



Hewlett Packard
Enterprise

HPE Media Server

Software Version: 11.0

Media Server Administration Guide

Document Release Date: March 2016
Software Release Date: March 2016

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise Development LP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2016 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation Updates

HPE Big Data Support provides prompt and accurate support to help you quickly and effectively resolve any issue you may encounter while using HPE Big Data products. Support services include access to the Customer Support Site (CSS) for online answers, expertise-based service by HPE Big Data support engineers, and software maintenance to ensure you have the most up-to-date technology.

To access the Customer Support Site

- go to <https://customers.autonomy.com>

The Customer Support Site includes:

- **Knowledge Base.** An extensive library of end user documentation, FAQs, and technical articles that is easy to navigate and search.
- **Support Cases.** A central location to create, monitor, and manage all your cases that are open with technical support.
- **Downloads.** A location to download or request products and product updates.
- **Requests.** A place to request products to download or product licenses.

To contact HPE Big Data Customer Support by email or phone

- go to <http://www.autonomy.com/work/services/customer-support>

Support

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, visit the Knowledge Base on the HPE Big Data Customer Support Site. To do so, go to <https://customers.autonomy.com>, and then click **Knowledge Base**.

The Knowledge Base contains documents in PDF and HTML format as well as collections of related documents in ZIP packages. You can view PDF and HTML documents online or download ZIP packages and open PDF documents to your computer.

About this PDF Version of Online Help

This document is a PDF version of the online help. This PDF file is provided so you can easily print multiple topics from the help information or read the online help in PDF format. Because this content was originally created to be viewed as online help in a web browser, some topics may not be formatted properly. Some interactive topics may not be present in this PDF version. Those topics can be successfully printed from within the online help.

Contents

Part I: Getting Started	11
Chapter 1: Introduction	13
Media Server	13
Ingest Video	14
Encode and Stream Video	14
Analyze Video	14
Event Stream Processing	15
Output Information	15
OEM Certification	15
Media Server Architecture	16
HPE IDOL	17
Related Documentation	18
Chapter 2: Install Media Server	19
System Requirements	19
Supported Platforms	19
Install Media Server on Windows	20
Install an IDOL Component as a Service on Windows	21
Install Media Server on UNIX	22
Install an IDOL Component as a Service on Linux	24
Install a Component as a Service for a systemd Boot System	24
Install a Component as a Service for a System V Boot System	25
Set up a Database to Store Training Data	26
Use the Internal Database	26
Use an External Database	26
Supported External Databases	27
Set Up a PostgreSQL Database on Windows	27
Set Up a PostgreSQL Database on Linux	31
Set Up a MySQL Database on Windows	35
Set Up a MySQL Database on Linux	38
Configure Media Server	41
Upgrade the Database Schema	42
Licenses	43
Display License Information	44
Configure the License Server Host and Port	45
Revoke a Client License	45
Troubleshoot License Errors	46
Chapter 3: Configure Media Server	48
The Media Server Configuration File	48
Modify Configuration Parameter Values	48
Include an External Configuration File	49

Include the Whole External Configuration File	50
Include Sections of an External Configuration File	50
Include a Parameter from an External Configuration File	50
Merge a Section from an External Configuration File	51
Encrypt Passwords	52
Create a Key File	52
Encrypt a Password	52
Decrypt a Password	53
Specify Modules to Enable	54
Customize Logging	55
Validate the Configuration File	56
Chapter 4: Start and Stop Media Server	57
Start Media Server	57
Stop Media Server	57
Verify that Media Server is Running	58
GetStatus	58
GetLicenseInfo	58
Access IDOL Admin	59
Display Online Help	59
Chapter 5: Send Actions to Media Server	60
Send Actions to Media Server	60
Send Actions by Using a GET Method	60
Send Data by Using a POST Method	61
Application/x-www-form-urlencoded	61
Multipart/form-data	62
Use Asynchronous Actions	63
Store Action Queues in an External Database	64
Prerequisites	64
Configure Media Server	65
Event Handlers	66
Configure an Event Handler	67
Use XSL Templates to Transform Action Responses	68
Part II: Process Video	70
Chapter 6: Introduction	72
Configuration Overview	72
Tasks	72
Tracks	73
Records	74
Analysis Task Output Tracks	75
Create a Configuration	76
Ingestion	77
Analysis	78
Transform	78
Encoding	79

Output	79
Example Configuration	80
Example Configuration - Advanced	81
Start Processing Video	84
Verify Media Server is Processing Video	85
Stop Processing Video	85
Chapter 7: Ingest Video	86
Supported Codecs and Formats	86
Choose the Rate of Ingestion	87
Ingest a Video File or Stream	88
Ingest Video from a DirectShow Device	89
Ingest Video from Milestone XProtect	89
Chapter 8: Analyze Video	92
Analyze Speech	92
Transcribe Speech	92
Identify Speakers	94
Detect QR Codes	95
Detect, Recognize, and Analyze Faces	96
Detect Faces	96
Train Media Server to Recognize Faces	97
Select Images for Training	98
Create a Database to Contain Faces	99
Add a Face to a Database	99
Add a Face to a Database (Using Separate Steps)	99
Add a Face to a Database (Using a Single Action)	101
List the Faces in a Database	102
Update or Remove Faces and Databases	103
IDOL Admin	103
Recognize Faces	104
Obtain Demographic Information	105
Analyze Facial Expression	105
Extract Keyframes	106
Recognize Objects	107
Introduction	107
2D Object Detection	108
Train Media Server to Recognize Objects	109
Select Images for Training	110
Create a Database to Contain Objects	111
Add an Object to a Database	111
Add an Object to a Database (Using Separate Steps)	112
Add an Object to a Database (Using a Single Action)	113
Object Training Options	114
List the Objects in a Database	115
Update or Remove Objects and Databases	116
IDOL Admin	116
Recognize Objects	116

Classify Objects	118
Train Media Server to Classify Images	118
Classifier Types	119
Training Requirements	120
Create a Classifier	120
Classifier Training Options	121
Add Classes to a Classifier	121
List Classifiers and Object Classes	123
Update or Remove Object Classes and Classifiers	124
Import a Classifier	124
Classify Objects	125
Perform Optical Character Recognition (OCR)	125
Perform Color Analysis	126
Analyze Vehicles	127
Requirements for ANPR	128
Detect and Read Number Plates	128
Train Media Server to Recognize Vehicles	129
Obtain Training Images	129
Train Media Server	130
Identify Vehicle Models	131
Identify Vehicle Colors	132
Scene Analysis	133
Train the Scene Analysis Engine	134
Run Scene Analysis	134
Chapter 9: Event Stream Processing	136
Introduction to Event Stream Processing	136
Filter a Track	137
Deduplicate Records in a Track	138
Combine Tracks	141
Identify Time-Related Events in Two Tracks–And Engine	142
Identify Time-Related Events in Two Tracks–AndThen Engine	144
Identify Isolated Events–AndNot Engine	145
Identify Isolated Events–AndNotThen Engine	147
Write a Lua Script for an ESP Engine	148
Chapter 10: Transform Data	150
Resize Images	150
Change the Format of Images	151
Chapter 11: Encode Video	153
Introduction	153
Encode Video to MPEG Files or a UDP Stream	153
Encode Images	154
Store Video in a Rolling Buffer	156
Calculate Storage Requirements	156
Set Up Rolling Buffers	157
Pre-Allocate Storage for a Rolling Buffer	158

Write Video to a Rolling Buffer	158
View the Rolling Buffer Contents	160
Retrieve an HLS Playlist	160
Create a Clip from a Rolling Buffer	161
Create an Image from a Rolling Buffer	162
Use Multiple Media Servers	162
Chapter 12: Output Data	163
Introduction	163
Select Input Records	164
Combine Records into Documents	165
XSL Transformation	165
Send the Data to the External System	165
Choose How to Output Data	165
Single Record Mode	166
Time Mode	166
Event Mode	168
Bounded Event Mode	170
Output Data to Files	172
Output Data to the Process Action Response	174
Send Information over HTTP	175
Index Documents into IDOL Server	176
Insert Data into a Vertica Database	178
Insert Records into a Database	180
Before You Begin	180
Configure the Output Task	181
Troubleshooting	183
Send Documents to Connector Framework Server	183
Index Records into Broadcast Monitoring	185
Send Data to Milestone XProtect	187
Before You Begin	187
Configure Media Server	187
Configure Milestone	188
Part III: Distribute Processing	190
Chapter 13: Use Multiple Media Servers	192
Distribute Media Server Operations	192
Install DAH	192
Appendixes	193
Appendix A: Supported Languages	194
Appendix B: Pre-Trained Classifiers	195
Glossary	197

Send Documentation Feedback200

Part I: Getting Started

This section describes how to set up, configure, and run Media Server.

- ["Introduction"](#)
- ["Install Media Server"](#)
- ["Configure Media Server"](#)
- ["Start and Stop Media Server"](#)
- ["Send Actions to Media Server"](#)

Chapter 1: Introduction

This section provides an overview of Media Server.

• Media Server	13
• Media Server Architecture	16
• HPE IDOL	17
• Related Documentation	18

Media Server

Video is an example of unstructured information. Media Server enables you to use this information by analyzing video and extracting information about its content.

Media Server can be deployed for broadcast monitoring purposes, to identify when individuals or organizations appear in television broadcasts and extract information about these appearances. You might also use Media Server to catalog an existing archive of audio and video clips.

In security and surveillance deployments, human operators can become overwhelmed by the amount of data available from CCTV cameras and other sensors. Media Server provides automatic processing that identifies important events and reduces the operator's workload, helping them respond to suspicious events more effectively.

Media Server identifies events that occur in a video and automatically extracts information about those events. For example, Media Server can:

- identify the exact time of a scene change
- determine whether the audio track contains speech, and convert the speech into text
- read text that appears on-screen, including background text and subtitles
- identify logos and other objects when they appear in the video
- identify a person who appears in a video by matching their face to a database of known faces
- read number plates that appear in the video
- detect suspicious events in video from security cameras

Media Server can transform the metadata that it extracts into many output formats, including custom XML and documents that can be indexed into IDOL Server.

After the information has been indexed into IDOL Server, you can search within your video for specific events, such as the appearance of a particular speaker or discussion of a particular subject. If you are running Media Server for broadcast monitoring you can run sentiment analysis to determine whether the appearance of an individual or organization contained positive sentiment. If you are cataloging clips, you can use IDOL to search the clips and categorize them into a custom taxonomy.

The following sections provide more information about Media Server features.

Ingest Video

Media Server can ingest video from files and IP streams. Many devices, for example IP cameras, network encoders, and IPTV devices can produce IP streams.

Media Server can also request video from the following third-party video management systems:

- Milestone XProtect

Encode and Stream Video

Media Server can write the video that it has ingested to disk, as MPEG files. You can also configure rolling buffers, fixed-size storage areas on disk where the oldest content is discarded to make space for the latest. For example, if you deploy Media Server for video surveillance, you could configure a rolling buffer to store the last seven days of video from a camera.

Media Server supports encoding in *evidential* mode, and *non-evidential* mode:

- In *evidential* mode, Media Server does not encode the ingested video. The video is written to the file or rolling buffer unaltered. You might need an evidential copy of the video ingested by surveillance deployments, for legal reasons.
- In *non-evidential* mode, Media Server encodes the ingested video. This means that the video is optimized for storage and playback. You can resize the video to a lower resolution and choose to balance quality against disk space requirements.

You can also create a live UDP stream of the content ingested by Media Server.

Analyze Video

Media Server can run many types of analysis, including:

- automatic number plate recognition
- barcode recognition
- color analysis
- face detection, face recognition, demographic analysis and expression analysis
- intelligent scene analysis
- keyframe extraction
- object detection
- object classification
- optical character recognition
- speaker identification
- speech-to-text

For more information about the types of analysis that you can run, see ["Analyze Video" on page 92](#).

Event Stream Processing

You can configure Media Server to process and filter the events that it detects in video. For example, you could use Event Stream Processing (ESP) rules to identify events where the text "Breaking News" appears in a television broadcast *and* the newsreader speaks the words "election results" within 10 seconds. You can add custom logic by writing scripts using the Lua scripting language.

