

The Art and Science of Backup

Data that doesn't have a backup copy, by definition, is unimportant.

Data protection plays a more central role today in production than ever before, thanks to our dependence on data. An increasing amount of data, combined with patchy understanding of professional data security planning, creates a dangerous discrepancy that often leads to a position of avoidance. Taking basic ground rules into consideration allows the gradual professionalizing of data security. Thus, the threat of data loss destroying a company is prevented. With a combination of current technologies, a new level of professional data security can be reached, also offering transparency in data storage.

Backups are one area where "paranoia is prudent" – James Pond

The process for copying (from Latin "copiare", meaning to transcribe multiple times) is nearly as old as writing. Monks in the Middle Ages took care that literature from the antiquity was spread through writing as much as possible, so that enough copies would survive fires, wars and other threats. This is the same theory of a modern day backup. But is data security still necessary, when today's systems are more reliable than ever? The increasing complexity of workflows, combined with the use of individual components, creates a higher possibility of failure. The greatest risk is posed by the users themselves: in the course of a hectic production schedule, errors and mistakes are bound to happen. Data security is necessary to protect against unforeseen consequences and errors.

What is the cost of not securing data?

What are the incurred costs when losing data? In this case, if production data is inaccessible, this impacts the production process down to the customers. Depending on the type of business, this means either a longer or shorter grace period. The longer the data remains inaccessible, the more customers are impacted. Additionally, workers are left with nothing to do until production can continue. Repetitive failures affect the work atmosphere and setting at work with a negative vibe. If data gets lost, the damage can be extremely expensive and range to incalculable, putting the future of the company at risk. American studies show that the majority of companies, which experience a large data loss had to file for bankruptcy later (1).

Planned Scare

The securing of data can only be as good as the data loss situation, which creates the need. Thus it makes sense to plan for the absolute worst case. A glance at the local server installation is a good place to start. For example, a defect in calculating cooling requirements and total heat load in a server room can allow a server to overheat and raises the possibility of fire and data loss. A neighboring system can also be affected, if it is in the vicinity. The primary storage and its local copy of the production data could potentially be lost simultaneously. The information regarding the configuration of the system also gets lost in these cases. This underscores the seriousness of having access to important information related to configuring the whole setup in an emergency situation. It is of crucial importance. Important information worth keeping, i.e. data to backup, includes the configuration of the servers and their respective services, SAN systems and network. In case of emergency none of these systems will be available to check for IP addresses, configurations and other specs. Therefore thorough documentation as well as a detailed step-by-step procedure is necessary to combine all information that is critical for the restoration. The following questions need to be answered:

- What happens when the central storage is destroyed?
- What happens when databases with addresses, contracts and financial data is destroyed?
- How far away data has to be taken offsite to survive this scenario?

A conclusive disaster recovery planning addresses these pressing questions (2).

Only an automatic backup secures data reliably. Any process that has to be initiated manually can be forgotten and postponed, especially under time

pressure. That is the time when user errors peak and a securing of data is especially necessary.

Multiple, respective multi-stage backups build additional security. Multiple backups enable off-site storing of data. A multi-stage backup can make use of the strengths of different methods.

With Disk2Disk2Tape the respective advantages of Disk and Tape are used. Disk-to-Disk offers fast backup and fast restore of files.

The securing of data to tape, runs from the disk backup. With the disk backup as the source the backup to tape can use its superior streaming capability and access the entire backup. The retention time can be prolonged considerably because of the cost advantage of tape per TB over disk. Additionally relocation of written tapes ensures a professional level of security. (3)

To save everything is not quite as easy as it sounds. To start with, one needs a list of all servers and workstations as well as their disks and partitions. Network and client management software can provide lists of computers and their respective software. A centralized administration of software licenses as well as disk images for fast installation, facilitate the recovery process. The documentation should be saved redundantly and off-site and permanently accessible. Workstations that have an elaborate configuration with lots of tools, plugins and applications require a lot of time to be reconstructed. Therefore it pays off to save these completely. For Mac OS X machines a bootable disk clone can speed up resuming of production considerably.

