

SEP sesam

Backup in distributed Environments

**An introduction for partners and customers,
backup architects, backup designers and CIOs**

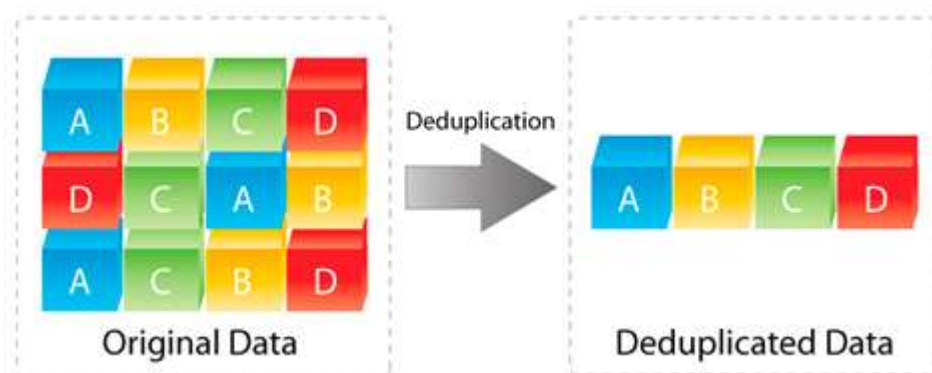
Contents

Introduction.....	2
Overview	3
Separate Data Zones	5
Backup of Clients directly via WAN.....	6
Local Backup at Remote Office	8
Replication.....	9
Source Deduplication	11
Replication with Deduplication	14
Backup of Laptops of Field Service.....	16
Network Structures	17
The SEP sesam Ports.....	17
Firewall Environment	18
Backup over an alternative Network	18
IPv6.....	19

Introduction

Because of the present possibilities for interconnection and increasing bandwidth company environments are nowadays almost 100% distributed. This means that in most cases a powerful company network is connected only via public networks to other network islands or to its field workers. In particular globalization and distribution of companies worldwide is continuously enlarging the problem of interconnectivity. Connecting a network island is on one hand a cost problem because renting broad bandwidth and transmitting big data is still very expensive. On the other hand it is technically not feasible everywhere at the moment or requires large efforts and investment depending on how isolated the remote branch office is in fact. In addition for data transmission not only the mere amount of data is important you also have to consider stability and security. Regarding the value and classification of the data to be transmitted out of the secure company network you also have to consider encryption and integrity.

In order to reduce the data to be transmitted via WAN, for a long time compression has been the only choice. Starting with Single Instancing (transmitting identical files only once e.g. with email attachments) deduplication has become a mainstream technology on the market. Deduplication means that the data stream and therefore also the files will be cut into chunks and a hash value will be matched with an index on the backup server. Already existing chunks of data will not be transmitted again and only a pointer to the existing block will be stored.



Overview

There is a wide spectrum of possibilities to backup distributed environments depending on the given infrastructure, amount of data or time frames. Two important criteria to define the exact requirements for backup and restore are the values RPO and RTO:

- **Recovery Time Objective (RTO)**

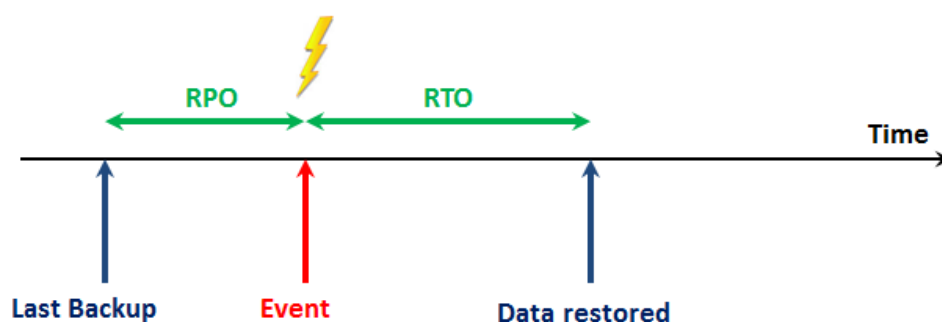
For how long you can afford a business/system to be out of order?

