



Before you start

In order to use this BackBlaze archiving plugin you need:

- CatDV Server 8.0.4 or later
- CatDV Pegasus 13.0.5 or later
- CatDV Plugin license 'ARCBB' (Rest API license with multiple sessions)

To trigger BackBlaze file transfers from the Worker you need:

- BackBlaze Worker Plugin 2.2.1, included in this installation as BackBlazeWorker.catdv

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

- CatDV Service Control Panel 1.3 or later

New Installation

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

- 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:
 - 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 always be the default account, used to verify that BackBlaze is available and used as a fallback when necessary. See [Manage service command](#) for information on the details and settings for a service account:
 - In the client, run Tools->Manage BackBlaze 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.
 - 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 BackBlaze 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 BackBlaze to include the path of the storage volume
- If running the 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 BackBlaze Archive Service again. The service status should be 'Running (online)'. The status may be 'Running (offline)' if BackBlaze is not currently accessible.
- IF the Worker IS NOT being installed
 - Delete BackBlazeWorker.catdv from the BackBlaze plugin directory (see first step)
- IF the Worker IS being installed:
 - Move BackBlazeWorker.catdv FROM the BackBlaze plugin directory installed in 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 (Edit Config / CatDV Server / Reconnect) so that the archive plugin fields are available.
- Verify that the BackBlaze worker actions are available. It may be necessary to quit and re-open the worker to pick up the new worker actions.

Upgrade Installation

- 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 BackBlaze 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 BackBlaze Archive Service again. The service status should be 'Running (online)'. The status may be 'Running (offline)' if BackBlaze is not currently accessible.
- Update (or delete) the BackBlaze worker plugin by following the instructions from the last step for a new installation. If there is an older BackBlaze 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.backblaze_archive.licence_code = <generated licence code for plugin>
```

When running the archive service standalone, the property file for the archive service must include:

```
catdv.backblaze_archive.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.backblaze_archive.debug = true
```

Running archive service standalone

By default, the service that handles archive jobs runs inside the plugin. 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
- Replace the Java 'cacerts' certificates file in the service control panel JVM with the one from the server JVM
 - Replace: <CatDV Service Home>/jre/lib/security/cacerts
 - With: <CatDV Server Home>/jre/lib/security/cacerts
- 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.

BackBlaze Configuration

In order to use the plugin, a BackBlaze account is required. If you need to set up a BackBlaze account for testing purposes, got to <https://www.backblaze.com/b2/sign-up.html> to create an account with 10GB free.

For details of pricing see: <https://help.backblaze.com/hc/en-us/articles/217667478-Understanding-B2-Pricing-Structure>

Once you are logged into a BackBlaze account, the 'My Account' link will take you to the BackBlaze console: https://secure.backblaze.com/b2_buckets.htm

In order for the plugin to copy to / move to / restore from your BackBlaze account, you will need to enable B2, set up an application key and create at least one bucket.

To enable B2 cloud storage:

- To enable B2 on an existing BackBlaze account, first login to the BackBlaze account with the registered email address and password.
- Once logged in, select the 'My Settings' link from the left side navigation.
- At the bottom of the 'My Settings' page B2 can be enabled under the 'Enabled Products' section.

To get a new application key:

- From the 'My Account' page (https://secure.backblaze.com/b2_buckets.htm), click 'App Keys'
- Click 'Add a New Application Key', enter the appropriate details and click 'Create New Key'.
- Use 'Copy to Clipboard' to copy the application key and store it safely as it will only be displayed once

To create a bucket:

- Go to the 'My Account' page (https://secure.backblaze.com/b2_buckets.htm)
- Click the link 'Create a bucket'
- Enter the bucket name, choose private or public and click 'Create a Bucket'.
 - Files in a private bucket can only be accessed via the console or using the application key.
 - Files in a public bucket can be downloaded by anyone, but the downloads will be charged against your account.
 - Note bucket names must be globally unique across all B2 cloud storage accounts.

The values required to configure the service in the BackBlaze plugin are:

- Account ID - account id from a BackBlaze account
- Application key - application key from a BackBlaze account
- Default bucket - name of an accessible bucket on BackBlaze which can be used to verify the connection, used as the default bucket the first time a user attempts to copy / move files to BackBlaze

Media Store Service Path Mappings

Media Stores should be used to automatically map file storage archive locations by adding a service path to each media store. The service path determines the service account, container 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 BackBlaze a service path has the format:

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

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

```
/Volumes/dance-videos/drafts  
backblaze://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 buckets or distinct folders within the same 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 copy / move / 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 BackBlaze Archive Service	View full status details of the archive service and manage the service account(s) required to access the B2 storage. Service account details include account id, application key and default bucket name. If a connection to BackBlaze cannot be made an error is displayed and the values are not saved.
View BackBlaze Archive Service 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.
Schedule copy to BackBlaze	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 BackBlaze is accessible). When the copy job is run, the source media associated with the clip is copied from the file system to BackBlaze and the original file is preserved.
Schedule move to BackBlaze	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 BackBlaze	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 BackBlaze is accessible). When the restore job is run, the source media associated with the clip is copied from BackBlaze to the file system.
Purge files archived to BackBlaze	Deletes the source media associated with the selected clips if they have been successfully archived.

