

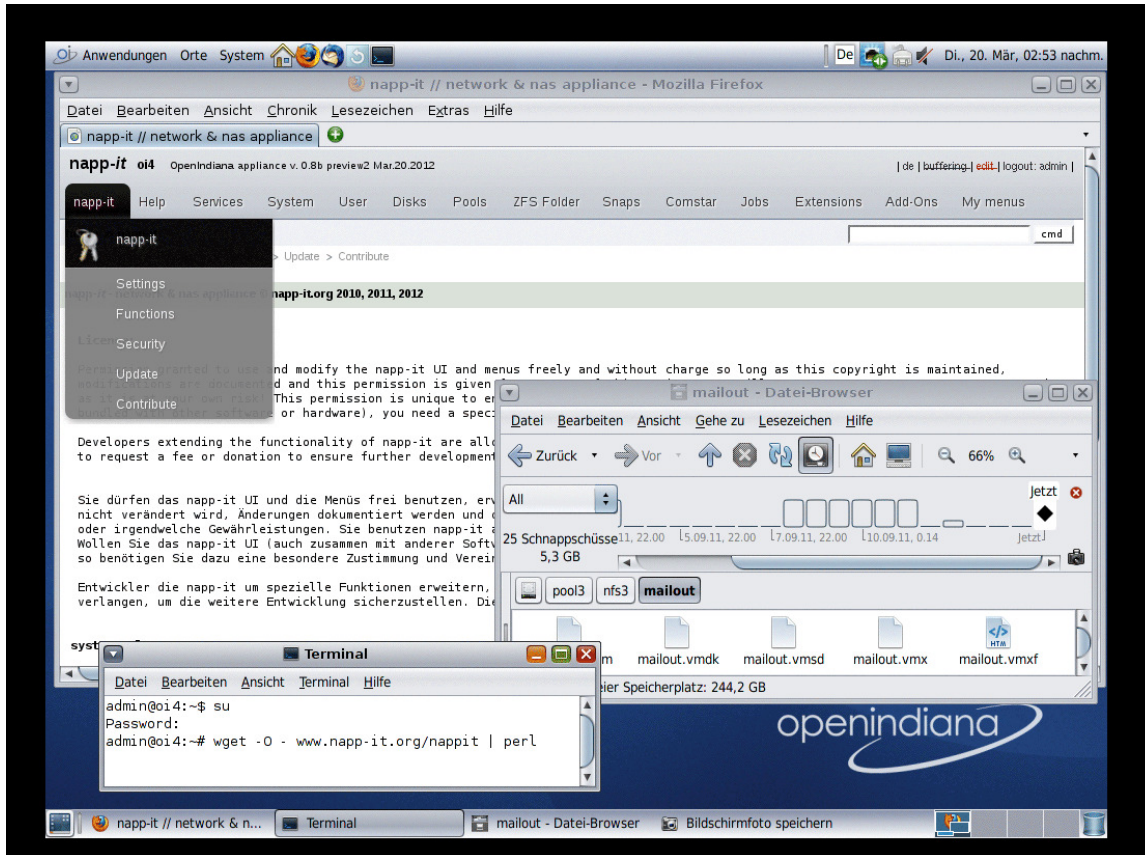
NAPP-IT ZFS SERVER: THE MISSING MANUAL

published: 2012-03-20

Licence: CC-BY-SA see <http://creativecommons.org/licenses/by-sa/2.0/>

always keep a reference to: <http://www.napp-it.org/doc/downloads/napp-it.pdf>
for newest versions

NAPP-IT ZFS SERVER UI



OTHER MANUALS

All-In-One

ESXi-Server (also free ESXi V. 4/5) with an embedded virtualized NAS/SAN ZFS storage server:
<http://www.napp-it.org/doc/downloads/all-in-one.pdf>

About this book

This book is based on the effort of all wanting a base documentation on OpenSolaris derived operating systems like Illumian, SmartOS, Solaris 11 and the napp-it ZFS server.

Currently most documents are from Oracle regarding Solaris 11. While this is a very valuable source, differences are increasing, so its often more confusing than helpful. On the other side, the OpenSolaris bible, the other often used source is more and more out of date.

Writing this book is done mostly in spare time and from non native english speakers with different skills. All are invited to improve this book regarding language, content or errors. If you like to translate these documents into another language you can start your own group and book-series. I suggest you name it like Illumian & OpenIndiana + napp-it (fr): System Management (translated name)

All documents should be named like "system_management.pdf" independant of the language. This will help to create a consistent folder with all documents, even if there are untranslated parts. A index.html file within the folder may link between the document. Please keep the look and feel (fonts, headings and structure also consistent).

If you want to collaborate you can either contribute by co-editing a book or request to start a new under this theme like:

Illumian & OpenIndiana + napp-it: Index (Content and overview of all books)

Illumian & OpenIndiana + napp-it: System Management
Illumian & OpenIndiana + napp-it: Networking
Illumian & OpenIndiana + napp-it: Media Server
Illumian & OpenIndiana + napp-it: Comstar
Illumian & OpenIndiana + napp-it: LDAP
Illumian & OpenIndiana + napp-it: Media Server
Illumian & OpenIndiana + napp-it: KVM

If you like to contribute, send a mail to community@napp-it.org

This book is published under a CC-BY-SA licence, which means practically you can do whatever you want unless you keep a reference to the origin/author und publish under same condition.

Gea
www.napp-it.org

INTRODUCING THE OPENSOLARIS DERIVED OPERATING SYSTEMS

Developed by Sun with its initial release in 2008, based on Unix System 5, Release 4 with revolutionary features like ZFS (<http://en.wikipedia.org/wiki/ZFS>), Comstar (enterprise ready iSCSI and FC technology), DTrace, Crossbow virtual switches, virtualization with Zones (http://en.wikipedia.org/wiki/Solaris_Zones) and ZFS integrated kernel based CIFS/SMB and NFS server among other features.

