How to recover a failed Storage Spaces

ReclaiMe Storage Spaces Recovery User Manual

Contents

Overview

Storage Spaces concepts and fundamentals

Storage Spaces, a new capability introduced by Microsoft to improve the utilization of disk space, uses a range of terms which an operator carrying out Storage Spaces recovery should be familiar with.

Pool

Pool is a set of physical drives, on top of which you can create various virtual disks, known as *spaces*.

Slab

Slab is a quantum of disk space of 256 MB. All the physical drives in the pool are "sliced" into slabs which are then given to the Storage Spaces driver on request. The driver then assigns slabs to a particular virtual disk (known as *space*).

Space

A virtual disk created from a set of slabs on the physical drives of the pool is called *Space*. The main characteristics of a space are:

- **layout** - there are simple, mirror, three-way mirror, and parity spaces;
- **type** - fixed or thinly provisioned;
- if a single logical volume is created on the space, the **filesystem type** (ReFS or NTFS) is also relevant.

Fixed vs. thinly provisioned spaces. Advertised capacity vs. physical capacity.
A space can be either fixed or thinly provisioned.

In a fixed provisioned space, the storage capacity for a space is allocated immediately when the space is created. Even if you do not use all the space capacity for the data, allocated disk space on the physical drives is still reserved for the fixed provisioned space and cannot be given to any other space.

With thin provisioning, space capacity can exceed the amount of physically available disk space in the pool because disk space for such a space is allocated on an as-needed basis. The total capacity of a thinly provisioned space is called *advertised size*. If at some time disk space is exhausted, a thinly provisioned space is automatically taken offline until user either adds more physical drives to the pool or frees disk space by deleting data from the other spaces of the same pool.

Thus, a space has two sizes - advertised capacity and physical capacity which depending on the provisioning type used can either be the same (fixed provisioning) or different (thin provisioning).

Simple, mirror, three-way mirror, and parity layouts

When creating a space, you should specify what type of redundancy it will utilize. Note that Microsoft uses the term *resiliency* instead of more common *redundancy*.

Simple means no redundancy; slabs are distributed among the drives in the pool in a way similar to a technique used in RAID0. If any drive assigned to such a space is missing, the entire space becomes unrecoverable.

In the **mirror** storage layout data from each slab is duplicated on some other drive of the pool. A mirrored space, similar to a RAID1, survives a single disk failure.

The **three-way mirror** storage layout can only be created on no less than five physical drives. In such a space there are three copies of each slab stored on different pool member disks. A three-way mirrored space survives a failure of up to two physical drives.

In a **parity** space, as in a RAID 5, data and parity information are striped across pool member disks. Similar to a RAID5, parity spaces survive a single disk failure.

Main components and failure modes in Storage Spaces

Storage Spaces consists of three main components:

1. **Filesystems on virtual spaces**. Strictly speaking, a filesystem has nothing to do with Storage Spaces, but a filesystem can fail as well. So it is needed to distinguish between filesystem failure and failure of a Storage Spaces pool or a virtual space.

2. **Physical drive identifiers**. Storage Spaces writes a unique identifier to each physical drive of the pool and then uses it to identify the drive. Drive identification is required, say, when the drives are moved between SATA ports, or between USB and SATA or vice-versa. Each drive stores two copies of its own identifier.

3. **Storage Spaces database** which contains a map defining a physical location of each slab in each virtual space. The map refers to physical drives only using their unique identifiers rather than SATA port numbers or disk serial numbers. Each drive in the pool stores an exact copy of the database.

Case I - File system failure and recovery

Case I deals with filesystem failure when underlying virtual disks (called spaces) work as they should. In this case you need to use regular data recovery tools which will operate in the same way as with a regular hard disk.

If the filesystem fails, regular data recovery software should be able to provide recovery solution. However, you may face the following problems:

- Abuse of "thin provisioning". People often create virtual disks of ridiculous sizes, using much larger sizes than available physical disk space is. Usually data recovery tools scan all the advertised volume capacity without considering "thin provisioning". This can lead to unexpectedly long scan time.
- The size of storage pool may easily reach dozens of terabytes. Since all data recovery tools only copy data to a new storage device, you will need free space of the same size as the original pool. Additionally, the more capacity of the pool, the more time copying will take.
- If the pool comes from Windows Server, it may use ReFS filesystem. As of January 2013, ReFS data recovery is not very well understood and there are few tools available on the market for it (although we do have one).

