



Deployment Manual

CoreMedia Content Cloud – v13

Copyright CoreMedia GmbH © 2026

CoreMedia GmbH

Altes Klöpperhaus, 5. OG

Rödingsmarkt 9

20459 Hamburg

International

All rights reserved. No part of this manual or the corresponding program may be reproduced or copied in any form (print, photocopy or other process) without the written permission of CoreMedia GmbH.

Germany

Alle Rechte vorbehalten. CoreMedia und weitere im Text erwähnte CoreMedia Produkte sowie die entsprechenden Logos sind Marken oder eingetragene Marken der CoreMedia GmbH in Deutschland. Alle anderen Namen von Produkten sind Marken der jeweiligen Firmen.

Das Handbuch bzw. Teile hiervon sowie die dazugehörigen Programme dürfen in keiner Weise (Druck, Fotokopie oder sonstige Verfahren) ohne schriftliche Genehmigung der CoreMedia GmbH reproduziert oder vervielfältigt werden. Unberührt hiervon bleiben die gesetzlich erlaubten Nutzungsarten nach dem UrhG.

Licenses and Trademarks

All trademarks acknowledged.
July 22, 2026 (Release 2512.1)

1. Preface	1
1.1. Audience	2
1.2. Typographic Conventions	3
1.3. CoreMedia Services	5
1.3.1. Registration	5
1.3.2. CoreMedia Releases	6
1.3.3. Documentation	7
1.3.4. CoreMedia Training	10
1.3.5. CoreMedia Support	10
1.4. Changelog	12
2. Docker Setup	13
2.1. Container Images Build Process	14
2.1.1. Google Jib	14
2.1.2. Maven Structure	14
2.1.3. The java-application-base Image	17
2.2. Docker Container Startup	24
2.2.1. Startup Entrypoint and Command chain	24
2.2.2. Health Check	24
2.3. Container Configuration	26
2.3.1. Build Time Configuration	26
2.3.2. Start Time Configuration	26
2.3.3. Runtime Configuration	27
2.4. Containerized Tools	28
2.4.1. Running the Tools	28
2.4.2. Configuring the Tools	30
3. CoreMedia Properties Overview	33
3.1. Content Application Engine Properties	34
3.1.1. General CAE Properties	34
3.1.2. Delivery CAE Properties	42
3.1.3. Http Cache Control Properties	43
3.2. Content Server Properties	47
3.2.1. General Content Server Properties	47
3.2.2. CORBA Properties	59
3.2.3. Properties for the Publisher	64
3.2.4. Properties for the Connection to the Database	68
3.2.5. Properties for Replicator Configuration	79
3.2.6. Properties for Timezone and IOR	84
3.3. Headless Server Properties	85
3.3.1. Headless Server Spring Boot Properties	85
3.3.2. Persisted Query Properties	96
3.3.3. Metadata Properties	98
3.3.4. Remote Service Adapter Properties	99
3.3.5. Headless Server Cache Control Properties	99
3.3.6. Headless Server Cache Key Properties	101
3.3.7. Properties of External Frameworks	102
3.4. Studio Properties	103
3.4.1. Studio Configuration	103
3.4.2. Available Locales Configuration	118
3.4.3. Shared Multi-Site Configuration	118

3.4.4. Content Configuration	120
3.4.5. Navigation Validator Configuration	121
3.4.6. Preview URL Service Properties	122
3.4.7. Content Security Policy Configuration	123
3.4.8. Content Hub Configuration	125
3.4.9. Feedback Hub Configuration	126
3.4.10. Editorial Comments Configuration	127
3.4.11. Notifications SQL Persistence Configuration	130
3.4.12. Projects/To-Dos SQL Persistence Configura- tion	134
3.4.13. Projects/To-Dos Migration Configuration	140
3.4.14. Commerce Related Configuration	141
3.5. User Changes Application Properties	143
3.6. Workflow Server Properties	145
3.6.1. General	145
3.6.2. Shared Multi-Site Configuration	158
3.7. Commerce Hub Properties	161
3.8. My Edited Content and Workflow Lists Properties	162
3.8.1. General Properties	162
3.8.2. Data Source Properties	162
3.8.3. Event Handling Properties	163
3.8.4. Data Migration Properties	165
3.9. Elastic Social Properties	168
3.9.1. General Elastic Social Properties	168
3.9.2. MongoDB Properties	170
3.9.3. Counter Properties	171
3.9.4. Task Queue Properties	172
3.9.5. Elastic Social Link Building Properties	173
3.9.6. Elastic Social Solr Properties	174
3.10. Content Feeder Properties	178
3.11. CAE Feeder Properties	206
3.12. UAPI Client Properties	233
3.12.1. Unified API Spring Boot Client Properties	233
3.13. Cache Properties	240
3.14. Web Security Properties	242
3.14.1. CORS Properties	242
3.15. Plugin Manager Properties	245
3.16. Blob Transformer Properties	246
3.17. Image Transformation Properties	250
3.18. Engagement Cloud Integration Properties	252
3.19. gRPC Properties	256
4. Encryption Service Setup	258
4.1. Plugin Based Encryption Service Setup	259
4.2. Java Keystore based encryption service	260
4.2.1. Prerequisites	260
4.2.2. Configuration Properties Setup	261
4.2.3. Properties File Setup	262
Index	264

List of Tables

1.1. Typographic conventions	3
1.2. Pictographs	4
1.3. CoreMedia manuals	7
1.4. Changes	12
3.1. Configuration Properties with Prefix cae	34
3.2. Delivery Properties	42
3.3. Configuration Properties with Prefix cache.control	43
3.4. Content Server Properties	47
3.5. CORBA Properties	59
3.6. Publisher Properties	64
3.7. SQL Properties	68
3.8. Replicator Properties	79
3.9. capclient.properties	84
3.10. Headless Server Properties	85
3.11. Persisted Query Properties	96
3.12. Metadata Root Properties	98
3.13. Headless Server Remote Properties	99
3.14. Headless Server Cache Control Properties	99
3.15. Headless Server Cache Key Properties	101
3.16. Headless Server External Framework Properties	102
3.17. Studio Properties	103
3.18. Available Locales Properties	118
3.19. Shared Multi-Site Properties	158
3.20. Content Properties	120
3.21. Navigation Validators Properties	121
3.22. Preview URL Service Properties	122
3.23. Content Security Policy Related Studio Properties	123
3.24. Content Hub Properties	125
3.25. Feedback Hub Properties	126
3.26. Editorial Comments Properties	127
3.27. Notification SQL Properties	130
3.28. Project SQL Properties	134
3.29. Project Migration Properties	140
3.30. Commerce Related Properties	141
3.31. User Changes App Properties	143
3.32. Workflow Server Properties	145
3.33. Shared Multi-Site Properties	158
3.34. Commerce Hub Properties	161
3.35. General Configuration Properties for CapLists	162
3.36. Data Source Configuration Properties for CapLists	163
3.37. Configuration Properties for Event Handling of CapLists	163
3.38. Configuration Properties for Migration of CapLists from Elastic Core to SQL Persistence	166
3.39. Elastic Social Properties	168
3.40. MongoDB Properties	170
3.41. Counters Properties	171
3.42. Task-Queues Properties	172

3.43. Elastic Social Link Building Properties	173
3.44. Elastic Solr Properties	174
3.45. Content Feeder Configuration Properties	178
3.46. Content Feeder Solr Configuration Properties	192
3.47. Feeder Batch Configuration Properties	229
3.48. Feeder Tika Configuration Properties	220
3.49. Feeder Core Configuration Properties	232
3.50. CAE Feeder Properties	206
3.51. CAE Feeder Data Source Properties	216
3.52. CAE Feeder Database Locking Properties	218
3.53. Feeder Tika Configuration Properties	220
3.54. CAE Feeder Solr Configuration Properties	224
3.55. Feeder Batch Configuration Properties	229
3.56. Feeder Core Configuration Properties	232
3.57. UAPI Spring Boot Client Properties	233
3.58. Cache Properties	240
3.59. Configuration Properties with Prefix cors	242
3.60. Plugin Manager Properties	245
3.61. Blob Transformer Properties	246
3.62. Further Configuration Properties	248
3.63. Image Transformation Properties	250
3.64. Configuration Properties for the CoreMedia Engagement Cloud In- tegration	252
3.65. gRPC Client Properties	256
4.1. Plugin Based Encryption Service Properties	259
4.2. Key Store Encryption Service Properties	261

1. Preface

1.1 Audience

This manual is intended for architects and developers who want to learn about the concepts of CoreMedia's flexible deployment. You should be familiar with configuration management and application deployment concepts.

Additionally, you should be familiar with *CoreMedia CMS*, *Spring Framework* and *Apache Maven*.

1.2 Typographic Conventions

CoreMedia uses different fonts and types in order to label different elements. The following table lists typographic conventions for this documentation:

Element	Typographic format	Example
Source code	Courier new	cm systeminfo start
Command line entries		
Parameter and values		
Class and method names		
Packages and modules		
Menu names and entries	Bold, linked with	Open the menu entry Format Normal
Field names	Italic	Enter in the field <i>Heading</i>
CoreMedia Components		The <i>CoreMedia Component</i>
Applications		Use <i>Chef</i>
Entries	In quotation marks	Enter "On"
(Simultaneously) pressed keys	Bracketed in "<>", linked with "+"	Press the keys <Ctrl>+<A>
Emphasis	Italic	It is <i>not</i> saved
Buttons	Bold, with square brackets	Click on the [OK] button
Code lines in code examples which continue in the next line	\	cm systeminfo \ -u user

Table 1.1. Typographic conventions

In addition, these symbols can mark single paragraphs:

Pictograph	Description
	Tip: This denotes a best practice or a recommendation.
	Warning: Please pay special attention to the text.
	Danger: The violation of these rules causes severe damage.

Table 1.2. Pictographs

1.3 CoreMedia Services

This section describes the CoreMedia services that support you in running a CoreMedia system successfully. You will find all the URLs that guide you to the right places. For most of the services you need a CoreMedia account. See [Section 1.3.1, "Registration" \[5\]](#) for details on how to register.

NOTE

CoreMedia User Orientation for CoreMedia Developers and Partners

Find the latest overview of all CoreMedia services and further references at:

<http://documentation.coremedia.com/new-user-orientation>



- [Section 1.3.1, "Registration" \[5\]](#) describes how to register for the usage of the services.
- [Section 1.3.2, "CoreMedia Releases" \[6\]](#) describes where to find the download of the software.
- [Section 1.3.3, "Documentation" \[7\]](#) describes the CoreMedia documentation. This includes an overview of the manuals and the URL where to find the documentation.
- [Section 1.3.4, "CoreMedia Training" \[10\]](#) describes CoreMedia training. This includes the training calendar, the curriculum and certification information.
- [Section 1.3.5, "CoreMedia Support" \[10\]](#) describes the CoreMedia support.

1.3.1 Registration

In order to use CoreMedia services you need to register. Please, start your [initial registration via the CoreMedia website](#). Afterwards, contact the CoreMedia Support (see [Section 1.3.5, "CoreMedia Support" \[10\]](#)) by email to request further access depending on your customer, partner or freelancer status so that you can use the CoreMedia services.

1.3.2 CoreMedia Releases

Downloading and Upgrading the Blueprint Workspace

CoreMedia provides its software as a Maven based workspace. You can download the current workspace or older releases via the following URL:

<https://releases.coremedia.com/cmcc-13>

Refer to our [Blueprint Github mirror repository](#) for recommendations to upgrade the workspace either via Git or patch files.

NOTE

If you encounter a 404 error then you are probably not logged in at GitHub or do not have sufficient permissions yet. See [Section 1.3.1, "Registration" \[5\]](#) for details about the registration process. If the problems persist, try clearing your browser cache and cookies.



Maven artifacts

CoreMedia provides parts of its release artifacts via Maven under the following URL:

<https://repository.coremedia.com>

You have to add your CoreMedia credentials to your Maven settings file as described in [section 3.1, "Prerequisites"](#) in *Blueprint Developer Manual*.

npm packages

CoreMedia provides parts of its release artifacts as npm packages under the following URL:

<https://repository.coremedia.com/nexus/repository/coremedia-npm/>

The `.npmrc` file must be configured to be able to use the registry (see [Section 3.1, "Prerequisites"](#) in *Blueprint Developer Manual*).

License files

You need license files to run the CoreMedia system. Contact the support (see [Section 1.3.5, "CoreMedia Support" \[10\]](#)) to get your licences.

1.3.3 Documentation

CoreMedia provides extensive manuals, how-tos and Javadoc as PDF files and as online documentation at the following URL:

<https://documentation.coremedia.com>

The manuals have the following content and use cases:

Manual	Audience	Content
Blueprint Developer Manual	Developers, architects, administrators	This manual gives an overview over the structure and features of <i>CoreMedia Content Cloud</i>. It describes the content type model, the <i>Studio</i> extensions, folder and user rights concept and many more details. It also describes administrative tasks for the features. It also describes the concepts and usage of the project workspace in which you develop your CoreMedia extensions. You will find a description of the Maven structure, the virtualization concept, learn how to perform a release and many more.
Connector Manuals	Developers, administrators	This manual gives an overview over the use cases of the eCommerce integration. It describes the deployment of the Commerce Connector and how to connect it with the CoreMedia and eCommerce system.
Content Application Developer Manual	Developers, architects	This manual describes concepts and development of the <i>Content Application Engine (CAE)</i>. You will learn how to write Freemarker templates that access the other CoreMedia modules and use the sophisticated caching mechanisms of the CAE.
Content Server Manual	Developers, architects, administrators	This manual describes the concepts and administration of the main CoreMedia component, the <i>Content Server</i>. You will learn about the content type model which lies at the heart of a CoreMedia system, about user and rights management, database configuration, and more.

Manual	Audience	Content
Deployment Manual	Developers, architects, administrators	This manual describes the concepts and usage of the CoreMedia deployment artifacts. That is the deployment archive and the Docker setup. You will also find an overview of the properties required to configure the deployed system.
Elastic Social Manual	Developers, architects, administrators	This manual describes the concepts and administration of the <i>Elastic Social</i> module and how you can integrate it into your websites.
Frontend Developer Manual	Frontend Developers	This manual describes the concepts and usage of the Frontend Workspace. You will learn about the structure of this workspace, the CoreMedia themes and bricks concept, the CoreMedia Freemarker facade API, how to develop your own themes and how to upload your themes to the CoreMedia system.
Headless Server Developer Manual	Frontend Developers, administrators	This manual describes the concepts and usage of the <i>Headless Server</i> . You will learn how to deploy the Headless Server and how to use its endpoints for your sites.
Multi-Site Manual	Developers, Multi-Site Administrators, Editors	This manual describes different options to design your site hierarchy with several languages. It also gives guidance to avoid common pitfalls during your work with the multi-site feature.
Operations Basics Manual	Developers, administrators	This manual describes some overall concepts such as the communication between the components, how to set up secure connections, how to start application.
Search Manual	Developers, architects, administrators	This manual describes the configuration and customization of the <i>CoreMedia Search Engine</i> and the two feeder applications: the <i>Content Feeder</i> and the <i>CAE Feeder</i> .
Studio Developer Manual	Developers, architects	This manual describes the concepts and extension of <i>CoreMedia Studio</i> . You will learn about the underlying concepts, how to use the devel-

Manual	Audience	Content
		opment environment and how to customize <i>Studio</i> to your needs.
Studio User Manual	Editors	This manual describes the usage of <i>CoreMedia Studio</i> for editorial and administrative work. It also describes the usage of the <i>Native Personalization</i> and <i>Elastic Social</i> GUI that are integrated into <i>Studio</i> .
Studio Benutzerhandbuch	Editors	The Studio User Manual but in German.
Supported Environments	Developers, architects, administrators	This document lists the third-party environments with which you can use the CoreMedia system, Java versions or operation systems for example.
Unified API Developer Manual	Developers, architects	This manual describes the concepts and usage of the <i>CoreMedia Unified API</i> , which is the recommended API for most applications. This includes access to the content repository, the workflow repository and the user repository.
Utilized Open Source Software & 3rd Party Licenses	Developers, architects, administrators	This manual lists the third-party software used by CoreMedia and lists, when required, the licence texts.
Workflow Manual	Developers, architects, administrators	This manual describes the <i>Workflow Server</i> . This includes the administration of the server, the development of workflows using the XML language and the development of extensions.

Table 1.3. CoreMedia manuals

If you have comments or questions about CoreMedia's manuals, contact the Documentation department:

Email: documentation@coremedia.com

1.3.4 CoreMedia Training

CoreMedia's training department provides you with the training for your CoreMedia projects either live online, in the CoreMedia training center or at your own location.

You will find information about the CoreMedia training program, the training schedule and the CoreMedia certification program at the following URL:

<https://www.coremedia.com/training>

Contact the training department at the following email address:

Email: training@coremedia.com

1.3.5 CoreMedia Support

CoreMedia's support is located in Hamburg and accepts your support requests between 9 am and 6 pm MET. If you have subscribed to 24/7 support, you can always reach the support using the phone number provided to you.

To submit a support ticket, track your submitted tickets or receive access to our forums visit the CoreMedia Online Support at:

<https://support.coremedia.com/>

Do not forget to request further access via email after your initial registration as described in [Section 1.3.1, "Registration" \[5\]](#). The support email address is:

Email: support@coremedia.com

Create a support request

CoreMedia systems are distributed systems that have a rather complex structure. This includes, for example, databases, hardware, operating systems, drivers, virtual machines, class libraries and customized code in many different combinations. That's why CoreMedia needs detailed information about the environment for a support case. In order to track down your problem, provide the following information:

Support request

- Which CoreMedia component(s) did the problem occur with (include the release number)?
- Which database is in use (version, drivers)?
- Which operating system(s) is/are in use?
- Which Java environment is in use?

- Which customizations have been implemented?
- A full description of the problem (as detailed as possible)
- Can the error be reproduced? If yes, give a description please.
- How are the security settings (firewall)?

In addition, log files are the most valuable source of information.

To put it in a nutshell, CoreMedia needs:

Support checklist

1. a person in charge (ideally, the CoreMedia system administrator)
2. extensive and sufficient system specifications
3. detailed error description
4. log files for the affected component(s)
5. if required, system files

An essential feature for the CoreMedia system administration is the output log of Java processes and CoreMedia components. They're often the only source of information for error tracking and solving. All protolling services should run at the highest log level that is possible in the system context. For a fast breakdown, you should be logging at debug level. See [Section 4.7, "Logging"](#) in *Operations Basics* for details.

Log files

Which Log File?

In most cases at least two CoreMedia components are involved in errors: the *Content Server* log files together with the log file from the client. If you know exactly what the problem is, solving the problem becomes much easier.

Where do I Find the Log Files?

By default, application containers only write logs to the console output but can be accessed from the container runtime using the corresponding command-line client.

For the *docker* command-line client, logs can be accessed using the **docker logs** command. For a detailed instruction of how to use the command, see [docker logs](#). Make sure to enable the timestamps using the `--timestamps` flag.

```
docker logs --timestamps <container>
```

For the *kubectrl* command-line client in a Kubernetes environment you can use the **kubectrl logs** command to access the logs. For a detailed instruction of how to use the command, see [kubectrl logs](#). Make sure to enable the timestamps using the `--timestamps` flag.

```
kubectrl logs --timestamps <pod>
```

1.4 Changelog

In this chapter you will find a table with all major changes made in this manual.

Section	Version	Description
Chapter 4, <i>Encryption Service Setup</i> [258]	Core-Media Content Cloud v11 – 11.2307.1	Added a new chapter about configuring the encryption service with the keystore based encryption service.

Table 1.4. Changes

2. Docker Setup

CoreMedia Content Cloud offers an application modularity that supports individual application lifecycles, including a modularized development, release and deployment process. Using Spring Boot standards and container technologies at its core, our goal is to enable our applications to become more cloud native and easier to integrate with available cloud technologies.

The docker setup provided with *CoreMedia Content Cloud* consists of:

- A Maven build process to create the images
- A Docker compose setup to start a development environment

To use *CoreMedia Content Cloud* in production with Docker images, you have one of the following choices:

- Use the Docker images to start containers and benefit from the process abstraction containers provide.
- Implement a Docker Swarm setup to benefit from the simplicity of a Swarm orchestration layer.
- Implement a Kubernetes Deployment to benefit from the flexibility and power, Kubernetes provides to create a cloud native cluster deployment.

Obviously Kubernetes is the most promising technology, but the initial complexity and the steep learning curve in addition with the requirement to build and maintain the knowledge of this technology may not fit each ones requirements.

To help you with this choice, this chapter will provide you with an overview of the build process of our Docker images, the conventions and configuration options the Docker images provide and how to use them for a production deployment using the containers as a simple process abstraction.

2.1 Container Images Build Process

To build the container images for Spring Boot applications with Maven, there are currently two plugins to consider:

- [Spring-Boot Maven plugin](#)
- [Google Jib Maven plugin](#)

Both provide a tight integration into Maven but are completely different from their approach. The Spring Boot Maven plugin uses [CloudNative Buildpacks](#) whereas the Jib plugin provides a more direct approach but with the benefit of not requiring Docker at all.

We chose the Jib plugin, because it allows us to provide more customizability without the need of a custom buildpack. It also allows us to use our custom base image to provide a seamless migration. See [java-application-base image \[17\]](#) for more information.

2.1.1 Google Jib

Google Jib integrates with Spring Boot applications and the image will be build by the same Maven module that builds the Spring-Boot application. To activate the image build process, the Maven profile `default-image` needs to be activated.

2.1.2 Maven Structure

2.1.2.1 Properties

```
<properties>
<jib.goal>dockerBuild</jib.goal> □
<application.image-base>gcr.io/distroless/java@11</application.image-base>
□
<application.image-prefix>coremedia</application.image-prefix> □
<application.image-suffix>my-app</application.image-suffix> □
<application.image-tag>${project.version}</application.image-tag> □
<application.image-arch>amd64</application.image-arch> □
</properties>
```

- ❶ The goal to build. Choose between either `dockerBuild` or `build`, visit the [jib documentation](#) for more details.

- ② The base image, here a Google distroless image. For deterministic build, the sha256 digest should be used. If you want to retrieve the base image from your local Docker daemon, you need to prefix the image source with `docker://`.
- ③ The image prefix. You should set your registry here.
- ④ The image suffix, that is the application name.
- ⑤ The tag of the image.
- ⑥ The architecture of the base image. In case a multi-arch OCI manifest is referenced, this determines the target platform image architecture.

2.1.2.2 Plugin

The Jib plugin has an extension mechanism and by default there are some extensions active. One of them is the Spring Boot extension, which takes care of reading the layer manifest of the Spring Boot plugin. The layered JARs mechanism was introduced with version 2.3 of Spring Boot to add application resources in layers, grouped by their nature to change. For more information, see the section about [Layered Jar](#) in the official plugin documentation.

Another extension, we make use of in our setup is the `jib-ownership-extension`. This extension can be used to set the ownership of files, see the [jib ownership extension documentation](#) for more details.

```
<plugin>
  <groupId>com.google.cloud.tools</groupId>
  <artifactId>jib-maven-plugin</artifactId>
  <version>3.4.0</version>
  <dependencies>
    <dependency>
      <groupId>com.google.cloud.tools</groupId>
      <artifactId>jib-ownership-extension-maven</artifactId>
      <version>0.1.0</version>
    </dependency>
  </dependencies>
  <configuration>
    <pluginExtensions>
      <pluginExtension>

<implementation>com.google.cloud.tools.jib.maven.extension.ownership.JibOwnershipExtension

      </implementation>
      <configuration>
implementation="com.google.cloud.tools.jib.maven.extension.ownership.Configuration">

        <rules>
          <rule>
            <glob>/coremedia**</glob>
            <ownership>1000:1000</ownership>
          </rule>
        </rules>
      </configuration>
    </pluginExtension>
  </pluginExtensions>
</configuration>
</plugin>
```

```
</configuration>
</plugin>
```

The configuration of the plugin is split up into two parts:

- A generic configuration
- The configuration to build the default image.

The Generic Configuration

```
<configuration>
  <from>
    <image>${application.image-base}</image> □
    <platforms>
      <platform>
        <os>linux</os>
        <architecture>${application.image-arch}</architecture>
      </platform>
    </platforms>
  </from>
  <to>
    <image>${application.image-prefix}/${application.image-suffix}</image>
  </to>
  <tags>
    <tag>${application.image-tag}</tag> □
  </tags>
  <container>
    <appRoot>/coremedia</appRoot> □
    <workingDirectory>/coremedia</workingDirectory> □
    <ports>
      <port>8080</port>
      <port>8081</port>
    </ports>
  </container>
</configuration>
```

- ❶ The base image to start from.
- ❷ The image to build. If that is a remote registry and goal is set to `build`, jib will build remotely without docker.
- ❸ The image tag.
- ❹ The application root, by default Jib will assign all files to root without write rights to others.
- ❺ The working directory all relative paths are calculated to.
- ❻ The ports that can be exposed.

The java-application-base Configuration

```
<configuration>
  <container>
    <entrypoint>INHERIT</entrypoint> □
    <args>
```

```
<arg>application</arg> □
</args>
<user>coremedia:coremedia</user> □
<environment> □
  <JAVA_HEAP>256m</JAVA_HEAP>
  <JAVA_PARALLEL_GC>true</JAVA_PARALLEL_GC> □
<SPRING_BOOT_EXPLODED_MAINCLASS>@jib-main-class-file</SPRING_BOOT_EXPLODED_MAINCLASS>
□
<SPRING_BOOT_EXPLODED_CLASSPATH>@jib-classpath-file</SPRING_BOOT_EXPLODED_CLASSPATH>
□
  </environment>
  <ports>
    <port>8080</port>
    <port>8081</port>
  </ports>
</container>
</configuration>
```

- ❶ With this directive, we tell Jib to inherit the entrypoint information from the base image. In this case it is `/coremedia/entrypoint.sh`.
- ❷ The argument to the entrypoint script. The `application` script encapsulates the startup logic. Visit the DockerHub page of the [coremedia/java-application-base](#) image for more details.
- ❸ The user to start the application. With this base image, we need to use the `coremedia` user to start the application.
- ❹ The environment variables the `application` entrypoint chain script can be configured.
- ❺ Set to `true` to activate the Parallel GC. See [Section 2.1.3.1, “Garbage Collector Configuration” \[18\]](#) for details.
- ❻ The application starter class we defined in the properties section above.
- ❼ The classloader classpath. This mimics the classloading, the Spring-Boot Jib extension automatically sets.

Extra files and directories

Some images require extra files and directories. A good example is the `workflow-server`. In this case and when using the `jib-compat` profile, we want to add the server-tools to start the `workflowconverter` in advance of starting the server itself. To do so we use the `extraDirectories` directive of the Jib plugin. For more details about how to copy additional files to the image visit the [official Jib documentation](#) about it.

2.1.3 The java-application-base Image

One of the major inventions Docker brought to software development is the ability to package applications as layers of filesystems, that when stacked upon

each other result in the final application image. This technique allows docker to reuse and share common parts between multiple images at build and startup time and reduces the disk and network footprint drastically.

In the CoreMedia Docker setup, this technique is used by providing a common base image on the public DockerHub repository. The `java-application-base` image is the common base image for all our Spring-Boot application based images. It encapsulates our best practice to configure, start and manage Core-Media applications in most secure way. It is provided in different flavors for different JDK base images on DockerHub and it is possible to fork the source repository on [GitHub](#), to build a customized image.

To summarize the capabilities of the image, it provides:

- A JVM to run our applications
- A startup script chain to integrate arbitrary initialization steps before starting the Spring-Boot application
- A set of convenient script chain steps to let applications delay their startup until service dependencies are available
- A monitoring integration for [Prometheus](#)

For a detailed description of the base image, please visit the `java-application-base` [DockerHub](#) page.

2.1.3.1 Garbage Collector Configuration

Default: Parallel GC

CoreMedia configures the **Parallel GC** (`-XX:+UseParallelGC`) as the default garbage collector for all containerized Spring Boot applications. This is expressed via the `JAVA_PARALLEL_GC=true` environment variable that the `java-application-base` image entrypoint evaluates at startup.

This default was introduced as a deliberate workaround for a [JDK bug](#) that causes G1GC to fail with an `OutOfMemoryError` under specific memory allocation patterns. The bug affects all JDK versions prior to JDK 25 and will not be backported to earlier release trains.

Why delivery components are primarily affected: The image transformation pipeline allocates large, contiguous heap regions whose size depends on the source image resolution and the requested transformation operations. Part of the transformation is executed via JNI (native image-processing libraries). While a JNI call is active, G1GC cannot free that memory region — if heap pressure is

high enough at that moment, the JVM raises an `OutOfMemoryError`. Parallel GC does not have this limitation and is therefore the safer default.

Parallel GC vs. G1GC

The two collectors make different trade-offs that operators should weigh against their deployment’s traffic profile and operational capabilities:

	Parallel GC (default)	G1GC (opt-in)
OOM risk	Low — JDK bug does not apply	Elevated on JDK < 25 (see above)
GC pause duration	Long — seconds on large heaps	Short and bounded by <code>MaxGC PauseMillis</code>
Responsiveness under load	Degraded during Full GC	Consistently good
Raw GC throughput	Higher	Slightly lower due to concurrent GC threads
Large image transformation	Safe	Risk of heap fragmentation (see tuning note below)
CoreMedia support	Full (default configuration)	Customer responsibility; not actively supported

G1GC is worth considering for deployments that run with a well-provisioned heap, have active GC monitoring in place, and can react quickly to an OOM event should one occur. It is especially relevant for:

- **Content Application Engine (CAE)** — live instances serving high-traffic sites
- **Headless Server** — live instances under sustained load

Switching to G1GC

To switch an application to G1GC, set `JAVA_PARALLEL_GC` to `false` and pass the desired G1 tuning flags via the `JAVA_OPTS` environment variable (or the equivalent `jvmFlags` in the distroless profile). The following example, derived from field experience with high-traffic CAE deployments, is a reasonable starting point:

```
<environment>
  <JAVA_HEAP>16g</JAVA_HEAP> □
  <JAVA_PARALLEL_GC>false</JAVA_PARALLEL_GC> □
  <JAVA_OPTS>
    -XX:+UseG1GC
    -XX:MaxGCPauseMillis=200
    -XX:+ParallelRefProcEnabled
    -XX:InitiatingHeapOccupancyPercent=45
    -XX:ParallelGCThreads=2
    -XX:ConcGCThreads=1
  </JAVA_OPTS> □
</environment>
```

- ❶ Heap size for a high-traffic deployment (4 cores, 16 GB). This represents the upper end of observed production configurations. Most CAE and Headless Server deployments operate in the 4–12 GB range.
- ❷ Disables the Parallel GC activation in the entrypoint script.
- ❸ Activates G1GC and tunes pause-time target, concurrent GC threads, and heap occupancy threshold. Adjust `ParallelGCThreads` and `ConcGCThreads` to match the actual number of available CPU cores.

Tuning G1GC for Image-Heavy Workloads

When the CAE or Headless Server processes high-resolution images, G1GC may allocate *Humongous Objects* — objects larger than half the G1 region size — directly in the old generation. Frequent Humongous allocations can cause back-to-back concurrent GC cycles and old-generation heap fragmentation, eventually leading to an `OutOfMemoryError`.

To understand the scale of the problem, consider a typical large source image: a 4000×2000 px PNG at 8-bit colour depth decompresses to approximately **23 MB** (RGB) or **30 MB** (RGBA) of raw pixel data. During a transformation the pipeline holds source, working, and output buffers concurrently, so peak allocation for a single image operation can reach **50–90 MB**.

The standard mitigation is to increase the G1 region size via `-XX:G1HeapRegionSize`. An object is only considered Humongous when it exceeds *half* the region size, so a larger region raises the threshold. Valid values are `1m`, `2m`, `4m`, `8m`, `16m`, and `32m` — **32m is the maximum the JVM permits**, which gives a Humongous threshold of only 16 MB. A single decompressed 4000×2000 RGB buffer already exceeds that limit, so **for deployments that regularly process high-resolution images, Humongous allocations cannot be avoided regardless of region size tuning**. In those cases the Parallel GC default is the more appropriate choice.

For moderate image resolutions (roughly $\leq 2000 \times 1000$ px), increasing the region size can still help. Use `-XX:G1HeapRegionSize=32m` as a starting point and monitor GC logs for Humongous allocation events.

The following mitigation is effective regardless of image resolution:

```
com.coremedia.trans  
form.throttle.per  
mits=<n>
```

Limits the number of image transformations that may run concurrently, capping the peak native memory demand (in megabytes). The default is 1/4 of the configured heap size. Reducing this value lowers peak memory pressure at the cost of transformation throughput.

Heap Sizing

Correct heap sizing is as important as the GC algorithm choice. Two common mistakes lead to OOM incidents in opposite directions:

- **Heap too small** — ParallelGC thrashes and triggers `GC overhead limit exceeded` under load.
- **Heap too large** — the JVM heap plus off-heap memory (native buffers, Metaspace, thread stacks) exceeds the container memory limit, causing the container to be killed by the OS or Kubernetes.

As a guideline:

- Allocate **70–80 % of the container memory limit** as Java heap via `JAVA_HEAP` or `-XX:MaxRAMPercentage`. The remaining 20–30 % covers off-heap memory.
- A heap of **4 GB is a tested baseline** for a delivery CAE or Headless Server without heavy customisation. Scale up for higher traffic or more complex business logic.