If you want to follow the way from Sun OpenSolaris to current Illumos based distributions, you should look at Bryan Cantrill's (Joyent) slideshow at <http://www.slideshare.net/bcantrill/fork-yeah-the-rise-and-development-of-illumos> or <http://www.youtube.com/watch?v=zRN7XLCRhc>

Some essentials from the slideshow:

- "In the mid-1990 it becomes a foregone conclusion that UNIX would die at the hands of Windows NT
- Hard to believe, but SUN was the only computer vendor whose Windows NT strategy was to beat Windows NT
- Sun's dedication to this vision - the operating system as a nexus of innovation - attracted an entire new generation of engineers to the company
- Development started on more radical ideas, each of it would become revolutionary in its own right (ZFS, DTrace, Zones, Crossbow, Comstar, Fault Management, Service Management, Least Privilege)
- These were technologies invented and initiated by engineers, not managers or marketers
- These projects reflected the people behind them..Organizations don't innovate - people do
- As the rise of Linux forced the market price of OS acquisition to zero, that open sourcing the (Solaris) operating system was the right business decision (in 2005. Unfortunately, not all elements of the OS could be opensourced, some contracts prevented some small but important bits from being opensourced. To allow such proprietary drivers, Sun developed a file based copy left licence (CDDL) ... this was not done to deliberately GPL incompatible!
- Ailing Sun was bought by Oracle in 2009. Over 2010, it becomes clear that Oracle had absolutely no interest in OpenSolaris... There was.. a move to close the system (OpenSolaris)
- Starting in the summer of 2010, Garrett D'Amore at Nexenta - with the help of Rich Lowe, Jason King and others - began the process of either writing the closed bits from scratch or porting from BSD.
- Dubbed "Illumos" (from illuminate, Latin for illuminate) and made available on August 3, 2010
- Illumos was not designed to be a fork, but rather an entirely open downstream repository of OpenSolaris
- Solaris 11 was released on November 9, 2011 - and there was no source release
- The entire DTrace team, all primary ZFS inventors and primary engineers for zones and networking had left Oracle.. nearly all of these engineers went to companies betting on Illumos.
- In Illumos, we have seen critical innovations and bug fixes in ZFS, DTrace, Zones and other core technologies. These innovations will never be in Oracle Solaris.
- Joyent team ported KVM from Linux. Illumos distributions SmartOS and OpenIndiana have KVM support by default."

Currently there are the following distributions derived from the last free OpenSolaris (Unix, OpenSource, <http://en.wikipedia.org/wiki/OpenSolaris>), no longer developed, but the base of all others due to the **illumos** fork of its base (www.illumos.org, <http://en.wikipedia.org/wiki/Illumos>)

Mainstream distributions (OpenSource or free to use)

- **Illumian** (OpenSource), Base of NexentaStor and maintained by Nexenta, something like CLI based-OpenIndiana with Debian-like package management (apt-get), server only installation, capable of creating mirrored system-disks during initial installation. Illumian is CLI only but can be setup and managed via napp-it. The Nexenta Web-GUI is not available for Illumian.
- **NexentaStor CE**, www.nexentastor.org, free version of commercial NexentaStor Storage Appliance with up to 18 TB of storage but without support or plugins from the commercial edition, not for general use but storage only, includes the Web-GUI from Nexenta. Napp-it is not running together with the Nexenta GUI.
- **OpenIndiana** (OpenSource), www.openindiana.org, Server and Desktop use, Solaris like package management (pkg.). Very similar in most functions beside encryption to Solaris 11 Express (Those manuals are quite correct). <http://en.wikipedia.org/wiki/OpenIndiana>. OpenIndiana comes in a serverlike CLI version or in a more user friendly (Desktop) live edition. Both can be setup and managed via napp-it.

If you do not have special demands, start with this distribution.
If you are not an experienced Unix admin, use live edition with a GUI

Commercial distributions

- **NexentaStor Appliance**, www.nexenta.com, Web managed Storage server appliance with support and plugins like HA and Vmware support. <http://en.wikipedia.org/wiki/NexentaStor>. Napp-it is not running together with the Nexenta GUI.
- **Solaris 11**, www.oracle.com. Solaris 11 comes in a serverlike CLI version or in a more user friendly (Desktop) live edition. Both can be setup and managed via napp-it.

Special distributions (OpenSource or free to use)

- **Eon**, based on Illumos, <http://sites.google.com/site/eonstorage/>, server like core install, intended to run from RAM and booted from USB
- **SmartOS**, based on Illumos, maintained by Joyent, www.smartos.org, mainly intended to be a base for virtualization with KVM and to host other Operating Systems on a ZFS, zones and dtrace capable machine

Less known distributions like **Schillix** or **Belenix**, <http://schillix.berlios.de/> and <http://en.wikipedia.org/wiki/Belenix>

OS-INSTALLATION AND FIRST STEPS

Care about:

If you use Windows from Microsoft nearly all hardware will work. If you use Linux a lot of hardware will work. If you use Apple or Solaris, you must care about your hardware. The good news is that most of the hardware one wants to use to build a stable ESXi, Linux or Windows server will work with Solaris with the following specifics.

Mainboards

Solaris and Co are running quite well on modern mainboards (AMD and Intel). But if you can afford, buy always server-grade mainboards. They are mostly more stable, have IPMI remote management, ECC-RAM (detects errors in RAM), more fast PCI-slots, stable server grade NICs and vt-d if you want to try virtualizing a storage server now or later (Virtual Server with included NAS/SAN).

CPU and RAM

While the minimum like a decent 32 Bit CPU, even Atom and 1 GB RAM is capable to run Solaris derived systems, its not enough beside a simple single user backup or media-server at home. ZFS was developed with really heavy use cases in mind. The limitations of ZFS are designed to be so large that they would NEVER be reached. Data security is on of the best one could reach currently with real data checksumming and selfhealing of errors on access or per scrubbing the entire pool. These extras are not cost-free, they cost performance. A ZFS server with a single disk is usually slower than a Windows NTFS or Linux ext system. But it scales perfectly with the number of vdevs a ZFS pool is build from. So the best rule is, the faster you want your server to be, the more RAM and CPU power you should have. Every free Byte RAM is used for caching and performance without the need of any tuning. A small business server for Office documents and 10 users is well with 4-8 GB RAM and a Dual Core 64 Bit CPU. If you want the same feeling with media files, you may need 10x more RAM.