Output Information

Media Server can output the metadata that it extracts to many formats and systems, including:

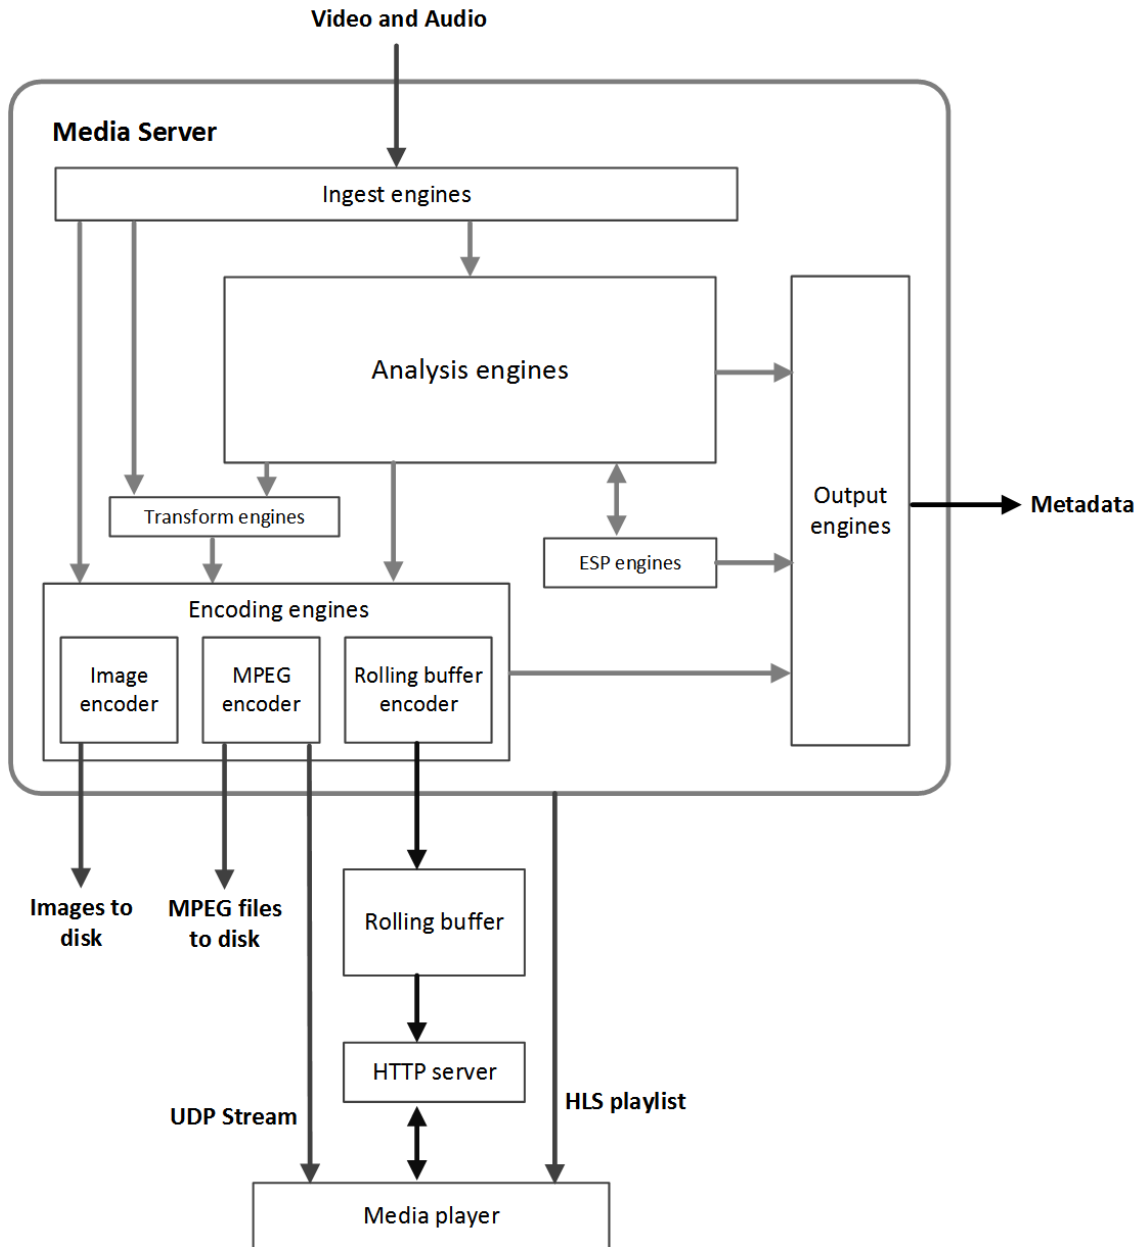
- IDOL Server
- Connector Framework Server (CFS)
- HPE Broadcast Monitoring
- Vertica
- XML
- Milestone XProtect

OEM Certification

Media Server works in OEM licensed environments.

Media Server Architecture

The following diagram shows the architecture of Media Server.



A Media Server *engine* performs a single function such as ingestion or analysis. There are several types of engine:

- **Ingest engines** demux and decode video into audio, video, and metadata.
- **Analysis engines** run analysis on ingested video, and create records that describe events in the video. Each engine performs a different type of analysis.

- **Event stream processing engines** can be used to introduce additional custom logic into video analysis. For example, you can filter or deduplicate the information produced during analysis. Event stream processing can change the schema of the data.
- **Transform engines** transform data, but do not modify its schema. For example, you can change the size of video frames before sending them to the image encoder.
- **Encoding engines** make an evidential copy of ingested video, or encode it into different formats. Some encoding engines, such as the MPEG encoder, can produce a live UDP stream that you can view in a media player.
- **Output engines** convert records into different formats and send the information to other systems such as IDOL Server or a Vertica database.

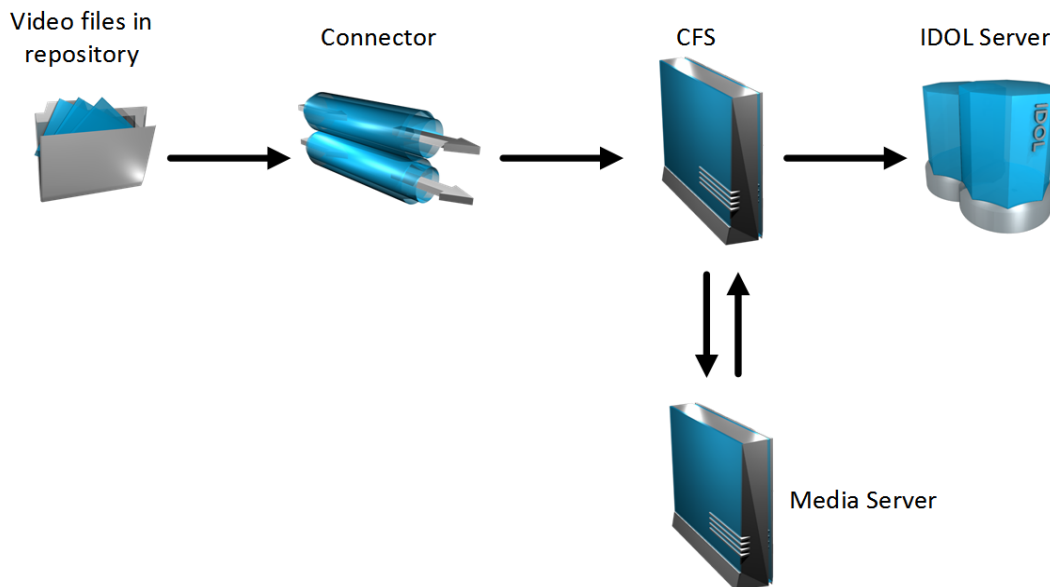
Information is passed between engines in the server. For example, an ingest engine decodes video and produces video frames and uncompressed audio for the other engines to use. Analysis engines run analysis on the decoded data. Output engines take information from the analysis engines and index it into other systems, such as IDOL Server.

HPE IDOL

Media Server is one of the components in HPE's *Intelligent Data Operating Layer* (IDOL). You can use Media Server independently or as part of a larger IDOL system.

You can use Media Server independently by writing a custom application that communicates with Media Server. Media Server accepts commands over HTTP and returns responses in XML format. You can also use the Autonomy Content Infrastructure (ACI) Client API to develop a custom application.

In a typical IDOL deployment, IDOL Connectors retrieve information from your data repositories for indexing into IDOL Server. You can configure your Connector Framework Server (CFS) to send video to Media Server and request one or more analysis operations. Media Server returns the results of the analysis operations to CFS, which enriches the information indexed into IDOL Server.



For example, a repository might contain video clips that you want to search or categorize. You could configure CFS to send the video to Media Server and request analysis such as face detection, face recognition, object detection, keyframe extraction and optical character recognition. Media Server returns information about the video content to CFS, which might perform additional operations, such as Education, before indexing the information into IDOL Server.

For more information about IDOL, refer to the *IDOL Getting Started Guide*.

Related Documentation

The following documents provide more details on Media Server.

- *Media Server Reference*
The *Media Server Reference* describes the ACI actions and configuration parameters that you can use with Media Server. For information about how to view the reference, see "[Display Online Help](#)" on page 59.
- *Media Server Scene Analysis Training Utility User Guide*
The *Media Server Scene Analysis Training Utility User Guide* describes how to train Media Server to perform Scene Analysis.
- *Media Management and Analysis Platform Installation Guide*
The HPE Media Management and Analysis Platform (MMAP) is a media analytics platform designed for viewing, searching and analyzing video footage coming from a variety of sources, typically CCTV surveillance camera footage and broadcast footage from IP streams. The *Media Management and Analysis Platform Installation Guide* provides more information about MMAP.

Chapter 2: Install Media Server

This section describes how to install Media Server.

- System Requirements 19
- Supported Platforms 19
- Install Media Server on Windows 20
- Install an IDOL Component as a Service on Windows 21
- Install Media Server on UNIX 22
- Install an IDOL Component as a Service on Linux 24
- Set up a Database to Store Training Data 26
- Upgrade the Database Schema 42
- Licenses 43

System Requirements

The system requirements for Media Server depend on the data that is being ingested and the processing tasks that are configured. HPE recommends the following minimum hardware specifications:

- A minimum of two dedicated CPUs
- 4 GB RAM

If you deploy multiple Media Servers across multiple machines, ensure that all of the machines synchronize their clocks with a time server.

Supported Platforms

- Windows x86 64
- Linux x86 64

The most fully tested versions of these platforms are:

Windows

- Windows Server 2012
- Windows Server 2008
- Windows 7

Linux

- Ubuntu 14.04
- Ubuntu 12.04
- CentOS 6

Install Media Server on Windows

Use the following procedure to install Media Server on Microsoft Windows operating systems, by using the IDOL Server installer.

The IDOL Server installer provides the major IDOL components. It also includes License Server, which Media Server requires to run.

To install Media Server

1. Double-click the appropriate installer package:

`IDOLServer_VersionNumber_Platform.exe`

where:

`VersionNumber` is the product version.

`Platform` is your software platform.


The Setup dialog box opens.

2. Click **Next**.

The License Agreement dialog box opens.

3. Read the license agreement. Select **I accept the agreement**, and then click **Next**.


The Installation Directory dialog box opens.

4. Specify the directory to install Media Server (and optionally other components such as License Server) in. By default, the system installs on `C:\HewlettPackardEnterprise\IDOLServer-VersionNumber`. Click  to choose another location. Click **Next**.

The Installation Mode dialog box opens.

5. Select **Custom**, and then click **Next**.

The License Server dialog box opens. Choose whether you have an existing License Server.

- To use an existing License Server, click **Yes**, and then click **Next**. Specify the host and ACI port of your License Server, and then click **Next**.
- To install a new instance of License Server, click **No**, and then click **Next**. Specify the ports that you want License Server to listen on, and then type the path or click  and navigate to the location of your HPE license key file (`licensekey.dat`), which you obtained when you purchased Media Server. Click **Next**.

The Component Selection dialog box opens.

6. Click **Next**.

7. Select the check boxes for the components that you want to install, and specify the port information for each component, or leave the fields blank to accept the default port settings.

For Media Server you must specify the following information:

ACI Port The port that you want Media Server to listen on, for ACI actions.

Service Port The port that you want Media Server to listen on, for service actions.