Managing Service Accounts

The accounts tab on Tools->Manage BackBlaze Archive Service provides facilities to manage the service account(s) used to connect to BackBlaze 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](#)).

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.

Delete	<p>Deletes the selected account.</p> <p>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.</p>
--------	--

Manage service command

The following can be configured from the Tools->Manage BackBlaze 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-numeric and hyphens. Note it is not possible to change the account identifier on an in-use account.
Account ID	Account ID for connecting to this service account
Application Key	Application 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.
Concurrent Request Count	The number of parts (of a multi-part upload) that the BackBlaze SDK will attempt to upload simultaneously. The default value is 10. Increase this value in order to upload a single large file quickly.
ARCHIVING TAB:	
Location mapping	<p>Note that this is mainly for legacy purposes as the preferred approach is now to use Media Store Service Path Mappings.</p> <p>Fallback method for generating the archive file location on BackBlaze if no media store service mapping applies. The default is 'mirror'. Valid values are:</p> <p><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.</p> <p><i>batchMirror</i>: The file(s) selected for archive are batched together in a date and time stamped directory (format /yyyy-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 BackBlaze as a new copy is written for each transfer.</p> <p><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.</p> <p><i>mediaStoreAbsolute</i>: As 'mediaStoreRelative' but also prepends the path with the name of the media store.</p>
Purge directories	Determines whether or not empty directories are deleted when purging (or moving) files.
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 in milliseconds)	Limits the maximum time period between retries of waiting jobs. The default value is blank.
Loop delay (time period in milliseconds)	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 in milliseconds)	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.

Stalled job delay	Determines the time period since a running job's last progress update, after which its status is flagged as stalled. Defaults to 1m
UI TAB:	
Restrict command access	Restricts the specified plugin commands to sys admin users. Can be used to hide 'config' commands only (i.e. Manage Service), 'config / archive' (i.e. to restrict copy or move to archive but not restore) or 'all' commands. Note that a server restart is required to ensure that this change is picked up under all circumstances. Default is 'none'.
Days to display	The number of days into the past for which jobs are listed in the job queue. Any job updated in this time period is listed. The default value is 10.
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.

**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 "BackBlaze Archive" containing the clip metadata which describes it's BackBlaze archive state, including:

Field	Description
squarebox.catdv.archive.BackBlaze.serviceType	Type of service responsible for file transfer
squarebox.catdv.archive.BackBlaze.serviceName	Name of service responsible for file transfer
squarebox.catdv.archive.BackBlaze.status	Archive status
squarebox.catdv.archive.BackBlaze.location	Location file will be (if job pending) or was last restored to
squarebox.catdv.archive.BackBlaze.restoreLocation	Location of file in storage
squarebox.catdv.archive.BackBlaze.date	Date (timestamp) of latest change in archive status
squarebox.catdv.archive.BackBlaze.dateLastArchived	Date last archived
squarebox.catdv.archive.BackBlaze.dateLastRestored	Date last restored
squarebox.catdv.archive.BackBlaze.numArchives	The number of times the clip has been successfully archived
squarebox.catdv.archive.BackBlaze.archiveKey	Identifier of file in storage
squarebox.catdv.archive.BackBlaze.batchID	Identifies the batch of files with which the clip was archived
squarebox.catdv.archive.BackBlaze.jobID	Identifier of current / latest archive job
squarebox.catdv.archive.BackBlaze.parentJobID	n/a for BackBlaze (related to bulk archives)
squarebox.catdv.archive.BackBlaze.userId	ID of user initiating current / latest transfer
squarebox.catdv.archive.BackBlaze.historyJson	Record of all archive activity in jsonFormat
squarebox.catdv.archive.BackBlaze.history	Record of all archive activity
squarebox.catdv.archive.BackBlaze.purgeError	Details of purge failure
squarebox.catdv.archive.BackBlaze.accountIdentifier	Identifier of the BackBlaze Storage Account in CatDV
squarebox.catdv.archive.BackBlaze.bucketName	Name of bucket to transfer file to / from

Known Issues

- Restore jobs don't give progress updates. This is a limitation of the B2 Java SDK.
- A network outage causes the archive service to block for 2 minutes. This is a limitation of the B2 Java SDK;

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.