Pools, Controller and disks

ZFS is designed to use disc-controller without any RAID functionality like pure AHCI SATA or HBA disc-controller, The Raid and Volumemanagement is build into ZFS together with error detection and repairing functions. Knowing that CPU power has increased more heavily than Hardware-Raid power over the last years, this ZFS Raid in Software is often faster than hardware-Raid solutions and avoids the write whole problem of Hardware-Raid. Minimum needed disks are a 20 GB bootdisk (Sata or SSD, never USB!!) and a pair of datadisks to build a datapool from a mirrored vdev (virtual storage device like a Raidset). A pool can grow nearly without limit by adding more vdevs. If you design your pool-layout you can optimize a pool in the triangle capacity, performance and security where mirrored vdevs gives the best performance while Raid-Z gives the best capacity and 3 way mirror and Raid-Z3 the best data-security.

Install Operating System and setup your ZFS Server

1. Download bootable System ISO:

- Illumian (OpenSource), very small server installation, successor of NexentaCore
<http://illumian.org>
- OpenIndiana (OpenSource, I suggest Live Edition, more user friendly, Time-Slider)
<http://openindiana.org/download/>
- Solaris Express 11 (I suggest Live Edition, more user friendly, Time-Slider)
(only free for non-commercial and demo usage
<http://www.oracle.com/technetwork/server-storage/solaris11/downloads/index.html>

Nexenta*/ Illumian supports mirrored boot-drives during installation. With OI and SE11 see:
<http://constantin.glez.de/blog/2011/03/how-set-zfs-root-pool-mirror-oracle-solaris-11-express>

2. Prepare your Boot-disc - set SATA mode to AHCI. With Illumian, you can install your OS as a ZFS mirror on both discs during first setup. With others you can mirror afterwards. USB is not an option with Solaris.

3. Insert Boot-CD/DVD and reboot

Nothing complicated now. You must answer questions about the disk to install on, your language and time-zone. With the text-versions you need to enter only a root password. With the live versions you boot the OS from DVD and install it afterwards to a disc. You must enter an additional user account. Reboot after Setup.

4. Setup your ZFS Server ready to use

Run the napp-it online installer from a console with root permissions. (Working Internet needed)
Text-only installations: login as root and enter from home-directory (look exact at the spaces, the '|' is a pipe character)

If you are using a live edition, login as a user, open a console, get root permissions with su and call the installer.

- `wget -O - www.napp-it.org/nappit | perl`
- `reboot`

Your ZFS server is now ready to use.

5. Open a browser and enter `http://serverip:81`

If you do not know your serverip, go back to console, login and enter: `ifconfig -a`
If you need to manually start/ restart nappit in case of problems: `/etc/init.d/napp-it restart [start,stop,restart]`

Manage your server, create a datapool with menu pool, create datasets (ZFS Folder, partitions) with menu ZFS folder and share these datasets from menu ZFS folder by clicking to the entry below SMB or NFS.

Options:

In the top-level menu, you can enable **edit mode** or **help pages**

NAPP-IT ZFS SERVER

The screenshot shows the napp-it web interface for monitoring SAS2 controllers. A context menu is open over the table, listing actions like Help, Details, Hotswap, Replace, Add, Remove, Volumes, SAS2 extension, Smartinfo, and Diskinfo. The table displays columns for id, sas2_slot, product, smart_sn, smart_health, smart_temp, pool, vdev, state, and busy.

id	sas2_slot	product	smart_sn	smart_health	smart_temp	pool	vdev	state	busy
c6150014EE000A679B6d0	c0:e1:s7	WDC WD3000GLFS-0	WDWXL508008267	PASSED	29 °C	pool3	mirror	ONLINE	configured
c6150014EE000DC32B5d0	c0:e1:s6	WDC WD3000GLFS-0	WDWXL808093323	PASSED	30 °C	pool3	mirror	ONLINE	configured
c6150014EE00276D0E63d0	c0:e1:s4	WDC WD3000GLFS-0	WDWXV1E302T133	PASSED	28 °C	pool3	mirror	ONLINE	configured
c6150014EE05631A41Bd0	c0:e1:s5	WDC WD3000GLFS-0	WDWXL808008473	PASSED	29 °C	pool3	mirror	ONLINE	configured
c6150014EE0AB4F10E2d0	c0:e1:s1	WDC WD3000GLFS-0	WDWXL408044552	PASSED	28 °C	pool3	mirror	ONLINE	configured
c6150014EE0AB511B1Cd0	c0:e1:s3	WDC WD3000GLFS-0	WDWXL508003373	PASSED	28 °C	pool3	mirror	ONLINE	configured
c6150014EE0AB61D458d0	c0:e1:s0	WDC WD3000GLFS-0	WDWXL608047513	PASSED	28 °C	pool3	mirror	ONLINE	configured
c6150014EE0AB86F82Ed0	c0:e1:s2	WDC WD3000GLFS-0	WDWXL808111008	PASSED	30 °C	pool3	mirror	ONLINE	configured

NAPP-IT ZFS STORAGE SERVER

Napp-it is a installer to setup a ready to use ZFS storage server based on OpenSolaris derived Systems like Illumian, OpenIndiana and Solaris 11 together with a Web-UI to manage the server remotely via Web-Browser as a software appliance. Its relation to the underlying OS is similar to other Storage distributions like FreeBSD or Openfiler to their base systems like FreeBSD or Linux. Napp-it is free to use even for commercial use cases but without support. Due to its open modular structure all menus and actions are user editable and expandable from the GUI via browser or external script editors. You can add private menus and actions that are independant from napp-it and napp-it updates. Beside free Add-Ons like the Comstar iSCSI Add-On or the XAMPP Webserver Module, developers can extend napp-it with extensions under their own nonfree licence. Distribution and updates of such extensions can happen together with the regular napp-it distribution like the ACL, monitoring or replication extension or completely independant from napp-it.