Click **Next** or **Back** to move between components.

8. After you have specified your settings, the Summary dialog box opens. Verify the settings you made and click **Next**.

The Ready to Install dialog box opens.

9. Click **Next**.

The Installing dialog box opens, indicating the progress of the installation. If you want to end the installation process, click **Cancel**.

10. After installation is complete, click **Finish** to close the installation wizard.

Install an IDOL Component as a Service on Windows

On Microsoft Windows operating systems, you can install any IDOL component as a Windows service. Installing a component as a Windows service makes it easy to start and stop the component, and you can configure a component to start automatically when you start Windows.

Use the following procedure to install Media Server as a Windows service from a command line.

To install a component as a Windows service

1. Open a command prompt with administrative privileges (right-click the icon and select **Run as administrator**).
2. Navigate to the directory that contains the component that you want to install as a service.
3. Send the following command:

```
Component.exe -install
```

where *Component.exe* is the executable file of the component that you want to install as a service.

The `-install` command has the following optional arguments:

<code>-start</code> {[auto] [manual] [disable]}	The startup mode for the component. Auto means that Windows services automatically starts the component. Manual means that you must start the service manually. Disable means that you cannot start the service. The default option is Auto .
<code>-username</code> <i>UserName</i>	The user name that the service runs under. By default, it uses a local system account.
<code>-password</code> <i>Password</i>	The password for the service user.
<code>-servicename</code> <i>ServiceName</i>	The name to use for the service. If your service name contains spaces, use quotation marks (") around the name. By default, it uses the executable name.
<code>-displayname</code> <i>DisplayName</i>	The name to display for the service in the Windows services manager. If your display name contains spaces, use quotation marks (") around the name. By default, it uses the service

```
name.  
  
-depend Dependency1  
 [,Dependency2 ...]
```

A comma-separated list of the names of Windows services that Windows must start before the new service. For example, if you are installing a Community component, you might want to add the Agentstore component and Content component as dependencies.

For example:

```
content.exe -install -servicename ContentComponent -displayname "IDOL Server  
Content Component" -depend LicenseServer
```

After you have installed the service, you can start and stop the service from the Windows Services manager.

When you no longer require a service, you can uninstall it again.

To uninstall an IDOL Windows Service

1. Open a command prompt.
2. Navigate to the directory that contains the component service that you want to uninstall.
3. Send the following command:

```
Component.exe -uninstall
```

where *Component.exe* is the executable file of the component service that you want to uninstall.

If you did not use the default service name when you installed the component, you must also add the *-servicename* argument. For example:

```
Component.exe -uninstall -servicename ServiceName
```

Install Media Server on UNIX

Use the following procedure to install Media Server in text mode on UNIX platforms.

To install Media Server on UNIX

1. Open a terminal in the directory in which you have placed the installer, and enter the following command:

```
./IDOLServer_VersionNumber_Platform.exe --mode text
```

where:

VersionNumber is the product version

Platform is the name of your UNIX platform

Note: Ensure that you have execute permission for the installer file.

The console installer starts and displays the Welcome screen.

2. Read the information and then press the `Enter` key.

The license information is displayed.

3. Read the license information, pressing `Enter` to continue through the text. After you finish reading the text, type `y` to accept the license terms.
4. Type the path to the location where you want to install the servers, or press `Enter` to accept the default path.

The Installation Mode screen is displayed.

5. Press `2` to select the Custom installation mode.

The License Server screen opens. Choose whether you have an existing License Server.

- To use an existing License Server, type `y`. Specify the host and port details for your License Server (or press `Enter` to accept the defaults), and then press `Enter`. Go to Step 7.
- To install a new instance of License Server, type `n`.

6. If you want to install a new License Server, provide information for the ports that the License Server uses.

- a. Type the value for the ACI Port and press `Enter` (or press `Enter` to accept the default value).

ACI Port The port that client machines use to send ACI actions to the License Server.

- b. Type the value for the Service Port and press `Enter` (or press `Enter` to accept the default value).

Service Port The port by which you send service actions to the License Server. This port must not be used by any other service.

- c. Type the location of your HPE license key file (`licensekey.dat`), which you obtained when you purchased Media Server. Press `Enter`.

7. The Component Selection screen is displayed. Press `Enter`. When prompted, type `y` for the components that you want to install. Specify the port information for each component, and then press `Enter`. Alternatively, leave the fields blank and press `Enter` to accept the default port settings.

For Media Server you must specify the following information:

ACI Port The port that you want Media Server to listen on, for ACI actions.

Service Port The port that you want Media Server to listen on, for service actions.

Note: These ports must not be used by any other service.

The Init Scripts screen is displayed.

8. Type the user that the server should run as, and then press `Enter`.

Note: The installer does not create this user. It must exist already.

9. Type the group that the server should run under, and then press `Enter`.

Note: If you do not want to generate init scripts for installed components, you can simply

press `Enter` to move to the next stage of the installation process without specifying a user or group.

The Summary screen is displayed.

10. Verify the settings that you made, then press `Enter` to begin installation.

The Installing screen is displayed.

This screen indicates the progress of the installation process.

The Installation Complete screen is displayed.

11. Press `Enter` to finish the installation.

Install an IDOL Component as a Service on Linux

On Linux operating systems, you can configure a component as a service to allow you to easily start and stop it. You can also configure the service to run when the machine boots. The following procedures describe how to install Media Server as a service on Linux.

Note: To use these procedures, you must have `root` permissions.

The procedure that you must use depends on the operating system and boot system type.

- For Linux operating system versions that use `systemd` (including Centos 7, and Ubuntu version 15.04 and later), see "[Install a Component as a Service for a systemd Boot System](#)" below.
- For Linux operating system versions that use System V, see "[Install a Component as a Service for a System V Boot System](#)" on the next page.

Install a Component as a Service for a systemd Boot System

To install an IDOL component as a service

1. Run the appropriate command for your Linux operating system environment to copy the init scripts to your `init.d` directory.

- Red Hat Enterprise Linux (and Centos)

```
cp IDOLInstalLDir/scripts/init/systemd/componentname /etc/systemd/system/
```

- Debian (including Ubuntu):

```
cp IDOLInstalLDir/scripts/init/systemd/componentname /lib/systemd/system/
```

componentname is the name of the init script that you want to use, which is the name of the component executable (without the file extension).

For other Linux environments, refer to the operating system documentation.

2. Run the following commands to set to appropriate access, owner, and group permissions for the component:

- Red Hat Enterprise Linux (and Centos)

```
chmod 755 /etc/systemd/system/componentname  
chown root /etc/systemd/system/componentname  
chgrp root /etc/systemd/system/componentname
```

- Debian (including Ubuntu):

```
chmod 755 /lib/systemd/system/componentname  
chown root /lib/systemd/system/componentname  
chgrp root /lib/systemd/system/componentname
```

For other Linux environments, refer to the operating system documentation.

where *componentname* is the name of the component executable that you want to run (without the file extension).

3. (Optional) If you want to start the component when the machine boots, run the following command:

```
systemctl enable componentname
```

Install a Component as a Service for a System V Boot System

To install an IDOL component as a service

1. Run the following command to copy the init scripts to your `init.d` directory.

```
cp IDOLInstallDir/scripts/init/systemv/componentname /etc/init.d/
```

where *componentname* is the name of the init script that you want to use, which is the name of the component executable (without the file extension).

2. Run the following commands to set to appropriate access, owner, and group permissions for the component:

```
chmod 755 /etc/init.d/componentname  
chown root /etc/init.d/componentname  
chgrp root /etc/init.d/componentname
```

3. (Optional) If you want to start the component when the machine boots, run the appropriate command for your Linux operating system environment:

- Red Hat Enterprise Linux (and CentOS):

```
chkconfig --add componentname  
chkconfig componentname on
```

- Debian (including Ubuntu):

```
update-rc.d componentname defaults
```

For other Linux environments, refer to the operating system documentation.

Set up a Database to Store Training Data

Media Server uses a database to store information that it requires for recognition operations, such as face recognition, object detection, or object classification. Media Server can be configured to use an internal database file or an external database hosted on a database server.

The default configuration supplied with Media Server uses an internal database and using this type of database requires no additional configuration.

You might want to use a database hosted on an external database server for the following reasons:

- **Better performance.** A database hosted on an external database server is likely to achieve significantly higher performance when your Media Server is used by multiple users and many training actions are sent to the Media Server simultaneously.
- **Sharing training data.** If you analyze large numbers of videos you can spread the load across multiple Media Servers. Multiple Media Servers can share a database hosted on an external database server so that all of the Media Servers use the same training data and you only need to maintain one database. Sharing an internal database is not supported.
- **Improved handling of concurrent requests.** When Media Server modifies data in an internal database file, it can lock the file. Any requests that need to modify the database at the same time might fail, especially if the file is locked for a significant amount of time (such as when you add large amounts of training). An external database hosted on a database server is able to handle concurrent requests.

Use the Internal Database

The default configuration supplied with Media Server stores training data in a database file (`mediaserver.db`, in the installation directory). If this file does not exist, Media Server creates it when you use an action that requires a database.

You can move or rename the database file but you will then need to change your configuration file to match the new file path.

To use an internal database

1. Open the Media Server configuration file in a text editor.
2. In the [Database] section, check that the `DatabaseType` configuration parameter is set to `internal`. This is the default setting and specifies that Media Server uses an internal database.
3. In the [Paths] section, set the `DatabasePath` parameter to the path and file name of the database file. If this file does not exist, Media Server creates an empty database at this location.
4. Save and close the configuration file.
5. Restart Media Server for your changes to take effect.

Use an External Database

This section describes how to set up an external database and configure Media Server to use that database.

Supported External Databases

The following table lists supported platforms and databases that you can use to store training data:

Media Server Operating System	Supported Databases
Windows (64-bit)	<ul style="list-style-type: none">• PostgreSQL version 9.1.x with PostgreSQL ODBC driver 09.03.0300 or later.• MySQL versions 5.x with MySQL ODBC Connector 5.3.x.
CentOS 6	<ul style="list-style-type: none">• PostgreSQL version 9.1.x with PostgreSQL ODBC driver 09.03.0300 or later and unixODBC version 2.2.14 or later.• MySQL versions 5.x with MySQL ODBC Connector 5.1.x and unixODBC version 2.2.14 or later.

Note: Other platforms might work but have not been tested.

Tip: You can also store asynchronous action queues in a PostgreSQL database. If you want to use the same PostgreSQL server to store training data and action queues, review the database requirements in the section "[Store Action Queues in an External Database](#)" on page 64, because the requirements for storing action queues might be different.

Set Up a PostgreSQL Database on Windows

To use Media Server with a PostgreSQL database, you must download and install a PostgreSQL server and ODBC driver, and configure Media Server to connect to the database through the driver.

The procedure describes setting up the database server using the psql command-line tool. If you prefer, you can use the pgAdmin graphical user interface. For more information, refer to the pgAdmin documentation on www.pgadmin.org.

To set up a PostgreSQL Media Server database on Windows

1. Download and install a PostgreSQL server. For instructions, refer to the PostgreSQL documentation on www.postgresql.org.
 - Ensure that the installation includes the PostgreSQL Unicode ODBC driver.
 - During installation, set up a user account with superuser privileges.

Note: Once installed, the PostgreSQL server appears in the Services tab in Windows Task Manager.

2. Add the PostgreSQL bin directory path to the PATH environmental variable.

Note: This step enables you to use the command `psql` to start the PostgreSQL command-line tool (psql) from the Windows Command Prompt. If the directory path is not added to the PATH variable, you must specify the `psql.exe` file path in the Command Prompt to start psql.

3. Open the psql command-line tool:
 - a. In the Windows Command Prompt, run the command:
4. Run a CREATE DATABASE command to create a new database. Specify the following database settings.

Database name	Any name.
Encoding	Must be Unicode—either UTF8 or UCS2.
Collation	Any that is compatible with the encoding.
Locale	Any that is compatible with the encoding.

For example:

```
CREATE DATABASE myDatabase WITH ENCODING 'UTF8' LC_COLLATE='English_United Kingdom' LC_CTYPE='English_United Kingdom';
```

5. Connect to the new database using the command:
6. Run the postgres.sql script provided in the Media Server installation directory. This script sets up the database schema that Media Server requires. The schema is inserted inside the public schema.
 - a. HPE recommends running the following command to ensure that the script stops running if it encounters an error:

```
\set ON_ERROR_STOP on
```

- b. Run the script using the command:

```
\i 'path/postgres.sql'
```

where *path* is the script file path.