2.1.3.2 Building Distroless Images

For Kubernetes deployments, it might be interesting to build distroless images. These are images, that contain only the JVM and the application but no shell. These images are, by nature, more secure. To build a distroless image, a second Maven profile should be added:

```
<profile>  
  <id>distroless-image</id>  
  <properties>  
    <jib.skip>false</jib.skip>  
    <application.image-base>  
gcr.io/distroless/java@sha256:65aa73135827584754f1f1949c59c3e49f1fed6c35a918fadba8b4638ebc9c5d  
    </application.image-base> ☐  
  </properties>  
</build>
```

```
<plugins>
<plugin>
  <groupId>com.google.cloud.tools</groupId>
  <artifactId>jib-maven-plugin</artifactId>
  <configuration>
    <container>
      <jvmFlags>
        <jvmFlag>-XX:+UseParallelGC</jvmFlag>
        <jvmFlag>-XX:+ExitOnOutOfMemoryError</jvmFlag>
        <jvmFlag>-XX:InitialRAMPercentage=45.0</jvmFlag>
        <jvmFlag>-XX:MaxRAMPercentage=70.0</jvmFlag>
        <!--
        | <jvmFlag>-XX:+HeapDumpOnOutOfMemoryError</jvmFlag>
        | <jvmFlag>-XX:HeapDumpPath=/tmp/app.hprof</jvmFlag>
        -->
      </jvmFlags>
    </container>
  </configuration>
</plugin>
</plugins>
</build>
</profile>
```

- ❶ Go to <https://console.cloud.google.com/gcr/images/distroless/GLOBAL/java> and pick a hash
- ❷ When building a distroless based image, there is only the JVM included and only environment variables and jvm flags are supported. `-XX:+UseParallelGC` is set explicitly here because the `JAVA_PARALLEL_GC` environment variable is not available without the `java-application-base` entrypoint scripts. See [Section 2.1.3.1, "Garbage Collector Configuration"](#) [18] for background.
- ❸ In Kubernetes the JVM memory is controlled by Kubernetes. In order to allow that, we need to define the upper and lower bound of the allocated pod memory using `InitialRAMPercentage` and `MaxRAMPercentage`.

2.1.3.3 Building ARM Images

To build ARM images with maven, you need to set the maven property `application.image-arch` to `arm64`. If the referenced OCI manifest has a matching image in its list, the build process will be based on that architecture. Be advised, that jib is currently not able to create multi image manifests. For that reasons only one architecture can be build at a time.

2.1.3.4 Building Dockerless

Jib allows building images without Docker being installed. To do so, you need to replace the build goal from `dockerBuild` to `build`. This can be done by setting the property `jib.goal`.

Docker Setup | The java-application-base Image

When building dockerless directly into a registry, make sure your registry supports the OCI format. If not, you need to set the `<format>OCI</format>` settings in the build poms of the app modules to `<format>docker</format>`.

To start a registry locally you can start with:

```
docker run -d \  
-e REGISTRY_HTTP_ADDR=0.0.0.0:5001 \  
-p 5001:5001 \  
--name registry \  
registry:2.8
```

The port here is switched to `5001` to not collide with Apple AirPlay listener on MacOS, a common mistake you can run into.

To build the images into this registry, run:

```
mvn install -Pdefault-image \  
-Djib.goal=build \  
-Djib.allowInsecureRegistries=true \  
-Dapplication.image-prefix=localhost:5001
```

2.2 Docker Container Startup

2.2.1 Startup Entrypoint and Command chain

When the Docker container is started, the command of the process is created by inspecting the `Entrypoint` and the `Cmd` metadata fields of the image. If you build the images will all defaults and you run

```
docker image inspect --format='{{.Config.Entrypoint}} - {{.Config.Cmd}}'
coremedia/cae-live
```

You will get a response like this:

```
[./entrypoint.sh] - [application]
```

This is the entrypoint script chain that will be processed at startup.

2.2.2 Health Check

Docker has a built-in health check mechanism that, if defined for an image, probes a defined command regularly to determine the health status of the container. Because containerization is all about automation of processes, health checks are not only important for monitoring purposes but even more important for triggering events in an automation chain in your cluster. Many tools depend on the status information the container daemon provides. The Traefik reverse proxy uses in our Docker development setup for example automatically removes a routing to a backend, when the backends health status is not `healthy`.

If you for example run

```
docker ps --format '{{.Names}} - {{.Status}}'
```

you will get a status report for all running containers like this:

```
elastic-worker - Up 15 hours (healthy)
user-changes - Up 15 hours (healthy)
cae-preview - Up 15 hours (healthy)
cae-live - Up 15 hours (healthy)
studio-server - Up 15 hours (healthy)
overview - Up 15 hours
content-feeder - Up 15 hours (healthy)
cae-feeder-live - Up 15 hours (healthy)
commerce-adapter-wcs - Up 15 hours (healthy)
headless-server-preview - Up 15 hours (healthy)
master-live-server - Up 15 hours (healthy)
```

```
cae-feeder-preview - Up 15 hours (healthy)
headless-server-live - Up 15 hours (healthy)
studio-client - Up 15 hours (healthy)
workflow-server - Up 15 hours (healthy)
content-management-server - Up 15 hours (healthy)
```

The health check command is defined in the `Dockerfile` of the application, that is

```
HEALTHCHECK --start-period=30s --interval=30s --timeout=3s \
  CMD curl -Lf http://localhost:8081/actuator/health || exit 1
```

2.3 Container Configuration

There are three ways to configure an application:

- Build time configuration
- Runtime configuration
- Start time configuration

Please review the official [Spring-Boot](#) documentation for further details how to externalize configuration.

2.3.1 Build Time Configuration

Build time configuration can be achieved by simply predefining properties in one of the configuration files in the workspace. The location is `application.properties` in the `apps/<app name>/spring-boot/<app name>-app/src/main/resources` directory

2.3.2 Start Time Configuration

Start time configuration can be achieved using

- environment variables
- mounted configuration files

Configuring Spring-Boot applications using environment variables is probably the most common case and is the default choice in a container based environment. With Docker you simply add each environment variable on the command-line or pass in a path to a file with the environment variables. With Spring environment variables are defined in upper case and transformed by Spring to the java properties syntax.

- `REPOSITORY_URL` is transformed to `repository.url`
- `LINK_URLPREFIXTYPE` is transformed to `link.urlPrefixType`

To pass in environment variables use the `-e` or the `--env-file` options.

Passing environment variables on the command-line:

```
docker run \  
-e REPOSITORY_URL=http://my-host:40180/ior coremedia/cae-preview
```

Configuring the application by mounting a configuration file should only be used if passing variables in does not work or if replacing an existing configuration file is desired. To mount a file into the container, the path to the file or directory on the host and in the container must be defined. I.e.

Mounting configuration files:

```
docker run \  
-v $PWD/override.properties:/coremedia/config/application.properties \  
coremedia/cae-live
```

2.3.3 Runtime Configuration

To support runtime (re)configuration, the application's configuration had to be backed by a k/v store and the application would have to support it. This kind of configuration is not supported by CoreMedia applications except for some special properties that can be configured using struct content items in the content repository.

2.4 Containerized Tools

With the Docker setup all CoreMedia tools are packaged within the `core media/management-tools` image. You will find the Maven module to build the image at:

```
global/management-tools/management-tools-image
```

The image contains the following tools:

- content-management-server tools
- master-live-server tools
- replication-live-server tools
- workflow-server tools
- cae-feeder tools
- theme-importer

2.4.1 Running the Tools

To run the tools, you can either:

- Start the `management-tools` container in interactive mode and run the tools *from within*. This approach has the advantage, that it provides a familiar user experience as if the tool is installed locally.
- Start a tool from the `management-tools` container directly, using the tool as the command argument and remove the container when the tool finishes. This approach has the advantage, that it can be used in the context of a script as it requires no interaction.

Regardless of the approach, you need to keep in mind the following common issues with containers:

- In case you need to read from or write to files on your host, you need to mount volumes or create a bind mount.

You can mount a file system path to a docker container by using the `--volume <local-path>:<container-path>` command-line option.

- In case the containers are placed within a user-defined network, you need to add the container to that network to access the servers you want to connect. This network configuration depends on the network plugins being used. For the localhost `docker compose` development setup, this network is named

`backend` and prefixed with either the value of the `COMPOSE_PROJECT` environment variable or the name of the directory, where the compose files lies. In the development setup this will be `compose` but you can find out what networks are available, by running:

```
docker network list
NETWORK ID   NAME                DRIVER  SCOPE
181flaf50e6e bridge              bridge  local
2788cdf4bd7a compose_backend     bridge  local ❶
3227b513840e compose_web         bridge  local ❷
494f6bdaafa8 host                host    local
3d8b08335372 none               null    local
```

- ❶ The network, where all services are put in
- ❷ The network, where all services with Traefik ingress are put in

You can attach the docker container to a network using the `--network` `<network name>` command-line option.

- In case you need to adjust logging for debugging purposes, you can tune the tools' message log level and pattern with `stdout.log.level` and `stdout.log.pattern` as `docker run` options. Example:

```
-e stdout.log.pattern="%d{yyyy-MM-dd HH:mm:ss} %-7([%level]) %logger -
%message%n%xEx{5}"
-e stdout.log.level=DEBUG
```

- Per default, all log output is printed to the Docker console and root logger level is set to `ERROR`. Should you need to change that, for specific log levels or output to a file, you can overwrite the default logback configuration file. Make a copy of `global/management-tools/management-tools-image/src/main/image/coremedia/tools/properties/corem/tools-logback.xml` as a starting point and mount it with this `docker run` option (see also [Section 2.4.2.1, "Mount Configuration Files" \[30\]](#)):

```
--volume
$PWD/tools-logback.xml:/coremedia/tools/properties/corem/tools-logback.xml
```

Remember to use an external volume (see above) when writing log output to a file.

2.4.1.1 Start an Interactive Shell in the Container

To start the `management-tools` container with an interactive shell, you need to run:

```
docker run -it --rm coremedia/management-tools
```

2.4.1.2 Start a Tool Directly With a Custom Entrypoint

To start a tool, for example the `version` tool directly, you need to run:

```
docker run --rm \
  coremedia/management-tools \
  tools/bin/cm version
```

Here the tool and its arguments are passed to the entrypoint. Because the default working directory is set to `/coremedia` by the image, you only need to pass the relative path to the `cm` executable. All arguments that follow will be handled as arguments to the `cm` executable. The results of the tools are written to stdout and can easily be processed using standard shell means.

Parse an IOR URL:

```
docker run --rm \
  coremedia/management-tools \
  tools/bin/cm ior <ior string>
```

2.4.2 Configuring the Tools

Many tools have special configuration files and their values cannot be passed into the tool using environment variables. In order to configure the tools, you have two options:

- Create configuration files on your host and mount them into the container when executing the tool
- set environment variables if supported

2.4.2.1 Mount Configuration Files

1. Create the config file

```
cat << EOF > capclient.properties
cap.client.server.ior.url=http://<host>:<port>/ior
EOF
```

2. Execute the tool with mounted configuration file

```
docker run --rm \
--entrypoint /coremedia/tools/bin/cm \
--volume \
    $PWD/capclient.properties:\
    /coremedia/tools/properties/corem/capclient.properties
coremedia/management-tools \
dump -u <user> -p <password> <id>
```

With mounted configuration files, you can create a separate configuration folder for each content server and for each environment you want to connect to and mount the whole folder to switch a tool from one environment to another.

Examples with Mounted Configuration

Create config:

```
cat << EOF > capclient.properties
cap.client.server.ior.url=http://<host>:<port>/ior
EOF

cat <<EOF > resetcafeeder.properties
cafeeder.datasource.url=jdbc:mysql://<host>:3306/cm_mcafeeder
cafeeder.datasource.username=cm_mcafeeder
cafeeder.datasource.password=cm_mcafeeder
EOF

cat <<EOF > sql.properties
sql.store.driver=com.mysql.cj.jdbc.Driver
sql.store.url=jdbc:mysql://localhost:3306/cm_management
sql.store.user=cm_management
sql.store.password=cm_management
EOF
```

Export content to your host:

```
docker run --rm \
--entrypoint /coremedia/tools/bin/cm \
--volume \
    $PWD/capclient.properties:\
    /coremedia/tools/properties/corem/capclient.properties \
--volume \
    $PWD/export:/export \
coremedia/management-tools \
serverexport -u admin -p admin \
-r --basedir /export /
```

Import content to your host:

```
docker run --rm \
--entrypoint /coremedia/tools/bin/cm \
--volume \
    $PWD/capclient.properties:\
    /coremedia/tools/properties/corem/capclient.properties \
--volume \
    $PWD/import:/import \
coremedia/management-tools \
```

Docker Setup | Configuring the Tools

```
serverimport -u admin -p admin \  
-r /import
```

Reset the cae-feeder-preview:

```
docker run --rm \  
--entrypoint /coremedia/tools/bin/cm \  
--volume \  
  $PWD/resetcafeeder.properties:\   
  /coremedia/tools/properties/corem/resetcafeeder.properties \  
coremedia/management-tools \  
resetcafeeder reset
```

Unlock the content-management-server:

```
docker run --rm \  
--entrypoint /coremedia/tools/bin/cm \  
--volume \  
  $PWD/sql.properties:\   
  /coremedia/tools/properties/corem/sql.properties \  
coremedia/management-tools \  
unlockcontentserver
```

3. CoreMedia Properties Overview

Properties can be configured via Spring Boot. You can use `application.properties`, system properties, environment variables in uppercase and many more. See the [official Spring Boot documentation](#) for details.

Spring Boot configuration

For details on specific property data types like `Duration` or `DataSize`, see Spring Boot documentation, section [Properties Conversion](#)

Spring's relaxed binding also allows for different notations of property names like snake or camel case, but the default is 'kebab case' (separating words with dashes). Generally a dot in a property name reflects some kind of logical hierarchy. List-valued properties are zero-based and use bracket notation (`x.y.1.*` -> `x.y[0].*`)

Spring relaxed binding

NOTE

Configuration properties that are defined in the Blueprint may be missing here. Therefore, also check the Blueprint sources for classes annotated with `ConfigurationProperties` and use their JavaDoc.



3.1 Content Application Engine Properties

Find more related properties here:

- [Section 3.14, “Web Security Properties” \[242\]](#)

3.1.1 General CAE Properties

cae.cookie.force-http-only	
Type	Boolean
Default	true
Description	Whether or not to force the 'HttpOnly' attribute on all cookies.
cae.cookie.force-secure	
Type	Boolean
Default	true
Description	Whether or not to force the 'Secure' attribute on all cookies.
cae.cookie.same-site	
Type	String
Description	The value of the cookie's 'SameSite' attribute. Valid values are the ones as defined by the spec. In addition, the value 'Unset' can be used to indicate that the attribute should not be set.
cae.csrf.ignore-paths	
Type	List<String>
Description	Path Patterns to ignore for CSRF prevention.

These must be specified relative to any servlet path prefix (meaning you should exclude the context path and any servlet path prefix in stating your pattern).

```
cae.hashing.secret
```

Type String

Description A Secret which is used for url parameter hashing. Needs to be at least 32 characters long. If not configured a secret will be generated and exposed via warn log on application startup. If multiple CAEs are used, ensure to set the secret instead of trusting a generated one.

```
cae.http-firewall.allow-semicolon
```

Type Boolean

Default false

Description Determines if semicolon is allowed in the URL (i.e. matrix variables).

```
cae.http-firewall.allow-url-encoded-double-slash
```

Type Boolean

Default false

Description Determines if a double slash (//) that is URL encoded (%2F%2F) should be allowed in the path or not.

```
cae.http-firewall.allow-url-encoded-percent
```

Type Boolean

Default false

Description Determines if a percent (%) that is URL encoded (%25) should be allowed in the path or not.

```
cae.http-firewall.allow-url-encoded-period
```

Type Boolean

Default	false
Description	Determines if a period (.) that is URL encoded (%2E) should be allowed in the path or not.
<code>cae.http-firewall.allow-url-encoded-slash</code>	
Type	Boolean
Default	false
Description	Determines if a slash (/) that is URL encoded (%2F) should be allowed in the path or not.
<code>cae.http-headers.csp.directives</code>	
Type	String
Default	<code>default-src 'self'; connect-src 'self' *.engagement.coremedia.cloud wss://*.engagement.coremedia.cloud; script-src 'self' *.engagement.coremedia.cloud 'unsafe-inline' 'unsafe-eval'; style-src 'self' *.engagement.coremedia.cloud 'unsafe-inline'; img-src 'self' *.engagement.coremedia.cloud data;</code>
Description	The CSP directives to be set. Defaults to "self". Set to empty to omit the CSP header. For CAE-Preview use the following CSP directives and adapt the STUDIO-FQDN. <code>"default-src 'self'; script-src 'self' STUDIO-FQDN 'unsafe-inline'; style-src 'self' 'unsafe-</code>
<code>cae.http-headers.csp.report-only</code>	
Type	Boolean
Default	false
Description	Flag to control if the content security policy is to be reported only.
<code>cae.http-headers.frame-options</code>	
Type	<code>com.coremedia.cae.security.CaeHttpHeadersConfiguration</code> <code>Properties\$FrameOptions</code>

Default	disable
Description	Configure the X-Frame-Options header.
<code>cae.http-headers.xss</code>	
Type	<code>com.coremedia.cae.security.CaeHttpHeadersConfigurationProperties\$XSS</code>
Default	disabled
Description	Configure the X-XSS-Protection header.
<code>cae.link-transformer.include-params-appender.uri-paths</code>	
Type	<code>List<String></code>
Default	<code>/dynamic/</code>
Description	uriPaths the IncludeParamsAppendingLinkTransformer should be applied to.
<code>cae.link-transformer.params-appender.skip-for-type</code>	
Type	<code>Map<String,List<Class<?>>></code>
Description	List of bean classes which should be ignored by the ParamsAppendingLinkTransformer.
<code>cae.link-transformer.serializer-classes</code>	
Type	<code>List<Class<?>></code>
Description	A list of fully qualified class names for which a <code>tools.jackson.databind.ValueSerializer</code> should be registered for view parameter conversion. Every class which is configured here, should have a proper <code>com.coremedia.id.IdScheme</code> implementation being registered at the <code>com.coremedia.id.IdProvider</code> bean.
<code>cae.preview.metadata-enabled</code>	
Type	Boolean

Default	false
Description	Whether to disable metadata rendering. Disabled by default.
<code>cae.preview.pbe.include-jquery</code>	
Type	Boolean
Default	false
Description	Configures if jquery should be included when rendering the preview related scripts.
<code>cae.preview.pbe.studio-url-whitelist</code>	
Type	List<String>
Description	Configures a list of valid Studio URLs. The Studio Preview integration does only work for listed Studio instances. If left blank, any Studio instance is considered valid.
<code>cae.set-unknown-mime-type</code>	
Type	Boolean
Default	false
Description	This property controls if an instance of <code>com.coremedia.blueprint.cae.filter.UnknownMimeTypeCharacterEncodingFilter</code> is registered to fix unknown encoding errors in WebSphere versions up to and including 8.5.5010.20160721_0036. The <code>UnknownMimeTypeCharacterEncoding</code> filter will be used when <code>cae.set-unknown-mime-type</code> is set to true. The default is suitable when using Tomcat or WebSphere starting from 8.5.5011.20161206_1434.
Deprecation	This property has been deprecated and will be removed in a future version.
<code>cae.single-node</code>	
Type	Boolean
Default	false

Description	This property is used in com.coremedia.blueprint.cae.handlers.HandlerBase#doCreateModelWithView to control if a possibly outdated resource is served or if a redirect is sent. The redirect is only a valid response when cae.single-node is set to true.
cae.table-section-filter.footer-class	
Type	String
Default	tr--footer
Description	CSS class to identify footer rows. Defaults to "tr--footer".
cae.table-section-filter.header-class	
Type	String
Default	tr--header
Description	CSS class to identify header rows. Defaults to "tr--header".
cae.table-section-filter.max-nesting-depth	
Type	Integer
Default	0
Description	Configures the maximum allowed nesting depth of tables for the table section filter. If this limit is exceeded, additional nested tables will be ignored and passed through by the table section filter. If 0, the filter is effectively disabled. Defaults to 0.
cae.view.cycle-check	
Type	Boolean
Default	true
Description	Check for cyclic inclusions. You should not disable the check, unless for some good reason, e.g.:

- You use a custom ViewDispatcher, whose getView method is not determined only by its arguments.
- Your View#render implementation invokes ViewUtils#render with the same bean + viewName, but with an other (delegate) View.

```
cae.view.debug-enabled
```

Type	Boolean
Default	false
Description	If set to true, html comments will be written to the rendered pages around included fragments. This is a development feature. With these comments you can easily see which JSP, bean and view was used to render a fragment.

```
cae.view.errorhandler.enabled
```

Type	Boolean
Default	true
Description	Enables/disables the view exception handler.

```
cae.view.errorhandler.output
```

Type	Boolean
Default	false
Description	If handler is enabled and set to true, exceptions will be displayed in the current page.

```
cae.view.filter-lookup-by-predicate
```

Type	Boolean
Default	false
Description	By convention, templates are written for bean interfaces, but views may be named after any type. If set to true, viewlookup will only be done for views named

after interfaces, not classes, with configurable excludes and includes. This option is turned off by default.

```
cae.view.max-depth
```

Type	Integer
Default	200
Description	Maximum depth of inclusions.

```
cae.viewdispatcher.cache.enabled
```

Type	Boolean
Default	true
Description	Defines if the caching of view lookups is enabled. Disabling might be useful when developing templates. Shouldn't be disabled when in production mode!

```
cae.viewdispatcher.expose-spring-macro-helpers
```

Type	Boolean
Default	true
Description	Set whether to expose a RequestContext for use by Spring's macro library, under the name "springMacroRequestContext". Default is "true". Currently needed for Spring's Velocity and FreeMarker default macros. Note that this is <i>not</i> required for templates that use HTML forms <i>unless</i> you wish to take advantage of the Spring helper macros.

```
cae.viewdispatcher.fallback-to-default-view
```

Type	Boolean
Default	true

Description Fallback to default view if requested view name raises view exception.

Table 3.1. Configuration Properties with Prefix cae

3.1.2 Delivery CAE Properties

delivery.developer-mode	
Type	Boolean
Default	false
Description	Setting to enable the developer mode. If set to true, merging of resources such as JavaScript and CSS is disabled, so that each corresponding script and stylesheet will be rendered as a separate link. If set to false (the default), these resources will be merged during delivery and only a single link will be rendered to each type of resource. Additionally, when developer mode is enabled, a developer mode indicator is rendered on the page, the livereload script is included, and missing resource bundle keys are highlighted.
delivery.image-settings-aware-etags	
Type	Boolean
Default	false
Description	Links to transformed images contain eTags in order to be unique if the image changes. If this property is set to true, the eTags do not only reflect the actual image data, but also the transformation settings like jpegQuality. This is useful for development purposes, if you want to see changes in the image transformation configuration immediately. But it makes link building very expensive, because the image data must always be streamed to compute the eTag that reflects the current settings. Therefore this flag should usually be false, esp. for production systems, where image transformation settings are changed rarely. If it is false, a CAE restart is required after changing the image transformation settings.
delivery.local-resources	
Type	Boolean

Default	false
Description	true if links to resources such as JavaScript and CSS should be generated to point to files in the application container (e.g. in modules under META-INF/resources/themes/) instead of the content repository.
delivery.outdated-transformed-blob-cache-duration	
Type	Duration
Default	0s
Description	Duration after which outdated transformed blobs should be removed from the cache. This value is used to set the cache control headers for transformed blobs in case they are requested but either yet unknown to this CAE instance or outdated meanwhile. Only takes effect if cae.single-node=false, otherwise a redirect to the latest version of the blob is returned.
delivery.preview-mode	
Type	Boolean
Default	false
Description	Indicates whether this is a preview cae.
delivery.standalone	
Type	Boolean
Default	true
Description	True when this is a "standalone" CAE in a servlet container serving requests directly or false when behind some URL rewriting reverse proxy (e.g. Apache).

Table 3.2. Delivery Properties

3.1.3 Http Cache Control Properties

cae.cache-control.cache-size	
------------------------------	--

Type	Integer
Default	10000
Description	Maximum count of cache configuration entries. Cache cleans up automatically considering LRU strategy. Default is 10000 entries.

```
cae.cache-control.for-type
```

Type	Map<String, com.coremedia.cms.delivery.configuration.Cachecontrol>
Description	Map of initial cache control configuration objects for Http cache control Header. See WebProperties.Resources.Cache.Cachecontrol for further details. The configuration is type specific, but also can benefit from abstraction along the class hierarchy. That means that you may simply configure a common value for a super class instead of taking care about all the child classes. The type information part of the property name should be the simple name of the class (no package information) and in lowercase. Configuration of classes with the same name in different packages is not supported.

Example Configuration:

```
cae.cache-control.for-type.cmlinkable.max-age=1m
cae.cache-control.for-type.blob.max-age=365d
```

Default/Fallback:

```
cae.cache-control.for-type.object.max-age=1m
```

```
cae.cache-control.for-url-pattern
```

Type	Map<String, com.coremedia.cms.delivery.configuration.Cachecontrol>
Description	Map of initial cache control configuration objects for Http cache control Header. See Resources.Cache.Cachecontrol for further details. The configuration is URL pattern specific.

```
cae.cache-control.initial-max-age-for-type
```

Type	Map<String, Duration>
-------------	-----------------------

Description Initial max-age value for Cache Control Header. Does not have to be adjusted necessarily. The initial max-age value should be small, since the effective cache duration is provided with the subsequent request. The configuration is type specific, but also can benefit from abstraction along the class hierarchy. That means that you may simply configure a common value for a super class instead of taking care about all the child classes. The type information part of the property name should be the simple name of the class (no package information) and in lowercase. Configuration of classes with the same name in different packages is not supported.

Example Configuration:

```
cae.cache-control.initial-max-age-for-type.blob=365d
```

Default/Fallback:

```
cae.cache-control.initial-max-age-for-type.object=1m
```

```
cae.cache-control.s-max-age-factor
```

Type	Double
Default	0

Description This value is used to compute a s-maxage header in relation to a recognized max-age value. If the value is 0.0 then the s-maxage value is not automatically adjusted. In particular this means the s-maxage value isn't written in a cache-control header unless it was otherwise configured, for example with:

```
cae.cache-control.for-type.cmlinkable.s-max-age=1m
```

But please note, a value other than 0.0 takes precedence over a static configured value. Then s-maxage takes the value of max-age * sMaxAgeFactor.

```
cae.cache-control.url-path.remove-semicolon-content
```

Type	Boolean
Default	false

Description Set if everything after a ";" (semicolon) should be cut from the request URI for path matching. E.g., to see all matrix parameters of fragment urls, this should be set to "false".

cae.cache-control.url-path.url-decode	
Type	Boolean
Default	true
Description	Set if the request URI should be decoded for path matching.
cae.cache-control.url-path.use-always-full-path	
Type	Boolean
Default	false
Description	Set if the full path within the current web application context is used for path matching. By default, this is set to "false" which means "/servlet" is cut off.

Table 3.3. Configuration Properties with Prefix cache.control

3.2 Content Server Properties

3.2.1 General Content Server Properties

Protocol for the connection of Publisher and Replicator to the Master Live Server can be either CORBA (default) or gRPC. Property `repository.use-grpc` (see [Section 3.12.1, “Unified API Spring Boot Client Properties” \[233\]](#)) activates gRPC.

```
cap.server.allow-synthetic-replay
```

Type	Boolean
Default	true
Description	Whether it is allowed for clients to request a synthetic replay of the content repository, for example using the constant <code>Timestamp.SYNTHETIC_REPLAY</code> Unified API. This is a very expensive operation that is rarely used except when setting up a Replication Live Server from scratch.

```
cap.server.base-dir
```

Type	String
Default	user.dir
Description	Used as base directory for handling relative file-paths in the content-server configuration. Defaults to the value of system-property <code>user.dir</code> .

```
cap.server.blob-channel-timeout
```

Type	Duration
Default	1m
Description	This property sets the timeout duration for streaming blobs to and from the database. (Default unit is seconds) In general, you don't have to change the default value. It is provided for exceptional cases, when the connection to the database is unreliable.

```
cap.server.blob-download-session-validation
```

Type	Boolean
Default	true
Description	Blob downloads from content servers are secured by verification of a valid login session on each call, if true. If set to false, the session validation for downloads is disabled. This is useful on Replication Live Servers to allow for load-balanced blob download. It must only be set to false if corresponding servers are appropriately secured by firewalls, private networks, or the like. Otherwise, public unsecured read access to blobs could be possible. Blob uploads are always secured by session validation, regardless of the option's setting. @since 2512.0.0

```
cap.server.blob-md5-permission-check
```

Type	Boolean
Default	true
Description	Blob access by clients is protected against URL guessing. If you want to use clients from releases < CMCC 2110 with content servers of CMCC 2110 or newer, you must set this flag to false for the new content servers.

```
cap.server.blob-url-pattern
```

Type	String
Default	https?:.*
Description	A regular expression that must match the entire URL string for URL blobs, that is, blobs that are only referenced by their URL in the content repository and are resolved at the client side on access. Using URL blobs can significantly reduce the storage requirements of the Content Server. The pattern must match the entire URL string, starting with the protocol. Before matching, the path component of the URL is normalized according to <code>java.net.URI#normalize()</code> where applicable. URLs with the schemes <code>s3:</code> and <code>classpath:</code> are not normalized. By default, only <code>http:</code> and <code>https:</code> URLs are allowed. Allowing too many URLs may cause security problems. For example, allowing <code>file:.*</code> would also grant access to all configuration files. Instead, a single path like <code>file:///share/blobData/.*</code> should be sufficient in most cases. Alternative patterns can be given according to the

Java regular expression syntax as implemented by `java.util.regex.Pattern`:
`file:///share/blobData/.*(http://blobstorage.internal/.*`

```
cap.server.blobstore.s3.bucketname
```

Type	String
-------------	--------

Description	The name of the S3 bucket that is used to store blobs. If this property is not specified, the S3 media store is not enabled.
--------------------	--

```
cap.server.blobstore.s3.encryption-mode
```

Type	String
-------------	--------

Description	The server side encryption mode to use for storing blobs in S3. Only valid for the enhanced S3 store implementation. If not set, the default encryption mode of the bucket is used. Valid values are AES256, aws_kms and aws_kms_dsse.
--------------------	--

```
cap.server.blobstore.s3.force-path-style
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description	Flag indicating whether to use path style URLs for S3 requests. Only valid for the enhanced S3 store implementation. If set to true, legacy path style URLs are used for all requests. If not set or set to false, virtual host style URLs are used, which is the AWS recommended way to access S3 buckets.
--------------------	---

```
cap.server.blobstore.s3.rootdir
```

Type	String
-------------	--------

Description	The name of an S3 path prefix, used to disambiguate multiple media stores in one S3 bucket. If possible, it is recommended to use separate buckets instead, improving performance by avoiding constant path prefixes.
--------------------	---

```
cap.server.blobstore.s3.tags
```

Type	Map<String,String>
-------------	--------------------

Description A map of additional tags to apply to all objects stored in the S3 bucket. Only valid for the enhanced S3 store implementation.

```
cap.server.cache.group-cache-size
```

Type Integer

Default 500

Description This property defines the size of the group cache. It limits the maximum number of groups which can be found in one search for groups by CMS clients. Set the property so that all groups connected to rights can be cached in memory.

```
cap.server.cache.group-cache-status-interval
```

Type Duration

Default 0

Description The duration between two log messages reporting the current state of the group cache. (Default unit is seconds)

The minimum value is 10 seconds, the maximum value is 1 hour.

```
cap.server.cache.member-folder-rights-cache-size
```

Type Integer

Default 1000

Description This property defines the size of the folder-specific rights cache. This cache stores the results of right calculations per folder and member, aggregating the results for all content types. This cache might help custom code using APIs other than the Unified API. Change this setting only if you observe the method `getRights(MemberKey)` in thread dumps of a slow Content Server.

```
cap.server.cache.resource-cache-size
```

Type Integer

Default 60000

Description	The capacity of the resource cache of the Content Server; the maximum value is 1000000, the minimum value is 100. This property defines the resource cache size, that is, the number of resources the server holds in memory. This value should sometimes be adapted to the increasing number of resources in the actual working set. If the value is too small, the server does not perform well. One resource needs about 2kB of heap space.
--------------------	--

cap.server.cache.resource-cache-status-interval

Type	Duration
Default	5m
Description	The duration between two log messages reporting the current state of the resource cache. (Default unit is seconds) The minimum value is 10 seconds, the maximum value is 1 hour.

cap.server.cache.rights-cache-size

Type	Integer
Default	3000
Description	This property defines the size of the rights cache. This cache stores the results of right calculations per resource, content type and member. If you have lots of different resources, content types and users you might need to adapt the value of the property. Check the proper size of the cache by examining the cache misses and faults in the log. To activate the log output of the rights cache set the cap.server.cache.rights-cache-status-interval property to a value larger than zero.

cap.server.cache.rights-cache-status-interval

Type	Duration
Default	0
Description	This property defines the interval at which log output of the rights cache is written. (Default unit is seconds) "0" means, that no log output is written.

```
cap.server.cache.user-cache-size
```

Type	Integer
Default	500
Description	This property defines the size of the user cache. It limits the maximum number of users which can be found in one search for users by CMS clients. Set the property to the size of the largest user search you want to perform, or the number of concurrently working users, whichever is greater.

```
cap.server.cache.user-cache-status-interval
```

Type	Duration
Default	0
Description	The duration between two log messages reporting the current state of the user cache. (Default unit is seconds) The minimum value is 10 seconds, the maximum value is 1 hour.

```
cap.server.case-sensitive-member-search
```

Type	Boolean
Default	false
Description	Determines whether users and groups should be searched case-sensitively. Affects only non-exact searches like "begins with" and "contains". Accessing a particular member by its exact name is always case-sensitive. This information is passed to all UserProviders. A UserProvider may not be able to fulfil this policy, e.g. if an underlying LDAP server supports only case-sensitive (or case-insensitive) queries, or the distinction is just not implemented. Default is false (i.e. case-insensitive search).

```
cap.server.check-unique-db-access
```

Type	Boolean
Default	true

Description	This property determines whether to check for another server that is running concurrently on the same database on server startup. By default, this property is set to false in the Blueprint to prevent issues with the IP-based lock value in containerized environments and virtual networks.
--------------------	---

```
cap.server.document-types
```

Type	List<String>
Default	classpath*:framework/doctypes/**/*.xml
Description	This property defines where the server finds the XML file(s) containing the content type definitions. The pattern supports resources from classpath or the filesystem. Relative file-paths are resolved against cap.server.base-dir). Examples:

```
classpath*/framework/doctypes/**/*.xml (Default),
file:///coremedia/doctypes/**/*.xml, config/contentserver/doctypes/**/*.xml
```

```
cap.server.encrypt-passwords-key-file
```

Type	String
Description	The location of the key generated by cm encryptpasswords (absolute or relative to cap.server.base-dir). If empty, defaults to etc/keys/DATABASE_NAME.DATABASE_USER.rijndael

```
cap.server.http-port
```

Type	Integer
Default	0
Description	Defines the HTTP(S) port of the application container containing the Content Server. The entry /Server/Service/Connector@port in server.xml has to have the same value as this property.