Featurelist

- One-Klick Online Installer script to setup a ready to run NAS within minutes for NexentaCore, OpenIndiana and Solaris Express 11
- Concept to include a virtualized ZFS-SAN with NFS datastore into a vmware virtual esxi Server (esxi 4 + SAN + virtual switch, all-in-one)
- Comstar iSCSI Management (LU, Targets with Auth, Target Groups, Portal Groups, Target Portal Groups, FCOE)
- Grouping with ZFS host-host replication (autojob)
- Integrated & editable Helppages
- Jobmanagement (autosnap, autoscrub, email status & alert mails; perl scriptable)
- Pool-Benchmarks via bonnie+
- Property Management of ZFS-folders
- Service Management (Start, Stop, Restart, Status)
- Share Management of ZFS-folders with optimized settings (ex. ACL and properties)
- Snapshot Management (Systemsnaps and Datasnaps, delete based on date or regex)
- System + Network Settings (hostname, Router, DNS, IP etc)
- User and Group-Management (Solaris and SMB)
- User extendable/ editable menus and actions (language: perl)
- ZFS Management: Disks, Pools, ZFS Folders, Raid Z3, Deduplication, unlimited Snapshots without extra space or time-lost
- SSD Hybrid-Storage, Compress, Online Pool-Check with data-refresh =(scrubbing)
- Data Encryption (currently only Solaris 11)

Services

- AFP Fileserver with Time Machine + Avahi Zero Config with extra-installer
- CIFS/SMB Fileserver (Workgroup/ AD-Domain with ACL-Support and Snaps via previous version)
- Comstar iSCSI Server
- Database Server with mySQL
- FTP Server
- Ipf Firewall
- NFS Server V.3 and V.4 (Datastore for Vmware esxi)
- SSH Server
- Webserver with Apache2 + PHP, Perl with extra-installer

Target users

- University and Educational use
- Enterprise use without external support needs
- Developer
- SOHO
- Home

USE CASE: STORAGE SERVER

This is the main use case of a ZFS server and base of all other use cases. The features of a ZFS Storage-Server are quite unique and responsible for the success of ZFS. Datapools that can grow nearly unlimited in size and can reach nearly any wanted performance numbers. Datasecurity at highest levels with data checksums and selfhealing capability on errors, nearly unlimited snapshots without initial space consumption or time delay and very flexible assignment of pool capacity to filesystems. You can setup a file based Storage Server NAS and a blockbased Storage SAN without needed extras or hassles.

NAS FILER

SMB/CIFS. (Solaris Kernel Based server) This is the most used one and available for all platforms like Apple OSX, Linux, Windows, Solaris and Unix. It has a quite good performance and Windows-ACL support. It is fully integrated as a ZFS property so settings keep intact after a pool export/import. You can integrate it in a Windows Active Directory Domain for centralized user-management. ZFS Snapshots are available from Windows folder property "Previous Version". Use it for file-sharing in your LAN.

SAMBA is an alternative to the kernel based SMB/CIFS server. Its also used in the Linux world. It has some more options like nested shares. But in general, the kernel based option is faster, easier to setup and has better Windows integration. SAMBA is not supported by napp-it. If you use, you must setup via CLI.

NFS V3 is a file sharing alternative used in Linux/ Unix environments or for deploying Storage for ESXi or XEN (where it is the preferred option). If you use NFS and SMB in parallel, be aware of problems between the user/ACL based access from Windows and the more host-based restrictions from NFS. It is mostly a little bit faster than SMB. To avoid problems, open the share for everyone or allow root from your NFS hosts.

NFS V4 is also supported and can use ACL's (not used with ESXi).

AFP is the filing protocol used by Apple and is supported via the **netatalk** package. It is quite fast on Macs and needed for Time Machine functionality. I would not use together with other protocols. It adds a lot of files visible to other clients and can irritate therefor. For my own (we have lots of Macs), i use SMB with Active Directory for file-services and netatalk for single Macs and Time Machine (which was really bad compared to ZFS because it gets a Snap by copying files instead of just freezing current state like ZFS).

SAN FILER

Beside file based sharing, you can share blockbased storage either via iSCSI and FC with the Comstar Add-ON. A Raw disk, a Volume or file on a ZFS pool is used like a local disk on your host system.

<http://en.wikipedia.org/wiki/iSCSI> and <http://en.wikipedia.org/wiki/Fiberchannel>

NAME	TYPE	DESCRIPTION	SIZE	UNIT	POOL
c4t1d0	VMware	ssd-write-cache-mirror	5.37 GB	tank	
c4t2d0	VMware	Virtual disk	5.37 GB	tank	

zpool list:

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
rpool	14.9G	6.76G	8.12G	45%	1.00x	ONLINE	-
tank	9.94G	118K	9.94G	0%	1.00x	ONLINE	-

Good to know:

In the triangle of data security, data capacity and performance you can optimize your pool layout, where Raid-Z vdevs give the best capacity, mirrored vdevs the best performance and 3-way mirror or Raid-Z3 the best security.

Optimizing performance:

In general a Raid-Z has the same I/O performance like a single disk (all discs are involved in every read or write), so use mirrors. The overall performance is going with number of vdevs, so use as many as possible.

Read performance

All reads are cached in RAM (ARC) not used by other applications. Use as much as you can afford. You can add a SSD as L2ARC to increase your Read cache. This can improve a lot with many concurrent users or a lot of different reads.

Write performance

All writes are going to RAM. If the writing application does not request sync write that must be confirmed after committed to disk, they are collected in RAM and written together after some time to improve performance. In case of a sync-write request (mostly via NFS or iSCSI) these writes are logged additionally in a log device. This Log is also on your (slow) pool and can reduce write performance dramatically compared no non-sync writes to RAM only. If you need performance you can disable these sync writes at all with the possibility of data-loss in case of problems or you can use a fast log device (SSD or DRAM based). But this log device must be really fast. It must not be greater than 50% the RAM.

CREATE ZFS FOLDERS

ZFS filesystems are often called **ZFS folders, datasets** or volumes or just folders because they are mounted below a pool and look like a folder under your root directory, but they are more. They are completely independent filesystems with their own properties. They are more like partitions on conventional systems and are used to organize the whole pool capacity in a very flexible and effective manner.