Note: Replace backslashes in the file path with forward slashes. The psql command-line tool does not recognize backslashes in file paths.

7. Grant privileges to the user that Media Server will connect as. Required privileges are:

Database	Create Temporary Tables
All tables	Select, Insert, Update, Delete
All functions and stored procedures	Execute

If security is not a consideration, grant all privileges.

For example:

```
GRANT TEMP ON DATABASE databaseName TO userName;
```

```
GRANT SELECT, INSERT, UPDATE, DELETE ON ALL TABLES IN SCHEMA public TO  
userName;
```

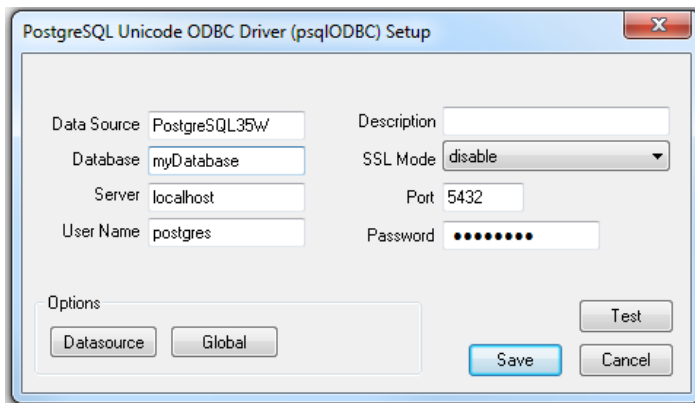
```
GRANT EXECUTE ON ALL FUNCTIONS IN SCHEMA public TO userName;
```

where,

databaseName is the name of the database that you created.

userName is the user name that Media Server will connect as.

8. Open the Data Sources (ODBC) program:
 - a. In the Windows Control Panel, click **System and Security**.
The System and Security window opens.
 - b. Click **Administrative Tools**.
The Administrative Tools window opens.
 - c. Double-click **Data Sources (ODBC)**.
The ODBC Data Source Administrator dialog box opens.
9. In the User DSN tab, click **Add...** .
The Create New Data Source dialog box opens.
10. Select the PostgreSQL Unicode driver from the list and click **Finish**.
The PostgreSQL Unicode ODBC Driver (psqlODBC) Setup dialog box opens.



11. Complete the data source information fields:
 - Data Source** The data source name (DSN). Media Server uses this string to connect to the database server.
 - Database** The name of the database that you created in [Step 2](#).
 - Server** The IP address or hostname of the server that the database server is installed on.
 - User Name** The user name to connect to the database server with.
 - Description** An optional description for the data source.
 - SSL Mode** Whether to use SSL to connect to the database server.

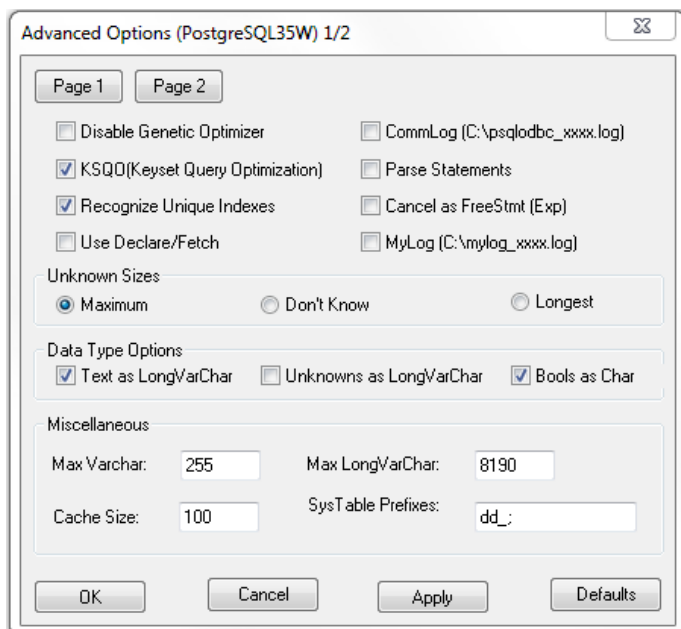
Note: To enable SSL mode, you must also configure the database server to support SSL. For instructions, refer to the PostgreSQL documentation.

Port The port to use to communicate with the database server.

Password The password for the user account that connects to the database server.

12. Click **Datasource**.

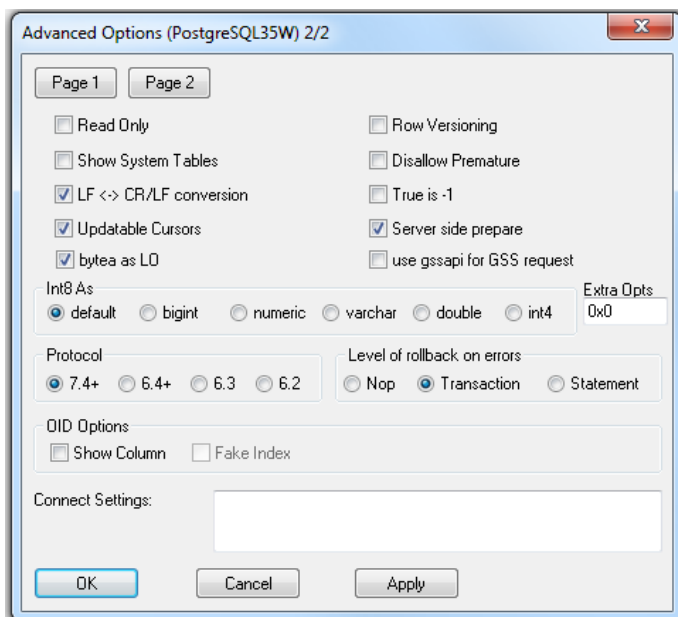
The Advanced Options (*driverName*) 1/2 dialog box opens.



13. (Optional) HPE recommends that you select the **Use Declare/Fetch** check box, to reduce memory use.

14. Click **Page 2**.

The Advanced Options (*driverName*) 2/2 dialog box opens.



15. Select the **bytea as LO** check box.
16. Click **Apply** and then **OK**.
The Advanced Options (*driverName*) 2/2 dialog box closes.
17. In the PostgreSQL Unicode ODBC Driver (psqlODBC) Setup dialog box, click **Test** to test the connection.
The Connection Test box opens containing a message describing whether the connection was successful. If the connection failed, use the information in the message to resolve any issues.
18. Click **OK** to close the Connection Test box.
19. In the PostgreSQL Unicode ODBC Driver (psqlODBC) Setup dialog box, click **Save** to close the dialog box.
20. In the ODBC Data Source Administrator dialog box, click **OK** to close the dialog box.
21. You can now configure Media Server to connect to the database (see "[Configure Media Server](#)" on page 41).

Set Up a PostgreSQL Database on Linux

To use Media Server with a PostgreSQL database, you must install a PostgreSQL server and ODBC driver, and configure Media Server to connect to the database through the driver.

The procedure describes how to set up a PostgreSQL database on a CentOS 6 distribution.

To set up a PostgreSQL Media Server database on Linux

1. Edit the .repo file to exclude PostgreSQL:
 - a. Open the CentOS-Base.repo file with a text editor. The file is usually located in `/etc/yum.repos.d`.
 - b. Add the following line to the `[base]` and `[updates]` sections:
`exclude=postgresql*`

2. Download the PostgreSQL 9.x RPM file for your Linux distribution from the PostgreSQL Yum repository on www.postgresql.org. For example:

```
curl -O http://yum.postgresql.org/9.3/redhat/rhel-5-x86_64/pgdg-centos93-9.3-1.noarch.rpm
```

3. Install the PostgreSQL RPM file by running the command:

```
sudo rpm -i RPM
```

where *RPM* is the name of the downloaded RPM file.

4. Install the required packages from the RPM file. Ensure that these include the ODBC driver. For example:

```
sudo yum install postgresql93 postgresql93-odbc
```

5. Add the PostgreSQL bin directory path to the PATH environmental variable by running the command:

```
export PATH=$PATH:binDirectoryPath
```

Note: This step enables you to use the command `psql` to start the PostgreSQL command-line tool (`psql`) from the terminal. If the directory path is not added to the PATH variable, you must specify the `psql.exe` file path in the terminal to start `psql`.

6. Initialize and start PostgreSQL.
 - a. Initialize the server by running the command:

```
sudo service postgresql-9.3 initdb
```

- b. Start the server by running the command:

```
sudo service postgresql-9.3 start
```

7. Log on to the `psql` command-line tool by running the command:

```
sudo -u postgres psql
```

8. Run a `CREATE DATABASE` command to create a new database. Specify the following database settings.

Database name	Any name.
Encoding	Must be Unicode—either UTF8 or UCS2.
Collation	Any that is compatible with the encoding.
Locale	Any that is compatible with the encoding.

For example:

```
CREATE DATABASE myDatabase WITH ENCODING 'UTF8' LC_COLLATE='en_US.UTF-8' LC_CTYPE='en_US.UTF-8';
```

9. Connect to the new database using the command:

```
\c databaseName
```


10. Run the `postgres.sql` script provided in the Media Server installation directory. This script sets up the database schema that Media Server requires. The schema is inserted inside the `public` schema.

- a. HPE recommends running the following command to ensure that the script stops running if it encounters an error:

```
\set ON_ERROR_STOP on
```

- b. Run the script using the command:

```
\i 'path/postgres.sql'
```

where *path* is the script file path.

11. Grant privileges to the user that Media Server will connect as. Required privileges are:

Database	Create Temporary Tables
All tables	Select, Insert, Update, Delete
All functions and stored procedures	Execute

If security is not a consideration, grant all privileges.

For example:

```
GRANT TEMP ON DATABASE databaseName TO userName;  

GRANT SELECT, INSERT, UPDATE, DELETE ON ALL TABLES IN SCHEMA public TO  

userName;  

GRANT EXECUTE ON ALL FUNCTIONS IN SCHEMA public TO userName;
```

where,

databaseName is the name of the database that you created.

userName is the user name that Media Server will connect as.

12. Install unixODBC driver manager version 2.2.14 or later. If using the Yum package manager, run the command:

```
sudo yum install unixODBC
```

13. Configure the data source.

- a. Open the `odbc.ini` file with a text editor. This file is usually stored in the `/etc` directory.
- b. Add the data source name in square brackets. The name can be any string. For example:

```
[PostgreSQL_1]
```

- c. Under the data source name, set the following parameters.

Parameter	Description
Driver	The driver to use.
ServerName	The IP address or hostname of the server that the database server is installed on.

Port	The port to use to communicate with the database server.
UserName	The user name to connect to the database server with.
Password	The password for the user account that connects to the database server.
Database	The name of the database that you created in Step 2 .
ByteAsLongVarBinary	You must set this parameter to 1. <div style="background-color: #f0f0f0; padding: 5px;">Caution: If this value is not set to 1, Media Server fails to start.</div>
UseDeclareFetch	(Optional) HPE recommends setting this parameter to 1, to reduce memory use.

For example:

```
[PostgreSQL_1]
Driver=PostgreSQL
ServerName=localhost
Port=5432
UserName=postgres
Password=password
Database=myDatabase
ByteAsLongVarBinary=1
UseDeclareFetch=1
```

Note: You can set other parameters in this file, but these have not been tested with Media Server.

- d. Save and close the file.
14. Configure the ODBC driver.
- a. Open the odbcinst.ini file with a text editor. This file is usually stored in the /etc directory.
 - b. If not already present, add the database server name in square brackets. For example:

```
[PostgreSQL]
```

- c. Under the database server name, set the following parameters.

Parameter	Description
Description	A description of the driver instance.
Driver64	The location of the PostgreSQL driver library file.
Setup64	The location of the driver installer file.
FileUsage	Set this parameter to 1.

For example:

```
[PostgreSQL]
Description=ODBC for PostgreSQL
Driver64=/usr/pgsql-9.3/lib/psqlodbc.so
Setup64=/usr/lib64/libodbcpsqlS.so
FileUsage=1
```

Note: You can set other parameters in this file, but these have not been tested with Media Server.

