napp-it

napp-it cs
Client Server Edition

Storageserver Web-GUI for Windows Storage Spaces and ZFS

published: 2024-Jul-28 (c) napp-it.org

Howto

1. About napp-it cs

Since around 2008, napp-it SE (Solaris Edition) was the first feature complete and easy to manage ZFS Web-Gui for (Open)Solaris and NexentaCore, then for the free Solaris fork Illumos with Open-ZFS around OpenIndiana and OmniOS. Now napp-it SE is ported to napp-it CS (Client/Server) to manage nearly any ZFS Server or Serverfarm on Free-BSD, Illumos, Linux, OSX, Solaris or Windows. Support for the upcoming ZFS on Windows requires integration in the Windows Storage Spaces concept to allow a seamless cooperation.

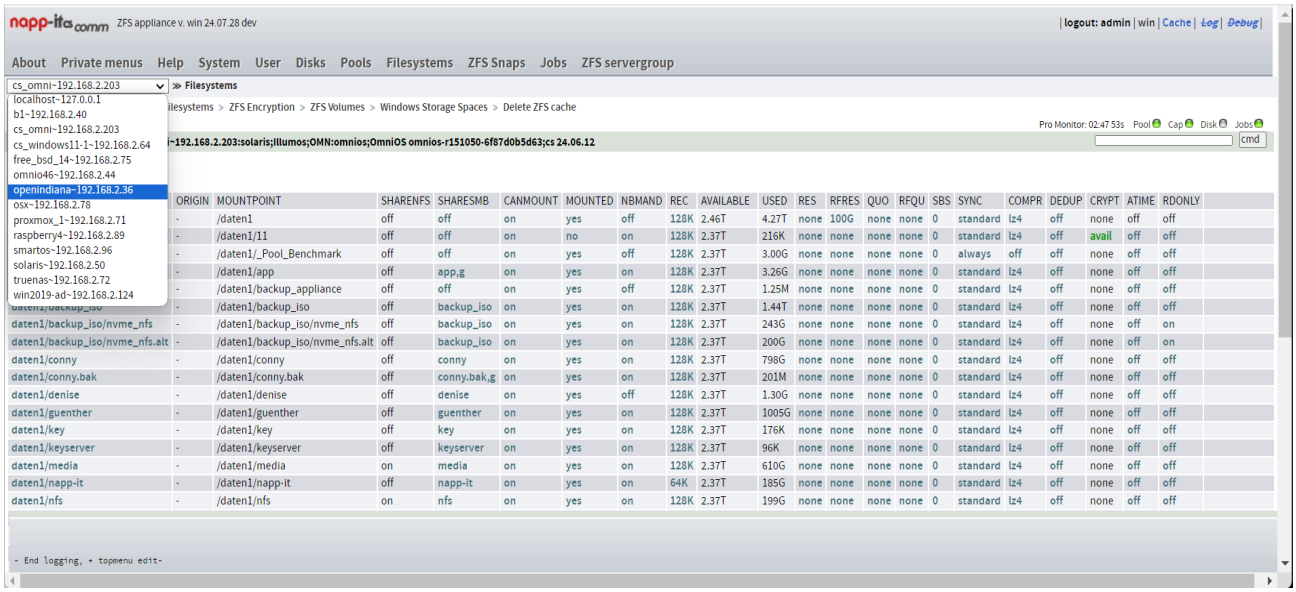

Initially only a functioning ZFS in a Windows environment was intended but then we decided to add full Storage Spaces support to enable Windows as a NAS option for everyone, not only the skilled Microsoft Storage Experts.

