

d.velop

d.ecs storage manager:
Administrator

Table of Contents

1. d.ecs storage manager	4
1.1. Introduction	4
1.1.1. About d.ecs storage manager	4
1.1.2. Prerequisites	4
1.1.3. Licensing	4
1.2. Installation	5
1.2.1. Installing d.ecs storage manager with d.velop software manager	5
1.2.2. Installing updates for d.ecs storage manager with d.velop software manager	5
1.2.3. Uninstalling d.ecs storage manager with d.velop software manager	8
1.2.4. Rollback of an installation of d.ecs storage manager using d.velop software manager	8
1.2.5. Adopting an old document index in the database	8
1.2.6. d.3 server preparations	8
1.2.7. Database recovery	8
1.3. Configuration	9
1.3.1. Create/delete configuration	11
1.3.2. Tab: Database	11
1.3.3. Tab: General	12
1.3.4. Tab: Systems	13
1.3.5. Tab: Pool Mappings	15
1.3.6. Tab: Database logging	17
1.3.7. Tab: HTTP / API Interface	18
1.3.8. Tab: Cache	19
1.3.9. Tab: Logging	19
1.3.10. Tab: Single Instance Store	20
1.3.11. Tab: Miscellaneous	21
1.3.12. Tab: Deletion	22
1.3.13. Tab: Retention process	23
1.4. d.ecs storage manager modules	24
1.4.1. d.ecs storage manager Celerra	25
1.4.2. d.ecs storage manager Centera	28
1.4.3. d.ecs storage manager cloud storage	31
1.4.4. d.ecs storage manager Data Domain	35
1.4.5. d.ecs storage manager FileLock	39
1.4.6. d.ecs storage manager HCP	42
1.4.7. d.ecs storage manager iCAS	43
1.4.8. d.ecs storage manager NAS	44
1.4.9. d.ecs storage manager NetApp	47
1.4.10. d.ecs storage manager Silent Cubes	53
1.4.11. d.ecs storage manager TSM	62
1.4.12. d.ecs storage manager Governikus LZA/DATA Aeonix	64
1.4.13. d.ecs storage manager system connect	68
1.4.14. d.ecs storage manager Isilon	69
1.4.15. d.ecs storage manager S3	72
1.4.16. d.ecs storage manager proNEXT	76
1.4.17. d.ecs storage manager Azure	77
1.5. Opening d.ecs storage manager	79
1.6. Web-Interface	80
1.6.1. Document information	81
1.6.2. Driver information	82
1.6.3. Job overview	83
1.6.4. Pool overview	84
1.6.5. Process overview	92

1.6.6. Support Package	93
1.6.7. System check	94
1.6.8. System information	96
1.6.9. System mirroring on document level	97
1.6.10. License information	106
1.7. Additional hints	107
1.7.1. Jukebox directories	107
1.7.2. Multiple instances of d.ecs storage manager	107
1.8. Deleting systems	107
1.9. d.ecs storage manager Status overview	108
1.10. Monitoring	109
1.10.1. Workload	109
1.10.2. Speeds	109
1.10.3. Job status	109
1.10.4. Pool status	110
1.11. Appendix	110
1.11.1. Troubleshooting	110
1.11.2. Meaning of the different file names of files in the Jukebox-directories	112
1.11.3. Optimize the network access	114
1.11.4. Create SSL-certificates with OpenSSL create	114
1.11.5. Retention periods/Retentiontimes	114
1.12. Glossary	120
1.12.1. Store	120
1.12.2. Jukebox directory/Jukebox public directory/Jukebox archive directory	120
1.12.3. References/Reference data/Reference entries/Reference table/Reference index/Index data	120
1.12.4. Secondary storage/Secondary storage I/O	120
1.12.5. Storage system/Storage module/Storage driver	120
1.12.6. System table	121
1.12.7. Retrieve	121
1.13. Additional information sources and imprint	121

1. d.ecs storage manager

1.1. Introduction

In this chapter you will find general product information.

1.1.1. About d.ecs storage manager

The application d.ecs storage manager is used for long-term stable archiving of documents on an external storage system.

As external storage systems, you can connect various systems of well-known vendors of audit-proof document storage systems (not-editable, not-deletable).

1.1.2. Prerequisites

This manual is mainly targeted at d.3 and Microsoft SharePoint/ecspand administrators.

To fully understand its content, you need some knowledge of the d.3 / Microsoft SharePoint or ecspand system as well as profound knowledge of operating systems and network systems.

1.1.3. Licensing

For d.ecs storage manager, three license types are available which are requested and booked via d.ecs license server.

1. **Synchronization**To be able to synchronize from one system or pool to another, you need the license for the synchronization.
2. **d.ecs storage manager modules**
You need a license for each d.ecs storage manager module used. Quotas are booked from the d.ecs license server for the license. For each quota, 100 GByte can be written to the respective secondary storage system.
You get a notification via the monitoring, if the quota is has been booked by 90% and 100%. If the quota is used up to 100%, you have 31 days to install a new license or else all pools containing systems using this module are disabled for secondary storage.

Warning

As the licensing of the d.ecs storage manager has changed with version 3.1, the d.velop AG provides an application to determine the data volume to be licensed before the update to d.ecs storage manager version 3.1.

3. **Privileged deletion**To delete documents prematurely before the retention period expires, you need a license for the "privileged delete" option.

1.2. Installation

Warning

As part of the document handover to an audit-proof storage, the documents are usually stored with a "retention time" (retention period). The retention period is passed from d.3 or ecspand to d.ecs storage manager together with the document.

The configuration of the retention periods in d.3 or ecspand must correspond to those of the storage systems as errors can occur otherwise, which may lead to a document that cannot be stored. Usually a minimum, a maximum, and a default retention time must be configured on these systems.

Additional information can be found in the respective documentation for the system to be connected.


Warning

Some external storage systems with respective modules available for d.ecs storage manager only ensure non-editable/non-deletable document storage only then, if certain software options or licenses are installed on the respective system.

The supported storage systems are listed in the chapter [d.ecs storage manager modules](#).

Furthermore, d.ecs storage manager features an interface which can be used to add additional storage drivers to the application.

Thus, different storage systems using different storage technologies can be added to an existing system later.



d.ecs storage manager has two possible interfaces for communication with the leading systems (d.3/ecspand). d.3 and d.ecs storage manager use a combination of a directory structure and a database table for the communication, while ecspand uses the API interface provided by d.ecs storage manager.

1.2.1. Installing d.ecs storage manager with d.velop software manager

You install the software exclusively using d.velop software manager. If an application is required for different products, the corresponding software packages are also installed automatically.

For further information on installing the software, see the d.velop software manager manual.

1.2.2. Installing updates for d.ecs storage manager with d.velop software manager

You can only update the software using d.velop software manager.

For further information on updates, see the d.velop software manager manual.

Note

If an older version of d.3 or d.ecs storage manager is already installed, create a backup of your configuration files before installing.

Running a database update after installing an update

As the table structures of the d.ecs storage manager tables may have changed during updates of the software, d.ecs storage manager may update the table structures when starting a new version for the first time and continue the start process afterwards.

If the automatic update is aborted due to an error, this will be logged in the d.3 logfile and additionally sent to d.ecs monitor. In this case, the program for updating the table structures must be run manually to perform further analysis.


The update program named **DBUpdater.exe** checks the table structure and the indexes and adapts them for the new version if necessary.

You can find the program in the d.ecs storage manager application directory. It uses the configuration file (ini file) of d.ecs storage manager to log in to the database.

The program has to be executed as follows:

```
DBUpdater.exe <ini-configuration file of d.ecs storage manager>
```

The program then returns messages like this:



```
Administrator: C:\Windows\system32\cmd.exe
d.ecs storage manager Database Updater
Version 2.7.0.0
(c) d.velop AG. All rights reserved.

20.10.2015 09:47:28: Initializing DB-Connection...
20.10.2015 09:47:28: Connected

Checking pool ids...
20.10.2015 09:47:28: Updating "sourced_out"-fields...
0 sourced_out fields updated
20.10.2015 09:47:28: Updating "pool_id"-fields...
0 records updated
20.10.2015 09:47:28: Changing d3sm_references.size_original to BIGINT
Dropping constraint DF_d3sm_refe_size_58D1301D
20.10.2015 09:47:28: d3sm_references.size_original changed!
20.10.2015 09:47:28: Changing d3sm_references.size_compressed to BIGINT
Dropping constraint DF_d3sm_refe_size_59C55456
20.10.2015 09:47:28: d3sm_references.size_compressed changed!

Checking "deleted"-column for NULL values
0 deleted fields updated
20.10.2015 09:47:29: Creating new indexes. This could take several minutes...
20.10.2015 09:47:29: Skipping index creation / update for d3sm_SYS_HASH. Index already exists and has right fields!
20.10.2015 09:47:29: Skipping index creation / update for d3sm_SYS_DOCID. Index already exists and has right fields!
20.10.2015 09:47:29: Update successful
The update took 0 day(s), 00:00:00
```

If you are prompted to select an archive ID, then you must select the main archive ID of the archive here.

To determine the steps to be executed as part of a table update, the **DBUpdater** can be started with the parameter **/CHECKONLY**.

```
DBUpdater.exe <ini-configuration file of d.ecs storage manager> /CHECKONLY
```

The program then returns messages like this:

```

Administrator: C:\Windows\system32\cmd.exe
d.ecs storage manager Database Updater
Version 2.7.0.0
(c) d.velop AG. All rights reserved.

20.10.2015 09:46:17: INFO: Only update checks will be executed!
20.10.2015 09:46:17: Initializing DB-Connection...
20.10.2015 09:46:17: Connected

Checking pool ids...
Counting documents...
Total docs: 676317
Checking "sourced_out"-fields...
"sourced_out"-fields to be updated: 0
Checking "pool_id"-fields...
"pool_id"-fields to be updated: 0
Field "size_original" has to be changed to BIGINT
Field "size_compressed" has to be changed to BIGINT

Checking "deleted"-fields...
"deleted"-fields to be updated: 0

Database check complete!

```

Note

Adjusting the d.ecs storage manager tables may take a while for large volumes of data as new fields may have to be added and (for existing data) may have to be filled with values. Please plan for a down-time of the d.ecs storage manager for this respect. Should you still have reference entries of old, no longer existing systems in the reference table of the d.ecs storage manager, then these should be deleted before the update to reduce the data volume pending to be updated. Please contact the d.velop support if necessary.

Note

Also make sure that you have sufficient table space for the changes. If the DBUpdater is aborting due to database errors (such as lack of disk space), it can be started again after having resolved the problem.

The **DBUpdater** executes the following steps on the table **d3sm_references**:

- Check of the column **sourced_out**
- Check of the column **sourced_out_time**
- Check of the column **pool_id**
- Check of the columns **size_original** and **size_compressed**
- Check of the column **reference**
- Check of the column **signature_cid**
- Check of the column **volume_id**
- Check of the column **system_check_result**
- Checking the column **evidence_state**
- Check of the column **delete_token**
- Check of the column **reference_index**

MSSQL: The field is created and directly populated with values, a unique index is generated.

DB2: The field is created with default 0, the default is dropped, the default is set to **IDENTITY**, the field is set to DEFAULT in blocks of 100,000, a unique index is generated.

Oracle: The field is created, a sequence is created, the field is populated in blocks of 100,000 via the sequence, a unique index is generated.

- Check of the column **retention_on_storage**
- Check of the column **deleted**
- Change of the column **job_date** in the table **decssm_jobs** from DATE to **TIMESTAMP** (only DB2)
- Check of the index **d3sm_sys_hash**

- Check of the index **d3sm_sys_docid**
- Check of the index **d3sm_sys_delete**
- Check of the index **d3sm_sys_ref**
- Check of the index **d3sm_sourced_out**
- Check of the table **d3sm_version**

Note

Please execute the following additional steps:

1. Open the administration interface and save the configuration there. This is necessary to confirm the now configurable **DB-Filegroups**.
2. At the first update via d.velop software manager: Adjust the call parameters for d.ecs storage manager in the d.3 process manager.

1.2.3. Uninstalling d.ecs storage manager with d.velop software manager

The software you installed using d.velop software manager can only be uninstalled with d.velop software manager. If the software to be uninstalled has dependencies with other software packages, you must resolve these conflicts accordingly.

For further information on uninstallation, see the d.velop software manager manual.

1.2.4. Rollback of an installation of d.ecs storage manager using d.velop software manager

You can restore an earlier version of the software that you installed with d.velop software manager. During this process, the software is only reset to a previous version.

For further information on rolling back to an earlier version, see the d.velop software manager manual.

1.2.5. Adopting an old document index in the database

From version 2.0 of d.ecs storage manager, the d.3 database is used to store the index data.

To migrate the file based index (d.ecs storage manager < version 2.0, originally: d.3 storage manager) to the database, the d.velop AG provides the tool d.ecs storage manager Reference Converter.

This tool reads an index directory and uses the information to create a file which then can be transferred to the d.ecs storage manager for further processing. This tool can be obtained from the d.velop support.

1.2.6. d.3 server preparations

For the d.3 server to work with the d.ecs storage manager, some parameters have to be configured in the d.3 administration:

- The use of a secondary storage system has to be enabled in the d.3 configuration under Secondary Storage (parameter **jukebox_daemon**).
- The parameter **Automatic Use of Secondary Storage** must be enabled under **Secondary Storage (AUTO_DOCS_ON_SEC_STORAGE)**.
- Check the secondary storage public directory and archive-directory (**SEC_STORAGE_PUBLIC_DIR**, **SEC_STORAGE_ARCHIVE_DIR**) in the d.3-Konfiguration section **Secondary storage**.
- To store documents, the parameter **Secondary storage I/O** has to be activated for the respective document types.
- A d.3 storage async process must be set up.

1.2.7. Database recovery

d.ecs storage manager logs the entries it makes in a database in chronological order. If entries in the database are lost, they can be recovered using these logs. This requires the following steps:

Prerequisite:

- The *.IR files were requested from the storage system by the recovery mode of the respective module (see module-manual)
- alternative: the *.IR files are directly accessible in the storage path of the storage.
- alternative: the d.ecs storage manager was configured to store the *.IR files in a backup directory and they are thus directly accessible.
- Copy the *.IR files which are to be processed into a directory.
- Start the d.ecs storage manager via the command line as follows:

```
decssm.exe <Name of the configuration file> recoverdb="<Path to the files>"
```

Thus a call could look as follows:

```
decssm.exe D3P recoverdb="c:\recoverdbfiles\"  
In this case the *.IR files must have the following filenames:  
D3SM_D3P-JJJJMMTT-HHMMSS.IR
```

During the processing, the Log Viewer shows entries signaling which file is currently being processed.

Warning


The recovery via the d.ecs storage manager must not be started via the d.3 process manager. This way, the data might be entered into the database twice. This could lead to problems, later.

1.3. Configuration

You can find the menu option **d.ecs storage manager administration** under **Start > Programs > d.velop > d.ecs storage manager**.

With this menu item you start the configuration mask of d.ecs storage manager.

The program automatically searches for existing configurations on start, selects the first one as the default and loads its settings. If no configuration exists yet, the program prompts for a name of a new configuration to be created. Then, the following dialog is displayed:



Note

The individual configurations / configuration files for storage drivers are written to the sub-directory **configs** in the application directory of d.ecs storage manager by default.

Backup copies of the configuration files should be generated after finishing the configuration to be prepared if the files got lost or were accidentally overwritten.

To use a different directory you have to provide the path when starting the configuration tool. This is important, if you need to configure multiple instances of the d.ecs storage manager for different archives.

By specifying the different directories, the configuration files of the different instances are separated from each other and faulty configurations are prevented. This can, for example, happen if the configurations of the different instances are stored in the same directory and system names are used more than once in the configurations.

Warning

During configuration, folders may need to be created and specified accordingly. An example of this is the log folder for database logging. These folders and also other folders or files must not be created in the program directory of d.ecs storage manager. During an update, these files and folders are removed from the program directory. The directory "configs" is the only directory that is taken over during an update.

Warning

After a change in the configuration d.ecs storage manager must be restarted.

1.3.1. Create/delete configuration



: Add a new configuration. Via the input field you are prompted to specify a name for the new configuration.

Note

In a d.3 environment it is recommended to use the name of the d.3 instance (e.g. D3P) as a name for the configuration as well.



: Delete a configuration.

Note

After an according user confirmation, the active configuration is deleted, i.e. the configuration file is removed. At the same time, all systems defined in the configuration are removed from the module configuration files. By deleting a configuration, no changes are applied to database tables of the d.ecs storage manager.

1.3.2. Tab: Database

Since the d.ecs storage manager has to make large amounts of reference data available, these data are stored in a database. Thus an installation of a database system is presumed. The system table is created by the configuration program. The reference table needed will be created when you start the d.ecs storage manager for the first time.

DSN: DSN (Data Source Name) to access the d.3 database or the database created by the d.ecs storage manager. In a database created by the d.ecs storage manager you must make sure that the template SQL-template-standalone is selected during the installation.

User: User name which is required to establish the connection

Password: Password which is required to establish the connection

Confirm password: Password which is required to establish the connection Repeat here to avoid typing errors.

Use Windows login: Instead of the specified user and password, the Windows user with which the d.ecs storage manager is executed is used to login to the database (only for Microsoft SQL Server).

Database system: The database system used has to be selected here as the necessary tables are created according to the system.

Test connection: Tests the database connection with the specified information.

Database template: Specifies the database template to be used to create the tables and indexes.

sql.template.d3: Template for the d.ecs storage manager in the d.3 context.

sql.template.lza.d3: Template for the downstream d.ecs storage manager in the d.3 context (only used when operating Governikus LZA systems).

sql.template.standalone: Template for the d.ecs storage manager in the ecspand context or when operated stand-alone.

sql.template.lza.standalone: Template for the downstream d.ecs storage manager in the ecspand context or when operated stand-alone (only used when operating Governikus LZA systems).

Store DB connections: Specifies the maximum number of database connections allowed to be established to secondary storage of documents.

Restore DB connection count: Specifies the maximum number of database connections allowed to be established to restore documents.

Tablespace Data/Indexes: Assignment to tablespaces for data and indexes within the database. Only displayed if a database template has been configured for the d.3 environment.

1.3.3. Tab: General

Jukebox path: Specify the Jukebox directory for the archive here (e.g. C:\d3\D3A.dok\jukebox). This directory is checked for restore jobs. When using d.ecs storage manager in connection with ecspan, the directory must be created manually.

Public path: Specify the Jukebox public directory for the archive here (e.g. C:\d3\D3A.dok\jukebox\public). This directory is checked for documents in status **Release** (formerly status "Public") which are to be stored. When using d.ecs storage manager in connection with ecspan, the directory must be created manually.

Archive path: Specify the Jukebox archive directory for the archive here (e.g. C:\d3\D3A.dok\jukebox\archiv). This directory is checked for documents in status **Archive**. When using d.ecs storage manager in connection with ecspan, the directory must be created manually.

Store job count: You can determine how many secondary storage jobs are to be processed at the same time here. When using ecspan, the number of simultaneous jobs depend on the project and must be determined separately.

Note

When using the module TSM or Centera, you should note that usually the number of module connections should at least be configured matching the job number.

Restore job count: You can determine how many restore jobs are to be processed at the same time here. The specified number should comply with the number of jukebox jobs configured in d.3 gateway.

Job repetitions: This parameter defines how many times a storage job is to be continued before being considered as erroneous. The default value is 20.

Note

If you specify a number >0 for the number of job retries, and this number of failed attempts is reached, then the secondary storage is aborted and the job receives an error state. The job overview of the web interface lists these jobs with the error text "The maximum number of secondary storage attempts was reached" and the error number 16.

Asynchronous repetitions: This parameter defines how many times a storage job is to be continued by the storage system during asynchronous processing before being considered as erroneous. The default value is 100. Asynchronous processing is supported only for Silent Cubes and Silent Bricks. For more information see **Wait for SAFE-file** in chapter [Configuration](#).

Max. wait jobs: This specifies the maximum number of storage jobs may be put on hold due to processing errors. If this value is exceeded, no more jobs are processed until the number falls below the maximum again.

Note

In combination with the Silentcube module and the option **Wait for Safe file** you should, deviating from the default value 100, set the value of **Max. Wait-Jobs** to e.g. 300. This leads to a faster secondary storage on a Silent Cube system.

Note

If the d.3 server is running on a separate system, the jukebox directories on this system have to be shared and addressable via UNC path. Note that d.ecs storage manager requires both read and write access to the shared directories used.

1.3.4. Tab: Systems

Every storage system used by the d.ecs storage manager has to be defined. These systems are later assigned to the individual pools in the [Pool mapping](#). You can determine which systems are to be used for storage and which systems should only be used for document retrieval.

Add system: Opens the dialog to define a new storage system.

Configure system: Opens the dialog to configure the selected storage system.

Set system parameters: Opens the parameter configuration dialog of the module which the storage system is connected to.

Index recovery: Opens the dialog for the recovery of the reference index of the selected storage system

Delete system: Removes the selected system from the configuration. The system is marked for deletion and removed from the configuration. The database references will remain in the database until the application `decssmdelsys.exe` is executed. For more information refer to the chapter [Deleting systems](#).

Warning

After the system has been marked for deletion the configuration has to be saved and the d.ecs storage manager has to be restarted.

System list (in database): This parameter is only relevant for a recovery of the configuration of the d.ecs storage manager and shows all storage systems registered in the d.3 database / ecspand database.

Note

After a system has been added, it has to be configured.

- For this purpose, click **Set system parameters**.

For more information on how to configure the module refer to the section [d.ecs storage manager modules](#).

Add system

System konfigurieren [X]

Bezeichnung:

Schnittstelle:

☒ **Auslagern**

Dokumenttypen:

☒ **Wiederherstellen**

☐ Wiederherstellen bei Fehler

☐ Überprüfung bei der Wiederherstellung

Name: Define a name for the new system here. The name has to be unique. Valid characters are A-Z , 0-9 , _ (underscore). The name can have up to 20 characters.

Note

The name for a system must also be distinct across several different configurations running on the same system. If, for example, you create a configuration for a D3A and for a D3P repository and you want to store on the same system with different volumes, then the system names must be different. The background for this is that for all systems using the same d.ecs storage manager module only one single configuration file (CFG file) is written. The configuration file is structured into sections per system. This structure is created based on the system name.

As an alternative, you can use different directories for the configuration per instance.

Interface: Here, you can define which storage module the system is supposed to work with. The respective module designations can be taken from the [Module overview](#).

Note

Only those modules are listed which have been installed during the setup. Moreover, a valid license is required for the module to be used.

Secondary storage This parameter has to be activated if data are to be stored on this system.

Warning

For modules, only supporting the **read-only** mode, this parameter must be disabled.

Document types: Specifies the document types to be stored on the system. The possible options are:

- all documents
- only documents with (embedded/external) signatures
- only documents without signatures (this includes: workflow-protocols, dependent documents, attribute files)

Restore: This parameter has to be activated if data are to be retrieved from this system.

Restore in case of an error: If this parameter is enabled, then d.ecs storage manager attempts to load documents from this system, even if it currently has the state ERROR in the d.ecs storage manager. Enabling this parameter can lead to longer delays when restoring documents, if for example a system is not physically available and a timeout has to be kept.

Note

If you only configure one system per pool, you should not disable this parameter.

Check on restore: If this parameter is enabled, a hash-check is applied after the restoring of a document. Enabling this parameter can lead to longer delays and higher CPU-load when restoring documents, due to the hash-value validation. This parameter only affects the general provision of documents. The processes **Synchronization**, **System check** and **Document check** are not affected.