- d. Save and close the file.
15. You can now configure Media Server to connect to the database (see "[Configure Media Server](#)" on page 41).

Set Up a MySQL Database on Windows

To use Media Server with a MySQL database, you must download and install a MySQL server and ODBC driver, and configure Media Server to connect to the database through the driver.

To set up a MySQL Media Server database on Windows

1. Download and install a MySQL server and MySQL Connector/ODBC 5.x (which contains the Unicode driver). For instructions, refer to the MySQL documentation on www.mysql.com.
 - During installation, set up a user account with superuser privileges.

Note: Once installed, the MySQL server appears in the Services tab in Windows Task Manager.

2. Configure the database server for use with Media Server:
 - a. Open the configuration or options file for the MySQL server (usually named `my.ini`).
 - b. So that Media Server can send large amounts of binary data (images) to the database, set the configuration parameter `max_allowed_packet=67108864`.
 - c. Save and close the configuration file.
3. Add the MySQL `bin` directory path to the `PATH` environmental variable.

Note: This step enables you to use the command `mysql` to start the `mysql` command-line tool from the Windows Command Prompt. If the directory path is not added to the `PATH` variable, you must specify the `mysql.exe` file path in the Command Prompt to start `psql`.

4. Open the `mysql` command line tool:
 - a. In the Windows Command Prompt, run the command:

```
mysql -u userName -p
```
 - b. Enter your password when prompted.
5. Run a `CREATE DATABASE` command to create a new database. Specify the following database settings.

Database name	Any name.
Character set	Must be Unicode—either UTF8 or UCS2.
Collation	Any that is compatible with the encoding.

For example:

```
CREATE DATABASE myDatabase CHARACTER SET utf8 COLLATE utf8_unicode_ci;
```

6. Run the `my.sql` script provided in the Media Server installation directory. This script sets up the database schema that Media Server requires.

- a. Close the `mysql` command-line tool:

```
quit
```

- b. In the Windows Command Prompt, run the following command:

```
mysql -u userName -p -v -D databaseName -e "source path/my.sql"
```

where,

userName is the MySQL user name.

databaseName is the name of the database you created in [Step 2](#).

path is the path to the `my.sql` file.

Note: Running the script non-interactively from the terminal ensures that the script terminates if an error occurs.

- c. Enter your password when prompted.

7. Grant privileges to the user that Media Server will connect as. Required privileges are:

Database	Create Temporary Tables
All tables	Select, Insert, Update, Delete
All functions and stored procedures	Execute

If security is not a consideration, grant all privileges.

- a. Start the `mysql` command-line tool:

```
mysql
```

- b. Run the GRANT commands:

```
GRANT CREATE TEMPORARY TABLES ON databaseName.* TO userName;
```

```
GRANT SELECT, INSERT, UPDATE, DELETE ON databaseName.* TO userName;
```

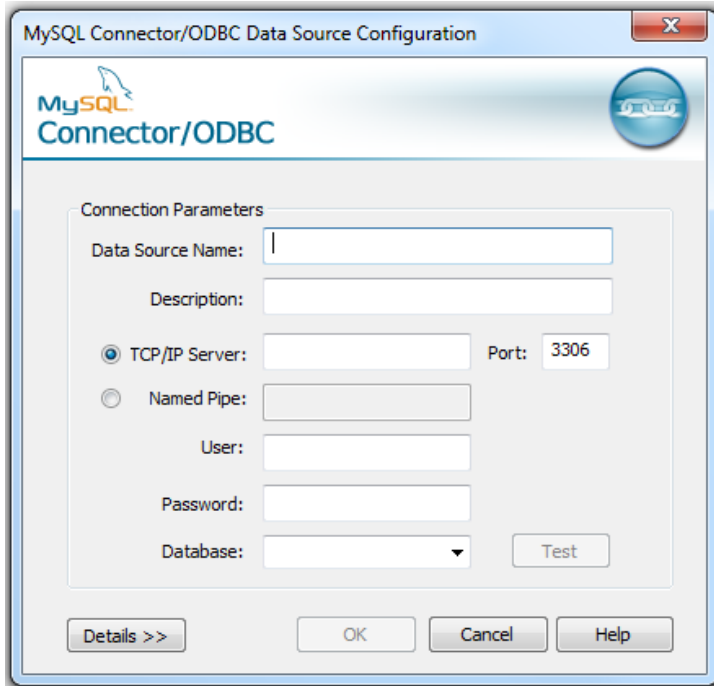
```
GRANT EXECUTE ON databaseName.* TO userName;
```

where,

databaseName is the name of the database you created in [Step 2](#).

userName is the user name that Media Server will connect as.

- c. Close the mysql command-line tool:
quit
8. Open the Data Sources (ODBC) program:
 - a. In the Windows Control Panel, click **System and Security**.
The System and Security window opens.
 - b. Click **Administrative Tools**.
The Administrative Tools window opens.
 - c. Double-click **Data Sources (ODBC)**.
The ODBC Data Source Administrator dialog box opens.
9. In the User DSN tab, click **Add...** .
The Create New Data Source dialog box opens.
10. Select the MySQL ODBC Unicode driver from the list and click **Finish**.
The MySQL Connector/ODBC Data Source Configuration dialog box opens.



11. Complete the Connection Parameters fields:
 - Data Source Name** The data source name (DSN). Choose any string. Media Server can use this string to connect to the database server.
 - Description** An optional description for the data source.
 - TCP/IP Server** The IP address or hostname of the server that the database server is installed on.

Port	The port to use to communicate with the database server.
User	The user name to connect to the database server with.
Password	The password for the user account that connects to the database server.
Database	The name of the database that you created in Step 2 .

12. Click **Test** to test the connection.

The Connection Test box opens containing a message describing whether the connection was successful. If the connection failed, use the information in the message to resolve any issues.

13. Click **OK** to close the Connection Test box.
14. In the MySQL Connector/ODBC Data Source Configuration dialog box, click **OK** to close the dialog box.
15. In the ODBC Data Source Administrator dialog box, click **OK** to close the dialog box.
16. You can now configure Media Server to connect to the database (see "[Configure Media Server](#)" on [page 41](#)).

Set Up a MySQL Database on Linux

To use Media Server with a MySQL database, you must install a MySQL server and ODBC driver, and configure Media Server to connect to the database through the driver.

The procedure describes how to set up a MySQL database on a CentOS 6 distribution.

To set up a MySQL Media Server database on Linux

1. Install a MySQL server. (Ensure that the package includes the mysql command-line tool.)
For instructions, refer to the MySQL documentation on www.mysql.com.
2. Configure the database server for use with Media Server:
 - a. Open the configuration or options file for the MySQL server (usually named `my.ini`).
 - b. So that Media Server can send large amounts of binary data (images) to the database, set the configuration parameter `max_allowed_packet=67108864`.
 - c. Save and close the configuration file.
3. Add the MySQL bin directory path to the PATH environmental variable by running the command:

```
export PATH=$PATH:binDirectoryPath
```

Note: This step enables you to use the command `mysql` to start the mysql command-line tool from the terminal. If the directory path is not added to the PATH variable, you must specify the `mysql.exe` file path in the terminal to start mysql.

4. Start the mysql command-line tool. In the terminal, run the command:

```
mysql
```
5. Run a `CREATE DATABASE` command to create a new database. Specify the following database settings.

Database name	Any name.
Character set	Must be Unicode—either UTF8 or UCS2.
Collation	Any that is compatible with the encoding.

For example:

```
CREATE DATABASE myDatabase CHARACTER SET utf8 COLLATE utf8_unicode_ci;
```

6. Run the `my.sql` script provided in the Media Server installation directory. This script sets up the database schema that Media Server requires.
 - a. Close the `mysql` command-line tool:

```
quit
```

- b. In the terminal, run the command:

```
mysql -u userName -p -v -D databaseName -e "source path/my.sql"
```

where,

userName is the MySQL user name.

databaseName is the name of the database you created in [Step 3](#).

path is the script file path.

Note: Running the script non-interactively from the terminal ensures that the script terminates if an error occurs.

7. Grant privileges to the user that Media Server will connect to the MySQL server as. Required privileges are:

Database	Create Temporary Tables
All tables	Select, Insert, Update, Delete
All functions and stored procedures	Execute

If security is not a consideration, grant all privileges.

- a. Start the `mysql` command-line tool:

```
mysql
```

- b. Run the GRANT commands:

```
GRANT CREATE TEMPORARY TABLES ON databaseName.* TO userName;
```

```
GRANT SELECT, INSERT, UPDATE, DELETE ON databaseName.* TO username;
```

```
GRANT EXECUTE ON databaseName.* TO username;
```

where,

databaseName is the name of the database you created in [Step 2](#).

userName is the user name that Media Server will connect as.

- c. Close the mysql command-line tool:

```
quit
```

- 8. Install unixODBC driver manager version 2.2.14 or later. If you have the relevant Yum repository, you can run the command in the terminal:

```
sudo yum install unixODBC
```

- 9. Install the MySQL driver. If you have the relevant Yum repository, you can run the command in the terminal:

```
sudo yum install mysql-connector-odbc
```

- 10. Configure the data source.

- a. Open the `odbc.ini` file with a text editor. This file is usually stored in the `/etc` directory.
- b. Add the data source name in square brackets. The name can be any string. For example:

```
[MySQL_1]
```

- c. Under the data source name, set the following parameters.

Parameter	Description
Driver	The driver to use.
Server	The IP address or hostname of the server that the database server is installed on.
Port	The port to use to communicate with the database server.
User	The user name to connect to the database server with.
Password	The password for the user account that connects to the database server.
Database	The name of the database that you created in Step 3 .

For example:

```
[MySQL_1]  
Driver=MySQL  
Server=localhost  
Port=5432  
User=mysql  
Password=password  
Database=myDatabase
```

Note: You can set other parameters in this file, but these have not been tested with Media Server.

- d. Save and close the file.

- 11. Configure the ODBC driver.

- a. Open the `odbcinst.ini` file with a text editor. This file is usually stored in the `/etc` directory.
- b. If not already present, add the database server name in square brackets. For example:

[MySQL]

- a. Set the following parameters.

Parameter	Description
Description	A description of the driver instance.
Driver64	The location of the MySQL driver library file.
Setup64	The location of the driver installer file.
FileUsage	Set this parameter to 1.

For example:

```
[MySQL]
Description=ODBC for MySQL
Driver64=/usr/lib64/libmyodbc5.so
Setup64=/usr/lib64/libodbcmyS.so
FileUsage=1
```

Note: You can set other parameters in this file, but these have not been tested with Media Server.

- b. Save and close the file.
12. You can now configure Media Server to connect to the database (see ["Configure Media Server" below](#)).

Configure Media Server

To configure Media Server to use an external database

1. Stop Media Server, if it is running.
2. Open the Media Server configuration file (`mediaserver.cfg`) with a text editor.
3. Find the [Database] section of the configuration file. Create this section if it does not exist.
4. Set the following parameters:

DatabaseType	The type of database to use, either <code>mysql</code> or <code>postgres</code> .
ODBCConnectionString	The ODBC Connection string to use to connect to the database. For example: <pre>DSN=MyDatabase; Driver = {PostgreSQL UNICODE}; Server = <i>IPAddress</i>; Port = <i>port</i>; Database = <i>myDatabase</i>; Uid = <i>myUsername</i>; Pwd = <i>myPassword</i>; Driver = {MySQL ODBC 5.x UNICODE Driver}; Server = <i>IPAddress</i>; Database = <i>myDatabase</i>; User = <i>myUsername</i>; Password = <i>myPassword</i>; Option = 3;</pre>

For example:

```
[Database]  
DatabaseType=postgres  
ODBCConnectionString=DSN=MyDatabase;
```

5. If you are running Media Server on Linux, set the following parameter:

ODBCDriverManager The unixODBC Driver Manager shared object file.

For example:

```
ODBCDriverManager=libodbc.so
```

6. Save and close the configuration file.
7. Start Media Server.

Upgrade the Database Schema

Sometimes the schema of the Media Server database must change in order to provide new features or enhancements. If you are using an internal database, any schema changes are applied automatically. If you are using a database that is hosted on an external database server, you must run an upgrade script when you upgrade Media Server.