The Recovery Time Objective (RTO) for an application is the goal for how quickly you need to have that application's information back available after downtime has occurred.

- **Recovery Point Objective (RPO)**

How much data you can accept to lose?

The Recovery Point Objective (RPO) for an application describes the point in time to which data must be restored to successfully resume processing (often thought of as time between last backup and when an "event" occurred).




Each configuration has its pros and cons. Here you will find an overview of the options to backup distributed environments which will be explained in more details in the next chapters. In the first four chapters you will find the classical possibilities while in the last two chapters the alternatives based on deduplication will be looked at closer.








Method	Assessment
Separate data zones	Flexible, costly, no synergy effects by centralization, but central control possible via MasterGUI
Backup of clients directly via WAN	Cheap, simple, only few clients, small amount of data
Local backup at remote office	Big amount of data, separate HW and local data storage with central remote administration via adminGUI
Replication (rsync, GlusterFS, Ceph)	Medium amount of data, not under control of backup SW
Source deduplication	Low bandwidth possible, initial backup and big restores causes problems
Replication with deduplication	Big amount of data possible, full flexibility

Separate Data Zones

This most simple case often comes first. Originating either from historical evolution over time by growth, from acquisitions or because remote offices would like to keep their autonomy. Completely separated data zones could operate totally independent and react flexible according to their individual requirements (e.g. different license models, backup strategies, HW, et al.). If necessary they can even run different backup software products.

If multiple data zones have a separate SEP sesam backup server each all backup servers can be administered centrally from one MasterGUI. Logfiles and backup stati can be checked centrally and follow-up actions can be triggered (e.g. migrations). The host with the MasterGUI does not necessarily have to be one of the backup servers. It could be any host in the network with an installed SEP sesam client and Java.




Pros	Cons
 independent and flexible	 high costs from duplicate administration, licensing and backup infrastructure
 no problems caused by data transmission	 no synergy effects
 Local data storage, but central control of multiple SEP data zones via MasterGUI	 No MasterGUI possible when using backup software from different vendors
 Usable for private cloud structures	

Backup of Clients directly via WAN

From a configuration view it is not a problem to add any client to a backup server as long as the clientname can be resolved from the DNS and it is reachable in the network. This is independent from how close or remote this client is located (firewall, another LAN segment, WAN connection, et al.). It is only important that the backup server can resolve clients' name and reach it via LAN and also vice versa.

The real problem pops up when the data will be transmitted. The lower the bandwidth of the network connection especially with WAN, the smaller is the amount of data which can be transmitted reasonably. Reasonably in this case means hitting the right backup time frames and restore SLAs. But also allocating bandwidth long-term for backup and affecting the smooth operation of other applications could raise problems. That is the reason why customers very often define specific bandwidth or wires dedicated for backup exclusively.


If single clients will be backed up directly remote to a backup server, because of its restrictions in data transmission, this concept is only usable with few clients and small amounts of data.



Example:

- 1 remote office with 2 clients and a total of 1TB of data
 - Change rate of 10%
 - WAN connection of 155Mbit/s (ATM)
- ⇒ Duration for a full backup = $1.000.000\text{MB} * 8\text{bits} / 155\text{Mbit/s} = \sim 14$ hours
- ⇒ Duration for an incremental backup = $\sim 1,4$ hours