Note

There is a d.3 Shortcuts video on this topic: [Configuration of the d.ecs storage manager with FILELOCK](#)

1.3.5. Tab: Pool Mappings

Within the d.3 repository, every document is assigned to a repository ID. The pool ID usually equals the repository ID, unless a different pool ID has been explicitly defined for a document type.

In the ecspand context, the Pool-IDs to be used are configured separately. Using the Pool mappings in the d.ecs storage manager you can define which systems the documents from a pool are to be stored on and in which order. The order is relevant here since the d.ecs storage manager retrieves the documents from the systems in that order during a recovery process.

Pool-Zuordnungen			
Konfigurierte Pools			
Designation	Pool ID	Systems	Pool hinzufügen
C	C	CELERRA	Pool Konfigurieren
CAS-1	CA	CAS1	
D	D	EMCDD_1	
H	H	HCP503	Pool entfernen
HC	HC	HCP60	
IC	IC	ICAS_	
K	K	TSM_CERT	
N	N	NETAPP81	
S	S	NAS-1,NAS_1_STATISTIC	
SC	SC	SC_RET	
SD	SD	SD1	

Add pool: Opens the dialog where you can add a pool.


Configure pool: Opens the dialog where you can configure the selected pool.

Delete pool: Deletes the selected pool assignment from the list.

Note

Consider the speed of the systems before determining the order. The fastest system should be first and the slowest last. This guarantees that the document can be retrieved as quickly as possible.

Mapping dialog



Name: Specify a freely selectable name (50 characters) for the pool here.

Note

It is recommended to use a descriptive name as a pool name such as **POOL-A**. In the d.3 context, the documents written to the secondary storage are reported back to the d.3 system based on the pool name. In the d.3 smart explorer, the document property dialog (tab **Versions**) lists the Pool name in the column Medium.

Pool ID: Specify the pool ID of the pool whose documents are to be stored.

Note

The pool ID is equivalent to the d.3 repository ID that can be associated to a document type and which can consist of one or two characters (letters A-Z).

Chosen systems: Select the systems on which documents are to be stored and from which the documents are to be restored.

Warning

If you are associating more than one system to a pool, then the order should be defined in descending speed of the storage systems, since the documents are stored and restored in this order.

Available systems: In this list, only those systems are displayed which have previously been created via **Add system**.



: Move a system from the list of **Available systems** into the list of **Chosen systems**.



: Remove a system from the list of **Chosen systems**. The system is then moved to the list **Available systems** and the pool is thus no longer used.

: Move a system upwards in the list **Chosen systems**. This affects the order in which the systems are addressed during the secondary storage or recovery of documents.


: Move a system downwards in the list **Chosen systems**. This affects the order in which the systems are addressed during the secondary storage or recovery of documents.

1.3.6. Tab: Database logging

Warning

d.ecs storage manager creates log files where the entries written to the reference table are logged in chronological order. As it is not possible to restore documents without the d.ecs storage manager reference table, these log files must be backed up regularly. An integrated, time-controlled backup function is provided in d.ecs storage manager for this purpose. This stores the log files (in compressed form) on the storage system and optionally in a directory (Copy logs to). The target directory must then also be backed up by a backup application.

Since this kind of logging is no replacement for a database backup, we recommend to create regular backups of the reference database of d.ecs storage manager.



Log path: Specify the path for the log files required for rebuilding the document reference table here.

Specify the path for the log files required for rebuilding the document reference table here.

Warning

Please enter a local path here. Saving the log-information via the network could result in problems, if the access to the network share is interrupted. If an entry cannot be written into the log file, then the storage is ended immediately. Additionally, this path must not be located under any storage-path (SnapLock area, CAS-areas etc.).

Destination pool: The log files are automatically stored in this pool. Enter the Pool ID of the storage pool where the log files are to be stored here.

Export log files: Specify the time interval when the log files should be written to the secondary storage here. The following time intervals can be selected

- Hourly
- Twice a day or
- Once a day

Note

For security reasons you should select the shortest interval (**Hourly**). Log files are written to secondary storage only, if documents have been written to the secondary storage in the last time interval and thus new entries were created in the reference table.

Copy logs to: If this option is enabled, then the log files are not only stored on the storage system but also copied to the specified directory.

Note

If you are using a storage system without direct file access (such as TSM, Centera, iCAS), we recommend enabling this option to be able to access this information quickly in case of errors. You should create a backup of this directory regularly.

Retention time for logs (in years): Specifies the retention time for the IR files (index-recovery files)

1.3.7. Tab: HTTP / API Interface

d.ecs storage manager provides a web interface for monitoring purposes. This interface can be enabled or disabled. It can be protected with a definable password.

Activate web monitoring: This enables or disables monitoring of the d.ecs storage manager via the web interface.

Note

This parameter must be enabled to allow further settings to be changed in this dialog. In combination with an ecspand connection, this parameter must be enabled.

Port: This defines the port over which the web interface is accessible. By default, the port 6122 is used. If multiple instances of the d.ecs storage manager are running on one server, each instance must be assigned to a different port.

Password: This defines the password to be requested for the web interface.

Confirm password: This defines the password to be requested for the web interface. Repeat here to avoid typing errors.

Note

The password is also used for the authentication via the API-interface (ecspand).

Web monitoring auto refresh (in sec.): This optionally specifies a time interval after which the surface of the web interface is automatically updated. If you do not want any automatic refresh, then enter a value ≤ 0 .

Allow API access: Enable this option to allow the API interface to access the d.ecs storage manager. This parameter must for example be enabled, if d.ecs storage manager is to be used in connection with ecspand.

Note

Authentication via d.ecs identity provider is an additional authentication method used to access web monitoring via d.ecs http gateway. The existing login using the user name "admin" and the specified password is still available and is used when accessing the web monitoring port directly.

Access group: Enter the name (or alternatively the ID) of the d.ecs identity provider group that should be granted access to web monitoring and the API.

Allow API access: Enable this option to allow the API interface to access the d.ecs storage manager. This parameter must for example be enabled, if d.ecs storage manager is to be used in connection with ecspan.

SSL encryption: This specifies, if the data transmission of the web interface should be encrypted by SSL. This also applies to the data transfer when using the d.ecs storage manager API interface. The access via HTTPS is only possible with an external access to the web interface. For local access, HTTP is used. This option is enabled by default when creating a new configuration and requires a certificate and key file to be specified. Furthermore, a password for the key file can be specified.

Certificate: This defines the SSL certificate to be used.

Key: This defines the SSL key file to be used.

Password: This specifies the passphrase used for the creation of the certificate.

Note

The user name for the web interface is defined as admin and cannot be changed.

If the web interface is called from the machine hosting the d.ecs storage manager, then no login is required.

To open the web interface, enter the following URL in the internet browser: `http://<host>:<Port>`

If the SSL encryption is used then you must enter `https://<host>:<Port>`.

`http://localhost:6122`

1.3.8. Tab: Cache

d.ecs storage manager has an internal document cache which enables containers that have been restored to be temporarily stored in a cache for a defined period of time. This has the advantage that the container does not need to be restored repeatedly from the storage system whenever a document in it has to be retrieved.

You can define how long the container should be retained in the cache. The retention time is specified in minutes. The container is deleted as soon as the difference between the current time and the last access exceeds the maximum retention time.

Furthermore, you can specify the maximum time that a document may remain in the jukebox directory. If the creation date of a document is too old then the document is removed from the Jukebox directory and retrieved from the storage system on demand.

Retention period for documents (in minutes): This specifies the time that a document is retained in the Jukebox directory before it is deleted there. The time is specified in minutes.

Activate container cache: This defines if stored containers should be cached.

Retention period for containers (in minutes): This specifies the time for which a container is stored in the container cache. The time is specified in minutes.

1.3.9. Tab: Logging

This tab allows to configure which event/errors/hints etc. are to be logged in the d.3 log file. Furthermore, you can define, if messages should also be sent to the d.ecs monitor or to an SNMP server by SNMP traps. If d.velop metrics analyzer is used, logging of syslog information can optionally be activated.

Note

Please note that the logging information is essential for the troubleshooting.

The d.3 log file can be displayed with the d.3 logview via **Start > All programs > d.velop > d.3 gateway > Log-Viewer**.

Note

The debug logging can also be enabled or disabled using the web interface. This activation/deactivation then only applies to the current session.

1.3.10. Tab: Single Instance Store

The Single Instance Store function makes it possible to save storage space on the long-term memory when archiving a document multiple times in d.3/ecspand, since the document is physically stored only once.

During the check for duplicates, the RipeMD 256 hash procedure is applied, where a 256-bit hash value is created to compare the binary data.

Warning


In the following cases, documents are not considered by the Single Instance Store function:

- The documents are stored together with signatures.
- The target system is Governikus LZA.

Single Instance Store

☒ Single Instance Store aktivieren

When the Single Instance Store function is activated, a message appears that must be accepted. Otherwise, the Single Instance Store function cannot be enabled.



Note

Since the storage systems Centera (object-based), NetApp (block-based), HCP (object-based), Celerra (object-based), Silent Cube (object-based) and TSM/TIAM (object-based) already internally support the single instancing, it is recommended to disable single instancing in d.ecs storage manager and to use the functions of the storage system. This saves you storage space. Especially for block-based Single Instancing.

1.3.11. Tab: Miscellaneous

The tab **Miscellaneous** allows the configuration of various parameters:

Disable disk space check: Disables the check of the free space in the Jukebox-director.

In %: This defines the percentage of memory to be kept free in the Jukebox directory. If the value falls below this threshold, then warning messages are written to the d.3 log and sent to d.ecs monitor. This option is enabled by default. We recommend this option to prevent the Jukebox directory from running full.

Absolute (in MB): Specifies the absolute value of the free space in the Jukebox-directory (secondary storage) in MB. If the value falls below this threshold, then warning messages are written to the d.3 log and sent to d.ecs monitor.

d.log identifier: This defines the name of the d.ecs storage manager under which the messages are sent to the d.log server and the d.ecs monitor. This option must be configured, if more than one d.ecs storage manager instance is running on one server.

Note

From version 2.4 of d.ecs storage manager the identifier **DSM** is used for log entries (previously **D3SM**) in the d.3 log file. You may need to modify d.ecs monitor objects.

Terminate jobs on exit: If this option is enabled, secondary storage/recovery jobs are aborted when shutting down d.ecs storage manager. This allows a faster shutdown of d.ecs storage manager for a cluster failover.

Retention period for logs (in years): Specifies the retention period for d.ecs storage manager logs (synchronization-/health-check logs).

Allow synchronization of signature systems: By enabling this parameter, you can perform a synchronization from or to a signature system. If the parameter is disabled, signature systems are not available for the synchronization.

Warning

With the activation of the synchronization between signature systems you confirm that you have manually corrected the assignment between document and signature.

Else, document with signatures cannot be synchronized correctly.

If you have been using d.3 version 8.0 and d.ecs storage manager 3.0 from the start, you can enable this option without a correction.

1.3.12. Tab: Deletion

The application d.ecs storage manager supports deleting documents from the secondary storage (if the secondary storage permits this) after expiry of the specified retention periods. Deletion takes place, if the document has also been deleted from d.3 (logically and physically) and a respective delete-job has been received from the d.3 system. Furthermore, the privileged deletion of documents is supported on selected systems.

Löschung

☒ Löschen aktivieren

Automatisches Löschen

	Mo:	Di:	Mi:	Do:	Fr:	Sa:	So:	Startzeit:	Laufzeit:
Startzeit 1:	☒	☒	☒	☒	☒	☒	☒	00:00	24
Startzeit 2:	☐	☐	☐	☐	☐	☐	☐	00:00	24
Startzeit 3:	☐	☐	☐	☐	☐	☐	☐	00:00	24

☐ Erlaube privilegiertes Löschen

Protokolltabelle

Maximale Übertragungen pro Durchgang: 500000

☒ Löschung alter Protokolltabellen aktivieren

Lösche Protokolltabelle nach x Jahren: 3

Enable deletion: This enables the deletion procedure.

Automatic deletion: This defines the day of the week and the time when the inventory procedure is to be started. Furthermore, you can define the maximum runtime of the process.

Allow privileged deletion: If this parameter is set, then the d.ecs storage manager accepts jobs for privileged deletion. With these special delete jobs, the documents are removed from the storage system with special privileges (this means, the retention are ignored by the storage system and the deletion is executed). The d.ecs storage manager supports privileged deletion for the systems NetApp, HCP, iCAS, Centera, Governikus LZA and Silent Brick.

Note

This option requires a special license. If you do not have this license, then this option is not available. Please note that the respective storage systems may also require a special license for the privileged deletion. You may also have to enable options on the storage systems or the deletion before expiry of the retention period may only be possible in certain configurations. The privileged deletion is always performed immediately (subject to a respective job by the d.3 system) and is thus independent from the settings applied under Automatic deletion.

Maximum transfers per cycle: Specifies how many old references in the protocol tables are transferred per cycle. A cycle starts after the deletion is complete. As soon as all old references have been transferred to the protocol tables, the process has no function any more. New references are transferred to the protocol table immediately on deletion.

Warning

The transfer of the references to the protocol tables leads to an increased number of transactions thus making the transaction log larger.

Enable deletion of old protocol tables: Enable this parameter so that the protocol tables are automatically deleted from the database after the specified number of years. The protocol tables can be deleted only after at least one year.

Delete protocol table after x years: This defines the number of years after which the protocol tables are to be removed.

Since the inventory procedure reads and modifies a large amount of data from the database, it should be started at times of low utilization of the d.3 system (e.g. over the weekend).

Note

On Silent Cube systems, a document can only be deleted from a share of the head unit. The deletion from a storage unit is not possible. Thus, the used disk space is not released by the deletion.

1.3.13. Tab: Retention process

d.ecs storage manager automatically extends the retention period for documents on the secondary storage whose retention period is about to expire soon. This is only done for documents not marked as deleted. The Retention process can be time-controlled. The runtime can also be defined. For security reasons, the process must run at least 8 hours per week.

Retention Prozess

Maximale Job-Anzahl:

Retention um x Monate erweitern:

Startzeit für den Retention Prozess

	Mo:	Di:	Mi:	Do:	Fr:	Sa:	So:	Startzeit:	Laufzeit:
Startzeit 1:	☐	☐	☐	☐	☐	☒	☒	00:00	4
Startzeit 2:	☐	☐	☐	☐	☐	☐	☐	00:00	0
Startzeit 3:	☐	☐	☐	☐	☐	☐	☐	00:00	0

Maximum number of jobs: Maximum number of documents to be processed (maximum number of retention jobs in the d.ecs storage manager job table).

Extend retention by x months: Number of months by which the retention is to be extended.

Start time for the retention process: This defines the day of the week and the time when the inventory procedure is to be started. Furthermore, you can define the maximum runtime of the process.

Since the retention process reads and modifies a large amount of data from the database, it should be started at times of low utilization of the d.3/ecspan system (e.g. over the weekend).

1.4. d.ecs storage manager modules

The following modules are currently available:

- AZURE = Azure Blob Storage
- CELERRA = EMC² VNX systems with FLR option
- CENTERA = EMC² Centera and EMC² Atmos systems
- CLOUDSTORAGE = d.velop cloud storage
- DATADOMAIN = EMC² Data Domain
- FILELOCK = Grau Data FileLock
- GOVERNIKUSLZA = Governikus LZA/DATA Aeonis signature system
- HCP = Hitachi Content Platform
- ICAS = iTernity iCAS
- ISILON = Dell EMC Isilon
- NAS = NAS (Network Attached Storage) (Only for test and development environments; if necessary for backups)
- NETAPP = Network Appliance NetApp (with SnapLock)
- PRONEXT = Procilon proNEXT signature system
- S3 = General S3 memory
- SILENCUBES = Fast LTA Silent Cubes
- SYSTEMCONNECT = System Connect for 3rd party connectors (read only)
- TSM = Tivoli Storage Manager (with data retention)

Warning

All storage systems (except for simple NAS systems) are equipped with technical protections against the modification and deletion of documents. However, we nevertheless recommend to integrate the storage environment into a backup scenario.

Warning

the d.ecs storage manager only receives one lifetime date per document ID from the d.3 system. This date is the same for all document versions.

The systems configured through d.ecs storage manager allow to reduce this lifetime retroactively and thus to delete documents from the secondary storage before their originally configured lifetime expires.

In this process, all related older document versions of the same document are also deleted that were previously subject to other retention periods.

Note

Please note the vendor information on the existing certificates and the respective vendor documentations.

1.4.1. d.ecs storage manager Celerra

The Celerra module allows d.ecs storage manager to store documents on a VNX system with FLR option. It supports the Filelevel Retention-technology which protect documents from being modified or deleted for a certain period of time.

Warning

In order to ensure audit-proof storage, the documents must be written to a volume protected by the "File Retention"-technology. This volume created by the Filelevel Retention option is addressed by the d.ecs storage manager via a network share.

Furthermore, a new user should be created in the Windows structure. On the side of d.ecs storage manager, this is the only user who requires the right **Full Control** for the VNX area on the VNX system. d.3 process manager, in which the d.ecs storage manager runs, is then started under this user.

If several users have **Full Control** access to the File Retention volume, it could happen that documents are incorrectly written to the VNX system and are not extended when starting the d.ecs storage manager under a different user.

If other d.3 processes run in this d.3 process manager, you must also consider the rights required by these other processes.

Warning

To store the documents with the d.ecs storage manager on an FLR protected file system, it is essential that the default, minimum, and maximum retention period for the FLR file system are set equivalent to the retention periods in the leading system (d.3/ecspand).

Configuration

The following parameters are configurable on the main tab of the system:

Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Upper storage limit: Specifies how much storage space may be used for storing data.

Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.

Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.

Force CAS creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. (During the retrieval it is checked, if the data in a CAS container are unchanged before a document is passed on to the d.3 server / d.3 gateway / API). Should be enabled in connection with ecspand.

Validate data: If this option is enabled, a file comparison based on the **RipeMD256** mechanism is performed after having copied a file to the Celerra volume. This ensures that the file is completely transferred.

Unity mode: This switch must be enabled if the used storage system is a Unity system.

The following parameters can be configured on the **HealthCheck** tab:

File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration

Alias: Alias name for the volume; if a name is specified here, it is displayed instead of the volume ID in the Web interface for a better overview.

Storage path: Specify a path here where the documents are stored

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Upper storage limit (in MB): Here you can define in MB how much storage space may be used for the storage of data.

Lower storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Minimum/maximum retention period (in months) Specifies the range in which the retention time passed by the delivering process may be located for a document. These values have to be set and must be equivalent to the settings of the Celerra/VNX volume.

Automatically correct retention period: If a retention-period outside the specified range is passed by a delivering process, then the time is automatically set to the minimum or maximum, if this parameter is enabled. Else the job is blocked and the document is not written to the secondary storage.

Log in with the following user: User to be used for the login to the volume.


Password/Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.

- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.

Index recovery

The d.ecs storage manager Celerra module allows to recover the internal document index after a loss. For this effect, it browses the specified target directory for documents and thus rebuilds the index. This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (see also [Tab: Database logging](#)).



Celerra directory: Directory on the VNX system to be searched for documents.

Index directory: Directory where the index of the documents for the VNX system is to be stored.

Preparation of the VNX volume

The volume on the VNX system must be defines as follows for d.ecs storage manager to operate with it properly:

File System		Quota Settings		Deduplication Settings		FLR Settings	
File System Name:		vnx81fs1					
File-level Retention:		Enterprise with protected files					
		FLR Clock Time:		Wednesday, December 10, 2014 10:01:19 AM EST			
		Last Currently Locked File Will Expire On:		25 years, 10 days (Thursday, December 15, 2039 8:53:18 AM EST)			
Minimum Retention Period:		☐ Unlimited ☒ Limited 0 Days					
Default Retention Period:		☐ Unlimited ☒ Limited 1 Days					
Maximum Retention Period:		☒ Unlimited ☐ Limited 1 Days					
Auto-lock new files:		☐					
Auto-lock policy interval:		1 Hours					
Auto-delete files when retention ends:		☐					
Retention Date Range:		☒ Default (2003-2070) ☐ Custom Earliest valid retention year: 2003 Latest valid retention year: 2070					
		OK Apply Cancel Help					

1.4.2. d.ecs storage manager Centera

Using the Centera module, the d.ecs storage manager can store documents on an EMC² Centera, on an EMC² Atmos system or an EMC² Elestic Cloud System.

The correct communication between the d.ecs storage manager and the Centera requires a Centera with CenteraStar version ≥ 3.1 .

Configuration

The screenshot shows a configuration window titled 'Konfiguration (CENTERA_1)'. It has several input fields and buttons. The 'Nodes' field contains '192.168.1.10'. The 'Filespace' field contains 'D3A'. The 'Verbindungen' field is a spinner set to '4'. The 'Aufbewahrung' section has radio buttons for 'Standard', 'Nie', 'Eventbasiert' (selected), 'Immer', and 'Nach x Sekunden'. Below 'Eventbasiert' is a text field for 'Aufbewahrungszeit (in Sek.):'. The 'Benutzer' and 'Passwort' fields are empty. The 'PEA-Datei' field contains 'c:\d3storagemanager\d3a.pea'. On the right, there are buttons for 'Node hinzufügen', 'Node bearbeiten', and 'Node löschen'. At the bottom are 'Info', 'OK', and 'Abbruch' buttons.

Nodes: Enter the IP addresses (the individual nodes) of the Centera here. Only the nodes of a Centera should be specified here.

Filespace: Specify a file space name for the documents here, e.g. **d3store**.

Warning

If you operate several d.3 servers (with the same Repository IDs), which are saving their data on one Centera, then you must make sure that the filespace for each d.3 server has a different name. Example: for server 1: **d3store1**, for server 2 **d3store2**. Otherwise this will result in problems later with the recovery of the index structure!

Connections: This defines the maximum number of simultaneous connections to the Centera.

Retention: This defines when a document clip may be deleted (Clip = Container where the Centera stores the data).

- **Default:** Standard setting of the Centera. (CE, GE)
- **Never:** Documents cannot be deleted at all except "Privileged Delete" being assigned to a Centera user. (GE)
- **Eventbased:** Documents are protected using the Advanced Retention Management of the Centera and can only be deleted after a Retention Event was triggered. (CE and GE with license)
- **Always:** Documents can always be deleted. (CE, GE)
- **After x seconds:** Documents can be deleted after the specified time (in seconds) or after the time passed by d.3. (CE, GE)

Note

This setting is only considered by d.ecs storage manager, if the delivering process (e.g. d.3 or ecspand) passes no retention period.

Using event-based deletion:

The d.3 system works with variable retention times. Since the Centera does not support this function, there are two alternatives regarding the event-based deletion:

- **Government Edition (GE)**

d.ecs storage manager stores the document on the Centera and protects it with the retention setting eternal. Only, if the document retention period in d.3 has expired, the document is removed from the Centera by d.ecs storage manager. The d.ecs storage manager is the only process allowed to access the document pool. The deletion is performed using the Privileged Deletion function on the Centera. This permission must be assigned to the user under which d.ecs storage manager is accessing the Centera.

- **Compliance Edition (CE) / Governance Edition (GE)**

d.ecs storage manager stores the document on the Centera and protects it using the Advanced Retention Management of the Centera. In the Compliance Edition of the Centera, the license for Event-based Retention is required since the function of the Privileged Deletion is not available here.