Case II - Partial Storage Spaces metadata failure and recovery

Failure of Storage Spaces metadata when maps of spaces, defining where a particular slab is physically located, are in good condition, but identifiers of physical drives are lost. In this case it is enough to match physical drives to their corresponding identifiers and then use existing maps.

In case of the Storage Spaces metadata failure, all drive identifiers or some of them may be lost. The possibility and time to find missing identifiers depends on:

- the total number of physical drives,
- the number of physical drives with no identifiers,
- the number and type of virtual spaces,
- certain patterns of data placement when the "thin provisioning" is used.

Case III - Complete Storage Spaces metadata failure and recovery

Complete loss of Storage Spaces metadata is when maps of spaces are lost. To rebuild the map you need to try *Deep scan* procedure. However, it is time consuming and not very accurate.

System requirements and hardware preparations

Regardless of the Storage Spaces failure type you will need:

- A lot of free disk space. You need at least as much as there was data on the Storage Spaces pool.
- A lot of SATA ports. You need at least as many as there were disks in the pool. In the worst case you can use USB, but should be aware that it will be significantly slower and somewhat less reliable.
- Internet connection. Part of the analysis is done on the server so a PC running ReclaiMe Storage Spaces Recovery must have a connection to the Internet.

- Plenty of memory. The exact requirements depend on the recovery type.

With the increasing complexity of recovery, the system requirements rise as well. For any type of Storage Spaces failure it will be better to use 64-bit (x64) Windows rather than 32-bit (x86) Windows.

Case I - Slightly damaged database

In this case ReclaiMe Storage Spaces Recovery displays the spaces as *Good* or *Readable*. You can then save such a space to the appropriate image file using options of the *Save* button or launch ReclaiMe File Recovery software - click *File recovery* button - to perform file-by-file recovery on the space.

Saving disk image files

Depending on the type of the image file, you need either as much disk capacity as the data stored on the space or as much free disk capacity as the advertised size of the space.

File-by-file recovery

If you launch ReclaiMe File Recovery to retrieve files and documents from the space, system requirements are quite typical: dual core CPU and at least 4 GB RAM.

Case II - Finding lost drive identifiers

Find drives procedure requires at least 128 MB RAM per each drive of the pool. Additionally, in some cases the speed of the analysis is proportional to the number of cores; anyway, the overall speed shouldn't be too slow even if you have just two cores. However, you probably should not launch the *Find drives* procedure on a laptop.

Case III - Reconstructing a totally damaged Storage Spaces database

In this case you need at least $(2 + N/2)$ GB RAM for N drives in the failed pool. Additionally, the optimal core/disk ratio is one CPU core for every two or three drives from the pool. Practically, however, you don't need more than eight CPU cores.

You also need free disk space of about 3-5 GB per each drive of the pool to store the saved-state files of ReclaiMe Storage Spaces Recovery software.

Overall workflow

Detected spaces can be in one of the following states:

- **Good**, if all of the physical drives in this space are present and recognized. You can immediately copy the *good* spaces or recover files from them. This is a *Case I* recovery.

- **Readable**, if some physical drives are missing but you still can copy the space content thanks to redundancy. *Readable* spaces, just like *Good* spaces, can be immediately saved or transmitted to ReclaiMe File Recovery for file-by-file recovery. With *Readable* parity spaces, expect that the overall performance will be worse as compared to *Good* parity spaces. This is because parity calculation is required to compensate for the lack of the drives. Note that simple spaces cannot be in *Readable* state since they are not redundant. This is still a *Case I* recovery.
- **Missing drives**, if more drives are missing than can be reconstructed based on redundancy. You should try *Find drives* procedure to find the missing drives. This is a *Case II* recovery.

In addition, if Storage Spaces metadata is severely damaged, the entire pools or spaces may be missing. In this case, the slab maps have to be reconstructed, making it *Case III* recovery.