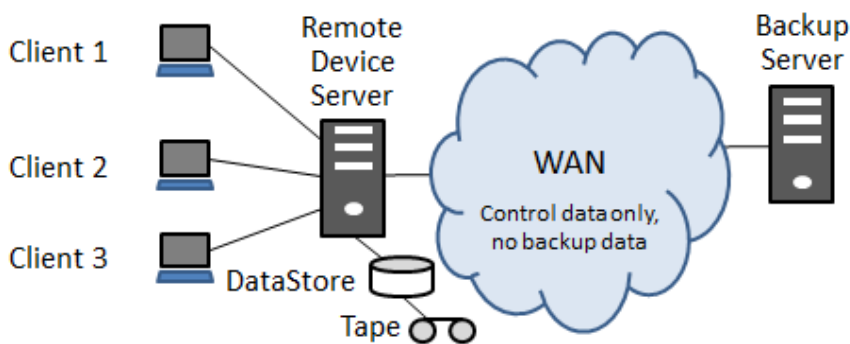
As you can easily see from the calculation above this concept is useful only with very fast WAN connections.






Pros	Cons
 centralized backup	 small amount of data
 Very simple configuration	 few clients
 no extra costs	 fast WAN recommended
 direct backup to tape possible	

Local Backup at Remote Office

If the amount of data will be too big to be transmitted via WAN, the backup data has to be stored locally. This can be accomplished by setting up a remote device server (RDS) in the remote office who is able to write data directly to locally attached target media (disk and/or tape) on behalf of the backup server. When using one or more RDS (competitors call it a Media Server or Storage Node) full control stays with the backup server while the data transmission for backup, restore or migration will take place in the RDS.

Even if this concept requires additional local HW and administration and contradicts to central data storage concepts, it has significant advantages with respect to restore performance because the backup data is accessible locally.



Pros	Cons
 high performance from local data storage	 no centralized data storage concept
 central control of the backups	 additional costs for HW and administration
 big amount of data, lots of clients	

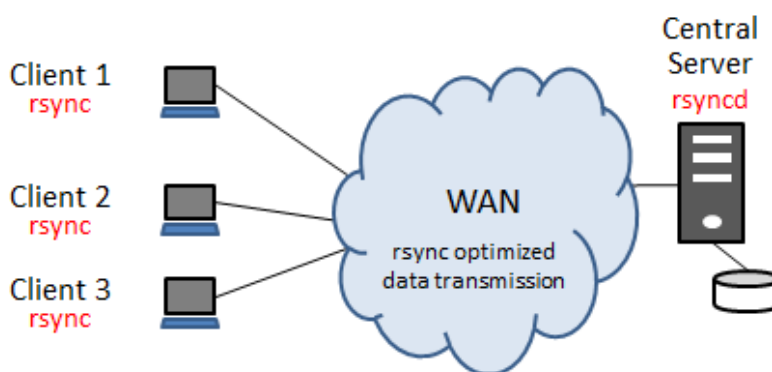
Replication

There was always the problem to transmit data to other locations especially for disaster recovery reasons. If there happens to be a disaster with a single computer or a complete location, all data should be available somewhere else to enable a full functioning environment by simply restoring the data on new hardware. The target for the data could be another data center or just a remote but disaster-proof location for disaster recovery. In the past and today e.g. tape copies have been physically relocated to old bunkers or mines. As bandwidth is permanently increasing, also disk-to-disk data replication or mirroring is possible nowadays, either synchronous or asynchronous.

As a very simple example on UNIX/Linux it is possible to define a `rsync` command in a crontab table to replicate data in regular intervals to a remote location. Rsync is also available for Windows from Cygwin but there are also a lot of rsync-based GUI tools. Rsync is optimized for data transmission via WAN because it cuts files in chunks, generates checksums and submits only changed blocks. Encrypted (ssh) or compressed data transmission is possible as well. As rsync is essentially working file-based for replication of complete directories, it is recommended to use some manual scripting or rsync-based tools. With rsync it is possible to replicate not only original data of a client but also local backup data (e.g. a datastore of a RDS). In this manner even deduplicated data could be replicated.

Another possibility of data replication is to use distributed file systems like GlusterFS, Ceph or DFS which provide replication as integrated functionality to be configured by the user.

With all the options it is absolutely recommended to consider scalability and even doing some testing in specific cases.