Note

If this option is enabled, the d.ecs storage manager ignores the retention period passed by the delivering process (e.g. d.3 or ecspond). The object is in this case for Event and not a fixed Retention bound.

User: If you are using Centera profiles, then you can specify the respective user name here.

Password: If you are using Centera profiles, then you can specify password here.

Note

When using the **PEA file**, this always has priority over the user name and password specified here.

PEA file: If you created a PEA file (containing the login information; created via the Centera; see Centera manual) for the login, then specify the file path here.

Warning

For a trouble-free operation, the following rights must be configured for a profile: **read, write, exists, delete, query**. If the **Privileged Delete** function is to be used, this right must also be enabled for this profile.

Index recovery

The d.ecs storage manager Centera module allows to search a Centera for stored documents thus recovering the internal document index after a loss of the index. This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (also see [Tab: Database logging](#)).

In order to search for documents on the pool "B", the Pool ID "B" must be specified. The index directory and the filesystem are automatically assigned. Sometimes, an IP address, the user name and the password must be entered. Then the recovery process is started with a click on the button. Depending on the number of documents located on the document pool, the search can take several minutes or even hours. All clips are searched for documents.

1.4.3. d.ecs storage manager cloud storage

The d.ecs storage manager cloud storage module allows d.ecs storage manager to store documents on the d.velop cloud storage.

Optionally, the documents can be protected from unauthorized access using AES256 encryption.

The module supports setting retention periods stored in the d.velop cloud storage and considered when for the deletion of documents.

Configuration

Access data

On the configuration page **Access data** you can enter the access data for your d.velop cloud storage.

Access key: Enter the access key for your d.velop cloud storage here.

Secret access key: Enter the secret access key for your d.velop cloud storage here.

Bucket: Enter the bucket for your d.velop cloud storage here.

Bucket region: Select here the region where your bucket is stored. It is also possible to specify the region as free text if the desired region is not available in the list.

- Confirm your access data by clicking on **Test connection**.

Warning

During configuration, make sure that two different instances of d.ecs storage manager never use the same bucket.

If no encryption is configured, a warning is displayed in the lower part of the dialog. Disable encryption via **Disable encryption** only if you are sure that you want to store the documents unencrypted in the d.velop cloud storage.

Encryption

The configuration page **Encryption** allows you to enable the encryption of the data objects so that objects are stored on the d.velop cloud storage in encrypted form.

Depending on the extent to which the encryption is already configured and initialized on your system, different configuration views may be displayed when you switch to the configuration page.

Initializing and activating the encryption function

To enable the encryption of the data objects, you must first generate a security ID.

The movement of the mouse is also considered during the generation thus providing you with an even more random key.

Konfiguration (CLOUDSTORAGE01)

Zugangsdaten | **Verschlüsselung**

Um die Verschlüsselung für den Bucket zu aktivieren generieren Sie einen Zufallsschlüssel (bewegen Sie während der Schlüsselgenerierung die Maus um einen noch zufälligeren Schlüssel zu erzeugen) und klicken Sie auf "Generieren". Drucken Sie den erzeugten Schlüssel aus und verwahren Sie ihn sicher. Geben Sie dann den auf dem Ausdruck abgebildeten Bestätigungscode ein um die Verschlüsselung zu aktivieren.

Bei Verlust des Sicherheitsschlüssels können die Dokumente nicht mehr entschlüsselt werden

Zufallsschlüssel generieren

Sicherheitsschlüssel

Bestätigung

Code:

- Click **Start** to begin with the generation.
- Click on the button again after a few seconds to stop the generation.
- Then, you must print the generated security ID. The printed copy contains a confirmation code to be entered in the section **Confirmation**.

Warning


The printed copy must be kept in a safe place. If this and thus the security ID is lost, it is not possible to restore the master key, if this should be lost.

When exchanging the software components, you may have to enter the security ID again.

Encryption can be activated only once the security ID has been entered correctly.

It may take some time to activate encryption (approximately 15 to 20 seconds).


Having successfully enabled the encryption, the following message is displayed:



The encryption was enabled now and the d.ecs storage manager can store the documents and data objects in the d.velop cloud in encrypted form.

Activating encryption when the encryption function has already been initialized

If the encryption has already been initialized, it can simply be enabled for another bucket.



- Click **Enable** to activate it.

It may take some time to enable the encryption (approximately 2 to 5 seconds).

Having successfully enabled the encryption, the following message is displayed:




The encryption was enabled now and the d.ecs storage manager can store the documents and data objects in the d.velop cloud in encrypted form.

Renewed initialization of the encryption function after the master key was lost

If the initialization of the encryption function is lost, e.g. as a result of a clean installation of the operating system, a server migration or changed hardware, the encryption function must be enabled again.

This is done using the security key on the printed copy generated at the time of initializing the encryption.

The following dialog is displayed to re-enable the encryption function:



The screenshot shows a Windows-style dialog box titled 'Konfiguration (CLOUDSTORAGE01)'. It has two tabs: 'Zugangsdaten' and 'Verschlüsselung'. The 'Verschlüsselung' tab is active. The text inside the dialog reads: 'Die Verschlüsselungsfunktion wurde für das konfigurierte Bucket bereits aktiviert, allerdings ist kein gültiger Schlüssel für die Verschlüsselung auf dem Rechner hinterlegt. Um die Verschlüsselungsfunktion erneut zu aktivieren geben Sie den Sicherheitsschlüssel ein, den Sie bei der erstmaligen Aktivierung der Verschlüsselung für dieses Bucket ausgedruckt haben.' Below this text is a section titled 'Verschlüsselung wiederherstellen' containing a label 'Sicherheitsschlüssel:' and a text input field with a dashed line indicating where to enter the key. At the bottom left is a button labeled 'Reaktivieren'. At the bottom right are two buttons: 'Speichern' and 'Abbrechen'.

- Enter the security key that appears on the printed copy.
- Click **Re-enable**.

It may take some time to activate encryption (approximately 15 to 20 seconds).


Having successfully enabled the encryption, the following message is displayed:



The encryption was enabled now and the d.ecs storage manager can store the documents and data objects in the d.velop cloud in encrypted form.

Proxy settings

It is possible to use a proxy server for communication with d.velop cloud storage. Enter the login credentials for the proxy server available in your network in the corresponding fields. The username and password are optional fields and only need to be filled if you are using a proxy server with user authentication.



Zugangsdaten	Verschlüsselung	Proxy-Einstellungen
☒ Proxy-Server verwenden		
Proxy-Server:	10.95.1.55	
Proxy-Port:	3128	
Benutzername:	proxyuser	
Passwort:	*****	

1.4.4. d.ecs storage manager Data Domain

With the Data Domain module, the d.ecs storage manager can store documents in an EMC² Data Domain system. It supports file retention technology, which protects documents from being modified or overwritten for a defined period of time.

Warning

In order to ensure audit-compliant storage, the documents must be written to a volume protected by file retention technology.

Furthermore, a new user should be created in the Windows structure. From the perspective of d.ecs storage manager, this is the only user who requires **Full Control** rights to the storage area of the Data Domain system. d.3 process manager, in which the d.ecs storage manager runs, is then started under this user.

If several users have **Full Control** access to the Data Domain volume, documents may end up being incorrectly written to the volume when d.ecs storage manager is started under a different user, and it may not be subsequently possible to extend the documents due to missing rights.


If other d.3 processes run in this d.3 process manager, you must also consider the rights required by these other processes.

To use the "File Retention" function, you must activate a separate license in the Data Domain system. This license must be purchased separately. If necessary, consult the documentation for the Data Domain system or contact EMC² support for the correct settings of the Data Domain system.

Warning

To store the documents with the d.ecs storage manager on a Data Domain file system, it is essential that the default, minimum and maximum retention periods for the Data Domain file system are set according to the retention periods used in the leading system (d.3/ecspand).

Configuration



Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Upper storage limit: Specifies how much storage space may be used for storing data.


Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.

Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.


Force CAS creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. (During the retrieval it is checked, if the data in a CAS container are unchanged before a document is passed on to the d.3 server/d.3 gateway/API). Should be enabled in connection with ecspand.

Validate data: If this option is enabled, then a file comparison based on the RipeMD256 hash procedure after copying a file to the Data Domain volume to make sure that the file was transferred completely.



File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration



Alias: Alias name for the volume. If a name is specified here then this is used instead of the volume-ID in the web interface. This is used for a better overview.

Storage path: Specify a path here where the documents are stored

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Minimum/maximum retention period (in months) Specifies the range in which the retention time passed by the delivering process may be located for a document. These values have to be set and must be equivalent to the settings of the DataDomain system.

Automatically correct retention period: If a retention-period outside the specified range is passed by a delivering process, the time is automatically set to the minimum or maximum, if this parameter is enabled. Else the job is blocked and the document is not written to the secondary storage.

Log in with the following user: If this parameter is enabled, the d.ecs storage manager automatically tries to log in to the respective volume. This parameter only has to be enabled, if the d.ecs storage manager is executed under a user account which does not have access to the volume.

User: User to be used for the login to the volume.


Password/Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.

- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.

Index recovery

The d.ecs storage manager Data Domain module allows to recover the internal document index after a loss. For this effect, it browses the specified target directory for documents and thus rebuilds the index. This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (also see [Tab: Database logging](#)).



Storage-directory: Directory on the Data Domain system to be searched for documents.

Index directory: Directory where the index of the documents for the Data Domain system is to be stored.

1.4.5. d.ecs storage manager FileLock

The FileLock module allows d.ecs storage manager to store documents on a partition protected with GrauData FileLock.


Note

FileLock must be installed in the version 2.3.

Warning

To store the documents with d.ecs storage manager on a FileLock volume, it is essential that the default, minimum, and maximum retention period for the FileLock volume are set equivalent to the retention periods in the leading system (d.3/ecspand).

Configuration



Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Upper storage limit: Specifies how much storage space may be used for storing data.

Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.


Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.

CAS-creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. (During the


retrieval it is checked, if the data in a CAS container are unchanged before a document is passed on to the d.3 server / d.3 gateway / API). Should be enabled in connection with ecspond.

Validate data: If this option is enabled, then a file comparison based on the "RipeMD256" mechanism is performed after having copied a file to the FileLock volume. This ensures that the file is completely transferred.



File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration



Alias: Alias name for the volume. If a name is specified here then this is used instead of the volume ID in the web interface. This is used for a better overview.

Storage path: Specify a path here where the documents are stored

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Upper storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Lower storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Minimum/maximum retention period (in months) This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Automatically correct retention period: If a retention-period outside the specified range is passed by a delivering process, then the time is automatically set to the minimum or maximum, if this parameter is enabled. Else the job is blocked and the document is not written to the secondary storage.

Log in with the following user: If this parameter is enabled, the d.ecs storage manager automatically tries to log in to the respective volume. This parameter only has to be enabled, if the d.ecs storage manager is executed under a user account which does not have access to the volume.

User: User to be used for the login to the volume.

Password/Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.

- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.


Configuration of the FileLock volume

The configuration of the FileLock volume has various requirements which will be described below.

Set the partition to WORM mode with the setting **1st Directory Level**.


- Set the folder to be enabled for FileLock to the **SnapLock Functionality Mode**.
- It is recommended that the values **Minimum Retention** and **Default Retention** are set to 0 and Maximum Retention to the highest retention period defined in d.3/ecspand. This avoids possible issues when setting the retention period through the d.ecs storage manager. Else, the FileLock overwrites the retention periods set by the d.ecs storage manager automatically with the specified defaults.
- The parameter **Auto Commit** must not be enabled, as the FileLock then automatically locks documents and under certain conditions this interferes with the procedures of the d.ecs storage manager.

A configuration could thus look as follows. In this example the maximum retention period is set to 70 years.



Index recovery

The d.ecs storage manager FileLock module allows to recover the internal document index after a loss. For this effect, it browses the specified target directory for documents and thus rebuilds the index. This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (also see [Tab: Database logging](#)).



FileLock directory: Directory in the FileLock system to be searched for documents.

Index directory: Directory where the index should store the documents for the FileLock system.

1.4.6. d.ecs storage manager HCP

The HCP module allows the d.ecs storage manager to store documents on an HCP system (Hitachi Content Platform). The HCP version 4 to 9 are supported.

Configuration

Namespaces: Specify the namespaces of the HCP systems to be used here. A namespace is composed of [namespace].[tenant].[HCP-hostname]

Warning

Only the first specified namespace is written into. All other namespaces are used read-only and serve for automatic error-handling, if the first namespace is not available due to a system failure. If you specify more than one namespace, all namespaces must exist on HCP systems which are mirrored with each other.

HCP Port: Specify the port for the HTTP-access to the HCP-system here (default value: 80). If SSL is enabled, then this must be changed to the port for the SSL connection (default: 443).

User: Specify the user name for the tenant to be used here.

Password: Specify the password for the tenant to be used here.

Use SSL Encryption: This specifies, if the SSL encryption is to be used for the connection to HCAP.

Proxy settings

Use proxy server: Activates the use of a proxy server for communication with the HCP system

Proxy server: IP address or host name of the proxy server

Proxy port: Port of the proxy server

User name: User name for the authentication on the proxy server (optional)

Password: Password for the authentication on the proxy server (optional)

Index recovery

Follow the steps below to load the index-recovery files from the HCP:

- Open the Tenant of the HCP in a web-browser. The directory /rest/[POOLID]/D3SMLOGS contains the index-recovery-files sorted by months which can be reimported into the database as described in the chapter Database-Recovery.

1.4.7. d.ecs storage manager iCAS

The d.ecs storage manager iCAS module allows d.ecs storage manager to store documents on an iCAS system by iTernity.

Warning

For a proper processing, an iTernity iCAS system of version 3.7 SP4 is required.

The user to be used to access the iCAS system, should have the right **ADVIRW** to be able to execute all API calls executed by d.ecs storage manager without any problems.

Configuration

Konfiguration (ICAS)

Verbindung

Host: 10.100.153.37

Port: 443

Benutzer: user

Passwort: xxxxxxxxxxxx

Passwort wiederholen: xxxxxxxxxxxx

Verbindung prüfen

Modus

Verschlüsselung: None

Kompression: Store

☒ SSL verwenden

☒ CSC überprüfen

TLSv12

Info OK Abbruch

Host: Hostname or IP of the iTernity-system.

Port: Port of the iTernity-system.

User: Users eligible to sign in to the iTernity-system who has the respective rights for the secondary storage or reading of the documents.

Password: Password of the user.

Compression: Optionally, select the type of compression for the documents here.

Encryption: Optionally, select the type of encryption for the documents here.

Use SSL: This allows you to enable the encryption for the data transmission between d.ecs storage manager and the secondary storage system. This also changes the default port to 443. Please check, if your storage system uses the same port. Furthermore, the SSL/TLS version to be used can be defined. This depends on the IIS (Internet Information Service) version of the iCAS system used.

Check CSC: If this parameter is enabled, then it is checked on uploading a file to the iTernity system, if the file exists on the iTernity system.

1.4.8. d.ecs storage manager NAS

The NAS module allows d.ecs storage manager to store documents on a NAS system (NAS = Network Attached Storage).


Warning

Technically speaking, the NAS module always stores documents / content on an unprotected area which can be manipulated and which is thus not audit-proof.

When using the NAS-variant for the audit-proof storage, then respective measures and protective mechanisms must be implemented to ensure and warrant the completeness, unchangeability and genuine storage on the storage media. These measures and protective mechanisms must be documented in a process manual.

In contrast to the NAS-solution, the other systems supported by d.ecs storage manager (([Centera](#), [Data Domain](#), [NetApp](#), [TSM / Jukebox](#) , [Silent Cubes](#) , [FileLock](#) , [iCAS](#) , [HCP](#)) ensures audit compliance from a technical point of view when used correctly..

Configuration



Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Upper storage limit: Specifies how much storage space may be used for storing data.

Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.

Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.

Force CAS creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. During the recovery it is checked, if the data in a CAS container are unchanged. This should only be enabled in connection with ecspan or if the consistency of the data on the NAS should be verifiable.

Validate data: If this option is enabled, a file comparison based on the **RipeMD256** mechanism is performed after having copied a file to the NAS-volume. This ensures that the file is completely transferred.

File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration

Alias: Alias name for the volume. If a name is specified here then this is used instead of the volume ID in the web interface. This is used for a better overview.

Storage path: Specify a path here where the documents are stored

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Upper storage limit (in MB): Here you can define in MB how much storage space may be used for the storage of data.

Lower storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Log in with the following user: If this parameter is enabled, the d.ecs storage manager automatically tries to log in to the respective volume. This parameter only has to be enabled, if the d.ecs storage manager is executed under a user account which does not have access to the volume.

User: User to be used for the login to the volume.

Password/Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.


- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.

Index recovery

The d.ecs storage manager NAS module allows to recover the internal document index after a loss.

For this effect, it browses the specified target directory for documents and thus rebuilds the index.

This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (also see [Tab: Database logging](#)).



NAS directory: Directory on the NAS system to be searched for documents.

Index Directory: Directory where the index of the documents for the NAS system is to be stored.

1.4.9. d.ecs storage manager NetApp

The NetApp module allows the d.ecs storage manager to store documents on a NetApp system. It supports the NetApp SnapLock technology protecting the documents from deletion or overwriting for a certain amount of time.

The NetApp module requires an installation of the Visual C++ 2013 Redistributable Package in the 32-bit version. This is installed during the installation of d.ecs storage manager.

Since the d.ecs storage manager access the storage area of the NetApp via the CIFS (Windows-network protocol), is necessary to activate the CIFS-protocol via a special license on the NetApp side.

Warning

In order to ensure audit-proof storage, the data must be written on a volume protected with the SnapLock technology.

Furthermore, a new user should be created in the Microsoft Windows structure. From the point of view of d.ecs storage manager, this user is the only one who needs the right **Full Control** to the SnapLock area on the Netapp. d.3 process manager, in which the d.ecs storage manager runs, is then started under this user.

If several users have **Full Control** access to the SnapLock volume, it could happen that documents are incorrectly written to the NetApp and are not extended when starting the d.ecs storage manager under a different user.

If other d.3 processes run in this d.3 process manager, you must also consider the rights required by these other processes.

A volume on the NetApp must not be set to read-only mode, as otherwise some processes of the d.ecs storage manager cannot work.

Warning


To store the documents with the d.ecs storage manager on a SnapLock volume, it is essential that the default, minimum, and maximum retention period for the SnapLock volume are set equivalent to the retention periods in the leading system (d.3/ecspand). Furthermore, the **Autocommit** option must not be activated for a volume, otherwise problems may occur when processing jobs.

Warning

The NetApp allows with the function privileged delete to remove documents from the NetApp before the configure lifetime expires. For this effect, the function must be licensed and enabled on the NetApp. Each such delete action is logged by the NetApp. The right **Privileged delete** must be assigned to the user in the **Advanced properties**.

Please note the respective information on the function and the vendor certification on audit-proof storage.

Configuration



Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Volume: Volume path on the NetApp.

Upper storage limit: Specifies how much storage space may be used for storing data.

Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.

Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.

Autosize: Automatically resize the volume.

Force CAS creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. (During the retrieval it is checked, if the data in a CAS container are unchanged before a document is passed on to the d.3 server/d.3 gateway/API). This should only be enabled in connection with ecspan.

Validate data: If this option is enabled, a file comparison based on the **RipeMD256** mechanism is performed after having copied a file to the NetApp-volume. This ensures that the file is completely transferred.

File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration

Alias: Alias name for the volume. If a name is specified here then this is used instead of the volume ID in the web interface. This is used for a better overview.

Storage path: Specify a path here where the documents are stored

Volume: Specify the volume-identifier of the NetApp-volume here. The directory must be located on the same directory level as the directory specified under Storage path. The volume-path always starts with **/vol/** followed by the volume name as configured in the NetApp. The volume name can be found in FilerView for NetApp. The volume names are provided under the option **Volumes > Manage in the FilerView**. Note that the volume names in the NetApp are case-sensitive.

Automatically resize the volume: If the used volume is automatically extended by NetApp, this parameter must be enabled. Enable this to remove the storage limits.

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Upper storage limit (in MB): Here you can define in MB how much storage space may be used for the storage of data.

Lower storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Minimum/maximum retention period (in months) Specifies the range in which the retention time passed by the delivering process may be located for a document. These values have to be set and must be equivalent to the settings of the NetApp volume. With a click on **Load times** it is attempted to retrieve the settings from the NetApp.

Automatically correct retention period: If a retention-period outside the specified range is passed by a delivering process, then the time is automatically set to the minimum or maximum, if this parameter is enabled. Else the job is blocked and the document is not written to the secondary storage.

Log in with the following user: If this parameter is enabled, the d.ecs storage manager automatically tries to log in to the respective volume. This parameter only has to be enabled, if the d.ecs storage manager is executed under a user account which does not have access to the volume.

User: User to be used for the login to the volume.

Password/Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.

- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.

Connection

The NetApp module checks the settings of the SnapLock volume on the NetApp when starting d.ecs storage manager. To be able to determine this information, the module requires a user defined on the NetApp with the appropriate permissions for access.

API mode: Defines which API mode should be used. From OnTAP version 9.10.1, select **REST** as the API mode. In this mode, the setting of the retention periods has been adapted to the changes in the OnTAP version. Note that the specified users have access to the REST interface.

Filer IP: Host name or IP of the NetApp.

Note

Enter the IP of the management interface of the corresponding storage VM here.

User: User who is allowed to log on to the NetApp and has the appropriate rights to query the volume information.

Password/Confirm password: Password of the user.

Use HTTPS: If HTTPS is to be used for the connection, enable this switch. This requires that HTTPS is enabled on the NetApp. **Privileged Delete** always requires an HTTPS connection.

User (Privileged Delete): User who has the right to delete files privileged.

Password/Repeat password (Privileged Delete): Password of the user.

Note

The **Privileged Delete** options will be visible only if you have purchased the **Privileged Delete** license.

Create the required users on the NetApp (console)

To create a user with the necessary rights in the NetApp, open the console and follow these steps.

API-user for querying the volume-information

OnTAP Version >= version 9

- Create a user for the application **ontapi**, specify a password and assign the role **vsadmin-readonly**.

OnTAP version >= version 9

- Create the role **r_decsapi** with the specific rights **login-http-admin**, **api-volume-list-info** and **api-volume-options-list-info** for the API access on the NetApp:

```
useradmin role add r_decsapi -c "d.ecs storage manager API access role" -a
login-http-admin,api-volume-list-info,api-volume-options-list-info
```

- Create a local group **g_decsapi** on the NetApp and assign the role **r_decsapi**:

```
useradmin group add g_decsapi -c "group for d.ecs storage manager API
access" -r r_decsapi
```

- Create a local user **decsapiaccess** on the NetApp and assign the role **g_decsapi**:

Note

When creating the user **decsapiaccess** you must define a password. The password should be configured not to expire.

```
useradmin user add decsapiaccess -c "d.ecs storage manager API access user"
-g g_decsapi
```

Check the user settings:

```
useradmin user list decsapiaccess
```

The following should be displayed:

```
-----
Name: decsapiaccess
Info: d.ecs storage manager API access user
Rid: 131079
Groups: g_decsapi
Full Name:
Allowed Capabilities: login-http-admin,api-volume-list-info,api-volume-
options-list-info
Password min/max age in days: 0/4294967295
Status: enabled
-----
```

API-user for privileged deletion

OnTAP Version >= version 9

- Create a user for the application **ontapi**, specify a password and assign the role **vsadmin-snaplock**.