The use of different storage media with their different technologies like disk and tape serves as risk reduction. While a disk ("online storage") can be hit by user errors, malware and defects while in use so that even the whole storage can be wiped, tape ("offline storage") can only be affected by physical damage and inadequate warehousing.

Relocation (off-site storage) is part of every professional security strategy. Major local damage at a company very likely will impact the local backup. Additionally damage that is limited to the local IT setup can hurt or destroy the backup. User errors like misconfiguration, update or modification of the hardware as well as local damage like fire, lightning and power blackouts all belong to this category. For all those cases and more the off-site storage is the key to rescuing data. This relocation helps to achieve a professional security level but requires the establishing of a media rotation. In this respect tapes are a lot easier to relocate than disks. Tape libraries allow for easy exchange of complete tape magazines with multiple tapes in one step.

Professional media rotation requires at least three tape sets. One resides at the off-site location, one is being written at the company and one is being transported to/from the remote location. This way any possible damage during transport is covered.

How about Backup in the Cloud?

Cloud computing is one of the most recent technology trends. A multitude of different technologies is hidden behind the term. Only some of them are related to securing of data. So what is a professional backup in the cloud, what needs to be taken into account and what is it suitable for?

Initially the whole transmission chain needs to be looked at. For a backup sufficient upload bandwidth is required. Typical DSL lines, being asymmetric, have their weak spot exactly there with a smaller upload than download bandwidth, depending on type between 10:1 and 24:1 (4).

Data lines with higher upload bandwidth respectively symmetrical lines still carry a considerable cost factor. Additionally there might be restrictions by the provider. This could be a throughput throttle after a certain amount

of data, the exclusion of specific kinds of data (e. hg. video files) and more. In case of emergency getting data back as fast as possible is the point. That means testing the maximum constant download rate. Depending on time of day, load at the provider and network in between, there can be significant variations. With a data volume of several TB those variations alone can result in a number of hours (or days) of additional waiting time. Using multiple cloud service providers requires specific software for each of them. For ease of operation use the free tool "WingFS" by Archiware. It offers the option of using multiple providers of the big cloud services and also works together with P5 Synchronize to sync to, from and in between cloud services (5).

Also of importance is the question how safe the data is at the provider. All US companies and their European branches have to hand over data to US authorities on request ("Patriot Act" and "FISA Amendments Act of 2008"). They are not allowed to even tell their customers after the fact.

This opens the door to all kind of uses and human error. Data protection rights of the country of

origin have to be obeyed by every company that uses these services and might be affected by this. This responsibility has to be taken into account. Microsoft and Google have already confirmed this practice (6). How far the interests of users and cloud services differ and contradict is shown in the doctoral dissertation of Christopher Soghoian (7).

Conclusion

Cloud backup cannot replace a traditional on-site backup. If the aforementioned criteria can be met in a useful way and the amount of data is limited, it can serve as an additional off-site backup (8).

The restore process should be tested regularly. The backup software is only one part of this process. More critical are other aspects, especially organizational ones.

What steps are important for the restore process?

To begin with someone triggers the restore of files, directories or even complete storage areas or