Pros	Cons
 central copy for disaster	 big manual efforts
 very cheap	 minimum support only
 medium amount of data, some clients	 enhanced storage consumption
 original data formats in readable form	

Source Deduplication

To be able to stick with the simple configuration concept of directly backing up clients at a central backup server and also being able to backup bigger amounts of data or larger numbers of clients, it is recommended to use source deduplication. In opposite to target deduplication this means that the data source/client cuts the data to be backed up in chunks, generates a hash value for each chunk, matches the hash with a central index at the backup server and only transmits the chunks which are not yet available on the backup server. This can reduce the data to be transmitted significantly because not only changed files are transmitted but moreover only changed chunks and chunks which have not been transmitted earlier. Even if this technology increases the workload on the client a bit, the amount of data can be reduced by a factor of 10 or more depending on the current change rate. And by this it also reduces the duration of a backup. Moreover the amount of data will be reduced further more by compressing the chunks to be transmitted.

Because of the restore process having to cope with associated blocks being distributed all over the backup media, deduplication is a purely disk-based technology. Furthermore a deduplicated backup logically always represents a full backup making additional full backups obsolete. However with each backup only a few blocks will be submitted.

A dedup rate (ratio) will mostly be given in relations e.g. 10:1 which corresponds to a reduction of data of 90%. A reasonable dedup rate should be 3:1 at least to differ significantly from simple compression. In addition the extra costs for dedup licenses should be accompanied by a significant cost saving from disk storage. To estimate your own dedup rate it is recommended to set up a demo installation with your own mixture of data.

A sensible usage of this technology – i.e. a good dedup rate – is depending on several factors:

- Change rate (should be <5%)
- Retention (longer retention periods increases the dedup rate, but also the amount of data on the backup target)
- Global Deduplication (the more backup data is handled in the central index the better the dedup rate will be)
- Data formats (e.g. picture or video data are not suitable for deduplication)
- Encryption (data should be encrypted or compressed after deduplication)
- The deduplication algorithm (SEP sesam uses an own developed and four times patented, highly efficient algorithm with variable block size)

Issues strongly recommended to be considered with source deduplication:

1. Initial Backup

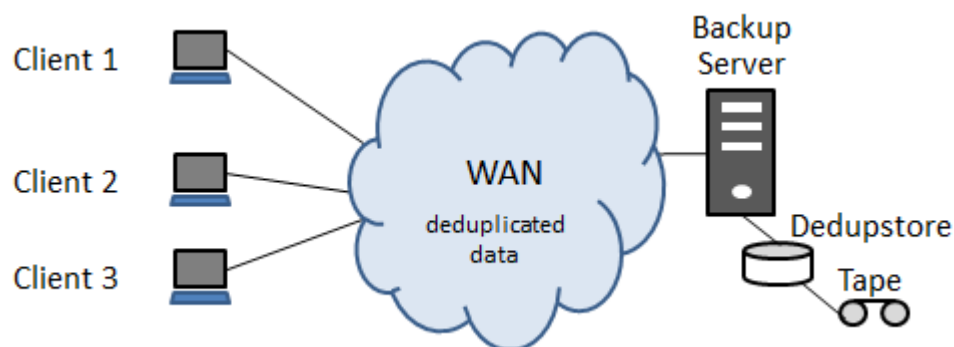
The disadvantage of a WAN connection will particularly emerge with the first backup because the full amount of data has to be transmitted at least once. This is of course depending on what is already in the dedup store. For instance the amount of data to be transmitted will be much less for the initial backup of a second Windows VM. It is recommended to allocate a dedicated time slot for the initial backup (e.g. weekend).

2. Restore

Because of the WAN connection the restore of single files or small amount of data is possible of course, however to restore bigger amounts of data have to be treated separately. Especially with disaster recovery when it is necessary to set up a complete server, specific measures have to be considered. A method well-known from customers is to set up a server in the data center and then send this server physically via express to the remote office. This concept of course requires considering the reachability of the remote location.











