

# **Oracle Endeca Commerce**

## **Deployment Template Usage Guide**

### **Version 11.0 • January 2014**





# Contents

<b>Preface</b> .....	<b>7</b>
About this guide.....	7
Who should use this guide.....	7
Conventions used in this guide.....	7
Contacting Oracle Support.....	8
<b>Chapter 1: Deploying and initializing an EAC Application</b> .....	<b>9</b>
Deployment prerequisites.....	9
About deploying EAC applications.....	9
Deploying and initializing an EAC application.....	9
Directories created by the Deployment Template.....	11
Configuring automated/file-based deployment.....	12
Modifying the template files to support custom applications.....	13
Custom application descriptors.....	13
Configuring an automated/file-based deployment for a custom application.....	15
Communicating with SSL-enabled Oracle Endeca components.....	16
Displaying the Deployment Template version.....	19
<b>Chapter 2: Configuring an EAC Application</b> .....	<b>21</b>
About configuring an EAC application.....	21
About the application configuration files.....	21
About the schema for AppConfig.xml.....	22
Configuring the application configuration files.....	22
Global application settings.....	24
Hosts.....	24
Lock Manager.....	25
Fault tolerance and polling interval properties.....	25
CAS Server.....	27
Forges.....	30
Dgidxs.....	31
Dgraphs.....	32
Log server.....	37
Report Generators.....	38
IFCR.....	38
Workbench Manager.....	39
Configuration Manager.....	40
Configuring the BeanShell scripts.....	42
Configuration overrides.....	44
<b>Chapter 3: Replacing the Default Forge Pipeline</b> .....	<b>47</b>
About the sample pipelines.....	47
Sample pipeline overview.....	47
Specifying a pipeline.....	47
Creating a new project.....	48
Modifying an existing project.....	50
Configuring a record specifier.....	51
Forge flags.....	52
Input record adapters.....	52
Dimension adapters.....	53
Indexer adapters.....	53
Output record adapters.....	53
Dimension servers.....	54
Common errors.....	54
<b>Chapter 4: Modifying Index Configuration for an Application</b> .....	<b>57</b>
Overview of the Index Configuration Command-line Utility.....	57

About index configuration ownership .....	58
About the schema for the index configuration file .....	59
Schema for an Endeca property, derived property, or dimension.....	59
Schema for precedence rules.....	64
Schema for global index configuration.....	66
Getting the index configuration for an application .....	67
Getting the merged index configuration for an application .....	69
Setting the index configuration for an application .....	71
Deleting the index configuration associated with an owner.....	72
An example of changing multi-select on a dimension.....	73
An example of changing a product.brand.name property to a dimension.....	74
An example of setting dimension display order.....	75

## **Chapter 5: Managing Data Operations.....77**

Running a baseline update with test data .....	77
Running a baseline update with production data .....	78
Running a partial update with production data .....	79
Running CAS crawls .....	80

## **Chapter 6: Script Reference.....81**

Analyzing Deployment Script Errors.....	81
Deployment Template script reference.....	81
Provisioning scripts.....	83
Forge-based data processing.....	83
Dgraph baseline update script using Forge.....	84
Dgraph partial update script using Forge.....	86
Dgraph baseline update script using Forge and a CAS full crawl script.....	88
Dgraph partial update script using Forge and a CAS incremental crawl script.....	90
Multiple CAS crawls and Forge updates.....	91
CAS-based data processing.....	91
Dgraph baseline update script using CAS.....	91
Dgraph partial update script using CAS.....	93
CAS crawl scripts for Record Store output.....	94
CAS crawl scripts for record file output.....	95
Report generation.....	97

## **Appendix A: EAC Development Toolkit.....101**

EAC Development Toolkit distribution and package contents.....	101
EAC Development Toolkit usage.....	102

## **Appendix B: Application Configuration File.....103**

Spring framework.....	103
XML schema.....	103
Application elements.....	103
Hosts.....	104
Components.....	104
Utilities.....	108
Customization/extension within the toolkit's schema.....	109
Customization/extension beyond the toolkit's schema.....	110

## **Appendix C: BeanShell Scripting.....111**

Script implementation.....	111
BeanShell interpreter environment.....	111
About implementing logic in BeanShell.....	113

## **Appendix D: Command Invocation.....115**

Invoke a method on an object.....	115
Identify available methods.....	115
Update application definition.....	117
Remove an application.....	117
Display component status.....	117

---

## Copyright and disclaimer

Copyright © 2003, 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

---

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

# Preface

Oracle Endeca Commerce is the most effective way for your customers to dynamically explore your storefront and find relevant and desired items quickly. An industry-leading faceted search and Guided Navigation solution, Oracle Endeca Commerce enables businesses to help guide and influence customers in each step of their search experience. At the core of Oracle Endeca Commerce is the MDEX Engine™, a hybrid search-analytical database specifically designed for high-performance exploration and discovery. The Endeca Content Acquisition System provides a set of extensible mechanisms to bring both structured data and unstructured content into the MDEX Engine from a variety of source systems. Endeca Assembler dynamically assembles content from any resource and seamlessly combines it into results that can be rendered for display.

Oracle Endeca Experience Manager is a single, flexible solution that enables you to create, deliver, and manage content-rich, cross-channel customer experiences. It also enables non-technical business users to deliver targeted, user-centric online experiences in a scalable way — creating always-relevant customer interactions that increase conversion rates and accelerate cross-channel sales. Non-technical users can determine the conditions for displaying content in response to any search, category selection, or facet refinement.

## About this guide

This guide describes how to configure, run, and customize the Deployment Template that is included with Tools and Frameworks.

The Deployment Template is a utility that you run to create a new Endeca application with the complete directory structure required for deployment, including Endeca Application Controller (EAC) control scripts, configuration files, and batch files or shell scripts that wrap common script functionality.

Some scripts created by the Deployment Template are documented in the *Assembler Application Developer's Guide*, rather than this guide, because the scripts are very closely associated with Assembler features. Similarly, some scripts are documented in the *Oracle Endeca Commerce Administrator's Guide* because the scripts are very closely associated with administrative tasks such as backing up or restoring site configuration.

## Who should use this guide

This guide is for developers or administrators who create and maintain Oracle applications using the Deployment Template.

## Conventions used in this guide

This guide uses the following typographical conventions:

Code examples, inline references to code elements, file names, and user input are set in `monospace` font. In the case of long lines of code, or when inline monospace text occurs at the end of a line, the following symbol is used to show that the content continues on to the next line: ↵

When copying and pasting such examples, ensure that any occurrences of the symbol and the corresponding line break are deleted and any remaining space is closed up.

## Contacting Oracle Support

Oracle Support provides registered users with important information regarding Oracle Endeca software, implementation questions, product and solution help, as well as overall news and updates.

You can contact Oracle Support through Oracle's Support portal, My Oracle Support at <https://support.oracle.com>.



## Deploying and initializing an EAC Application

This section describes how to deploy and initialize an EAC application using the Deployment Template.

### Deployment prerequisites

You must have installed Tools and Frameworks on the machine running the EAC Central Server (part of the Platform Services package) and set environment variables used by the Oracle Endeca software (including `ENDECA_ROOT`).

### About deploying EAC applications

The Deployment Template (`deploy`) script is available for both Windows and UNIX platforms. The prompts for the `deploy.sh` script are exactly the same as the `deploy.bat` script.

In every deployment environment, one server serves as the primary control machine and hosts the EAC Central Server, while all other servers act as agents to the primary server and host EAC Agent processes that receive instructions from the Central Server.

Both the EAC Central Server and the EAC Agent run as applications inside the Endeca HTTP Service. (As mentioned in the prerequisites, Tools and Frameworks only needs to be installed on the machine that hosts the EAC Central Server.)



**Note:** Mixed-platform deployments may require customization of the default Deployment Template scripts and components. For example, paths are handled differently on Windows and on UNIX, so paths and working directories are likely to require customization if a deployment includes servers running both of these operating systems.

### Deploying and initializing an EAC application

The `deploy` script in the `bin` directory creates, configures, and distributes the EAC application files into the deployment directory structure.

To deploy an EAC application on Windows:

1. Start a command prompt (on Windows) or a shell (on UNIX).

2. Navigate to `<installation path>\ToolsAndFrameworks\<version>\deployment_template\bin` or the equivalent path on UNIX.

3. From the `bin` directory, run the `deploy` script.  
For example, on Windows:

```
C:\Endeca\ToolsAndFrameworks\4.0.0\deployment_template\bin>deploy
```

4. If the path to the Platform Services installation is correct, press `Enter`.  
(The template identifies the location and version of your Platform Services installation based on the `ENDECA_ROOT` environment variable. If the information presented by the installer does not match the version or location of the software you plan to use for the deployment, stop the installation, reset your `ENDECA_ROOT` environment variable, and start again. Note that the installer may not be able to parse the Platform Services version from the `ENDECA_ROOT` path if it is installed in a non-standard directory structure. It is not necessary for the installer to parse the version number, so if you are certain that the `ENDECA_ROOT` path points to the correct location, proceed with the installation.)
5. Specify a short name for the application.  
The name should consist of lower- or uppercase letters, or digits between zero and nine.
6. Specify the full path into which your application should be deployed.  
This directory must already exist. The `deploy` script creates a folder inside of the deployment directory with the name of your application and the application directory structure.  
For example, if your application name is `MyApp`, and you specify the deployment directory as `C:\Endeca\apps`, the `deploy` script installs the template for your application into `C:\Endeca\apps\MyApp`.
7. Specify the port number of the EAC Central Server.  
By default, the Central Server host is the machine on which you are running `deploy` script and that all EAC Agents are running on the same port.
8. Specify the port number of Oracle Endeca Workbench, or press **Enter** to accept the default of 8006.
9. Specify the port number of the Live Dgraph, or press **Enter** to accept the default of 15000.
10. Specify the port number of the Authoring Dgraph, or press **Enter** to accept the default of 15002.
11. Specify the port number of the Log Server, or press **Enter** to accept the default of 15010.  
If the application directory already exists, the `deploy` script time stamps and archives the existing directory to avoid accidental loss of data.
12. Specify the path to the Oracle Wallet configuration file, `jps-config.xml`, or press **Enter** to accept the default path of `../../server/workspace/credential_store/jps-config.xml`.
13. Specify the path to the location where you can export your application content to, or press **Enter** to accept the default path of `../../server/workspace/state/repository`.
14. Navigate to the `control` directory of the newly deployed application.  
This is located under your application directory. For example: `C:\Endeca\apps\<app dir>\control`.
15. From the `control` directory, run the `initialize_services` script.
  - On Windows:  

```
<app dir>\control\initialize_services.bat
```
  - On UNIX:  

```
<app dir>/control/initialize_services.sh
```

The script initializes each server in the deployment environment with the directories and configuration required to host your application. This script removes any existing provisioning associated with this application in the EAC and then adds the hosts and components in your application configuration file to the EAC. Use

caution when running this script. The script forces any components that are defined for this application to stop, which may lead to service interruption if executed on a live environment.

Once deployed, an EAC application includes all of the scripts and configuration files required to create an index and start an MDEX Engine.

If no script customization is required, the application is now ready for use. Go on to [Managing Data Operations](#) on page 77.

However, if you need to configure an EAC application (the scripts and files) to reflect your environment and data processing requirements, go on to [Configuring an EAC Application](#) on page 21 before [Managing Data Operations](#) on page 77.

## Directories created by the Deployment Template

The Deployment Template creates the following default directory structure. For each Endeca implementation that is deployed with the Deployment Template, look into these directories to identify currently used configuration options and scripts.

The Deployment Template is designed to support operations with the MDEX Engine in the production environment. This means it must support a variety of possible configurations and their modifications. Therefore, its `AppConfig.xml` file contains all the possible blocks and directories that you may need on your production servers.

For example, the Deployment Template has separate directories to ensure that the MDEX Engine operations are safely accessing only the information they need. Further, the default Deployment Template allows for configuring multiple Dgraphs, so additional directories are created to facilitate this task.

Directory	Contents
<code>config/lib</code>	Subdirectories to store any custom scripts or code for your Deployment Template project.
<code>config/pipeline</code>	The Developer Studio pipeline file and XML configuration files.
<code>config/report_templates</code>	Files required to generate an application's reports.
<code>config/script</code>	The <code>AppConfig.xml</code> file and related Deployment Template scripts responsible for defining the baseline update workflow and communication of different Endeca components with the EAC Central Server.
<code>control</code>	Shell (UNIX) or batch (Windows) scripts responsible for running different operations defined within <code>AppConfig.xml</code> .
<code>data/incoming</code>	The premodified incoming data files that are ready acquisition by the Endeca pipeline and should be processed.
<code>data/processing</code>	Temporary data and configuration files created and stored during the baseline update process.
<code>data/forge_output</code>	The data and configuration files that are output from the Forge process to the Dgidx process.
<code>data/dgidx_output</code>	The index files that are output from the Dgidx process.
<code>data/dgraphs</code>	The copy of the index files used by an instance of the MDEX Engine.
<code>data/state</code>	Autogenerated dimensions files.

Directory	Contents
data/complete_index_config	Merged configuration (that is, Developer Studio files from <code>config/pipeline</code> , with any Workbench- maintained files specified in the Deployment Template's <code>ConfigManager</code> component overwritten by files downloaded from the Workbench instance).
data/web_studio/config	Configuration files extracted from Workbench by the Deployment Template's <code>ConfigManager</code> component.
logs	Various log files within subdirectories, such as Dgidx logs.
reports	Generated reports.

## Configuring automated/file-based deployment

The Deployment Template supports a file-based configuration option to simplify the deployment of an EAC Application. This automation may be especially useful during development, when the same deployment process must be repeated many times.

You can create a deployment configuration file that contains name/values that satisfy the `deploy` script prompts, so you do not have to respond to the prompts manually. You specify the deployment configuration file as an argument to the `--install-config` flag when you run the `deploy` script.

The deployment configuration file should specify the application name, deployment path, deployment type, and all ports. The following example specifies the installation of a Dgraph deployment named `Discover`:

```
<install app-name="Discover">
  <deployment-path>/localdisk/endeca/apps</deployment-path>
  <base-module type="dgraph" />
  <options>
    <option name="eac-port">8888</option>
    <option name="workbench-port">8006</option>
    <option name="dgraph1Port">15000</option>
    <option name="authoringDgraphPort">15002</option>
    <option name="logserverPort">15010</option>
    <option name="jps-config-location">ToolsAndFrameworks/11.0.0/server/workspace/credential_store/jps-config.xml</option>
  </options>
</install>
```

To configure automated/file-based deployment:

1. Start a text editor, create a new text file, and copy/paste the example above.
2. If necessary, modify the default port values for the EAC Central Server, Workbench, Live Dgraph, Authoring Dgraph, and the Log Server to new values.
3. Save and close the file.
4. Run the `deploy` script and specify the `--install-config` flag and the location of the deployment configuration file.

The following example specifies the deployment descriptor (`deploy.xml`) for a version of the Discover Electronics reference application, then the `--install-config` flag with an argument to the deployment configuration file (`pci-app-install-config.xml`):

```
./deploy.sh --app /localdisk/endeca/ToolsAndFrameworks/*/reference/discover-
data-pci/deploy.xml --install-config /localdisk/infrontSetupScripts/config/pci-
app-install-config.xml --no-prompt
```

When a configuration file is specified for the Deployment Template, the deployment attempts to retrieve and validate required information from the document before proceeding. If any information is missing or invalid, the Deployment Template prompts for that information, as described in previous sections. To truly automate the install process, the `--no-prompt` flag may be passed to the installer, instructing it to fail (with error messages) if any information is missing and to bypass interactive verification of the Oracle Endeca version.

## Modifying the template files to support custom applications

This section provides information about deploying custom applications.

### Custom application descriptors

The Deployment Template deploys new applications based on application descriptor XML documents. The documents describe the directory structure associated with an application as well as the files to distribute during the deployment process.

By default, the Deployment Template ships with application descriptor files named `base_descriptor.xml` located in `<installation path>\ToolsAndFrameworks\<version>\deployment_template\app-templates`.

This document describes the directory structure of the deployment as well as the copying that is done during the deployment to distribute files into the new directories. Additionally, this document describes whether files are associated with a Windows or UNIX deployment, and whether copied files should be updated to replace tokens in the format `@@TOKEN_NAME@@` with text strings specified to the installer.

The following tokens are handled by the base descriptor:

- `@@WORKBENCH_PORT@@` - Oracle Endeca Workbench port.
- `@@DGRAPH_1_PORT @@` - Live Dgraph port.
- `@@AUTHORING_DGRAPH_PORT @@` - Authoring Dgraph port.
- `@@LOGSERVER_PORT@@` - Log Server port.
- `@@JPSCONFIG_LOCATION@@` - Path to Oracle Wallet configuration file.

The following tokens are handled by the Deployment Template:

- `@@EAC_PORT@@` - EAC Central Server port.
- `@@HOST@@` - Hostname of the server on which the deploy script is invoked.
- `@@PROJECT_DIR@@` - Absolute path of the target deployment directory.
- `@@PROJECT_NAME@@` - Name of the application to deploy.
- `@@ENDECA_ROOT@@` - Absolute path of the ENDECA\_ROOT environment variable.
- `@@SCRIPT_SUFFIX@@` - ".bat" for Windows, ".sh" for Linux installs.

In addition to these tokens, you can specify custom tokens to substitute in the files. Tokens are specified in the application descriptor file, including the name of the token to substitute as well as the question with which to prompt the user or the installer configuration option to parse to retrieve the value to substitute for the token.

The default application descriptors use this functionality to request the port number for Dgraphs, Log Servers and Forge servers.

If a project deviates from the Deployment Template directory structure, it may find it useful to create a custom application descriptor document, so that the default Deployment Template can continue to be used for application deployment.

Custom deployment descriptors may also be used to define add-on modules on top of a base install. For example, sample applications (such as the Sample Term Discovery and Clustering application) are shipped with a custom deployment descriptor file, which describes the additional files and directories to install on top of a base Dgraph deployment. Modules may be installed using the deploy batch or shell script, specifying the `--app` argument with the location of the application descriptor document. For example:

```
deploy.bat --app \
C:\Endeca\Solutions\sampleTermDiscovery-[VERSION]\data\deploy.xml
```

The installer prompts you to specify whether it should install the module as a standalone installation or if it should be installed on top of the base Dgraph deployment. Multiple add-on modules may be specified to the installer script, though only one of them may be a base install (that is, all but one of them should specify an attribute of `update="true"`).

The following excerpt from the Dgraph deployment application descriptor identifies the document's elements and attributes:

```
<!--
  Deployment Template installer configuration file. This file defines the directory
  structure to create and the copies to perform to distribute files into the
  new directory structure.

  The update attribute of the root install element indicates whether this is a
  core installation or an add-on module. When set to false or unspecified, the
  installation requires the removal of an existing target install directory (if
  present). When update is set to true, the installer preserves any existing
  directories, adding directories as required and distributing files based on the
  specified copy pattern.
-->
<app-descriptor update="false" id="Dgraph">

  <custom-tokens>
    <!-- Template custom token:
      <token name="MYTOKEN">
        <prompt-question>What is the value to substitute for token MYTO-
KEN?</prompt-question>
        <install-config-option>myToken</install-config-option>
        <default-value>My Value</default-value>
      </token>

      This will instruct the installer to look for the "myToken" option
      in a specified install config file (if one is specified) or to
      prompt the user with the specified question to submit a value. If a
      value is entered/retrieved, the installer will substitute instances
      of @@MYTOKEN@@ with the value.
    -->
    </custom-tokens>

  <dir-structure>
    <!-- Template directory:
      <dir platform="unix" primary="true"></dir>

      primary          builds directory only on primary server installs
```

```

    platform          builds directory only on specified platform.
                     Valid values: "win" and "unix"
-->
</dir-structure>

<!--
Copy source directory is specified relative to this file's directory
-->
<copy-pattern src-root="../data ">
  <!-- Template copy pattern:
    <copy clear-dest-dir="true" recursive="true"
          preserve-subdirs="true" filter-files="true"
          primary="true" platform="win" Endeca-version="480">
      <src-dir></src-dir>
      <src-file></src-file>
      <dest-dir></dest-dir>
    </copy>

src-dir             source directory, relative to root of deployment
                    template package.

src-file            source filename or pattern (using '*' wildcard
                    character) to copy from source dir

dest-dir            destination directory, relative to root of target
                    deployment directory.

clear-dest-dir      removes all files in target dir before copying

recursive           copies files matching pattern in subdirectories
                    of the specified source dir

preserve-subdirs    copies files, preserving dir structure. Only
                    applicable to recursive copies

filter-files        filters file contents and file names by replacing
                    tokens (format @@TOKEN@@) with specified
                    strings.

mode                applies the specified permissions to the files
                    after the copy. Mode string should be 3 octal
                    digits with an optional leading zero to
                    indicate octal, e.g. 755, 0644. Not relevant
                    for Windows deployments.

platform            applies copy to specified platform. Valid
                    values: "win" "unix"

Endeca-version      applies copy to specified Oracle Endeca version Valid
                    values: "460" "470" "480" "500"
-->
</copy-pattern>
</app-descriptor>

```

## Configuring an automated/file-based deployment for a custom application

The configuration file discussed in previous sections may be used to specify the location of custom application descriptor documents in place of the `--app` command line argument to the installer.

The following example shows how to install the Sample Term Discovery and Clustering application on top of the base Dgraph deployment.

```
<install app-name="MyApp" >
  <deployment-path>C:\Endeca</deployment-path>
  <base-module type="dgraph" />
  <additional-module type="custom">
    C:\Endeca\Solutions\sampleTermDiscovery-[VERSION]\data\deploy.xml
  </additional-module>
  <options>
    <option name="eac-port">8888</option>
    <option name="dgraph1Port">15000</option>
    <option name="logserverPort">15010</option>
    <option name="jps-config-location">ToolsAndFrameworks/11.0.0/server/workspace/credential_store/jps-config.xml</option>
  </options>
</install>
```

## Communicating with SSL-enabled Oracle Endeca components

The Deployment Template supports enabling SSL to communicate securely with the EAC Central Server and with the Content Acquisition System version 3.0.x and later. (Secure communication between the Deployment Template and CAS is not supported in CAS 2.2.x.)

For details about enabling SSL in the EAC Central Server or Agent, refer to the *Oracle Endeca Security Guide*. For details about enabling SSL in CAS, refer to the *CAS Developer's Guide*.

To use the template with an SSL-enabled Central Server:

1. Update `runcommand.bat` / `.sh` to load your SSL keystore and truststore.



**Note:** To enable secure communication, you must have already followed the documentation to create a Java keystore and truststore, containing your generated certificates. Upload a copy of these certificates to the server on which your Deployment Template scripts will run. Edit the `runcommand` file to specify the locations of these files.

- On Windows, edit `runcommand.bat` to add the following lines:

```
...
set JAVA_ARGS=%JAVA_ARGS% "-Djava.util.logging.config.file=%~dp0..\config\script\logging.properties"

if exist [%path%\to\truststore] (
  set TRUSTSTORE=[%path%\to\truststore]
) else (
  echo WARNING: Cannot find truststore at [%path%\to\truststore]. Secure EAC
  communication may fail.
)

if exist [%path%\to\keystore] (
  set KEYSTORE=[%path%\to\keystore]
) else (
  echo WARNING: Cannot find keystore at [%path%\to\keystore]. Secure EAC com-
  munication may fail.
)
```



```

set JAVA_ARGS=%JAVA_ARGS% "-Djavax.net.ssl.trustStore=%TRUSTSTORE%" "-
Djavax.net.ssl.trustStoreType=JKS" "-Djavax.net.ssl.trustStorePassword=[trust-
store password]"

set JAVA_ARGS=%JAVA_ARGS% "-Djavax.net.ssl.keyStore=%KEYSTORE%" "-
Djavax.net.ssl.keyStoreType=JKS" "-Djavax.net.ssl.keyStorePassword=[keystore
password]"

set CONTROLLER_ARGS=--app-config AppConfig.xml

...

```

Note that the final two new lines (beginning with "set JAVA\_ARGS" are wrapped to fit the page size of this document, but each of those two lines should have no line breaks. Also note that you need to fill in the locations and passwords of your keystore and truststore files in the locations indicated by the placeholders in italics.

- On UNIX, edit `runcommand.sh` as follows:

```

...

JAVA_ARGS="{JAVA_ARGS} -Djava.util.logging.config.file={WORKING_DIR}/../con-
fig/script/logging.properties"

if [ -f "[/path/to/truststore]" ] ; then
  if [ -f "[/path/to/keystore]" ] ; then
    TRUSTSTORE=[/path/to/truststore]
    KEYSTORE=[/path/to/keystore]
    JAVA_ARGS="{JAVA_ARGS} -Djavax.net.ssl.trustStore={TRUSTSTORE}"
    JAVA_ARGS="{JAVA_ARGS} -Djavax.net.ssl.trustStoreType=JKS"
    JAVA_ARGS="{JAVA_ARGS} -Djavax.net.ssl.trustStorePassword=[truststore
password]"
    JAVA_ARGS="{JAVA_ARGS} -Djavax.net.ssl.keyStore={KEYSTORE}"
    JAVA_ARGS="{JAVA_ARGS} -Djavax.net.ssl.keyStoreType=JKS"
    JAVA_ARGS="{JAVA_ARGS} -Djavax.net.ssl.keyStorePassword=[keystore pass-
word]"
  else
    echo "WARNING: Cannot find keystore at [/path/to/keystore]. Secure EAC
communication may fail."
  fi
else
  echo "WARNING: Cannot find truststore at [/path/to/truststore]. Secure EAC
communication may fail."
fi

CONTROLLER_ARGS="--app-config AppConfig.xml"

...

```

2. In the `app` element of the `AppConfig.xml` document, update the `sslEnabled` attribute to `true`.  
The `sslEnabled` attribute is a application-wide setting that applies to the EAC and to CAS (if used in your application).
3. Specify the SSL-enabled port for the EAC.

The Endeca HTTP Service uses a separate port to communicate securely. For example, the default non-SSL connector is on port 8888 and the default SSL connector listens on port 8443. The SSL port should be specified in the `eacPort` attribute of the `app` element in the `AppConfig.xml` document.

4. If you are using CAS in your application, specify the SSL-enabled port for CAS.

The Endeca CAS Service uses a separate port to communicate securely. For example, the default non-SSL port is 8500 and the default SSL port is 8505. The SSL port should be specified in the `value` attribute of `casPort`.

5. Specify the non-SSL connector for hosts.

Internally, the EAC Central Server always initiates communication with Agents by communicating with the non-SSL connector. When the Agent is SSL-enabled, the non-secure port redirects communication to the secure port. In both cases, the appropriate configuration is to specify the non-secure port for provisioned hosts.

6. Specify the non-SSL connector for Oracle Endeca Workbench.

In the `ConfigManager` component, the property `webStudioPort` should specify the non-secure connector for the Endeca Tools Service, as communication with Oracle Endeca Workbench configuration store always uses the unsecured channel.