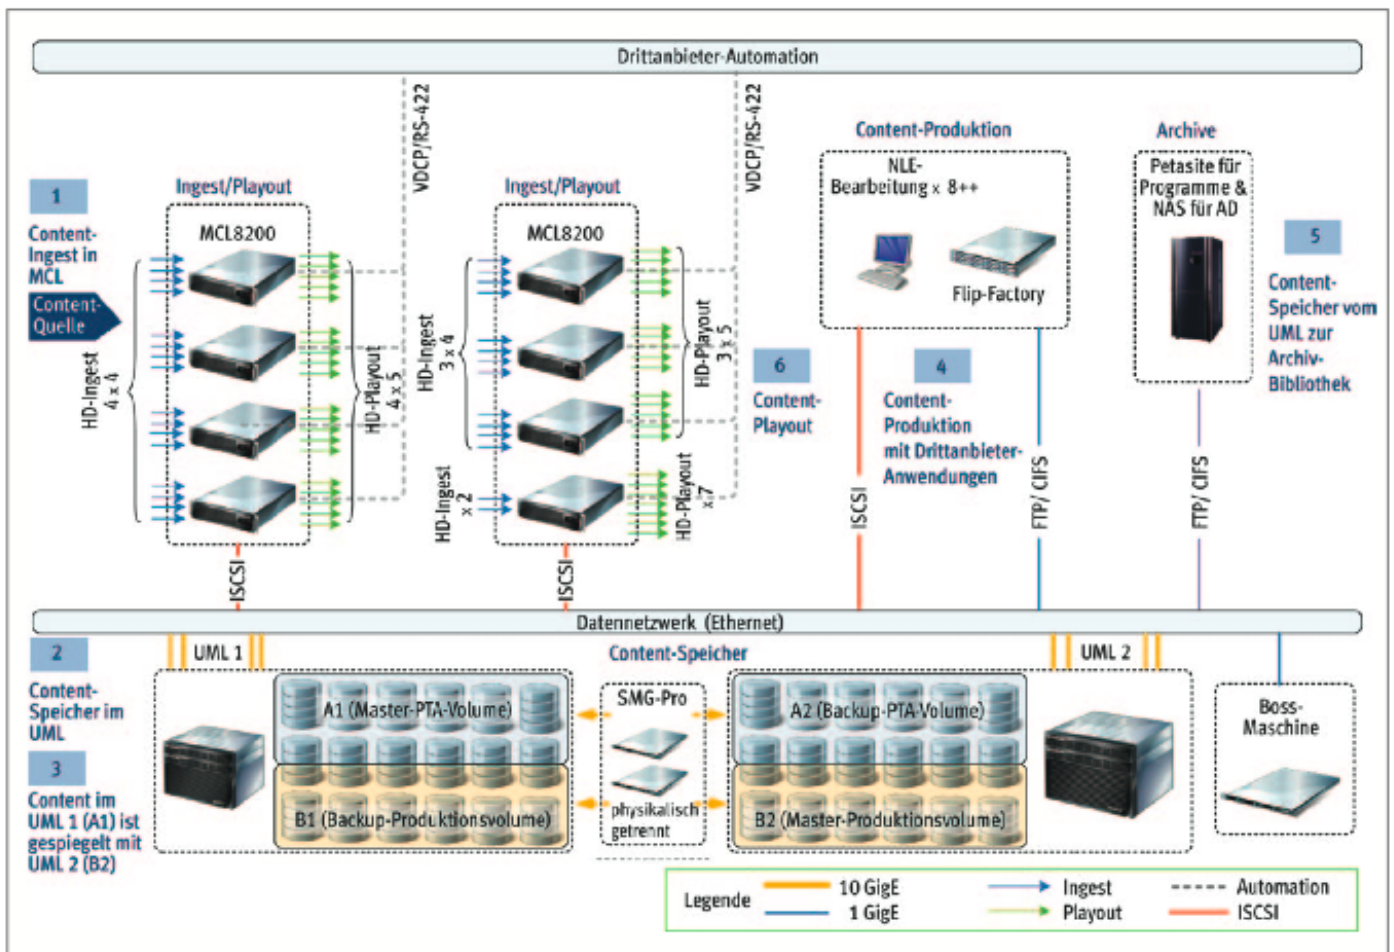


Image 1: Example of using a data availability solution to clone between storage A1 and A2 respectively B1 and B2 at MTV 3 Finland using XOR storage and P4 Synchronize.

servers. *Who in a company is allowed to do that?* The next step is deciding where to something is going to be restored. Of course there is a big difference here, between restoring a file, a server directory or a complete SAN:

- Is there sufficient spare storage?
- If a server fails is there virtual or real replacement?
- Who has access to the necessary resources, passwords and accounts?
- What is the exact procedure for what failure?
- How regularly is this tested?

