

# NetarchiveSuite System Design

Printer friendly version

**Contents**

1. Introduction
2. Overall Systems Design
  1. Access (Viewerproxy)
  2. Archive
  3. Common
  4. Deploy
  5. Harvester
  6. Monitor
  7. Wayback
3. Specific Design in the Code
  1. Settings
  2. JMS Channels
    1. Placement in NetarchiveSuite software
    2. Channel Behavior
    3. Naming Conventions
    4. Design Notes
  3. Localization
  4. JSP
    1. The SiteSection system
    2. Processing updates
    3. I18n
  5. Pluggable parts
    1. How pluggability works
    2. RemoteFile
    3. JMSConnection
    4. ArcRepositoryClient
    5. IndexClient
    6. DBSpecifics
    7. Notifications
    8. HeritrixController
    9. ActiveBitPreservation
  6. XML handling by Deploy
4. Detailed Archive Design Description
  1. Archive Overview
    1. Architecture
    2. Repository State Machine
    3. Extra functionality
  2. Indexes and caching
  3. CrawlLogIndexCache
5. Detailed Common Design Description
6. Detailed Deploy Design Description
7. Detailed Harvester Design Description
  1. Database
    1. Database overview
    2. Table and Column Descriptions
  2. Harvesting roundtrip
    1. Initial steps
    2. Creating a selective harvest
    3. Scheduling and splitting
    4. Talking to the harvesters
    5. Harvest setup
    6. Running Heritrix
    7. Creating metadata
    8. Uploading
    9. Storing harvest statistics
    10. Old jobs
    11. Deduplication in NetarchiveSuite
8. Detailed Monitor Design Description

# 1. Introduction

edit

*This System Design description is still undergoing restructuring. It is intended that those parts which are complete should be accurate.*

This document only describes the underlying design of the NetarchiveSuite software, i.e. it does not describe how to install, run, or use NetarchiveSuite, for that see the Installation Manual and the User Manual.

The first section gives an overview, and the remainder of the document gives more details about the design.

The reader is expected to be familiar with Java programming and have an understanding of the core issues involved in large-scale web harvesting. Having used Heritrix before is a definite plus, and an elementary understanding of SQL databases is required for some parts.

The code is available in the downloaded package (see Release Overview) or from our [subversion repository](#).

## 2. Overall Systems Design

edit

This section includes an overall description of the NetarchiveSuite modules. Additional information can be found in the Overview document.

There are seven modules in the NetarchiveSuite software. This section gives an overview of what each module contains. All Java sourcefiles are found in the `SRC` directory, and all packages start with `dk.netarkivet`. Units tests are similarly arranged, but under `tests` instead of `SRC`. The web interface definitions are found in the `webpages` directory. The `lib` directory contains all the libraries necessary to compile and run the code.

More detailed descriptions are given later in this document.

### 2.1. Access (Viewerproxy)

The `dk.netarkivet.viewerproxy` package implements a simple access client to the archived data, based on web-page proxying. For more details please refer to Detailed Access Design.

### 2.2. Archive

The `dk.netarkivet.archive` package and its subpackages provide redundant, distributed storage primarily for ARC files as well as Lucene indexing of same. The `arcrepository` subpackage contains the logic of keeping multiple bit archives synchronized.

The **bitarchive** subpackage contains the application that stores the actual files and manages access to them. The **indexserver** subpackage handles merging CDX files and crawl.log files into a Lucene index used for deduplication and for viewerproxy access. The **checksum** subpackage contains the checksum replica code. For more details please refer to Detailed Archive Design

## 2.3. Common

The **dk.netarkivet.common** package and its subpackages provide module-neutral code partly of a generic nature, partly specific to NetarchiveSuite, e.g. settings and channels. For more details please refer to Detailed Common Design

## 2.4. Deploy

The **dk.netarkivet.deploy** module contains software for installing NetarchiveSuite on multiple machines. This module is only used in the deployment phase. For more details please refer to Detailed Deploy Design

## 2.5. Harvester

The **dk.netarkivet.harvester** package and its subpackages handle the definition and execution of harvests. Its main parts are the database containing the harvest definitions (the **datamodel** subpackage), the webinterface that the user can access the database with, the **scheduler** subpackage which handles scheduling and splitting into jobs, and the **harvesting** subpackage which encapsulates running Heritrix and sending the results off to the archive. For more details please refer to Detailed Harvester Design

## 2.6. Monitor

The **dk.netarkivet.monitor** package provides web-access to JMX-packaged information from all NetarchiveSuite applications. For more details please refer to Detailed Monitor Design

## 2.7. Wayback

The **dk.netarkivet.wayback** package provides tools for integrating NetarchiveSuite with the open-source wayback machine for browsing webarchives. These are described in the Additional Tools Manual.

# Specific Design in the Code

*The following is to be split between the this document (incooperated in the below detailed module design descriptions) and the Coding Guidelines document, so that coding guidelines of how to code according to this design is put into the Coding Guidelines, while description of how the design works*

*is put into the this document.*

edit

## 2.8. Settings

*Also include relevant parts to the design decision input document (is now implemented) settings structure*

Until release 3.8, all our settings were defined in the class `dk.netarkivet.common.Settings`, and all access to the settings was done using static methods in the same class. Furthermore, you needed to have all settings defined in a `settings.xml`.

As of version 3.8, the settings are read using methods in the new class `dk.netarkivet.common.utils.Settings`. The declaration of the settings themselves, however, have been moved to new settings classes in each of the modules. Every module except for the Deploy module now have its own Settings class and a `settings.xml` file containing the default values for the module settings.

In addition to that, every plug-in is intended to declare its own settings inside itself, and be associated with an xml-file containing the default values of these settings placed in the same directory as the plugin itself. For more information, please refer to the Configuration Basics description in the Configuration Manual.

Associated with most of the NetarchiveSuite plug-ins, there are also factory classes that hides the complexity behind selecting the correct plugin according to the chosen settings. The names of these classes all ends on *Factory*, e.g. `JMSConnectionFactory`, `RemoteFileFactory`.

Almost all configuration of NetarchiveSuite is done through the main module `Settings` classes as for example the `dk.netarkivet.common.Settings` class. It provides a simple interface to the `settings.xml` file as well as definitions of all current configuration default settings. The `settings.xml` file itself is an XML file with a structure reminiscent of the package structure.

Settings are referred to inside NetarchiveSuite by their path in the XML structure. For instance, the `storeRetries` setting in `arcrepositoryClient` under `common` is referred to with the string `"settings.common.arcrepositoryClient.storeRetries"`. However, to avoid typos, each known setting has its path defined as a String constant in the `Settings` class, which is used throughout the code (also deploy in checks of whether the settings are known). The `Settings` class file also includes description of the setting in their javadoc.

To add a new general setting, the following steps need to be taken:

1. The `Settings` class should get a definition for the path of the setting. This String, although a constant, **must not be declared final** since settings are initialised by a static initialiser in the class
2. Javadoc for the definition must including the path name of the setting as well as description of the

setting.

3. all default `settings.xml` files must be updated (including those in the unit tests).
4. `examples/settings_example.xml` must be updated.

Note that there are no XML Schema to be updated, because the use of default settings means that a settings file does not need to be there, and in the case of settings for plug-ins we generally do not know which setting will be used. Note also the Configuration Manual includes a description of how deploy can be used to validate whether the settings in a settings file are valid (not necessarily exhaustive!).

## 2.9. JMS Channels

edit

### 2.9.1. Placement in NetarchiveSuite software

Every channel is named after the set of applications instances that are expected to receive messages sent to it. All channel names are constructed privately in the `dk.netarkivet.distribute.ChannelID` class. To get a channel, you must use one of the public methods in `dk.netarkivet.distribute.Channels`.

### 2.9.2. Channel Behavior

There are used to types of channels:

- **Queue** where only one listener takes a message from the queue. For example a queue where request for doing a harvest is only received by one harvester.
- **Topic** where all listeners takes a copy of the message. For example a batch job which has to be executed by all bitarchive applications.

The type of channel is not affecting the channel name directly. It is indicated by the Channel Prefix, since only channel names starting with **ALL** are topics, while the rest are queues.

The channel type for different queues is given in a table in the next section.

### 2.9.3. Naming Conventions

The structure of channel names are as follows:

`<Environment Name>_<Channel Prefix>_<Replica Annotation>[_<Machine>][_<Application Instance>]`

- **Environment Name** value must be the same for all channels on all JVMs belonging to a single NetarchiveSuite installation. It identifies the environment for the queue, e.g.: "DEV", "TEST", "RELEASESTEST", "CLOVER", "PROD" or initials on developer running personal environment for e.g. sanity tests. It is read from the setting `settings.common.environmentName` of the NetarchiveSuite installation.
- **Channel Prefix** is constructed by a convention for channel behavior together with denotation of the

concerned application(s). For example **THE\_SCHED** uses the convention **THE** for **SCHED** denoting the scheduler of the harvester.

The conventions for channel behavior are:

- **THE** - communicate with singleton in the system.
- **THIS** - communicate with one uniquely identified instance of a number of distributed listeners.
- **ANY** - has all instances listening to the same queue.
- **ALL** - is used for topics only, i.e. topics are sent to *all* listeners on the channel.
- **Use Replica Id** will result in a channel name with the replica identifier in question (which must match possible replicas in settings). In case of no use of replica id the **COMMON** will be used.
- **Use IP Node** will result in a channel name with the IP address of the machine where the application in question is placed. It is read directly from the running system.

In case of no use of IP Node, nothing will be written.

- **Use Application Instance Id** is used, if only one process is expected to listen to the queue. It will result in a channel name with an identification of the individual application instance. It consists of the application abbreviation (uppercase letters of application name) and the application instance identifier from the settings. It is read from the common settings:

`settings.common.applicationName` and  
`settings.common.applicationInstanceId`.

In case of no use of Application Instance Id, nothing will be written.

- **Channel Type** is only here to complete the picture. Please refer to the section on channel behavior.

Channel names are constructed as described in the below table (columns described after the table).

<i>Channel Prefix</i>	<i>Use Replica Id (for Replica Annotation)</i>	<i>Use IP Node (for Machine)</i>	<i>Use Application Instance Identification</i>	<i>Channel Type</i>	<i>Example</i>
THE_SCHED	No	No	No	Queue	PROD_THE_SCHE
ANY_HIGHPRIORITY_HACO	No	No	No	Queue	PROD_ANY_HIGH
ANY_LOWPRIORITY_HACO	No	No	No	Queue	PROD_ANY_LOW
THIS_REPOS_CLIENT	No	Yes	Yes	Queue	PROD_THIS_REPC
THE_REPOS	No	No	No	Queue	PROD_THE_REPO
THE_BAMON	Yes	No	No	Queue	PROD_THE_BAMC
ALL_BA	Yes	No	No	Topic	PROD_ALL_BA_T
ANY_BA	Yes	No	No	Queue	PROD_ANY_BA_T
THIS_INDEX_CLIENT	No	Yes	Yes	Queue	PROD_THIS_INDE
INDEX_SERVER	No	No	No	Queue	PROD_INDEX_SEI
MONITOR	No	No	No	Queue	PROD_MONITOR_
ERROR	No	No	No	Queue	PROD_ERROR_CO

THE_CR	Yes	No	No	Queue	PROD_THE_CR_TI
--------	-----	----	----	-------	----------------

The examples are using values


- Environment name **PROD**
- Possible replica identifiers **ONE** or **TWO**
- IP on machine **130.226.258.7**
- Application instances
  - **HCA\_HIGH** (for HarvestControllerApplication with instance id "**HIGH**" )
  - **ISA** (for IndexServerApplication with instance id " " )



### 2.9.4. Design Notes

Note that creation of channel names for the ANY xxx\_HACO-queues are designed in a way so extension to more priorities is easy.

## 2.10. Localization

edit

The NetarchiveSuite web pages are internationalized, that is they are ready to be translated into other languages. The default distribution contains a default (English) version and Danish, Italian, French and German versions, but adding a new language does not take any coding. All translatable strings are collected in five  resource bundles, one for each of the five main modules mentioned above. The default translation files are `src/dk/netarkivet/common/Translations.properties`, `src/dk/netarkivet/archive/Translations.properties`, `src/dk/netarkivet/harvester/Translations.properties`, `src/dk/netarkivet/viewerproxy/Translations.properties`, and `src/dk/netarkivet/monitor/Translations.properties`.

To translate to a new language, first copy each of these files to a file in the same directory, but with `_XX` after `Translations`, where `XX` is the  Unicode language code for the language you're going to translate into, e.g. if you're translating into Limburgish, use `Translations_li.properties`. If you're translating into a language that has different versions for different countries, you may need to use `_XX_YY`, where `XX` is the language code and `YY` is the  ISO country code, e.g.

`Translations_fr_CA.properties` for Canadian French. Then edit each of the new files to have your translation instead of the English translation for each line. Most of the important syntax should be evident from the original, but for details consult the XXX. According to the Java documentation (specifically the Javadoc of the `Properties` class) resource bundles should use iso-8859-1 with escaped Unicode for all other characters. It is good practice to use escaped Unicode for *all* non-ASCII characters as this results in files which are more-easily shared between different text-editing environments.

The translation has not been done throughout the code, only in the web-related parts. Thus log messages and unexpected error messages are in English and cannot be translated through the resource bundles.

## 2.11. JSP



edit

The webpages in NetarchiveSuite are written using JSP (Java Server Pages) with Apache I18N Taglib for internationalization. To support a unified look across pages from different modules, we have divided the pages into SiteSections as described in the next section. Any processing of requests happens in Java code before the web page is displayed, such that update errors can be handled with a uniform error page. Internationalization is primarily done with the taglib tags `<fmt:message>`, `<fmt:date>` etc.

The main feature of JSP is that ordinary Java (not JavaScript) can be used at server-side to generate HTML. The special tags `<%...%>` indicate a piece of Java code to run, while the tags `<%=...>` indicates a Java expression to run whose value will be inserted (as is, see escape mechanisms below) in the HTML. While it is possible to output to HTML from Java code using `out.println()`, it is discouraged as it a) is confusing to read, and b) does not allow for using taglibs for internationalization.

We use a number of standard methods defined in `dk.netarkivet.common.webinterface.HTMLUtils`. Of particular note are the following methods:

#### **generateHeader()**

This method takes a `PageContext` and generates the full header part of the HTML page, including the starting `<body>` tag. It should always be used to create the header, as it also creates the menu and language selection links. After this method has been called, redirection or forwarding is no longer possible, so any processing that can cause fatal errors must be done before calling `generateHeader()`. The title of the page is taken from the `SiteSection` information based on the URL used in the request.

#### **generateFooter()**

This closes the HTML page and should be called as the last thing on any JSP page.

#### **setUTF8()**

This method must be called at the very start of the JSP page to ensure that reading from the request is handled in UTF-8.

#### **encode()**

This encodes any character that is not legal to have in a **URL**. It should be used whenever an unknown string (or a string with known illegal characters) is made part of a URL. Note that it is not idempotent, calling it twice on a string is likely to create a mess.

#### **escapeHTML()**

This escapes any character that has special meaning in **HTML** (such as `<` or `&`). It should be used any time a unknown string (or a string with known special characters) is being put into HTML. Note that it is **not** idempotent: If you escape something twice, you get a horrible-looking mess.

#### **encodeAndEscape()**

This method combines `encode()` and `escapeHTML()` in one, which is useful when you're putting unknown strings directly into URLs in HTML.

