Storage Dreamteam

Services that play together best
Something we are asked quite often

published: 2020-Mar-14 (c) napp-it.org

Storage Dreamteam:
A possible functional specification document

Storage Dreamteam

Between small Soho setups and large enterprises, there are many who just look for a storage configuration that allows a single sign on and role based usermananagement, one ore some secure filer for data, backup and secure external access.

A possible functional specification document for many small or mid sized companies or schools and universities may look like

1.) Usermanagement
- Rolebased user management user, groups, groups in groups
- Granular permission management with inheritance to folders for easy settings
- Moving profiles, multiuser accounts on clients
- Single sign on for services like Filer, VPN, Mail, Cloud access

2.) Filer
- Management via Browser (storage appliance)
- High Capacity with option to grow on the fly
- Resilient agains power outages (no corrupt filesystem)
- Protection against Ransomware and silent data errors
- Access restrictions based on Windows ntfs ACL from 1.)
- Readonly versioning (hourly, daily, weekly, monthly or longtime)
- Services FC/iSCSI, NFS and SMB optional others like S3
- Secure Write behaviour for VM Storage and databases
- Encryption per project or department with centralised key management

3.) Backup
- Protection against Ransomware and silent data errors
- Backup and optionally data transfer from filers must be encrypted.
- Inhouse backup (high performance, backup open files, must work under high load with short delays like a few minutes and readonly versioning)
- Backup must preserve file ACL after a restore on another server
- Instant access for filebased restore or rollback.

4.) Internet Access to filer data
- Access based on usermanagement from1.)
- Access exclusively possible after authentication and respects filer authorisation
- Access to selected folders from any filer based on user group or user
- Access protocols sftp, ftps and https with browserbased up/download
- optional 2 Factor Authentication

5.) General
- low technical complexity (Keep it simple)
- „Standard“ configurations, no vendor lockin regarding hardware or software
- manageable costs with support for 1-4

1.) Rolebased Usermanagement

The default solution for this is Windows Active Directory by far. It plays perfect with Windows clients that can be part of the domain security but you can use the userbase for all sorts of authentication demands from filer, mail, webserver, vpn or external access methods on any client OS.