OnTAP version >= version 9

```
useradmin user add decsprivdel -c "d.ecs storage manager privdel" -g
"Compliance Administrators"
```

Check the user settings:

```
useradmin user list decsprivdel
```

The following should be displayed:

```
-----
Name: decsprivdel
Info: d.ecs storage manager privdel
Rid: 131080
Groups: Compliance Administrators
Full Name:
Allowed Capabilities: cli-cifs*,cli-exportfs*,cli-nfs*,cli-useradmin*,api-
cifs-*,api-nfs-*,login-telnet,login-http-admin,login-rsh,login-ssh,api-
system-api-*,cli-snaplock*,api-snaplock-*,api-file-*,compliance-*
Password min/max age in days: 0/4294967295
Status: enabled
-----
```

Create REST API users for the Storage VM

With the release of version 3.9.0, d.ecs storage manager supports the new REST API of NetApp systems.

To use the REST API, users with different roles are required.

Creating users for REST API - This is how it works

1. Create a user with the **vsadmin** role in the appropriate storage VM on the NetApp.
2. Assign the login method **HTTP** with the authentication method **Password** to the user.

Optional for Privileged Delete:

To use the **Privileged Delete** function, create a separate user.

Creating users for Privileged deletion - This is how it works

1. Create a user with the **vsadmin-snaplock** role in the appropriate storage VM on the NetApp.
2. Assign the login method **HTTP** with the authentication method **Password** to the user.


Usage with OnTAP

If the d.ecs storage manager is operated with the OnTAP version 8 and the function **Privileged Delete** is used, the option **tls.enable** on the NetApp must be set to "on". This can be done via the command line of the NetApp. The command looks like this:

```
options tls.enable on
```

Index recovery

The d.ecs storage manager NetApp module allows to recover the internal document index after a loss. For this effect, it browses the specified target directory for documents and thus rebuilds the index. This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (also see [Tab Database logging](#)).



NetApp Directory: Directory on the NetApp to be searched for documents.

Index Directory: Directory where the index of the documents for the NetApp system is to be stored.

1.4.10. d.ecs storage manager Silent Cubes


The d.ecs storage manager Silent Cube module allows the d.ecs storage manager to store documents on a Fast LTA Silent Cube or Silent Brick system.

The Compliance-option must be enabled for the Silent Cubes/Silent Brick system. Moreover, the user using d.ecs storage manager to access the share must have the right **Full Control**.

Warning

To store the documents with the d.ecs storage manager on a Silent Cube/Silent Brick volume, it is essential that the default and minimum retention period for the Silent Cubes volume is set equivalent to the retention periods in the leading system (d.3/ecspand).

Configuration



Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Upper storage limit: Percentage specification of how much storage space may be used for storing data.

Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.

Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.


Enforce CAS-creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. (During the retrieval it is checked, if the data in a CAS container are unchanged before a document is passed on to the d.3 server/d.3 gateway/API). Should be enabled in connection with ecspan.

Validate data: If this option is enabled, a file comparison based on the **RipeMD256** mechanism is performed after having copied a file to the Silent Cubes volume. This ensures that the file was completely transferred.

Wait for SAFE-file: If this parameter is enabled, then the d.ecs storage manager only reports an OK to the d.3 server/ecspan, if the document was completely processed by the Silent Cube/Silent Brick system.

Warning

For this procedure, the generation of SAFE files for the volume must be enabled on the Silent Cube/Silent Brick system! The format of SAFE-files must be configured to XML for this effect.



File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration

SILENTCUBES-1

Volume-Konfiguration | Privilegiertes Löschen

Alias: SILENTCUBES-1

Pfad:

☐ Automatische Vergrößerung des Volumes

Speicherüberprüfung

☒ Obere Speichergrenze (in %) 77

☐ Obere Speichergrenze (in MB) 0

☐ Untere Speichergrenze (in %) 33

☐ Untere Speichergrenze (in MB) 0

☒ Wiederbeschreibbar

Retention Zeitspanne

Minimale Retention-Zeit (in Monaten) 6

Maximale Retention-Zeit (in Monaten) 600

☐ Retention-Zeit automatisch korrigieren

☐ Mit folgendem Benutzer anmelden

UNC Benutzer

Benutzer:

Passwort:

Passwort wiederholen:

Verbindung prüfen

Status: Wartend

Volumenstatus zurücksetzen

OK Abbrechen

Alias: Alias name for the volume. If a name is specified here then this is used instead of the volume-ID in the web interface. This is used for a better overview.

Storage path: Specify a path here where the documents are stored

Automatically resize the volume: If the used volume is automatically extended by Silent Cubes/Silent Brick, this parameter must be enabled. Enable this to remove the storage limits.

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Upper storage limit (in MB): Here you can define in MB how much storage space may be used for the storage of data.

Lower storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Minimum/maximum retention period (in months) Specifies the range in which the retention time passed by the delivering process may be located for a document. These values have to be set and must be equivalent to the settings of the Silent Cubes volume.

Automatically correct retention period: If a retention-period outside the specified range is passed by a delivering process, then the time is automatically set to the minimum or maximum, if this parameter is enabled. Else the job is blocked and the document is not written to the secondary storage.

Log in with the following user: If this parameter is enabled, the d.ecs storage manager automatically tries to log in to the respective volume. This parameter only has to be enabled, if the d.ecs storage manager is executed under a user account which does not have access to the volume.

User: User to be used for the login to the volume.

Password/Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.

- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.

Privileged deletion

The FAST LTA Silent Brick systems with software version 2.22 provide the function **Privileged Deletion**. If this function is to be used, it must be activated in the configuration of a volume. The API access must also be configured for the Silent Brick. This is possible in the volume configuration on the tab **Privileged deletion**.

SILENTCUBES-1

Volume-Konfiguration Privilegiertes Löschen

☒ Privilegiertes Löschen aktivieren

Silent Bricks-API-Zugriff

Host:

Port:

Benutzer:

Passwort:

Volume-Name:

Pfad-Präfix:

Verbindung prüfen

OK Abbrechen

Enable Privileged deletion: Enables the function **Privileged deletion** and provides a configuration area for the Silent Brick API.

Host: The hostname or IP of the Silent Brick.

Port: Port where the API can be reached. The default is here 443 (HTTPS).

User name: User name of the on-premises user of the Silent Brick System configured for privileged deletion on the Silent Brick System.

Password: Password of the user.

Volume name: Name of the volume on which the privileged deletion is to be executed.

Path prefix: Prefix for the directory structure below the share of the volume used. This specification only has to be made if the d.ecs storage manager does not write directly to the root directory of the volume share.

Example:


Silent Bricks volume share: \\silentbricks\\volume1

Target path within this share: \\silentbricks\\volume1\\ArchivA\\

Path prefix to configure: ArchiveA

Index recovery

The d.ecs storage manager Silent Cube module allows to recover the internal document index after a loss. For this effect, it browses the specified target directory for documents and thus rebuilds the index. This type of index recovery should only be used, if no index-recovery-files (*.IR) exist to restore the index (also see [Tab Database logging](#)).



Silent Cubes directory: Directory on the Silent Cube/Silent Brick system to be searched for documents.

Index directory: Directory where the index of the documents for the Silent Cubes/Silent Brick-System system is to be stored.

Preparation of the Silent Cube/Silent Brick volume

Configuration Silent Cube

When creating the volume on the Silent Cube system, you must comply with the following preconditions to ensure the proper collaboration of the d.ecs storage manager with the Silent Cube system:

- The volume-type must be configured as WORM with lifetime (RWRO triggered).
- The lifetime should be 0 years and 0 days.

Note

The d.ecs storage manager passes the lifetime for every document to the Silent Cube system.

Bearbeiten eines Volumes

Name: vol_d3p Quota [MB]: 3096576 Autoinkrement [MB]: 102400 Schwelle [%]: 90

Volume-Typ: WORM mit Lebensdauer (RWRO getriggert) Jahre: 0 Tage: 0

Archivierung Safe-Dateien Sonstiges

☒ Wahlfreier Zugriff innerhalb von Dateien

☐ Daten komprimieren

☒ Dateien nach dem Archivieren aus dem Cache löschen

☒ Benachrichtigung wenn noch nicht archivierte Dateien älter sind als:

1 Tage 0 Stunden 0 Minuten

☒ Benachrichtigung wenn noch nicht getriggerte Dateien älter sind als:

1 Tage 0 Stunden 0 Minuten

Speichern Abbrechen

If you enabled the option **Wait for SAFE-file** in the d.ecs storage manager configuration, then you must also enable the option **SAFE files** in the Silent Cube **Volume definition**.

Konfiguration (SC_RET)

SC_RET | HealthCheck

Ziel	Alias	Speicherpfad	Obere Speichergre...	Wiederb...	Untere Speichergre...	Volumd ID
SC_RE...		\\10.100.150...	90 %	True	50 %	011F3CCB

☐ Erzwinge CAS-Erzeugung

☐ Überprüfe Daten

☒ Auf SAFE-Datei warten

Verbindung prüfen

Info OK Abbrechen

The SAFE-files option **List** must be enabled for the d.ecs storage manager to process SAFE-files.

Set the file extension of SAFE-files to the default **safe**.

From d.ecs storage manager version 2.4, the extension for the Safe-files must be set to **xml**.

Bearbeiten eines Volumes

Name: vol_d3p Quota [MB]: 3096576 Autoinkrement [MB]: 102400 Schwelle [%]: 90

Volume-Typ: WORM mit Lebensdauer (RWRO getriggert) Jahre: 0 Tage: 0

Archivierung Safe-Dateien Sonstiges

☒ Safe-Dateien

☒ Auflistung

Dateierweiterung: safe

Format: xml

Speichern Abbrechen

The default settings in the section **Miscellaneous** should not be changed when you **Create of a new volume**. You can only enable "Prioritize the archiving", if you want to preconfigure a priority when using multiple volumes.

Bearbeiten eines Volumes

Name: vol_d3p Quota [MB]: 3096576 Autoinkrement [MB]: 102400 Schwelle [%]: 90

Volume-Typ: WORM mit Lebensdauer (RWRO getriggert) Jahre: 0 Tage: 0

Archivierung Safe-Dateien Sonstiges

Volume-Typ zu "WORM mit Lebensdauer (RWRO getriggert, Autosave)" ändern

☒ Priorisiere die Archivierung

☐ Optimierungen für Robocopy

☐ Optimierungen für Quantum StorNext

☐ 0-Byte Dateien nicht archivieren

☐ Umbenennen von leeren Verzeichnissen verbieten

☐ Umbenennen von Dateien verbieten


Speichern Abbrechen

You should not enable the Button **Change volume-type to "WORM with lifetime (RWRO triggered)"**.

Configuration Silent Brick

On a Silent Brick system, you must create a so-called "Compliant archive" via the web user interface before you can create a volume. Select a WORM brick as an archive medium for the "Compliant archive". If you are using the WORM Brick as a mere archive volume or as an archive medium and stage, depends on the bricks existing in the Silent Brick system. You can select a WORM brick as a mere archive medium and e.g. an SSD brick for the stage area. Detailed information can be found in the respective documentation on the Silent Brick system.

Under the created “Compliant archive”, you must create a so-called “Sub volume”. From software version 2.22 of the Silent Brick system, you must decide whether you want to create a sub volume in **Compliance** or in **Enterprise** mode. The following selection appears:



Choose Sub Volume Mode


Attention:
Once created the chosen volume type cannot be modified.
Please make sure to choose the right configuration meeting your needs. In most cases „Compliance” will be the desired feature.
When mode is set to Enterprise files may be deleted even before their retention is reached.

☒ **Compliance**
Files will be hardware locked against deletion for a given time period.
By using the compliance mode there will be no option to delete or modify files prior to the end of the defined retention time.
This volume mode is commonly used for lawful and compliant long-term archival where retention time is rated higher than personal rights due to legal requirements.

☐ **Enterprise**
Files will be hardware locked against deletion for a given time period but there will be an option to remove single files by overriding the retention using special API calls.
This volume mode is commonly used for the archival of enterprise data where rights of file owners may override legal requirements.

Cancel Continue

For the standard storage via the d.ecs storage manager you should choose the subvolume mode **Compliance**. If you need to delete documents before the retention time expires, you must use the Sub Volume Mode **Enterprise**. For a subvolume in Compliance or Enterprise Mode, choose the following settings:



Add Sub Volume

Name:
D3P

Description:
Volume for D3P

Volume type:
WORM with retention (RWRO triggered)

Years: 0 **Days:** 0 **Hours:** 0 **Minutes:** 0

Safe files **Miscellaneous**

☒ **List**

File extension:
safe

File format:
xml

Cancel Add

When creating the sub volume on the Silent Brick system, you must comply with the following preconditions to ensure the proper collaboration of the d.ecs storage manager with the Silent Brick system:

- The volume-type must be configured as **WORM with lifetime (RWRO triggered)**.
- The lifetime should be 0 years, 0 days, 0 hours and 0 minutes.

Note

The d.ecs storage manager passes the lifetime for every document to the Silent Brick system.

- Enable the option **List** so that the SAFE files in the file system are visible for the d.ecs storage manager.
- The **File extension** of the **Safe files** should not be changed and remain **safe** as default. Else, a proper collaboration with d.ecs storage manager is not ensured.
- Please select **xml** as the **File format** for the SAFE-files as this is expected by d.ecs storage manager.

The **Miscellaneous** settings for the “Sub volume” to be created should not be changed. Only the setting **Prioritize ingests** can optionally be enabled.

Add Sub Volume

Name:
D3P

Description:
Volume for D3P

Volume type:
WORM with retention (RWRO triggered)

Years: 0 **Days:** 0 **Hours:** 0 **Minutes:** 0

Safe files **Miscellaneous**

☒ **Prioritize ingests**

☐ **Robocopy optimizations**

☐ **Do not archive 0-byte files**

☐ **Prevent renaming of empty directories**

☐ **Prevent renaming of files**

Cancel **Add**

Note

To enable the d.ecs storage manager to store documents on a Silent Cube volume or a Silent Brick sub volume, a respective share to the volume must be configured. The share can be done for a domain user or a local user of the Silent Brick System.

Setting up a user for privileged deletion associated with Enterprise Sub Volumes.

The privileged deletion of documents, i.e. the deletion of documents before the retention time has expired, via the d.ecs storage manager can only be performed by a local user of the Silent Brick system. Create a corresponding user and assign him the role "ComplianceAdmin" for the corresponding "Compliance Archive".

Add User

User Name: PrivDel

Description: Compliance Admin User for privileged delete

Role: ComplianceAdmin - d3_Archiv

Password:

Retype Password:

To activate '2-factor authentication' please save and then edit this new user.

Cancel Save

1.4.11. d.ecs storage manager TSM

With the d.ecs storage manager TSM module, d.ecs storage manager can store documents on a Tivoli Storage Manager/IBM Storage Protect server.

A 64-bit operating system is required to use the d.ecs storage manager-TSM module due to its dependency on the IBM Storage Protect client.

In order to ensure a proper cooperation between d.ecs storage manager and the Tivoli Storage Manager/IBM Storage Protect server, an installed IBM Storage Protect client with version 8.2.2.0 is required.

The compatibility to a higher API version cannot be guaranteed due to a change in the API by IBM. Before updating d.ecs storage manager, ensure that the IBM Storage Protect client version to be used is compatible with the version you are using on the server.

Configuration

DSM-API directory: Specify the API-directory of the TSM client here.

Warning

After upgrading from an older d.ecs storage manager version (< 3.3) where the 32 bit TSM client was still used, the API directory has to be changed to the 64 bit variant.

By default, the API directory is located at:

C:\Program Files\Common Files\Tivoli\TSM\api64

OPT file (Global): Specify the OPT file (Options file of the TSM client) to be used by default. A global OPT file must be specified.

OPT file: Enter the explicit OPT file to be used, if the default OPT file is not used, e.g. to establish a connection to a separate TSM server.

Management class: Enter the management class based on which the TSM server should store a document.

Note

Note that an event-based management class is required here.

Filespace: Enter a name for the storage area on the TSM server.

Warning

If you operate several d.3 server (with the same repository IDs), which are saving their data on one TSM server, then you must make sure that the filespace for each d.3 server has a different name. Example: for server 1: d3store1, for server 2 d3store2. Otherwise this will result in problems later with the recovery of the index structure!

Connections: This defines the maximum number of simultaneous connections to the TSM server. If a connection is not used, it is automatically disconnected from the d.ecs storage manager after five minutes.

User: Enter the user name to be used for the TSM connection here.

Password: Specify the password for the TSM connection here.

Prevent closing connections: If this parameter is enabled, then the TSM-connection is kept open after a retrieval.

Warning

Only enable this option, if you are sure that all objects do not have to be read from a sequential storage medium.

You can obtain information about how to use a management class for the further processing of the documents by the Tivoli Storage Manager in the documentation of the Tivoli Storage Manager.

Since the passwords of a TSM user expire after certain intervals and must be reset, it is recommended to use the login procedure `passwordaccess generate` of the TSM client. This procedure automatically resets the passwords automatically after their expiry.

To activate this procedure, enter the following line in the OPT file:

```
passwordaccess generate
```

The OPT-file could then look as follows:

TCPServeraddress	192.168.0.245
Nodename	d3user
TCPPORT	1500
passwordaccess	generate
enablearchiveretentionprotection	yes


Please consult the TSM client manual for information on the respective entries of the file.

Having changed the login procedure in the OPT file, please delete the user name in the configuration mask of the TSM system. After a successful connection to the TSM, the password can also be deleted from the configuration mask. The new password saves the TSM client encrypted in the Microsoft Windows registry.

Index recovery

The d.ecs storage manager TSM module offers the possibility to restore index recovery files (*.IR) stored on the TSM in order to rebuild the reference index of d.ecs storage manager (see also [Tab Database logging](#)).

Enough disk space must be available in the specified index directory so that the downloaded index recovery files can be stored there. The filespace name should be assigned automatically. In order to search for Index recovery file on the pool "B", the Pool ID **B** must be specified. Then the recovery process is started with a click on the button **Start recovery**. Depending on the number of documents located on the document pool, the search can take several minutes or even hours.



1.4.12. d.ecs storage manager Governikus LZA/DATA Aeonix

The d.ecs storage manager Governikus LZA module enables you to store documents in a Governikus LZA/DATA Aeonix signature system. For final storage on a separate storage system, you can use a second (downstream) d.ecs storage manager that stores the data on the supported storage systems.

To use this module, you require a Governikus LZA system with version 3.5 or a DATA Aeonix system with version 10.1.2, with the latest hotfixes installed in each case. The most recent tested version of LZA is 3.5.3.2. The most recent tested version of DATA Aeonix is 10.1.2. Configure the LZA/DATA Aeonix system so that the data is stored in TR-ESOR 1.2 format.

Warning


Since the Governikus LZA/DATA Aeonix signature system uses special automated signature functions that require access to the original document content (for re-signing), the Single Instancing function is not supported when using a Governikus LZA/DATA Aeonix signature system.

Warning

Due to current restrictions in the d.ecs storage manager API, the maximum document size that can be sent to LZA/DATA Aeonix is 1.2 GB.

Communication chart

The following diagram shows the technical procedure when there is a second d.ecs storage manager downstream from the LZA/DATA Aeonix system.



Secondary storage (red)

- d.3/ecspand sends the document and signatures to d.ecs storage manager.
- d.ecs storage manager Governikus LZA generates a XAIP container from it and sends it to the LZA system.
- The LZA system checks the XAIP container, adds a timestamp, generates the AOID and passes it to the LZA storage adapter (the LZA must be configured in such a way that evidence records are not also stored in the XAIP to prevent duplicate storage later).
- The LZA storage adapter sends the XAIP container to a second downstream d.ecs storage manager instance and uses the passed AOID as a DOCID. The LZA storage adapter then waits for the downstream d.ecs storage manager instance to store the container in the secondary storage.
- The downstream d.ecs storage manager instance stores the document on the configured secondary storage system.
- If the document has been stored on the secondary storage, the LZA storage adapter returns an OK to the LZA system. The LZA system passes the OK and the AOID to the first d.ecs storage manager, d.ecs storage manager stores the AOID for the document and signatures and returns the OK to d.3/ecspand.
- The first d.ecs storage manager asynchronously queries the LZA system for the evidence records of the documents stored on the secondary storage and stores them via the downstream d.ecs storage manager.

Unproven restore (green)

For the ordinary (unproven) restore (green), the first d.ecs storage manager instance loads the XAIP container using the stored AOID directly via the API from the second downstream d.ecs storage manager instance and extracts the document/signatures.

Proven restore (blue)

In the proven restore (blue), the first d.ecs storage manager instance requests the XAIP container from the LZA system. The LZA system then loads the XAIP container through the LZA storage adapter via the d.ecs storage manager API from the second downstream d.ecs storage manager instance. The LZA storage adapter then passes the XAIP container to the LZA system. The LZA system checks the XAIP container and passes it to the first d.ecs storage manager instance. This can now extract the document / the signatures and deliver them to d.3/ecspand.

Configuration

LZA connection settings

Server: The IP or host name of the DATA Aeonix server.

Port: The port for the ArchiSafe module of the LZA/DATA Aeonix server.

Certificate file: The certificate for client-side authentication (in PEM format).

Key file: The Key file for client-side authentication (in PEM format).

Test connection: Tests the connectivity to the LZA/DATA Aeonix server.

Signature validation mapping

Select the signatures to be checked by the Governikus LZA/DATA Aeonix system.

Attribute settings

Properties of d.3 can only be adopted in the XAIP container from d.3 server version 8.1.

In order to transfer the document properties from d.3 to the Governikus DATA Aeonix system, the credentials of a d.3 user with access to all documents and document types must be specified in the

configuration. The user must be configured in such a way that their access never expires, otherwise secondary storage may fail due to login issues.

Store properties in XAIP container: Enables saving the properties from d.3 to the swapped XAIP container.

Base address: Specifies the base address under which the d.ecs http gateway is accessible.

Archive ID: The main ID of the d.3 repository.

User: The user to be used for access to the d.3 repository. Leave this parameter empty if you are using a technical user.

Password / API token, Confirm password / API token: The password of the d.3 user in question or the API token if using a technical user.

Test connection: Checks the access to the d.3 API using the specified user and password.

Secondary d.ecs storage manager instance

If the option **Use secondary d.ecs storage manager** is enabled, an additional page is displayed for the configuration of the second d.ecs storage manager instance.

Server: The IP or host name of the server on which d.ecs storage manager is running.

Port: The port of the web interface or the API of d.ecs storage manager.

Password/Confirm password: Password for access to the web interface or API.

Pool ID: Pool ID of the pool where the documents are to be stored.

Use SSL encryption: This enables the SSL encryption.

Proven restore: If this option is enabled, the documents are always requested via the Governikus DATA Aeonia system. This option should only be enabled in case of doubt for performance reasons. If this option is enabled the documents are always requested from the downstream d.ecs storage manager by API.

Test connection: Checks the connection to the d.ecs storage manager.

Preparing the Governikus LZA/DATA Aeonia system

A certificate is required to establish a connection to the Governikus LZA/DATA Aeonia system. The d.ecs storage manager Governikus LZA module uses this certificate to log in to the Governikus LZA/DATA Aeonia system. The certificate consists of a certificate file and a matching key. Both must be provided in PEM format.

Note

The following settings are required only if there is a second d.ecs storage manager instance downstream from the LZA/DATA Aeonia system.