### **2.11.1. The SiteSection system**

Each part of the web site (as identified by the top-level menu items on the left side) is defined by one subclass of the `SiteSection` class. These sections are loaded through the `<siteSection>` settings, each of

which connect one SiteSection class with its WAR file and the path it will appear under in the URL.

Each SiteSection subclass defines the name used in the left-hand menu, the prefix of all its pages, the number of pages visible in the left-hand menu when within this section, a suffix and title for each page in the section (including hidden pages), the directory that the section should be deployed under, and a resource bundle name for translations. Furthermore, the SiteSections have hooks for code that needs to be run on deployment and undeployment. If you want to add a new page to the section, you will only need to add a new line to the list of pages with a unique (within the SiteSection) suffix and a key for the page title, plus a default translation in the corresponding Translation.properties file. If you want it to appear in the left-hand menu, update the number of visible pages to n+1 and put your new pages as one of the first n+1 lines.

This is an example of what a simple SiteSection can look like. Note that only the first two pages from the list have entries in the left-hand menu. This class does no special initialisation and shutdown.

```
public HistorySiteSection() {
    super("sitesection;history", "Harveststatus", 2,
        new String[][]{
            {"alljobs", "pagetitle;all.jobs"},
            {"perdomain", "pagetitle;all.jobs.per.domain"},
            {"perhd", "pagetitle;all.jobs.per.harvestdefinition"},
            {"perharvestrun", "pagetitle;all.jobs.per.harvestrun"},
            {"jobdetails", "pagetitle;details.for.job"}
        }, "History",
        dk.netarkivet.harvester.Constants.TRANSLATIONS_BUNDLE);
}

public void initialize() {}
public void close() {}
```

### 2.11.2. Processing updates

Some JSP sites cause updates when posted with specific parameters. Such parameters should always be specified in the beginning of the JSP file. All updates of underlying file systems, databases etc should happen before `generateHeader()` is called, so processing errors can be properly redirected. The preferred way to process updates is to create a method `processRequest()` in a class corresponding to the web page, but under the `webinterface` package of the corresponding module. This method should take the `pageContext` and `I18N` parameters from the JSP page, together they contain all the information needed from there.

In case of processing errors, the processing method should call `HTMLUtils.forwardToErrorPage()` and then throw a `ForwardedToErrorPage` exception. The JSP code should always enclose the `processRequest()` call in a try-catch block and return immediately if `ForwardedToErrorPage` is thrown. This mechanism should be used for "expected" errors, mainly illegal parameters. Errors of the "this can never happen" nature should just cause normal exceptions. Like in other code, the `processRequest()` method should check its parameters, but it should also check the parameters posted in the request to check that they conform to the requirements. Some methods for that purpose can be found in `HTMLUtils`.

### 2.11.3. I18n

We use the Apache I18n taglib for most internationalization on the web pages. This means that instead of

writing different versions of a web page for different languages, we replace all natural language parts of the page with special formatting instructions. These are then used to look up translations to the language in effect in translation resource bundles.

Normal strings can be handled with the `<fmt:message/>` tag. If variable parameters are introduced, such as object names or domain names, they can be passed as parameters using `<fmt:message key="translation.key"><fmt:param value="<%myVal>"/></fmt:message>`. Note that while the message retrieved for the key gets any HTML-specific characters escaped, the values do not and should be manually escaped. It is possible if necessary to pass HTML as parameters.

Dates should in general be entered using `<fmt:formatDate type="both">`, though a few places use a more explicit handling of formats. This lets the date be expressed in the native language's favorite style.

Integers and longs are handled in Java properties files with `{0, number, integer}` or `{0, number, long}` as described in [MessageFormat](#). In JSP files we handled integer/long to strings with `HTMLUtils.localiseLong(long, PageContext)` or `HTMLUtil.localiseLong(long, Locale)`.

Note the boilerplate code at the start of every page that defines output encoding, taglib usage, translation bundle, and a general-purpose I18N object. It is important that the translation bundles from the `Constants` class for the module you're in is used, or incomprehensible errors will occur.

```

pageEncoding="UTF-8"
%><%@taglib uri="http://java.sun.com/jsp/jstl/fmt" prefix="fmt"
%><fmt:setLocale value="<%=HTMLUtils.getLocale(request)%>" scope="page"
/><fmt:setBundle scope="page" basename="
<%=dk.netarkivet.archive.Constants.TRANSLATIONS_BUNDLE%>"/><%!
    private static final I18n I18N
        = new I18n(dk.netarkivet.archive.Constants.TRANSLATIONS_BUNDLE);
%><%

```

## 2.12. Pluggable parts

edit

Some points in NetarchiveSuite can be swapped out for other implementations, in a way similar to what Heritrix uses.

*Also include relevant parts of design document that was basis for implementation of plug-ins*

[To be introduced more]

### 2.12.1. How pluggability works

Factories [To be described more]

...request for suggestions on pluggability areas [To be described more]

### 2.12.2. RemoteFile

The RemoteFile interface defines how large chunks of data are transferred between machines in a NetarchiveSuite installation. This is necessary because JMS has a relatively low limit on the size of

messages, well below the several hundred megabytes to over a gigabyte that is easily stored in an ARC file. There are two current implementations available in the default distribution:

- **FTPRemoteFile** - this implementation uses one or more FTP servers for transfer. While this requires more setup and causes extra copying of data, the method has the advantage of allowing more protective network configurations.
- **HTTPRemoteFile** - this implementation uses an embedded HTTP server in each application that wants to send a RemoteFile. Additionally, it will detect when a file transfer happens within the same machine and use local copying or renaming as applicable. For single-machine installations, this is the implementation to use. In a multi-machine installation, it does require that all machines that can send RemoteFile objects (including the bitarchive machines) must have a port accessible from the rest of the system, which may go against security policies.
- **HTTPSRemoteFile** - This is an extension of HTTPRemoteFile that ensures that the communication is secure and encrypted. It is implemented with a shared certificate scheme, and only clients with access to the certificate will be able to contact the embedded HTTP server.

All three implementations will detect when 0 bytes are to be transferred and avoid creating unnecessary file in this case.

The RemoteFile interface is defined by the Java interface *dk.netarkivet.common.distribute.RemoteFile*:

```
/**
 * RemoteFile: Interface for encapsulating remote files.
 * Enables us to transmit large files between system components situated
 * on different machines.
 * Our current JMS broker(s) does not allow large message
 * (i.e. messages > 70 MB).
 */
public interface RemoteFile extends Serializable {
    /**
     * Copy remotefile to local disk storage.
     * Used by the data recipient
     * @param destFile local File
     * @throws IOException on communication trouble.
     * @throws IllegalArgumentException on null parameter or non-writable file
     */
    public void copyTo(File destFile);
    /**
     * Write the contents of this remote file to an output stream.
     * @param out OutputStream that the data will be written to. This stream
     * will not be closed by this operation.
     * @throws IOException If append operation fails
     * @throws IllegalArgumentException on null parameter
     */
    public void appendTo(OutputStream out);
    /**
     * Get an inputstream that contains the data transferred in this RemoteFile.
     * @return A stream object with the data in the object. Note that the
     * close() method of this may throw exceptions if e.g. a transmission error
     * is detected.
     * @throws IOException on communication trouble.
     */
    public InputStream getInputStream();
    /**
     * Return the file name.
     * @return the file name
     */
}
```

```

public String getName();
/**
 * Returns a MD5 Checksum on the file. May return null, if checksums not
 * supported for this operation.
 * @return MD5 checksum
 */
public String getChecksum();
/**
 * Cleanup this remote file. The file is invalid after this.
 */
public void cleanup();
/** Returns the total size of the remote file.
 * @return Size of the remote file.
 */
public long getSize();
}

```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/common/distribute/RemoteFile.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/common/distribute/RemoteFile.java?root=netarchivesuite) (requires access to our code repository).

### 2.12.3. JMSConnection

The JMSConnection provides access to a specific JMS connection. The default NetarchiveSuite distribution contains only one implementation, namely JMSConnectionSunMQ which uses Sun's OpenMQ. We recommend using this implementation, as other implementations have previously been found to violate some assumptions that NetarchiveSuite depends on.

The JMSConnection interface is defined by the abstract Java class *dk.netarkivet.common.distribute.JMSConnection*

Implementations of this interface needs to implement the three abstract methods in this interface

```

ConnectionFactory getConnectionFactory() throws JMSEException;
Destination getDestination(String destinationName) throws JMSEException;
void onException(JMSEException e);

```

There are more fully explained in the following transcript of the interface.

```

/**
 * Handles the communication with a JMS broker. Note on Thread-safety: the
 * methods and fields of JMSConnection are not accessed by multiple threads
 * (though JMSConnection itself creates threads). Thus no synchronization is
 * needed on methods and fields of JMSConnection. A shutdown hook is also added,
 * which closes the connection. Class JMSConnection is now also a
 * exceptionhandler for the JMS Connections
 */
public abstract class JMSConnection implements ExceptionListener, CleanupIF {
    /** The log. */
    private static final Log log = LogFactory.getLog(JMSConnection.class);
    /**
     * Separator used in the consumer key. Separates the ChannelName from the
     * MessageListener.toString().
     */
    protected static final String CONSUMER_KEY_SEPARATOR = "##";
    /** The number to times to (re)try whenever a JMSEException is thrown. */
    protected static final int JMS_MAXTRIES = 3;

```

```
/** The JMS Connection. */
protected Connection connection;
/**
 * The Session handling messages sent to / received from the NetarchiveSuite
 * queues and topics.
 */
protected Session session;
/** Map for caching message producers. */
protected final Map<String, MessageProducer> producers
    = Collections.synchronizedMap(
        new HashMap<String, MessageProducer>());
/**
 * Map for caching message consumers (topic-subscribers and
 * queue-receivers).
 */
protected final Map<String, MessageConsumer> consumers
    = Collections.synchronizedMap(
        new HashMap<String, MessageConsumer>());
/**
 * Map for caching message listeners (topic-subscribers and
 * queue-receivers).
 */
protected final Map<String, MessageListener> listeners
    = Collections.synchronizedMap(
        new HashMap<String, MessageListener>());
/**
 * Lock for the connection. Locked for read on adding/removing listeners and
 * sending messages. Locked for write when connection, releasing and
 * reconnecting.
 */
protected final ReentrantReadWriteLock connectionLock
    = new ReentrantReadWriteLock();
/** Shutdown hook that closes the JMS connection. */
protected Thread closeHook;
/**
 * Singleton pattern is be used for this class. This is the one and only
 * instance.
 */
protected static JMSConnection instance;
/**
 * Should be implemented according to a specific JMS broker.
 *
 * @return QueueConnectionFactory
 *
 * @throws JMSEException If unable to get QueueConnectionFactory
 */
protected abstract ConnectionFactory getConnectionFactory()
    throws JMSEException;
/**
 * Should be implemented according to a specific JMS broker.
 *
 * @param destinationName the name of the wanted Queue
 *
 * @return The destination. Note that the implementation should make sure
 *         that this is a Queue or a Topic, as required by the
 *         NetarchiveSuite design. {@link Channels#isTopic(String)}
 *
 * @throws JMSEException If unable to get a destination.
 */
protected abstract Destination getDestination(String destinationName)
    throws JMSEException;
/**
 * Exceptionhandler for the JMSConnection. Implemented according to a
```

```
* specific JMS broker. Should try to reconnect if at all possible.
*
* @param e a JMSEException
*/
public abstract void onException(JMSEException e);
/** Class constructor. */
protected JMSConnection() {
}
/**
* Initializes the JMS connection. Creates and starts connection and
* session. Adds a shutdown hook that closes down JMSConnection. Adds this
* object as ExceptionListener for the connection.
*
* @throws IOFailure if initialization fails.
*/
protected void initConnection() throws IOFailure {
    log.debug("Initializing a JMS connection " + getClass().getName());
    connectionLock.writeLock().lock();
    try {
        int tries = 0;
        JMSEException lastException = null;
        boolean operationSuccessful = false;
        while (!operationSuccessful && tries < JMS_MAXTRIES) {
            tries++;
            try {
                establishConnectionAndSession();
                operationSuccessful = true;
            } catch (JMSEException e) {
                closeConnection();
                log.debug("Connect failed (try " + tries + ")", e);
                lastException = e;
                if (tries < JMS_MAXTRIES) {
                    log.debug("Will sleep a while before trying to
                        + " connect again");
                    TimeUtils.exponentialBackoffSleep(tries,
                        Calendar.MINUTE);
                }
            }
        }
        if (!operationSuccessful) {
            log.warn("Could not initialize JMS connection "
                + getClass(), lastException);
            cleanup();
            throw new IOFailure("Could not initialize JMS connection "
                + getClass(), lastException);
        }
        closeHook = new CleanupHook(this);
        Runtime.getRuntime().addShutdownHook(closeHook);
    } finally {
        connectionLock.writeLock().unlock();
    }
}
/**
* Submit an object to the destination queue. This method cannot be
* overridden. Override the method sendMessage to change functionality.
*
* @param msg The NetarkivetMessage to send to the destination queue (null
* not allowed)
*
* @throws ArgumentNotValid if nMsg is null.
* @throws IOFailure if the operation failed.
*/
public final void send(NetarkivetMessage msg) {
```

```
        ArgumentNotValid.checkNotNull(msg, "msg");
        log.trace("Sending message (" + msg.toString() + ") to " + msg.getTo());
        sendMessage(msg, msg.getTo());
    }
    /**
     * Sends a message msg to the channel defined by the parameter to - NOT the
     * channel defined in the message.
     *
     * @param msg Message to be sent
     * @param to The destination channel
     */
    public final void resend(NetarkivetMessage msg, ChannelID to) {
        ArgumentNotValid.checkNotNull(msg, "msg");
        ArgumentNotValid.checkNotNull(to, "to");
        log.trace("Resending message (" + msg.toString() + ") to "
            + to.getName());
        sendMessage(msg, to);
    }
    /**
     * Submit an object to the reply queue.
     *
     * @param msg The NetarkivetMessage to send to the reply queue (null not
     *         allowed)
     *
     * @throws ArgumentNotValid if nMsg is null.
     * @throws PermissionDenied if message nMsg has not been sent yet.
     * @throws IOFailure if unable to reply.
     */
    public final void reply(NetarkivetMessage msg) {
        ArgumentNotValid.checkNotNull(msg, "msg");
        log.trace("Reply on message (" + msg.toString() + ") to "
            + msg.getReplyTo().getName());
        if (!msg.hasBeenSent()) {
            throw new PermissionDenied("Message has not been sent yet");
        }
        sendMessage(msg, msg.getReplyTo());
    }
    /**
     * Method adds a listener to the given queue or topic.
     *
     * @param mq the messagequeue to listen to
     * @param ml the messagelistener
     *
     * @throws IOFailure if the operation failed.
     */
    public void setListener(ChannelID mq, MessageListener ml) throws IOFailure {
        ArgumentNotValid.checkNotNull(mq, "ChannelID mq");
        ArgumentNotValid.checkNotNull(ml, "MessageListener ml");
        setListener(mq.getName(), ml);
    }
    /**
     * Removes the specified MessageListener from the given queue or topic.
     *
     * @param mq the given queue or topic
     * @param ml a messagelistener
     *
     * @throws IOFailure On network trouble
     */
    public void removeListener(ChannelID mq, MessageListener ml)
        throws IOFailure {
        ArgumentNotValid.checkNotNull(mq, "ChannelID mq");
        ArgumentNotValid.checkNotNull(ml, "MessageListener ml");
        removeListener(ml, mq.getName());
    }
}
```



```
}
/**
 * Clean up. Remove close connection, remove shutdown hook and null the
 * instance.
 */
public void cleanup() {
    connectionLock.writeLock().lock();
    try {
        //Remove shutdown hook
        log.info("Starting cleanup");
        try {
            if (closeHook != null) {
                Runtime.getRuntime().removeShutdownHook(closeHook);
            }
        } catch (IllegalStateException e) {
            //Okay, it just means we are already shutting down.
        }
        closeHook = null;
        //Close session
        closeConnection();
        //Clear list of listeners
        listeners.clear();
        instance = null;
        log.info("Cleanup finished");
    } finally {
        connectionLock.writeLock().unlock();
    }
}
/**
 * Close connection, session and listeners. Will ignore trouble, and simply
 * log it.
 */
private void closeConnection() {
    // Close terminates all pending message received on the
    // connection's session's consumers.
    if (connection != null) { // close connection
        try {
            connection.close();
        } catch (JMSEException e) {
            // Just ignore it
            log.warn("Error closing JMS Connection.", e);
        }
    }
    connection = null;
    session = null;
    consumers.clear();
    producers.clear();
}
/**
 * Unwraps a NetarkivetMessage from an ObjectMessage.
 *
 * @param msg a javax.jms.ObjectMessage
 *
 * @return a NetarkivetMessage
 *
 * @throws ArgumentNotValid when msg in valid or format of JMS Object
 *         message is invalid
 */
public static NetarkivetMessage unpack(Message msg)
    throws ArgumentNotValid {
    ArgumentNotValid.checkNotNull(msg, "msg");
    ObjectMessage objMsg;
    try {
```

```

        objMsg = (ObjectMessage) msg;
    } catch (ClassCastException e) {
        log.warn("Invalid message type: " + msg.getClass());
        throw new ArgumentNotValid("Invalid message type: "
            + msg.getClass());
    }
    NetarkivetMessage netMsg;
    String classname = "Unknown class"; // for error reporting purposes
    try {
        classname = objMsg.getObject().getClass().getName();
        netMsg = (NetarkivetMessage) objMsg.getObject();
        // Note: Id is only updated if the message does not already have an
        // id. On unpack, this means the first time the message is received.
        netMsg.updateId(msg.getJMSMessageID());
    } catch (ClassCastException e) {
        log.warn("Invalid message type: " + classname, e);
        throw new ArgumentNotValid("Invalid message type: " + classname, e);
    } catch (Exception e) {
        String message = "Message invalid. Unable to unpack "
            + "message: " + classname;
        log.warn(message, e);
        throw new ArgumentNotValid(message, e);
    }
    log.trace("Unpacked message '" + netMsg + "'");
    return netMsg;
}
/**
 * Submit an ObjectMessage to the destination channel.
 *
 * @param nMsg the NetarkivetMessage to be wrapped and send as an
 *             ObjectMessage
 * @param to   the destination channel
 *
 * @throws IOException if message failed to be sent.
 */
protected void sendMessage(NetarkivetMessage nMsg, ChannelID to)
    throws IOException {
    Exception lastException = null;
    boolean operationSuccessful = false;
    int tries = 0;
    while (!operationSuccessful && tries < JMS_MAXTRIES) {
        tries++;
        try {
            doSend(nMsg, to);
            operationSuccessful = true;
        } catch (JMSEException e) {
            log.debug("Send failed (try " + tries + ")", e);
            lastException = e;
            if (tries < JMS_MAXTRIES) {
                onException(e);
                log.debug("Will sleep a while before trying to send again");
                TimeUtils.exponentialBackoffSleep(tries,
                    Calendar.MINUTE);
            }
        } catch (Exception e) {
            log.debug("Send failed (try " + tries + ")", e);
            lastException = e;
            if (tries < JMS_MAXTRIES) {
                reconnect();
                log.debug("Will sleep a while before trying to send again");
                TimeUtils.exponentialBackoffSleep(tries,
                    Calendar.MINUTE);
            }
        }
    }
}