HPE provides a script to upgrade to the latest version of the database schema from each of the earlier versions. The following table describes the schema changes for the Media Server database.

Schema version	Media Server version	Script to run to upgrade to latest schema
2	10.11.x	You are using the latest database schema
1	10.10.x	mysql-upgrade_from_v1.sql (for MySQL databases) postgres-upgrade_from_v1.sql (for PostgreSQL databases)

Running one of these scripts copies your training data and adds it to the database using the latest schema. These scripts do not remove training data stored using earlier schema versions. This means that you can continue to use the database with an earlier version of Media Server. Be aware that if you use multiple versions of Media Server, any new training you perform is only added to the database using the schema for the Media Server that performs the training. For example, if you upgrade from schema version 1 to schema version 2, you can perform training with Media Server 10.10.x and Media Server 10.11.x. However, any training you perform with Media Server 10.10.x is available only to Media Server version 10.10.x, and any training that you perform using Media Server 10.11.x is available only to Media Server version 10.11.x.

After a successful schema upgrade you can remove data stored using earlier schema versions. This saves storage space on the database server. HPE provides scripts to remove the data, named `mysql-purge_before_vX.sql` (for MySQL databases) or `postgres-purge_before_vX.sql` (for PostgreSQL databases), where *X* is the oldest schema version you want to retain in the database. For example, if you upgrade from schema version 1 to schema version 2, and do not want to use Media Server 10.10.x

again, you can run `mysql-purge_before_v2.sql` or `postgres-purge_before_v2.sql` to remove schema version 1 from the database.

To upgrade the database schema

1. In the table above, find the version of Media Server that you are upgrading from.
2. Run the corresponding upgrade script for your database, using the same command syntax as used to create the database (see the following topics):
 - ["Set Up a PostgreSQL Database on Windows" on page 27](#)
 - ["Set Up a PostgreSQL Database on Linux" on page 31](#)
 - ["Set Up a MySQL Database on Windows" on page 35](#)
 - ["Set Up a MySQL Database on Linux" on page 38](#)

Note: Run the upgrade script using the `psql` command-line tool (for PostgreSQL databases) or the `mysql` command-line tool (for MySQL databases). The script contains instructions that are only supported when the script runs through these tools.

3. Start Media Server 11.0, and run training and analysis as normal.
4. (Optional) When you have confirmed that the upgrade was successful, remove the training data stored using earlier schema versions by running the relevant script, either `mysql-purge_before_vX.sql`, or `postgres-purge_before_vX.sql`, where *X* is the oldest schema version you want to retain in the database.

Licenses

To use HPE IDOL solutions, you must have a running HPE License Server, and a valid license key file for the products that you want to use. Contact HPE Big Data Support to request a license file for your installation.

License Server controls the IDOL licenses, and assigns them to running components. License Server must run on a machine with a static, known IP address, MAC address, or host name. The license key file is tied to the IP address and ACI port of your License Server and cannot be transferred between machines. For more information about installing License Server and managing licenses, see the *License Server Administration Guide*.

When you start Media Server, it requests a license from the configured License Server. You must configure the host and port of your License Server in the Media Server configuration file.

You can revoke the license from a product at any time, for example, if you want to change the client IP address or reallocate the license.

Caution: Taking any of the following actions causes the licensed module to become inoperable.

You **must not**:

- Change the IP address of the machine on which a licensed module runs (if you use an IP address to lock your license).

- Change the service port of a module without first revoking the license.
- Replace the network card of a client without first revoking the license.
- Remove the contents of the `license` and `uid` directories.

All modules produce a `license.log` and a `service.log` file. If a product fails to start, check the contents of these files for common license errors. See "[Troubleshoot License Errors](#)" on page 46.

Display License Information

You can verify which modules you have licensed either by using the IDOL Admin interface, or by sending the `LicenseInfo` action from a web browser.

To display license information in IDOL Admin

- In the **Control** menu of the IDOL Admin interface for your License Server, click **Licenses**. The **Summary** tab displays summary information for each licensed component, including:
 - The component name.
 - The number of seats that the component is using.
 - The total number of available seats for the component.
 - (Content component only) The number of documents that are currently used across all instances of the component.
 - (Content component only) The maximum number of documents that you can have across all instances of the component.

The **Seats** tab displays details of individual licensed seats, and allows you to revoke licenses.

To display license information by sending the `LicenseInfo` action

- Send the following action from a web browser to the running License Server.

```
http://LicenseServerHost:Port/action=LicenseInfo
```

where:

LicenseServerHost is the IP address of the machine where License Server resides.

Port is the ACI port of License Server (specified by the `Port` parameter in the `[Server]` section of the License Server configuration file).

In response, License Server returns the requested license information. This example describes a license to run four instances of IDOL Server.

```
<?xml version="1.0" encoding="UTF-8" ?>
<autnresponse xmlns:autn="http://schemas.autonomy.com/aci/">
  <action>LICENSEINFO</action>
  <response>SUCCESS</response>
  <responsedata>
```

```
<LicenseDiSH>
  <LICENSEINFO>
    <autn:Product>
      <autn:ProductType>IDOLSERVER</autn:ProductType>
      <autn:TotalSeats>4</autn:TotalSeats>
      <autn:SeatsInUse>0</autn:SeatsInUse>
    </autn:Product>
  </LICENSEINFO>
</LicenseDiSH>
</responsedata>
</autnresponse>
```

Configure the License Server Host and Port

Media Server is licensed through HPE License Server. In the Media Server configuration file, specify the information required to connect to the License Server.

To specify the license server host and port

1. Open your configuration file in a text editor.
2. In the [License] section, modify the following parameters to point to your License Server.

LicenseServerHost The host name or IP address of your License Server.

LicenseServerACIPort The ACI port of your License Server.

For example:

```
[License]
LicenseServerHost=licenses
LicenseServerACIPort=20000
```

3. Save and close the configuration file.

Revoke a Client License

After you set up licensing, you can revoke licenses at any time, for example, if you want to change the client configuration or reallocate the license. The following procedure revokes the license from a component.

Note: If you cannot contact the client (for example, because the machine is inaccessible), you can free the license for reallocation by sending an `AdminRevokeClient` action to the License Server. For more information, see the *License Server Administration Guide*.

To revoke a license

1. Stop the HPE solution that uses the license.
2. At the command prompt, run the following command:

```
InstallDir/ExecutableName[.exe] -revokelicense -configfile cfgFilename
```

This command returns the license to the License Server.

You can send the LicenseInfo action from a web browser to the running License Server to check for free licenses. In this sample output from the action, one IDOL Server license is available for allocation to a client.

```
<autn:Product>
  <autn:ProductType>IDOLSERVER</autn:ProductType>
  <autn:Client>
    <autn:IP>192.123.51.23</autn:IP>
    <autn:ServicePort>1823</autn:ServicePort>
    <autn:IssueDate>1063192283</autn:IssueDate>
    <autn:IssueDateText>10/09/2003 12:11:23</autn:IssueDateText>
  </autn:Client>
  <autn:TotalSeats>2</autn:TotalSeats>
  <autn:SeatsInUse>1</autn:SeatsInUse>
</autn:Product>
```

Troubleshoot License Errors

The table contains explanations for typical licensing-related error messages.

License-related error messages

Error message	Explanation
Error: Failed to update license from the license server. Your license cache details do not match the current service configuration. Shutting the service down.	The configuration of the service has been altered. Verify that the service port and IP address have not changed since the service started.
Error: License for <i>ProductName</i> is invalid. Exiting.	The license returned from the License Server is invalid. Ensure that the license has not expired.
Error: Failed to connect to license server using cached licensed details.	Cannot communicate with the License Server. The product still runs for a limited period; however, you should verify whether your License Server is still available.
Error: Failed to connect to license server. Error code is SERVICE: <i>ErrorCode</i>	Failed to retrieve a license from the License Server or from the backup cache. Ensure that your License Server can be contacted.
Error: Failed to decrypt license keys. Please contact Autonomy support. Error code is SERVICE:<i>ErrorCode</i>	Provide HPE Big Data Support with the exact error message and your license file.
Error: Failed to update the license from the license server. Shutting down	Failed to retrieve a license from the License Server or from the backup cache. Ensure that your License Server can be contacted.
Error: Your license keys are invalid. Please	Your license keys appear to be out of sync.

License-related error messages, continued

Error message	Explanation
contact Autonomy support. Error code is <code>SERVICE:ErrorCode</code>	Provide HPE Big Data Support with the exact error message and your license file.
Failed to revoke license: No license to revoke from server.	The License Server cannot find a license to revoke.
Failed to revoke license from server <code>LicenseServer Host:LicenseServerPort</code>. Error code is <code>ErrorCode</code>	Failed to revoke a license from the License Server. Provide HPE Big Data Support with the exact error message.
Failed to revoke license from server. An instance of this application is already running. Please stop the other instance first.	You cannot revoke a license from a running service. Stop the service and try again.
Failed to revoke license. Error code is <code>SERVICE:ErrorCode</code>	Failed to revoke a license from the License Server. Provide HPE Big Data Support with the exact error message.
Your license keys are invalid. Please contact Autonomy Support. Error code is <code>ACISERVER:ErrorCode</code>	Failed to retrieve a license from the License Server. Provide HPE Big Data Support with the exact error message and your license file.
Your product ID does not match the generated ID.	Your installation appears to be out of sync. Forcibly revoke the license from the License Server and rename the <code>license</code> and <code>uid</code> directories.
Your product ID does not match this configuration.	The service port for the module or the IP address for the machine appears to have changed. Check your configuration file.

Chapter 3: Configure Media Server

This section describes how to configure Media Server.

- The Media Server Configuration File 48
- Modify Configuration Parameter Values 48
- Include an External Configuration File 49
- Encrypt Passwords 52
- Specify Modules to Enable 54
- Customize Logging 55
- Validate the Configuration File 56

The Media Server Configuration File

The configuration file is named `mediaserver.cfg`, and is located in the Media Server installation directory. You can modify the configuration file to customize the operation of Media Server.

The configuration file must include some parameters, such as those that specify the ports to use and those that configure the connection to the License Server.

The Media Server configuration file includes the following sections:

- [License] Contains parameters that configure the connection to your License Server.
- [Logging] Contains parameters that determine how messages are logged. You can configure *log streams* to send different types of message to separate log files.
- [Paths] Contains parameters that specify the location of files required by Media Server.
- [Server] Contains general settings for Media Server. Specifies the ACI port of the Media Server and contains parameters that control the way the connector handles ACI requests.
- [Service] Contains settings that determine which machines can use and control the Media Server service.

For a complete list of parameters that you can use in the configuration file, refer to the *Media Server Reference*.

Modify Configuration Parameter Values

You modify Media Server configuration parameters by directly editing the parameters in the configuration file. When you set configuration parameter values, you must use UTF-8.

Caution: You must stop and restart Media Server for new configuration settings to take effect.

This section describes how to enter parameter values in the configuration file.

Enter Boolean Values

The following settings for Boolean parameters are interchangeable:

TRUE = true = ON = on = Y = y = 1

FALSE = false = OFF = off = N = n = 0

Enter String Values

To enter a comma-separated list of strings when one of the strings contains a comma, you can indicate the start and the end of the string with quotation marks, for example:

```
ParameterName=cat,dog,bird,"wing,beak",turtle
```

Alternatively, you can escape the comma with a backslash:

```
ParameterName=cat,dog,bird,wing\,beak,turtle
```

If any string in a comma-separated list contains quotation marks, you must put this string into quotation marks and escape each quotation mark in the string by inserting a backslash before it. For example:

```
ParameterName="<font face=\"arial\" size=\"+1\"><b>\",<p>"
```

Here, quotation marks indicate the beginning and end of the string. All quotation marks that are contained in the string are escaped.

Include an External Configuration File

You can share configuration sections or parameters between ACI server configuration files. The following sections describe different ways to include content from an external configuration file.

You can include a configuration file in its entirety, specified configuration sections, or a single parameter.

When you include content from an external configuration file, the `GetConfig` and `ValidateConfig` actions operate on the combined configuration, after any external content is merged in.