These are many questions, but it is necessary to look at and answer them to be able to take precise action in case of emergency. The arrival of a new workstation or server is a good opportunity for a restore test. With minimum effort this machine can be used as replacement hardware to be loaded from the backup before deployed for the original purpose.

The whole IT landscape should be solidly configured and run. Details like temperature and humidity need to be controlled to avoid corrosion and ageing by raised humidity and electrostatic charge through dry air. All components like servers and RAID's only carry parts certified for them. The administrator is trained for his systems and documents the setup and changes as well as failure plans. The combination of these professional conditions minimizes risks and provides a high level of dependability. This way, the likelihood, that a backup needs to be accessed, is reduced.

Backup Planning Begins at the End

The first topic when deciding on a backup procedure is data availability. The question is: *How long can I be without my data?* In data security jargon this is called Recovery Time Objective (RTO). There are three main categories:

- Time critical, that means minutes
- Hours and
- One or several days

After that one needs to decide on the retention time, meaning how far one wants to be able to go back in the data set:

- Short (days)
- Medium (weeks) or
- Long (month to years).

This has consequences for the decision of the storage medium or a combination like:

- Disk = short to medium (days to month)
- Tape = long (months to years)

How long to be without data?

The backup method is mostly determined by how long one can be without a respective data set (RTO). Time critical data that are required within minutes after a failure can only be saved with a data availability solution. This is also called a fail-over solution and allows "switching over" to the cloned or mirrored storage – no restore process is necessary. Action by an administrator is needed though to bring data online again. Such a fail-over solution is very attractive especially for SAN systems where multiple users access storage with a high workload. Sufficient secondary storage is necessary as cloning target as well as its connection to the SAN. P5 Synchronize by Archiware is such a solution and is in use in many production facilities (9). It fills the requirement of transferring all file attributes, ACLs, extended finder attributes, Xsan attributes etc. (Picture 1). Starting with a setup modernization at MTV3, SD and HD storage areas were combined, hi-performance XOR-storage was deployed redundantly and kept available with P4 Synchronize by Archiware. Additionally file versions and snapshots of the whole storage can be kept.

With lower requirements, when recovery time can be several hours, a traditional backup can be used. In this case a runtime for the restore process is necessary and is dependent on the amount of data being restored. In case of many TB it can take several hours. It can be helpful to give users (even limited) access to the backup. Minor mishaps can then be corrected without the need for an administrator.

Backup Methods

There are different kinds of backup to provide the best fit for the respective requirements. Since terms differ from one solution to the other it makes sense to thoroughly deal with the specific terms when planning.

To begin with there is the **full backup**. As the name suggests all data is saved here. The **incremental backup** saves everything that has changed or was created since the last backup. The **differential backup** saves everything that has been changed since the last full backup. Additionally there are some special cases like the **permanent incremental backup** that does without a full backup. Within P5 backup by Archiware this is called Progressive Backup. By applying additional mechanisms, sufficient redundancy of data is achieved without a full backup, that means when the amount of data is too big for a given time frame. This can happen and does so with increasing frequency.

Depending on throughput and time frame during one weekend (Friday 6pm – Monday 6am = 60 hrs) at 100 MB/s about 21TB can be saved. In today's storage world that is a rather limited amount. Since the throughput does not grow at the same rate as the storage size more and more backup time frames will hit their limitations.

When are different backup procedures applied, and in what frequency, and in what combination, depends on the specific requirements. The following questions need to be addressed:

- **When is the best time to run a backup without compromising production?** With production times expanding and being fed from different time zones, deciding on a backup time frame, can be challenging. Maybe the setup needs to be upgraded to support backup and production simultaneously. Establishing subnets or exclusive communication pathways might also be necessary. In general it is better to save during the night since most files are in a stable state and the infrastructure has less load.
- **How frequently the backup has to run to keep possible data loss in a tolerable range?** The frequency of backup runs determines the maximum data loss (RPO = Recovery Point Objective). Different data sets mostly have different levels of tolerance for data loss. Segmenting storage areas can support efficient coverage on the respective levels of requirements. A daily backup is the absolute minimum in any production environment.
- **What is the best combination of backup methods?** Multiple factors contribute to this decision. To begin with the size of the storage to be saved determines the runtime of a full backup. The rate of data change determines the same for an incremental backup. Combining these and adding the RPO one can build a schematic and coordinate with the production flow to distribute load.