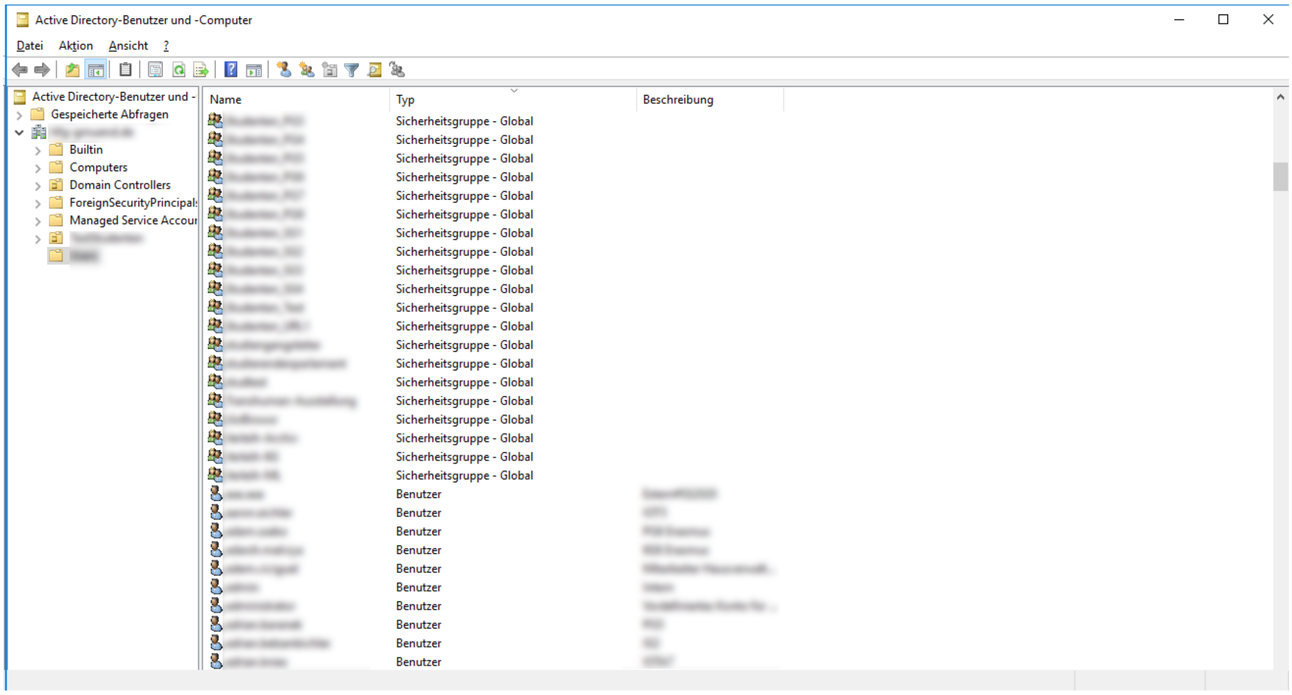

Name	Typ	Beschreibung
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Sicherheitsgruppe - Global	
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]
[redacted]	Benutzer	[redacted]

Costs of a Windows AD Server depends either on number of client computers or number of users. In environments with higher demands on availablity, you want two AD servers, the second as backup and failover if the first is down for maintenance.

You can run a newer Windows server not only on hardware but virtualize ex on ESXi that gives you the comfort bootable snapsots to go back to a state prior last update. On a hardware setup you can use tools like Aomei to make bootable
backups to an external disk or SMB share that you can restore from an Aomei USB stick with Windows PE on it.

Costs: Windows Server 2019 + user or client licenses

1.) Filer

The ultimate solution when it comes to data security and expandability is ZFS. If you use a Solaris based ZFS, you get the best of all ZFS and Windows integration, newest features like encryption and the „it just works". You will hardly find a solution with less trouble on setup and maintenance like bugfixes or setup. Even a re-install after a serious crash is easy. Reinstall the OS, import the pool and optionally napp-it settings and you are on again as Solarish is storage pur with everything what makes ZFS superiour is there even after setup of a minimal ZFS storage distribution like OmniOS.

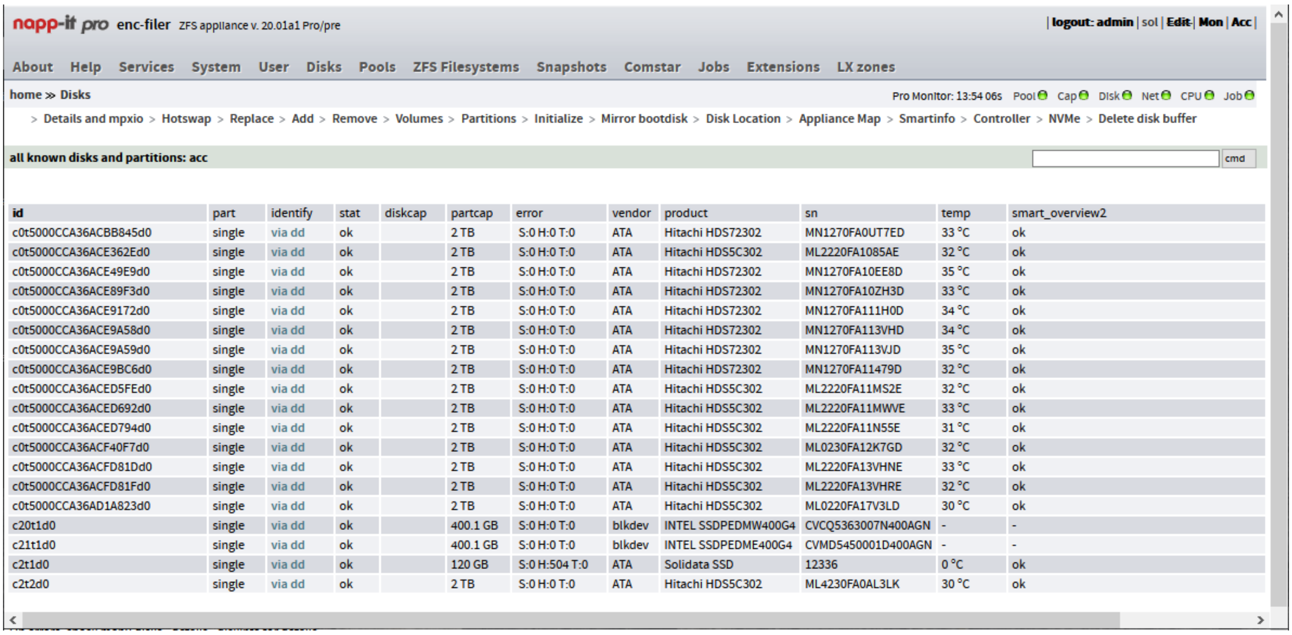
napp-it pro enc-filer ZFS appliance v. 20.01a1 Pro/pre
| logout: admin | sol | Edit | Mon | Acc |
About Help Services System User Disks Pools ZFS Filesystems Snapshots Comstar Jobs Extensions LX zones
home ≫ Disks
Pro Monitor: 13:54 06s Pool Cap Disk Net CPU Job
> Details and mpxio > Hotswap > Replace > Add > Remove > Volumes > Partitions > Initialize > Mirror bootdisk > Disk Location > Appliance Map > Smartinfo > Controller > NVMe > Delete disk buffer
cmd