March 2020

Release notes

2.2.6 (2020-03-11)

Upgrade (READ IN FULL FIRST)

- Now require CatDV Server 8.0.4 or later and CatDV Pegasus 13.0.5 or later (for transfer progress notifications)
- If running service standalone, update CatDV Service Control Panel to 1.3
- If required, update worker plugin to 2.2.1
- If required, need to set an undocumented config property to enable multiple accounts, as it is an add on product
- backup the service table from the DB if it contains multiple rows
- run `catdvbackblazeplugin2.2.6.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' BackBlaze, 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).
- The plugin will automatically migrate the connection details from the service definition currently in use (whether in server or standalone) 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 identifier of an account which is 'in use' cannot be changed (NB - The specified account identifier may contain only alphanumeric characters and hyphens):
`catdv.backblaze_archive.migration_account_identifier = <my-account-identifier>`
- **NB It is recommended to at least set this to 'master', as the 1.x versions were limited to the master application key, whereas 2.x can use other keys, including those restricted to specific bucket(s). Customers are therefore likely to want to make a new account the default. 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 and/or service panel config.
- **NB 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.backblaze_archive.licence_code`
`catdv.backblaze_archive.debug`
- In the BackBlaze panel, ensure "BackBlaze Archive History" (identifier `squarebox.catdv.archive.BackBlaze.history`) is below the new "BackBlaze Archive History (json)" (identifier `squarebox.catdv.archive.BackBlaze.historyjson`). Ensure these fields are the last two in the panel as the json history field displays as a table.

Changes (MAJOR UPDATE)

- Upgrade to BackBlaze API V2 / Java SDK 3.1.0, which enables multiple application keys. (Changed method to validate bucket, as cannot list buckets with a restricted key)
- Rename BackBlaze "Account ID" to "Application Key ID" in plugin user interface
- 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. `backblaze://squarebox-test-2/vcvideotest/dance`
- Implement restore path mapping
- 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.
- Add ability to restrict access to (config and) archive but allow restore
- Send job progress notifications, using new notification system (requires server >= 7.4.2 to make use of them)
- Include file size in job data / job queue list
- Support multi-select for cancelling jobs from queue
- Run another job (if available) immediately after completion / failure of a running job
- Add json version of archive history to clip archive metadata
- Update service fields / panel creation to include all fields in field group but exclude some from the panel
- 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)
- Change default location mapping for new installs to 'mirror', as batching needs to be phased out with media store mapping
- On job queue ensure job details/results fields are cleared when job table empties due to refresh or cancel
- Fix to prevent job exceeding maxRetries

1.5.0p2 (2019-07-29)

Upgrade

- run `catdvarchiveplugins1.5.sql` against the CatDV DB to fix the field type for archive date fields
- If the `'catdv.backblaze_archive.allow_override'` server property is explicitly set to include `'location:Location:archive'`, modify it to `archiveLocation:Location:archive'`

Changes

- 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
- Fix NPE on cancel job, when associated clip is not found
- Changes to purge (apply to move to archive as well as direct purge): clear date last restored, log 'File purged' to clip archive params history, add option of "Yes to all" confirmation for purging multiple altered files.
- Fix blocking mechanism for ensuring that multiple processes cannot update a job's status simultaneously
- Prevent spurious 'Network outage'/NPE error when no restore directory can be extracted from the restore location
- 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
- Improve performance of Job Queue command for large lists of jobs
- Fix to cause job creation to fail if no source media are updated with the archive details for the job
- Fix for issue of jobs stuck in running / stalled state
- Trim trailing path separators from archive / restore location overrides
- 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
- 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.backblaze_archive.max_retry_delay' to limit the delay period between retries of waiting jobs.
- Add config param 'catdv.backblaze_archive.purge_directories' to turn off purging of directories when moving / purging files.
- Update README to cover installing and updating the archive worker plugin, now included as part of the plugin installation.
- 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
- Improve the error reporting when attempting to archive offline files

1.4.8betap4 (2018-07-25)

- Fixed bug causing purge commands to fail from the web UI

1.4.8betap3 (2018-07-11)

- Fix build / jar files

1.4.8betap2

- Fix location override for archives triggered from worker
- Record archive failure in clip archive status if clip has never been archived

1.4.8beta

- Integrate with ServicePlugin framework
- Manage service command: preserve new connection details even if they are not valid / updated, to ease entry
- Manage service command: obfuscate the application key value

1.4.6p1 (2018-05-29)

- Minimize class C BackBlaze transactions

1.4.6 (2018-05-09)

- 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
- Fix for NPE when clip has media file but no metadata (introduced by fix in 1.4.5)

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 (2018-04-5)

- Changed verify BackBlaze connection to use downloadByName with invalid bucket/file name, to reduce api calls/cost.
- Update rest API wrapper to use version 7 of the API rather than version 5
- 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

- Initial beta version (2018-03-29)