Disk or Tape

If short access times are needed a disk based solution is preferred. If the focus is on retention time, that means how long one can go back within the backup, tape is definitely superior. Combining disk and tape establishes an additional level of security since both technologies have different vulnerabilities. Price per TB, energy cost, space requirement and durability put tape one level above disk. LTO tape is an established standard in IT and is used by large corporations, banks and insurance companies worldwide.

Besides its cost advantages, the performance (throughput LTO-6 = 160 MB/s, native), capacity (2,5TB, native), security mechanisms (read-after-write, error correction, auto speed, servo tracks) and the longevity (30 years readability) are convincing. Additionally, tape scales easily. Putting new tapes in a library to expand storage without configuration change and setup modification. For large data sets LTO tape right now is without comparison (10).

Disk and Tape

Disk2Disk2Tape is a method that makes use of the respective strength of disk and tape in an optimized combination. Disk-to-Disk offers fast saving and fast restore of data. Saving to tape takes place from the disk backup volume. With this as source tape can make use of its high throughput and access the complete backup. Tape throughput can even be multiplied by parallelization as offered by Archiware's P5 Backup. Multiple drives can be used to write simultaneously and thus secure a very fast data stream (11).

The cost advantage of tape can be used to extend the retention time considerably compared to disk. Additionally written tapes that are stored off-site provide a professional level of security.

Another combination of disk and tape is often advantageous in production – when MAM systems (Media Asset Management) are in use. Most of the time these use SAN storage. Here availability is of increased importance because the MAM is the central hub of the workflow. Combining it with a tape Archive to relieve the SAN and move finished productions or rarely used assets to the archive, is a natural choice. This can also be automated (12).

Especially convenient is it when users can trigger the Archive and restore process from within the MAM system. A recent example of this integration of MAM and Archive is offered by moosystems.com in Cologne. Their "CP Archive App" connects Cantemo Portal MAM with P5 Archive by Archiware (13). The investment in an Archive pays off from the capital saved holding off expanding the SAN storage alone. Additionally, drastically reduced energy and cooling costs increase the gain. When calculated over a longer period significant savings result when using tape instead of disk systems (10,14).

Balance between Backup and Archive

- Backup = secondary copy of primary data
- Archive = primary copy of secondary data
(Curtis Preston)

Contrary to backup the Archive serves as (mostly) open-ended saving of data. It forms the entirety

of completed productions. There is no way of deleting a file on a tape. Only the whole tape can be erased. This means an archive is constantly growing in size while a backup is renewed in a cyclic fashion and stays at roughly the same size. Since the amount of data on the backed up source is growing, the backup also grows in size. Often this is not necessary since most data that is created is not often or hardly ever used again. Nobody likes to admit it but it seems to be true in most environments (15).

A considerable potential for cost savings exists here since these data could be moved to an archive. From there, they could be restored if need be without blocking costly main storage. At the same time the size of the backup would be reduced and thus the archive would pay off in a few years.

Files that are used in production, currently likely to be modified, belong on the production storage and integrated in its backup. Any completed production should be moved to the archive. There it will be safeguarded by archive security mechanisms like tape cloning and should be deleted from the main storage.

If files are archived too soon and have to be restored frequently the archive grows disproportionately since deleting files from the tape archive is not possible for technical reasons. The decision when to archive is essential for an economic operation of an archive.

Incremental Archiving, that can only archive new and changed files of a specific file system, was recently introduced by Archiware to its P5 Archive software. Depending on setup this can improve efficiency and limit growth of the archive (16).