The following excerpt from the `AppConfig.xml` document shows a sample configuration for an SSL-enabled application.

```
<!--
#####
# EAC Application Definition
#
-->
<app appName="test" eacHost="localhost" eacPort="8888"
  dataPrefix="test" sslEnabled="true" lockManager="LockManager">
  <working-dir>${ENDECA_PROJECT_DIR}</working-dir>
  <log-dir>./logs</log-dir>
</app>

<!--
#####
# Lock Manager - Used to set/remove/test flags and obtain/release locks
#
-->
<lock-manager id="LockManager" releaseLocksOnFailure="true" />

<!--
#####
# Content Acquisition System Server
#
<custom-component id="CAS" host-id="CASHost" class="com.Oracle Endeca.eac.toolk-
it.component.cas.ContentAcquisitionServerComponent">
  <properties>
    <property name="casHost" value="localhost" />
    <property name="casPort" value="8505" />
  </properties>
</custom-component>

-->
```

## Displaying the Deployment Template version

You can print out the version number of the Deployment Template from the command line.

The `runcommand` script has a `--version` flag that prints the version number of the Deployment Template and exits. The command actually prints the version number of the EAC Development Toolkit.

Displaying the version is important for troubleshooting purposes.

To display the version of the Deployment Template:

1. From a command prompt, navigate to the `<app dir>\control` directory on Windows (`<app dir>/control` on UNIX).
2. Run the `runcommand` script with the `--version` flag, as in this Windows example:

```
C:\Endeca\Apps\control>runcommand --version
```

The command prints the version, as in this sample output:

```
Deployment Template: 4.0.0
```



# Configuring an EAC Application

This section provides an overview of the elements defined in `AppConfig.xml`.

## About configuring an EAC application

The standard processing and script operations of the Deployment Template are sufficient to support the operational requirements of most projects. Some applications require customization to enable custom processing steps, script behavior, or even directory structure changes.

Developers are encouraged to use the template as a starting point for customization. The scripts and modules provided with the template incorporate Oracle's best practice recommendations for synchronization, archiving, and update processing. The Deployment Template is intended to provide a set of standards on which development should be founded, while allowing the flexibility to develop custom scripts to meet specific project needs.

There are two ways to configure an EAC application:

- Configure `AppConfig.xml` files. The simplest form of configuration consists of editing the `AppConfig.xml` configuration file and its associated configuration files to change the behavior of components or to add or remove components.

This type of configuration includes the addition or removal of Dgraphs to the main cluster or even the creation of additional clusters. In addition, this category includes adjustment of process arguments (for example, adding a Java classpath for the Forge process in order to enable the use of a Java Manipulator), custom properties and directories (for example, changing the number of index archives that are stored on the indexing server).

- Change behavior of existing BeanShell scripts. Scripts are written in the Java scripting language BeanShell. Scripts are defined in the `AppConfig.xml` document and are interpreted at runtime by the BeanShell interpreter. This allows developers and system administrators to adjust the behavior of the baseline, partial, and configuration update scripts by simply modifying the configuration document.

## About the application configuration files

The application configuration file `<app_dir>/config/script/AppConfig.xml` and its associated files define the hosts, components, and scripts that make up an EAC application and the that orchestrate updates by executing the defined components.

The Deployment Template provides a single `AppConfig.xml` file that contains pointers to refer to other files that define distinct parts of an application, separate scripts from component provisioning, and are used for other purposes. The full set of application configuration files are as follows:

- `AssemblerConfig.xml` - Specifies the application server cluster that allows for quick updating of all Assemblers.
- `AuthoringDgraphCluster.xml` - Specifies the Dgraphs used in the authoring environment and a script that pushes configuration from Workbench to each Dgraph in the authoring cluster.
- `DataIngest.xml` - Specifies data processing scripts, including the baseline update script, partial update script, and the components to perform data processing such as CAS or Forge and Dgidx.
- `DgraphDefaults.xml` - Specifies default values that are inherited by all Dgraph components. These values include host IDs, data processing paths, and Dgraph flags.
- `InitialSetup.xml` - Specifies scripts to perform initial setup tasks, such as uploading initial configuration to Workbench.
- `LiveAppServerCluster.xml` - Specifies your application server clusters, the servers within each cluster, and the applications running on a given Assembler.
- `LiveDgraphCluster.xml` - Specifies the Dgraphs used in the live environment and a script that pushes configuration from Workbench to each Dgraph in the live cluster.
- `ReportGeneration.xml` - Specifies the hosts used for logging and report generations, and several scripts that produce log files at different time intervals.
- `UsageCollectionConfig.xml` - Specifies the Dgraph clusters and application server clusters from which usage is collected.
- `WorkbenchConfig.xml` - Specifies the Endeca Configuration Repository component, the Workbench Manager component, and a script that promotes content from the authoring environment to the live environment.

In addition to these files, any number of `--app-config` arguments may be specified to the Controller class in the EAC development toolkit. All of the objects in the files will be read and processed and scripts can refer to components, hosts, or other scripts defined in other files.

## About the schema for AppConfig.xml

The `eacToolkit.xsd` schema determines the valid syntax within `AppConfig.xml`.

The `eacToolkit.xsd` file is located at the top level of the `eacHandlers.jar` archive file. If any of your Deployment Template scripts fail due to XML syntax errors, you can look at the schema to learn which syntax options for attributes and values are allowed. You may decide to modify the schema to allow you to specify the options you need.

This archive file resides in the `config/lib/java` sub-directory of a deployed application. It also resides in the `data/eac-java/common/config/lib/java` directory of the Deployment Template installation.

To explore this file, use the following command at a prompt from the directory containing the `eacHandlers.jar` archive file:

- On UNIX: `$ENDECA_ROOT/j2sdk/bin/jar xvf eacHandlers.jar eacToolkit.xsd`
- On Windows: `%ENDECA_ROOT%\j2sdk\bin\jar xvf eacHandlers.jar eacToolkit.xsd`

## Configuring the application configuration files

This topic guides you through the process of configuring an EAC application.

1. Edit the `AppConfig.xml` file in `<app dir>/config/script` to reflect the details of your environment. Specifically, set the following values:
  - Specify the `eacHost` and `eacPort` attributes of the `app` element with the correct host and port of the EAC Central Server.
  - Specify the `host` elements with the correct host name or names and EAC ports of all EAC Agents in your environment.
  - Specify the `WorkbenchManager` component with the correct host and port for Oracle Endeca Workbench.
2. If necessary, edit the `InitialSetup.xml` file in `<app dir>/config/script`. This file does not usually require any modifications.
3. Edit the `DataIngest.xml` file in `<app dir>/config/script` to reflect your data processing requirements. Specifically, ensure that the baseline update script and partial update script are correct and that the CAS or Forge and Dgidx components are correctly configured.
4. Edit the `DgraphDefaults.xml` file in `<app dir>/config/script` with the default values that are inherited by all Dgraph components in both the authoring cluster and live cluster.
5. Edit the `AuthoringDgraphCluster.xml` file in `<app dir>/config/script` to ensure the authoring Dgraph, the authoring cluster and post-startup script is correct for your environment.
6. Edit the `LiveDgraphCluster.xml` file in `<app dir>/config/script` to ensure the live Dgraph, the live cluster and post-startup script is correct for your environment.
7. Edit the `LiveAppServerCluster.xml` file in `<app dir>/config/script` to ensure that the application server clusters, the servers within each cluster, and the applications running on the servers are correct for your environment.

- a) For each server cluster, create an `<app-server-cluster>` element with an `id` attribute that corresponds to the cluster name.

For example:

```
<app-server-cluster id="LiveAppServerCluster">
</app-server-cluster>
```

- b) For each server within the cluster, create an `<app-server>` element with the following attributes:

- `id` — The name of the server.
- `hostName` — The DNS name or IP address of the server hosting the Assembler.
- `port` — The port on which the Assembler Web application is running.

For example:

```
<app-server id="LiveDiscover" hostName="assemblerHost.example.com"
port="8006">
</app-server>
```

- c) For each application running on a given Assembler, create a `<web-app>` element with the following attributes:

- `id` — The name of the Assembler application.
- `contextPath` — The path to the application relative to the Assembler server.
- `sslEnabled` — Optionally, whether the application is SSL-enabled.

For example:

```
<web-app id="DiscoverWebApp" contextPath="/discover" sslEnabled="true" />
<web-app id="DiscoverAsService" contextPath="/discoverAsService" />
```

- d) Add the `<web-app>` elements to their respective `<app-server>`s as referenced elements.

For example:

```
<app-server id="LiveDiscover" hostName="assemblerHost" port="8006">
  <web-app ref="DiscoverWebApp" />
  <web-app ref="DiscoverAsService" />
</app-server>
```

- e) Add the `<app-server>` elements to their respective `<app-server-cluster>`s as referenced elements.

For example:

```
<app-server-cluster id="LiveAppServerCluster">
  <app-server ref="LiveDiscover" />
</app-server-cluster>
```

8. Edit the `AssemblerConfig.xml` file in `<app dir>/config/script` to ensure that it references the application server clusters that are correct for your environment.
9. Edit the `WorkbenchConfig.xml` file in `<app dir>/config/script` to ensure the Workbench Manager and IFCR components are correct for your environment.
10. Edit the `UsageCollectionConfig.xml` file in `<app dir>/config/script` to ensure that the Dgraph clusters and application server clusters from which usage is collected are correct for your environment.
11. If necessary, edit the `ReportGeneration.xml` file in `<app dir>/config/script`. This file does not usually require any modifications.

The following topics describe the components that you can define in the application configuration files.

## Global application settings

This first section of the application configuration file defines global application-level configuration, including the host and port of the EAC Central Server, the application name and whether or not SSL is to be used when communicating with the EAC Central Server.

In addition, a default working and log directory are specified and a default `lockManager` is specified for use by other elements defined in the document. All elements inherit these settings or override them.

```
<!--
#####
# Global variables
#
-->
<app appName="MyApp" eacHost="myhost1.company.com" eacPort="8888"
  dataPrefix="MyApp" sslEnabled="false" lockManager="LockManager">
  <working-dir>C:\Endeca\MyApp</working-dir>
  <log-dir>./logs/baseline</log-dir>
</app>
```

## Hosts

All servers in a deployment are enumerated in the host definition portion of the document.

Each host must be given a unique ID. The port specified for each host is the port on which the EAC Agent is listening, which is the Endeca HTTP Service port on that server. This example shows a host defined to run CAS and a host to run the MDEX Engine.

```
<!--
#####
# Servers/hosts
#
```



```
-->
<host id="CASHost" hostName="myhost1.company.com" port="8888" />
<host id="MDEXHost" hostName="myhost2.company.com" port="8888" />
```

## Lock Manager

The LockManager component is used to obtain and release locks and to set and remove flags using the EAC's synchronization Web service.

A LockManager object is associated with the elements in the application to enable a centralized access point to locks, allowing multiple objects to test for the existence of locks and flags. When a script or component invocation fails, the Deployment Template attempts to release all locks acquired during the invocation for a LockManager configured to release locks on failure. Multiple LockManager components may be configured, if it is appropriate for some locks to be released on failure while others remain.

```
<!--
#####
# Lock manager, used to set/remove/test flags and obtain/release
# locks
#
-->
<lock-manager id="LockManager" releaseLocksOnFailure="true" />
```

## Fault tolerance and polling interval properties

Two sets of configurable properties set the behavior of the Deployment Template fault tolerance mechanism and the frequency of status checks for components.

### Fault tolerance property

You can now configure fault tolerance (i.e., retries) for any component (such as Forge, Dgidx, and Dgraph) when invoked through the EAC. This functionality also extends to the CAS server when running a crawl with the CAS component. The name of the fault-tolerance property is `maxMissedStatusQueriesAllowed`.

When components are run, the Deployment Template instructs the EAC to start a component, then polls on a regular interval to check if the component is running, stopped, or failed. If one of these status checks fails, the Deployment Template assumes the component has failed and the script ends. The `maxMissedStatusQueriesAllowed` property allows a configurable number of consecutive failures to be tolerated before the script will end.

The following is an example of a Forge component configured to tolerate a maximum of ten consecutive failures:

```
<forge id="Forge" host-id="ITLHost">
  <properties>
    <property name="numStateBackups" value="10"/>
    <property name="numLogBackups" value="10"/>

    <property name="maxMissedStatusQueriesAllowed" value="10"/>
  </properties>
  ...
</forge>
```

The default number of allowed consecutive failures is 5. Note that these status checks are consecutive, so that every time a status query returns successfully, the counter is reset to zero.

Keep in mind that you can use different fault-tolerance settings for your components. For example, you could set a value of 10 for the Forge component, a value of 8 for Dgidx, and a value of 6 for the Dgraph.

## Polling interval properties

As described in the previous section, the Deployment Template polls on a regular interval to check if a started component is running, stopped, or failed. A set of four properties is available to configure each component for how frequently the Deployment Template polls for status while the component is running. Because each property has a default value, you can use only those properties that are important to you.

The polling properties are as follows:

- `minWaitSeconds` specifies the threshold (in seconds) when slow polling switches to standard (regular) polling. The default is `-1` (i.e., no threshold, so the standard polling interval is used from the start).
- `slowPollingIntervalMs` specifies the interval (in milliseconds) that status queries are sent as long as the `minWaitSeconds` time has not elapsed. The default slow polling interval is **60** seconds.
- `standardPollingIntervalMs` (specified in milliseconds) is used after the `minWaitSeconds` time has passed. If no `minWaitSeconds` setting is specified, the `standardPollingIntervalMs` setting is always used. The default standard polling interval is **1** second.
- `maxWaitSeconds` specifies the threshold (in seconds) when the Deployment Template gives up asking for status and assumes that it has failed. The default is `-1` (i.e., no threshold, so the Deployment Template will keep trying indefinitely).

Here is an example configuration for a long-running Forge component that typically takes 8 hours to complete:

```
<forge id="Forge" host-id="ITLHost">
  <properties>
    <property name="numStateBackups" value="10"/>
    <property name="numLogBackups" value="10"/>

    <property name="standardPollingIntervalMs" value="60000"/>
    <property name="slowPollingIntervalMs" value="600000"/>
    <property name="minWaitSeconds" value="28800"/>
    <property name="maxMissedStatusQueriesAllowed" value="10"/>
  </properties>
  ...
</forge>
```

The result of this configuration would be that for the first 8 hours (`minWaitSeconds=28800`), Forge's status would be checked every 10 minutes (`slowPollingIntervalMs=600000`), after which time the status would be checked every minute (`standardPollingIntervalMs=60000`). If a status check fails, a maximum of 10 consecutive retries will be attempted, based on the `standardPollingIntervalMs` setting.

Keep in mind that these values can be set independently for each component.

## Fault tolerance and polling interval for utilities

Fault tolerance and polling interval values can also be set for these utilities:

- copy
- shell
- archive
- rollback

You set the new values by adjusting the BeanShell script code that is used to construct and invoke the utility. You adjust the code by using these setter methods from the EAC Toolkit's `Utility` class:

- `Utility.setMinWaitSeconds()`
- `Utility.setMaxWaitSeconds()`
- `Utility.setMaxMissedStatusQueriesAllowed()`
- `Utility.setPollingIntervalMs()`
- `Utility.setSlowPollingIntervalMs()`

- `Utility.setMaxMissedStatusQueriesAllowed()`

If you do not use any of these methods, then the utility will use the default values listed in the two previous sections.

For example, here is a default utility invocation in the CAS crawl scripts:

```
// create the target dir, if it doesn't already exist
mkdirUtil = new CreateDirUtility(CAS.getAppname(),
    CAS.getEacHost(), CAS.getEacPort(), CAS.isSslEnabled());
mkdirUtil.init(Forge.getHostId(), destDir, CAS.getWorkingDir());
mkdirUtil.run();
```

You would then add these methods before calling the `run()` method, so that the code would now look like this:

```
// create the target dir, if it doesn't already exist
mkdirUtil = new CreateDirUtility(CAS.getAppname(),
    CAS.getEacHost(), CAS.getEacPort(), CAS.isSslEnabled());
mkdirUtil.init(Forge.getHostId(), destDir, CAS.getWorkingDir());
mkdirUtil.setMinWaitSeconds(30);
mkdirUtil.setMaxWaitSeconds(120);
mkdirUtil.setMaxMissedStatusQueriesAllowed(10);
mkdirUtil.setPollingIntervalMs(5000);
mkdirUtil.setSlowPollingIntervalMs(30000);
mkdirUtil.run();
```

Alternatively, if your utility was defined in your `AppConfig.xml` like this:

```
<copy id="MyCopy" src-host-id="ITLHost" dest-host-id="MDEXHost" recursive="true">
    <src>./path/to/files</src>
    <dest>./path/to/target</dest>
</copy>
```

You would add the same type of lines as above, before calling the `run()` method; for example:

```
MyCopy.setMaxMissedStatusQueriesAllowed(10);
MyCopy.run();
```

For more information on the `Utility` methods, see the Javadocs for the EAC Toolkit package.

## CAS Server

The Deployment Template provides support for running CAS crawls with the CAS Server Component. A CAS Server component is implemented as a `custom-component`. You configure the component according to the output type of a crawl. The sections below describe the common configuration properties, the output-type configuration properties, and then provide examples for each output type including Record Store output, MDEX-compatible output, and record file output.



**Note:** The Deployment Template cannot create a new CAS crawl. You create a crawl using CAS and run it using the Deployment Template. For details about creating a crawl, see the *CAS Developer's Guide*.

### The `custom-component` configuration properties

The `custom-component` configuration properties identify the CAS server in the Servers/hosts section of `AppConfig.xml`. The properties are defined as follows:

- `id` - Assigns a unique ID to a specific CAS Server.

- `host-id` - Points back to the `id` attribute of the `host` global configuration element.
- `class` - Specifies the class that implements the `ContentAcquisitionServerComponent`. Specify `class="com.endeca.eac.toolkit.component.cas.ContentAcquisitionServerComponent"`.

### Common configuration properties

The common configuration properties describe the host and port running CAS. The properties are defined as follows:

- `casHost` - Host name of the server on which the Content Acquisition System is running.
- `casPort` - Port on which the Endeca CAS Service listens. If the application is running in SSL mode, the `casPort` is the SSL port of the Endeca CAS Service. The port number must match the `com.endeca.cas.port` value that is used in the CAS Service configuration script. Or, if the Endeca CAS Service is configured for SSL, then the port number must match `com.endeca.cas.ssl.port` value. The configuration script is in `<install path>\CAS\workspace\conf\jetty.xml`.
- `httpSocketTimeout` is the maximum period of inactivity in milliseconds between two consecutive data packets before http times out.

### Configuration properties specific to MDEX-compatible output

The configuration properties for MDEX-compatible output are defined as follows:

- `numPartialsBackups` - Indicates the number of backups to keep for the cumulative partials directory (`cumulativePartialsDir`). If this property is not configured, then no backups are retained.
- `cumulativePartialsDir` - Indicates the directory on the CAS host where partial MDEX output should be accumulated. This allows partial updates to be reapplied in the event of a failure while applying partial updates.
- `numDvalIdMappingsBackups` - Indicates the number of backups to keep for the dimension value ID mappings file. This allows you to restore dimension value ID mappings if the CAS host fails. If this property is not configured, then five backups are retained. If set to zero, then no backing up is performed.
- `dvalIdMappingsArchiveDir` - Indicates the directory where the dimension value ID mappings files are stored. If this property is not configured, then mappings are written to `./data/dvalid_mappings_archive`. However, to provide more secure backups, Oracle recommends that you specify a network drive that is available to CAS but not the same as the CAS host.

### Example

This example CAS Server component is configured for MDEX-compatible output:

```
<!--
#####
# Content Acquisition System Server
#
-->
<custom-component id="CAS" host-id="ITLHost" class="com.endeca.eac.toolkit.com-
ponent.cas.ContentAcquisitionServerComponent">
  <properties>
    <property name="casHost" value="localhost" />
    <property name="casPort" value="8500" />
    <property name="httpSocketTimeout" value="180000" />
    <property name="numPartialsBackups" value="5" />
    <property name="numDvalIdMappingsBackups" value="5" />
  </properties>
  <directories>
    <directory name="cumulativePartialsDir">./data/partial/cumulative_par-
tials</directory>
    <directory name="dvalIdMappingsArchiveDir">./data/dvalid_mappings_archive</di-
```

```
rectory>
  </directories>
</custom-component>
```

### Configuration properties specific to Record Store output

There are no additional configuration properties required for crawls that write to a Record Store instance. Only the `custom-component` and common configuration properties are required.

#### Example

This example CAS Server component is configured for Record Store output:

```
<!--
#####
# Content Acquisition System Server
#
<custom-component id="CAS" host-id="CASHost" class="com.endeca.eac.toolkit.compo-
nent.cas.ContentAcquisitionServerComponent">
  <properties>
    <property name="casHost" value="localhost" />
    <property name="casPort" value="8500" />
    <property name="httpSocketTimeout" value="180000" />
  </properties>
</custom-component>
-->
```

### Configuration properties specific to record file output

The configuration properties are defined as follows:

- `casCrawlFullOutputDestDir` - Indicates the destination directory to which the crawl output file will be copied after a baseline crawl. Note that this is not the directory to which the CAS crawl writes its output; that output directory is set as part of the crawl configuration.
- `casCrawlIncrementalOutputDestDir` - Indicates the destination directory to which the crawl output file will be copied after an incremental crawl. As with the previous property, this is not the directory to which the CAS crawl writes its output. If you run incremental crawls, the default settings assume that the output format will be compressed binary files.
- `casCrawlOutputDestHost` - Indicates the ID of the host on which the destination directories (specified by the previous two properties) reside.

#### Example

This example CAS Server component is configured for a record file output:

```
<!--
#####
# Content Acquisition System Server
#
-->
<custom-component id="CAS" host-id="CASHost" class="com.endeca.soleng.eac.toolk-
it.component.ContentAcquisitionServerComponent">
  <properties>
    <property name="casHost" value="localhost" />
    <property name="casPort" value="8500" />
    <property name="httpSocketTimeout" value="180000" />
    <property name="casCrawlFullOutputDestDir" value="./data/com-
plete_cas_crawl_output/full" />
    <property name="casCrawlIncrementalOutputDestDir" value="./data/com-
plete_cas_crawl_output/incremental" />
```

```

    <property name="casCrawlOutputDestHost" value="CASHost" />
  </properties>
</custom-component>

```

## Forges

One or many Forge components are defined for baseline update processing and partial update processing depending on the deployment type you choose.

If necessary, you can define a Forge cluster component to apply actions to an entire cluster of Forges, rather than manually iterating over a number of Forges. You could use this feature to run several instances of Forge in parallel to process large joins.

In addition, the object contains logic associated with executing Forges in parallel based on Forge groups, which are described below. Multiple Forge clusters can be defined, with no restriction around which Forges belong to each cluster or how many clusters a Forge belongs to.

A Forge cluster is configured with references to all Forges that belong to that cluster. In addition, the cluster can be configured to copy data in parallel or serially. This setting applies to copies that are performed to retrieve source data and configuration to each server that hosts a Forge component. By default, the template sets this value to true.

```

<!--
#####
# Forge Cluster
#
-->
<forge-cluster id="ForgeCluster" getDataInParallel="true">
  <forge ref="ForgeServer" />
  <forge ref="ForgeClient1" />
  <forge ref="ForgeClient2" />
</forge-cluster>

```

In addition to standard Forge configuration settings and process arguments, the Deployment Template uses several configurable properties and custom directories during processing:

- `numLogBackups` - Number of log directory backups to store.
- `numStateBackups` - Number of autogen state directory backups to store.
- `numPartialsBackups` - Number of cumulative partials directory backups to store. It is recommended that you increase the default value of 5. The reason is that the files in the updates directory for the Dgraph are automatically deleted after partials are applied to the Dgraph. The number you choose depends on how often you run partial updates and how many copies you want to keep.
- `incomingDataHost` - Host to which source data files are extracted.
- `incomingDataDir` - Directory to which source data files are extracted.
- `incomingDataFileName` - Filename of the source data files that are extracted.
- `configHost` - Host from which configuration files and dimensions are retrieved for Forge to process.
- `configDir` - Directory from which configuration files and dimensions are retrieved for Forge to process.
- `cumulativePartialsDir` - Directory where partial updates are accumulated between baseline updates.
- `wsTempDir` - Temp Oracle Endeca Workbench directory to which post-Forge dimensions are copied to be uploaded to the Workbench.
- `skipTestingForFilesDuringCleanup` - Used for directory-cleaning operations. If set to "true", will skip the directory-contents test and instead proceed directly to cleaning the directory. The default behavior is to test the directory contents and skip cleanup if the directory is not empty.
- The properties documented in the "Fault tolerance and polling interval properties" topic.

This excerpt combines properties from both the baseline and partial update Forge to demonstrate the use of all of these configuration settings.

```
<properties>
  <property name="forgeGroup" value="A" />
  <property name="incomingDataHost">ITLHost</property>
  <property name="incomingDataFileName">project_name-part0-*</property>
  <property name="configHost">ITLHost</property>
  <property name="numStateBackups" value="10" />
  <property name="numLogBackups" value="10" />
  <property name="numPartialsBackups" value="5" />
  <property name="skipTestingForFilesDuringCleanup" value="true" />
</properties>
<directories>
  <directory name="incomingDataDir">./data/partials/incoming</directory>
  <directory name="configDir">./config/pipeline</directory>
  <directory name="cumulativePartialsDir">
    ./data/partials/cumulative_partials
  </directory>
  <directory name="wsTempDir">./data/web_studio/temp</directory>
</directories>
```

In addition to standard Forge configuration and process arguments, Forge processes add a custom property used to define which Forge processes run in parallel with each other when they belong to a Forge cluster.

`forgeGroup` - Indicates the Forge's membership in a Forge group. When the run method on a Forge cluster is executed, Forge processes within the same Forge group are run in parallel. Forge group values are arbitrary strings. The Forge cluster iterates through the groups in alphabetical order, though non-standard characters may result in groups being updated in an unexpected order.

## Dgidxs

One or many Dgidx components are defined depending on the deployment type you choose.

If necessary, you can define a Dgidx cluster to apply actions to an entire cluster of Dgidxs, rather than manually iterating over a number of Dgidxs. In addition, the object contains logic associated with executing Dgidxs in parallel based on Dgidx groups, which are described below. Multiple indexing clusters can be defined, with no restriction around which Dgidx belongs to each cluster or how many clusters a Dgidx belongs to.

An indexing cluster is configured with references to all Dgidxs that belong to that cluster. In addition, the cluster can be configured to copy data in parallel or serially. This setting applies to copies that are performed to retrieve source data and configuration to each server that hosts a Dgidx component. By default, the template sets this value to true.

```
<!--
#####
# Indexing Cluster
#
-->
<indexing-cluster id="IndexingCluster" getDataInParallel="true">
  <dgidx ref="Dgidx1" />
  <dgidx ref="Dgidx2" />
</indexing-cluster>
```

In addition to standard Dgidx configuration settings and process arguments, the Deployment Template uses several configurable properties and custom directories during processing:

- `numLogBackups` - Number of log directory backups to store.
- `numIndexbackups` - Number of index backups to store.

- `incomingDataHost` - Host to which source data files are extracted.
- `incomingDataDir` - Directory to which source data files are extracted.
- `incomingDataFileName` - Filename of the source data files that are extracted.
- `configHost` - Host from which configuration files and dimensions are retrieved for Dgidx to process.
- `configDir` - Directory from which configuration files and dimensions are retrieved for Dgidx to process.
- `configFileName` - Filename of the configuration files and dimensions that are retrieved for Dgidx to process.
- `skipTestingForFilesDuringCleanup` - Used for directory-cleaning operations. If set to "true", will skip the directory-contents test and instead proceed directly to cleaning the directory. The default behavior is to test the directory contents and skip cleanup if the directory is not empty.
- The properties documented in the "Fault tolerance and polling interval properties" topic.

