



SEP sesam:

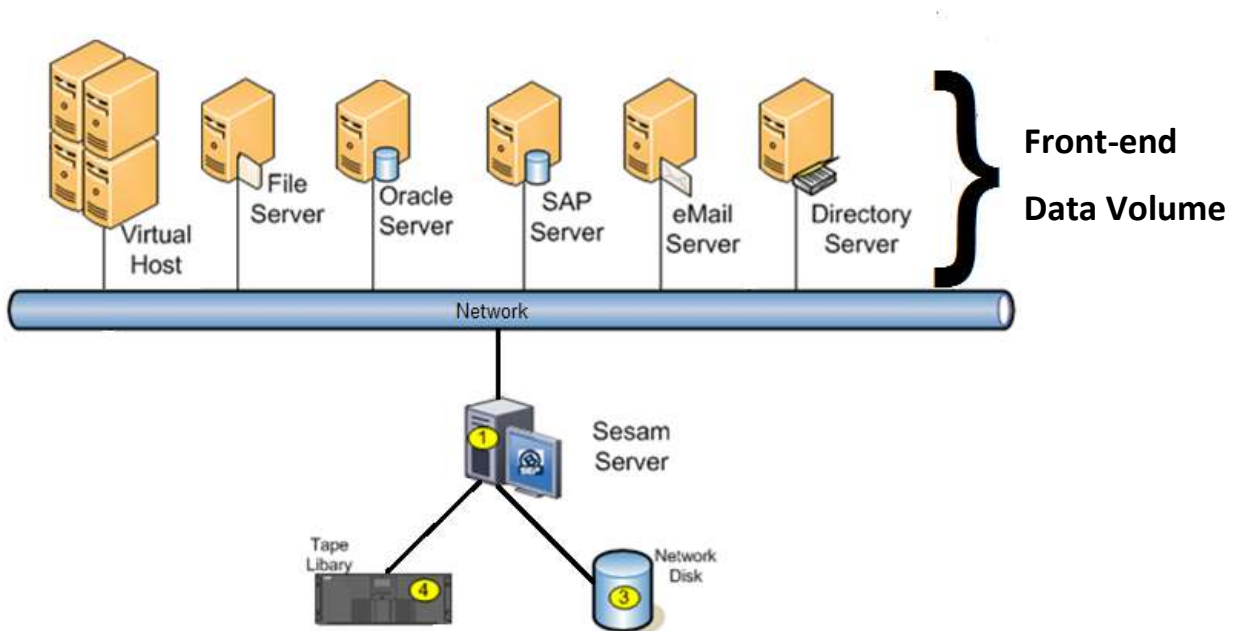
Description of Capacity Based Licensing

Calculation of Front-end TB (FET)

Introduction

Capacity based licensing has become increasingly popular over the past several years. This capacity based licensing calculates original, source-side, front-end backup data measured in TB (FET). Of course most companies continue to offer their classic component-based licensing model as well, the main focus however has shifted to FET.







Even if the biggest benefit of FET is its simplicity only a few vendors offer only one single TB license to cover all possible functionality while others have several TB models differentiated by supported features and databases. There are even models on the market combining FET with a component license model to bring together the benefits of both license worlds.



Overview

The front-end TB model is very popular because of its simplicity and extremely high flexibility with regards to the implemented backup infrastructure or backup strategy. On the other side it is directly linked with the data growth which could generate unplanned impact on the budget.

FET also has the advantage that only really used capacity has to be paid.


Advantages of FET	Disadvantages of FET
 Simple license administration and set up of quotes	
 High flexibility with regards to backup infrastructure and backup strategy.	
 „Pay as you use“	 Unplanned data growth can create unplanned costs
 Attractive especially in complex environments with the need for lots of agents	 Expensive in simple configurations with big amounts of data e.g. file server

Possible rule of thumb:

If in a complex environment (data zone) the average amount of backup data per client is more than 250GB most probably the classic component-based licensing model will be the cheaper solution.

Licenses

SEP sesam offers 2 types of front-end TB licenses graded by supported databases and number of parallel streams. A third license is provided for both editions to enhance the number of parallel streams in order to boost backup performance.



Contains essentially the same functionality as our competitors and even more:

- Good performance and support of a lot of databases and applications
- MS SQL, MS Exchange, PostgreSQL, EnterpriseDB, Open LDAP, IMAP, Zarafa, Novell Groupwise, OpenXchange, Scalix

Combines Professional Edition plus enterprise agents for

- MS Sharepoint
- Oracle
- Informix
- SAP (R/3, NetWeaver, HANA)
- Lotus Domino
- IBM DB2

Could be used to enhance both editions by 25 additional parallel streams

Volume Calculation

The FET model is very simple indeed because you only have to know the amount of data to be backed up. To determine this amount of data however is not trivial. Depending on the calculation algorithm of the backup software there are several possibilities which even more could yield to different results.