Special Case Laptop Backup

For creation and modification of files laptops are in use in many cases. Changing locations and freelancers contribute to this fact. Since the production depends on the completeness of the files on mobile workstations the backup of those machines should be centralized. Data loss, even of freelancers, can jeopardize the whole production. A flexible, automatic and laptop optimized backup that saves all data centrally in the company is necessary to minimize risks. The administrator should have minimal workload from this additional task. Smaller workgroups can use open policies to decide what gets saved when and how frequently. Larger corporations use closed policies that define every parameter and provide a consistent security level. A solution like "P5 backup2Go" (Archiware) covers both policies using templates. New users can be integrated in minutes. Again a user restore

function to restore single files or directories eases the work for the administrator.

Especially with the laptop backup it may be necessary to encrypt business data. Thus management and business documents can also be integrated with the backup routine. It is of critical importance that the encryption key is accessible if need be. This can be achieved by printing it and putting it in the safe, in a sealed envelope. In most cases encrypting of media files will hinder production and be omitted.

Challenge on Set

When in production with tapeless cameras, a multitude of security issues arises already on set. A mobile setup has to take care and provide maximum security for the just recorded data. Initially the contents of the storage cards needs to be copied to, at a minimum, hard drives. Using different technologies like one SSD and one hard disk drive further increases security. Maximum security can be achieved by delaying the re-use of the card by one day. In case there are any inconsistencies they can be corrected using the original data.

A new profession has been established for this responsible task, called the data wrangler. There are specialized procedures like the labeling and marking of the memory cards that are related to procedures in data management. The whole production depends on the work of this person since one mishap can destroy the work of the whole team. Using checksums is common but they in themselves provide limited benefits to security and should be used wisely. One copy of the data should be transferred on an LTO tape that in turn should immediately be removed from the set and be stored off-site. Relocation of LTO tapes is required by many set insurances in the US.

The decision what software is used for securing data is of importance for the whole data management. Cost for the software and maintenance is to be taken into account as well as ease of use. Sooner or later it will be necessary that multiple colleagues will have to work with the backup software, some only with minimal training. Using complex tools, modified scripts or open source software introduces unnecessary risks. In general data management software should be designed in a way that it is easy to navigate even after months where there was no need for modifications. Support for all relevant platforms as well as storage technologies increases flexibility. Scalability, that means the ability to grow with the needs, is an important aspect. Another aspect is the option of giving users access to their backup.

When purchasing new equipment and designing new workflows the impact on the backup system is too seldom taken into account. A comprehensive planning integrates data management from the beginning as well as maintenance intervals, consumable materials and other consequential costs. An automatic procedure that checks all new hardware for its backup relevance is already a good step towards a disaster prevention strategy.

Starting with classic informatics principles like minimal data collection is an important step to keep a production securable and manageable. Checking the production formats and their respective benefit is advisable. What codec can deliver the necessary quality while reducing the amount of data? What resolution is actually necessary? The temptation to record in 4K with a 4K camera is big even when there is almost no delivery option yet.

As conclusion one can say that the biggest risk consists of a missing or inadequate backup. This opens the door to compromising the production through user errors. One increasingly frequent attitude of decision makers – when deciding on new technologies – is making decisions from their own (consumer) experience, like using USB disks for their private backup. To make decisions in this way means professional requirements and necessities are disregarded.

With implemented backups the biggest problems arise from not testing the actual restore procedure. The difference between the estimated and the actual time needed to complete the whole procedure can run up as high as several days. A detailed step-by-step documentation for emergency situations is extremely important. Especially under pressure one cannot count on the experience and adequate judgment of all involved.

Dangerous Misconceptions

Numerous misjudgments prevent adequate backups or obstruct the view on existing risks.

“We are safe – we use RAID.” This is a common misconception that RAID manufacturers have contributed to with exaggerated messages. On the one hand 70% of data loss can be attributed to users and their involuntarily deleting or renaming of files, directories or volumes, which RAID cannot prevent them from. On the other hand the probability of failure rises with the number of components involved. A RAID consists of highly critical components like RAID controllers and a larger number of disks. Their interplay is subject to very tight tolerances. A RAID can only protect against the failure of one or, depending on RAID level, several disks. There is no protection against

file system corruption, deletion, moving of files, user errors etc. Without a backup the complete RAID system has to be brought to a data recovery service in case of emergency. The inquiries for data recovery have increased dramatically with RAID 6 as the assumed security led to omitted backup implementations.

