group appliances to execute commands and recorded
                  sessions on each appliance in the group.
                  For background execution use '&' at the end of the
                  commandline.
  iscsi           Setup iSCSI configuration
  network         Administer network interfaces and services
  recording       Record and re-play NMC sessions
  reporter        Enable, disable, and run appliance reporters
  snapshot        Create and destroy snapshots, rollback folders
                  to previous snapshots
  trigger         Enable, disable, run, show fault triggers
  usage           Usage Guide: full collection of usage texts
                  for all specific 'setup' commands in the system
  volume          Administer volumes: create, destroy, import, export,
                  add and replace LUNs, set properties, snapshot
nmc@ahost:/$
```

NMC is a CLI, and as such it provides numerous commands. The important, however, is that you never need to remember those commands. To configure appliance, simply enter **setup**, and then keep pressing TAB-TAB and making selection. See Section "Navigation" (or type '**help navigation**') for details.

Command Reference

show	display any given object, setting or status
setup	create or destroy any given object; modify any given setting
query	advanced query and selection
switch	manage another Nexenta Appliance or a group of appliances
destroy	destroy any given object: volume, folder, snapshot, storage service, etc.
create	create any given object: volume, folder, snapshot, storage service, etc.
run	execute any given runnable object, including storage services: auto-snap auto-scrub auto-sync auto-tier
share	share (via NFS, CIFS, RSYNC, FTP and WebDAV [24]) a volume or a folder. Share zvol (Section "Terminology") via iSCSI.
unshare	unshare a volume or a folder
record	start and stop NMC recording sessions
help	NexentaStor manual pages

As a general rule of thumb, verbs such as **show**, **create**, **destroy**, etc. take as arguments the service, dataset, or equivalent to act upon. Detailed NMC command reference follows below.

show Display any given object, setting or status.

In particular:

- show “all” appliance at a glance: appliance's network and storage services, interfaces, shares, volumes, and hardware
- display system information, versioning, memory, mailer settings, storage services, network collectors, fault triggers
- show Nexenta Management Console (NMC) - settings and logs
- show Nexenta Management Server (NMS) - settings and logs hostname and domainname, UUID, saved configurations
- show storage volumes, folders/filesystems, and snapshots
- show selected folders, their properties, shares, and snapshots
- show snapshots: verbose/non-verbose, recursive/non-recursive
- show **auto-scrub** storage services
- show **auto-snap** storage services
- show **auto-sync** storage services
- show **auto-tier** storage services
- show **auto-cdp** storage services
- display all outstanding faults: count, severity, description, timestamps
- display groups of appliances
- show iSCSI initiators and iSCSI configuration
- show LUNs[36]: directly and remotely attached, removable media, IO statistics, general information
- show network interfaces, routes, and services
- show reporters: daily, weekly and custom-defined
- show fault triggers
- show statistic collectors
- show version

For complete collection of '**show**' sub-commands and man pages, run:

```
nmc@myhost:/$ show usage
```

36. See http://en.wikipedia.org/wiki/Logical_Unit_Number. Often, the terms “LUN”, “hard drive” and “disk” are used

setup Configure, create or destroy any given object; modify any given setting. In particular:

- setup appliance configuration (general settings)
- configure and deploy **auto-scrub** storage services
- configure and deploy **auto-snap** storage services
- configure and deploy **auto-sync** storage services
- configure and deploy **auto-tier** storage services
- configure and deploy **auto-cdp** storage services
- administer volumes and **zvols**: create, destroy, import, export; add, replace, attach and detach LUNs, set properties, snapshot
- create, destroy, share, unshare, snapshot, and configure folders
- create and destroy snapshots, rollback folders to previous snapshots
- setup iSCSI configuration, discover iSCSI targetrs, create and destroy iSCSI sessions
- setup group appliances to execute commands and recorded sessions on each appliance in the group
- administer network interfaces and services
- record and re-play NMC sessions
- configure, enable, disable, run, show fault triggers
- configure, enable, disable, run and show appliance collectors
- configure, enable, disable, run, and show appliance reporters

Run '**setup**' to set a property, enable, disable, share, unshare, import, export, create, destroy, run a service, grow a volume – generally, to perform any configuration management operation.

Complete collection of '**setup**' sub-commands and man pages is available via:

```
nmc@myhost:/$ setup usage
```

interchangeably.

query Advanced query and selection.

To quickly locate and display appliance's objects using arbitrary complex search criteria, NMC provides a powerful '**query**' command.

The syntax of the command is '**query -e <expression>**', where **expression** is any logical formula with names of objects' properties used as expression's variables, e.g.:

"compression eq off" - to locate folders (for which compression is switched off),

"quota < 1M" - to query data sets limited to 1 MB,

"type eq hourly" - to locate all hourly auto-services (including **auto-sync, auto-scrub, auto-tier**, and **auto-snap**),

Use regular expression syntax (canonical symbols) such as:

eq, ne	**for string comparison,**
<, >, ==, <=, >=	**for number comparison,**
=~	**for regular expressions,**
and, or, not	**to combine all those elements into one compound logical formula**

NMC supports an extension of the expression syntax that allows to disregard property values and query based on the presence of properties with matching names. For instance,

```
nmc@myhost:/$ query -e size and type
```

will locate all objects having both 'size' and 'type' properties.

Example: check health of all volumes in the system:

```
nmc@myhost:/$ query -e health
```

or, more specifically:

```
nmc@myhost:/$ query volume -e health
```

The result may look like:

```
NAME PROPERTY
vol1 health=ONLINE
vol4 health=DEGRADED
...
```

Run '**query -h**' or '**query usage**' for an overview, usage, and examples.

switch Manage another Nexenta Storage Appliance, or a group of appliances.

Nexenta Management Console supports two different switch-to-remote-appliance mechanisms:

1. regular (non-fast), and
2. fast switch, or '**fastswitch**'.

Once "switched" to the remote appliance, you can inspect its storage and networks settings, setup services, generate reports, etc. There is no need to run separate ssh sessions and/or open separate management terminals - you can switch back and forth between appliances by simply running '**switch appliance <hostname>**'.

Managing a remote appliance via non-fast version of switching requires previously defined ssh-binding (see "Note on SSH Binding"). In that sense the non-fast **switch** is similar to **auto-tier** and **auto-sync** services

The "fast" switching is built upon the ability of management consoles to communicate with each other via remote procedure call transport. The '**fastswitch**' alias (see Sections "", "NMC shortcuts and aliases") could be used as a shortcut instead of '**switch appliance -f**'. The '**fastswitch**' does not require appliances on the network to be ssh-bound. The added benefit of -f '**fastswitch**' option, is that, as the name implies, the "switching" itself from an appliance to another appliance is fast. In fact, it is instantaneous.

The '**switch**' command has two variations:

- '**switch group**' – switch to a group, or in other words, enter a "group mode" to perform operations on a given group of appliances
- '**switch appliance**' – switch to a remote appliance. All NMC operations including interactive ones, will be communicated to a remote appliance, and the results will be transferred back.

See also: '**fastswitch**' below, in Section "".

See also:

- Section "Centralized Management" - for general background and introduction
- Section "Tutorial" - for usage examples and commentary
- Section "Management GUI – Quick Introduction" - for similar NMV functionality

destroy

destroy any given object: volume, folder, snapshot, storage service, etc.

For the complete collection of 'destroy' man pages, run:

```
nmc@myhost:/$ destroy usage
```

create

Create any given object: volume, folder, snapshot, storage service, group of appliances, network interface.

See also: Section “”

For a quick list of objects that can be created, type '**create**' and press TAB-TAB. For the full list of '**create**' man pages, run:

```
nmc@myhost:/$ create usage
```

This will list all '**create**' subcommands, with their corresponding manual pages.

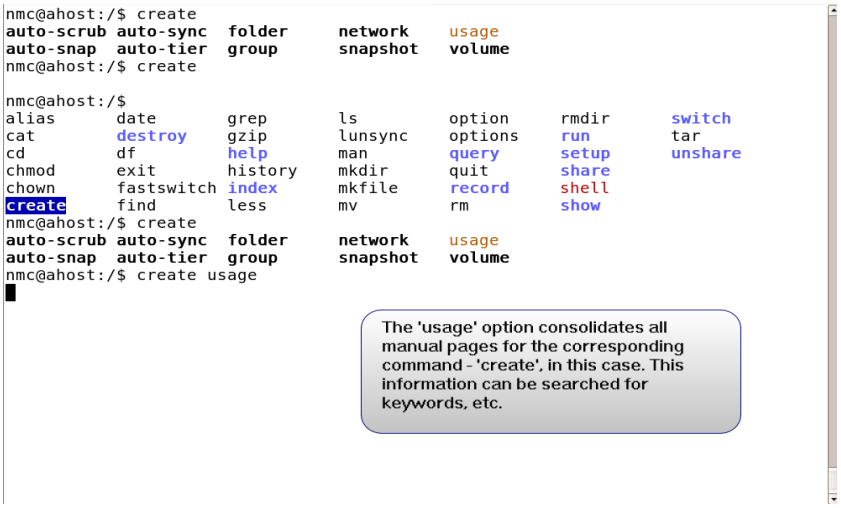

run	Execute any given *runnable* object, including storage services: **auto-snap, auto-scrub, auto-sync, auto-tier**
	and NexentaStor runners:
	fault triggers, statistic collectors, reporters
	The command facilitates on-demand execution.
	For an overview and usage of any one of the listed services, run
	`nmc@myhost:/$ create <service-name> -h`
	Complete collection of '**run**' manual pages is available via:
	`nmc@myhost:/$ run usage`
	See also: Section "Tutorial".

share

Share a volume or a folder via NFS, CIFS, RSYNC, or FTP.

Share zvol (see Section "Terminology") via iSCSI.

Makes local volume or (folders) pathname available for mounting by remote systems.

Makes local zvol available for iSCSI Initiators for using it as iSCSI-attached block device.

For NFS, the command supports '**rw**', '**ro**', and '**root**' options – each option contain zero or more colon-delimited netgroups and/or hosts.

Example – **share** '**vol1/a/b**' folder via NFS:

```
nmc@myhost:/$ share vol1/a/b nfs
rw              : group-engineering:10.16.16.92
ro              : group-marketing
root            : admin
extra-options:
```

Any existing share(s) can be listed, for instance:

```
nmc@myhost:/$ share vol1/a/b nfs show
```

See also: Section "Tutorial"

Or run:

```
nmc@myhost:/$ share usage
```

This will list all '**share**' subcommands and corresponding manual pages.

unshare

Unshare volume, folder, or zvol.

For more information, run:

```
nmc@myhost:/$ unshare usage
```

This will list all '**unshare**' manual pages.

record Record NMC session. Make the "recording", start, stop, and view NMC recording sessions.

With NMC, management operations often require specification of multiple options. A combination of NMC recording facility and NMC multiple-selection and other dialogs gives the best of both:

- interactive (guided) filling-out of the required parameters
- ability to execute the "recorded" command many times, on many different Nexenta appliances.

Usage: <name> | stop | view-current

Once recorded, the NMC session can be re-played any number of times.

The recording facility enhances capability of the Nexenta Management Console to execute batch (non-interactive) jobs. While recording is active, all commands are stored in their "final" form, even though you may be entering some (or all) of the options interactively.

For instance, run:

- **record abc**
- **create auto-tier**

At this point NMC will prompt you to start filling out all required parameters. Once completed, run:

- **record view-current**

You will see a complete command populated, with all the right options. This command, and the entire recorded session, can now run without user interaction.

To start a new session, simply state the (new) name in the command line, for instance: '**record new-batch**'. To stop recording, simply type '**record stop**'.

While the recording is active all commands are stored with the session. An active recording is denoted by **(r)** in the prompt, for instance:

```
nmc@myhost:/$ record aaa

Started recording, session name: 'aaa'. To stop
recording, type 'record stop'.

(r)nmc@myhost:/$
```

A recording session may be closed (stopped) and re-opened multiple times. Each next time new records are appended to the session file.

4. Tutorial

This Tutorial demonstrates only a small part of appliance's functionality. Please use NMC and NMV to discover its full power and flexibility.

Additional tutorial is available on the website, in a form of playback movies, at http://www.nexenta.com/tutorials

Here are the specific URLs:

Initial Configuration Wizard will guide you through the first steps of setting up the appliance. Just point your Internet browser to the Wizard's URL (see Quick Start Guide for details on that), and follow online instructions.

- Initial Configuration Wizard

Nexenta Management View (NMV) provides appliance's web management interface. NMV combines intuitive graphical interface with ease of use, ability to monitor all appliances on the network with ability to easily execute the most essential administrative operations.

- Nexenta Management View

Nexenta Management Console (NMC) can be used universally to view and configure every single aspect of the appliance: volumes and folders, storage and network services, any given appliance on the network and groups of appliances. The following 5 tutorials introduce NMC and demonstrate very basic operations. Please use NMC to discover its full power!

- Nexenta Management Console: Introduction
- Nexenta Management Console: Volumes and Folders
- Nexenta Management Console: Automated Snapshots
- Nexenta Management Console: Replication (Tiering) Basics
- Nexenta Management Console: Continuous Data Protection (remote mirroring)

The first 10 sub-sections of this section demonstrate how to:

1. Create a volume
2. Create folders
3. Share a folder
4. Take periodic snapshots
5. Set up a tiering service (a) over NFS and (b) over SSH

6, 7, 8. Upgrade appliance software

9. Setup **`auto-cdp`** service (see "Terminology")
10. Re-register the appliance with a new license key.

All operations are accompanied by many examples, are screenshot-illustrated and heavily commented, with comments being generally applicable to all the rest NMC operations..

Creating a volume

NexentaStor utilizes the ZFS filesystem for flexible, unlimited storage growth. Volume management has been made simple, with only a basic knowledge of performance versus redundancy useful to best utilize your raw storage. A good way to start is to attempt "dry-run" volume creation. NMC dry-run is one of many supported options that one can find, along with examples, but trying out the "**`create volume -h`**" command. In almost all cases, you will want to set up either mirrored volume pool or parity based RAID pool[37].

The following NMV screenshot demonstrates functionality similar to NMC's '**`create volume`**':

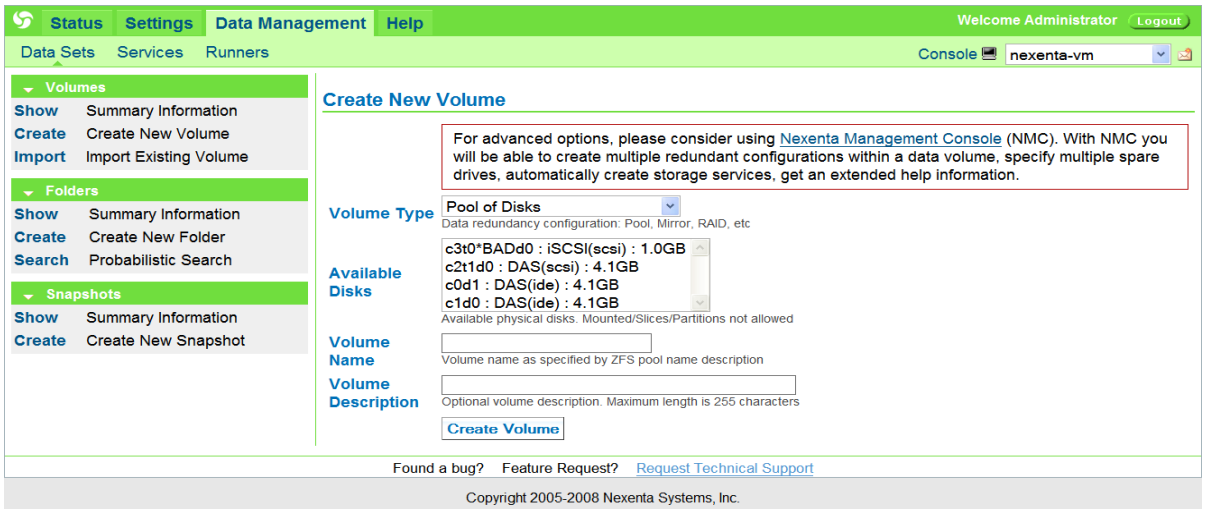

Notice an "advanced option" comment on top of this picture.

37. For details and discussion regarding ZFS setup please refer to http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide which serves as an excellent source of in-depth information for both beginners and storage experts.

A Note On Redundant Configurations

A mirrored volume pool consists of matched drives or drive groups, where by data always has a redundant copy on the mirrored set of disks. Mirroring can make use of other pooled technologies such as parity, allowing multiple groups of disks to be setup each with one primary array and one secondary, mirrored array. In most cases, for best reliability and performances, administrators would setup a combined or striped set of mirrored devices (sometimes referred to as **RAID 10**). In the case of two-way mirrors RAID 10 will halve your overall storage capacity, but will provide the best read/write performance, as reads are striped across all of the primary disks, and writes only require a single duplication of each write to a secondary drive. At any time, any number of failed drives are permitted, as long as no two drives in a paired set fail at the same time.

Parity based RAID volumes make use of one or two dedicated drives to maximize capacity without reducing redundancy of stored data. Each write is committed across all drives in a group, including the parity devices, and they further take some penalty in calculating the parity. The reverse is equally true, as reads must combine data and parity across all devices in a group. To improve performance, it is generally recommended to also stripe multiple parity based RAID groups together to allow parallel reads/writes to the disk. This is commonly referred to as RAID 50. Up to one drive in a RAIDZ1 group, or two drives in a RAIDZ2 group can fail at a time without losing data. In the RAID50 setup, you both allow for future expansion with new parity groups, as well as allow for more drive failures, limited still to at most two per group.

In both mirrored and parity based RAID volumes, you should establish multiple spare devices equal to the size of each member drive. Redundant, striped arrays of either variety, with sufficient spare disks, allow one to achieve the greatest level of reliability on commodity disks. As disk capacity grows and gets ever cheaper, you can expand on these striped volumes. The ZFS based filesystem allows for continuous volume growth, but consistent disk group sizing across a striped array is recommended. Therefore, as disk sizes increase, it is considered a good practice to create disk sub-groups of as close to an equal size as possible.

Redundant configurations improve not only reliability of your NexentaStor system but performance as well. For mirrored configurations:

- Random reads scale linearly with the number of disks; writes scale linearly with the number of mirror sets.
- Read throughput scales linearly with the number of disks; write throughput scales linearly with the number of mirror sets.

For parity (RAID-Z, RAID-Z2) configurations:

- Random I/O reads and writes scale linearly with the number of RAID sets.
- Sequential I/O throughput scales linearly with the number of data (non-parity) disks.

All of the above may sound a bit complicated but in fact it is not.. NexentaStor hides the complexity and makes the volume configuration management simple and intuitive. So, read on.

To create a new Nexenta Appliance volume, you need spare drives a. k. a. LUNs ("Terminology", see also footnote on page #35 of this document).

To see all appliance's disks:

```
nmc@myhost:/$ show lun
```

As always, detailed manual page for this command (as well as all the rest NMC commands) is available via -h (help) option:

```
nmc@myhost:/$ show lun -h
```

If there are no spare LUNs, you could still go through the steps below, but a volume obviously won't be created..

For more information on LUN management, see Section "Managing LUNs".

The corresponding NMV function is located under Settings/Disks, example:

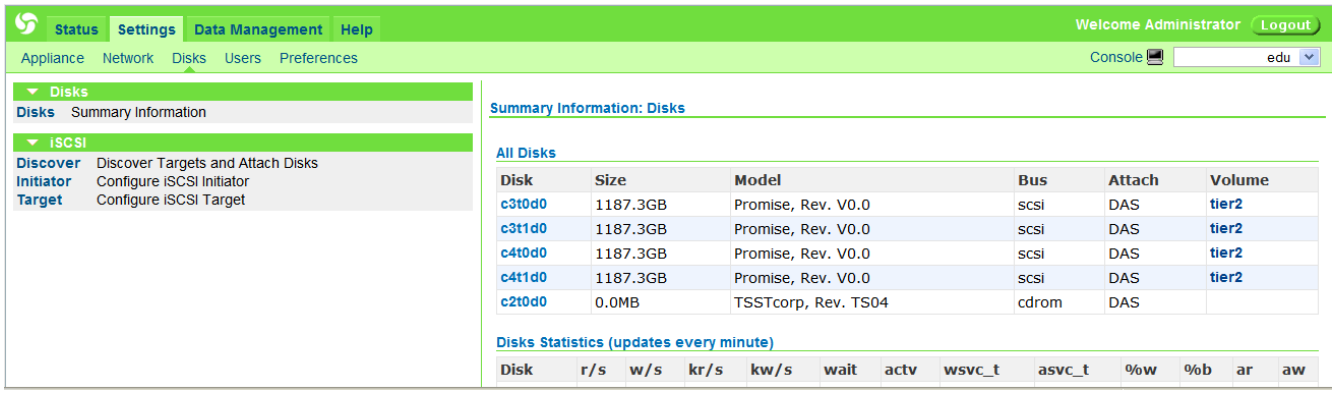

To create a volume, run:

```
nmc@myhost:/$ create volume
```

or, the same:

```
nmc@myhost:/$ setup volume create
```

and press TAB-TAB. This will display a list of completion options explained in detail in Section "Navigation" ('**help navigation**'). Notice the **<-h>** and **<help>** options – their presence indicates the availability of extended usage information. Thus, for the man page run:

```
nmc@myhost:/$ create volume -h
```

One of the supported options allows to try (or "dry-run") volume creation without actually creating the volume and making any changes. You could try this option to play with various redundancy configurations: mirror, raid, etc.

```
  -n            Dry-run. Display the resulting configuration without
                actually performing the operation.

        Please consider your data redundancy and IO performance
        requirements carefully. For instance, a simple mirror
        configuration provides both the benefits of data redundancy
        and read performance. If you need even better data protection,
        3-way has significantly greater MTTDL (Mean Time To Data Loss)
        than a 2-way mirror, while going to a 4-way mirror offers only
        marginal improvement. More guidelines can be found at
        www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide

        * A raidz1 configuration maximizes disk space and generally
          performs well when data is written and read in large chunks
nmc@ahost:/$ create volume -n
Volume                : vdemo
Group of devices      : c2t0d0, c2t1d0
Group redundancy type : mirror

would create 'vdemo' with the following layout:

        vdemo
          mirror
            c2t0d0
            c2t1d0
nmc@ahost:/$ show volume all
no volumes available
nmc@ahost:/$
```

Run:

nmc@myhost:/$ create volume

and press Enter. NMC will guide you through the interactive process of creating a new volume.

```
nmc@ahost:/$ create volume
Volume                : vdemo
Group of devices      : (Use SPACEBAR for multiple selection)
 c1d1  c2t0d0  c2t1d0
  ------------------------------------------------------------------------------
  Select one or more LUNs to form a new group of devices in the volume vdemo.
  (Note that next you'll be able to specify type of the device configuration:
  mirror, raidz, etc.). Navigate with arrow keys (or hjkl) or Ctrl-C to exit.
  lun id: c2t0d0, device: sd0, size: 2047MB, SCSI
```

At any step of the way you can break without making any changes by pressing Ctrl-C. Note helpful hints at the bottom of all dialogs. To create a volume non-interactively, supply the volume name and configuration in the command line, e.g.:

nmc@myhost:/$ setup volume create vol1 mirror c2t0d0 c2t1d0 -y

This will create a two-way mirror 'vol1' based on two SCSI drives. If unsatisfied, you can destroy the newly created volume and repeat the steps above. To destroy the volume, run:

nmc@myhost:/$ destroy volume

or even, skipping the confirmation dialogs:

nmc@myhost:/$ destroy volume vol1 -y

Creating a folder

This and the following subsections assumes that we have a volume 'vol1' in the system.

```
nmc@myhost:/$ create folder vol1/a/b
```

This will create a folder vol1/a and its subfolder vol1/a/b - in one shot. See '**create folder -h**' for detailed discussion and examples.

```
chmod       exit        history     mkdir       quit        share
chown       fastswitch  index       mkfile      record      shell
create      find        less        mv          rm          show
nmc@ahost:/$ destroy folder vdemo/tmp
Destroy folder 'vdemo/tmp' and its subfolders ?  Yes
folder 'vdemo/tmp' destroyed.
nmc@ahost:/$ setup volume vdemo
clear-errors export       replace-lun  show         folder      snapshot
destroy      grow         share        unshare      property    <?>
nmc@ahost:/$ setup volume vdemo
clear-errors export       replace-lun  show         folder      snapshot
destroy      grow         share        unshare      property    <?>
nmc@ahost:/$ setup volume vdemo folder
create show   <?>
nmc@ahost:/$ setup volume vdemo folder
create show   <?>
nmc@ahost:/$ setup volume vdemo folder create
<Enter>   <help|-h> <?>
nmc@ahost:/$ setup volume vdemo folder create tmp
NAME             TYPE   USED  AVAIL  MOUNTED  QUOTA  COMPRESS
vdemo/tmp  filesystem    18K  1.95G      yes   none       off

nmc@ahost:/$ create folder vdemo/other-folder/subfolder
NAME                              TYPE   USED  AVAIL  MOUNTED  QUOTA  COMPRESS
vdemo/other-folder             filesystem    36K  1.95G      yes   none       off
vdemo/other-folder/subfolder   filesystem    18K  1.95G      yes   none       off

nmc@ahost:/$
```

In NMV, similar functionality is provided via Data Management tab: click on **Data Sets** under this tab, and then **Create** under Folders:

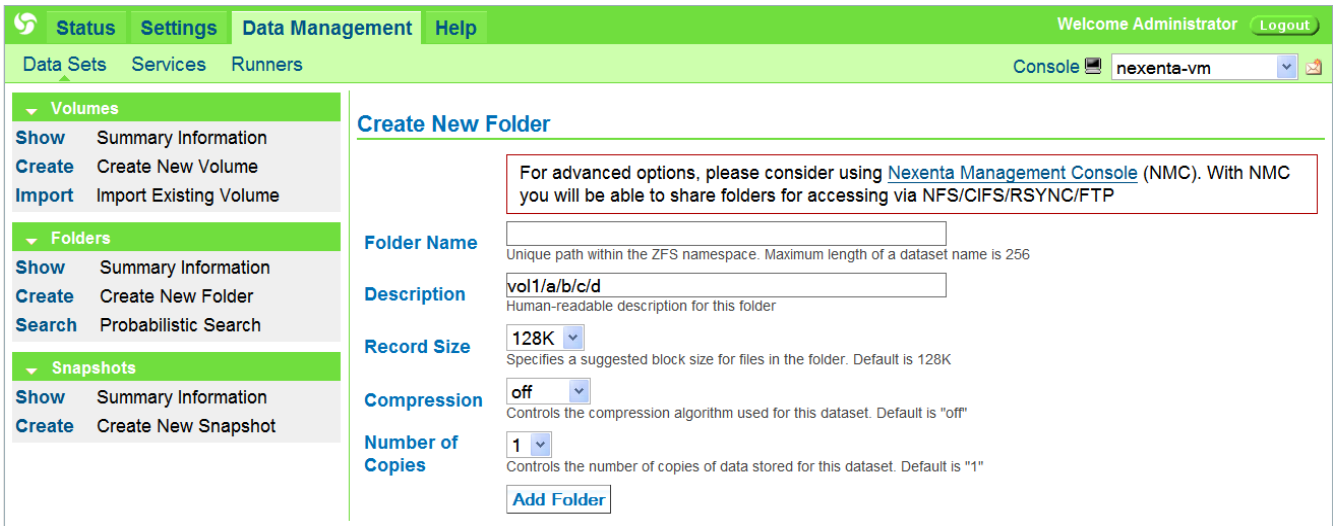

Similar to NMC, here you can create several nested folders in one shot.

Obviously, the volume and folder names are given just as an example, here and throughout the rest of this section.

Use:

```
nmc@myhost:/$ setup volume vol1 folder a

nmc@myhost:/$ show volume vol1 folder a/b
```

to view and change folder properties.

The following screenshot illustrates folder-property-update operation[38]:

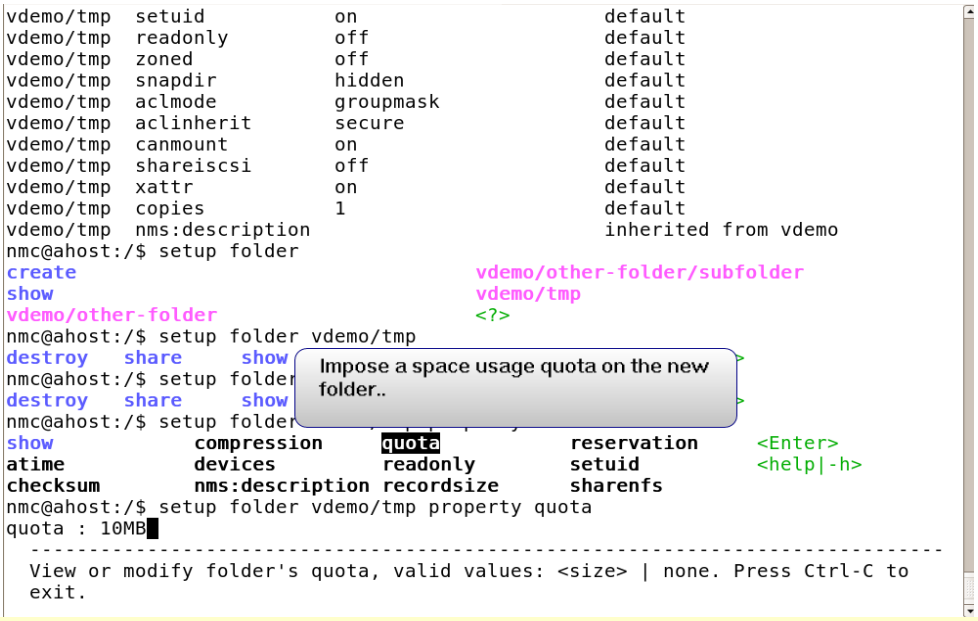

Most NMC operations that deal with folders and volumes can be typed in a shorter form, by omitting 'volume' and 'folder' keywords.

For instance:

```
nmc@myhost:/$ setup volume vol1 folder a
```

is equivalent to:

```
nmc@myhost:/$ setup vol1/a
```

and:

```
nmc@myhost:/$ show volume vol1 folder a snapshot
```

is the same as:

```
nmc@myhost:/$ show vol1/a snapshot
```

and the latter in turn is the same as:

```
nmc@myhost:/$ show snapshot vol1/a
```

For more information, please refer to Section "Tips and Guidelines".

38. For similar NMV functionality, please see Section "NMV Configuration and Fault Management"

Sharing a folder via NFS

Any network attached file server on the market today is expected to service both UNIX and Windows clients with the NFS and CIFS protocols respectively. Where they generally differ is in their levels of robustness, scalability, and ability to join the various authentication domains available for the two protocols.

NexentaStor provides one of the best existing NFS stacks. NFSv4 is used by default; optionally you could change the default (to NFSv3 for instance) via NMS configuration. Its multi-threaded design allows for great scalability, with users of both v3 and v4 able to access file system snapshots through the protocol.

First, make sure NFS is up and running:

```
nmc@myhost:/$ show network service nfs-server
```

Notice that '**`show network service`**' and '**`setup network service`**' can be used to display and configure network services, as the names imply..

The corresponding NMV function is located under Settings/Network. For instance, to configure NFS, go to Settings => Network => NFS server and click on **Configure**. A sample screenshot follows:

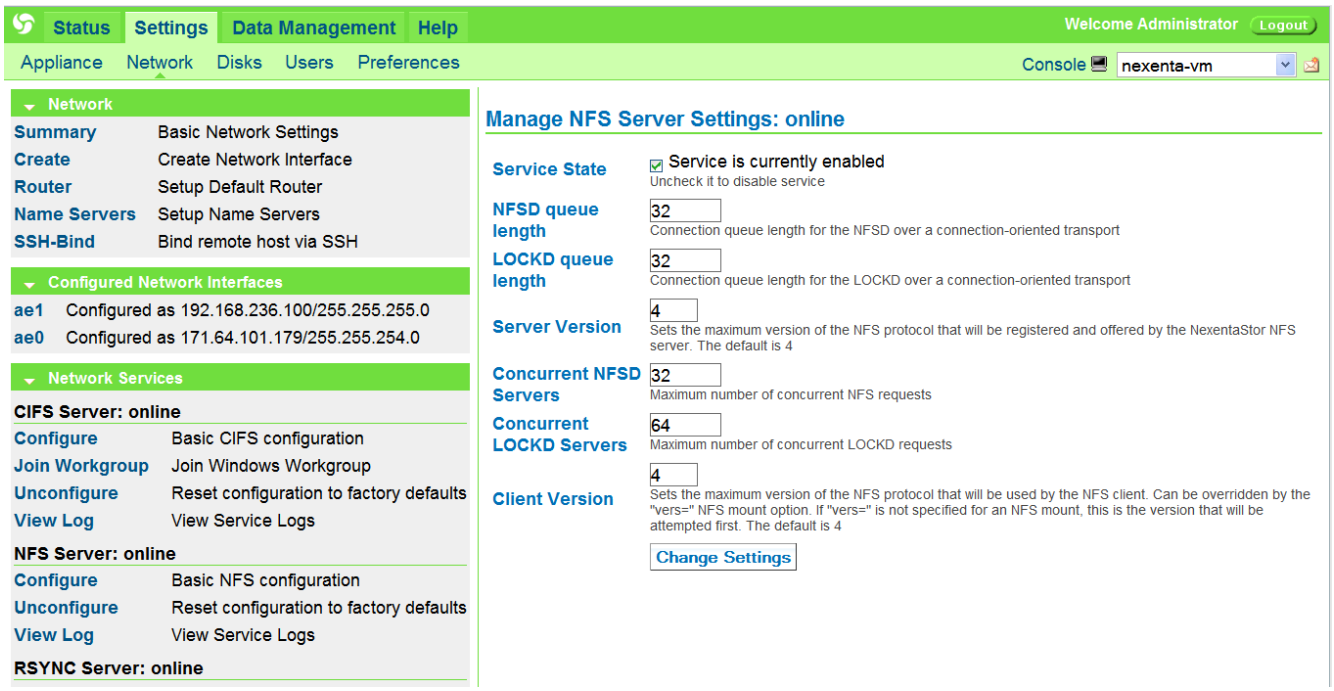

In NMC, type:

```
nmc@myhost:/$ share
```

and press TAB-TAB or Enter. The completion choices will include '**folder**', '**usage**', '**volume**', and '**vol1/**' - that is, assuming '**vol1**' is the only volume in the system.

Let's now share the entire volume 'vol1':

```
nmc@myhost:/$ setup volume vol1 share nfs -o
                rw=primary,root=primary
```

Alternatively, since '**share**' is one of the top-level commands (see Section 'Top-Level Commands' or run '**help commands**'), you could use a shorter variant:

```
nmc@myhost:/$ share vol1/ nfs -o rw=primary,root=primary
```

Here we are also using the fact that '**volume**' and '**folder**' keywords can be omitted, as explained in the previous subsection.

This will share the entire '**vol1**', with all its subfolders. More exactly, by default the subfolders inherit the "sharing" property from their parent. It is possibly to override this default behavior - see '**setup volume vol1 share nfs -h**'

So, the above would be mounted as **myhost://volumes/vol1** if our hostname was **myhost**.

As always with NMC, you could create a share interactively, by running:

```
nmc@myhost:/$ share vol1/ nfs
```

NMC will then prompt you to specify read-write access and other share qualifiers, and will guide you through the process.

It is important to note that all shares for NFS are relative to **/volumes**, or whatever the '**volroot**' property is set to.

Run '**show volume <name> share nfs**' to show the actual mountpoint of any given shared folder.

To see the current settings:

- run '**show appliance nms config**',
- and/or refer to Section "Preferences",
- and/or run '**help options**'.

To see all shares in the appliance, simply type 'show share', for instance:

```
nmc@ahost:/$ show share
FOLDER                    CIFS   NFS    RSYNC FTP    SAMBA
vdemo/a/b/c               -      Yes    -     -      -
vdemo/new                 -      Yes    -     -      -
```

Here's another 'show share' example, that will display NFS shares a verbose form:

```
nmc@myhost:/$ show vol1/ share nfs -v
```

This will display all 'vol1' shares recursively, including subfolders. See the corresponding man page for details:

```
nmc@myhost:/$ show vol1/ share nfs -h
```

All NexentaStor shares for NFS are relative to **/volumes**, or whatever the NMS '**volroot**' property is set to.

NexentaStor uses NFSv4 by default. To force the server to use NFSv3, set NFS_SERVER_MAX variable to 3 in the server configuration file.

Edit the server configuration via:

nmc@myhost:/$ setup network service nfs-server edit-settings

This is particularly useful if *only* NFSv3 clients are used with the NexentaStor NFS server.

However, if you have a mix of v3 and v4 clients, it would then make sense not to change the default on NexentaStor side. Instead, to force NFSv4 clients to use v3, mount using the following syntax:

```
server:/share /localmount bg,intr,vers=3,noacl
```

NFS shares can be easily created and administered using NexentaStor's GUI: NMV. Go to Data Management => Data Sets => Folders, and click on **Show**. The corresponding summary screen

will list all existing folders in the appliance, and allow you to share the folders via variety of access methods. For instance:

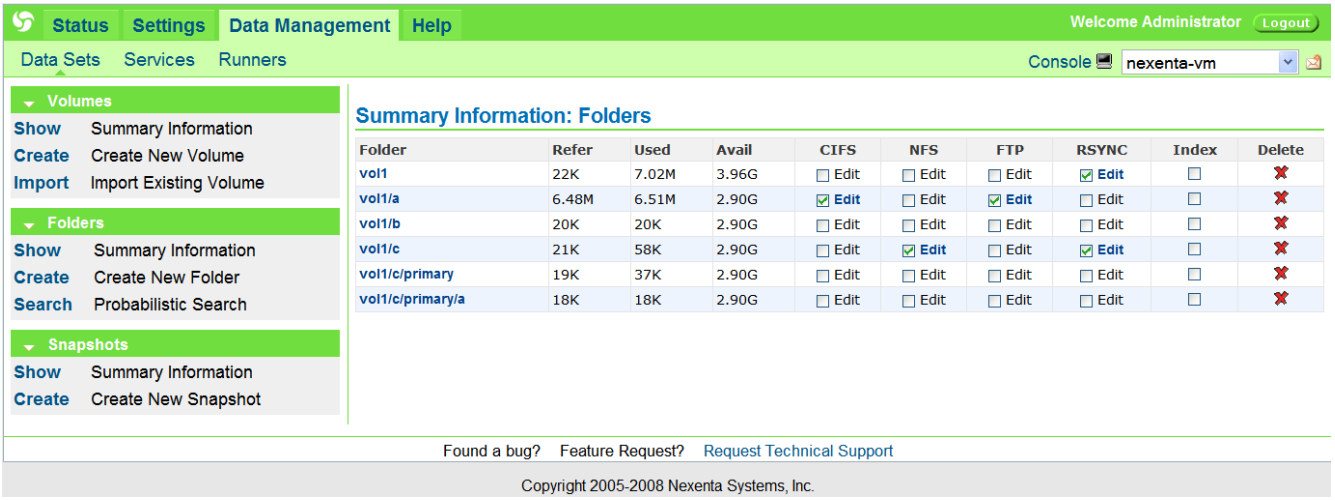

To share a folder, simply check the corresponding checkbox. In this example, we are NFS-sharing folder '**vol1/c**'. The operation is recursive – it'll share the folder and its sub-folders.

Note also that anonymous read/write access is enabled by default. The anonymous username is: **nfs** (user id = 60001).

The anonymous NFS username is: **nfs** (user id = 60001). For anonymous access, the client must have the same user ID defined. NexentaStor client will have this ID defined.

To view or change the default settings, click on the **Edit** link to the right of the corresponding checkbox. The following screen will show up:

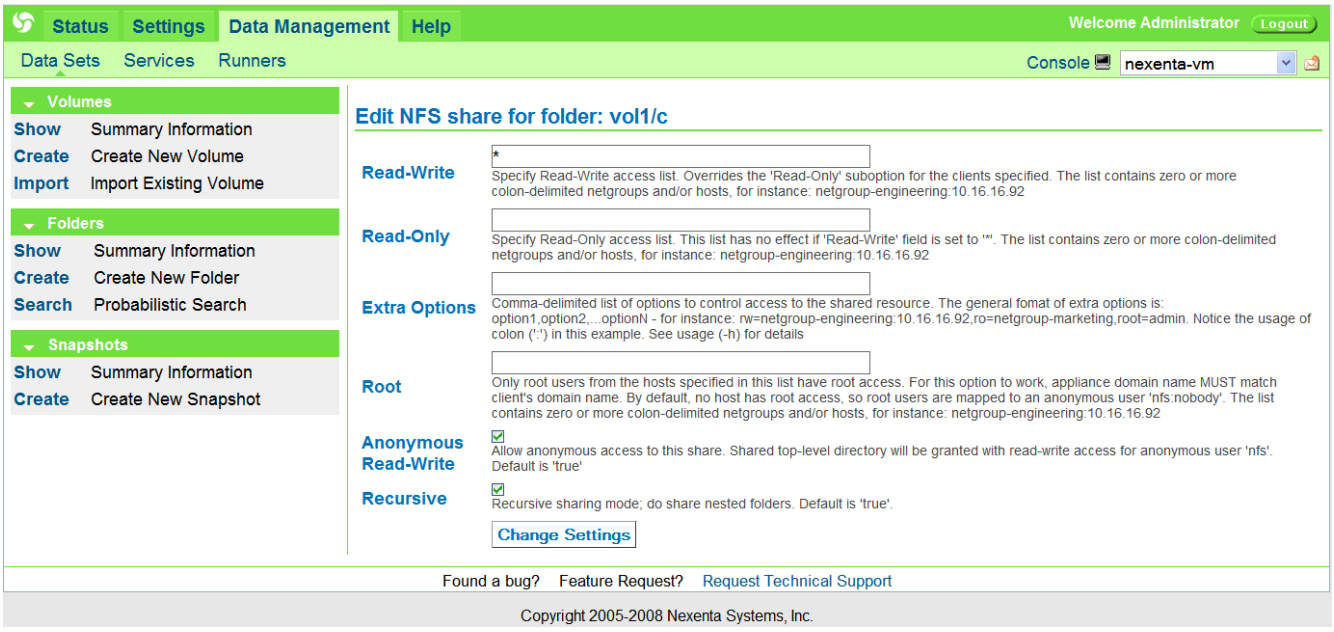

Here you can change the defaults.

For more information, please refer to Section "Management GUI – Quick Introduction".

Unless you are an experienced NFS user or system administrator, here are a few important guidelines:

1. Make sure that domain name on the client and the server are the same. This is particularly critical for NFSv4 only environments. Different domain names is one common reason for getting 'Permission denied' when accessing a shared location from NFS client. Run '**`show appliance domainname`**' to find out.
2. If the client and server domain names are not matching (and cannot be changed to match), try using '**`anon=0`**' setting via Extra-Options.
3. Do not leave the '**`Root`**' field empty. The safest option is to explicitly set Root to client's hostname. Note: specify not the client's IP address but the actual hostname in a form <host.domain> that can be resolved. Use colon (':') delimiter to specify more that one client host and/or netgroup.
4. Otherwise, if you want to open root permissions for all NFSv4 and NFSv3 clients without explicitly specifying the client hostnames in the '**`Root`**' field, use '**`anon=0`**' setting via Extra-Options.

NexentaStor and MAC OS X

Here's yet another example that demonstrates mounting NexentaStor shares from Mac OS X version 10.5 ("Leopard"). The following assumes that **`nstor.mydomain.com`** is the appliance's hostname, and **`machost.mydomain.com`** is 10.5 client. The example includes NMC output and commentary.

```
nstor:/$ share tier2/test nfs
    Read-Write           : machost.mydomain.com
    Read-Only            :
    Root                 : machost.mydomain.com
    Extra Options        : root
    Anonymous Read-Write : false
    Recursive            : false
    Modifed NFS share for folder 'tier2/test'
```

On machost, the 10.5 client:

```
machost:/$ sudo vi /etc/auto_home
```

(adding the following line to **/etc/home**)

```
test    nstor.mydomain.com:/volumes/tier2/test
```

(note that the above left is makes it **/home/test** and the above right is the source volume.)

```
machost:/$ sudo automount -vc
   automount: /net updated
   automount: /home updated
   automount: no unmounts
```

Sharing a folder via CIFS

NexentaStor provides one of the best existing kernel and ZFS integrated CIFS stacks. CIFS support in NexentaStor is a new addition. The underlying system allows for defining flexible share points ('**`setup network service cifs`**'). NexentaStor management UI provides a similar functionality. This section includes both NMC and NMV based examples to illustrate NexentaStor/CIFS operation.

The first step is to configure CIFS server. Or, at least, review the built-in defaults.

NMV provides Network page under its Settings tab, to configure all network services, including CIFS:

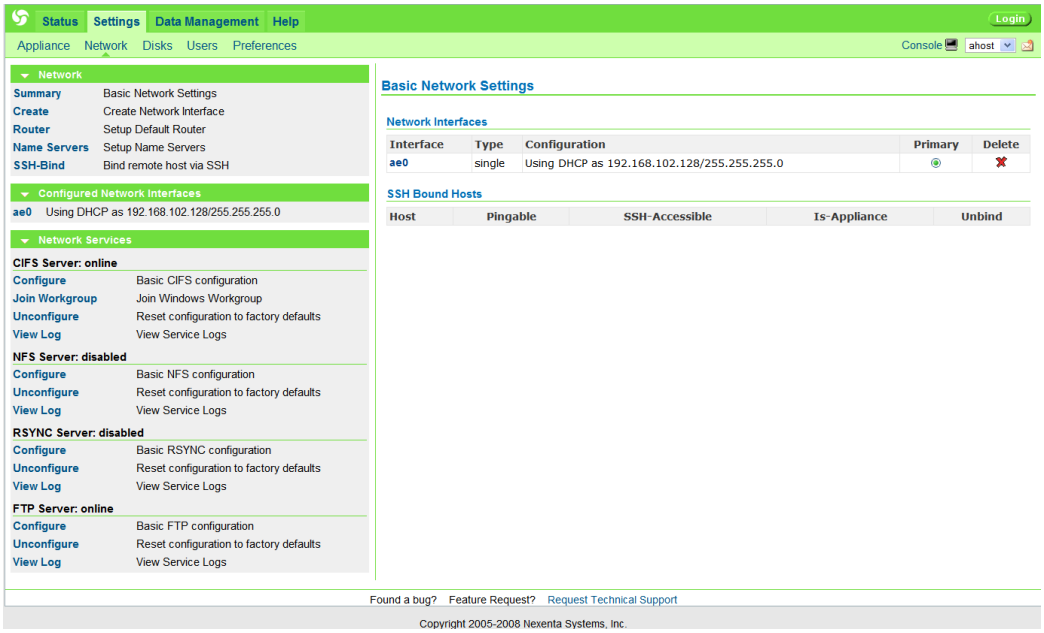

In NMC, network services are configured via '**`setup network service`**'.

In NMV, you will find a number of links to configure, join workgroup, unconfigure, and view the log file (see above). The following screenshot illustrates viewing CIFS logfile:

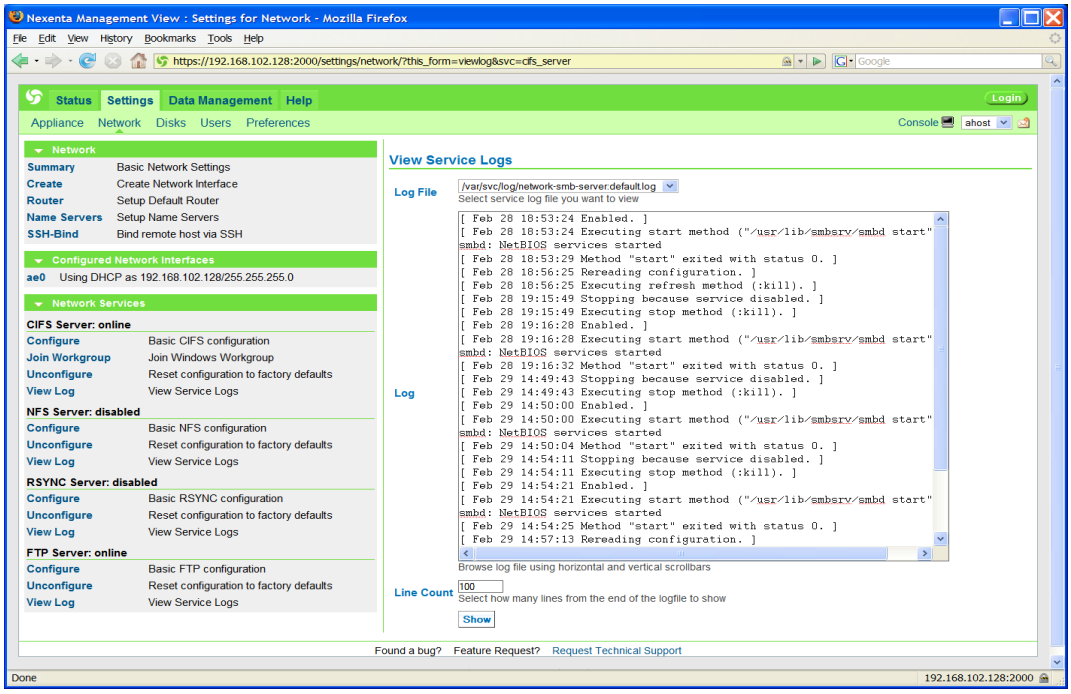

The corresponding NMC command would be '**`show network service cifs-server log`**'.

The important screen, however, is CIFS Server Settings, which you get by clicking on the link denoted as **Configure**. In NMC, the corresponding command would be '**`setup network service cifs-server configure`**' (see footnote[39]).

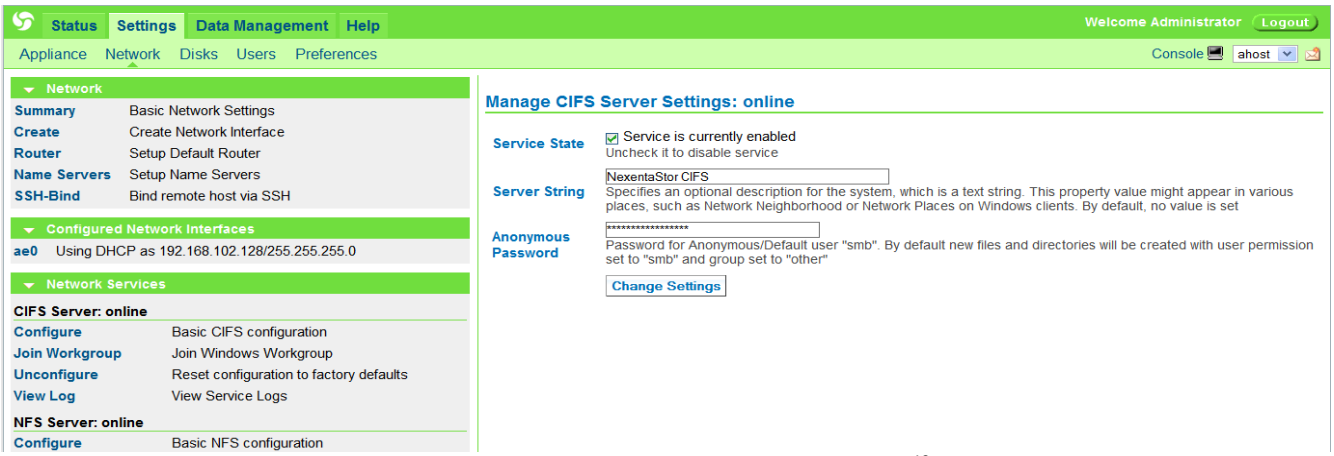

Here, make sure that the server is enabled, and specify a password[40] for anonymous access.

At this point you may create CIFS shares. In NMV, click on the Data Management tab, go to Data Sets, and list all existing folders by clicking on the **Show** link under Folders:

39. As always, there is no need to remember this command. Simply enter **`setup`**, and then keep pressing TAB-TAB and making selection.
40. The default password is sent to you in email, along with the product Registration Key. For more information please see NexentaStor Quick Start Guide at http://www.nexenta.com/docs.

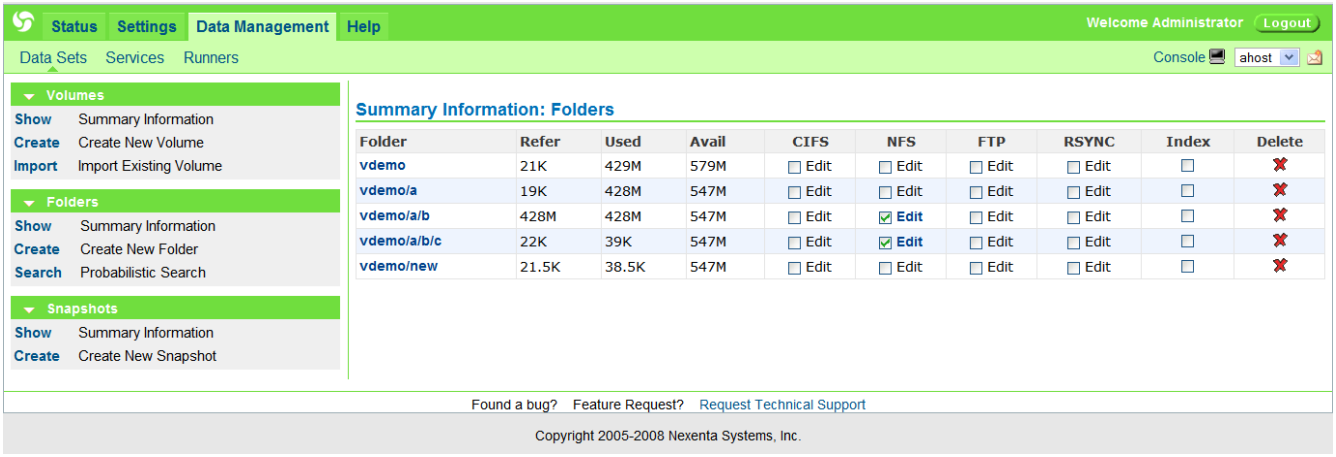

In NMC, the corresponding commands are '**show share**' and '**show folder**' (or '**df**'), for instance:

```
nmc@ahost:/$ show share

FOLDER                  CIFS   NFS    RSYNC  FTP    SAMBA

vdemo/a/b/c             -      Yes    -      -      -

vdemo/new               -      Yes    -      -      -
```

To share a folder, use '**share**' command (NMC) or simply check the corresponding checkbox (NMV). In this example, we are sharing folder '**vdemo/a/b**':

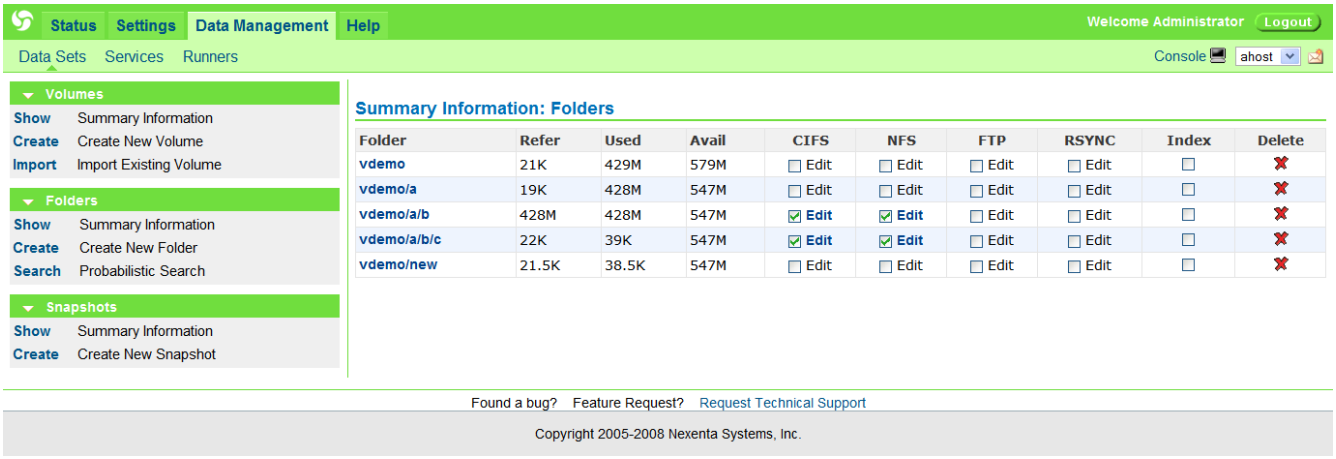

The operation is recursive – it'll share the folder and its sub-folders. Note that in the example above '**vdemo/a/b/c**' got shared as well.

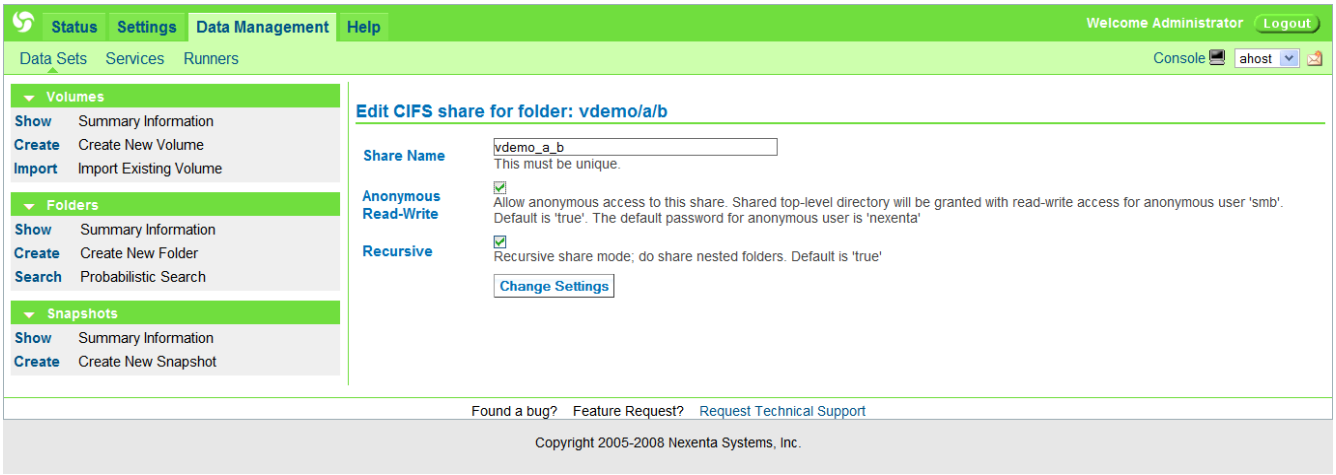

This screenshot (see above) contains several important pieces of information:

1) Anonymous username

   The built-in anonymous username is: **smb**. Unless you are using Active Directory (Section "Using Active Directory"), this is the name you will need to specify to access the share.

   Note that anonymous read/write access is enabled by default. To view or change the default settings, click on the **Edit** link to the right of the corresponding checkbox (see picture above).

2) Default group of CIFS users

   By convention, the pre-configured group of CIFS users is: WORKGROUP. To change the default, go to Setting => Network and click on **Join Workgroup** link.

3) Share name

   By convention, a folder named '**volume-name/folder-name**' becomes a CIFS share named '**volume-name_folder-name**'.

   That fact is reflected on the previous screenshot: '**vdemo/a/b**' will be visible on CIFS clients under name '**vdemo_a_b**' (see above).

Next, we simply point Windows Explorer to the appliance's hostname or IP address. The very first time, Login will be required:

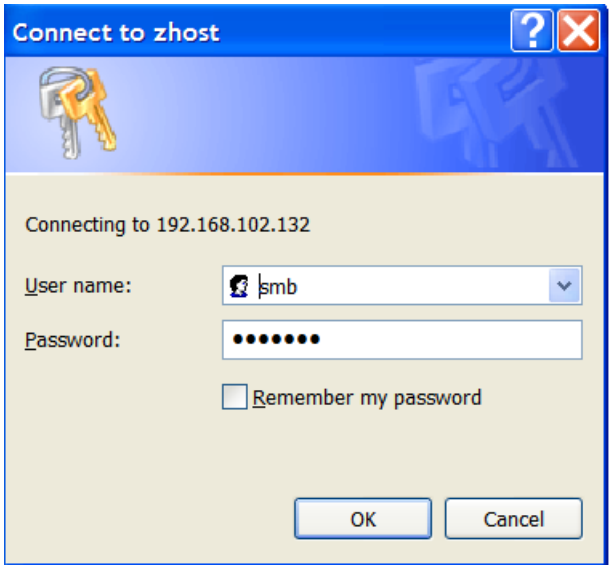

If you forgot the password, please go to CIFS Server Settings (under Settings => Network), click on **Configure**, and re-enter the password. In NMC, the corresponding command would be '**`setup network service cifs-server configure`**'

After successful authentication, the shared folders will show up:

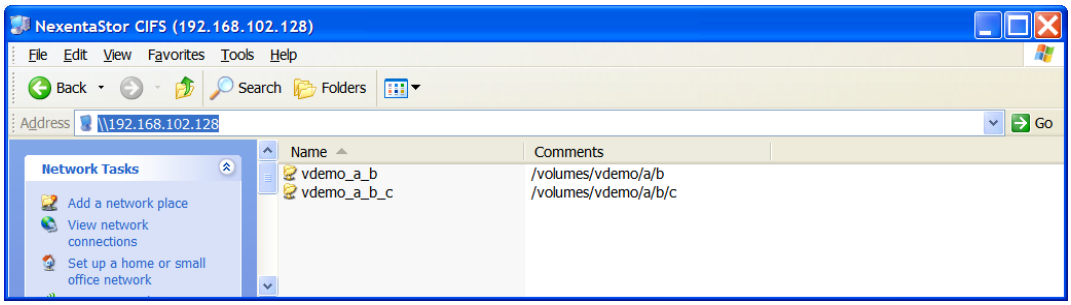

Assuming anonymous access is enabled, we can now start using the NexentaStor folders as Windows directories:

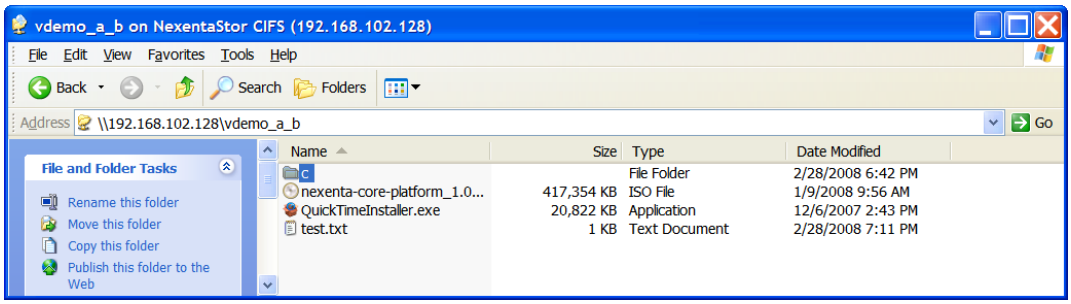

Note on the picture above that NexentaStor folder '**`vdemo/a/b/c`**' shows up here as a directory named '**`c`**'. Since '**`vdemo/a/b`**' was shared recursively, you can access '**`vdemo/a/b/c`**' via its own share named '**`vdemo_a_b_c`**'. Accessing it from '**`vdemo_a_b`**' requires setting up ACL on the CIFS server side that by default NexentaStor currently does not provide.

Sharing a folder via WebDAV

NexentaStor provides an easy way to access (read, write, create, delete) its files and directories via Web-based Distributed Authoring and Versioning (WebDAV).

WebDAV is a popular set of extensions to the HTTP protocol allowing users to collaboratively edit and manage files on remote web servers. NexentaStor certainly includes web server functionality, which you do not have to manage. The web server is managed behind the scenes by the appliance itself. This section explains how to share a NexentaStor folder in a few easy steps.

Actually, there are only two (easy) steps.

Similarly to "Sharing a folder via NFS" and "Sharing a folder via CIFS", go to NMV's Data Management => Shares, select a folder to share, and check the corresponding checkbox:

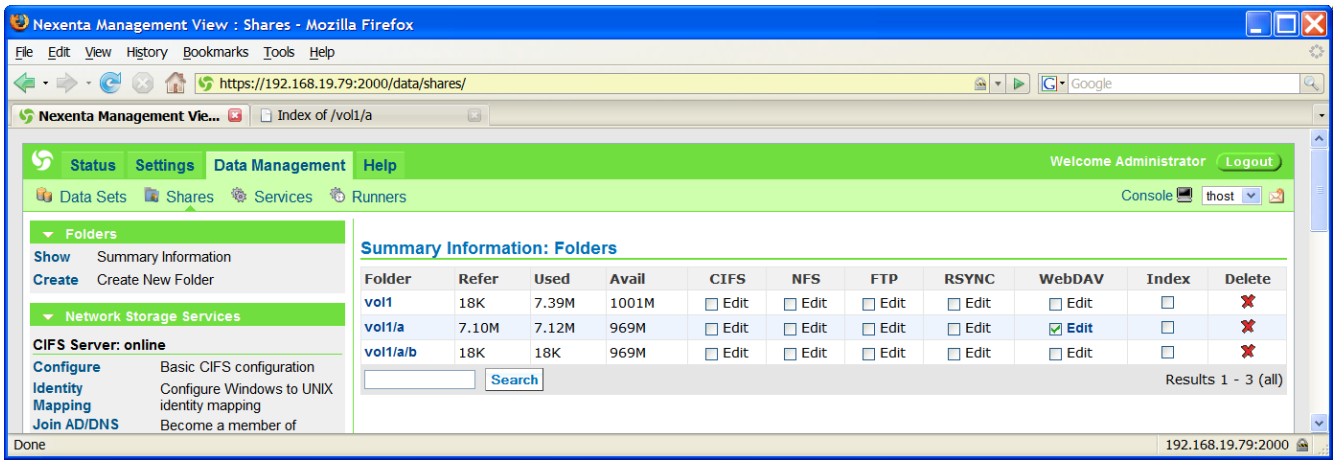

To edit share settings, click on the **Edit** link to the right of the corresponding checkbox. As always, you can do the same via management CLI. For instance, to show existing shares in NMC, the corresponding commands are '**show share**' and '**show share -v**' (for verbose), for instance:

```
nmc@thost:/$ show share
FOLDER                    CIFS   NFS    RSYNC  FTP    WEBDAV  SAMBA
vol1/a                    -      -      -      -      Yes     -
nmc@thost:/$ show share  -v
PROPERTY                  VALUE
folder                    vol1/a
share_name                /volumes/vol1/a
comment                   ""
auth_type                 Basic
anonymous_rw              true
deny_hosts
allow_hosts
```

At this point, the share is already accessible via Internet Browser. For anonymous access, use the the default (anonymous) username:

The built-in anonymous username is: **webdav**. To change it, in NMV go to Data Management => Shares. In NMC, run '**`setup network service webdav-server`**', or simply '**`setup network service`**' and press TAB-TAB

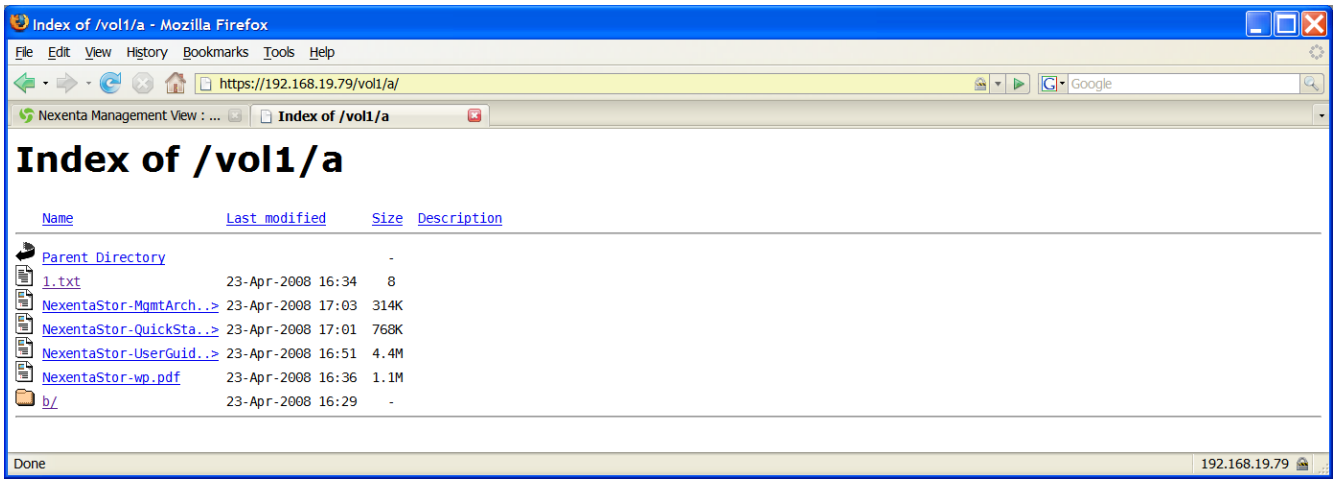

Notice the WebDAV URL, with the following outline: **https://<hostname>/folder**

At this point you can already read-access the WebDAV shared folder, via Internet Browser. For easy access from Windows, you will also need to Map Network Drive.

During our experiments with WebDAV we've found that Windows ability to use WebDAV depends both on the OS version (Windows 2003, XP, Vista), and installed Service Pack. Some of the issues are documented by Microsoft, e. g.:

http://support.microsoft.com/?kbid=841215

The easiest option seems to be: using **`webdrive`**, as demonstrated below.

To map appliance's folder onto Windows network drive, run webdrive application[41], and connect:

41. Downloadable from http://www.webdrive.com/, with trial version available at the time of this writing

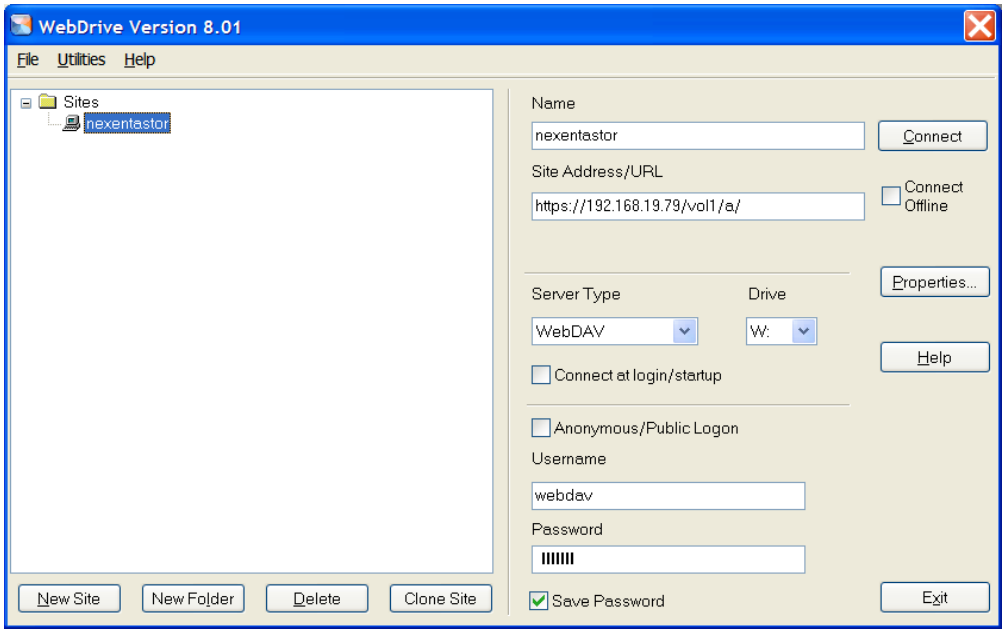

Use the same URL to map the Windows drive letter as used for the Internet Browser (see above).

Once the Windows drive letter (W: in the example above) is created, you have a full native access to appliance's folder(s) from Windows file Explorer, as shown below:

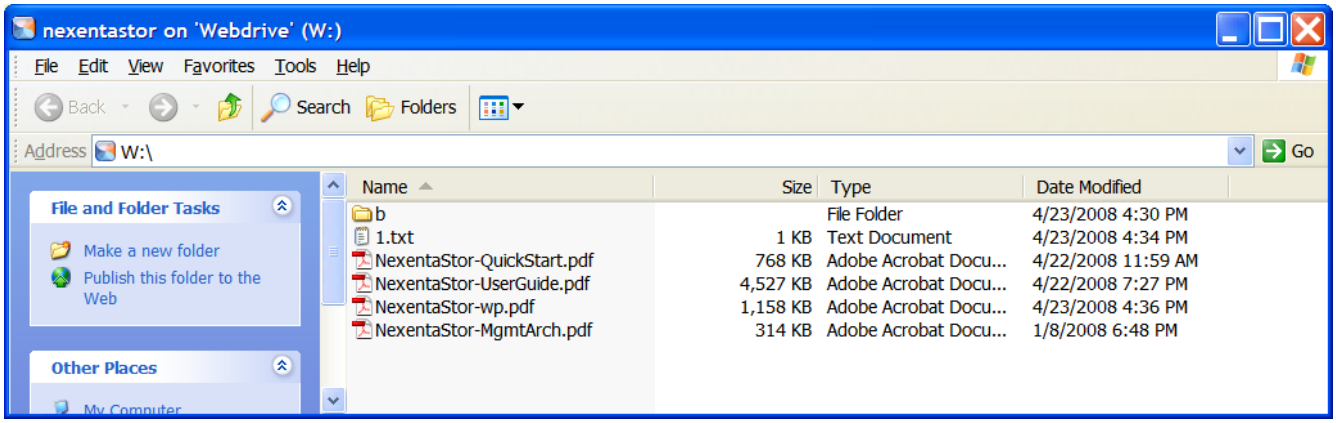

Similarly to "Sharing a folder via NFS" and "Sharing a folder via CIFS" (above), you can limit access to read-only, remove or allow the "anonymous" access option.

With WebDAV, you can use secure connection (the appliance's default, via HTTPS), or unsecured access via plain HTTP. To change these settings, in NMV go to Data Management => Shares. In NMC, run '**`setup network service webdav-server`**', or simply '**`setup network service`**' and press TAB-TAB

Taking snapshots

A major feature of this appliance is virtually unlimited point in time snapshots of the filesystem. The underlying filesystem is of the copy-on-write variety, that further uses transactions with checksums to ensure that all data on disk is uncorrupted. At the same time, a reference to the filesystem state can be made at any time to serve as a checkpoint that can be either referenced or reverted to. These are snapshots, and provide the most efficient online backup solution to date.

Naturally, one can manually create snapshots per volume, such as "**setup volume vol1 snapshot create mysnap**" which would create a point-in-time snapshot for "**vol1**" named "**mysnap**". More examples follow below:

1) **nmc@myhost:/$ create snapshot vol1/a@today**

Take snapshot of **vol1/a** folder and name it '**vol1/a@today**'

2) **nmc@myhost:/$ create volume vol1 folder a snapshot today**

Same as above.

3) **nmc@myhost:/$ create vol1/a snapshot today**

Same as above.

4) **nmc@myhost:/$ setup vol1/a snapshot create today**

Same as above.

One can view snapshots relative to that volume (but not subfolders) using "**show volume data snapshot**". One can also use the **query** command (Section "Command Reference - query") to find snapshots based on a qualifier, such as all snapshots with "**data**" in its name: **query snapshot -e "name=~vol1/a"** - see the **query** command man page ('**query -h**') for details.

Appliance's **auto-snap** service enables easy management of snapshots, providing regular multiple period scheduling of any filesystem, as well as retention policy enforcement. You can selectively choose per snapshot type to keep as little or as many snapshots as you wish, always deleting older snapshots to make way for new ones. Snapshots can be kept for years, and can be generated frequently throughout the day.

You can create snapshot services on a per volume or per folder basis. In many cases a global per-appliance definition for snapshots will do.

For instance, to snapshot all filesystems daily, at 12am:

nmc@myhost:/$ create auto-snap

Prompted by NMC, select 'daily' and then choose 12am time. Select some reasonable number of snapshots to keep (a.k.a. retention policy), for instance 30 (this will provide one month of history).

The same requirements can be executed in one shot, non-interactively:

nmc@myhost:/$ create auto-snap -r -p 1 -T 12am -k 30 -i daily

A "**`show auto-snap`**" would then list this enabled service. As always, the "-h" flag will give full explanation of command-line options and further examples.

The corresponding NMV example follows below:

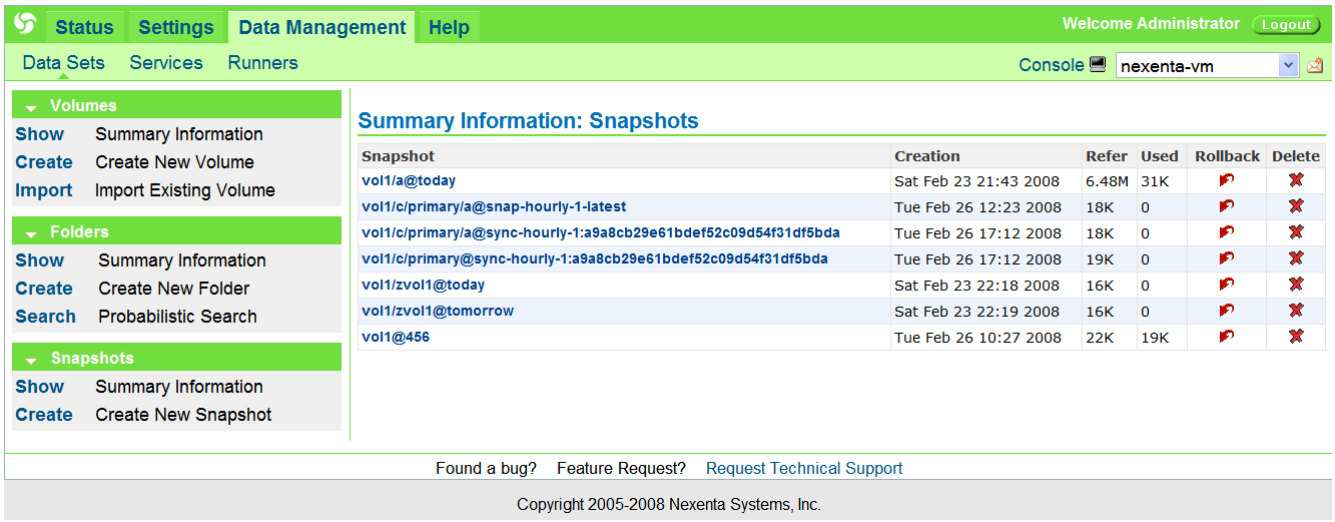

Setting up a tiering service

First, to show existing storage services, use NMC's '**`show`**' command, or NMV's Data Management/ Services pages (below):

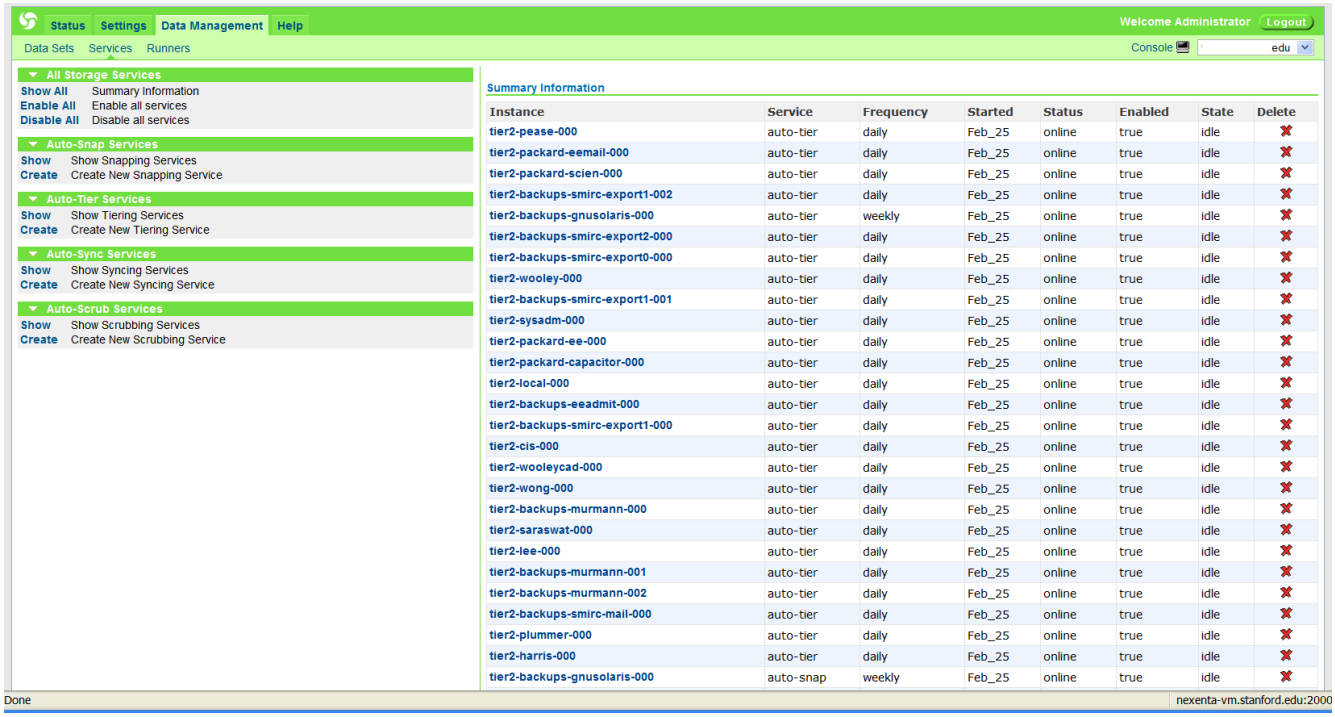

The appliance supports a number of storage services, including tiering service a.k.a. **`auto-tier`**

– for the basics and terminology please refer to Section "Terminology", or run '**help terms**'. In large arrays where the appliance encompasses both first tier and second tier storage, you'll often see local-to-local tiering. Tiering is accomplished by taking a given filesystem or share, breaking into smaller manageable chunks, and replicating that data at that point in time to another volume. Using snapshots at the target end, one can maintain a full efficient backup of the primary storage at unique intervals typical of backups. For more background please refer to Section "".

A simple example of tiering data from an NFS file server to our example volume would be to first create a filesystem to tier to, such as:

```
nmc@myhost:/$ create folder vol1/cad
```

and then to setup an auto-tier from our source NFS server (for our CAD tools in this case)[42]:

```
nmc@myhost:/$ create auto-tier -s rsync+nfs://hd/cad/* -d
              /vol1/cad -T 3am -i daily -p 1
```

The **auto-tier** creation command above will pull data via NFS, breaking up the synchronization into chunks determined by the number of subdirectories found in cad. The **-i** flag determines the type of periodic interval (**'minute', 'hourly', 'daily', 'weekly', 'monthly'**), the **-T** details the time of day, etc.

As always, the options are documented and available via -h:

```
nmc@myhost:/$ create auto-tier -h
```

The corresponding man page provides extensive details and more tiering examples.

Irrespectively of the man pages, it is fair to assume that the first-time and maybe even second-time users will never want to type a complete command in its entirety... As always, for the an interactive, dialog-driven way, run:

```
nmc@myhost:/$ create auto-tier
```

or, the same:

```
nmc@myhost:/$ setup auto-tier create
```

And follow the prompts.

Tiering over ssh would be another often-used example (see screenshot below).

42. For similar web GUI functionality, please see Section "NMV Configuration and Fault Management"

```
nmc@ahost:/$ setup auto-tier
create  destroy disable enable  run     show    volume  vdemo/  <?>
nmc@ahost:/$ setup auto-tier create

nmc@ahost:/$ create auto-tier
Frequency             : hourly
Period                : 1
Enable auto-snap      :
Protocol              : rsync+ssh
Source host           : morioka.stanford.edu
Tiering snapshot      :
Source directory      : data/tmp/*
Destination directory : vdemo/tmp
 ../ ./
  Choose existing local auto-tier destination to where to replicate. Select
  directory by pressing ENTER on './'. Navigate with arrow keys (or hjkl). Press
  Ctrl-C to exit.
```

Any **auto-tier** that use ssh as a transport also needs to have the source and destination hosts pre-authorized through "**ssh-bind**" . To tier over ssh, the local and remote hosts must be *ssh-bound*.

The auto-tier service is not limited just to the first two tiers, as tertiary tiering for more critical data is also common. As legal and business drivers dictate, tiering will also include access policy enforcement, limiting data access to restricted personnel to over longer periods of time. The service addresses Disaster Recovery, Continuous Data Protection, and other emerging technology solutions as every aspect of business goes digital, and must be protected, preserved, or potentially expired as required.

A Note on SSH Binding

Tiering over SSH, syncing (that is, running NexentaStor **auto-sync** service – see Section "Terminology") over SSH, and managing remote appliance over SSH – these are the features in use by the appliance that rely on pre-existing trusts being established between multiple devices, whether source or destination data volumes, or other appliances. SSH binding allows for remote management as well as strong cryptographic verification of multi-host trusts, forming a network of ssh-bound hosts.

One binds and unbinds hosts via:

- "**setup network ssh-bind**"

  and

- "**setup network ssh-unbind**".

To view currently bound hosts, type "**show network ssh-bindings**".

A simple example to bind to server2 is "**setup network ssh-bind server2**" resulting in a prompt for that server's root password. You'll be prompted in a similar fashion to unbind a host. If the host is no longer up, one can force an unbind with a flag, as in "**setup network ssh-unbind server2 -f**".

For more information, run "**setup network ssh-bind -h**"

Auto-sync: the service that replicates both data and metadata

Designed from ground up for consistency and ease of use, NexentaStor provides the same general interfaces to manage all storage services. All NMC commands, NMV UI pages, and all general principles discussed in the previous section with regard to **auto-tier** service apply to **auto-sync** as well.

Yet, there are important differences between the appliance's replication storage services. One size does not in fact fits all – if it weren't true a single replication service would suffice.

When deciding which replication service to deploy in your environment, please consider the discussion presented above, in Section "Data Replication", or the following FAQ entry on the website support page:

What is the difference between 'auto-sync' and 'auto-tier' storage services?

Both **auto-sync** and **auto-tier** can be monitored in real time, in terms of bytes transferred and bandwidth. Here's an example that effectively shows auto-sync generated traffic at one-second interval:

```
nmc:/$ show auto-sync :vol1-a-000 stats -i 1
TCP CONNECTION                                  SNEXT      RNEXT      TRANSFERED
192.168.37.128.39305=>192.168.37.134.22         2176247586 4217818800 -
192.168.37.128.39305=>192.168.37.134.22         2176255378 4217819152 8.14 KB
192.168.37.128.39305=>192.168.37.134.22         2176266306 4217819632 11.41 KB
192.168.37.128.39305=>192.168.37.134.22         2176273922 4217819952 7.94 KB
...
```

Upgrading appliance software – Live

NexentaStor software upgrades and rollbacks are managed using the following easy-to-use NMC commands:

```
1) show appliance syspool
2) show appliance upgrade
   and, symmetrically:
3) setup appliance syspool
4) setup appliance upgrade
```

This section demonstrates live-upgrading appliance software, and then rolling back the system to the pre-upgrade checkpoint. (For a quick introduction and background, please refer to Section "Appliance Software Upgrade").

Note that some of the screenshots in this section may be outdated – see footnote on page #30.

The first screenshot below results from running '**show appliance syspool**'. The particular example shows an active system folder '**rootfs-nmu-002**' and a single system checkpoint '**rootfs-nmu-001**':

```
nmc
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE              CURRENT DEFAULT
rootfs-nmu-002  Sep 14  0:39 2007 active-upgrade-checkpoint  Yes     Yes
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint        No      No
nmc@myhost:/$
```

Notice a few things:

The **current** system folder **rootfs-nmu-002** appears to be **active**, which means that the next reboot will "bring" us into this same system folder. At any point in time there can be only one **current** folder and only one **active** folder. "Activating" a system folder simply means that the next reboot will make it **current**.

NexentaStor supports two types of checkpoints:

1. '**rollback-checkpoint**' - as the name implies, can be used to rollback the changes and restore the system to its pre-upgrade state.
2. '**upgrade-checkpoint**' – a system folder that contains upgrades.

To demonstrate the usage, we will first try to upgrade.

Command '**setup appliance upgrade**' does it:

```
nmc
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE              CURRENT DEFAULT
rootfs-nmu-002  Sep 14  0:39 2007 active-upgrade-checkpoint  Yes     Yes
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint        No      No
nmc@myhost:/$ setup appliance upgrade
This operation could potentially break client connections. Proceed?  Yes
Checking repository sources. Please wait...
Found new upgrades!
Verifying upgrades...
Trying to gain exclusive access to the appliance.
This operation might take up to 30 seconds. Please wait...
Exclusive access granted.
Downloading upgrades and checking if reboot will be required.
This may take a few minutes. Please wait...
Verifying free space...
Success. Upgrade requires ~22.73MB
Upgrade is in progress. Please DO NOT interrupt...
Creating Rollback Checkpoint...

Rollback Checkpoint has been created: rootfs-nmu-003

Use NMC's 'show appliance syspool' command to list all available
upgrade/rollback checkpoints

(Reading database ... 28093 files and directories currently installed.)
Preparing to replace nms 0.98 (using .../nms_0.98_solaris-i386.deb) ...
Stopping NMS service... done
Waiting for NMS to shutdown... done
Clearing from maintenance... done
Unpacking replacement nms ...
Preparing to replace nmc 0.98 (using .../nmc_0.98_solaris-i386.deb) ...
Unpacking replacement nmc ...
Preparing to replace nmv 0.98 (using .../nmv_0.98_solaris-i386.deb) ...
Unpacking replacement nmv ...
Setting up nms (0.98) ...
Starting NMS service... done
Waiting for NMS to come up... done

Setting up nmc (0.98) ...
Re-starting NMCD service... done

Setting up nmv (0.98) ...

Appliance was successfully upgraded. Press any key to re-login...
```

The upgrade is done: Nexenta management software has got upgraded. Post-upgrade NMC will come up[43] with its familiar Welcome screen:

```
****************************************************************
*                    Nexenta Management Console                *
*                        Version 0.98                          *
*                                                              *
*  press     TAB-TAB to list and/or complete available options  *
*                                                              *
*  type      help   for comprehensive guide to NMC             *
*            exit   to exit local NMC, remote NMC, or group mode *
*            q[uit] or Ctrl-C  to exit NMC dialogs             *
*            q[uit] or Ctrl-C  to exit NMC Viewer              *
*                                                              *
*            option bg=dark     to set dark background         *
*                                                              *
*  use       command history, command usage guides (-h option)  *
*            command completions, output redirection, | pipes  *
*            reverse search (Ctrl-R), background execution (&), *
*            and more.                                         *
*                                                              *
*  type      help   and press TAB-TAB                          *
*                                                              *
*  Management GUI: https://10.16.16.64:2000                    *
*                                                              *
****************************************************************
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE             CURRENT DEFAULT
rootfs-nmu-003  Sep 28 12:02 2007 rollback-checkpoint         No      No
rootfs-nmu-002  Sep 14  0:39 2007 active-upgrade-checkpoint   Yes     Yes
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint         No      No
nmc@myhost:/$ _
```

As a result of this particular upgrade transaction the system created a new checkpoint '**rootfs-nmu-003**' – which *is the* pre-upgrade snapshot of the system (picture below). To demonstrate NexentaStor rollback capability, let us roll the changes back. This is achieved simply by activating the corresponding rollback checkpoint – the '**rootfs-nmu-003**' checkpoint in the case.

Notice the usage of '**setup appliance syspool**' command (below):

```
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE             CURRENT DEFAULT
rootfs-nmu-003  Sep 28 12:02 2007 rollback-checkpoint         No      No
rootfs-nmu-002  Sep 14  0:39 2007 active-upgrade-checkpoint   Yes     Yes
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint         No      No
nmc@myhost:/$ setup appliance syspool rootfs-nmu-003
activate  destroy   property
nmc@myhost:/$ setup appliance syspool rootfs-nmu-003 activate
Activate rollback checkpoint 'rootfs-nmu-003' ?  Yes
Checkpoint 'rootfs-nmu-003' has been activated. You can reboot now.
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE             CURRENT DEFAULT
rootfs-nmu-003  Sep 28 12:02 2007 active-rollback-checkpoint No      Yes
rootfs-nmu-002  Sep 14  0:39 2007 upgrade-checkpoint         Yes     No
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint         No      No
nmc@myhost:/$ setup appliance reboot
Reboot appliance 'myhost' ?  Yes
Rebooting now...
```

Upon reboot, the activated system folder shows up as **current** (see below). This system folder is the exact pre-upgrade snapshot of the system. The '**002**' checkpoint is preserved, which makes it possible to boot into it at any point in the future to take advantage of the new (upgraded) functionality. Notice also the SYSTEM NOTICE below – it tells us that upgrades are available, which makes sense because we just now rolled them back..

```
Last login: Fri Sep 28 12:36:37 2007
                      * * *
                   SYSTEM NOTICE

     New upgrades available!
     Run 'setup appliance upgrade' to upgrade your system.
     Use 'option' not to display this message again.

nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE             CURRENT DEFAULT
rootfs-nmu-003  Sep 28 12:02 2007 active-rollback-checkpoint Yes     Yes
rootfs-nmu-002  Sep 14  0:39 2007 upgrade-checkpoint         No      No
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint         No      No
nmc@myhost:/$ _
```

43. Only the very first login. To configure NMC to always display its Welcome screen on startup, see Section "Preferences"

On the picture above, notice the '**rootfs-nmu-001**' checkpoint that was present at the beginning of this section. Activating the '**001**' checkpoint instead of the currently active '**003**' would reset the system even further back in time, in this example – to September 14 0:33 2007.

NexentaStor supports a configurable number of system checkpoints, with full flexibility to go back and forth in time.

On the off-chance that live-upgraded appliance does not boot, you can always use GRUB boot menu to explicitly select one of the available checkpoints. Note that rollback checkpoints are **always** created prior to live upgrades, which makes iit ultimately safe to upgrade even a live running system.

For safe upgrade, please see the next section.

Upgrading appliance software – Safe

Safe upgrading means updating not a live-and-running system (see previous Section), but its exact clone. This section demonstrates the safe-upgrading facility. For a quick introduction and a background, please refer to Section "Appliance Software Upgrade".

First, initiate safe upgrade by explicitly specifying -s (safe) option in the command line. Note that upgrades that require reboot are automatically detected by the system and performed with the -s (safe) option enabled. For demonstration purposes, we will first safe-upgrade a single selected package, and later in this section upgrade the rest packages.

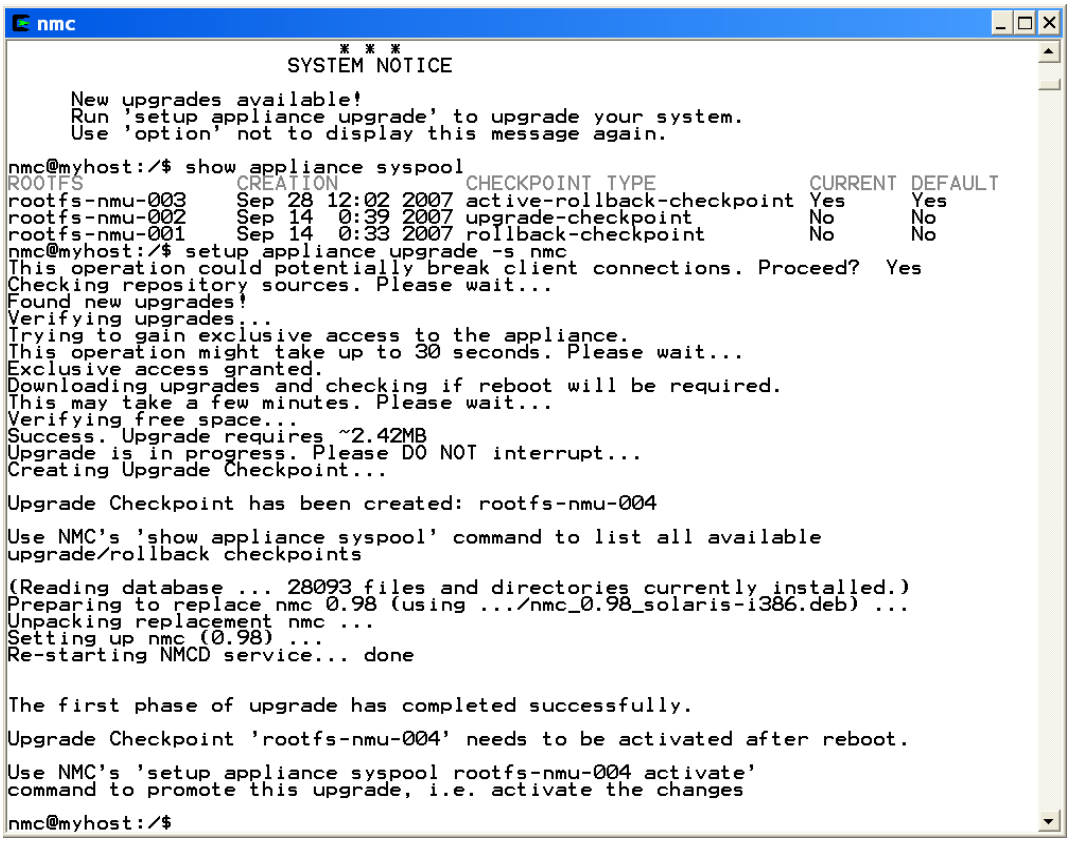

Notice the '**`setup appliance upgrade`**' command and creation of a new ('**`004`**') checkpoint.

Similar to live upgrade (previous section), the safe upgrade transaction creates a checkpoint – 'rootfs-nmu-004' in this case (see above).

This particular checkpoint *is the* upgraded system – the current root filesystem has remained intact, and the appliance can **continue operating** (all services running, all users connected, etc.) until administrator finds it **convenient/feasible/necessary** to switch to the upgraded state.

More exactly, at this point there are 3 (three) options:

1. You can reboot now, make sure that system is healthy and then activate the current (i. e., newly created) checkpoint;
2. You can activate the newly created (upgraded) checkpoint right now, and then reboot;

3. Or, you can simply continue using the appliance as is and do (1) or (2) later.

The safest option is certainly the (1), and the argument here is two-fold. First, your current working system folder is not modified by the upgrade and therefore available for usage after we reboot into the new checkpoint. Second, on the off-chance that the the newly created (upgraded) checkpoint is not good, the next subsequent reboot will automatically bring you into the current (unmodified) system folder, which would be ('**003**') in this particular case.

The same 3 available options are displayed on the management console:

```
                        * * *
                    SYSTEM NOTICE

    The first phase of upgrade has completed successfully:
      - created Upgrade Checkpoint 'rootfs-nmu-___'
      - the appliance is ready to reboot into the new checkpoint

   +-----------------------------------------------------------------+
   |                                                                 |
   |  At this point you have three options:                          |
   |                                                                 |
   |    1. You can reboot now, make sure that system is healthy and  |
   |       then activate the current (i.e., newly created) checkpoint;
   |    2. You can activate the newly created (upgraded) checkpoint  |
   |       right now, and then reboot;                               |
   |    3. Or, you can simply continue using the appliance as is and |
   |       do (1) or (2) later.                                      |
   |                                                                 |
   |  See also: 'setup appliance syspool'                            |
   |  See also: 'show appliance syspool'                             |
   |  See also: 'setup appliance reboot'                             |
   |                                                                 |
   +-----------------------------------------------------------------+
Would you like to follow the option (1) above and reboot now ? (y/n) _
```

The default option is (1) – reboot into the upgraded system. To take the default, simply answer 'Yes' to the posted question (above).

Rebooting at this point simply means making use of the safely upgraded system bits - which typically brings us new features, security fixes, bug fixes, and new functionality. Note that the reboot into **upgrade-checkpoint** is 100% safe in a sense that it is always possible to go back to the pre-upgraded state – the '**003**' in this case.

In this tutorial for demonstration purposes we'll say 'No' to the question posted on the console above, and use the 2nd option – activate the new checkpoint ('**004**') right away:

```
The first phase of upgrade has completed successfully.

Upgrade Checkpoint 'rootfs-nmu-004' needs to be activated after reboot.

Use NMC's 'setup appliance syspool rootfs-nmu-004 activate'
command to promote this upgrade, i.e. activate the changes

nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE            CURRENT DEFAULT
rootfs-nmu-004  Sep 28 12:45 2007 upgrade-checkpoint         No      No
rootfs-nmu-003  Sep 28 12:02 2007 active-rollback-checkpoint Yes     Yes
rootfs-nmu-002  Sep 14  0:39 2007 upgrade-checkpoint         No      No
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint        No      No
nmc@myhost:/$ setup appliance syspool rootfs-nmu-004 activate
Activate upgrade checkpoint 'rootfs-nmu-004' ?  Yes
Checkpoint 'rootfs-nmu-004' has been activated. You can reboot now.
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE            CURRENT DEFAULT
rootfs-nmu-004  Sep 28 12:45 2007 active-upgrade-checkpoint  No      Yes
rootfs-nmu-003  Sep 28 12:02 2007 rollback-checkpoint        Yes     No
rootfs-nmu-002  Sep 14  0:39 2007 upgrade-checkpoint         No      No
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint        No      No
nmc@myhost:/$
```

And then reboot. Upon reboot, notice the new ('004') entry in GRUB menu:

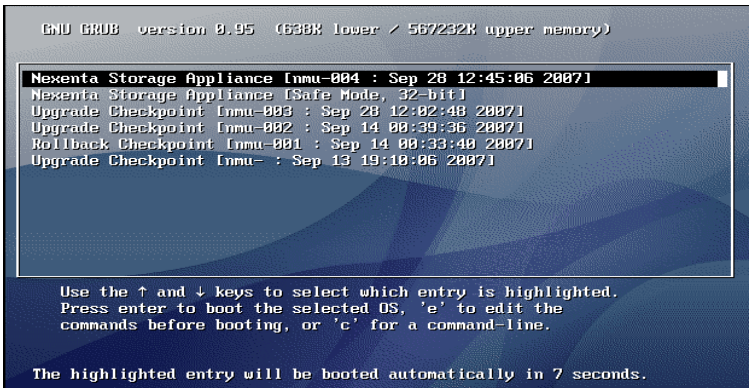

... and the familiar NMC Welcome (see also footnote on page #70):

```
*****************************************************************
*                    Nexenta Management Console                 *
*                         Version 0.98                          *
*                                                               *
*  press     TAB-TAB to list and/or complete available options  *
*                                                               *
*  type      help   for comprehensive guide to NMC              *
*            exit   to exit local NMC, remote NMC, or group mode *
*            q[uit] or Ctrl-C  to exit NMC dialogs              *
*            q[uit] or Ctrl-C  to exit NMC Viewer               *
*                                                               *
*            option bg=dark     to set dark background          *
*                                                               *
*  use       command history, command usage guides (-h option)  *
*            command completions, output redirection, | pipes   *
*            reverse search (Ctrl-R), background execution (&),  *
*            and more.                                          *
*                                                               *
*  type      help   and press TAB-TAB                           *
*                                                               *
*  Management GUI: https://10.16.16.66:2000                     *
*                                                               *
*****************************************************************
                        * * *
                    SYSTEM NOTICE

     New upgrades available!
     Run 'setup appliance upgrade' to upgrade your system.
     Use 'option' not to display this message again.

nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION                 CHECKPOINT TYPE             CURRENT DEFAULT
rootfs-nmu-004  Sep 28 12:45 2007 active-upgrade-checkpoint  Yes     Yes
rootfs-nmu-003  Sep 28 12:02 2007 rollback-checkpoint        No      No
rootfs-nmu-002  Sep 14  0:39 2007 upgrade-checkpoint         No      No
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint        No      No
nmc@myhost:/$
```

Safe upgrade – another example

Next, we'll safe-upgrade the appliance without specifying a particular package. This is the default (and recommended) operation – in general, updating only parts of appliance's software is not recommended.

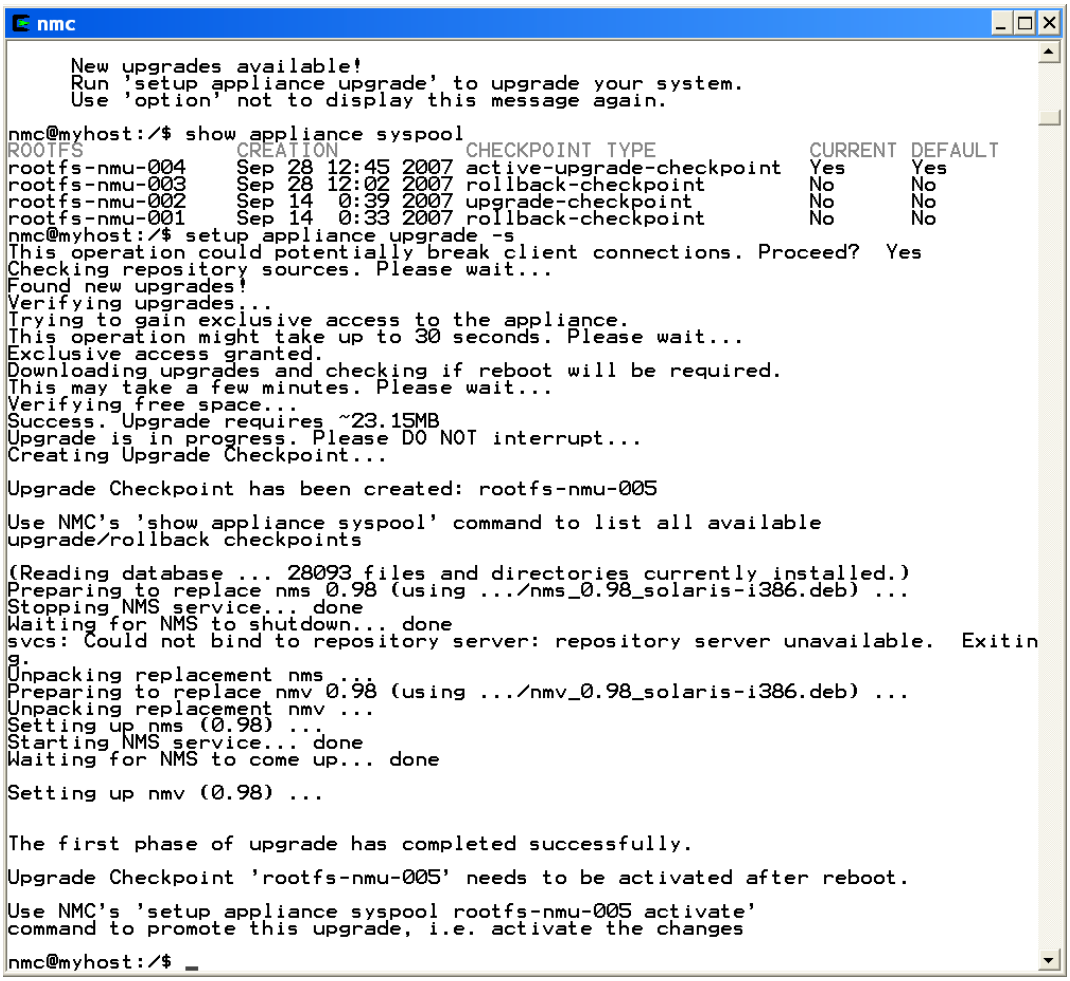

Done. And another ('005') system checkpoint created:

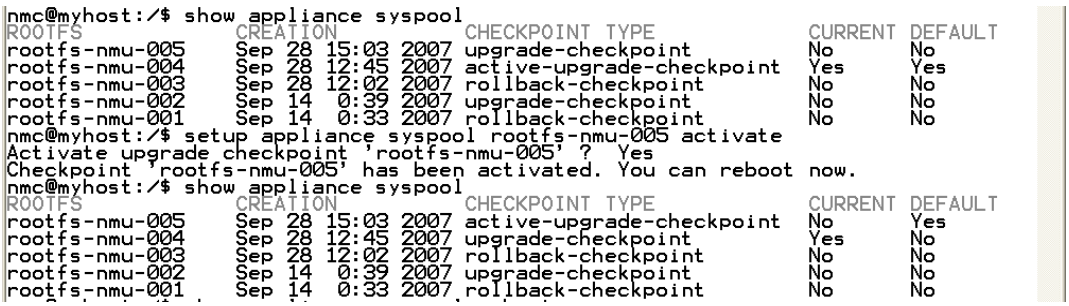

Rebooting at this point simply means making use of the safely upgraded system bits - which brings new features, security fixes, bug fixes, and new upgraded functionality. Note that the reboot into **upgrade-checkpoint** is 100% safe in a sense that it is always possible to go back to the pre-upgraded state – the '**004**' in this case.

Note that there is no limitation on a number of upgrade checkpoints – each upgrade transaction automatically adds another unique checkpoint; only one of those checkpoints can be active at any given time.

And the final picture (below). Upon reboot into the '**005**' the latter shows up as **current** (because we are effectively *in it*) and **default** (because it has been activated – see above), while all the previous system checkpoints still remain available for possible usage. The checkpoints retention policy works similar to retention policy associated with '**auto-sync**' and '**auto-snap**' generated snapshots – the appliance keeps only so many[44] most recent checkpoints (for NMS properties, their default values and descriptions see for instance Section "Server Options").

Asking for upgrades at this point will result in "No new upgrades available" (see below). The appliance is safely upgraded.