```

```
    }
    if (!operationSuccessful) {
        log.warn("Send failed", lastException);
        throw new IOFailure("Send failed.", lastException);
    }
}
/**
 * Do a reconnect to the JMSbroker. Does absolutely nothing, if already in
 * the process of reconnecting.
 */
protected void reconnect() {
    if (!connectionLock.writeLock().tryLock()) {
        log.debug("Reconnection already in progress. Do nothing");
        return;
    }
    try {
        log.info("Trying to reconnect to jmsbroker");
        boolean operationSuccessful = false;
        Exception lastException = null;
        int tries = 0;
        while (!operationSuccessful && tries < JMS_MAXTRIES) {
            tries++;
            try {
                doReconnect();
                operationSuccessful = true;
            } catch (Exception e) {
                lastException = e;
                log.debug("Reconnect failed (try " + tries + ")", e);
                if (tries < JMS_MAXTRIES) {
                    log.debug("Will sleep a while before trying to
                        + " reconnect again");
                    TimeUtils.exponentialBackoffSleep(tries,
                        Calendar.MINUTE);
                }
            }
        }
        if (!operationSuccessful) {
            log.warn("Reconnect to JMS broker failed",
                lastException);
            closeConnection();
        }
    } finally {
        // Tell everybody, that we are not trying to reconnect any longer
        connectionLock.writeLock().unlock();
    }
}
/**
 * Helper method for getting the right producer for a queue or topic.
 *
 * @param queueName The name of the channel
 *
 * @return The producer for that channel. A new one is created, if none
 *         exists.
 *
 * @throws JMSEException If a new producer cannot be created.
 */
private MessageProducer getProducer(String queueName) throws JMSEException {
    // Check if producer is in cache
    // If it is not, it is created and stored in cache:
    MessageProducer producer = producers.get(queueName);
    if (producer == null) {
        producer = getSession().createProducer(getDestination(queueName));
    }
}
```

```
        producers.put(queueName, producer);
    }
    return producer;
}
/**
 * Get the session. Will try reconnecting if session is null.
 *
 * @return The session.
 *
 * @throws IOException if no session is available, and reconnect does not
 *         help.
 */
private Session getSession() {
    if (session == null) {
        reconnect();
    }
    if (session == null) {
        throw new IOException("Session not available");
    }
    return session;
}
/**
 * Helper method for getting the right consumer for a queue or topic, and
 * message listener.
 *
 * @param channelName The name of the channel
 * @param ml           The message listener to add as listener to the
 *                   channel
 *
 * @return The producer for that channel. A new one is created, if none
 *         exists.
 *
 * @throws JMSEException If a new producer cannot be created.
 */
private MessageConsumer getConsumer(String channelName, MessageListener ml)
    throws JMSEException {
    String key = getConsumerKey(channelName, ml);
    MessageConsumer consumer = consumers.get(key);
    if (consumer == null) {
        consumer = getSession().createConsumer(getDestination(channelName));
        consumers.put(key, consumer);
        listeners.put(key, ml);
    }
    return consumer;
}
/**
 * Generate a consumerkey based on the given channel name and
 * messageListener.
 *
 * @param channel      Channel name
 * @param messageListener a messageListener
 *
 * @return the generated consumerkey.
 */
protected static String getConsumerKey(
    String channel, MessageListener messageListener) {
    return channel + CONSUMER_KEY_SEPARATOR + messageListener;
}
/**
 * Get the channelName embedded in a consumerKey.
 *
 * @param consumerKey a consumerKey
 */
```

```
* @return name of channel embedded in a consumerKey
*/
private static String getChannelName(String consumerKey) {
    //assumes argument consumerKey was created using metod getConsumerKey()
    return consumerKey.split(CONSUMER_KEY_SEPARATOR)[0];
}
/**
 * Helper method to establish one Connection and associated Session,
 *
 * @throws JMSEException If some JMS error occurred during the creation of
 * the required JMS connection and session
 */
private void establishConnectionAndSession() throws JMSEException {
    // Establish a queue connection and a session
    connection = getConnectionFactory().createConnection();
    session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE);
    connection.setExceptionListener(this);
    connection.start();
}
/**
 * Sends an ObjectMessage on a queue destination.
 *
 * @param msg the NetarkivetMessage to be wrapped and send as an
 * ObjectMessage.
 * @param to the destination topic.
 *
 * @throws JMSEException if message failed to be sent.
 */
private void doSend(NetarkivetMessage msg, ChannelID to)
    throws JMSEException {
    connectionLock.readLock().lock();
    try {
        ObjectMessage message = getSession().createObjectMessage(msg);
        synchronized (msg) {
            getProducer(to.getName()).send(message);
            // Note: Id is only updated if the message does not already have
            // an id. This ensures that resent messages keep the same ID
            msg.updateId(message.getJMSMessageID());
        }
    } finally {
        connectionLock.readLock().unlock();
    }
    log.trace("Sent message '" + msg.toString() + "'");
}
/**
 * Method adds a listener to the given queue or topic.
 *
 * @param channelName the messagequeue to listen to
 * @param ml the messagelistener
 *
 * @throws IOFailure if the operation failed.
 */
private void setListener(String channelName, MessageListener ml) {
    log.debug("Adding " + ml.toString() + " as listener to "
        + channelName);
    String errMsg = "JMS-error - could not add Listener to queue/topic: "
        + channelName;

    int tries = 0;
    boolean operationSuccessful = false;
    Exception lastException = null;
    while (!operationSuccessful && tries < JMS_MAXTRIES) {
        tries++;
        try {
```

```

        connectionLock.readLock().lock();
        try {
            getConsumer(channelName, ml).setMessageListener(ml);
        } finally {
            connectionLock.readLock().unlock();
        }
        operationSuccessful = true;
    } catch (JMSEException e) {
        lastException = e;
        log.debug("Set listener failed (try " + tries + ")", e);
        if (tries < JMS_MAXTRIES) {
            onException(e);
            log.debug("Will sleep a while before trying to set listener"
                + " again");
            TimeUtils.exponentialBackoffSleep(tries, Calendar.MINUTE);
        }
    } catch (Exception e) {
        lastException = e;
        log.debug("Set listener failed (try " + tries + ")", e);
        if (tries < JMS_MAXTRIES) {
            reconnect();
            log.debug("Will sleep a while before trying to set listener"
                + " again");
            TimeUtils.exponentialBackoffSleep(tries, Calendar.MINUTE);
        }
    }
}
}
if (!operationSuccessful) {
    log.warn(errMsg, lastException);
    throw new IOFailure(errMsg, lastException);
}
}
/**
 * Remove a messagelistener from a channel (a queue or a topic).
 *
 * @param ml          A specific MessageListener
 * @param channelName a channelname
 */
private void removeListener(MessageListener ml, String channelName) {
    String errMsg = "JMS-error - could not remove Listener from "
        + "queue/topic: " + channelName;

    int tries = 0;
    Exception lastException = null;
    boolean operationSuccessful = false;
    log.info("Removing listener from channel '" + channelName + "'");
    while (!operationSuccessful && tries < JMS_MAXTRIES) {
        try {
            tries++;
            connectionLock.readLock().lock();
            try {
                MessageConsumer messageConsumer = getConsumer(channelName,
                    ml);

                messageConsumer.close();
                consumers.remove(getConsumerKey(channelName, ml));
                listeners.remove(getConsumerKey(channelName, ml));
            } finally {
                connectionLock.readLock().unlock();
            }
            operationSuccessful = true;
        } catch (JMSEException e) {
            lastException = e;
            log.debug("Remove listener failed (try " + tries + ")", e);
            onException(e);
        }
    }
}

```

```

        log.debug("Will and sleep a while before trying to remove"
            + " listener again");
        TimeUtils.exponentialBackoffSleep(tries, Calendar.MINUTE);
    } catch (Exception e) {
        lastException = e;
        log.debug("Remove listener failed (try " + tries + ")", e);
        reconnect();
        log.debug("Will and sleep a while before trying to remove"
            + " listener again");
        TimeUtils.exponentialBackoffSleep(tries, Calendar.MINUTE);
    }
}
if (!operationSuccessful) {
    log.warn(errMsg, lastException);
    throw new IOFailure(errMsg, lastException);
}
}
/**
 * Reconnect to JMSBroker and reestablish session. Resets senders and
 * publishers.
 *
 * @throws JMSEException If unable to reconnect to JMSBroker and/or
 *         reestablish sessions
 */
private void doReconnect()
    throws JMSEException {
    closeConnection();
    establishConnectionAndSession();
    // Add listeners already stored in the consumers map
    log.debug("Re-add listeners");
    for (Map.Entry<String, MessageListener> listener
        : listeners.entrySet()) {
        setListener(getChannelName(listener.getKey()),
            listener.getValue());
    }
    log.info("Reconnect successful");
}
}

```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/common/distribute/JMSConnection.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/common/distribute/JMSConnection.java?root=netarchivesuite) (requires access to our code repository).

#### 2.12.4. ArcRepositoryClient

The ArcRepositoryClient handles access to the Archive module, both upload and low-level access. There are two implementations in the default distribution:

- JMSArcRepositoryClient - this is a full-fledged distributed implementation using JMS for communication, allowing multiple locations with multiple machines each.
- LocalArcRepositoryClient - An ARC repository implementation that stores all files in a local directories.

The ArcRepositoryClient interface is defined by the Java interface *dk.netarkivet.common.distribute.arcrepository.ArcRepositoryClient*:

```

/**
 * Generic interface defining all methods that an ArcRepository provides.

```

```

* Typically, an application using this will only see one of the restricted
* superinterfaces.
*
*/
public interface ArcRepositoryClient extends
    HarvesterArcRepositoryClient, ViewerArcRepositoryClient,
    PreservationArcRepositoryClient {
    /** Call on shutdown to release external resources. */
    void close();
    /**
     * Gets a single ARC record out of the ArcRepository.
     *
     * @param arcfile The name of a file containing the desired record.
     * @param index The offset of the desired record in the file
     * @return a BitarchiveRecord-object, or null if request times out or object
     * is not found.
     * @throws IOException If the get operation failed.
     * @throws IllegalArgumentException if the given arcfile is null or empty string,
     * or the given index is negative.
     */
    BitarchiveRecord get(String arcfile, long index) throws IllegalArgumentException;
    /**
     * Retrieves a file from an ArcRepository and places it in a local file.
     * @param arcfilename Name of the arcfile to retrieve.
     * @param replica The bitarchive to retrieve the data from. On
     * implementations with only one replica, null may be used.
     * @param toFile Filename of a place where the file fetched can be put.
     * @throws IOException if there are problems getting a reply or the file
     * could not be found.
     */
    void getFile(String arcfilename, Replica replica, File toFile);
    /**
     * Store the given file in the ArcRepository. After storing, the file is
     * deleted.
     *
     * @param file A file to be stored. Must exist.
     * @throws IOException thrown if store is unsuccessful, or failed to clean
     * up files after the store operation.
     * @throws IllegalArgumentException if file parameter is null or file is not an
     * existing file.
     */
    void store(File file) throws IOException, IllegalArgumentException;
    /**
     * Runs a batch job on each file in the ArcRepository.
     *
     * @param job An object that implements the FileBatchJob interface. The
     * initialize() method will be called before processing and the finish()
     * method will be called afterwards. The process() method will be called
     * with each File entry.
     *
     * @param replicaId The archive to execute the job on.
     * @return The status of the batch job after it ended.
     */
    BatchStatus batch(FileBatchJob job, String replicaId);
    /** Updates the administrative data in the ArcRepository for a given
     * file and replica.
     *
     * @param fileName The name of a file stored in the ArcRepository.
     * @param bitarchiveId The id of the replica that the administrative
     * data for fileName is wrong for.
     * @param newval What the administrative data will be updated to.
     */
    void updateAdminData(String fileName, String bitarchiveId,

```



```

        ReplicaStoreState newval);
/** Updates the checksum kept in the ArcRepository for a given
 * file. It is the responsibility of the ArcRepository implementation to
 * ensure that this checksum matches that of the underlying files.
 *
 * @param filename The name of a file stored in the ArcRepository.
 * @param checksum The new checksum.
 */
void updateAdminChecksum(String filename, String checksum);
/** Remove a file from one part of the ArcRepository, retrieving a copy
 * for security purposes. This is typically used when repairing a file
 * that has been corrupted.
 *
 * @param fileName The name of the file to remove.
 * @param bitarchiveId The id of the replica from which to remove the file.
 * @param checksum The checksum of the file to be removed.
 * @param credentials A string that shows that the user is allowed to
 * perform this operation.
 * @return A local copy of the file removed.
 */
File removeAndGetFile(String fileName, String bitarchiveId,
        String checksum, String credentials);
}

```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/common/distribute/arcrepository/ArcRepositoryClient.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/common/distribute/arcrepository/ArcRepositoryClient.java?root=netarchivesuite) (requires access to our code repository).