Some hard disk math might illustrate the risk associated with RAID. The calculation of “Mean Time to Data Loss” (MTDL) of RAID 5 and RAID 6 includes the bit error probability. This approach is necessary to include the possibility of data loss during a rebuild after the failure of a disk. For this all sectors have to be available and absolute errorless. A rebuild process with a bigger RAID can take more than a day (!). During that time the performance is reduced and the danger of failure is increased and not protected by RAID. With a RAID 5 and 40 disks this value is only one (!) year (17).

Summary: RAID is no replacement for a backup but needs a backup of its own.

“Tape is obsolete.” This misjudgment has its root in the numerous predecessor technologies of LTO tape. Those were to some extent slow, unpredictable and insecure. LTO tape, though, is one of the rare strokes of luck of the IT history where the best features of precursor technologies were combined to form the new technology. The LTO consortium consists of IBM, HP and Quantum as well as associated members. Furthermore LTO is right now in its 6th generation with a roadmap up to generation 8 and presents a robust and reliable long-term solution. At the same time it is the only medium that supports relocation of large data sets (18).

“We have a snapshot.” The snapshot feature of virtual environments is used to re-establish earlier stages of the system. This feature is no backup, though. It needs the complete virtualization structure to work flawlessly as a prerequisite. In case of emergency it will not restore anything in general.

“We backup when we have time” or *“Everything runs fine”* fit the same category: saving data is under estimated and not considered to be part of the business processes. Since today’s production, contact, financial and email data of a company represent the actual core of a company, securing these is an existentially necessary part of the business itself. There are legal requirements to do so and companies and managers are responsible. Guidelines of several sources help to limit the risks (19).

“We use a free/bundled software.” Most of the time this is software for a single user. A professional backup and respective requirements never were the scope here. This can result in correspondent restrictions: obscure processes, no report on what was saved, deleting of previous backups without prior warning, lack of configuration options, lack of scalability, restricted to one OS, no overview about the condition of the setup and others more.

“The backup went through.” Frequently ignored is the fact that a backup is only valid if the restore process is tested regularly. The focus here is not only the saved data, but mainly the necessary organizational procedures. Where can the data be restored to, what storage is available in case of emergency, how can a replacement server be set up, who has the necessary privileges, how long does the restore really take, what has to be adjusted in the network or infrastructure, where is the documentation of the necessary steps?

“We have a NAS.” Today’s networks can carry considerably more data than ten years ago. At the same time the amount of data grew tremendously. In most cases the data growth surpassed the throughput growth by far. This arises when large data sets have to be saved across a network. Besides restricting the usage for other purposes at the same time the duration for the transfer in general is too long. With growing data sets it increases even more.

Important and large storage areas should be connected to the backup with a dedicated connection or placed within the FC (Fibre Channel) network of the backup storage. Furthermore most NAS systems are black boxes with respect to knowing what Kernel version, what RAIDs and what hardware is built in. Sometimes a disaster shows what hurdles need to be overcome or if data will be lost completely.

LTO Tape with LTFS

Beginning with the fifth generation of LTO tapes (LTO-5) LTFS was introduced. Here an option is available of using the tape with partitions, one for the data and one for the metadata. In theory the tape can be used like a CD-RW/DVD-RW where the content of the tape is displayed when put in the drive and files can be read and written. Tape manufacturers often compare this to usage like a hard disk but that is strongly exaggerated.

As a format for securing data LTFS right now is not suitable. The format cannot keep all attributes of all files. This is the reason for all Backup software vendors using their own format that is optimized for backup.