Storage Spaces recovery flowchart

```
                        ┌─────────────────────┐
                        │   Select drives     │
                        │   Click Start       │
                        └─────────────────────┘
                                  │
                                  ▼
              ┌──────────────────────────────────────────┐
              │      Examine list of pools and spaces     │
              └──────────────────────────────────────────┘
         │                        │                        │
      [Case I]                [Case II]                [Case III]
         ▼                        ▼                        ▼
  Spaces listed as        Spaces listed as         Spaces or pools
  Good or Readable        Missing drives           missing (not listed)
         │                        │                        │
         │                        ▼                        │
         │                  Use Find drives                │
         │                 │            │                  │
         │                 ▼            ▼                  │
         │        Spaces now listed   Spaces still listed  │
         │        as Good or Readable as Missing drives    │
         │                 │            │                  │
         │                 │            └──────┐           ▼
         │                 │                   └──► Use Deep scan
         │                 ▼                          │
         └────► Use Save config or Recover files ◄────┘
```

| Actions |

Observations
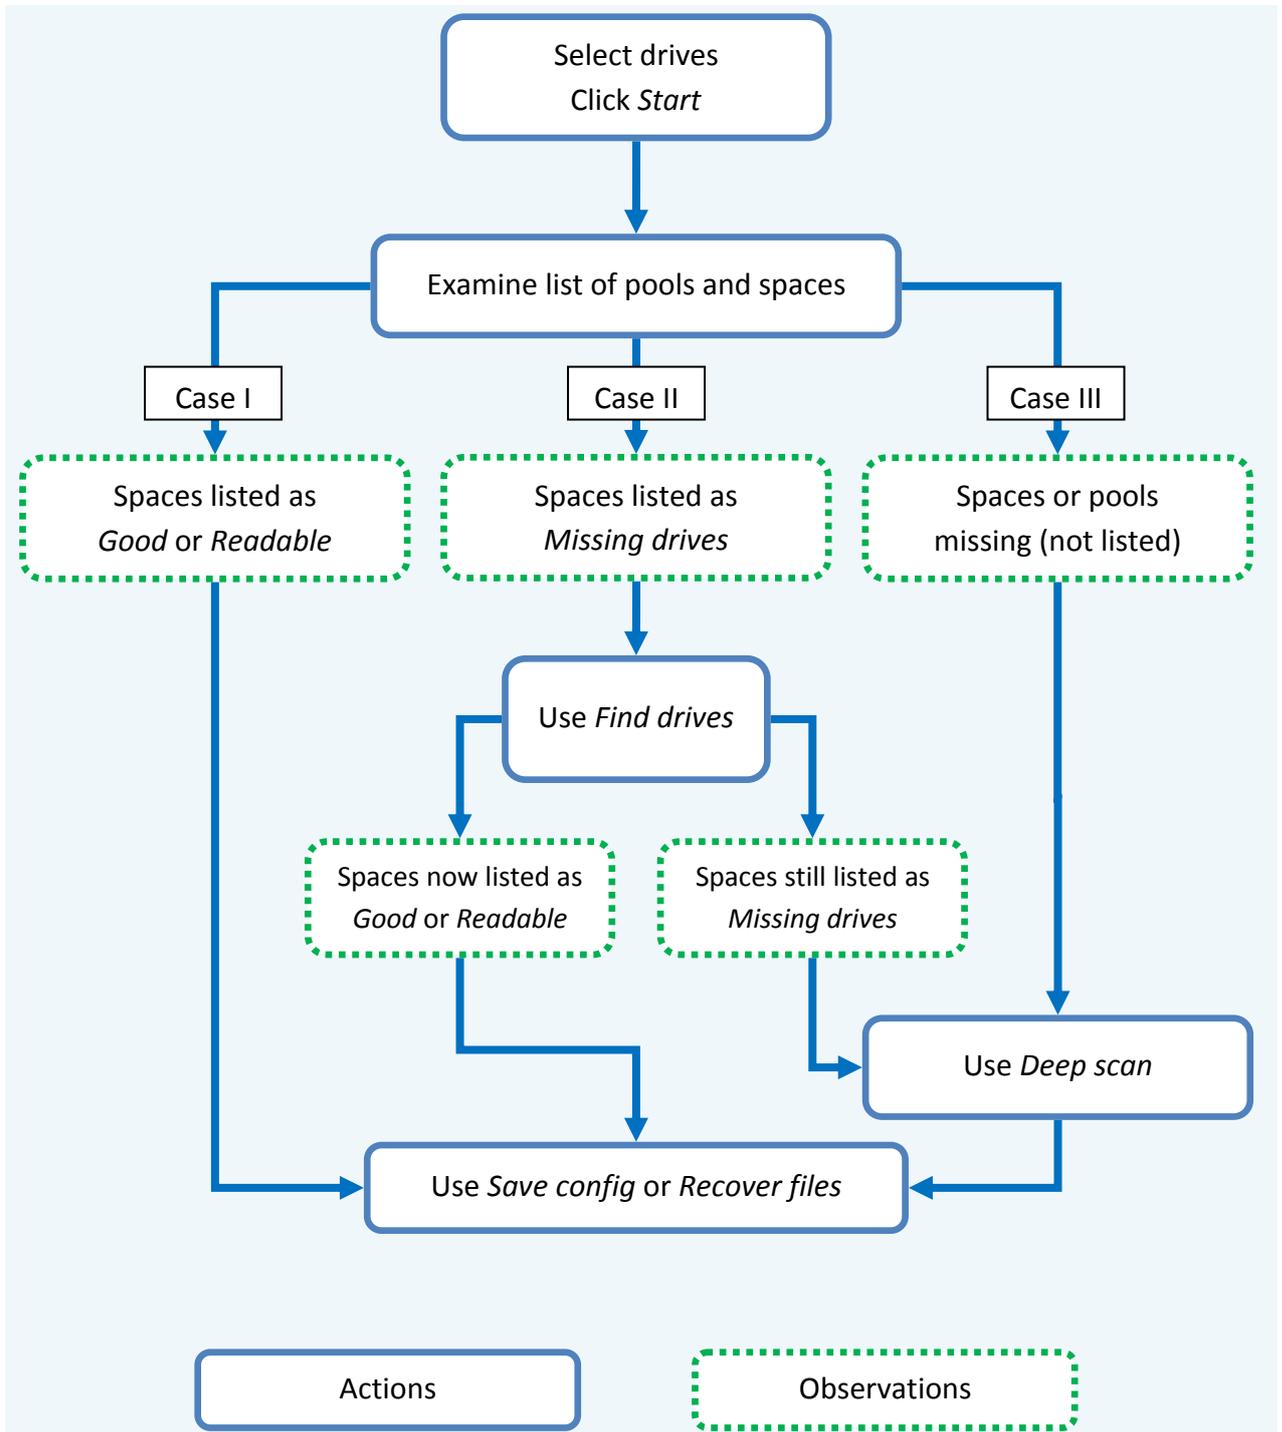
Performing the analysis

Case II - finding drive identifiers

To start searching for matches between physical drives and their identifiers click the *Find drives* button. The *Find drives* procedure works only with those drives that you selected on the first screen before clicking the *Start* button.