Napp-it cs is not only a solution to manage a single Windows or ZFS server.
It was build to manage Serverfarms with ZFS on any OS

1.1 Start napp-it CS web-gui on Windows

There is no setup. Just download https://www.napp-it.de/doc/downloads/xampp.zip
and unzip the xampp folder to c:\xampp. Then start the web-gui (mouse right click as admin):
„C:\xampp\web-gui\data\start_server_only_as_admin.bat“ and open a Browser at https://localhost

2. Windows Storage Concept

Windows was for years known for weak software raid concepts compared to Linux or Unix. This is why many are still using hardwareraid on Windows while on Linux/Unix superiour modern software raid like btrfs or ZFS is common. In a current Windows (desktop or server edition) there is not only the weak/old software raid available but also the modern Storage Spaces concept that allows a flexible pooling of disks (any type, any size) with realtime data redundancy despite disks of different sizes or automatical data tiering between HDD, SSD or SCM/NVMe storage Tiers. This is a flexibility far above normal realtime raid including ZFS. With the new ReFS filesystem, this flexbility is paired on Windows with realtime checksums and Copy on Write ideas from ZFS and Hyper-V, one of the best Hypervisors.

Why is everyone talking about btrfs or ZFS webmanaged storage appliances and ESXi/ Proxmox VM servers instead the Windows machine you already own with Storage Spaces and Hyper-V that anyone knows and is able to manage?

Like Sun with ZFS in 2008, MicroSoft has a quite superiour storage technology but is not able to reach the mass market. While Sun had a strong focus with ZFS on „big enterprise“ that failed to achieve enough revenues, the Windows problem seems simply weak usability with high complexity and lack of easy or full featured management tools outside the Server editions and a missing web-gui for easy remote management. We can change that with more than 15 years of experience in webbased storage management. The Best: You can decide when to use Storage Spaces and when the upcoming Open-ZFS on Windows that has reached a quite final release candidate state.