all known disks and partitions: acc

id	part	identify	stat	diskcap	partcap	error	vendor	product	sn	temp	smart_overview2
c0t5000CCA36ACBB845d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA0UT7ED	33 °C	ok
c0t5000CCA36ACE362Ed0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML2220FA1085AE	32 °C	ok
c0t5000CCA36ACE49E9d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA10EE8D	35 °C	ok
c0t5000CCA36ACE89F3d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA10ZH3D	33 °C	ok
c0t5000CCA36ACE9172d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA111H0D	34 °C	ok
c0t5000CCA36ACE9A58d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA113VHD	34 °C	ok
c0t5000CCA36ACE9A59d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA113VJD	35 °C	ok
c0t5000CCA36ACE9BC6d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS72302	MN1270FA11479D	32 °C	ok
c0t5000CCA36ACED5FEd0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML2220FA11MS2E	32 °C	ok
c0t5000CCA36ACED692d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML2220FA11MWVE	33 °C	ok
c0t5000CCA36ACED794d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML2220FA11N55E	31 °C	ok
c0t5000CCA36ACF40F7d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML0230FA12K7GD	32 °C	ok
c0t5000CCA36ACFD81Dd0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML2220FA13VHNE	33 °C	ok
c0t5000CCA36ACFD81Fd0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML2220FA13VHRE	32 °C	ok
c0t5000CCA36AD1A823d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML0220FA17V3LD	30 °C	ok
c20t1d0	single	via dd	ok		400.1 GB	S:0 H:0 T:0	blkdev	INTEL SSDPEDMW400G4	CVCQ5363007N400AGN	-	-
c21t1d0	single	via dd	ok		400.1 GB	S:0 H:0 T:0	blkdev	INTEL SSDPEDME400G4	CVMD5450001D400AGN	-	-
c2t1d0	single	via dd	ok		120 GB	S:0 H:504 T:0	ATA	Solidata SSD	12336	0 °C	ok
c2t2d0	single	via dd	ok		2 TB	S:0 H:0 T:0	ATA	Hitachi HDS5C302	ML4230FA0AL3LK	30 °C	ok

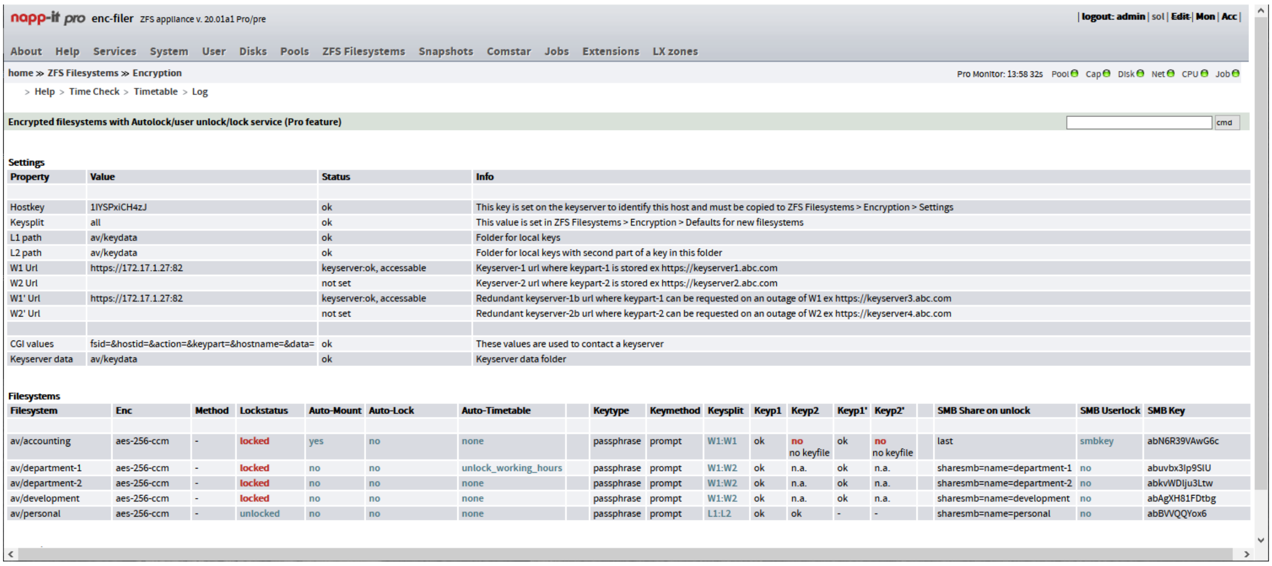
napp-it pro enc-filer ZFS appliance v. 20.01a1 Pro/pre
| logout: admin | sol | Edit | Mon | Acc |
About Help Services System User Disks Pools ZFS Filesystems Snapshots Comstar Jobs Extensions LX zones
home ≫ ZFS Filesystems ≫ Encryption
Pro Monitor: 13:58 32s Pool Cap Disk Net CPU Job
> Help > Time Check > Timetable > Log
cmd