### 2.12.5. IndexClient

The IndexClient provides the Lucene indices that are used for deduplication and for viewerproxy access. It makes use of the ArcRepositoryClient to fetch data from the archive and implements several layers of caching of these data and of Lucene-indices created from the data. It is advisable to perform regular clean-up of the cache directories.

The IndexClient interface is defined by the Java interface *dk.netarkivet.common.distribute.indexserver.JobIndexCache*:

```

/** An interface to a cache of data for jobs. */
public interface JobIndexCache {
    /** Get an index for the given list of job IDs.
     * The resulting file contains a suitably sorted list.
     * This method should always be safe for asynchronous calling.
     * This method may use a cached version of the file.
     *
     * @param jobIDs Set of job IDs to generate index for.
     * @return An index, consisting of a file and the set this is an index for.
     * This file must not be modified or deleted, since it is part of the cache
     * of data.
     */
    Index<Set<Long>> getIndex(Set<Long> jobIDs);
}

```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/common/distribute/indexserver/JobIndexCache.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/common/distribute/indexserver/JobIndexCache.java?root=netarchivesuite) (requires access to our code repository).

### 2.12.6. DBSpecifcs

This DBSpecifics interface allows substitution of the database used to store harvest definitions. There are three implementations, one for MySQL, one for Derby running as a separate server, and one for Derby running embeddedly. Which is these to choose is mostly a matter of individual preference. The embedded Derby implementation has been in use at the Danish web archive for over two years.

The interface is defined by the abstract Java class *dk.netarkivet.harvester.datamodel.DBSpecifics*:

```

/**
 * Abstract collection of DB methods that are not standard SQL. This class
 * is a singleton class whose actual implementation is provided by a subclass
 * as determined by the DB_SPECIFICS_CLASS setting.
 *
 */
public abstract class DBSpecifics extends SettingsFactory<DBSpecifics> {

    /** The instance of the DBSpecifics class. */
    private static DBSpecifics instance;

    Log log = LogFactory.getLog(DBSpecifics.class);
    /** Get the singleton instance of the DBSpecifics implementation class.
     *
     * @return An instance of DBSpecifics with implementations for a given DB.
     */
    public static synchronized DBSpecifics getInstance() {
        if (instance == null) {
            instance = getInstance(CommonSettings.DB_SPECIFICS_CLASS);
        }
        return instance;
    }
    /**
     * Shutdown the database system, if running embeddedly. Otherwise, this
     * is ignored.
     *
     * Will log a warning on errors, but otherwise ignore them.
     */
    public abstract void shutdownDatabase();
    /** Get a temporary table for short-time use. The table should be
     * disposed of with dropTemporaryTable. The table has two columns
     * domain_name varchar(Constants.MAX_NAME_SIZE)
     * + config_name varchar(Constants.MAX_NAME_SIZE)
     * All rows in the table must be deleted at commit or rollback.
     *
     * @param c The DB connection to use.
     * @throws SQLException if there is a problem getting the table.
     * @return The name of the created table
     */
    public abstract String getJobConfigsTmpTable(Connection c)
    throws SQLException;
    /** Dispose of a temporary table gotten with getTemporaryTable. This can be
     * expected to be called from within a finally clause, so it mustn't throw
     * exceptions.
     *
     * @param c The DB connection to use.
     * @param tableName The name of the temporarily created table.
     */
    public abstract void dropJobConfigsTmpTable(Connection c, String tableName);
    /**
     * Backup the database. For server-based databases, where the administrator
     * is expected to perform the backups, this method should do nothing.
     * This method gets called within one hour of the hour-of-day indicated
     * by the DB_BACKUP_INIT_HOUR settings.
     */
}

```

```
*
* @param backupDir Directory to which the database should be backed up
* @throws SQLException On SQL trouble backing up database
* @throws PermissionDenied if the directory cannot be created.
*/
public abstract void backupDatabase(File backupDir) throws SQLException;
/** Get the name of the JDBC driver class that handles interfacing
* to this server.
*
* @return The name of a JDBC driver class
*/
public abstract String getDriverClassName();
/** Update a table to a newer version, if necessary. This will check the
* schemaversions table to see the current version and perform a
* table-specific update if required.
*
* @param tableName The table to update
* @param toVersion The version to update the table to.
* @throws IllegalStateException If the table is an unsupported version, and
* the toVersion is less than the current version of the table
* @throws NotImplementedException If no method exists for migration from
* current version of the table to the toVersion of the table.
* @throws IOException in case of problems in interacting with the database
*/

public synchronized void updateTable(String tableName, int toVersion) {
    ArgumentNotValid.checkNotNullOrEmpty(tableName, "String tableName");
    ArgumentNotValid.checkPositive(toVersion, "int toVersion");
    int currentVersion = DBUtils.getTableVersion(
        DBConnect.getDBConnection(), tableName);
    log.info("Trying to migrate table '" + tableName + "' from version '"
        + currentVersion + "' to version '" + toVersion + "'.");
    if (currentVersion == toVersion) {
        // Nothing to do. Version of table is already correct.
        return;
    }
    if (currentVersion > toVersion) {
        throw new IllegalStateException("Database is in an illegalState: "
            + "The current version of table '" + tableName
            + "' is not acceptable "
            + "(current version is greater than requested version).");
    }
    if (tableName.equals("jobs")) {
        if (currentVersion < 3) {
            throw new IllegalStateException("Database is in an illegalState: "
                + "The current version " + currentVersion
                + " of table '" + tableName + "' is not acceptable. "
                + "(current version is less than open source version).");
        }
        if (currentVersion == 3 && toVersion >= 4) {
            migrateJobsv3tov4();
            currentVersion = 4;
        }
        if (currentVersion == 4 && toVersion >= 5) {
            migrateJobsv4tov5();
        }
        if (currentVersion == 5 && toVersion >= 6) {
            throw new NotImplementedException(
                "No method exists for migrating table '" + tableName
                + "' from version " + currentVersion
                + " to version " + toVersion);
        }
    }
    // future updates of the job table are inserted here
}
```

```
        if (currentVersion > 5) {
            throw new IllegalStateException("Database is in an illegalState: "
                + "The current version (" + currentVersion
                + ") of table '" + tableName
                + "' is not an acceptable/known version. ");
        }
    } else if (tableName.equals("fullharvests")) {
        if (currentVersion < 2) {
            throw new IllegalStateException("Database is in an illegalState: "
                + "The current version " + currentVersion
                + " of table '" + tableName + "' is not acceptable. "
                + "(current version is less than open source version).");
        }
        if (currentVersion == 2 && toVersion >= 3) {
            migrateFullharvestsv2tov3();
            currentVersion = 3;
        }
        if (currentVersion == 3 && toVersion >= 4) {
            throw new NotImplementedException(
                "No method exists for migrating table '" + tableName
                + "' from version " + currentVersion
                + " to version " + toVersion);
        }
        // future updates of the job table are inserted here
        if (currentVersion > 4) {
            throw new IllegalStateException("Database is in an illegalState: "
                + "The current version (" + currentVersion
                + ") of table '" + tableName
                + "' is not an acceptable/known version. ");
        }
    } else if (tableName.equals("configurations")) {
        if (currentVersion < 3) {
            throw new IllegalStateException("Database is in an illegalState: "
                + "The current version " + currentVersion
                + " of table '" + tableName + "' is not acceptable. "
                + "(current version is less than open source version).");
        }
        if (currentVersion == 3 && toVersion >= 4) {
            migrateConfigurationsv3ov4();
            currentVersion = 4;
        }
        if (currentVersion == 4 && toVersion >= 5) {
            throw new NotImplementedException(
                "No method exists for migrating table '" + tableName
                + "' from version " + currentVersion
                + " to version " + toVersion);
        }
        // future updates of the job table are inserted here
        if (currentVersion > 5) {
            throw new IllegalStateException("Database is in an illegalState: "
                + "The current version (" + currentVersion
                + ") of table '" + tableName
                + "' is not an acceptable/known version. ");
        }
    } else if (tableName.equals("global_crawler_trap_lists")) {
        if (currentVersion == 0 && toVersion == 1) {
            createGlobalCrawlerTrapLists();
            currentVersion = 1;
        }
        if (currentVersion > 1) {
            throw new NotImplementedException(
                "No method exists for migrating table '" + tableName
                + "' from version " + currentVersion
```

```

        + " to version " + toVersion);
    }
} else if (tableName.equals("global_crawler_trap_expressions")) {
    if (currentVersion == 0 && toVersion == 1) {
        createGlobalCrawlerTrapExpressions();
        currentVersion = 1;
    }
    if (currentVersion > 1) {
        throw new NotImplementedException(
            "No method exists for migrating table '" + tableName
            + "' from version " + currentVersion
            + " to version " + toVersion);
    }
} else {
    // This includes cases where currentVersion < toVersion
    // for all tables that does not have migration functions yet
    throw new NotImplementedException(
        "No method exists for migrating table '" + tableName
        + "' to version " + toVersion);
}
}
}
/** Migrates the 'jobs' table from version 3 to version 4
 * consisting of a change of the field forcemaxbytes from int to bigint
 * and setting its default to -1.
 * Furthermore the default value for field num_configs is set to 0.
 * @throws IOException in case of problems in interacting with the database
 */
protected abstract void migrateJobsv3tov4();

/** Migrates the 'jobs' table from version 4 to version 5
 * consisting of adding new fields 'resubmitted_as_job' and 'submittedDate'.
 * @throws IOException in case of problems in interacting with the database
 */
protected abstract void migrateJobsv4tov5();

/** Migrates the 'configurations' table from version 3 to version 4.
 * This consists of altering the default value of field 'maxbytes' to -1.
 */
protected abstract void migrateConfigurationsv3ov4();

/** Migrates the 'fullharvests' table from version 2 to version 3.
 * This consists of altering the default value of field 'maxbytes' to -1.
 */
protected abstract void migrateFullharvestsv2tov3();
/**
 * Creates the initial (version 1) of table 'global_crawler_trap_lists'.
 */
protected abstract void createGlobalCrawlerTrapLists();
/**
 * Creates the initial (version 1) of table
 * 'global_crawler_trap_expressions'.
 */
protected abstract void createGlobalCrawlerTrapExpressions();
}
}

```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/harvester/datamodel/DBSpecifics.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/harvester/datamodel/DBSpecifics.java?root=netarchivesuite) (requires access to our code repository).

## 2.12.7. Notifications

The Notifications interface lets you choose how you want important error notifications to be handled in your system. Two implementations exist, one to send emails, and one to print the messages to `System.err`. Adding more specialised plugins should be easy.

The Notifications interface is defined by the abstract class `dk.netarkivet.common.utils.Notifications`:

```
/**
 * This class encapsulates reacting to a serious error message.
 *
 */
public abstract class Notifications {
    /**
     * Notify about an error event. This is the same as calling
     * {@link #errorEvent(String, Throwable)} with null as the second parameter.
     *
     * @param message The error message to notify about.
     */
    public void errorEvent(String message) {
        errorEvent(message, null);
    }
    /**
     * Notifies about an error event including an exception.
     *
     * @param message The error message to notify about.
     * @param e        The exception to notify about.
     * May be null for no exception.
     */
    public abstract void errorEvent(String message, Throwable e);
}
```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/common/utis/Notifications.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/common/utis/Notifications.java?root=netarchivesuite) (requires access to our code repository).

### 2.12.8. HeritrixController

The HeritrixController interface defines our interface for initialize a running Heritrix instance and communicate with this instance. We have an implementation that starts heritrix as its own process and then communicates with it using JMX (JMXHeritrixController), and a deprecated implementation with heritrix embedded inside NetarchiveSuite (DirectHeritrixController), which controls Heritrix using a CrawlController instance.

The HeritrixController interface is defined by the Java interface `dk.netarkivet.harvester.harvesting.HeritrixController`:

```
/**
 * This interface encapsulates the direct access to Heritrix, allowing
 * for accessing in various ways (direct class access or JMX). Heritrix is
 * expected to perform one crawl for each instance of an implementing class.
 *
 */
public interface HeritrixController {
    /** Initialize a new CrawlController for executing a Heritrix crawl.
     * This does not start the crawl.
     *
     */
}
```

```
void initialize();
/** Request that Heritrix start crawling.  When this method returns,
 * either Heritrix has failed in the early stages, or the crawljob has
 * been successfully created.  Actual crawling will commence at some
 * point hereafter.
 * @throws IOException If something goes wrong during startup.
 */
void requestCrawlStart() throws IOException;
/** Tell Heritrix to stop crawling.  Heritrix may take a while to actually
 * stop, so you cannot assume that crawling is stopped when this method
 * returns.
 */
void beginCrawlStop();
/** Request that crawling stops.
 * This makes a call to beginCrawlStop(), unless the crawler
 * is already shutting down.  In that case it does nothing.
 *
 * @param reason A human-readable reason the crawl is being stopped.
 */
void requestCrawlStop(String reason);
/** Query whether Heritrix is in a state where it can finish crawling.
 * Returns true if no uris remain to be harvested, or it has met
 * either the maxbytes limit, the document limit,
 * or the time-limit for the current harvest.
 *
 * @return True if Heritrix thinks it is time to stop crawling.
 */
boolean atFinish();
/** Returns true if the crawl has ended, either because Heritrix finished
 * or because we terminated it.
 *
 * @return True if the CrawlEnded event has happened in Heritrix,
 * indicating that all crawls have stopped.
 */
boolean crawlIsEnded();
/**
 * Get the number of currently active ToeThreads (crawler threads).
 * @return Number of ToeThreads currently active within Heritrix.
 */
int getActiveToeCount();
/** Get the number of URIs currently on the queue to be processed.  This
 * number may not be exact and should only be used in informal texts.
 *
 * @return How many URIs Heritrix have lined up for processing.
 */
long getQueuedUriCount();
/**
 * Get an estimate of the rate, in kb, at which documents
 * are currently being processed by the crawler.
 * @see
 * org.archive.crawler.framework.StatisticsTracking#currentProcessedKBPerSec()
 * @return Number of KB data downloaded by Heritrix over an undefined
 * interval up to now.
 */
int getCurrentProcessedKBPerSec();
/** Get a human-readable set of statistics on the progress of the crawl.
 * The statistics is
 * discovered uris, queued uris, downloaded uris, doc/s(avg), KB/s(avg),
 * dl-failures, busy-thread, mem-use-KB, heap-size-KB, congestion,
 * max-depth, avg-depth.
 * If no statistics are available, the string "No statistics available"
 * is returned.
 * Note: this method may disappear in the future.
 */
```

```

*
* @return Some ascii-formatted statistics on the progress of the crawl.
*/
String getProgressStats();
/** Returns true if the crawler has been paused, and thus not
* supposed to fetch anything. Heritrix may still be fetching stuff, as
* it takes some time for it to go into full pause mode. This method can
* be used as an indicator that we should not be worried if Heritrix
* appears to be idle.
*
* @return True if the crawler has been paused, e.g. by using the
* Heritrix GUI.
*/
boolean isPaused();
/** Release any resources kept by the class.
*/
void cleanup();
/**
* Get harvest information. An example of this can be an URL pointing
* to the GUI of a running Heritrix process.
* @return information about the harvest process.
*/
String getHarvestInformation();
}

```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/harvester/harvesting/HeritrixController.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/harvester/harvesting/HeritrixController.java?root=netarchivesuite) (requires access to our code repository).