```
cap.server.init-runlevel
```

Type	String
Default	online
Description	The initial runlevel that the server will try to reach on startup. Possible runlevels are: online, administration, maintenance. This property does not override the default behavior of the Replication Live Server for the initial replication.
<code>cap.server.initial-password</code>	
Type	Map<String,String>
Description	The initial password to set for the default user with the indicated name. This password is set when the server is started for the first time. You can change the passwords later on at any time.
<code>cap.server.license</code>	
Type	String
Default	license.zip
Description	Defines where the server finds the license file (url or file-path absolute or relative to <code>cap.server.base-dir</code>).
<code>cap.server.login-service-webserver-privileged</code>	
Type	Boolean
Default	false
Description	This property specifies whether client connections of the login service 'webserver' are privileged and may log in as different users without further authentication. The default is false. It may be set to true to make the 'webserver' login service a privileged service as it was the case in releases before 1907.
<code>cap.server.login.authentication</code>	
Type	String
Default	classpath:coremedia-jaas.conf

Description	Defines where the server finds the JAAS login configuration file (url or file-path absolute or relative to cap.server.base-dir). Will be set on the java.security.auth.login.config System property (in case of a URL, a temporary file is created).
--------------------	--

```
cap.server.login.bouncers
```

Type	String
-------------	--------

Description	This property points to the optional login bouncer configuration file (absolute or relative to cap.server.base-dir). A login bouncer can grant or deny access to the Content Server based on the characteristics of the user and the set of currently logged in users.
--------------------	--

```
cap.server.login.password-hash-algorithm
```

Type	String
-------------	--------

Default	bcrypt:10
----------------	-----------

Description	A specification of the hash algorithm used for storing passwords. Allowed values are md5 for MD5-based password hashing and bcrypt:N with N being an integer between 4 and 30 (inclusive) for bcrypt-based password hashing. In the latter case N denotes the work factor which should adapted to the available CPU resources. This parameter applies to passwords of users defined in the built-in user repository of the Content Server, only. The value md5 is discouraged, because it makes brute-force attacks on passwords of low and medium strength possible. It should only be used if passwords need to be changed by clients (Studio, cm changepassword) that have not been updated to a CMS release that supports configurable password hashing. Old clients can login even after a password change without any restrictions. After changing this property, it is recommended to update the passwords of all users to ensure that all hashes have been computed according to the desired algorithm.
--------------------	---

```
cap.server.maximum-startup-delay
```

Type	Integer
-------------	---------

Default	120
----------------	-----

Description	Configured maximum time for Content Server startup in seconds
--------------------	---

This is the maximum time the startup of the server's Spring-Boot component is delayed if the configured initial runlevel has not been reached. After this timeout expires or the initial runlevel is reached the following things will happen:

- The Spring-Boot component will resume its startup and open the web application and actuator ports.
- The server will listen to 'health', 'liveness' and 'readiness' actuator requests. However, the 'health' and 'readiness' probes will block until the initial runlevel is reached.
- The server will serve the CORBA IOR via the IOR-Servlet. As the server may not be fully initialized yet, clients should not connect to the server before the 'readiness' probe returns a positive result.

As it is possible that the server is not fully initialized after this timeout expires, e.g. when running a database schema upgrade, it is recommended to delay the startup of dependent services until the 'health' or 'readiness' probes return a positive result. Additionally, health checks for the server are best based on the 'liveness' probe but must take the configured startup delay into account as the server may be unresponsive for this time.

cap.server.multiple-live-servers	
Type	Boolean
Default	false
Description	This property defines whether the server publishes to multiple live servers. Note that this flag cannot be easily changed after the first start of the Content Management Server.
cap.server.naming-policy-allow-at	
Type	Boolean
Default	false
Description	Allow '@' in member names (may cause confusion with user domains)
cap.server.persistent-property-writers	
Type	List<String>
Default	*

Description	A list of names of groups that may write or delete persistent properties. The magic group "*" grants rights to all users. This is the default for compatibility reasons. Connections using the publisher and replicator login service are always allowed to write persistent properties.
--------------------	--

```
cap.server.repository-home
```

Type	String
-------------	--------

Default	/Home
----------------	-------

Description	Defines the folder which will be used to store the home folders of the users. The whole folder hierarchy of the home folders is only visible to the administrator. Other user will only see one home folder with the path defined in cap.server.repository-home containing his personal files, such as the preferences. The default folder is /Home which will be automatically created by the system. If you define another folder, you need to create this folder by your own.
--------------------	--

```
cap.server.repository-system
```

Type	String
-------------	--------

Default	/System
----------------	---------

Description	Defines the system folder. It contains for example the public dictionary of the spell checker. The default folder is /System which will be automatically created by the system. If you define another folder, you need to create this folder by your own.
--------------------	---

```
cap.server.restore-inactive-users
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description	Behaviour for re-appearing users When a known user disappears from the underlying user provider, the user is marked as inactive in the CMS. This flag controls what happens if the user re-appears later on (recognized by the UUID). If true, the old CMS-user is reactivated, incl. his home folder. If false, a new CMS-user with a new empty home folder is created.
--------------------	---

```
cap.server.search.enable
```

Type	Boolean
Default	false
Description	If true full text search is enabled.

```
cap.server.session-ping-interval
```

Type	Duration
Default	1m
Description	The maximum duration that a ping is delayed when there are no available events. (Default unit is seconds) The session ping interval must not be less than 10 seconds to avoid possible session loss and ping flooding and not more than half an hour so that dead sessions are detected after an hour.

```
cap.server.termination-timeout-seconds
```

Type	Duration
Default	30s
Description	Timeout for waiting on running threads on shutdown

```
cap.server.unique-db-access-write-interval
```

Type	Duration
Default	0
Description	If unique DB access is checked and if this property is positive, this property determines the duration between two writes of the current timestamp to the database to indicate the liveliness of the server. (Default unit is seconds) If 0 or negative the server neither writes a timestamp regularly nor expects a timestamp to be written.

```
cap.server.use-strict-workflow
```

Type	Boolean
Default	false
Description	This property enforces the strict workflow mode. That is, the approver of a resource must be different from the editor. This is checked independently of the workflow engine, and should only be used in cases where a custom workflow definition is not an option.
<code>cap.server.userproviders</code>	
Type	List<hox.corem.server.ServerConfigurationProperties\$User providers>
Description	Configurations for UserProviders
<code>spring.grpc.server.port</code>	
Description	gRPC server port.
<code>spring.grpc.server.port</code>	
Description	gRPC server port.

Table 3.4. Content Server Properties

3.2.2 CORBA Properties

<code>com.coremedia.corba.client.local-socket</code>	
Type	Boolean
Default	false
Description	Redirect all outgoing connections to localhost. This is useful for running UAPI clients locally, while forwarding HTTP and CORBA connections to the servers through an SSH tunnel. A typical SSH invocation that allows the use of this socket factory would look like this, forwarding all ports of the <i>Content Management Server</i> and the <i>Workflow Server</i>:

```
ssh -L 40180:SERVER:40180 -L 40183:SERVER:40183 -L 40380:SERVER:40380
-L 40383:SERVER:40383 HOST
```

When using this option, you should also activate `com.coremedia.corba.server.set-NoSocket(boolean)` in order to reduce the attack surface of the JVM.

```
com.coremedia.corba.client.redirect.original-host
```

Type	String
Description	Redirect IIOp invocations to other addresses. Useful in tunnelling scenarios: imagine a server X on port P running in a network you cannot reach; but you are able to create a tunnel that forwards your address Y:Q to X:P, e.g. you do <code>user@Y % ssh -g -LQ:X:P</code> Then you can configure your clients by the <code>com.coremedia.corba.client.redirect</code> property group: • <code>originalHost</code>: The original server's name (X). If set, only IORs pointing to that server will be redirected. • <code>redirectHost</code>: The new server name (Y). If set, this name will be placed in the new target address. • <code>originalPort</code>: The original server's port (P). If set, only IORs pointing to that port will be redirected. • <code>redirectPort</code>: The new server port (Q). If set, this port will be placed in the new target address.

```
com.coremedia.corba.client.redirect.original-port
```

Type	Integer
Default	-1
Description	Redirect IIOp invocations to other addresses. For details see <code>com.coremedia.corba.client.redirect.originalHost</code>.

```
com.coremedia.corba.client.redirect.redirect-host
```

Type	String
Description	Redirect IIOp invocations to other addresses. For details see <code>com.coremedia.corba.client.redirect.originalHost</code>.

```
com.coremedia.corba.client.redirect.redirect-port
```

Type Integer

Default -1

Description Redirect IIOP invocations to other addresses.
For details see `com.coremedia.corba.client.redirect.originalHost`.

```
com.coremedia.corba.client.ssl.clear-text-ports
```

Type String

Description The Corba clear text ports
Usually one, maybe two (comma separated) values for *Content Server* and *Workflow Server*. In case of multiple values, the order must correspond to the `com.coremedia.corba.client.ssl.sslPorts`.

```
com.coremedia.corba.client.ssl.keystore
```

Type String

Description The keystore for SSL encrypted communication

```
com.coremedia.corba.client.ssl.passphrase
```

Type String

Description The passphrase for SSL encrypted communication

```
com.coremedia.corba.client.ssl.ssl-ports
```

Type String

Description The Corba SSL text ports
Usually one, maybe two (comma separated) values for *Content Server* and *Workflow Server*. In case of multiple values, the order must correspond to the `com.coremedia.corba.client.ssl.clearTextPorts`.

```
com.coremedia.corba.server.host
```

Type	String
Description	The host of the ORB. This hostname is exposed to the client via the IOR. Normally, you do not need to care about this, but in runtime environments with special DNS configurations it may be necessary to set a certain name that the client is able to resolve.
<code>com.coremedia.corba.server.no-socket</code>	
Type	Boolean
Default	true
Description	Disable ORB server sockets. This is useful when no incoming requests are expected, so that the server socket would only pose a security risk. The default is true, i.e. server sockets are disabled. This is appropriate for all client applications and must be set to false only for <i>Content Server</i> (CMS, MLS and RLS) and <i>Workflow Server</i> applications.
<code>com.coremedia.corba.server.port</code>	
Type	Integer
Default	-1
Description	The clear text port of the ORB.
<code>com.coremedia.corba.server.single-ip</code>	
Type	String
Description	Bind ORB sockets to this IP address. Useful to keep control on multi-homed hosts.
<code>com.coremedia.corba.server.ssl.enforce</code>	
Type	Boolean
Default	false

Description	Enforce SSL communication
	By default, the ORB opens an SSL port additionally to the clear text port. With this flag you can suppress the clear text port and thus ensure that clients use SSL connections.
<code>com.coremedia.corba.server.ssl.keystore</code>	
Type	String
Description	The keystore for SSL encrypted communication
<code>com.coremedia.corba.server.ssl.passphrase</code>	
Type	String
Description	The passphrase for SSL encrypted communication
<code>com.coremedia.corba.server.ssl.ssl-port</code>	
Type	Integer
Default	-1
Description	The port of the ORB for SSL encrypted communication.
	The orb will open the SSL socket <i>in addition</i> to the clear text port. There is no way to disable this.

Table 3.5. CORBA Properties

3.2.3 Properties for the Publisher

In the properties given below, some properties take a different form when using *CoreMedia Multi-Master Management*. In that case, an index number is used to indicate the publication target that is configured using this property value. The index numbers are consecutive integers starting with 1 and running up to the number of publication targets.

The protocol used for the connection to the *Master Live Server* can be either CORBA (default) or gRPC. Set property `repository.use-grpc` on the

Content Management Server for gRPC connection to the *Master Live Server* (see [Section 3.12.1, "Unified API Spring Boot Client Properties" \[233\]](#)).

```
publisher.autoextend.latest-approved-version
```

Type Boolean

Default true

Description This property affects the automatic extension of publication sets. The default behavior ("true") is as follows: If the referenced content item is not published yet, publish its latest (= newest) approved version. If the referenced content item is already published, do nothing. If set to "false", not the latest but the earliest approved version is published.

```
publisher.autoextend.max-failures
```

Type Integer

Default 0

Description The maximum number of implicitly added content items and folders whose publication may fail before no further content items and folders may be added implicitly. If set to 0 (the default), the publisher never stops adding content items and folders.

```
publisher.autoextend.update-linked-documents
```

Type Boolean

Default false

Description This property affects the automatic extension of publication sets. By default, ("false") only direct linked and approved content items are added to the publication set. If set to "true", all linked content items are recursively added to the publication set. The recursion stops when a version is already published and may lead to surprisingly large publication sets. Setting `updateLinkedDocuments` implicitly also sets `latestApprovedVersion` to "true".

```
publisher.connection-timeout-seconds
```

Type Duration

Default	5m
Description	The timeout duration for the connection to the Master Live Server ior.
<code>publisher.destroy-intermediate-versions</code>	
Type	String
Default	strict
Description	Whether intermediate content item versions between two publications will be destroyed or not on the Content Management Server. Example for dumb mode: Version 1 of content item A was published. In the meantime, the versions 2, 3, 4 and 5 have been created. When you now publish version 5, the versions 2, 3, 4 are destroyed and only version 1 and 5 remain on the Content Server. Old flags true and false are supported, too. false maps to off while true maps to strict. Available Modes • off Turns destruction of intermediate versions off. • dumb Simply destroys all versions between the currently published one and the previously published one; if you run into performance issues on publication or you do not use a multi-site set up this should be chosen. Mind that for multi-site set up this setting might break translation states. • strict Destroys all versions between the currently published one and the previously published one, but only if the versions are not referenced by master/masterVersion properties. This ensures that for a multi-site set up your translation state will be valid. This setting is recommended for multi-site set ups.
<code>publisher.destroy-older-versions-on-live-servers</code>	
Type	Boolean
Default	true

Description	Whether older published versions should be destroyed ("true") on the Master Live Server. That is, only two published versions (the current and the one before) of each content item remain on the Master Live Server. Only change if you have a valid reason.
--------------------	---

```
publisher.enable-bypass-previews
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	Whether publication previews bypass ("true") or not ("false") the publication queue for faster response times.
--------------------	--

```
publisher.local.domain
```

Type	String
-------------	--------

Description	The domain for logging in locally to the Content Management Server.
--------------------	---

```
publisher.local.password
```

Type	String
-------------	--------

Default	publisher
----------------	-----------

Description	The password for logging in locally to the Content Management Server.
--------------------	---

```
publisher.local.user
```

Type	String
-------------	--------

Default	publisher
----------------	-----------

Description	The user name for logging in locally to the Content Management Server.
--------------------	--

```
publisher.priority
```

Type	Map<String, Integer>
-------------	----------------------

Description	Priorities of publications requested through:
--------------------	---

- guiEditor: 60
- uapi: 40
- generator: 40
- jpython: 20
- utility: 20
- unknown: 20

publisher.send-session-token-in-url	
Type	Boolean
Default	true
Description	If true, the session token for connection to Master Live Server servlets (for blob up/download and CPU usage) will be sent as a URL query parameter. Otherwise, it will be sent as a request field. While sending the token as a query parameter is inherently insecure, it is the default to keep backward compatibility. Unless connection to an older server is required, it is recommended to set this property to false. The property is active for <i>all</i> publication targets.
publisher.status-interval	
Type	Duration
Default	5m
Description	The duration in between statistics log entries. (Default unit is seconds)
publisher.target	
Type	List<hox.corem.server.publish.PublisherConfiguration Properties\$Target>
Description	target is a list-valued property. You configure it with publisher.target[0].name, for example. Values other than "O" are only required for the multi-master feature, that is publication to different Master Live Servers.

Table 3.6. Publisher Properties

3.2.4 Properties for the Connection to the Database

NOTE

sql.properties for Command Line Tools

Be aware that the command line tools that access the database directly (like `cm unlockcontentserver` or `cm schemaaccess`) still need the `sql.properties` file in their `properties/corem` directory, as the tools are not (yet) configurable via Spring.



sql.pool.check-connection-max-threads

Type	Integer
Default	-1
Description	The maximum number of threads for checking database connections, or -1 for twice the number of sql.pool.max-connections.

sql.pool.check-timeout

Type	Duration
Default	5s
Description	Maximum duration the check is allowed to take. (Default unit is seconds)

sql.pool.log-queries

Type	Boolean
Default	false
Description	If the property is "true", messages concerning queries (search of content item versions) will be generated.

sql.pool.log-query-statements

Type	Boolean
Default	false
Description	If the property is "true", SQL statements concerning queries will be written to the log.
<code>sql.pool.log-schedule-messages</code>	
Type	Boolean
Default	false
Description	If the property is "true", write messages of the connection pool to the log file.
<code>sql.pool.log-verbose</code>	
Type	Boolean
Default	false
Description	If the property is "true", more debug messages will be generated.
<code>sql.pool.max-connections</code>	
Type	Integer
Default	4
Description	Maximum number of general connections to the database. Queries use their own connection pool, so that the max total number of database connections is the sum of max-connections and max-queries.
<code>sql.pool.max-queries</code>	
Type	Integer
Default	4
Description	Maximum number of connections used for queries, that is, the maximum number of parallel queries. Queries use their own connection pool, so that the max total

number of database connections is the sum of max-connections and max-queries.

sql.pool.max-retries	
Type	Integer
Default	15
Description	Maximum number of retries on job that failed due to read-only DB or connection loss.
sql.pool.min-connections	
Type	Integer
Default	2
Description	Minimum number of general connections to the database. Cannot be set to values below 1 to ensure timely detection of database connection issues.
sql.pool.open-timeout	
Type	Duration
Default	30s
Description	Maximum duration the opening of connections is allowed to take. (Default unit is seconds)
sql.pool.reaper-interval	
Type	Duration
Default	2m
Description	The interval in which it is checked if connections can be closed. (Default unit is seconds)
sql.pool.reaper-timeout	

Type	Duration
Default	3m
Description	The duration a connection must be idle before it will be closed. (Default unit is seconds)
<code>sql.pool.retry-delay</code>	
Type	Duration
Default	15s
Description	Delay interval before a retry is attempted due to a read-only DB or connection loss.
<code>sql.pool.retry-on-connection-loss</code>	
Type	Boolean
Default	false
Description	If "true", will retry failed transactions due to db connection failure.
<code>sql.pool.retry-on-read-only-db</code>	
Type	Boolean
Default	false
Description	If "true", will retry failed transactions due to a read-only database.
<code>sql.pool.scheduler-pull-up-ratio</code>	
Type	Double
Default	0.1
Description	Ratio of non-exclusive elements to pull up in scheduler queue in case of a potential deadlock (in the range]0.0,1.0]).

sql.pool.scheduler-watch-interval	
Type	Duration
Default	10s
Description	The interval between two checks for potential scheduler deadlocks.
sql.pool.scheduler-watchdog-enabled	
Type	Boolean
Default	true
Description	Flag to enable scheduler watchdog.
sql.pool.validator-interval	
Type	Duration
Default	5m
Description	The interval in which existing connections will be checked for function. (Default unit is seconds)
sql.pool.validator-timeout	
Type	Duration
Default	2m
Description	The duration a connection must be idle before it will be checked for function. (Default unit is seconds)
sql.schema.alter-table	
Type	Boolean
Default	false

Description	Setting the property to "true", causes the Content Server to automatically add database columns for newly defined content type properties and to fix wrong widths of String properties. This will only work if you set <code>sql.schema.check-Columns</code> to "true".
--------------------	---

```
sql.schema.check-columns
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	Setting the property <code>sql.schema.checkColumns</code> to "true", causes the Content Server to compare the content type definition with the existing database schema for missing columns and matching widths of String properties. If there are differences, it depends on the setting of <code>sql.schema.alterTables</code> if the Content Server refuses to start ("false") or if it adds and changes the columns automatically ("true"). Checking the tables consumes a considerable amount of time, so that the server starts up more slowly. If <code>sql.schema.checkColumns</code> is set to "false", the Content Server will not check the columns. If there are differences, you will run into Content Server exceptions later on.
--------------------	---

```
sql.schema.create-drop-indexes
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description	Setting the property to "true", causes the Content Server to automatically create and drop indexes on content type properties, according to the Index attribute in the <code>doctypes.xml</code> . This flag only affects existing columns. For newly added columns and tables, an index is always created if the Index attribute is set.
--------------------	---

```
sql.schema.create-table
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	The Content Server always checks if tables for the content types are missing. Setting the property to "true", causes the Content Server to create missing tables for new content types. If the setting is "false" and there are missing tables the Content Server refuses to start
--------------------	--

```
sql.store.collector.blob-preservation-period
```

Type	Duration
Default	1d
Description	The duration, a blob, that has no reference to a resource, will be preserved. (Default unit is seconds)

```
sql.store.collector.delay
```

Type	Duration
Default	1d
Description	The delay between two collections. (Default unit is seconds)

```
sql.store.collector.initial-delay
```

Type	Duration
Default	10m
Description	The delay before starting the first collection. (Default unit is seconds)

```
sql.store.collector.start-time
```

Type	Duration
Default	-1s
Description	The duration after 0:00h in the default time zone, when the collector should run every day. If set to -1, it does not run at a fixed time. The initial delay is always enforced as a minimal delay after server startup. (Default unit is seconds)

```
sql.store.collector.suspend
```

Type	Boolean
Default	false

Description	If set to true, unused blobs will not be deleted in the blob store. This can be used during backup if you have a non-transactional blob store. Alternatively, the blobcollector actuator endpoint must be used to suspend deletion of unused blobs at runtime before starting a backup of a non-transactional blob store.
--------------------	---

```
sql.store.convert
```

Type	Map<String,List<Class<? extends <code>hox.text.sax.SaxFilter</code> >>>
-------------	---

Description	Converters which are used to convert custom XML formats which have been defined for the obsolete <code>coremedia-sgmltext.dtd</code> . Example: <code>sql.store.convert.DocumentType.PropertyType=com.customer.XML-Converter</code>
--------------------	---

```
sql.store.convert-correct-rich-text
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	The editor in versions before SCI 4.1.38 created XML text, which was not valid according to the <code>coremedia-richtext-1.0.dtd</code> . If you have created content with versions before SCI 4.1.38, set the property to "true" to correct these errors (CoreMedia recommends to always use "true").
--------------------	--

```
sql.store.driver
```

Type	String
-------------	--------

Description	The JDBC driver used to connect to the database. Example: <code>oracle.jdbc.driver.OracleDriver</code>
--------------------	--

```
sql.store.folder-index-timeout
```

Type	Duration
-------------	----------

Default	1m
----------------	----

Description	Maximum time to wait for an enabled folder index to become up-to-date when a content query with BELOW-clause is executed. Folder index updates may take
--------------------	---

some time, if large folders were moved. If the maximum time has been reached, the query will be executed as if the folder index was offline. Set to negative value for unlimited.

sql.store.generate-blob-ids	
Type	Boolean
Default	true
Description	Whether the store generates blob ids on its own. Defaults to "true". If set to "false", the store relies on externally provided blob ids. This allows the use of shared media stores for MLS and RLS. This flag must not be changed after the first start of a Content Server. NOTE: this feature switch is currently reserved exclusively for integration of the Image Transformation Service and Blob Service and must not be used in projects otherwise.
sql.store.isolation	
Type	String
Description	Define the transaction isolation level. An empty value uses the driver default setting.
sql.store.log-driver-messages	
Type	Boolean
Default	false
Description	Write messages of the JDBC driver to the logs • <code>hox.corem.server.sql.SQLStore</code> and • <code>com.coremedia.cotopaxi.server.DatabaseProperties</code>.
sql.store.password	
Type	String
Description	The password of the user at the database.

```
sql.store.prepared-statement-cache-size
```

Type	Integer
Default	-1
Description	If set, denotes the maximum number of prepared statements that is kept open per database connection.

```
sql.store.read-only-media-store-names
```

Type	List<String>
Description	A list of names of read-only media stores. Such stores receive their binary data from another Content Server which also generates the blob ids that must then be used when creating blobs in this Content Server. Because the blob ids are provided externally, <code>sql.store.generate-blob-ids=false</code> is required if this property is non-empty. This makes it possible to share blobs between MLS and RLS. The blob collector will not remove any blobs from shared stores. NOTE: this feature is currently reserved exclusively for integration of the Image Transformation Service and Blob Service and must not be used in projects otherwise.

```
sql.store.replace-substitute
```

Type	Boolean
Default	false
Description	There are problems with zOS DB2 databases to store characters that are not in the databases character set. Such characters are converted to the character with the code "0x001A" upon read. If the property is set to "true", this character will be replaced with the character defined in <code>sql.store.substituteCharacter</code> , thereby avoiding the replacement character that is illegal in XML texts.

```
sql.store.sgml-cache-interval-seconds
```

Type	Duration
Default	0

Description	Duration between cache statistics log entries to facility cap.server.store.sgml-cache
--------------------	---

```
sql.store.sgml-cache-size
```

Type	org.springframework.util.unit.DataSize
-------------	--

Default	10MB
----------------	------

Description	Total size of XML objects cached in memory. Unit defaults to bytes.
--------------------	---

```
sql.store.sgml-cache-size-bytes
```

Type	org.springframework.util.unit.DataSize
-------------	--

Deprecation	This property has been deprecated since 2404.1 and will be removed in a future version.
--------------------	---

Use `sql.store.sgml-cache-size` instead.

Reason:

Aligning name with all other capacity configuration options.

```
sql.store.substitute-character
```

Type	String
-------------	--------

Default	?
----------------	---

Description	The character, with which "0x001A" should be replaced.
--------------------	--

```
sql.store.url
```

Type	String
-------------	--------

Description	The URL of the database to connect to.
--------------------	--

Example: jdbc:oracle:thin@HostName:Port:CM Replace HostName and Port with the appropriate values of the database host. Don't replace HostName with "localhost", this may cause problems with some JDBC drivers.

```
sql.store.use-ctes-for-below-queries
```

Type	Boolean
Default	true
Description	Whether to use recursive common table expressions (aka "WITH RECURSIVE") in SQL statements for content queries with BELOW-clauses (aka descendantOf clauses) to find contents below some folder, if the folder index is not online.
<code>sql.store.user</code>	
Type	String
Description	The username at the database.

Table 3.7. SQL Properties

3.2.5 Properties for Replicator Configuration

The protocol used for the connection to the *Master Live Server* can be either CORBA (default) or gRPC. Set properties `repository.use-grpc`, `replicator.publication-http-base-uri`, and `spring.grpc.client.channel.cap.target` on the *Replication Live Server* for gRPC connection to the *Master Live Server* (see [Section 3.12.1, "Unified API Spring Boot Client Properties" \[233\]](#)).

<code>replicator.auto-restart</code>	
Type	Boolean
Default	true
Description	Restarts the Replicator sessions if lost (true).
<code>replicator.check-blob-hashes</code>	
Type	Boolean
Default	true

Description	Whether the replicator checks that the hashes of uploaded blobs are the same as the hashes stored on the Master Live Server. A failed check indicates that the blob store configuration differs between the Master Live Server and the Replication Live Server. Defaults to true. Only set to false to force replication to resume in case of an emergency. The resulting inconsistent blobs ids may affect blob retrieval and the affected Replication Live Server should be recreated afterwards.
--------------------	---

```
replicator.check-stream
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	Defines if the Replication Live Server checks the event queue for connection (true). The event stream is not checked during the initial replication.
--------------------	--

```
replicator.check-timeout
```

Type	Duration
-------------	----------

Default	5m
----------------	----

Description	The duration after which the Replication Live Server checks the event queue for connection. (Default unit is seconds)
--------------------	---

```
replicator.chunking-threshold
```

Type	Integer
-------------	---------

Default	10000
----------------	-------

Description	Sets the maximum number of events that is fetched from the Master Live Server in one chunk during startup; 0 indicates no limit. Lowering the threshold will reduce main memory consumption at the cost of startup times.
--------------------	---

```
replicator.enable
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description Defines if the Replicator should be started on start of the Replication Live Server (true).

```
replicator.force-online-switch
```

Type Boolean

Default false

Description Switch to runlevel online even if the repository contents have not yet been replicated from the Master Live Server. Ignored if not cap.server.init-runlevel=online.

```
replicator.log-events
```

Type Boolean

Default false

Description All repository events will be logged (true).

```
replicator.max-accepted-lag
```

Type Long

Default 100

Description If the Replication Live Server is offline after a consistent replication but it is more than the given number of events behind the current timestamp, then the Replication Live Server will not go online until it catches up.

```
replicator.packager-flush-size
```

Type Integer

Default 500

Description Sets the maximum number of events that the packager holds in main memory before flushing to disk. Lowering the flush size will reduce main memory consumption at the cost of write performance.

```
replicator.password
```

Type	String
Default	replicator
Description	Configures the password to authenticate against the Master Live Server.
<code>replicator.publication-http-base-uri</code>	
Type	URI
Description	HTTP base URI of the Master Live Server. Only required if gRPC is used, i.e., "repository.use-grpc=true".
<code>replicator.publication-ior-url</code>	
Type	String
Description	Sets the location where the Replicator can find the IOR of the Master Live Server. Only effective if CORBA is used, i.e., "repository.use-grpc" is unset or "false".
<code>replicator.publication-send-session-token-in-url</code>	
Type	Boolean
Default	true
Description	If true, the session token for connection to Master Live Server servlets (for blob up/download and CPU usage) will be sent as a URL query parameter. Otherwise, it will be sent as a request field. While sending the token as a query parameter is inherently insecure, it is the default to keep backward compatibility. Unless connection to an older server is required, it is recommended to set this property to false.
<code>replicator.restart-replicator-on-error</code>	
Type	Boolean
Default	true
Description	Restart the Replicator if an error occurs on Replication Live Server side (true). Otherwise, the Replicator will be stopped.

replicator.shutdown-server-on-error	
Type	Boolean
Default	false
Description	Stop the Replication Live Server if an error occurs on Replication Live Server side.
replicator.startup-timeout	
Type	Boolean
Default	false
Description	Defines if the Replication Live Server waits on start for the Replicator to connect to the Master Live Server.
replicator.tmp-dir	
Type	String
Description	The path to the folder where the Replication Live Server should write temporary data during replication. The path is relative to cap.server.base-dir of the replicator application, so you will probably use an absolute path. During initial replication of a complete repository, a huge amount of data will be written to this directory, in the order of the repository size. So, this directory should be located on a suitable file system with enough space left. During normal operation, only newly published blob content will be buffered on local disk.
replicator.user	
Type	String
Default	replicator
Description	Configures the user name to authenticate against the Master Live Server.

Table 3.8. Replicator Properties

3.2.6 Properties for Timezone and IOR

cap.client.server.ior.url	
Value	URL format http://<server>:<port> /ior
Default	
Description	This property determines where to get the IOR of the <i>contentserver</i> . <server> must be the name of the <i>Content Server</i> host. For <port> you have to set the server's web server HTTP port.
cap.client.timezone.default	
Value	TimeZone
Default	Europe/Berlin
Description	This parameter determines the used timezone. The standard value is <i>Europe/Berlin</i> . More time zones are documented in the class <code>java.util.TimeZone</code> .

Table 3.9. capclient.properties

3.3 Headless Server Properties

Find more related properties here:

- [Section 3.14, “Web Security Properties” \[242\]](#)

3.3.1 Headless Server Spring Boot Properties

<code>caas.bypass-filter-predicates</code>	
Default	false
Description	If enabled, FilterPredicates may be ignored via request header 'X-Ignore-Filters'. If disabled, the header is ignored.
<code>caas.cache-specs</code>	
Type	Map<String,String>
Description	(Caffeine Cache) configuration: Cache specs by cache name. Size based eviction: • <code>maximumSize=[long]</code> - eviction occurs when the configured size limit of the cache is exceeded • <code>maximumWeight=[long]</code> - The values are removed from the cache when the weight is exceeded Time based eviction: • <code>expireAfterWrite</code> = entry is expired after period is passed since the last write occurs • <code>expireAfterAccess</code> = entry is expired after period is passed since the last read or write occurs • <code>refreshAfterWrite</code> = refresh entries after a defined period automatically Durations are represented by an integer, followed by one of "d", "h", "m", or "s", representing days, hours, minutes, or seconds respectively.
<code>caas.doctype.doc-type-collection</code>	

Type	String
Default	CMCollection
Description	The name of the doctype for Collections.
<code>caas.doctype.doc-type-linkable</code>	
Type	String
Default	CMLinkable
Description	The name of the doctype for Linkables.
<code>caas.doctype.doc-type-location-taxonomy</code>	
Type	String
Default	CMLocTaxonomy
Description	The name of the doctype for Location Taxonomies.
<code>caas.doctype.doc-type-navigation</code>	
Type	String
Default	CMNavigation
Description	The name of the doctype for Navigations.
<code>caas.doctype.doc-type-picture</code>	
Type	String
Default	CMPicture
Description	The name of the doctype for Pictures.
<code>caas.doctype.doc-type-taxonomy</code>	
Type	String

Default	CMTaxonomy
Description	The name of the doctype for Taxonomies.
<code>caas.exception-error-map</code>	
Type	Map<Class<? extends Throwable>, graphql.ErrorType>
Default	[RuntimeException.class:ErrorType.DataFetchingException]
Description	Map of exception classes to GraphQL error types to be resolved by the data fetcher exception resolver.
<code>caas.expose-extensions</code>	
Type	Boolean
Default	false
Description	Set to true to expose extensions as described by the GraphQL spec .
<code>caas.forward-cookies</code>	
Type	Boolean
Default	false
Description	Enables http cookie forwarding of incoming requests on endpoint 'graphql' and for persisted queries on '/caas/v1/*' towards outgoing, secondary requests to 'graphql' (triggered by '/preview') and cae remote link resolving.
<code>caas.forward-header-names</code>	
Type	List<String>
Description	Comaseparated enumeration of http headernames to be forwarded for incoming requests on endpoint 'graphql' and for persisted queries on '/caas/v1/*' towards outgoing, secondary requests to 'graphql' (triggered by '/preview') and cae remote link resolving.
<code>caas.graphql-restmapping-controller.enabled</code>	

Deprecation	This property has been deprecated and will be removed in a future version. Use <code>spring.autoconfigure.exclude</code> instead. Reason: The property was removed due to changes in the implementation. The rest mapping is now auto configured by 'com.coremedia.caas.web.rest.RestMappingAutoConfiguration' and is enabled by default. To disabled it, add the auto configuration class to the corresponding Spring property.
--------------------	---

`caas.graphql.expose-secrets`

Type	Boolean
Default	false
Description	Expose encrypted values from settings in their encrypted version. Defaults for security reasons to false.

`caas.graphql.introspectionEnabled`

Deprecation	This property has been deprecated and will be removed in a future version. Use <code>spring.graphql.schema.introspection.enabled</code> instead. Reason: This property was dropped and replaced by the Spring-GraphQL property. The new property is set to 'false' in headless built-in application properties for security reasons. It should be set to 'true' if necessary, e.g. in a preview environment.
--------------------	---

`caas.graphql.max-execution-timeout`

Type	Duration
Default	0
Description	Limits the allowed execution time for a query, set in milliseconds. 0 = no timeout.

`caas.graphql.max-query-complexity`

Type	Integer
Default	0

Description Limits the complexity of a graphql query if set to a value greater than 0. 0 = off.