You can create a ZFS folder with menu ZFS folder - create.
The menu allows to setup some basic properties and others to useful defaults.

The screenshot shows the napp-it web interface for creating a ZFS folder/dataset. The browser address bar shows `http://172.16.16.46:81/cgi-bin/napp-it/admin.pl`. The page title is "napp-it // network & nas appliance". The navigation menu includes "ZFS Folder" and "Create". The main content area is titled "ZFS folder/ Datasets (independant file-systems)".

basic settings

- pool: tank
- name of new zfs-folder/ filesystem: data
- Encryption: on
- Strenghht (default=aes-128-ccm): aes-128-ccm
- Passphrase min 8 char: [password field]
- repeat Passphrase: [password field]
- use insensitive for smb/win/mac case sensitivity: insensitive

share settings

- smb share: on
- smb guest access: off

other settings

- atime to log last read of files: off
- nbmand (must be off for AFP shares): off

A "submit" button is at the bottom left. A warning message reads: "without this passphrase, you will never be able to reaccess your data, DO NOT FORGET !. -> Use the first characters of the words of a rhyme and include numbers or special characters".

create a dataset under Solaris 11

If you want to use a dataset especially for NAS use, you should just enter a name and keep the other option at defaults.

- optional: encryption: available from ZFS pool version 30 (currently Solaris11 only)
- compression: integrated in ZFS
- casesensitivity: Windows is case insensitive (file "Name"="name"), on Unix you can have different files "name" and "Name". If you want to use it with Windows/ Mac only set it to insensitive, otherwise to mixed
- smb share: share this ZFS folder with Kernel based CIFS Server
- optimize acl: Solaris defaults are very restrictive. Optimized smb-settings will change this to root=full, everybody=modify access
- smb guest: when enabled, you do not need a password to access
- nfs share: NFS is a file-based sharing used especially to share the folder as a SAN datastore for ESXi. If used with ESXi, I recommend to share it in parallel via SMB to have easy access to VM's and VM snaps and easy and fast Copy/ Move/ Clone/ Backup.
- atime: Set to off for performance. Atime will update last read access to a file and can cause performance degrees on high load
- other properties, you can set from menue folder or via CLI

Each dataset has its own set of ZFS properties. They are inherited from the parent dataset or pool. Some of them like casesensitivity can be set only during creation. Most of them can be changed at any time. Per default a dataset can eat the whole pool capacity. You can limit the used space with quotas and ensure an amount of space with reservations.

NFS sharing and SMB/CIFS sharing are also ZFS properties (unlike iSCSI or AFP sharing). If you export/ import a pool, all ZFS properties are kept. Non-ZFS properties are lost after export and import on another system.

If you do a snapshot of a dataset, this snap is stored within the dataset. This is important. If you delete a dataset, all associated snaps are lost.

Share a ZFS folder

Open menu ZFS folder and at the dataset under SMB or NFS to share it. If you have installer AFP Service for Macintosh, you can share a dataset also via AF

The screenshot shows the napp-it web interface with the "ZFS Folder" menu open. The browser address bar shows `http://172.16.16.46:81/cgi-bin/napp-it/admin.pl`. The page title is "napp-it // network & nas appliance". The navigation menu includes "ZFS Folder" and "key". The main content area is titled "ZFS folder/ Datasets (independant file-systems)".

ZFS (all properties)	SMB	NFS	AFP	NBMAND	AVAILABLE	USED	RES	RFRES	QUO	RFQU	SYNC	COMPR	DEDUP	CRYPT	FOLDER-ACL (O,G,E)	SMB-SHARE-all	PERM	RDONLY
backup (pool)	-	-	-	off	178G [7%]	2.48T	none	none	none	none	standard	on	off	off	o=owner_default, g=crea, e=crea	-	777	off
backup/conny	conny=/backup/conny ++ add share	off	off	off	178G [7%]	2.30T	none	none	none	none	standard	on	off	off	e=modi	full_set	777+	off
backup/guenther	guenther=/backup/guenther ++ add share	on	off	off	178G [50%]	176G	none	none	none	none	standard	on	off	off	o=full, e=modi	full_set	777+	off
backup/marco	marco=/backup/marco ++ add share	on	off	off	178G [97%]	4.89G	none	none	none	none	disabled	on	off	off	o=full, e=modi	full_set	777+	off
vm (pool)	-	-	-	off	77.2G [71%]	32.1G	none	none	none	none	standard	off	off	off	o=owner_default, g=readxs, e=readxs	-	755	off
vm/nfs	nfs=/vm/nfs ++ add share	on	off	on	77.2G [71%]	32.0G	none	none	none	none	standard	on	off	off	e=modi	full_set	777+	off

on problems with buffering=on, you can reload list with menu ZFS folder - reload

Only Solaris 11 can share a dataset multiple times under different names. But you can not nest shares (share it under another share) like you can with Windows or SAMBA because there is currently no mechanism to handle different dataset properties. Per default you should start with file properties everybody=modify and root =full or 777 with NFS or AFP (AFP needs always write access to the share itself).

USERMANAGEMENT

Basics of user management

On Linux/Unix user and groups are identified by UID and GID numbers. Permissions are usually based on simple owner/group/everybody settings. Windows use SID (security id numbers) to identify user, groups or computers. Permissions are more fine granular based on ACL's. SMB passwords have a different format than unix passwords and stored separately on Unix in separate files.

This are bad requirements if you want to use a UNIX server to replace a Windows server. In the past; SMB sharing was often a pain due to the different user and pw management and missing or inconsistent ACL - behaviours. Usually you had to define all users on the Unix Machine - even in environments where you have a common user database like MS Active Directory. In a second step you had to define a user mapping between a Windows SID and a Unix UID because of the inability of a Unix filesystem to store ACL based on a Windows SID directly.

With Solaris derived systems this is now handle the problem. Solaris use NFSv4 ACL. They are mostly compatible to Microsoft Windows, more than Posix ACL's used for ex. on Linux. But the main reason of the easyness of a Solaris SMB Server is a fact, that you have a fully Windows compatible built-in user and group management with automated mappings.