```
nmc@myhost:/$ show appliance syspool
ROOTFS          CREATION              CHECKPOINT TYPE             CURRENT DEFAULT
rootfs-nmu-005  Sep 28 15:03 2007 active-upgrade-checkpoint  Yes     Yes
rootfs-nmu-004  Sep 28 12:45 2007 upgrade-checkpoint         No      No
rootfs-nmu-003  Sep 28 12:02 2007 rollback-checkpoint         No      No
rootfs-nmu-002  Sep 14  0:39 2007 upgrade-checkpoint         No      No
rootfs-nmu-001  Sep 14  0:33 2007 rollback-checkpoint         No      No
nmc@myhost:/$ setup appliance upgrade
This operation could potentially break client connections. Proceed?  Yes
Checking repository sources. Please wait...
No new upgrades available.
nmc@myhost:/$
```

44. The corresponding NMS property is called '**upgrade_rootfs_keep**', with default value equal 16 at the time of this writing

Setting up remote mirroring – CDP

Setting up remote mirroring a. k. a. **auto-cdp** for a given local appliance's volume is easy, and takes a single command:

```
nmc:/$ create auto-cdp
```

and the following steps:

1) selecting a local (primary) volume to replicate
2) selecting remote appliance
3) selecting disks (or LUNs) on the remote appliance to serve as block-level replicas of the disks (LUNs) of the local volume

As always, to carry out the 1-2-3 steps, NMC provides a guided multiple-choice interactive environment. As always, the same steps can be executed via command line, using the options specified above.

Once initiated, **auto-cdp** will start initial synchronization, including transferring the local (primary) volume's metadata, which will effectively create a secondary (remote) volume out of the corresponding remote disks. The appliance's **auto-cdp** will keep both data and ZFS metadata on the replicated disks in-sync, at all times.

Note that **auto-cdp** requires using either DNS hostname for the local and remote appliances, or their "replacements" via local host tables.

See '**setup appliance hosts -h**' for more information.

It is recommended **not** to use the primary (active) volume during the initial CDP synchronization. Any updates on the primary during this period of time may considerably delay the initial synchronization.

You can disable (stop) any **auto-cdp** instance at any time, and re-enable it at any time. One possible scenario of usage would be:

a) Disable a given **auto-cdp** instance, for instance:

```
nmc@myhost:/$ setup auto-cdp :vol1 disable
```

Note that 'vol1' here is a name of the volume that is being **auto-cdp** replicated. By convention, **auto-cdp** instance is named after the corresponding primary volume.

b) Import secondary volume on the remote appliance, and start using it:

```
nmc@rem_host:/$ setup volume import vol1
```

It is recommended to export the primary 'vol1' out of the **myhost** appliance prior to importing its duplicate on the remote (**rem_host**) appliance.

c) Use the volume on the remote appliance. After some period of usage – re-enable the service and make sure that the changes are synchronized back to the **myhost** appliance from **rem_host** appliance:

```
nmc@myhost:/$ setup auto-cdp :vol1 enable -r
```

Note the -r (reverse) option. See '**setup auto-cdp <name> enable -h**' for details.

d) To monitor auto-cdp progress, use either one of the following commands:

```
nmc:/$ show auto-cdp :vol1 stats

or

nmc:/$ show auto-cdp :vol1 iostat
```

To troubleshoot, execute either one of the following two commands:

1) '**setup auto-cdp <name> enable**' - to re-enable the service
2) '**setup auto-cdp <name> lun <LUN> enable**' - to re-enable specific <LUN> pair.

The latter is especially useful when one specific pair of syncing LUNs appears to be in a so called "logging" mode, and will not change states for the rest of LUN pairs in the selected syncing <service> group.

In all cases use the traffic monitoring commands:

'**show auto-cdp <name> stats**'

and

'**show auto-cdp <name> iostat**'

to see the state and progress of the service. The former shows service-generated network traffic, the latter shows I/O statistics.

Creating auto-cdp – example

In the following example appliance '**testbox1**' is a primary, '**testbox2**' is a secondary. This example includes all NMC prompts – it is a complete demonstration of **auto-cdp** creation:

```
nmc@testbox1:/$ create auto-cdp
Remote appliance  : 192.168.37.128
Remote for c2t1d0 : c2t1d0
Remote for c2t0d0 : c2t0d0
Creating new Auto CDP service 'auto-cdp:vol1', please wait...
Successfully created service 'auto-cdp:vol1'
Enable it now?  Yes
Enabling service, please wait...
PROPERTY                VALUE
name                    :vol1
max_q_fbas              16384
autosync                off
max_q_writes            4096
async_threads           2
state                   syncing
to_host                 testbox2
from_host               testbox1
type                    active
TESTBOX1          TESTBOX2
c2t1d0      =>      c2t1d0
c2t0d0      =>      c2t0d0
The local host is 'active' auto-cdp node.
```

Once the initial synchronization between a pair of active (primary) and passive (secondary) volumes commences, you can monitor it either via '**show auto-cdp <name> stats**' or '**show auto-cdp <name> iostat**' NMC commands.

In fact, these two commands are always useful, in terms of monitoring the data replication traffic, whether this is **auto-cdp**, **auto-sync** or **auto-tier** service. However, **auto-cdp** traffic monitoring is particularly useful at the time of the initial block-level syncing:

It is recommended not to use the primary (active) volume during the very first (the initial) CDP synchronization. Any updates on the primary during this period of time may considerably delay the initial synchronization. Note also that during this phase a major part of the available I/O bandwidth is used by the auto-cdp service, which is yet another reason to let it run through as soon as possible.

See '**show auto-cdp <name> stats**' for more information.

Following is an example of '**show auto-cdp <name> stats**' output:

```
nmc@testbox1:/$ show auto-cdp :vol1 stats -i 1
TCP CONNECTIONS                                SNEXT      RNEXT      TRANSFER
192.168.37.128.1022-192.168.37.134.121         1313611534 3140553278 1.60 MB
192.168.37.128.1022-192.168.37.134.121         1314180374 3140554114 569.68 KB
192.168.37.128.1022-192.168.37.134.121         1314838374 3140554994 658.88 KB
192.168.37.128.1022-192.168.37.134.121         1316976874 3140557854 2.14 MB
192.168.37.128.1022-192.168.37.134.121         1321352574 3140563706 1.25 MB
192.168.37.128.1022-192.168.37.134.121         1322109274 3140564718 757.71 KB
192.168.37.128.1022-192.168.37.134.121         1322536974 3140565290 331.77 KB
192.168.37.128.1022-192.168.37.134.121         1323030474 3140565950 494.16 KB
192.168.37.128.1022-192.168.37.134.121         1323820074 3140567006 790.66 KB
192.168.37.128.1022-192.168.37.134.121         1327471974 3140571890 955.38 KB
192.168.37.128.1022-192.168.37.134.121         1328722174 3140573562 1.25 MB
...
```

Once the traffic stops, you'll be able to see the block-level replicated volume on the remote side:

```
nmc@testbox2:/$ show auto-cdp :vol1 -v
PROPERTY                VALUE
name                    :vol1
max_q_fbas              16384
autosync                off
max_q_writes            4096
async_threads           2
state                   logging
to_host                 testbox2
from_host               testbox1
type                    passive
TESTBOX1          TESTBOX2
c2t1d0       =>     c2t1d0
c2t0d0       =>     c2t0d0
The local host is 'passive' auto-cdp node.
```

Reverse mirroring – example

In the following 6-steps example appliance '**testbox1**' is again a primary, and '**testbox2**' is a secondary. The reverse mirroring starts from exporting a volume from the primary appliance (Step #1)... In short, several preparation steps need to be performed before actually enabling reverse mirroring from '**testbox2**' to '**testbox1**' (Step #5 below):

```
Step #1. testbox1 (primary): first, export vol1

nmc@testbox1:/$ setup volume vol1 export
Export volume 'vol1' and destroy all associated shares ?  Yes

Step #2. testbox2 (secondary): import vol1

nmc@testbox2:/$ setup volume import vol1
volume: vol1
 state: ONLINE
 scrub: none requested
config:
        NAME         STATE     READ WRITE CKSUM
        vol1         ONLINE       0     0     0
          mirror     ONLINE       0     0     0
            c2t0d0   ONLINE       0     0     0
            c2t1d0   ONLINE       0     0     0

errors: No known data errors

Step #3. Use volume till problem with primary is resolved...

Step #4. testbox2 (secondary): export vol1

nmc@testbox2:/$ setup volume vol1 export
Export volume 'vol1' and destroy all associated shares ?  Yes

Step #5. testbox1 (primary): reverse syncing

nmc@testbox1:/$ setup auto-cdp :vol1 enable -r
Enable reverse synchronization for auto CDP servce 'vol1'?  Yes
Enabling service, please wait...
Auto CDP service ':vol1' enabled.

Step #6. testbox1 (primary): import vol1

nmc@testbox1:/$ setup volume import vol1
volume: vol1
 state: ONLINE
 scrub: none requested
config:
        NAME         STATE     READ WRITE CKSUM
        vol1         ONLINE       0     0     0
          mirror     ONLINE       0     0     0
            c2t0d0   ONLINE       0     0     0
            c2t1d0   ONLINE       0     0     0
errors: No known data errors
```

Upgrading evaluation license

You can start using NexentaStor Evaluation Product, and some time later decide to upgrade it to a commercial version, with better support options, extra features, and fewer (or none at all) usage restrictions. The following assumes that the appliance is already registered with Evaluation License key, as described in the Quick Start Guide document. Upgrading existing appliance's Evaluation license is definitely possible.

The steps are:

1) Find out your appliance's **`Machine Signature`**. To find out your appliance's **`Machine Signature`**, in Nexenta Management View, go to Help => About, for instance:

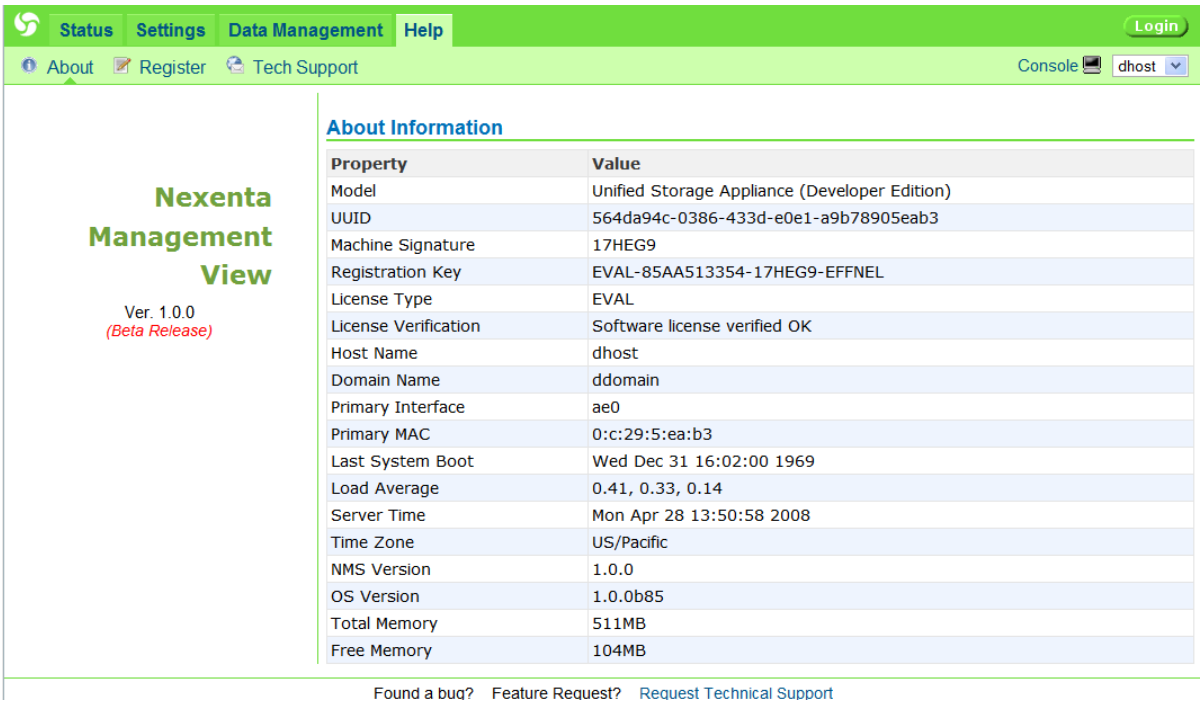

Alternatively, use NMC command:

```
nmc:/$ show appliance license
```

2) Contact Nexenta Systems, Inc. sales at sales@nexenta.com. Make sure to include your **`Machine Signature`** in the correspondence.
3) In response, Nexenta Systems will send you a pair of a new license key and a new licensing module. Copy the latter onto the appliance.
4) Finally, in the Nexenta Management Console, run:

```
nmc:/$ setup appliance register -u <licensing module>
```

For more information on registering (and re-registering) appliance, see '**setup appliance register -h**' manual page.

```
nmc:/$ setup appliance register -h
```

The process of upgrading product's license is safe – it is protected by NexentaStor's system checkpoints/transactional upgrades functionality. If anything goes wrong during the process (e. g., bad network connection, invalid license key) – the appliance will automatically rollback to the previous known good state.

Following is a detailed example:

```
nmc:/$ scp user@remotehost:~/nlm-com.deb .

nmc:/$ setup appliance register -u ./nlm-com.deb

Created rollback checkpoint 'rootfs-nmu-002'
New licensing module successfully installed!
Trying to gain exclusive access to the appliance.
This operation might take up to 30 seconds. Please wait...
Exclusive access granted.
This operation will make Nexenta Management Server temporarily unavailable.
Proceed?  Yes
Restarting NMS. Please wait...
NMS was successfully restarted

nmc:/$ show appliance license
Model             : Unified Storage Appliance (Enterprise Edition)
Machine Signature : 1G9I3C
License           : ENS4-845E2F1819-1G9I3C-CKBHKQ
Key Type          : ENS4
Days Left         : 365
```

Pluggable Modules – Runners

For a background on NexentaStor Fault and Performance Management, Reporting and Indexing facilities, see Sections "Terminology", "Indexing NexentaStor Archives", "Fault Management".

To show the existing "runners", run:

```
nmc:/$ show appliance runners
```

or, with more man pages:

```
nmc:/$ show appliance runners -h
```

Here's a sample output:

```
nmc:/$ show appliance runners
RUNNER                  STATUS          STATE           SCHEDULED
memory-check            enabled         ready           every 12 minutes
runners-check           enabled         ready           every 12 hours
cpu-utilization-check   enabled         ready           every 15 minutes
nms-check               enabled         ready           not schedulable
services-check          enabled         ready           every 15 minutes
volume-check            enabled         ready           every 10 minutes
hosts-check             enabled         running         hourly
nms-fmacheck            enabled         ready           not schedulable
network-collector       enabled         ready           hourly
nfs-collector           enabled         ready           hourly
volume-collector        enabled         ready           hourly
volume-reporter         enabled         ready           weekly on Sat 04:00am
services-reporter       enabled         ready           daily at  05:00am
nfs-reporter            enabled         ready           weekly on Sat 03:00am
network-reporter        enabled         ready           weekly on Sat 02:00am
indexer:vtest/a         enabled         ready           daily at  01:00am
indexer:vtest/b         enabled         ready           daily at  01:00am
```

This shows:

a) several fault triggers (all with extension "check", Section "Fault Management"), followed by
b) statistic collectors, followed by
c) storage and network service reporters, followed by
d) two specific indexers with their associated folders (Section "Indexing NexentaStor Archives").

The corresponding NMV function is located under Data Management/Runners, and is illustrated below:

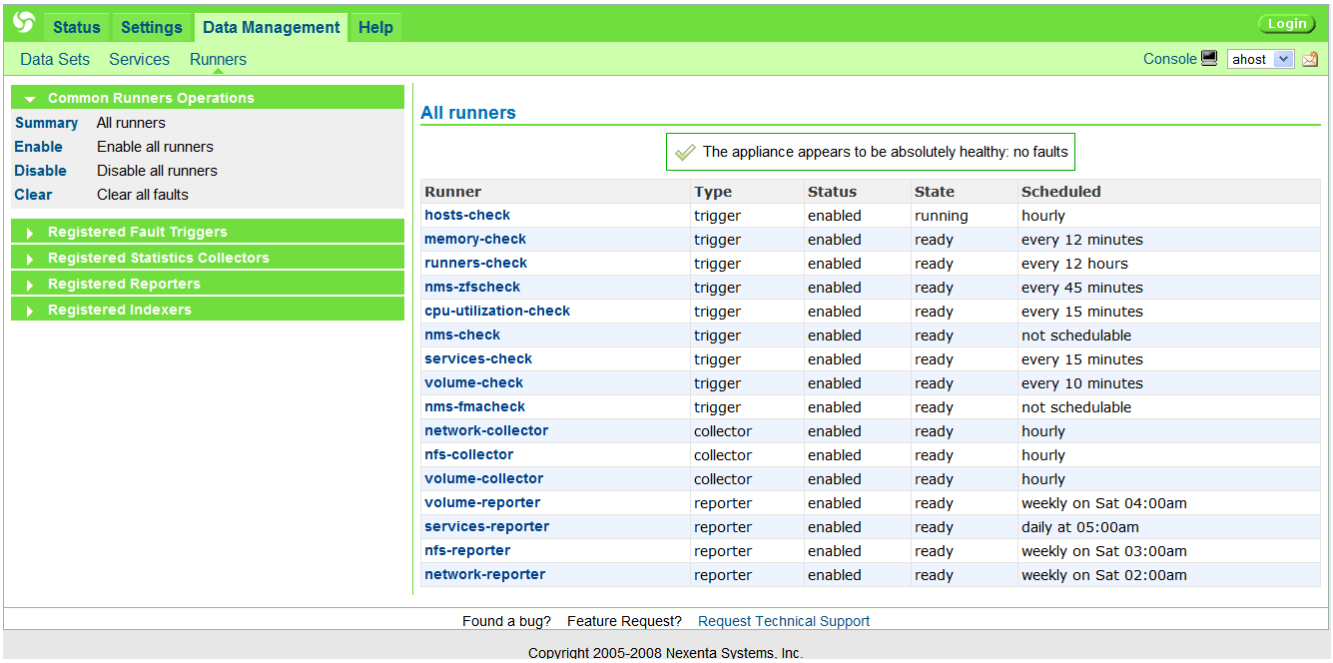

The appliance's framework allows to quickly deploy additional runners. NexentaStor runners have the advantage of exercising the entire NMS-provided SA-API (Section "Terminology"; see also [3], Section "References"), execute periodically, and/or on event, and/or run constantly in the background.

NexentaStor runners rely on the appliance's mailing facility – please do not forget to configure the mailing:

```
nmc:/$ setup appliance mailer
```

To show the mailer configuration:

```
nmc:/$ show appliance mailer
```

All appliance's runners are runtime-configurable. Runners' times-to-run and other properties can be changed via:

```
nmc:/$ setup trigger

nmc:/$ setup collector

nmc:/$ setup reporter

nmc:/$ setup indexer
```

As always, each of the **setup** commands listed above has its **show** "counterpart", to show the existing configuration and runtime status:

```
nmc@myhost:/$ show trigger

nmc@myhost:/$ show collector

nmc@myhost:/$ show reporter

nmc@myhost:/$ show indexer
```

For instance:

```
nmc:/$ setup trigger cpu-utilization
```

This can be used to disable, enable, run, and configure standard fault trigger that monitors CPU utilization. For instance, press TAB-TAB or Enter, type or select '**property**', and view all 'cpu-utilization' properties available for tuning. You could change the alarm-generating thresholds (in this case - low and critically low idle CPU), make it run more or less frequent, etc.

```
nmc:/$ show trigger cpu-utilization -v
```

This will show the trigger's current runtime state, status and existing configuration in detail (notice the verbose -v option).

Fault Management[45]

Part of the NexentaStor Fault Management facility is realized through Fault Triggers. A fault trigger, or simply, a trigger, is a special kind of a pluggable runner module ('**help runners**') that performs a certain fault management and monitoring operation(s). Each trigger monitors one, or a few related conditions.

If any of the monitored conditions are violated, a fault trigger raises an alarm, which manifests itself in several ways:

1) email notification to the administrator, with detailed description of the fault, including: severity, time, scope, suggested troubleshooting action, and often an excerpt of a related log with details.

2) red color showing up via of the following NMC 'show' operations:
   (a) show trigger all-faults

   (b) show trigger <name>

   (c) show appliance runners

   (d) show faults all-appliances

Notifications of hardware faults are *immediate.* Unlike many other potentially faulty conditions that are getting periodically "polled", any hardware fault itself triggers the appliance's fault management logic, that in turn includes email notification.

45. For the fault management functionality provided by the NexentaStor web GUI, see Section "NMV Configuration and Fault Management".

To see all available fault triggers, run:

```
nmc:/$ show trigger all
```

or simply run '**show trigger**' and select one of the available completion options.

In all cases a trigger that "carries" the alarm will be shown in red, assuming NMC colors are enabled.

In addition, the faulted trigger will try to notify system administrator via appliance's mailing facility. Therefore, as already noted elsewhere, it is important to setup the appliance's mailer:

```
nmc:/$ setup appliance mailer
```

Trigger counts the fault conditions every time it runs. Typically, the trigger will send email once the faulty condition is observed a certain configurable number of times. Typically, after that the trigger itself goes into 'maintenance' state - it will still run and count the faulty conditions but it will not send email notification anymore - that is, until system administrator clears it from its maintenance state:

```
nmc:/$ setup trigger <name> clear-faults
```

Similar to the rest appliance's runners, triggers are flexible, in terms of their runtime behavior and trigger-specific conditions they monitor. For details on any specific fault trigger, run:

```
nmc:/$ show trigger <name> -v
```

where <name> stands for the trigger's name, and -v (verbose) is used to display details

The appliance includes one special fault trigger – 'nms-check'. This trigger performs fault management/monitoring function for the Fault Management facility itself. Nms-check tracks NMS connectivity failures and internal errors.

```
nmc:/$ show trigger nms-check -v
```

In presence of network failures, this will show all alarms (in detail) that the appliance failed to report.

```
nmc:/$ show faults all-applinaces
```

This generates Fault Management summary report that includes all known (explicitly ssh-bound and dynamically discovered) Nexenta appliances.

Upon generating the summary, use a combination of NMC '**switch**' operation () (Section "Command Reference - switch") and '**show faults**' - to "zoom-in" into a particular ("faulted") appliance for details.

The next section describes an uncorrectable I/O error scenario and provides troubleshooting tips.

Recovering from uncorrectable I/O errors

Hardware failures are statistically inevitable; it is only prudent to have a plan. Part of such plan could be: prevention. It is always better (and much less expensive!) to prevent uncorrectable hardware failures in the first place. The appliance's underlying filesystem – ZFS – implements state of the art protection and self-healing mechanism. In particular, the system is capable of utilizing disk redundancy – at runtime, and without loosing serviceability.

However, the redundancy on the volume level must be administratively provided. Quoting "ZFS Best Practices Guide":

For all production environments, configure ZFS so that it can repair data inconsistencies. Use ZFS redundancy such as raidz, raidz2, mirror, or copies > 1, regardless of the RAID level implemented on the underlying storage device.

To that end, see also A Note On Redundant Configurations, Section "Creating a volume". Rest of this Section presents a case of an unrecoverable I/O error[46], with commentary and troubleshooting tips. Here's an example:

46. which certainly has a "better" chance to occur with non-redundant configurations..

```
FMA EVENT: ======= START =======
FMA EVENT: SOURCE: zfs-diagnosis
FMA EVENT: PROBLEM-IN: zfs://pool=iscsidisk/vdev=56f1d7932d4c039d
FMA EVENT: AFFECTS: zfs://pool=iscsidisk/vdev=56f1d7932d4c039d
FMA EVENT: ======== END ========
fault trigger 'nms-fmacheck (E1)' reached the configured maximum
of 1 failure

FAULT:
*************************************************************
 FAULT: Appliance    : ups-nxstor1 (OS v0.99.5b82, NMS v0.99.5)
 FAULT: Machine SIG  : 1CG5KI
 FAULT: Primary MAC  : 0:15:17:a:d1:fc
 FAULT: Time         : Tue Mar 11 14:10:47 2008
 FAULT: Trigger      : nms-fmacheck
 FAULT: Fault Type   : ALARM
 FAULT: Fault ID     : 1
 FAULT: Fault Count  : 1
 FAULT: Severity     : CRITICAL
 FAULT: Description  : FMA Module: zfs-diagnosis, UUID:
 FAULT:              : 5bbb38fb-f518-4aa2-9018-8c2fe7e70360
 FAULT:
*************************************************************

 Fault class : fault.fs.zfs.vdev.io
 Description : The number of I/O errors associated with a ZFS
 device exceeded acceptable levels. Refer to
 http://sun.com/msg/ZFS-8000-FD for more information.
```

Printout above is an example of a fault notification delivered by the appliance's mailer – **if** the mailer is configured (see Sections "5. Initial Configuration Wizard", "Fault Management and Reporting").

An unrecoverable I/O error scenario presents only two options:

(1) Manually recover the faulted device. As specified in the fault report, it make sense to review the posted URL (http://sun.com/msg/ZFS-8000-FD in this case) for the latest tips and guidelines.

In case of FC/iSCSI/USB attached drives, please verify connectivity to the corresponding target(s).

Next, ssh into the appliance as root. At this point NMC will automatically determine the presence of faulted condition and will prompt you to execute correcting action (you will simply need to press Enter).

(2) The second option is simple: power cycle the appliance. This may cause an unrecoverable loss of data: the in-flight data that was not committed to stable storage at the time of the hardware failure will be lost. However, the existing data on the affected volume will **not** be corrupted.

After power cycle the entire faulted volume (that is, the volume that contains the faulted drive) will be marked 'offline' and unaccessible.

Search Engine and NexentaStor Indexers

NexentaStor provides built-in Search Engine capability, realized via a special kind of NexentaStor runners – indexers[47]. For an introduction, background, features and overview of this facility, please see Section "Indexing NexentaStor Archives".

Being specialized NexentaStor runners, indexers do support all the generic conventions applicable to other runners. That is, as far as management console is concerned, the descriptions and guidelines outlines in the Section "Pluggable Modules – Runners" above do apply.

For details on any specific indexer, run:

```
nmc:/$ show indexer <name>
```

As always, you can simply type 'show indexer', press TAB-TAB, and make a selection.

Unlike all the rest NexentaStor runners, indexers are dynamically "created" and "destroyed". To create an indexer means to associate the latter with a storage folder (for the subsequent indexing of this folder snapshots and the most recent content). Symmetrically, to destroy an indexer means to disassociate it from the corresponding folder, and in addition, to remove from storage all the corresponding indexing data.

```
nmc:/$ create indexer
```

Similar to the rest appliance's runners, indexers are flexible and fully customizable, at runtime. To configure a specific indexer, run:

```
nmc:/$ setup indexer <name> property
```

Finally, it is possible to use any given indexer to perform a search – via NMC. This capability cannot be considered very user-friendly and is provided only as a quick "backup" for the corresponding Nexenta Management View capability (Sections "Indexing NexentaStor Archives", ""). To search an archive, run:

```
nmc:/$ setup indexer <name> search <expression>
```

Managing LUNs

NexentaStor provides services and tools to administer physical and logical disks, commonly called LUNs (for the term "LUN", see Section "Terminology"). The following NMV screenshot is an example that shows appliance's LUNs and I/O statistics:

47. "Indexer" as a term is defined in Section "Terminology".

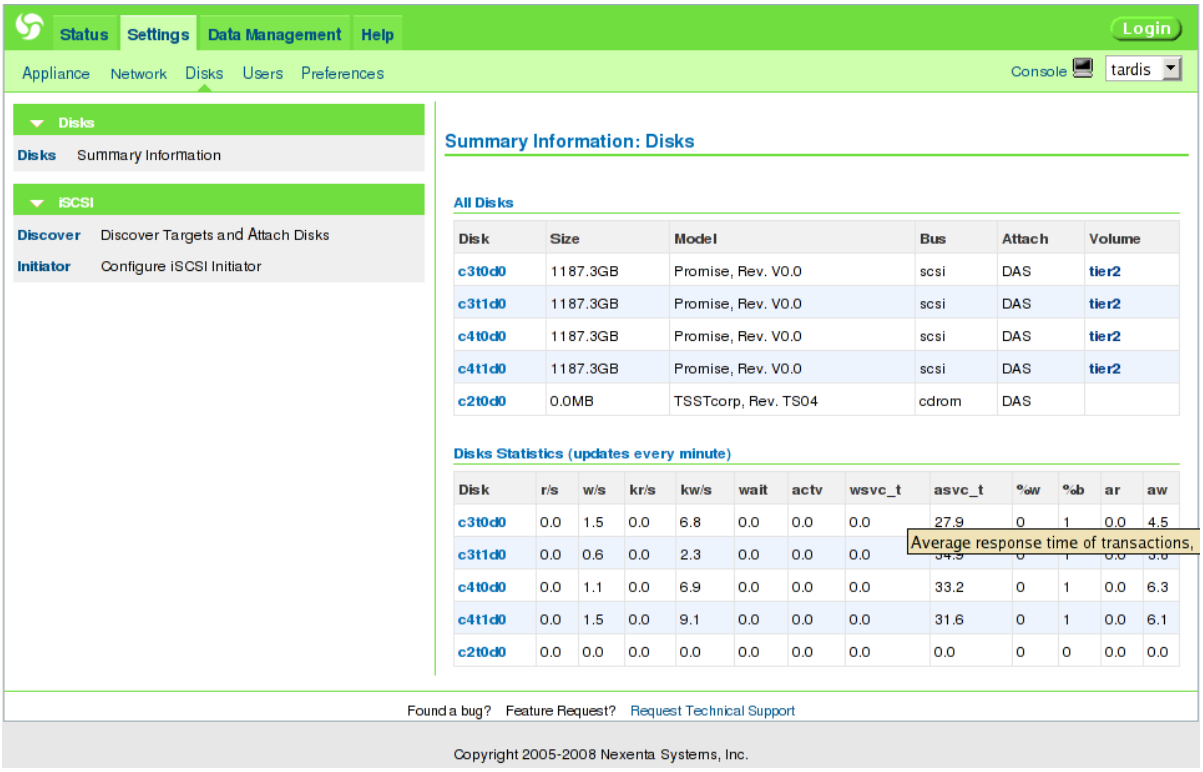

From this NMV page you can also discover iSCSI targets and provision new iSCSI attachments.

Appliance's management console provides a similar functionality, and more. Most importantly, faulted drive can be administered without losing the service. In NMC, the set of related commands includes:

• **attach-lun**	Acts on an existing mirror, or creates a new mirror. Attaching to a non-mirrored device automatically transforms the pair into a two-way mirror. See NMC's '**setup volume <name> attach-lun -h**' for more information.
• **detach-lun**	Acts on an existing mirror only: detaches device from a mirror.
• **replace-lun**	Can be used with redundant mirror and raidz configurations. Effectively, is a combination of attaching a new device to a redundant configuration, waiting for it to resilver, followed by subsequent detaching of the specified "old" device
• **offline-lun**	Takes the specified LUN offline.
• **online-lun**	Brings the specified LUN online.
• **remove-lun**	Currently only supports removing hot spares and cache devices from a volume. As far as spares and cache devices, the affect of **remove-lun** is opposite of '**setup volume <name> grow**'.

Following are a few usage examples, with comments:

nmc:/$ setup volume vol1 replace-lun -o c0d1 -n c1d0

This will replace IDE drive c0d1 with a new IDE drive c1d0 in the volume 'vol1'.

nmc:/$ setup volume vol1 replace-lun -o c2t0d1 -n c2t0d1

Notice that the old and the new drives have the same name. This will replace a faulted drive c2d0d1 with a new drive. The scenario here is that, after c2t0d1 had faulted, it was removed and another drive was attached to the same controller (c2, in the case) and placed into the same slot.

nmc:/$ setup volume syspool replace-lun -o c0t0d1 -n c0t0d1

Same as above, performed on a system volume. As any other volume, syspool may fail. It is generally advised to always provision system volumes with mirrored or RAIDZ configurations. You could do this at installation time (the NexentaStor Installer provides an option to create syspool out of two mirrored drives), or later, using '**attach-lun**' command (see next). And then, you could always use '**replace-lun**' operation to handle disk faults, with the appliance remaining online and operational.

nmc:/$ setup volume vol1 attach-lun -o c0d1 -n c1d0

This will attach a new drive c1d0 to an existing drive c0d1 in the volume 'vol1'. If the existing drive c0d1 was used as a standalone (non-mirrored) drive, this command will automatically create a mirrored configuration out of (c0d1, c1d0). Alternatively, if the old drive c0d1 was already a part of the mirror, the command will simply add another drive to the existing mirror, thus converting, say, two-way mirror into three-way mirror, three-way into four-way, etc.

As always, you could simply enter '**setup volume vol1 attach-lun**' or '**setup volume**', and then use NMC dialog mode to enter the required parameters. NMC will let you select the drives, with the selection process simplified by NMC-displayed tips and prompts.

For additional information, please refer to Sections "iSCSI", "Lunsync – rediscovering devices", and "Using ZVOLs".

Using LDAP

The Lightweight Directory Access Protocol (LDAP) is a common protocol interface to Network Directory Services. For a quick introduction and background, please see Section "LDAP integration" above.

To make use of Unix based LDAP service, use the following NMC commands:

- **show network service ldap-client**
- **setup network service ldap client**

For instance:

```
nmc@myhost:/$ show network service ldap-client
PROPERTY                    VALUE
info                      : ldap-client
name                      : svc:/network/ldap/client:default
start_pid                 : 14280
state_timestamp           : 11:56:59
start_method_timestamp    : 11:56:59
state                     : online
enabled                   : true
```

```
nmc@myhost:/$ setup network service ldap-client
Base DN                      : dc=cs,dc=stanford,dc=edu
LDAP Servers                 : 172.24.100.246
Authentication Method        : none
Proxy DN                     :
Proxy Password               :
Users Service Descriptor     :
Groups Service Descriptor    :
Netgroup Service Descriptor : cn=netgroup,dc=cs,dc=stanford,dc=edu
```

Once LDAP client is configured (the command '**setup network service ldap-client**' above), you'll be able to list LDAP-managed netgroups (groups of hosts), users, and groups (of users). The corresponding functionality is structured under 'show network': you could simply type '**show network**' and press TAB-TAB. The resulting list will include '**netgroup**' (that is, groups of hosts), '**group**' (that is, groups of users), and '**user**'.

For instance:

```
nmc@myhost:/$ show network netgroup
cadengineering   csengineering    hwengineering
pchr             pccorpexec       ...
```

NMC provides easy to use consistent interface to administer all network service, LDAP including. To configure any network service, type '**setup network service**' and press TAB-TAB. After selecting '**ldap-client**' from the list, press TAB-TAB again. The following list of choices will show up:

```
nmc@myhost:/$ setup network service ldap-client
confcheck    disable      edit-settings enable      restart      show       configure
unconfigure  <?>
```

To unconfigure local LDAP client (and stop using Unix based LDAP service with a given appliance), run:

```
nmc@myhost:/$ setup network service ldap-client unconfigure
```

NexentaStor GUI manager provides a similar functionality. The following NMV screen (under Settings => Network) allows to configure LDAP settings. Notice the similarity with the '**setup network service ldap-client**' command, in terms of configuration parameters:

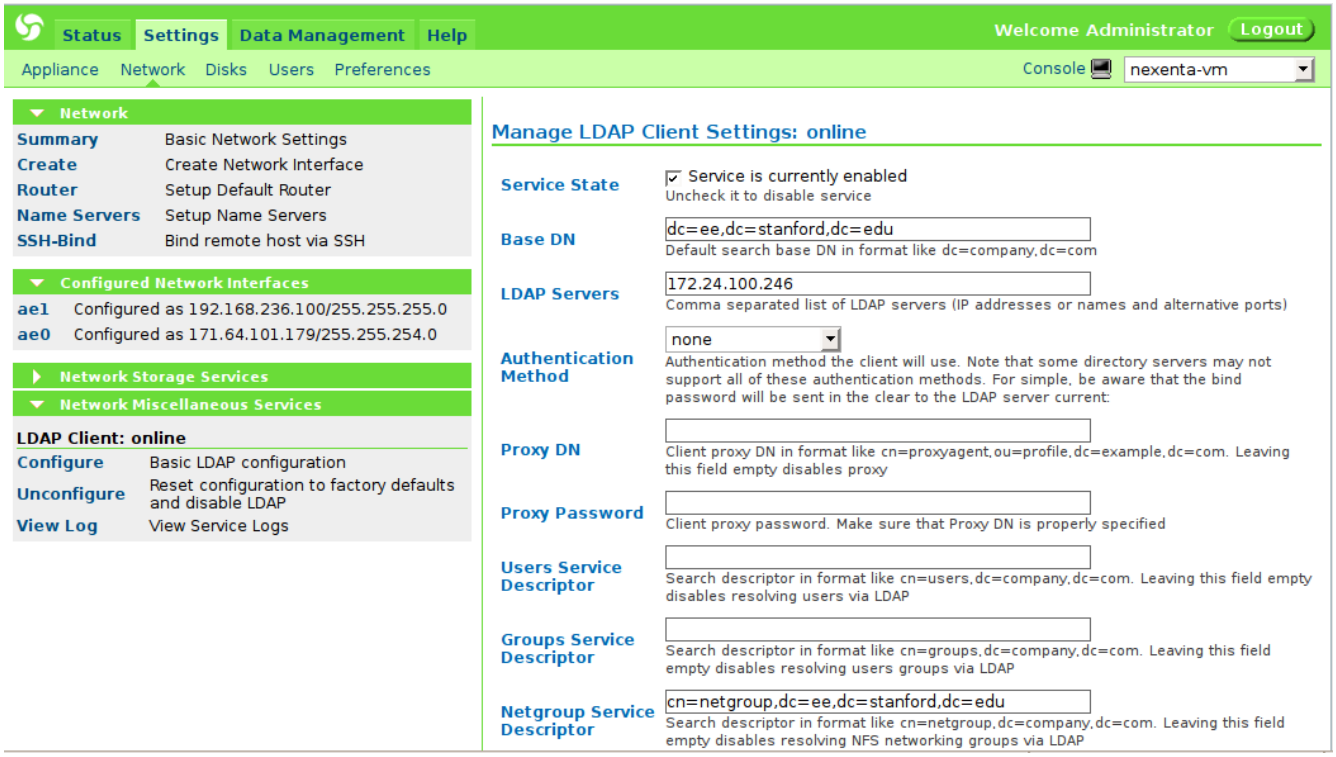

NexentaStor LDAP (integrated) client provides integrated ability to authenticate itself using X.509 certificates. Management console and management UI both provide the corresponding interfaces.

The corresponding NMC man page follows below:

```
nmc:/$ setup network ssl-bind

Usage: [-s hostname:port] [-a alias] [-c certfile]

Create SSL-LDAP "binding" between the appliance and remote LDAP
server. This interface allows to authenticate the appliance with LDAP
server using X.509 certificates.

There are two principal ways to create SSL-LDAP "binding":

1) retrieve SSL certificate from a known LDAP server (Usage #1 below)

2) use a local certificate file (see Usage #2)

Usage #1: [-s hostname:port] [-a alias]
  -s <hostname:port>    LDAP server hostname and port number.
                        Default port: 636
  -a <alias>            Certificate alias. If not specified, hostname
                        will be used
   Example #1:
   nmc@thost:/$ setup network ssl-bind -s ee-od1.mycorp.com:636


Usage #2: [-c certfile] [-a alias]
  -c <certfile>         Trusted remote host's certificate file
  -a <alias>            Certificate alias. Optional - if omitted,
                        the alias will be retrieved from the
                        certificate file
   Example #2:
   nmc:/$ setup network ssl-bind -c 0.cert -a abc
   This example assumes that '0.cert' certificate was previously
   copied (for instance, using NMC 'scp' command).


See also: 'setup network ssl-unbind'
See also: 'show network ssl-bindings'
```

This NMC interface allows to authenticate the appliance (or more exactly, the appliance's LDAP client) with a remote LDAP server using X.509 certificates.

As always, both NMC and NMV provide the corresponding interfaces – to retrieve, use and display X.509 certificates. The following screenshot is an example that shows LDAP client – its NMV

summary page. Notice X.509 certificate details below:

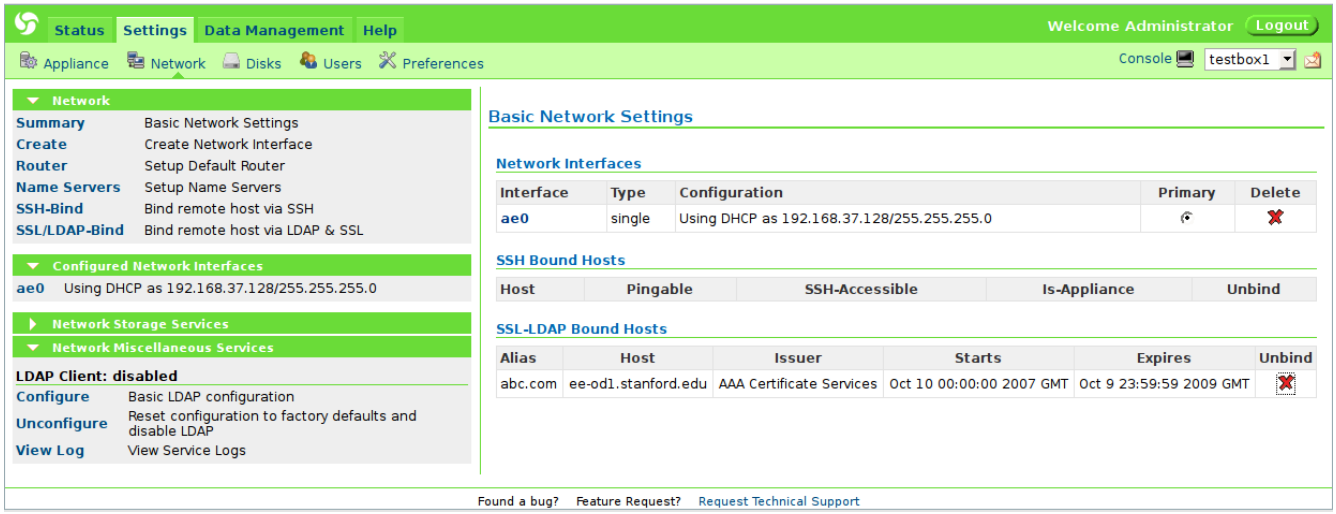

In addition to Unix based LDAP, NexentaStor provides Active Directory integration (next Section).

Using Active Directory

Active Directory (AD[48]) is an implementation of LDAP directory services (Sections "LDAP integration", "Using LDAP") by Microsoft for use primarily in Windows environments. AD purpose is providing central authentication and authorization services for Windows-based computers. In addition, Active Directory supports deploying software and assigning policies on a level of organizations.

To start using AD, you first need to make NexentaStor appliance to become a *member server*. In AD terms that particular operation is often called *join* or *join-ads*.

The second step requires *identity mapping*. Rest of this section illustrates both steps.

To join Active Directory, and subsequently get access to centralized authentication and authorization information, go to NMV's Settings => Network and click on **Join AD/DNS Server**:

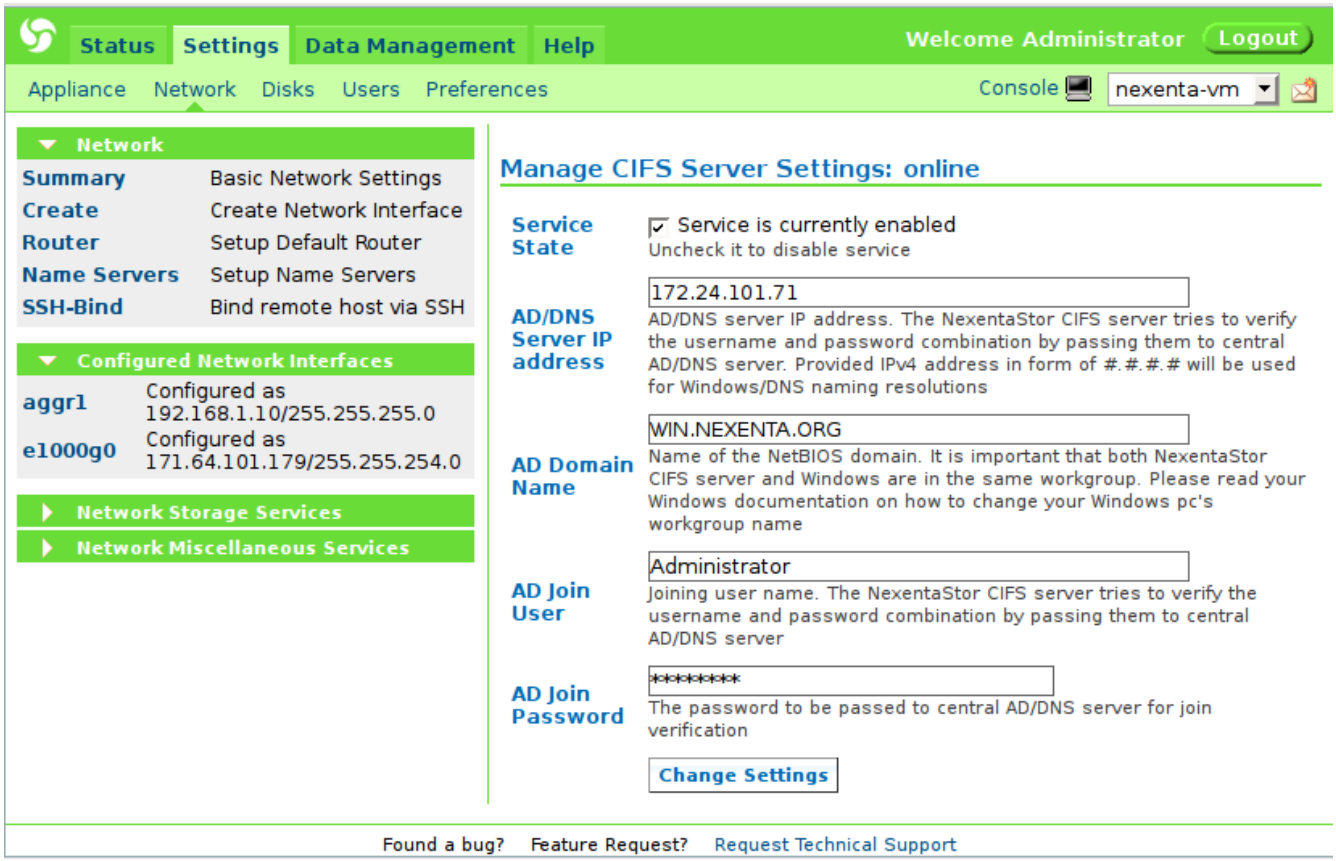

NMC provides a similar functionality, via '**`setup network service cifs-server join-ads`**':

48. http://en.wikipedia.org/wiki/Active_Directory

```
nmc@myhost:/$ setup network service cifs-server join-ads
AD/DNS Server IP address :
--------------------------------------------------------------------------------
AD/DNS server IP address. The NexentaStor CIFS server tries to verify the username
and password combination by passing them to central AD/DNS server. Provided IPv4
address in form of #.#.#.# will be used for Windows/DNS naming resolutions. This
field required to be filled in. Press Ctrl-C to exit.
```

Secondly, as stated in the beginning of this Section, we'd need to establish name-based mappings – that is, name equivalence between Windows users and groups and their counterparts in the UNIX. These mappings persist across reboots. To map Windows users/groups onto UNIX users/groups, go to NMV's Settings => Network and click on the **Identity Mapping** link:

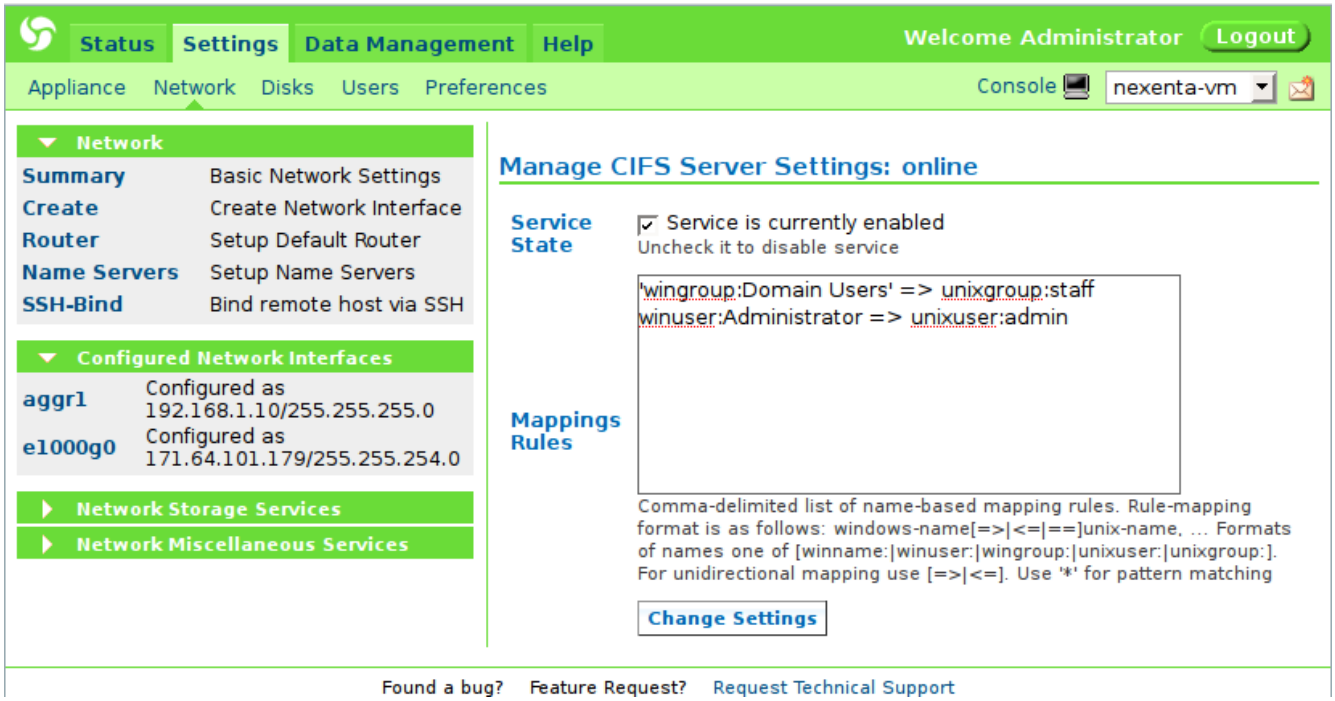

The example above shows two identity mappings: for a group of Windows users called "Domain Users" and for Windows administrator. The latter is mapped on appliance's pre-configured 'admin' account.

NMC provides a similar functionality, via '**setup network service cifs-server idmap**':

```
nmc@myhost:/$ setup network service cifs-server idmap
Mappings Rules :
--------------------------------------------------------------------
Comma-delimited list of name-based mapping rules. Rule-mapping format is as follows:
windows-name[=>|<=|==]unix-name, ... Formats of names one of [winname:|winuser:|
wingroup:|unixuser:|unixgroup:]. For unidirectional mapping use [=>|<=]. Use '*' for
pattern matching. This field required to be filled in. Press Ctrl-C to exit.
```

Keep-Alive

To show all appliances on the network, run:

```
nmc@myhost:/$ show network appliance
```

The list is updated in real time.

For more details, see:

```
nmc@myhost:/$ show network appliance -h
```

This particular keep-alive function is carried out by a special fault trigger called 'hosts-check'. Along with all the rest fault triggers and other pluggable modules (see Section "Fault Management and Reporting"), hosts-check is runtime-configurable and can be tuned for a particular environment (unless the defaults are good - which is expected):

```
nmc@myhost:/$ setup trigger hosts-check
```

Advanced Usage – Quick Reference

This section merely names some of the "advanced" topics and gives the corresponding pointers. This includes:

1) To manage remote appliance :
   (a) switch appliance
   (b) switch appliance -f
   (c) setup network ssh-bind
   (d) show network appliance
2) To manage a group of appliances:
   (a) setup group
   (b) switch group
   (c) show group
3) To record and replay NMC sessions (Section "Command Reference"):
   (a) record <name>
   (b) record view-current
   (c) record stop
   (d) show recording
4) To upgrade or rollback appliance software, manage system checkpoints:
   (a) setup appliance upgrade
   (b) setup appliance syspool
5) To save/restore appliance's configuration:
   (a) setup appliance configuration
   (b) show appliance saved-configs

For in-depth information and usage examples, see Sections:

- "Command Reference"
- "Centralized Management"
- "Note on SSH Binding"
- "Tutorial"
- "Appliance Software Upgrade"

I/O Benchmark[49]

```
nmc@thost:/$ run benchmark

Usage: [-p numprocs] [-b blocksize] [-q] [-s]

   -p <numprocs>     Number of process to run. Default is 2.

   -b <blocksize>    Block size to use. Default is 32k

   -s                No write buffering. fsync() after every write

This benchmark is using well known Bonnie++ tool, it is based on the
Bonnie benchmark written originally by Tim Bray.

Sequential Write (SEQ-WRITE):

    1. Block. The file is created using write(2). The CPU overhead

       should be just the OS file space allocation.

    2. Rewrite. Each <blocksize> of the file is read with read(2),

       dirtied, and rewritten with write(2), requiring an lseek(2).

 Sequential Read (SEQ-READ):

    Block. The file is read using read(2). This should be a very pure

    test of sequential input performance.

Random Seeks (RND-SEEKS):

    This test runs SeekProcCount processes (default 3) in parallel,

    doing a total of 8000 lseek()s to locations in the file specified

    by random(). In each case, the block is read with read(2).

    In 10% of cases, it is dirtied and written back with write(2).

Example:

  nmc@thost:/$ run volume vol1 benchmark -p 2 -b 8192

   Testing 'vol2'. Optimal mode. Using 1022MB files and 8192 blocks.

   SEQ-WRITE CPU    S-REWRITE  CPU    SEQ-READ  CPU    RND-SEEKS

   162MB/s    8%    150MB/s    6%     188MB/s   9%     430/sec

   158MB/s    7%    148MB/s    8%     184MB/s   7%     440/sec

   --------- ----   ---------  ----   --------- ----   ---------

   160MB/s    8%    149MB/s    7%     186MB/s   8%     435/sec
```

49. Distributed as a separate plugin

More NMC Commands and Aliases

In addition to the top-level commands (Section "Top-Level Commands") NMC provides a few more built-in commands, and a set of essential ubiquitous Unix commands.

The extra built-ins are:

```
alias cd df exit history lunsync option
```

For instance, run '**df [folder-name]**' to display volumes and folder, or '**df -h**' for usage.

The available Unix commands include:

```
tar rm date clear mv ls find chmod chown grep

rmdir cat less mkfile mkdir gzip
```

For instance: run '**find**' or '**grep**' to search files and directories, '**edit**' or '**less**' to edit/view files, and so on.

NMC supports a so called "expert mode" operation, one of the differences between "NMC experts" and "non-experts" being that the former are permitted to execute more Unix commands. To enable expert mode, run:

```
nmc:/$ option expert_mode = 1
```

For more information on NMC options, see Section "Preferences", or run:

```
nmc:/$ option -h
```

For the complete and most updated list of NMC commands, including Unix commands, simply press TAB-TAB at NMC prompt.

The following two screenshots show NMC commands.

```
nmc@ahost:/$
alias      date       grep       ls         option     rmdir      switch
cat        destroy    gzip       lunsync    options    run        tar
cd         df         help       man        query      setup      unshare
chmod      exit       history    mkdir      quit       share
chown      fastswitch index      mkfile     record     shell
create     find       less       mv         rm         show
nmc@ahost:/$
```

Screenshot above shows NMC commands with **expert_mode** option switched off. Screenshot below results from TAB-TAB at the prompt with **expert_mode** enabled[50].

50. The set of NMC commands shown above may not necessarily correspond to the latest NexentaStor release.

```
nmc@ahost:/$ option exp
expansion    expert_mode
nmc@ahost:/$ option expert_mode=1
nmc@ahost:/$
alias      date       grep       lunsync    ping       setup      unshare
awk        dd         gzcat      man        query      share      uptime
cal        destroy    gzip       mkdir      quit       show       vi
cat        df         help       mkfile     record     switch     vim
cd         du         history    more       rm         sync       wc
chmod      egrep      joe        mv         rmdir      tail       xargs
chown      exit       less       netstat    rsync      tar
clear      fastswitch ln         option     run        touch
create     find       ls         options    sed        traceroute
nmc@ahost:/$
```

The following are NMC built-in aliases:

```
alias quit='exit'

alias options='option'

alias fastswitch='switch appliance -f'
```

As always, aliases and the commands they alias can be used interchangeably. Run 'alias' to see the current aliases and assign new ones.


```
nmc@ahost:/$ option exp
expansion    expert_mode
nmc@ahost:/$ option expert_mode=1
nmc@ahost:/$
alias      date       grep       lunsync    ping       setup      unshare
awk        dd         gzcat      man        query      share      uptime
cal        destroy    gzip       mkdir      quit       show       vi
cat        df         help       mkfile     record     switch     vim
cd         du         history    more       rm         sync       wc
chmod      egrep      joe        mv         rmdir      tail       xargs
chown      exit       less       netstat    rsync      tar
clear      fastswitch ln         option     run        touch
create     find       ls         options    sed        traceroute
nmc@ahost:/$
```


And a final note on any and all NMC commands:

For manual pages, simply run:

```
nmc:/$ help <command>
```

For instance:

```
nmc:/$ help find
```

Notes on User Management and Access Control

The appliance supports full management of per-user, per-group, per-folder Access Control Lists (ACLs) in its user interface, while also populating the system with accounts and groups that you may have already defined in Active Directory or other LDAP-based directory service. For a background on this integrated functionality, introduction, and screenshot examples, please see Section "User, Group and ACL Management" above.

Following are more screenshots and examples that demonstrate, first, how to edit a folder's extended ACL, and the management GUI capability to create new users and groups that can be later associated with specific access permissions.

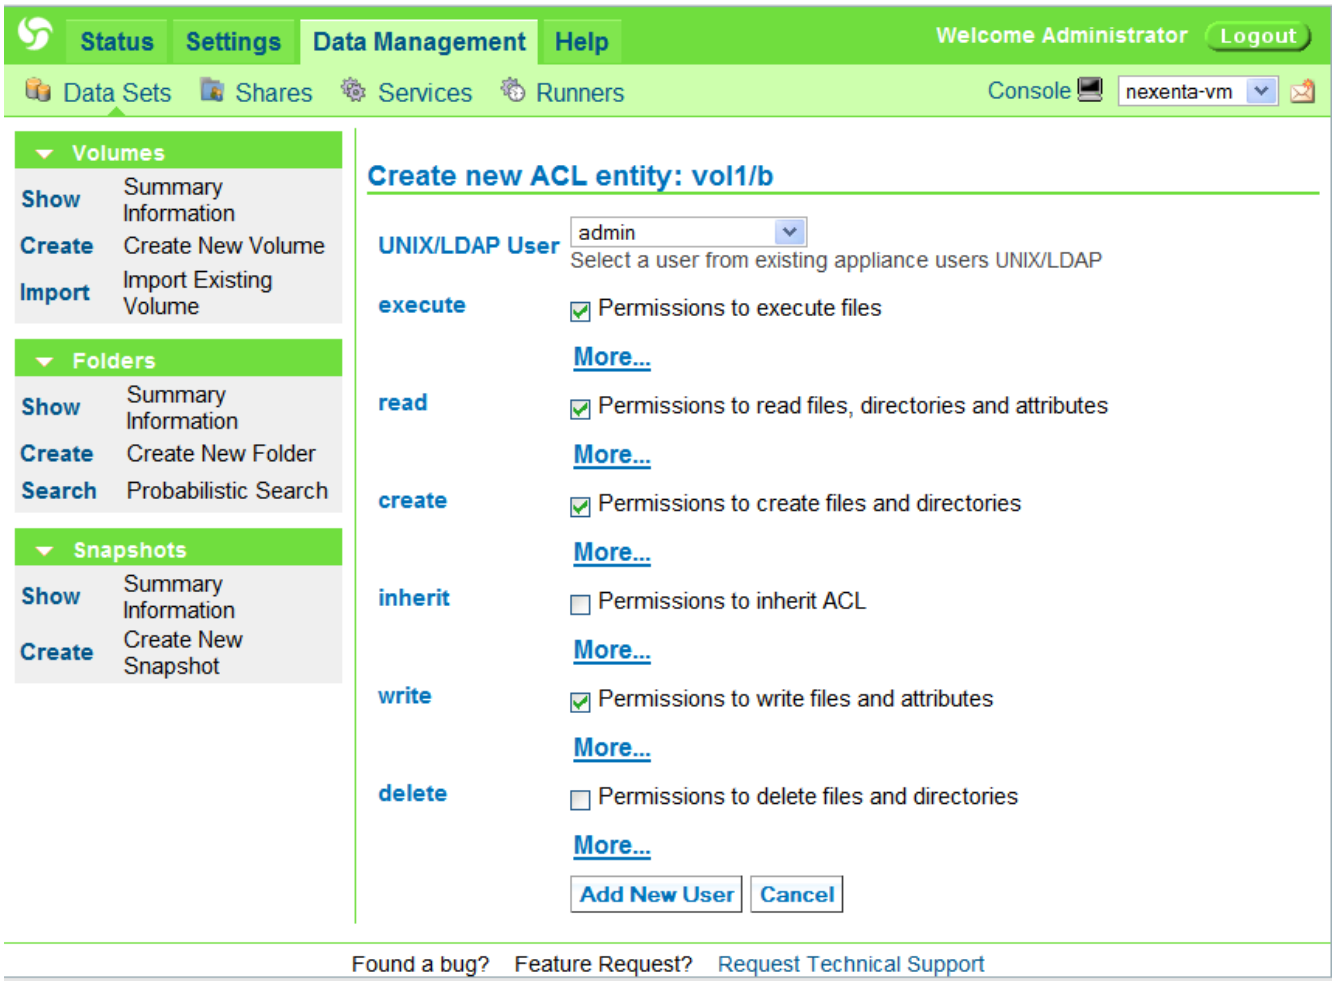

Notice on the picture above that, for convenience, NMV groups all available extended access permissions into the following 6 groups:

execute, read, create, inherit, write, and delete.

The 'read' group for instance, if checked, will automatically include all read permissions: list_directory, read_data, read_acl, etc. For any of the groups of permissions, you can further click on **More** and uncheck some of the corresponding individual permissions.

The corresponding NMC interfaces are:

```
setup folder <name> acl
setup folder <name> ownership
show folder <name> acl
```

See also related NMC commands:

```
setup appliance user
setup appliance usergroup
show appliance user
show appliance usergroup
create appliance user
create appliance usergroup
```

To create a new group of local users, in NMV go to Settings => Users:

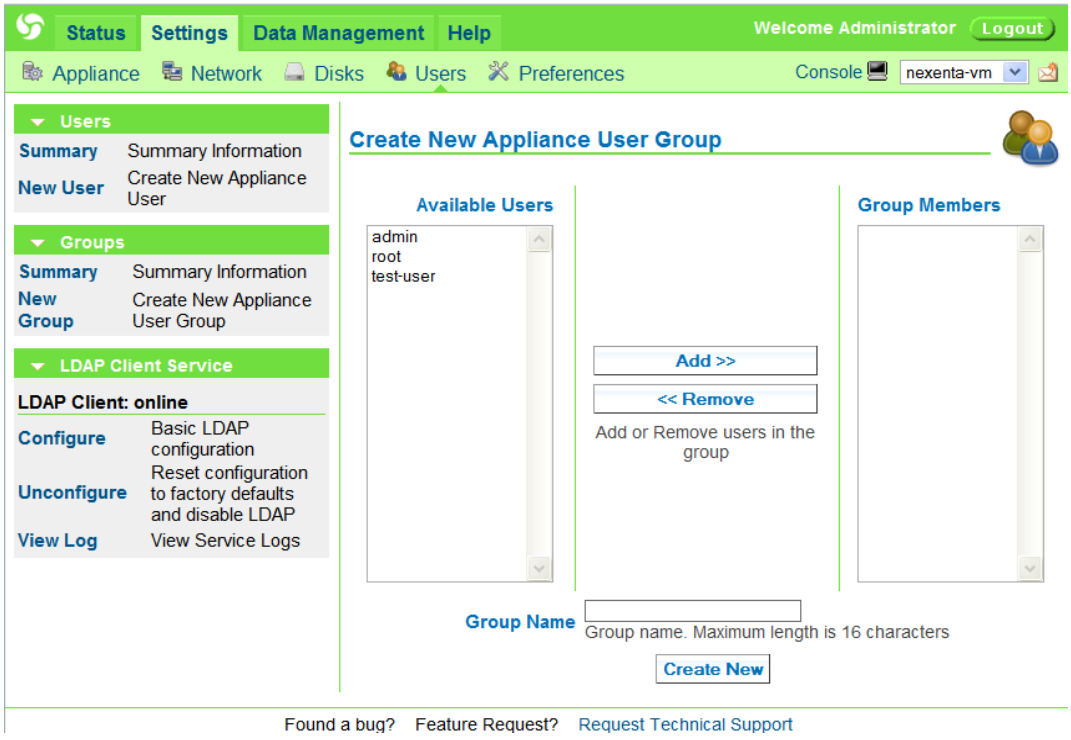

Users, whether they are locally or centrally defined (via LDAP/AD), can be further associated with not only per-folder access permissions but NMV permissions as well – notice "NMV GUI Permissions" on the picture below:

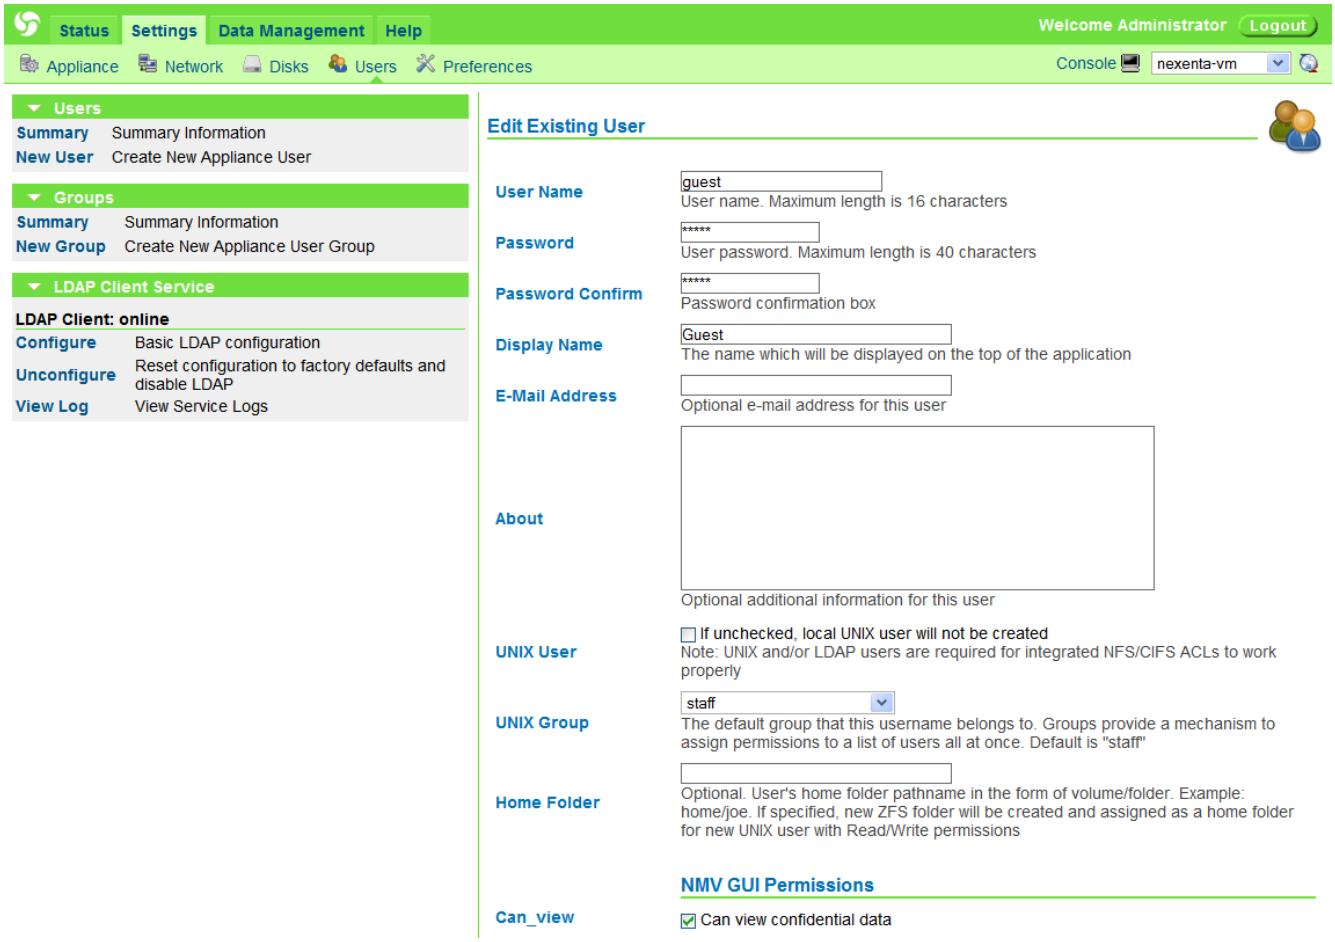

NMC Preferences

For current NMC settings, run:

```
nmc:/$ option
```

For manual page, complete description of available options, their settings, and examples:

```
nmc:/$ option -h
```

To change any option value for the current NMC session, type 'option <name>=<value>', for instance:

```
nmc:/$ option bg=dark
```

To make this change persistent, use -s option:

```
nmc:/$ option bg=dark -s
```

This will store the new setting in NMC configuration file. Alternatively:

```
nmc:/$ setup appliance nmc
```

This will open NMC configuration file for editing. See also:

Section "Tips and Guidelines" below (or type '**help tips**')

To configure Nexenta Management Server (NMS):

```
nmc:/$ setup appliance nms
```

Note that changes to the server configuration are global in a sense that the server is shared by all clients, which, in addition to possibly multiple console clients (NMCs), include storage services, fault triggers, reporters and statistic collectors, and custom runners, and certainly, GUI-based web management (NMV) clients.

All those clients listed above access the same universal NMS-provided public API (Section "Storage Appliance API").

See also NMV Section "Server Options" that illustrates similar preferences-view-and-update functionality provided via web interface.

NMC Tips and Guidelines

NMC colors

NMC currently supports the following screen backgrounds:

none	- none (no colors)
1 or 'white'	- white background (default)
2 or 'gray'	- gray background
3 or 'dark'	- dark or black background

If your screen background is dark, run:

```
nmc:/$ option bg=dark
```

To make the pallete change permanent:

```
nmc:/$ setup appliance nmc
```

This will open NMC configuration file in the internal NMC editor.

NMC man pages

The NMC manual pages (a.k.a. command usage guides) are universally available via -h command line option. Typing '?' in the command line typically produces a short summary on the available option. The system manual pages are also available - simply type '**help <option>**'. For instance, run '**help zfs**' to display ZFS manual page. For more information and examples, see Section "Navigation" ('**help navigation**').

NMC editors

VIM is a default NMC text editor - see http://en.wikipedia.org/wiki/Vim_(text_editor)

VIM supports 3 modes of operations:

- browsing/scrolling (Up, Down, PgUp, PgDn, Home, End)
- editing (press 'i', and ESC to get back to browsing)
- command mode (press ':', and ESC to exit the command mode, 'w' to save changes, 'q' to quit)

In addition to VIM, the appliance provides an alternative text editor called 'joe' a.k.a. joe-editor. You can enable its usage instead of the default via '**option alt_editor**'. See Section "Preferences" ('**help options**') for details.

Quitting, exiting, aborting

Ctrl-C	to abort any command, terminate any dialog, terminate NMC Viewer.
Type 'exit' or 'quit'	to exit remote appliance (that is, return back to local appliance), NMC, or group execution mode.
Type 'q'	to quit NMC Viewer, confirmation dialogs.

'show all'

```
nmc:/$ show all
```

This will '**show**' the entire appliance[51], in one shot.

NMC viewer

NMC uses a well-known '**less**' utility to view the files and command output. Once in 'less', use cursor keys to browse, '/' to search forward, '?' to search backward, 'q' to quit. Press 'h' for help and more options, or quit 'less' (with 'q') and run 'help less'.

Man pages

There are two ways to get manual pages, usage examples, and general information on any one of the following commands:

```
show setup

query switch help

destroy create run share unshare record
```

These are:

a) **'help <command>'**

b) **<command> usage**

For instance, to list the entire collection of '**show**' sub-commands with their respective usage & manual pages, run:

```
nmc:/$ show usage
```

or, same:

```
nmc:/$ help show
```

51. More exactly, not the "entire" appliance but substantial part of it anyway..

NMC shortcuts and aliases

NMC is a *power shell* providing many ways to speed up manual operations - and make it easy. This includes various shortcuts and aliases, namely:

1. All NMC operations that deal with folders and volumes can be typed in a shorter form, by omitting '**volume**' and '**folder**' keywords. For instance, the following two 'setup' operations are identical:

```
nmc:/$ setup volume vol1 folder a

nmc:/$ setup vol1/a
```

2. The following is an example of two identical ways to run a service:

```
nmc:/$ run auto-tier :vol1-a-000

nmc:/$ setup auto-tier :vol1-a-000 run
```

By design, **setup** is a "super-command" (Section "Top-Level Commands") that can be used to modify, create, restore, run, and preform other "active" management operations on the appliance. However, a number of the most important "verbs" are defined explicitly, and provided largely for convenience. The set includes: create destroy run share unshare

3. Aliasing

See section "More NMC Commands and Aliases".

Command history

To quickly repeat a command that was already executed, press Ctrl-R and type a few first letters (of the command). Or, simply press Up and Down arrows at the prompt. Type '**history**' for the the list of executed commands.

Batch mode

Management operations can be assembled in batches and run non-interactively. Here's an example:

```
nmc:/$ record abc

nmc:/$ create volume

... (guided by NMC, fill out all the volume-creation options -

     see "Tutorial", Section "Creating a volume")

nmc:/$ record view-current

... optionally, view the currently recorded content

nmc:/$ record stop

... and finally, stop the recording.
```

Now, at this point the entire volume-creation operation is stored in session 'abc' in a fully-expanded form. You could destroy the volume and later re-run (that is, replay) the session:

```
nmc:/$ run recording abc
```

The result will be that the volume is created again, but this time without user interaction.

NMC recorded sessions can contain any number of valid operations and can be used to execute in a group mode (that is, on a group of appliances), and/or on a single appliance, remote or local.

For more information, please see Section "Command Reference – record", or run '**record -h**'.

Finding NMC commands

NMC '**help keyword**' instruction supports yet another way to quickly locate and execute NMC operations.

Nexenta Management Console supports many commands... NMC supports an ever increasing number of commands, and in that sense the '**help keyword**' facility may come handy.

Simply type '**help keyword <some-keyword>**' - and NMC will list all available operations that mention the specified keyword.

For instance:

```
nmc:/$ help keyword checkpoint
```

results in the following printout:

```
* show appliance checkpoint
* setup appliance syspool checkpoint
* setup appliance checkpoint
```

You do not need to remember the exact command line option - a few letters will suffice. For instance:

```
nmc:/$ help keyword checkpo
```

will have the same result as shown above, while

```
nmc:/$ help keyword chec
```

will print an additional command:

```
* show appliance checkpoint
* setup appliance syspool checkpoint
* setup appliance checkpoint
* setup network service <name> confcheck
```

Force operation

The -f (force) flag is used consistently throughout NMC to force operations to proceed in spite of certain exceptional conditions. Few examples:

Example #1:

> The **auto-sync** service expects that its destination is not tampered with between auto-sync runs. If prior to the next scheduled the **auto-sync** determines that the destination is not in-sync anymore, it will fail to run – unless the -f (force) option is specified. See '**create auto-sync -h**' for details.

Example #2:

> Nexenta Management Server will not restart if any of the scheduled services or fault triggers or statistic collectors or reporters are currently executing. More exactly, if you try to restart NMS (via '**setup appliance nms restart**') while it executes anything on behalf of its clients, the restart operation will fail – unless the -f (force) option is specified. See '**setup appliance nms restart -h**' for details.

Locating objects and services

Use **query** to find objects, locate services, and optionally, display their properties. For instance, the quickest way to list all existing objects:

```
nmc:/$ query -e name
```

The **query** allows to execute arbitrary logical expressions. See Section "Command Reference - query", or run '**query -h**' (or '**query usage**') for detailed description and many examples.

iSCSI

On iSCSI initiator side, to discover iSCSI targets and attach (via iSCSI) remote disks and disk arrays, use **show iscsi** and **setup iscsi** commands.

In addition to iSCSI initiator, NMC provides an extremely easy iSCSI target integration. A quick check on the iSCSI target capabilities can be done via NMC 'help keyword' facility (see Section "Finding NMC commands"). The following is an at-a-glance view of the supported iSCSI target capabilities[52]. As always, each of the listed commands has its own manual page available via -h option.

52. Needless to say, 'help keyword' can be used to find any and all NMC commands, not only those listed in this section. See Section "Finding NMC commands" for more information and examples.

```
nmc:/$ help keyword target
* create iscsi target
* show iscsi target
* setup iscsi target
* destroy iscsi target
nmc:/$ help keyword zvol
* create zvol
* create auto-snap zvol
* create volume <name> zvol
* show zvol
* show lun zvol
* show volume <name> zvol
* unshare zvol
* unshare volume <name> zvol
* setup zvol
* setup auto-snap zvol
* setup volume <name> zvol
* run auto-snap zvol
* share zvol
* share volume <name> zvol
* destroy zvol
* destroy auto-snap zvol
* destroy volume <name> zvol
nmc:/$ help keyword tpgt
* create iscsi target tpgt
* show iscsi target tpgt
* setup zvol <name> tpgt
* setup iscsi target tpgt
* setup volume <name> zvol <name> tpgt
* destroy iscsi target tpgt
```

You can convert the Nexenta NAS into a Target in a matter of seconds, thus combining the flexibility and power of both file and block level access to the same storage. Simply, use **`setup zvol`** and **`create zvol`** commands. For more information, run:

```
nmc:/$ create zvol -h
```

For more information and use cases, please see Section "Using ZVOLs".

Lunsync – rediscovering devices

NMC command '**lunsync** ' re-enumerates devices in the appliance. The command is particularly useful in extremely rare cases when (and if) the system does not discover newly attached LUNs (physical or logical disk drives) automatically. To troubleshoot, run:

```
nmc:/$ show lun
```

This will show all physical and logical disks, commonly called LUNs. If a disk is missing in this output, simply run:

```
nmc:/$ lunsync
```

The process of re-enumerating devices may take some time, especially if you deploy remote attachments via iSCSI and/or FC SAN. See also Section "Managing LUNs".

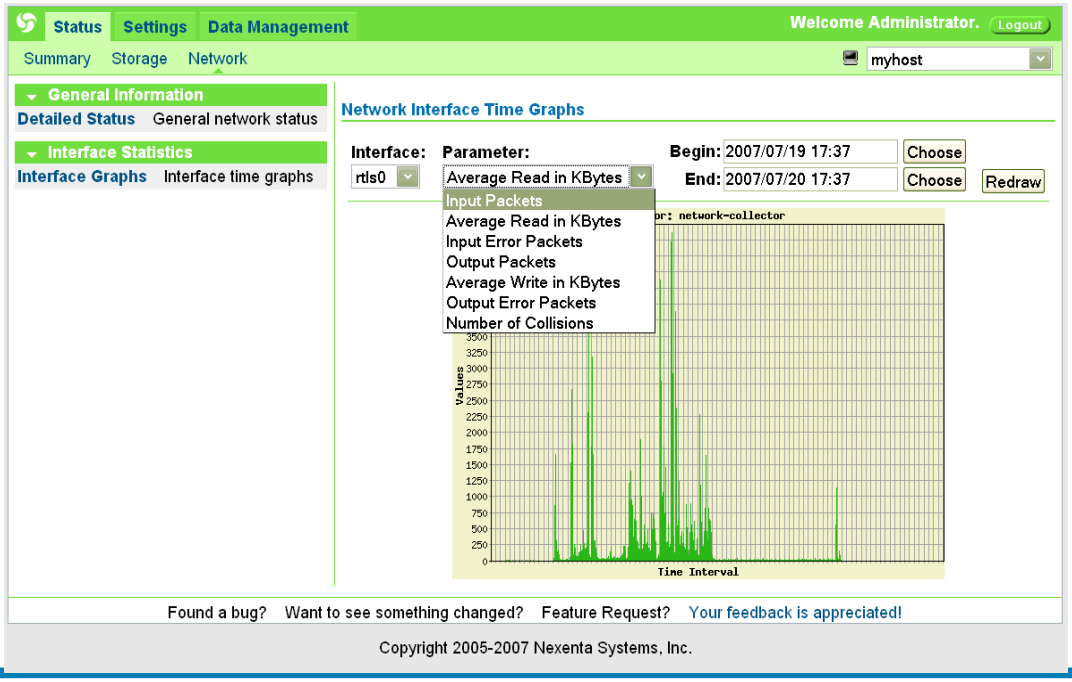

5. Initial Configuration Wizard

The very first time a URL for NexentaStor Initial Configuration Wizard shows up when you register the appliance software:

```
 Model            : Storage Appliance (Digital Archive Edition)
 Software Version : 0.96b70
 Release Date     : Sat Sep  8 20:48:52 2007
 UUID S/N         : 564df0a7-0e34-7214-1b4b-835ac0c95b7f

 Product registration
 ---------------------------------------------------------------------------
Machine Signature : DB8FCL
Registration Key  : EVAL-CDB5220753-DB8FCL-BINMEL


 Review/modify primary network interface
 ---------------------------------------------------------------------------
Interface e1000g0 (PRIMARY): Configured as 192.168.1.1/255.255.255.0
Default Gateway           :
DNS #1                    :
DNS #2                    :
DNS #3                    :