Possibilities to determine the amount of backup data:

- Manual cumulation of data: files + databases + VMs gives a rough estimation
- External applications e.g. another already installed backup software
- From the reporting of SEP sesam.

When using front-end TB licensing the amount of used data is visible in the daily protocol (here 1,5TB of used volume):

```

2014-12-08 08:00:22 I001-LICENSE[ 19873]:
2014-12-08 08:00:22 I001-LICENSE[ 19873]:
2014-12-08 08:00:22 I001-LICENSE[ 19873]: Volume Based License:
2014-12-08 08:00:22 I001-LICENSE[ 19873]:
2014-12-08 08:00:22 I001-LICENSE[ 19873]: 3.160 TB of 50 Volume Basis 50 TB + Volume Extension 0 TB Storage
2014-12-08 08:00:22 I001-LICENSE[ 19873]: {
2014-12-08 08:00:22 I001-LICENSE[ 19873]: FrontSide / Storage 1.500 TB / 3.160 TB ( 1 : 2.11 )
2014-12-08 08:00:22 I001-LICENSE[ 19873]: }
2014-12-08 08:00:22 I001-LICENSE[ 19873]:
2014-12-08 08:00:22 I001-LICENSE[ 19873]: - Automatic Media Read Check is enabled
  
```

When using the classic component-based licensing e.g. if SEP sesam is running in demo mode after a download, the amount of used data is visible in the daily protocol in the following format:

```
„Volume used: 1.500 TB (1536000 GB) FrontSide, 3.160 TB (3235840 GB) Storage“
```

As most of the backup software vendors are only starting step by step to introduce license enforcement for backup data cumulation, the value given by the reporting and the implemented algorithm of the respective software is gaining significance.

In SEP sesam too the amount of backup data is enforced by licensing.

Short description of the algorithm:

All SEP sesam capacity licenses are based on how much data is protected on the source side (Front-Side-TB). The maximum volume of all backup jobs is cumulated as long as it is stored on any SEP sesam media (retention policy).
An example: There is a system full backup, an image backup on hypervisor level and a database backup with the database backup agent. If the database is part of all three backup jobs this data is counted three times. The real data is cumulated; before any deduplication and before any compression.

Hint:

There will be cumulated the maximum size (`data_size`) of each executed backup task (INC/DIFF/FULL/COPY) – normally this is a COPY or FULL –, but this could mean in very specific cases that if a DIFF or INC backup task is bigger than a COPY or FULL (e.g. after changing the exclude list), then the DIFF or INC will contribute to the total amount and not the COPY/FULL.

Exceeding the limit

Currently SEP sesam cumulates the data and will check against the installed licenses after all the backup processes are finished. If the limit is exceeded the usual message will appear that the license will expire in 15 days and that you should immediately contact your local sales representative.

Examples

Presumptions for the examples:

- One full backup per week and incremental backups during the week
- Retention Time of the full backup: 3 weeks
- Incremental counts about 10% of the respective full

Example 1

We only have one task for file backup.

	FULL	INC1	INC2	INC3	INC4	INC5	Σ	TARGET	FET
1. Week	A: 100GB	A1: 10GB	A2: 10GB	A3: 10GB	A4: 10GB	A5: 10GB	150GB	150GB	100GB
2. Week	B: 150GB	B1: 15GB	B2: 15GB	B3: 15GB	B4: 15GB	B5: 15GB	225GB	375GB	150GB
3. Week	C: 110GB	C1: 11GB	C2: 11GB	C3: 11GB	C4: 11GB	C5: 11GB	165GB	540GB	150GB
4. Week	D: 80GB	D1: 8GB	D2: 8GB	D3: 8GB	D4: 8GB	D5: 8GB	120GB	510GB	150GB
5. Week	E: 100GB	E1: 10GB	E2: 10GB	E3: 10GB	E4: 10GB	E5: 10GB	150GB	435GB	110GB