To allow Governikus LZA/DATA Aeonia to communicate with the downstream d.ecs storage manager, you must manually install a special storage adapter in combination with an adjusted configuration file in Governikus LZA/DATA Aeonia. Following the installation of d.ecs storage manager, the storage adapter is located in the sub-directory **LZA-StorageAdapter** (**storage_decssm_lza_storage_adapter.rar**).

Installing the LZA storage adapter on the DATA Aeonia system

Note

To install the adapter, you must have a DATA Aeonix system running version 11.1.5 or later.

To install the LZA storage adapter

- Copy the file **storage_decsm_lza_storage_adapter.rar** from the d.ecs storage manager sub-directory **LZA-StorageAdapter** to the DATA Aeonix installation directory (for example, to a deployments sub-directory).
- Start the JBOSS CLI and connect to the JBOSS server.
- Make the storage adapter known to JBOSS using the following command:

```
deploy [Installation directory]
\deployments\storage_decsm_lza_storage_adapter.rar
```

- Optional: Check whether the storage adapter is available in the file **standalone.xml** in the directory [Installation directory]\jboss\standalone\configuration\ by searching for **storage_decsm_lza_storage_adapter.rar**.
- Log in to the DATA Aeonix WebAdmin and select the long-term storage section.
- Select the option **General configuration** in the menu on the left.
- Under **Archiving type**, select the **d.ecs storage manager LZA storage adapter**.

An input form appears for the configuration data of the d.ecs storage manager to be used for the secondary storage of the DATA Aeonix system data.

The parameters have the following meaning:

Server-Name / IP: Host name or IP of the server on which d.ecs storage manager is running.

Port: The port of the web interface/API of the d.ecs storage manager.

IDP-Token verwenden (Use IDP token): Indicates whether the password is the API token of an IDP user. Possible values: true, false.

Password: Password for the d.ecs storage manager web interface/API. If the value **true** is specified under **IDP-Token verwenden** (Use IDP token), you need the API token of an IDP user who has access to the d.ecs storage manager web interface or API.

Pool ID: Pool ID of the pool where the documents are to be stored.

Use SSL: Specifies whether the connection to the d.ecs storage manager is to be established via SSL. Possible values: true, false.

After specifying the parameters:

- Select **Apply** in the lower section.
- Select the option **Long-term storage** in the menu on the left.
- Save the configuration.

In addition, please note that saving of the Evidence Records in a XAIP container (container format for documents and signatures) in the Governikus DATA Aeonix system must be disabled to prevent the duplicate storage of data. To change the settings, log in to the DATA Aeonix WebAdmin, go to **ArchiSig modules > General configuration** and disable the option **Insert ERs into XAIPs** in the **Embedding EvidenceRecords** category.


1.4.13. d.ecs storage manager system connect

The d.ecs storage manager system connect module allows access to third-party repositories using special connectors. The connectors provide a read-only access to the third party repositories. The secondary

storage of documents via the d.ecs storage manager system connect module and the respective connectors is not possible.

Configuration

The d.ecs storage manager system connect module has a minimalistic configuration form only allowing you to select the connector to be used. The then following configuration dialog is provided by the selected connector and must be explained in the connector's documentation.



Connector: This defines the connector to be used.

Configure: Opens the configuration dialog of the selected connector.

Note

A connector cannot be changed later. If another connector is to be used, a new system must be created in d.ecs storage manager.


1.4.14. d.ecs storage manager Isilon

The Isilon module allows the d.ecs storage manager to store documents on a Dell EMC Isilon storage system.

Warning

To store the documents with the d.ecs storage manager on an Isilon volume, it is essential that the default, minimum, and maximum retention period for the Isilon volume are set equivalent to the retention periods in the leading system (d.3/ecspand).

Configuration



Objective: List of the target paths where the documents are stored and their properties.

Alias: Alias name of the volume.

Storage path: Target path where the documents are stored.

Upper storage limit: Specifies how much storage space may be used for storing data.


Rewritable: If this option is set, the data will be re-written as soon as the available storage space falls below the Lower storage limit.

Lower storage limit: Specify the percentage of disk space which may be occupied before documents can be stored again.

Volume ID: Every storage path gets a unique volume ID.

CAS-creation: If this option is enabled, then the documents are stored in CAS containers. These containers have a unique ID which is used to verify the consistency of a document during restore. (During the retrieval it is checked, if the data in a CAS container are unchanged before a document is passed on to the d.3 server/d.3 gateway/API). This should only be enabled in connection with ecspand products.

Validate data: If this option is enabled, then a file comparison based on the **RipeMD256** hash procedure after copying a file to the Isilon volume to make sure that the file was transferred completely.



File name for HealthCheck: File name to be used during the system check by the HealthCheck process.

Volume configuration

Alias: Alias name for the volume. If a name is specified here then this is used instead of the volume ID in the web interface. This is used for a better overview.

Storage path: Specify a path here where the documents are stored

Upper storage limit (in %): Here you can define a percentage of how much storage space may be used for the storage of data.

Lower storage limit (in %): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Upper storage limit (in MB): Here you can define in MB how much storage space may be used for the storage of data.

Lower storage limit (in MB): Specify the percentage of hard disk space which may be occupied before documents can be stored again.

Rewritable: This option is only available, if more than one volume is configured and specifies, if a volume should be rewritten as soon as the lower storage threshold of the volume is exceeded (e.g. by deleting old documents).

Minimum/maximum retention period (in months) Specifies the range in which the retention time passed by the delivering process may be located for a document. These values have to be set and must be equivalent to the settings of the Isilon volume.

Automatically correct retention period: If a retention-period outside the specified range is passed by a delivering process, then the time is automatically set to the minimum or maximum, if this parameter is enabled. Else the job is blocked and the document is not written to the secondary storage.

Log in with the following user: If this parameter is enabled, d.ecs storage manager automatically tries to log in to the respective volume. This parameter only has to be enabled, if d.ecs storage manager is executed under a user account which does not have access to the volume.

User: User to be used for the login to the volume.

Password / Confirm password: Password of the user to be used for the login to the volume.

Status: One of three possible values of the volume is displayed. Use the **Reset volume status** button to reset the status of a volume to **Waiting**.

- **Waiting:** There is currently no writing on the volume.
- **Writing:** Currently writing on the volume.
- **Full:** The volume was written on until it exceeded the configured memory limits. The volume cannot and will no longer be written to.

Preparation of the Isilon volume

The configuration of the Isilon volume has various requirements which will be described below.

- It is recommended that the values **Minimum Retention** and **Default Retention** are set to one day and **Maximum Retention** to the highest retention time defined in d.3/ecspand. This avoids possible issues when setting the retention time through d.ecs storage manager. Else, Isilon overwrites the retention periods set by d.ecs storage manager automatically with the specified defaults.
- The parameter **Auto Commit** must not be enabled, as the Isilon then automatically locks documents and under certain conditions this interferes with the procedures of d.ecs storage manager.
- Since d.ecs storage manager does not support the function **Privileged Delete** of Isilon, this option can be disabled.

Create a WORM Domain

* = Required field

Domain Settings

Privileged Delete
Off

* Path
/ifs/data/SmartLock_JE Browse...

Apply Retention Settings to WORM Protected Files

☒ Apply a default retention span
Specific
1 Days

☒ Enforce a minimum retention time span
Specific
1 Days

☒ Enforce a maximum retention time span
50 Years

☐ Automatically commit files after a specific period of time

☐ Override retention periods and protect all files until a specific date

Cancel Create Domain

1.4.15. d.ecs storage manager S3

With the module d.ecs storage manager S3, d.ecs storage manager can store documents on an S3-compatible data storage device. The S3 data storage itself can be provided by the customer.

Currently supported systems:

- Standard S3 (**without revision security**)
- DELL EMC/ECS
- d.velop cloud storage OTC
- S3 Object Lock Governance/Compliance
- more systems upon request

The documents can optionally be protected against unauthorized viewing with AES256 encryption. When storing data with an Internet S3 provider, enabling encryption is recommended.

The operation of d.ecs storage manager S3 requires an installed .NET Core Runtime in version 3.1.5 or higher. This is installed during the installation of d.ecs storage manager.

Note

The d.ecs storage manager S3 module is only to be used in conjunction with on-premises storage solutions. For use with Amazon S3 or Open Telekom Cloud, d.velop cloud storage (AWS) or d.velop storage service (OTC) can be booked with d.velop AG. If you want to use the S3 module with other S3 cloud systems, you must license the d.ecs storage manager S3 for cloud module.

Configuration

The configuration of an S3 system in d.ecs storage manager consists of three sections: access data, encryption and proxy settings. The individual sections are explained below.

Access data

Access key: Enter the access key with which access to the S3 memory is possible.

Secret access key: Enter the secret access key with which access to the S3 memory is possible.

Bucket: Enter the name of the target bucket here.

Endpoint: Enter the service endpoint address of your S3 memory here.

Region: Specify the region to use in your S3 memory.

Tenant ID (optional): A tenant ID (max. 32 characters) can be specified here. This specification is optional and only necessary if several clients are to be stored in the same bucket. Specifying a tenant ID causes the client data to be divided into different subdirectories below the root directory of the target bucket.

Test connection: Tests the connection to the specified service endpoint and displays information about the configured encryption.

Encryption

The setup of the encryption is described in more detail in the subchapter [Encryption](#).

Proxy

It is possible to use a proxy server to communicate with the S3 memory. Enter the login credentials for the proxy server available in your network in the corresponding fields. The username and password are optional fields and only need to be filled if you are using a proxy server with user authentication.

New functions

Some S3 storage systems support advanced S3 features such as the use of retention periods for data objects. This functions can be enabled optionally.

Activate the advanced functions and select the type of the connected storage system.

Note

If only standard S3 is to be used, then this must be explicitly confirmed. In this case, documents are not stored in an audit-proof manner (no protection by setting retention periods).

Miscellaneous

Under the section Miscellaneous there are various options that are important for one or the other S3-based system.

Force "path style": Depending on how a bucket is addressed in the addressed system, path style or host style (default) must be used.

Disable SSL checking: Disables the validity check of SSL certificates.

Chunk size in MB: Specifies the chunk size for uploading data objects. Minimum is 5 MB, maximum is 100 MB.

Connection timeout: Specifies the maximum time in seconds for a connection attempt. The default value is 5 seconds.

Read/write timeout: Specifies the maximum time in seconds for a read or write attempt. The default value is 300 seconds.

Number of trials: Specifies the maximum attempts for an S3 API call. The default value is 2.

Encryption

On the **Encryption** configuration page, you can activate the encryption of the data objects so that the objects are stored encrypted in the S3 memory.

Depending on the extent to which the encryption is already configured and initialized on your system, different configuration views may be displayed when you switch to the configuration page.

Initializing and activating the encryption function

To enable the encryption of the data objects, you must first generate a security ID.

The movement of the mouse is taken into account during generation, resulting in an even more random key.

- Click **Start** to begin with the generation.
- Click on the button again after a few seconds to stop the generation. So that the mouse is not necessarily at the position of the button, the button can also be operated by pressing the space bar.
- Then, you must print the generated security ID. The printout contains a confirmation code that you must enter in the **Confirmation** section.

Encryption can be activated only once the security ID has been entered correctly.

It may take some time to activate encryption (approximately 15 to 20 seconds).

After successful enabling of the encryption, the message "The encryption is active" appears.

The encryption has now been enabled and d.ecs storage manager can store the documents/data objects encrypted in the connected S3 memory.

Activating encryption when the encryption function has already been initialized

If the encryption has already been initialized, it can simply be enabled for another bucket.

- Click **Enable** to activate it.

It may take some time to enable the encryption (approximately 2 to 5 seconds).

After successful enabling of the encryption, the message "The encryption is active" appears.

The encryption has now been enabled and d.ecs storage manager can store the documents/data objects encrypted in the connected S3 memory.

Renewed initialization of the encryption function after the master key was lost

If the initialization of the encryption function is lost, e.g. as a result of a clean installation of the operating system, a server migration or changed hardware, the encryption function must be enabled again.

This is done using the security key on the printed copy generated at the time of initializing the encryption.

- Enter the security key that appears on the printed copy.
- Click **Re-enable**.

It may take some time to activate encryption (approximately 15 to 20 seconds).

After successful enabling of the encryption, the message "The encryption is active" appears.

The encryption has now been enabled and d.ecs storage manager can store the documents/data objects encrypted in the connected S3 memory.

More information about supported storage systems

d.velop cloud storage OTC

Notes on storage and deletion of data objects

The data objects stored in d.velop cloud storage OTC each have a retention period. This retention period is given by the leading system when the data objects are stored. The retention period is taken into account during a deletion process and the deletion is stopped if the retention period has not yet expired.

Documents that can be deleted due to an expired retention period are deleted directly during a deletion process. This deletion process cannot be undone. Deleted data objects cannot be restored.

During configuration, note that the endpoint is set to "<https://obs.eu-de.otc.t-systems.com>" and the region is set to "eu-de". The accessibility of the endpoint may have to be enabled in the existing firewall. Furthermore, the accessibility of the end point "<https://otc-storage.service.d-velop.cloud>" must be guaranteed.

DELL/EMC ECS

If you are using d.ecs storage manager S3 in conjunction with an ECS (Elastic Cloud Storage), then you must enable the **Enforce path style** option on the **Miscellaneous** tab.

S3 Object Lock

The following S3-compatible systems were successfully tested with the Object Lock option activated:

System	Version	Object Lock mode	Tested on
Hitachi Content Platform (HCP)	9.6.0.214	Governance	13.10.2023
Hitachi Content Platform (HCP)	9.6.0.214	Compliance	13.10.2023
iTernity iCAS FS	1.18.4	Governance	06.08.2025
iTernity iCAS FS	1.18.4	Compliance	06.08.2025

System	Version	Object Lock mode	Tested on
NetApp Storage Grid	11.6.0	Compliance	06.01.2023

1.4.16. d.ecs storage manager proNEXT

The d.ecs storage manager proNEXT module enables you to store documents in a Procilon proNEXT signature system. At least version 2.12.0 of proNEXT ArchiveManager is required for operation.

Warning

Since the Procilon proNEXT signature system includes special automated signature functions that need to revert to the original document contents (for oversigning), the singlein-standing function is not supported when using a Procilon proNEXT signature system.

Configuration

proNEXT connection settings

ArchiSafe: The IP address or the host name and the port of the ArchiSafe service of the proNEXT server.

UserManager: The IP address or the host name and the port of the UserManager service of the proNEXT server.

IDP: The IP address or the host name and the port of the IDP service of the proNEXT server.

User: The user name for accessing the proNEXT server.

Password: The password to access the proNEXT server.

Keycloak: Specifies whether Keycloak should be used for login.

Test connection: Tests connectivity to the proNEXT server.

XAIP mode: Select the mode to be used for compiling an XAIP.

- **XAIP:** When the XAIP container is created, the data objects are also saved in the XAIP container.
- **LXAIP:** When the LXAIP container is created, the data objects in the XAIP container are only referenced. This makes the container considerably smaller.

Signature validation mapping: Select the signatures that can or should be verified by the proNEXT system.

Attribute settings

Properties of d.3 can only be adopted in the XAIP container from d.3 server version 8.1.

To be able to transfer the document properties from d.3 to the proNEXT system, save the access data of a d.3 user in the configuration.

Ensure the following:

- The user has access to all documents and document types.
- The access never expires. Otherwise, it may not be possible to perform the secondary storage due to login problems.

Store properties in XAIP container: Enables saving the properties from d.3 to the swapped XAIP container.

Base address: The basic address at which d.ecs http gateway can be reached.

Archive ID: The main ID of the d.3 repository.

User: The user to be used for access to the d.3 repository. The password of the d.3 user in question or the API token if using a technical user.

Password / API token, Confirm password / API token: The password of the d.3 user in question or the API token if using a technical user.

Test connection: Checks the access to the d.3 API using the specified user and password.

Routes: Here you can configure the endpoints you use for the proNEXT services. To configure the default routes, you can enable the corresponding option for the default routes. If you use Keycloak, there are various routes for the IDP service.

Important notes on the use of LXAIP containers

When using LXAIP containers, in contrast to normal XAIP containers, the data objects are not transferred to the proNEXT system. The transferred containers only contain metadata and object references with which the proNEXT system can request the data objects from d.ecs storage manager. Omitting the data objects leads to considerably smaller XAIP containers in the proNEXT system. The advantage is that changes to the metadata or evidence information of the LXAIP container do not lead to the user data being recreated.

Please note the following:

- To enable the proNEXT system to request the data objects from d.ecs storage manager, the data objects must be stored on a separate storage system. When configuring the storage pools, make sure that at least two systems are assigned to the respective pool. The first system must correspond to the system to which the data objects of a storage job are initially to be stored. Only then are the data objects transferred to the proNEXT system as a reference.
- You can deactivate the retrieval for this system in the system settings of the proNEXT system in d.ecs storage manager, as the d.ecs storage manager proNEXT module does not support the retrieval of documents from an LXAIP container. In this scenario, the data objects are always retrieved from the first system in the storage pool. If XAIP containers are already stored in the respective proNEXT system, retrieval must be activated. Otherwise, documents from the XAIP containers that have already been stored cannot be retrieved.
- Only run the system check from a proNEXT system with LXAIP if the retrieval test remains deactivated. Otherwise, retrieval errors will occur. The mere existence check of the containers is possible.
- When synchronizing the data to a new system, the data must be synchronized from the separate storage system. You can then perform a synchronization from the proNEXT system to the new system. Synchronization of the proNEXT system is only necessary if there are also XAIP containers in the proNEXT system in addition to the LXAIP containers.
- Do not select a system that is configured for LXAIP as the destination for the database logs (IR files). IR files are blocked by the proNEXT module for technical reasons if LXAIP mode is configured.

1.4.17. d.ecs storage manager Azure

With the d.ecs storage manager Azure module, d.ecs storage manager can store documents in an Azure Blob Storage container.

Optionally, the documents can be protected from unauthorized access using AES256 encryption.

Configuration

The configuration of an Azure system in d.ecs storage manager consists of three sections: access data, encryption and proxy settings. The individual sections are explained below.

Access data

Account Name: Enter the account name for accessing the Azure memory.

Account Key: Enter the account key for accessing the Azure memory.

Container: Enter the name of the target container here.

Tenant ID (optional): You can enter a tenant ID (max. 32 characters) here. This entry is optional and only necessary if several tenants are to be stored in the same container. Specifying a tenant ID causes the tenant data to be divided into different subdirectories below the root directory of the target container.

Test connection: Tests the connection to the specified service endpoint and displays information about the configured encryption.

Encryption

For more information on setting up encryption, see [Azure encryption](#).

Proxy

You can use a proxy server to communicate with the Azure memory. Enter the login credentials for the proxy server that is available in your network in the corresponding fields. The username and password are optional fields and only need to be filled if you are using a proxy server with user authentication.

Miscellaneous

Under the **Miscellaneous** section, there are various options that are important for communication with the Azure system.

Disable SSL checking: Disables the validity check of SSL certificates.

Number of trials: Specifies the maximum attempts for an Azure API call. The default value is **2**.

Encryption

On the **Encryption** configuration page, you can activate the encryption of data objects. This means that the objects are stored in the Azure memory in encrypted form.

Depending on the extent to which the encryption is already configured and initialized on your system, different configuration views may be displayed when you switch to the configuration page.

Initializing and activating the encryption function

To enable the encryption of data objects, you must first generate a security ID.

The movement of the mouse is taken into account during generation, resulting in a more random key.

1. Click **Start** to begin the generation process.
2. After a few seconds, click **Stop** to stop the generation. You can also operate the button by pressing the space bar so that the mouse is not necessarily at the position of the button.
3. Print out the security ID that is generated. The printout contains a confirmation code that you must enter in the **Confirmation** section.

Warning

Keep the printout safe. If the printout and thus the security ID are lost, it will not be possible to restore the master key.

When replacing the hardware components, you may have to enter the security ID again.

Encryption can be activated only once the security ID has been entered correctly.

It may take some time to activate encryption (approximately 15 to 20 seconds). Once encryption has been successfully activated, the following message is displayed: **Encryption is active**.

Encryption has now been activated and d.ecs storage manager can store the documents and data objects in the connected Azure memory in encrypted form.

Activating encryption when the encryption function has already been initialized

If encryption has already been initialized, you can simply activate encryption for another container.

Click **Enable** to activate it.

It may take some time to activate encryption (approximately two to five seconds). Once encryption has been successfully activated, the following message is displayed: **Encryption is active**.

Encryption has now been activated and d.ecs storage manager can store the documents and data objects in the connected Azure memory in encrypted form.

Re-initialization of the encryption function after the master key was lost

If the initialization of the encryption function is lost, e.g. as a result of a clean installation of the operating system, a server migration or changed hardware, you must reactivate the encryption function.

To do so, use the security key on the printed copy that was generated when encryption was initialized.

1. Enter the security key that appears on the printed copy.
2. Click **Re-enable**.

It may take some time to activate encryption (approximately 15 to 20 seconds). Once encryption has been successfully activated, the following message is displayed: **Encryption is active**.

Encryption has now been activated and d.ecs storage manager can store the documents and data objects in the connected Azure memory in encrypted form.

Preparing the Azure Blob Storage container

For the storage account under **Data Protection**, activate the option **Enable Versioning of Blobs** and for the container during creation, activate the option **Enable version-level immutability support**. This is to enable you to store data via d.ecs storage manager in a container in Azure Blob Storage.

1.5. Opening d.ecs storage manager

Start the d.ecs storage manager as follows:

```
decssm.exe [option] <configfile>
```

You can pass one of the following start options to the d.ecs storage manager:

/SA: Starts the storage process only

/SR: Starts the restore process only

/SAR: Starts the storage and restore process

If none of these parameters is passed, then the d.ecs storage manager starts without activating the storage or restore process. In this case, it is then possible to activate the storage/retrieval process via the [Web interface](#). If the secondary storage- and/or the restore process is started via the web interface, this only applies until the next start of the d.ecs storage manager.

Other internal d.ecs storage manager processes are automatically started independent from the passed start parameters.

Note

The name of the configuration file must always be passed as a parameter or the d.ecs storage manager will automatically stop after the start.

The configuration files are saved in the sub-directory `configs` in the program directory of the d.ecs storage manager and can be passed to the d.ecs storage manager. A configuration always consists of a configuration file (INI file) and one or more module configuration file (CFG file). The name of a module configuration file is equivalent to the name of the module used (e.g. `netapp.cfg`, `centera.cfg`, `tsm.cfg`, ...). The configuration file (INI) to be used must be passed when calling the d.ecs storage manager.

If the d.ecs storage manager is started via the d.3 process manager, then you must create a new process there entering the command line described above under Target, e.g.:

```
c:\decs\decsstoragemanager\decsm.exe /SAR
c:\decs\decsstoragemanager\configs\d3FU.ini
```

Furthermore, the **Working path** in the d.3 process manager must be set to the d.ecs storage manager directory (e.g. `c:\decs\decsstoragemanager`).

1.6. Web-Interface

d.ecs storage manager offers a web interface for the monitoring of the d.ecs storage manager. The individual processes (storage/retrieval) can be enabled or disabled.

Furthermore, you can obtain information on the currently installed storage drivers. For troubleshooting purposes the debug logging can be enabled during runtime.