3. Backup to Tape

To fulfill the requirement for a backup to tape (e.g. archiving for compliance purposes) then this has to be accomplished in a follow-up task. It makes sense to 'rehydrate' the deduplicated data before tape-out i.e. to reverse the deduplication process.



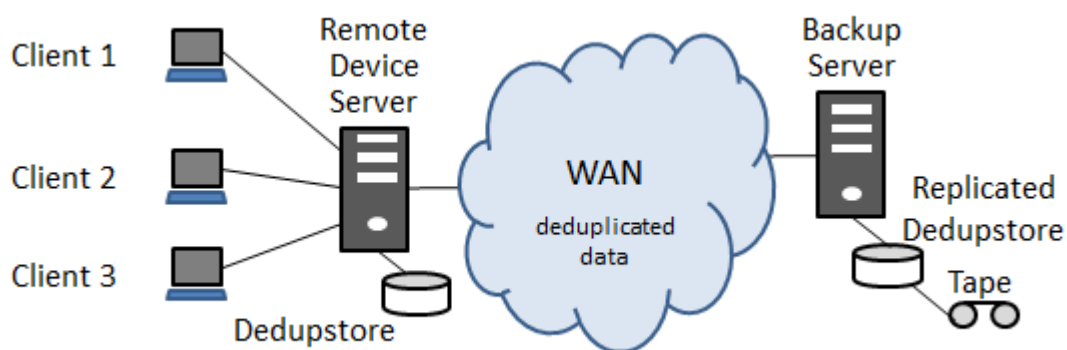
Example:

- 1 remote office with 2 clients in total 1TB of data
 - Change rate for blocks 2%
 - A WAN connection with 155Mbit/s (ATM)
 - Dedup rate of 10:1
- ⇒ Duration for a logical full backup =
 $1.000.000\text{MB} * 8\text{bits} * 0,02 * 0,1 / 155\text{Mbit/s} = \sim 1,7 \text{ minutes}$








Pros	Cons
 central backup	 initial backup (amount of data almost the same as without dedup)
 simple configuration	 big restores, disaster recovery
 permanent logical full backups	 must use disk, migration on tape as follow-up task
 medium amount of data, more than a few clients	 global deduplication recommended
	 puts a bit more workload on the clients
	 Increased risk of data loss in case of inconsistencies with the data storage

Replication with Deduplication

To make a replication more efficient this is done best by means of deduplication. The amount of data to be transmitted will be reduced so efficiently that with the current WAN technology (ATM) even a synchronous replication of data of almost any size is feasible. This technology in the past being a focus of the storage hardware vendors has now been well-established as a pure software feature. Replication via software has the advantage that the backup software keeps control over all the backups and in case of disaster recovery is able to restore from the replication at once. Hardware vendors tend to underline the software transparent replication as hardware benefit. On the other side they are trying to meet the requirements for cooperation of hardware and software vendors by supporting interfaces like Symantec OST (Open Storage Technology) or DD Boost for NetWorker with EMC Data Domain. However these interfaces are all proprietary to lock-in to a specific software vendor. Replication in SEP sesam will be accomplished by means of a remote device server (RDS) which can be located in a remote office with a dedupstore for the backups of the clients out there. The replication will be configured and triggered centrally at the backup server. Software replication allows very simple disk arrays to be used for deduplication.



Depending on the infrastructure there are several replication modes. n:1, a number of dedupstores will be replicated to a data center or 1:m, a dedupstore will be replicated to several locations. But also any combination n:m is possible. All modes will be available step by step with SEP sesam Si3R replication. Today the replication with SEP sesam is only asynchronous and configurable n:1 via the GUI.

Pros	Cons
 one or more data copies for disaster recovery	 costs for licenses in classic license model
 almost any data size	 costs for additional storage
 simple and cheap disk arrays can be used	 more powerful HW for RDS
 different replication methods are possible	

Backup of Laptops of Field Service

A special kind of remote offices are employees working not in an office but in the field. While using mobile laptops it is difficult to plan a scheduled backup. Also at locations, when employees take their laptops home overnight, there is no time window for full backups.