In addition to standard Dgidx configuration and process arguments, Dgidx processes add a custom property used to define which Dgidx processes run in parallel with each other when they belong to an indexing cluster.

`dgidxGroup` - Indicates the Dgidx's membership in a Dgidx group. When the run method on an indexing cluster is executed, Dgidx processes within the same Dgidx group are run in parallel. Dgidx group values are arbitrary strings. The indexing cluster iterates through the groups in alphabetical order, though non-standard characters may result in groups being updated in an unexpected order.

## Dgraphs

If a Dgraph deployment type is chosen, a Dgraph cluster component is defined.

This object is used to apply actions to an entire cluster of Dgraphs, rather than manually iterating over a number of Dgraphs. In addition, the object contains logic associated with Dgraph restart strategies, which are described below. Multiple Dgraph clusters can be defined, with no restriction around which Dgraphs belong to each cluster or how many clusters a Dgraph belongs to.

A Dgraph cluster is configured (via the `dgraph-cluster` element) with references to all Dgraphs that belong to that cluster. In addition, the cluster can be configured to copy data in parallel or serially. This setting applies to copies that are performed to distribute a new index, partial updates or configuration updates to each server that hosts a Dgraph. By default, the template sets this value to true.

```
<!--
#####
# Dgraph Cluster
#
-->
<dgraph-cluster id="DgraphCluster" getDataInParallel="true">
  <dgraph ref="Dgraph1" />
  <dgraph ref="Dgraph2" />
</dgraph-cluster>
```

Two Dgraphs are defined by the template by default.

### Global Dgraph settings

In order to avoid defining shared configuration for multiple Dgraphs in each Dgraph's XML configuration, the document provides the `dgraph-defaults` element, where shared settings can be configured and inherited (or overridden) by each Dgraph defined in the document. This defaults object specifies a number of custom configuration properties that are used by the update scripts to define operational functionality.

- `numLogBackups` - Number of log directory backups to store.
- `shutdownTimeout` - Number of seconds to wait for a component to stop (after receiving a stop command).



- `numIdleSecondsAfterStop` - Number of seconds to pause/sleep after a component is stopped. Typically, this will be used to ensure that log file locks are release by the component before proceeding.
- `srcIndexDir` - Location from which a new index will be copied to a local directory on the Dgraph's host.
- `srcIndexHostId` - Host from which a new index will be copied to a local directory on the Dgraph's host.
- `localIndexDir` - Local directory to which a single copy of a new index is copied from the source index directory on the source index host.
- `srcPartialsDir` - Location from which a new partial update will be copied to a local directory on the Dgraph's host.
- `srcCumulativePartialsDir` - Location from which all partial updates accumulated since the last baseline update will be copied to a local directory on the Dgraph's host.
- `srcPartialsHostId` - Host from which partial updates will be copied to a local directory on the Dgraph's host.
- `localCumulativePartialsDir` - Local directory to which partial updates are copied from the source (cumulative) partials directory on the source partials host.
- `srcDgraphConfigDir` - Location from which Dgraph configuration files will be copied to a local directory on the Dgraph's host.
- `srcDgraphConfigHostId` - Host from which Dgraph configuration files will be copied to a local directory on the Dgraph's host.
- `localDgraphConfigDir` - Local directory to which Dgraph configuration files are copied from the source Dgraph config directory on the source Dgraph config host.
- `srcXQueryHostId` - Host from which XQuery modules will be copied to a local directory on the Dgraph's host.
- `srcXQueryDir` - Location from which XQuery modules will be copied to a local directory on the Dgraph's host.
- `localXQueryDir` - Local directory to which XQuery modules are copied from the source Dgraph XQuery directory on the source Dgraph XQuery modules host.
- `skipTestingForFilesDuringCleanup` - Used for directory-cleaning operations. If set to "true", will skip the directory-contents test and instead proceed directly to cleaning the directory. The default behavior is to test the directory contents and skip cleanup if the directory is not empty.
- The properties documented in the "Fault tolerance and polling interval properties" topic.

```

<!--
#####
# Global Dgraph settings, inherited by all dgraphs
#
-->
<dgraph-defaults>
  <properties>
    <property name="srcIndexDir" value="./data/dgidx_output" />
    <property name="srcIndexHostId" value="ITLHost" />
    <property name="srcPartialsDir" value="./data/partial/forge_output" />
    <property name="srcPartialsHostId" value="ITLHost" />
    <property name="srcCumulativePartialsDir" value="./data/partial/cumulative_partials" />
    <property name="srcCumulativePartialsHostId" value="ITLHost" />
    <property name="srcDgraphConfigDir" value="./data/web_studio/dgraph_config" />
  />
  <property name="srcDgraphConfigHostId" value="ITLHost" />
  <property name="srcXQueryHostId" value="ITLHost" />
  <property name="srcXQueryDir" value="./config/lib/xquery" />
  <property name="numLogBackups" value="10" />
  <property name="shutdownTimeout" value="30" />
  <property name="numIdleSecondsAfterStop" value="0" />
</properties>
<directories>

```

```

    <directory name="localIndexDir">./data/dgraphs/local_dgraph_input</directory>
    <directory name="localCumulativePartialsDir">./data/dgraphs/local_cumulative_partials</directory>
    <directory name="localDgraphConfigDir">./data/dgraphs/local_dgraph_config</directory>
  </directories>
  <args>
    <arg>--threads</arg>
    <arg>2</arg>
    <arg>--spl</arg>
    <arg>--dym</arg>
    <arg>--xquery_path</arg>
    <arg>./data/dgraphs/local_xquery</arg>
  </args>
  <startup-timeout>120</startup-timeout>
</dgraph-defaults>

```

Each Dgraph defined in the document (via the `dgraph` element) inherits from the settings defined in the `dgraph-defaults` element, and also specifies settings that are unique to the Dgraph.



**Note:** As of version 3.1 of the Deployment Template, the `numCacheWarmupSeconds` and `offlineUpdate` properties are ignored (and warning messages generated) because they are not supported in the 6.1.x MDEX Engine.

### Restart and update custom properties

In addition to standard Dgraph configuration and process arguments, the `dgraph` element adds two custom properties that define restart and update strategies:

- `restartGroup`
- `updateGroup`

The `restartGroup` property indicates the Dgraph's membership in a restart group. When applying a new index or configuration updates to a cluster of Dgraphs (or when updating a cluster of Dgraphs with a provisioning change such as a new or modified process argument), the Dgraph cluster object applies changes simultaneously to all Dgraphs in a restart group.

Similarly, the `updateGroup` property indicates the Dgraph's membership in an update group. When applying partial updates, the Dgraph cluster object applies changes simultaneously to all Dgraphs in an update group.

This means that a few common restart strategies can be applied as follows:

- To restart/update all Dgraphs at once: specify the same `restartGroup/updateGroup` value for each Dgraph.
- To restart/update Dgraphs one at a time: specify a unique `restartGroup/updateGroup` value for each Dgraph, or omit one or both of the custom properties on all Dgraphs (causing the template to assign a unique group to each Dgraph).
- To restart/update Dgraphs on each server simultaneously: specify the same `restartGroup/updateGroup` value for each Dgraph on a physical server.
- To restart Dgraphs one at a time but apply partial updates to all Dgraphs at once: specify a unique `restartGroup` value for each Dgraph and specify the same `updateGroup` value for each Dgraph.

```

<dgraph id="Dgraph1" host-id="MDEXHost" port="15000">
  <properties>
    <property name="restartGroup" value="A" />
    <property name="updateGroup" value="a" />
  </properties>

```

```
<log-dir>./logs/dgraphs/Dgraph1</log-dir>
<input-dir>./data/dgraphs/Dgraph1/dgraph_input</input-dir>
<update-dir>./data/dgraphs/Dgraph1/dgraph_input/updates</update-dir>
</dgraph>
```

Restart and update group values are arbitrary strings. The DgraphCluster will iterate through the groups in alphabetical order, though non-standard characters may result in groups being updated in an unexpected order.

## Running scripts

Dgraph components can specify the name of a script to invoke prior to shutdown and the name of a script to invoke after the component is started. These optional attributes must specify the ID of a Script defined in the XML file(s). These BeanShell scripts are executed just before the Dgraph is stopped or just after it is started. The scripts behave identically to other BeanShell scripts, except that they have an additional variable, `invokingObject`, which holds a reference to the Dgraph that invoked the script. This functionality is typically used to implement calls to a load balancer, adding or removing a Dgraph from the cluster as it is updated.

The following example shows two dummy scripts (which just log a message, but could be extended to call out to a load balancer) provisioned to run pre-shutdown and post-startup for Dgraph1.

```
<dgraph id="Dgraph1" host-id="MDEXHost" port="15000"
  pre-shutdown-script="DgraphPreShutdownScript"
  post-startup-script="DgraphPostStartupScript">
  <properties>
    <property name="restartGroup" value="A" />
  </properties>
  <log-dir>./logs/dgraphs/Dgraph1</log-dir>
  <input-dir>./data/dgraphs/Dgraph1/dgraph_input</input-dir>
  <update-dir>./data/dgraphs/Dgraph1/dgraph_input/updates</update-dir>
</dgraph>

<script id="DgraphPreShutdownScript">
  <bean-shell-script>
    <![CDATA[
      id = invokingObject.getElementId();
      hostname = invokingObject.getHost().getHostName();
      port = invokingObject.getPort();
      log.info("Removing dgraph with id " + id + " (host: " + hostname +
        ", port: " + port + ") from load balancer cluster.");
    ]]>
  </bean-shell-script>
</script>

<script id="DgraphPostStartupScript">
  <bean-shell-script>
    <![CDATA[
      id = invokingObject.getElementId();
      hostname = invokingObject.getHost().getHostName();
      port = invokingObject.getPort();
      log.info("Adding dgraph with id " + id + " (host: " + hostname +
        ", port: " + port + ") to load balancer cluster.");
    ]]>
  </bean-shell-script>
</script>
```

The following log excerpt shows these scripts running when a new index is being applied to the dgraph:

```
[03.10.08 10:03:28] INFO: Applying index to dgraphs in restart group 'A'.
[03.10.08 10:03:28] INFO: [MDEXHost] Starting shell utility 'mkpath_dgraph-input-new'.
```

```
[03.10.08 10:03:30] INFO: [MDEXHost] Starting copy utility 'copy_in-
dex_to_temp_new_dgraph_input_dir_for_Dgraph1'.
[03.10.08 10:03:35] INFO: Removing dgraph with id Dgraph1 (host: mdex1.mycompa-
ny.com, port: 15000) from load balancer cluster.
[03.10.08 10:03:35] INFO: Stopping component 'Dgraph1'.
[03.10.08 10:03:37] INFO: [MDEXHost] Starting shell utility 'move_dgraph-in-
put_to_dgraph-input-old'.
[03.10.08 10:03:39] INFO: [MDEXHost] Starting shell utility 'move_dgraph-input-
new_to_dgraph-input'.
[03.10.08 10:03:40] INFO: [MDEXHost] Starting backup utility 'back-
up_log_dir_for_component_Dgraph1'.
[03.10.08 10:03:42] INFO: [MDEXHost] Starting component 'Dgraph1'.
[03.10.08 10:03:45] INFO: Adding dgraph with id Dgraph1 (host: mdex1.mycompany.com,
port: 15000) to load balancer cluster.
[03.10.08 10:03:45] INFO: [MDEXHost] Starting shell utility 'rmdir_dgraph-input-
old'.
```

Note that the `dgraph-default` element can also specify the use of pre-shutdown and post-startup scripts as attributes, allowing all Dgraphs in an application to execute the same scripts. For example:

```
<dgraph-defaults pre-shutdown-script="DgraphPreShutdownScript"
  post-startup-script="DgraphPostStartupScript">
  ...
</dgraph-defaults>
```

## Deploying XQuery modules

The Deployment Template supports the distribution of XQuery modules to each Dgraph in the group. The `<app dir>config/lib/xquery` directory is provided for users to store their XQuery modules. In addition, a `LoadXQueryModules` script (in the `AppConfig.xml` file) distributes the XQuery modules to Dgraph servers and instructs the Dgraphs to load the modules.

The procedure to deploy the XQuery modules is:

1. Make certain that the `dgraph-defaults` section of the `AppConfig.xml` file has the XQuery properties set. These global Dgraph setting properties are `srcXQueryHostId`, `srcXQueryDir`, and `localXQueryDir`.
2. Make certain that the Dgraph `--xquery_path` flag is specified as an argument in the `dgraph-defaults` section.
3. Place all the XQuery code in the `<app dir>/config/lib/xquery` and `<app dir>/config/lib/xquery/lib` directories.
4. Execute the `runcommand` script with the `LoadXQueryModules` argument, as in this Windows example:

```
C:\Endeca\Apps\control>runcommand LoadXQueryModules
```

The XQuery modules are distributed to the Dgraphs in the deployment and they are instructed to reload/compile the modules.

## Specifying arguments for the Dgraphs

Both the `dgraph` and `dgraph-defaults` elements allow you to use the `args` sub-element to pass command-line flags to the Dgraphs. However, if you use an `args` section in both the `dgraph` and `dgraph-defaults` configurations, the results are not cumulative.

Instead, the `args` section for an individual Dgraph completely overrides the `dgraph-defaults` definition (i.e., it does not inherit the parameters that are specified in the `dgraph-defaults` section and then add the ones that are unique for that Dgraph).

## Enabling SSL for the Dgraph

You can configure the Dgraph for SSL by using the following elements to define the certificates to use for SSL:

- `cert-file` specifies the path of the `eneCert.pem` certificate file that is used by the Dgraph to present to any client. This is also the certificate that the Application Controller Agent should present to the Dgraph when trying to talk to the Dgraph.
- `ca-file` specifies the path of the `eneCA.pem` Certificate Authority file that the Dgraph uses to authenticate communications with other Oracle Endeca components.
- `cipher` specifies one or more cryptographic algorithms, one of which Dgraph will use during the SSL negotiation. If you omit this setting, the Dgraph chooses a cryptographic algorithm from its internal list of algorithms. See the *Endeca Commerce Security Guide* for more information

All three elements are first-level children of the `<dgraph-defaults>` element.

The following example shows the three SSL elements being used within the `dgraph-default` element:

```
<dgraph-defaults>
...
  <cert-file>
    C:\Endeca\PlatformServices\workspace\etc\eneCert.pem
  </cert-file>
  <ca-file>
    C:\Endeca\PlatformServices\workspace\etc\eneCA.pem
  </ca-file>
  <cipher>AES128-SHA</cipher>
</dgraph-defaults>
```

## Log server

A LogServer component is defined.

In addition to standard LogServer configuration settings and process arguments, the Deployment Template uses a configurable property for log archiving.

- `numLogBackups` - Number of log directory backups to store.
- `shutdownTimeout` - Number of seconds to wait for a component to stop (after receiving a stop command).
- `numIdleSecondsAfterStop` - Number of seconds to pause/sleep after a component is stopped. Typically, this will be used to ensure that log file locks are release by the component before proceeding.
- `targetReportGenDir` - Directory to which logs will be copied for report generation.
- `targetReportGenHostId` - Host to which logs will be copied for report generation.
- `skipTestingForFilesDuringCleanup` - Used for directory-cleaning operations. If set to "true", will skip the directory-contents test and instead proceed directly to cleaning the directory. The default behavior is to test the directory contents and skip cleanup if the directory is not empty.
- The properties documented in the "Fault tolerance and polling interval properties" topic.

```
<logserver id="LogServer" host-id="ITLHost" port="15010">
  <properties>
    <property name="numLogBackups" value="10" />
    <property name="targetReportGenDir" value="./reports/input" />
    <property name="targetReportGenHostId" value="ITLHost" />
  </properties>
  <log-dir>./logs/logservers/LogServer</log-dir>
  <output-dir>./logs/logserver_output</output-dir>
  <startup-timeout>120</startup-timeout>
  <gzip>>false</gzip>
</logserver>
```

## Report Generators

Four report generator components are defined.

In addition to standard Report Generator configuration settings and process arguments, the Deployment Template uses a configurable property for log archiving, as well as these configurable properties:

- `skipTestingForFilesDuringCleanup` - Used for directory-cleaning operations. If set to "true", will skip the directory-contents test and instead proceed directly to cleaning the directory. The default behavior is to test the directory contents and skip cleanup if the directory is not empty.
- The properties documented in the "Fault tolerance and polling interval properties" topic.

The configuration file includes the name of an output file for each report generator, which defaults to `report.html` or `report.xml`. This file name is never used when the report generation scripts in the `AppConfig.xml` file are used. During execution, the script re-provisions the report generator to output a file named with a date stamp. This means that the provisioning in the file will always be "out of synch" with the provisioning in the EAC. This will result in the Report Generator's definition changing repeatedly as scripts are executed.

```
<report-generator id="WeeklyReportGenerator" host-id="ITLHost">
  <log-dir>./logs/report_generators/WeeklyReportGenerator</log-dir>
  <input-dir>./reports/input</input-dir>
  <output-file>./reports/weekly/report.xml</output-file>
  <stylesheet-file>
    ./config/report_templates/tools_report_stylesheet.xsl
  </stylesheet-file>
  <settings-file>
    ./config/report_templates/report_settings.xml
  </settings-file>
  <time-range>LastWeek</time-range>
  <time-series>Daily</time-series>
  <charts-enabled>true</charts-enabled>
</report-generator>
```

## IFCR

The IFCR is a custom component that specifies user information for an Endeca Configuration Repository that is running inside Oracle Endeca Workbench. The deployment template scripts use the information to connect to an Endeca Configuration Repository and move configuration used by Authoring and Live Dgraphs, the media MDEX reference application, and the IFCR Backup Utility.

You define an IFCR component in the `WorkbenchConfig.xml` file which is then referenced by `AppConfig.xml`.

### The custom-component configuration properties

The `custom-component` configuration properties identify the IFCR in the Data Ingest Hosts section of `DataIngest.xml`.

The properties are defined as follows:

- `id` - Assigns a unique ID to a specific IFCR instance.
- `host-id` - Points back to the `id` attribute of the `host` global configuration element.
- `class` - Specifies the class that implements the `IFCRComponent`. Specify `class="com.endeca.soleng.eac.toolkit.component.IFCRComponent"`.

## IFCR configuration properties

The configuration properties are defined as follows:

- `repositoryUrl` - Specifies host, port, and ifcr directory as `http://<workbench host>:<port>/ifcr`.
- `username` - Name of the user logging in to Oracle Endeca Workbench where the Endeca Configuration Repository is hosted.
- `password` - Corresponding password for the user name.
- `numExportBackups` - Indicates the number of backups to keep for exported configuration of the Endeca Configuration Repository. If this property is not configured, then no backups are retained. The default value is 5.

## Example

This example shows a typical configuration:

```
<!--
#####
# IFCR - A component that interfaces with the Workbench repository.
-->
<custom-component id="IFCR" host-id="ITLHost" class="com.endeca.soleng.eac.toolkit.component.IFCRComponent">
  <properties>
    <property name="repositoryUrl" value="http://localhost:8006/ifcr" />
    <property name="username" value="admin" />
    <property name="password" value="admin" />
    <property name="numExportBackups" value="3" />
  </properties>
</custom-component>
```

## Workbench Manager

The Workbench Manager is a custom component that specifies connection information for Oracle Endeca Workbench and also a configuration directory for Oracle Endeca Workbench. The deployment template scripts use the information to connect to Workbench and update shared configuration contained in the configuration directory.

You define a Workbench Manager component in the `WorkbenchConfig.xml` file which is then referenced by `AppConfig.xml`.

### The custom-component configuration properties

The `custom-component` configuration properties identify the Workbench Manager in the Data Ingest Hosts section of `DataIngest.xml`.

The properties are defined as follows:

- `id` - Assigns a unique ID to a specific Workbench instance.
- `host-id` - Points back to the `id` attribute of the `host` global configuration element.
- `class` - Specifies the class that implements the `WorkbenchManagerComponent`. Specify `class="com.endeca.soleng.eac.toolkit.component.WorkbenchManagerComponent"`.

### Workbench configuration properties

The configuration properties are defined as follows:

- `workbenchHost` - Host name of the server on which Oracle Endeca Workbench is running.

- `workbenchPort` - Port on which Workbench listens. This is the port of the Endeca Tools Service on the Oracle Endeca Workbench host. If the application is running in SSL mode, the `workbenchPort` is the SSL port of Workbench.
- `configDir` - Directory to which Workbench configuration files are uploaded or downloaded by other components in the implementation.
- `workbenchTempDir` - Temporary directory used for Workbench interaction. Post-Forge dimensions are uploaded or downloaded from this directory by other components in the implementation.

### Example

This example shows a typical configuration:

```
<!--
#####
# WorkbenchManager - A component that interfaces with the legacy
# 'web studio' configuration repository. It is used primarily during
# data ingest to load post-forge dimensions into Workbench.
-->
<custom-component id="WorkbenchManager" host-id="ITLHost" class="com.endeca.soleng.eac.toolkit.component.WorkbenchManagerComponent">
  <properties>
    <property name="workbenchHost" value="localhost" />
    <property name="workbenchPort" value="8006" />
  </properties>
  <directories>
    <directory name="configDir">./config/pipeline</directory>
    <directory name="workbenchTempDir">./data/workbench/temp</directory>
  </directories>
</custom-component>
```

## Reporting

Oracle Endeca Workbench provides an interface for viewing and analyzing reports produced by the Report Generator.

In order for Oracle Endeca Workbench to display these reports, report files and associated charts need to be created and delivered to a directory in Oracle Endeca Workbench's workspace. Alternatively, a "webstudio" host can be provisioned with a "webstudio-report-dir" custom directory, which indicates to Oracle Endeca Workbench where it should read reports for the application. In addition, the files need to be named with a date stamp to conform to Oracle Endeca Workbench's naming convention. The Deployment Template includes report generation scripts that perform these naming and copying steps to deliver reports for Oracle Endeca Workbench to read. Common extension or customization of this functionality may occur when one or more of the components in the reporting lifecycle run in different environments. The `AppConfig.xml` allows components to work independently of each other. Specifically, the `LogServer` can be configured to deliver files to an arbitrary directory, from where the files can be copied to another environment for report generation. Similarly, the Report Generator's output report can be delivered to an arbitrary target directory, from where the files can be copied to another environment for display in Oracle Endeca Workbench.

## Configuration Manager

The Configuration Manager component is a custom component that does not correlate to an Oracle Endeca process.



**Note:** In Tools and Frameworks 3.0, the Configuration Manager was deprecated and replaced by Workbench Manager.



Instead, this object implements logic used to manage configuration files. Specifically, the current implementation supports retrieving and merging configuration from Developer Studio with files maintained in Oracle Endeca Workbench.

The following configuration properties and custom directories are used to implement the logic of the Config Manager component.

- `webStudioEnabled` - "true" or "false," indicating whether integration with Oracle Endeca Workbench is enabled.
- `webStudioHost` - Hostname of the server on which Oracle Endeca Workbench is running.
- `webStudioPort` - Port on which Oracle Endeca Workbench listens. This is the port of the Endeca Tools Service on the Oracle Endeca Workbench host.
- `webStudioMaintainedFile*` - Specifies the name of a file that will be maintained in Oracle Endeca Workbench. The ConfigManager respects all properties prefixed with "webStudioMaintainedFile" but requires that all properties have unique names. When configuring files, each should be given a unique suffix. Note that the names of files specified may use wildcards (e.g. `<property name="webStudioMaintainedFile1" value="merch_rule_group_*.xml" />`).
- `devStudioConfigDir` - Directory from which Developer Studio configuration files are retrieved.
- `webStudioConfigDir` - Directory to which Workbench configuration files are downloaded.
- `webStudioDgraphConfigDir` - Directory from which Workbench configuration files are retrieved.
- `mergedConfigDir` - Directory to which merged configuration is copied.
- `webStudioTempDir` - Temporary directory used for Workbench interaction. Post-Forge dimensions are uploaded from this directory to the Workbench.
- `skipTestingForFilesDuringCleanup` - Used for directory-cleaning operations. If set to "true", will skip the directory-contents test and instead proceed directly to cleaning the directory. The default behavior is to test the directory contents and skip cleanup if the directory is not empty.
- The properties documented in the "Fault tolerance and polling interval properties" topic.

```
<!--
#####
# Config Manager. Manages Dev Studio and Workbench config sources.
#
-->
<custom-component id="ConfigManager" host-id="ITLHost"
  class="com.Endeca.soleng.eac.toolkit.component.ConfigManagerComponent">
  <properties>
    <property name="webStudioEnabled" value="true" />
    <property name="webStudioHost" value="localhost" />
    <property name="webStudioPort" value="8006" />
    <property name="webStudioMaintainedFile1"
      value="thesaurus.xml" />
    <property name="webStudioMaintainedFile2"
      value="merch_rule_group_default.xml" />
    <property name="webStudioMaintainedFile3"
      value="merch_rule_group_default_redirects.xml" />
  </properties>
  <directories>
    <directory name="devStudioConfigDir">
      ./config/pipeline
    </directory>
    <directory name="webStudioConfigDir">
      ./data/web_studio/config
    </directory>
    <directory name="webStudioDgraphConfigDir">
      ./data/web_studio/dgraph_config
    </directory>
    <directory name="mergedConfigDir">
```

```

    ./data/complete_index_config
  </directory>
  <directory name="webStudioTempDir">
    ./data/web_studio/temp
  </directory>
</directories>
</custom-component>

```

## Configuring the BeanShell scripts

The following list describes a number of customization approaches that you can implement to extend the existing functionality or add new functionality to the template.

- For example, if a deployment uses JDBC to read data into the Forge pipeline instead of using extracted data files, the following changes would be implemented in the BaselineUpdate script:

1. Remove the line that retrieves data and configuration for Forge: `Forge.getData()`;
2. Insert a new copy command to retrieve configuration for Forge to process:

```

...
// get Workbench config, merge with Dev Studio config
ConfigManager.downloadWsConfig();
ConfigManager.fetchMergedConfig();

// fetch extracted data files, run ITL
srcDir = PathUtils.getAbsolutePath(Forge.getWorkingDir(),
    Forge.getConfigDir()) + "\\*";
destDir = PathUtils.getAbsolutePath(Forge.getWorkingDir(),
    Forge.getInputDir());

dimensionCopy = new CopyUtility(Forge.getAppname(),
    Forge.getEacHost(), Forge.getEacPort(), Forge.isSslEnabled());
dimensionCopy.init("copy_dimensions", Forge.getHostId(),
    Forge.getHostId(), srcDir, destDir, true);
dimensionCopy.run();

Forge.getData();
Forge.run();
Dgidx.run();
...

```

Note that this amended BeanShell script imports two classes from the classpath, references variables that point to elements in the `AppConfig.xml` document (e.g. `Forge`, `Dgidx`) and defines new variables without specifying their type (e.g. `srcDir`, `destDir`). Details about BeanShell scripting can be found in Appendix A of this guide.

- Write new BeanShell scripts - Some use cases may call for greater flexibility than can easily be achieved by modifying existing BeanShell scripts. In these cases, writing new BeanShell scripts may accomplish the desired goal. For example, the following BeanShell script extends the previous example by pulling the new functionality into a separate script:

```

<script id="CopyConfig">
  <bean-shell-script>
    <![CDATA[

// fetch extracted data files, run ITL
srcDir = PathUtils.getAbsolutePath(Forge.getWorkingDir(),

```

```

Forge.getConfigDir() + "\\*";
destDir = PathUtils.getAbsolutePath(Forge.getWorkingDir(),
Forge.getInputDir());

dimensionCopy = new CopyUtility(Forge.getAppName(),
Forge.getEacHost(), Forge.getEacPort(), Forge.isSslEnabled());
dimensionCopy.init("copy_dimensions", Forge.getHostId(),
Forge.getHostId(), srcDir, destDir, true);
dimensionCopy.run();

]]>
</bean-shell-script>
</script>

```

Once the new script is defined, the BaselineUpdate script simplifies to the following:

```

...
// get Workbench config, merge with Dev Studio config
ConfigManager.downloadWsConfig();
ConfigManager.fetchMergedConfig();

// fetch extracted data files, run ITL
CopyConfig.run();
Forge.getData();
Forge.run();
Dgidx.run();
...