3. Windows Storage Pools

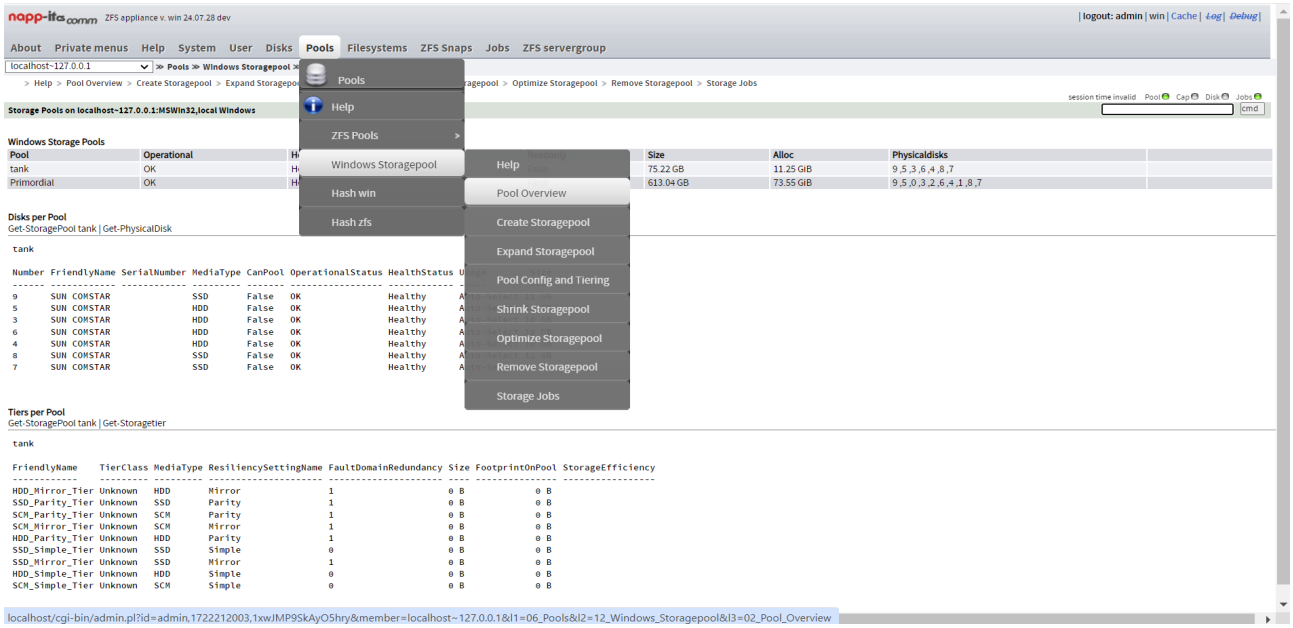

Unlike a classic Raid array or ZFS Pool, a Windows Pool is not the place where data is organized. It is just like a basket where you throw in whatever you have, Harddisks, SSD, NVMe or SCM disks. As there is no redundancy on Pool disks, size, type or number of disks does not matter in a first step. To create a Windows Storage Pool, just assign any of your disks to one or more Storage Pools.

To destroy a pool, just delete all Storage Spaces, then the Storage Pool.

4. Windows Storage Disks

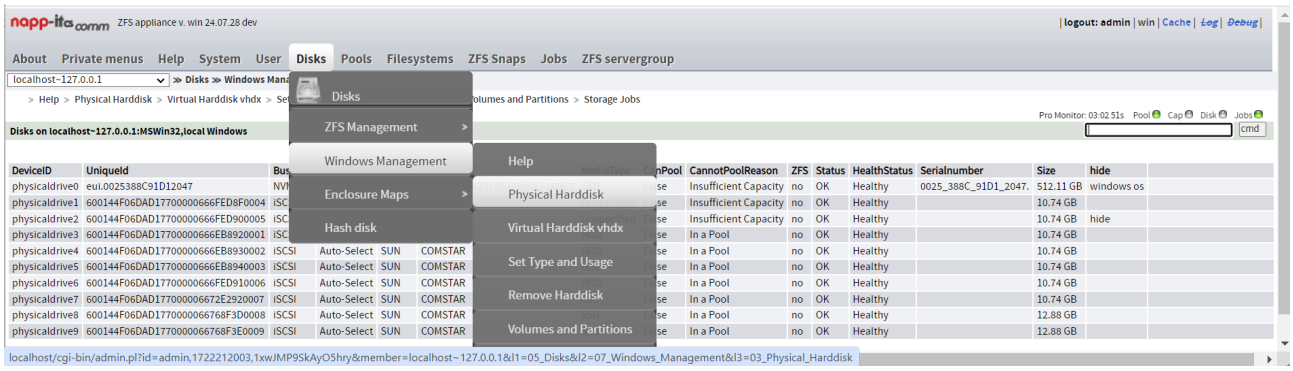

Windows uses „Physical Disks" as members of a Storage Pool. These are not only „real" disks like HDD, NVMe or SCM but also „Virtual HardDisks" like iSCSI targets or „Virtual VHDX Harddisks" that come from Hyper as high performance filebased devices with a size up to around 64000GB that can be located locally or remote on SMB. You can set the Mediatype HDD, SSD or SCM per „Physical Disk". This allows to place Storage Spaces on disks of a dedicated performance level or a Tiered Storage Space where dedicated/hot/newest/often requested data is automatically placed on the fast storage Tier and cold/old/not so often requested on the slower Tier. You can use three Tiers (HDD, SSD, NVME/SCM). If you use remote Virtual VHDX Harddisks over SMB direct (requires Windows Server), you can build a high performance Storage Cluster over Nodes even with desktop Windows as nodes.

You can add new disks at any time to a Pool. To remove a disk, set it to state „retired", then remove if Resiliency allows a removal. You can also use disks for Journal or Hotspare. A Writecache can be defined per Virtual Disk on a faster MediaType disk.

