

This plugin supports the transfer of files between CatDV and Amazon S3, (simple storage service). In addition, it supports the restore of files that were automatically transferred to Glacier cold storage from S3 via bucket lifecycle policies.

This S3 plugin does NOT support the direct transfer of files between CatDV and a Glacier Vault. Glacier vaults are a separate type of storage which is entirely independent of S3 and would need to be handled by a different plugin.



Before you start

For this version of the Amazon S3 archiving plugin, the following are required:

- CatDV Server 10.6.3 or later
- CatDV Pegasus or Enterprise Client 14.1 or later
- CatDV Plugin licence 'ARCS3' (Rest API licence with multiple sessions)

To trigger Amazon S3 cloud file transfers from the Worker you need:

- S3 Worker Plugin 3.5.3, included in this installation as AmazonS3Worker.catdv

In addition, to run the archive service standalone (outside the server):

- CatDV Service Control Panel 2.1.0 or later

New Installation

IMPORTANT: If you don't already have an Amazon S3 account or you don't know the access keys see [S3 Configuration](#) for further details.

- BEFORE PROCEEDING, please ensure the prerequisites from [Before you start](#) are satisfied.
- BEFORE PROCEEDING, please read the release notes for this version of the plugin.
 - IMPORTANT: Third party S3 implementations have not been tested with this release, which includes a significant upgrade to the AWS SDK. If you are using the plugin to archive to a service other than S3, please do not install unless you have a test environment where you can test this release prior to installing it in production.
 - IMPORTANT: Snowball Edge is NO LONGER SUPPORTED. See [Snowball Edge](#).
 - IMPORTANT: KMS Encryption is NOT SUPPORTED. See [Using KMS Encryption](#)

- Copy the whole directory extracted from this zip to the plugins directory:

```
e.g. on Unix systems: /usr/local/catdvServer/plugins
e.g. on Windows: C:\Program Files (x86)\Square Box\CatDV Server\plugins
```

- Set optional server properties for the plugin, if desired:
 - Open 'Server Config' from the control panel
 - Enter required properties into the CatDV Server tab, 'Other' field. See [Plugin server properties](#).
- Restart the CatDV server from the control panel
- Open the CatDV client (or log the client back in to the server)
- Configure one or more service accounts (see [Managing Service Accounts](#)). The first account will initially be the default account, used to verify that S3 is available and used as a fallback when necessary. See [Manage service command](#) for further information on the details and settings for a service account:
 - In the client, run Tools->Manage Amazon S3 Archive Service
 - On the accounts tab enter the service account details. Mandatory details are flagged with an asterisk. Note that once a service account has been used to archive a clip, it will no longer be possible to update its identifier.
 - Optionally modify the defaults on the settings tab. These are mostly applicable to 3rd party S3 compatible services.
 - If the account will be used to archive to a [Snowball Edge](#), make sure "Amazon Snowball" is set to yes, or archive will not work. (This option is only visible if Snowball support has been purchased and enabled).
 - If optional KMS encryption applies to the bucket, set KMS settings appropriately. See [Using KMS encryption](#).
 - Click 'Add'.
- Configure service settings (See [Manage service command](#))
 - Optionally configure archiving, processing and UI settings for the plugin via the corresponding tabs on Tools->Manage Amazon S3 Archive Service.
 - If S3 plugin is being used as a bridge for a 3rd party service, e.g. Dell ECS, ensure that the Service Name and Service Description are set appropriately via the UI tab of Tools->Manage Amazon S3 Archive Service.
- Set up [Media Store Service Path Mappings](#) to automatically map drives/folders to archive locations. If you plan to always use a single service account and a single bucket, you may skip this step and use default location mapping. However, it is still recommended to use this approach if you don't want the file 'paths' on S3 to include the path of the storage volume.
- If [Running archive service standalone](#), then configure and start the archive service via the service control panel.
- Verify the service setup: In the client, run Tools->Manage Amazon S3 Archive Service again. The service status for the default account should be 'Running (online)'. The status may be 'Running (offline)' if Amazon S3 is not

currently accessible.

- IF the Worker IS NOT being installed
 - Delete AmazonS3Worker.catdv from the Amazon S3 plugin directory (see first step)
- IF the Worker IS being installed:
 - Move AmazonS3Worker.catdv FROM the Amazon S3 plugin directory (see first step) TO the worker plugins directory:

```
e.g. on a Mac: /Library/Application Support/Square Box/Extensions  
e.g. on Windows: %ProgramData%\Square Box\Extensions  
e.g. on Linux: /usr/local/catdvWorker/extensions
```

- If the worker is running, re-connect to the server so that the archive plugin fields are available.
- Verify that the Amazon S3 worker actions are available. It may be necessary to quit and re-open the worker to pick up the new worker actions.

Upgrade Installation

- BEFORE PROCEEDING, please ensure the prerequisites from [Before you start](#) are satisfied.
- BEFORE PROCEEDING, please read the [Release Notes](#) for this version of the plugin, and for each plugin version listed in the release notes above the last installed version.
 - IMPORTANT: Third party S3 implementations have not been tested with this release, which includes a significant upgrade to the AWS SDK. If you are using the plugin to archive to a service other than S3, please do not upgrade unless you have a test environment where you can test this release prior to upgrading in production.
 - IMPORTANT: Snowball Edge is NO LONGER SUPPORTED. See [Snowball Edge](#).
 - IMPORTANT: KMS Encryption is NOT SUPPORTED in this release. If you have been using this feature with a previous release, do not upgrade and contact support.
- Make a note of the current plugin version from this document (or see latest version in [Release Notes](#))
- If running the archive service standalone, use the service control panel to stop the service.
- Stop the CatDV server from the control panel
- Copy the whole directory extracted from this zip to the plugins directory:

```
e.g. on unix systems: /usr/local/catdvServer/plugins  
e.g. on Windows: C:\Program Files (x86)\Square Box\CatDV Server\plugins
```

- Remove or move the directory / files for the prior plugin from the plugins directory.
- Carry out any 'Upgrade' instructions for each plugin version listed in the [Release Notes](#) above the last installed version, working back from the oldest applicable version.
- Start the CatDV server from the control panel.
- Read through the 'Changes' for each plugin version listed in the [Release Notes](#) above the last installed version, and go to Tools->Manage Amazon S3 Archive Service to verify that the details / settings for each account and the archiving / processing / UI settings for the plugin are correct for this installation. See [Manage service command](#) for more details.
- If [Running Archive Service Standalone](#), then update and start the archive service via the service control panel.
- Open the CatDV client (or log the client back in to the server).
- Verify the service setup: In the client, run Tools->Manage Amazon S3 Archive Service again. The service status should be 'Running (online)'. The status may be 'Running (offline)' if Amazon S3 is not currently accessible.
- Update (or delete) the Amazon S3 worker plugin by following the instructions for the worker plugin in a new installation. If there is an older Amazon S3 worker plugin file in the extensions directory with a versioned name, delete it.

Plugin server properties

When running the archive service "InPlugin", the following server properties must be set:

```
catdv.s3archive_service.licence_code = <generated licence code for plugin>
```

When running the archive service standalone, the properties for the archive service must include:

```
catdv.s3archive_service.licence_code = <generated licence code for plugin>
catdv.rest_api.host = <CatDV server hostname>
catdv.rest_api.port = <CatDV API port>
catdv.rest_api.user = <CatDV API user>
catdv.rest_api.password = <CatDV API user password>      (** exclude if blank, for dev)
```

Typical rest API values for a development machine would be (password property omitted for a blank password):

```
catdv.rest_api.host = localhost
catdv.rest_api.port = 8080
catdv.rest_api.user = Administrator
```

In addition, the following optional property may be set to turn on debug logging on both the server and standalone service:

```
catdv.s3archive_service.debug = true
```

Running Archive Service Standalone

By default, the service that handles archive jobs runs inside the CatDV server. The following is required to run the plugin's archive service on a separate machine from the CatDV server:

- CatDV Service Control Panel 1.3 or later

To configure and start the standalone archive service using the service control panel:

- Open the [Manage service command](#) in the client / web plugin UI
- Verify that at least one service account has been configured
- In the Processing tab, set the 'Service mode' to 'standalone' and click 'Save'
- Re-open the Manage Service Command and verify that the service is 'Stopped'. If not, restart the CatDV server.
- Install the CatDV Service Control Panel
- Open the CatDV Service Control Panel
- Click 'Service Config' to open the service configuration
- Fill in license details on the Licensing tab
- Optionally, fill in details of a REST endpoint on the REST Service tab
- Fill in server connection details on the Server Connection tab (connection can be checked on the control panel)
- Enter required service properties on the Settings tab. See [Plugin server properties](#) above. (**Note the server connection details are duplicated here for now, pending a plugin update to remove them)
- On the Installation tab, note the install location ('Server Install Path')
- On the file system, navigate to the install location and copy the plugin files into the plugins sub-directory.
- Click 'Start' to start the archive service.

To update the standalone archive service:

- The service should have been stopped via the Service Control Panel 'Stop' prior to stopping the CatDV Server.
- Open the CatDV Service Control Panel
- Add / edit service properties on the Settings tab if applicable.
- Go to the install location (see Installation tab) and replace the plugin files in the plugins sub-directory with those for the latest version of the plugin.
- Click 'Start' to start the archive service.