```

- Define utilities in `AppConfig.xml` - A common use case for customization is to add or adjust the functionality of utility invocation. Our previous example demonstrates the need to invoke a new copy utility when the Forge implementation changes. Other common use cases involve invoking a data pre-processing script from the shell and archiving a directory. In order to enable this, the Deployment Template allows utilities to be configured in the `AppConfig.xml` document. To configure the copy defined above in the document, use the copy element:

```

<copy id="CopyConfig" src-host-id="ITLHost" dest-host-id="ITLHost"
recursive="true">
  <src>./data/complete_index_config/*</src>
  <dest>./data/processing</dest>
</copy>

```

Once configured, this copy utility is invoked using the same command that was previously added to the BaselineUpdate to invoke the custom BeanShell script: `CopyConfig.run()`;

- Extend the Java EAC Development Toolkit - In rare cases, you may need to implement complex custom functionality that would be unwieldy and difficult to maintain if implemented in the `AppConfig.xml` document. In these cases, you can extend objects in the toolkit to create new Java objects that implement the desired custom functionality. Staying with the previous example, the developer might implement a custom Forge object to change the behavior of the `getData()` method to simply copy configuration without looking for extracted data files.

```

package com.Endeca.soleng.eac.toolkit.component;

import java.util.logging.Logger;
import com.Endeca.soleng.eac.toolkit.exception.*;

public class MyForgeComponent extends ForgeComponent
{
  private static Logger log =
    Logger.getLogger(MyForgeComponent.class.getName());

  protected void getData() throws AppConfiguratonException,

```

```

        EacCommunicationException, EacComponentControlException,
        InterruptedException
    {

        // get dimensions for processing
        getConfig();
    }
}

```

Obviously, this trivial customization is too simple to warrant the development of a new class. However, this approach can be used to override the functionality of most methods in the toolkit or to implement new methods.

In order to use the new functionality, the developer will compile the new class and ensure that it is included on the classpath when invoking scripts. The simplest way to do this is to deploy the compiled `.class` file to the `<app dir>/config/script` directory. Once on the classpath, the new component can be loaded in place of the default Forge component by making the following change to the Forge configuration in `AppConfig.xml`:

```

<forge class="com.Endeca.soleng.eac.toolkit.component.MyForgeComponent"
    id="Forge" host-id="ITLHost">
    ...
</forge>

```

Some types of customization will require more complex configuration. Refer to Appendix A ("EAC Development Toolkit") for information about configuring custom Java classes using the Spring Framework namespace in the `AppConfig.xml` document.

## Configuration overrides

The Deployment Template allows the use of one or more configuration override files.

These files can be used to override or substitute values into the configuration documents. For example, developers may want to separate the specification of environment-specific configuration (e.g. hostnames, ports, etc.) from the application configuration and scripts. This may be useful for making configuration documents portable across environments and for dividing ownership of configuration elements between system administrators and application developers.

Override files are specified by using the `--config-override` flag to the EAC development toolkit's controller. For example, the `runcommand` script in the template includes an `environment.properties` file by default, though this file only contains examples of overrides and does not specify any active overrides.

Two types of properties can be specified in an override file:

1. `[object].[field] = [value]` - This style of override specifies the name of an object and field and sets the value for that field, overriding any value specified for that field in the XML configuration document or documents. For example:

```

Dgraph1.port = 16000
Dgraph1.properties['restartGroup'] = B
ITLHost.hostName = itl.mycompany.com

```

2. `[token] = [value]` - This style of override specifies the name of a token defined in the XML config file and substitutes the specified value for that token. For example, if the `AppConfig.xml` defines the following host:

```

<host id="ITLHost" hostName="${itl.host}" port="${itl.port}" />

```

The override can specify the values to substitute for these tokens:

```
itl.host = it.mycompany.com
itl.port = 8888
```

It is important to note that both styles of substitution are attempted for every value defined in the override file. When a token fails to match, a low-severity warning is logged and ignored. This is required because most tokens will only match one of the two styles of substitution. It may be important to avoid using token names that coincide with object names. For example, defining the token `${Forge.tempDir}` will cause the corresponding value to substitute for both the token as well as the `tempDir` field of the Forge component.



# Replacing the Default Forge Pipeline

This chapter describes how to modify or create a Forge pipeline that is designed for use within the deployment template operational structure. This includes pipeline naming requirements, common errors encountered, etc.

**Note:** This chapter only applies to applications that use Forge to process source data. If your application uses CAS to produce MDEX-compatible output, this chapter does not apply.

## About the sample pipelines

For testing purposes, the Deployment Template includes a Developer Studio project with two Forge pipelines (a baseline and a partial). The sample pipelines facilitate testing the deployment template; however, the files should be replaced with project-specific files immediately after a deployed application has been properly configured.

The pipelines are located in `<app_dir>/config/pipeline`. The pipeline for a baseline update processes 10 records, and the pipeline for a partial update that adds 2 more records.

## Sample pipeline overview

This section describes the high-level steps that are necessary to integrate a new/existing pipeline with a deployment template.

Additional detail on each of these steps is provided in later sections.

1. Ensure that the application name and pipeline configuration prefix match the data prefix configured in the deployment template.
2. Place pipeline configuration files in the `<app_dir>/config/pipeline/` directory of the primary server.
3. In order to enable partial updates, ensure that the project is configured with a record spec (i.e., a unique record identifier property).
4. Ensure that any input Record Adapters requiring filenames specify the file location relative to the `<app_dir>/data/processing/` (or `<app_dir>/data/partials/processing`) directory.

## Specifying a pipeline

By default, the Deployment Template checks the `<app_dir>/config/pipeline` for the pipeline to run. This includes baseline updates and partial updates. It is simplest to put your pipeline files in this directory.

Alternatively, the `devStudioConfigDir` attribute in the `ConfigManager` custom component specifies the pipeline to run.

To specify a pipeline to run in `AppConfig.xml`:

1. Ensure that your pipeline files are located in `<app dir>/config/pipeline`.
2. Alternatively, modify the `devStudioConfigDir` property in the `ConfigManager` custom component to reference the pipeline directory.

In this example, the pipeline is stored in the `pipeline` directory:

```
<custom-component id="ConfigManager" host-id="ITLHost"
  class="com.Endeca.soleng.eac.toolkit.component.ConfigManagerComponent">
  <properties>
  ...
  </properties>
  <directories>
    <directory
      name="devStudioConfigDir">./config/pipeline
    </directory>
    ...
  </directories>
```

3. If you modified the value in step 2, also modify the value of the `configDir` attribute in the `Partial update Forge` section to reference the `config/pipeline` directory.

For example:

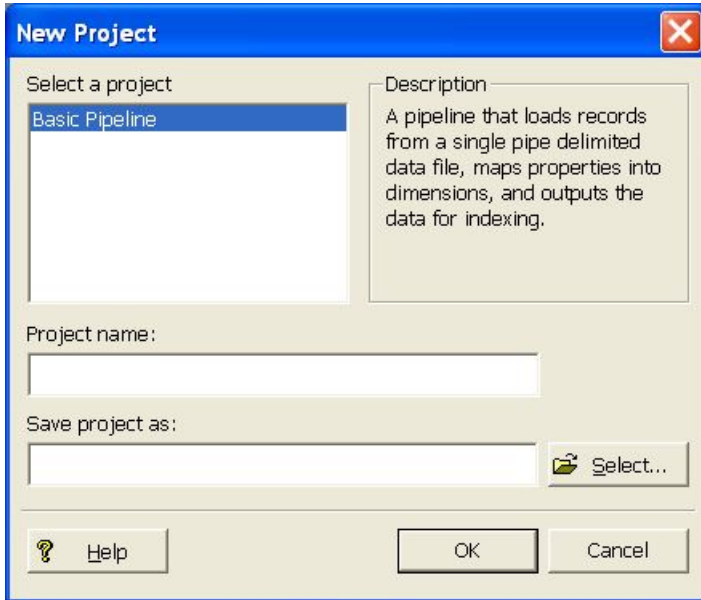
```
<!--
#####
# Partial update Forge
-->
<forge id="PartialForge" host-id="ITLHost">
  <properties>
  ...
  </properties>
  <directories>
    ...
    <directory name="configDir">./config/pipeline</directory>
    ...
  </directories>
```

## Creating a new project

Once the reference configuration files have been deleted, a new pipeline configuration project can be created.

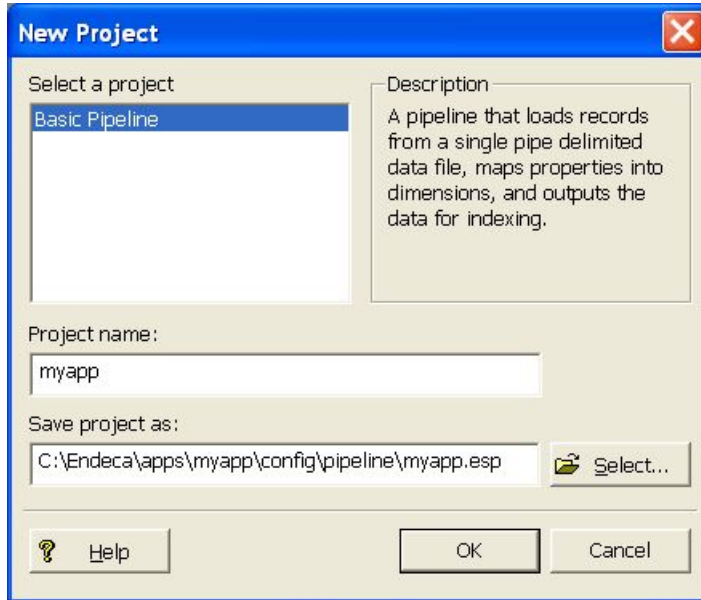
When creating a new project using the Oracle Endeca Developer Studio, you are prompted with the following dialog box:





To create a new project:

1. In order for a new pipeline to be run properly within the deployment template, the following must be properly specified:
  - a) The **Project Name** field must be the same as the data prefix specified for the "app" element in `<app dir>/config/script/AppConfig.xml`. By default, this data prefix will have been set to the name of the application that was specified when running `deploy.bat` or `deploy.sh`.
  - b) Recall that the `[appname]` specified was also used to create the base `<app dir>` directory. For example, if "myapp" was supplied as the `[appname]`, and "c:\Endeca\apps" was supplied as the Deployment Directory, then `<app dir>` would be `c:\Endeca\apps\myapp`. In this example, the Project Name should also be specified as "myapp".
2. The **Save Project As** field should be `<app dir>\config\pipeline\[appname].esp`  
 In the example above, the **Save Project As** field would be `c:\Endeca\apps\myapp\config\pipeline\myapp.esp`.



After clicking the "OK" button, a number of files are created in the `<app_dir>/config/pipeline/` directory. The primary files to be concerned with are listed below:

File name	Description
<code>pipeline.epx</code>	This is the main pipeline file that the deployment template will reference when running forge.
<code>[appname].esp</code>	This is the Developer Studio project file that will be used whenever reopening the project. Although this file does not actually require the <code>[appname]</code> prefix, it is good practice to keep it consistent with other project files.
<code>[appname].*.xml</code>	These are the various configuration files that will be used later by the indexer and MDEX Engine processes. It is important that they have the same prefix as the deployment template Application Name.
<code>dimensions.xml</code>	This is the dimension file referenced by the default Dimension Adapter.

## Modifying an existing project

Modifying an existing Developer Studio project to match a new deployment template application is a somewhat tedious task. In fact, it is often easier to simply create a new deployment template application instead.

The important key is that the `[appname].*.xml` files share the same `[appname]` as the deployment template project. Since there are 30+ XML files, you can either:

- Rename each of the XML files with a new prefix, and update the `[appname].esp` file to reference each new file.

- Update the deployed application's `AppConfig.xml` file to specify the `[appname]` of your configuration files. For example, if your configuration files are named `myapp.*.xml`, update the configuration as follows:

```
<app appName="myapp" eacHost="host1.company.com" eacPort="8888"
  dataPrefix="myapp" sslEnabled="false"
  lockManager="LockManager">
  <working-dir>C:\Endeca\apps\myapp</working-dir>
  <log-dir>./logs/baseline</log-dir>
</app>
```

In most cases, the `appName` attribute and the `dataPrefix` attribute will be identical. However, this is not required and an application can be configured to support files with a data prefix other than the application name. If the data prefix is not specified, the application defaults to using the application name.

Note that opening an existing project in the Oracle Endeca Developer Studio and using the **Save As** feature will not rename the corresponding `*.xml` files. It will only rename the `[appname].esp` file. The prefix for the XML files can only be specified when a new project is created.

### Related Links

[Common errors](#) on page 54

This section provides troubleshooting information for commonly received errors.

## Configuring a record specifier

The deployment includes support for both baseline and partial index updates. In order to support partial updates, an application must include a record specifier, which is a property marked as the unique identifier of records in the index.

For details about the record specifier property, refer to the *Platform Services Forge Guide*.

When configuring your application, identify a property for which each record will have a unique assigned value.

To enable the use of that property as a record spec:

- Open the **Property** dialog box in Developer Studio.
- Check the box labeled "Use for record spec."



## Forge flags

In order to reduce the amount of configuration required to integrate a pipeline into a deployment template, a standard deployment template application runs the primary and partial update Forge processes with an abbreviated set of flags.

Since the deployment template already specifies directory structures and file prefixes, the following flags are used to override a pipeline's input and output components, specifying the appropriate directories and prefixes for either reading or writing data.

### Primary Forge flags

Flag	Description
<code>--inputDir</code>	<code>&lt;app dir&gt;/data/processing</code>
<code>--stateDir</code>	<code>&lt;app dir&gt;/data/state</code>
<code>--tmpDir</code>	<code>&lt;app dir&gt;/data/forge_temp</code>
<code>--logDir</code>	<code>&lt;app dir&gt;/logs/baseline</code>
<code>--outputDir</code>	<code>&lt;app dir&gt;/data/forge_output</code>
<code>--outputPrefix</code>	<code>[dataPrefix]</code>

### Partial update Forge flags

Flag	Description
<code>--inputDir</code>	<code>&lt;app dir&gt;/data/partials/processing</code>
<code>--stateDir</code>	<code>&lt;app dir&gt;/data/state</code>
<code>--tmpDir</code>	<code>&lt;app dir&gt;/data/forge_temp</code>
<code>--logDir</code>	<code>&lt;app dir&gt;/logs/partial</code>
<code>--outputDir</code>	<code>&lt;app dir&gt;/data/partials/forge_output</code>
<code>--outputPrefix</code>	<code>[dataPrefix]</code>

## Input record adapters

The record adapters load the source data.

To start, here is a quick review of how sample data included with the deployment template is processed. The sample application includes a sample dataset in `<app dir>/test_data/baseline` directory. When processing the sample data, the `load_baseline_test_data` script copies the contents of this directory into the `<app dir>/data/incoming/` directory and sets a flag in the EAC.

This flag, named `baseline_data_ready`, indicates to the deployment template scripts that the data extraction process is complete and data is ready for processing. Once that has occurred, the baseline update process copies these files into the `<app dir>/data/processing` directory before running the Forge process.

When using a default deployment template application, it is therefore necessary for all input record adapters to look in the `<app dir>/data/processing` directory for incoming data extracts. The deployment template

handles this automatically by specifying the `--inputDir` flag when running the primary forge process. This flag overrides any absolute path specified for specific input adapters with the proper deployment template path: `<app_dir>/data/processing`. However, the `--inputDir` flag respects relative paths, resolving them relative to the path specified as the input directory.

The URL property of any record adapter component therefore only needs to specify the relative path to a specific file or subdirectory within the `<app_dir>/data/incoming` directory. (Remember that files and subdirectories in the incoming directory are copied to the processing directory by the deployment template before Forge is run.)

For example, if a single extract file called `data.txt` is copied into the `<app_dir>/data/incoming` directory before running a baseline, the URL property of that data's input record adapter should specify a URL of `data.txt`.

For a more complex deployment where, for instance, multiple text extract files are copied into the `<app_dir>/data/incoming/extracted_data` directory before running a baseline update, the URL property of a single input record adapter configured to read these files should be set to `extracted_data/*.txt`.

### Related Links

[Output record adapters](#) on page 53

Output record adapters are often used to generate debug or state information. By default, the location to which this data is written will be overridden by the `--outputDir` flag.

## Dimension adapters

The `--inputDir` flag specified to forge overrides the input URL for dimension adapters.

Since the dimensions for a project are usually stored in the `<app_dir>/config/pipeline` directory along with other configuration files, the deployment template copies these files into the `<app_dir>/data/processing/` directory before running the Forge process. The URLs specified in dimension adapters should follow the same rules as those described for input record adapters, specifying dimension XML file URLs relative to the `--inputDir` directory. In most cases, this is as simple as specifying the URL for the main dimension adapter as `Dimensions.xml`, which is the value used by the default "Dimensions" adapter created by Developer Studio's project template.

More complex deployments that include multiple dimension adapters or external delivery of dimension files should ensure that the dimension XML files are copied into the `<app_dir>/data/incoming/` directory before the forge process runs.

## Indexer adapters

Because the `--outputPrefix` and `--outputDir` flags are both included, the deployment template will override any values specified for the Indexer Adapter "URL" and "Output prefix" properties.


Therefore, it is unnecessary to modify these properties in most cases.

## Output record adapters

Output record adapters are often used to generate debug or state information. By default, the location to which this data is written will be overridden by the `--outputDir` flag.

In most cases, however, it is undesirable for these files to be written to the same location as the Forge output files.

In these cases, an output record adapter can be configured to instead respect the `--stateDir` flag by selecting the **"Maintain State"** checkbox.



Now any files generated by this output record adapter will be written to the `<app_dir>/data/state/` directory.

Note that the output file name must still be specified in the "URL" property of the record adapter. The `--outputPrefix` flag only overrides the indexer adapter output file names, not output record adapter file names.

### Related Links

[Input record adapters](#) on page 52

The record adapters load the source data.

## Dimension servers

The `--stateDir` flag will override the URL value for all Dimension Server components, and place any autogen state files in the `<app_dir>/data/state/` directory.

## Common errors

This section provides troubleshooting information for commonly received errors.

### Unable to Find Pipeline.epx

If Forge fails, check the logs (`<app_dir>/logs/baseline/err.forge`) to make sure that Forge was able to find the `pipeline.epx` file in its proper location. Remember that a basic deployment template application assumes that it will find the project's `pipeline.epx` file in `<app_dir>\config\pipeline\`.

On UNIX platforms, file names are case sensitive. The deployment template expects the primary pipeline file to be named `pipeline.epx` and the partial update pipeline (if one is required for the deployed application) to be named `partial_pipeline.epx`. Ensure that the files in your deployment use this capitalization.

### Missing Configuration Files

This more common error is also more difficult to detect. Since all pipelines created by the Oracle Endeca Developer Studio typically contain a `Pipeline.epx` file, it is unlikely that the Forge process will be unable to find the file, unless it was placed in the wrong directory. If the XML configuration files, however, have a different prefix from the deployment template [`appname`], these files will not be copied into the `<app dir>/data/forge_output/`, `<app dir>/data/dgidx_output/`, and `<app dir>/data/dgraphs/*/dgraph_input/` directories. All processes will likely complete successfully, but any configuration information specified by these XML files, such as search interfaces, business rules, sort keys, etc. will be missing from the resulting MDEX Engine. To correct this problem, check the XML files located in `<app dir>/config/pipeline/` and make sure they have the correct prefix. Also check the directories mentioned above to make sure that these XML files are being properly copied.

### MDEX Engine Fails to Start

If an MDEX Engine fails to start, check the log for the appropriate Dgraph in `<app dir>/logs/dgraphs/[dgraph]/[dgraph].log`. If the log indicates that the Dgraph failed to start because no record specifier was found, follow the steps in this document to create a unique record specifier property for you project.

### Record Adapter Unable to Open File

Another common error may occur if a record adapter is unable to find or open a specified file for either input or output. In this case, the Forge error log (`<app dir>/logs/baseline/err.forge`) should specify which file or directory could not be found. To correct this problem, make sure the files or directories specified by the record adapters correspond to the directory structure established by the deployment template application. Note that this error may be masked if the "Require Data" property is not checked for a given input adapter, since Forge will only log a warning instead of a fatal error.





## Chapter 4

---

# Modifying Index Configuration for an Application

This section describes how to modify index configuration using the Index Configuration Command-line Utility in a CAS-based processing model. If you are using Forge to process updates, this chapter does not apply to your deployment.

## Overview of the Index Configuration Command-line Utility

The Index Configuration Command-line Utility modifies index configuration stored in the Endeca Configuration Repository for an application. This utility is typically used to modify data after it has been exported from a product catalog system and modify the search configuration settings for the data. In many cases, the utility is also used to manually create index configuration that is not part of a product catalog system.

The Index Configuration Command-line Utility is a script named `index-config-cmd` that you run from a command prompt. After you deploy a new application, the `index-config-cmd` script is available in the `<app name>\control` directory.

### Help options

The Index Configuration Command-line Utility has two help options that display the usage syntax. The `--help` option displays a summary of the tasks. The `--help-detail` option displays detailed usage information for a specified task. For example:

```
C:\Endeca\apps\Discover\control>index_config_cmd.bat --help
usage: index-config-cmd <task-name> [options]
where <task-name> is one of the following:
    get-config
    set-config
    delete-owner
    get-merged-config
```

For detailed usage information for individual task options, use `<task-name> --help`

### Command-line options

The command syntax for executing the tasks is:

```
index_config_cmd <task-name> [options]
```

The `<task-name>` argument is the task to be performed by the utility, such as the `get-config` task. The task options vary, depending on the task. However, the following option can be used with any task:

- `-o` (or `--owner`) specifies an import owner for a task. If you specify the `-o` option, the task applies only to the owner specified. The option can have an argument of `all`, `system`, or user-specified owner name. The `all` owner includes both the `system` owner and all user-specified owners. If you omit this option, the task applies to the `system` owner.

## About index configuration ownership

The index configuration for an application is associated with one or more *import owners*. An import owner provides a way of indicating that a portion of an application's configuration came from one source rather than another source. The name of the import owner typically identifies the source. For example, the name of an import owner could reflect a product catalog system.

In addition to any number of import owners, there is also one default `system` owner. The `system` owner is typically a developer who uses the utility to augment index configuration and troubleshoot data issues as part of update processing.

### Creating and deleting import owners

An import owner is typically created during the data import operation using the Endeca Configuration Import API. The owner name is specified as argument to the `ConfigRepositoryImporter` constructor. For details, see the *Endeca Configuration Import API Reference (Javadoc)*. You can also create an import owner using the `-o` option of the `set-config` task. If the owner does not already exist, the utility creates it.

The index configuration associated with an owner is removed using the `delete-owner` task.

### Examining index configuration for an owner

You can retrieve and examine index configuration for an owner using the `get-config` task and the `-o` option. The task returns a JSON file with the index configuration for the specified owner. If the file contains configuration from multiple import owners, each import owner's configuration is represented as a node within the file.

### Overwriting index configuration per owner

As described in [Setting the index configuration for an application](#) on page 71, the `set-config` task overwrites all previous index configuration for a specified owner.

### Setting global configuration

Only the `system` owner can set global index configuration.

### Merging index configuration from multiple owners

During baseline update processing, the Content Acquisition System merges and processes index configuration from all owners into a consolidated set of MDEX-compatible output files.

If multiple import owners modify the same attribute, the configuration from the `system` owner always overrides other import owners during the merge process.

For example, suppose an import owner named `ATG` creates an attribute that represents an Endeca property. Then the `system` owner updates the Endeca property to become an Endeca dimension with `isAutogen` set to `true`. The merged configuration processes the attribute and updates it to become an autogen dimension.

## About the schema for the index configuration file

The Index Configuration Command-line Utility writes and reads index configuration as JSON. The schema for the JSON file varies depending on whether you retrieve configuration for one owner or more than one owner and whether you restrict the types of configuration that you retrieve.

Types of configuration include:

- Endeca properties, derived properties, and dimensions. These are specified under the `attributes` node.
- Precedence rules. These are specified under the `precedenceRules` node.
- Search configuration. These are specified under the `searchIndexConfig` node.

See the topics and examples below that illustrate the schema for each index configuration type.

## Schema for an Endeca property, derived property, or dimension

You specify an Endeca property or dimension as an `attribute` in the index configuration file. Each attribute has a `jcr:primaryType` property with one of the following values:

- `endeca:property` - indicates that the attribute is an Endeca property.
- `endeca:derivedProperty` - indicates that the attribute is a derived property.
- `endeca:dimension` - indicates that the attribute is an Endeca dimension.

Properties, dimensions, and precedence rules must be named uniquely across the index configuration. The only exception to this is where you update a property by setting the `mergeAction` property to `UPDATE`.

### Schema for an attribute node that defines an Endeca property

An attribute that is an `endeca:property` can have the following schema properties:

Name	Data Type	Description
<code>isEnabled</code>	BOOLEAN	Optional. Indicates whether the property is processed by CAS when CAS writes MDEX-compatible output. A value of <code>true</code> includes the property during processing; <code>false</code> excludes the property. This setting is useful when troubleshooting data issues for specific attributes. If omitted, the default value is <code>true</code> .
<code>isRecordFilterable</code>	BOOLEAN	Optional. Indicates whether the property can be used to filter records. Record filtering presents a subset of the data to the end-user. If omitted, the default value is <code>false</code> .
<code>isRecordSearchEnabled</code>	BOOLEAN	Optional. Specifies whether or not record search should be enabled for this property. Record search finds all records in an Endeca application that are tagged with an Endeca property that matches a term the user provides. You must enable each property that you want available for record search. If omitted, the default value is <code>false</code> .
<code>isRollupKey</code>	BOOLEAN	Optional. Indicates whether the property can be used as a rollup key. This allow aggregated records to be based on this Endeca property. If omitted, the default value is <code>false</code> .