Encrypted filesystems with Autolock/user unlock/lock service (Pro feature)

Settings

Property	Value	Status	Info
Hostkey	1IYSPxiCH4zJ	ok	This key is set on the keyserver to identify this host and must be copied to ZFS Filesystems > Encryption > Settings
Keysplit	all	ok	This value is set in ZFS Filesystems > Encryption > Defaults for new filesystems
L1 path	av/keydata	ok	Folder for local keys
L2 path	av/keydata	ok	Folder for local keys with second part of a key in this folder
W1 Url	https://172.17.1.27:82	keyserver:ok, accessable	Keyserver-1 url where keypart-1 is stored ex https://keyserver1.abc.com
W2 Url		not set	Keyserver-2 url where keypart-2 is stored ex https://keyserver2.abc.com
W1' Url	https://172.17.1.27:82	keyserver:ok, accessable	Redundant keyserver-1b url where keypart-1 can be requested on an outage of W1 ex https://keyserver3.abc.com
W2' Url		not set	Redundant keyserver-2b url where keypart-2 can be requested on an outage of W2 ex https://keyserver4.abc.com
CGI values	fsid=&hostid=&action=&keypart=&hostname=&data=	ok	These values are used to contact a keyserver
Keyserver data	av/keydata	ok	Keyserver data folder

Filesystems

Filesystem	Enc	Method	Lockstatus	Auto-Mount	Auto-Lock	Auto-Timetable		Keytype	Keymethod	Keysplit	Keyp1	Keyp2	Keyp1'	Keyp2'		SMB Share on unlock	SMB Userlock	SMB Key
av/accounting	aes-256-ccm	-	locked	yes	no	none		passphrase	prompt	W1:W1	ok	no no keyfile	ok	no no keyfile		last	smbkey	abN6R39VAwG6c
av/department-1	aes-256-ccm	-	locked	no	no	unlock_working_hours		passphrase	prompt	W1:W2	ok	n.a.	ok	n.a.		sharesmb=name=department-1	no	abuvbx3Ip9SIU
av/department-2	aes-256-ccm	-	locked	no	no	none		passphrase	prompt	W1:W2	ok	n.a.	ok	n.a.		sharesmb=name=department-2	no	abkvWDIju3Ltw
av/development	aes-256-ccm	-	locked	no	no	none		passphrase	prompt	W1:W2	ok	n.a.	ok	n.a.		sharesmb=name=development	no	abAgXH81FDtbg
av/personal	aes-256-ccm	-	unlocked	no	no	none		passphrase	prompt	L1:L2	ok	ok	-	-		sharesmb=name=personal	no	abBVQQYox6

Install: easy
Maintenance: easy, perfect AD integration with ntfs alike ACL and Windows sid
Disaster backup: not needed (or via replication of current bootenvironment)
Costs hardware: A standard storage server from SuperMicro, Dell, HP etc
OmniOS: Opensource with a commercial support option (500$/year), opt. Solaris
napp-it: Free version and Pro versions up from 120Euro/ year, Cluster optional

3.) Backup

If your filer is ZFS you can use ZFS replication. This is a datastream based method based on snaps and not a file compare. This means that you can backup even open files and keep a high load Petabyte server in sync with a backup system over ethernet down to a minute delay.