### 2.12.9. ActiveBitPreservation

The ActiveBitpreservaton interface defines our interface for initializing bitpreservation actions from our GUI. We have a filebased and a database based implementations. Both these implementations communicate with the archive through the ArcRepository interface.

The ActiveBitPreservation interface is defined by the Java interface *dk.netarkivet.archive.arcpository.bitpreservation.ActiveBitPreservation*:

```

/**
* All bitpreservation implementations are assumed to have access to admin data
* and bitarchives. Operations may request information from the bitarchive by
* sending batch jobs, reading admin data directly, or reading from cached
* information from either.
*/
public interface ActiveBitPreservation {
    // General state
    /**
    * Get details of the state of one or more files in the bitarchives
    * and admin data.
    *
    * @param filenames the list of filenames to investigate
    * @return a map ([filename]-> [FilePreservationState]) with the
    * preservationstate of all files in the list.
    * The preservationstates in the map will be null for all filenames,
    * that are not found in admin data.
    */
    Map<String, PreservationState> getPreservationStateMap(
        String... filenames);
}

```



```
* Get the details of the state of the given file in the bitarchives
* and admin data.
*
* @param filename A given file
* @return the FilePreservationState for the given file. This will be null,
* if the filename is not found in admin data.
*/
PreservationState getPreservationState(String filename);

// Check state for bitarchives
/**
 * Return a list of files marked as missing on this replica.
 * A file is considered missing if it exists in admin data, but is not known
 * in the bit archives. Guaranteed not to recheck the archive, simply
 * returns the list generated by the last test.
 *
 * @param replica The replica to get missing files from.
 * @return A list of missing files.
 */
Iterable<String> getMissingFiles(Replica replica);
/**
 * Return a list of files with changed checksums on this replica.
 * A file is considered changed if checksum does not compare to
 * admin data. Guaranteed not to recheck the archive, simply returns the
 * list generated by the last test.
 *
 * @param replica The replica to get a list of changed files from.
 * @return A list of files with changed checksums.
 */
Iterable<String> getChangedFiles(Replica replica);
/** Update the list of files in a given bitarchive. This will be used for
 * the next call to getMissingFiles.
 *
 * @param replica The replica to update list of files for.
 */
void findMissingFiles(Replica replica);
/** Update the list of checksums in a given replica. This will be used
 * for the next call to getChangedFiles.
 *
 * @param replica The replica to update list of files for.
 */
void findChangedFiles(Replica replica);
/**
 * Return the number of missing files for replica. Guaranteed not to
 * recheck the archive, simply returns the number generated by the last
 * test.
 *
 * @param replica The replica to get the number of missing files from.
 * @return The number of missing files.
 */
long getNumberOfMissingFiles(Replica replica);
/**
 * Return the number of changed files for replica. Guaranteed not to
 * recheck the archive, simply returns the number generated by the last
 * test.
 *
 * @param replica The replica to get the number of changed files from.
 * @return The number of changed files.
 */
long getNumberOfChangedFiles(Replica replica);
/**
 * Return the total number of files for replica. Guaranteed not to
 * recheck the archive, simply returns the number generated by the last
```

```
* update.
*
* @param replica The replica to get the number of files from.
* @return The number of files.
*/
long getNumberOfFiles(Replica replica);
/**
 * Return the date for last check of missing files for replica. Guaranteed
 * not to recheck the archive, simply returns the date for the
 * last test.
 *
 * @param replica The replica to get date for changed files from.
 * @return The date for last check of missing files.
 */
Date getDateForMissingFiles(Replica replica);
/**
 * Return the date for last check of changed files for replica. Guaranteed
 * not to recheck the archive, simply returns the date for the
 * last test.
 *
 * @param replica The replica to get date for changed files from.
 * @return The date for last check of changed files.
 */
Date getDateForChangedFiles(Replica replica);
// Update files in bitarchives
/**
 * Check that files are indeed missing on the given replica, and present
 * in admin data and reference replica. If so, upload missing files from
 * reference replica to this replica.
 *
 * @param replica The replica to restore files to
 * @param filenames The names of the files.
 */
void uploadMissingFiles(Replica replica, String... filenames);
/**
 * Check that the checksum of the file is indeed different to the value in
 * admin data and reference replica. If so, remove missing file and upload
 * it from reference replica to this replica.
 *
 * @param replica The replica to restore file to
 * @param filename The name of the file
 * @param credentials The credentials used to perform this replace operation
 * @param checksum The known bad checksum. Only a file with this bad
 * checksum is attempted repaired.
 */
void replaceChangedFile(Replica replica, String filename,
                        String credentials, String checksum);
// Check state for admin data
/**
 * Return a list of files represented in replica but missing in AdminData.
 *
 * @return A list of missing files.
 */
Iterable<String> getMissingFilesForAdminData();
/**
 * Return a list of files with wrong checksum or state in admin data.
 *
 * @return A list of files with wrong checksum or state.
 */
Iterable<String> getChangedFilesForAdminData();
// Update admin data
/**
 * Add files unknown in admin.data to admin.data.
```

```
*
 * @param filenames The files to add.
 */
void addMissingFilesToAdminData(String... filenames);
/**
 * Reestablish admin data to match bitarchive states for file.
 *
 * @param filename The file to reestablish state for.
 */
void changeStateForAdminData(String filename);
}
```

The above code may differ from the latest version in our repository: [https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/\\*checkout\\*/trunk/src/dk/netarkivet/archive/arcrepository/bitpreservation/ActiveBitPreservation.java?root=netarchivesuite](https://gforge.statsbiblioteket.dk/plugins/scmsvn/viewcvs.php/*checkout*/trunk/src/dk/netarkivet/archive/arcrepository/bitpreservation/ActiveBitPreservation.java?root=netarchivesuite) (requires access to our code repository).

## 2.13. XML handling by Deploy

edit

Deploy needs to create a settings file in the `xml` fileformat for each application. The creation of these files involves the ability to manipulate the `xml` datastructure, create new `xml` instances based on previous ones and save the result to a file.

The existing `XMLTree` class (under `dk.netarkivet.common.utils`) did not support these abilities, and it would require a larger structural change of the class to meet these demands.

A new class (`dk.netarkivet.deploy.XmlStructure`) has therefore been implemented to meet the requests for handling the `xml` datastructure. This class uses the `com.dom4j` structure for handling the `xml` datastructure.

The `dk.netarkivet.deploy.XmlStructure` class has the ability to inherit and overwrite, which is used in the deploy configuration structure. The ability to inherit is implemented as creating a new instance identical to a current instance, thus inheritance can only occur during creation of a new instance of this class.

The overwrite function merges the current `XmlStructure` with a new tree. This means that the leafs which are present in both new tree and the current one, the value in the current leaf will be overwritten by the leaf on the new tree. Branches which only exists in the new tree will be appended to the current tree. The branches in both trees are recursively overwritten

## 3. Detailed Archive Design Description

edit

The Archive Design Description contains the description of overview of how the archive works, describes Indexing and caching and describes `CrawlLogIndexCache`.

*This section will be evolved further over time*

## 3.1. Archive Overview

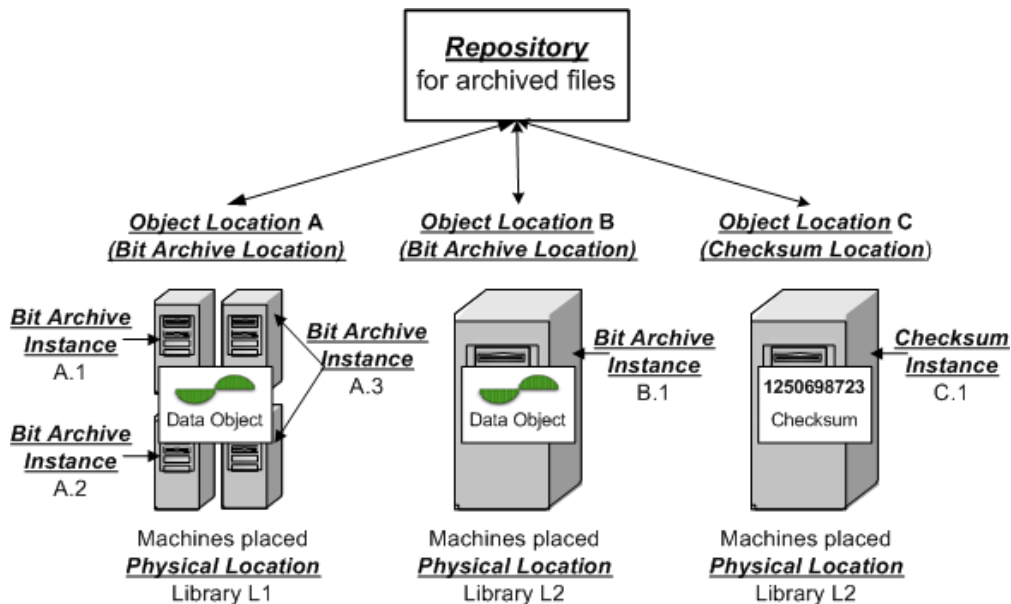
The NetarchiveSuite archive component allows files to be stored in a replicated and distributed environment.

Basically, the storage nodes are replicated over a number of *bitarchive replicas*. On storage each copy is checked to have equal checksum, before a store is accepted. Besides that, functionality for automatically checking that each bitarchive replica holds the same files with the same checksum, ensures that ever losing a bit is highly improbable.

For each *bitarchive replica*, the files may be stored in a distributed manner on several *bitarchives instances*. The philosophy is that you can use off-the-shelf hardware with ordinary hard disks for one bitarchive replica. This allows you to use cheaper hardware, and for more CPU-power per byte. This may be important if you regularly want to perform large tasks on the given bits, like indexing or content analysis. Beware however, that the power usage may be higher in such a setup, and that the maintenance needs may be higher.

### 3.1.1. Architecture

The archive architecture is shown in the following figure:



where

- Repository is handled by the ArcRepository application
- Replica X for bitarchives is handled by the BitarchiveMonitor application
- Bitarchive Instance X.N is handled by the Bitarchive application

So there is

- An ArcRepository: **Exactly One**
- BitarchiveMonitors: **One per bitarchive replica**

- Bitarchives: **As many as you like per bitarchive replica**

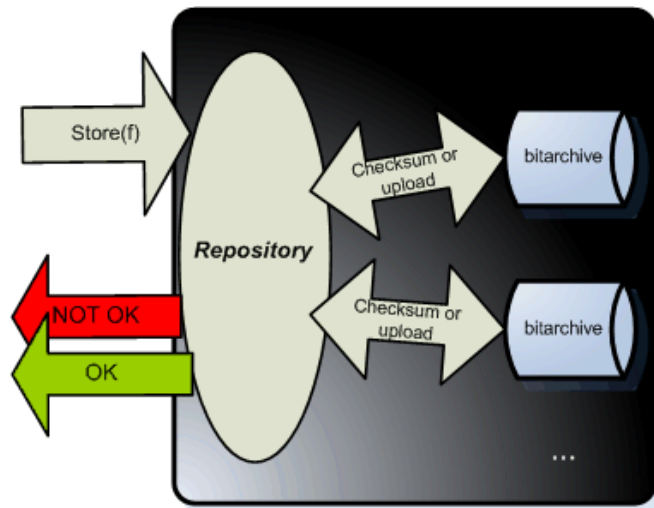
The components communicate using JMS, and a file transfer method (currently FTP, HTTP and HTTPS methods are implemented).

The public interface is through the ArcRepository. This interface allows you to send JMS messages to the ArcRepository and get responses, using the JMSArcRepositoryClient. It has methods to store a file, get a file, get one record from an ARC file, and run batch jobs on the archive. Batch jobs allow you to submit a job to be run on all files, or selected individual files, on a given location, and return the results.

Additionally, there are some methods for recovering from an error scenario which we will cover shortly under bit preservation.

### 3.1.2. Repository State Machine

When files are uploaded to the repository, it uploads the file by sending upload requests and checking the uploads as sketched in the below figure.

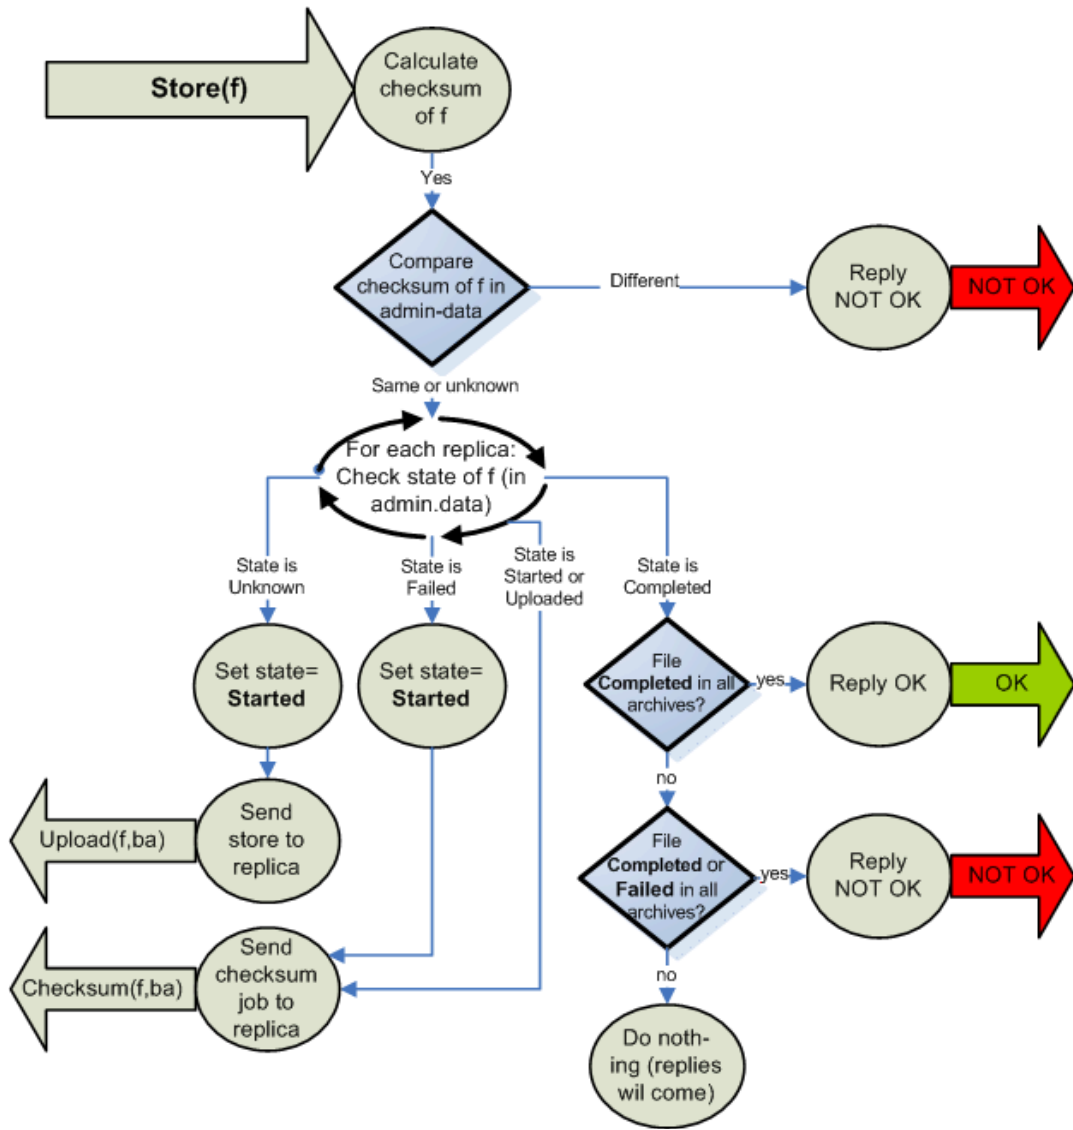


Internally this is handled following a state machine based on the messages it receives. These can either be