Outlook

The outlook for data growth is quite clear. A further increase in data volume results from larger production formats and more delivery platforms. Today most productions have data sets of a size that only few years ago only data centers had to deal with. This also results in a discrepancy between experience and requirements in data management and backup. Data loss will occur increasingly often until a reasonably professional data management culture is established.

The usage of SSD storage will increase and also be pushed by the IT industry. The failure risk here is open and remains to be seen. Disks will become even cheaper probably while reducing quality and increasing vulnerability (except in the expensive category). LTO tape will remain the professional archiving and long-term backup medium (20).

Since the network bandwidth will grow a lot slower than the amount of data, saving large quantities of data via the Internet will remain inadequate. A professional local backup stays the best solution for media production in the foreseeable future.

Basic Backup Rules:


1. Automatic Backup Process
2. Multiple Stages/Copies
3. Include All Required Files
4. Change Media/Technology = Disk/Tape
5. Relocate/Off-site Storage
6. Test Restore Process

References:

- [1] http://www.bisimplified.com/_pdf/_newsletters/BI_NewsletterApril2010.pdf
- [2] http://en.wikipedia.org/wiki/Disaster_recovery
- [3] <http://www.infoworld.com/d/storage/wp/long-term-data-protection-and-retention-finding-the-correct-balance-570>
<http://www.archiware.com/disk-to-disk-to-tape-d2d2t.33.1.html>
- [4] http://de.wikipedia.org/wiki/Asymmetric_Digital_Subscriber_Line
- [5] <http://www.wingfs.com/>
- [6] <http://www.zdnet.com/blog/igeneration/microsoft-we-can-hand-over-office-365-data-without-your-permission/11041>
<http://news.softpedia.com/news/Google-Admits-Handing-over-European-User-Data-to-US-Intelligence-Agencies-215740.shtml>
<http://www.bigbrotherawards.de/2012/comm>

- [7] Doctoral dissertation: "The Spies We Trust: Third Party Service Providers and Law Enforcement Surveillance" <http://files.dubfire.net/csoghoian-dissertation-final-8-1-2012.pdf>
- [8] <http://www.techrepublic.com/blog/five-apps/five-tips-for-backing-up-your-data-to-the-cloud/697?>
- [9] <http://www.archiware.com/p5-synchronize.301.1.html>
- [10] <http://www.spectralogic.com/blog/index.cfm/2012/8/24/Economics-of-Tape-Indicate-Warm-Waters-for-Glacier>
<http://www.spectralogic.com/index.cfm?fuseaction=members.docContactInfoForm&DocID=4255>
- [11] <http://www.archiware.com/support/download/Parallelizing.pdf>
- [12] <http://www.andre-aulich.de/en/perm/use-presstore-archive-to-automatically-move-data-to-tape>
- [13] <http://moosystems.com/products/moofs-archive-app/>
- [14] <http://www.clipper.com/research/TCG2010054.pdf>
- [15] <http://www.ssrc.ucsc.edu/Papers/leung-usenix08.pdf>
- [16] <http://www.archiware.com/archiware-p5-new-features.310.1.html>
- [17] Calculations mentioned on page 21 of presentation in German
http://www.heinlein-support.de/upload/slac08/Heinlein-RAID_Mathematik_fuer_Admins.pdf
- [18] <http://www.lto.org/technology/roadmap.html>
- [19] https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Minimum_information/SecurityRecommendationsCloudComputingProviders.pdf?__blob=publicationFile
<http://www.symantec.com/business/support/index?page=content&tid=TECH91705>
- [20] http://www.theregister.co.uk/2011/05/06/soirage_trifecta

Marc M. Batschkus
MD, PhD is medical informatics specialist, scientist and lecturer. At Archiware he is in charge of global business development Science/Medicine/Media. He can be reached at mmb@archiware.com



Originally published in German.
FKT is the magazine of German "Fernseh- und Kinotechnische Gesellschaft" (TV and Cinema Technology Society)
All rights reserved – © Copyright 2013 by Fachverlag Schiele & Schön GmbH, Berlin
Translation: Ty Eriksen, Marc M. Batschkus
Proofreading: Mat X