For a easy to use SMB-Server based on the Kernel CIFS modul you have to do the following three steps:

1.) After installing napp-it via the online installer, your NAS is ready to use.

You have at least one user (root) with unix passwords. To use root account for SMB, you have to create an additional SMB password with passwd root after setup napp-it the first time at console via passwd root (or passwd „other already created user“) You are then able to mount a SMB share as root with full permissions.

2.) Create a SMB-share from napp-it menu ZFS-folder. If you enable guest=on all users can modify files without entering name+pw and without restrictions.

3.) Connect your Windows with this share \\servername or \\ip
You can now connect as root (and opt. assign ACL for other users from Windows). If you need, create more Solaris users. They can connect after you have set the needed ACL to the shared folder like user or everyone@=any wished permission. If you create the the same user like on your Windows machine (user names are case sensitive) with the same PW, this user should connect without entering name+pw.

Create new users:

smb-groups	about_and_members	SID	option
administrators	(Members can fully administer the computer/domain)	S-1-5-32-544	No members
backup operators	(Members can bypass file security to back up files)	S-1-5-32-551	No members
power users	(Members can share directories)	S-1-5-32-547	No members

Attention: do not set any id-mappings in workgroup mode or you may get problems connected as a user as well as connecting with guest access enabled.

Usermapping

ZFS is a Unix filesystem. It can unfortunately only use Unix UID/GUID as a reference for permission or ACL settings and not Windows SID's directly like a NTFS filesystem. The solution is a extra database to store user-mappings. You can set id-mappings manually via idmap. If a Windows client stores a file or assigns ACL, such a mapping is generated automatically (ephemeral). If you want to set ACL from Solaris, you can use only Windows Users/groups with a known mapping. Acl-extension shows Domain-Users therefore only if the user had accessed a share or the user was used in a ACL assignment from Windows.

idmap dump

shows the relation between Unix UID/GID and Windows SID

```
gsid:S-1-5-2 == gid:2147483653
gsid:S-1-5-32-546 == gid:2147483654
usid:S-1-5-21-3079207246-1797441517-2176294450-501 == uid:2147483651
usid:S-1-5-21-3079207246-1797441517-2176294450-513 == gid:2147483652
usid:S-1-5-21-3952067331-766569302-2624163198-501 == uid:2147483649
..
```

idmap -n

same, but use names for known Windows SIDs

```
wingroup:Network == gid:2147483653
wingroup:Guests@BUILTIN == gid:2147483654
usid:S-1-5-21-3079207246-1797441517-2176294450-501 == uid:2147483651
usid:S-1-5-21-3079207246-1797441517-2176294450-513 == gid:2147483652
winuser:Guest@S11-san == uid:2147483649
wingroup:Domain Users@S11-san == gid:2147483650
usid:S-1-5-21-2156763062-2982440606-2313789320-501 == uid:2147483650
usid:S-1-5-21-2156763062-2982440606-2313789320-513 == gid:2147483651
```

idmap -v

same but shows if the mapping was automatically generated or via a manual idmapping

```
gsid:S-1-5-2 == gid:2147483653
Method: Ephemeral
gsid:S-1-5-32-546 == gid:2147483654
Method: Ephemeral
```

```
usid:S-1-5-21-3079207246-1797441517-2176294450-501 == uid:2147483651
Method: Ephemeral
usid:S-1-5-21-3079207246-1797441517-2176294450-513 == gid:2147483652
Method: Ephemeral
usid:S-1-5-21-3952067331-766569302-2624163198-501 == uid:2147483649
Method: Ephemeral
gsid:S-1-5-21-3952067331-766569302-2624163198-513 == gid:2147483650
Method: Ephemeral
usid:S-1-5-21-2156763062-2982440606-2313789320-501 == uid:2147483650
Method: Ephemeral
usid:S-1-5-21-2156763062-2982440606-2313789320-513 == gid:2147483651
Method: Ephemeral
```

This mapping can become a problem if you import a pool from another server.

ACTIVE DIRECTORY

Your Illumian Nexenta, OpenIndiana or Solaris NAS can either be a member of a Windows Workgroup or a Windows Active Directory domain.

In the first case you authenticate against a local NAS user. Only root has full access per default. If you want to manage a share and set permissions for other users, you must set ACL on your Solaris or remotley from Windows as root. It is also possible to enable guest access. In this case, you can access a share without a password.

If you want to use Active-Directory accounts, your NAS has to join the Active Directory first, (example according to this manual:
http://blogs.sun.com/timthomas/entry/configuring_the_opensolaris_cifs_server)

With napp.it, you just need to open menu services-smb-active directory, fill out the form and join the domain.

The screenshot shows the napp-it web interface for Solaris 11. The browser address bar shows `http://172.16.16.46:81/cgi-bin/napp-it/admin.pl`. The page title is "napp-it 5111 Solaris 11 appliance v. 0.7c preview Jan.27.2012". The navigation menu includes "napp-it", "Help", "Services", "System", "User", "Disks", "Pools", "ZFS Folder", "Snaps", "Comstar", "Jobs", "Extensions", "Add-Ons", and "My menus". The "Services" menu is expanded, showing "AFP", "Apache", "Comstar", "FTP", "MySQL", "NFS", "Rsync", "SMB", "SSH", and "XAMPP". The "SMB" service is selected, and the "Active Directory" sub-menu is open, showing "help", "Active Directory", "check DNS", "check krb5.conf", "user mapping", and "properties". The "Active Directory" sub-menu is selected, and the "properties" sub-menu is open, showing "CID:" and "cmd".