Status Prozess	Aktuelle Nachricht
Auslagerung	Wartet auf auszulagernde Dokumente
1. Auslagerung	Wartet auf Job
2. Auslagerung	Wartet auf Job
3. Auslagerung	Wartet auf Job
4. Auslagerung	Wartet auf Job
Wiederherstellung	Wartet auf zu verarbeitende Jobs
1. Wiederherstellung	Wartet auf Job
2. Wiederherstellung	Wartet auf Job
3. Wiederherstellung	Wartet auf Job
4. Wiederherstellung	Wartet auf Job
5. Wiederherstellung	Wartet auf Job
6. Wiederherstellung	Wartet auf Job

The links at the left and upper border of the web-interface allow to retrieve and display certain information as well as apply respective configuration settings for the system check or synchronization.

Open the web interface

Optionally, open the web-interface from the server hosting d.ecs storage manager via the address

```
http://localhost:[Port]/
```

must be called. The port is defined in the d.ecs storage manager configuration (see [Tab: HTTP / API Interface](#)).

Note

If the web interface is opened from another client, a login dialog is displayed. In this case, the user name is "Admin" and the password is equivalent to the web password configured in the d.ecs storage manager administration.

For a smooth operation of the web interface, you must enable Javascript (active scripting) in the web-browser.

1.6.1. Document information

The hyperlink **Document-information** provides you with information on a document written to the secondary storage by the d.ecs storage manager. The displayed information is determined from the d.ecs storage manager database table **d3sm_references** at runtime. In combination with d.3, you define the d.3 document ID and the respective system. Using these restrictions, all document-versions of the specified document ID are displayed. In a second step, you can select detailed information on a selected document-version and finally execute a check, if the selected document can be restored from the storage system.

Defaults for the **Document Information**

 **Dokument-Informationen:**

Dokument:

Speichersystem:

CENTERA01 (CENTERA)
 HCP_6 (HCP)
 NAS01 (NAS)
 NETAPP_SLE (NETAPP)
NETAPP01 (NETAPP)
 SSAM01 (TSM)

Display of the document versions

Gefundene Dokumente:

<u>Document ID</u>	<u>Document Extension</u>	<u>Archiv Index</u>	<u>Overwrite Counter</u>
A0000013		1	1

- Select the document version with a click on the row with the document ID.

Detailed information a document version

```

reference_index: 26169
system_id: 3
doc_id: A0000013
doc_ext:
doc_idx: 1
overwrite_counter: 1
deleted: 0
phys_filename:
container:
reference: A0000013.1
hash_value: RMD256:SKvqFr4l6MKbZEAfMAtFI5ak1tK+AzDvd115LOcjGs=
ttl: 2014-12-31
size_original: 1024
size_compressed: 1024
pool_id: A
sourced_out: 2014-07-18
volume_id: 031DCCE8
delete_token:
retention_on_storage: 2015-01-18
system_check_result: 0

```

Prüfe Wiederherstellung

Check, if the document can be restored from the storage system

- To do so, select the link **Check recalling**

Note

The hyperlink **Check recalling** is only displayed, if the restore process is enabled. Also compare [Process overview](#).

Dokument A0000013.1_1 wurde erfolgreich wiederhergestellt

If it is not possible to restore a document, a respective error message is displayed. The error message is displayed in red.

1.6.2. Driver information

Use the hyperlink **Driver-information** to get information on the installed and configured d.ecs storage manager modules. You also see the respective version number for the respective modules. In connection with the **System-information**, this can be essential information in a support case.

**Treiber-Informationen:**

Modulname	Hersteller	Speichersystem	Version	Modulinformationen
celerra.smd	d.velop AG	CELERRA	3.3.0.0	This module includes the d.ecs storage manager Celerra addon, that adds support for Celerra Storage Solutions to the d.ecs storage manager.
centera.smd	d.velop AG	CENTERA	3.3.0.0	This module includes the d.ecs storage manager Centera addon, that adds support for EMC ² Centera Storage Solutions to the d.ecs storage manager.
cloudstorage.smd	d.velop AG	CLOUDSTORAGE	3.3.0.0	This module includes the d.ecs storage manager cloud storage addon, that adds support for using the d.velop cloud storage to the d.ecs storage manager.
datadomain.smd	d.velop AG	DATADOMAIN	3.3.0.0	This module includes the d.ecs storage manager Data Domain addon, that adds support for EMC ² Data Domain Storage Solutions to the d.ecs storage manager.

1.6.3. Job overview

Since version 2.5 of the d.ecs storage manager, the d.3 jobs (from d.3 v7.1.1) and the API-jobs for the job-types **STORE**, **DELETE**, **PRIVILEGED DELETE**, **RETENTION** and **COPYJOB** are managed via a job table. The job-overview allows to view and filter the currently pending with various criteria.. Jobs that encountered errors can be optionally continued, restarted or removed.

The Job-Overview is divided in the sections **Status**, **Filter** and **Result list**.

The section **Status** displays the number of jobs grouped by job types. The display is refreshed automatically.

With a click on one of the jobs, the respective job type is selected as a filter. The row is highlighted accordingly and the job type is displayed in the title of the section **Filter**.



Status

Job-Typ
Auslagerungs-Jobs
Aufbewahrungsfrist-Jobs

The section **Filter** allows you to initially filter the job list. For this effect, you can specify a job ID or restrict the date range. Click on **Start search** to begin searching the jobs. The results are listed in the section **Result list**.

Filter > Store jobs

Suche starten

Job-Id:

Datumsfilter:

☒ aktivieren

Startdatum:

26.11.2016

Enddatum:

01.12.2016

The section **Result list** displays the result of the search. One element is displayed per row. The result list can be further filtered with several options. Click the search field **Search for...** to search for job properties. A separate filtering per property can be applied entering the filter criteria in the respective input fields displayed after having enabled the option **Advanced filtering**.

A click on an entry shows all job properties. The result list allows you to select one or more jobs and apply an action to them. The available options are displayed as buttons in the column **Options**. If multiple jobs are selected, then all available options are displayed as buttons in the title bar. The options are only available, if a job has the status Error.

Trefferliste (1 / 2)

- ☐ Alle auswählen
☐ Erweiterte Filterung

Suchen nach...

Status	Dokument-Id	Dokument-Index	Dokument-Erweiterung	Fehlermeldung	Pool-Id	Job-Typ	Optionen
☐		P000000011	0		P	RETENTION	
☐		P000000011	6		P	STORE	
Eigenschaften							
Dokument-Hash	RIPEMD256:EIMadPy9BNd7e6waB7HempaqT5biVPsJtzipw+zJHx3E=						
Dokument-Id	P000000011						
Dokument-Index	6						
Dokumenttyp	TXT						
Job-Datum	6/3/2019 12:06:05 PM						
Job-Id	35925						
Job-Typ	STORE						
Pool-Id	P						
Aufbewahrungsfrist	03.06.2020						

1.6.4. Pool overview

The Pool-overview lists information on the defined pools, the systems, and optionally the volumes per system.

By default, the number of stored and searched documents as well as the number of encountered errors is displayed.

The numbers shown refer to a time range calculated from the difference between the current date and time and the date and time displayed at the lower end of the Pool overview next to Reset pool statistics.

The numbers can be reset with the link **Reset pool statistics**.

Pool Übersicht

Pool	Verwendete Systeme	Verwendete Volumes	Dokumente ausgelagert	Dokumente ausgelagert (Fehler)	Dokumente wiederhergestellt	Dokumente wiederhergestellt (Fehler)
▼ P			0	0	0	0
	 NAS		0	0	0	0
		 Volume01	0	0	0	0
		 Volume02	0	0	0	0
		 Volume03	0	0	0	0
	 HCP		0	0	0	0
Gesamt			0	0	0	0

Pool-Statistik zurücksetzen (07.11.2016 11:13:39)

Pool: This displays all Pool IDs for all pools defined in the d.ecs storage manager administration.

Used systems: This displays only the pools associated to the respective pool and their status. The different icons preceding the system names represent the system status.



NAS

: Okay, color green: The status of the system is OK, i.e. the system and the connected volumes can be used for the configured tasks.



NAS

: Unknown, Color grey: The current status of the system is not known. This can be a result of a currently running HealthCheck process regularly checking the availability of the systems or a system used in the pool is neither configured for storage nor for search.



NAS

: Error, Color red: The system has an error state which could be caused by a system or volume being unavailable. This also leads to a state change of the associated pool. The row with the pool ID is highlighted in red. If a system only has one associated volume and this has reached the maximum fill level, then the system and the pool are also set to error.

Used volumes: On systems addressed via a file system interface (NetApp, Filelock, NAS, Celerra, Silent Cubes, DataDomain), the configured alias names of the volumes and their status are displayed. If no alias names were defined, then the volume ID is displayed instead. The colored highlighting of the respective alias name or Volume ID represents the current volume status.



: The volume has reached the storage limit configured via the d.ecs storage manager administration and the status was set to "Full". This means, no documents can be stored on this volume any more. A search on the volume is still possible.

Note

If only one volume is defined for a system, an extension of the volume on the storage system or a change of the settings relating to the disk space (Lower storage limit (in %) / (in MB)) via the d.ecs storage manager administration has the effect that the volume again gets the status "Pending" (color black). Please remember that changes via the d.ecs storage manager administration only take effect after a restart.

If more than one volume is defined for a system, then the system automatically switches to the next volume with status **Pending** when the volume status **Full** is reached for one. The status **Pending** then changes to **Writing** with the next document written. Before the next volume with the status **Pending** is used, it is checked if a volume on the system that had the status **Full** has meanwhile got sufficient storage space again (**Lower storage limit (in %) / (in MB)**). If this is the case, this volume is set to the status **Pending**. With the next document written to secondary storage, the status may immediately change to status **Writing** again.

If all volumes have the status **Full**, this also results in an error status on the system.



: The volume currently has the status **Writing**. New documents are written to the volume according to the pool/system association. Documents can be searched on this volume. When reaching the configured storage limit, a volume is automatically set to the status **Full**.



: The volume currently has the status **Pending**. The volume is currently not used for the secondary storage of documents. Usually, a volume has the status **Pending**, if it has been newly created and was adopted in the d.ecs storage manager configuration. If the current write volume () reaches the status **Full** () , the next volume associated to the system with the status **Pending** is used for writing, i.e. the status automatically changes from **Pending** to **Writing**.

Documents written to secondary storage: The number of documents written to the secondary storage is displayed per pool, system, optionally per volume and pool-wide (total). The displayed values refer to the time range described above.

Documents written to secondary storage (error): The number of documents written to the secondary storage is displayed where errors have occurred. This is displayed per pool, system, optionally per volume and pool-wide. The displayed values refer to the time range described above.

Documents restored: The number of documents written to the secondary storage is displayed per pool, system, optionally per volume and pool-wide. The displayed values refer to the time range described above.

Documents restored (error): The number of documents retrieved from the secondary storage is displayed where errors have occurred. This is displayed per pool, system, optionally per volume and pool-wide. The displayed values refer to the time range described above.

The alias names of the volumes, the volume ids and the displayed system names are hyperlinked allowing to display additional information on a volume, a system or a pool. The displayed system information differ depending on the storage system.

Volume-Information

If you click on the alias name of a volume in the pool-overview of d.ecs storage manager or if this was not defined, configuration and status-information on the selected volume is displayed. Depending on the system to which the selected volume is associated, some additional system-specific information may also be displayed.



Informationen zu Volume: 026B08EF (NETAPP_SLE [NETAPP])

Volume ID: 026B08EF

State:	1 (Writing)
Alias:	VOL_SLE
Storage path:	\\10.100.151.100\vol_SLE\jerrn2_6
Volume path:	Vol\vol_SLE\jerrn2_6
Upper disk space limit:	10 MB
Rewritable:	True
Autosizing:	False
Lower disk space limit:	4 MB
Total MB:	890 MB
Free MB:	886 MB
Used disk space:	3 MB
Total Inodes:	27053
Used Inodes:	537
Autocorrect retention time:	False
Minimum retention time:	6 Months
Maximum retention time:	360 Months

Upper disk space limit

Used disk space

Lower disk space limit

System information

If you click on the system name of a volume in the pool-overview of d.ecs storage manager, information on the selected system is displayed. If the selected system is a system working file-system based then you will also get all volume information on the associated volumes.

Example system information of a NetApp system



System-Informationen: NETAPP_SLE [NETAPP]

Force CAS-creation: False

Verify data: True

Volume ID: 026B08EF

State:	1 (Writing)
Alias:	VOL_SLE
Storage path:	\\10.100.151.100\vol_SLE\jerrn2_6
Volume path:	Vol\vol_SLE\jerrn2_6
Upper disk space limit:	10 MB
Rewritable:	True
Autosizing:	False
Lower disk space limit:	4 MB
Total MB:	890 MB
Free MB:	886 MB
Used disk space:	3 MB
Total Inodes:	27053
Used Inodes:	537
Autocorrect retention time:	False
Minimum retention time:	6 Months
Maximum retention time:	360 Months

Upper disk space limit

Used disk space

Lower disk space limit

Example system information of a Centera system

**System information: CEN D3P IWB [CENTERA]**

Cluster Name:	Centera-eFile-IWB
Cluster ID:	b51abeb2-1dd1-11b2-b85b-d2c367ffc60a
Cluster Version:	4.2.2-4554-1-0
Cluster Replica Address:	
Cluster Capacity:	31.457.280 MB
Cluster Free Space:	10.702.663 MB
Cluster Time:	2014.07.09 09:45:13 GMT
Blobnaming:	MD5, MG
Read:	TRUE
Write:	TRUE
Delete:	TRUE
Privileged Delete:	FALSE
Eventbased Retention:	UNSUPPORTED
Monitoring:	TRUE
Deletion Logging (Reflect Clips):	TRUE
Centera Edition:	CE (Governance Edition)

The displayed information shows that the option **Eventbased retention** is not installed/enabled on the Centera system. Moreover, the option **Privileged delete** is disabled.

Example system information of a TSM system



System-Informationen: SSAM01 [TSM]

Client version:	6.2.3.1
Server name:	MHOL_TSMSEVER1_SERVER1
Server version:	6.2.1.0
Default MGMT class:	STANDARD
Domain name:	STANDARD
ID:	NODE1
Node type:	d3smanagertsm
Owner:	
Policyset name:	STANDARD
Server host:	10.100.150.32
Server type:	Windows

Management class infos

Management class name:	MGMT_EVENT1
Management class description:	
Copygroup name:	STANDARD
Copygroup destination:	DISKPOOL
Copygroup RetainInit:	EVENT
Retain min:	0
Retain version:	0

Copygroup name (Backup):	
Copygroup destination (Backup):	
Retain only versions:	0
Retain extra versions:	0

The displayed information for example shows, which TSM archive/backup-client version is installed.

Example system information of an HCP system



System-Information: HCP_6 [HCP]

Tenant information

Tenant Hostname:	mhol1.hcp6.d-velop.de
HTTP Scheme:	https
Namespace Name:	nsmhol1
Name IDNA:	nsmhol1
Versioning Enabled:	false
Search Enabled:	false
Retention Mode:	enterprise
Default Shred Value:	false
Default Index Value:	true
Default Retention Value:	0
Hash Scheme:	SHA-256
DPL:	2

Namespace information

Total Capacity:	2147483648 (Bytes)
Used Capacity:	507904 (Bytes)
SoftQuota:	85 %
Object Count:	62
Shred Object Count:	0
Shred Object Bytes:	0
Custom Metadata Object Count:	0
Custom Metadata Object Bytes:	0

Open Namespace-Interface

Pool information


You can get information on a pool clicking on the pool ID in the Pool-overview.

For the selected pool, you will see information on the number and memory consumption of the stored and retrieved documents related to the last 24 hours.

The information is displayed in two charts.

The values required for this representation (number/volume of stored/retrieved documents) are obtained from the database table constantly updated by the d.ecs storage manager at runtime.

The first chart shows the stored and retrieved documents while the second chart shows the volume (storage space consumption) of the stored/retrieved documents.



Below the two charts, there are configuration options to specify the time range for the analysis. Furthermore, you can configure, if you want to see information on the storage and search.

Startdatum:
Stunde:
Enddatum:

☒ **Zeige Auslagerungsinformationen**
☒ **Zeige Wiederherstellungsinformationen**
☒ **Zeige Minimalwerte**
☒ **Zeige Maximalwerte**

The end of the pool information optionally contains two lists for the job queue, pending jobs).

The lists show the currently running storage jobs and those storage jobs in a Wait state.

If no current storage or wait jobs are found, only the respective captions are displayed.

The path and name of the respective job file is displayed for secondary storage jobs.

For wait jobs, you additionally see the time when the job will next be addressed and continued.

Job-Warteschlange**Job Datei**

C:\d3\D3A.dok\jukebox\archiv\WORK\A0002000.1.job
 C:\d3\D3A.dok\jukebox\archiv\WORK\A0002001.1.job
 C:\d3\D3A.dok\jukebox\archiv\WORK\A0002002.1.job
 C:\d3\D3A.dok\jukebox\archiv\WORK\A0002004.T11.job
 C:\d3\D3A.dok\jukebox\archiv\WORK\A0002005.1.job

Wartende Jobs**Job Datei**

C:\d3\D3A.dok\jukebox\archiv\WORK\A0003275.T11.wait
 C:\d3\D3A.dok\jukebox\archiv\WORK\A0003492.1.wait

Wiederaufnahme

09.07.2014 11:31:40
 09.07.2014 11:32:04

1.6.5. Process overview

Having started the d.ecs storage manager web interface, the process-overview is displayed by default. This overview provides a summary over the currently running storage and retrieval processes. Moreover, you can start or stop the storage and restore process.

Note

Depending on the start parameters passed on calling the d.ecs storage manager, the secondary storage- and (or) the restore is enabled or disabled after the start.

Each of the two processes (**Storage/Restore**) works with a configurable number of sub-processes (see [Tab General](#)). If a process (sub-process) is active or inactive is represented by a green or red icon in the column **Status**. The sub-processes are furthermore numbered. In the following example, the storage process with its four sub-processes as well as the restore process with its sub-processes are active.



Prozess-Übersicht:

Status Prozess



Auslagerung



1. Auslagerung



2. Auslagerung



3. Auslagerung



4. Auslagerung

Aktuelle Nachricht

Wartet auf auszulagernde Dokumente

Schreibe Datenbankeintrag: NAS01 (NAS)

Bearbeitet Job: A0000139.T11.tempjob

Schreibe Datenbankeintrag: NAS01 (NAS)

Schreibe Datenbankeintrag: NAS01 (NAS)



Wiederherstellung



1. Wiederherstellung



2. Wiederherstellung



3. Wiederherstellung



4. Wiederherstellung



5. Wiederherstellung



6. Wiederherstellung

Wartet auf zu verarbeitende Jobs

Bearbeite Jobzeile: A0000137.T11 ARCHIV

Wartet auf Job

Wartet auf Job

Wartet auf Job

Wartet auf Job

Wartet auf Job

To shut down the storage or restore process, click on the green icon (column **Status**) of the respective parent process. A message is displayed, that you have to confirm with **Yes** to stop the respective process. The storage or restore process is shut down as soon as the running sub-processes have been completed. In the same way, the storage and restore process can be restarted again.



Prozess-Übersicht:

Status Prozess



Auslagerung



Wiederherstellung

Aktuelle Nachricht

Wird nicht ausgeführt

Wird nicht ausgeführt

For this effect, click on the respective red icon to start the storage and / or restore process again. The applied settings remain active until they are explicitly changed or until d.ecs storage manager is restarted. After the restart, the settings passed as start parameters apply.

1.6.6. Support Package

d.ecs storage manager offers from version 2.6 the possibility to create a support package via the web interface.

You can specify, which information should be added to the package. Such a package can be useful for a support case.



Support-Paket erstellen

Folgende Informationen dem Support-Paket hinzufügen:

d.ecs storage manager Konfiguration ☒

System Informationen ☒

d.3 Log-Datei (68 MB) ☒

Jukebox-Inhalt (Auflistung) ☒

decssm_jobs Inhalt ☒

Job-Typ: ☐ ERROR ☐ STORE ☐ RETENTION ☐ AUTORETENTION ☐ DELETE ☐ PRIVDELETE ☐ COPY ☐ JOBINPRO

Max. Datensätze:

Dokument-Referenz Informationen ☒

d.ecs storage manager configuration: If this option is enabled, the entire configuration of d.ecs storage manager instance is added to the package

System information: If this option is enabled, the detailed information (e.g. used / available memory, installed program files of the main d.ecs storage manager incl. version information etc.) are added to the package.

d.3 logfile: If this option is enabled, the d.3 logfile is added to the package. This is only possible, if the d.3 log server is running on the same system. Else, d.ecs storage manager does not have access to the file.

Jukebox-content (list): If this option is enabled, up to three files are added to the package listing the content of the handover directories **Jukebox**, **Public** and **Archive**. If the directories **Public** and **Archive** are located under the Jukebox-directory, then only one file is created.

decssm_jobs Content: If this option is enabled, a file is added to the package listing the jobs of the selected job-types. If no job-type is selected, all jobs are listed. The information **Max. data** specifies the maximum number of records to be written to the file.

Document-reference information: If this option is enabled, a file is added to the package listing all reference records for the specified document IDs. Document IDs must be specified without the archive index and be separated by a comma.

1.6.7. System check

The **System check** allows you to check the data of an entire secondary storage system for the complete secondary storage time range (or a part of it) and to create a test report. You can select various options influencing the check procedure. The options available in this context are described below.



Systemüberprüfung:

Zu überprüfendes System: NAS1 (NAS) ▼

Prüfe Wiederherstellung: ☐

Retention-Zeit prüfen: ☐

Retention-Zeit automatisch anpassen: ☐

Startdatum: 2019-12-05 Auswählen

Enddatum: 2019-12-05 Auswählen

Abbruch nach x Fehlern: ☐

Maximale Fehleranzahl: 1000

Parallele Arbeiter: 3

Dokumente pro Arbeiter: 1000

Ausführungszeiten:

☒ Mo	☒ Di	☒ Mi	☒ Do	☒ Fr	☐ Sa	☐ So	00:00	-	23:59
☐ Mo	☐ Di	☐ Mi	☐ Do	☐ Fr	☒ Sa	☒ So	00:00	-	23:59
☐ Mo	☐ Di	☐ Mi	☐ Do	☐ Fr	☐ Sa	☐ So	00:00	-	00:00

Drosselzeiten:

☒ Mo	☒ Di	☒ Mi	☒ Do	☒ Fr	☐ Sa	☐ So	08:00	-	17:00
☐ Mo	☐ Di	☐ Mi	☐ Do	☐ Fr	☐ Sa	☐ So	00:00	-	00:00
☐ Mo	☐ Di	☐ Mi	☐ Do	☐ Fr	☐ Sa	☐ So	00:00	-	00:00

Drosselpausenzeit: 1000 (ms)

Sicherheitscode: (1792606100)

Überprüfung starten

Verfügbare Prüfprotokolle:

CHECKREPORT-20191205-095241.htm

System to check: Specifies the system to be validated.

Check recalling: By default only an check for existence is performed for the documents on the selected storage system. If this option is enabled, the document is temporarily restored from the storage system and its consistency is validated using the stored hash value (if available).

Verify retention period: This option is only available for systems supporting retention periods for which the d.ecs storage manager does not work event-based. If this option is enabled, the retention period set on the storage system for a document is compared to that from the database.

Automatically correct retention period: This option is only available, if the parameter **Verify retention period** is enabled. If this option is enabled, then it is automatically attempted to adjust the retention

period, if this does not match the retention period stored in the database. Incorrectly set retention periods can occur as results of incorrect configuration etc.

Start date/End date: Date range for which the documents are to be verified in the format **YYYY-MM-DD**. In case that a complete system is to be checked, the start and stop dates must be set to "0". This causes the area to be checked to be determined automatically using the reference information from the database.