Name	Data Type	Description
<code>isWildcardEnabledInRecordSearch</code>	BOOLEAN	Optional. Indicates whether wildcard search is enabled for this Endeca property. Wildcard searching allows user queries that contain a wildcard character (*) to match against fragments of words in a property value. You must enable each property that you want available for wildcard searching. If <code>isWildcardEnabledInRecordSearch</code> is set to <code>true</code> , then <code>isRecordSearchEnabled</code> must also be set to <code>true</code> for a property. (Enabling wildcard record search depends on first enabling record search.) If omitted, the default value is <code>false</code> .
<code>mergeAction</code>	STRING	Optional. The <code>mergeAction</code> specifies how to merge the attribute into the index configuration. Valid enumerations are <code>ADD</code> and <code>UPDATE</code> . Specify a value of <code>ADD</code> to merge new attributes that are not already in the system. Specify a value of <code>UPDATE</code> to merge changes to an existing attribute. If omitted, the default value of <code>mergeAction</code> is <code>ADD</code> .
<code>propertyDataType</code>	STRING	Optional. The <code>propertyDataType</code> enumerates the valid values for the data type of an Endeca property. The valid enumerations are <code>ALPHA</code> , <code>INTEGER</code> , <code>DOUBLE</code> , <code>GEOCODE</code> , <code>DATETIME</code> , <code>DURATION</code> , and <code>TIME</code> . Such data types have several uses. Non-alpha properties can be used for range filtering. Temporal properties can be used for record sorting and analytics. If omitted, the default value of <code>propertyDataType</code> is <code>ALPHA</code> .
<code>sourcePropertyNames</code>	STRING (multi-valued)	Optional. Specifies an explicit mapping between one or more source properties and an Endeca property. An Endeca property is populated with data from the source property that it is mapped to. If specified, the <code>sourcePropertyNames</code> value can be empty, single-valued, or multi-valued.  If empty (a zero-length list), no source property is mapped to an Endeca property. This allows you to define an Endeca property but not populate it with any data.  If single-valued, then the source property has its value mapped to the Endeca property.  If multi-valued, then each source property in this list has its value mapped to the Endeca property.  If omitted, the source property has its value mapped to an Endeca property of the same name. In other words, a source property with a name that is identical to an Endeca property is automatically mapped to that Endeca property. This is the default behavior.

### Schema for an attribute node that defines an Endeca derived property

An attribute that is an `endeca:derivedProperty` can have the following schema properties:

Name	Data Type	Description
derivedPropertyFunction	STRING	Required. The <code>derivedPropertyFunction</code> enumerates the the valid functions that can be applied to the <code>derivedPropertySource</code> property. The valid enumerations are <code>MIN</code> , <code>MAX</code> , <code>SUM</code> , and <code>AVG</code> .
derivedPropertySource	STRING	Required. Specifies the Endeca property from which the derived property is calculated.
isEnabled	BOOLEAN	Optional. Indicates whether the property is processed by CAS when CAS writes MDEX-compatible output. A value of <code>true</code> includes the property during processing; <code>false</code> excludes the property. This setting is useful when troubleshooting data issues for specific attributes.
mergeAction	STRING	Optional. The <code>mergeAction</code> specifies how to merge the attribute into the index configuration. Valid enumerations are <code>ADD</code> and <code>UPDATE</code> . Specify a value of <code>ADD</code> to merge new attributes that are not already in the system. Specify a value of <code>UPDATE</code> to merge changes to an existing attribute. If omitted, the default value of <code>mergeAction</code> is <code>ADD</code> .

### Schema for an attribute node that defines an Endeca dimension

An attribute that is an `endeca:dimension` can have the following schema properties:

Name	Data Type	Description
displayOrder	INTEGER	Optional. Specifies the display order of a dimension relative to other dimensions in refinement results. Dimensions with lower values display before dimensions with higher values. Valid values are integers between 0 and 2147483647. If omitted, the dimension displays lower than dimensions with specified display orders. If dimensions have the same display order value (a tie), the dimensions are ordered alphabetically by dimension name.
isAutogen	BOOLEAN	Optional. Specifies whether the dimension values for a dimension are automatically generated during a CAS crawl. A value of <code>true</code> generates dimension values for a dimension. If omitted, the default value is <code>false</code> . (An error results if you set this to <code>true</code> and also specify dimension values for the dimension.)
isEnabled	BOOLEAN	Optional. Indicates whether the dimension is processed by CAS when CAS writes MDEX-compatible output. A value of <code>true</code> includes the dimension during processing; <code>false</code> excludes the dimension. This setting is useful when troubleshooting data issues for specific attributes.
isHierarchicalDimensionSearchEnabled	BOOLEAN	Optional. Specifies whether a dimension search also considers ancestor dimension values in this dimension when matching a dimension search query. If omitted, the default value is <code>false</code> .

Name	Data Type	Description
<code>isHierarchicalRecordSearchEnabled</code>	BOOLEAN	Optional. Specifies whether a record search also considers ancestor dimension values in this dimension when matching a record search query.  If <code>isHierarchicalRecordSearchEnabled</code> is set to <code>true</code> , then <code>isRecordSearchEnabled</code> must also be set to <code>true</code> for a dimension. (Enabling hierarchical search depends on first enabling record search.) If omitted, the default value of <code>isHierarchicalRecordSearchEnabled</code> is set to the value of <code>isRecordSearchEnabled</code> .
<code>isRecordSearchEnabled</code>	BOOLEAN	Optional. Specifies whether or not record search should be enabled for this dimension. Record search finds all records in an Endeca application that are tagged with a dimension value that matches a term the user provides. You must enable each property that you want available for record search. If omitted, the default value is <code>false</code> .
<code>isWildcardEnabledInRecordSearch</code>	BOOLEAN	Optional. Indicates whether wildcard search is enabled for this dimension. Wildcard searching allows user queries that contain a wildcard character (*) to match against fragments of words in a dimension. You must enable each dimension that you want available for wildcard searching.  If <code>isWildcardEnabledInRecordSearch</code> is set to <code>true</code> , then <code>isRecordSearchEnabled</code> must also be set to <code>true</code> for a dimension. (Enabling wildcard dimension search depends on first enabling dimension search.) If omitted, the default value is <code>false</code> .
<code>mergeAction</code>	STRING	Optional. The <code>mergeAction</code> specifies how to merge the attribute into the index configuration. Valid enumerations are <code>ADD</code> and <code>UPDATE</code> . Specify a value of <code>ADD</code> to merge new attributes that are not already in the system. Specify a value of <code>UPDATE</code> to merge changes to an existing attribute. If omitted, the default value of <code>mergeAction</code> is <code>ADD</code> .
<code>multiSelectType</code>	STRING	Optional. The <code>multiSelectType</code> enumerates the valid values for specifying multiselect dimensions. When <code>AND</code> is specified on a dimension, the MDEX Engine returns all records from all the select dimension values. The result set is expanded with each additional dimension value that a user selects. The <code>OR</code> enumeration returns the records from one selected dimension value. The result set is reduced with each dimension value that a user selects. The valid enumerations are <code>NONE</code> , <code>OR</code> , and <code>AND</code> .
<code>rangeComparisonType</code>	STRING	Optional. The <code>rangeComparisonType</code> enumerates the types that can be used to map source properties to dimension values that represent ranges. The valid enumerations are <code>STRING</code> , <code>INTEGER</code> , and <code>FLOAT</code> .
<code>sourcePropertyNames</code>	STRING (multi-valued)	Optional. Specifies an explicit mapping between one or more source properties and an Endeca dimension. A dimension is

Name	Data Type	Description
		<p>populated with data from the source property that it is mapped to. If specified, the <code>sourcePropertyNames</code> value can be empty, single-valued, or multi-valued.</p> <p>If empty (a zero-length list), no source property is mapped to a dimension. This allows you to define a dimension but not populate it with any data. This is useful when creating trigger dimension values for content spotlighting cartridges.</p> <p>If single-valued, then the source property has its value mapped to the Endeca dimension.</p> <p>If multi-valued, then each source property in this list has its value mapped to the Endeca dimension.</p> <p>If omitted, the source property has its value mapped to an Endeca dimension of the same name. In other words, a source property with a name that is identical to a dimension is automatically mapped to that dimension. This is the default behavior.</p>

### Example index configuration for two owners and all configuration types

In this example, the utility returns index configuration for the `system` owner. The configuration in this case is made up of attributes and global index configuration settings:

```
{
  "indexConfig" : {
    "system" : {
      "attributes" : {
        "product.price" : {
          "propertyDataType" : "DOUBLE",
          "jcr:primaryType" : "endeca:property"
        },
        "product.brand.name" : {
          "isRecordSearchEnabled" : true,
          "isAutogen" : true,
          "jcr:primaryType" : "endeca:dimension",
          "multiSelectType" : "OR"
        },
        "product.review.count" : {
          "propertyDataType" : "INTEGER",
          "jcr:primaryType" : "endeca:property"
        },
        "product.sku" : {
          "isRecordSearchEnabled" : true,
          "propertyDataType" : "ALPHA",
          "jcr:primaryType" : "endeca:property"
        },
        "product.id" : {
          "isRecordFilterable" : true,
          "isRecordSearchEnabled" : true,
          "propertyDataType" : "ALPHA",
          "jcr:primaryType" : "endeca:property"
        },
        "camera.color" : {
          "sourcePropertyNames" : [ "camera.Colour of product" ],

```

```

    "isAutogen" : true,
    "jcr:primaryType" : "eneca:dimension"
  },
  "product.category" : {
    "sourcePropertyNames" : [ "product.category_id" ],
    "isRecordSearchEnabled" : true,
    "jcr:primaryType" : "eneca:dimension"
  },
  "product.name" : {
    "isRecordSearchEnabled" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "eneca:property"
  },
  "product.features" : {
    "isAutogen" : true,
    "jcr:primaryType" : "eneca:dimension",
    "multiSelectType" : "AND"
  },
  "product.min_price" : {
    "derivedPropertySource" : "product.price",
    "derivedPropertyFunction" : "MIN",
    "jcr:primaryType" : "eneca:derivedProperty"
  },
  "product.price_range" : {
    "sourcePropertyNames" : [ "product.price" ],
    "rangeComparisonType" : "FLOAT",
    "jcr:primaryType" : "eneca:dimension"
  },
  "common.id" : {
    "isRecordFilterable" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "eneca:property"
  },
  ...
},
"precedenceRules" : {
},
"searchIndexConfig" : {
  "spellingDictMinNumWordOccurrences" : 4,
  "spellingDictMaxWordLength" : 16,
  "isWildcardEnabledInDimensionSearch" : true,
  "spellingDictMinWordLength" : 3
}
}
}
}

```

## Schema for precedence rules

You specify a precedence rule in the `precedenceRules` node of the index configuration file. The `precedenceRules` node is a sibling of the `attributes` node.

### Schema for a precedenceRule node to define a precedence rule

Each node representing a precedence rule under `precedenceRules` can have the following schema properties:



Name	Data Type	Description
isEnabled	BOOLEAN	Optional. Indicates whether the precedence rule is processed by CAS when CAS writes MDEX-compatible output. A value of <code>true</code> includes the precedence rule during processing; <code>false</code> excludes the precedence rule. This setting is useful when troubleshooting data issues for specific attributes.
isLeafTrigger	BOOLEAN	Optional. Specifies a Boolean to indicate if the trigger is a leaf trigger or not. If set to <code>true</code> , the rule only triggers on leaf dimension values.
mergeAction	STRING	Optional. The <code>mergeAction</code> enumerates the valid values that describe how to merge the precedence rule into the index configuration. Valid enumerations are <code>ADD</code> and <code>UPDATE</code> . Specify a value of <code>ADD</code> to merge new precedence rules that are not already in the system. Specify a value of <code>UPDATE</code> to merge changes to an existing precedence rule. If omitted, the default value of <code>mergeAction</code> is <code>ADD</code> .
triggerDimension	STRING	Required. Specifies the trigger dimension for a precedence rule. Recall that a user's selection of the trigger dimension reveals the previously unavailable target dimension to the user.
triggerDimensionValueSpec	STRING	Optional. Specifies the dimension value specification of the <code>triggerDimension</code> . If omitted, the precedence rule fires for any selection from the trigger dimension.
targetDimension	STRING	Required. Specifies the target dimension for a precedence rule.

### Example index configuration for precedence rules

In this example, the utility returns index configuration from one owner, named `ATG`, and the `precedenceRules` configuration type.

```
C:\Endeca\apps\Discover\control>index_config_cmd.bat get-config -o ATG -t precedenceRules
[07.26.12 12:57:19] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/ifcr with username admin
```

The output contains a `precedenceRules` root node, because that was the specified configuration type, and then child nodes for two precedence rules:

```
{
  "precedenceRules" : {
    "aspectRatioDigitalCamerasRule" : {
      "targetDimension" : "camera.aspect_ratio",
      "triggerDimensionValueSpec" : "575",
      "triggerDimension" : "product.category",
      "isLeafTrigger" : false
    },
    "digitalZoomDigitalCamerasRule" : {
      "targetDimension" : "camera.digital_zoom",
      "triggerDimensionValueSpec" : "575",
      "triggerDimension" : "product.category",
      "isLeafTrigger" : false
    }
  }
  ...
}
```

```
}
}
```

## Schema for global index configuration

You specify search configuration in the `searchIndexConfig` node of the index configuration file. In this release, the settings control spelling dictionary configuration and wildcard search. The `searchIndexConfig` node is a sibling of the `attributes` and `precedenceRules` nodes.

### Schema for a `searchIndexConfig` node

Each property under `searchIndexConfig` represents an index configuration setting. The following properties are available:

Name	Data Type	Description
<code>spellingDictMinNumWordOccurrences</code>	LONG	Optional. Specifies the minimum number of times the word must appear in the source data before the word should be included in the spelling dictionary. This setting applies to record search only. If omitted, the default value is 4.  For dimension search, this setting is always set to 1. (All dimension value names are included in the spelling dictionary by default.)
<code>spellingDictMaxWordLength</code>	LONG	Optional. Specifies the maximum length of a word that should be included in the spelling dictionary. Words longer than this value are excluded. This setting applies to both dimension search and record search. If omitted, the default value is 16.
<code>isWildcardEnabledInDimensionSearch</code>	BOOLEAN	Optional. Specifies a Boolean to indicate that a query can contain a wildcard character (*) to match against fragments of words in a dimension value. If omitted, the default value is true.
<code>spellingDictMinWordLength</code>	LONG	Optional. Specifies the minimum character length for a word to be included in the spelling dictionary. This setting applies to both dimension search and record search. If omitted, the default value is 3.

### Example index configuration for global settings

In this example, the utility returns index configuration for the `system` owner. The index configuration is restricted to only the `searchIndexConfig` type.

```
C:\Endeca\apps\Discover\control>index_config_cmd.bat get-config -o system -t
searchIndexConfig
[07.26.12 12:57:19] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/
ifcr with username admin
```

The output contains a `searchIndexConfig` root node, because that was the specified configuration type, and then properties for each configuration setting:

```
{
  "searchIndexConfig" : {
    "spellingDictMinNumWordOccurrences" : 4,
    "spellingDictMaxWordLength" : 16,
```

```

    "isWildcardEnabledInDimensionSearch" : true,
    "spellingDictMinWordLength" : 3
  }
}

```

## Getting the index configuration for an application

The `get-config` task retrieves the index configuration for an application.

The syntax for this task is:

```

index_config_cmd get-config [-o OwnerName] [-f FileName]
[-r] [-t precedenceRules|attributes|searchIndexConfig]

```

Where:

- `-o` (or `--owner`) specifies an import owner for a task. If you specify the `-o` option, the task applies only to the owner specified. The option can have an argument of `all`, `system`, or user-specified owner. The `all` owner includes both the `system` owner and all import owners. If you omit this option, the task applies to the `system` owner. Optional.
- `-f` (or `--file`) specifies a path to a JSON output file that contains the index configuration. Omitting the `-f` option prints the index configuration to standard out. Optional.
- `-r` (or `--repositoryMetadata`) specifies whether to return metadata about each attribute value in the index configuration. Metadata for an attribute includes properties such as `jcr:lastModifiedBy`, `jcr:createdBy`, `jcr:created`, `jcr:lastModified`, and so on. Optional.
- `-t` (or `--type`) specifies the type of index configuration you want the task to return. The arguments are `precedenceRules`, `attributes`, and `searchIndexConfig`. Specifying `precedenceRules` returns only precedence rules in the index configuration, or none. Specifying `attributes` returns the attributes in the index configuration. Specifying `searchIndexConfig` returns only the global index configuration settings. Omitting the `-t` option returns all types of index configuration. Optional.



**Note:** There is a size limit on the total number of attributes and precedence rules the task can retrieve. If the index configuration that you are retrieving contains more than approximately 10,000 attributes and precedence rules, the `get-config` task returns a `Multiple Choices (300)` error.

To get the index configuration for an application:

1. Start a command prompt and navigate to `<app dir>\control` (for Windows) or `<app dir>/control` (for UNIX).
2. Type `index_config_cmd` and specify the `get-config` task.



**Note:** This task name is case sensitive.

### Example of getting the index configuration for an application

```

C:\Endeca\apps\Discover\control>index_config_cmd.bat get-config
[07.23.12 15:50:54] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/
ifcr with username admin
{
  "indexConfig" : {
    "system" : {
      "attributes" : {
        "product.price" : {
          "propertyDataType" : "DOUBLE",

```

```

    "jcr:primaryType" : "endeca:property"
  },
  "product.brand.name" : {
    "isRecordSearchEnabled" : true,
    "isAutogen" : true,
    "jcr:primaryType" : "endeca:dimension",
    "multiSelectType" : "OR"
  },
  "product.review.count" : {
    "propertyDataType" : "INTEGER",
    "jcr:primaryType" : "endeca:property"
  },
  "product.sku" : {
    "isRecordSearchEnabled" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "endeca:property"
  },
  "product.id" : {
    "isRecordFilterable" : true,
    "isRecordSearchEnabled" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "endeca:property"
  },
  "camera.color" : {
    "sourcePropertyNames" : [ "camera.Colour of product" ],
    "isAutogen" : true,
    "jcr:primaryType" : "endeca:dimension"
  },
  "product.category" : {
    "sourcePropertyNames" : [ "product.category_id" ],
    "isRecordSearchEnabled" : true,
    "jcr:primaryType" : "endeca:dimension"
  },
  "product.name" : {
    "isRecordSearchEnabled" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "endeca:property"
  },
  "product.features" : {
    "isAutogen" : true,
    "jcr:primaryType" : "endeca:dimension",
    "multiSelectType" : "AND"
  },
  "product.min_price" : {
    "derivedPropertySource" : "product.price",
    "derivedPropertyFunction" : "MIN",
    "jcr:primaryType" : "endeca:derivedProperty"
  },
  "product.price_range" : {
    "sourcePropertyNames" : [ "product.price" ],
    "rangeComparisonType" : "FLOAT",
    "jcr:primaryType" : "endeca:dimension"
  },
  "common.id" : {
    "isRecordFilterable" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "endeca:property"
  },
  ...
},

```

```

    "precedenceRules" : {
    },
    "searchIndexConfig" : {
      "spellingDictMinNumWordOccurrences" : 4,
      "spellingDictMaxWordLength" : 16,
      "isWildcardEnabledInDimensionSearch" : true,
      "spellingDictMinWordLength" : 3
    }
  }
}
}
}

```

## Getting the merged index configuration for an application

The `get-merged-config` task retrieves the merged index configuration for an application. In other words, the output of this task is the consolidated index configuration for all import owners.

This task is primarily used as a debugging tool to troubleshoot configuration and data issues in the MDEX-compatible output files produced by a CAS crawl.

In some ways, the `get-merged-config` task is logically similar to `get-config` task with the `owner` option (`-o`) set to `all`. However, there are several important differences between the two tasks:

- `get-config` outputs configuration that is grouped in nodes by the owner name. `get-merged-config` outputs the consolidated configuration with no distinction for ownership.
- `get-config` does not remove copies of attributes from other owners. For example, if an ATG owner adds attribute A and `system` owner updates attribute A, `get-config` returns attribute A in the node for the ATG owner and also the node for the `system` owner. Whereas, `get-merged-config` merges the copies of the attribute and returns only one instance of attribute A which is from the `system` owner.

The syntax for this task is:

```

index_config_cmd get-merged-config [-f FileName]
[-t precedenceRules|attributes|searchIndexConfig]

```

Where:

- `-f` (or `--file`) specifies a path to a JSON output file that contains the index configuration. Omitting the `-f` option prints the index configuration to standard out. Optional.
- `-t` (or `--type`) specifies the type of index configuration you want the task to return. The arguments are `precedenceRules`, `attributes`, and `searchIndexConfig`. Specifying `precedenceRules` returns only precedence rules in the index configuration, or none. Specifying `attributes` returns the attributes in the index configuration. Specifying `searchIndexConfig` returns only the global index configuration settings. Omitting the `-t` option returns all types of index configuration. Optional.

To get the merged index configuration for an application:

1. Start a command prompt and navigate to `<app dir>\control` (for Windows) or `<app dir>/control` (for UNIX).
2. Type `index_config_cmd` and specify the `get-merged-config` task.



**Note:** This task name is case sensitive.

**Example of getting the merged index configuration for an application**

```

C:\Endeca\apps\Discover\control>index_config_cmd.bat get-merged-config
[08.17.12 11:48:05] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/
ifcr with username admin
{
  "indexConfig" : {
    "attributes" : {
      "product.price" : {
        "propertyDataType" : "DOUBLE",
        "jcr:primaryType" : "endeca:property"
      },
      "product.brand.name" : {
        "isRecordSearchEnabled" : true,
        "isAutogen" : true,
        "jcr:primaryType" : "endeca:dimension",
        "multiSelectType" : "OR"
      },
      "product.sku" : {
        "isRecordSearchEnabled" : true,
        "propertyDataType" : "ALPHA",
        "jcr:primaryType" : "endeca:property"
      },
      "product.id" : {
        "isRecordFilterable" : true,
        "isRecordSearchEnabled" : true,
        "propertyDataType" : "ALPHA",
        "jcr:primaryType" : "endeca:property"
      },
      "camera.megapixel_range" : {
        "sourcePropertyNames" : [ "camera.Megapixel" ],
        "rangeComparisonType" : "FLOAT",
        "jcr:primaryType" : "endeca:dimension"
      },
      "product.category" : {
        "sourcePropertyNames" : [ "product.category_id" ],
        "isRecordSearchEnabled" : true,
        "jcr:primaryType" : "endeca:dimension"
      },
      "product.name" : {
        "isRecordSearchEnabled" : true,
        "propertyDataType" : "ALPHA",
        "jcr:primaryType" : "endeca:property"
      },
      "product.features" : {
        "isAutogen" : true,
        "jcr:primaryType" : "endeca:dimension",
        "multiSelectType" : "AND"
      },
      "product.min_price" : {
        "derivedPropertySource" : "product.price",
        "derivedPropertyFunction" : "MIN",
        "jcr:primaryType" : "endeca:derivedProperty"
      },
      "product.code" : {
        "isRollupKey" : true,
        "isRecordSearchEnabled" : true,
        "jcr:primaryType" : "endeca:property",
        "propertyDataType" : "ALPHA"
      },
      "product.price_range" : {

```

```

    "sourcePropertyNames" : [ "product.price" ],
    "rangeComparisonType" : "FLOAT",
    "jcr:primaryType" : "endeca:dimension"
  },
  "product.max_price" : {
    "derivedPropertySource" : "product.price",
    "derivedPropertyFunction" : "MAX",
    "jcr:primaryType" : "endeca:derivedProperty"
  },
  "product.long_desc" : {
    "isRecordSearchEnabled" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "endeca:property"
  },
  "product.short_desc" : {
    "isRecordSearchEnabled" : true,
    "propertyDataType" : "ALPHA",
    "jcr:primaryType" : "endeca:property"
  },
  ...
},
"precedenceRules" : {
},
"searchIndexConfig" : {
  "spellingDictMinNumWordOccurrences" : 4,
  "spellingDictMaxWordLength" : 16,
  "isWildcardEnabledInDimensionSearch" : true,
  "spellingDictMinWordLength" : 3
}
}
}

```

## Setting the index configuration for an application

The `set-config` task sets the index configuration for a specified owner. You provide the index configuration in a JSON file.

Running this task overwrites any previous index configuration for the owner. Oracle recommends that developers who modify the index configuration, use the default `system` owner. This usage separates index configuration that comes from the `system` owner from configuration that comes from import operations which are owner by a user-specified owner. If the JSON configuration file contains index configuration from multiple owners, you must specify the `-o` option with a value of `all`.

If desired, you can also update an Endeca property to become a dimension by modifying the `jcr:primaryType` from `endeca:property` to `endeca:dimension`. However, you cannot modify it from dimension to property.

The syntax for this task is:

```
index_config_cmd set-config [-o OwnerName] -f FileName
```

Where:

- `-o` (or `--owner`) specifies an import owner for a task. If you specify the `-o` option, the task applies only to the owner specified. The option can have an argument of `all`, `system`, or a user-specified import owner. The `all` owner includes both the `system` owner and all import owners. If you omit this option, the task applies to the `system` owner. Optional.
- `-f` (or `--file`) specifies a path to a JSON file that contains the index configuration. Required.

To set the index configuration for an application:

1. Start a command prompt and navigate to <app dir>\control (for Windows) or <app dir>/control (for UNIX).
2. Type `index_config_cmd` and specify the `set-config` task and the `-f` option with a path to the JSON file.



**Note:** This task name is case sensitive.

### Examples of setting the index configuration for an application

This example sets index configuration from three owners.

```
C:\Endeca\apps\Discover\control>index_config_cmd.bat set-config -f C:\temp\index-
ConfigAllOwners.json -o all
[07.24.12 15:53:58] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/
ifcr with username admin
You are attempting to write schema configuration that will be overwritten in the
event of a fresh import.
Are you sure you want to continue? (y/n)
y
```

This example sets index configuration from an owner named ATG.

```
C:\Endeca\apps\Discover\control>index_config_cmd.bat set-config -f C:\temp\index-
ConfigATGOwner.json -o ATG
[07.24.12 16:23:06] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/
ifcr with username admin
You are attempting to write schema configuration that will be overwritten in the
event of a fresh import.
Are you sure you want to continue? (y/n)
y
```

## Deleting the index configuration associated with an owner

The `delete-owner` task removes index configuration from an application that is associated with an owner that you specify.

The syntax for this task is:

```
index_config_cmd delete-owner -o OwnerName
```

Where:

- `-o` (or `--owner`) specifies an import owner for a task. If you specify the `-o` option, the task applies only to a user-specified owner. You cannot delete the `system` or `all` owners. Required.

To delete the index configuration associated with an owner:

1. Start a command prompt and navigate to <app dir>\control (for Windows) or <app dir>/control (for UNIX).
2. Type `index_config_cmd` and specify the `delete-owner` task with an argument for the owner's index configuration that you want to remove.





**Note:** This task name is case sensitive.

### Example of deleting the index configuration associated with an owner

This example deletes the index configuration for the ATG owner.

```
C:\Endeca\apps\Discover\control>index_config_cmd.bat delete-owner -o ATG
[07.24.12 17:14:50] INFO: Using site Discover at URL http://JSMITH-WIN7:8006/
ifcr with username admin
```

## An example of changing multi-select on a dimension

This topic provides a simple example of using the Index Configuration Command-line Utility to update index configuration. In this example, suppose an import owner named ATG has added index configuration to an Endeca application. You want to update the index configuration by adding `multiSelectType` to the `product.category` dimension.