S3 Configuration

To use the plugin, at least one Amazon AWS account is required which owns the bucket(s) to which media will be archived. The following JSON security policy indicates the minimum permissions which must be configured for a user to use the core plugin functionality:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ExampleStatement0",
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3::my-bucket/*",
        "arn:aws:s3::my-bucket",
      ]
    },
    {
      "Sid": "ExampleStatement1",
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "*"
    }
  ]
}
```

S3 Test Account

IMPORTANT: The instructions in this section are only applicable to setting up a development account for testing purposes. They are not meant for setting up a production ready S3 configuration which is outside the scope of this documentation. For the minimum S3 permissions required to use the plugin in a production setting, see the preceding section.

If you need to set up an Amazon AWS account for testing purposes, follow the link 'Get started with Amazon S3' from <http://aws.amazon.com/s3/>. Then the most straightforward way to proceed is to use an existing Amazon account.

NB - The account is only free for *limited usage for a one year period*, see <https://aws.amazon.com/s3/pricing/>

Once you are logged into an AWS account, <http://aws.amazon.com/s3/> will take you to the S3 console. In order for the plugin to copy to / move to / restore from your AWS account, you will need to set up a user with access keys and at least one bucket:

To set up a user:

- Open the AWS Identity and Access Management Console: <https://console.aws.amazon.com/iam/home#home>
- Follow the steps under security status to create a User AND create a Group containing that user which has the permission AmazonS3FullAccess.

To get access keys for a user - these are the access key and secret key required to configure the plugin:

- Go to <https://console.aws.amazon.com/iam/home?#users>
- Click on the user name
- Click on the 'Security Credentials' tab
- Click 'Create Access Key'
- Either Click 'Download Credentials' or click 'Show User Security Credentials' and copy the key values.

To create a bucket:

- Go to <http://aws.amazon.com/s3/>
- Click 'Create Bucket'
- Enter bucket name and click 'Create'
- This is the bucket name required to configure the plugin. The value for the region can be seen in the URL when viewing the bucket. This must be a region code from the following list (default is us-east-1):

US East (N. Virginia) / US Standard:	us-east-1
US West (Oregon):	us-west-2
US West (N. California):	us-west-1
EU (Ireland):	eu-west-1
EU (Frankfurt):	eu-central-1
Asia Pacific (Singapore):	ap-southeast-1
Asia Pacific (Tokyo):	ap-northeast-1
Asia Pacific (Sydney):	ap-southeast-2
Asia Pacific (Seoul):	ap-northeast-2
South America (Sao Paulo):	sa-east-1

The values required to configure the service in the Amazon S3 plugin are:

- region/endpoint - region code (above) OR endpoint for an S3 compatible archive
- access key - access key from S3 credentials of Amazon S3 user
- secret key - secret key from S3 credentials of Amazon S3 user
- default bucket - name of an accessible bucket on Amazon S3 which can be used to verify the connection, used as the default bucket the first time a user attempts to copy / move files to Amazon S3
- CMK ID - customer master key for authorising client side KMS encryption, if applicable
- KMS Key ID (optional) - KMS key ID for authorising server side KMS encryption, if applicable

Snowball Edge

Support for Snowball Edge has been removed from the S3 plugin for the following reasons combined:

- AWS Snowball Edge is no longer available to new customers. According to Amazon, new customers should explore [AWS DataSync](#) for online transfers, [AWS Data Transfer Terminal](#) for secure physical transfers, or AWS Partner solutions. For edge computing, explore [AWS Outposts](#)."
- Amazon mandated upgrading the AWS SDK for Java used by the plugin from 1.x to 2.x, as 1.x is going out of support. The 2.x SDK is a complete rewrite which means the Snowball support would need re-engineering in conjunction with a customer that has a Snowball Edge available for testing.

Using KMS Encryption

KMS Encryption is not supported in this release of the plugin, following the upgrade to AWS SDK for Java 2. If you have been using this feature with a previous release, do not upgrade and please contact support.

~~NB - The support for KMS encryption only works with Amazon S3, not with S3 compatible interfaces to other cloud services.~~

~~The Amazon S3 plugin supports the use of KMS encryption to protect your data either at rest on S3 (server) OR both in-transit to/from S3 and at rest on S3 (client). If you turn on server encryption, the data will be encrypted and decrypted by and on S3. If you turn on client KMS encryption, the plugin will use the Amazon S3 encryption client to interact with S3, rather than the standard client, and all data will be encrypted in the plugin before it is sent to S3 and decrypted in the plugin when it is retrieved from S3.~~

For more details of Amazon's implementation, see "Using the Amazon S3 Encryption Client" in:

<https://docs.aws.amazon.com/kms/latest/developerguide/services-s3.html>

To use KMS encryption:

- ~~Create a service account, or update an unused account (see [Managing Service Accounts](#) and [Manage Service Command](#)), setting the 'KMS Encryption' to client or server, as required:~~
 - ~~The 'client' option encrypts your data locally to ensure its security as it passes to the Amazon S3 service, where it remains encrypted. The S3 service receives your encrypted data and does not play a role in encrypting or decrypting it. The 'server' option enables data to be encrypted in storage on S3 but not in-transit. The S3 server encrypts and decrypts the data.~~
 - ~~For 'client', a customer master key (CMK) ID for encryption must be provided and the user associated with the S3 credentials must have permission to use the CMK for encryption and decryption.~~
 - ~~For 'server', a kms key ID may optionally be provided, otherwise the default kms key will be used.~~
 - ~~IMPORTANT: If the encryption settings for a service account are changed after the account has been used to archive files, subsequent restores may be unable to decrypt those files. Care must be taken when doing this.~~
- ~~For 'client' client encryption, install JCE (Java Cryptography Extension) unlimited strength policy files~~
 - ~~Download the files from <http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html>~~

- ~~Unzip download~~
- ~~Copy the two jar files from the unzipped directory to jre/lib/security for the plugin installation:~~
- ~~i.e. for archive service run 'InPlugin' (default):~~

~~e.g. on unix systems: /usr/local/catdvServer/jre/lib/security/
e.g. on Windows: C:\Program Files (x86)\Square Box\CatDV Server\jre\lib\security~~

- ~~i.e. for archive service run 'Standalone':~~

~~e.g. on unix systems: /usr/local/catdvService/jre/lib/security/
e.g. on Windows: C:\Program Files (x86)\Square Box\CatDV Service\jre\lib\security~~

- ~~Create Media Store service mappings to determine which clips will be archived using the encrypted service account~~

Media Store Service Path Mappings

Media Stores should be used to automatically map file storage to archive locations by adding a service path to each media store. The service path determines the service account, bucket and (optional) folder path for the archived files from that file storage. Service accounts are set up via the Manage Service screen (see [Managing Service Accounts](#)) and are given an identifier when they are created.

For Amazon S3 a service path has the format:

```
<serviceType>://<accountIdentifier>/<bucketName>  
or  
<serviceType>://<accountIdentifier>/<bucketName>/<folderPrefix>
```

For example, if a Media Store were set up with the following two entries:

```
/Volumes/dance-videos/drafts  
amazons3://squarebox-test/video-test/dance
```

Then a clip located at:

```
/Volumes/dance-videos/drafts/greatest-show/this-is-me
```

Would be archived:

- using the credentials from the 'squarebox-test' account
- into the bucket 'video-test'
- with a path of '/dance/greatest-show/this-is-me'

Each Media Store is intended to represent a distinct segment of physical storage so, to avoid ambiguity, no path should be present in multiple media stores and paths should not overlap (e.g. there should not co-exist media stores for /Volumes/dance-videos/drafts and /Volumes/dance-videos/drafts/special).

To avoid any clash or overwrite of files with the same path and name from different drives, make sure they are mapped to either distinct S3 buckets or distinct folders within the same S3 bucket.

The 'folder path' can be set or overridden on archive if 'location' is set as an allowed override. See [Manage service command](#) (UI TAB).

Restore path mapping

If a restore location is not writable when a restore job is run, the plugin will use Media Stores to attempt to find an alternative writable restore location. If there is a Media Store with a path matching the restore location, the plugin will check whether any of the other paths in that Media Store are writable and substitute the root directory of the restore location accordingly.

The restore location can be overridden when a restore job is scheduled, if the option to override it is turned on. In earlier versions of the plugin, in this case no file hierarchy was preserved – all files were written directly to the specified directory. Now the plugin will use Media Stores to try and preserve file hierarchy for restored files. If any Media Store path is found that matches the default restore location, it will be used to substitute the root directory of the default restore location with the override.

Note that if many files are restored together to an overridden location, some files may have hierarchy preserved and some may not, depending on whether matching Media Store paths are found.

Using the plugin

The plugin comprises various commands that are available from the Tools menu in the client.

File transfers are initiated by the schedule backup / archive / restore commands but are carried out by a separate process. This means that the client does not automatically reflect updates to a clip on completion of a transfer. Use 'Server->Refresh Window' in the client to pick up any changes during / after transfer.

The plugin includes the following commands:

Command	Description
Manage Amazon S3 Archive Service	View full status details of the archive service and manage the service account(s) required to access S3 archives. Service account details include region / endpoint, access key, secret key, default bucket name and optional CMK ID. If a connection to Amazon S3 cannot be made an error is displayed and the values are not saved.
View Amazon S3 Job Queue	Lists jobs, which can be filtered by status (Pending, Complete, Failed, All). It provides the capability to view job contents, view job results and cancel waiting or queued jobs. IMPORTANT: The number of jobs displayed in any view is limited to the 'Max no jobs to display' configured on the UI tab of manage service. To see all jobs use the web admin job queue.
View Amazon S3 Jobs for Clip(s)	Provides similar functionality to the job queue but displays all jobs / jobs of a given status for the selected clip(s) only, including clip members where applicable. For folder type clips (folder / playlist / version set / multicam) this incorporates all clips & members in the hierarchy.
Schedule copy to Amazon S3	Adds a copy job for each clip selected (if multiple clips reference the same source media only one job is created). Copy jobs can be scheduled when the archive service is not running or is offline but will only be run when the archive service is online (running and Amazon S3 is accessible). When the copy job is run, the source media associated with the clip is copied from the file system to Amazon S3 and the original file is preserved.
Schedule move to Amazon S3	As copy but on successful transfer of the file to storage the job attempts to delete the original file. If for some reason the deletion fails, this does not affect the job status (or archive status of the clip) but "Purge failed" is recorded in the job data.
Schedule restore from Amazon S3	Adds a restore job for each clip selected (if multiple clips reference the same source media only one job is created). Restores can be scheduled when the archive service is not running or is offline but will only be run when the archive service is online (running and Amazon S3 is accessible). When the restore job is run, the source media associated with the clip is copied from Amazon S3 to the file system. Unless the force restore option is used, clips will not be scheduled / run if they've been updated (date or size) since the last archive.
Schedule purge of local files copied to Amazon S3	Adds a purge job(s) for each clip selected if they have been successfully archived (if multiple clips reference the same source media only one job is created). Purges can be scheduled when the archive service is not running or is offline but will only be run when the archive service is online (running and Amazon S3 is accessible). When a purge job is run, the source media associated with the clip is deleted. Unless the force purge option is used, clips will not be scheduled / run if they've been updated (date or size) since the last archive. If there is an urgent need to free up space when there are already many jobs in the queue, the priority of purge jobs can be adjusted to 'Highest', which will push them through the Highest priority lane regardless of how many other jobs are running.

Recover Amazon S3 archive data	Attempts to recover the archive data for the selected clip(s) so that they can subsequently be restored and/or purged. If the archive key is missing, the same approach as copy / move is used to determine the location of the file on S3, including media store mappings and overrides where applicable. The location is used to try and locate the file on S3. If a clip has the 'Last archived parameters' filled in then those will be re-instated. Otherwise: A combination of the source media level archive status (if present), last complete archive job (if present) and archive key are then used to recover as much of the archive data as possible.
---------------------------------------	---

Managing Service Accounts

The accounts tab of Tools->Manage Amazon S3 Archive Service provides facilities to manage the service account(s) used to connect to Amazon S3 archives. A single account is usually sufficient to operate the plugin. However, if required, there is an add-on which enables multiple accounts to be set up for security purposes etc. This is enabled as an add-on via a special server config property.

When there are multiple accounts, the account used may be determined automatically via [Media Store Service Path Mappings](#) or selected by the user on archive / restore if allow override of Account ID is turned on from the UI tab. See [Manage Service Command](#) (UI Tab).

An account which is in use, i.e. has been used to archive files, must be unlocked before it can be updated or deleted. This is to ensure that accidental changes cannot be made to an account such that files can no longer be restored. Note it is not possible to change the account identifier on an in-use account.

The accounts tab provides the following operations:

Button	Description
Clear	Clears the current selection / details so that only default values are filled in.
Set Default	Updates the default account to the current selected account.
Add	Creates a new account with the specified details.
Unlock	Unlocks an in-use account in preparation for update / delete.
Update	Updates the selected account with the specified details. If an account is in use, it must be unlocked before it can be updated. Please take care when updating an account which is in use, as changing some settings could break the restore of files archived with that account. For example, if the account key is switched to one which does not have the same access or if any encryption settings are changed. Note it is not possible to change the account identifier on an in-use account. IMPORTANT: Changes to service accounts will not take effect until the archive service has been restarted. If operating 'in-server', then the CatDV server must be re-started. If operating 'standalone' then re-start via the service control panel on each data mover.
Delete	Deletes the selected account. Please take care when deleting an account which is in use, as it will no longer be possible to restore any clips archived with that account. This feature is intended for removing test accounts and any applicable clips should be deleted or restored and re-archived.

Manage service command

The following can be configured from Tools->Manage Amazon S3 Archive Service in the web or desktop UI for the plugin:

ACCOUNTS TAB / DETAILS (for an account):	
Field	Description
Account Identifier	Identifying name for this service account. This may contain only alpha- numerics and hyphens. Note it is not possible to change the account identifier on an in-use account.
Region	Region for S3 or endpoint for a 3rd party S3 compatible service IMPORTANT: An endpoint must now be prefixed with either https:// or http://
Access Key	Access key for connecting to this service account
Secret Key	Secret key for connecting to this service account
Default Bucket Name	The bucket that will be used as the archive fallback for this account by default. This is not applicable when a media store service mapping applies. Otherwise, if the UI tab settings allow the user to override the bucket then this will only be used as the default the first time the user does an archive. Subsequent archives by the same user will default to the last value they entered.
Default Cold Storage Restore Priority	The restore priority that will be applied to restores from cold storage (Glacier) by default. If the UI tab settings allow the user to override the restore priority, then this will only be used as the default the first time a user does a restore. Subsequent restores by the same user will default to the last value they selected.
KMS Encryption	Determines whether client, server or no KMS encryption is used. See section on 'Using KMS Encryption' above.
KMS Key ID / CMK ID	Must be provided if client side KMS encryption has been selected (KMS Key ID) OR may be specified if KMS encryption is configured for a bucket in S3 and you need to use a CMK ID other than the default.
ACCOUNTS TAB / SETTINGS (for an account):	
Target throughput in Gbps	NB This is typically used in place of max concurrency (set that to blank). The target throughput for transfer requests. Higher value means more connections will be established with S3. Whether the transfer manager can achieve the configured target throughput depends on various factors such as the network bandwidth of the environment, whether max concurrency is configured and the amount of available memory. By default, it is 10 gigabits per second. If users want to transfer as fast as possible, it's recommended to set it to the maximum network bandwidth on the host that the application is running on. If you are running into out of file descriptors errors, consider using max concurrency to limit the number of connections. Note: This setting affects the native memory usage used by CRT; a higher throughput value will result in a larger memory usage. Typically, a range of throughput values maps to a discrete memory limit value in CRT, with a maximum upper limit.