Limitations

The *Find drives* procedure can be executed only if the pool has at least one fault-tolerant space - mirror (regular or three-way) or parity. In some cases it is required to find all the missing drives, while in other cases part of the drives may be left missing. How many drives can be missing depends on a particular pattern according to which slabs are placed in the space.

Performance considerations

Speed of the *Find drives* procedure depends mainly on the pool configuration and to a lesser extent on CPU power of the PC running Storage Spaces Recovery. The more drives to find and the less CPU power, the more recovery time. Typically it takes about several minutes; the worst-case scenario is up to a couple of hours. If the procedure does not complete within three hours, then no more result can be produced for a given set of disks, and you should abort the process.

If during the *Find drives* procedure the Internet connection to the server is lost, the analysis will stop. All the data accumulated by the time of disconnection will be saved. Once the connection to the server is re-established, you can restart the analysis by clicking *Find drives* again and it will pick up where it left off.

For faster speed, try not to specify the drives which don't belong to the pool and especially not to specify the drives from another pool.

If matching fails, nothing can be improved within the limits of case II and you should try case III algorithms.

If this didn't work out

If you have already tried the *Find drives* on a particular set of drives and this didn't help (needed spaces have not changed to *Good* or *Readable*) the *Find drives* button becomes disabled and all you can do next is to try the *Deep scan*.