```
caas.graphql.max-query-depth
```

Type Integer

Default 30

Description Limits the depth of a graphql query if set to a value greater than 0. 0 = off.

```
caas.graphql.max-search-limit
```

Deprecation This property has been deprecated and will be removed in a future version.
Use `caas.search.max-search-limit` instead.

```
caas.graphql.repository-path-exclude-patterns
```

Type List<String>

Default [/Settings/Options/Settings/Internal/**, /Sites/**/Options/Settings/Internal/**, /Settings/Options/Settings/Feedback Hub, /Sites/**/Options/Settings/Feedback Hub, /Settings/Options/Settings/Content Hub, /Sites/**/Options/Settings/Content Hub, /Settings/Options/Settings/GlobalLink, /Settings/Options/Settings/Translation Services, /Sites/**/Options/Settings/Translation Services, /Settings/Options/Settings/Elastic Social Credentials]

Description Repository paths excluded from retrieval via GraphQL endpoint. The paths may contain glob style expressions. Note, that the document name must be part of the path, if you do not use glob style expressions and want to protect a distinct document.

Examples:

```
caas.graphql.repository-path-exclude-patterns[0]=/A/B/C/D
caas.graphql.repository-path-exclude-patterns[1]=/A/*C/*
caas.graphql.repository-path-exclude-patterns[2]=/A/B/**
caas.graphql.repository-path-exclude-patterns[3]=/A/**C/*
```

```
caas.log-requests
```

Type Boolean

Default	false
Description	Enables logging for HTTP requests (excluding OPTIONS)
<code>caas.media-controller.blob-validator-enabled</code>	
Type	Boolean
Default	true
Description	Enables the blob validator. If enabled, the blob of the media is checked for existence. This validator is executed after the crop name validator.
<code>caas.media-controller.crop-name-validator-enabled</code>	
Type	Boolean
Default	true
Description	Enables the crop name validator. If enabled, the given crop name and width is checked for validity. This validator is executed after the filename validator.
<code>caas.media-controller.filename-validator-enabled</code>	
Type	Boolean
Default	true
Description	Enables the filename validator. If enabled, the given filename is checked vs the resolved filename. This validator is executed after the validity date validator.
<code>caas.media-controller.hash-validator-enabled</code>	
Type	Boolean
Default	true
Description	Enables the hash validator. If enabled, the given hash of the blob is checked for validity. This validator is executed after the blob validator.
<code>caas.media-controller.path-validator-enabled</code>	

Type	Boolean
Default	true
Description	Enables the path validator. If enabled, the given path is checked for validity. This validator is executed first.
<code>caas.media-controller.validity-date-validator-enabled</code>	
Type	Boolean
Default	true
Description	Enables the validity date validator. If enabled, the validity date of the media is checked. This validator is executed after the path validator.
<code>caas.media.hash-property-names</code>	
Type	List<String>
Default	[width, height, widthRatio, heightRatio]
Description	List of global or site specific transformation option names defined in image transformation settings whose values are to be considered to compute image hashes.
<code>caas.preview</code>	
Type	Boolean
Default	false
Description	Enables preview mode
<code>caas.querylist-search-cache-for-seconds</code>	
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>caas.search.cache.querylist-search-cache-for-seconds</code> instead.
<code>caas.search.cache.querylist-search-cache-for-seconds</code>	

Type	Duration
Default	-1s
Description	Duration to cache query list search results. Set it to "-1" to deactivate the query list cache.
<code>caas.search.cache.seconds</code>	
Type	Duration
Default	-1s
Description	Duration to cache search query results. Set it to "-1" to deactivate the search-query cache.
<code>caas.search.enabled</code>	
Type	Boolean
Default	true
Description	Enable graphql search extensions.
<code>caas.search.load-search-schema</code>	
Type	Boolean
Default	true
Description	Load the GraphQL search schema extensions.
<code>caas.search.max-search-limit</code>	
Type	Integer
Default	200
Description	Limit the count of search results hits. Defaults to 200.
<code>caas.search.solr.collection</code>	

Type	String
Default	preview
Description	Solr collection to use for CaaS search.
<code>caas.site-filter-config-properties.excluded-field-names</code>	
Type	Collection<String>
Default	[localizedVariant, localizedVariants, derivedSites]
Description	Fields to be excluded from site restrictions check. To disable all fields (thus prohibiting all objects not belonging to the filtered site), add the config property without any field names (=empty list). Adding additional fields requires to add the defaults as mentioned above to the config property! Defaults to "localizedVariant", "localizedVariants" and "derived-Sites".
<code>caas.solr.collection</code>	
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>caas.search.solr.collection</code> instead.
<code>caas.stax.context-trace-enabled</code>	
Type	Boolean
Default	false
Description	When 'true', wraps ContextHandlers and OutputHandlers while parsing in order to tracelog all parsing events. This is for debugging purposes only. Defaults to false. Additionally this requires the loglevel for 'com.coremedia.caas.richtext.stax.handler.output' to be set to trace level: <code>logging.level.com.core-media.caas.richtext.stax.handler.output = trace</code>
<code>caas.stax.max-aliases-for-collections</code>	
Type	Integer

Default	50
Description	Configures the maximum number of allowed yaml aliases used by snakeyaml instances while parsing rich text configurations. Defaults to 50.
<code>caas.stax.suppress-root-tag</code>	
Type	Boolean
Default	true
Description	When 'false', the root tag in the string result of a rich text transformation is not suppressed. Defaults to true.
<code>caas.stax.table-section-xml-reader.footer-class</code>	
Type	String
Default	tr--footer
Description	CSS class to identify footer rows. Defaults to "tr--footer".
<code>caas.stax.table-section-xml-reader.header-class</code>	
Type	String
Default	tr--header
Description	CSS class to identify header rows. Defaults to "tr--header".
<code>caas.stax.table-section-xml-reader.max-nesting-depth</code>	
Type	Integer
Default	0
Description	Configures the maximum allowed nesting depth of tables for the table section filter. If this limit is exceeded, additional nested tables will be ignored and passed through by the table section filter. If 0, the filter is effectively disabled. Defaults to 0.

caas.strict-blob-hash-validation	
Type	Boolean
Default	true
Description	If true, the hash provided in media URLs is strictly validated. If false, an outdated or wrong hash will result in an HTTP redirect (301) to the URL with the correct hash. Default is true.
caas.strict-transport-security-header.include-subdomains	
Type	Boolean
Default	true
Description	If true, adds the 'includeSubdomains' flag to the Strict-Transport-Security HTTP response header.
caas.strict-transport-security-header.max-age	
Type	Duration
Default	730d
Description	Max-Age of the Strict-Transport-Security HTTP response header. (Default unit is seconds)
caas.strict-transport-security-header.preload	
Type	Boolean
Default	false
Description	If true, adds the 'preload' flag to the Strict-Transport-Security HTTP response header.

Table 3.10. Headless Server Properties

3.3.2 Persisted Query Properties

caas.persisted-queries.allow-list	
Type	Boolean
Default	false
Description	Query allow listing. Set this to true to disallow any queries not loaded by some of the persisted queries pattern above.
caas.persisted-queries.apollo-query-map-resources-pattern	
Type	String
Default	classpath*:graphql/queries/apollo*.json
Description	Resource pattern for persisted queries in Apollo Persisted Queries format
caas.persisted-queries.automatic	
Type	Boolean
Default	true
Description	Enable Apollo Automatic Persisted Queries
caas.persisted-queries.exclude-file-name-pattern	
Type	String
Default	.*Fragment(s)?.graphql
Description	Resource pattern for persisted queries in Relay Compiler format
caas.persisted-queries.query-resources-pattern	
Type	String
Default	classpath*:graphql/queries/*.graphql

Description	Resource pattern for persisted queries, one query per resource file. The filename w/o extension serves as the query id. See also: Spring PathMatchingResourcePatternResolver class.
--------------------	---

```
caas.persisted-queries.relay-query-map-resources-pattern
```

Type	String
-------------	--------

Default	graphql/queries/relay*.json
----------------	-----------------------------

Description	Resource pattern for persisted queries in Relay Compiler format
--------------------	---

```
caas.persisted-queries.whitelist
```

Deprecation	This property has been deprecated and will be removed in a future version. Use <code>caas.persisted-queries.allow-list</code> instead.
--------------------	---

Reason:

Renamed property name to a political correct name.

```
caas.rest.jslt-enabled
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	En/disables the JSLT transformation processor on REST endpoints
--------------------	---

```
caas.rest.jslt-transformations-pattern
```

Type	String
-------------	--------

Default	classpath*:transformations/*.jslt
----------------	-----------------------------------

Description	Returns the path pattern where to find the jslt transformation resources within headless.
--------------------	---

```
caas.rest.query-mapping-pattern
```

Type	String
-------------	--------

Default	classpath*:graphql/rest-mapping/*.properties
Description	Returns the path pattern where to find rest mapping resources within headless. <code>com.coremedia.caas.web.persistedqueries.impl.PersistedQueryAutoConfiguration</code>
Type	List
Description	To disable the feature 'persisted-queries' add the class 'com.core-media.caas.web.persistedqueries.impl.PersistedQueryAutoConfiguration' to 'spring.autoconfigure.exclude'. <code>com.coremedia.caas.web.rest.RestMappingAutoConfiguration</code>
Type	List
Description	To disable the feature 'persisted-query-rest-mapping' add the class 'com.coremedia.caas.web.rest.RestMappingAutoConfiguration' to 'spring.autoconfigure.exclude'.

Table 3.11. Persisted Query Properties

3.3.3 Metadata Properties

<code>caas.metadata.enabled</code>	
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>spring.autoconfigure.exclude</code> instead. Reason: The property was removed due to changes in the implementation. The metadata schema is now auto configured by 'com.core-media.caas.web.metadata.MetadataAutoConfiguration' and is enabled by default. To disabled it, add the auto configuration class to the corresponding Spring property.
<code>caas.metadata.property-mapping-default-filename</code>	
Type	String

Default	propertyMapping.json
Description	Name of the default property mapping definition file
	<code>caas.metadata.property-mapping-location</code>
Type	String
Default	classpath*:graphql/metadata/propertyMapping*.json
Description	Location of the property mapping file(s)
	<code>com.coremedia.caas.web.metadata.MetadataAutoConfiguration</code>
Type	List
Description	To disable the feature 'metadata-query-root' add the class 'com.core-media.caas.web.metadata.MetadataAutoConfiguration' to 'spring.autoconfigure.exclude'.

Table 3.12. Metadata Root Properties

3.3.4 Remote Service Adapter Properties

	<code>caas.remote.base-url</code>
Type	String
Default	http://localhost:8080/webapp/servlet/service-endpoint
Description	Base URL to the remote service handler

Table 3.13. Headless Server Remote Properties

3.3.5 Headless Server Cache Control Properties

	<code>caas.cache-control.for-url-pattern</code>
--	---

Type	<code>Map<String,org.springframework.boot.autoconfigure.web.WebProperties\$Resources\$Cache\$Cachecontrol></code>
-------------	---

Description	Map of cache control configuration objects for HTTP Cache-Control header. The configuration is URL pattern specific. Patterns are matched using
--------------------	---

```
org.springframework.util.AntPathMatcher
```

. For every URL the content validity dates are taken into account to compute the max-age value. A negative configured max-age value disables cache control for the URL pattern. A configured positive max-age value is taken as maximum value for the URL pattern. Thus, if X is configured and Y is computed, then $\min(X,Y)$ is returned as max-age.

Note that max-age values must not be higher than 31536000 (365 days in seconds).

Also note that, by consequence, a response including properties of a content that is valid for more than a year according to its content settings, and for whose URL pattern no max-age value is defined, will include a max-age value of a year.

```
caas.cache-control.s-max-age-factor
```

Type	Double
-------------	--------

Default	0
----------------	---

Description	This value is used to compute a s-maxage header in relation to a recognized max-age value. If the value is 0.0 then the s-maxage value is not automatically adjusted. In particular this means the s-maxage value isn't written in the Cache-Control header unless it was otherwise configured, for example with:
--------------------	---

```
caas.cache-control.for-url-pattern[/example/**].s-max-age=1m
```

But please note, a value other than 0.0 takes precedence over a static configured value. Then s-maxage takes the value of $\text{max-age} * \text{sMaxAgeFactor}$.

```
caas.cache-control.url-path.remove-semicolon-content
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description	Set if everything after a ";" (semicolon) should be cut from the request URI for path matching. E.g., to see all matrix parameters of fragment urls this should be set to "false".
<code>caas.cache-control.url-path.url-decode</code>	
Type	Boolean
Default	true
Description	Set if the request URI should be decoded for path matching.
<code>caas.cache-control.url-path.use-always-full-path</code>	
Type	Boolean
Default	false
Description	Set if the full path within the current web application context is used for path matching. By default, this is set to "false" which means that the servlet path is cut off.

Table 3.14. Headless Server Cache Control Properties

3.3.6 Headless Server Cache Key Properties

<code>cache.capacities.com.coremedia.caas.search.solr.SolrQueryCacheKey</code>	
Type	Long
Default	5000
Description	Sets the cache size for the solr query cache.

Table 3.15. Headless Server Cache Key Properties

3.3.7 Properties of External Frameworks

graphql.enabled	
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>spring.graphql.graphiql.enabled</code> instead. Reason: The GraphiQL-UI is now part of Spring GraphQL, thus configured by Spring properties.
graphql.url	
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>spring.graphql.path</code> instead. Reason: All transport aspects are now configured by Spring GraphQL.

Table 3.16. Headless Server External Framework Properties

3.4 Studio Properties

Find more related properties here:

- [Section 3.8, “My Edited Content and Workflow Lists Properties” \[162\]](#)

3.4.1 Studio Configuration

The following list contains configuration properties for the *Studio*.

studio.auto-logout.delay	
Type	Duration
Default	1800s
Description	The delay before a Studio session is automatically disconnected after inactivity. Inactivity is checked on the client. Set to 0 to disable auto logout (not recommended). For clients that cannot reach the Studio web application, the timeout of the server session determines the timeout of the Studio session.
studio.bulk-operation-thread-pool.max-threads	
Type	Integer
Default	10
Description	Maximum number of threads to use for the bulk operation thread pool.
studio.cloud.enable	
Type	Boolean
Default	false
Description	Whether to enable cloud support.
studio.core.json.pretty-print	
Type	Boolean

Default	false
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>spring.jackson.serialization.indent_output</code> instead. Reason: Not used anymore.
<code>studio.dashboard-refresh-interval</code>	
Type	Duration
Default	30000ms
Description	The interval at which the dashboard is refreshed when visible (0 to disable automatic refresh).
<code>studio.dashboard.refresh.interval</code>	
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>studio.dashboard-refresh-interval</code> instead.
<code>studio.default-picture.content.paths</code>	
Type	Map<String,String>
Description	Configure <code>com.coremedia.cap.content.Content</code> property paths that point to picture blobs. Paths are given per <code>com.coremedia.cap.content.ContentType</code> , e.g. <code>'studio.default-picture.content.paths.CMPicture=data'</code> . The mechanism is recursive: it is possible to configure a path to other <code>com.coremedia.cap.content.Content</code> items, e.g. <code>'studio.default-picture.content.paths.CMTeasable=pictures'</code> .
<code>studio.default-picture.resolving.enabled</code>	
Type	Boolean
Default	true
Description	Configure whether server-side default picture resolving is enabled.
<code>studio.default-time-zone</code>	

Type	String
Default	Europe/Berlin
Description	Default time zone in CoreMedia Studio. Make sure that the default time zone is included in the studio.timeZones list. Defaults to 'Europe/Berlin'.
<code>studio.differencing.max-markup-size</code>	
Type	org.springframework.util.unit.DataSize
Default	300000B
Description	The maximum size of a markup object (in bytes as default unit) for which differences with other markup can be computed. By default, this value is set to 300,000 bytes, which amounts to approximately 15,000 words.
<code>studio.expose-license-info</code>	
Type	Boolean
Default	false
Description	Whether to expose license information along with the system information.
<code>studio.jobs.max-threads</code>	
Type	Integer
Default	10
Description	Maximum number of threads to use for the jobs service thread pool.
<code>studio.locales</code>	
Type	List<String>
Default	en,de

Description	A comma-separated list of locales from which the user can choose the locale in which the Studio is displayed. If unset, the list locales defaults to English or German.
--------------------	---

```
studio.login-url
```

Type	String
-------------	--------

Description	If non-empty, a URL to which the Studio client should redirect a user who is not yet logged in. This supports single-sign-on solutions using an external login page. If empty, Studio shows a login form itself.
--------------------	--

```
studio.preview-controller-pattern
```

Type	String
-------------	--------

Default	preview?id={0}
----------------	----------------

Description	Get the preview controller pattern. If it is empty or not defined, then use the default preview controller pattern.
--------------------	---

```
studio.preview-url-prefix
```

Type	String
-------------	--------

Description	The URI prefix of the CAE preview web app. This prefix is used to compose preview URLs for contents, assuming that a standard CAE preview controller is mapped at the path 'preview'. Since studio and cae are usually deployed as independent spring boot application or as a containerized application (docker/kubernetes), it is necessary to provide an absolute URL to load the preview, like 'https://your-host-name.com:40980/blueprint/servlet'
--------------------	---

```
studio.preview-url-whitelist
```

Type	List<String>
-------------	--------------

Description	Configures a list of valid preview origins. The preview integration does only work for contents from listed origins. Wildcards (*) may be used for valid origin entries. If left blank (or if the property is not listed at all), the origin of studio.previewUrl-Prefix is the only accepted origin. The URLs of the preview origins must not contain a trailing slash. NOTE: Once a whitelist is configured, the Studio has no chance to set a target origin in outgoing messages anymore. Be aware that this
--------------------	---

is a minor security drawback. EXAMPLE: studio.previewUrlWhitelist=https://host1:port1,\ https://host2:port2,\ http://localhost.*,\ *.company.com Enabling Elastic Social tenants in the embedded preview requires including a placeholder CoreMedia Studio then replaces the token with the current tenant. In a Blueprint related project, this could be: studio.previewUrlPrefix=http://{O}.local-host:40081/blueprint/servlet

studio.publication-priority	
Type	Integer
Default	60
Description	The priority of publications started in Studio. For an overview of relevant priorities take a look at "com.coremedia.cap.content.publication.Publication".
studio.rest.changeset-max-iterations	
Type	Integer
Default	10
Description	Maximum number of iterations to complete a change set.
studio.rest.changeset-max-updated-links-recursion-depth	
Type	Integer
Default	10
Description	Maximum depth to follow updated links while completing a change set.
studio.rest.invalidation-source.caplist.capacity	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
studio.rest.invalidation-source.cmec.capacity	

Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.content.capacity</code>	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.model.capacity</code>	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.notification.capacity</code>	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.project.capacity</code>	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.system.capacity</code>	
Type	Integer

Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.timeout</code>	
Type	Duration
Default	20000ms
Description	The timeout to wait for new events to be propagated via the invalidation sources.
<code>studio.rest.invalidation-source.todo.capacity</code>	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.invalidation-source.workflow.capacity</code>	
Type	Integer
Default	10000
Description	The capacity of the invalidation source.
<code>studio.rest.max.iterations.complete.changeset</code>	
Type	Integer
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>studio.rest.changeset-max-iterations</code> instead.
<code>studio.rest.network.max-requests</code>	
Type	Integer
Default	15

Description	The maximum amount of parallel pending requests executed to one host.
<code>studio.rest.search-service.default-search-limit</code>	
Type	Integer
Default	5000
Description	The default search limit;
<code>studio.rest.search-service.max-search-limit</code>	
Type	Integer
Default	80000
Description	The maximum search limit
<code>studio.rest.search-service.post-method-threshold</code>	
Type	Integer
Default	2048
Description	The threshold for switching from GET to POST requests for search queries. If the query string exceeds this length, a POST request is used instead of GET. Default is 2048 characters, which is a common recommendation for maximum URL length.
<code>studio.rest.search-service.replace-special-characters-in-prefix-or-wildcard-query</code>	
Type	Boolean
Default	true
Description	Whether special characters in prefix and wildcard queries (queries ending with * or containing *) are replaced by spaces. Replacing characters makes sense because Solr does not analyze prefix and wildcard queries in the same way as it indexes text when configured with the StandardTokenizer. By replacing these characters, indexed text can still be matched and queries return the expected

hits. However, it can make sense to disable this setting, if Solr is configured to preserve special characters in indexed text.

studio.rest.search-service.semantic.available-content-types	
Type	List<String>
Description	The types of content items which are available for semantic search. Available means, that those content types (or subtypes) are fed to the solr with embeddings.
studio.rest.search-service.semantic.default-search-handler	
Type	String
Default	vectorSimilarity
Description	The default search handler for semantic search.
studio.rest.search-service.semantic.enabled	
Type	Boolean
Default	false
Description	Whether semantic search is enabled.
studio.rest.search-service.semantic.media-ranking-boost	
Type	Float
Description	Boosting value added to the ranking score for media vector similarity. Value must be between 0 and 1. Default is -0.88, which means that the media similarity score is decreased by 8%. Default was evaluated with model "amazon.nova-2-multimodal-embeddings-v1:0".
studio.rest.search-service.semantic.min-return-threshold	
Type	Float
Default	0.57

Description	Defines the minimum score to be returned when using semantic search. Value must be between 0 and 1. Default is 0.57, which means that only documents with a similarity score higher than 0.57 are returned. Default was evaluated with model "amazon.nova-2-multimodal-embeddings-v1:0".
--------------------	--

```
studio.rest.search-service.semantic.top-k
```

Type	Integer
-------------	---------

Default	100
----------------	-----

Description	The number of top K results to retrieve from the knn vector search. We consider a higher default, since the filtering after the vector search might remove some results otherwise.
--------------------	--

```
studio.rest.search-service.semantic.vector-dimension
```

Type	Integer
-------------	---------

Default	1024
----------------	------

Description	The dimension of the semantic field vector. Required to identify fields containing semantic vectors. Default is 1024, which is the dimension also used by aws titan embedding models.
--------------------	---

```
studio.struct-as-blob-property-names
```

Type	List<String>
-------------	--------------

Description	List of struct property names which are to be handled as struct properties by the client.
--------------------	---

```
studio.supported-locales-variable-name
```

Type	String
-------------	--------

Default	joo.localization.supportedLocales
----------------	-----------------------------------

Description	The variable name of the supported locales javascript snippet as expected by the client. Defaults to 'joo.localization.supportedLocales'.
--------------------	---

```
studio.system-info.enable
```

Type	Boolean
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>studio.expose-license-info</code> instead.
<code>studio.time-zones</code>	
Type	List<String>
Default	Europe/Berlin
Description	Supported time zones in CoreMedia Studio. Provide a comma separated list of supported Java time zone IDs. Defaults to 'Europe/Berlin'.
<code>studio.translation.max-dependent-content-iterations</code>	
Type	Integer
Default	100
Description	Defines how deep links should be followed, when calculating dependent content. Can be overridden via content setting.
<code>studio.translation.max-dependent-contents</code>	
Type	Integer
Default	100
Description	Defines how many dependent contents will be returned in "Translation Workflow Start window" when starting a translation workflow. Can be overridden via content setting.
<code>studio.translation.show-pull-translation-start-window</code>	
Type	Boolean
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>studio.translation.show-start-window</code> instead.

```
studio.translation.show-start-window
```

Type Boolean

Default true

Description Configures Pull Translation (translation into preferred site) behavior: When set to true the workflow dialog is displayed, otherwise the workflow is started automatically.

```
studio.upload.default-blob-property
```

Type String

Default data

Description The default blob property name of the default content to be used for bulk uploads, used if no mapping is found in the upload settings defined in content. Defaults to 'data'.

```
studio.upload.default-content-type
```

Type String

Default CMDownload

Description The default content type to be used for bulk uploads, used if no mapping is found in the upload settings defined in content. Defaults to 'CMDownload'.

```
studio.upload.upload-settings-struct-property
```

Type String

Default settings

Description The name of the struct property to be used when the UploadSettings document is read. Defaults to 'settings'.

```
studio.usermanager.enable-content-live-groups
```

Type Boolean

Default	false
Description	Enable this flag to edit the content server and live server group properties
	<code>studio.usermanager.external-domain-search-limit</code>
Type	Integer
Default	10
Description	Maximum number of search results returned when searching external user domains. Defaults to 10.
	<code>studio.usermanager.manager-groups</code>
Type	String[]
Description	List of groups that are allowed to open the user manager. The direct subgroups are allowed as well. By default, admin users are allowed to open the manager, but if this property is set, this list will be checked instead. Not that this list is only used to hide the UI from other admin users and isn't checked by the backend (yet).
	<code>studio.usermanager.manager-users</code>
Type	String[]
Description	List of users that are allowed to open the user manager. By default, admin users are allowed to open the manager, but if this property is set, this list will be checked instead. Not that this list is only used to hide the UI from other admin users and isn't checked by the backend (yet).
	<code>studio.usermanager.min-search-characters</code>
Type	Integer
Default	3
Description	Minimum number of characters required before user or group search is performed. When set to 0, all users or groups are shown by default.
	<code>studio.usermanager.offer-delete-home-folder</code>

Type	Boolean
Default	true
Description	When enabled, a checkbox option is offered to delete the user's home folder when deleting a user account. Set to false to hide this option from the user interface.
<code>studio.usermanager.protect-internal-domain</code>	
Type	Boolean
Default	false
Description	Enabling this flag will completely hide the internal domain from the user manager
<code>studio.usermanager.protected-group-names</code>	
Type	String[]
Default	[administratoren, editor, importer, compose, composer-role, approver, approver-role, publisher, publisher-role, staff, system, system-write]
Description	List of protected group names, the 'administratoren' group is part of that list, no matter if enlisted here or not
<code>studio.usermanager.protected-user-names</code>	
Type	String[]
Default	[feeder, studio, replicator, webserver, workflow, publisher, importer, admin]
Description	List of protected usernames
<code>studio.validate-before</code>	
Type	com.coremedia.rest.cap.config.ValidatedActionLevel
Default	approve

Description	Defines which content actions are not allowed if the content has error issues. As some actions are in a dependency relationship (e.g. APPROVE depends on CHECKIN), only the least inclusive actions need to be declared (e.g. CHECKIN entails the check of CHECKIN and APPROVE actions). Currently, the only supported OPTIONS are CHECKIN, APPROVE or nothing for no restriction. Defaults to APPROVE, i.e. content approval and publication is not allowed in case of error issues.
<code>studio.workflow.translation.extended-workflow</code>	
Type	Integer
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>studio.translation.max-dependent-contents</code> instead.
<code>studio.workflow.translation.max-dependent-content-iterations</code>	
Type	Integer
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>studio.translation.max-dependent-content-iterations</code> instead.
<code>studio.xml-filters.enable-multi-class-span-filter</code>	
Type	Boolean
Default	false
Description	Enables a filter for CoreMedia RichText 1.0, which splits up spans with multiple classes to dedicated spans, each with a single class only. This filter should be enabled for CKEditor 4 and should be disabled for all later versions of CKEditor or alternative rich text editing component frameworks.

Table 3.17. Studio Properties

3.4.2 Available Locales Configuration

The following list contains configuration properties for the available locales in *Studio*.

available-locales.content-path	
Type	String
Default	/Settings/Options/Settings/LocaleSettings
Description	Path of the content that defines the available locales. Defaults to '/Settings/Options/Settings/LocaleSettings'.
available-locales.property-path	
Type	String
Default	settings.availableLocales
Description	Property path to the Struct-StringListProperty containing the locales. Defaults to 'settings.availableLocales'.

Table 3.18. Available Locales Properties

3.4.3 Shared Multi-Site Configuration

The following list contains configuration properties for the multi-site feature, that are typically shared between different applications (or they must be kept in sync).

multisite.transfer.relations.enabled	
Type	Boolean
Default	true
Description	Whether referrer relation augmentation is enabled globally. When set to false, no relation resolvers are created, regardless of individual relation settings.

multisite.transfer.relations.referrer	
Type	Map<String, com.coremedia.cap.multisite.transfer.RelationsConfigurationProperties\$ReferrerSlot>
Description	Named referrer relations. Each map entry defines how referencing content items are discovered for a given content type. Default relations: • taxonomy — content-type: CMTaxonomy, source-property: children • taxonomy-root — content-type: CMTaxonomy, source-property: settings, source-content-type: CMSettings • navigation — content-type: CMNavigation, source-property: children • navigation-root — content-type: CMNavigation, source-property: root, source-content-type: CMSite Customization can add, override, or disable individual entries. To disable a specific relation, set its enabled property to false.
translatable.predicate.property-path	
Type	Map<String, Boolean>
Description	Marks properties or nested (Struct) properties as translatable. Typically applied to nested properties in Structs, which cannot be individually marked as translatable in the content type model. The map entry's key represents the property path starting from the content type that declares the property. A map entry's value of true marks the property selected by the property path as translatable. Example Configuration: <pre>translatable.predicate.property-path[CMLinkable.localSettings.callToActionCustomText]=true</pre> For Struct properties containing dots in their names, you may use bracket syntax instead:

```
translatable.predicate.property-path[QMLinkable.localSettings['with.dots.in.name']]=true
```

Note, that a value of false cannot veto a translatable decision by another predicate. The parsed configuration will just ignore those property paths marked with false from this configuration.

```
translate.item.include-unchanged-properties
```

Type	Boolean
Default	true
Description	Whether properties that have not changed since last translation are also included in items for translation.

```
translate.item.transform.failure.mode
```

Type	String
Default	fail
Description	Controls behavior when no matching transformers are available to eventually represent a property in XLIFF. Available values: fail, warn, and none.

```
translate.xliff.export.excludeContentName
```

Type	Boolean
Default	true
Description	Only if set to false, content names are included in the XLIFF export.

Table 3.19. Shared Multi-Site Properties

3.4.4 Content Configuration

The following list contains configuration properties for content repository paths with special meaning in *Studio*.