The main content area is titled "SMB/CIFS Service" and contains the following text:

```
Your napp-it Server s111 could be either a member of a workgroup (use local user) or a domain (use domain user).
If you not only want to rejoin a domain, please verify that your domain does not already have a computer named s111
if so, you will get UNSUCCESSFULLY. Please check if smb is online and you have a already shared folder.

Current state of SMB/CIFS Server: online
Current membership: domain Primary : RZ Primary : hfg-gmuend.de Domain controller: lizenzi62.hfg-gmuend.de Local : S111 [5-1-5-21-4205560778-667208428-2415761105]
Primary : RZ [5-1-5-21-2685157207-3384782735-186247241]

Become a member of a [domain]
name of domain ex university.org [napp-it.org]
IP of ad-domain-server ex 192.168.1.1 [192.168.1.1]
NTLM authentication, default 4, on Server 2008 problems try 2
LM authentication level [2]
domainadmin [administrator]
password [.....]

if you are member of a domain, you can use domain users.
in case of problems, read http://wiki.genunix.org/wiki/index.php/CIFS_Service_Troubleshooting submit
```

Messages:

Unsuccesfull: happens mostly if there is already a host with your hostname in your domain. delete this host and retry. Check also DNS and krb5.conf settings.

A domain controller could not be found: check if its online and try all other lmauth settings
Mostly it works with lmauth = 3 or 2

internal error: occurs mostly if SMB service is not running. Currently you can use only one AD-Server with some Solaris versions. If this server is offline, no user can login - even if you have a Backup-AD-Server. You have to manually rejoin.

Even if your NAS is a domain member, you can use local user. If you want to manage your NAS as a domain-user, you have to set the following idmapping via idmap add

```
winuser:domainadmin = unixuser: root
wingroups: domainadmins = unixgroup: root
```

In Workgroup Mode

(also after leaving a Domain) check mappings and delete them all. Never try to map a local Solaris User to another Solaris user)

ACL

Solaris use NFSv4 ACL's. They are more Windows-like than Posix ACL's used on Linux. Together with the Kernel based CIFS Server, you can build a Windows compatible SMB-Fileserver with Windows ACL-support.

If you SMB-share a ZFS-folder via napp-it, these settings are used as default:

Unix permission

- set unix permissions to 777 (or share options will not work)

File and folder based ACL

- root=full access
- everybody@=modify

Share level ACL

- everybody=full access (default value)

If you want to change permissions:

- Open **menu ZFS folder - acl extension - ACL on folders**
- Select a pool and a ZFS folder (dataset). ACLs of this folders are listed
- Add or modify trivial or user/group based ACL of this folder

i suggest to use only allow settings and keep the default inheritance=on settings. They are used for newly created files or directories in a folder

- Modify Share-level ACL (**ACL on shares**)

Like a real Windows server, the Kernel based CIFS-server supports share-level ACL to adjust (default=full) permissions for a share without changing folder or file ACL.

Be aware:

- The kernel **SMB/CIFS server is ACL only**, just like a real Windows server. If you want to set permissions, look only at ACL entries.
- use **trivial ACL** like everyone@, group@ or owner@, they work similar to Unix permissions but differ in inheritance
- SMB passwords are different to Unix passwords. They must be created additional (The reason you must reenter root-pw after setup of napp-it)
- Windows and OI ACL behaves quite similar with one difference. Windows looks first on deny rules, then on allow while OI respect the Order of ACL where the first matching allow or deny entry does the job. From Windows, you cannot set ACL order. Avoid deny rules when not needed.
- If you modify Unix permissions or ACL's the other was reduced to a "compatible" setting

about acl-extension

Base ACL settings like add trivial ACL or user ACL. Other settings mostly needed in enterprises like Domainuser or detailed settings requires registration and are not free. (60 day trial key available). Without the acl-extension you are able to modify ACL via CLI or remote from Windows when connected as root (beware of the order problem)

USE CASE: WEBSERVER

we are using the XAMP package see
<http://www.apachefriends.org/en/xampp.html>

Installation:
wget -O - www.napp-it.org/amp | perl

included is:

- Apache 2.2.21
- MySQL 5.5.16,
- PHP 5.3.8 & PEAR + multibyte (mbstring) support,
- Perl 5.12.1,
- ProFTPD 1.3.3e,
- phpMyAdmin 3.4.5,
- OpenSSL 1.0.0c,
- Freetype2 2.4.2,
- libjpeg 6b,
- libpng 1.2.39,
- zlib 1.2.5,
- expat 1.95.7,
- pdf class 009e,
- IMAP C-Client 2007e,
- OpenLDAP 2.4.23,
- libiconv 1.13.1,
- FreeTDS 0.82,
- libgdbm 1.8.3,
- libxslt 1.1.26,
- mcrypt 2.5.8,
- cURL 7.21.1,
- eAccelerator 0.9.6.1,
- Webalizer 2.21-2,
- libxml 2.7.7,
- mod_perl 2.0.4,
- SQLite 2.8.17 + 3.7.2

Base settings: Menu services - Xampp
Manuals see Xampp.org or apache.org

7. AFP SERVICES FOR MAC'S

Installation of netatalk:
wget -O - www.napp-it.org/amp | perl

OpenIndiana problems:
see <http://napp-it.org/downloads/openindiana.html>

Manuals:
<http://netatalk.sourceforge.net/>

REMOTE MANAGEMENT

1. Local management

If you use a live version, you can manage your server locally via GUI tools or CLI and with the included Firefox browser by calling `http://localhost:81`
If you need your ip-setting, call at console: `ifconfig -a`

Remote management can be done by:

2. IPMI (If supported by your mainboard)

This is something like the content of the local monitor to a remote host. Even Bios settings and cold boot can be done.

3. VNC (If using a live version)

Login and enable via menu System - setting -remote desktop

I use tightVNC client to connect the desktop remotely from Windows

4. Putty (Windows) remote console

If you use a live version, root is only a role and cannot login.

Enable root via napp-it menu services - ssh or login or connect as user and enter `su` to get root permissions

With putty or at a local console you can start **midnight commander**, a quite easy to use file browser and editor

Start with: `mc`

5. WinSCP

If you like to browser your server or edit files, use WinSCP (Windows).