In the procedures in the following sections, you can specify external configuration file locations by using absolute paths, relative paths, and network locations. For example:

```
../sharedconfig.cfg  
K:\sharedconfig\sharedsettings.cfg  
\\example.com\shared\idol.cfg  
file://example.com/shared/idol.cfg
```

Relative paths are relative to the primary configuration file.

Note: You can use nested inclusions, for example, you can refer to a shared configuration file that references a third file. However, the external configuration files must not refer back to your original configuration file. These circular references result in an error, and Media Server does not start.

Similarly, you cannot use any of these methods to refer to a different section in your primary configuration file.

Include the Whole External Configuration File

This method allows you to import the whole external configuration file at a specified point in your configuration file.

To include the whole external configuration file

1. Open your configuration file in a text editor.
2. Find the place in the configuration file where you want to add the external configuration file.
3. On a new line, type a left angle bracket (<), followed by the path and name to the external configuration file, in quotation marks (""). You can use relative paths and network locations. For example:

```
< "K:\sharedconfig\sharedsettings.cfg"
```

4. Save and close the configuration file.

Include Sections of an External Configuration File

This method allows you to import one or more configuration sections from an external configuration file at a specified point in your configuration file. You can include a whole configuration section in this way, but the configuration section name in the external file must exactly match what you want to use in your file. If you want to use a configuration section from the external file with a different name, see ["Merge a Section from an External Configuration File" on the next page](#).

To include sections of an external configuration file

1. Open your configuration file in a text editor.
2. Find the place in the configuration file where you want to add the external configuration file section.
3. On a new line, type a left angle bracket (<), followed by the path and name to the external configuration file, in quotation marks (""). You can use relative paths and network locations. After the configuration file name, add the configuration section name that you want to include. For example:

```
< "K:\sharedconfig\extrasettings.cfg" [License]
```

Note: You cannot include a section that already exists in your configuration file.

4. Save and close the configuration file.

Include a Parameter from an External Configuration File

This method allows you to import a parameter from an external configuration file at a specified point in your configuration file. You can include a section or a single parameter in this way, but the value in the external file must exactly match what you want to use in your file.

To include a parameter from an external configuration file

1. Open your configuration file in a text editor.
2. Find the place in the configuration file where you want to add the parameter from the external configuration file.
3. On a new line, type a left angle bracket (<), followed by the path and name to the external configuration file, in quotation marks (""). You can use relative paths and network locations. After the configuration file name, add the name of the configuration section name that contains the parameter, followed by the parameter name. For example:

```
< "license.cfg" [License] LicenseServerHost
```

To specify a default value for the parameter, in case it does not exist in the external configuration file, specify the configuration section, parameter name, and then an equals sign (=) followed by the default value. For example:

```
< "license.cfg" [License] LicenseServerHost=localhost
```

4. Save and close the configuration file.

Merge a Section from an External Configuration File

This method allows you to include a configuration section from an external configuration file as part of your Media Server configuration file. For example, you might want to specify a standard SSL configuration section in an external file and share it between several servers. You can use this method if the configuration section that you want to import has a different name to the one you want to use.

To merge a configuration section from an external configuration file

1. Open your configuration file in a text editor.
2. Find or create the configuration section that you want to include from an external file. For example:

```
[SSLOptions1]
```

3. After the configuration section name, type a left angle bracket (<), followed by the path and name to the external configuration file, in quotation marks (""). You can use relative paths and network locations. For example:

```
[SSLOptions1] < "../sharedconfig/ssloptions.cfg"
```

If the configuration section name in the external configuration file does not match the name that you want to use in your configuration file, specify the section to import after the configuration file name. For example:

```
[SSLOptions1] < "../sharedconfig/ssloptions.cfg" [SharedSSLOptions]
```

In this example, Media Server uses the values in the [SharedSSLOptions] section of the external configuration file as the values in the [SSLOptions1] section of the Media Server configuration file.

Note: You can include additional configuration parameters in the section in your file. If these parameters also exist in the imported external configuration file, Media Server uses the values

in the local configuration file. For example:

```
[SSLOptions1] < "ssloptions.cfg" [SharedSSLOptions]  
SSLCACertificatesPath=C:\IDOL\HTTPConnector\CACERTS\  

```

4. Save and close the configuration file.

Encrypt Passwords

HPE recommends that you encrypt all passwords that you enter into a configuration file.

Create a Key File

A key file is required to use AES encryption.

To create a new key file

1. Open a command-line window and change directory to the Media Server installation folder.
2. At the command line, type:

```
outpassword -x -tAES -oKeyFile=./MyKeyFile.ky
```

A new key file is created with the name MyKeyFile.ky

Caution: To keep your passwords secure, you must protect the key file. Set the permissions on the key file so that only authorized users and processes can read it. Media Server must be able to read the key file to decrypt passwords, so do not move or rename it.

Encrypt a Password

The following procedure describes how to encrypt a password.

To encrypt a password

1. Open a command-line window and change directory to the Media Server installation folder.
2. At the command line, type:

```
outpassword -e -tEncryptionType [-oKeyFile] [-cFILE -sSECTION -pPARAMETER]  
PasswordString
```

where:

Option	Description
-t <i>EncryptionType</i>	The type of encryption to use: <ul style="list-style-type: none">• Basic• AES For example: -tAES

Option	Description
	Note: AES is more secure than basic encryption.
<code>-oKeyFile</code>	AES encryption requires a key file. This option specifies the path and file name of a key file. The key file must contain 64 hexadecimal characters. For example: <code>-oKeyFile=./key.ky</code>
<code>-cFILE -sSECTION -pPARAMETER</code>	(Optional) You can use these options to write the password directly into a configuration file. You must specify all three options. <ul style="list-style-type: none"> • <code>-c</code>. The configuration file in which to write the encrypted password. • <code>-s</code>. The name of the section in the configuration file in which to write the password. • <code>-p</code>. The name of the parameter in which to write the encrypted password. For example: <code>-c./Config.cfg -sMyTask -pPassword</code>
<code>PasswordString</code>	The password to encrypt.

For example:

```
autpassword -e -tBASIC MyPassword
```

```
autpassword -e -tAES -oKeyFile=./key.ky MyPassword
```

```
autpassword -e -tAES -oKeyFile=./key.ky -c./Config.cfg -sDefault -pPassword MyPassword
```

The password is returned, or written to the configuration file.

Decrypt a Password

The following procedure describes how to decrypt a password.

To decrypt a password

1. Open a command-line window and change directory to the Media Server installation folder.
2. At the command line, type:

```
autpassword -d -tEncryptionType [-oKeyFile] PasswordString
```

where:

Option	Description
<code>-t EncryptionType</code>	The type of encryption: <ul style="list-style-type: none"> • Basic

Option	Description
	<ul style="list-style-type: none">• AES For example: <code>-tAES</code>
<code>-oKeyFile</code>	AES encryption and decryption requires a key file. This option specifies the path and file name of the key file used to decrypt the password. For example: <code>-oKeyFile=./key.ky</code>
<code>PasswordString</code>	The password to decrypt.

For example:

```
autopassword -d -tBASIC 9t3M3t7awt/J8A
```

```
autopassword -d -tAES -oKeyFile=./key.ky 9t3M3t7awt/J8A
```

The password is returned in plain text.

Specify Modules to Enable

You can disable Media Server modules to reduce the amount of memory used and the time required for the server to start.

If licensed, the following modules are enabled by default, but you can disable them:

Barcode	Barcode detection module
FaceAnalyze	Face analysis module
FaceDetect	Face detection module
FaceRecognize	Face recognition module
Object	Object detection module
ObjectClass	Object class recognition module
OCR	Optical Character Recognition (OCR) module

To specify the modules to enable

1. Open the Media Server configuration file in a text editor.
2. Find the `[Modules]` configuration section. If the section does not exist, add it by typing `[Modules]` on a new line.
3. Add the `Enable` parameter and specify a comma-separated list of modules to enable. For example:

```
Enable=barcode,ocr,object
```

4. Save and close the configuration file.

When you restart Media Server, only the specified modules are enabled, in addition to the modules that are always enabled.

Customize Logging

You can customize logging by setting up your own *log streams*. Each log stream creates a separate log file in which specific log message types (for example, action, index, application, or import) are logged.

To set up log streams

1. Open the Media Server configuration file in a text editor.
2. Find the [Logging] section. If the configuration file does not contain a [Logging] section, add one.
3. In the [Logging] section, create a list of the log streams that you want to set up, in the format *N=LogStreamName*. List the log streams in consecutive order, starting from 0 (zero). For example:

```
[Logging]
LogLevel=FULL
LogDirectory=logs
0=ApplicationLogStream
1=ActionLogStream
```

You can also use the [Logging] section to configure any default values for logging configuration parameters, such as `LogLevel`. For more information, see the *Media Server Reference*.

4. Create a new section for each of the log streams. Each section must have the same name as the log stream. For example:

```
[ApplicationLogStream]
[ActionLogStream]
```

5. Specify the settings for each log stream in the appropriate section. You can specify the type of logging to perform (for example, full logging), whether to display log messages on the console, the maximum size of log files, and so on. For example:

```
[ApplicationLogStream]
LogTypeCSVs=application
LogFile=application.log
LogHistorySize=50
LogTime=True
LogEcho=False
LogMaxSizeKBs=1024

[ActionLogStream]
LogTypeCSVs=action
LogFile=logs/action.log
LogHistorySize=50
LogTime=True
LogEcho=False
LogMaxSizeKBs=1024
```

6. Save and close the configuration file. Restart the service for your changes to take effect.

Validate the Configuration File

You can use the `ValidateConfig` service action to check for errors in the configuration file.

Note: For the `ValidateConfig` action to validate a configuration section, Media Server must have previously read that configuration. In some cases, the configuration might be read when a task is run, rather than when the component starts up. In these cases, `ValidateConfig` reports any unread sections of the configuration file as unused.

To validate the configuration file

- Send the following action to Media Server:

```
http://Host:ServicePort/action=ValidateConfig
```

where:

Host is the host name or IP address of the machine where Media Server is installed.

ServicePort is the service port, as specified in the `[Service]` section of the configuration file.

Chapter 4: Start and Stop Media Server

The following sections describe how to start and stop Media Server.

- [Start Media Server](#)57
- [Stop Media Server](#)57
- [Verify that Media Server is Running](#)58
- [Access IDOL Admin](#)59
- [Display Online Help](#)59

Start Media Server

Note: Your License Server must be running before you start Media Server.

To start Media Server

- Start Media Server from the command line using the following command:

```
mediaserver.exe -configfile configname.cfg
```

where the optional `-configfile` argument specifies the path of a configuration file that you want to use.

- On Windows, if you have installed Media Server as a service, start the service from the Windows Services dialog box.
- On UNIX, if you have installed Media Server as a service, use one of the following commands:
 - On machines that use `systemd`:

```
systemctl start mediaserver
```
 - On machines that use `system V`:

```
service mediaserver start
```

Tip: On both Windows and UNIX platforms, you can configure services to start automatically when you start the machine.

Stop Media Server

You can stop Media Server by using one of the following procedures.

To stop Media Server

- Send the Stop service action to the service port.

```
http://host:ServicePort/action=stop
```

where:

host is the host name or IP address of the machine where Media Server is installed.

ServicePort is the Media Server service port (specified in the [Service] section of the configuration file).

- On Windows platforms, if Media Server is running as a service, stop Media Server from the Windows Services dialog box.
- On UNIX platforms, if Media Server is running as a service, use one of the following commands:
 - On machines that use systemd:
`systemctl stop mediaserver`
 - On machines that use system V:
`service mediaserver stop`

Verify that Media Server is Running

After starting Media Server, you can run the following actions to verify that Media Server is running.

- [GetStatus](#)
- [GetLicenseInfo](#)

GetStatus

You can use the `GetStatus` service action to verify the Media Server is running. For example:

```
http://Host:ServicePort/action=GetStatus
```

Note: You can send the `GetStatus` action to the ACI port instead of the service port. The `GetStatus` ACI action returns information about the Media Server setup.

GetLicenseInfo

You can send a `GetLicenseInfo` action to Media Server to return information about your license. This action checks whether your license is valid and returns the operations that your license includes.