Cancel after x errors: If this parameter is enabled, the process aborts as soon as it has reached the number of errors configured under **Maximum error count**.

Maximum error count: Specifies the number of errors from which the process should abort with the option **Cancel after x errors** enabled.

Execution times: The times at which the check process is executed are defined here. If the check process is not currently running, no new workers are started - but the remaining workers finish their work.

Throttle times / throttle pause time: Here time ranges can be defined, when the check process should be slowed down. Then, the process pauses after each checked document. The duration of the break can be specified in milliseconds under **Throttle pause time**.

Parallel workers: Specifies how many threads should be initially started for the check (the number can be adjusted during execution). Please note that each thread establishes its own connection to the database. It must be ensured that the database allows a sufficient number of connections.

Documents per worker: Specifies the maximum number of documents to be processed per worker.

Security code: Specify the displayed security code here.

Start verification: Starts the verification. The button is only enabled after the security code has been entered correctly.

Available logs: This displays all logs of previous system checks. These can be opened and viewed with a single click on the respective link.

As soon as the check is started, the view changes to the following statistics overview.

Systemüberprüfung:

Informationen

Startdatum: 2019-11-25
 Enddatum: 2019-12-05
 Dokumente: 61
 Geprüft: 0 (0%)
 Aufgetretene Fehler: 0

Ausführungszeiten

	Mo	Di	Mi	Do	Fr	Sa	So
00:00 - 23:59	✓	✓	✓	✓	✓		
00:00 - 23:59						✓	✓
00:00 - 00:00							

Drosselzeiten

	Mo	Di	Mi	Do	Fr	Sa	So
08:00 - 17:00	✓	✓	✓	✓	✓		
00:00 - 00:00							
00:00 - 00:00							

NAS1: Systemüberprüfung wird ausgeführt

Status: Gedrosselt

Parallele Arbeiter: 3

ID	Startreferenz	Endreferenz	Gesamt	Geprüft	Fehler	Aktuelles Dokument
1	1	61	61	40	0	CHECKREPORT-20191203-140227.htm (1)

Verfügbare Prüfprotokolle:

CHECKREPORT-20191205-095241.htm

The verification process can be aborted here with **Cancel**. Furthermore there is the possibility to adjust the number of worker processes to be used. This view is automatically updated at a certain interval.

1.6.8. System information

The hyperlink **System-information** provides you with general information on the d.ecs storage manager installation. You will see the version of the currently running d.ecs storage manager, the date and time

of the last start and the runtime expired since then. Moreover, the current memory consumption of the d.ecs storage manager and some important settings are displayed. This is essential information for the support case.



System-Informationen:

Software:	d.ecs storage manager
Version:	3.3.8.8
Hostname:	TEST-D3-SRV
Host-IP:	10.95.6.110
Virtualisierung:	VMware, Inc. - VMware Virtual Platform
Datenbank (DSN):	D3P (Microsoft SQL Server 12.00.4213 - sqlncli11.dll 11.00.2100)
Jukebox-Pfad:	D:\d3\D3P.dok\Jukebox\
Public-Pfad:	D:\d3\D3P.dok\Jukebox\public\
Archiv-Pfad:	D:\d3\D3P.dok\Jukebox\archiv\
Single Instance Store:	False
StartUp:	1/19/2018 11:32:00 PM
Laufzeit:	19 Day(s), 11:23:39
Speicherverbrauch:	69,956 MB
System:	Windows Server 2012 R2 Professional (6.3.9600.17415)
Konfigurationsdatei:	C:\decs\decsstoragemanager\configs\Config.ini

1.6.9. System mirroring on document level

The menu option **System synchronization** in the [Web interface](#) of d.ecs storage manager allows you to perform a synchronization of the documents from one system to another one. Use this synchronization to synchronize the difference between two systems or to prepare the migration to another system.

The following settings can be applied in the dialog:

 Ziel-/Quelle

Pool

P ▼

☐ **Ganzen Pool synchronisieren**

☐ **Alle Dokumente synchronisieren**

Quellsystem

NAS (Nur Dokumente ohne Signature ▼)

Zielsystem

CELERRA (Alle Dokumente) ▼

Startdatum

1.1.1970



Enddatum

16.11.2017



Pool: This specifies the pool to be mirrored.

Source system: This specifies the system used as a document source.


Synchronize entire pool: If this parameter is enabled, then the selected source system is synchronized with all systems of the same pool. This only applies to target systems for which the secondary storage is enabled. Moreover, the secondary storage of the documents depends on the settings of the systems regarding the document types (see chapter "Add system") to be stored on the respective system.

Target system: This is the system to be used as the document destination. Which documents are synchronized on the target system depends on the settings of the target system regarding the storage of the document types (see chapter "Add system"). The destination system cannot be selected, if the option **Synchronize entire pool** is enabled.

Synchronize all documents: If this option is enabled, the entire time period of the source system is synchronized.

Start date: The start date defines from which time the documents are to be stored. If this has the value "0", the d.ecs storage manager determines the earliest storage date and uses it as the start date.

End date: The end date defines up to which time the documents are to be stored. If this has the value "0", the d.ecs storage manager determines the latest storage date and uses it as the.



Synchronisierungstyp

Synchronisierungstyp

Synchronisieren ▼

☐ Physikalische Überprüfung der Dokumente

☒ CAS-Container übernehmen


Synchronization type: The possible synchronization types are described [here](#).

Physical verification of the documents: This option is only available for the synchronization types **Synchronize** and **Check only**.

Adopt CAS container: This option specifies, if existing CAS containers are to be adopted or dissolved.

Note

It is recommended to always enable this option when using d.ecs storage manager in an ecspand context.



Auslastung

Anzahl Threads

1

Dokumente pro Thread

5000

Thread count: Specifies the number of processed to be used for the parallel Synchronization.

Note

- To prevent the Synchronization from occupying all restore processes, the maximum number of threads is limited to 2/3 of the configured job threads.
- The higher the number of threads is configured, the more load the Synchronization creates on the application-server, on the network and on the storage systems.
- At runtime, the number of threads cannot be changed. To do so, the synchronization must be aborted and restarted.
- Each process (thread) creates three connections to the database. The maximum number of connections to the database must be high enough so that the processes are not forced to abort the synchronization due to database errors.
- Each process restores up to 100 document per restore-job from the source system. The more processes are used, the more space is required for restoring the documents. Make sure that the Jukebox-directories have sufficient disk space for the synchronization.

Documents per thread: Specifies the maximum number of documents to be processed per thread per cycle.

24

Laufzeiten

Wochentag	Startzeit	Endzeit	
Jeden Tag	08:00	17:00	+
Montag	00:00	23:59	
Dienstag	00:00	23:59	
Mittwoch	00:00	23:59	
Donnerstag	00:00	23:59	
Freitag	00:00	23:59	
Samstag	00:00	23:59	
Sonntag	00:00	23:59	

Weekday: Enter the day or the days when the synchronization is to be performed.

Start time: Enter the time when the synchronization is to be performed on the specified day.

Stop time: Enter the time when the synchronization is to be paused again on the specified day. You can specify a time beyond midnight.

Gedrosselte Zeiten

Gedrosselte Zeit [ms]

1000

Wochentag	Startzeit	Endzeit	
Jeden Tag	08:00	17:00	+

Throttled time [ms]: At the configured times the document synchronization is slowed down. This reduces the load of the d.ecs storage manager for faster processing of the queries at times of high load. The time can be up to 1000 milliseconds. If a higher interval is entered then it is automatically set to 1000.

Weekday: Enter the day or the days when the synchronization is to be slowed down.

Start time: Enter the time when the synchronization is to be slowed down on the specified day.

Stop time: Enter the time when the synchronization is not to be slowed down on the specified day. You can specify a time beyond midnight.

 Fehlerbehandlung

☐ Nur fehlerfreie Dokumente synchronisieren

☒ Abbruch nach x Fehlern

Maximale Fehleranzahl (Logging)

Maximal Fehleranzahl

Synchronize only documents without errors: If you enable this option, documents identified as erroneous during the synchronization check are ignored. These documents are accordingly **not** synchronized. The number of documents that have been skipped per error type can then be viewed in the synchronization report.

Maximum error count (logging): Specifies the maximum number of error to be written to the synchronization report.

Cancel after x errors: If this option is enabled and the number defined under **Maximum error count** is reached, the synchronization is aborted.

Maximum error count: Specifies the maximum number of errors which may occur during the synchronization, before the synchronization is aborted automatically (depending on the option **Cancel after x errors**)

Note

You need a separate license to execute the synchronization (d.ecs storage manager synchronization). If this license does not exist, then the configuration dialog for the synchronization is not displayed.


To start the synchronization, the processes **Store** and **Retrieve** must be active.

Please keep in mind that a large number of retrieval-jobs with large files may take some time to complete. Optionally, set the time for the maximum document retention time (tab **Document cache**) to a higher value to prevent the premature deletion of the retrieved files by the housekeeping process of the d.ecs storage manager.

Synchronization types

The system synchronization of the d.ecs storage manager can be performed in three different ways.


1. Logical synchronization (logical test)



The screenshot shows a settings panel titled "Synchronisierungstyp" with a refresh icon. Below the title, there is a dropdown menu labeled "Synchronisierungstyp" with the option "Synchronisieren" selected. To the right of the dropdown is a radio button labeled "Physikalische Überprüfung der Dokumente", which is currently unselected.

With the logical synchronization only the reference table is checked for missing document entries and documents are duplicated if required.


2. Physical synchronization (logical as well as physical test)



The screenshot shows the same settings panel as above, but the radio button "Physikalische Überprüfung der Dokumente" is now selected, indicated by a blue checkmark.

During the physical synchronization, the reference table is checked for missing documents entries and documents are duplicated if required. Additionally, the accessibility of the document is verified which are located on the storage according to the reference table. Should any document be inaccessible in this process, then the reference entry is removed and the document is copied to the storage again.


3. Synchronization test (logical test) (nothing is mirrored)



The screenshot shows the settings panel with the dropdown menu "Synchronisierungstyp" set to "Nur Überprüfen". The radio button "Physikalische Überprüfung der Dokumente" remains unselected.

With a synchronization test, the reference table is checked for missing document entries. The synchronization test only performs a verification but no synchronization. This only writes a log file containing information about what would be mirrored, if the synchronization was executed in the **Non-test mode**.

4. Synchronization test (logical and physical test) (nothing is mirrored)



 Synchronisierungstyp


Synchronisierungstyp

Nur Überprüfen ▼

 **Physikalische Überprüfung der Dokumente**

With a synchronization test, the reference table is checked for missing document entries and an existence check is applied on the storage system. The synchronization test only performs a verification but no synchronization. This only writes a log file containing information about what would be mirrored, if the synchronization was executed in the **Non-test mode**.

5. Check completeness only (logical, database-based verification)



 Synchronisierungstyp

Synchronisierungstyp

Nur Vollständigkeitsprüfung ▼


During the completeness check, it is only verified with a database query, if the information for the specified systems is consistent. No physical verification is performed on the storage system and no physical synchronization is performed.

Synchronization overview

Synchronisierung

⏸ Pausieren

✖ Abbrechen



Progress bar: The progress of the synchronization is displayed as a progress bar.

Total This specifies the overall number of documents checked for the specified time slot.

Checked: This specifies how many documents were synchronized so far.

Synchronized: This specifies how many documents were synchronized so far.

Erroneous: Indicates the number of documents for which synchronization problems have occurred so far.

Runtime: Shows how long the current synchronization process has been running so far.

Expected time remaining: Shows how the estimated time that the current synchronization process will be running. This value is only an estimate and may change during the execution of the synchronization.

Current reference scope: Specifies the reference scope that has last been passed to one of the worker processes for processing.

Status message: Shows the status of the synchronization:

⚙️ Einstellungen
🔗 NAS1 → NAS10
👤 P
📅 1. Jan. 1970 → 25. Okt. 2019
🕒 25. Okt. 2019, 13:52:05
🔔 Nur Überprüfung
✖ Synchronisations-Wiederaufnahme
✓ CAS-Container übernehmen

Source and destination: Specifies the source and target systems for synchronization.

Pool: Specifies the pool ID for which a synchronization is performed.

Synchronization period: Specifies the time period in which the synchronization takes place.

Start time: Specifies, when the synchronization was started.

Synchronization type: Specifies the mode in which the synchronization is performed.

Synchronization retry: Shows whether the synchronization is an initial run or has been resumed.

Adopt CAS container: Shows whether CAS containers are transferred during synchronization.

📶 Auslastung				
Anzahl Threads  4				
ID	Überprüft	Synchronisiert	Fehlerhaft	Fortschritt
1	85		0	

The large progress bar specifies, how much of the overall verification and synchronization has been finished.

For every worker process, a separate progress bar is displayed showing the progress of the currently processed part. While a worker process checks its part to be processed, the percentage on the progress bar remains 0%. As soon as the worker-process starts synchronizing data, the percentage changes according to the current progress. The individual progress bars of the worker processes show the reference scope currently processed by the worker process as a tooltip.

During the synchronization process, you can adjust the number of worker processes, for example, to react to changes in the workload without having to cancel the synchronization process. If the number of worker processes is reduced, running worker processes are completely completed and terminated, new worker processes are only started if the configured number is not exceeded. An increase in the number has a direct effect.

Note

The systems Governikus LZA and d.velop cloud storage can currently not be selected for a synchronization process.

1.6.10. License information

This page displays information about the licenses currently in use.

The area **General** provides basic information about the licenses.

Allgemein

Name	d.velop AG
Straße	Schildarpstraße 6-8
Stadt	48712 Gescher
Seriennummer	
Hostname	d-velop-license
Lizenzbibliothek	C:\d3\d.ecs storage manager\d.3 storage manager\liccheck.dll

On the one hand, there is information about the owner of the license, such as **Name**, **Street**, **City** and **Serial number**.

Furthermore, the currently used **Host name** is listed, to which the application has connected, and at **License library** which library was loaded by d.ecs storage manager for connection.

The area **Available licenses** lists all licenses which are available for d.ecs storage manager. A range of detailed information is displayed for each of these licenses.

Verfügbare Lizenzen

ID	Name	Status	Ablaufdatum	Verwendete Menge	Gesamtmenge
decs-storage-manager-NAS-prod	d.ecs storage manager NAS (productive-environment)	Lizensiert	12/31/9999	25	9999
decs-storage-manager-isilon	d.ecs storage manager Isilon	Lizensiert	12/31/9999	0	9999
d.3 storage manager Celerra	d.ecs storage manager Celerra	Lizensiert	12/31/9999	0	9999
d.3 storage manager Centera	d.ecs storage manager Centera	Lizensiert	12/31/9999	0	9999
d.3 storage manager FileLock	d.ecs storage manager FileLock	Lizensiert	12/31/9999	1	9999
d.3 storage manager HCP	d.ecs storage manager HCP	Lizensiert	12/31/9999	0	9999
d.3 storage manager NetApp	d.ecs storage manager NetApp	Lizensiert	12/31/9999	0	9999
d.3 storage manager Silent Cubes	d.ecs storage manager Silent Cubes	Lizensiert	12/31/9999	33	9999
d.3 storage manager TSM	d.ecs storage manager TSM	Lizensiert	12/31/9999	1	9999
d.connect easy	d.connect EASY	Lizensiert	12/31/9999	0	1
d-connect-igstore	-	Nicht lizensiert	-	-	-
d.ecs storage manager Data Domain	d.ecs storage manager Data Domain	Lizensiert	12/31/9999	0	1
d.ecs storage manager governikus lza	d.ecs storage manager Governikus LZA	Lizensiert	12/31/9999	0	9999
d.ecs storage manager iCAS	d.ecs storage manager iCAS	Lizensiert	12/31/9999	0	1
d.ecs storage manager synchronization	d.ecs storage manager synchronization	Lizensiert	12/31/9999	0	1
d.ecs storage manager system connect	d.ecs storage manager system connect	Lizensiert	12/31/9999	0	9999
Privileged Delete	d.3 privileged delete	Lizensiert	12/31/9999	0	1

The unique ID and, if available, the display name entered in the license is listed for each license.

The **Status** of a license is either **Licensed** or **Not licensed** - unlicensed entries are indicated by a colored highlighting.

For licensed entries, the expiry date and the quantities used and available are listed.

In case of full or over licensing an entry is marked by an error color. If the booking of a license reaches a usage of 90% of the total quantity, an entry is marked by a warning color.

The area **License overview** lists all currently executed license bookings. Several values per booking can also be viewed here.

Lizenzübersicht

ID	Systemname	Volume	Größe [Bytes]	Anzahl Dateien	Ablaufdatum
1	NAS1	01833CFC	422912	413	-
5	DATADOMAIN1	0228D44E	3072	3	-

ID, **System name** and **Volume** allow an entry to be assigned to a configured module or volume.

Size [bytes] indicates the amount of space currently booked for an entry, while **Number of files** gives information about how many files are responsible for that storage space.

An entry contains a value for **Expiry date** if a license could not be booked.

1.7. Additional hints

1.7.1. Jukebox directories

Do not copy any important data to the **Jukebox** directories manually. d.ecs storage manager contains several processes automatically deleting no longer required data from the **Jukebox** directories.

Note

The term jukebox in this context generally refers to secondary storage systems.

The directories used for this purpose can be configured in the d.3 admin/ecspand.

For additional information, please consult the manual for the d.3 admin/ecspand.

1.7.2. Multiple instances of d.ecs storage manager

Warning

You can start several instances of the d.ecs storage manager. However, you must make sure that two d.ecs storage managers must never manage the same **Jukebox** directories since this would negatively affect the functionality.

Furthermore, you should define separate configuration directories for each d.ecs storage manager instance.

As a precaution, starting the d.ecs storage manager more than once with the same configuration file is rejected. This rejection is based on the name of the configuration file. The content of the configuration file is not checked.

1.8. Deleting systems

The deletion of a system can have various reasons.

This could, for example, be a change of the storage-system, a cancellation of the storage space occupied by a client etc.

Having marked a system for deletion in the d.ecs storage manager administration, the program **decssm-delsys.exe** must be called.

This tool logs all references of the system to be deleted in a logfile. On exceeding a size of 50 MB a new log-file is created. The log-files are then compressed and written to the new storage system. Then all system references of the respective system are deleted from the database. The deletion is performed in blocks of 1000 records.

The program must be started as follows:

```
decssmdelsys.exe <Path to the configuration file of the d.ecs storage manager>
```

Example:

```
decssmdelsys.exe c:\decs\decsstoragemanager\configs\D3A.ini
```

Note

If the execution of the program **decssmdelsys.exe** is aborted, be it manually or through an interrupted network connection, then the process can be continued at the same point again with a later restart.


The program only deletes the reference entries of a system from the database. No data is removed from the storage system.

The deletion of data from the storage system must either be triggered before the system deletion from d.3/ecspand and then be executed by the d.ecs storage manager or via the administrative interface of the storage system if this provides a delete function.

Note that deleting a large number of records may result in a large number of entries in the log files of the database used.

1.9. d.ecs storage manager Status overview

The status overview is displayed after a direct program start but not, if it is started using the d.3 process manager (see [Opening d.ecs storage manager](#)).



Secondary storage: This displays, if the storage process is started or stopped.

Restore: This displays, if the retrieval processes are started or stopped.

Runtime: This displays the d.ecs storage manager runtime.

Stored (xxx (yyy)): This displays, how many documents were stored (xxx) and how many are queued for storage (yyy).

Retrieved (xxx (yyy)). This displays, how many documents were restored (xxx) and how many are queued to be restored (yyy).

Server: This displays the server name and the port for an active [Web interface](#). A click on the link opens the web interface in the default browser.

1.10. Monitoring

The monitoring allows you to monitor several relevant metrics of the d.ecs storage manager. The determined metrics are sent to the d.ecs monitor.

The following metrics are sent to the d.ecs monitor:

- [Workload](#)
- [Speeds](#)
- [Job status](#)
- [Pool status](#)

d.ecs storage manager configures predefined alert thresholds in d.ecs monitor. Using these alert thresholds, notifications can be sent as soon as one of the metrics exceeds a critical or alarming value.

Apart from the metrics, the overall status of d.ecs storage manager is also sent to the d.ecs monitor.

The monitoring function requires a d.ecs monitor agent on the system with the d.ecs storage manager.

1.10.1. Workload

Workload of retrieval processes

To determine the workload of the retrieval processes, the current utilization is transmitted every ten seconds. Every two minutes, the average is sent to the d.ecs monitor.

If the workload of the retrieval processes exceeds 80%, d.ecs monitor sends an alert. With a workload of more than 95%, an error is sent.

Workload of secondary storage processes

To determine the workload of the storage processes, the current utilization is determined every minute. Every 30 minutes, the average is sent to the d.ecs monitor.

If d.ecs monitor receives a value exceeding 95%, an alert is sent. By default, no threshold exists for an error message.

1.10.2. Speeds

d.ecs storage manager send information on various speeds and timings to the d.ecs monitor. Among these speeds are the secondary storage and retrieval speed as well as the storage and retrieval times.

Every minute, the average speed and times of the last 100 storage transactions and retrievals is sent to the d.ecs monitor.

The secondary storage and retrieval times are specified in megabytes per second.

The secondary storage and retrieval times are specified in milliseconds.

These values are determined for each system and volume and then sent to the d.ecs monitor.

Note

d.ecs storage manager sends information on various speeds and timings to the d.ecs monitor.

1.10.3. Job status

d.ecs storage manager sends information on the different job types in a five minute interval.

Storage jobs:

- Alert from: 1000 jobs
- Error from: 5000 jobs

Retention jobs:

- Alert from: 5000 jobs
- Error from: 10000 jobs

Successfully processed storage jobs:

- Alert from: 5000 jobs
- Error from: 10000 jobs

Erroneous storage jobs:

- Alert from: 1 Job
- Error from: 50 jobs

Deletion jobs:

- Alert from: 1000 jobs
- Error from: 5000 jobs

Privileged deletion jobs:

- Alert from: 1000 jobs
- Error from: 5000 jobs

1.10.4. Pool status

The pool status is determined from the erroneous retrieved documents in the last 24 hours. The value is sent to the d.ecs monitor every hour.

If the metrics exceed a value of 50, d.ecs monitor sends a warning. An error is returned after 100 erroneous document restore jobs.

1.11. Appendix

1.11.1. Troubleshooting

d.ecs storage manager does not provide (all) documents

Error state:

d.ecs storage manager does not provide (all) documents.

Notes:

- Check in the web interface, if the retrieval process of the d.ecs storage managers is active. If not, activate it.
- Check in the web interface, if all retrieval processes are busy.
- Check if the d.3 gateway creates a restore job in the Jukebox directory.
- Check, if the ecspan service for the access to the d.ecs storage manager is correctly configured and running.
- Check the connection of the d.ecs storage manager to the database. Usually, the d.ecs storage manager re-establishes its connection after an interruption. Should this not be the case, then start the d.ecs storage manager again in the d.3 process manager.
- Check the provision of the documents with the test function of the web-interface of the d.ecs storage manager (option: **Document information**).

- Use the option **System status** in the web-interface of the d.ecs storage manager, if the storage-system encounters any connection problems.

d.ecs storage manager does not store any documents

Error state:

d.ecs storage manager does not store any documents.