```
content.global-configuration-path
```

Type	String
Default	/Settings
Description	A global folder where additional settings are read from. Several Studio plugins lookup settings content from here. Defaults to 'Settings'.
<code>content.site-configuration-path</code>	
Type	String
Default	Options/Settings
Description	A local folder where additional settings are read from. Several Studio plugins lookup settings content from here. Defaults to 'Options/Settings'.

Table 3.20. Content Properties

3.4.5 Navigation Validator Configuration

The following list contains configuration properties for validating the navigation structure in *Studio*.

<code>validators.navigation.ignore-path</code>	
Type	List<String>
Default	[/Settings, /Home, /System]
Description	If a content path matches one of these values or is a subfolder of it, the corresponding validator won't continue it's validation. Wildcard annotations are allowed here, e.g. '/Home/*' matches '/Home/Adam' or '/Sites/*/Navigation/*' matches '/Sites/Chef Corp./English/Navigation/Content ABC'.

Table 3.21. Navigation Validators Properties

3.4.6 Preview URL Service Properties

The following list contains configuration properties related to the Multi Preview Menu in *Studio*.

preview.urlservice.content-type	
Type	String
Default	CMSettings
Description	Defines the content type that is used to configure previews.
preview.urlservice.global-path	
Type	String
Default	/Settings/Options/Settings/Multi Preview
Description	Defines the repository folder path with contents that configure previews for all sites and contents without a site. Subfolders are ignored.
preview.urlservice.headless-preview-host	
Type	String
Description	Defined the host of the headless preview server.
preview.urlservice.preview-url-allow-list	
Type	List<String>
Description	Defines the commercePreviewUrlAllowList. In order to limit the urls in the preview frame and prevent CSRF.
preview.urlservice.property	
Type	String
Default	settings

Description	Defines the name of the struct property in the configured content type.
<code>preview.urlservice.site-path</code>	
Type	String
Default	Options/Settings/Multi Preview
Description	Defines the repository folder path below a site root folder with contents that configure previews for one site. Subfolders are ignored.

Table 3.22. Preview URL Service Properties

3.4.7 Content Security Policy Configuration

The following list contains configuration properties related to Content Security Policy (CSP) in the *Studio*.

<code>studio.security.csp.child-src</code>	
Type	List<String>
Description	List of values for the 'child-src' policy directive. This directive is only applied if extended by plugins or extensions.
<code>studio.security.csp.connect-src</code>	
Type	List<String>
Description	List of values for the 'connect-src' policy directive. Defaults to 'self'.
<code>studio.security.csp.font-src</code>	
Type	List<String>
Description	List of values for the 'font-src' policy directive. Defaults to 'self'.
<code>studio.security.csp.frame-ancestors</code>	

Type	List<String>
Description	List of values for the 'frame-ancestors' policy directive. Defaults to 'self'.
Deprecation	This property has been deprecated and will be removed in a future version.
<code>studio.security.csp.frame-src</code>	
Type	List<String>
Description	List of values for the 'frame-src' policy directive. The hierarchy of default values for this directive is as follows 1. studio.previewUrlWhitelist values if specified 2. schema and authority of studio.previewUrlPrefix if specified 3. 'self' To allow YouTube videos inside the external preview, add the Youtube URL: studio.security.csp.frameSrc=http://localhost:40980,*.coremedia.vm:40980,*.coremedia.vm,*.coremedia.com,*.coremedia.com:8000,*.coremedia.vm:8000,'self',www.youtube.com
<code>studio.security.csp.img-src</code>	
Type	List<String>
Description	List of values for the 'img-src' policy directive. Defaults to 'self'.
<code>studio.security.csp.manifest-src</code>	
Type	List<String>
Description	List of values for the 'manifest-src' policy directive. Defaults to 'self'.
<code>studio.security.csp.media-src</code>	
Type	List<String>
Description	List of values for the 'media-src' policy directive. Defaults to 'self'.
<code>studio.security.csp.object-src</code>	

Type	List<String>
Description	List of values for the 'object-src' policy directive. Defaults to 'self'.
<code>studio.security.csp.report-uri</code>	
Type	List<String>
Description	List of values for the 'report-uri' policy directive. If no custom list is provided the directive is not included.
<code>studio.security.csp.script-src</code>	
Type	List<String>
Description	List of values for the 'script-src' policy directive. Defaults to 'self','unsafe-eval'.
<code>studio.security.csp.style-src</code>	
Type	List<String>
Description	List of values for the 'style-src' policy directive. Defaults to 'self','unsafe-inline'.

Table 3.23. Content Security Policy Related Studio Properties

3.4.8 Content Hub Configuration

The following table lists the configuration properties for the *CoreMedia Content Hub*.

<code>contenthub.studio.global-configuration-path</code>	
Type	String
Default	/Settings/Options/Settings/Content Hub
Description	Defines the global configuration path. The Content Hub will look up 'Connection' settings document in this folders. The default value is /Settings/Options/Set-tings/Content Hub.

contenthub.studio.site-configuration-path	
Type	String
Default	/Options/Settings/Content Hub
Description	Defines the site specific configuration path. If a Content Hub connection should only be available sites, the 'Connection' settings document can be but in this subfolder of a site. The default value is /Options/Settings.

Table 3.24. Content Hub Properties

3.4.9 Feedback Hub Configuration

The following table lists the configuration properties for the *CoreMedia Feedback Hub*.

feedbackhub.bindings.content-type	
Type	String
Default	CMSettings
Description	Defines the content type that is used to configure Feedback Hub bindings.
feedbackhub.bindings.global-path	
Type	String
Default	/Settings/Options/Settings/Feedback Hub
Description	Defines the repository folder path with contents that configure Feedback Hub bindings for all sites and contents without a site. Subfolders are ignored.
feedbackhub.bindings.property	
Type	String
Default	settings

Description	Defines the name of the struct property in the configured content type with the Feedback Hub binding configuration.
<code>feedbackhub.bindings.site-path</code>	
Type	String
Default	Options/Settings/Feedback Hub
Description	Defines the repository folder path below a site root folder with contents that configure Feedback Hub bindings for one site. Subfolders are ignored.

Table 3.25. Feedback Hub Properties

3.4.10 Editorial Comments Configuration

The following table lists the configuration properties for the *CoreMedia Editorial Comments* feature, which establishes a connection to the relational database.

The most important property is `editorial.comments.datasource.url` which sets the URL to connect to.

In case another *schema*, *username*, or *password* has to be configured use the properties `editorial.comments.db.schema`, `editorial.comments.db.username`, `editorial.comments.db.password`.

<code>editorial.comments.datasource.driver-class-name</code>	
Type	String
Description	Required to be set according to your database. See official spring documentation 'spring.datasource.driver-class-name' for detailed information
<code>editorial.comments.datasource.hikari.connection-timeout</code>	
Type	Integer
Default	20000ms
Description	Value must not be greater then the studio request timeout. See official spring documentation 'spring.datasource.hikari.connection-timeout' for more information

`editorial.comments.datasource.password`

Type String

Default `cm_editorial_comments`

Description See official spring documentation 'spring.datasource.password' for more information

`editorial.comments.datasource.url`

Type String

Description Required to be set according to your database. See official spring documentation 'spring.datasource.url' for detailed information

`editorial.comments.datasource.username`

Type String

Default `cm_editorial_comments`

Description See official spring documentation 'spring.datasource.username' for more information

`editorial.comments.db.password`

Type String

Default `cm_editorial_comments`

Description Use to set the password for editorial-comments data.

`editorial.comments.db.schema`

Type String

Default `cm_editorial_comments`

Description Use to set the schema for editorial-comments data.

CoreMedia Properties Overview | Editorial Comments Configuration

```
editorial.comments.db.username
```

Type	String
------	--------

Default	cm_editorial_comments
---------	-----------------------

Description	Use to set the username for editorial-comments data.
-------------	--

```
editorial.comments.jpa.database-platform
```

Type	String
------	--------

Description	See official spring documentation 'spring.jpa.properties.hibernate.database-platform' for detailed information
-------------	--

```
editorial.comments.jpa.properties.hibernate.jdbc.time_zone
```

Type	String
------	--------

Default	UTC
---------	-----

Description	Configuration of the default timezone. See official spring documentation of 'spring.jpa.properties.hibernate.jdbc.time_zone' for more information.
-------------	--

```
editorial.comments.notification-strategies.created-comment-on-content
```

Type	Boolean
------	---------

Default	true
---------	------

Description	This property enables the CreatedCommentOnContentCollectUsersToNotifyS-strategy which collects all users who wrote a comment at this content in the last 30 days. True if the strategy is enabled, false if the strategy is disabled.
-------------	---

```
editorial.comments.notification-strategies.edited-content-in-last-thirty-days
```

Type	Boolean
------	---------

Default	true
---------	------

Description This property enables the EditedContentInLastThirtyDaysCollectUsersToNotifyStrategy which collects all users who edited the given content in last 30 days. True if the strategy is enabled, false if the strategy is disabled.

`editorial.comments.notification-strategies.my-edited-contents`

Type Boolean

Default true

Description This property enables the MyEditedContentsCollectUsersToNotifyStrategy, which notifies every user that has the Content where an EditorialComment was created for in his "myEditedContent" List. true if the strategy is enabled, false if the strategy is disabled.

Table 3.26. Editorial Comments Properties

3.4.11 Notifications SQL Persistence Configuration

The following table lists the configuration properties for *CoreMedia Notifications* with SQL persistence.

The most important property is `notification.datasource.url` which sets the URL to connect to.

In case another *username*, or *password* has to be configured use the properties `notification.datasource.username`, `notification.datasource.password`.

`notification.datasource`

Type `org.springframework.boot.jdbc.autoconfigure.DataSourceProperties`

Description Configuration of the data source. See official spring documentation of 'spring.datasource' for more information

`notification.datasource.hikari`

Type `com.zaxxer.hikari.HikariDataSource`

Description	Configuration of the Hikari data source. See official spring documentation of 'spring.datasource.hikari' for more information
notification.datasource.hikari.connection-timeout	
Type	Integer
Default	20000
Description	Hikari connection timeout in ms. Value must not be greater than the studio request timeout. See official spring documentation of 'spring.datasource.hikari.connection-timeout' for more information.
notification.datasource.password	
Type	String
Default	cm_notifications
Description	DB password.
notification.datasource.url	
Type	String
Description	DB jdbc connection url. See official spring documentation of 'spring.datasource.url' for more information.
notification.datasource.username	
Type	String
Default	cm_notifications
Description	DB username.
notification.deletion.enabled	
Type	Boolean
Default	true

Description	Defines if task is enabled.
--------------------	-----------------------------

<code>notification.deletion.interval</code>

Type	Duration
-------------	----------

Default	1h
----------------	----

Description	The interval for task execution. Default unit: seconds.
--------------------	---

<code>notification.deletion.limit</code>
--

Type	Integer
-------------	---------

Default	1000
----------------	------

Description	The batch size limit for deletion.
--------------------	------------------------------------

<code>notification.deletion.ttl</code>
--

Type	Duration
-------------	----------

Default	180d
----------------	------

Description	The time to live after the task deletes DB entries. Default unit: days.
--------------------	---

<code>notification.invalidation.enabled</code>
--

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	Defines if task is enabled.
--------------------	-----------------------------

<code>notification.invalidation.error-delay</code>
--

Type	Duration
-------------	----------

Default	1m
----------------	----

Description	The sleep duration in case of an exception. The default unit is seconds.
--------------------	--

```
notification.invalidation.interval
```

Type Duration

Default 5s

Description The interval for task execution. The default unit is seconds.

```
notification.invalidation.limit
```

Type Integer

Default 200

Description The limit of event entities to retrieve.

```
notification.invalidation.max-retries-for-missing-ids
```

Type Integer

Default 2

Description Event fetching is based on an auto-incrementing id column and only works correctly when the id increases for each new notification. As this is not guaranteed by the database, the event task will keep trying to obtain notifications for missing ids for this number of task executions.

```
notification.jpa
```

Type org.springframework.boot.jpa.autoconfigure.JpaProperties

Description Configuration of the JPA. See official spring documentation of 'spring.jpa' for more information.

```
notification.jpa.properties.hibernate.jdbc.time_zone
```

Type String

Default UTC

Description	Configuration of the default timezone. See official spring documentation of 'spring.jpa.properties.hibernate.jdbc.time_zone' for more information.
<code>notification.persistence</code>	
Type	String
Default	sql
Description	Persistence mode for notifications, possible values are 'sql' (default) and 'elastic-core' (deprecated). Using 'elastic-core' is deprecated since 2506.0. Please use 'sql' instead.

Table 3.27. Notification SQL Properties

3.4.12 Projects/To-Dos SQL Persistence Configuration

The following table lists the configuration properties for *CoreMedia Projects/To-Dos* with SQL persistence.

The most important property is `project.datasource.url` which sets the URL to connect to.

In case another *username*, or *password* has to be configured use the properties `project.datasource.username`, `project.datasource.password`.

<code>project.datasource</code>	
Type	<code>org.springframework.boot.jdbc.autoconfigure.DataSourceProperties</code>
Description	Configuration of the data source. See official spring documentation of 'spring.datasource' for more information
<code>project.datasource.hikari</code>	
Type	<code>com.zaxxer.hikari.HikariDataSource</code>
Description	Configuration of the Hikari data source. See official spring documentation of 'spring.datasource.hikari' for more information

```
project.datasource.hikari.connection-timeout
```

Type	Integer
------	---------

Default	20000
---------	-------

Description	Hikari connection timeout in ms. Value must not be greater than the studio request timeout. See official spring documentation of 'spring.datasource.hikari.connection-timeout' for more information.
-------------	--

```
project.datasource.password
```

Type	String
------	--------

Default	cm_projects
---------	-------------

Description	DB password.
-------------	--------------

```
project.datasource.url
```

Type	String
------	--------

Description	DB jdbc connection url. See official spring documentation of 'spring.datasource.url' for more information.
-------------	--

```
project.datasource.username
```

Type	String
------	--------

Default	cm_projects
---------	-------------

Description	DB username.
-------------	--------------

```
project.entity.deletion.enabled
```

Type	Boolean
------	---------

Default	true
---------	------

Description	Defines if task is enabled.
-------------	-----------------------------

```
project.entity.deletion.interval
```

Type Duration

Default 1h

Description The interval for task execution. Default unit: seconds.

```
project.entity.deletion.limit
```

Type Integer

Default 1000

Description The batch size limit for deletion.

```
project.entity.deletion.ttl
```

Type Duration

Default 1d

Description The time to live after the task deletes project and todo DB entries, that are marked for deletion. Default unit: days.

```
project.event.deletion.enabled
```

Type Boolean

Default true

Description Defines if task is enabled.

```
project.event.deletion.interval
```

Type Duration

Default 1h

Description The interval for task execution. Default unit: seconds.

project.event.deletion.limit	
Type	Integer
Default	1000
Description	The batch size limit for deletion.
project.event.deletion.ttl	
Type	Duration
Default	1d
Description	The time to live after the task deletes DB entries. Default unit: days.
project.invalidation.enabled	
Type	Boolean
Default	true
Description	Defines if event processing is enabled.
project.invalidation.error-delay	
Type	Duration
Default	1m
Description	The sleep duration in case of an exception. The default unit is seconds.
project.invalidation.interval	
Type	Duration
Default	1s
Description	The interval for task execution. The default unit is seconds.
project.invalidation.limit	

Type	Integer
Default	200
Description	The limit of event entities to retrieve.
<code>project.invalidation.max-retries-for-missing-ids</code>	
Type	Integer
Default	5
Description	Event fetching is based on an auto-incrementing id column and only works correctly when the id increases for each new project event. As this is not guaranteed by the database, the event task will keep trying to obtain project events for missing ids for this number of task executions.
<code>project.invalidation.remote-enabled</code>	
Type	Boolean
Default	true
Description	Defines if remote event processing is enabled. Required for a setup with multiple studio-server instances.
<code>project.jpa</code>	
Type	<code>org.springframework.boot.jpa.autoconfigure.JpaProperties</code>
Description	Configuration of the JPA. See official spring documentation of 'spring.jpa' for more information.
<code>project.jpa.properties.hibernate.jdbc.time_zone</code>	
Type	String
Default	UTC
Description	Configuration of the default timezone. See official spring documentation of 'spring.jpa.properties.hibernate.jdbc.time_zone' for more information.

project.migration.elastic-core-persistence	
Type	String
Default	mongodb
Description	The elastic-core persistence for the source repositories; either mongodb or memory. Alternative to setting elastic.core.persistence.
project.migration.enabled	
Type	Boolean
Default	false
Description	Whether to enable migration from Elastic Core to SQL persistence for Projects and Todos.
project.migration.log-report	
Type	Boolean
Default	false
Description	Whether to log a report after migration (the report is also available via GET on actuator/projectmigration).
project.migration.mute-notifications	
Type	Boolean
Default	true
Description	Whether to mute all notifications for Projects and Todos during a running migration.
project.migration.on-startup	
Type	Boolean
Default	true

Description	Whether to start migration on application startup (migration can also be triggered via POST to <code>actuator/projectmigration</code>).
--------------------	--

```
project.persistence
```

Type	String
-------------	--------

Default	sql
----------------	-----

Description	Persistence mode for projects, possible values are 'sql' (default) and 'elastic-core' (deprecated). Using 'elastic-core' is deprecated since 2506.0. Please use 'sql' instead.
--------------------	--

Table 3.28. Project SQL Properties

3.4.13 Projects/To-Dos Migration Configuration

The following table lists the configuration properties for the Migration of *CoreMedia Projects/To-Dos* from *Elastic Core* (MongoDB or In-Memory) to SQL persistence.

The most important property is `project.migration.enabled` which enables and by default starts the migration on application start-up.

```
project.migration.elastic-core-persistence
```

Type	String
-------------	--------

Default	mongodb
----------------	---------

Description	The elastic-core persistence for the source repositories; either mongodb or memory. Alternative to setting <code>elastic.core.persistence</code> .
--------------------	--

```
project.migration.enabled
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description Whether to enable migration from Elastic Core to SQL persistence for Projects and Todos.

`project.migration.log-report`

Type Boolean

Default false

Description Whether to log a report after migration (the report is also available via GET on actuator/projectmigration).

`project.migration.mute-notifications`

Type Boolean

Default true

Description Whether to mute all notifications for Projects and Todos during a running migration.

`project.migration.on-startup`

Type Boolean

Default true

Description Whether to start migration on application startup (migration can also be triggered via POST to actuator/projectmigration).

Table 3.29. Project Migration Properties

3.4.14 Commerce Related Configuration

The following table lists the commerce related *CoreMedia Studio* properties.

`studio.commerce.preload-child-categories`

Type java.lang.String

Default ALL

Description	The default behavior of the Studio library catalog tree is to load the next level of categories no matter if they are displayed. This is done to determine if a child category is virtual or not. All occurrences of a category that are not in the primary location in the catalog tree are considered as virtual. Set this property to <i>ALL_EXCEPT_TOP_LEVEL</i> if top level categories should be excluded from pre-loading. It can be useful if there is a huge number of top level categories and if you are sure they are not virtual. In commerce systems where no physical root category exists it must be this way. Moreover, if you are sure there is no virtual category at all you can use the value <i>NONE</i> . If a child category is not pre-loaded, its state is assumed to be non-virtual.
-------------	---

Table 3.30. Commerce Related Properties

3.5 User Changes Application Properties

The following list contains configuration properties related to the *User Changes Application*.

Find more related properties here:

- [Section 3.8, “My Edited Content and Workflow Lists Properties” \[162\]](#)
- [Notifications SQL Persistence Configuration for CoreMedia Studio \[130\]](#)

userchanges.excluded-paths	
Type	List<String>
Default	[/Home/*/EditorPreferences, /Home/*/KioSavedPrompts, /Home/*/My Preferences, /Home/*/My Dictionary, /System/Public Dictionary, /Home/*/EditorProfile]
Description	Paths to be excluded from tracking of user changes.
userchanges.excluded-user-names	
Type	List<String>
Default	[admin, publisher, workflow, importer, translation-workflow-robot]
Description	Users to be excluded from tracking of user changes.
userchanges.listener.enabled	
Type	Boolean
Default	false
Description	Set to 'true' to let the application register the user changes listener for content events.
userchanges.max-length	
Type	Integer

Default	1000
Description	The maximum number of contents in a user's "My Edited Content" list. The User Changes App will not add further contents to a list, if the number of contents has reached this maximum. It is recommended to configure a maximum that can still be handled by editors. If MongoDB persistence is used, it should be significantly lower than the maximum number of contents that can be stored in one MongoDB document, which depends on the length of stored document IDs but can be estimated to something around 600.000 contents.
<code>userchanges.notifications.workflow.enabled</code>	
Type	Boolean
Default	false
Description	Set to 'true' to let the application generate notifications for workflow events.

Table 3.31. User Changes App Properties

3.6 Workflow Server Properties

The *Workflow Server* sets up a connection to the *Content Management Server*. The protocol used for this connection can be either CORBA (default) or gRPC. See properties `repository.use-grpc`, `repository.http-base-uri` and `spring.grpc.client.channel.cap.target` at [Section 3.12.1, “Unified API Spring Boot Client Properties” \[233\]](#) for gRPC connection settings.

Find more related properties here:

- [Section 3.8, “My Edited Content and Workflow Lists Properties” \[162\]](#)

3.6.1 General

<code>spring.grpc.server.port</code>	
Description	gRPC server port.
<code>spring.grpc.server.port</code>	
Description	gRPC server port.
<code>workflow.archive.endpoint-url</code>	
Type	String
Default	http://localhost:41080/studio
Description	Endpoint of the process archive in the in-memory setup.
Deprecation	This property has been deprecated since 2512.0.0 and will be removed in a future version. Reason: In-Memory persistence is deprecated and will be removed in a future release.
<code>workflow.archive.password</code>	
Type	String

Description Password to use for process archiving at the endpoint in the in-memory setup.

Deprecation This property has been deprecated since 2512.0.0 and will be removed in a future version.

Reason:

In-Memory persistence is deprecated and will be removed in a future release.

`workflow.archive.username`

Type String

Description Username to use for process archiving at the endpoint in the in-memory setup.

Deprecation This property has been deprecated since 2512.0.0 and will be removed in a future version.

Reason:

In-Memory persistence is deprecated and will be removed in a future release.

`workflow.blob-cache-size`

Type `org.springframework.util.unit.DataSize`

Default 32MB

Description The size of the main disk cache of the UAPI

`workflow.blob-streaming-size-threshold`

Type `org.springframework.util.unit.DataSize`

Default 128KB

Description The minimum size of streamed blobs

`workflow.blob-streaming-threads`

Type Integer

Default 2

Description The number of streaming threads

`workflow.domain`

Type String

Description The server's domain name

`workflow.heap-cache-size`

Type `org.springframework.util.unit.DataSize`

Default 20MB

Description The size of the main memory cache of the UAPI

`workflow.ior-url`

Type String

Description The *Content Server's* IOR URL

`workflow.localization.auto-merge.legacy-list-merge`

Type Boolean

Default false

Description Configures whether an old but backwards-compatible merge algorithm for merging list properties should be used. This also applies to struct lists. If true, merge conflicts are reported if there are any changes in overlapping regions of list values. This also includes changes in nested property values, for example if there are changes in translatable properties and those properties aren't excluded with `workflow.localization.auto-merge.translatable=false`, or if the `AutoMerge-StructListMapKeyFactory` is configured to merge a struct list as map, and some non-key properties have changed. An example for the latter would be a conflict caused by removing a link in a master content and a changed annotation property for that link in the derived content. If false, an improved merge algorithm is used for list values, which tries to avoid merge conflicts in more cases.

`workflow.localization.auto-merge.translatable`

Type	Boolean
Default	true
Description	Configures if translatable properties are automatically merged in translation workflows. If true, changes of translatable properties in a master content are also merged to derived content items by translation workflows, if there are no conflicting changes and the properties haven't been excluded otherwise. Opposed to non-translatable properties, merge conflicts for translatable properties are typically resolved silently in favor of the derived content and do not generate a warning for the user. There are some exceptions, especially when using the legacy list merge algorithm (see property workflow.localization.auto-merge.legacy-list-merge). Even if set to true, translatable properties can still be excluded from auto-merge, for example with extensions:auto-merge="false" in the document type definition or with a custom implementation of com.coremedia.translate.workflow.AutoMergePredicateFactory.
workflow.map-role	
Type	Map<String,String>
Description	Role mappings are used to replace group names in a process definition file by other group names without having to actually modify the textual definition
workflow.max-cached-blob-size	
Type	org.springframework.util.unit.DataSize
Description	The maximum size of cached blobs
workflow.navigatethrough	
Type	Boolean
Default	true
Description	Whether to enable navigate through cache of the UAPI. This cache is used for computing read rights on folders in case no explicit read rights are defined and read rights on subfolders force read rights on folders on the path to these folders. Because these computations are expensive and sometimes unnecessary, this flag allows you to ignore navigate-through rights in the Workflow Server.

```
workflow.password
```

Type String

Default workflow

Description Define a password for the connection with the *Content Management Server*.

```
workflow.pool.queue-size
```

Type Integer

Default 0

Description Pool queue limit.

```
workflow.pool.threads-max
```

Type Integer

Description Pool thread max limit.

```
workflow.pool.threads-min
```

Type Integer

Default 0

Description Pool thread min limit.

```
workflow.server.allow-live
```

Type Boolean

Default false

Description Whether to allow connecting to a live *Content Server*

```
workflow.server.archive.migration.enabled
```

Type Boolean

Default	false
Description	Whether to enable migration of the workflow archive from Elastic to SQL persistence.
<code>workflow.server.archive.migration.retry-delay</code>	
Type	Duration
Default	10s
Description	The delay for retrying the migration of a process after an Exception.
<code>workflow.server.archive.migration.skip-processes</code>	
Type	List<Integer>
Description	A list of numeric process ids to skip during migration (if for some reason a specific process cannot or should not be migrated).
<code>workflow.server.archive.persistence</code>	
Type	String
Default	sql
Description	Persistence option where archived workflow processes are stored. By default, SQL is used for storage. Using "elastic-core" is deprecated since 2506.0. Please use "sql" instead.
<code>workflow.server.archive.retry-exceptions</code>	
Type	List<String>
Default	[java.io.IOException, com.mongodb.MongoTimeoutException]
Description	List of exception classes that indicate temporary problems when archiving a workflow process. If one of these exceptions is thrown directly or as cause of another exception in <code>com.coremedia.workflow.common.actions.ArchiveProcessFinalAction</code> , then the Workflow Server will retry archiving with some delay as configured in properties <code>workflow.server.retry-final-actions.*</code> .

```
workflow.server.enable-workflow-converter
```

Type	Boolean
Default	false
Description	Whether to convert all enabled workflow definitions automatically on Workflow Server start. This is similar to manually executing "cm workflowconverter --convert", with the difference that only enabled workflow definitions are converted to reduce impact on startup time.

```
workflow.server.final-actions-max
```

Type	Integer
Default	10
Description	The maximum number of threads executing FinalActions of finished processes. This setting can be changed at runtime with the JMX attribute 'FinalActionsMax'.

```
workflow.server.remote-action-handler
```

Type	String
Default	com.coremedia.cotopaxi.workflow.BuiltInRemoteActionHandler
Description	The remote action handler.

```
workflow.server.retry-final-actions.delay.initial
```

Type	Duration
Default	1s
Description	The initial time interval to wait before retrying the execution of a workflow final action after it failed for the first time. The value must be ≥ 1 millisecond. The initial interval is multiplied with the configured multiplier for each retry attempt (exponential back-off) until the configured maximum has been reached. This setting is ignored if the exception provides a different value in method <code>RetryableActionException#getDelayBeforeRetry</code>.

```
workflow.server.retry-final-actions.delay.max
```

Type	Duration
Default	1h
Description	The maximum time to wait before retrying a failed execution of a workflow final action. The duration must be larger or equal to the specified initial time to wait. This setting is ignored if the exception provides a different value in method <code>RetryableActionException#getDelayBeforeRetry</code> .

```
workflow.server.retry-final-actions.delay.multiplier
```

Type	BigDecimal
Default	2
Description	Multiplier for exponential back-off. The value must be ≥ 1.0 . This setting is ignored if the exception provides a different value in method <code>RetryableActionException#getDelayBeforeRetry</code> .

```
workflow.server.retry-final-actions.delay.random-factor
```

Type	BigDecimal
Default	0.1
Description	Factor to randomize the effective delay. The value must be ≥ 0.0 and < 1.0 . Use 0.0 for no randomization. For example, a value of 0.1 creates effective delays in the range of plus/minus 10 percent of the computed delay. This also means that the actual delay can be up to 10 percent above the configured max value. This setting is ignored if the exception provides a different value in method <code>RetryableActionException#getDelayBeforeRetry</code> .

```
workflow.server.retry-final-actions.enabled
```

Type	Boolean
Default	true
Description	Whether the execution of workflow final actions will be retried if a <code>com.core-media.cap.workflow.plugin.RetryableActionException</code> was thrown. If set to false,

retries after these exceptions are disabled completely and this cannot be overridden in the implementation of the `RetryableActionException`. This setting can be changed at runtime with the JMX attribute 'FinalActionsRetryEnabled'.

<code>workflow.server.retry-final-actions.max-retries</code>	
Type	Integer
Default	-1
Description	The maximum number of retry attempts, after a final action has failed with a <code>com.coremedia.cap.workflow.plugin.RetryableActionException</code> for a finished process instance. When the maximum number of retry attempts is reached, then the action will be skipped and the process instance will be deleted without successful execution of the failed final action. A value of -1 means unlimited retries. A value of 0 disables retries. This setting is ignored if the exception provides a different value in method <code>RetryableActionException#getMaxRetries</code> .

<code>workflow.server.retry-register-pending-process.delay.initial</code>	
Type	Duration
Default	1s
Description	The initial time interval to wait before retrying the execution of the <code>RegisterPendingProcess</code> action after it failed for the first time. The value must be ≥ 1 milli-second. The initial interval is multiplied with the configured multiplier for each retry attempt (exponential back-off) until the configured maximum has been reached.

<code>workflow.server.retry-register-pending-process.delay.max</code>	
Type	Duration
Default	1m
Description	The maximum time to wait before retrying the execution of the <code>RegisterPendingProcess</code> action. The duration must be larger or equal to the specified initial time to wait.

<code>workflow.server.retry-register-pending-process.delay.multiplier</code>	
--	--

Type	BigDecimal
Default	2
Description	Multiplier for exponential back-off. The value must be ≥ 1.0 .
<code>workflow.server.retry-register-pending-process.delay.random-factor</code>	
Type	BigDecimal
Default	0.1
Description	Factor to randomize the effective delay. The value must be ≥ 0.0 and < 1.0 . Use 0.0 for no randomization. For example, a value of 0.1 creates effective delays in the range of plus/minus 10 percent of the computed delay. This also means that the actual delay can be up to 10 percent above the configured max value.
<code>workflow.server.retry-register-pending-process.enabled</code>	
Type	Boolean
Default	true
Description	Whether the execution of the RegisterPendingProcess action will be retried if an exception occurred.
<code>workflow.server.session.reaper-timeout</code>	
Type	Duration
Default	1m
Description	The interval of the session reaper checks. (Default unit is seconds)
<code>workflow.server.session.timeout</code>	
Type	Duration
Default	5m
Description	The timeout until an inactive session is closed. (Default unit is seconds)

```
workflow.server.standalone
```

Type	String
------	--------

Description	Possible values: 'true' 'false' 'embedded'
-------------	--

```
workflow.server.sweeper.concurrent
```

Type	Boolean
------	---------

Default	false
---------	-------

Description	The concurrent sweep mode.
-------------	----------------------------

```
workflow.server.termination-timeout-seconds
```

Type	Duration
------	----------

Default	30s
---------	-----

Description	Timeout for waiting on running threads on shutdown
-------------	--

```
workflow.server.tx.idletimeout
```

Type	Duration
------	----------

Default	30s
---------	-----

Description	Transaction handler idle timeout. (Default unit is seconds)
-------------	---

```
workflow.server.tx.max
```

Type	Integer
------	---------

Default	20
---------	----

Description	Limits the maximum number of database connections.
-------------	--

```
workflow.server.tx.retry-on-connection-loss
```

Type	Boolean
------	---------

Default	false
Description	If true, will retry failed transactions due to a database connection failure. Before retrying, the DB connection will be renewed. This can also be useful in clustered databases to cope with failover switch.
<code>workflow.server.tx.retry-on-connection-loss-delay</code>	
Type	Duration
Default	15s
Description	Duration to wait before a new attempt is made after a transaction failed due to database connection failure. Must be larger than zero.
<code>workflow.server.tx.retry-on-read-only-db</code>	
Type	Boolean
Default	false
Description	If true, will retry failed transactions due to a read-only database. Before retrying, the DB connection will be renewed. This can be useful when running on a clustered database (e.g., Amazon Aurora) that dynamically assigns writer and reader instances to cope with failover switch. This is currently only supported on MySQL and MariaDB.
<code>workflow.server.tx.retry-on-read-only-db-delay</code>	
Type	Duration
Default	5s
Description	Duration to wait before a retry attempt is made after a transaction failed due to a read-only database. Must be larger than zero.
<code>workflow.server.uapi.session-pool.max-idle-count</code>	
Type	Integer
Default	1000

Description The maximum number of idle sessions to keep in the *Content Server* session pool.

```
workflow.server.uapi.session-pool.max-idle-time-m-s
```

Type Duration

Default 1m

Description The duration after which an idle session of the *Content Server* session pool should be closed.

```
workflow.server.uapi.session-pool.max-open-count
```

Type Integer

Default 1000

Description The maximum number of sessions to hold at any time (busy + idle) in the *Content Server* session pool.

```
workflow.server.use-uploaded-jar
```

Type Boolean

Default false

Description Whether it is possible to upload workflow process definitions with JAR files containing executable code, and previously uploaded workflow definitions can still use classes and resources from uploaded JARs.

It is recommended to keep this disabled, because uploaded executable code can be a security risk. With different versions of custom code in the database, it can become hard to track which code is actually used in production, and to detect and fix security issues in all variants of uploaded custom code. Furthermore, in case attackers have gained write access to the database, they could use this feature to execute code in the Workflow Server.

Uploading JAR files with process definitions is a legacy feature of the Workflow Server that will be removed in a future release. Required classes should be put directly on the classpath of the Workflow Server.

Deprecation	This property has been deprecated since 2512.0.0 and will be removed in a future version. Reason: The feature to upload and use JAR files is a legacy feature that will be removed in a future release.
<code>workflow.serviceindicator</code>	
Type	String
Description	The server's service indicator
<code>workflow.user</code>	
Type	String
Default	workflow
Description	Define a user for the connection with the <i>Content Management Server</i> .
<code>workflow.working-dir</code>	
Type	String
Default	var/tmp
Description	Working directory of the workflow-server used to handle custom classes up-loaded to the server. Path can be absolute or relative to user.dir system-property.

Table 3.32. Workflow Server Properties

3.6.2 Shared Multi-Site Configuration

The following list contains configuration properties for the multi-site feature, that are typically shared between different applications (or they must be kept in sync).

<code>multisite.transfer.relations.enabled</code>	
Type	Boolean

Default	true
Description	Whether referrer relation augmentation is enabled globally. When set to false, no relation resolvers are created, regardless of individual relation settings.
<code>multisite.transfer.relations.referrer</code>	
Type	<code>Map<String, com.coremedia.cap.multisite.transfer.RelationsConfigurationProperties\$ReferrerSlot></code>
Description	Named referrer relations. Each map entry defines how referencing content items are discovered for a given content type. Default relations: • taxonomy — content-type: CMTaxonomy, source-property: children • taxonomy-root — content-type: CMTaxonomy, source-property: settings, source-content-type: CMSSettings • navigation — content-type: CMNavigation, source-property: children • navigation-root — content-type: CMNavigation, source-property: root, source-content-type: CMSite Customization can add, override, or disable individual entries. To disable a specific relation, set its enabled property to false.
<code>translatable.predicate.property-path</code>	
Type	<code>Map<String, Boolean></code>
Description	Marks properties or nested (Struct) properties as translatable. Typically applied to nested properties in Structs, which cannot be individually marked as translatable in the content type model. The map entry's key represents the property path starting from the content type that declares the property. A map entry's value of true marks the property selected by the property path as translatable. Example Configuration: <pre>translatable.predicate.property-path[CMLinkable.localSettings.callToActionCustomText]=true</pre> For Struct properties containing dots in their names, you may use bracket syntax instead:

```
translatable.predicate.property-path[CMLinkable.localSettings['with.dots.in.name']]=true
```

Note, that a value of false cannot veto a translatable decision by another predicate. The parsed configuration will just ignore those property paths marked with false from this configuration.

```
translate.item.include-unchanged-properties
```

Type	Boolean
Default	true
Description	Whether properties that have not changed since last translation are also included in items for translation.

```
translate.item.transform.failure.mode
```

Type	String
Default	fail
Description	Controls behavior when no matching transformers are available to eventually represent a property in XLIFF. Available values: fail, warn, and none.

```
translate.xliff.export.excludeContentName
```

Type	Boolean
Default	true
Description	Only if set to false, content names are included in the XLIFF export.

Table 3.33. Shared Multi-Site Properties

3.7 Commerce Hub Properties

This section lists the Commerce Adapter client properties. For the configuration of the Commerce Adapter properties themselves, see the vendor specific Commerce Adapter manuals.

commerce.hub.caching.enabled	
Type	Boolean
Default	true
Description	De-/activates the commerce cache. If activated the caching is adjusted by the related property keys in com.coremedia.blueprint.base.livecontext.commerce-cache-defaults.properties.
commerce.hub.data.custom-entity-params	
Type	Map<String,String>
Description	The custom entity parameters.