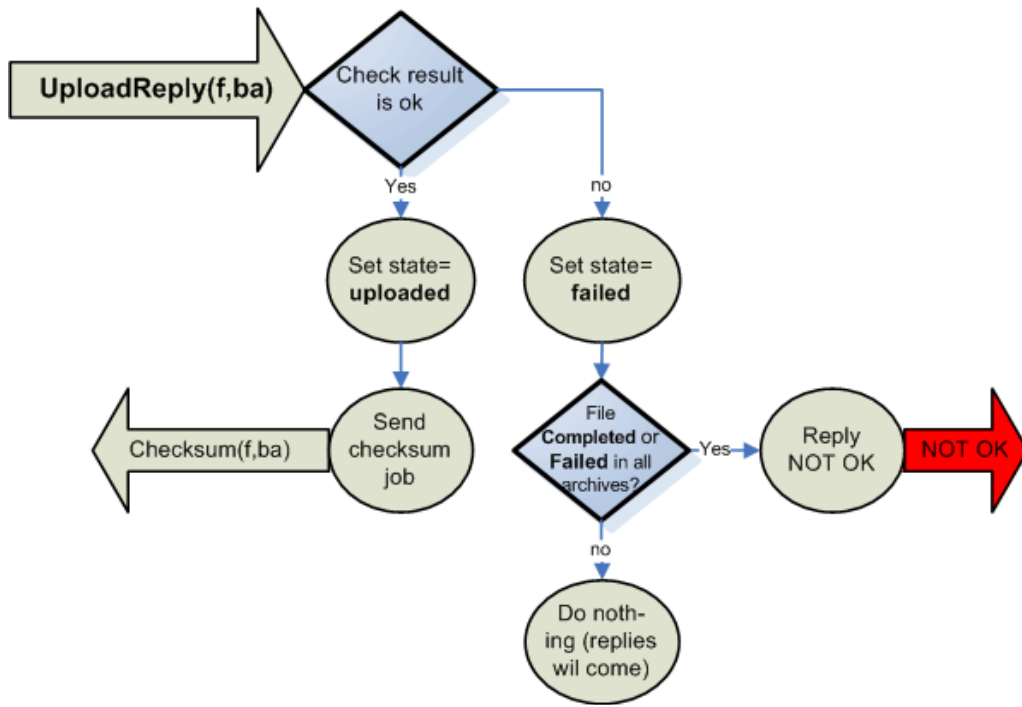
- a store message for a file to be stored
- a Upload Reply message from a bitarchive replica that was requested to upload a file (in storing process).
- a Checksum Reply message from a bitarchive replica that was requested find checksum as part of checking a file status (in storing process).

The state diagram for each message is given in the below figures:

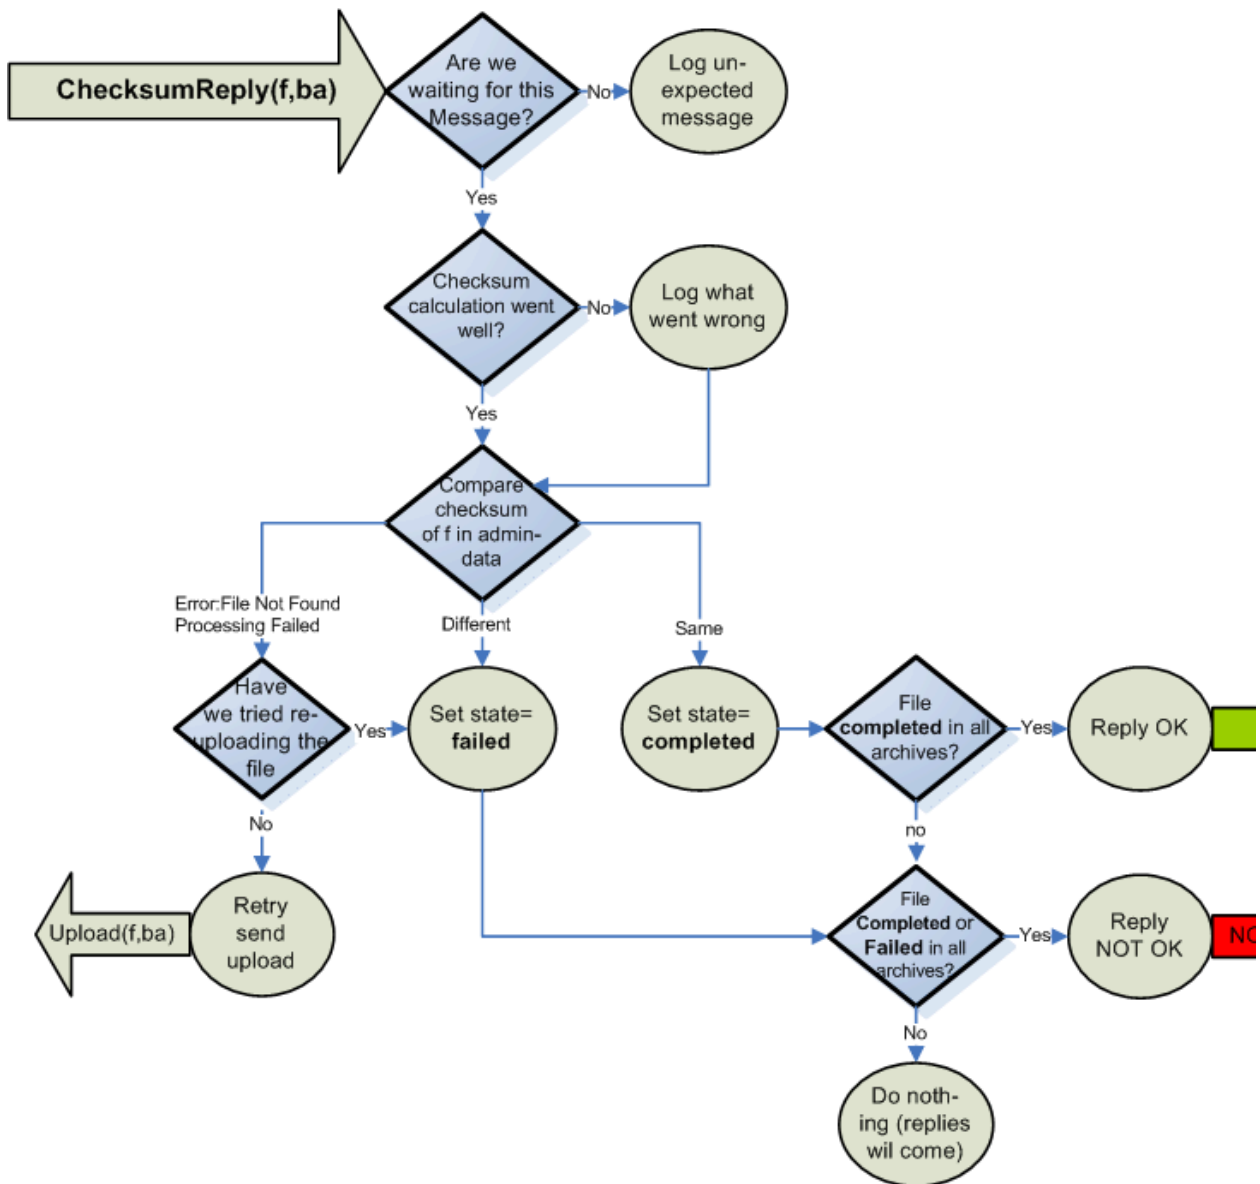
#### *Store message*



*Upload Reply message*



*Checksum Reply message*

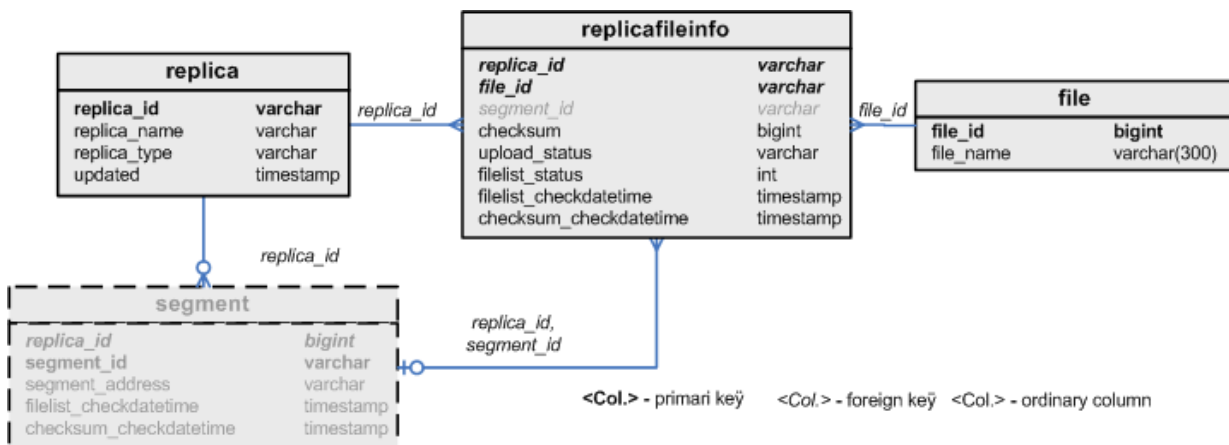


### 3.1.2.1. Admin for the Repository

The ArcRepository keep a record of the upload status of all the files for all the replicas. This is by default stored in the admin.data file. Though it is possible to store this in a database instead. This database can also be used for the bitpreservation.

It has the following table diagram:



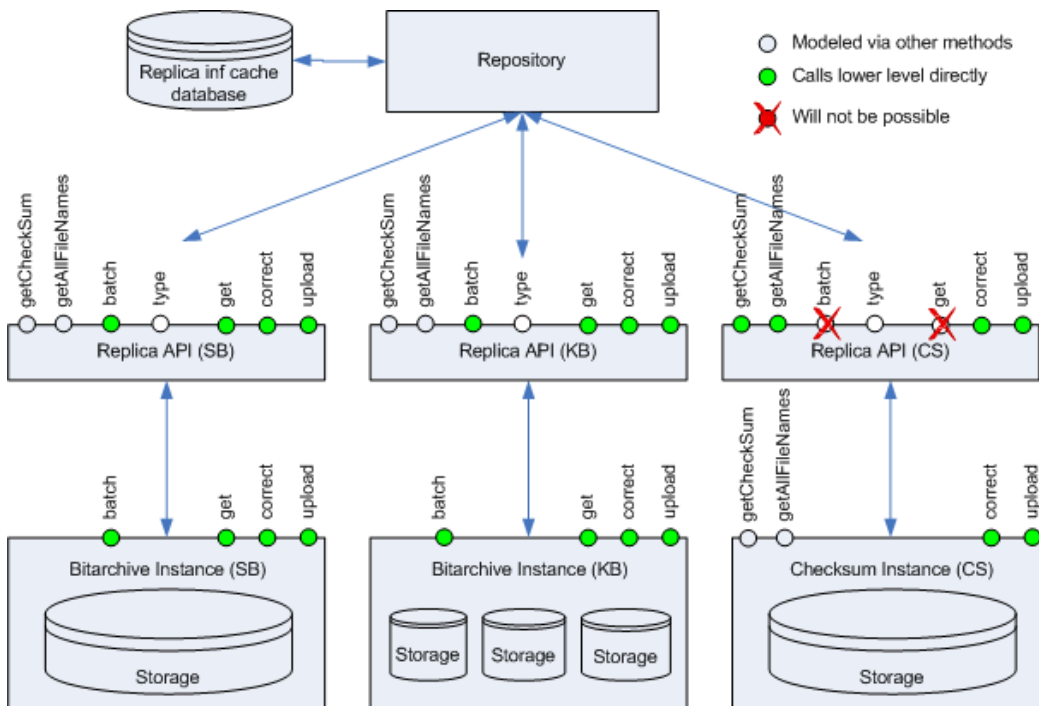


A few databases indices are also needed:

```
create index fileandreplica on replicafileinfo (file_id, replica_id);
create index replicaandfileliststatus on replicafileinfo (replica_id, filelist_status);
create index replicaandchecksumstatus on replicafileinfo (replica_id, checksum_status);
create index fileindex on file (filename);
```

### 3.1.2.2. Communication between ArcRepository and the replicas

The following drawing shows the message interaction between the ArcRepository and the replicas:



The 'replica inf cache database' is the database described above.

### 3.1.3. Extra functionality

Besides the basic architecture, the archive component contains the following extras:

- An index server, providing three kinds of indexing capabilities of a location.
- A bitpreservation web interface, providing an interface for monitoring bit integrity and handling error scenarios(missing/corrupt files in the archive)
- Command line tools for uploading a file, getting a file, getting an arc record or running a batch job.

### 3.1.3.1. The index server

The index server allows you to build an index over your data.

It does this by using the batch method defined in the arc repository.

It assumes you have an archive containing only ARC files, and that arcfiles are named <<job-number>>-\*.\*.arc(.gz), and that for each job number, there are arc files with the names <<job-number>>-metadata-\*.\*.arc, containing the crawl.log for the harvest that generated the arc files, and a cdx file for all the arc files. These files will be generated by the NetarchiveSuite harvester component.

Using the IndexRequestClient, you may request indexes over a number of jobs, either of type CDX, or as a lucene index. The lucene index comes in two different flavours,

- One is used for the deduplication feature of the harvesting component, which only contains the objects that are not of mime-type text/\*
- The other is used by the access component, and contains all objects.

### 3.1.3.2. The bit preservation interface

For monitoring the bit integrity of your archive, and for performing actions in case of bit errors, a user interface is available. This is installed as a site section of the NetarchiveSuite webserver.

This will basically give you a user interface with the following possibilities:

- Perform a check that all files are present on a given bitarchive replica (can take hours to complete this check)
- Perform a check that all files have correct checksum on a given bitarchive replica (can take hours to complete this check)
- Reestablish files that are missing on one bitarchive replica, but available in another replica
- Replace a file with bad checksum (ie. the file is corrupt) on one bitarchive replica, where a healthy copy is available in another replica

## 3.2. Indexes and caching

The deduplication code and the viewer proxy both make use of an index generating system to extract Lucene indexes from the data in the archive. This system makes extensive use of caching to improve index generation performance. This section describes the default index generating system implemented as the IndexRequestClient plugin.