Max concurrency	NB This is typically used in place of target throughput (set that to blank). The maximum number of S3 connections that should be established during a transfer. If not provided, the TransferManager will calculate the optimal number of connections based on target throughput In Gbps. If the value is too low, the TransferManager might not achieve the specified target throughput. Note, there is a separate TransferManager for each S3 account set up in the plugin, so this should be taken into consideration.
Multipart upload threshold (MB)	The file size at which the AWS SDK will start breaking files into multiple parts for upload. Uploads/copies over this size will automatically use a multipart upload strategy, while uploads/copies smaller than this threshold will use a single connection to upload/copy the whole object.

	Multipart uploads are easier to recover from and potentially faster than single part uploads, especially when the upload parts can be uploaded in parallel. Because there are additional network API calls, small objects are still recommended to use a single connection for the upload. See Uploading and copying objects using multipart upload
Minimum part size (MB)	The minimum size of each part for a multiple part upload / download. Decreasing the minimum part size causes multipart transfers to be split into a larger number of smaller parts. Setting this value too low has a negative effect on transfer speeds, causing extra latency and network communication for each part. By default, it is 8MB. See Amazon S3 multipart upload limits for guidance.
Native memory limit	The amount of native memory that CRT is allowed to use when making requests to S3. If not provided (blank), the CRT attempts to limit native memory usage in an optimal way, based on a number of parameters such as target throughput. Therefore, only configure the memory limit explicitly when needed. Supported range: Min: 1 GB, Max: The lowest value of the supplied value and the SIZE_MAX of the system
Cold storage restore expiration (days)	The expiration period, as a number of days, for which a file restored from cold storage (Glacier) will be accessible directly from S3 (restore from cold storage is always temporary). The default is 2 days.
VPC endpoint	Amazon Virtual Private Cloud endpoint. This must be used in conjunction with a region in the region/endpoint field on the details tab. IMPORTANT: This value must now be prefixed with either https:// or http://
Host addressing	Determines whether the plugin uses virtual host addressing or resource path addressing to connect with the S3 service. Valid values are 'virtual' or 'pathStyle'. The default value is 'virtual'.
Signer override	Sets the name of the signature algorithm to use for signing requests. If not set or explicitly set to null, the Amazon S3 SDK will choose a signature algorithm to use based on a configuration file of supported signature algorithms for the service and region. The default is not set, which should ideally not be overridden for the Amazon S3 service. However, other archive services which provide an S3 interface may need to set this property, e.g "S3SignerType".
Disable parallel downloads	Disables parallel download of multiple parts of an object. This might reduce performance for large files. Default is "no".
Manually create target path	Ensures that the 'folders' in the target path are created prior to archiving a file. This is not applicable to S3 itself, which has no true concept of 'folders'.
Service Name	Enables the override of the plugin service name in job details and clip archive details using this service account. The service name used in plugin menu items and labels in the UI is the one from the UI tab. The only time this override should be used is when service accounts are set up to access multiple different S3 compatible services.
ARCHIVING TAB:	
Location mapping	Note that this is mainly for legacy purposes as the preferred approach is now to use Media Store Service Path Mappings. Fallback method for generating the archive file location on S3 if no media store service mapping applies. The default is 'mirror'. Valid values are: <i>mirror</i>: The archive location always mirrors the source file path which means a file is replaced each time it is archived. The downside of this approach is that it does not cater for files with the same path on different file systems. <i>batchMirror</i>: The file(s) selected for archive are batched together in a date and time stamped directory (format /yyy-mm-dd/hh:mm:ss.mmm) and mirror their source file path(s) within the batch directory. This has the effect of versioning uploads to S3 as a new copy is written for each transfer. <i>mediaStoreRelative</i>: Generates relative file paths from the relevant media store(s). If no matching media store path is found, the archive location is

	generated using the 'mirror' approach. <i>mediaStoreAbsolute</i>: As 'mediaStoreRelative' but also prepends the path with the name of the media store.
Exclude archived files from archive	Skip files that have already been archived when scheduling jobs to copy / move to archive, unless they currently have a copy / move failed status. Default is "yes" for new installs. Note, if the file(s) were previously copied, an attempt to move the file will simply be skipped, it will not purge the file. When this option is turned on, use the purge operation to purge files which have already been copied. This option can be temporarily disabled if it is necessary to resubmit files for some reason.
Exclude existing files from restore	Skip files that already exist when scheduling restore jobs. Default is "yes" for new installs.
Use media path as restore location	Always use the media path as the restore location when a job is scheduled (unless it is explicitly overridden), mapping the restore path at job runtime as usual if it is not writable.
Force purge (if updated since last archive) default	Default value for forcing purge of files updated since last archive, used if not overridden when submitting purge jobs. Default is no.
Force restore (if updated since last archive) default	Default value for forcing restore of files updated since last archive, used if not overridden when submitting restore jobs. Default is no.
Purge directories	Determines whether or not empty directories are deleted when purging (or moving) files.
Recover on restore	Determines whether a restore will automatically check first whether it can find the file on the archive using the current archive metadata and, if not, automatically attempt to recover the archive data prior to restore. Default is no due to the additional call to s3 for each restore. Can be overridden on restore from the UI (if enabled) and worker.
Files accessible from	Indicates whether files are accessible to the 'server & service' or the standalone 'service only'. In the case of 'service only' file related checks are only done when jobs are run by the service, not when files are submitted to the queue.
Metaclip size threshold	Threshold for number of members in a metaclip that can be handled by the plugin. If the number of members exceeds this, processing will be increasingly slow, and the clip should be ingested as a lightweight clip. Default is 500.
PROCESSING TAB:	
Service mode	NB Always ensure that the standalone service is NOT running prior to changing this value. Determines whether the job processing service runs in-server (as a thread in the CatDV Server) or standalone (via the Service Control Panel). The initial/default value is in-server. When this value is switched, the plugin will attempt to start or stop the in-server service as appropriate. If this does not appear to have been successful, as reported on the Overview tab, then it will be necessary to restart the CatDV server to apply the change.
Concurrent job limit	Determines the number of transfers that the archive service will attempt to run concurrently. The default value is 4.
Number Delayed retries	The number of times a waiting job is retried at an interval which increases exponentially. The default value is 10, which is equivalent to roughly 34 hours elapsed running time of the archive service.
Maximum retry delay (time period)	Limits the maximum time period between retries of waiting jobs. The default value is blank.
Loop delay (time period)	Determines the frequency with which the archive service checks the Job queue and attempts to run outstanding Jobs. Defaults to 30s.
Retry running job delay (time period)	Determines the time period after which a Job which is running will be restarted if it has not been updated during that period. Defaults to 1h.

UI TAB:	
Service Name	This property may be used to customise the name for an S3 compatible service such as Google Cloud Storage. Note that plugin menu items will not be updated until the UI has re-connected to the server (i.e. the plugin has been reloaded) and that it will not change the service name in any existing jobs or clip archive details.
Service Description	This property may be used to customise the description for an S3 compatible service such as Google Cloud Storage.
Restrict command access	Restricts the specified plugin commands to sys admin users. Can hide: - 'config' commands only (i.e. Manage Service) - 'config / archive' (i.e. to restrict copy and move to archive but not restore and purge) - 'config / archive / purge' (i.e. to restrict copy, move and purge but not restore) - 'all (excl queue)' - 'all' commands. NB a server restart is required to ensure that this change is picked up under all circumstances. Default is 'none'.
Max no jobs to display	The maximum number of jobs which will be listed in the job queue. This overrides days_to_display. The default value is 1000.
Allow overrides	Enables the facility for users to override one or more parameters at the point of archive or restore. Turning on any of these options adds a field to the Archive or Restore dialog as applicable, so that the value can be specified at the point of archive or restore: Archive: Account ID / Bucket name / Location / Job priority Restore: Account ID / Bucket name / Cold storage restore priority / Location / Job priority / Recover on restore

**Any duration can be set in h (hours), m (minutes), s (seconds) or any combination in that order, e.g. 2h, 2h 30m, 5m, 1m 30s, 1s.

Archive details pane

The plugin automatically creates a panel entitled "Amazon S3 Archive" containing the clip metadata which describes it's S3 archive state, including:

Field	Description
squarebox.catdv.archive.AmazonS3.serviceType	Type of service responsible for file transfer
squarebox.catdv.archive.AmazonS3.serviceName	Name of service responsible for file transfer
squarebox.catdv.archive.AmazonS3.status	Archive status
squarebox.catdv.archive.AmazonS3.location	Location of file in storage
squarebox.catdv.archive.AmazonS3.restoreLocation	Location file will be (if job pending) or was last restored to
squarebox.catdv.archive.AmazonS3.date	Date (timestamp) of latest change in archive status
squarebox.catdv.archive.AmazonS3.dateLastArchived	Date last archived
squarebox.catdv.archive.AmazonS3.dateLastRestored	Date last restored
squarebox.catdv.archive.AmazonS3.numArchives	The number of times the clip has been successfully archived