Table 3.34. Commerce Hub Properties

3.8 My Edited Content and Workflow Lists Properties

The features *My Edited Content* and *Workflow Lists* use CapLists for maintaining their state. This section lists the properties required to configure the data source and event handling for CapLists.

Applications where these properties are available:

- *Studio Server*,
- *User Changes Application*
- *Workflow Server*

3.8.1 General Properties

caplist.persistence	
Type	String
Description	Persistence mode for CapList, set to 'sql' for SQL persistence. Using 'elastic-core' is deprecated since 2506.0. Please use 'sql' instead.

Table 3.35. General Configuration Properties for CapLists

3.8.2 Data Source Properties

The properties in this section are used to configure the data source. They are only an excerpt of the properties available for the data source. For more details consult the documentation of `spring.datasource` properties ([SQL Databases :: Spring Boot](#), [Spring Boot Appendix: Data Properties](#)), which we wrap into application specific properties `caplist.datasource`. `spring.datasource` is backed by class [org.springframework.boot.autoconfigure.jdbc.DataSourceProperties](#).

As connection pool, we use HikariCP. If there is a need to adapt these, you may also refer to the official Spring Boot documentation for properties `spring.datasource.hikari` (again, here referenced as `caplist.datasource.hikari`).

The most important property is `caplist.datasource.url` which sets the URL to connect to.

In case another *username*, *password* has to be configured use the properties `caplist.datasource.username`, `caplist.datasource.password`.

caplist.datasource.password	
Type	String
Default	cm_caplist
Description	Login password of the database, possibly encrypted.
caplist.datasource.url	
Type	String
Description	JDBC URL of the database connection for the cm_caplist schema.
caplist.datasource.username	
Type	String
Default	cm_caplist
Description	Login username of the database.

Table 3.36. Data Source Configuration Properties for CapLists

3.8.3 Event Handling Properties

caplist.event.deletion.batch-size	
Type	Integer
Default	1000
Description	The maximum number of old events to delete with a single SQL DELETE statement. If more old events need to be deleted, multiple statements will be used.

```
caplist.event.deletion.cron
```

Type	String
Default	@midnight
Description	The cron specification for when old events get deleted. For the supported format, see the Javadoc of Spring method <code>org.springframework.scheduling.support.CronExpression#parse(java.lang.String)</code> .

```
caplist.event.deletion.enabled
```

Type	Boolean
Default	false
Description	Whether this application deletes old events from the database, when they are no longer needed according to the configured time-to-live (TTL). For cleaning up old events, it is sufficient to enable this option for only one application like the 'User Changes App'.

```
caplist.event.deletion.ttl
```

Type	Duration
Default	7d
Description	The time to live for events, after which they are considered old, no longer needed, and become subject to deletion.

```
caplist.event.fetch-interval
```

Type	Duration
Default	1s
Description	The time to wait between fetching events from the database.

```
caplist.event.max-delay
```

Type	Duration
-------------	----------

Default	2s
Description	The maximum delay between fetching an event from the database, and publishing it to listeners. An event is delayed, if another event with smaller ID could not be fetched from the database yet. This may happen if there's a gap in the auto-incremented IDs of already fetched events. If the missing event is received within this maximum delay, then events can still be published to listeners in the order of their event IDs.
<code>caplist.event.max-lag-missing-id</code>	
Type	Integer
Default	100
Description	The maximum number of fetched events, after which an event with smaller missing ID will not be fetched anymore, even if the time in 'caplist.event.max-wait-missing-id' has not yet elapsed. This limit is used to avoid repeated fetching of large amounts of events.
<code>caplist.event.max-wait-missing-id</code>	
Type	Duration
Default	30s
Description	The maximum time to wait for a missing event to still appear and to retry fetching it, after events with higher event ID have already been published to listeners. Such a gap in the auto-incremented IDs of fetched events can be temporary if stored events become visible in different order, or it can be permanent if IDs are skipped by the database. A gap is considered permanent after this configured time, or if the number of fetched events with higher ID has exceeded the value configured in 'caplist.event.max-lag-missing-id'.

Table 3.37. Configuration Properties for Event Handling of CapLists

3.8.4 Data Migration Properties

The properties in this section are used to configure the migration of CapList data from *Elastic Core* (MongoDB or In-Memory replacement) to SQL persistence.

As prerequisite, SQL persistence needs to be configured with `caplist.persistence=sql` and [Section 3.8.2, “Data Source Properties” \[162\]](#).

The most important property is `caplist.migration.enabled` which has to be set to true for the *Workflow Server*, *User Changes Application*, and *Studio Server* to enable migration.

Data migration is performed by the *Studio Server*. If present, additional *Studio Server* instances must be configured to not also start the migration by setting their property `caplist.migration.on-startup=false`.

caplist.migration.enabled	
Type	Boolean
Default	false
Description	Set to true to enable migration of CapList data from Elastic Core (MongoDB or In-Memory replacement) to SQL persistence. This property needs to be set to the same value for Workflow Server, User Changes App, and Studio Servers. This property is only used if <code>caplist.persistence=sql</code> .
caplist.migration.force	
Type	Boolean
Default	false
Description	Set to true to force a restart of CapList data migration from Elastic Core (MongoDB or In-Memory replacement) to SQL persistence, after it has already completed. Note, that this will not remove already migrated data from SQL persistence. This property is only used if <code>caplist.persistence=sql</code> , <code>caplist.migration.enabled=true</code> , and <code>caplist.migration.on-startup=true</code> .
caplist.migration.on-startup	
Type	Boolean
Description	Whether this application performs the migration of CapList data from Elastic Core (MongoDB or In-Memory replacement) to SQL persistence. This property defaults to true for the Studio Server and to false for all other applications. You only have to change this, if you start multiple Studio Servers. In that case, all but one Studio Server must be explicitly configured with value false. This property is only used if <code>caplist.persistence=sql</code> and <code>caplist.migration.enabled=true</code> .

caplist.migration.skip	
Type	Boolean
Default	false
Description	Set to true to skip CapList data migration from Elastic Core (MongoDB or In-Memory replacement) to SQL persistence, as if it had completed. This property is only used if caplist.persistence=sql, caplist.migration.enabled=true, and caplist.migration.on-startup=true.

Table 3.38. Configuration Properties for Migration of CapLists from Elastic Core to SQL Persistence

3.9 Elastic Social Properties

3.9.1 General Elastic Social Properties

elastic.core.event-collection-ttl-sec	
Type	Duration
Default	5m
Description	Time-To-Live index (TTL) for the collection event. Set it to 0 to disable automatic deletion.
elastic.social.comments.auto-reject-interval-ms	
Type	Duration
Description	Interval in milliseconds in which comments will be automatically rejected when the number of complaints reaches the limit elastic.social.comments.auto-reject-limit.
elastic.social.comments.auto-reject-limit	
Type	Integer
Default	0
Description	Number of complaints after which a comment will automatically be rejected. Set to 0 to disable automatic rejection.
elastic.social.mails.on-profile-changes	
Type	Boolean
Default	true
Description	Flag which states if a mail should be send on profile changes

`elastic.social.mails.on-registration`

Type Boolean

Default true

Description Flag which states if a mail should be send on registration

`elastic.social.password-hash-algorithm`

Type String

Description The password hashing algorithm for elastic social user passwords. Allowed values are: sha1 and bcrypt:N. The algorithm sha1 is not recommended as it is vulnerable to brute-force attacks. The work factor N for bcrypt salt generation can be changed from the default value to an integer value between 4 and 30 (inclusive). This should be adapted to the available CPU resources.

`elastic.social.users.auto-block-interval-ms`

Type Duration

Description Interval in milliseconds in which users will be automatically blocked when the number of complaints reaches the limit `elastic.social.users.auto-block-limit`

`elastic.social.users.auto-block-limit`

Type Integer

Default 0

Description A number of complaints after which a user will automatically be blocked. Set to 0 to disable automatic rejection.

`elastic.social.users.pre-moderation-properties`

Type List<String>

Description Properties for the pre-moderation of users.

`elastic.social.users.token-expiration-time-ms`

Type	Duration
Description	Token expiration time in milliseconds, defines the lifetime of the user token used for user activation and password reset. After expiration, the token becomes invalid

Table 3.39. Elastic Social Properties

3.9.2 MongoDB Properties

mongodb.client-uri	
Type	String
Default	mongodb://localhost:27017
Description	The standard MongoDB connection string URI is used to configure your MongoDB connection, for example, it allows you to configure read preferences and write concerns. The format of a client URI is documented under the following link: http://docs.mongodb.org/manual/reference/connection-string/ .
mongodb.host	
Type	String
Default	localhost
Description	Set the hostname of the mongoDb server. Only relevant for development purposes.
mongodb.models.create-indexes	
Type	Boolean
Default	true
Description	Set to false to skip index creation on startup. This can be useful to speed up initial data import. Make sure to set this true before production use.
mongodb.prefix	

Type	String
Default	elastic
Description	Prefix for MongoDB database names. Change this when sharing a MongoDB installation with other Elastic Core applications.

Table 3.40. MongoDB Properties

3.9.3 Counter Properties

<code>counters.aggregation-interval-milliseconds.daily</code>	
Type	Duration
Description	Interval in milliseconds in which aggregated values for counter events of the last 24 hours will be updated. Defaults to 300000 (5 min).
<code>counters.aggregation-interval-milliseconds.for-all</code>	
Type	Duration
Description	Interval in milliseconds in which aggregated values for all historic counter events will be updated. Defaults to 86400000 (24 h).
<code>counters.aggregation-interval-milliseconds.monthly</code>	
Type	Duration
Description	Interval in milliseconds in which aggregated values for counter events of the last 30 days will be updated. Defaults to 86400000 (24 hours).
<code>counters.aggregation-interval-milliseconds.weekly</code>	
Type	Duration
Description	Interval in milliseconds in which aggregated values for counter events of the last 7 days will be updated. Defaults to 10800000 (3 hours).
<code>counters.aggregation-interval-milliseconds.yearly</code>	

Type	Duration
Description	Interval in milliseconds in which aggregated values for counter events of the last 365 days will be updated. Defaults to 86400000 (24 hours).

Table 3.41. Counters Properties

3.9.4 Task Queue Properties

taskqueues.execution-timeout-milli-seconds	
Type	Duration
Description	Specifies the task queue execution timeout in milliseconds. Defaults to 600000 (10 m).
taskqueues.execution-warning-timeout-milli-seconds	
Type	Duration
Description	Specifies duration in milliseconds after log warning for all long running tasks should be shown. Defaults to 300000 (5 m).
taskqueues.number-of-retries	
Type	Integer
Default	10
Description	Specifies the number of retries for a MongoDB task.
taskqueues.polling-interval-milli-seconds	
Type	Duration
Description	Specifies the duration of the task queue polling interval in milliseconds. Used by the scheduler in the MongoDBTaskQueueService. Defaults to 500 (0.5 s).
taskqueues.recovery-interval-milli-seconds	

Type	Duration
Description	Specifies the duration of the task queue recovery interval in milliseconds. Used by the scheduler in the MongoDBTaskQueueService. Defaults to 60000 (1 m).
<code>taskqueues.retry-interval-milli-seconds</code>	
Type	Duration
Description	Specifies task queue retry interval in milliseconds. Defaults to 300000 (5 m).
<code>taskqueues.worker-node</code>	
Type	Boolean
Default	false
Description	Set to false to disable execution of background tasks. This setting can be used to differentiate applications nodes into worker nodes which only execute background tasks and rendering nodes which serve requests.

Table 3.42. Task-Queues Properties

3.9.5 Elastic Social Link Building Properties

<code>elastic.social.links.context-path</code>	
Type	String
Default	/blueprint
Description	The context path to use for the social links.
<code>elastic.social.links.protocol</code>	
Type	String
Default	https
Description	The protocol to use for the social links.

elastic.social.links.servlet-path	
Type	String
Default	/servlet
Description	The servlet path to use for the social links.
elastic.social.links.use-original-request	
Type	Boolean
Default	false
Description	If set to true, the original request will be used to build the social links.

Table 3.43. Elastic Social Link Building Properties

3.9.6 Elastic Social Solr Properties

elastic.solr.cloud	
Type	Boolean
Description	Whether to connect to SolrCloud. Do not set this property to let Elastic Core use the standard CoreMedia setting. If true, connect to a SolrCloud cluster. SolrCloud connection details must be set either as ZooKeeper addresses (elastic.solr.zookeeper.addresses) or, if the former is unset or empty as HTTP URLs (elastic.solr.url). If false, connect to stand-alone Solr nodes via HTTP URLs (elastic.solr.url).
elastic.solr.connection-timeout	
Type	Integer
Description	Connection timeout in milliseconds, or 0 for no timeout, or a negative value to use SolrClient default. Do not set this property to let Elastic Core use the standard CoreMedia setting.
elastic.solr.index-config	

Type	String
Default	elastic
Description	Name of the Apache Solr config set for Elastic Core applications. This config set must exist in the Solr server, typically as subdirectory of "\$SOLR_HOME/config-sets".
<code>elastic.solr.index-data-directory</code>	
Type	String
Default	data
Description	Value for the "dataDir" parameter of the Solr CoreAdmin API / Collection API request to create a Solr index.
<code>elastic.solr.index-prefix</code>	
Type	String
Default	elastic
Description	Prefix for Apache Solr index names. Change this when sharing an Apache Solr installation with other Elastic Core applications.
<code>elastic.solr.lazy-index-creation</code>	
Type	Boolean
Default	false
Description	Whether all indices are created lazily. The default is false, and all indices are created eagerly except indices for tenants configured with elastic.solr.tenants-with-lazy-index-creation.
<code>elastic.solr.password</code>	
Type	String
Description	Password for HTTP basic authentication, used if a non-empty solr.username has been specified. The value may have been encrypted with the tool "cm encrypt-

passwordproperty". Do not set this property to let Elastic Core use the standard CoreMedia setting.

elastic.solr.socket-timeout	
Type	Integer
Description	Socket timeout in milliseconds, or 0 for no timeout, or a negative value to use SolrClient default. Do not set this property to let Elastic Core use the standard CoreMedia setting.
elastic.solr.tenants-with-lazy-index-creation	
Type	List<String>
Description	List of tenants for which indices are not eagerly created when the application starts, but only upon first access.
elastic.solr.url	
Type	String
Description	The Solr URL to connect to, or null to let Elastic Core use the standard CoreMedia setting. In a Solr leader/follower setup, this must be the URL of the Solr leader. For SolrCloud, it's recommended to set elastic.solr.zookeeper.addresses instead.
elastic.solr.username	
Type	String
Description	Username for HTTP basic authentication, or empty string for no authentication. Do not set this property to let Elastic Core use the standard CoreMedia setting.
elastic.solr.zookeeper.addresses	
Type	List<String>
Description	ZooKeeper addresses for connecting to SolrCloud. Only used if elastic.solr.cloud=true. Do not set this property to let Elastic Core use the standard CoreMedia setting.
elastic.solr.zookeeper.chroot	

Type	String
Description	Optional ZooKeeper chroot path for Solr. ZooKeeper chroot support makes it possible to isolate the SolrCloud tree in a ZooKeeper instance that is Only used if elastic.solr.cloud=true and elastic.solr.zookeeper.addresses is set to non-empty value. Do not set this property to let Elastic Core use the standard CoreMedia setting.
<code>elastic.solr.zookeeper.client-timeout</code>	
Type	Integer
Description	Client-timeout for ZooKeeper in milliseconds, or a negative value to use SolrClient default. Only used if elastic.solr.cloud=true and elastic.solr.zookeeper.addresses is set to non-empty value. Do not set this property to let Elastic Core use the standard CoreMedia setting.
<code>elastic.solr.zookeeper.connect-timeout</code>	
Type	Integer
Description	Connect-timeout for ZooKeeper in milliseconds, or a negative value to use SolrClient default. Only used if elastic.solr.cloud=true and elastic.solr.zookeeper.addresses is set to non-empty value. Do not set this property to let Elastic Core use the standard CoreMedia setting.

Table 3.44. Elastic Solr Properties

3.10 Content Feeder Properties

Properties for the Content Feeder

feeder.blob.enabled

Type Map<org.springframework.util.MimeType, Boolean>

Description The MIME-types of indexed blobs.

This property maps included MIME-types to true, and excluded MIME-types to false. Entries for more specific types like text/xml override mappings for less specific types like text/*.

Example Configuration:

```
feeder.blob.enabled[application/pdf]=true
```

If the map does not contain any enabled MIME-type, then the catch-all type */* is enabled implicitly.

See also property feeder.blob.max-size, which restricts blobs based on their size. Blobs are only included if allowed by both feeder.blob.enabled and feeder.blob.max-size.

feeder.blob.max-size

Type Map<org.springframework.util.MimeType, org.springframework.util.unit.DataSize>

Description The maximum size of indexed blobs per MIME-type. Larger blob values will be skipped.

This property maps MIME-types to their maximum size. Values must be non-negative. Entries for more specific types like text/xml override mappings for less specific types like text/*.

Example Configuration:

```
feeder.blob.max-size[application/pdf]=5MB
```

If the map does not contain a configuration for an enabled MIME-type, then blobs of that MIME-type are indexed regardless of their size.

See also property `feeder.blob.enabled`, which restricts blob values based on type. Blobs are only included if allowed by both `feeder.blob.enabled` and `feeder.blob.max-size`.

`feeder.content.background-feed`

Type `Map<String, com.coremedia.cms.feeder.content.BackgroundFeedProperties>`

Description Optional properties for named background feeds.

This property maps background feed identifiers to properties of that background feed.

The following background feed properties can be configured as component of the map value:

`priority` (0 if unset) – the priority compared to other background feeds; smaller possibly negative values represent higher priority.

`max-time-percentage` (100 if unset) – the maximum percentage of time the background feed is active; allowed values are 0 to 100 (0 effectively disables the feed)

Example: Change priority for background feed ISSUES

```
feeder.content.background-feed.ISSUES.priority=5000
```

Note, that predefined background feeds have different default priorities. Configured background feeds are logged with their effective priority on application start at level INFO. The log message contains the text "Configured background feeds".

`feeder.content.background-feed-delay`

Type `Duration`

Default `3s`

Description The minimum time after editorial changes were sent to the Search Engine and before background feeding takes place. This is used to prioritize feeding of editorial changes over background feeding, for example to process rights-rule changes or for periodic issue reindexing. It should not be necessary to change the default setting.

```
feeder.content.background-feed-max-time-window
```

Type	Duration
Default	10m
Description	The time window used with <code>feeder.content.background-feed[id].max-time-percentage</code> . Values less than one second are not allowed. Larger values for the time window lead to less but longer pauses.

```
feeder.content.index-deleted
```

Type	Boolean
Default	true
Description	Whether contents in the trash should be indexed. If you do not need to find contents in the trash and want to keep your index smaller, you can change this to false.

```
feeder.content.index-groups
```

Type	Boolean
Default	true
Description	Whether the IDs of groups with potential rights to read the content are indexed in the field "groups". This set of groups is then used to narrow a user's search to the contents where he might have read rights to. This is an optimization to get smaller search results for some queries and content structures and to get more accurate search suggestion counts. The client has to check for read rights anyway. For details, see also the description of the field "groups" in Solr schema.xml. If set to false, then you must also configure Studio Server and Content Server to not add a query condition for the indexed groups. To this end, set the Studio property <code>"studio.rest.search-service.use-groups-filter-query"</code> and the Content Server <code>"solr.use-groups-filter-query"</code> to "false".

```
feeder.content.index-name-in-textbody
```

Type	Boolean
Default	true

Description	Whether the content name should be indexed in field "textbody". It can make sense to disable this if lots of content names contain unique identifiers (from third-party systems, for example) to avoid problems with too many unique terms in field "textbody".
feeder.content.index-referrers	
Type	Boolean
Default	false
Description	Whether a content is reindexed after its referrers have changed.
feeder.content.index-translation-state	
Type	Boolean
Default	true
Description	Whether the translation state should be indexed. Computing the translation state can be an expensive operation when there are many derived sites and a lot of content is changed, for example when many derived sites that contain many content items are deleted.
feeder.content.initial-feeding-prefetch-size	
Type	Integer
Default	0
Description	The number of content items whose data is preloaded into the cache at a time when iterating over contents during initial feeding. A value of 0 (the default) disables prefetching, meaning content data is not proactively loaded into the cache. Configuring a positive value may improve initial feeding performance by reducing the number of round-trips to the Content Server, if the data fits into the cache. For prefetching to be effective, the Unified API heap cache must be large enough to hold the prefetched data. Use the property repository.heap-cache-size to size it accordingly. If the cache is too small, prefetched data may be evicted before it can be used, which can result in worse performance compared to not prefetching at all. It is therefore recommended to measure the effect of any

change to this setting in a representative environment before applying it to production.

feeder.content.issues.index	
Type	Boolean
Default	true
Description	Whether to index content issues.
feeder.content.issues.initial-feeding	
Type	Boolean
Default	false
Description	Whether content issues are already part of the initial feeding of an empty index. This property does not have any effect if feeder.content.issues.index is set to false. If true, initial feeding may take longer. If false, feeding of content issues starts after initial feeding has been completed.
feeder.content.issues.reindex-after	
Type	Duration
Default	1d
Description	The duration after which indexed issues are considered outdated and become subject to periodic reindexing. This property does not have any effect if feeder.content.issues.index or feeder.content.issues.reindex-periodically are set to false.
feeder.content.issues.reindex-periodically	
Type	Boolean
Default	true
Description	Whether content issues are reindexed periodically. Note that issue reindexing is performed with low priority, and will not block feeding of editorial changes. Issue reindexing will be paused as long as editorial changes need to be processed.

This property does not have any effect if `feeder.content.issues.index` is set to `false`.

`feeder.content.issues.reindex-time-max-percentage`

Type	Integer
Default	100
Description	The maximum percentage of time used to trigger issue reindexing. If set to a value below 100, periodic issue reindexing will try to pause and stay inactive for some time, so that it does not use more than the configured percentage of a time window, even if issues are older than configured in <code>feeder.content.issues.reindex-after</code> . This only applies to issue reindexing and the Content Feeder may still perform other tasks. The configured value must be in the range of 1 to 100. Note that issue reindexing is always performed with low priority, and will be paused as long as editorial changes need to be processed, even if this property is set to 100. This property does not have any effect if <code>feeder.content.issues.index</code> or <code>feeder.content.issues.reindex-periodically</code> are set to <code>false</code> .
Deprecation	This property has been deprecated since 2512.0.0 and will be removed in a future version. Use <code>feeder.content.background-feed.ISSUES.max-time-percentage</code> instead. Reason: This property and its companion property <code>'feeder.content.issues.reindex-time-window'</code> were replaced by the <code>'feeder.content.background-feed'</code> and <code>'feeder.content.background-feed-max-time-window'</code> properties.

`feeder.content.issues.reindex-time-window`

Type	Duration
Default	10m
Description	The time window used with <code>feeder.content.issues.reindex-time-max-percentage</code> . Larger values for the time window lead to less but longer pauses. This property does not have any effect if <code>feeder.content.issues.index</code> or <code>feeder.content.issues.reindex-periodically</code> are set to <code>false</code> , or if <code>feeder.content.issues.reindex-time-max-percentage</code> is 100.

Deprecation This property has been deprecated since 2512.0.0 and will be removed in a future version.

Use `feeder.content.background-feed-max-time-window` instead.

Reason:

This property and its companion property 'feeder.content.issues.reindex-time-max-percentage' were replaced by the 'feeder.content.background-feed-max-time-window' and 'feeder.content.background-feed' properties.

`feeder.content.management.password`

Type	String
Default	feeder
Description	The password to be used in the HTTP authentication of the administration page of the Content Feeder.

`feeder.content.management.user`

Type	String
Default	feeder
Description	The username to be used in the HTTP authentication of the administration page of the Content Feeder. This is not an account from the user management of the Content Server.

`feeder.content.partial-update-aspects`

Type	List<String>
Default	*
Description	Configures the aspects of index documents that can be updated with a partial update, provided that the connected Indexer supports partial updates (for example, <code>feeder.solr.partial-updates.enabled=true</code> for Solr). Multiple values are separated by comma. Use the special value "*" to use partial updates for all aspects, if possible. An empty value means that partial updates are not used. See the API documentation of <code>Feedable.isPartialUpdate</code> , <code>FeedableAspect</code> and <code>ContentFeedableAspect</code> in package <code>com.coremedia.cap.feeder</code> for more details.

```
feeder.content.populator-threads
```

Type	Integer
Default	1
Description	The number of threads used to call FeedablePopulator implementations. For effective thread usage, the queue capacity configured in <code>feeder.core.executor-queue-capacity</code> (default 100) must be significantly larger than this number of threads. If you increase the number of threads, consider increasing the mentioned queue size as well. When doing that, keep an eye on heap memory usage. If the "Semantic Search" feature is enabled, and this option is set to more than one thread, then you must also limit the number of threads that may call Apache Tika to extract text from binaries. For example, set <code>feeder.tika.threads=2</code> to limit semantic search feeding to two threads.

```
feeder.content.property-type.blob-max-size
```

Type	<code>org.springframework.util.unit.DataSize</code>
Description	Configure the maximum size of indexed blob properties. Larger blob values will be skipped. This configuration can be overridden for specific MIME-types by customizing Spring bean "feederContentBlobMaxSizePerMimeType".
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>feeder.blob.max-size[*/*]</code> instead. Reason: The property was replaced by 'feeder.blob.max-size', which allows configuration by MIME-type and also limits blobs that are added by custom FeedablePopulator implementations.

```
feeder.content.property-type.blob-mime-type.excludes
```

Type	<code>List<String></code>
Description	List of MIME-types of "Blob" properties excluded from indexing. You can exclude a more specific type (e.g. <code>text/xml</code>) while including the corresponding primary type (e.g. <code>text/*</code>).

Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>feeder.blob.enabled</code> instead. Reason: The property was replaced by 'feeder.blob.enable', which takes a map value instead of a comma-separated list and also limits blobs that are added by custom FeedablePopulator implementations.
feeder.content.property-type.blob-mime-type.includes	
Type	List<String>
Default	*/*
Description	List of MIME-types of indexed "Blob" properties. If you don't configure any MIME-types in the includes property, no blob properties will be indexed.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>feeder.blob.enabled</code> instead. Reason: The property was replaced by 'feeder.blob.enable', which takes a map value instead of a comma-separated list and also limits blobs that are added by custom FeedablePopulator implementations.
feeder.content.property-type.date	
Type	Boolean
Default	false
Description	Whether properties of type "Date" are indexed.
feeder.content.property-type.integer	
Type	Boolean
Default	false

Description Whether properties of type "Integer" are indexed.

```
feeder.content.property-type.link-list
```

Type Boolean

Default false

Description Whether properties of type "LinkList" are indexed.

```
feeder.content.property-type.string
```

Type Boolean

Default true

Description Whether properties of type "String" are indexed.

```
feeder.content.property-type.struct
```

Type Boolean

Default false

Description Whether properties of type "Struct" are indexed.

```
feeder.content.property-type.xml-grammars
```

Type List<String>

Default coremedia-richtext-1.0

Description The list of grammars of indexed "Markup" properties (as used in the document type definition as attribute "Name" of element "XmlGrammar").

```
feeder.content.retry-connect-to-index-delay
```

Type Duration

Default 10s

Description The time to wait between retries to connect to the search engine on startup.

```
feeder.content.semantic-search.allowed-path-prefixes
```

Type List<String>

Description The paths of content items to feed for semantic search. Paths start with a slash. A change to this property will not trigger automatic re-indexing of formerly included or excluded content items. See the Search Manual, chapter "Reindexing" for details on re-indexing.

```
feeder.content.semantic-search.enabled
```

Type Boolean

Default false

Description Enable the semantic search feature.

```
feeder.content.semantic-search.media-property-paths
```

Type Map<String, List<String>>

Description Configure com.coremedia.cap.content.Content property paths that point to picture blobs. Paths are given per com.coremedia.cap.content.ContentType.

Example Configuration:

```
feeder.content.semantic-search.mediaPropertyPaths.CMPicture=data
```

```
feeder.content.semantic-search.mime-types
```

Type Map<org.springframework.util.MimeType, Boolean>

Description The MIME-types that are supported for semantic search indexing of blobs.

This property maps included MIME-types to true, and excluded MIME-types to false. Entries for more specific types like image/jpeg override mappings for less specific types like image/*.

Example Configuration:

```
feeder.content.semantic-search.mime-types[image/*]=true
```

If the map does not contain any MIME-type, then the catch-all type `*/*` is enabled implicitly.

```
feeder.content.semantic-search.model-version-media
```

Type	String
Default	1
Description	The embedding model version to use for media inputs (e.g. images). Default is "1".

```
feeder.content.semantic-search.model-version-text
```

Type	String
Default	1
Description	The embedding model version to use for text inputs. Default is "1".

```
feeder.content.semantic-search.query-min-interval
```

Type	Duration
Default	1m
Description	The minimum time to wait between starting queries (for example, after empty query result). It does not delay multiple queries used for paginating large results.

```
feeder.content.semantic-search.semantic-vector-dimension
```

Type	Integer
Default	1024
Description	The dimension of the semantic vectors. Has to match the dimension of the embedding model. Has to match the size of the solr field type "vector_embeddings_{dimension}".

feeder.content.semantic-search.skip-on-initialize	
Type	Boolean
Default	true
Description	Skip processing existing contents during initial feeding.
feeder.content.semantic-search.text-body-max-length	
Type	Integer
Default	50000
Description	The max length of the text that can be used for embedding Default is 50_000, because Bedrock Nova is supporting this.
feeder.content.semantic-search.text-body-threshold	
Type	Map<String, Integer>
Description	Configure minimum text body lengths per com.coremedia.cap.content.Content-Type. Content items with shorter text bodies are ignored for semantic search indexing. If not configured, all text bodies are indexed if they are longer than the configured textBodyThresholdDefault. Example Configuration:
feeder.content.semantic-search.textBodyThreshold.CMQueryList=120	
feeder.content.semantic-search.text-body-threshold-default	
Type	Integer
Default	150
Description	The default minimum text body length (in characters) for semantic search indexing. Content items with text bodies shorter than this value are ignored unless a specific threshold is configured for their content type via textBodyThreshold. The unit is characters.

This property ensures that only sufficiently descriptive content is considered for semantic indexing when no type-specific threshold is set.

feeder.content.semantic-search.text-property-paths

Type Map<String,List<String>>

Description Configure com.coremedia.cap.content.Content property paths that point to text properties. Paths are given per com.coremedia.cap.content.ContentType.

Example Configuration:

```
feeder.content.semantic-search.textPropertyPaths.CMPicture=caption,alt,teaserTitle,teaserText,detailText
```

feeder.content.semantic-search.type

Type Map<String,Boolean>

Description The types of content items to feed. This property maps included content types to the value 'true'. Note, that entries with value 'false' are ignored and cannot be used to exclude a subtype of an included type. A change to this property will not trigger automatic re-indexing of formerly included or excluded content items. See the Search Manual, chapter "Reindexing" for details on re-indexing.

feeder.content.type.excludes

Type List<String>

Default [Preferences, EditorPreferences, Dictionary, Query]

Description List of abstract or concrete content types excluded from feeding. With the configuration of some type, all of its subtypes are excluded implicitly, if not configured otherwise. Note that it is an error to configure the same content type in this property and in feeder.content.type.includes. Rules for more specific types override rules for less specific types. Regular expressions are not supported.

feeder.content.type.includes

Type List<String>

Default	Content_
Description	List of abstract or concrete content types included for feeding. With the configuration of some type, all of its subtypes are included implicitly, if not configured otherwise. Note that it is an error to configure the same content type in this property and in <code>feeder.content.type.excludes</code> . Rules for more specific types override rules for less specific types. Regular expressions are not supported.
<code>feeder.content.update-groups-immediately</code>	
Type	Boolean
Default	false
Description	If <code>feeder.content.index-groups</code> is true, configures whether the field "groups" is updated immediately after a change of a folder's right rule. It is recommended to keep this set to false, and let the Content Feeder update the index field in the background with lower priority than updates for editorial changes. It is quite expensive to set this to true because all contents below the folder would be reindexed.