The steps to accomplish this are as follows:

1. Retrieve the index configuration by running:

```
C:\Endeca\apps\\control>index_config_cmd.bat get-config -o ATG -f
C:\temp\indexConfig.json
```

2. Open the resulting JSON file and locate the `product.category` attribute:

```
{
  "indexConfig" : {
    "attributes" : {
      "product.category" : {
        "sourcePropertyNames" : [ "product.category_id" ],
        "isRecordSearchEnabled" : true,
        "jcr:primaryType" : "endeca:dimension"
      },
    },
  },
  ....
}
```

3. Add `multiSelectType` and set the value to `OR`, and also add the `mergeAction` with a value of `UPDATE`. You can delete other properties of the attribute because they are not changing as part of the update:

```
{
  "indexConfig" : {
    "attributes" : {
      "product.category" : {
        "mergeAction" : "UPDATE",
        "multiSelectType" : "OR",
      },
    },
  },
  ....
}
```

4. Set the revised index configuration by running:

```
C:\Endeca\apps\\control>index_config_cmd.bat set-config -o ATG
-f C:\temp\indexConfig.json
```

5. If desired, examine the merged configuration by running:

```
C:\Endeca\apps\\control>index_config_cmd.bat get-merged-config
-f C:\temp\indexConfig.json
```

You see the following:

```
{
  "indexConfig" : {
    "ATG" : {
      "attributes" : {
        "product.category" : {
          "sourcePropertyNames" : [ "product.category_id" ],
          "isRecordSearchEnabled" : true,
          "multiSelectType" : "OR",
          "jcr:primaryType" : "endeca:dimension"
        },
      },
    },
  },
  ....
}
```

## An example of changing a product.brand.name property to a dimension

This topic provides a simple example of using the Index Configuration Command-line Utility to update index configuration. In this example, suppose an import owner named ATG has added index configuration to an Endeca application. You want to update the index configuration by changing the `product.brand.name` attribute from an Endeca property to an Endeca dimension.

The steps to accomplish this are as follows:

1. Retrieve the index configuration by running:

```
C:\Endeca\apps\<<app dir>\control>index_config_cmd.bat get-config -o ATG -f
C:\temp\indexConfig.json
```

2. Open the resulting JSON file and locate the `product.brand.name` attribute:

```
{
  "indexConfig" : {
    "attributes" : {
      "product.brand.name" : {
        "isRecordSearchEnabled" : true,
        "jcr:primaryType" : "endeca:property"
      },
    },
  },
  ....
}
```

3. Change `jcr:primaryType` from `endeca:property` to `endeca:dimension`, add the `mergeAction` with a value of `UPDATE`, and also add `isAutogen` with a value of `true`:

```
{
  "indexConfig" : {
    "attributes" : {
      "product.brand.name" : {
        "mergeAction" : "UPDATE",
        "isAutogen" : true,
        "jcr:primaryType" : "endeca:dimension",
        "multiSelectType" : "OR"
      },
    },
  },
  ....
}
```

4. Set the revised index configuration by running:

```
C:\Endeca\apps\<>app dir>\control>index_config_cmd.bat set-config -o ATG
-f C:\temp\indexConfig.json
```

5. If desired, examine the merged configuration by running:

```
C:\Endeca\apps\<>app dir>\control>index_config_cmd.bat get-merged-config
-f C:\temp\indexConfig.json
```

You see the following:

```
{
  "indexConfig" : {
    "attributes" : {
      "product.brand.name" : {
        "isHierarchicalDimensionSearchEnabled" : true,
        "isRecordSearchEnabled" : true,
        "isAutogen" : true,
        "jcr:primaryType" : "endeca:dimension",
        "multiSelectType" : "OR"
      },
      ....
    }
  },
  ....
}
```

## An example of setting dimension display order

This topic provides an example of how the `displayOrder` property sets the display order of dimensions in the `discover-data-cas` application.

The following JSON snippet shows the `displayOrder` property for the Category dimension, Price Range dimension, and Brand Name dimension, where `displayOrder` is set to 0, 1, and 2, respectively.

```
...
"product.category" : {
  "displayOrder" : 0,
  "sourcePropertyNames" : [ "product.category_id" ],
  "isRecordSearchEnabled" : true,
  "jcr:primaryType" : "endeca:dimension"
},
"product.price_range" : {
  "sourcePropertyNames" : [ "product.price" ],
  "displayOrder" : 1,
  "rangeComparisonType" : "FLOAT",
  "jcr:primaryType" : "endeca:dimension"
},
...
"product.brand.name" : {
  "isHierarchicalDimensionSearchEnabled" : true,
  "displayOrder" : 2,
  "isRecordSearchEnabled" : true,
  "isAutogen" : true,
  "jcr:primaryType" : "endeca:dimension",
  "multiSelectType" : "OR"
},
....
```

When the dimensions are rendered in the Discover Electronics reference application, they render in the order specified by the property value. Category displays first, Price Range second, and Brand Name third:

**Narrow Your Results**

**Category**  
[Bags & Cases \(665\)](#)  
[Cameras \(4951\)](#)

---

**Price Range**  
[Under 25 \(81\)](#)  
[25 - 50 \(96\)](#)  
[50 - 100 \(170\)](#)  
[100 - 250 \(548\)](#)  
[250 - 500 \(915\)](#)  
[500 - 1000 \(1912\)](#)  
[Over 1000 \(1934\)](#)

---

**Brand Name**  
[A4Tech \(13\)](#)  
[Alecto \(5\)](#)  
[Approx \(3\)](#)  
[Asus \(1\)](#)  
[Belkin \(2\)](#)  
[Benq \(38\)](#)  
[Canon \(2306\)](#)  
[Canyon \(13\)](#)  
[Case Logic \(176\)](#)  
[Cisco \(24\)](#)

## Managing Data Operations

This section describes how to incorporate test data and production data into an application.

### Running a baseline update with test data

A deployed application includes test data that you can process with baseline update scripts, baseline test data, and a baseline Forge pipeline. Because this task describes test data, not production data, you use the `load_baseline_test_data` script to simulate the data extraction process (or to set the data readiness signal, in the case of an application that uses a non-extract data source).

The `load_baseline_test_data` script loads the test data stored in `<app_dir>/test_data/baseline` and runs the `set_baseline_data_ready_flag` script which sets a flag in the EAC indicating that data has been extracted and is ready for baseline update processing.



**Note:** This script is not required in applications that use CAS to produce MDEX-compatible output.

When you are done familiarizing yourself with the data processing steps and the test data, see [Running a baseline update with production data](#) on page 78. Processing production data requires the following changes to an application's configuration:

- Replace the step to run `load_baseline_test_data` with a data extraction process that delivers production data into the `<app_dir>/test_data/baseline` directory. Delete the `data.txt` file from `<app_dir>/test_data/baseline`. This step is not necessary if your application does not use data extracts: for example, if your application retrieves data directly from a database via ODBC or JDBC or from a CAS crawl.
- Set the `baseline_data_ready` flag in the EAC. You set the `baseline_data_ready` flag by making a Web service call to the EAC or by running the `set_baseline_data_ready_flag` script.

To run a baseline update with test data:

1. Ensure that the Endeca HTTP Service is running on each server in the deployment environment and that you have already deployed and initialized an application.
2. Start a command prompt (on Windows) or a shell (on UNIX).
3. Navigate to the `control` directory of deployed application.

This is located under your application directory. For example: `C:\Endeca\apps\<app_dir>\control`.

4. Run the `load_baseline_test_data` script.

- On Windows:
 

```
<app dir>\control\load_baseline_test_data.bat
```
  - On UNIX:
 

```
<app dir>/control/load_baseline_test_data.sh
```
5. Run the `baseline_update` script.
    - On Windows:
 

```
<app dir>\control\baseline_update.bat
```
    - On UNIX:
 

```
<app dir>/control/baseline_update.sh
```
  6. Examine the indexed data in an Endeca front-end application.  
For example, start a Web browser and open the JSP reference application at `http://localhost:8006/endeca_jspref`.  
You should see 10 records.

## Running a baseline update with production data

You run the `baseline_update` script to process production data and distribute the resulting index files to one or more Dgraphs. Production data may come from any number of sources including data extracts, CAS crawls, or direct calls to a database via ODBC or JDBC.

To run a baseline update with production data:

1. Ensure that the Endeca HTTP Service is running on each server in the deployment environment and that you have already deployed and initialized an application.
2. Replace the default Forge pipeline (Developer Studio configuration files) in `<app dir>/config/pipeline` with the Developer Studio configuration files for your application. For details, see [Replacing the Default Forge Pipeline](#) on page 47.
3. Replace the baseline test data stored in `<app dir>/test_data/baseline` with production data for the application. This step varies depending on your application requirements. It can include any of the following approaches:
  - Add a data extract file to the `<app dir>/test_data/baseline` and delete the test data extract.
  - Set up a CAS crawl to run as part of the `baseline_update` script.
  - Make a direct call to a database via ODBC or JDBC.

4. Start a command prompt (on Windows) or a shell (on UNIX).
5. Navigate to the `control` directory of deployed application.

This is located under your application directory. For example: `C:\Endeca\apps\<app dir>\control`.

6. Set the `baseline_data_ready` flag in the EAC by running the `set_baseline_data_ready_flag` script.
  - On Windows:
 

```
<app dir>\control\set_baseline_data_ready_flag.bat
```

- On UNIX:

```
<app_dir>/control/set_baseline_data_ready_flag.sh
```



**Note:** This script is not required in applications that use CAS to produce MDEX-compatible output.

7. Run the `baseline_update` script.

- On Windows:

```
<app_dir>\control\baseline_update.bat
```

- On UNIX:

```
<app_dir>/control/baseline_update.sh
```

8. Examine the indexed data in an Endeca front-end application.

For example, start a Web browser and open the JSP reference application at `http://localhost:8006/endeca_jspref`.

## Running a partial update with production data

You run the `partial_update` script to process *incremental changes* in production data and distribute the resulting index files to one or more Dgraphs. Production data may come from any number of sources including data extracts, CAS crawls, or direct calls to a database via ODBC or JDBC.

For more information on partial updates, see the *MDEX Engine Partial Updates Guide*.

To run a partial update with production data:

1. Ensure that the Endeca HTTP Service is running on each server in the deployment environment and that you have already deployed and initialized an application.
2. Replace the default Forge pipeline (Developer Studio configuration files) in `<app_dir>/config/pipeline` with the Developer Studio configuration files for your application. For details, see [Replacing the Default Forge Pipeline](#) on page 47.
3. Provide the partial data (incremental data changes since the last baseline update). This step varies depending on the application requirements. It can include any of the following approaches:
  - Add a data extract file to the `<app_dir>/test_data/partial`.
  - Set up a CAS crawl to run as part of the `baseline_update` script.
  - Make a direct call to a database via ODBC or JDBC.
4. Start a command prompt (on Windows) or a shell (on UNIX).
5. Navigate to the `control` directory of deployed application.

This is located under your application directory. For example: `C:\Endeca\apps\<app_dir>\control`.

6. Set the `partial_data_ready` flag in the EAC by running the `set_partial_data_ready_flag` script.

- On Windows:

```
<app_dir>\control\set_partial_data_ready_flag.bat
```

- On UNIX:

```
<app_dir>/control/set_partial_data_ready_flag.sh
```

7. Run the `partial_update` script.

- On Windows:

```
<app dir>\control\partial_update.bat
```

- On UNIX:

```
<app dir>/control/partial_update.sh
```

8. Examine the indexed data in an Endeca front-end application.

For example, start a Web browser and open the JSP reference application at `http://localhost:8006/endeca_jspref`.

## Running CAS crawls

In your `DataIngest.xml` code, you can run baseline or partial updates that include CAS crawls using the methods available in `ContentAcquisitionServerComponent`.

For details about `ContentAcquisitionServerComponent`, see the *EAC Component API Reference for CAS Server (Javadoc)* installed in `CAS\<version>\doc\cas-dt-javadoc` and see the CAS examples in [Script Reference](#) on page 81.



## Script Reference

This section describes scripts that are included with the Deployment Template, provides additional sample scripts, and provides information about running and configuring them.

### Analyzing Deployment Script Errors


When errors occur during the execution of a deployment template script, consult the error messages in the log files of the Endeca Application Controller (EAC) or the Workbench for information about the errors. These messages can help you analyze the cause of the errors by revealing the server state, operations performed, and exceptions encountered by Workbench or EAC. Note that deployment template scripts rely primarily on EAC and Workbench web services, invoking their operations in sequence to accomplish the overall task.

### Deployment Template script reference

The Deployment Template includes a set of utility scripts with deployed applications.

The following scripts are available in the `control` directory of a deployed application:

Script	Purpose
<code>baseline_update</code>	Runs a baseline update.
<code>export_site</code>	Takes a path to an XML file as an argument and exports the content in the Endeca Configuration Repository to the specified XML file.  If no file is specified, site data is exported to <code>&lt;App_Name&gt;-&lt;timestamp&gt;.xml</code> , where the timestamp format is <code>YYYY-MM-DD_HH-MM-SS</code> .
<code>get_editors_config</code>	Exports editor configuration to the <code>&lt;app dir&gt;\config\editors\config</code> directory.
<code>get_media</code>	Exports media configuration to the <code>&lt;app dir&gt;\config\media</code> directory.

Script	Purpose
get_templates	Exports template configuration to the <app dir>\config\cartridge_templates directory.
import_site	Takes a path to an XML file and imports the content to the Endeca Configuration Repository. Optionally, you can use the <code>--force</code> flag to override the confirmation prompt for overwriting site content that already exists.
load_baseline_test_data	Copies data from the <app dir>\test_data\baseline\ directory to <app dir>\data\incoming for a baseline update and calls the <code>set_baseline_data_ready_flag</code> script.
load_partial_test_data	Copies data from the <app dir>\test_data\partial\ directory to <app dir>\data\partials\incoming for a partial update and calls the <code>set_partial_data_ready_flag</code> script.
partial_update	Runs a partial update.
promote_content	Promotes content and configuration in the authoring environment to the live environment.
runcommand	Provides a means of invoking methods in <code>AppConfig.xml</code> against specified instances of objects.  You can run <code>runcommand</code> with the <code>--help</code> flag for a list of command line arguments and flags.
set_baseline_data_ready_flag	Sets the <code>baseline_data_ready</code> flag in the EAC.   <b>Note:</b> This script is not required in applications that use CAS to produce MDEX-compatible output.
set_editors_config	Imports editor configuration from <app dir>\config\editors\config to the Endeca Configuration Repository.
set_media	Imports media from <app dir>\config\media to the Endeca Configuration Repository.
set_partial_data_ready_flag	Sets the <code>partial_extract</code> flag in the EAC.
set_templates	Imports templates from <app dir>\config\cartridge_templates to the Endeca Configuration Repository.

Script	Purpose
initialize_services	<p>This script should be run once after deploying an application. It does the following:</p> <ul style="list-style-type: none"> <li>• Removes existing application provisioning</li> <li>• Sets new EAC provisioning and performs initial setup</li> <li>• Calls <code>set_editors_config</code></li> <li>• Calls <code>set_media</code></li> <li>• Calls <code>set_templates</code></li> </ul>

## Provisioning scripts

The EAC allows scripts to be provisioned and invoked via Web service calls. A script is provisioned by specifying a working directory, a log directory into which output from the script is recorded, and a command to execute the script.

The `AppConfig.xml` document allows defined scripts to be provisioned by specifying the command used to invoke the script from the command line. When the provisioning configuration information is included, the script is provisioned and becomes available for invocation via Web service calls or from the EAC Admin console in Oracle Endeca Workbench. When excluded, the script is not provisioned.

```
<script id="BaselineUpdate">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>
    ./control/baseline_update.bat
  </provisioned-script-command>
  <bean-shell-script>
    <![CDATA[
...
    ]]>
  </bean-shell-script>
</script>
```

The command line used to invoke scripts can always be specified in this form, relative to the default Deployment Template working directory:

```
./control/runcommand.[sh|bat] [script id]
```

## Forge-based data processing

The Deployment Template supports running baseline and partial updates using Forge. In this processing model, an update essentially runs a CAS crawl (if applicable), Forge, Dgidx, and then updates the Dgraphs in an application.

## Dgraph baseline update script using Forge

The baseline update script defined in the `DataIngest.xml` document for a Dgraph deployment is included in this section, with numbered steps indicating the actions performed at each point in the script.

```
<script id="BaselineUpdate">
  <![CDATA[
    log.info("Starting baseline update script.");
```

1. Obtain lock. The baseline update attempts to set an "update\_lock" flag in the EAC to serve as a lock or mutex. If the flag is already set, this step fails, ensuring that the update cannot be started more than once simultaneously, as this would interfere with data processing. The flag is removed in the case of an error or when the script completes successfully.

```
    // obtain lock
    if (LockManager.acquireLock("update_lock")) {
```

2. Validate data readiness. Check that a flag called "baseline\_data\_ready" has been set in the EAC. This flag is set as part of the data extraction process to indicate that files are ready to be processed (or, in the case of an application that uses direct database access, the flag indicates that a database staging table has been loaded and is ready for processing). This flag is removed as soon as the script copies the data out of the `data/incoming` directory, indicating that new data may be extracted.

```
    // test if data is ready for processing
    if (Forge.isDataReady()) {
```

3. Clean processing directories. Files from the previous update are removed from the `data/processing`, `data/forge_output`, `data/temp`, `data/dgidx_output` and `data/partials/cumulative_partials` directories.

```
    // clean directories
    Forge.cleanDirs();
    PartialForge.cleanCumulativePartials();
    Dgidx.cleanDirs();
```

4. Copy data to processing directory. Extracted data in `data/incoming` is copied to `data/processing`.

```
    // fetch extracted data files to forge input
    Forge.getIncomingData();
```

5. Release Lock. The "baseline\_data\_ready" flag is removed from the EAC, indicating that the incoming data has been retrieved for baseline processing.

```
    LockManager.releaseLock("baseline_data_ready");
```

6. Copy config to processing directory. Configuration files are copied from `data/complete_index_config` to `data/processing`.

```
    // fetch config files to forge input
    Forge.getConfig();
```

7. Archive Forge logs. The `logs/forges/Forge` directory is archived, to create a fresh logging directory for the Forge process and to save the previous Forge run's logs.

```
    // archive logs
    Forge.archiveLogDir();
```

8. Forge. The Forge process executes.

```
    Forge.run();
```

9. Archive Dgidx logs. The `logs/dgidxs/Dgidx` directory is archived, to create a fresh logging directory for the Dgidx process and to save the previous Dgidx run's logs.

```
// archive logs
Dgidx.archiveLogDir();
```

10. Dgidx. The Dgidx process executes.

```
Dgidx.run();
```

11. Distribute index to each server. A single copy of the new index is distributed to each server that hosts a Dgraph. If multiple Dgraphs are located on the same server but specify different `srcIndexDir` attributes, multiple copies of the index are delivered to that server.

12. Update MDEX Engines. The Dgraphs are updated. Engines are updated according to the `restartGroup` property specified for each Dgraph. The update process for each Dgraph is as follows:

- a. Create `dgraph_input_new` directory.
- b. Create a local copy of the new index in `dgraph_input_new`.
- c. Stop the Dgraph.
- d. Archive Dgraph logs (e.g. `logs/dgraphs/Dgraph1`) directory.
- e. Rename `dgraph_input` to `dgraph_input_old`.
- f. Rename `dgraph_input_new` to `dgraph_input`.
- g. Start the Dgraph.
- h. Remove `dgraph_input_old`.

This somewhat complex update functionality is implemented to minimize the amount of time that a Dgraph is stopped. This restart approach ensures that the Dgraph is stopped just long enough to rename two directories.

```
// distributed index, update Dgraphs
DistributeIndexAndApply.run();
```

```
<script id="DistributeIndexAndApply">
  <bean-shell-script>
    <![CDATA[
      DgraphCluster.cleanDirs();
      DgraphCluster.copyIndexToDgraphServers();
      DgraphCluster.applyIndex();
    ]]>
  </bean-shell-script>
</script>
```

13. If Workbench integration is enabled, upload post-Forge dimensions to Oracle Endeca Workbench. The latest dimension values generated by the Forge process are uploaded to Oracle Endeca Workbench, to ensure that any new dimension values (including values for autogen dimensions and external dimensions) are available to Oracle Endeca Workbench for use in, for example, dynamic business rule triggers.



**Note:** This action does not add new dimensions or remove existing dimensions. These changes can be made by invoking the `update_web_studio_config.[bat|sh]` script.

```
// if Workbench is integrated, update Workbench with latest
// dimension values
if (ConfigManager.isWebStudioEnabled()) {
  ConfigManager.cleanDirs();
  Forge.getPostForgeDimensions();
  ConfigManager.updateWsDimensions();
}
```

14. Archive index and Forge state. The newly created index and the state files in Forge's state directory are archived on the indexing server.

```
// archive state files, index
Forge.archiveState();
Dgidx.archiveIndex();
```

15. Cycle LogServer. The LogServer is stopped and restarted. During the downtime, the LogServer's error and output logs are archived.

```
// cycle LogServer
LogServer.cycle();
```

16. Release Lock. The "update\_lock" flag is removed from the EAC, indicating that another update may be started.

```
// release lock
LockManager.releaseLock("update_lock");

log.info("Baseline update script finished.");
} else {
log.warning("Failed to obtain lock.");
}
}]>
</bean-shell-script>
</script>
```

## Related Links

[Dgraph partial update script using Forge](#) on page 86

The partial update script defined in the `DataIngest.xml` document for a Dgraph deployment is included in this section, with numbered steps indicating the actions performed at each point in the script.

## Dgraph partial update script using Forge

The partial update script defined in the `DataIngest.xml` document for a Dgraph deployment is included in this section, with numbered steps indicating the actions performed at each point in the script.

```
<script id="PartialUpdate">
  <bean-shell-script>
    <![CDATA[
```

1. Obtain lock. The partial update attempts to set an "update\_lock" flag in the EAC to serve as a lock or mutex. If the flag is already set, this step fails, ensuring that the update cannot be started more than once simultaneously, as this would interfere with data processing. The flag is removed in the case of an error or when the script completes successfully.

```
log.info("Starting partial update script.");
// obtain lock
if (LockManager.acquireLock("update_lock")) {
```

2. Validate data readiness. Test that the EAC contains at least one flag with the prefix "partial\_extract::". One of these flags should be created for each successfully and completely extracted file, with the prefix "partial\_extract::" prepended to the extracted file name (e.g. "partial\_extract::adds.txt.gz"). These flags are deleted during data processing and must be created as new files are extracted.

```
// test if data is ready for processing
if (PartialForge.isPartialDataReady()) {
```

3. Archive partial logs. The `logs/partial` directory is archived, to create a fresh logging directory for the partial update process and to save the previous run's logs.

```
// archive logs
PartialForge.archiveLogDir();
```

4. Clean processing directories. Files from the previous update are removed from the `data/partials/processing`, `data/partials/forge_output`, and `data/temp` directories.

```
// clean directories
PartialForge.cleanDirs();
```

5. Move data and config to processing directory. Extracted files in `data/partials/incoming` with matching "partials\_extract::" flags in the EAC are moved to `data/partials/processing`. Configuration files are copied from `config/pipeline` to `data/processing`.

```
// fetch extracted data files to forge input
PartialForge.getPartialIncomingData();

// fetch config files to forge input
PartialForge.getConfig();
```

6. Forge. The partial update Forge process executes.

```
// run ITL
PartialForge.run();
```

7. Apply timestamp to updates. The output XML file generated by the partial update pipeline is renamed to include a timestamp, to ensure it is processed in the correct order relative to files generated by previous or following partial update processes.

```
// timestamp partial, save to cumulative partials dir
PartialForge.timestampPartials();
```

8. Copy updates to cumulative updates. The timestamped XML file is copied into the cumulative updates directory.

```
PartialForge.fetchPartialsToCumulativeDir();
```

9. Distribute update to each server. A single copy of the partial update file is distributed to each server specified in the configuration.

```
// distribute partial update, update Dgraphs
DgraphCluster.copyPartialUpdateToDgraphServers();
```

10. Update MDEX Engines. The Dgraph processes are updated. Engines are updated according to the `update-Group` property specified for each Dgraph. The update process for each Dgraph is as follows:

- a. Copy update files into the `dgraph_input/updates` directory.
- b. Trigger a configuration update in the Dgraph by calling the URL `admin?op=update`.

```
DgraphCluster.applyPartialUpdates();
```

11. Archive cumulative updates. The newly generated update file (and files generated by all partial updates processed since the last baseline) are archived on the indexing server.

```
// archive partials
PartialForge.archiveCumulativePartials();
```

12. Release Lock. The "update\_lock" flag is removed from the EAC, indicating that another update may be started.

```
// release lock
LockManager.releaseLock("update_lock");
```

```

    log.info("Partial update script finished.");
  }
  else {
    log.warning("Failed to obtain lock.");
  }
  ]]>
</bean-shell-script>
</script>

```

### Preventing non-nullable element exceptions

When running the partial updates script, you may see a Java exception similar to this example:

```

INFO: Starting copy utility 'copy_partial_update_to_host_MDEXHost1'.
Oct 20, 2008 11:46:37 AM org.apache.axis.encoding.ser.BeanSerializer serialize
SEVERE: Exception:
java.io.IOException: Non nullable element 'fromHostID' is null.
...

```

If this occurs, make sure that the following properties are defined in the `AppConfig.xml` configuration file:

```

<dgraph-defaults>
  <properties>
    ...
    <property name="srcPartialsDir" value="./data/partials/forge_output" />
    <property name="srcPartialsHostId" value="ITLHost" />
    <property name="srcCumulativePartialsDir" value="./data/partials/cumula-
tive_partials" />
    <property name="srcCumulativePartialsHostId" value="ITLHost" />
    ...
  </properties>
  ...
</dgraph-defaults>

```

The reason is that the script is obtaining the `fromHostID` value from this section.

### Related Links

[Dgraph baseline update script using Forge](#) on page 84

The baseline update script defined in the `DataIngest.xml` document for a Dgraph deployment is included in this section, with numbered steps indicating the actions performed at each point in the script.

## Dgraph baseline update script using Forge and a CAS full crawl script

After running a full CAS crawl, you can run a baseline update using Forge to incorporate the records from a Record Store instance.

This example runs a *baseline* update that includes a *full* CAS crawl. The crawl writes output to a Record Store instance and then Forge incorporates the records from the crawl. To create this sequential workflow of CAS crawl and then baseline update, you can do the following:

- Remove the default `Forge.isDataReady` check from the baseline update script. This call handles concurrency control around Forge input files. The Record Store has built-in logic to handle concurrency between read and write operations, so no external concurrency control is required. Removing this call means that the lock manager does not check the flag or wait on the flag to be cleared before running a CAS crawl.
- Add a call to `runBaselineCasCrawl()` to run the full CAS crawl.
- Remove the call to `Forge.getIncomingData()` that fetches extracted data files.



For example, this baseline update script calls `CAS.runBaselineCasCrawl("MyCrawl")` which runs a full CAS crawl that writes output to a Record Store instance. Then the script continues with baseline update processing.

```
<!--
#####
# Baseline update script
#
-->
<script id="BaselineUpdate">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/baseline_update.bat</provisioned-script-
command>
  <bean-shell-script>
    <![CDATA[
log.info("Starting baseline update script.");
// obtain lock
if (LockManager.acquireLock("update_lock")) {

    // call the baseline crawl script to run a full CAS
    // crawl.
    CAS.runBaselineCasCrawl("MyCrawl");

    // clean directories
    Forge.cleanDirs();
    PartialForge.cleanCumulativePartials();
    Dgidx.cleanDirs();

    // fetch extracted data files to forge input
    Forge.getIncomingData();
    LockManager.removeFlag("baseline_data_ready");

    // fetch config files to forge input
    Forge.getConfig();

    // archive logs and run ITL
    Forge.archiveLogDir();
    Forge.run();
    Dgidx.archiveLogDir();
    Dgidx.run();

    // distributed index, update Dgraphs
    DistributeIndexAndApply.run();

    WorkbenchManager.cleanDirs();
    Forge.getPostForgeDimensions();
    WorkbenchManager.updateWsDimensions();

    // archive state files, index
    Forge.archiveState();
    Dgidx.archiveIndex();

    // (start or) cycle the LogServer
    LogServer.cycle();

    // release lock
    LockManager.releaseLock("update_lock");
    log.info("Baseline update script finished.");
} else {
    log.warning("Failed to obtain lock.");
}
    ]]>
  </bean-shell-script>
</script>
```

```

    ]]>
    </bean-shell-script>
</script>

```

You run the baseline update by running `baseline_update` in the `apps/<app_dir>/control` directory.