5. Windows Storage Spaces

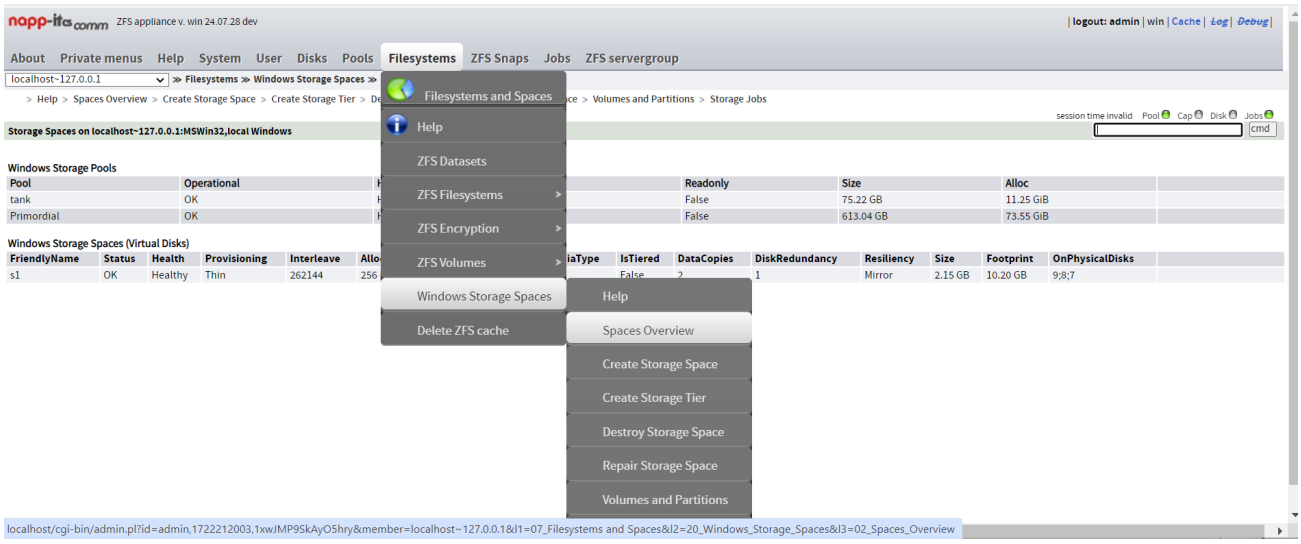

A Storage Space is a thin or thick provisioned „Virtual Disk" with a flexible size that you can create on a Storage Pool with an expandable size. You can create Storage Spaces with different redundancy levels ex Simple, 2way Mirror, 3way Mirror, Parity or Dual Parity, on a dedicated media type or a tiered Space over different media types with automatic tiering. Realtime Redundancy (Resiliency) is not based on Disk Raid but Data Copies on one or more disks. When you create a Storage Space with NTFS or ReFS, you can set AUS (Allocation Unit Size). The default 4K is slow but space efficient with small files. For a good performance you should use larger values ex 64K with ReFS or 64-512K with NTFS, read https://wasteofserver.com/storage-spaces-with-parity-very-slow-writes-solved/