squarebox.catdv.archive.AmazonS3.archiveKey	Identifier of file in storage
squarebox.catdv.archive.AmazonS3.batchID	Identifies the batch of files with which the clip was archived
squarebox.catdv.archive.AmazonS3.jobID	Identifier of current / latest archive job
squarebox.catdv.archive.AmazonS3.parentJobID	n/a for Amazon S3 (related to bulk archives)
squarebox.catdv.archive.AmazonS3.userId	ID of user initiating current / latest transfer
squarebox.catdv.archive.AmazonS3.historyJson	Record of all archive activity in json format
squarebox.catdv.archive.AmazonS3.history	Record of all archive activity (including before historyjson added)
squarebox.catdv.archive.AmazonS3.purged	Indicates whether or not file has been purged by plugin (reset on restore)
squarebox.catdv.archive.AmazonS3.purgeError	Details of purge failure
squarebox.catdv.archive.AmazonS3.accountIdentifier	Identifier of the Amazon S3 Storage Account in CatDV
squarebox.catdv.archive.AmazonS3.regionName	Region of the Amazon S3 archive
squarebox.catdv.archive.AmazonS3.bucketName	Name of bucket to transfer file to / from

Known Issues

- Restore can overwrite read only files.

License Code

IMPORTANT: You may install and use this software only in accordance with the terms of the CatDV Server 8 license agreement.

Square Box Systems Ltd.

July 2026

Release notes

IMPORTANT: Third party S3 implementations have not been tested with this release, which includes a significant upgrade to the AWS SDK. If you are using the plugin to archive to a service other than S3, please do not upgrade unless you have a test environment where you can test this release prior to upgrading in production.

IMPORTANT: Snowball Edge is NO LONGER SUPPORTED. See [Snowball Edge](#).

IMPORTANT: KMS Encryption is NOT SUPPORTED in this release. If you have been using this feature with a previous release, do not upgrade and contact support.

3.5.3p2 (2026-07-31)

- S3 Account settings (for consistency with latest AWS documentation):
 - Renamed "thread pool size" to max concurrency and amended UI handling for it
 - Renamed "minimum upload part size" to "minimum part size"
 - Added target throughput option
- Fixes / workarounds for issues with the AWS Java SDK:
 - Upgrade AWS Java SDK from 2.41.18 to 2.44.4 to resolve an issue with the AWS SDK using excessive memory
 - Fix for intermittent fault which can result in inflated memory usage if the first two jobs are started almost simultaneously and are both transferring large files.
 - Workaround for large S3 transfers blocking the service loop
 - Workaround for apparent bug in AWS SDK causing transfers to wait indefinitely due to contention in calls to S3 between plugin threads and AWS event handling threads. There is now one S3 client for file transfers and a separate one for all other calls to S3.
 - Changed mechanism for aborting an S3 transfer, as thread interrupts of an active transfer sometimes hang
- Fix for requests for 'restore from cold storage to S3' failing silently as they were asynchronous.
- Add more details to debug output for S3 cold storage restores

3.5.3p1 (2026-02-09)

- Add S3 account setting for 'Native memory limit'
- Fix bug causing the max concurrency (thread pool size) for an S3 account to always be set to the cutoff value of 48.
- Fix recover archive data to only attempt to map an offline media path if the server has access to files

3.5.3 (2026-01-27)

IMPORTANT See notes at the top of the release notes section, which apply to this and subsequent releases until explicitly stated otherwise, regarding support for: 3rd party S3 interfaces, Snowball Edge and KMS encryption.

Upgrade

- Worker archiving plugin to 3.5.3 (in release)
- For S3 accounts configured with an endpoint or VPC endpoint, amend them to add a prefix of https:// or http://

Changes

- Upgrade AWS SDK For Java from 1.x to 2.x: This is a significant upgrade of the embedded AWS SDK, requiring changes to the interaction between the plugin and the SDK. This upgrade is mandated as the 1.x SDK is going out of support. It should improve transfer speeds to and from S3.
- Misc
 - Fix for highest priority lane
 - Fix recover archive data overrides
 - Ensure that the default override value for job priority is set correctly when there's no prior user value
 - Add option on archiving tab to always "Use media path as restore location" when a job is scheduled (unless it is explicitly overridden), mapping the restore path at job runtime as usual if it is not writable.
 - Fix service loop to prevent looping with no exponential backoff after some errors
 - Fix for infinite loop iterating clips if the REST API throws an error retrieving the members of a lightweight clip
 - Improve efficiency of service loop housekeeping.
- Purge
 - Enhance mechanism for cleaning up ancestor directories on purge to include .\$\$\$ temp restore files and the MacOS system files .DS_Store and ._ prefixed extended file info
 - Include detail when a purge fails with "File not found on archive" because the bucket name or archive key is empty, or some other unexpected error.
 - Fixes to setting purged flag and purge / delete error on complete / fail of various job types
- Recover
 - Fix recover archive data to attempt to map an offline media path to a writable path, as restore does, in order to try and map to the correct archive location
- Error Handling:
 - Ensure that trashed clips (catalog id 0) are handled as 'not found'

- Ensure that jobs give a fatal error when attempting to put / get a clip not found
- Improve error handling for clips with no media path
- Ensure that service loop automatically fails down a (re-)queued job with associated clip 'not found'
- Fix for an NPE processing a waiting job with associated clip 'not found'
- Logging:
 - Add server plugin debug switch on Manage Service Overview (only server debug, NOT standalone service)
 - Ensure logged path for "no source media found matching path" is no longer incorrectly reported as "null"
 - Log S3 transfer parameters as info, rather than debug
 - Log S3 transfer when bytes start to be transferred
 - Log failure to rename a restored temporary file as plugin debug

3.5.2 (2025-05-20)

- See 2.4.3p1 - changes from that release incorporated into this release
- Amend copy / move / purge jobs to map file paths at job runtime, in case there is no server access to files (in which case the server cannot map file paths passed to the plugin at queue time)
- If a restore job gets a runtime 'file updated since last archive' error, fail it with no retries
- Fix to ensure purged flag is always set correctly on lightweight clip members
- Fix to ensure clip archive fields that are specific to restore (restore location, date last restored) and purge (purged flag, purge error) are not overwritten/cleared by reinstating the last archived details on recover.
- Fix to ensure that JobResult output cannot exceed the max permitted length for a SQL text field
- Simplify error message returned to UI for an AWS signature mismatch, as it now has a long postscript relating to Java 12

3.5.1 (2025-03-10)

- Map path prior to running an archive job to cater for mapped paths with service only file access
- Add S3 account option to disable chunked encoding

3.5 (2024-12-03)

WAIT TO UPGRADE / CONTACT SUPPORT TO CONFIRM COMPATIBILITY IF:

You are using the S3 plugin with any 3rd party S3 interface, rather than direct to S3.

IMPORTANT

Note that all purges are now carried out via scheduled jobs, which means they will no longer be executed immediately from either the UI or the worker. If any worker configuration relies on the assumption that a completed purge action means that the file(s) have been purged, it will need to be adjusted to work along the same lines as handling for copy / move / restore. Note this does not affect move jobs, which continue to incorporate purge in the same way as before.

If there is an urgent need to free up space when there are already many jobs in the queue, the priority of purge jobs can be adjusted to 'Highest', which will put them through the Highest priority lane even if max jobs are already running.

Upgrade

- Worker archiving plugin to 3.5 (in release)
- Amend worker workflows to take account of the introduction of purge jobs

Major Changes

- Purge now runs as a job like copy/move/restore, to facilitate running with no server access to files
- Addition of option to specify 'Files accessible from' service & server OR service only
- Addition of force purge option, to skip checks for files updated since last archive
- Addition of force restore option, to skip checks for files updated since last archive
- Provision of Highest priority 'lane', whereby one Highest priority job can always run in addition to the max jobs
- Add config (with override option) on Archiving tab for 'Recover on restore' of externally archived files

Minor changes and fixes

- Prevent processing of outsized metaclips which ought to be lightweight, configurable threshold defaults to 500
- Amend default filter option to pending for "View ... Jobs for Clip(s)" command
- Amend archive plugin job queue menu item name to remove "Archive Service" for consistency
- Amend recover archive data to always save reinstated last archived params if present (and can retrieve object metadata), rather than returning a 'skipping' warning.
- On job queue, suspend refresh while cancelling / updating priority on jobs
- Log in the clip archive history as 'skipped' rather than 'failed to queue': already purged (for purge), clip not found
- Ensure that there is a fatal job error when attempting to put / get a clip not found or with a missing media path
- Ensure that trashed clips (catalog id 0) are handled as 'not found' (*NB doesn't apply to lite clip members)

- Enable cancel of clips missing a media path
- Fix recover archive data for empty directories in metaclips
- Fix status order in aggregated archive status for lightweight clips, to be consistent with metaclips
- Fix to minimise clashes where two threads pick up the same job and one reports "current status is xxx, expected yyy", to improve efficiency
- Fix for open file leak purging ancestor directories.
- Fix to reset flag to bypass loop delay if there is an unexpected error processing jobs immediately after a notification of new jobs.
- Fix purge confirmation that can cause archive history to balloon on 'No to all' option
- Fix counting of duplicate clips submitted for metaclip members

2.4.3p1 (2025-04-18)

- Update fail/complete handling to try and address the issue of transfers hanging due to network errors
- On running job timeout, attempt to abort plugin job and associated file transfer
- Dump threads after a partially transferred job has not been updated for 1 hour, regardless of running job timeout
- Log file size of jobs when queueing and starting them

2.4.3 (2025-03-24)

- Additional debug, mainly to help spotting in logs if the AWS notifier thread is stuck on a file lock
- When retrieving S3 object metadata, return a more concise error if the object is not found on the archive
- Improve abort of incomplete multi-part uploads on failure, to reduce unnecessary S3 storage costs
- For S3 client config, ensure that max connections is always higher than the thread pool size
- Check file is readable immediately before submitting file for S3 transfer
- Upgrade AWS Java SDK to latest version 1.12.753, including switch to Glacier restoreObjectV2 method

2.4.2p1 (2024-05-15)