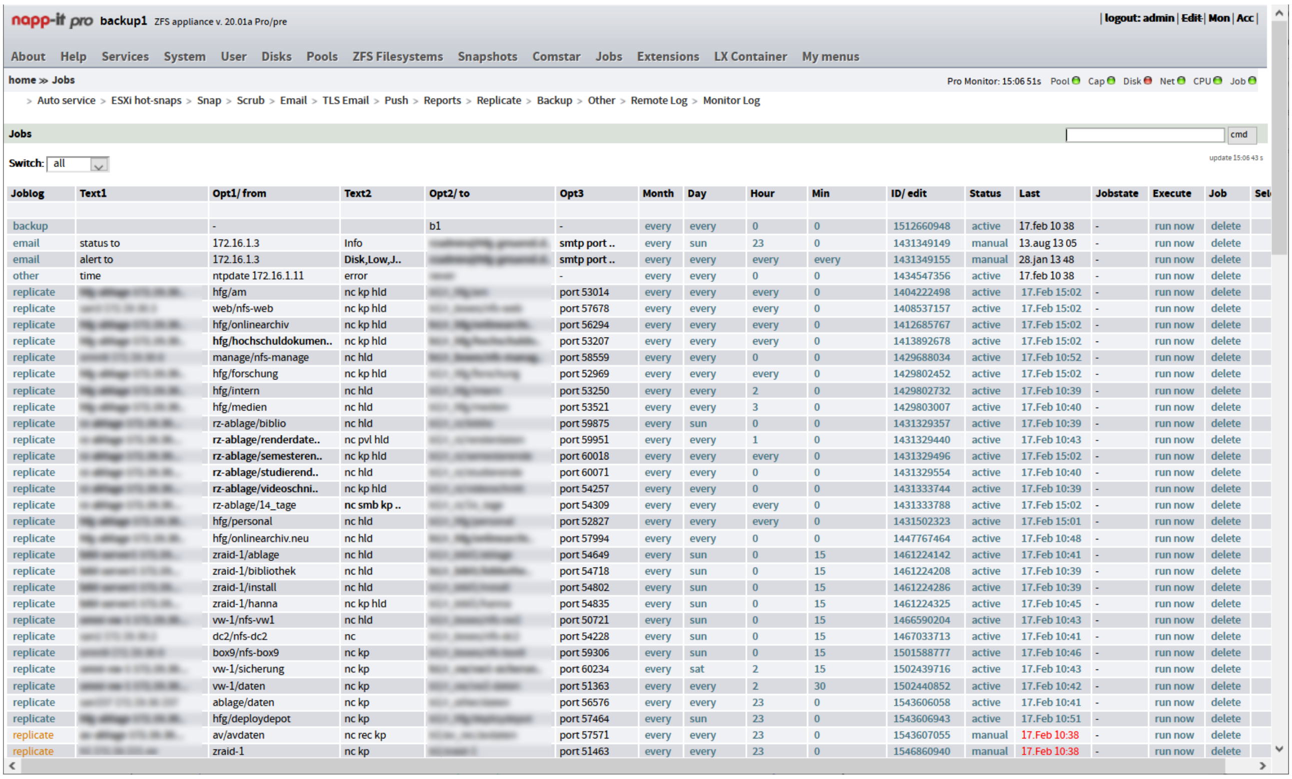

napp-it pro backup1 ZFS appliance v. 20.01a Pro/pre
| logout: admin | ~~Edit~~ | Mon | Acc |
About Help Services System User Disks Pools ZFS Filesystems Snapshots Comstar Jobs Extensions LX Container My menus
home ≫ Jobs
Pro Monitor: 15:06 51s Pool Cap Disk Net CPU Job
> Auto service > ESXi hot-snaps > Snap > Scrub > Email > TLS Email > Push > Reports > Replicate > Backup > Other > Remote Log > Monitor Log
Jobs
cmd
update 15:06 43 s