For example:

```
C:\Endeca\apps\DocApp\control>baseline_update.bat
```

## Dgraph partial update script using Forge and a CAS incremental crawl script

After running an *incremental* CAS crawl, you can run a *partial* update that incorporates the records from a Record Store instance.

To create this sequential workflow of incremental CAS crawl and then partial update, you can do the following:

- Remove the default `PartialForge.isPartialDataReady` check from the partial update script. This call handles concurrency control around Forge input files. The Record Store has built-in logic to handle concurrency between read and write operations, so no external concurrency control is required. Removing this call means that the lock manager does not check the flag or wait on the flag to be cleared before running a CAS crawl.
- Add a call `runIncrementalCasCrawl()` to run the incremental CAS crawl.
- If the pipeline does not read from sources in the Forge incoming directory, remove the call to `PartialForge.getPartialIncomingData()` that fetches extracted data files.

For example, this partial update script calls `CAS.runIncrementalCasCrawl("MyCrawl")` which runs an incremental CAS crawl named `MyCrawl`. Then the script continues with partial update processing.

```

<!--
#####
# Partial update script
#
-->
<script id="PartialUpdate">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/partial_update.bat</provisioned-script-
command>
  <bean-shell-script>
    <![CDATA[
log.info("Starting partial update script.");

// obtain lock
if (LockManager.acquireLock("update_lock")) {

    // call the partial crawl script to run an incremental
    // CAS crawl.
    CAS.runIncrementalCasCrawl("MyCrawl");

    // archive logs
    PartialForge.archiveLogDir();

    // clean directories
    PartialForge.cleanDirs();

    // fetch config files to forge input
    PartialForge.getConfig();

    // run ITL
    PartialForge.run();

```

```

// timestamp partial, save to cumulative partials dir
PartialForge.timestampPartials();
PartialForge.fetchPartialsToCumulativeDir();

// distribute partial update, update Dgraphs
DgraphCluster.cleanLocalPartialsDirs();
DgraphCluster.copyPartialUpdateToDgraphServers();
DgraphCluster.applyPartialUpdates();

// archive partials
PartialForge.archiveCumulativePartials();

// release lock
LockManager.releaseLock("update_lock");
log.info("Partial update script finished.");
} else {
log.warning("Failed to obtain lock.");
}
}]>
</bean-shell-script>
</script>

```

You run the partial update by running `partial_update` in the `apps/<app_dir>/control` directory. For example:

```
C:\Endeca\apps\DocApp\control>partial_update.bat
```

## Multiple CAS crawls and Forge updates

There are more complicated cases where multiple CAS crawls are running on their own schedules, and Forge updates are running on their own schedules. To coordinate this asynchronous workflow of CAS crawls and baseline or partial updates, you add code that calls methods in `ContentAcquisitionServerComponent`.

In your `DataIngest.xml` code, the main coordination task is one of determining how you time running CAS crawls and how you time running baseline or partial updates that consume records from those crawls. For example, suppose you have an application that runs three full CAS crawls and those records are consumed by a single baseline update. In that scenario, each of the three full crawls has its own full crawl script in `DataIngest.xml` that runs on a nightly schedule. And the `DataIngest.xml` file contains a baseline update that runs nightly to consume the latest generation of records from each of the three crawls. The `Forge.isDataReady` check is not required in the baseline update script because the source data is not locked.

## CAS-based data processing

The Deployment Template supports running baseline and partial updates using CAS as a replacement for Forge. In this processing model, the update runs a CAS crawl to produce MDEX-compatible output. This is the step that removes the need for Forge. Then the update runs `Dgidx` and updates the `Dgraphs` in an application.

### Dgraph baseline update script using CAS

You do not need to run Forge if you run a CAS crawl that is configured to produce MDEX-compatible output as part of your update process.

This example runs a *baseline* update that includes a *full* CAS crawl. The crawl writes MDEX compatible output and then the update invokes Dgidx to process the records, dimensions, and index configuration produced by the crawl. To create this sequential workflow of CAS crawl and then baseline update, you add a call to `runBaselineCasCrawl()` to run the CAS crawl.

For example, this baseline update script calls `CAS.runBaselineCasCrawl("${lastMileCrawlName}")` which runs a CAS crawl that writes MDEX-compatible output. Then the script continues with baseline update processing by running Dgidx and distributing the index files.

```
<!--
#####
# Baseline update script
#
-->
<script id="BaselineUpdate">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/baseline_update.bat</provisioned-script-command>
  <bean-shell-script>
    <![CDATA[
log.info("Starting baseline update script.");
// obtain lock
if (LockManager.acquireLock("update_lock")) {
  // clean directories
  CAS.cleanCumulativePartials();
  Dgidx.cleanDirs();

  // run crawl and archive any changes in the dvalId mappings
  CAS.runBaselineCasCrawl("${lastMileCrawlName}");
  CAS.archiveDvalIdMappingsForCrawlIfChanged("${lastMileCrawlName}");

  // archive logs and run the Indexer
  Dgidx.archiveLogDir();
  Dgidx.run();

  // distributed index, update Dgraphs
  DistributeIndexAndApply.run();

  // Upload the generated dimension values to Workbench
  WorkbenchManager.cleanDirs();
  CAS.copyOutputDimensionsFile("${lastMileCrawlName}", WorkbenchManager.getWorkbenchTempDir());
  WorkbenchManager.updateWsDimensions();

  // Upload the generated config to Workbench
  WorkbenchManager.updateWsConfig();

  // archive state files, index
  Dgidx.archiveIndex();

  // (start or) cycle the LogServer
  LogServer.cycle();
  // release lock
  LockManager.releaseLock("update_lock");
  log.info("Baseline update script finished.");
} else {
  log.warning("Failed to obtain lock.");
}
]]>
</bean-shell-script>
</script>
```

You run the baseline update by running `baseline_update` in the `apps/<app_dir>/control` directory.

For example:

```
C:\Endeca\apps\DocApp\control>baseline_update.bat
```

## Dgraph partial update script using CAS

You do not need to run Forge if you run a CAS crawl that is configured to produce MDEX-compatible output as part of your update process.

This example runs an *incremental* CAS crawl that writes MDEX compatible output and then runs a *partial* update to process data records. Remember that in an incremental CAS crawl, the index configuration and dimension value records are not processed.

To create this sequence of CAS crawl and then partial update, you add a call to `runIncrementalCasCrawl()` to run the CAS crawl. For example, this partial update script calls `CAS.runIncrementalCasCrawl("${lastMileCrawlName}")` which runs a CAS crawl that writes MDEX-compatible output. Then the script continues with update processing by running `Dgidx` and distributing the index files.



**Note:** In some applications, the `archiveDvalIdMappingsForCrawlIfChanged` call can take modest amounts of processing time (for example, typically less than 10 seconds). This method is recommended in all but the most time-sensitive partial update scenarios.

```
<!--
#####
# Partial update script
#
-->
<script id="PartialUpdate">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/partial_update.bat</provisioned-script-
command>
  <bean-shell-script>
    <![CDATA[
log.info("Starting partial update script.");
// obtain lock
if (LockManager.acquireLock("update_lock")) {

    // run crawl and archive any changes in the dvalId mappings
    CAS.runIncrementalCasCrawl("${lastMileCrawlName}");
    CAS.archiveDvalIdMappingsForCrawlIfChanged("${lastMileCrawlName}");

    // Copy the partial to the master cumulative directory
    CAS.fetchPartialsToCumulativeDir("${lastMileCrawlName}");

    // copy from srcPartials to localCumulative for authoring
    AuthoringDgraphCluster.copyPartialUpdateToDgraphServers();

    // copy from local to mdex's update-dir and trigger the update for authoring
    AuthoringDgraphCluster.applyPartialUpdates();

    // copy from srcPartials to localCumulative for live
    LiveDgraphCluster.copyPartialUpdateToDgraphServers();

    // copy from localCumulative to mdex's update-dir and trigger the update
    LiveDgraphCluster.applyPartialUpdates();
```

```

// Archive accumulated partials
CAS.archiveCumulativePartials();

// release lock
LockManager.releaseLock("update_lock");
log.info("Partial update script finished.");
} else {
log.warning("Failed to obtain lock.");
}
}]>
</bean-shell-script>
</script>

```

You run the baseline update by running `partial_update` in the `apps/<app_dir>/control` directory.

For example:

```
C:\Endeca\apps\DocApp\control>partial_update.bat
```

## CAS crawl scripts for Record Store output

This topic provides an example CAS crawl script with a crawl that is configured to write to Record Store output. To create a similar CAS crawl script in your application, add code to `AppConfig.xml` that specifies the CAS crawl to run locks the crawl (to wait for any running crawls to complete), runs the crawl, and releases the lock. Depending on your environment, you may need a script that runs a full CAS crawl and a script that runs an incremental CAS crawl.

This example `AppConfig.xml` code runs a *full* crawl that writes to a Record Store instance:

```

<!--
#####
# full crawl script
#
-->

<script id="MyCrawl_fullCrawl">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/runcommand.bat MyCrawl_fullCrawl
run</provisioned-script-command>
  <bean-shell-script>
    <![CDATA[
crawlName = "MyCrawl";

log.info("Starting full CAS crawl '" + crawlName + "'.");

// obtain lock
if (LockManager.acquireLock("crawl_lock_" + crawlName)) {

  CAS.runBaselineCasCrawl(crawlName);

  LockManager.releaseLock("crawl_lock_" + crawlName);
}
else {
  log.warning("Failed to obtain lock.");
}

log.info("Finished full CAS crawl '" + crawlName + "'.");
}]>

```

```

</bean-shell-script>
</script>

```

This example runs an *incremental* crawl that writes to a Record Store instance:

```

<!--
#####
# incremental crawl script
#
-->
<script id="MyCrawl_IncrementalCrawl">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/runcommand.bat MyCrawl_IncrementalCrawl
run</provisioned-script-command>
  <bean-shell-script>
    <![CDATA[
      crawlName = "MyCrawl";

      log.info("Starting incremental CAS crawl '" + crawlName + "'.");

      // obtain lock
      if (LockManager.acquireLock("crawl_lock_" + crawlName)) {

        CAS.runIncrementalCasCrawl(crawlName);

        LockManager.releaseLock("crawl_lock_" + crawlName);
      }
      else {
        log.warning("Failed to obtain lock.");
      }

      log.info("Finished incremental CAS crawl '" + crawlName + "'.");
    ]]>
  </bean-shell-script>
</script>

```

## CAS crawl scripts for record file output

This topic provides an example CAS crawl script with a crawl that is configured to write to record file output. To create a similar CAS crawl script in your application, add code to `DataIngest.xml` that specifies the CAS crawl to run locks the crawl (to wait for any running crawls to complete), runs the crawl, and releases the lock. Depending on your environment, you may need a script that runs a full CAS crawl and a script that runs an incremental CAS crawl.

This example `DataIngest.xml` code runs a *full* crawl that writes to record file output:

```

<!--
#####
# full crawl script
#
-->

<script id="MyCrawl_fullCrawl">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/runcommand.bat MyCrawl_fullCrawl
run</provisioned-script-command>
  <bean-shell-script>
    <![CDATA[
      crawlName = "MyCrawl";

```

```

log.info("Starting full CAS crawl '" + crawlName + "'.");

// obtain lock
if (LockManager.acquireLock("crawl_lock_" + crawlName)) {

    if (!CAS.isCrawlFileOutput(crawlName)) {
        throw new UnsupportedOperationException("The crawl " + crawlName +
            " does not have a File System output type. The only supported output
type for this script is File System.");
    }

    log.info("Starting full CAS crawl '" + crawlName + "'.");
    // Remove all files from the crawl's output directory
    CAS.cleanOutputDir(crawlName);
    CAS.runBaselineCasCrawl(crawlName);
    // Rename the output to files to include the crawl name
    // so they do not collide with the output from other crawls
    CAS.renameBaselineCrawlOutput(crawlName);

    destDir = PathUtils.getAbsolutePath(CAS.getWorkingDir(),
        CAS.getCasCrawlFullOutputDestDir());

    // create the target dir, if it doesn't already exist
    mkdirUtil = new CreateDirUtility(CAS.getAppname(),
        CAS.getEacHost(), CAS.getEacPort(), CAS.isSslEnabled());
    mkdirUtil.init(CAS.getCasCrawlOutputDestHost(), destDir, CAS.getWork-
ingDir());
    mkdirUtil.run();

    // clear the destination dir of full crawl from previous crawls
    CAS.clearFullCrawlOutputFromDestinationDir(crawlName);

    // remove previously collected incremental crawl files,
    // which are expected to be incorporated in this full crawl
    CAS.clearIncrementalCrawlOutputFromDestinationDir(crawlName);

    // copy the full crawl output to destination directory
    CAS.copyBaselineCrawlOutputToDestinationDir(crawlName);
    LockManager.releaseLock("crawl_lock_" + crawlName);
}

else {
    log.warning("Failed to obtain lock.");
}

log.info("Finished full CAS crawl '" + crawlName + "'.");
]]>
</bean-shell-script>
</script>

```

This example `DataIngest.xml` code runs an *incremental* crawl that writes to record file output:

```

<!--
#####
# incremental crawl script
#
-->
<script id="MyCrawl_IncrementalCrawl">
  <log-dir>./logs/provisioned_scripts</log-dir>
  <provisioned-script-command>./control/runcommand.bat MyCrawl_IncrementalCrawl
run</provisioned-script-command>

```



```

<bean-shell-script>
  <![CDATA[
crawlName = "MyCrawl";

log.info("Starting incremental CAS crawl '" + crawlName + "'.");

// obtain lock
if (LockManager.acquireLock("crawl_lock_" + crawlName)) {

    if (!CAS.isCrawlFileOutput(crawlName)) {
        throw new UnsupportedOperationException("The crawl " + crawlName +
            " does not have a File System output type. The only supported output
type for this script is File System.");
    }

    log.info("Starting incremental CAS crawl '" + crawlName + "'.");
    // Remove all files from the crawl's output directory
    CAS.cleanOutputDir(crawlName);
    CAS.runIncrementalCasCrawl(crawlName);
    // Timestamp and rename the output to files to include the
    // crawl name so they do not collide with the output from
    // previous incremental output from this crawl or incremental
    // output from other crawls
    CAS.renameIncrementalCrawlOutput(crawlName);

    destDir = PathUtils.getAbsolutePath(CAS.getWorkingDir(),
        CAS.getCasCrawlIncrementalOutputDestDir());

    // create the target dir, if it doesn't already exist
    mkDirUtil = new CreateDirUtility(CAS.getAppname(),
        CAS.getEacHost(), CAS.getEacPort(), CAS.isSslEnabled());
    mkDirUtil.init(CAS.getCasCrawlOutputDestHost(), destDir, CAS.getWork-
ingDir());
    mkDirUtil.run();

    // copy crawl output to destination directory
    // Note: We assume a downstream process removes incremental crawl output
    // from this directory that has already been processed.
    CAS.copyIncrementalCrawlOutputToDestinationDir(crawlName);

    LockManager.releaseLock("crawl_lock_" + crawlName);
}

else {
    log.warning("Failed to obtain lock.");
}

log.info("Finished incremental CAS crawl '" + crawlName + "'.");
  ]]>
</bean-shell-script>
</script>

```

## Report generation

Four report generation scripts are defined in the `DataIngest.xml` document.

Two of the scripts are used to generate XML reports for Oracle Endeca Workbench and two generate HTML reports that can be viewed in a browser. All scripts share similar functionality, so only one is included below, with numbered steps indicating the actions performed at each point in the script.

```
<script id="DailyReports">
  <bean-shell-script>
    <![CDATA[
      log.info("Starting daily Workbench report generation script.");
```

1. Obtain lock. The report generation script attempts to set a "report\_generator\_lock" flag in the EAC to serve as a lock or mutex. If the flag is already set, this step fails, ensuring that the report generator cannot be started more than once simultaneously, as the default report generators share input directories and working directories. The flag is removed in the case of an error or when the script completes successfully.

```
      if (LockManager.acquireLock("report_generator_lock")) {
```

2. Clean working directories. Clear any files in the report generator's input directory.

```
        // clean report gen input dir
        DailyReportGenerator.cleanInputDir();
```

3. Distribute configuration files to each server. A single copy of the Dgraph configuration files is distributed to each server specified in the configuration.

```
        DgraphCluster.copyDgraphConfigToDgraphServers();
```

4. Roll LogServer. If the LogServer is actively writing to a file and the file is required for the specified time range, the LogServer needs to be rolled in order to free up the log file. This code handles that test and invokes the roll administrative URL command on the LogServer, if necessary.

```
        // roll the logserver, if the report requires the active log file
        if (LogServer.isActive() &&
            LogServer.yesterdayIncludesLatestLogFile()) {
            LogServer.callLogserverRollUrl();
        }
```

5. Retrieve logs for specified report. The LogServer identifies log files in its output directory that are required to generate a report for the requested date range. Those files are copied to the target directory configured for the LogServer. Note that this step could be modified to include retrieving logs from multiple LogServers, if more than one is deployed.

```
        // retrieve required log files for processing
        LogServer.copyYesterdayLogFilesToTargetDir();
```

6. Update Report Generator to the appropriate time range and output file name. Oracle Endeca Workbench requires reports to be named according to a time stamp convention. The Report Generator component's provisioning is updated to specify the appropriate time range, time series and output filename. The output file path in the existing provisioning is updated to use the same path, but to use the date stamp as the filename. Files default to a ".xml" extension, though the component will attempt to retain a ".html" extension, if specified in the AppConfig.xml.

```
        // update report generator to the appropriate dates, time series
        // and to output a timestamped file, as required by Workbench
        DailyReportGenerator.updateProvisioningForYesterdayReport();
```

7. Archive logs. If one or more files were copied into the report generator's input directory, report generation will proceed. Start by archiving logs associated with the previous report generator execution.

```
        if (DailyReportGenerator.reportInputDirContainsFiles()) {
            // archive logs
            DailyReportGenerator.archiveLogDir();
```

8. Run report generator. Execute the report generation process.

```
// generate report
DailyReportGenerator.run();
```

9. Copy report to Oracle Endeca Workbench report directory. By default, Oracle Endeca Workbench reads reports from a directory in its workspace. Typically, the directory is `[ENDECA_TOOLS_CONF]/reports/[appName]/daily` or `[Endeca_TOOLS_CONF]/reports/[appName]/weekly`. Starting in Oracle Endeca Workbench 1.0.1, this location can be configured by provisioning a host named "webstudio" with a custom directory named "webstudio-report-dir." The Deployment Template provisions this directory and delivers generated reports to that location for Workbench to read. The report file (and associated charts) will be copied to this directory, as specified in the `AppConfig.xml`, which defaults to `<app dir>/reports`. Note that this step is not necessary for HTML reports, as those reports are not viewed in Oracle Endeca Workbench.

```
// copy generated report and charts
// defined in "webstudio" host and its "webstudio-report-dir"
// directory
reportHost = "webstudio";
absDestDir = PathUtils.getAbsolutePath(webstudio.getWorkingDir(),
    webstudio.getDirectory("webstudio-report-dir"));
isDaily = true;
DailyReportGenerator.copyReportToWebStudio(reportHost,
    absDestDir, isDaily);
}
else {
    log.warning("No log files for report generator to process.");
}

LockManager.releaseLock("report_generator_lock");
log.info("Finished daily Workbench report generation.");
}
else {
    log.warning("Failed to obtain lock.");
}
]]>
</bean-shell-script>
</script>
```



## Appendix A

---

# EAC Development Toolkit

The EAC Development Toolkit provides a common set of objects, a standard and robust configuration file format and a lightweight controller implementation that developers can leverage in order to implement operational controller applications. The toolkit is designed to enable quick deployment, while providing complete flexibility for developers to extend and override any part of the implementation to create custom, project-specific functionality.

## EAC Development Toolkit distribution and package contents

The EAC Development Toolkit is distributed as a set of JAR files bundled with the Deployment Template.

The toolkit consists of three JAR files and depends on two others that are distributed with this package. The following sections describe the JAR files. Details about classes and methods can be found in Javadoc distributed with the EAC Development Toolkit. These JAR files must be on the classpath of any application built using the EAC Development Toolkit.

### `eacToolkit.jar`

This JAR contains the source and compiled class files for the core EAC Development Toolkit classes. These classes encompass core EAC functionality, from which all component implementations extend. Included are low-level classes that access the EAC's central server via SOAP calls to its Web Service interface as well as higher level objects that wrap logic and data associated with hosts, components, scripts and utilities. In addition, this JAR includes the controller implementation used to load the Toolkit's application configuration file, and to invoke actions based on the configuration and the user's command line input.

### `eacComponents.jar`

This JAR contains the source and compiled class files for common implementations of Oracle Endeca components. These classes extend core functionality in `eacToolkit.jar` and implement standard versions of Forge, Dgidx, Dgraph and other components of an Oracle Endeca deployment.

### `eacHandlers.jar`

This JAR contains the source and compiled class files for parsing application configuration documents. In addition, the EAC Dev Toolkit's application configuration XML document format is defined by an XSD file packaged with this JAR. Finally, the JAR includes files required to register the schema and the toolkit's namespace with Spring, the framework used to load the toolkit's configuration.

### **spring.jar**

The toolkit uses the Spring framework for configuration management.

### **bsh-2.0b4.jar**

The toolkit uses BeanShell as the scripting language used by developers to write scripts in their application configuration documents.

## **EAC Development Toolkit usage**

The EAC Development Toolkit provides a library of classes that developers can use to develop and configure EAC scripts.

Classes in the library expose low level access to the EAC's web services and implement high level functionality common to many EAC scripts. Developers may implement applications by simply configuring functionality built in the toolkit or by extending the toolkit at any point to develop custom functionality.

This document discusses the toolkit's configuration file format, BeanShell scripting, command invocation and logging. This document does not provide a reference of the classes in the toolkit, or the functionality implemented in various objects and methods. Developers should refer to Javadoc or Java source files distributed with this package for details about the implementation.

## Appendix B

---

# Application Configuration File

The EAC toolkit uses an XML configuration file to define the elements that make up an application. In most deployments, this document will serve as the primary interface for developers and system administrators to configure, customize, and maintain a deployed application.

## Spring framework

The EAC Development Toolkit uses the Spring Framework's Inversion of Control container to load an EAC application based on configuration specified in an XML document.

A great deal of functionality and flexibility is provided in Spring's IoC Container and in the default bean definition XML file handled by Spring's `XmlBeanDefinitionReader` class. For details about either of these, refer to Spring Framework documentation and JavaDoc.

The EAC Development Toolkit uses a customized document format and includes a schema and custom XML handlers to parse the custom document format. It uses Spring to convert this customized configuration metadata into a system ready for execution. Specifically, the toolkit uses Spring to load a set of objects that represent an EAC application with the configuration specified for each object in the configuration document.

## XML schema

A customized document format is used to provide an intuitive configuration format for EAC script developers and system administrators.

However, this customization restricts the flexibility of the configuration document. The following sections describe elements available in the custom namespace defined by the `eacToolkit.xsd` XML schema. Each element name is followed by a brief description and an example configuration excerpt. For details, refer to the `eacToolkit.xsd` schema file distributed within the file `eacHandlers.jar`.

## Application elements

This section describes the application elements available in the custom namespace defined by the `eacToolkit.xsd` XML schema.

For more details, refer to the `eacToolkit.xsd` schema file distributed within the file `eacHandlers.jar`.

Element	Description
app	<p>This element defines the global application settings inherited by all other objects in the document, including application name, EAC central server host and port, data file prefix, the lock manager used by the application and whether or not SSL is enabled. In addition, this object defines global defaults for the working directory and the logs directory, which can be inherited or overridden by objects in the document.</p> <pre>&lt;app appName="myApp" eacHost="devhost.company.com" eacPort="8888"   dataPrefix="myApp" sslEnabled="false" lockManager="LockManager"   &gt;   &lt;working-dir&gt;C:\Endeca\apps\myApp&lt;/working-dir&gt;   &lt;log-dir&gt;./logs/baseline&lt;/log-dir&gt; &lt;/app&gt;</pre>
lock-manager	<p>This element defines a LockManager object used by the application to interact with the EAC's synchronization web service. Lock managers can be configured to release locks when a failure is encountered, ensuring that the system returns to a "neutral" state if a script or component fails. Multiple lock managers can be defined.</p> <pre>&lt;lock-manager id="LockManager" releaseLocksOnFailure="true" /&gt;</pre>

## Hosts

This section describes the host element available in the custom namespace defined by the `eacToolkit.xsd` XML schema.

The `host` element defines a host associated with the application, including the ID, hostname and EAC agent port of the host. Multiple host elements can be defined.

```
<host id="ITLHost" hostName="itlhost.company.com" port="8888" />
```

## Components

This section describes the component elements available in the custom namespace defined by the `eacToolkit.xsd` XML schema.

For more details, refer to the `eacToolkit.xsd` schema file distributed within the file `eacHandlers.jar`.