There are four parts involved in getting an index, each of them having their own cache. The first part resides on the client side, in the `IndexRequestClient` class, which caches unzipped Lucene indexes and makes them available for use. The `IndexRequestClient` receives its data from the `CrawlLogIndexCache` in the form of gzipped Lucene indexes. The `CrawlLogIndexCache` generates the Lucene indexes based on Heritrix `crawl.log` files and CDX files extracted from the ARC files, and caches the generated indexes in gzipped form. The `crawl.log` files and CDX files are in turn received through two more caches, both of which extract their data directly from the archive using batch jobs and store them in sorted form in their caches.

All four caches are based on the generic `FileIndexCache` class, which handles the necessary synchronization to ensure that not only separate threads but also separate processes can access the cache simultaneously without corrupting it. When a specific cache item is requested, the cache is first checked to see if it already exists. If it doesn't, a file indicating that work is being done is locked by the process. If this lock is acquired, the actual cache-filling operation can take place, otherwise another thread or process must be working on it already, and we can wait until it finishes and take its data.

The `FileIndexCache` class is generic on the type of the identifier that indicates which item to get. The higher-level caches (`IndexRequestClient` and `CrawlLogIndexCache`) use a `Set<Long>` type to allow indexes of multiple jobs based on their IDs. The two low-level caches just use a `Long` type, so they operate on just one job at a time.

The two caches that handle multiple job IDs as their cache item ID must handle a couple of special scenarios: Their cache item ID may consist of hundreds or thousands of job IDs, and part of the job data may be unavailable. To deal with the first problem, any cache item with more than four job IDs in the ID set is stored in a file whose name contains the four lowest-numbered IDs followed by an MD5 checksum of a concatenation of all the IDs in sorted order. This ensures uniqueness of the cache file without overflowing operating system limits.

Every subclass to `FileBasedCache` uses its own directory, where the cache files are placed. The name of the final cache file is uniquely created from the id-set, which should be made into an index. Since the synchronization is done on the complete path to the cache file, then it must require two instances of the same class (e.g. `DedupCrawlLogIndexCache`), which is attempting to make cache on the same id-set at the same time, for a synkronisation block to occur. In this case the cache file would anyway only be made once, since the waiting instance will use the same cache file as the first instance creates.

A subclass of `CombiningMultiFileBasedCache` uses a corresponding subclass of `RawMetadataCache` to make sure that an cache file for every id exists (an id-cache file). If this file does not exist, then it will be created. Afterwards all the id-cache files will be combined to a complete file of the wanted id-set.

The id-cache files are blocked for other processes during their creation, but they are only created once since they can be used directly to create the Lucene cache for other id-sets, which contain this id.

### 3.3. CrawlLogIndexCache

The `CrawlLogIndexCache` guarantees that an index is always returned for a given request, regardless of whether part of the necessary data was available. This is done by performing a preparatory step where the data required to create the index is retrieved. If any of the data chunks are missing, a recursive attempt at generating an index for a reduced set is performed. Since the underlying data is always fetched from a cache, it is very likely that all the data for the reduced set is already available, so no further recursion is

typically needed. The set of job IDs that was actually found is returned from the request to cache data, while the actual data is stored in a file whose name can be requested afterwards. Note that future requests for the full set of job IDs will cause a renewed attempt at downloading the underlying data, which may take a while, especially if the lack of data is caused by a time-out.

The `CrawlLogIndexCache` is the most complex of the caches, but its various responsibilities are spread out over several superclasses.

- The top class is the generic `FileBasedCache` handles the locking necessary to have only one thread in one process at a time create the cached data. It also provides two helper methods: `getIndex()` is a forgiving cache lookup for complex cache items that handles the partial results described before, and the `get(Set<I>)` method allows for optimized caching of multiple simple cache requests.
- The `MultiFileBasedCache` handles the naming of files for caches that use sets as cache item identifiers.
- The `CombiningMultiFileBasedCache` extends the `MultiFileBasedCache` to have another, simpler cache as a data source, and providing an abstract method for combining the data from the underlying cache. It adds a step to the caching process of getting the underlying data, and only performs the combine action if all required data was found.
- The `CrawlLogIndexCache` is a `CombiningMultiFileBasedCache` whose underlying data is `crawl.log` files, but adds a simple CDX cache to provide data not found in the `crawl.log`. It also implements the combine method by creating a Lucene index from the `crawl.log` and CDX files, using code from Kristinn Sigurðsson. The other subclass of `CombiningMultiFileBasedCache`, which provides combined CDX indexes, is not currently used in the system, but is available at the `IndexRequestClient` level.
- The `CrawlLogIndexCache` is further subclasses into two flavors, `FullCrawlLogIndexCache` which is used in the viewer proxy, and `DedupCrawlLogIndexCache` which is used by the deduplicator in the harvester. The `DedupCrawlLogIndexCache` restricts the index to non-text files, while the `FullCrawlLogIndexCache` indexes all files.

The two caches used by `CrawlLogIndexCache` are `CDXDataCache` and `CrawlLogDataCache`, both of which are simply instantiations of the `RawMetadataCache`. They both work by extracting records from the archived metadata files based on regular expressions, using batch jobs submitted through the `ArcRepositoryClient`. This is not the most efficient way of getting the data, as a batch job is submitted separately for getting the files for each job, but it is simple. It could be improved by overriding the `get(Set<I>)` method to collect all the data in one batch job, though some care has to be taken with synchronization and avoiding refetching unnecessary data.

## 4. Detailed Common Design Description

edit

The Common Design Description is to contain parts of the Specific design in code section.

*This section will be evolved further over time*

## 5. Detailed Deploy Design Description

edit

The Deploy Design Description is to contain parts of the Specific design in code section.

*This section will be evolved further over time*

## 6. Detailed Harvester Design Description

edit

The Harvester Design Description contains the description of the database and gives a harvesting roundtrip.

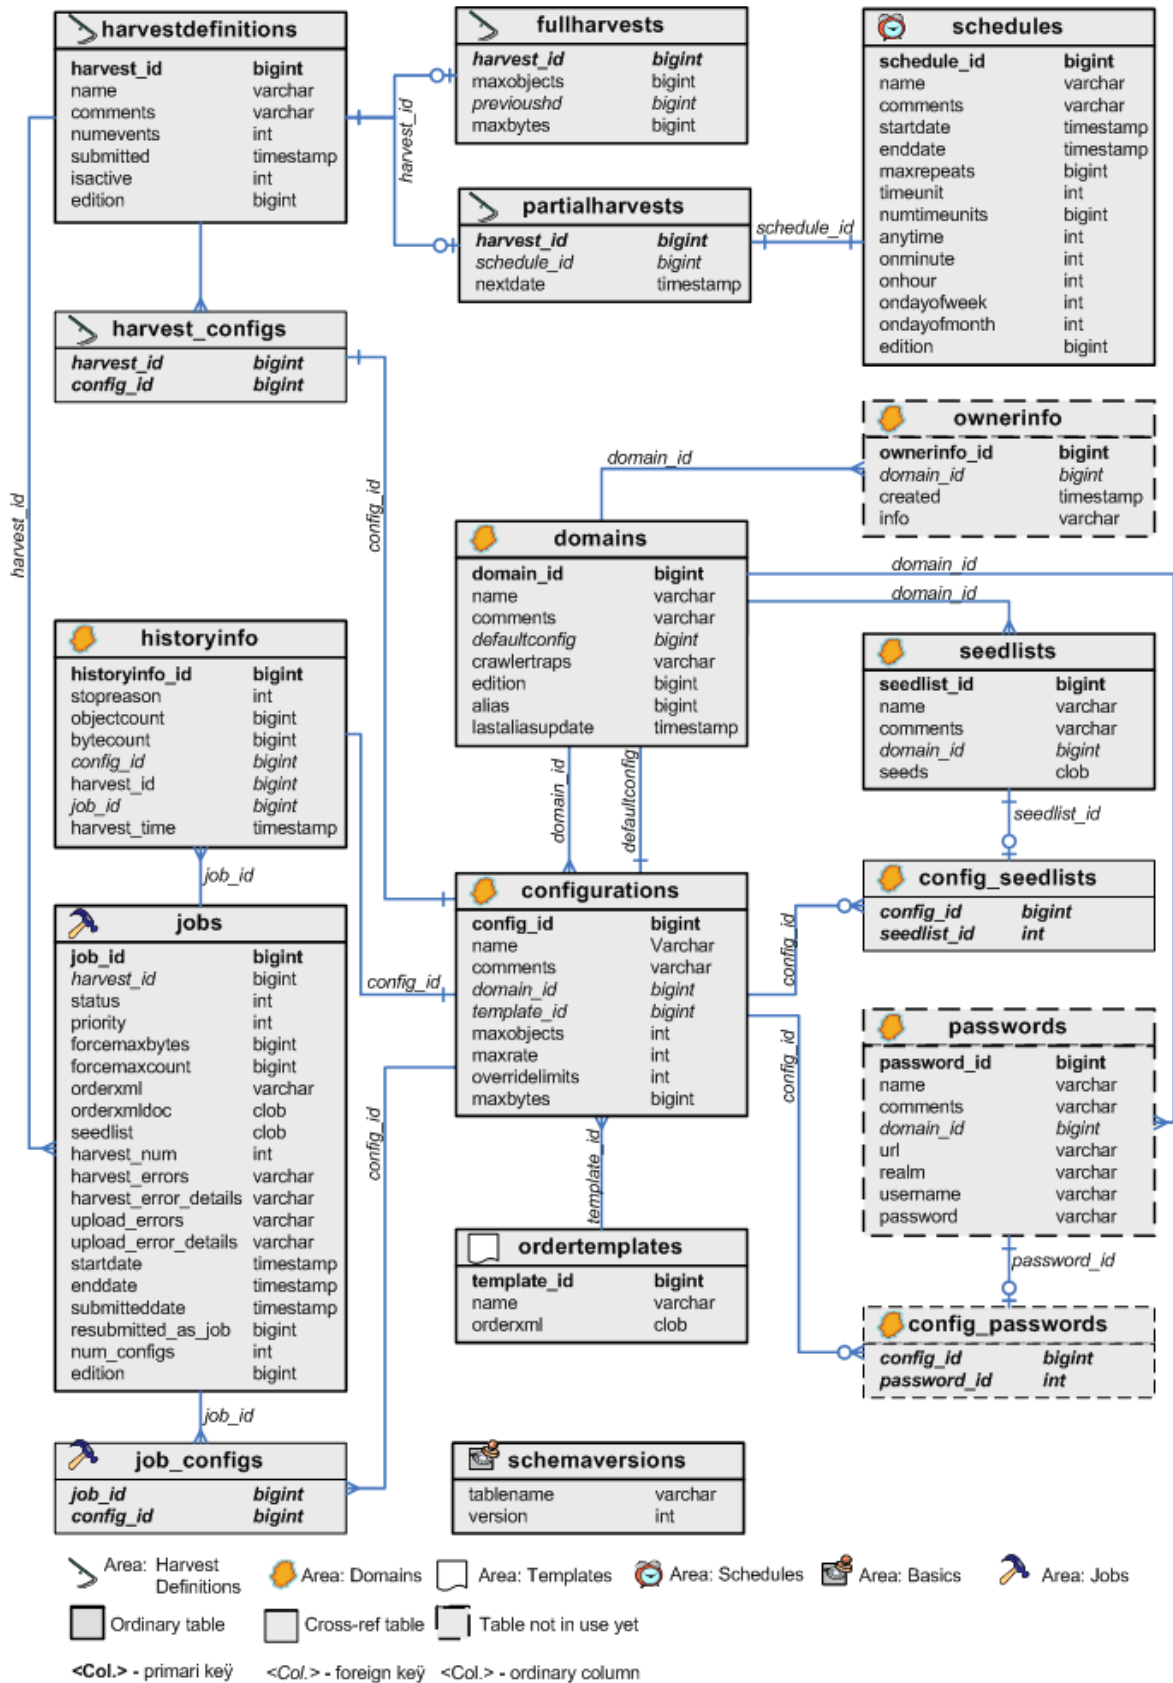
*This section will be evolved further over time*

### 6.1. Database

edit

#### 6.1.1. Database overview

Below you find a database diagram of the physical tables in the NetarchiveSuite database.



### 6.1.2. Table and Column Descriptions

Description of the individual tables and their columns can be now found in the Derby SQL create script [createfullhddb.sql](#). Note that Derby automatically generates indexes for unique columns, other types of databases may need to set these indices manually.

The table and column descriptions is given on form:

```

<ALL> ::= <AREA_WITH_TABLES>*

<AREA_WITH_TABLES> ::=
  --*****--
  -- Area: <Area Name> -- <Area description>
  --*****--
  <TABLE_DEFINITION>*

<TABLE_DEFINITION> ::=
  -----
  -- Name:      <Table name>
  -- Descr.:    <Table description>
  [-- Purpose: <Purpose description>]
  [-- Expected entry count: <Exp. count description>]
  <SQL_CREATE_TABLE_DEF>
  <Sql create index def>*
  <Sql insert data def>*

<SQL_CREATE_TABLE_DEF> ::=
  create table <Table name>(
    <COLUMN_DEF>+
    [<Primary key definition over more columns>]
  );

<COLUMN_DEF> ::=
  <Sql column def> -- <Column description>

```

## 6.2. Harvesting roundtrip

edit

This section describes what goes on during a harvest, from the templates are uploaded and harvests created till the ARC files are uploaded to the archive and historical information about the harvest is stored.

### 6.2.1. Initial steps

To create a harvest in the first place, we need to have a template to base it on. Additionally, we need a schedule for a selective harvest or some domains to harvest for a snapshot harvest.

#### 6.2.1.1. Uploading templates

Templates (Heritrix order.xml files) are uploaded using the **Definitions-upload-harvest-template.jsp** page. Templates are internally represented with the HeritrixTemplate object, which as part of its constructor verifies that certain elements - later modified programmatically - exist in the template. The template is then stored in the **templates** table in the database.

### 6.2.1.2. Creating domains

Domains can either be auto-created through selective harvest definitions or mass upload, or they can be created manually through the interface. Domains are represented with a Domain object, which in turn contains a list of SeedList objects and a list of DomainConfiguration objects. These are stored in the `domains`, `seedlists`, and `domainconfigurations` tables respectively. New domains are created with a single configuration, a single seedlist containing `http://www.<domain>` as its only seed, and a limit on number of bytes to download.

### 6.2.1.3. Creating schedules

NetarchiveSuite comes with four schedules, repeating respectively hourly, daily, weekly and monthly. More schedules can be created using the web interface. Schedules are represented with the Schedule object and stored in the `schedules` table in the database.

## 6.2.2. Creating a selective harvest

Harvests are created using the web-based user interface, as described in the User Manual.

### 6.2.3. Scheduling and splitting

The scheduler, a part of the harvest definition application, is responsible for creating new jobs for harvests, including splitting jobs into manageable chunks and sending them off to the harvesters.

#### 6.2.3.1. Scheduling

The HarvestScheduler class runs a TimerTask that executes every minute (which is the finest interval available for defining harvests anyway). It explicitly makes sure that only one scheduling thread runs at a time, as scheduling can take a while when a snapshot harvest is started.