Table 3.45. Content Feeder Configuration Properties

Solr specific properties for Content Feeder

<code>feeder.content.issues.solr.fetch-size</code>	
Type	Integer
Default	1000
Description	The maximum number of results to fetch with a single paginated Solr query when retrieving content items with outdated issues. If more results are available, multiple queries with Solr cursor pagination will be used, and each one will be restricted to this configured maximum number of results.
<code>feeder.content.issues.solr.filter</code>	
Type	String
Default	types:Document_

Description Solr filter query to restrict the content items for which outdated issues are reindexed.

```
feeder.content.issues.solr.query-min-delay
```

Type Duration

Default 10s

Description The minimum time to wait before Solr is queried again for content items with outdated issues after the last query. This delay is not used for paginated queries that just retrieve the next page for a previous query.

```
feeder.solr.nested-documents.enabled
```

Type Boolean

Default true

Description Whether storing nested feedables as nested documents is supported in Solr. This requires that the Solr schema contains a `_root_` field. Note that if you add that field to the schema, you have to recreate the index from scratch.

```
feeder.solr.nested-documents.skip-index-check
```

Type Boolean

Default false

Description If `feeder.solr.nested-documents.enabled` is true, the Solr index schema is checked whether it contains the `_root_` field. The Feeder will log a warning and not use nested documents, if feeding of nested documents is attempted but the index does not support it. You can set this property to true to skip checking the index schema.

```
feeder.solr.partial-updates.enabled
```

Type Boolean

Default true

Description	Whether partial updates are supported for updating content metadata in Solr. This requires that all fields in the Solr index are configured as stored="true" or docValues="true" except fields that are copyField destinations, which must be configured as stored="false". This is because partial updates are applied to the index document reconstructed from the existing stored field values.
--------------------	--

```
feeder.solr.partial-updates.skip-index-check
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description	If feeder.solr.partial-updates.enabled is true, the Solr index schema is analyzed whether fields are stored as required for partial updates. The Feeder will log a warning and not use partial update functionality if the index seems to not support it. You can set this property to true to skip the check.
--------------------	--

```
feeder.solr.send-retry-delay
```

Type	Duration
-------------	----------

Default	30s
----------------	-----

Description	The delay to wait before the Feeder retries to send data after failures from Solr.
--------------------	--

```
solr.cloud
```

Type	Boolean
-------------	---------

Default	false
----------------	-------

Description	Whether to connect to SolrCloud. If true, connect to a SolrCloud cluster. SolrCloud connection details must be set either as ZooKeeper addresses (solr.zookeeper.addresses) or, if the former is unset or empty as HTTP URLs (solr.url). If false, connect to stand-alone Solr nodes via HTTP URLs (solr.url).
--------------------	---

```
solr.connection-timeout
```

Type	Duration
-------------	----------

Default	0
Description	The connection timeout set on the SolrJ SolrClient. It determines how long the client waits to establish a connection without any response from the server. The default value 0 means, that it will wait forever. Set a negative value to use the SolrClient default. (Default unit is milliseconds)
<code>solr.content.collection</code>	
Type	String
Default	studio
Description	The name of the Solr collection for editorial search.
<code>solr.content.config-set</code>	
Type	String
Default	content
Description	The name of the Solr config set to use when creating the collection for editorial search. This property is used by the Content Feeder.
<code>solr.index-data-directory</code>	
Type	String
Default	data
Description	Value for the "dataDir" parameter of the Solr CoreAdmin API / Collection API request to create a Solr index.
<code>solr.password</code>	
Type	String
Description	Password for HTTP basic authentication, used if a non-empty solr.username has been specified. The value may have been encrypted with the tool "cm encrypt-passwordproperty".

<code>solr.proxy-host</code>	
Type	String
Description	Proxy host for Solr communication that needs to be set if a proxy should be used.
<code>solr.proxy-is-secure</code>	
Type	Boolean
Default	false
Description	Secure flag for Solr proxy.
<code>solr.proxy-is-socks4</code>	
Type	Boolean
Default	false
Description	SOCKS 4 flag for Solr proxy.
<code>solr.proxy-port</code>	
Type	Integer
Default	0
Description	Proxy port for Solr communication that needs to be set if a proxy should be used.
<code>solr.socket-timeout</code>	
Type	Duration
Default	10m
Description	The socket timeout set on the SolrJ SolrClient. It determines how long the client waits for a response from the server after the connection was established and the request was already sent.

Set to 0 for no timeout, or to a negative value to use SolrClient default. (Default unit is milliseconds)

`solr.url`

Type `List<String>`

Default `http://localhost:40080/solr`

Description The list of Solr URLs to connect to.

These URLs are ignored if connecting to SolrCloud (`solr.cloud=true`) and non-empty ZooKeeper addresses (`solr.zookeeper.addresses`) have been set. For a Feeder application that is not connected to a SolrCloud cluster, a single URL to the Solr leader must be configured.

`solr.use-http1`

Type `Boolean`

Default `false`

Description Whether HTTP/1 (true) or HTTP/2 (false) shall be used by Solr clients.

`solr.use-xml-response-writer`

Type `Boolean`

Default `false`

Description Whether SolrJ should use XML response format instead of Javabin format.

`solr.username`

Type `String`

Description Username for HTTP basic authentication, or empty string for no authentication.

`solr.zookeeper.addresses`

Type `List<String>`

Description	ZooKeeper addresses for connecting to SolrCloud. Only used if solr.cloud=true.
<code>solr.zookeeper.chroot</code>	
Type	String
Description	Optional ZooKeeper chroot path for Solr. ZooKeeper chroot support makes it possible to isolate the SolrCloud tree in a ZooKeeper instance that is Only used if solr.cloud=true and solr.zookeeper.ad-dresses is set to non-empty value.
<code>solr.zookeeper.client-timeout</code>	
Type	Duration
Default	10s
Description	Client-timeout duration for ZooKeeper. Set to a negative value to use SolrClient default. (Default unit is milliseconds) Only used if solr.cloud=true and solr.zookeeper.addresses is set to non-empty value.
<code>solr.zookeeper.connect-timeout</code>	
Type	Duration
Default	10s
Description	Connect-timeout duration for ZooKeeper. Set to a negative value to use SolrClient default. (Default unit is milliseconds) Only used if solr.cloud=true and solr.zookeeper.addresses is set to non-empty value.

Table 3.46. Content Feeder Solr Configuration Properties

Batch configuration properties for Content Feeder

With these properties you can configure the processing of batches.

feeder.batch.max-bytes	
Type	org.springframework.util.unit.DataSize
Description	The maximum batch size in bytes. The Feeder sends a batch to the search engine if its maximum size would be exceeded when adding more entries. Note, that byte computation is a rough estimate only. A smaller batch may be sent if the maximum number of index documents is reached before, or if configured delays are reached. The Content Feeder default value is 5 MB. The CAE Feeder default value is 20 MB.
feeder.batch.max-open	
Type	Integer
Default	5
Description	The maximum number of batches indexed in parallel. This setting is not used with the default integration of Apache Solr but only with custom implementations of the com.coremedia.cap.feeder.index.async.AsyncIndexer interface. The Feeder does not call the index method of the AsyncIndexer interface to index another batch if the maximum number of parallel batches has been reached. The method will not be called until a callback about the persistence of one of these batches has been received.
feeder.batch.max-processed	
Type	Integer
Default	1
Description	The maximum number of batches processed by the Indexer in parallel. This setting is not used with the default integration of Apache Solr but only with custom implementations of the com.coremedia.cap.feeder.index.async.AsyncIndexer interface. The Feeder does not call the index method of the AsyncIndexer interface to index another batch if the configured number of currently processed

batches has been reached. The method will not be called until a callback about completed processing or persistence of one of these batches has been received.

```
feeder.batch.max-size
```

Type	Integer
Default	500
Description	The maximum number of index documents in a batch. If the maximum number is reached, the Feeder sends the batch to the search engine. A smaller batch may be sent if the maximum byte size is reached before, or if configured delays are reached.

```
feeder.batch.retry-send-idle-delay
```

Type	Duration
Default	1m
Description	The time to wait before retrying to send index documents to the search engine after failures. This delay is used if the feeder is idle.

```
feeder.batch.retry-send-max-delay
```

Type	Duration
Default	10m
Description	The maximum time to wait before retrying to send index documents to the search engine after failures. This delay is used if the feeder is not idle. The setting is typically larger than retry-send-idle-delay.

```
feeder.batch.send-idle-delay
```

Type	Duration
Description	The time between adding an index document to a batch and sending that batch to the search engine, if the batch is not yet full according to the max-size and max-bytes configuration properties, and if the feeder is idle. If a change needs to be sent to the search engine, and no further changes were made within the specified time, then an index document for the change will be sent after that

time to the search engine. A small delay ensures low latency for changes to become visible in the search engine, as long as the system is not too busy.

The Content Feeder default value is 3 seconds.

The CAE Feeder default value is 10 seconds.

feeder.batch.send-max-delay	
Type	Duration
Description	The maximum time between adding an index document to a batch and sending that batch to the search engine. This setting is typically larger than send-idle-delay to allow batches to grow and increase throughput, for example when large amounts of content are created by an import process. The configured value may still be exceeded under high load, or if there are problems connecting to the search engine. The Content Feeder default value is 20 seconds. The CAE Feeder default value is 2 minutes.

Table 3.47. Feeder Batch Configuration Properties

Apache Tika Properties

You can customize text extraction with Apache Tika using the following properties:

feeder.tika.append-metadata	
Type	String
Description	Comma-separated list of metadata identifiers returned by Apache Tika to append to the extracted body text.
feeder.tika.config	
Type	org.springframework.core.io.Resource
Description	The location of a custom Tika Config XML, for example to customize the default Tika parsers. See Apache Tika documentation for details on configuring Tika. The value of this property must be a Spring Resource location (e.g. file:/path/to/local/file) or empty for defaults.

feeder.tika.copy-metadata	
Type	String
Description	Comma-separated list of metadata identifiers returned by Apache Tika and names of Feedable elements to copy the metadata to. Entries in the comma separated list have the following format: "metadata identifier"="element name". With Apache Solr, target index fields must be defined as multiValued="true" to avoid indexing errors if there are multiple metadata values with the same identifier.
feeder.tika.log-metrics-interval	
Type	Duration
Default	10m
Description	The time interval to log metrics about blob text and metadata extraction. Metrics are logged at level INFO, and only if changed. No metrics will be logged if the value is less than 1 second.
feeder.tika.parse	
Type	Map<org.springframework.util.MimeType, Boolean>
Description	The MIME-types of blobs in Feedable elements to parse with Apache Tika for plain text and/or metadata extraction. This property maps included MIME-types to true, and excluded MIME-types to false. Entries for more specific types like application/pdf override entries for less specific types like application/*. If the property does not map any MIME-type to true, then the catch-all type */* is included implicitly. Blobs of included MIME-types are parsed with Apache Tika, and are replaced with their extracted plain text and optional metadata in Feedable elements, so that the indexer receives string values. Blobs of excluded MIME-types are not parsed with Apache Tika, and are kept unchanged in Feedable elements, so that the indexer receives the original binary data. See also related property feeder.blob.enabled, which generally restricts feeding of blobs by MIME-type. Blobs that have been excluded with that property are neither parsed with Apache Tika, nor passed as binary data to the indexer.

Example configuration: Parse all blobs except JPEG images

```
feeder.tika.parse[image/jpeg]=false
```

Note, that MIME-type */* is included implicitly in this example, because no type is explicitly mapped to true.

```
feeder.tika.threads
```

Type	Integer
Description	The maximum number of threads used to extract text and metadata from binaries with Apache Tika. Note, that even with higher values, the Feeder does not automatically use more than one thread. The configuration limits Tika usage for the Content Feeder's semantic search feeding, which may call Tika internally.

```
feeder.tika.timeout
```

Type	Duration
Default	2m
Description	The maximum time after which text extraction from binary data with Apache Tika fails. If extraction fails, the binary data will be skipped for the index document. Lower values will avoid that the Feeder is blocked for a long time in text extraction.

```
feeder.tika.warn-time-lost-thread
```

Type	Duration
Default	10m
Description	The additional time after a text extraction timeout, after which another warning is logged if an interrupted parser thread is stuck and cannot be reused by the thread pool

```
feeder.tika.warn-time-threshold
```

Type	Duration
------	----------

Default	15s
Description	The time after which a warning is logged when text extraction from binary data with Apache Tika takes some time.
<code>feeder.tika.zip-bomb-prevention.enabled</code>	
Type	Boolean
Default	true
Description	Sets whether Apache Tika's "Zip bomb" prevention is enabled. When a "Zip bomb" is detected, no text will be extracted from the Blob, but a warning will be logged. Note that "Zip bombs" are not restricted to ZIP files but also apply to PDFs or other formats. Disabled "Zip bomb" prevention bears the risk of OutOfMemory-Error-s. Note that false positives are possible.
<code>feeder.tika.zip-bomb-prevention.maximum-compression-ratio</code>	
Type	Long
Default	-1
Description	Sets the ratio between output characters and input bytes for the Apache Tika "Zip bomb" prevention. If this ratio is exceeded (after the output threshold has been reached) then no text will be extracted and a warning will be logged. Set to -1 to use the default of Apache Tika.
<code>feeder.tika.zip-bomb-prevention.maximum-depth</code>	
Type	Integer
Default	-1
Description	Sets the maximum XML element nesting level for the Apache Tika "Zip bomb" prevention. If this depth level is exceeded then no text will be extracted, and a warning will be logged. Set to -1 to use the default of Apache Tika.
<code>feeder.tika.zip-bomb-prevention.maximum-package-entry-depth</code>	
Type	Integer

Default	-1
Description	Sets the maximum package entry nesting level for the Apache Tika "Zip bomb" prevention. If this depth level is exceeded then no text will be extracted, and a warning will be logged. Set to -1 to use the default of Apache Tika.

Table 3.48. Feeder Tika Configuration Properties

Feeder Core Properties

You can use the following properties to customize some internal settings of the *Content Feeder*.

feeder.core.executor-queue-capacity	
Type	Integer
Description	Maximum capacity of the Feeder's executor queue, which is internally used to transfer evaluated values. The value must be greater than 0. The Content Feeder default value is 100. The CAE Feeder default value is 2000. A larger value may lead to increased usage of heap memory.
feeder.core.executor-retry-delay	
Type	Duration
Default	1m
Description	The delay to wait before the Feeder retries to access the source data after failures.

Table 3.49. Feeder Core Configuration Properties

3.11 CAE Feeder Properties

General Properties

caefeeder.content.include-subtypes	
Type	Boolean
Default	true
Description	Whether the configuration of content types in property 'caefeeder.content.type' includes subtypes. A change to this property will not trigger automatic re-indexing of formerly included or excluded content items. See the Search Manual, chapter "Reindexing" for details on re-indexing.
caefeeder.content.path	
Type	Map<String, Boolean>
Description	The paths of content items to feed. This property maps included base folder paths to the value 'true'. Paths start with a slash. Note, that the configuration of overlapping paths is not supported, and that entries with value 'false' are ignored. It is not possible to exclude a subfolder of an included folder. A change to this property will not trigger automatic re-indexing of formerly included or excluded content items. See the Search Manual, chapter "Reindexing" for details on re-indexing. This property is ignored, if deprecated property 'feeder.contentSelector.basePath' is still set.
caefeeder.content.type	
Type	Map<String, Boolean>
Description	The types of content items to feed. This property maps included content types to the value 'true'. Note, that entries with value 'false' are ignored and cannot be used to exclude a subtype of an included type. A change to this property will not trigger automatic re-indexing of formerly included or excluded content items. See the Search Manual, chapter "Reindexing" for details on re-indexing. This property is ignored, if deprecated property 'feeder.contentSelector.contentTypes' is still set.

```
caefeeder.evaluation.delay
```

Type	Duration
------	----------

Default	0
---------	---

Description	The minimum delay between evaluations. Higher values lead to reduced throughput.
-------------	--

```
caefeeder.evaluation.error-retry-delay
```

Type	Duration
------	----------

Default	10m
---------	-----

Description	The delay before retrying the evaluation of data for feeding, after it threw an exception and triggered feeding of a placeholder document. If set to zero, then feeding will not be retried automatically after some time. Feeding may be retried earlier, when a dependency is invalidated that was accessed before the exception was thrown.
-------------	--

```
caefeeder.evaluation.idle-delay
```

Type	Duration
------	----------

Default	10s
---------	-----

Description	The delay between evaluations if the application is idle. Smaller values can be configured to reduce the latency but may lead to increased load on the database.
-------------	--

```
caefeeder.evaluation.log-interval
```

Type	Duration
------	----------

Default	10m
---------	-----

Description	The time interval to log some statistics about the progress, including the number of keys that are currently invalid and still need to be computed. (Default unit is seconds)
-------------	---

```
caefeeder.evaluation.threads
```

Type	Integer
Default	50
Description	The number of evaluate threads, which influences performance not only because evaluations can execute concurrently but also because higher values increase the probability that the CAE Feeder writes the state of multiple evaluations to the database in one database transaction.
<code>caefeeder.index-verifier.enabled</code>	
Type	Boolean
Default	true
Description	Whether the CAE Feeder checks that the configured database and search engine index were initialized by the same CAE Feeder instance. If enabled, the CAE Feeder stores the same random token in both the database and the search engine index initially, and compares stored values when restarted. Normally, this setting should not be changed. In case of problems or custom integrations of search engines that cannot store such a token, you can set this to false.
<code>caefeeder.invalidation.paused</code>	
Type	Boolean
Default	false
Description	Whether dependency invalidation is paused. Invalidations will not be lost. They will be handled when invalidation is resumed.
<code>caefeeder.invalidation.timed-dependency-interval</code>	
Type	Duration
Default	10s
Description	The interval to invalidate timed dependencies, which represents the resolution of timed dependencies.

```
caefeeder.store.dependency-job-queue-capacity
```

Type	Integer
Default	10000
Description	The maximum number of pending jobs to change stored dependencies. If the maximum is reached, scheduling more jobs for execution will fail.

```
caefeeder.store.dependency-names
```

Type	Boolean
Default	false
Description	Whether dependency names are stored as separate column in addition to dependency hashes in table pdcdependencies. This can be useful for debugging/monitoring but will decrease performance, increase the size of the table, and is not recommended for production usage.

```
caefeeder.store.dependency-transaction-max-weight
```

Type	Integer
Default	2500
Description	The maximum weight of a database transaction to change stored dependencies. The weight is interpreted as the number of changed keys, that is, a transaction with one deleted key has weight 1. Multiple transactions will be used to process an event that causes the invalidation of more keys.

```
caefeeder.store.reset
```

Type	Boolean
Default	false
Description	Whether the CAE Feeder should be reset, and database tables should be cleared on application startup. If true, previous data will be lost.

```
contentDependencyInvalidator.invalidatedStopped
```

Default	false
Description	Whether dependency invalidation is paused.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.invalidation.paused</code> instead. Reason: The property was renamed.
<code>dependencyStore.jobQueueCapacity</code>	
Default	10000
Description	The maximum number of pending jobs to change stored dependencies.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.store.dependency-job-queue-capacity</code> instead. Reason: The property was renamed.
<code>dependencyStore.maxTransactionWeight</code>	
Default	2500
Description	The maximum weight of a database transaction to change stored dependencies.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.store.dependency-transaction-max-weight</code> instead. Reason: The property was renamed.
<code>dependencyStore.storeDependencyNames</code>	
Default	false

Description	Whether dependency names are stored as separate column in addition to dependency hashes in table pdcdependencies. This can be useful for debugging/monitoring but will decrease performance, increase the size of the table, and is not recommended for production usage.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.store.dependency-names</code> instead. Reason: The property was renamed.
<code>feeder.beanMapping.mimeType.excludes</code>	
Description	A comma-separated list of excluded MIME types for blob values added by the BeanMappingFeedablePopulator.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>feeder.blob.enabled</code> instead. Reason: The property was replaced by 'feeder.blob.enabled', which takes a map value instead of a comma-separated list and also limits blobs that are added by custom FeedablePopulator implementations.
<code>feeder.beanMapping.mimeType.includes</code>	
Default	<code>/*/*</code>
Description	A comma-separated list of included MIME types for blob values added by the BeanMappingFeedablePopulator.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>feeder.blob.enabled</code> instead. Reason: The property was replaced by 'feeder.blob.enabled', which takes a map value instead of a comma-separated list and also limits blobs that are added by custom FeedablePopulator implementations.

feeder.beanPropertyMaxBytes	
Default	-1
Description	The maximum size in bytes for the value of a bean property or -1 for no limitation. Larger values are ignored and will not be sent to the Search Engine.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>feeder.blob.max-size[*/*]</code> instead. Reason: The property was replaced by 'feeder.blob.max-size', which allows configuration by MIME-type and also limits blobs that are added by custom FeedablePopulator implementations.

feeder.blob.enabled	
Type	Map<org.springframework.util.MimeType, Boolean>
Description	The MIME-types of indexed blobs. This property maps included MIME-types to true, and excluded MIME-types to false. Entries for more specific types like text/xml override mappings for less specific types like text/*. Example Configuration: <pre>feeder.blob.enabled[application/pdf]=true</pre> If the map does not contain any enabled MIME-type, then the catch-all type */* is enabled implicitly. See also property <code>feeder.blob.max-size</code>, which restricts blobs based on their size. Blobs are only included if allowed by both <code>feeder.blob.enabled</code> and <code>feeder.blob.max-size</code>.

feeder.blob.max-size	
Type	Map<org.springframework.util.MimeType, org.springframework.util.unit.DataSize>
Description	The maximum size of indexed blobs per MIME-type. Larger blob values will be skipped.

This property maps MIME-types to their maximum size. Values must be non-negative. Entries for more specific types like text/xml override mappings for less specific types like text/*.

Example Configuration:

```
feeder.blob.max-size[application/pdf]=5MB
```

If the map does not contain a configuration for an enabled MIME-type, then blobs of that MIME-type are indexed regardless of their size.

See also property `feeder.blob.enabled`, which restricts blob values based on type. Blobs are only included if allowed by both `feeder.blob.enabled` and `feeder.blob.max-size`.

```
feeder.contentSelector.basePath
```

Default	/Site
Description	A comma-separated list of base folder paths below which content beans are indexed.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeedер.content.path</code> instead. Reason: The property was replaced by '<code>caefeedер.content.path</code>', which takes a map value instead of a comma-separated list.

```
feeder.contentSelector.contentTypes
```

Default	Document_
Description	A comma-separated list of content types for which content beans are indexed.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeedер.content.type</code> instead. Reason: The property was replaced by '<code>caefeedер.content.type</code>', which takes a map value instead of a comma-separated list.

```
feeder.contentSelector.includeSubTypes
```

Default	true
Description	Whether the configuration of content types includes subtypes.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.content.include-subtypes</code> instead. Reason: The property was renamed.

```
persistentcache.clear
```

Default	false
Description	Whether the CAE Feeder should be reset, and database tables should be cleared on application startup. If true, previous data will be lost.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.store.reset</code> instead. Reason: The property was renamed.

```
proactiveengine.log.progress.interval.seconds
```

Default	600
Description	The time interval in seconds to log some statistics about the progress, including the number of keys that are currently invalid and still need to be computed.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.evaluation.log-interval</code> instead. Reason: The property was renamed.

proactiveengine.senders.delay	
Default	0
Description	The minimum delay in milliseconds between evaluations.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.evaluation.delay</code> instead. Reason: The property was renamed.
proactiveengine.senders.evaluators	
Default	50
Description	The number of evaluate threads.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.evaluation.threads</code> instead. Reason: The property was renamed.
proactiveengine.senders.idledelay	
Default	10000
Description	The delay in milliseconds between evaluations if the application is idle.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.evaluation.idle-delay</code> instead. Reason: The property was renamed.
timedDependencyInvalidator.timerResolutionMinimum	

Default	10000
Description	The interval in milliseconds to invalidate timed dependencies, which represents the resolution of timed dependencies.
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.invalidation.timed-dependency-interval</code> instead. Reason: The property was renamed.

Table 3.50. CAE Feeder Properties

Database Properties

The properties in this section are used to configure the database, also known as data source. This is only an excerpt of all available properties. For more, consult the Spring Boot documentation for `spring.datasource` properties ([SQL Databases :: Spring Boot](#), [Spring Boot Appendix: Data Properties](#)), which are available for the *CAE Feeder* under the application specific name `caefeed-er.datasource`, and are backed by Spring Boot class [org.springframework.boot.autoconfigure.jdbc.DataSourceProperties](#).

HikariCP is used as database connection pool. To fine-tune its settings, see the Spring Boot documentation for `spring.datasource.hikari` properties, which are available for the *CAE Feeder* under the name `caefeed-er.datasource.hikari`.

<code>caefeed-er.datasource.password</code>	
Type	String
Description	Login password of the database, possibly encrypted.
<code>caefeed-er.datasource.url</code>	
Type	String
Description	JDBC URL of the database connection for the CAE Feeder.

caefeeder.datasource.username	
Type	String
Description	Login username of the database.
jdbc.driver	
Description	Fully qualified name of the JDBC driver (ignored).
Deprecation	This property has been deprecated since 2412.0.0 and will be removed in a future version. Use <code>caefeeder.datasource.driver-class-name</code> instead. Reason: The database driver class name does not need to be specified anymore, because it gets auto-detected for the JDBC URL. If really needed, the auto-detected driver class name can still be overridden with property '<code>caefeeder.data-source.driver-class-name</code>'.
jdbc.login-user-name	
Description	Login username of the database.
Deprecation	This property has been deprecated since 2412.0.0 and will be removed in a future version. Use <code>caefeeder.datasource.username</code> instead. Reason: The property was renamed to use the '<code>caefeeder.datasource</code>' prefix, which exposes more data source properties from <code>org.springframework.boot.jdbc.autoconfigure.DataSourceProperties</code>. Use '<code>caefeeder.datasource.username</code>' instead with the full login username.
jdbc.password	
Description	Login password of the database, possibly encrypted.
Deprecation	This property has been deprecated since 2412.0.0 and will be removed in a future version. Use <code>caefeeder.datasource.password</code> instead.

Reason: The property was renamed to use the 'caefeeder.datasource' prefix, which exposes more data source properties from org.springframework.boot.jdbc.autoconfigure.DataSourceProperties.	
jdbc.url	
Description	JDBC URL of the database.
Deprecation	This property has been deprecated since 2412.0.0 and will be removed in a future version. Use caefeeder.datasource.url instead. Reason: The property was renamed to use the 'caefeeder.datasource' prefix, which exposes more data source properties from org.springframework.boot.jdbc.autoconfigure.DataSourceProperties.
jdbc.user	
Description	Login username of the database.
Deprecation	This property has been deprecated since 2412.0.0 and will be removed in a future version. Use caefeeder.datasource.username instead. Reason: The property was renamed to use the 'caefeeder.datasource' prefix, which exposes more data source properties from org.springframework.boot.jdbc.autoconfigure.DataSourceProperties. Use 'caefeeder.datasource.username' instead with the full login username.

Table 3.51. CAE Feeder Data Source Properties

Database Locking Properties

Each *CAE Feeder* needs its own dedicated database. The properties in this section configure the locking mechanism to detect configuration errors when the same database is accidentally used by multiple *CAE Feeder* instances.

caefeeder.lock.application-identifier

Type	String
Description	The application identifier that is used to identify the owning application of a lock stored in the database. If unset, the application identifier will be computed as [hostname]/[ip]:[current working directory]:[applicationName].
caefeeder.lock.application-name	
Type	String
Default	caefeeder
Description	The application name that will be used to compute an application identifier, if not set with 'caefeeder.lock.application-identifier'.
caefeeder.lock.check-interval	
Type	Duration
Default	1s
Description	The minimum time to wait between checking the database if a requested lock has been unlocked.
caefeeder.lock.enabled	
Type	Boolean
Default	true
Description	Whether locking is enabled.
caefeeder.lock.expiry	
Type	Duration
Default	30s
Description	The time after which a lock is considered expired, if not refreshed.
caefeeder.lock.refresh-interval	

Type	Duration
Default	10s
Description	The time between refreshing acquired locks in the database.
<code>caefeeder.lock.timeout</code>	
Type	Duration
Default	30s
Description	The time to wait for a lock to become available, or 0 to not wait.
<code>jdbc.lock.enable</code>	
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.lock.enabled</code> instead. Reason: The property was renamed.
<code>jdbc.lock.timeout</code>	
Deprecation	This property has been deprecated since 2506.0.0 and will be removed in a future version. Use <code>caefeeder.lock.timeout</code> instead. Reason: The property was renamed.

Table 3.52. CAE Feeder Database Locking Properties

Apache Tika Properties

You can customize text extraction with Apache Tika using the following properties:

```
feeder.tika.append-metadata
```

Type	String
Description	Comma-separated list of metadata identifiers returned by Apache Tika to append to the extracted body text.
<code>feeder.tika.config</code>	
Type	<code>org.springframework.core.io.Resource</code>
Description	The location of a custom Tika Config XML, for example to customize the default Tika parsers. See Apache Tika documentation for details on configuring Tika. The value of this property must be a Spring Resource location (e.g. <code>file:/path/to/local/file</code>) or empty for defaults.
<code>feeder.tika.copy-metadata</code>	
Type	String
Description	Comma-separated list of metadata identifiers returned by Apache Tika and names of Feedable elements to copy the metadata to. Entries in the comma separated list have the following format: "metadata identifier"="element name". With Apache Solr, target index fields must be defined as <code>multiValued="true"</code> to avoid indexing errors if there are multiple metadata values with the same identifier.
<code>feeder.tika.log-metrics-interval</code>	
Type	Duration
Default	10m
Description	The time interval to log metrics about blob text and metadata extraction. Metrics are logged at level INFO, and only if changed. No metrics will be logged if the value is less than 1 second.
<code>feeder.tika.parse</code>	
Type	<code>Map<org.springframework.util.MimeType, Boolean></code>
Description	The MIME-types of blobs in Feedable elements to parse with Apache Tika for plain text and/or metadata extraction.

This property maps included MIME-types to true, and excluded MIME-types to false. Entries for more specific types like application/pdf override entries for less specific types like application/*. If the property does not map any MIME-type to true, then the catch-all type */* is included implicitly.

Blobs of included MIME-types are parsed with Apache Tika, and are replaced with their extracted plain text and optional metadata in Feedable elements, so that the indexer receives string values.

Blobs of excluded MIME-types are not parsed with Apache Tika, and are kept unchanged in Feedable elements, so that the indexer receives the original binary data.

See also related property feeder.blob.enabled, which generally restricts feeding of blobs by MIME-type. Blobs that have been excluded with that property are neither parsed with Apache Tika, nor passed as binary data to the indexer.

Example configuration: Parse all blobs except JPEG images

```
feeder.tika.parse[image/jpeg]=false
```

Note, that MIME-type */* is included implicitly in this example, because no type is explicitly mapped to true.

feeder.tika.threads

Type	Integer
Description	The maximum number of threads used to extract text and metadata from binaries with Apache Tika. Note, that even with higher values, the Feeder does not automatically use more than one thread. The configuration limits Tika usage for the Content Feeder's semantic search feeding, which may call Tika internally.

feeder.tika.timeout

Type	Duration
Default	2m
Description	The maximum time after which text extraction from binary data with Apache Tika fails. If extraction fails, the binary data will be skipped for the index document. Lower values will avoid that the Feeder is blocked for a long time in text extraction.

```
feeder.tika.warn-time-lost-thread
```

Type	Duration
Default	10m
Description	The additional time after a text extraction timeout, after which another warning is logged if an interrupted parser thread is stuck and cannot be reused by the thread pool

```
feeder.tika.warn-time-threshold
```

Type	Duration
Default	15s
Description	The time after which a warning is logged when text extraction from binary data with Apache Tika takes some time.

```
feeder.tika.zip-bomb-prevention.enabled
```

Type	Boolean
Default	true
Description	Sets whether Apache Tika's "Zip bomb" prevention is enabled. When a "Zip bomb" is detected, no text will be extracted from the Blob, but a warning will be logged. Note that "Zip bombs" are not restricted to ZIP files but also apply to PDFs or other formats. Disabled "Zip bomb" prevention bears the risk of OutOfMemory-Error-s. Note that false positives are possible.

```
feeder.tika.zip-bomb-prevention.maximum-compression-ratio
```

Type	Long
Default	-1
Description	Sets the ratio between output characters and input bytes for the Apache Tika "Zip bomb" prevention. If this ratio is exceeded (after the output threshold has been reached) then no text will be extracted and a warning will be logged. Set to -1 to use the default of Apache Tika.

feeder.tika.zip-bomb-prevention.maximum-depth	
Type	Integer
Default	-1
Description	Sets the maximum XML element nesting level for the Apache Tika "Zip bomb" prevention. If this depth level is exceeded then no text will be extracted, and a warning will be logged. Set to -1 to use the default of Apache Tika.
feeder.tika.zip-bomb-prevention.maximum-package-entry-depth	
Type	Integer
Default	-1
Description	Sets the maximum package entry nesting level for the Apache Tika "Zip bomb" prevention. If this depth level is exceeded then no text will be extracted, and a warning will be logged. Set to -1 to use the default of Apache Tika.

Table 3.53. Feeder Tika Configuration Properties

Apache Solr Client Properties

The following properties are only used for a *CoreMedia Search Engine* based on Apache Solr:

feeder.solr.nested-documents.enabled	
Type	Boolean
Default	true
Description	Whether storing nested feedables as nested documents is supported in Solr. This requires that the Solr schema contains a <code>_root_</code> field. Note that if you add that field to the schema, you have to recreate the index from scratch.
feeder.solr.nested-documents.skip-index-check	
Type	Boolean
Default	false

Description	If <code>feeder.solr.nested-documents.enabled</code> is true, the Solr index schema is checked whether it contains the <code>_root_</code> field. The Feeder will log a warning and not use nested documents, if feeding of nested documents is attempted but the index does not support it. You can set this property to true to skip checking the index schema.
<code>feeder.solr.send-retry-delay</code>	
Type	Duration
Default	30s
Description	The delay to wait before the Feeder retries to send data after failures from Solr.
<code>solr.cae.collection</code>	
Type	String
Default	preview
Description	The name of the Solr collection for web-site search. This property is typically set to 'preview' or 'live' depending on the environment.
<code>solr.cae.config-set</code>	
Type	String
Default	cae
Description	The name of the Solr config set to use when creating the CAE collection. This property is used by the CAE Feeder.
<code>solr.cloud</code>	
Type	Boolean
Default	false
Description	Whether to connect to SolrCloud. If true, connect to a SolrCloud cluster. SolrCloud connection details must be set either as ZooKeeper addresses (<code>solr.zookeeper.addresses</code>) or, if the former is

unset or empty as HTTP URLs (solr.url). If false, connect to stand-alone Solr nodes via HTTP URLs (solr.url).

solr.connection-timeout	
Type	Duration
Default	0
Description	The connection timeout set on the SolrJ SolrClient. It determines how long the client waits to establish a connection without any response from the server. The default value 0 means, that it will wait forever. Set a negative value to use the SolrClient default. (Default unit is milliseconds)
solr.index-data-directory	
Type	String
Default	data
Description	Value for the "dataDir" parameter of the Solr CoreAdmin API / Collection API request to create a Solr index.
solr.password	
Type	String
Description	Password for HTTP basic authentication, used if a non-empty solr.username has been specified. The value may have been encrypted with the tool "cm encrypt-passwordproperty".
solr.proxy-host	
Type	String
Description	Proxy host for Solr communication that needs to be set if a proxy should be used.
solr.proxy-is-secure	
Type	Boolean

Default	false
Description	Secure flag for Solr proxy.
<code>solr.proxy-is-socks4</code>	
Type	Boolean
Default	false
Description	SOCKS 4 flag for Solr proxy.
<code>solr.proxy-port</code>	
Type	Integer
Default	0
Description	Proxy port for Solr communication that needs to be set if a proxy should be used.
<code>solr.socket-timeout</code>	
Type	Duration
Default	10m
Description	The socket timeout set on the SolrJ SolrClient. It determines how long the client waits for a response from the server after the connection was established and the request was already sent. Set to 0 for no timeout, or to a negative value to use SolrClient default. (Default unit is milliseconds)
<code>solr.url</code>	
Type	List<String>
Default	<code>http://localhost:40080/solr</code>
Description	The list of Solr URLs to connect to.

These URLs are ignored if connecting to SolrCloud (`solr.cloud=true`) and non-empty ZooKeeper addresses (`solr.zookeeper.addresses`) have been set. For a Feeder application that is not connected to a SolrCloud cluster, a single URL to the Solr leader must be configured.