Element	Description
forge	<p>This element defines a Forge component, including attributes that define the functionality of the Forge process as well as custom properties and directories used to configure the functionality of the Forge object's methods. Multiple <code>forge</code> elements can be defined.</p> <pre>&lt;forge id="Forge" host-id="ITLHost"&gt;   &lt;properties&gt;     &lt;property name="numStateBackups" value="10" /&gt;     &lt;property name="numLogBackups" value="10" /&gt;   &lt;/properties&gt;   &lt;directories&gt;     &lt;directory name="incomingDataDir"&gt;./data/incoming&lt;/directo- ry&gt;</pre>



Element	Description
	<pre data-bbox="495 252 1476 672"> &lt;directory name="configDir"&gt;./data/complete_config&lt;/directory&gt; &lt;directory name="wsTempDir"&gt;./data/web_studio_temp_dir&lt;/directory&gt; &lt;/directories&gt; &lt;args&gt;   &lt;arg&gt;-vw&lt;/arg&gt; &lt;/args&gt; &lt;input-dir&gt;./data/processing&lt;/input-dir&gt; &lt;output-dir&gt;./data/forge_output&lt;/output-dir&gt; &lt;state-dir&gt;./data/state&lt;/state-dir&gt; &lt;temp-dir&gt;./data/temp&lt;/temp-dir&gt; &lt;num-partitions&gt;1&lt;/num-partitions&gt; &lt;pipeline-file&gt;./data/processing/pipeline.epx&lt;/pipeline-file&gt; &lt;/forge&gt; </pre>
forge-cluster	<p data-bbox="495 724 1445 819">This element defines a Forge cluster, including a list of ID references to the Forge components that belong to this cluster. This object can be configured to distribute data to Forge servers serially or in parallel.</p> <pre data-bbox="495 840 1421 976"> &lt;forge-cluster id="ForgeCluster" getDataInParallel="true"&gt;   &lt;forge ref="ForgeServer" /&gt;   &lt;forge ref="ForgeClient1" /&gt;   &lt;forge ref="ForgeClient2" /&gt; &lt;/forge-cluster&gt; </pre>
dgidx	<p data-bbox="495 1029 1469 1155">This element defines a Dgidx component, including attributes that define the functionality of the Dgidx process as well as custom properties and directories used to configure the functionality of the Dgidx object's methods. Multiple dgidx elements can be defined.</p> <pre data-bbox="495 1176 1226 1428"> &lt;dgidx id="Dgidx" host-id="ITLHost"&gt;   &lt;args&gt;     &lt;arg&gt;-v&lt;/arg&gt;   &lt;/args&gt;   &lt;input-dir&gt;./data/forge_output&lt;/input-dir&gt;   &lt;output-dir&gt;./data/dgidx_output&lt;/output-dir&gt;   &lt;temp-dir&gt;./data/temp&lt;/temp-dir&gt;   &lt;run-aspell&gt;true&lt;/run-aspell&gt; &lt;/dgidx&gt; </pre>
indexing-cluster	<p data-bbox="495 1480 1469 1575">This element defines an indexing cluster, including a list of ID references to the Dgidx components that belong to this cluster. This object can be configured to distribute data to indexing servers serially or in parallel.</p> <pre data-bbox="495 1596 1469 1732"> &lt;indexing-cluster id="IndexingCluster" getDataInParallel="true"&gt;   &lt;dgidx ref="Dgidx1" /&gt;   &lt;dgidx ref="Dgidx2" /&gt; &lt;/indexing-cluster&gt; </pre>
dgraph	<p data-bbox="495 1785 1469 1879">This element defines a Dgraph component, including attributes that define the functionality of the Dgraph process as well as custom properties and directories used to configure the functionality of the Dgraph object's methods. Multiple dgraph elements</p>

Element	Description
	<p>can be defined. Each <code>dgraph</code> element inherits, and potentially overrides, configuration specified in the <code>dgraph-defaults</code> element (see below).</p> <pre data-bbox="496 352 1474 638"> &lt;dgraph id="Dgraph1" host-id="MDEXHost" port="15000"&gt;   &lt;properties&gt;     &lt;property name="restartGroup" value="A" /&gt;     &lt;property name="updateGroup" value="a" /&gt;   &lt;/properties&gt;   &lt;log-dir&gt;./logs/dgraphs/Dgraph1&lt;/log-dir&gt;   &lt;input-dir&gt;./data/dgraphs/Dgraph1/dgraph_input&lt;/input-dir&gt;   &lt;update-dir&gt;./data/dgraphs/Dgraph1/dgraph_input/updates&lt;/update-dir&gt; &lt;/dgraph&gt; </pre>
dgraph-defaults	<p>This element defines the default settings inherited by all <code>dgraph</code> elements specified in the document. This enables a single point of configuration for common Dgraph configuration such as command line arguments, and script directory configuration. Only one <code>dgraph-defaults</code> element can be defined.</p> <pre data-bbox="496 842 1474 1394"> &lt;dgraph-defaults&gt;   &lt;properties&gt;     &lt;property name="srcIndexDir" value="./data/dgidx_output" /&gt;   &lt;/properties&gt;   &lt;property name="srcIndexHostId" value="ITLHost" /&gt;   &lt;property name="numLogBackups" value="10" /&gt;   &lt;directories&gt;     &lt;directory name="localIndexDir"&gt;       ./data/dgraphs/local_dgraph_input     &lt;/directory&gt;   &lt;/directories&gt;   &lt;args&gt;     &lt;arg&gt;--threads&lt;/arg&gt;     &lt;arg&gt;2&lt;/arg&gt;     &lt;arg&gt;--spl&lt;/arg&gt;     &lt;arg&gt;--dym&lt;/arg&gt;   &lt;/args&gt;   &lt;startup-timeout&gt;120&lt;/startup-timeout&gt; &lt;/dgraph-defaults&gt; </pre>
dgraph-cluster	<p>This element defines a Dgraph cluster, including a list of ID references to the Dgraph components that belong to this cluster. This object can be configured to distribute data to Dgraph servers serially or in parallel.</p> <pre data-bbox="496 1570 1474 1675"> &lt;dgraph-cluster id="DgraphCluster" getDataInParallel="true"&gt;   &lt;dgraph ref="Dgraph1" /&gt;   &lt;dgraph ref="Dgraph2" /&gt; &lt;/dgraph-cluster&gt; </pre>
logserver	<p>This element defines a LogServer component, including attributes that define the functionality of the LogServer process as well as custom properties and directories</p>

Element	Description
	<p>used to configure the functionality of the LogServer object's methods. Multiple logserver elements can be defined.</p> <pre data-bbox="496 331 1472 667"> &lt;logserver id="LogServer" host-id="ITLHost" port="15002"&gt;   &lt;properties&gt;     &lt;property name="numLogBackups" value="10" /&gt;     &lt;property name="targetReportGenDir" value="./reports/input" /&gt;   &lt;/properties&gt;   &lt;log-dir&gt;./logs/logserver&lt;/log-dir&gt;   &lt;output-dir&gt;./logs/logserver_output&lt;/output-dir&gt;   &lt;startup-timeout&gt;120&lt;/startup-timeout&gt;   &lt;gzip&gt;&gt;false&lt;/gzip&gt; &lt;/logserver&gt; </pre>
report-generator	<p>This element defines a ReportGenerator component, including attributes that define the functionality of the ReportGenerator process as well as custom properties and directories used to configure the functionality of the ReportGenerator object's methods. Multiple report-generator elements can be defined.</p> <pre data-bbox="496 867 1472 1455"> &lt;report-generator id="WeeklyReportGenerator" host-id="ITLHost"&gt;   &lt;properties&gt;     &lt;property name="webStudioReportDir" value="C:\Endeca\MDEXEngine\workspace\reports\MyApp" /&gt;     &lt;property name="webStudioReportHostId" value="ITLHost" /&gt;   &lt;/properties&gt;   &lt;log-dir&gt;./logs/report_generators/WeeklyReportGenerator&lt;/log-dir&gt;   &lt;input-dir&gt;./reports/input&lt;/input-dir&gt;   &lt;output-file&gt;./reports/weekly/report.xml&lt;/output-file&gt;   &lt;stylesheet-file&gt;     ./config/report_templates/tools_report_stylesheet.xsl   &lt;/stylesheet-file&gt;   &lt;settings-file&gt;     ./config/report_templates/report_settings.xml   &lt;/settings-file&gt;   &lt;time-range&gt;LastWeek&lt;/time-range&gt;   &lt;time-series&gt;Daily&lt;/time-series&gt;   &lt;charts-enabled&gt;&gt;true&lt;/charts-enabled&gt; &lt;/report-generator&gt; </pre>
custom-component	<p>This element defines a custom component, including custom properties and directories used to configure the functionality of the custom component object's methods. Multiple custom-component elements can be defined, though each must specify the name of the implemented class that extends <code>com.Endeca.soleng.eac.toolkit.component.CustomComponent</code>.</p> <p>The custom component is also used to implement the Configuration Manager, Workbench Manager, and IFCR components.</p> <pre data-bbox="496 1770 1472 1873"> &lt;custom-component id="IFCR" host-id="ITLHost" class="com.endeca.soleng.eac.toolkit.component.IFCRComponent"&gt;   &lt;properties&gt;     &lt;property name="repositoryUrl" value="http://local- </pre>

Element	Description
	<pre>host:8006/ifcr" /&gt;   &lt;property name="username" value="admin" /&gt;   &lt;property name="password" value="admin" /&gt;   &lt;property name="numExportBackups" value="3" /&gt; &lt;/properties&gt; &lt;/custom-component&gt;</pre>

## Related Links

[Display component status](#) on page 117

The controller provides a convenience method for displaying the status of all components defined in the configuration document.

## Utilities

This section describes the utility elements available in the custom namespace defined by the `eacToolkit.xsd` XML schema.

For more details, refer to the `eacToolkit.xsd` schema file distributed within the file `eacHandlers.jar`.

Element	Description
copy	<p>This element defines a copy utility invocation, including the source and destination and whether or not the source pattern should be interpreted recursively. Multiple <code>copy</code> elements can be defined.</p> <pre>&lt;copy id="CopyData" src-host-id="ITLHost" dest-host-id="ITLHost"   recursive="true" &gt;   &lt;src&gt;./data/incoming/*.txt&lt;/src&gt;   &lt;dest&gt;./data/processing/&lt;/dest&gt; &lt;/copy&gt;</pre>
shell	<p>This element defines a shell utility invocation, including the command to execute and the host on which the command will be executed. Multiple <code>shell</code> elements can be defined.</p> <pre>&lt;shell id="ProcessData" host-id="ITLHost" &gt;   &lt;command&gt;perl procesDataFiles.pl ./data/incoming/data.txt&lt;/command&gt; &lt;/shell&gt;</pre>
backup	<p>This element defines a backup utility invocation, including the directory to archive, how many archives should be saved and whether the archive should copy or move the source directory. Multiple <code>backup</code> elements can be defined.</p> <pre>&lt;backup id="ArchiveState" host-id="ITLHost" move="true" num-back- ups="5"&gt;   &lt;dir&gt;C:\Endeca\apps\myApp\data\state&lt;/dir&gt; &lt;/backup&gt;</pre>

Element	Description
rollback	<p>This element defines a rollback utility invocation, including the directory whose archive should be recovered. Multiple <code>rollback</code> elements can be defined.</p> <pre>&lt;rollback id="RollbackState" host-id="ITLHost"&gt;   &lt;dir&gt;./data/state&lt;/dir&gt; &lt;/rollback&gt;</pre>

## Customization/extension within the toolkit's schema

Most configuration tasks are performed by simply altering an element in the configuration document, by adding elements to the document, or by removing elements from the configuration.

These three actions enable users to alter the behavior of objects in their application, change which objects make up their application and change the way scripts acts on the objects in their application.

In addition to these simple actions, users can customize the behavior of objects in their application or create new objects while continuing to use the EAC development toolkit's XML configuration document format. The following are examples of customization that are possible within the constructs of the XML schema defined in the `eacToolkit.xsd` schema file.

### Implement a custom component

Users can develop new custom components by extending the class `com.Endeca.soleng.eac.toolkit.component.CustomComponent`. This class and its associated XML element allow any number of properties and directories to be specified and accessed by methods in the object. This customization method may be appropriate for cases where functionality needs to be developed that is not directly associated with an Oracle Endeca process.

### Extend an existing object

Users can implement customizations on top of existing objects by creating a new class that extends an object in the toolkit. Most elements in the configuration document (with the notable exception of the "app" element, which specifies global configuration, but does not directly correspond to an object instance) can specify a class attribute to override the default class associated with each element. For example, a user could implement a `MyForgeComponent` class by extending the toolkit's `ForgeComponent` class.

```
package com.Endeca.soleng.eac.toolkit.component;

import java.util.logging.Logger;

import com.Endeca.soleng.eac.toolkit.exception.AppConfigurationException;
import com.Endeca.soleng.eac.toolkit.exception.EacCommunicationException;
import com.Endeca.soleng.eac.toolkit.exception.EacComponentControlException;

public class MyForgeComponent extends ForgeComponent
{
    private static Logger log =
        Logger.getLogger(MyForgeComponent.class.getName());

    protected void getIncomingData() throws AppConfigurationException,
        EacCommunicationException, EacComponentControlException,
        InterruptedException
    {
```

```
// custom data retrieval implementation
}
}
```

The new class can override method functionality to customize the behavior of the object. As long as the new object does not require configuration elements unknown to the `ForgeComponent` from which it inherits, it can continue to use the forge element in the XML document to specify object configuration.

```
<forge class="com.Endeca.soleng.eac.toolkit.component.MyForgeComponent"
  id="CustomForge" host-id="ITLHost">
...
</forge>
```

### Implement custom functionality in BeanShell scripts

Users can implement custom functionality by writing new code in the XML document in new or existing BeanShell scripts. This form of customization can be used to add new functionality or to override functionality that is built in to toolkit objects. While this customization approach is very flexible, it can become unwieldy and hard to maintain and debug if a large amount of custom code needs to be written.

## Customization/extension beyond the toolkit's schema

Customization approaches within the existing schema will be sufficient for the majority of applications, but some developers will require even greater flexibility than can be supported by the XML document exposed by the toolkit.

This type of customization can still be achieved, by switching out of the default `eacToolkit` namespace in the XML document and leveraging the highly flexible and extensible Spring Framework bean definition format.

As an example, a developer might implement a new class, `PlainOldJavaObject`, which needs to be loaded and accessed by EAC scripts. If the object is implemented, compiled and added to the classpath, it can be loaded based on configuration in the XML document by specifying its configuration using the "spr" namespace.

```
<spr:bean id="MyPOJO" class="com.company.PlainOldJavaObject">
  <spr:constructor-arg>true</spr:constructor-arg>
  <spr:property name="Field1" value="StringValue" />
  <spr:property name="Map1">
    <spr:map>
      <spr:key>one</spr:key>
      <spr:value>1</spr:value>
      <spr:key>two</spr:key>
      <spr:value>2</spr:value>
    </spr:map>
  </spr:property>
</spr:bean>
```

## Appendix C

---

# BeanShell Scripting

The EAC Development Toolkit uses BeanShell to interpret and execute scripts defined in the app configuration document. The following sections describe the toolkit's use of the BeanShell interpreter and provide sample BeanShell script excerpts.

## Script implementation

In the toolkit, the `com.Endeca.soleng.eac.toolkit.script.Script` class implements scripts.

This class exposes simple execution logic that either uses a BeanShell interpreter to execute the script specified in the configuration file or, if no BeanShell script is specified in the script's configuration, uses the Script object's `scriptImplementation` method. By default, the `scriptImplementation` method has no logic and must be overridden by an extending class to take any action. This allows developers to leverage BeanShell to implement their scripts or to extend the Script object, overriding and implementing the `scriptImplementation` method.

By implementing scripts as BeanShell scripts configured in the toolkit's XML configuration document, developers can quickly develop and adjust scripts, and system administrators can adjust script implementations without involving developers. The scripting language should be familiar to any Java developer, as it is a Java based scripting language that can interpret strict Java code (i.e. code that could be compiled as a Java class). BeanShell also provides a few flexibilities that are not available in Java; for example, BeanShell allows developers to import classes at any point in the script, rather than requiring all imports to be defined up front. In addition, BeanShell allows variables to be declared without type specification.



**Note:** For details about BeanShell and ways in which it differs from Java, developers should refer to BeanShell documentation and Javadoc.

## BeanShell interpreter environment

The most common use of BeanShell scripts in the EAC Development Toolkit is to orchestrate the elements defined in the application configuration document.

More precisely, BeanShell scripts are used to orchestrate the execution of methods on the objects that are loaded from the configuration document. In order to enable this, when the toolkit constructs the BeanShell Interpreter environment, it sets internal variables associated with each element defined in the configuration document. While additional variables can be declared at any point in a script, this allows scripts to immediately act on objects defined in the document without declaring any variables.

Take, for example, the following configuration document:

```
<app appName="myApp" eacHost="devhost.company.com" eacPort="8888"
  dataPrefix="myApp" sslEnabled="false" lockManager="LockManager" >
  <working-dir>C:\Endeca\apps\myApp</working-dir>
  <log-dir>./logs/baseline</log-dir>
</app>

<host id="ITLHost" hostName="itlhost.company.com" port="8888" />

<copy id="CopyData" src-host-id="ITLHost" dest-host-id="ITLHost"
  recursive="true" >
  <src>./data/incoming/*.txt</src>
  <dest>./data/processing/</dest>
</copy>

<backup id="ArchiveState" host-id="ITLHost" move="true" num-backups="5">
  <dir>C:\Endeca\apps\myApp\data\state</dir>
</backup>

<forge id="Forge" host-id="ITLHost">
  <properties>
    <property name="numStateBackups" value="10" />
    <property name="numLogBackups" value="10" />
  </properties>
  <directories>
    <directory name="incomingDataDir">./data/incoming</directory>
    <directory name="configDir">./data/processing</directory>
  </directories>
  <args>
    <arg>-vw</arg>
  </args>
  <input-dir>./data/processing</input-dir>
  <output-dir>./data/forge_output</output-dir>
  <state-dir>./data/state</state-dir>
  <temp-dir>./data/temp</temp-dir>
  <num-partitions>1</num-partitions>
  <pipeline-file>./data/processing/pipeline.epx</pipeline-file>
</forge>
```

A BeanShell script defined in this document will have five variables immediately available for use: ITLHost, CopyData, ArchiveState, Forge, and log. Note that there is no variable associated with the app element in the document, as this element does not correspond to an object instance. Each of the other elements is instantiated, loaded with data based on its configuration and made available in the BeanShell interpreter. In addition, a special variable called log is always created for each script with a java.util.Logger instance.

A simple BeanShell script can then be written without importing a single class or instantiating a single variable.

```
<script id="SimpleForgeScript">
  <bean-shell-script>
    <![CDATA[
      log.info("Starting Forge script.");
      CopyData.run();
      Forge.run();
      ArchiveState.setNumBackups(Forge.getProperty("numStateBackups"));
      ArchiveState.run();
      log.info("Finished Forge script.");
    ]]>
```



```
</bean-shell-script>
</script>
```

In addition to exposing objects defined in the document, the toolkit imports and executes a default script each time a BeanShell script is invoked. If a file named "beanshell.imports" is successfully loaded as a classpath resource, that file is executed each time a BeanShell script is executed. This allows a default set of imports to be defined. For example, the following default file imports all of the classes in the toolkit, exposing them to BeanShell scripts:

```
import com.Endeca.soleng.eac.toolkit.*;
import com.Endeca.soleng.eac.toolkit.application.*;
import com.Endeca.soleng.eac.toolkit.base.*;
import com.Endeca.soleng.eac.toolkit.component.*;
import com.Endeca.soleng.eac.toolkit.component.cluster.*;
import com.Endeca.soleng.eac.toolkit.exception.*;
import com.Endeca.soleng.eac.toolkit.host.*;
import com.Endeca.soleng.eac.toolkit.logging.*;
import com.Endeca.soleng.eac.toolkit.script.*;
import com.Endeca.soleng.eac.toolkit.utility.*;
import com.Endeca.soleng.eac.toolkit.utility.perl.*;
import com.Endeca.soleng.eac.toolkit.utility.webstudio.*;
import com.Endeca.soleng.eac.toolkit.utility.wget.*;
import com.Endeca.soleng.eac.toolkit.utils.*;
```

## About implementing logic in BeanShell

BeanShell scripts will typically be used to orchestrate method execution for objects defined in the configuration document.

However, scripts can also implement logic, instantiating objects to provide a simple point of extension for developers to implement new logic without compiling additional Java classes.

For example, the following script excerpt demonstrates how a method can be defined and referenced in a script:

```
<script id="Status">
  <bean-shell-script>
    <![CDATA[

      // define function for printing component status
      import com.Endeca.soleng.eac.toolkit.component.Component;
      void printStatus( Component component ) {
        log.info(component.getAppName() + "." +
          component.getElementId() + ": " +
          component.getStatus().toString() );
      }

      // print status of forge, dgidx, logserver
      printStatus( Forge );
      printStatus( Dgidx );
      printStatus( LogServer );

      // print status for dgraph cluster
      dgraphs = DgraphCluster.getDgraphs().iterator();
      while( dgraphs.hasNext() ) {
        printStatus( dgraphs.next() );
      }

    ]]>
  </bean-shell-script>
</script>
```

```
</bean-shell-script>  
</script>
```

## Appendix D

---

# Command Invocation

The toolkit provides a simple interface for invoking commands from the command line.

## Invoke a method on an object

By default, the controller tries to invoke a method called "run" with no arguments on the specified object.

The following simple command invokes the run method on the BaselineUpdate script object:

```
java Controller --app-config AppConfig.xml BaselineUpdate
```

If a method name is specified, the controller looks for a method with that name on the specified object and invokes it. For example, the following command executes the applyIndex method on the DgraphCluster object:

```
java Controller --app-config AppConfig.xml DgraphCluster applyIndex
```

In addition to no-argument method invocation, the controller allows any number of String arguments to be passed to a method. The following example shows the releaseLock method being invoked on the LockManager object with the single String argument "update\_lock" specifying the name of the lock to release:

```
java Controller --app-config AppConfig.xml LockManager releaseLock  
update_lock
```

## Identify available methods

In order to help users identify the objects and methods available for invocation, the controller provides a help argument that can be called to list all available objects or methods available on an object.

If specified with an app configuration document, the help command displays usage and available objects:

```
java Controller --app-config AppConfig.xml --help
```

```
...
```

```
The following objects are defined in document 'AppConfig.xml':  
[To see methods available for an object, use the --help command line argument and  
specify the name of the object.]
```

```
[com.Endeca.soleng.eac.toolkit.base.LockManager]  
  LockManager  
[com.Endeca.soleng.eac.toolkit.component.ConfigManagerComponent]
```

```

    ConfigManager
  [com.Endeca.soleng.eac.toolkit.component.DgidxComponent]
    Dgidx
  [com.Endeca.soleng.eac.toolkit.component.DgraphComponent]
    Dgraph1
    Dgraph2
  [com.Endeca.soleng.eac.toolkit.component.ForgeComponent]
    Forge
    PartialForge
  [com.Endeca.soleng.eac.toolkit.component.LogServerComponent]
    LogServer
  [com.Endeca.soleng.eac.toolkit.component.ReportGeneratorComponent]
    WeeklyReportGenerator
    DailyReportGenerator
  [com.Endeca.soleng.eac.toolkit.component.cluster.DgraphCluster]
    DgraphCluster
  [com.Endeca.soleng.eac.toolkit.host.Host]
    ITLHost
    MDEXHost
  [com.Endeca.soleng.eac.toolkit.script.Script]
    BaselineUpdate
    DistributeIndexAndApply
    PartialUpdate
    DistributePartialsAndApply
    ConfigUpdate

```

The name of each object loaded from the configuration document is printed along with the object's class. To identify the available methods, the help command can be invoked again with the name of an object in the document:

```

java Controller --app-config AppConfig.xml --help DgraphCluster
...

```

The following methods are available for object 'DgraphCluster':  
 [Excluded: private, static and abstract methods; methods inherited from Object; methods with names that start with 'get', 'set' or 'is'. For details, refer to Javadoc for class com.Endeca.soleng.eac.toolkit.component.cluster.DgraphCluster.]

```

start(), stop(), removeDefinition(), updateDefinition(), cleanDirs(),
applyIndex(), applyPartialUpdates(), applyConfigUpdate(),
cleanLocalIndexDirs(), cleanLocalPartialsDirs(),
cleanLocalDgraphConfigDirs(), copyIndexToDgraphServers(),
copyPartialUpdateToDgraphServers(),
copyCumulativePartialUpdatesToDgraphServers(),
copyDgraphConfigToDgraphServers(), addDgraph(DgraphComponent)

```

Note that not all methods defined for the class `com.Endeca.soleng.eac.toolkit.component.cluster.DgraphCluster` are displayed. As the displayed message notes, methods declared as private, static or abstract are excluded, as are methods inherited from Object, getters and setters, and a few reserved methods that are known not to be useful from the command line. These restrictions are intended to make the output of this help command as useful as possible, but there are likely to be cases when developers will need to refer to Javadoc to find methods that are not displayed using the help command.

## Update application definition

By default, the controller will test the application definition in the configuration document against the provisioned definition in the EAC and update EAC provisioning if the definition in the document has changed.

This will happen by default any time any method is invoked on the command line.

System administrators may find it useful to update the definition without invoking a method. To facilitate this, a flag has been provided to perform the described definition update and exit.

```
java Controller --app-config AppConfig.xml --update-definition
```

In addition, there may be a need to invoke a method without testing the application definition. This can be accomplished by using an alternate command line argument:

```
java Controller --app-config AppConfig.xml --skip-definition  
    BaselineUpdate
```

## Remove an application

The controller provides a convenience method for removing an application from the EAC's central store.

When invoked, this action checks whether the application loaded from the configuration document is defined in the EAC. If it is, all active components are forced to stop and the application's definition is completely removed from the EAC.

```
java Controller --remove-app --app-config AppConfig.xml
```

## Display component status

The controller provides a convenience method for displaying the status of all components defined in the configuration document.

When the following method is invoked, the controller iterates over all defined components, querying the EAC for the status of each one and printing it.

```
java Controller --print-status --app-config AppConfig.xml
```

### Related Links

[Components](#) on page 104

This section describes the component elements available in the custom namespace defined by the `eacToolkit.xsd` XML schema.



# Index

## A

- AppConfig.xml file
  - schema for 22
- Application configuration 22
- Application descriptors 13
- Application settings
  - Report Generator 38
  - CAS Server 27
  - Configuration Manager 40
  - Dgidx 31
  - Dgraphs 32
  - Forges 30
  - global 24
  - hosts 24
  - IFCR 38
  - Lock Manager 25
  - log server 37
  - WorkbenchManager 39
- Applications, custom 13
- Automated deployments 12
  - custom 16

## B

- Baseline update
  - Forge flags 52
  - running sample scripts 77
- BeanShell scripting
  - about implementing logic 113
  - interpreter environment 111
  - script implementation 111

## C

- CAS Server 27
- Command invocation
  - display component status 117
  - identify available methods 115
  - method on an object 115
  - remove an application 117
  - update application definition 117
- Configuration file, application 22
- Configuration Manager 40
- Configuration overrides 44
- Configuring an application 23
- customizations
  - commonly used 21
  - introduced 21

## D

- Deploying
  - EAC application 9
  - on UNIX 9
  - on Windows 9
- Deployment Template
  - directories 11
- Development Toolkit, See EAC Development Toolkit
- Dgraph
  - clusters 32
  - enabling SSL 37
  - partial update script 86
- Dimension adapters 53
- Dimension servers 54

## E

- EAC
  - applications 9
  - deploying an EAC application 9
  - SSL-enabled 16
- EAC Development Toolkit
  - application configuration file 103
  - BeanShell scripting 111, 113
  - command invocation 115, 117
  - distribution 101
  - package contents 101
  - Spring framework 103
  - usage 102
  - XML schema 103, 104, 108, 109, 110
- deployment scripts 81

## F

- fault tolerance for components, configuring 25
- File-based deployment 12
  - custom 16
- Forge cluster 30
- Forge flags 52

## G

- Global application settings 24

## I

- Indexer adapters 53
- Indexing cluster 31
- Installer tokens 13

## L

- LockManager
  - configuring 25
  - default 24
- Log directory, default 24

## O

- Oracle Endeca Deployment Template
  - automated deployment 12, 16
  - configuration overrides 44
  - deploying XQuery modules 36
  - Dgraph partial update script 86
  - displaying version 19
  - integration with Oracle Endeca Workbench, reporting 40
  - provisioning scripts 83
  - report generation script 98
  - sample pipelines 47
  - standard Forge flags 52
  - with SSL-enabled EAC 16
- Oracle Endeca Workbench
  - reporting 40
- Output record adapters 54

## P

- Partial updates
  - Dgraph scripts 86
  - Forge flags 52
- Pipeline configuration
  - creating a new project 48
  - modifying a project 50
  - record spec 51
- polling intervals for components, configuring 26

## R

- Report generation script 98
- Report Generator 38

## S

- Sample pipeline
  - common errors 54
  - creating a new project 48
  - dimension adapters 53
  - dimension servers 54
  - Forge flags 52
  - indexer adapters 53
  - modifying a project 50
  - output record adapters 54
  - overview 47
  - record spec 51
- sample scripts
  - baseline update script 77
- scripts 81
- Spring framework 103
- SSL-enabled deployments 16

## U

- utilities, setting fault tolerance and polling intervals for 26

## V

- version of Deployment Template, displaying 19

## W

- Working directory, default 24

## X

- XML schema 103
  - application elements 103
  - components 104
  - customization 109, 110
  - extension 109, 110
  - hosts 104
  - utility elements 108
- XQuery modules, deploying 36