- Fix for issue where running copy/move from the worker can result in intermittent failure with an "Incorrect number of arguments ..." error. NB - Issue caused by job priority feature added in 2.4.2

2.4.2 (2024-05-10)

Upgrade

- Worker archiving plugin to 2.4p1 (in release)

Changes

- Add job priority as an override for archive and restore
- Fix for issue queueing clips with a null media path
- Fix to ensure that duplicate metaclips are excluded from the batch of files for a single job

2.4.1 (2023-10-20)

Upgrade

- Check whether CatDV server DB has an index on jobId on the jobResult table. Add this index if it is missing.
- Check whether any worker actions are impacted by the change to compound archive status for clips with members

Changes

- Plugin job queue / job queue for clip(s) add the capability to adjust the priority of jobs
- Add created date to job table on plugin UI job queue (also applies to 'View ... jobs for clips'.
- Add pending option to job queue for clip(s)
- Apply 'max jobs for display' limit to job queue for clips
- Change status order in compound archive status for clips with members. The 'complete' status will now always appear last if failed or other are present, e.g. "Copy failed:1, other:2, complete:2". This means that the text "Copy complete" or "Move complete" will only be present when all the clip members have been archived. This is to prevent worker queries from matching queries for "xxx complete" when not all members have been archived.
- Improvements to recover archive data
- Fix for clip failing to purge after confirmation
- Fix for service stuck in stopped status after switching back from standalone to in_service
- Fix clip queries for "View ... Jobs for Clip(s)" to try to address out of memory issue.
- Update plugin helper DB scripts to include MSSQL versions
- Update plugin helper script for deletion of all plugin data to remove data from the clipField table
- Fix for occasional collation error running some helper DB scripts
- Fix for handling clip not found, as exception throw appears to have changed in server 10.5
- When updating service accounts, display account unlocked even if account already unlocked

2.4p5 (2023-10-20)

- Fix to enable files on Windows marked as read-only to be purged

2.4p4 (2023-10-19)

- Additional logging for purge of local files

2.4p3 (2023-09-15)

- Fix to “fix for archiving files from a mapped media path”, with respect to lightweight clips

2.4p2 (2023-08-11)

- Amend the mechanism to get the next job to process to try and prevent an infinite loop for clips with corrupt archive data. NB jobs for such clips will be consistently skipped and need to be cancelled. The only indication of this will be info entries in the logs (it isn't necessary for plugin debug to be on to see these).

2.4p1 (2023-08-11)

- Fix for archiving files from a mapped media path
- Fix to exclude service paths when looking for a writable path for restore

2.4 (2022-11-15)

Upgrade

- Check requirements in 'Before you start'
- Worker archiving plugin to 2.4 (in release)

Changes

- Add support for nested folder clips (folder / playlist / version-set / multicam)
- Add support for lightweight directory clips
- New (beta) command “View ... Jobs for Clip(s)”, to aid monitoring folder / lightweight clips and for troubleshooting.
- Minimise duplicate failure entries in archive history
- Remove days to display limitation and UI config option for job queue. Limit by max jobs only.
- Remove stalled job delay / stalled indicator on running jobs (thread dump now only on running time out)
- Improve error reporting by checking for file not found before passing a file to Amazon's TransferManager
- (Merged from 2.2.11p2/p3) Fix for NumberFormatException parsing clip JSON from server containing proxyOffset

2.3.x (n/a)

These release numbers were used for beta releases of multiple data movers, which has been released at 3.x

2.2.14p5 (2022-11-14)

- Add S3 account configuration option for a VPC endpoint in addition to a region

2.2.14p4 (2022-10-25)

- Enable purge of externally archived files which have no user metadata on S3
- Set last modified date on restore of externally archived files which have no user metadata on S3
- Dump threads when a job stalls or times out, for troubleshooting purposes

2.2.14p3 (2022-09-01)

- Minimise duplicate failure entries in archive history

2.2.14p2 (2022-07-11)

- Fix for out of bounds exception triggering recover archive data from the worker

2.2.14betap1 (2022-05-13)

- Fix restore of objects from Glacier Instant Retrieval

2.2.14beta (2022-04-12)

- Modify thread syncing for job processing to improve performance and attempt to resolve a thread sync issue that intermittently causes jobs to be re-run after completion, resulting in an MD5 error for a move job.

2.2.13beta

Upgrade

- Worker plugin to 2.2.13

Changes

- Replace references to 'glacier' in plugin UI with 'cold storage'
- Replace references 'restore tier' in plugin UI and doc with 'restore priority'
- Add option to restrict access to plugin commands for all but the job queue
- Add 'Recover Amazon S3 Archive Data' plugin command (UI and worker)
- Allow a clip with a failed archive status to be purged - if a re-archive failed, a size/date mismatch will trigger that error
- Fix outstanding job errors to include the archive status text from the outstanding job, rather than the clip.

2.2.11p1 (2022-03-10)

- Fix to ensure the plugin differentiates correctly between an Image Sequence (Metaclip) and Image Sequence Lite.

2.2.11 (2022-02-25)

Upgrade

- worker plugin to 2.2.11

Changes

- (Server 10 only) Add support for archive / purge / restore of Image Sequence Lite
- Update worker plugin to support bulk server queries
- Fix for restore issue where filename contains a comma followed by a space
- Add skipped / queue failed job history entry when transfer fails to queue

2.2.10p3 (2022-02-29) released simultaneously with 2.2.11

- Fix thread sync issue causing jobs to fail with an MD5 error after having completed

2.2.10p2 (2022-02-21)

- Fix endpoint configuration to work with BlackPearl Native S3

2.2.10p1 (2021-10-27)

- Fix update source media matching a media path to ensure that duplicate source media are only updated once
- Fix file transfer progress updates so limited to 4 per second when multiple threads are transferring parts
- Apply exclusion settings when running jobs, in case duplicate jobs for a file path are somehow scheduled
- Fix purge/restore of empty folders in complex media for server 10
- Prevent move job failing if file is not found for purge, which can occur on job retry after a DB connection failure
- Fix job stuck in running state due to server failure between allocating job to service and updating job to running state
- Enhancement to utility script (catdvfixarchivestatus.sql) to selectively update status of queued / running / waiting jobs
- Utility script (catdvreducejobresults.sql) that can be used to clear out intermediate job results
- Utility script (catdvdeleteallplugindata.sql) that can be used to clear out all plugin related data, e.g. to reset an installation set up with the wrong workflow

2.2.10 (2021-10-08)

- Include JAXB library required by AWS Java SDK in release, for compatibility with server 10, which uses JDK 11

2.2.9p1 (2021-09-20)

Upgrade

- worker plugin to 2.2.9

Changes

- Add option to override bucket on restore, to support migration of files between buckets

2.2.9 (2021-04-26)

Upgrade

- server to 9.3.3

Changes

- Add 'No to all' option for purge clip confirmation
- Modify handling of waiting jobs to limit the resources they consume, including throttling re-queue of waiting jobs.

2.2.8 (2021-04-08)

Upgrade

- If using worker, update worker plugin to 2.2.8 (included with this plugin)
- If there is any script in place which relies on the value of the cross-plugin source media 'archiveStatus' field (e.g. for archive status traffic light indicators), ensure it is updated to reflect that 'Archived' may now be either 'Archived' or 'Purged'.

Changes

- Add service config option to 'Exclude archived files from archive', now the default for new installs. Where files are not expected to change, this ensures there is no duplication of archive time / cost if files are picked up repeatedly.
- Add 'purged' flag to archive parameters
- Update generation of clip level archive status to include a 'purged' state
- Improve consistency of error reporting between server and worker plugins
- In worker plugin, return success if file(s) have a pending job of the same type, to reduce errors from picking up duplicate files
- Fix for plugin license expiration causing ballooning logs

2.2.7p1 (2020-06-29)

Upgrade

- Server to 8.0.7p3 or later OR 9.0.1p14 or later

Changes

- Fix for conflicting clip edits in worker archive plugin actions, following meta-clip status update
- Fix to try and ensure that duplicate clips are excluded from restores, based on media path rather than file path

2.2.7 (2020-05-29)

- Update summary archive status on meta-clips when meta-clip member archive status changes

2.2.6p3 (2020-05-21)

Upgrade

- If using worker, update worker plugin to 2.2.2 (included with this plugin)
- If running service standalone, update CatDV Service Control Panel to 1.3.1

Changes

- Add S3 account options for configuring parallelization of large uploads to optimise speed and reliability
- Fix for purge of restored file where mapped media path has changed since file was archived
- Additional manage service option to restrict access to purge command along with archive commands
- Add option to exclude existing files from restore jobs
- Fix default html format for json history field, to resolve issue of archive pane flipping / redrawing
- Fix purge of empty complex media folders from bulk jobs
- Fix for issue restoring files ingested with non-windows paths on Windows

2.2.6p2 (2020-03-25)

- Add 'disable parallel downloads' setting to service accounts, for 3rd party services
- Fix field type for archive fields 'media file path' and 'last archived params'
- Update log of archive history to include location restored to, rather than location restored from

2.2.6p1 (2020-03-6)

- Ignore trailing slashes in media store paths
- Fix worker plugin commands so that they don't update the allowed overrides settings for the service to all be enabled
- Update worker plugin to 2.2.1 (only for debug output / version consistency)

2.2.6 (2020-02-07)

Upgrade (whilst CatDV server is down)

- backup the service table from the DB if it contains multiple rows
- run catdvs3plugin2.2.6.sql against the CatDV DB to update the service UID to the new combined value
- **Note that if there are multiple service rows with 'serviceType' AmazonS3, in which case only the one most recently updated should be in use, this script will update the UID on the service with the most recent lastModifiedDate and delete the other(s).
- remove occurrences of catdv.s3archive_service.service_mode from the server and service properties, via their

respective control panels.

Changes

- significant update to plugin initialisation to prevent it blocking the server thread on start-up and simplify switching between in-server or standalone service for job processing. The same UID is now used for in-server and standalone services and the service mode is switched via a UI service config option, rather than system property. See [Manage Service Command](#) (Processing Tab – Service mode).

2.2.5 (2020-01-30)

Upgrade

- If running service standalone, update CatDV Service Control Panel to 1.3