6. Windows Storage Tiers

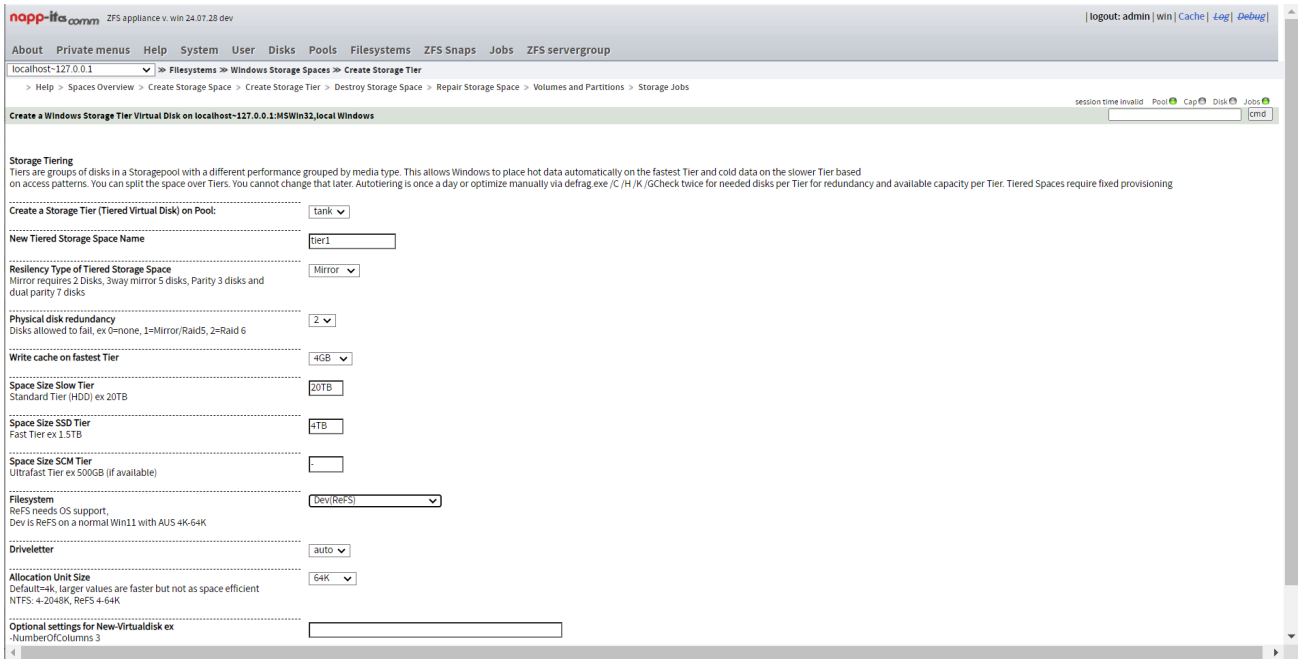

A Storage Tier is a special Storage Space that you can span over disks of different media types with an automatical move of files between slower and faster Tiers. To use Tiers you must set a media type per disk and you must prepare the Storage Pool by creating Tiers for different Resilency levels.

7. Windows Volumes and Partitions

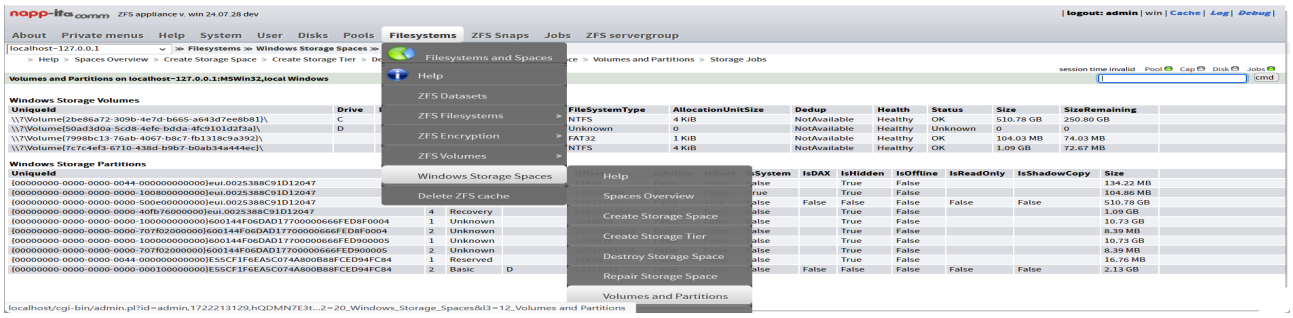

A partition is a logical division of a disk (physical or virtual). A volume is a logical assembly of one or more partitions which the operating system knows how to use as a mass storage container. When you create a Storage Space a partition and a volume is created anf formatted with ntfs or ReFS (Dev-Drive on Windows 11)

8. Windows Filesystems

On Window you can use different filesystems each with other advantages or disadvantages.
You can place a partition and a volume with any filesystem on a physical disk, a virtual disk or a virtual harddisk,
With the help of Storage Spaces even with tiering, thin provisioning or SSD write caches

FAT filesystems are the oldest. They are quite restricted in size and not so robust.
If you just unplug an USB stick with FAT during write, the filesystem is mostly damaged.
Advantage: They are perfect to move some data between any systems
Disadvantage: Low reliability