```
solr.use-http1
```

Type Boolean

Default false

Description Whether HTTP/1 (true) or HTTP/2 (false) shall be used by Solr clients.

```
solr.use-xml-response-writer
```

Type Boolean

Default false

Description Whether SolrJ should use XML response format instead of Javabin format.

```
solr.username
```

Type String

Description Username for HTTP basic authentication, or empty string for no authentication.

```
solr.zookeeper.addresses
```

Type List<String>

Description ZooKeeper addresses for connecting to SolrCloud.
Only used if `solr.cloud=true`.

```
solr.zookeeper.chroot
```

Type String

Description Optional ZooKeeper chroot path for Solr.

ZooKeeper chroot support makes it possible to isolate the SolrCloud tree in a ZooKeeper instance that is Only used if solr.cloud=true and solr.zookeeper.ad-
dresses is set to non-empty value.

solr.zookeeper.client-timeout	
Type	Duration
Default	10s
Description	Client-timeout duration for ZooKeeper. Set to a negative value to use SolrClient default. (Default unit is milliseconds) Only used if solr.cloud=true and solr.zookeeper.addresses is set to non-empty value.
solr.zookeeper.connect-timeout	
Type	Duration
Default	10s
Description	Connect-timeout duration for ZooKeeper. Set to a negative value to use SolrClient default. (Default unit is milliseconds) Only used if solr.cloud=true and solr.zookeeper.addresses is set to non-empty value.

Table 3.54. CAE Feeder Solr Configuration Properties

Feeder Batch Properties

With these properties you can configure the processing of batches.

feeder.batch.max-bytes	
Type	org.springframework.util.unit.DataSize
Description	The maximum batch size in bytes. The Feeder sends a batch to the search engine if its maximum size would be exceeded when adding more entries. Note, that byte computation is a rough estimate only. A smaller batch may be sent if the maximum number of index documents is reached before, or if configured delays are reached.

The Content Feeder default value is 5 MB.

The CAE Feeder default value is 20 MB.

```
feeder.batch.max-open
```

Type	Integer
Default	5
Description	The maximum number of batches indexed in parallel. This setting is not used with the default integration of Apache Solr but only with custom implementations of the <code>com.coremedia.cap.feeder.index.async.AsyncIndexer</code> interface. The Feeder does not call the <code>index</code> method of the <code>AsyncIndexer</code> interface to index another batch if the maximum number of parallel batches has been reached. The method will not be called until a callback about the persistence of one of these batches has been received.

```
feeder.batch.max-processed
```

Type	Integer
Default	1
Description	The maximum number of batches processed by the Indexer in parallel. This setting is not used with the default integration of Apache Solr but only with custom implementations of the <code>com.coremedia.cap.feeder.index.async.AsyncIndexer</code> interface. The Feeder does not call the <code>index</code> method of the <code>AsyncIndexer</code> interface to index another batch if the configured number of currently processed batches has been reached. The method will not be called until a callback about completed processing or persistence of one of these batches has been received.

```
feeder.batch.max-size
```

Type	Integer
Default	500
Description	The maximum number of index documents in a batch. If the maximum number is reached, the Feeder sends the batch to the search engine. A smaller batch may be sent if the maximum byte size is reached before, or if configured delays are reached.

```
feeder.batch.retry-send-idle-delay
```

Type	Duration
Default	1m
Description	The time to wait before retrying to send index documents to the search engine after failures. This delay is used if the feeder is idle.

```
feeder.batch.retry-send-max-delay
```

Type	Duration
Default	10m
Description	The maximum time to wait before retrying to send index documents to the search engine after failures. This delay is used if the feeder is not idle. The setting is typically larger than retry-send-idle-delay.

```
feeder.batch.send-idle-delay
```

Type	Duration
Description	The time between adding an index document to a batch and sending that batch to the search engine, if the batch is not yet full according to the max-size and max-bytes configuration properties, and if the feeder is idle. If a change needs to be sent to the search engine, and no further changes were made within the specified time, then an index document for the change will be sent after that time to the search engine. A small delay ensures low latency for changes to become visible in the search engine, as long as the system is not too busy. The Content Feeder default value is 3 seconds. The CAE Feeder default value is 10 seconds.

```
feeder.batch.send-max-delay
```

Type	Duration
Description	The maximum time between adding an index document to a batch and sending that batch to the search engine. This setting is typically larger than send-idle-delay to allow batches to grow and increase throughput, for example when large amounts of content are created by an import process. The configured value

may still be exceeded under high load, or if there are problems connecting to the search engine.

The Content Feeder default value is 20 seconds.

The CAE Feeder default value is 2 minutes.

Table 3.55. Feeder Batch Configuration Properties

Feeder Core Properties

You can use the following properties to customize some internal settings of the CAE Feeder.

feeder.core.executor-queue-capacity	
Type	Integer
Description	Maximum capacity of the Feeder's executor queue, which is internally used to transfer evaluated values. The value must be greater than 0. The Content Feeder default value is 100. The CAE Feeder default value is 2000. A larger value may lead to increased usage of heap memory.
feeder.core.executor-retry-delay	
Type	Duration
Default	1m
Description	The delay to wait before the Feeder retries to access the source data after failures.

Table 3.56. Feeder Core Configuration Properties

3.12 UAPI Client Properties

3.12.1 Unified API Spring Boot Client Properties

repository.blob-cache-path	
Type	String
Description	The directory in which cached blobs are stored. Make sure that the file system for this directory is large enough. Note that forced shutdowns of component's web application may result in leftover files in this directory, which should be cleared while the components are down. The configured directory is shared among components, because the actual cache content is placed in dynamically allocated subdirectories.
repository.blob-cache-size	
Type	org.springframework.util.unit.DataSize
Default	-1B
Description	The maximum allowed size of the disk cache for blobs. The blobs are cached temporarily and are garbage collected if no longer needed. Note that the file system overhead for storing the files does not count towards this value. So the physical space that has to be reserved on the disk for the cache has to be slightly higher than value of this configuration property. The value of "-1" means, that the default value of 32,000,000 bytes is used for disk caches. If several concurrent threads write large blobs at the same time, the deletion of the folder with the old unused files can be postponed for later. This is the second reason why the maximum allowed cache size can grow slightly higher than this configuration property. The size of such deviation depends on the blobs size as well as the amount of parallel threads.
repository.blob-streaming-size-threshold	
Type	org.springframework.util.unit.DataSize

Default	-1B
Description	The minimum size of streamed blobs. blobs less than or equal to this size will be downloaded completely to disk before the first byte can be read. Larger blobs will be downloaded in the background. Negative values trigger a fallback to the default of 128 KB.
<code>repository.blob-streaming-threads</code>	
Type	Integer
Default	-1
Description	The number of threads reserved for streaming blob. The value of "-1" means, that the default value "2" is used.
<code>repository.blob-upload.connect-timeout</code>	
Type	Duration
Default	60s
Description	The timeout used for establishing a connection to the server for blob uploads.
<code>repository.blob-upload.request-timeout</code>	
Type	Duration
Default	1h
Description	The timeout used for blob uploads. When uploading a blob, the data of the response must become available for reading before this timeout is exceeded.
<code>repository.caplist.connect</code>	
Type	Boolean
Description	Whether to connect the Cap list repository. Disabled by default.
<code>repository.caplist.mongo-db-client-uri</code>	

Type	String
Default	mongodb://localhost:27017
Description	The Cap List MongoDB connection string URI is used to configure your MongoDB connection. Property 'readpreference' must be 'primary'. See also mongoDb documentation .
<code>repository.caplist.mongo-db-prefix</code>	
Type	String
Default	elastic
Description	Prefix for Cap List MongoDB database names.
<code>repository.connect-retry-delay-seconds</code>	
Type	Duration
Default	10s
Description	The delay between UAPI connect retries.
<code>repository.domain</code>	
Type	String
Description	The domain of the <i>Content Server</i>
<code>repository.enable-blob-download-urls</code>	
Type	Boolean
Default	true
Description	A flag which allows to disable serving blob download URLs. The server side download URL computation is expensive, esp. if a client iterates over many contents, so that the download URLs are evicted from the cache. Clients may reduce server load by setting this flag to false, if they don't need the download URLs. Default is true.

repository.event-chunk-size	
Type	Integer
Default	-1
Description	Maximum number of events fetched from a server at once. If zero or negative, the default value 1000 will be used.
repository.force-immediate-login	
Type	Boolean
Default	false
Description	Configuration option that forces the UAPI connection to connect immediately instead of waiting for the <i>Content Server</i> to come up. This is useful for command line tools like the workflow converter that should provide quick feedback.
repository.heap-cache-size	
Type	org.springframework.util.unit.DataSize
Default	-1B
Description	The total number of bytes used by the main memory cache. For 32 bit JVMs this value is exact, for 64 bit JVMs, the actual memory consumption may be up to 2 times the configured value.
repository.http-base-uri	
Type	URI
Description	The base URI of the Content Server where the HTTP servlets are provided, e.g. blob servlet. Only required if gRPC is active (see property "repository.use-grpc").
repository.max-cached-blob-size	
Type	org.springframework.util.unit.DataSize
Default	-1B

Description	The maximum size of blobs that are cached on the local disk. Larger blobs are downloaded from the Content Server on every request. The value of "-1" means, that the default value <i>Integer#MAX_VALUE</i> is used.
--------------------	--

```
repository.password
```

Type	String
-------------	--------

Description	The password of the user.
--------------------	---------------------------

```
repository.send-session-token-in-url
```

Type	Boolean
-------------	---------

Default	true
----------------	------

Description	If true, the session token for connection to Content Server servlets (for blob up/download and CPU usage) will be sent as a URL query parameter. Otherwise, it will be sent as a request field. While sending the token as a query parameter is inherently insecure, it is the default to keep backward compatibility. Unless connection to an older server is required, it is recommended to set this property to false. @since 2512.0
--------------------	---

```
repository.url
```

Type	String
-------------	--------

Description	The URL of the <i>Content Server</i> . Only required if gRPC is inactive (see property "repository.use-grpc"). As of release 2512.0, it <i>must</i> still be specified even with gRPC activated if the client is to connect to the Workflow Server, too.
--------------------	--

This property determines where to get the IOR of the *Content Server* (format: http://<server>:<port>/ior).

- <server> must be the name of the *Content Server* host.
- <port> must be the server's web server HTTP port.

```
repository.use-grpc
```

Default	false
----------------	-------

Description	Activates gRPC for communication with the servers and de-activates CORBA communication. Requires gRPC endpoint target(s) for servers to be set (see "spring.grpc.client.channel.cap.target" and "spring.grpc.client.channel.wf.target"). In release 2512.O, this is not yet effective for communication with the Workflow Server as that does not yet offer gRPC services. Note that this property is _NOT_ a spring boot property but is only effective if set as system property or as environment variable 'REPOSITORY_USEGRPC'.
--------------------	--

```
repository.user
```

Type	String
-------------	--------

Description	Define the user which connects to the <i>Content Server</i> .
--------------------	---

```
repository.workflow.connect
```

Type	Boolean
-------------	---------

Description	Whether to connect the workflow repository. Workflow is disabled by default.
--------------------	--

```
repository.workflow.url
```

Type	String
-------------	--------

Description	The workflow server IOR URL.
--------------------	------------------------------

```
spring.grpc.client.channel.cap.target
```

Default	localhost:9090
----------------	----------------

Description	gRPC endpoint target of a Content Server, e.g., "localhost:40165". Only effective if gRPC is active (see property "repository.use-grpc").
--------------------	---

```
spring.grpc.client.channel.wf.target
```

Default	localhost:9090
----------------	----------------

Description	gRPC endpoint target of the Workflow Server, e.g., "localhost:40365". This property is not effective in release 2512.0 as the Workflow Server does not yet offer gRPC services to clients.
--------------------	--

Table 3.57. UAPI Spring Boot Client Properties

3.13 Cache Properties

cache.capacities	
Type	Map<String, Long>
Description	Number of cache entries per cache class until cache eviction takes place. The keys must match the cache classes as defined by the cache keys. Please refer to javadoc of com.coremedia.cache.CacheKey.
cache.capacities.contenthub.children	
Type	Long
Default	1000
Description	Sets the cache size for the getChildren call in content hub.
cache.capacities.contenthub.object	
Type	Long
Default	1000
Description	Sets the cache size for the content hub objects which includes items as well as folders.
cache.capacities.contenthub.rootfolder	
Type	Long
Default	20
Description	Sets the cache size for the content hub root folders (adapters).
cache.timeout-seconds	
Type	Map<String, Long>

Description	TTL in seconds until certain cache entries are invalidated.
--------------------	---

Table 3.58. Cache Properties

3.14 Web Security Properties

3.14.1 CORS Properties

`cors.allow-credentials-for-url-pattern`

Type	Map<String,Boolean>
Description	Map of whether user credentials are supported, based on URL patterns. Example: <pre>cors.allow-credentials-for-url-pattern[{path\:.*}]=true</pre> See Javadoc for more information on CORS configuration for the delivery applications.

`cors.allow-private-network-for-url-pattern`

Type	Map<String,Boolean>
Description	Map of whether private network access is supported, based on URL patterns. Example: <pre>cors.allow-private-network-for-url-pattern[{path\:.*}]=true</pre> See Javadoc for more information on CORS configuration for the delivery applications.

`cors.allowed-headers-for-url-pattern`

Type	Map<String,List<String>>
Description	Map of headers that a pre-flight request can list as allowed for use during an actual request, based on URL patterns. A header name is not required to be listed if it is one of: Cache-Control, Content-Language, Expires, Last-Modified or Pragma. Example:

```
cors.allowed-headers-for-url-pattern[{path\:.*}]=x-requested-with,x-csrf-token
```

See Javadoc for more information on CORS configuration for the delivery applications.

```
cors.allowed-methods-for-url-pattern
```

Type Map<String, List<String>>

Description Map of HTTP methods to allow, based on URL patterns.

Example:

```
cors.allowed-methods-for-url-pattern[{path\:.*}]=GET,POST,PUT
```

See Javadoc for more information on CORS configuration for the delivery applications.

```
cors.allowed-origin-patterns-for-url-pattern
```

Type Map<String, List<String>>

Description Map of origin patters to allow, based on URL patterns.

Alternative to #allowedOriginsForUrlPattern that supports more flexible origins patterns with "*" anywhere in the host name in addition to port lists.

Example:

```
cors.allowed-origin-patters-for-url-pattern[{path\:.*}]=https://*.domain1.com,https://domain2.com: [8080,8081]
```

See Javadoc for more information on CORS configuration for the delivery applications.

```
cors.allowed-origins-for-url-pattern
```

Type Map<String, List<String>>

Description Map of origins to allow, based on URL patterns.

Example:

```
cors.allowed-origins-for-url-pattern[{path\:.*}]=https://domain1.com,https://domain2.com
```

E.g. in the preview CAE, this property may be configured with the Studio host in order to allow AJAX requests from the Studio to the CAE.

See Javadoc for more information on CORS configuration for the delivery applications.

cors.exposed-headers-for-url-pattern

Type Map<String, List<String>>

Description Map of response headers other than simple headers (i.e. Cache-Control, Content-Language, Content-Type, Expires, Last-Modified or Pragma) that an actual response might have and can be exposed, based on URL patterns.

Example:

```
cors.exposed-headers-for-url-pattern[{path\:.*}]=x-requested-with,x-csrf-token
```

See Javadoc for more information on CORS configuration for the delivery applications.

cors.max-age-for-url-pattern

Type Map<String, Duration>

Description Map of how long, as a duration, the response from a pre-flight request can be cached by clients, based on URL patterns.

Example:

```
cors.max-age-for-url-pattern[{path\:.*}]=3m
```

See Javadoc for more information on CORS configuration for the delivery applications.

Table 3.59. Configuration Properties with Prefix cors

3.15 Plugin Manager Properties

plugins.directories	
Type	List<String>
Description	The directories from where plugins are loaded.
plugins.directory	
Type	String
Description	The directory from where plugins are loaded.
Deprecation	This property has been deprecated and will be removed in a future version. Use <code>directories</code> instead. Reason: Introduction of new list-valued property 'directories'
plugins.required-plugins	
Type	List<String>
Description	A list of ids for plugins that are required. If one of these plugins is missing or could not be started, the application startup will be aborted. This property is only effective when 'plugins.directories' is set.

Table 3.60. Plugin Manager Properties

3.16 Blob Transformer Properties

blobtransformer.enable-builtin-transformation	
Type	Boolean
Default	true
Description	If 'false', the Spring beans for the built-in Java-based image transformation are not instantiated. The setting is for example automatically set to 'false' in a CoreMedia Content Cloud Service instance that is using the cloud-only image transformation service with WebP support. In this case, customizations of the 'BlobTransformer' are not possible and as such it helps preventing illegal configurations. The setting may only be set to 'false', if there is a plugin which provides a 'BlobTransformer'. Otherwise, the application will fail to start with a bean creation error.
com.coremedia.transform.accepted-mime-types	
Type	Collection<jakarta.activation.MimeType>
Description	Only the listed MIME Types are accepted as input for the built-in blob transformer. If empty (default), all image types Java can handle natively are accepted. Some MIME types (such as BMP and TIFF, for example) cannot be correctly handled by the built-in blob transformer.
com.coremedia.transform.blob-cache.base-path	
Type	Path
Description	The path to the transformed image blobs cache. If not set (which is the default), then this cache is deactivated and the results of image transformations are stored using the UAPI cache. It is recommended to set this property to a folder that is not cleared during CAE restart, for example, /var/cache/coremedia/persistent-transformed-blobcache.
com.coremedia.transform.blob-cache.log-interval-seconds	
Type	Duration

Default	600s
Description	Delay for logging metrics of the transformed BLOB cache. The default unit is "seconds".
<code>com.coremedia.transform.blob-cache.size</code>	
Type	<code>org.springframework.util.unit.DataSize</code>
Default	1GB
Description	The maximum allowed size that the transformed image blob cache can occupy on the disk. This is a separate cache where results of blob transformations are stored persistently and can survive CAE restarts. Note that the file system overhead for storing the files does not count towards this value. So the physical space that has to be reserved on the disk for the cache has to be slightly higher than value of this configuration property. If several concurrent threads write large blobs at the same time, the deletion of the folder with the old unused files can be postponed for later, thus this is the second reason why the maximum allowed cache size can grow slightly higher than this configuration property. The size of such deviation depends on the blobs size as well as the amount of parallel threads. The default unit is "bytes".
<code>com.coremedia.transform.memory-guard.safety-factor</code>	
Type	Double
Default	1.2
Description	A safety factor of the memory guard configured for the image transformer. This memory guard checks available memory before performing individual operations on an image. It complements the <code>throttle.safetyFactor</code> in that it checks for memory shortage immediately before operations are performed. Setting this to zero disables memory checks.
<code>com.coremedia.transform.throttle.fallback-factor</code>	
Type	Integer
Default	1

Description	A factor to multiply the blob size with if the size of the in-memory representation needed for transformation cannot be determined. A warning is logged if this happens.
<code>com.coremedia.transform.throttle.hard-limit</code>	
Type	Boolean
Default	true
Description	Specifies whether the configured permits should be considered a hard limit for transformations. Setting this to <i>false</i> allows transformations needing more permits (serialized, one at a time). This would allow larger images to be transformed, but would make OutOfMemoryErrors during transformations more likely.
<code>com.coremedia.transform.throttle.permits</code>	
Type	Integer
Default	0
Description	Capacity of the ThrottlingBlobTransformer, representing the maximum heap memory that can be allocated for blob transformations (in megabytes; conforming to ImageSizePermitStrategy). Default is 0 (zero), which falls back to 1/4 of the total JVM memory size.
<code>com.coremedia.transform.throttle.safety-factor</code>	
Type	Double
Default	3
Description	A safety factor for the ImageSizePermitStrategy to multiply the memory size of an image with. The default number of 3 takes into account that some image operations need to make a copy of the image, plus one copy for the in-memory cache for loaded images.

Table 3.61. Blob Transformer Properties

<code>cache.capacities.com.coremedia.transform.image.java2d.LoadedImageCacheKey</code>
--

Value	Number
Default	100 MB
Description	The size of the loaded image cache in bytes. See imageTransformerInputAdapter in <i>Content Application Developer Manual</i> for more information.

Table 3.62. Further Configuration Properties

3.17 Image Transformation Properties

imagetransformation.default-avif-quality	
Type	Float
Default	0
Description	The default avif quality used for re-encoding. Initially, this is set to zero which suppresses adding the corresponding parameter to the transformation string. Choice of default quality is thus left to the image transformation component.
imagetransformation.default-jpeg-quality	
Type	Float
Default	0.8
Description	The default jpeg quality used for re-encoding.
imagetransformation.default-webp-quality	
Type	Float
Default	0
Description	The default webp quality used for re-encoding. Initially, this is set to zero which suppresses adding the corresponding parameter to the transformation string. Choice of default quality is thus left to the image transformation component.
imagetransformation.dynamic-variants	
Type	Boolean
Default	true
Description	If true, resolve transformations from content, otherwise only programmatically configured transformations are used.

imagetransformation.remove-metadata	
Type	Boolean
Default	true
Description	The fallback value for removeMetadata. Will be effective for transformations where neither the given Breakpoint nor the given Transformation have a remove-Metadata value. If true, the "rm" image operation will be applied to remove the image metadata
imagetransformation.sharpen	
Type	Boolean
Default	true
Description	The fallback value for sharpen. Will be effective for transformations where neither the given Breakpoint nor the given Transformation have a sharpen value. If true, the "usm" image operation will be applied and the image will be sharpened.

Table 3.63. Image Transformation Properties

3.18 Engagement Cloud Integration Properties

The following table lists the configuration properties for the *CoreMedia Engagement Cloud* integration, which is in particular used for the client-side personalization extension.

cmec.auth-connect-timeout	
Type	Duration
Description	The timeout used when connecting to the Engagement Cloud server for the purpose of authentication. Defaults to 10s.
cmec.auth-read-timeout	
Type	Duration
Description	The timeout used when reading the response from the server for the purpose of authentication. Defaults to 30s.
cmec.default-password	
Type	String
Description	The default Engagement Cloud password. It is recommended to use a distinct password for each webcare id.
cmec.default-region	
Type	String
Description	The default Engagement Cloud region.
cmec.default-webcare-id	
Type	String
Description	The default webcare id to use.

`cmec.enabled`**Type** Boolean**Default** false**Description** Whether the Engagement Cloud integration is enabled.`cmec.global-configuration-path`**Type** String**Default** /Settings/Options/Settings/Internal/Engagement Cloud**Description** The path below a site folder for the internal settings content item defining Engagement Cloud parameters. This content item may be used to override parameters specified through the Spring configuration.`cmec.group-mapping.cmcc`**Type** Map<String,String>**Description** A map from an id to the full name of a Content Cloud group, potentially including the domain. The same id is used in the property `cmec.group-mapping.cmec` to derive the name of a group of Engagement Cloud that is assigned to all users that are direct or transitive members of the mapped CC group. The group may contain a domain suffix, separated from the actual group name by an at-sign. Multiple ids may map to the same CC group to assign multiple Engagement Cloud groups to a single CC group. Ids are case-insensitive.`cmec.group-mapping.cmec`**Type** Map<String,String>**Description** A map from an id to the full name of an Engagement Cloud group. The same id is used in the map `cmec.group-mapping.cmcc` to derive the name of a CC group whose direct or transitive members are mapped to a given Engagement Cloud group. Multiple ids may map to the same Engagement Cloud group to assign the groups multiple times. Ids are case-insensitive.`cmec.password-by-site-id`

Type	Map<String,String>
Description	A map from site ids to webcare password. This property takes precedence over cmec.default-password.
<code>cmec.proxy-connect-timeout</code>	
Type	Duration
Description	The timeout used when connecting to the server for proxy calls. Defaults to 10s.
<code>cmec.proxy-read-timeout</code>	
Type	Duration
Description	The timeout used when reading the response from the server for proxy calls. Defaults to 300s.
<code>cmec.region-by-site-id</code>	
Type	Map<String,String>
Description	A map from site ids to Engagement Cloud regions. This property takes precedence over cmec.default-region.
<code>cmec.site-configuration-path</code>	
Type	String
Default	Options/Settings/Internal/Engagement Cloud
Description	The path in the content repository for the settings content item defining Engagement Cloud parameters. This content item may be used to override parameters specified through the Spring configuration.
<code>cmec.uri-by-api-and-region</code>	
Type	Map<String,String>

Description	A map from strings of the form "API.REGION" to base URIs of Engagement Cloud APIs. This map makes it possible to adjust the plugin if the URI of an API changes or if you need to hide the API behind a proxy.
<code>cmec.webcare-id-by-site-id</code>	
Type	<code>Map<String,String></code>
Description	A map from site ids to webcare id. Multiple sites may share the same webcare id. This property take precedence over <code>cmec.default-webcare-id</code> .

Table 3.64. Configuration Properties for the CoreMedia Engagement Cloud Integration

3.19 gRPC Properties

gRPC clients can be configured via Spring gRPC Configuration Properties (`spring.grpc.*`). Find out more about those properties in the [official Spring gRPC documentation](#). Additionally, the following configuration properties can be used for all gRPC Clients.

<code>grpc.client.default-shutdown-grace-period</code>	
Type	Duration
Description	The default duration gRPC client channels will wait for a graceful shutdown. The default value is 30 seconds.
<code>grpc.client.default-timeout</code>	
Type	Duration
Description	The default duration gRPC client services will wait for an answer from the server. If no unit is given the value is interpreted as seconds. The default value is 20 seconds.
<code>grpc.client.service-timeouts</code>	
Type	Map<String,Duration>
Description	Specifies timeouts for gRPC client service names or full method names individually. Any service or method that is present in this mapping, will terminate calls after the specified duration, overriding the default timeout for its calls. A negative value means no timeout. The client will wait indefinitely for an answer from the server. If no duration unit is given the value is interpreted as seconds.
<code>grpc.client.shutdown-grace-periods</code>	
Type	Map<String,Duration>

Description	Specifies the time to wait for the channel to gracefully shutdown. Any channel whose name is present in this mapping, will use the specified duration for its shutdown grace period, overriding the default shutdown grace period. If no duration unit is given the value is interpreted as milliseconds.
--------------------	---

Table 3.65. gRPC Client Properties

4. Encryption Service Setup

CoreMedia Content Cloud supports an encryption service that can be used to encrypt secret strings. This chapter covers the configuration of the mentioned service.

NOTE

For all features using an encryption service, it is required to use the same encryption service type. For example mixing a Java Keystore based encryption service with a plugin based encryption service, which is not Java Keystore based, is not supported.



4.1 Plugin Based Encryption Service Setup

The following table lists the properties that can be used to configure the plugin based encryption service. Plugin specific configuration properties are defined in the plugin documentation. A plugin based encryption service will be provided as a Spring bean.

<code>encryption.plugin-support.enabled</code>	
Type	Boolean
Default	false
Description	Whether loading of encryption service bean from plugins is enabled. Note that encryption service plugins must be 'independent'.

Table 4.1. Plugin Based Encryption Service Properties

Available Encryption Service Plugins

The following plugins are available for the plugin based encryption service. Access to the plugin documentation might be restricted to registered GitHub users.

- [AWS KMS Encryption Plugin](#)

4.2 Java Keystore based encryption service

This encryption service utilizes a single pair of public and private keys to encrypt and decrypt values. The keys are retrieved from a Java keystore located in the file system.

4.2.1 Prerequisites

Before you can use the keystore based encryption service, you have to create a keystore file using the Java `keytool` command. The keystore will contain the keys used for encryption and decryption. On the command prompt type:

```
keytool -genkeypair -keyalg RSA
-keystore <KEYSTORE_FILENAME>
-storepass <KEYSTORE_PASSWORD>
-alias <KEY_ALIAS>
-keypass <KEY_PASSWORD>
```

For secure usage at command line, it is not recommended to provide passwords directly on the command line. Use the secure input by not presenting a value for the password options. The `keytool` will prompt for secure password input. For details, see the documentation of `keytool`. Certain keystore types require identical passwords for the keystore and key. You will get an appropriate warning when generating the key in this case.

NOTE

For the above-mentioned command, the generated PKCS12 KeyStore needs to have the same `KEY_PASSWORD` and `KEYSTORE_PASSWORD`.



The `keytool` will prompt you for your username, organizational unit, organization, city, state/province and country code. This information (which goes into your self-signed certificate) is not relevant for the keystore to work properly. The resulting key/certificate will be valid for 90 days unless you specify another validity duration using the `validity` option. Note that you have to update the stored encrypted values when your keypair changes. Chose your keystore validity carefully.

The next step is to configure the keystore encryption service. This is described in the following sections. The recommended way is to use the Spring Boot configuration properties.

4.2.2 Configuration Properties Setup

The following table lists the properties that can be used to configure the keystore encryption service. They are all mandatory. Following the Configuration Properties Setup approach for the Java Keystore Encryption Service will provide the encryption service as a Spring bean.

encryption.keystore.cipher-transformation-name	
Type	String
Default	RSA
Description	The name of the cipher transformation which should be used for de- & encryption. It is strongly recommended to overwrite the default by setting this property to a more secure transformation with padding.
encryption.keystore.key-alias	
Type	String
Description	Alias of the key inside the key store specified by encryption.keystore.keystore-file-location, that should be used for de- & encryption. The property value must not be blank.
encryption.keystore.key-password	
Type	String
Description	The password of the key, specified by encryption.keystore.key-alias. The property value must not be blank.
encryption.keystore.keystore-file-location	
Type	String
Default	\${user.home}/.cmservices/.keystore
Description	Absolute path of the Java key store. The property value must not be blank.
encryption.keystore.keystore-password	

Type	String
Description	Password of the Java key store specified by encryption.keystore.keystore-file-location. The property value must not be blank.

Table 4.2. Key Store Encryption Service Properties

NOTE

The encryption service Spring bean takes precedence over the static encryption service instance configured using the `CM_KEYSTORE_LOCATION` and `CM_KEYSTORE_PASSWORD_FILE_LOCATION` properties.



4.2.3 Properties File Setup

After creating the keystore, servers and clients need to have access to the keystore credentials. For this, they have to be stored in a Java properties file. Then the servers and clients will be able to access the keystore without prompting for passwords. The password file has to contain the following entries:

```
CM_KEYSTORE_PASSWORD=<KEYSTORE_PASSWORD>
CM_KEY_ALIAS=<KEY_ALIAS>
CM_KEY_PASSWORD=<KEY_PASSWORD>
```

As the password file contains the clear text passwords for your keystore, the file has to be protected from unauthorized access. This could be done by setting reasonable access rights for the file, or by putting it on a removable device.

Cipher transformation: By default the service uses less secure RSA-transformation, which is known to be available on all systems. For enhanced security, it is recommended to switch to an RSA algorithm with padding. You may do so by providing an additional property `CM_CIPHER_TRANSFORMATION` in the password file mentioned above. For available `Cipher` transformations, have a look at your installed security providers. By default, your Java platform should support the following `Cipher` transformations:

- RSA (default and fallback; see below)
- RSA/ECB/PKCS1Padding
- RSA/ECB/OAEPWithSHA-1AndMGF1Padding
- RSA/ECB/OAEPWithSHA-256AndMGF1Padding

Example configuration for different cipher transformation:

```
CM_KEYSTORE_PASSWORD=<KEYSTORE_PASSWORD>
CM_KEY_ALIAS=<KEY_ALIAS>
CM_KEY_PASSWORD=<KEY_PASSWORD>
CM_CIPHER_TRANSFORMATION=RSA/ECB/OAEPWithSHA-256AndMGF1Padding
```

Cipher transformation migration: To ease migrating from the default RSA transformation to a more secure transformation with padding, RSA is always used as a fallback. In case decrypting a password failed for the configured cipher, RSA will be tried as well.

Configuring the location of keystore files: In order to use the keystore with the encryption service, the service needs to know the location of the keystore and keystore password files. For this, you have the following two options:

- By default, the service expects
 - the keystore file under the path `${user.home}/.cmservices/.keystore`
 - and the password file under `${user.home}/.cmservices/.keystore.properties`
- If you want to store the files under different paths, you have to provide the following two system properties:
 - `CM_KEYSTORE_LOCATION`: location of the keystore file
 - `CM_KEYSTORE_PASSWORD_FILE_LOCATION`: location of the password file

NOTE

The encryption service Spring bean takes precedence over the static encryption service instance configured using the `CM_KEYSTORE_LOCATION` and `CM_KEYSTORE_PASSWORD_FILE_LOCATION` properties.



Index

B

build time configuration, 26

C

configuration

CapList, 162

My Edited Content, 162

Notifications, 130

Projects/To-Dos, 134

User Changes Application, 143

Workflow Lists, 162

configure

jib-ownership-extension, 16

mount files, 30

tools, 30

container

build, 14

build time configuration, 26

configuration, 26

mount file system, 28

runtime configuration, 27

start time configuration, 26

user-defined network, 29

D

Docker

setup, 13

G

GIGC, 18

GC

Garbage Collector, 18

Google Jib

jib-ownership-extension, 15

Spring, 14

I

images

arm, 22

Distroless, 21

java-application-base, 17

management-tools, 28

interactive shell, 29

J

jib-ownership-extension

configuration, 16

K

Kubernetes, 13

M

management tools, 28

P

Parallel GC, 18

properties, 33

Blob Transformer, 246

Cache, 240

CAE, 34

CapList, 162

Commerce Hub, 161

Content Server, 47

Elastic Social, 168

Engagement Cloud, 252

Headless Server, 85

Image Transformation, 250

My Edited Content, 162

Notifications, 130

Plugin Manager, 245

Projects/To-Dos, 134

relaxed binding, 33

Spring Boot, 33

Studio, 103

Unified API, 233

Web Security, 242

Workflow Lists, 162

R

runtime configuration, 27

S

- start time configuration, 26
- Swarm, 13

T

- tools
 - configuration, 30
 - logging, 29
 - start, 28
 - start directly, 30
 - start interactively, 28

U

- User Changes Application
 - configuration, 143