When the scheduler activates, it selects all harvests that are ready to harvest, i.e. harvests that are activated and where the next harvest date is in the past (or, for snapshot harvest, they haven't run yet). For each harvest, a new thread is started to perform the operations required to send off the jobs -- this keeps snapshot harvests from blocking scheduling of selective harvests. Since the threads may run for a while, we keep a set of harvests currently undergoing scheduling and uses it to avoid that the same harvest gets scheduled several times concurrently.

As a side note, backups of embedded databases are performed by HarvestScheduler, too, as part of the regular scheduling check.

#### 6.2.3.2. Splitting

Splitting a harvest into jobs and sending those jobs off to the harvesters happens, as mentioned, in a separate thread. The first part of splitting is to reduce the harvests into chunks that can be handled by the scheduler itself -- since the data structures for domain configurations and their associated domains contain a fair amount of information, we cannot keep all of them in memory at the same time. For that reason alone, we split harvests into arbitrary chunks with no more domain configurations per chunk than the



configChunkSize setting allows. We use an iterator to avoid keeping all the domains and their configurations in memory for this operation, which iterates over all configurations in sorted order. The configurations are sorted by order.xml template, then by maximum number of bytes to harvest, and finally by the expected number of objects harvested. This makes sure that configurations that should be in the same job are sorted next to each other. A FilterIterator weeds out configurations in a snapshot harvest whose domains either are marked as aliases or were completely harvested in the snapshot harvest that the current snapshot harvest is based on.

For each chunk, we iteratively create new jobs by taking one domain configuration at a time and checking if it can be added to the job we're building. If it cannot, we store the job and start making a new one. Note that the jobs created are not submitted to the harvesters yet, that happens asynchronously as part of the scheduling check.

The check for whether a configuration can be added to a job is the most complex part of the scheduling system. It is based on the need to partition the domains into chunks such that all domains in a job take approximately the same amount of time to harvest and doesn't exceed memory limits of Heritrix. The estimation of the size of a domain is complicated by the facts that previously unharvested domains have an unknown size, and that domains can easily increase in size by several orders of magnitude by adding forums, image galleries or crawler traps. Furthermore, each Heritrix instance can only use one order.xml file.

Whether a domain configuration can be added to a job is a multi-stage check with the following stages:

1. The configuration must not already have been added to the job.
2. The job must not end up with more than configChunkSize configurations.
3. The configuration must use the same crawl template as the other configurations in the job.
4. If the byte limit for this job is determined by the harvest definition, the configuration must not have a smaller byte limit than the definition specifies. If the byte limit for the job is determined by the other configurations in the job, this configuration must have the same byte limit as the other configurations.
5. The expected number of objects harvested by all configurations in the job, based on previous harvests of the configurations, must not exceed the maxTotalSize setting.
6. The *relative* difference between the largest and smallest expected numbers of objects harvested by configurations in the job must be no more than the maxRelativeSizeDifference setting. Note that the default setting for this is 100, so expectations within a job differ by a factor 100, not just 100%. This prevents jobs from finishing many small configurations quickly and take a long time to finish a few, large configurations.
  - However, if the *absolute* difference between the largest and smallest expected numbers of objects harvested by configurations in the job is less than the .minAbsoluteSizeDifference setting, the relative difference is ignored. This allows the very smallest configurations to be lumped together in fewer jobs.

*Note: Check on overrides.*

The expected number of objects is found based on previous harvests of a given configuration and a few assumptions about the development of web sites. If a configuration hasn't been harvested before, defaults from the settings file are used. Expectations for previously harvested domains are calculated as follows:

1. The "best" previous harvest to estimate from is found by picking the most recent complete harvest

- using the configuration, or the harvest that harvested the most objects if the configuration never completed.
2. The expected size per object is found based on the average size in the "best" previous harvest, if that harvest got enough objects to be considered (at least 50), but at least as many as the `expectedAverageBytesPerObject` setting.
  3. A maximum number of objects is found based on the current limits of the configuration and the harvest and the expected size per object. If neither configuration nor harvest imposes any limits, an artificial limit for estimation purposes is taken from the `maxDomainSize` setting.
  4. A minimum number of objects is the number of objects found in the the "best" previous harvest, or is 0 if no previous harvest was found.
  5. If the configuration had previously been completed, the estimated number of objects is the difference of minimum and maximum divided by the `errorFactorPrevResult` setting plus the minimum.
    - Otherwise, the estimated number of objects is the difference of minimum and maximum divided by the `errorFactorBestGuess` setting plus the minimum.
  6. The expected number of objects is capped by the maximum based on the limits.

The `errorFactorBestGuess` setting should generally be smaller than the the `errorFactorPrevResult` setting, since there is more uncertainty about the actual number of objects when the harvest has never been completed. These two settings are best understood as the largest possible factor of error between our estimate and reality. If we use an error-factor of 10, we accept that while configurations could end up growing by as much as the hard limits allow, we split as if they only grow by one-tenth that amount. In most cases, growth will be limited, but it is likely that if a new archive or forum or somesuch is added to a site, the site can grow significantly between harvests. These settings determine the trade-off between the likelihood that some sites have grown a lot and the desire to keep similar-sized configurations in the same job.

Once the job does not get any more domain configurations added to it, it is added to the database with status 'New', and cannot change further except for status updates.

When all domain configurations for a harvest have been placed in jobs, the time for the next execution of the harvest is set. Note that the execution time is updated regardless of whether the jobs are actually successfull, or even have been run. Additionally, the counter of number of runs is updated.

If there are any errors in the scheduling process, including the creation of jobs, the harvest is deactivated to prevent the system from being overloaded with broken scheduling attempts.

#### **6.2.4. Talking to the harvesters**

Jobs created by the scheduler are sent to the harvesters as a `DoOneCrawlMessage`. This message contains not only the Job object, but also some metadata entries that are associated with the job. Currently, the metadata consists of a listing of the aliases that were used in the job creation and of a listing of the job IDs that should be used to get the deduplication index.

The `DoOneCrawlMessages` are placed on a JMS queue, either `ANY_HIGHPRIORITY_HACO` for selective/event harvests or `ANY_LOWPRIORITY_HACO` for snapshot harvests. At the same time, the job is set to status 'Submitted', indicating that it's in queue for being picked up by a harvester. The names of these queues is a historical artifact and does not indicate that "high priority" jobs can "get ahead" of "low priority" jobs, and there could potentially be just one or more than two queues. Notice that since the JMS messages are expected to be cleaned from the queues at system restart, we assume that any messages

about jobs in state "Submitted" are lost after a restart, and they are therefore automatically resubmitted at system startup.

Each HarvestControllerServer application listens to just one of the two queues (unless it is out of disk space). When it receives a message (remember that JMS guarantees exactly-once delivery for queues), it immediately sends a message back that tells the scheduler that the job has been picked up and can be put in state 'Started'. The HarvestControllerServer then spins off a new thread that as the first thing stops the HarvestControllerServer from listening for more jobs (it is done this way due to limitations on what the thread that JMS started can do to its listeners). There is also a bit of logic to ensure that no job messages are accepted between the start of the thread and the time that it stops listening.

At this point, the HarvestControllerServer has accepted that it will attempt to run the job and can start to set up the files necessary for running Heritrix.

### 6.2.5. Harvest setup

The directory used in a crawl is created by the HarvestControllerServer, using the job id and timestamp in the directory name. Details on what Heritrix reads and writes can be found in the [Heritrix "outside the GUI" page](#).

### 6.2.6. Running Heritrix

The HeritrixLauncher class sets the correct file paths in the Heritrix order.xml file and keeps an eye on the progress of the harvest. If Heritrix does not download any data for a period of time defined by the `noResponseTimeout` setting, HeritrixLauncher will stop the crawl. This is to avoid a single very slow web server from extending the crawl for very little gain. Also, if no crawler threads are active in Heritrix for a period of time defined by the `inactivityTimeout` setting, HeritrixLauncher will stop the crawl. This is a workaround for a subtle bug in Heritrix.

Since version 3.3.2, Heritrix is run by the harvester system as a standalone process. This allows access to Heritrix' web interface. The interfacing to the Heritrix process is controlled by JMXHeritrixController, an implementation of the HeritrixController interface. The old method of running Heritrix is implemented with the DirectHeritrixController, which is now deprecated. General documentation on JMX can be found as [part of the Java documentation](#), on the [Sun JMX Technologies pages](#), in the [JMX Accelerated Howto](#), and via the [JMX Wikipedia page](#). Heritrix' documentation of its JMX interface is partially described in the [JMX feedback page](#), but can also be investigated in more depth via the [Heritrix JMX command-line client](#), and in the source files [Heritrix.java](#) and [CrawlJob.java](#) (links for Heritrix version 1.12.1).

JMXHeritrixController starts a new process as part of its constructor, putting the jar files in `lib/heritrix/lib` and the NetarchiveSuite jar files in the classpath. The process is started in the directory created by the HarvestControllerServer, and all files created as part of the crawl are put into that directory. Stdout and stderr from Heritrix, along with a dump of the startup environment, are put in the `heritrix.out` file. The full command line used for running Heritrix is put in the log file.

Before the process is started, a [shutdown-hook](#) is added to attempt proper cleanup in case the harvest controller is shut down prematurely. Notice that this hook is removed if the process finishes normally.

After constructing the JMXHeritrixController object, HeritrixLauncher calls the `initialize()` method on the


JMXHeritrixController, which first checks that we're talking to the correct Heritrix instance (in case one was left over from earlier), then uses the addJob JMX command to create a job for the crawl. Before returning from initialize, we call getJobName() to extract from the job a unique name we can use to locate it by later. getJobName() also has the task to wait (using exponential back-off) until the job has actually been created, since the addJob command can return before the job actually exists.

After initialize() is done, the requestCallStart() method executes the JMX command requestCrawlStart to start the job, and we then enter a loop for the duration of the crawl. Inside the loop, we check for the two timeouts as well as for orderly termination of the job and log status reports every 20 seconds. These logs can be seen by the user in the System Overview web page.

Access to the Heritrix user interface can be had by connecting to the port specified by the `heritrixAdminGui` setting, using the admin name and password specified by the `heritrixAdminName` and `heritrixAdminPassword` settings, respectively.

The cleanup of the JMXHeritrixController involves issuing the shutdown JMX command to Heritrix, then waiting for a while (duration defined by the processTimeout setting) for Heritrix to end crawls and write its reports. If Heritrix doesn't stop within the timeout period, we forcibly kill it. After that, we collect the exit code and wait for the stdout/stderr collector processes to finish.

### 6.2.7. Creating metadata

After the heritrix has finished with the harvesting, the harvest is documented, and the result of this documentation is stored in a separate arcfile prefixed with the job id, and ending with "-metadata-1.arc". This metadata file contains all heritrix logs and reports associated with this harvest (crawl.log, local-errors.log, progress-statistics.log, runtime-errors.log, uri-errors.log, heritrix.out, crawl-report.txt, frontier-report.txt, hosts-report.txt, mimetype-report.txt, processors-report.txt, responsecode-report.txt, seeds-report.txt), some metadata about the job itself, and CDX'es of the contents of the arcfiles created by Heritrix. A  CDX line points to where an object is located in an arcfile, its length and mimetype. This metadata arcfile is uploaded along with the rest of the arcfiles.

### 6.2.8. Uploading

When Heritrix is finished, and the metadata arcfile created, all arcfiles are uploaded to the archive using a `ArcrepositoryClient`.

### 6.2.9. Storing harvest statistics

When uploading is done, a status message is sent back to the scheduler, containing error reports and harvest statistics. Errors are split into harvest errors and upload errors, since upload is attempted even if the harvest fails. For each, a short error description and a longer, detailed description are sent. The statistics sent are the following for each domain harvested:

- Number of objects harvested
- Number of bytes harvested
- Reason the harvest stopped, one of completed (no more objects found), object-limit (hit maximum allowed number of objects), size-limit (hit maximum allowed number of bytes, as specified by the harvest), config-size-limit (hit maximum allowed number of bytes, as specified by the configuration),

and unfinished (the harvest was interrupted before any of the other stop reasons applied).

...need to specify what gets counted within a domain...

...need to clarify the states of a harvest...

When the status message is received, the statistics from it is stored per domain in the database, along with the job number, the harvest number, the domain name, the configuration name, and a timestamp for receipt of the information.

After the harvest statistics have been sent to the database, the HarvestController application checks, if there is free space on the machine for a new harvestjob. If this is the case, it starts to listen on the job queue. If not, it goes into a dormant mode.

### 6.2.10. Old jobs

When a harvester application starts up, it checks whether any jobs are left from previous runs, in case the harvest or the upload was aborted. If there is, the last three steps described above are taken for the old jobs before the harvest application starts listening for new jobs.

### 6.2.11. Deduplication in NetarchiveSuite

deduplication is performed by using the DeDuplicator module developed by Kristinn Sigurdsson as the first of Heritrix write-processors:

```
<newObject name="DeDuplicator" class="is.hi.bok.deduplicator.DeDuplicator">
  <boolean name="enabled">true</boolean>
  <map name="filters"/>
  <string name="index-location">/home/test/JOLF/cache/DEDUP_CRAWL_LOG
/empty-cache</string>
  <string name="matching-method">By URL</string>
  <boolean name="try-equivalent">true</boolean>
  <boolean name="change-content-size">false</boolean>
  <string name="mime-filter">^text/.*</string>
  <string name="filter-mode">Blacklist</string>
  <string name="analysis-mode">Timestamp</string>
  <string name="log-level">SEVERE</string>
  <string name="origin"/>
  <string name="origin-handling">Use index information</string>
  <boolean name="stats-per-host">true</boolean>
</newObject>
```

This uses a Lucene (v. 2.0.0) index with information about previously harvested objects. This index may be empty.

In NetarchiveSuite, the index contains entries for objects fetched in an earlier harvest job, which this harvestjob is likely to revisit, i.e. this new harvestjob revisits some of same domains that the previously harvest jobs did.

The following arc record part of the metadata for every harvestjob mentions the NetarchiveSuite harvest jobs, that this job can be seen to continue. In the following record, only job with ID 2 is mentioned:

```
metadata://netarkivet.dk/crawl/setup/duplicatereductionjobs?majorversion=1&
```

```
minorversion=0&harvestid=3&harvestnum=0&jobid=4 130.226.228.7 20081215100759
text/plain 1
2
```

Each entry in the index contains the URL (unNormalized), content-digest, mimetype , Etag (if available), and the origin of the entry <arcfile>,<offset>.

Only non-text fetched URLs are indexed, as only URIs with non-text mimetypes are looked up by the DeDuplicator.

During the actual crawl-time, all non-text URIs are looked up in the index, and if a match (A URI is matched, if the URI is found, and it has the same digest as the current URI) is found in the index, the URI is ignored by the ARCWriter, but an entry is written to the crawl-log that contains the reference to the original stored URI. This reference is written to the Annotations part of the crawl-line (12th part of the crawl-line): deduplicate:arcfile,offset

Generation of deduplication indices is made by merging information in the crawl-log with information in the CDX-files generated for each job at the conclusion of the harvest.

The CDX'es contain information about the contents of the arc-files generated by Heritrix, but they lack information of the deduplicated objects. CDX-files containing references also to deduplicated objects can be generated from the crawl-logs by the tools provided in the wayback module. These tools are described in the Additional Tools manual.

The merging of this information in NetarchiveSuite was necessitated by the way we do Quality Assurance of the harvestJobs, which is done on a job by job basis, so we needed a way to refer to the deduplicated objects.

## 7. Detailed Monitor Design Description

edit

The Monitor Design Description is to contain parts of the Specific design in code section.

*This section will be evolved further over time*

System Design 3.12 (last edited 2010-08-16 10:24:31 by localhost)