NTFS (New Technology FileSystem) is the default since Windows NT. It is quite robust and good for large disks.
If you just unplug an USB stick with NTFS during write, the filesystem is mostly damaged.
Advantage: very stable and bugfree, many features like Compress or encryption
Disadvantage: Lack of state of the art features like realtime checksums or Copy on Write

ReFS (Resilent FileSystem) is a modern filesystem with Copy on Write and realtime checksums similar to ZFS.
If you just unplug an USB stick with NTFS during write, the filesystem remains mostly good due Copy on Write.
Advantage: more robust as NTFS
Disadvantage: not as mature as NTFS, partly incompatible versions, still not bootable, lacks some NTFS features

ZFS (Zetabyte FileSystem, last word regarding Filesystems, 128bit with checksums and Copy on Write). It combines Filesystem, Volume and Raid Management with a superiour feature set above ReFS like realtime dedup and compress, encryption per filesystem, sync write to protect a rambased writecache, Hybridpools based on small/large data, superior Raid concepts over multiple striped vdevs, Draid or Z3 where three disks are allowed to fail in a Parity Raid.
Advantages: Offers everything you want to protect even Petabytes of data, widely used on any OS, Replication
Disadvantage: Still beta on Windows (very near to be usable, ZFS is stable, Windows OS integration needs some work)

9. Open-ZFS vs NTFS/ReFS

In the moment ZFS on Windows reaches a state where you can use it with a similar stability to Open-ZFS on BSD, Illumos or Linux, you mostly prefer ZFS over NTFS/ReFS especially with large arrays where you want ZFS snaps and ZFS replication to keep even high load, high capacity filesystems in sync with a short delay. Large arrays with hundreds of disks are even easier to manage than a usual Windows fileserver with a few disks. Outside ZFS you miss Snaps, sync write protection in software (similar to hardware raid with BBU).

But ZFS still lacks real data Tiering, flexible integration of different disks (in size and mediatype) and is slower than NTFS or ReFS. A Windows SMB Direct Fileserver with 40-100Gb nics ex foe multiple user videocut is much faster than a ZFS solution. On Windows you can have both, select based on use case.

10. SMB multichannel and RDMA

You mainly use a storage server as fileserver via iSCSI when you need blockstorage (remote disks). If you need multi-user fileaccess or authorisation and authentication SMB is the matter of choice. SMB is originally the sharing protocol from Windows so any SMB server should behave like a Windows server for best compatibility.

If you need a good performance, you can use any SMB server with nics better 1G (2.5G, 5G or 10G). If possible use one or more 10G nics or nics 20-100G. This allows a performance comparable to local disks over network.

Beside SMB version (1,2,3) and support for Windows ntfs ACL permissions there are two SMB features that are performance relevant.

With a 10G nic, expect 300-500 MB/s network performance with some tunings and Jumbo frames up to 900 MB/s. In a multiuser/ multithreaded environment a single threaded SAMBA SMB server can limit against a Windows Server or multithreaded Solaris kernelbased SMB server. For really high performance demands > 1000 MB/s you need SMB multichannel with SMB direct/ RDMA.

10.1. SMB multichannel

SMB Multichannel enables file servers to use multiple network connections simultaneously. It facilitates aggregation of network bandwidth and network fault tolerance when multiple paths are available between the SMB 3 client and the SMB 3 server

10.2 SMB direct

SMB direct additionally allows supported networks adaptors to have Remote Direct Memory Access (RDMA) capability. This can reduce latency and CPU usage. If supported by both client and server, SMB direct can greatly increase performance.

SMB direct is an alternative to ip connections . For SMB direct, use SMB direct capable nics >= 20G, a Windows Server edition as SMB server (up from a free Server 2019) and Windows 11 as SMB clients. SMB direct/RDMA on SAMBA is not ready.

https://learn.microsoft.com/en-us/windows-server/storage/file-server/smb-direct

Any comments?
Send an email to support@napp-it.org