Switch: all

Joblog	Text1	Opt1/ from	Text2	Opt2/ to	Opt3	Month	Day	Hour	Min	ID/ edit	Status	Last	Jobstate	Execute	Job	Sel
backup		-		b1	-	every	every	0	0	1512660948	active	17.feb 10 38	-	run now	delete	
email	status to	172.16.1.3	Info	[redacted]	smtp port ..	every	sun	23	0	1431349149	manual	13.aug 13 05	-	run now	delete	
email	alert to	172.16.1.3	Disk,Low,J..	[redacted]	smtp port ..	every	every	every	every	1431349155	manual	28.jan 13 48	-	run now	delete	
other	time	ntpdate 172.16.1.11	error	[redacted]	-	every	every	0	0	1434547356	active	17.feb 10 38	-	run now	delete	
replicate	[redacted]	hfg/am	nc kp hld	[redacted]	port 53014	every	every	every	0	1404222498	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	web/nfs-web	nc kp hld	[redacted]	port 57678	every	every	every	0	1408537157	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	hfg/onlinearchiv	nc kp hld	[redacted]	port 56294	every	every	every	0	1412685767	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	hfg/hochschuldokumen..	nc kp hld	[redacted]	port 53207	every	every	every	0	1413892678	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	manage/nfs-manage	nc hld	[redacted]	port 58559	every	every	0	0	1429688034	active	17.Feb 10:52	-	run now	delete	
replicate	[redacted]	hfg/forschung	nc kp hld	[redacted]	port 52969	every	every	every	0	1429802452	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	hfg/intern	nc hld	[redacted]	port 53250	every	every	2	0	1429802732	active	17.Feb 10:39	-	run now	delete	
replicate	[redacted]	hfg/medien	nc hld	[redacted]	port 53521	every	every	3	0	1429803007	active	17.Feb 10:40	-	run now	delete	
replicate	[redacted]	rz-ablage/biblio	nc hld	[redacted]	port 59875	every	sun	0	0	1431329357	active	17.Feb 10:39	-	run now	delete	
replicate	[redacted]	rz-ablage/renderdate..	nc pvl hld	[redacted]	port 59951	every	every	1	0	1431329440	active	17.Feb 10:43	-	run now	delete	
replicate	[redacted]	rz-ablage/semesteren..	nc kp hld	[redacted]	port 60018	every	every	every	0	1431329496	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	rz-ablage/studierend..	nc hld	[redacted]	port 60071	every	every	0	0	1431329554	active	17.Feb 10:40	-	run now	delete	
replicate	[redacted]	rz-ablage/videoschni..	nc kp hld	[redacted]	port 54257	every	every	0	0	1431333744	active	17.Feb 10:39	-	run now	delete	
replicate	[redacted]	rz-ablage/14_tage	nc smb kp ..	[redacted]	port 54309	every	every	every	0	1431333788	active	17.Feb 15:02	-	run now	delete	
replicate	[redacted]	hfg/personal	nc hld	[redacted]	port 52827	every	every	every	0	1431502323	active	17.Feb 15:01	-	run now	delete	
replicate	[redacted]	hfg/onlinearchiv.neu	nc hld	[redacted]	port 57994	every	every	0	0	1447767464	active	17.Feb 10:48	-	run now	delete	
replicate	[redacted]	zraid-1/ablage	nc hld	[redacted]	port 54649	every	sun	0	15	1461224142	active	17.Feb 10:41	-	run now	delete	
replicate	[redacted]	zraid-1/bibliothek	nc hld	[redacted]	port 54718	every	sun	0	15	1461224208	active	17.Feb 10:39	-	run now	delete	
replicate	[redacted]	zraid-1/install	nc hld	[redacted]	port 54802	every	sun	0	15	1461224286	active	17.Feb 10:39	-	run now	delete	
replicate	[redacted]	zraid-1/hanna	nc kp hld	[redacted]	port 54835	every	sun	0	15	1461224325	active	17.Feb 10:45	-	run now	delete	
replicate	[redacted]	vw-1/nfs-vw1	nc hld	[redacted]	port 50721	every	sun	0	15	1466590204	active	17.Feb 10:43	-	run now	delete	
replicate	[redacted]	dc2/nfs-dc2	nc	[redacted]	port 54228	every	sun	0	15	1467033713	active	17.Feb 10:41	-	run now	delete	
replicate	[redacted]	box9/nfs-box9	nc kp	[redacted]	port 59306	every	sun	0	15	1501588777	active	17.Feb 10:46	-	run now	delete	
replicate	[redacted]	vw-1/sicherung	nc kp	[redacted]	port 60234	every	sat	2	15	1502439716	active	17.Feb 10:43	-	run now	delete	
replicate	[redacted]	vw-1/daten	nc kp	[redacted]	port 51363	every	every	2	30	1502440852	active	17.Feb 10:42	-	run now	delete	
replicate	[redacted]	ablage/daten	nc kp	[redacted]	port 56576	every	every	23	0	1543606058	active	17.Feb 10:41	-	run now	delete	
replicate	[redacted]	hfg/deploydepot	nc kp	[redacted]	port 57464	every	sun	23	0	1543606943	active	17.Feb 10:51	-	run now	delete	
replicate	[redacted]	av/avdaten	nc rec kp	[redacted]	port 57571	every	every	23	0	1543607055	manual	17.Feb 10:38	-	run now	delete	
replicate	[redacted]	zraid-1	nc kp	[redacted]	port 51463	every	every	23	0	1546860940	manual	17.Feb 10:38	-	run now	delete	