The first measure is a separate schedule for laptops. The backup shouldn't be planned overnight but rather during the day. 2 features might be helpful in these scenarios:

- **Wake on LAN (WoL)**

With this functionality which has to be switched on in the BIOS and has to be configured in the SEP sesam GUI the backup server can trigger a client to boot via MAC address over the connected LAN cable. After finishing the backup the client will shut-down depending on the local system settings.

- **Script to check the connection**

It might be useful to configure a PRE script to check the reachability of a client with a connection test before starting a backup. By using a counter it is possible to send an email with a warning to the owner or administrator and let him know that the last successful backup is long ago and that he should take care of this issue.

In addition there should be a possibility for laptops to start backups from the client if they have the adequate connection and enough time. The needed concept for multi-tenancy is currently under development for SEP sesam. The functionality to take over profiles from Active Directory or LDAP will be implemented in the WebUI in a second step. As a basic alternative in SEP sesam today field employees could install the adminGUI and do an immediate start from the access role "user".

Particularly efficient for doing full backups with laptops during the day is the usage of source deduplication. This reduces the duration of a full backup significantly making it acceptable for field workers.

Network Structures

While talking about backup in distributed environments it is necessary to consider the network structure. There are a lot of different cases and varieties and SEP sesam has to take care of each and all.

The SEP sesam Ports

SEP sesam communication is always a client-server-communication, that means a client connects a daemon. For the different SEP sesam daemons there are default TCP ports. The daemons are running on SEP sesam server, SEP sesam RDS or SEP sesam client. Which daemon runs on which machine depends on the installed module.

Service	Port	Description
sm_ctrlld	11301	unencrypted control communication (Ctrl)
sm_sshd	11322	control communication over SSH tunnel
sm_stpd	11001	port on device server for ftp data transfer
sm_stpd	11000	port on device server for http data transfer
sm_stpd	11443	port on device server for https data transfer
rmi GUI	11401	port on SEP sesam server for GUI connections
rmi Web	11403	webserver on SEP sesam server for RestAPI connections
sm_db	11201	database port on SEP sesam server (PostgreSQL)
Si3-T	117xx	port Si3-T Dedupe process on SEP sesam server and RDS

Most ports can be changed in the configuration files of the SEP sesam server, RDS or client. If entries or config files are not existing, add the file or entry manually.

To increase security communication with the client, it is also possible to establish an encrypted connection via SSH or HTTPS!

Firewall Environment

Increasing security requirements in businesses not only cause businesses to seal off externally from the internet, but also internally, between branches and areas.

In order to operate a backup client over such a firewall, two steps are necessary:


- Configuration of the firewall options in the Sesam client
- Giving access permission of the ports configured in the client in the firewall that you want to pass

Please notice that there are two possibilities for the control communication and three possibilities for the data transfer. Control communication and data transfer are working completely independent i.e. can operate in any combination.

In the default case SEP sesam uses multiple ports for communication. If company-internal security rules are very restrictive, using the http/https protocol allows to connect to the client through only one defined port.

Backup over an alternative Network

In order to backup files over a separate network, both the Sesam server, as well as the clients that are to be backed up, have to be given a 2nd network interface. After that, every 2nd network interface needs to be assigned its own IP address and a hostname. The figure below shows the network principle of using a backup LAN. Please do not mix up a separate LAN segment for backups with a SAN structure. In a SAN, the backup drives (shared drives) can only be used by a single host at a time, while in a 2nd LAN segment the parallelization of backups can also be used by Sesam.



IPv6

Because the current range for IP addresses is running short this implies a swap to the much bigger range of IPv6 in the near future. Software must be able to cope with longer and restructured IP addresses. This affects any software as like as operating systems, middleware or applications. And therefore backup software too.

For SEP sesam there is currently a lot of quality assurance for IPv6 in progress.