Backup Week 1

Sunday: Full A => 100GB
 Monday: INC A1 => 10GB
 next days similar INC

Valid Front Side Volume ==> 100GB

In the backup storage resides
 Full A = 100GB + all INCs

Backup Week 2

Sunday: Full B => 150 GB

Valid Front Side Volume ==> 150GB

In the backup storage resides
 Full A + Full B = 250GB + INCs

Backup Week 3

Sunday: Full C => 110 GB

Valid Front Side Volume ==> 150GB

In the backup storage resides
 Full A + Full B + Full C = 360GB + INCs

Backup Week 4

Sunday: Full D => 80 GB

Valid Front Side Volume ==> 150GB

Release (EOL) of Full A.
 In the backup storage resides
 Full B + Full C + Full D = 340GB + INCs

Backup Week 5

Sunday: Full E => 100 GB

Valid Front Side Volume ==> 110GB

Release (EOL) of Full B and hence the maximum
 volume comes from Full C (110GB).
 In the backup storage resides
 Full C + Full D + Full E = 290GB + INCs

Example 2

One physical server:

- 50GB of data: OS + file data
- SEP defines one backup task. In the scheduler there will be a full backup of this task on Sundays and incremental backups during the week

The FET volume will be **50GB** as it is the biggest size of this backup task this week.

Example 3

One physical server:

- 100GB of data: 10GB OS + file data 50GB + MS Exchange DB 40GB
- One backup task: OS + file data = 60GB
- One backup task: OS only for Bare Metal Recovery = 10GB
- One backup task: MS Exchange DB = 40GB

The FET volume will be **110GB** as this is the sum of the 3 biggest backup tasks this week.

Beispiel 4

Assuming the server from example 3 is running as a VM (VMware hypervisor):

- One backup task: OS + file data = 60GB
- One backup task: OS only for Bare Metal Recovery = 10GB
- One backup task: MS Exchange DB = 40GB
- In addition: One backup task for VMware snapshot backup of the entire VM = 100GB

The FET volume will be **210GB** as this is the sum of the 4 biggest backup tasks this week.

FAQ

1. Will SEP sesam also count all the incremental backups after a full?

No, only the biggest job will be taken into account. This will normally be a full or a copy backup, however in specific cases this could also be an incremental or differential backup (see hint in chapter „Volume Calculation“), but never the sum.

This is a major difference to the calculation of other vendors because the data volumes relevant for SEP sesam licensing will therefore be significantly smaller.

2. How will data of VMs be counted?

INC/DIFF/FULL/COPY? CBT? Null-byte reduced snapshots?

The sizes of the snapshots will be counted exactly the size how VMware provides the snapshots i.e. after CBT and null-byte reduction.

In case the snapshot backup will be defined as a separate backup task e.g. for Instant Recovery or Single File Restore, the biggest job of this task will be added to the other maximum backup jobs (see example 4).

3. When will data be counted twice e.g. with database and file agents, COPY-jobs, migration, replication?

Only the biggest job of a backup task will be counted (see above).

If identical data will be backed up in different tasks or via different agents they will be counted multiple times. A COPY job is just treated as a FULL backup.

Migration and replication tasks are transparent to the algorithm. They are not counted.

4. Is it possible that for front-end data reduction I could simply switch on compression on file system level or even deduplication with Linux (btrfs) or Windows 2012 (NTFS)? Both file systems provide post processing deduplication.

For a backup with SEP sesam agents (no matter if file or DB backup) the file system provides real data. The backup will be executed on file level.

Therefore compression or deduplication of the files on file system level does not change anything of the file sizes. A file of 20MB will still be a file of 20MB in the backup job.

With VMware backups however this can be different because VMware provides snapshots of the virtual system. This could have the effect that if the system compresses or deduplicates the data, the snapshot represents the compressed system and the backup volume of the front-side job will be smaller.

5. How can I delete data from the front-end data calculation?

If savesets from the biggest backup job are deleted the front-end data volume will be decreased. This happens automatically by means of the retention policy.

Please notice that the sheer expiration of the retention policy (EOL of a saveset is older than the current date) does not necessarily result in a data reduction. This happens only after physical deletion of the data.

This means in particular:

- a. In a datastore the expired savesets will be deleted not before reaching the HW mark. Instead of waiting for this event it is also possible to launch this event manually from the GUI via "Purge".

Please be aware that a purge will stop after reaching the LW mark. This has the result that savesets will still not be deleted from the datstore even if the retention has expired.

Example: HW 500 GB, LW 300 GB, EOL savesets take 100 GB, then about 200GB expired savesets will be kept in the datastore.

Solution: Reduce LW temporarily e.g. to 100GB.

This problem can be avoided generally if LW will be set to 0. Then always all expired savesets will be deleted with a purge.

- b. On tape media the savesets will be kept until the whole media will be erased.

Solution: Migrate savesets to other media. Reinitialize media (might be a complex task).