What do I really need to consider when setting up my NAS?

When creating a NAS (Network-Attached Storage), be it for the home or a small or home office environment, it is the initial choices made during its conception that are most crucial. Specifically, the choices regarding the RAID implementation and the drives used will have a lasting impact on the NAS that, once chosen, require a complete reconfiguration and loss of data to change.

Understanding the interdependencies of the triad of RAID, NAS and drives is critical. At the same time, it is easy to get lost in the minutiae of disk drive parameters that will have little impact on the final NAS implementation. Here we guide our readers through the forest of acronyms and provide advice on disk drive options for an optimal NAS solution.

Getting the initial choices right

NAS is all about storage, and providing plenty of it. In the home, the goal is probably to provide backups for family laptops, PCs and handheld devices, as well as music and movies. In the small business space, users are looking for backing up files, emails and providing shared storage to allow employees to share files with one-another. They may also want to have their own platform from which they can share files with off-site partners, clients and customers.

The NAS hardware will by physically limited to support for a specific number of bays – slots to insert hard drives – so there need to be enough bays to handle the quantity of drives you intend to use. Our goal here is large and reliable capacity, so spinning platters, rather than SSDs, are the choice here in 3.5” format. Such drives can provide up to 14TB storage for minimal outlay and, with spinning platter technology being improved all the time, it is likely that capacities will continue to grow in the future.

Small NAS systems targeting home users will offer as few as two bays. Small and medium sized businesses (SMB) will probably require something with at least four or five bays, while enterprise solutions are offered in a 19” rack format with space for 12 to 16 drives. It is the combination of available bays and the features each RAID (Redundant Array of Independent Disks) option offers that will ultimately determine the RAID approach selected.

Understanding RAID options

Prior to selecting the NAS system to be used, it is critical to review the RAID level you wish to implement. Once selected, changing it requires a reconfiguration the NAS that is coupled with a loss of the data stored. Additionally, it will not be possible to move the drives to a different NAS system and retain the contents of the drives.

The RAID level determines three key factors: the quantity of storage of the final system, the ease with which

Toshiba Electronics Europe GmbH
Hansaallee 181
40549 Düsseldorf, Germany

Further information
toshiba-storage.com

you can recover from a single drive failure, and how many drives will be required. The RAID level is typically written as RAIDn, where n denotes the level.

RAID0 provides lots of capacity without any redundancy using a method known as 'striping'. Put quite simply, the available capacity is the sum of the capacity of the drives in use; two 2TB drives provides 4TB of storage. There is, however, risk here since, should one of the drives fail at some point, everything on that drive will be lost forever.

RAID1 provides reliability by using a method known as 'mirroring'. This uses two drives and, as the name suggests, everything that is written to the first drive is also copied to the second. Thus, two 2TB drives provides 2TB of storage. However, should a drive fail, the NAS can still provide the stored data from the healthy drive.

For home use the likelihood of drives failing simultaneously is very low. The drive that fails can also be exchanged for a new drive and the NAS will assimilate it into the system. Many NAS solutions allow the replacement drive to be of higher capacity than the original as, years after original creation of the NAS, drives of the original capacity may no longer be available. This capability also allows the total capacity to be increased later should storage become tight.

RAID5 and RAID6 require more disks to support their redundancy approaches, provide more storage compared to the number of disks used than RAID1, but less total storage than RAID0. These RAID approaches spread single pieces of user data across several drives. As well as storing the user's data, an extra piece of information, known as parity data, is also stored. Should one disk fail, either the original data is still available or the parity information enables the NAS to calculate the missing information from the failed drive.

This results in a minimum of three drives being required for RAID5, and a minimum of four for the implementation of RAID 6. Once the failed drive is replaced its contents can be rebuilt. RAID5 can accommodate up to one drive failure, while RAID6 can survive the failure of two drives. As an example, three 2TB drives provide 4TB of RAID5 storage, while four 2TB drives provide 4TB of RAID6 storage.

Another aspect to consider is the system state should a drive fail. If a RAID5 drive fails, it cannot withstand a further drive failure. Bearing in mind that rebuilding the system after replacing the failed drive can take several days, and that it places significant disk access load on the remaining drives, should a further drive fail during the rebuild process all the data will be lost. Thus, RAID6 really provides the most robust approach for users with the highest reliability needs.

It should also be noted that the calculation of the parity information leads to an impact on write performance for RAID5 and RAID6. However, for SMB implementations it is unlikely that this will be noticeable in day-to-day use.

If the NAS provides eight or more bays, further RAID approaches are possible. These combine the capacity of RAID0 with the robustness of RAID1, RAID5 or RAID6. For example, RAID10 mirrors two RAID0 implementations, with eight 2TB drives providing 4TB of capacity with the robustness of RAID1. RAID50 and RAID60 provide data robustness coupled with the capacity striping of RAID0.

Selecting suitable hard drives

With the NAS hardware selected, the number of bays known, and the RAID approach chosen, the next step is to source the hard drives themselves. If you take the time to review their datasheets the list of parameters to

Toshiba Electronics Europe GmbH
Hansaallee 181
40549 Düsseldorf, Germany

Further information
toshiba-storage.com

consider can seem long. However, for SOHO (small office/home office) implementations, with the exception of drive capacity, these specifications have little impact on the overall performance. The choice is made simpler by suppliers, such as Toshiba, who provide drives such as the N300 series that are optimised for the demands of a NAS.

Although a NAS is expected to operate 24/7, most NAS systems implement some sort of power management. This means that it will idle or enter a low-power mode at times when there is no activity. In most cases this will be at night. The drives themselves have limited impact on overall power consumption as, nowadays, they are all very energy efficient.

For example, a conventional 7200 rpm air-filled drive consumes 10W when active. A 2 bay RAID1 implementation with two such 8TB drives would consume around 175kWh/year. A move to two 14TB N300 helium-filled drives reduces power consumption by half but, assuming 0.30€ per kWh, this reduces annual electricity costs from around 50€ to 25€. Considering that the NAS will also power the drives down for a few hours per day, the real savings end up being significantly lower.

Moving to the lower-power helium filled drives makes more sense when the need arises to increase capacity in a small two bay NAS.

Noise is often a consideration, especially in a work environment. Nobody wants the continual distraction of whirring disk drives while trying to concentrate. The 14TB HDWG21EEZSTA is ideal in such situations, generating just 20db (typical) noise when in idle mode. However, optimal location of the NAS should also be carefully considered in addition to sourcing low-noise drives. Placement in a basement or a dedicated room keeps it away from accidental knocks or being unplugged from the power source, while additionally resolving the noise issue.

Additionally, heat has the largest impact on hard drive longevity. Keeping drive temperatures under 40°C is essential as, above this level, reliability is significantly impacted. The NAS will have a fan built in, but this is of little help if it sits in direct sunlight for part of the day. Drives such as the N300 automatically adjust seek speed during high-temperature operation to reduce heat generation in demanding situations.

Summary

It is very easy to get caught up in the details of individual drive specification when defining a NAS solution. But the reality is that today's drives are so reliable, responsive and energy optimised that the real focus should be on overall capacity and robustness. Rather than focus on reducing drive noise in an office environment, focus should be placed upon installing the NAS itself in a safe, cool and ideally secure location, away from prying eyes and fiddling fingers.

With the exception of RAID0, all other RAID options result in less total capacity than the sum of the drives used. This is offset by increased robustness. Helium-filled drives mean that huge capacities can be constructed that will provide enough storage for many years of data growth.

Toshiba Electronics Europe GmbH
Hansaallee 181
40549 Düsseldorf, Germany

Further information
toshiba-storage.com