Reconfigure ?  Yes
Option ?
 <?> _dhcp  static  unconfigure
  ------------------------------------------------
  Choose one of the options above, 'q' or Ctrl-C to quit
  Summary info
```

and configure the appliance's primary IP interface:

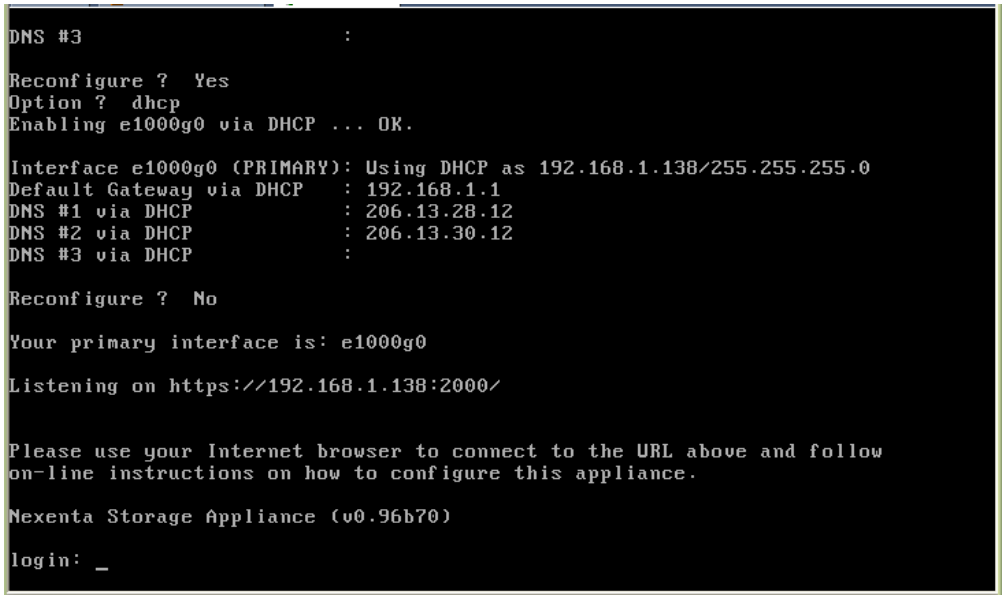

As per instruction on the screen above, you would need to point your Internet browser to the displayed URL (in this case that would be https://192.168.1.138:2000). That will connect your browser to the NexentaStor GUI-based **Initial Configuration Wizard**. The Wizard guides you through the most essential appliance's setup, which is described and illustrated in this section.

Starting from version 1.0.3, there is an added capability to choose transport protocol for NexentaStor Web GUI (NMV). You can select either HTTP – fast and (plain text) unsecured, or (the default) HTTPS – secured but potentially less responsive:

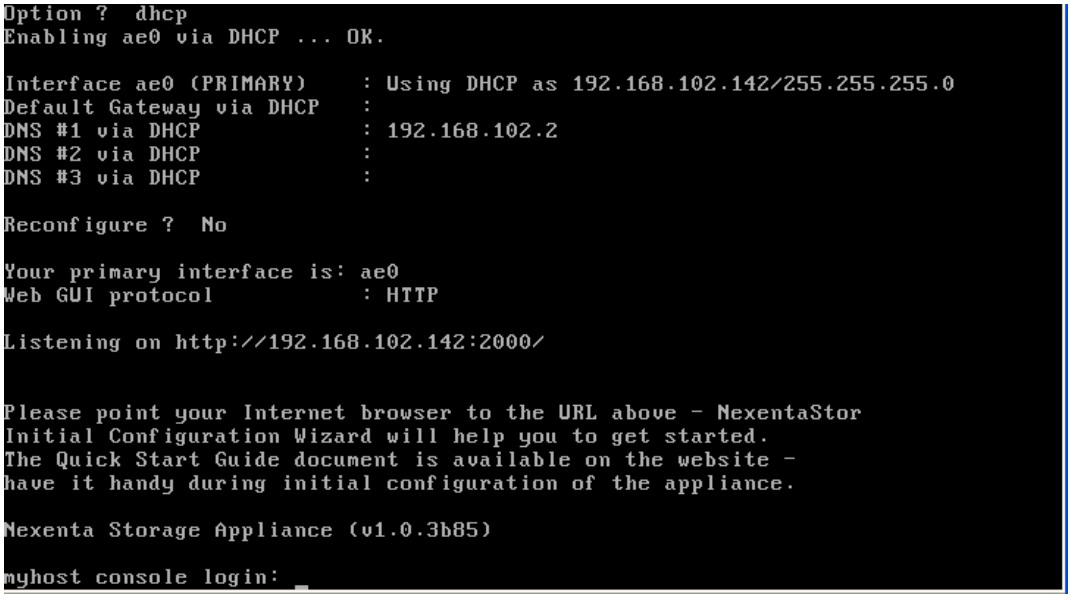

The screenshot shown above is in fact the very first NMC dialog, that can be invoked at any later time via '**setup appliance init**'

NMC '**`setup appliance init`**' functionality includes:

1. setting primary IP address, as shown on the picture above
2. selecting Web GUI (NMV) transport protocol.

The latter option is available starting NexentaStor version 1.0.3. You can choose to run NMV using HTTP unsecured, or HTTPS secured transport. The former is a good choice if you are behind a firewall, and at the same time you feel like NMV response time is slow. Running NMV over plain-text unsecured HTTP often provides for much better (faster) response times.

In any case, the NMV transport setting can be always adjusted and re-adjusted multiple times, via NMC '**`setup appliance init`**' command.

NexentaStor **Initial Configuration Wizard** is subdivided into two guided stages:

- Stage I - Basic Configuration (the first screenshot below)
- Stage II - Configuring Network and Storage

It is important to to point your Internet browser to the displayed URL and perform all Wizard-guided steps of the Basic Configuration.

Stage I – Basic Configuration

The first two screens allow you to set appliance's hostname, domain, localization settings:

and the administrative passwords:

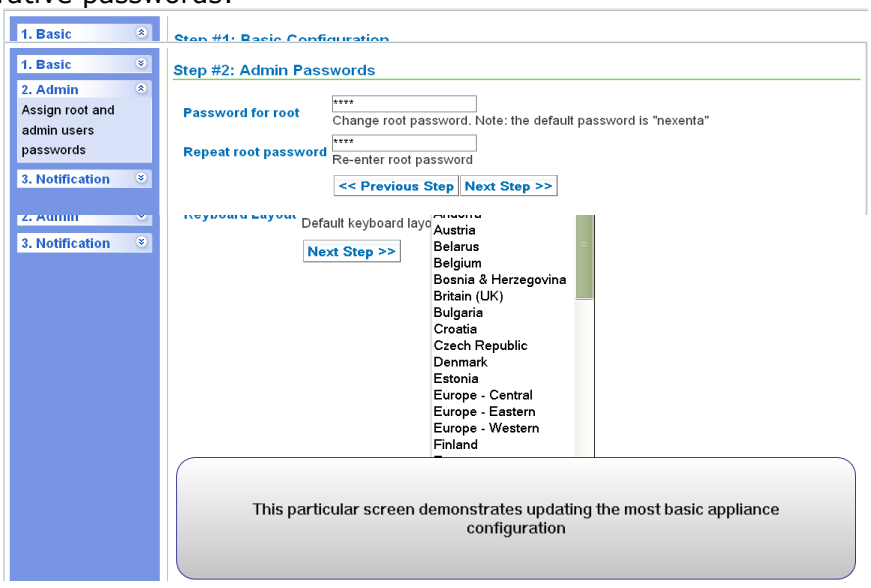

It is important to setup at least a mail server and mailing address to system notifications, reports, and faults to be reported. Configure as necessary for your site.

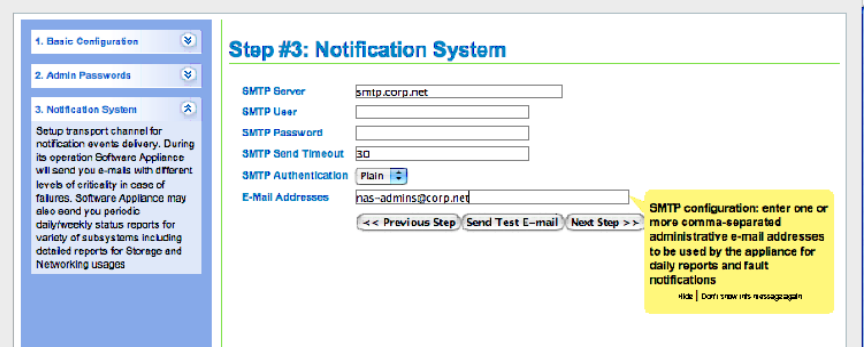

Part of the appliance's Fault Management and Reporting is realized via email notifications. It is strongly recommended to configure appliance's mailer. The appliance is designed to notify of important events – if the mailer is not configured, it will obviously fail, and all the vitally important information will silently end up in the log.

At this point, the current configuration should be applied. You can go back to edit.

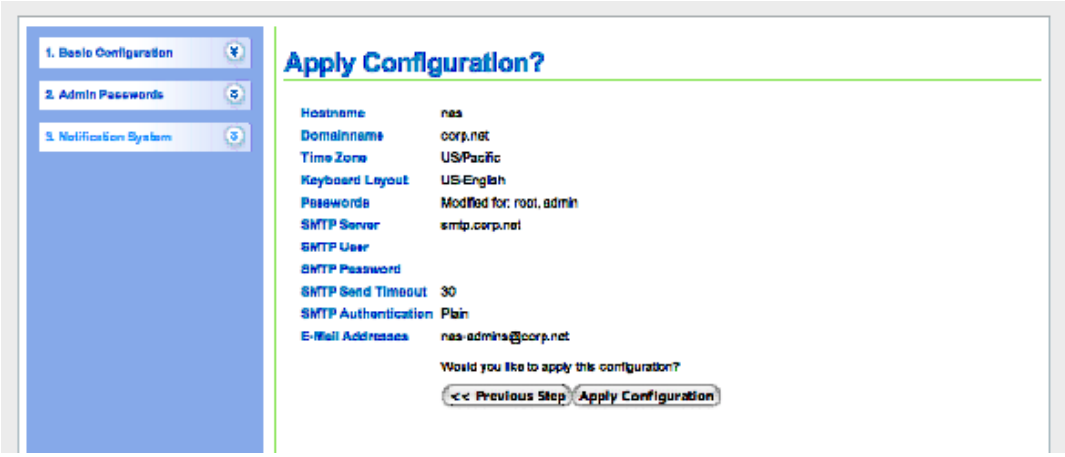

The Stage-I Basic Configuration Wizard is run only once, after the appliance is freshly installed.

The Stage-II Wizard (see next Section) can be run multiple times.

Stage II – Network, Disks and Volumes

The second stage wizard follows, letting you edit your network settings:

... view and configure available disks:

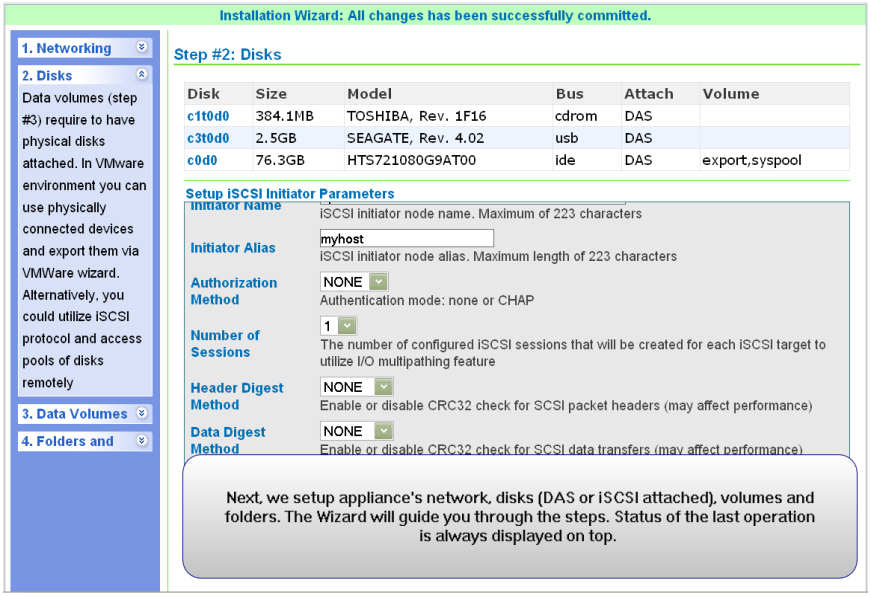

... define data volumes based on that disk storage:

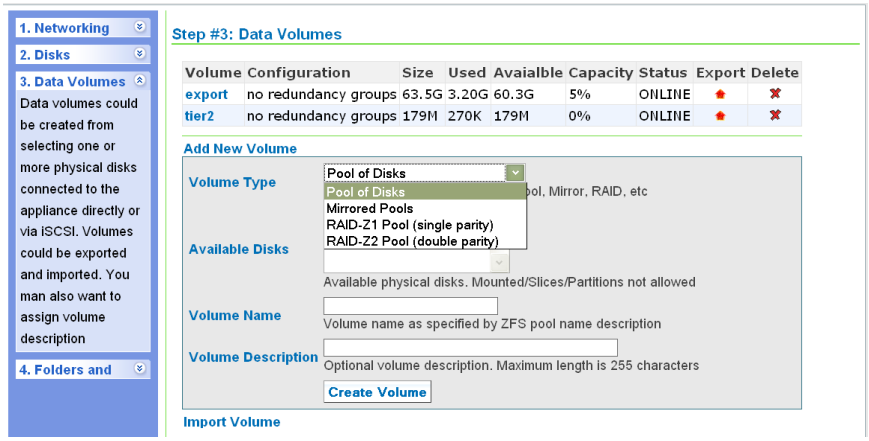

and finally, create folders within those new data volumes:

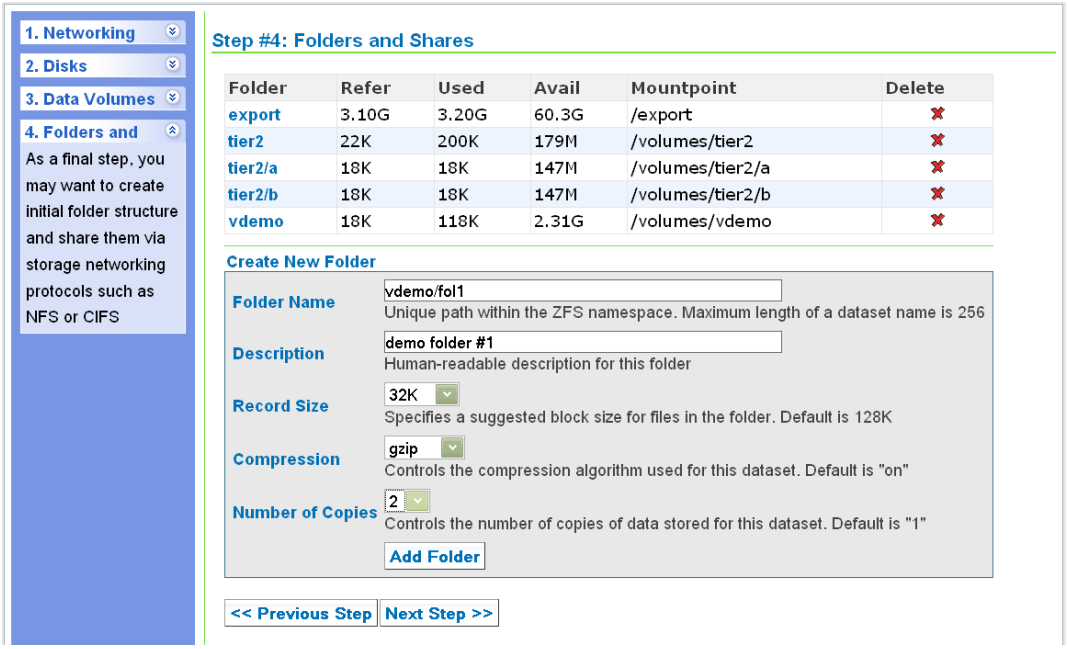

The summary of all second-stage changes can be reviewed...

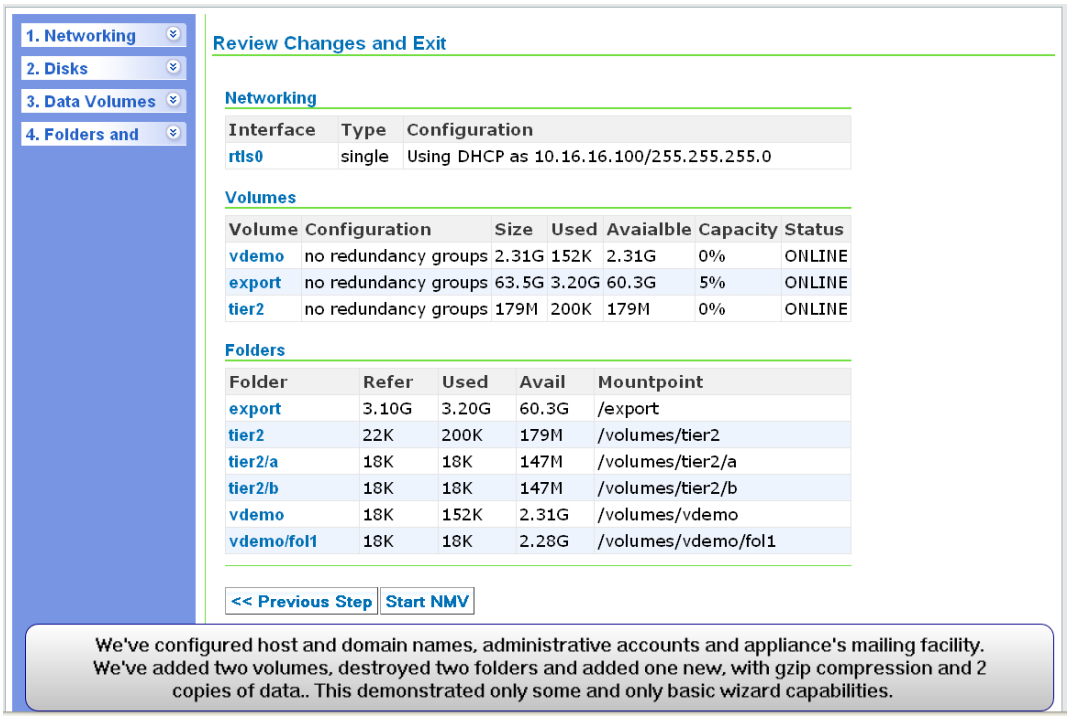

Upon exiting, the Wizard simply runs Nexenta Management View – the primary management GUI interface (next Section).

While "walking" Wizard-guided steps, pay attention to system notices, for example:

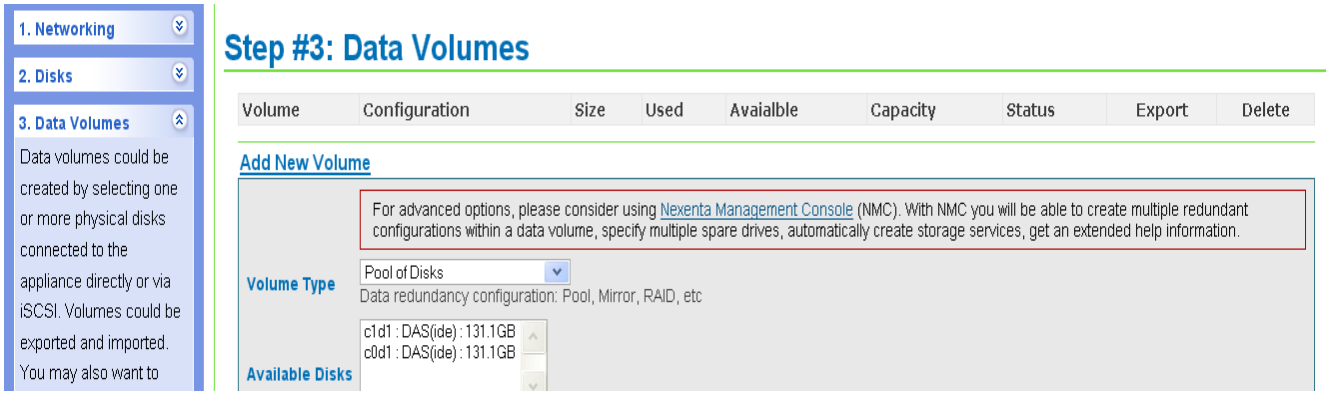

This particular notice recommends using NMC for advanced volume creation options[53]. If required, you could simply click on the provided link and invoke Nexenta Management Console right from the Wizard.

With its last screen the Wizard invites to review all changes, for instance:

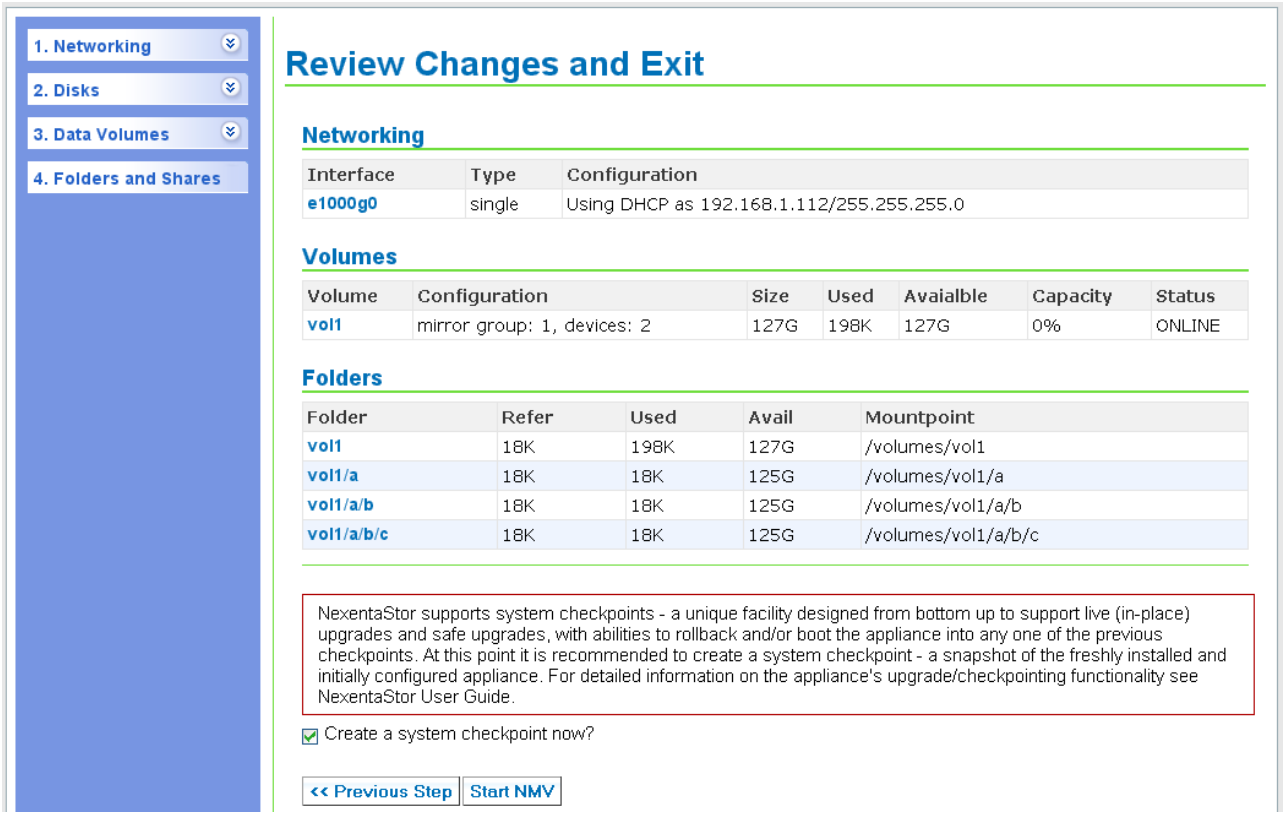

The Wizard will also recommend to create a system checkpoint (see above) – "a snapshot of the freshly installed and initially configured appliance". Please see Sections "Appliance Software Upgrade" and "Tutorial" for detailed information on appliance's upgrade/checkpoint functionality.

53. Some of the screenshots in this section may be outdated.

The Stage-II Installation Wizard can be run multiple times. Even though NMV (Section "6. Management GUI – Quick Introduction" below) provides a superset of functionality, you may still find it convenient to re-run the Wizard.

To do this, point your Internet browser to a the following URL:

<https://appliance's IP address:2000>/wizard2

or:

<https://appliance's DNS hostname>/wizard2

For instance, if the appliance's primary IP is 192.168.16.92, the actual Wizard's URL could look like:

https://192.168.16.92:2000/wizard2

6. Management GUI – Quick Introduction

Nexenta Management View (NMV) provides appliance's web management interface. NMV combines intuitive graphical interface with ease of use, ability to monitor all appliances on the network with ability to easily execute the most essential administrative operations.

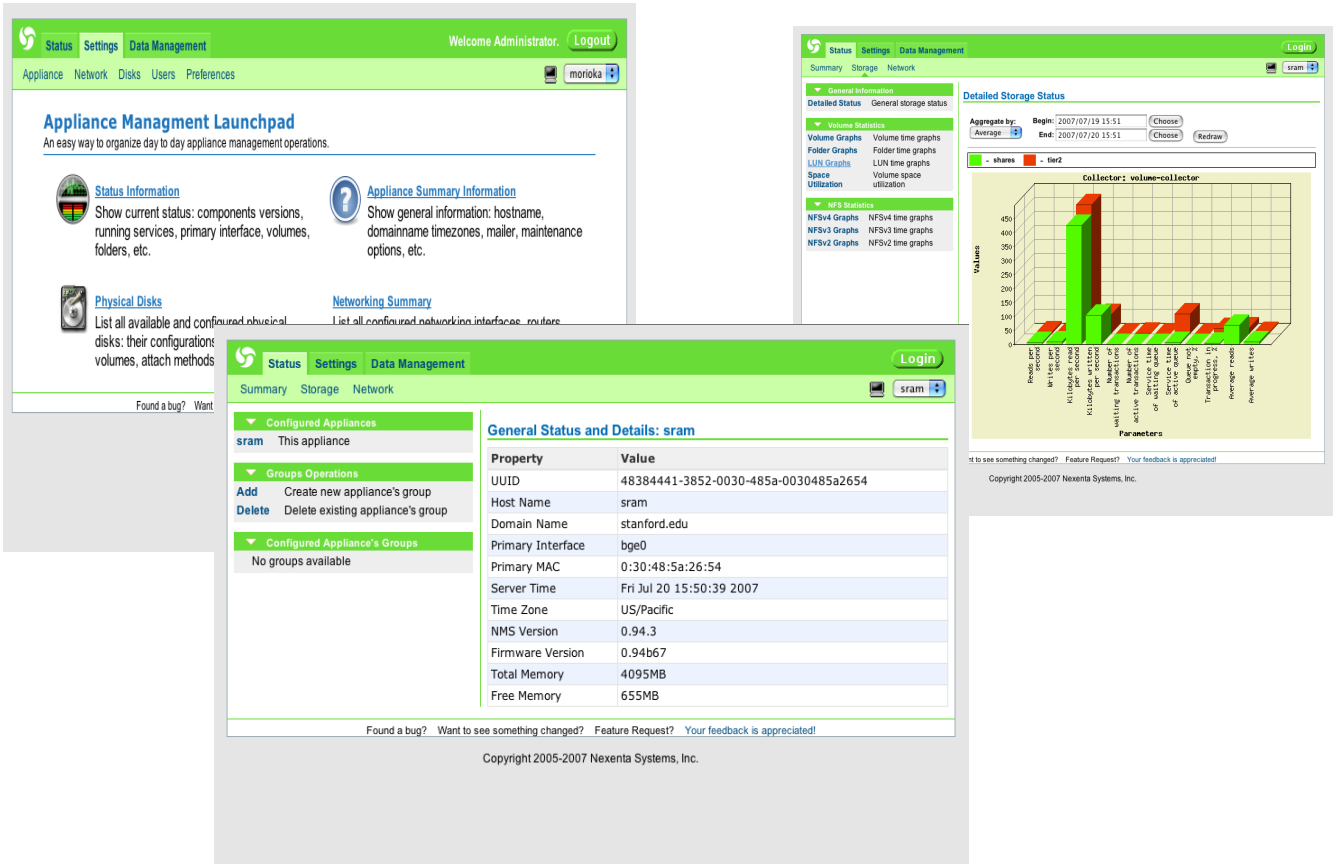

To start using NMV, point your Internet browser to a the following URL:

<https://appliance's IP address:2000>

or

<https://appliance's DNS hostname:2000>

For instance, if the appliance's primary IP is 192.168.16.92, the actual NMV URL could look like:

https://192.168.16.92:2000

This same URL will be displayed by both NexentaStor Wizard (see Section "Initial Configuration Wizard") and NMC – see Quick Start Guide for details ([1]).

Note that all communications with NMV are secure and encrypted over Secure Socket Layer (SSL).

Don't forget to specify the NMV port: **2000**

NMV is a single-login management client[54]. On the screenshots above and in the rest of this section, notice a little list box in the top right corner just below the Login button. That list box at any given moment contains the *current* appliance's hostname. You can switch between appliances by simply selecting another name from the list. The appliance's on the network are statically bound (see "Note on SSH Binding") and/or discovered dynamically – in either case the appliance name will show up in the list box. For similar management console functionality, please see NMC's '**switch**' operation in Section "Command Reference"; for general background please refer to Section "Centralized Management".

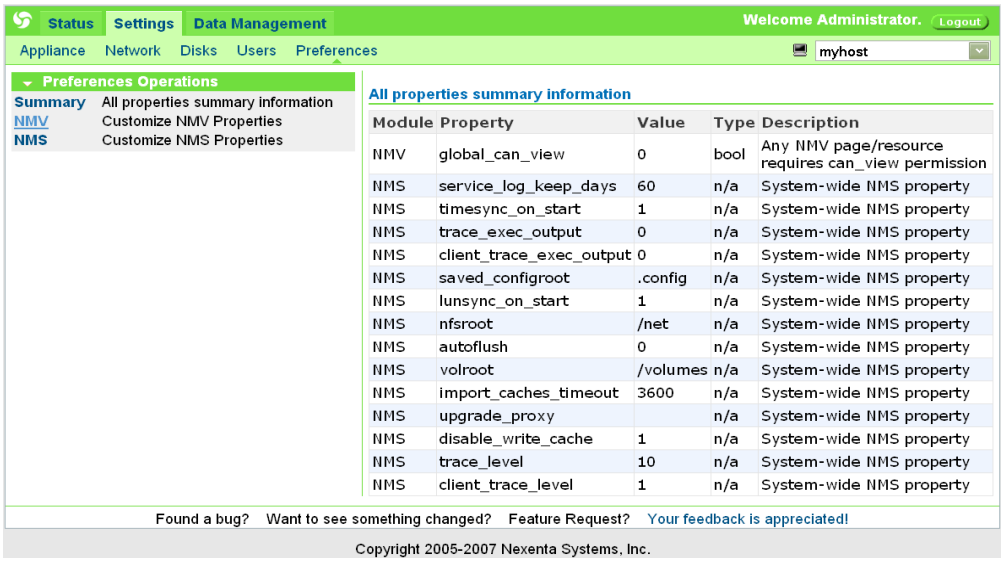

Once the appliance is initially configured (Section "Initial Configuration Wizard"), one of the first things you may want to decide is whether to allow unauthenticated read access to it. The screenshot above shows all configurable preferences[55] – the very first one called '**global_can_view**' determines whether global read access to the appliance is available. Please

54. How to login the very first time? Please see next Section..
55. In your version of NMV this and other screenshots listed in this section may slightly differ, reflecting updated functionality and new features.

pay attention to this particular setting, especially if the appliance's IP address is visible outside your firewall.

NMV is the second available NexentaStor management client, the first one being Nexenta Management Console (Section "NMC – Management Console"). Use NMV for user-friendly self-documented (easy) administration of all Nexenta appliances on the network. Note that you can always run management console from inside the GUI, by simply clicking on a small terminal icon in the top right corner.

NMV tutorials are available on the web, at http://www.nexenta.com/tutorials

Following is the URL:

- Nexenta Management View

NMV Login

Nexenta Management View is pre-configured with administrative super-user account: **admin**. If not logged in, or if the current login privileges are insufficient to perform the operation, NMV will display the following dialog:

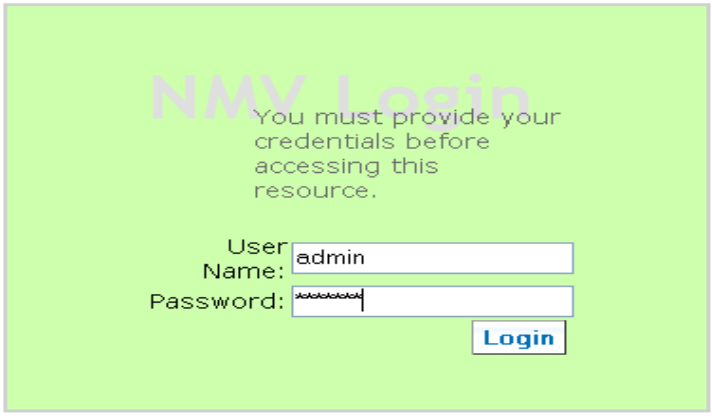

The default password is sent to you in email, along with the product Registration Key. Don't forget to change this password! For more information, please see NexentaStor Quick Start Guide at http://www.nexenta.com/docs.

Note that NMV uses secure SSL connection to communicate to the Internet Browser.

To use plain (non-secure!) HTTP:

1) go to raw shell and open /var/lib/nza/prod.cfg in VIM or editor of your choice
2) comment out lines starting with "server.ssl_"
3) restart nmv via 'svcadm restart nmv' or simply reboot the appliance

Note: using plain HTTP is NOT recommended.

The rest of this section is a brief screenshot based (walkthrough) introduction into Nexenta Management View. Please use NMV to discover its full power!

NMV Summary Page

The Summary pages show appliance at a glance, including the most important (summarized) configuration and fault management information, for instance:

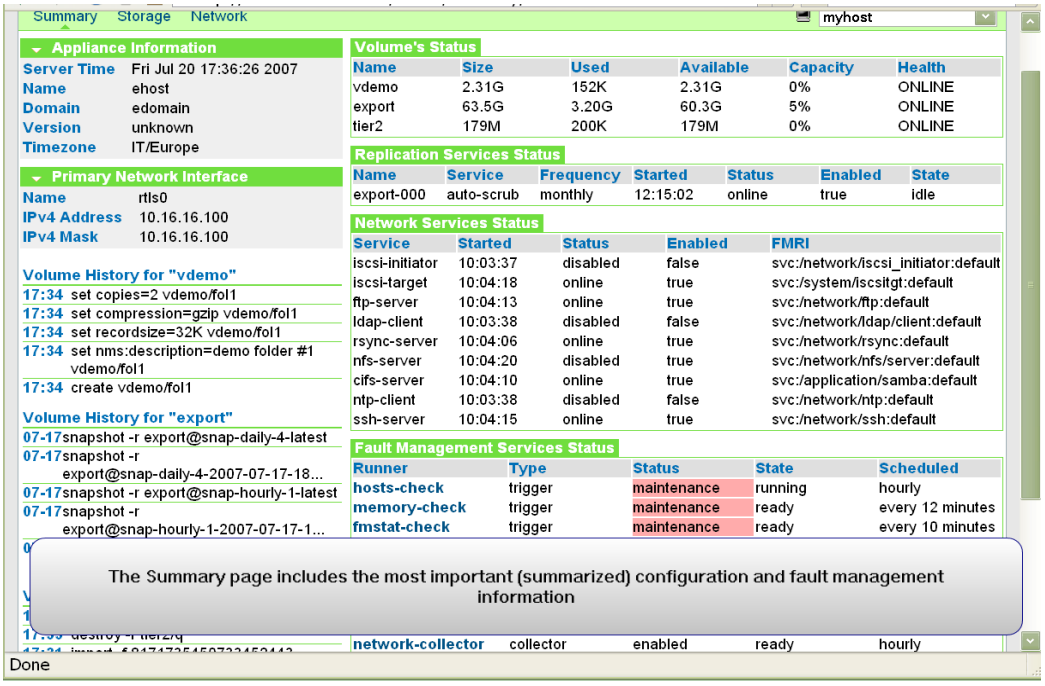

Graphical Statistics

Nexenta Management Server, along with pluggable "statistic collectors" (see Section "Terminology") collects – and NMV displays graphically – numerous statistics: network and storage, performance, utilization, and capacity related, graphs, bars, charts and much more. The following screenshot illustrates part of the work performed by one specific collector – '**network-collector**'. The picture shows input and output IP packets, input and output errors, and other related counters for a specified period of time. Note that all NexentaStor **collectors** do collect statistics continuously – a user can select any time interval in the past for an aggregated graphical representation.

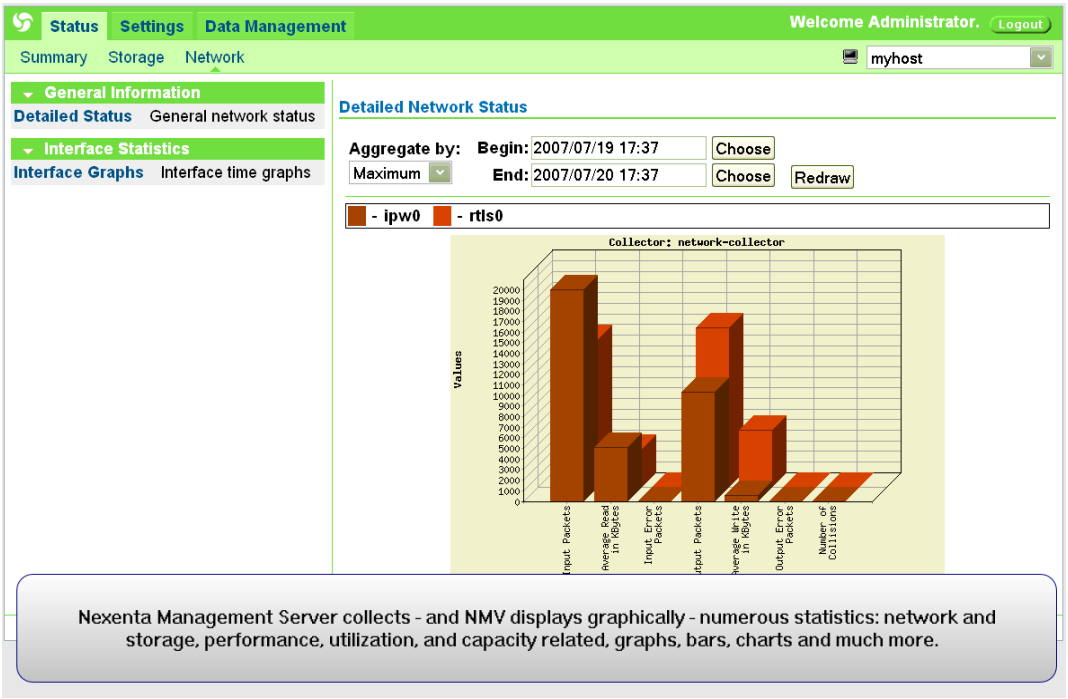

The concrete screenshot above shows statistics aggregated by Maximum for an interval of 24 hours. What that means in this particular case is that:

1. for each sampling interval (by default – one hour) during the 24 hours specified – take the maximum value of a statistic counter (say "input packets" for a given network interface)
2. sum up all these "maximums"
3. divide the result by the number of sampling intervals, 24 in the case.

NMV supports "aggregating" by Maximum, Minimum and Average, changing default sampling interval, selecting an interval of time, and more.

NMV supports a variety of disk space utilization and network I/O statistics (the next 2 screens):

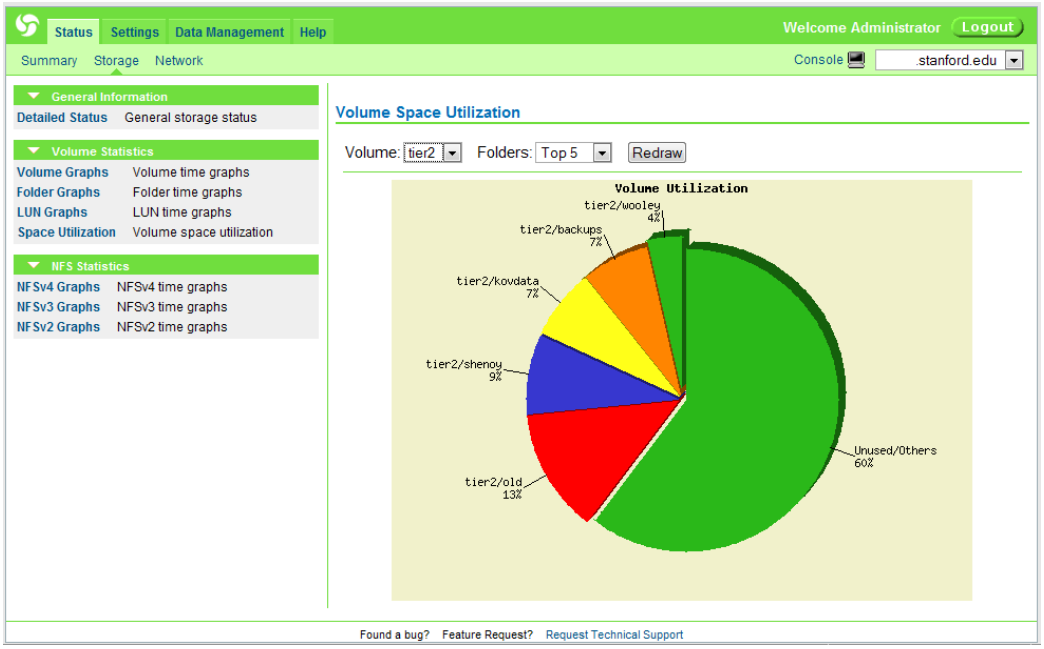

Folders and Shares

A set of screens under Data Management tab (below) allows to display and administer NexentaStor datasets: volumes, folders, and snapshots. The corresponding pages provide functionality to facilitate an easy management of the datasets. Appliance's folders, for instance, can be shared across different systems via a variety of access methods: CIFS, NFS, FTP, RSYNC, WebDAV ([24]).

New folders with a given set of properties can be created. Existing folders can be deleted and indexed (Section "Notes on Indexing and Searching"). Folders, as well as snapshots can be filtered by named – see the Search button at the bottom of the screenshot below:

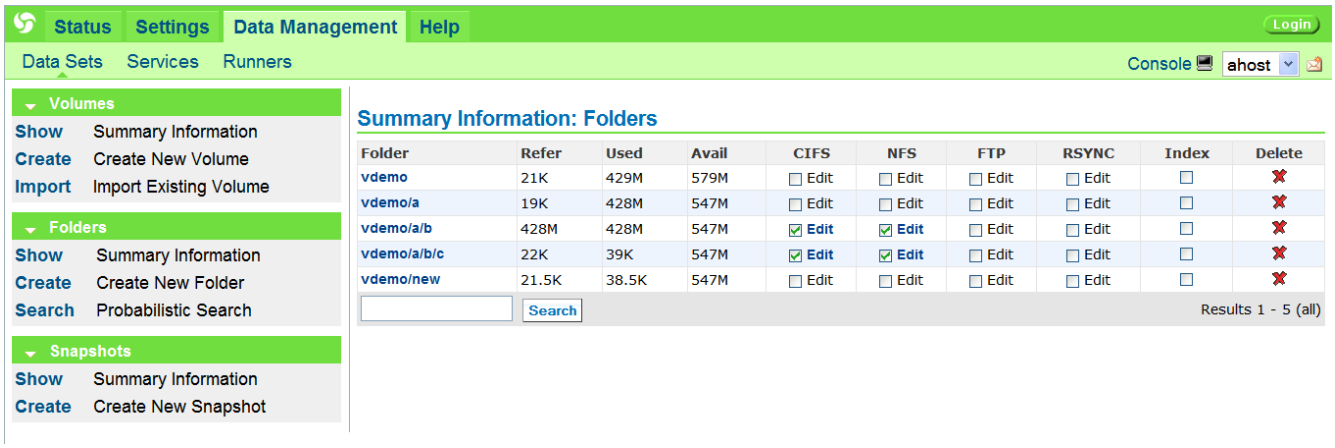

For detailed overview and illustrated examples of CIFS shares, please refer to Section "Tutorial", subsection "Sharing a folder via CIFS".

To share a given folder via NFS – that is, to make a local folder available for mounting by remote NFS clients, simply check a little checkbox in *this* folder's row under NFS, and then optionally click on Edit (see above) – to change the default sharing options. The following screen (in this example – for the folder **`primary/a`**) will show up:

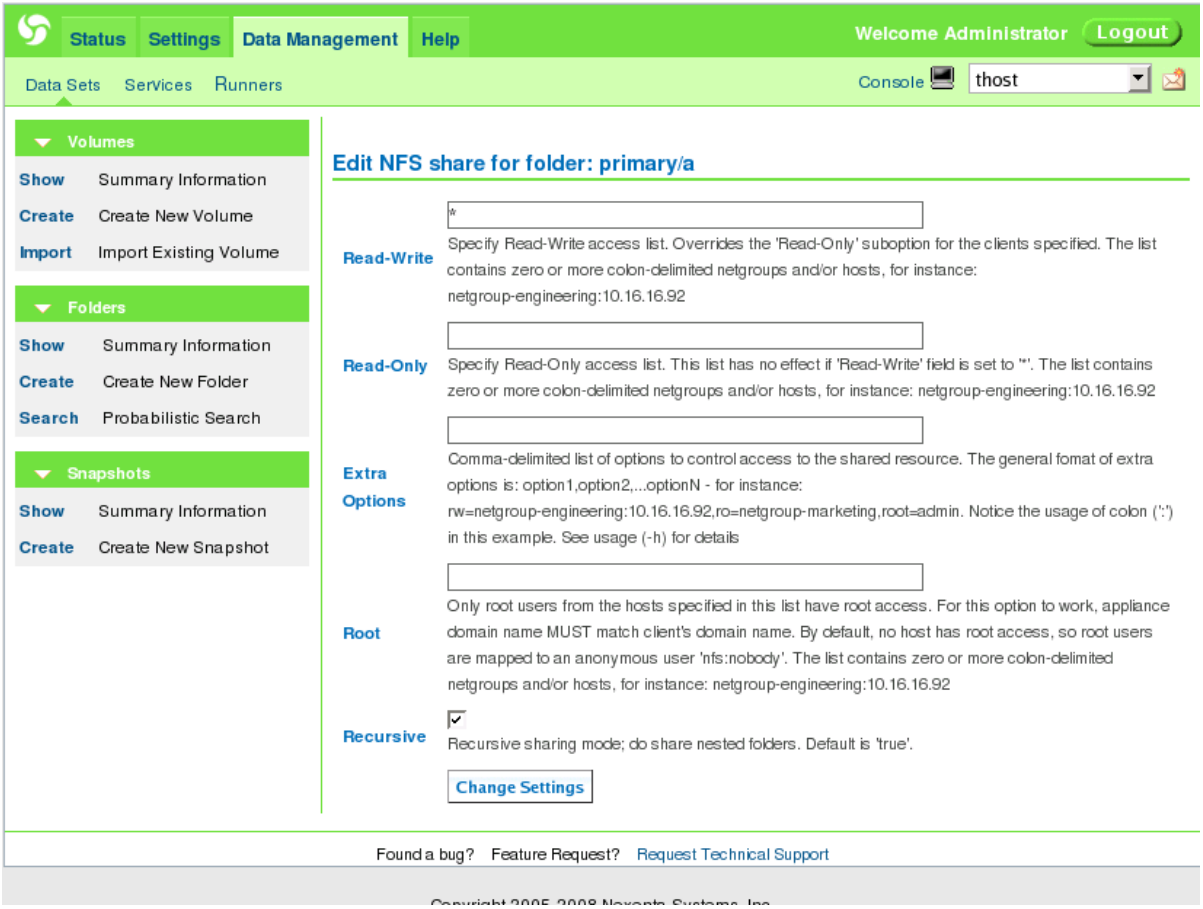

By default, NMV will create a share with a read/write access provided to all clients. Notice that the Root field remains empty – it is generally advisable to fill it out. For extended information and guidelines on NFS shares, and/or to administer NexentaStor folders via NMC, please see Section "Tutorial".

Configuration and Fault Management

With NMV, we can list:

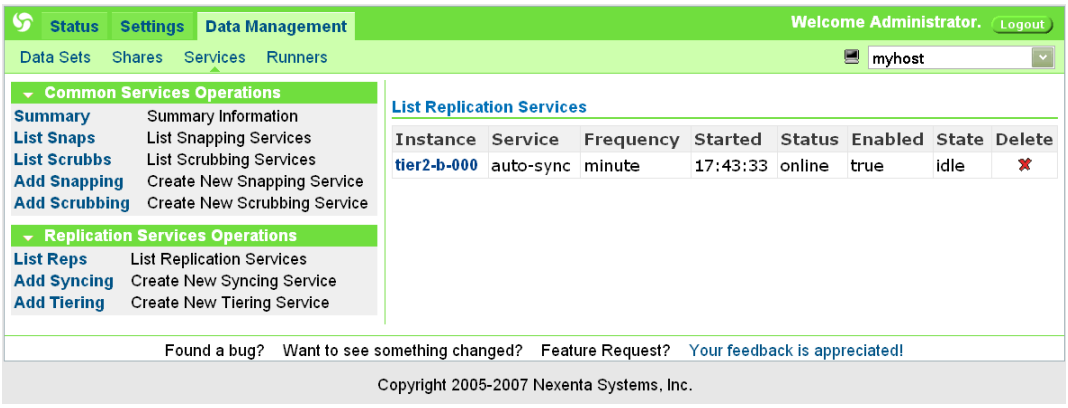

... create:

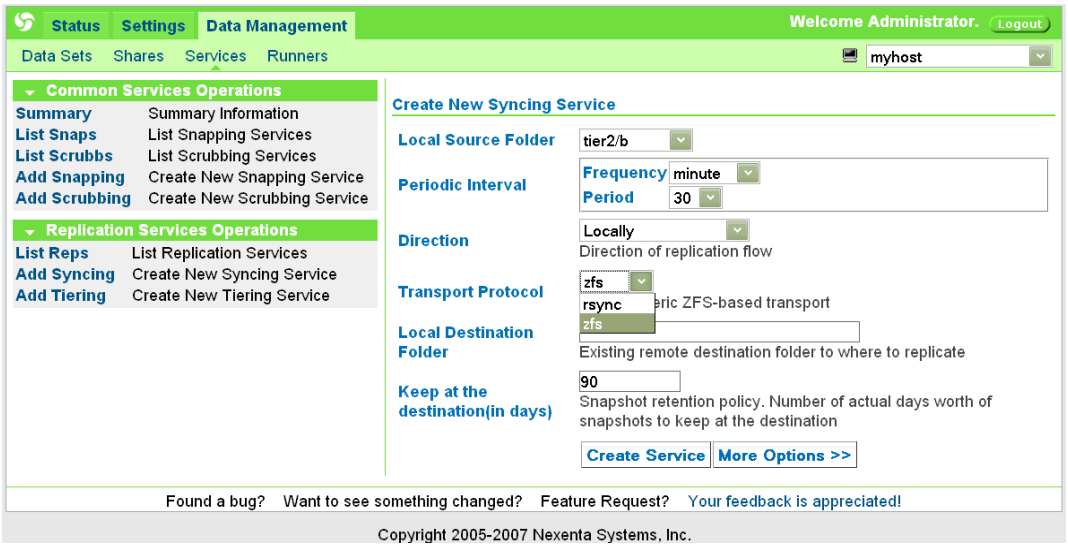

... and/or modify Nexenta storage services:

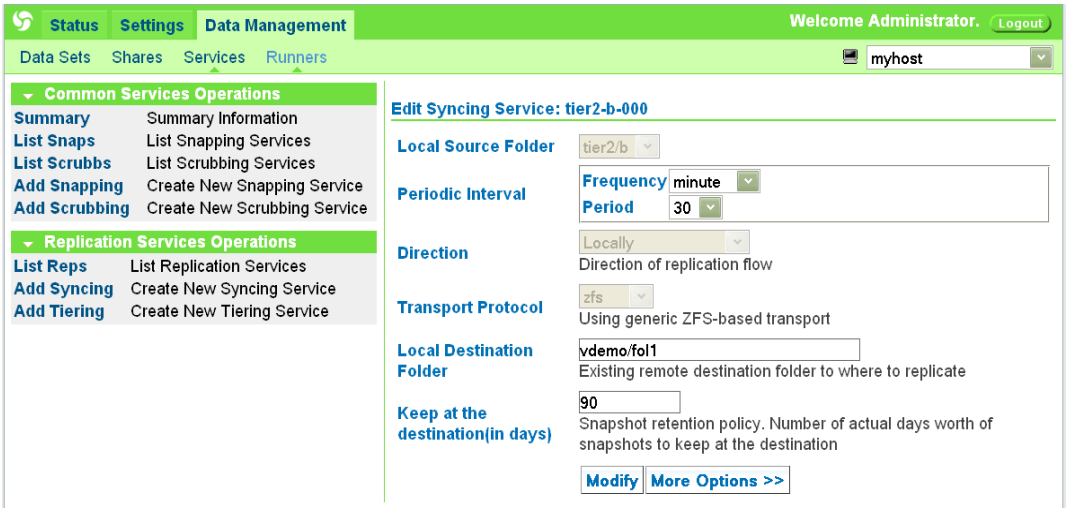

... view fault triggers, "zoom" in those that are in "red" state to get detailed reports, troubleshoot the reported faults:

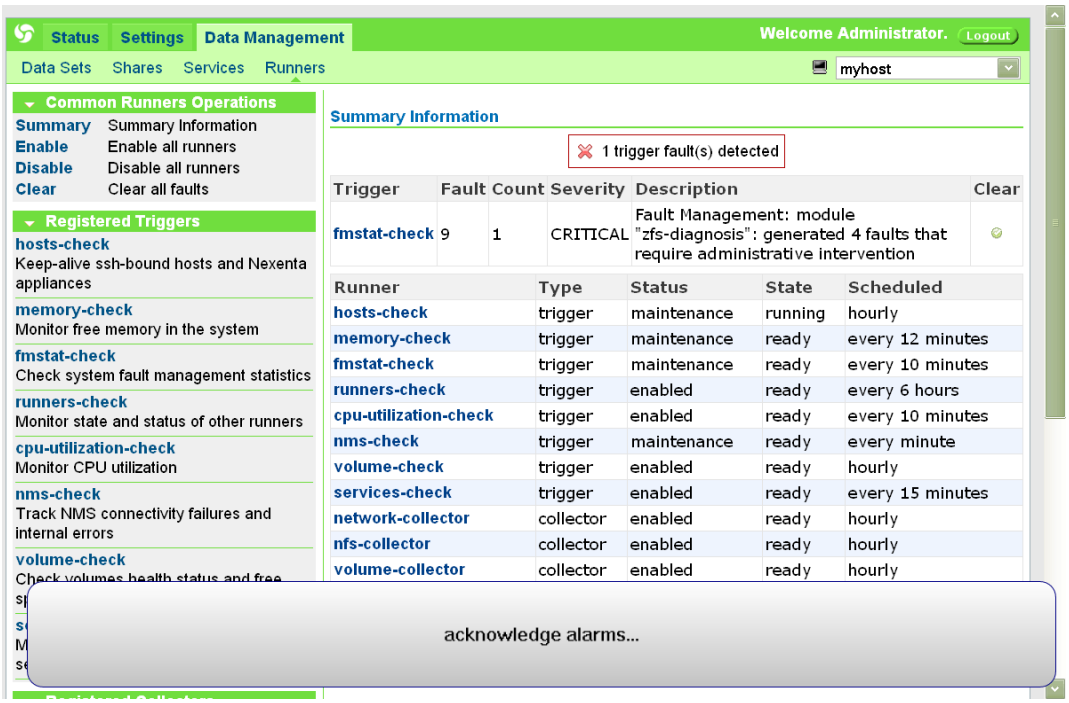

...acknowledge alarms (below), and more.

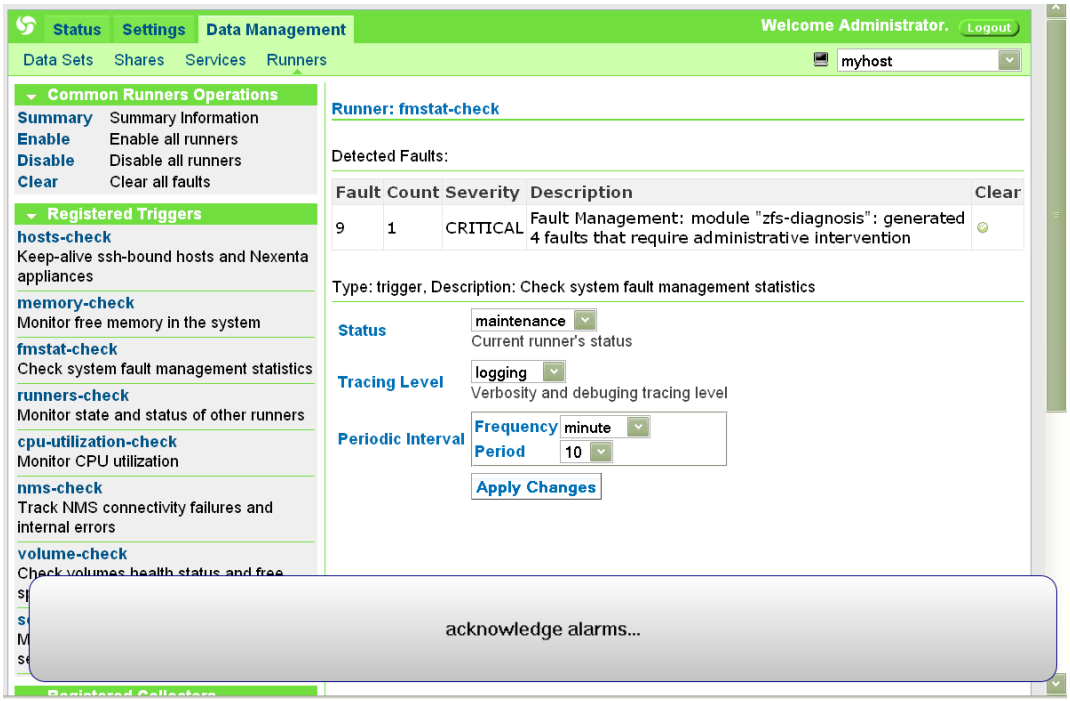

Terminal Access – NMC

The following two screenshots show NMV + management console integration. It is possible to run the management console from the web GUI interface, by clicking a little terminal icon on the NMV status bar.

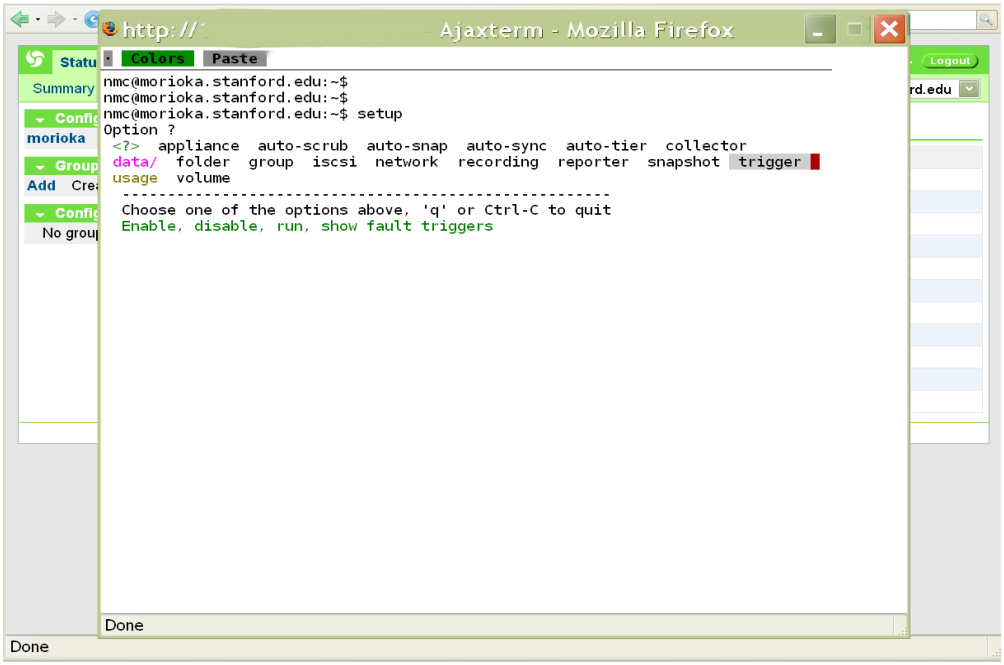

Which means, full power of the nmc (Section "NMC – Management Console") at your fingertips..

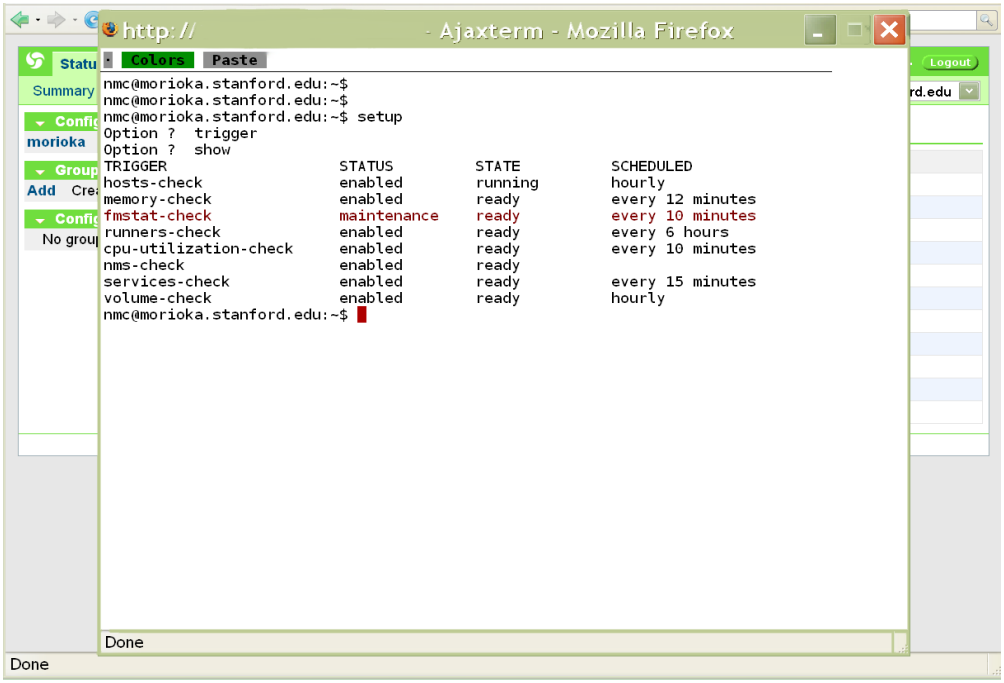

Server Options

The web management interface supports viewing and setting NMS preferences (compare with NMC – section "Preferences"):

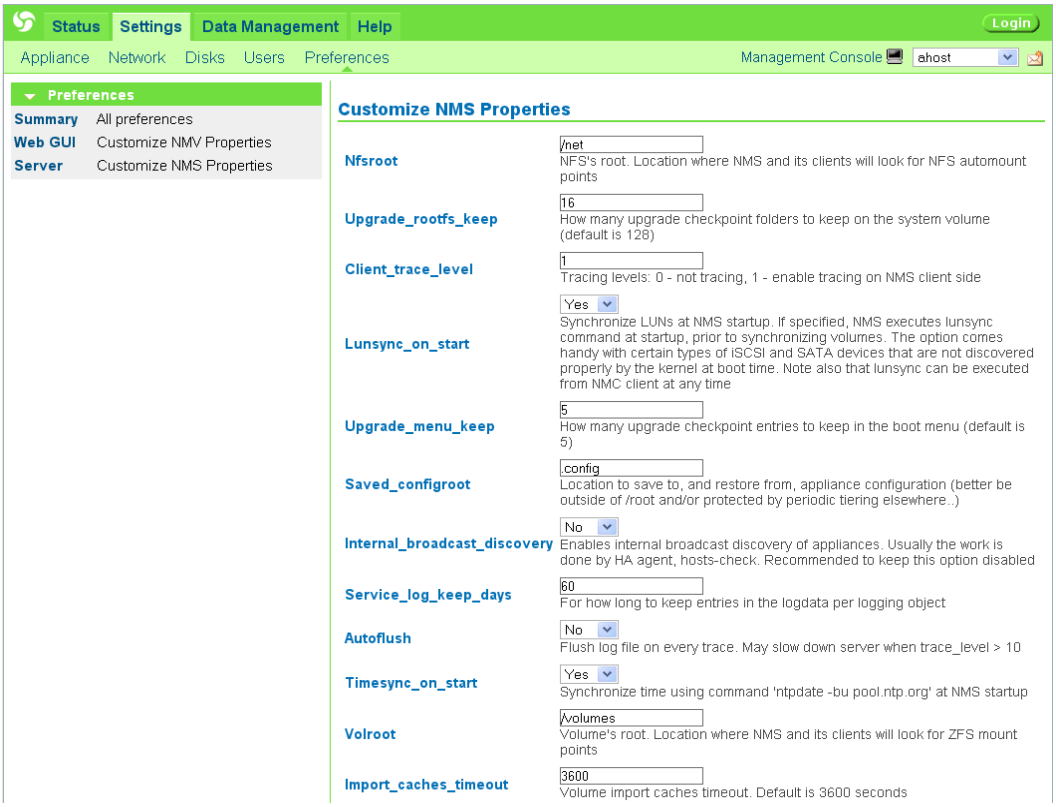

System Settings

System-wide settings (commented below, on the screenshot), can be modified to optimize performance or engage appliance's watchdog.

The following two system settings - **`sys_zfs_nocacheflush`** and **`sys_zil_disable`** can be set to optimize appliance's I/O performance. These configuration variables can be modified to disable ZFS cache flushing and ZIL (ZFS Intent Log), respectively.

While providing a considerable performance improvement in certain scenarios (in particular those involving CIFS, NFS or iSCSI) – this settings may be unsafe, in terms of application-level data integrity. It is strongly recommended to use this feature if and only if your storage is NVRAM protected, and the hardware platform is connected to Uninterrupted Power Supply (UPS). Default setting: unchecked (disabled)

Starting from version 1.0.5, the appliance supports setting the size of ZFS Adaptive Replacement Cache (ARC). The corresponding NMS property is called **sys_zfs_arc_max**. The ZFS adaptive replacement cache (ARC) uses most of a system's available memory to cache file system data. As memory pressure increases, the ARC relinquishes memory. It is recommended to reduce the **sys_zfs_arc_max** if the amount of RAM is less than 4GB[56]. In addition, as per ZFS Best Practices Guide, consider limiting the maximum ARC memory footprint in the following situations:

- A certain fixed amount of memory is always required by an application
- On platforms that support dynamic reconfiguration of memory boards

Notes on Indexing and Searching

To create an indexer[57] for a given folder (and then start running it periodically at configured intervals of time), simply check the corresponding checkbox on a NMV folders page:

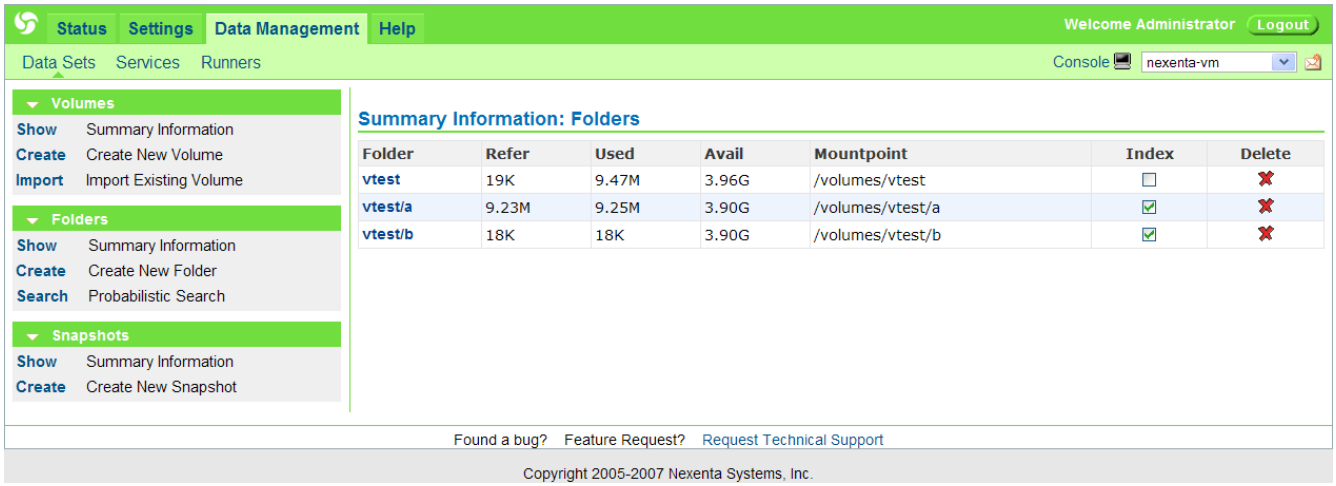

In the example above, two folders are getting indexed. The theory of operations behind it is specified in Section "Indexing NexentaStor Archives", along with one example of searching in folder **vtest/a**.

Unchecking an Index field of a folder table (see above) will effectively stop the corresponding indexer instance and cleanup (that is, remove from disk) all the accumulated indexing data.

Configuring an indexer is as easily done in NMV as it is done in the management console (Section "Search Engine and NexentaStor Indexers"). For instance, to configure an indexer associated with folder **vtest/a**, use the following NMV page:

56. Note that 4GB RAM is recommended for most bare-metal deployments (see http://www.nexenta.com/products for system requirements)
57. "Indexer" as a term is defined in Section "Terminology".

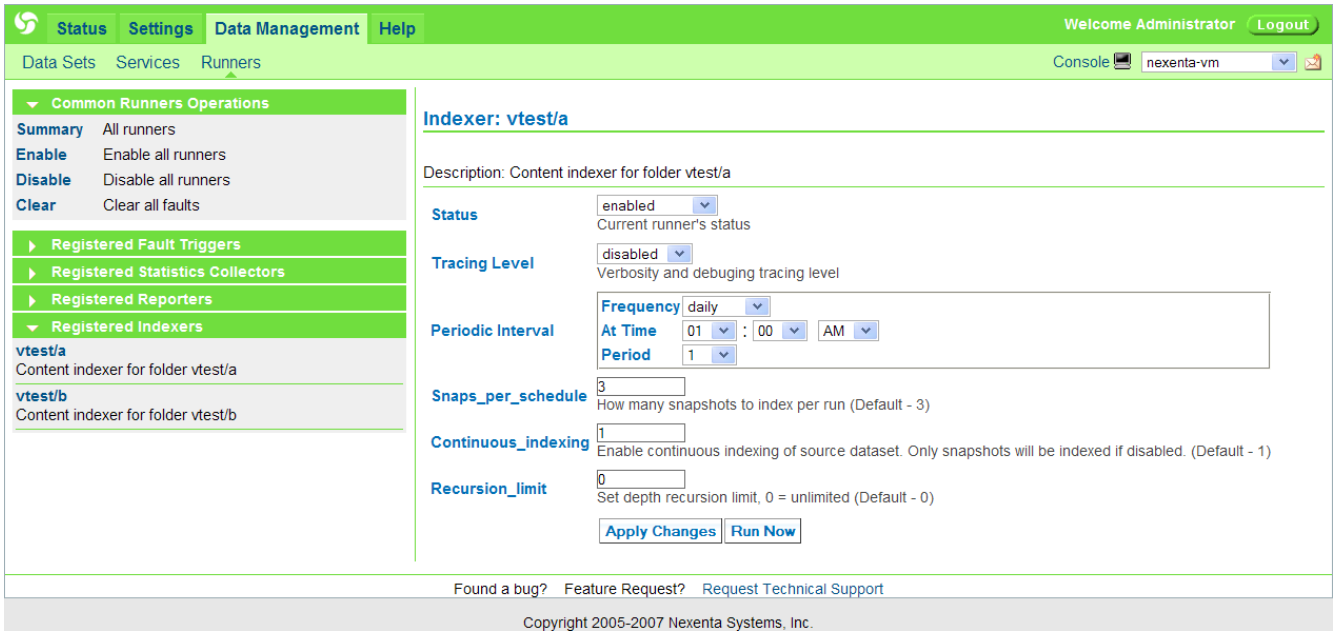

Documentation – Registration – Support

Complete product documentation is available online – from the Help menu simply click on the "NexentaStor User Guide" link (below):

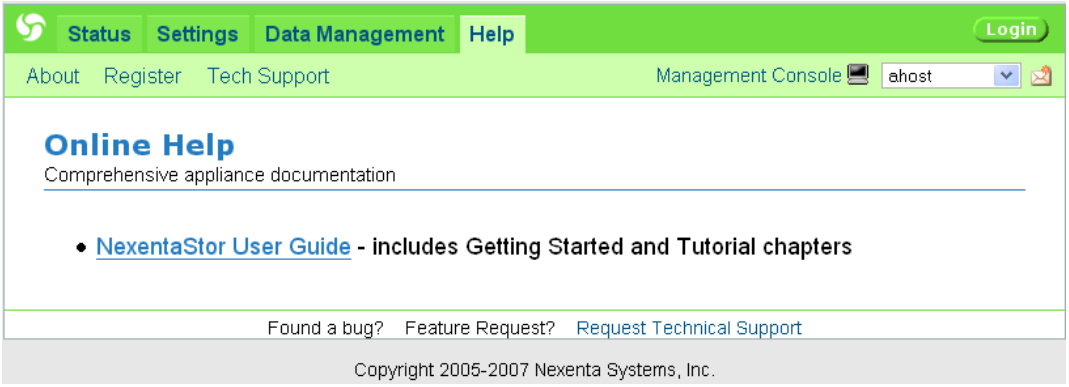

Here you can also register the product, extend or upgrade its license.

Finally, use the Technical Support form (below) to report any technical problems and/or request new features. From this page a simple email can be sent to support technicians. This email will include a snapshot of your system configuration, which reduces the time spent on support and allows technicians to begin troubleshooting immediately.

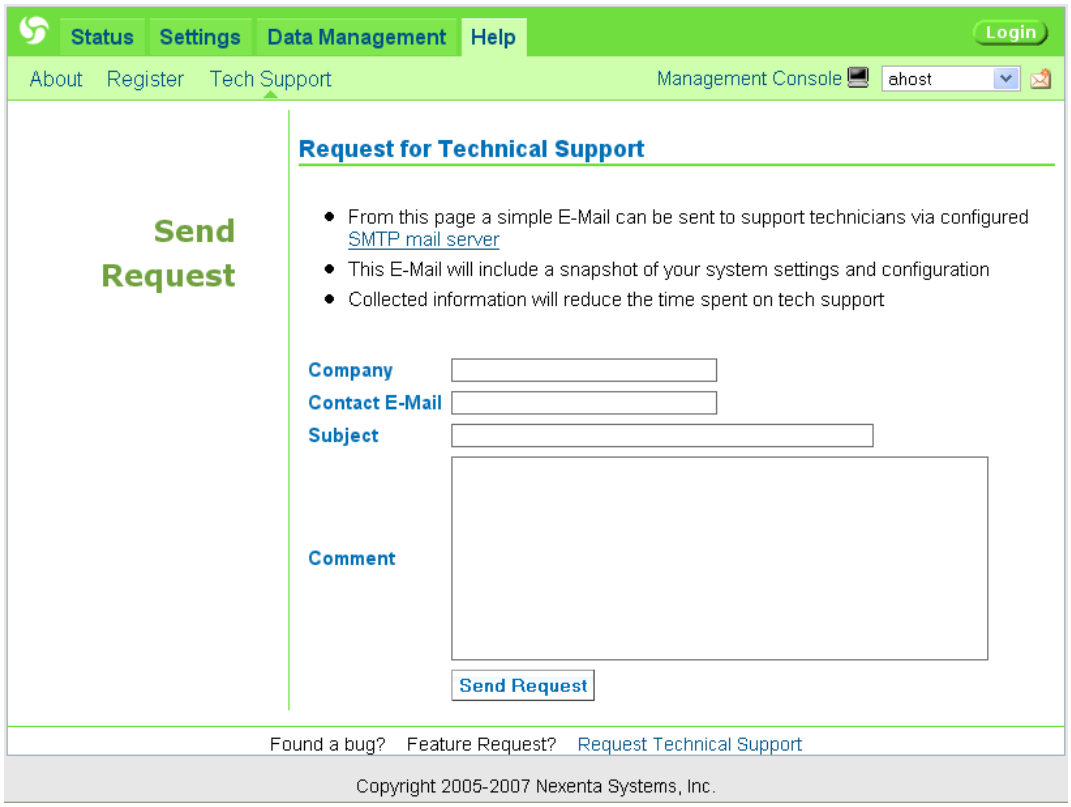
Status
Settings
Data Management
Help
Login
About
Register
Tech Support
Management Console
ahost
Send Request
Request for Technical Support
From this page a simple E-Mail can be sent to support technicians via configured SMTP mail server
This E-Mail will include a snapshot of your system settings and configuration
Collected information will reduce the time spent on tech support
Company
Contact E-Mail
Subject
Comment
Send Request
Found a bug?
Feature Request?
Request Technical Support
Copyright 2005-2007 Nexenta Systems, Inc.

7. More Use Cases

Configuring Thumper

NexentaStor is well suited for the X4500 product, a system seemingly designed from the beginning to exploit all the positives of a ZFS-based system. With a little guidance, this 48-drive, 24TB beast can be configured to operate in a similar fashion to a FAS3050. First, one needs to realize the target configuration. In the case of a FAS3050, you tend to have multiple RAID groups configured, generally covering disks spread across a disk shelf. These shelves combine, with redundancy up to and including the system disk, to minimize total system failure scenarios. In the case of the X4500, installation of NexentaStor is relatively straightforward to achieve this. The first pearl of wisdom is that even though it may be labeled Disk 0 and Disk 1, the actual system startup drives are on the 6th controller, disks 0 and 4.

After starting up the system with the Nexenta installer, one much unselect the first pre-selected drive and pick c6t0d0 and c6t4d0. This configuration will set aside two of the disks for the system images and all subsequent checkpointed configurations.

The usual configuration has 8 disks each on 6 controllers, numbered 0, 1, 5, 6, 7, and 8. The PCI bus labels the last two as 10 and 11 during the BIOS post, but software will list them as mentioned above. To achieve the most resilient configuration, its obvious that any configuration will need treat card 6 differently. Thinking of our goal again, it makes sense to stripe any array across all cards, which is mostly synonymous with the disk shelves in other NAS solutions.

One hitch with ZFS is that it would be otherwise wasteful and unreliable to define 8 RAID groups of only 5 drives each, as we'll avoid the somewhat pre-allocated card 6. As this configuration would be akin to a RAID 5, we'd only be able to survive one full card failure at a time for all RAID groups. Rather, the ideal configuration is a RAID 6 whereby any two drives can fail at a time. By configuring 4 RAID groups of 10 drives each, with 2 drives per card, we allow for both a full card failure as well as any two drives to fail at a time. The remainder 6 drives on card 6 I allocate as global spares for the NAS. This allows for recovery of up to 1 disk per RAID group in the case of a card failure, putting us in a safer situation.
The NMC interface was sufficient to define a volume of this configuration:

- run '**`create volume data`**' (see "Tutorial", Section "Creating a volume")
- select drives c0t0d0 ... c8t1d0 (skip c6t*d* drives)
- select redundancy type raidz2
- repeat the last two steps to add more groups
- add remaining c6 series drives as spares;
- execute. (NOTE: you can always use the -n (dry-run) option to simulate creation.)

The resulting configuration follows below:

```
NAME                STATE       READ WRITE CKSUM
tier2               ONLINE       0     0     0
      raidz2        ONLINE       0     0     0
            c0t0d0  ONLINE       0     0     0
            c0t1d0  ONLINE       0     0     0
            c1t0d0  ONLINE       0     0     0
            c1t1d0  ONLINE       0     0     0
            c5t0d0  ONLINE       0     0     0
            c5t1d0  ONLINE       0     0     0
            c7t0d0  ONLINE       0     0     0
            c7t1d0  ONLINE       0     0     0
            c8t0d0  ONLINE       0     0     0
            c8t1d0  ONLINE       0     0     0
      raidz2        ONLINE       0     0     0
            c0t2d0  ONLINE       0     0     0
            c0t3d0  ONLINE       0     0     0
            c1t2d0  ONLINE       0     0     0
            c1t3d0  ONLINE       0     0     0
            c5t2d0  ONLINE       0     0     0
            c5t3d0  ONLINE       0     0     0
            c7t2d0  ONLINE       0     0     0
            c7t3d0  ONLINE       0     0     0
            c8t2d0  ONLINE       0     0     0
            c8t3d0  ONLINE       0     0     0
      raidz2        ONLINE       0     0     0
            c0t4d0  ONLINE       0     0     0
            c0t5d0  ONLINE       0     0     0
            c1t4d0  ONLINE       0     0     0
            c1t5d0  ONLINE       0     0     0
            c5t4d0  ONLINE       0     0     0
            c5t5d0  ONLINE       0     0     0
            c7t4d0  ONLINE       0     0     0
            c7t5d0  ONLINE       0     0     0
            c8t4d0  ONLINE       0     0     0
            c8t5d0  ONLINE       0     0     0
      raidz2        ONLINE       0     0     0
            c0t6d0  ONLINE       0     0     0
            c0t7d0  ONLINE       0     0     0
            c1t6d0  ONLINE       0     0     0
            c1t7d0  ONLINE       0     0     0
            c5t6d0  ONLINE       0     0     0
            c5t7d0  ONLINE       0     0     0
            c7t6d0  ONLINE       0     0     0
            c7t7d0  ONLINE       0     0     0
            c8t6d0  ONLINE       0     0     0
            c8t7d0  ONLINE       0     0     0
      spares
          c6t1d0    AVAIL
          c6t2d0    AVAIL
          c6t3d0    AVAIL
          c6t5d0    AVAIL
          c6t6d0    AVAIL
          c6t7d0    AVAIL
```

NexentaStor and ESX

VMware ESX (R) also works well with NexentaStor in multitude of ways. Today, you can use Nexenta as your Data Storage pool for virtual machines, either using iSCSI or via NFS. To enable the easiest VMotion usage, high availability, and management, it is suggested to use the NFS NAS functionality. When defining Data Stores in the ESX management consoles Configuration tab, you simply need to list the Nexenta store by hostname with the correct share point, such as (for instance) **`/volumes/data/esx`**.

Prior to the above, one would need to:

- have a volume '**`data`**' defined,
- with '**`create folder data/esx`**' followed by
- '**`share data/esx nfs`**', with rw and root permissions being explicitly set to your ESX VMKernel IP address.

Migrating a Server

Here's a scenario where you are migrating away from one server to another for your second tier (target) storage.

If you want to copy data first to the new target, you'll need to setup **`auto-sync`** services for all volumes to the new system. After those run and complete (this can take many hours depending on storage capacity!), you'll also need to do some cleanup.

The cleanup is extremely simple: on the source of the syncs, destroy all the **`auto-syncs`**. The added benefit of this operation is that the synced folders on the target will immediately become mounted.

By default, **`auto-sync`** service keeps its destination folders/filesystems not mounted. Mounting a destination opens a possibility of changing its data and/or metadata, which means that the delta accumulated at the **`auto-sync`** source since the previous **`auto-sync`** run cannot be correctly applied.

However.

The default behavior can be overridden – see NMC's '**`create auto-sync -h`**'. The service logic will make sure that if the sync destination has changed since the previous service run, those changes will be reverted.

In addition, and this is relevant to the topic at hand, '**`destroy auto-sync`**' automatically mounts its destination.

Once the old (source) and the new (target) servers are completely in-sync, we can return to NMC or use NMV for the **`auto-tier`** creation to continue to tier those migrated volumes from their original sources over days and weeks to come..

Using ZVOLs

Zvol is an emulated (virtual) block device based on a given appliance's volume. **Zvol** can be used additional swap partition but the primary purpose of zvols is to provide an easy iSCSI integration.

Tight integration with the appliance's services amplifies the usability of zvols. Similar to NexentaStor volumes and folders, you could snapshot zvols, **auto-snap** and **auto-sync** zvols (see Sections "Terminology", "Data Replication"), and in general, apply the full power of the appliance's management framework.

To illustrate it, one could imagine the following use case:

1) Thinly provision a 5TB block device (zvol) based on a volume vol1:

```
nmc:/$ create zvol vol1/zvol1 -S -s 5TB
```

This will create a new thinly provisioned block device named **vol1/zvol1**, with volume **vol1** providing backend storage. As always, instead of specifying parameters in the command line, you could simply run '**create zvol**' and follow the prompts.

The term "thin provisioning" here means that you do not need to have the entire 5TB of storage available; the space is allocated on demand. For more information on the command '**create zvol**', see:

```
nmc:/$ create zvol -h
```

In NMV, a similar functionality is available via Settings => Disks (click on **Create**):

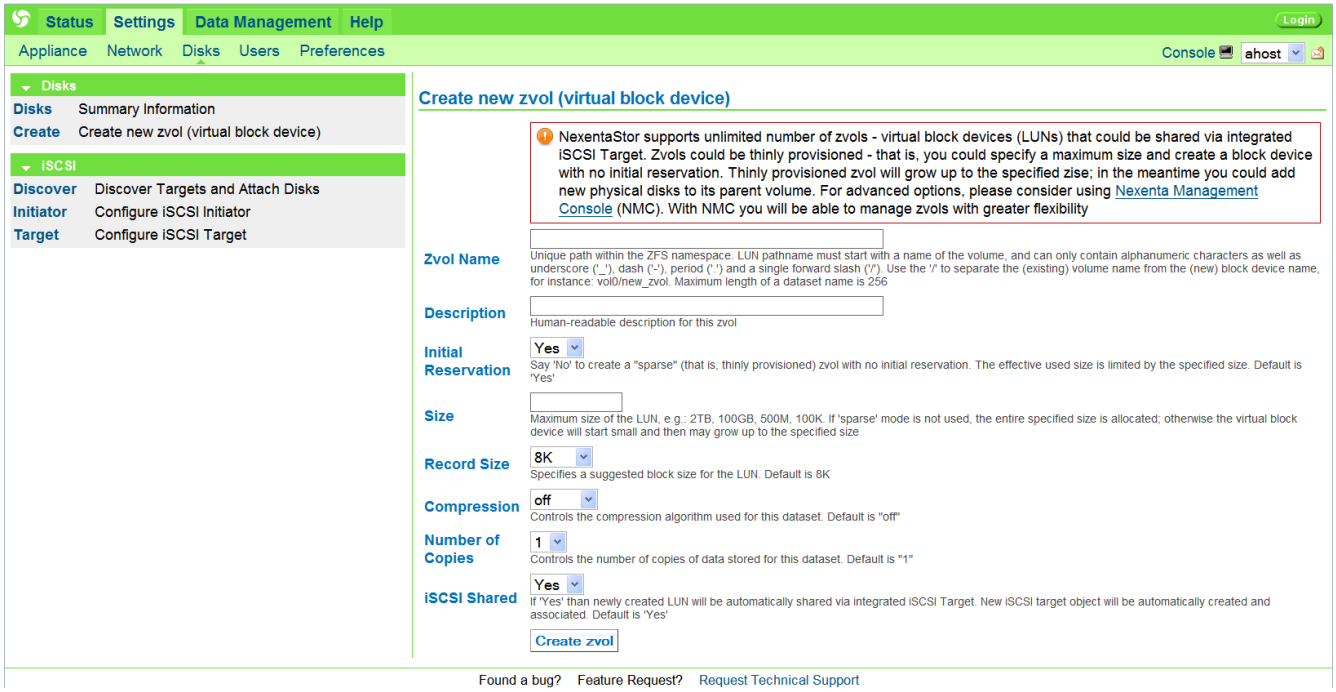

2) Share the new zvol via iSCSI:

```
nmc:/$ share zvol vol1/zvol1
```

This step could be skipped if you use -T option with '**create zvol**' command – see '**create zvol -h**' for details.

3) Discover this device (or, more exactly, the corresponding iSCSI target) from your Windows workstation or laptop. Open iSCSI connection, and effectively attach the device via iSCSI.

4) Format the device as NTFS or FAT, and start using it.

On the NexentaStor (backend) side, you could **auto-snap** and/or **auto-sync** the newly created zvol, thus maintaining a fully synchronized history of the NTFS filesystem, for days, weeks and months. For instance, run:

```
nmc:/$ create auto-snap zvol vol1/zvol1
```

and follow NMC prompts to specify the frequency of snapshots, retention policy, and other details. Later you could run '**setup auto-snap**' to modify the parameters, etc.

As usual, NexentaStor management GUI provides a similar support. For instance, the following screenshot shows a particular **zvol** attributes. The page is located under Settings => Disks: simply click on the **Show**, and then on a specific **zvol**:

8. Windows Backup and Restore

Nexenta RsyncShare is an open source implementation of rsync shares manager for Windows servers, workstations, and laptops. RsyncShare is a GUI based, CDDL licensed, and free of charge part of the NexentaStor product. RsyncShare makes Windows machines accessible via any rsync client. Use RsyncShare for fast incremental Windows backup and restore. Use **NexentaStor + RsyncShare** combination for cost-effective and high-performance disk-to-disk backup and recovery, with continuous data protection and a rich variety of daily/weekly/monthly archiving options (Section "Data Replication").

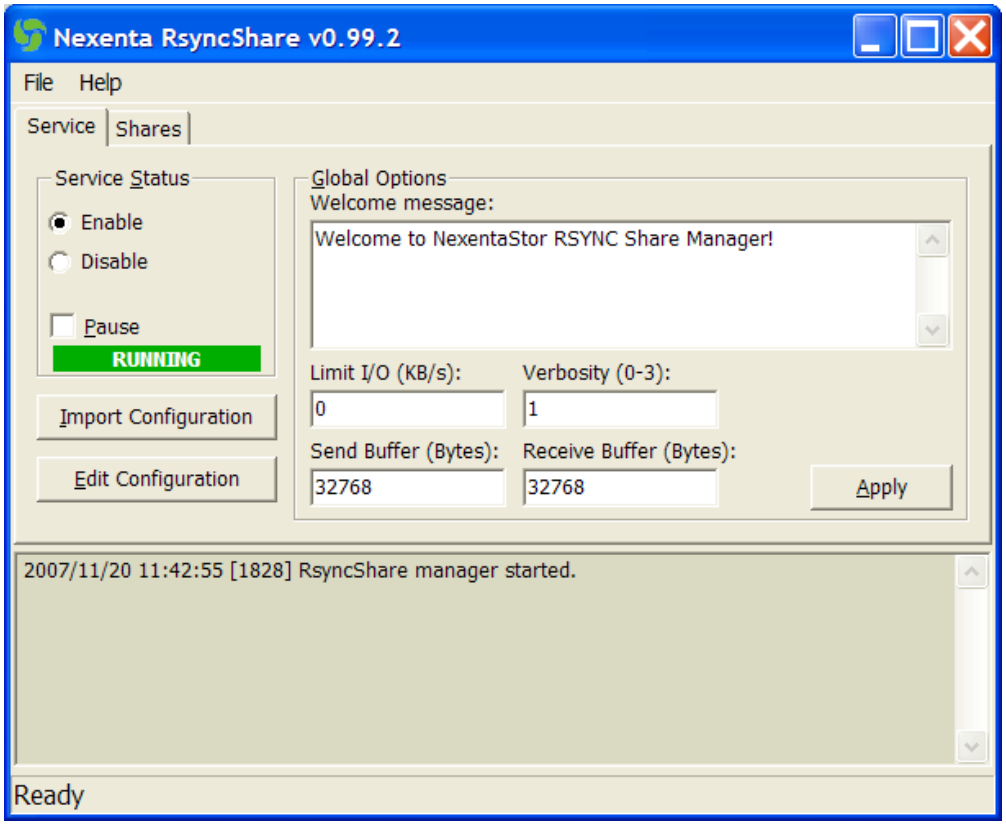

RsyncShare uses the rsync algorithm which enables a fast method for synchronizing remote and local files and directoies, by sending just the differences across the link.

Requirements:

- Windows NT/2000/2003/XP/VISTA
- 32MB of free memory
- 20MB of free disk space

RsyncShare is available for download here.

9. References

[1] NexentaStor Virtual Appliance: Quick Start Guide at http://www.nexenta.com/docs

[2] NexentaStor API and Software Development Kit, at http://www.nexenta.com/corp/index.php?option=com_jreviews&Itemid=122

[3] Tutorials on the web: at http://www.nexenta.com/tutorials

[4] Frequently Asked Questions (FAQ): at http://www.nexenta.com/nexentastor-faq

[5] Downloads, at http://www.nexenta.com/products

For comments, recommendations or information contact support@nexenta.com.