Notes:

- Use the web interface of the d.ecs storage manager to check, if the storage process of the d.ecs storage manager is active. If not, activate it.
- Use the web interface of the d.ecs storage manager to check, if possibly all storage processes are in use.
- Check the connection of the d.ecs storage manager to the database. Usually, the d.ecs storage manager re-establishes its connection after an interruption. Should this not be the case, then start the d.ecs storage manager again in the d.3 process manager.
- Validate the settings of the d.3 server and adjust the settings as described under d.3 server preparations (d.3). Check, if the ecspand service for the access to the d.ecs storage manager is correctly configured and running (ecspand).
- Check the option **Pool Information** in the d.ecs storage manager web interface to find out, if the pool was disabled due to a faulty system.

OK-files remain and are not processed

Error state:

OK-files remain in `jukebox\public` or `jukebox\archiv` and are not processed (only d.3).

Notes:

- **OK** files are job receipts of d.ecs storage manager which signal the d.3 server that a certain document was successfully stored. These files are further processed by the d.3 async.
1. Check if d.3 async is started and works properly.
 2. Check if d.3 async has read- and write access to the jukebox directories.
- The processing of OK files might be delayed, if d.3 async is currently busy with other tasks (e.g. creating dossier links) having a higher priority.
 - An additional d.3 async process can be configured which only processes the **OK**-files. For this effect, contact the d.velop support.

The d.ecs storage manager does not start

Error state:

d.ecs storage manager does not start.

Notes:

- Check if you specified the correct configuration file (as described under [Start the d.ecs storage manager](#)) and if d.ecs storage manager has reading access to the configuration files.
- Check if the messages in the d.3 logviewer indicate a missing license or if the access to the storage directories (`jukebox\*`) or to the database ways denied.
- Check if another instance of d.ecs storage manager with the same configuration file is already started. The duplicate start of d.ecs storage manager with the same configuration file is aborted.

d.ecs storage manager CAS reports that CAS Proof is not started

Error state:

d.ecs storage manager CAS reports that CAS Proof is not started.

Notes:

- Check, if the read/write permissions for the CAS mirror are set.
- Check if the servers running d.ecs storage manager and CAS Proof are time synchronized. The time difference must not exceed 4 minutes.

Windows error 112 when transferring to a NetApp volume.

The Windows error code 112 means "There is not enough space on the disk".

This error can occur although according to Windows there is still enough free space on the disk.

This error is often caused by the configured maximum number of available Inodes on the NetApp.

To fix the problem, the number of free Inodes on the NetApp must be checked and increased if necessary.

The web interface of the d.ecs storage manager cannot be started

Error state:

The web interface of the d.ecs storage manager cannot be started.

Notes:

- Check if the web interface is enabled in the d.ecs storage manager administration.
- Check the correct spelling of the password.
- Check if you enabled the SSL encryption. In this case, the web interface is called with https:// instead of http://.
- Check, if the selected port (default: 6122) is already used by another application or service. The d.3 log will report any issues on initializing the web-interface on starting d.ecs storage manager.

1.11.2. Meaning of the different file names of files in the Jukebox-directories

The list below shows the different file names that may occur in the Jukebox-directories. This list is continuously extended.

File name of the files in the Jukebox-directories:

XXXXXXXXXX.job: Retrieval job pending to be processed.

XXXXXXXXXX.tempjob: Retrieval job that is currently being processed.

XXXXXXXXXX.OK1: OK1-Receipt: This means that d.ecs storage manager has started to process the respective retrieval job.

XXXXXXXXXX.OK2: OK2-Receipt: This means that d.ecs storage manager has processed the respective retrieval job.

XXXXXXXXXX.ER1: Error-Receipt: This means that none of the requested files have been found.

XXXXXXXXXX.ER2: Error-Receipt: This means that only some of the requested files have been found.

D0000001.1: Document that has been retrieved from the storage system

D0000001.T11: Dependent document that has been retrieved from the storage system.

D0000001.1_X: Document that has been retrieved from the storage system as part of a synchronization.

D0000001.T11_X: Dependent document that has been retrieved from the storage system as part of a synchronization.

D0000001-2010-06-29091742.prot: Workflow protocol that has been retrieved from the storage system.

D0000001-2010-06-29091742.prot_X: Workflow protocol that has been retrieved from the storage system as part of a synchronization.

Meaning of the different file names of files in the directories jukebox\public / jukebox\archiv or their subdirectories:

D0000001.1: Base document transferred by d.3.

D0000001.1.job: Job file for a base document transferred by d.3.

D0000001.T11: Dependent document transferred by d.3.

D0000001.T11.job: Job-file for a dependent document transferred by d.3.

D0000001.1_X: Base document that is written to the secondary storage as part of a synchronization.

D0000001.1_X.job: Storage job for a base document that is written to the secondary storage as part of a synchronization.

D0000001.T11_X: Dependent document that is written to the secondary storage as part of a synchronization.

D0000001.T11_X.job: Storage job for a dependent document that is written to the secondary storage as part of a synchronization.

D0000001-2010-06-29091742.prot: Workflow protocol transferred by d.3.

D0000001-2010-06-29091742.prot.job: Job file for a workflow protocol transferred by d.3.

D0000001-2010-06-29091742.prot_X: Workflow protocol that is written to the secondary storage as part of a synchronization.

D0000001-2010-06-29091742.prot_X.job: Storage job for a workflow protocol that is written to the secondary storage as part of a synchronization.

D3SM_XXXXXXXX-20120829-111337.IR: d.ecs storage manager index recovery file.

D3SM_XXXXXXXX-20120829-111337.IR.job: Storage job for a d.ecs storage manager Index recovery file.

d3attrib-20130917-045150.xml: D3ATTRIB-file transferred by d.3.

d3attrib-20130917-045150.xml.job: Job file for a D3ATTRIB-file transferred by d.3.

d3attrib-20130917-045150.xml_X: D3ATTRIB-file written to the secondary storage as part of a synchronization.

d3attrib-20130917-045150.xml_X.job: Storage job for a D3ATTRIB-file written to the secondary storage as part of a synchronization.

Job files may contain a leading "!" in the file name. These jobs have either already been started once (and have gone to WAIT mode in the meantime etc.) or come from d.ecs storage manager itself (e.g. jobs for IR-files) and are thus processed with a higher priority.

1.11.3. Optimize the network access

Since the SMB2 Client Redirector Cache was introduced in SMB 2.0 (Windows-network layer), caching-issues can occur, if d.ecs storage manager is accessing the Jukebox-directories (e.g. in the d.3 document tree) via the network (UNC-share).

It can happen that files do not exist on the share according to Microsoft Windows, although they are visible on the server itself. As a result, STORE jobs may encounter the error "Document not found" during d.ecs storage manager secondary storage, because d.ecs storage manager expects a specific file in the directory but this is currently reported as non existent. This is due to the SMB2 Client Redirector Cache which was introduced with Microsoft Windows Vista / Microsoft Windows Server 2008.

This cache is by default only updated every 10 seconds which has the effect that the visible folder content is not always up-to-date. Since d.ecs storage manager is relying on the correct representation of the folder content, a specific parameter must be set in the registry.

To open the registry editor, select Start > Run and execute the command regedit. The parameter **DirectoryCacheLifetime** to be set is located under HKEY_LOCAL_MACHINE\System\CurrentControl-Set\Services\Lanmanworkstation\Parameters Set the parameter to the value "0" to disable this cache. It may happen that the key **DirectoryCacheLifetime** does not exist. This must then be created manually as a **DWORD-value (32-bit)**.

Having set the parameter in the registry, you must restart the machine for the change to take effect.

Detailed information on this subject by Microsoft can be found under the subject "SMB2 Client Redirector Caches Explained".

1.11.4. Create SSL-certificates with OpenSSL create

d.ecs storage manager allows you to encrypt the web-interface and the API via SSL.

This requires a valid SSL certificate. To create such a certificate, the sub-directory **SSLCertificate** in the installation-directory of d.ecs storage manager contains two batch files using OpenSSL to create such a certificate. To do so, follow the steps below:

- Start the file **create_key.bat** to create an SSL-certificate (the files **server.crt** and **server.key** are created). This certificate is already sufficient to enable an SSL encryption in d.ecs storage manager. Since the certificate, however, is not validated and marked as "trusted" by an official verification authority , Internet browsers may display respective warning messages.
- To have the certificate validated and signed by an official verification authority, start the file **create_csr.bat** and thus create the files required for this (**request.csr**). The official verification authority requires these files to create an officially signed certificate.

Warning

Consider the following when choosing the host name: The SSL certificate is only created for one specific host name (in this case this must be the host name of the machine on which the d.ecs storage manager is running). If the host name changes, the SSL certificate must be recreated and signed by the verification authority accordingly.

Please understand that this manual does not further explain SSL certificates and OpenSSL. For additional information refer to <http://www.openssl.org>.

1.11.5. Retention periods/Retentiontimes

Writing documents to the secondary storage, the leading system (d.3 / ecspand) passes retention periods for the documents to the d.ecs storage manager to write them to the storage system.

Please also consider that the d.3 system can store d.3 property and OCR files in addition to the documents. In current d.3 versions, you can set the retention period for these files via the configuration parameter

- **RETENTION_FOR_XML_ATTRIBUTE_FILES** and
- **RETENTION_FOR_XML_OCR_FILES**

d3addon.ini.

The retention time in these two configuration parameters is specified in months. If the retention period is not configured via these parameters, the d.3 property and OCR-files are by default passed to the d.ecs storage manager with a retention period of 50 years (600 months).

d.ecs storage manager itself regularly writes so-called database logs (see d.ecs storage manager administration ([Tab Database logging](#))). The retention period for these database logs could be influenced in d.ecs storage manager version 2.5 with a configuration parameter. The parameter had to be specified in the configuration file of the respective d.ecs storage manager instance (INI-file). The respective entry was:

```
[Special]
RetentionForIRFiles=50
```

The value for the retention period had to be specified in years.

With d.ecs storage manager version 2.6, the retention period can be configured via the d.ecs storage manager administration ([Tab Database logging](#)). Additionally, the d.ecs storage manager version 2.6 allows to define the retention period for synchronization and system check logs in the section Miscellaneous.

The passed retention periods must be allowed by the storage system, i.e. it is essential to know the current settings of the storage system regarding the minimum, maximum and default retention period or to have them configured on the storage system so that they cover the full range of retention periods with which the documents and additional files such as d.3 property files and OCR-files could be passed from the leading system.

The different storage systems have different names for the minimum, maximum and default retention period.

Netapp/SnaplockNetapp

- **snaplock minimum period:** Minimum
- **snaplock maximum period:** Maximum
- **snaplock default period:** Standard (Default)

EMC Celerra/VNX with FLR

- **Minimum Retention Period:** Minimum
- **Maximum Retention Period:** Maximum
- **Default Retention Period:** Standard (Default)

EMC Data Domain with RetentionlockData

- **Minimum Retention Period:** Minimum
- **Maximum Retention Period:** Maximum

Grau Data Filelock

- **MINIMUM RETENTION:** Minimum
- **MAXIMUM RETENTION:** Maximum
- **DEFAULT RETENTION:** Standard (Default)

Moreover, the different storage systems behave differently when documents with a retention period are passed with a retention period outside the range configured by the minimum and maximum retention

time. Usually, the minimum and maximum retention time are set by the storage system, if a document with a lower or higher retention period is to be stored. However, it may also happen that a storage system (e.g. Data Domain with Retentionlock) rejects the document storage, if the passed retention period is outside the configured limits.

The file system-based storage modules (NetApp, Celerra, FileLock, Silent Cubes, Data Domain), a time range for the retention period can be configured in the d.ecs storage manager administration (version 2.6) (**Minimum/maximum retention period in months**) in which the retention period passed by the leading system (d.3/ecspand) must be specified.

The settings regarding **Minimum retention period (in months)** and **Maximum retention period (in months)** should be equivalent to the settings applied on the storage system. The, again, should match the range of retention periods configured in the document type settings (d.3/ecspand).

Allowed values for the **Minimum retention period (in months)** must be between one and twelve months. The **Maximum retention time (in months)** can be between 120 and 600 months.

Note

The allowed values configurable for the maximum retention time may deviate on different storage systems. Thus, it may happen that a d.3/ecspand retention period of 100 years for a document cannot be implemented on the storage system. In such a case, the maximum retention time on the storage system should be set to the highest possible value. In the d.ecs storage manager administration, configure the maximum possible value of 600 months under **Maximum retention period (in months)**. Additionally, enable the option **Adjust retention period automatically**.

Note

If the system to be configured is a NetApp-system with SnapLock and you have configured the API connection to the NetApp system via the d.ecs storage manager administration, then the button **Load times** is displayed as illustrated above.


This button allows you to determine the current settings (**Snaplock minimum/maximum period**) from the NetApp-system for the volume to be configured and then automatically adopt these settings in the d.ecs storage manager configuration dialog. The **SnapLock minimum period** on a NetApp-volume is often set to the value 0 (zero) years. This value is not allowed in the d.ecs storage manager as the documents are at least stored for one month. The retrieved value of the NetApp volumes is thus set to the minimum value of 1 (one) by the d.ecs storage manager administration.

Note

The button **Load times** is only displayed when configuring a NetApp system as only a NetApp system can establish an API connection (http/https). The other file-system-based storage systems do not allow an automatic retrieval of the settings. For the configuration, you must read the values directly from the configuration application of the storage system or request them from the customer or computer center. Keep in mind that these values may also have to be changed on the storage system, if they contradict the retention periods of the leading system (d.3/ecspand).

If a document is passed with a retention period outside the range defined by **Minimum retention period (in months)** and **Maximum retention period (in months)**, the document storage on the storage system is rejected by the respective d.ecs storage manager module. The respective storage jobs thus encounter errors. With the configuration parameter **Automatically correct retention period**, this behavior can be changed so that a passed retention period outside the configured range is automatically set to the values configured under **Minimum/Maximum retention period (in months)**.

Before the parameter **Automatically correct retention period** is enabled, the following message must be confirmed with **Yes**.



With **Yes** you confirm and agree that the retention times passed by the leading system (d.3/ecspand) during the document storage or extension of retention times are automatically corrected by the d.ecs storage manager and that these adjusted retention periods are used for the secondary storage. Furthermore, by selecting yes you confirm that the settings you applied for the Minimum retention time (in months) and the Maximum retention time (in months) should be used for the automatic correction of the retention time.

Note

If possible, you should adjust the retention periods in the document type definition of the leading system (d.3/ecspand), so that no automatic correction by the d.ecs storage manager is required. Moreover, the retention periods must be configurable accordingly on the storage system.

The value for the retention period for a document passed by the leading system, is written to the DB-column TTL of the table **D3SM_REFERENCES**. The TTL value in the database is only changed by the d.ecs storage manager only via a so-called “Retention-job” from the leading system. Additionally, the date on which the document was actually written to the secondary storage is stored in the column **RETENTION_ON_STORAGE**.

Extending the retention period via retention-jobs

The d.3 system sends so called retention-jobs to the d.ecs storage manager, if the retention period for a document changes in d.3.

This can, for example, be the case, if documents are moved to another document type. You must consider in the d.3 configuration that a storage system typically only allows to extend the retention periods of a document and not to reduce them.

A retention-job contains a retention period for each document to be passed to the storage system by the d.ecs storage manager. This retention period must be between the limits (**Minimum retention period (in months)**, **Maximum retention period (in months)**) defined for the storage system for the retention period to be adjusted on the s storage system.

If the retention period passed with the retention-job is outside the limits configured in the d.ecs storage manager system configuration, the retention-job will encounter an error unless the parameter **Automatically correct retention period** was enabled.

If the retention period passed by the d.3 system is within the configured limits, the retention period for the respective document is extended on the storage system. Having successfully performed the extension, the values for **TTL** and **RETENTION_ON_STORAGE** are also adjusted in the table **D3SM_REFERENCES**.

If the passed retention period is outside the configured limits and the parameter **Automatically correct retention period** is set, then the retention period for the document on the storage system is set to the maximum possible value (today plus Maximum retention period (in month)). The value passed by d.3 for the retention period is set according to the column **TTL**. The adjusted (limited) value for the storage system is validated with a query to the storage system and is then written to the column **RETENTION_ON_STORAGE**.

Note

For a possible migration (synchronization by the d.ecs storage manager), only the value in the column **TTL** is used to set the retention period for the documents on the new storage system. If the migration is performed using functionality provided by the storage system vendor, only those values can be used that were also written to the old storage system.

Automatic extension of retention periods by the retention process

The so called retention process is a new function of the d.ecs storage manager version 2.6. This process automatically extends the retention periods on the storage system, if the remaining retention time falls below 60 days.

The retention time is extended to prevent that the document becomes deletable with expiry of the retention period. This may be necessary, if no automatic deletion takes place in d.3 after expiry of the retention period. Another use case, where the retention period for a document has to be extended automatically affects the documents where the document type was configured with the option **event-based deletion**.

Note

The automatic extension of the retention period is not applied to already deleted documents and for documents where a delete-job is pending but has not been executed on the storage system, yet.

The retention process can be configured in the d.ecs storage manager administration (see illustration below). The value selected for **Automatically extend the retention by x months** is of special significance here.

Startzeit für den Retention Prozess								Startzeit:	Laufzeit:
	Mo:	Di:	Mi:	Do:	Fr:	Sa:	So:		
Startzeit 1:	☒	☒	☒	☐	☐	☐	☐	00:00	4
Startzeit 2:	☐	☐	☐	☐	☐	☐	☐	00:00	0
Startzeit 3:	☐	☐	☐	☐	☐	☐	☐	00:00	0

Allowed values must be between 6 and 24 months. The default value is 6 months. The basis for the extension of the retention period by the d.ecs storage manager is the value stored in the column **RETENTION_ON_STORAGE** of the table **D3SM_REFERENCES**. If no value is stored in the column **RETENTION_ON_STORAGE**, the date value from the column **TTL** is used as a basis to extend the retention period.

If the date value in the column **RETENTION_ON_STORAGE** for a document is lower than the result of the calculation of the current date plus 60 days, then the retention period is extended automatically. For this effect, the retention period of the document on the storage system is set to the date value resulting from the calculation of the current date plus the value configured for the retention process in the d.ecs storage manager administration under Extend retention time by x months.

Having applied the extension of the retention period on the storage system, the value is also adjusted in the column **RETENTION_ON_STORAGE**. The date value in the column **TTL** is not affected by this.

If no value is stored for a document in the column **RETENTION_ON_STORAGE**, the check, if the retention period has to be extended uses the value in the column **TTL**. In this case, however, the value of the retention time is determined from the storage system, before the retention period is extended.

If the retrieved retention time of the storage system is lower than the calculated date (current date plus extension configured for the retention process), the retention period is extended on the storage system. The newly set value of the retention period is updated in the column **RETENTION_ON_STORAGE**.

If the determined retention time from the storage system is higher than the calculated date, no extension of the retention period has to take place. However, the value in the column **RETENTION_ON_STORAGE** is set to the value of the document retention time on the storage system.

Note

You should possibly not configure the value for **Extend retention period by x months** too low or otherwise the retention periods for the documents have to be extended again and again.

1.12. Glossary

This documentation frequently uses phrases and terminology in the context of long-term storage, document management and the control of external storage systems.

Some of these terms are therefore briefly explained in the following chapters.

1.12.1. Store

The term secondary storage is used in connection with the transfer of documents from d.3 to the d.ecs storage manager and the writing on the external storage system by the d.ecs storage manager. Typically, a document is handed over in a configurable directory structure and a respective job is passed to the d.ecs storage manager to store the document on the external storage system. The document handed over by d.3 is deleted from the handover directory structure after it has been written to the external storage system. The respective job is afterwards used to give feedback to the d.3 system. Having received the feedback, d.3 can delete the document from its document tree, provided it has been configured accordingly.

1.12.2. Jukebox directory/Jukebox public directory/Jukebox archive directory

The term “jukebox directory” goes back to a time when typically WORM (write once read many) media were used for the storage of documents. The WORM medium itself then safeguarded the immutability of the documents. Current d.3 and ecspan versions hand the documents over to the d.ecs storage manager for secondary storage via the sub-directories **Public** and **Archive** under the Jukebox directory. The default name if the Jukebox-directory is **Jukebox** but can optionally be changed. However, then all configurations accessing the Jukebox-directory must also be adjusted.

1.12.3. References/Reference data/Reference entries/Reference table/Reference index/Index data

The terms mentioned above are related to the document information required by the d.ecs storage manager to manage the documents handed over to it. Due to a large data volume, a database table is used for the management, in which the required document information is maintained. Information such as the document ID, the document version, the size, the hash value the retention period and the date of writing to the secondary storage are maintained in this table. New entries in this database table are usually created whenever a document was successfully written to a respective storage system. Documents that do not have an entry in this database table are not provided by the d.ecs storage manager and are treated as “not referenced”.

1.12.4. Secondary storage/Secondary storage I/O

The term secondary storage or secondary storage I/O is used in connection with the d.3 configuration or the document type definition in d.3. The setting **Secondary storage I/O** in the document type definition allows to control, if a document of a specific document type should also be handed over to the d.ecs storage manager. The section **Secondary storage** in the d.3 configuration allows you define the respective settings required for the communication between d.3 and the d.ecs storage manager.

1.12.5. Storage system/Storage module/Storage driver

The term storage system describes a system supported by the d.ecs storage manager typically consisting of hardware and software components, on which d.3 or ecspan documents are written for secondary storage. Among others, these could be a NetApp system with SnapLock or an EMC2 Centera system. Every supported storage system is addressed via a specific d.ecs storage manager module. The module(s) must be selected and installed during the d.ecs storage manager setup. Every storage module or storage driver is either designed for exactly one or for several storage systems by one specific vendor and must not be used for compatible systems by other vendors. Examples for this could be the d.ecs storage manager modules Celerra and NetApp. The module Celerra can be used for EMC2 Celerra systems with File Level Retention option as well as for EMC2 VNX systems with File Level Retention option. The d.ecs storage manager module NetApp is designed for storage systems by the company NetApp (so-called FAS systems) with SnapLock-Enterprise or SnapLock-Compliance and must not be used for other compatible systems (SnapLock-compatible) by other vendors.

1.12.6. System table

If respective systems are configured via the d.ecs storage manager administration, the system-specific settings are stored in respective module configuration files. The names of the module configuration files have the same name as the configured systems. Example: Module Netapp, name of the module configuration file **netapp.cfg**. In addition to the storage in the module configuration files, specific configuration information are also stored in a database table (system table).

1.12.7. Retrieve

Retrieval means that a document no longer located in the d.3 document tree or in the d.3 gateway cache is requested via the d.ecs storage manager. Typically this is done by the d.3 gateway passing a retrieval job (restore job) to the d.ecs storage manager and this provides the requested document(s) for the d.3 gateway. Depending on the d.3 gateway configuration, a document provided by the d.ecs storage manager is stored in the d.3 gateway cache. Retrieval does not mean that the document is written back to the d.3 document tree. This is only the case, if a new document version has to be created for a document only located on the secondary storage.

1.13. Additional information sources and imprint

If you want to deepen your knowledge of d.velop software, visit the d.velop academy digital learning platform at <https://dvelopacademy.keelelearning.de/>.

Our E-learning modules let you develop a more in-depth knowledge and specialist expertise at your own speed. A huge number of E-learning modules are free for you to access without registering beforehand.

Visit our Knowledge Base on the d.velop service portal. In the Knowledge Base, you can find all our latest solutions, answers to frequently asked questions and how-to topics for specific tasks. You can find the Knowledge Base at the following address: <https://kb.d-velop.de/>

Find the central imprint at <https://www.d-velop.com/imprint>.