Changes

- Implement restore path mapping
- Fix standalone service to run with updated Service Control Panel
- All service config duration values are now displayed, and may be entered, in h (hours), m (minutes), s (seconds) or any combination in that order, e.g. 2h, 2h 30m, 5m, 1m 30s, 1s.
- Fix job queue to ensure that job results and details are cleared when the job table empties due to refresh or cancel
- Make feedback message more prominent for account operations on the manage service command
- Fatally fail jobs where the associated clip is not found (deleted since job created) and include 'Clip not found' in status details.
- Fix to prevent job exceeding max retries

2.2.3beta (2019-10-22)

- Change default location mapping for new installs to 'mirror', as batching needs to be phased out with media store mapping
- Require server version 8.0.4 for fix to #190925-110544-112, issue with table row auto-selection on web
- Require server version 8.0.4 for fix to #190927-110649-832, cannot multi-select expandable rows on web
- Fix Job Queue so that an attempt to cancel a failed/complete job does not report as a job cancelled
- Fix likely cause of unexpected service account re-ordering in manage service command
- Fix occasional erroneous display of unlock warning for on update/delete for accounts which are not in use
- Updated documentation for Media Store Service Path Mapping

2.2.2beta (2019-09-10)

- Fix allow overrides checkboxes on manage service command in line with server change for ticket #181106-091441-139
- Retry service start if rest API is not ready
- Modify generation of clip level archive status to incorporate both current and pending status
e.g. 'Archived [Copy running] to ...' etc
- Use queued archive job notifications to expedite processing of newly queued archive jobs
- Send archive job status notifications for all applicable changes in archive job status
- Restore by default to current media file path from clip, rather than from last archived parameters, so that target restore location can be updated when e.g. media paths are updated with a different volume name (*merged from 1.5.0p2)

2.1.4beta (2019-07-8)

Upgrade

- In the S3 panel, ensure "Amazon S3 Archive History" (identifier squarebox.catdv.archive.AmazonS3.history) is below the new "Amazon S3 History (json)" (identifier squarebox.catdv.archive.AmazonS3.historyjson). Ensure these fields are the last two in the panel as the json history field displays as a table.

Changes

- Permit update of in use accounts via unlock/update/confirm
- Change deletion of in use accounts to to unlock/delete/confirm direct from accounts tab
- Update worker plugin to 2.1.2: add account parameter to restore, as required for snowball support
- Update service fields / panel creation to include all fields in field group but exclude some from the panel
- Add json version of archive history to clip archive metadata
- Fix for UnsupportedOperationException in Collections::UnmodifiableMap.put() on plugin init
- Merged in fix from 1.5.0p1

2.1.0betap3 (2019-04-04)

- Support multi-select for cancelling jobs from queue

- Updated job progress notifications to exclude percentage from Copied / Moved / Restored notifications
- Run another job (if available) immediately after completion / failure of a running job

2.1.0beta (2019-04-02)

Upgrade

- Now require CatDV Server 8.0.1 or later and CatDV Pegasus 13.0 or later (for transfer progress notifications)
- Now need to set an undocumented config property to enable multiple accounts, as it is an add on product

Changes

- Fix accounts tab to work in both web and desktop UI now that issues with tabs within a panel have been resolved
- Send job progress notifications, using new notification system (requires server >= 7.4.2 to make use of them)
- Add dev support for snowball edge - this is also an add on which is enabled via an undocumented config property
- Add capability to override / select the account when copying / moving a file to S3, or restoring a file from S3
- Add capability to delete an account, via flag for deletion combined with save service
- Add config param to enable support for multiple accounts
- Updated version of worker to accommodate new account parameter
- Terminate archive job with no retries if an S3 404 Not Found error occurs
- Include file size in job data / job queue list
- Merged in purge changes from 1.5
- Merged in fixes from 1.4.19
- Merged in fixes from 1.4.18
- Merged in fixes from 1.4.16

2.0.4beta (2019-02-11)

- Fix issue with display of account details tab content on the manage service command in the web UI
- Fix to ensure first account added (rather than migrated from previous installation) is set to be the default

2.0.3beta (2019-02-06)

- Add ability to set default account (purge/restore of pre-migration archives will still use the 'first', i.e. migrated account)
- Add ability to restrict access to (config and) archive but allow restore
- Display 'Retry' for status of unsuccessful job results if the job is being / will be retried
- Fix NPE listing jobs if user doesn't have permission to access selected clip

2.0.2beta

- Merged in fixes from 1.4.14p4 and 1.4.15

2.0.1beta (2019-01-25)

Upgrade

- **NB – run upgrade for 2.2.6 prior to this or the migration will not find the service definition as it's UID will be wrong
- The plugin will automatically migrate the connection details from the service definition currently in use to be the initial 'default' service account. There is a ONE TIME opportunity to customise the identifier of this account by setting the following configuration parameter in the server BEFORE re-starting for the first time after the upgrade, as the details of an account which is 'in use' cannot be changed (NB - The specified account identifier may contain only alpha-numeric and hyphens):

catdv.s3archive_service.migration_account_identifier = <my-account-identifier>

With a single account and no media store service mappings, the plugin will continue to operate as before, using the default account and configured location mapping for all archives. With multiple accounts and corresponding media store service mappings, any clips for which no service mapping is found will fall back to use the configured location mapping, as before.

- The plugin will automatically migrate any optional config values from the server properties to the service definition and they should subsequently be edited using the web or desktop Manage Service UI. You should be able to confirm that these have been set appropriately by going into the UI and comparing the values to the ones in the server config. For a standalone service, any additional properties set only in the service config will need to be either copied into the server config before starting the server with the new plugin for the first time OR set manually in the UI.
- After this migration, all optional service config properties should be deleted from the config for the server and (if applicable) service control panel. The following properties are not optional and should not be removed if present:

catdv.s3archive_service.licence_code
catdv.s3archive_service.debug

Changes (MAJOR UPDATE)

- Add support for multiple archive accounts
- Enhancements to manage service command, enabling most configuration to be done via the UI
- Support for automatic mapping from clips to archive locations, via Media Store service mappings which include a service identifier, account identifier, container name and (optional) 'folder'
e.g. amazons3://squarebox-test-2/vcvideotest/dance

1.5.0p2 (2019-07-09)

- Restore by default to current media file path from clip, rather than from last archived parameters, so that target restore location can be updated when e.g. media paths are updated with a different volume name

1.5.0p1 (2019-06-06)

- Fix NPE on cancel job, when associated clip is not found

1.5 (2019-03-28)

Upgrade

- run catdvarchiveplugins1.5.sql against the CatDV DB to fix the field type for archive date fields
- update worker plugin if applicable

Changes

- Changes to purge (these apply to move to archive as well as direct purge): clear date last restored, log 'File purged' or 'Purge failed' to clip archive params history, add purge error to clip archive params, add option of "Yes to all" confirmation for purging multiple altered files.

1.4.19 (2019-03-22)

- Fix blocking mechanism for ensuring that multiple processes cannot update a job's status simultaneously
- Terminate restore job with no retries if an S3 404 Not Found error occurs

1.4.18 (2019-03-06)

- prevent spurious 'Network outage'/NPE error when no restore directory can be extracted from the restore location
- update thread handling when processing jobs, to ensure a completed job will not be retried before it's status has been fully updated

1.4.16 (2019-02-06)

- Display 'Retry' for status of unsuccessful job results if the job is being / will be retried
- Fix NPE listing jobs if user doesn't have permission to access selected clip

1.4.15 (2019-01-31)

- Improve performance of Job Queue command for large lists of jobs

1.4.14p4 (2019-01-28)

- Fix to cause job creation to fail if no source media are updated with the archive details for the job

1.4.14p3 (2019-01-25)

- Fix for issue of jobs stuck in running / stalled state
- Trim trailing path separators from archive / restore location overrides

1.4.14p2 (2019-01-20)

- Fix to ensure plugin calls server rest API as /catdv/api/... instead of /api/... to prevent failed calls due to redirection of POST requests in some environments

1.4.14p1 (2019-01-11)

- Update README with correct version of required worker plugin

1.4.14 (2018-11-13)

- Add plugin command that can be triggered via the rest API to generate clip data for a file archived outside CatDV
- Add config param 'catdv.s3archive_service.max_retry_delay' to limit the delay period between retries of waiting jobs.

1.4.12

- Add config param 'catdv.s3archive_service.purge_directories' to turn off purging of directories when moving / purging files.
- Update README with note on Glacier support
- Update README to cover installing and updating the archive worker plugin, now included as part of the plugin installation.

1.4.11

Upgrade

- If the 'catdv.s3archive_service.allow_override' server property is explicitly set to include 'location:Location:archive', modify it to 'archiveLocation:Location:archive'

Changes

- Add option to restore files to a specified location / directory (can be enabled as an override for restore)

1.4.9 (2018-09-14)

- Fix for error message on cancel job contains NullPointerException
- Fix for issue which causes a new FieldGroup for an archive plugin to be created each time the server is restarted
- Fixed bug causing purge commands to fail from the web UI
- Improve the error reporting when attempting to archive offline files

1.4.8betap3 (2018-07-11)

- Fix build / jar files

1.4.8betap2 (2018-07-03)

- Fix location override for archives triggered from worker
- Log start / end of transfer jobs outside debug mode
- Record archive failure in clip archive status if clip has never been archived
- Update standalone service instructions in README
(** above changes merged from version 1.4.3p1)
- Integrate with ServicePlugin framework
(** above change merged from version 1.4.3)
- Additional location mapping option to determine location on archive from Media Store paths, includes mediaStoreRelative and mediaStoreAbsolute.