Case III - a totally damaged Storage Spaces database

The *Deep scan* procedure tries to match each slab of physical drives to a particular space on the pool. Like *Find drives*, the *Deep scan* procedure works only with those drives that you selected on the first screen before clicking the *Start* button. Note that you need to

specify a folder to place intermediate data to. This folder must be empty and there must be free disk space of 5 GB per drive.

Performance considerations

The overall speed depends greatly on CPU power and, to a lesser extent, on the internal structure of the pool (number of spaces, slabs allocation and etc.).

In any case, it is required to read and analyze data from all the drives in the pool. At average read speed of approximately 60 MB/s it takes ten hours or so just to read data from a 2 TB drive. Analysis time depends on the following:

- how many drives are in the original pool;
- the total size of the drives in the pool;
- how many signals for the analysis are found on the drives.

The estimated time of scanning and analysis depends on the density of data as well. Some data requires more complex processing and therefore more time to complete the analysis. If ReclaiMe Storage Spaces Recovery runs into area with a high data density, the estimated time may increase - this is normal.

Space content considerations

The quality of recovery depends greatly on the type of the data stored on the drives and also on the filesystems used on the pool spaces. Generally, ReFS-formatted space is better than NTFS-formatted one in terms of recovery.

If the quality of recovery is unacceptable, contact us and we will investigate the case.

Multiple spaces in a single pool

Typically, the more spaces in the pool, the more difficult the recovery task.

Note that *Deep scan* applied to the mirrored spaces brings several separate simple spaces - two for a regular mirror and three for a three-way mirror.

Deep scan output

When *Deep scan* is complete you will see spaces which ReclaiMe Storage Spaces Recovery was able to recover. Each space is characterized by the following:

- **Layout** - can be either simple or parity. If your pool contained a mirror space (or three-way mirror), look for its identical parts among simple spaces.
- **Contains starting slabs** - if a space contains first slabs, it is a good sign that an entire space has decent quality. First slabs store information about various important characteristics of space such as what partitions this space contained.

- **Logical size** - correlates with the latest slab which the tool was able to recover for this space. In case of thinly provisioned spaces, the logical size should be close to the advertised size.
- **Allocated size** - total size of all slabs reconstructed for a given space. Generally, it shows what maximum data capacity can be extracted from the space.

Typically, logical size of a space is larger than its allocated size meaning that either it is a thinly provisioned space or part of the slabs has not been identified.

You can sort the spaces according to any of the above parameters by clicking the appropriate column header.

Saving and loading *Deep scan* progress

To load the saved-sate files you need to click a downward arrow next to the *Deep scan* button and specify the folder with the saved-state files. Note that loading takes up to several minutes for an average size pool.

Once the saved-state files are loaded, ReclaiMe Storage Spaces Recovery starts to analyze the loaded data which means that you will see the progress bars indication again but this time the overall speed will be much faster. Certain data still has to be reprocessed even when the saved-state files are used. This causes a noticeable delay between Scan 1 and Scan 2 during load.

Note that even if you are going to recover Storage Spaces using the saved-state files, you still need to specify an empty folder to store the saved-state files from the current run. The saved-state files are being saved on an ongoing basis and it is impossible to turn the creation of these files off. This is by design.

Copying data. Output formats.

XML file

You can save the configuration of the recovered space to an XML file which then can be loaded to ReclaiMe File Recovery software to extract files and folders from the space.

Image files

ReclaiMe Storage Spaces Recovery can save a sector-by-sector image file of the space.

Raw sector-by-sector image file

- can be very large;
- doesn't know about thin provisioning and therefore requires you to provide as much free disk space as the advertised capacity of the space;
- can be loaded into any data recovery software;

- can be copied to a physical device.

File recovery

Clicking the *Recover files* button launches ReclaiMe File Recovery software transferring the pool parameters and immediately loading them.

Note that if ReclaiMe File Recovery has not been yet installed on your PC, you are offered to download and install it.

Licensing

ReclaiMe Storage Spaces Recovery is licensed on a per case basis. You can get the most current license terms and pricing at

www.storage-spaces-recovery.com/buynow.aspx