You must enable root via napp-it menu services - ssh or login to connect as root

while the above are mainly base management tools, you can configure your ZFS server via napp-it

6. napp-it Web-UI

Open your browser and call `http://serverip:81`

ILLUMIAN

COMPILING SMARTMONTTOOLS UNDER ILLUMIAN

Update repository
(browse at <http://apt.illumian.org/illumian/pool/main/>)

```
- apt-get update  
- apt-get upgrade
```

```
apt-get install unzip  
apt-get install gcc-46 developer-gcc-46 developer-gcc-44  
apt-get install library-c++-stdc++  
apt-get install make  
apt-get install library-security-openssl-compatibility-0.9.8
```

and set my CC env variable:
export PATH=\$PATH:/usr/gcc/4.4/bin

and was able to run

```
cd $HOME  
wget http://downloads.sourceforge.net/project/smartmontools/smartmontools/5.42/smartmontools-5.42.tar.gz  
tar -xvzf $HOME/smartmontools-5.42.tar.gz  
cd smartmontools-5.42/  
./configure --prefix=/usr --sysconfdir=/etc
```

some warnings but running

10.2 PROBLEMS ON OPENINDIANA

current problems on OpenIndiana: (Bug ?)

You need to reboot after each package that creates a boot environment.
The above wget installer installs napp-it but the former installed packages are missing after a reboot.

you must install them separately everytime followed by a reboot:

```
1. bonny++
pkg install bonnieplus
reboot

2. netcat
pkg install netcat
reboot

3. midnight commander
pkg install netcat
reboot

4. iperf
pkg install iperf
reboot

5. iscsi target
pkg install network/iscsi/target
reboot

6. gcc (needed to compile smartmontools and others)
pkg install gcc-dev
reboot

7. smartmontools
cd $HOME
wget !http://downloads.sourceforge.net/project/smartmontools/smartmontools/5.42/smartmontools-5.42.tar.gz
gzcat $HOME/smartmontools-5.42.tar.gz | tar xf -
cd smartmontools-5.42/
./configure --prefix=/usr --sysconfdir=/etc
make && make install && make clean
reboot

8. opt. netatalk (opt reboot needed after pkg action)
cd $HOME
pkg install SUNWgcc SUNWgmake SUNWxorg-headers SUNWGtk SUNWsfwhea
pkg install pkg:developer/gnome/gettext
echo y | pkgadd -d http://get.opensw.org/how
echo auto | /opt/csw/bin/pkgutil -i CSWlibgcrpt-dev
wget http://download.oracle.com/berkeley-db/db-4.8.30.tar.gz
mkdir -p $HOME/src
mkdir -p $HOME/src/db-build
cd $HOME/src
gzcat $HOME/db-4.8.30.tar.gz | tar xf -
cd $HOME/src/db-build
crle -u -l /usr/local/lib
$HOME/src/db-4.8.30/dist/configure --prefix=/usr/local
make
make install
pkg install system/network/avahi
pkg install service/network/dns/mdns
wget http://sourceforge.net/projects/netatalk/files/netatalk/2.2.2/netatalk-2.2.2.tar.bz2/download
cd $HOME/src
bzcat $HOME/netatalk-2.2.2.tar.bz2 | tar xf -
cd $HOME/src/netatalk-2.2.2
./configure --with-pam --prefix=/usr --sysconfdir=/etc
make
make install

now run the netatalk installer:
wget -O - www.nappit.org/afp | perl
```

10.3 PROBLEMS ON SOLARIS 11

POWER MANAGEMENT

- disk spin down: not working

NEEDED MODIFICATIONS TO RUN NAPP-IT

```
pkg install storage-server
ln -s /usr/lib/libssl.so.1.0.0 /usr/lib/libssl.so.0.9.8
ln -s /usr/lib/libcrypto.so /usr/lib/libcrypto.so.0.9.8
svcadm enable -r smb/server
```

11. LINKS

Some useful Links about ZFS:

excellent summary on practical tips about ZFS
<http://www.nex7.com/readme1st>

newest docs from Oracle
http://docs.oracle.com/cd/E23824_01/ Manuals from Oracle Solaris 11

<http://www.c0t0d0s0.org/archives/663...Sun-Books.html> (recommended books)
<http://hub.opensolaris.org/bin/view/Main/documentation> (overview about docs)
<http://www.c0t0d0s0.org/pages/lksfbook.html> (less known features)
<http://dlc.sun.com/pdf/819-5461/819-5461.pdf> (ZFS manual)
<http://www.solarisinternals.com/> (best source of Howtos)
http://www.solarisinternals.com/wiki...ractices_Guide (best practice guide)
http://www.solarisinternals.com/wiki...Tuning_Guide (best tuning guide)
<http://www.oracle.com/us/products/se...are/031857.htm> (oracle info)
<http://hub.opensolaris.org/bin/view/...fs-server/docs> (cifs server)
<http://www.zfsbuild.com/> (howto and performance tests)
[ten-ways-easily-improve-oracle-solaris-zfs-file-system-performance](http://www.oracle.com/technetwork/server-storage/zfs-1000000-10-ways-to-improve-oracle-solaris-zfs-file-system-performance-10-ways-to-improve-performance.html) (10 ways to improve performance)
[favorite-oracle-solaris-performance-analysis-commands](http://www.oracle.com/technetwork/server-storage/zfs-1000000-favorite-oracle-solaris-performance-analysis-commands-1000000-1.html) (analysis tools)
<http://constantin.glez.de/blog/2011/...ris-11-express> (mirrored bootdrive on Solaris/OpenIndiana)

Help forums, wiki and IRC

<http://hardforum.com/> (**OpenSolaris derived ZFS NAS/ SAN with napp-it**)
<http://forums.servethehome.com/forumdisplay.php?26-Solaris-Nexenta-OpenIndiana-and-napp-it>

<http://echelog.matzon.dk/logs/browse/openindiana/>
<http://echelog.matzon.dk/logs/browse/illumos>

<https://forums.oracle.com/forums/forumID=1321>
<http://wiki.openindiana.org/oi/OpenIndiana+Wiki+Home>
http://wiki.illumos.org/index.php/Main_Page

german ZFS manuals (Handbuch)

<http://dlc.sun.com/pdf/820-2313/820-2313.pdf> (ZFS Handbuch)
<http://www.hardwareluxx.de/community/f101/zfs-stammtisch-570052.html> (dt. Forum)

about silent data corruption (schleichende Datenfehler)

<http://www.solarisday.de/downloads/s...ver0609nau.pdf>
<http://www.fh-wedel.de/~si/seminare/...ionen.html#n17>