A napp-it backups server pulls data. This means that you can use a single backup system for more than one filer. Backup can be encrypted. You can replicate an unencrypted file-system to a encrypted destination. You can even replicate a encrypted locked filesystem in raw mode (OmniOS).

Configuration: similar to filer
Costs hardware: A standard storage server from SuperMicro, Dell, HP etc
OmniOS: Opensource with a commercial support option (500$/year), opt. Solaris
napp-it: Pro complete or Pro replication (up from 120 Euro/year)

4.) External Internet access

If you want external Internet access, security must be the main concern especially if personal data is involved. This means that anonymous access ex by a link like most cloud solutions are offering is forbidden outside personal use. In an environment with Windows AD and Windows or Solaris/ OmniOS filers as AD members it is not as easy to find a solution that integrates well into the AD concept and the permissions on filers. Most Linux and Cloud offerings does not.

A quite perfect solution is the Titan SFT server. This is a Windows software that offers SFTP, FTPs and browser access via https (Upload/download). Titan offers virtual directories from all filers based on AD groups or users what allows different levels of access. All Windows ACL are respected. For an additional security, webdrive and 2 FA is optional.

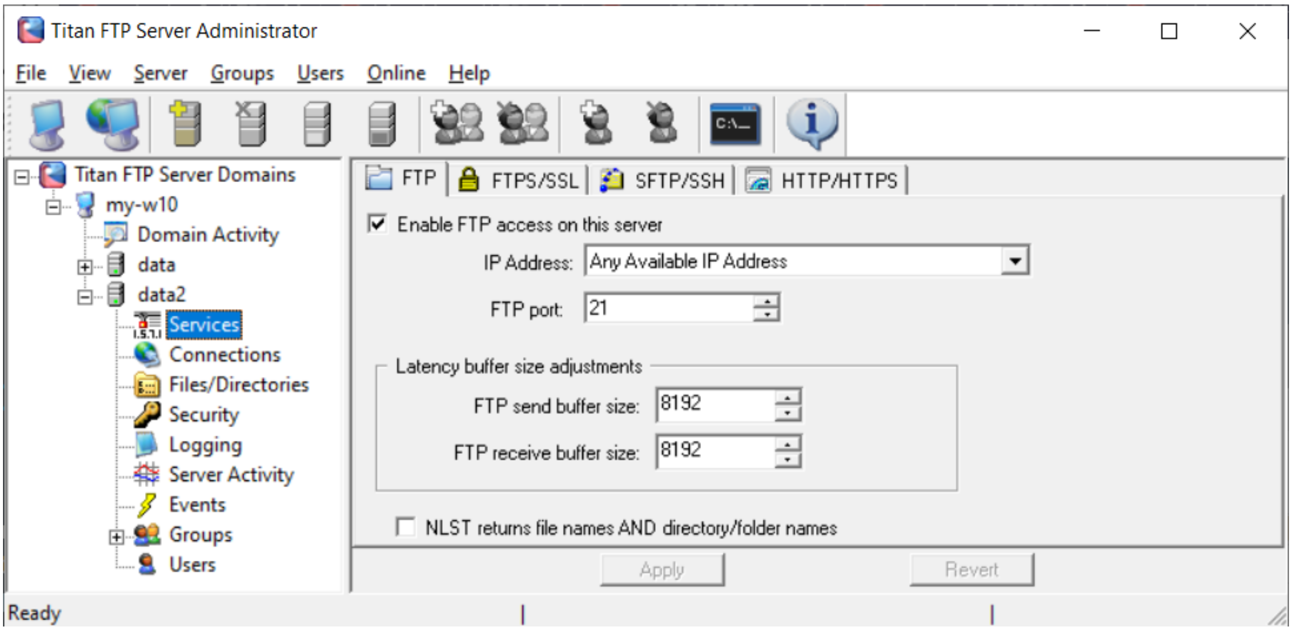

☑ Enable FTP access on this server

IP Address: Any Available IP Address

FTP port: 21

Latency buffer size adjustments

FTP send buffer size: 8192

FTP receive buffer size: 8192

□ NLST returns file names AND directory/folder names

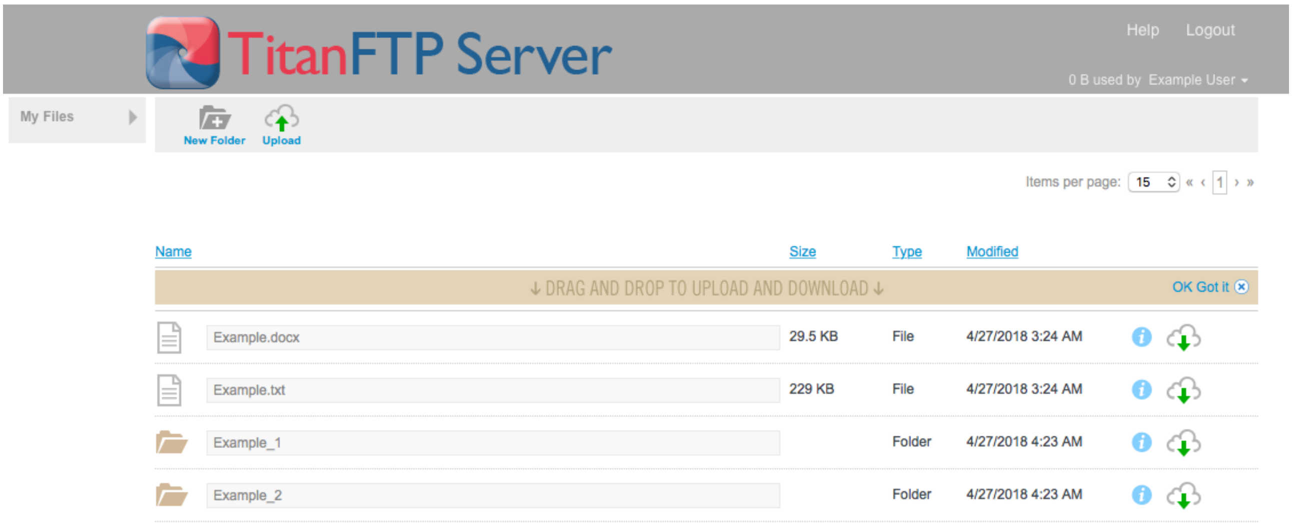

Name	Size	Type	Modified
Example.docx	29.5 KB	File	4/27/2018 3:24 AM
Example.txt	229 KB	File	4/27/2018 3:24 AM
Example_1		Folder	4/27/2018 4:23 AM
Example_2		Folder	4/27/2018 4:23 AM

Costs:
- Windows Server +
- Titan SFT/with webaccess: around 2000$ + 500$ support/year