- Fix to ensure that location on archive does not contain '\' separators, regardless of location mapping in use
(** above changes merged from patches 1.4.3betap3 and 1.4.3betap4)
- Manage service command: preserve new connection details even if they are not valid / updated, to ease entry
- Manage service command: obfuscate the secret key value
- Rename all packages written for plugins from squarebox.catdv to squarebox.plugins
- Fix for NPE when clip has media file but no metadata (introduced by fix in 1.4.5)
- Fix to eliminate spurious exception starting / stopping standalone service

1.4.5 (2018-05-03)

- Fix for rest.NotFoundException scheduling jobs: change method for checking job exists, to avoid server incompatibility

1.4.5beta

- Move packages in order to split out re-usable plugin SDK classes
- Fix to prevent NPE - on archiving / restoring, skip clips that have no source media (i.e. sequences)
- Fix for location override in WorkerBackupCommand

1.4.4beta (2018-03-29)

- Upgraded httpclient and httpcore libraries (required to support other archive plugin)

1.4.3p1

- Fix location override for archives triggered from worker
- Log start / end of transfer jobs outside debug mode
- Record archive failure in clip archive status if clip has never been archived
- Update standalone service instructions in README

1.4.3 (2018-06-08)

- Integrate with ServicePlugin framework

1.4.3betap4 (2018-05-02)

- Fix to ensure that location on archive does not contain '\ ' separators, regardless of location mapping in use
- Updated location mapping options to add mediaStoreRelative and mediaStoreAbsolute

1.4.3betap3 (2018-04-19)

- Additional location mapping option to determine location on archive from Media Store paths

1.4.3betap2 (2018-03)

- Fix to truncate job status before passing to Rest API

1.4.3betap1 (2018-03)

- Removed obsolete duplicate aws library
- Additional debug output for KMS encryption

1.4.3beta (2018-02)

- Add capability to use either client or server side KMS encryption. This is set via the config param 'catdv.s3archive_service.use_kms_encryption', which should now be set to 'client', 'server' or left unset for no encryption. For backwards compatibility, a value of 'true' sets server side encryption and a value of false no encryption.
- Add config param 'catdv.s3archive_service.allow_override' to enable override of glacier tier and archive location from the UI.
- Add config param 'catdv.s3archive_service.location_mapping' to enable batching of file archive locations to be turned off - i.e. subsequent archives replace the existing file(s), rather than creating new 'date / time stamped' copies.
- Fix to ensure server config values that affect the UI are picked up on startup when archive service is standalone.
- Fix exponential back off timing for archive job retries
- Add date last restored as an archive parameter

1.4.2betap1 (2018-01-22)

- Fix for archive service start

1.4.2beta (2018-01-09)

Upgrade

- run catdvarchiveplugins1.4.2.sql against the CatDV DB to rename a job data field and to update the textual job and clip archive status values for consistency

Changes

- Add capability to run plugin archive service standalone, via config param catdv.s3archive_service.service_mode
- Add capability to restore from Glacier deep storage. If the object being restored is in deep storage, a request is made to transfer the object from deep storage to S3, where it will be available for a specified number of days (2 days by default, can be configured via catdv.s3archive_service.glacier_restore_expiration). The archive service will attempt to periodically retry the restore job until the object is available on S3 and can be fully restored or the maximum number of retries has been attempted.
- Add capability to use AWS KMS to encrypt objects in transit to / from and at rest on Amazon S3. This requires AWS KMS to be set up so that the SDK can get the Customer Master Key(CMK) for encryption. Set the new config param 'catdv.s3archive_service.use_kms_encryption' to 'true' to turn KMS encryption on and ensure that a CMK ID is provided via the Manage service command.
- Provision for providing more detailed status information for jobs and (archiving) files
- Add detailed job status and job data to Job Details pane in the Service Job Queue UI
- Fix for archive failure due to special characters in the file path / name (such as ™), with an error message of "The request signature we calculated does not match the signature you provided". The original file path in the object metadata is now url encoded on archive and decoded on metadata retrieval.
- New config param 'catdv.s3archive_service.max_jobs_to_display' for configuring the maximum number of jobs which will be listed in the job queue. This overrides days_to_display and defaults to 1000
- Increased default value for 'catdv.s3archive_service.loop_delay' to 30 seconds
- Increased default value for 'catdv.s3archive_service.concurrent_transfer_limit' to 4
- Increased default value for 'catdv.s3archive_service.days_to_display' to 10
- Updated Amazon AWS for Java library to version 1.11.221

1.4.1betap6 (2017-12-22)

- Fixes for worker plugin commands

1.4.1betap5 (2017-12-20)

- Return json response message from WorkerPurgeCommand
- Include metaclipID where applicable in error details for worker backup / restore / purge commands

1.4.1betap4 (2017-12-12)

- Fix bulk backup and restore commands to capture and report unexpected errors queuing child jobs
- Update worker backup and restore commands to return a JSON representation of the archive result in the command response.

1.4.1betap3 (2017-12-01)

- Fix to ignore clips with duplicate mapped file paths, before processing to queue jobs

1.4.1betap2

- New hidden versions of backup (copy/move) and purge commands for use by worker

1.4.1betap1 (2017-11-16)

- Fix to ensure archive metadata for source media with matching media paths are updated simultaneously
- Additional service logging, new config param 'catdv.s3archive_service.debug' now turns on/off most logging after initial plugin startup

1.4.1beta (2017-06-09)

Upgrade

- remove 'catdv.rest_api.*' config properties in Server Config on CatDV control panel
- run catdvarchiveplugins1.3.5.sql against the CatDV DB to fix the field type for archive date fields

Changes

- major update to job queue command UI, utilising new features provided by version 3 of plugin framework
- make API calls in process (inside server) from plugin commands and from archive service when running inside server (depends on version 3 of plugin framework)

1.3.5beta (2017-06-09)

Upgrade

- change config property 'catdv.rest_api.client_key' to 'catdv.s3archive_service.licence_code' in Server Config on CatDV control panel
- run catdvarchiveplugins1.3.5.sql against the CatDV DB to fix the field type for archive date fields

Changes

- improve job processing loop so that job priorities are always respected: only take one job from the queue at a time, when waiting job delays have elapsed re-queue jobs rather than processing waiting jobs directly (note re-queue jobs of equal priority will be processed before newer jobs, as the lowest job id will be processed first)
- ensure that transfers to S3 never block the job processing loop, keeping the connection status up to date
- improve processing of multiple concurrent transfers plus config param max_jobs_running_delay for optimisation
- improve handling when S3 is unavailable: after about 30s the service status updates to "Running (offline)"
- restart orphaned "in progress" jobs (e.g. from a server restart) more quickly
- switch to using statusCode as primary status value (except job queries which requires latest server)
- add support for pausing a service (processing transfers)
- improve plugin licence handling
- update to version 1.11.127 of Amazon AWS Java SDK
- script to fix archive date fields in desktop client for existing installations
- fix NPE when clip has no metadata

1.3.4beta3 (2017-03-02)

- Option to specify destination location 'path' when copying / moving files to Amazon S3 from the worker
- Add config param "catdv.s3archive_service.restrict_command_access" which can be used to hide 'all' plugin commands or 'config' commands (currently Manage Service). See README for details.

- fix archive date fields in desktop client for new installations

1.3.4beta2 (2017-03-02)

- Option to enter bucket name when copying / moving files to Amazon S3. The bucket name on the Manage Service screen is now only the default value filled in the first time a user attempts to copy or move files. Once the user has entered a bucket, a subsequent move / copy will default to the last value entered. The bucket warning and "Apply bucket to queued jobs" checkbox are no longer on the Manage Service screen as changes there now only affect the default bucket for a new user of the plugin.
- to support HGST/StrongBox archives, add new config param `catdv.s3archive_service.signer_override` and upgrade to `aws-java-sdk-1.11.52` (along with dependencies: `joda-time-2.9.5.jar`, `httpcore-4.4.5.jar`, `httpclient-4.5.2.jar`)

1.3.3 (2016-12-07)

Upgrade

- generate a rest api licence code registered to the customer for client 'ARCS3' and add to the configured properties:
`catdv.rest_api.client_key=<generated_licence_code>`
- run `catdvs3plugin1.3.3.sql` against the CatDV DB to set status codes for existing Jobs
- check in web admin that there are not two Amazon S3 archive panels
 - if there are, delete the one that is not referenced by the data field of the service table entry for 'AmazonS3'

Changes

- use a separate license pool for each archive plugin's access to the Rest API
- fixes for archiving of complex media
- complex media archive status reflects the archive status of it's contents (clip media 'Archive Details' field now updated along with the archive values stored in clip media metadata)
- rename package `squarebox.util` to `squarebox.archiveplugin.util` to avoid name clashes in the server
- add more detail to successful output for job results
- change location of files on S3 archive: `<batch id>/<actual file path>`
 - the batch id is the date/time that a selection of files was scheduled for archive to S3
 - ** this means files on the archive will no longer be overwritten, but may have multiple copies
- propagate changes to archive metadata to all source media with the same value for `mediaPath`
- set status code on archive Jobs along with text status
- purge empty directories along with files when using the 'Schedule move...' or 'Purge files...' commands
 - the top two directory levels are always preserved even if empty, e.g. `'/volumes/media/'`
- automatically create field definition(s) for any new archive metadata and add them to the plugin's archive panel
- throw a meaningful error if any of the properties required by the RestAPI are not set

1.3.2p1 (2016-09-21)

- patch for `NullPointerException` when scheduling transfers for clips with no existing metadata

1.3.2 (2016-09-08)

Upgrade

- Run `catdvs3plugin1.3.2.sql` against the CatDV DB to update the Job types

Changes

- fix for archiving mapped paths
- changed plugin command names and job types for clarity

1.3.1 (2016-08-10)

Initial production version

** If you previously had a beta version installed:

- any existing archive jobs should be removed from the database
- any manually created archive panel should be removed
- any existing archive data will be ignored (field names now prefixed with service type)