5.) General demands

5.1 low technical complexity (Keep it simple)

If you can setup a AD Windows Server you will have zero problems tro setup a OmniOS or Solaris filer with napp-it, see http://www.napp-it.org/doc/downloads/setup_napp-it_os.pdf, a backup system or the Titan sftp/https server.

You can even use a single Windows Server with Titan without AD. This is ok for a few users. You only need to create users in the Windows Server and same users with same passowords on OmniOS.

5.2 „Standard" configurations, no vendor lockin regarding hardware or software

Windows is the default on dektops and in many server setups when it comes to authentication via Active Directory.

From hardware, you can use your preferred vendor. I prefer Supermicro as this gives more options than any other vendor and you can re-buy parts even after a few years.

5.3 Summary

A solution with AD Server, filer, backup and external access is from hardware in a range up from say 700 $ per system without disks (total around 3000 USD minimum). With 19" systems, even with a Petabyte capable system you are in the range 10k - 20k $ from hardware.

Software costs are in the range of a few thousand USD initial costs with additional annual support as an option.

6. Light Alternative: Amazon S3 compatible sharing option

If you do not need superior rolebased access methods with fine granular access restrictions but a reliable and ultra fast way to share a folder in the internet or to offer a cloud based (optionally inhouse) backup option with a name/password combo per filesystem/ user/project or want to share files anonymously via a link, look at S3 cloud services. You can add an Amazon S3 compatible sharing option via minIO for each ZFS filesystem easily to a basic OmniOS filer with a few clicks, see https://forums.servethehome.com/indexphp?threads/amazon-s3-compatible-client-server-minio.27524/

S3/miniO cloud sharing is supported in napp-it 19.12 and 20.dev