Send the `GetLicenseInfo` action to the Media Server ACI port. For example:

```
http://Host:ACIport/action=GetLicenseInfo
```

The following result indicates that your license is valid.

```
<autn:license>  
  <autn:validlicense>true</autn:validlicense>  
</autn:license>
```

As an alternative to submitting the `GetLicenseInfo` action, you can view information about your license, and about licensed and unlicensed actions, on the **License** tab in the Status section of IDOL Admin.

Access IDOL Admin

IDOL Admin is a utility that you can use to help monitor and configure Media Server. You can access IDOL Admin from a Web browser. For more information about IDOL Admin, refer to the *IDOL Admin User Guide*.

To access IDOL Admin

- Type the following URL into the address bar of your Web browser:

```
http://host:port/action=admin
```

where:

host is the name or IP address of the host that Media Server is installed on.

port is the Media Server ACI port.

Display Online Help

You can display the Media Server Reference by sending an action from your web browser. The Media Server Reference describes the actions and configuration parameters that you can use with Media Server.

For Media Server to display help, the help data file (`help.dat`) must be available in the installation folder.

To display help for Media Server

1. Start Media Server.
2. Send the following action from your web browser:

```
http://host:port/action=Help
```

where:

host is the IP address or name of the machine on which Media Server is installed.

port is the ACI port by which you send actions to Media Server (set by the `Port` parameter in the `[Server]` section of the configuration file).

For example:

```
http://12.3.4.56:9000/action=help
```

Chapter 5: Send Actions to Media Server

This section describes how to send actions to Media Server.

- [Send Actions to Media Server](#) 60
- [Use Asynchronous Actions](#) 63
- [Store Action Queues in an External Database](#) 64
- [Event Handlers](#) 66
- [Use XSL Templates to Transform Action Responses](#) 68

Send Actions to Media Server

You can make requests to Media Server by using either GET or POST HTTP request methods.

- A GET request sends parameter-value pairs in the request URL. GET requests are appropriate for sending actions that retrieve information from Media Server, such as the `GetStatus` action. For more information, see ["Send Actions by Using a GET Method" below](#).
- A POST request sends parameter-value pairs in the HTTP message body of the request. POST requests are appropriate for sending data to Media Server. In particular, you must use POST requests to upload and send files to Media Server. For more information, see ["Send Data by Using a POST Method" on the next page](#).

HPE recommends that you send video data to Media Server by providing a URL or a path to a file. Video files are very large and sending them over HTTP, as either multipart/form-data or a base64 encoded string, could cause Media Server to slow and potentially lead to interruptions in service. Image, audio, and text files are generally much smaller and so you can send these files over HTTP.

Note: Media Server accepts HTTP strings that contain a maximum of 64,000 characters, and uploaded files with a maximum size of 10,000,000 bytes. You can override the default settings by setting the `MaxInputString` and `MaxFileUploadSize` configuration parameters in the `[Server]` section of the configuration file.

Send Actions by Using a GET Method

You can use GET requests to send actions that retrieve information from Media Server.

When you send an action using a GET method, you use a URL of the form:

`http://host:port/?action=action¶meters`

where:

- `host` is the IP address or name of the machine where Media Server is installed.
- `port` is the Media Server ACI port.
- `action` is the name of the action that you want to run. The `action` (or `a`) parameter must be the

first parameter in the URL string, directly after the host and port details.

parameters are the required and optional parameters for the action. These parameters can follow the *action* parameter in any order.

You must:

- Separate each parameter from its value with an equals symbol (=).
- Separate multiple values with a comma (,).
- Separate each parameter-value pair with an ampersand (&).

For more information about the actions that you can use with Media Server, refer to the *Media Server Reference*.

GET requests can send only limited amounts of data and cannot send files directly. However, you can set a parameter to a file path if the file is on a file system that Media Server can access. Media Server must also be able to read the file.

Send Data by Using a POST Method

You can send files and binary data directly to Media Server using a POST request. One possible way to send a POST request over a socket to Media Server is using the cURL command-line tool.

The data that you send in a POST request must adhere to specific formatting requirements. You can send only the following content types in a POST request to Media Server:

- `application/x-www-form-urlencoded`
- `multipart/form-data`

Tip: Media Server rejects POST requests larger than the size specified by the configuration parameter `MaxFileUploadSize`.

Application/x-www-form-urlencoded

The `application/x-www-form-urlencoded` content type describes form data that is sent in a single block in the HTTP message body. Unlike the query part of the URL in a GET request, the length of the data is unrestricted. However, Media Server rejects requests that exceed the size specified by the configuration parameter `MaxFileUploadSize`.

This content type is inefficient for sending large quantities of binary data or text containing non-ASCII characters, and does not allow you to upload files. For these purposes, HPE recommends sending data as `multipart/form-data` (see "[Multipart/form-data](#)" on the next page).

In the request:

- Separate each parameter from its value with an equals symbol (=).
- Separate multiple values with a comma (,).
- Separate each parameter-value pair with an ampersand (&).
- Base-64 encode any binary data.
- URL encode all non-alphanumeric characters, including those in base-64 encoded data.

Example

The following example base-64 encodes the file `image.jpg` into a file `imagedata.dat` and sends it (using cURL) as `application/x-www-form-urlencoded` data to Media Server located on the `localhost`, using port `14000`. The example action adds the image to a new face in an existing face database named `politicians`.

1. Base-64 encode the image.
 - On Windows, create a Powershell script called `base64encode.ps1` that contains the following:

```
Param([String]$path)
[convert]::ToBase64String((get-content $path -encoding byte))
```

Then from the command line, run the script using Powershell:

```
powershell.exe base64encode.ps1 image.jpg > imagedata.dat
```

- On Linux:

```
base64 -w 0 image.jpg > imagedata.dat
```

2. Send the request to Media Server:

```
curl http://localhost:14000
  -d 'action=TrainFace&database=politicians&identifier=president'
  --data-urlencode imagedata@imagedata.dat
```

You can send multiple images in a single request by separating the binary data for each image by a comma (in `imagedata.dat`).

Multipart/form-data

In the `multipart/form-data` content type, the HTTP message body is divided into parts, each containing a discrete section of data.

Each message part requires a header containing information about the data in the part. Each part can contain a different content type; for example, `text/plain`, `image/png`, `image/gif`, or `multipart/mixed`. If a parameter specifies multiple files, you must specify the `multipart/mixed` content type in the part header.

Encoding is optional for each message part. The message part header must specify any encoding other than the default (7BIT).

Multipart/form-data is ideal for sending non-ASCII or binary data, and is the only content type that allows you to upload files. For more information about form data, see <http://www.w3.org/TR/html401/interact/forms.html>.

Note: In cURL, you specify each message part using the `-F` (or `--form`) option. To upload a file in a message part, prefix the file name with the `@` symbol. For more information on cURL syntax, see the cURL documentation.

Example

The following example sends `image.jpg` (using `cURL`) as `multipart/form-data` to Media Server located on the `localhost`, using port `14000`. The example action adds the image to a new face in an existing face database named `politicians`.

```
curl http://localhost:18000 -F action=TrainFace
                           -F database=politicians
                           -F identifier=president
                           -F imagedata=@image.jpg
```

Use Asynchronous Actions

When you send an asynchronous action, Media Server adds the task to a queue and returns a token that you can use to check its status. Media Server performs the task when a thread becomes available.

Check the Status of an Asynchronous Action

To check the status of an asynchronous action, use the token that was returned by Media Server with the `QueueInfo` action. For more information about the `QueueInfo` action, see the *Media Server Reference*.

To check the status of an asynchronous action

- Send the `QueueInfo` action to Media Server with the following parameters.

<code>QueueName</code>	The name of the action queue that you want to check.
<code>QueueAction</code>	The action to perform. Set this parameter to <code>GetStatus</code> .
<code>Token</code>	(Optional) The token that the asynchronous action returned. If you do not specify a token, Media Server returns the status of every action in the queue.

For example:

```
/action=QueueInfo&QueueName=...&QueueAction=getstatus&Token=...
```

Cancel an Asynchronous Action that is Queued

To cancel an asynchronous action that is waiting in a queue, use the following procedure.

To cancel an asynchronous action that is queued

- Send the `QueueInfo` action to Media Server with the following parameters.

<code>QueueName</code>	The name of the action queue that contains the action to cancel.
<code>QueueAction</code>	The action to perform . Set this parameter to <code>Cancel</code> .
<code>Token</code>	The token that the asynchronous action returned.

Stop an Asynchronous Action that is Running

You can stop an asynchronous action at any point.

To stop an asynchronous action that is running

- Send the `QueueInfo` action to Media Server with the following parameters.

<code>QueueName</code>	The name of the action queue that contains the action to stop.
<code>QueueAction</code>	The action to perform. Set this parameter to <code>Stop</code> .
<code>Token</code>	The token that the asynchronous action returned.

Store Action Queues in an External Database

Media Server provides asynchronous actions. Each asynchronous action has a queue to store requests until threads become available to process them. You can configure Media Server to store these queues either in an internal database file, or in an external database hosted on a database server.

The default configuration stores queues in an internal database. Using this type of database does not require any additional configuration.

You might want to store the action queues in an external database so that several servers can share the same queues. In this configuration, sending a request to any of the servers adds the request to the shared queue. Whenever a server is ready to start processing a new request, it takes the next request from the shared queue, runs the action, and adds the results of the action back to the shared database so that they can be retrieved by any of the servers. You can therefore distribute requests between components without configuring a Distributed Action Handler (DAH).

Note: You cannot use multiple servers to process a single request. Each request is processed by one server.

Tip: You can also store training data used by Media Server in an external database. You can use the same database server to store training data and asynchronous action queues. However, you must create separate databases for storing training data and action queues. As a result you must also create separate data source names (DSNs), and connection strings. For information about configuring Media Server to use an external database for training data, see ["Configure Media Server" on page 41](#).

Prerequisites

- PostgreSQL 9.0 or later.

Tip: This is the earliest version of PostgreSQL that you can use to store action queues in a PostgreSQL database. If you want to use the same PostgreSQL server to store training data, review the database requirements in the section ["Set up a Database to Store Training Data" on page 26](#), because a later version might be required.

- You must set the PostgreSQL ODBC driver setting `MaxVarChar` to 0 (zero). If you use a DSN, you can configure this parameter when you create the DSN. Otherwise, you can set the `MaxVarcharSize` parameter in the connection string.

Configure Media Server

To configure Media Server to use a shared action queue, follow these steps.

To store action queues in an external database

1. Stop Media Server, if it is running.
2. Open the Media Server configuration file.
3. Find the relevant section in the configuration file:
 - To store queues for all asynchronous actions in the external database, find the `[Actions]` section.
 - To store the queue for a single asynchronous action in the external database, find the section that configures that action.
4. Set the following configuration parameters.

`AsyncStoreLibraryDirectory` The path of the directory that contains the library to use to connect to the database. Specify either an absolute path, or a path relative to the server executable file.

`AsyncStoreLibraryName` The name of the library to use to connect to the database. You can omit the file extension. The following libraries are available:

- `postgresAsyncStoreLibrary` - for connecting to a PostgreSQL database.

`ConnectionString` The connection string to use to connect to the database. The user that you specify must have permission to create tables in the database. For example:

```
ConnectionString=Driver={PostgreSQL};  
Server=10.0.0.1; Port=9876;  
Database=SharedActions; Uid=user; Pwd=password;  
MaxVarcharSize=0;
```

or

```
ConnectionString=DSN=myDSN
```

Tip: Do not use the same database, data source name, and connection string that you use for storing training data. You must create a separate database, DSN, and connection string.

For example:

```
[Actions]
AsyncStoreLibraryDirectory=acidlls
AsyncStoreLibraryName=postgresAsyncStoreLibrary
ConnectionString=DSN=ActionStore
```

5. If you are using the same database to store action queues for more than one type of component, set the following parameter in the [Actions] section of the configuration file.

DatastoreSharingGroupName The group of components to share actions with. You can set this parameter to any string, but the value must be the same for each server in the group. For example, to configure several Media Servers to share their action queues, set this parameter to the same value in every Media Server configuration. HPE recommends setting this parameter to the name of the component.

Caution: Do not configure different components (for example, two different types of connector) to share the same action queues. This will result in unexpected behavior.

For example:

```
[Actions]
...
DatastoreSharingGroupName=ComponentType
```

6. Save and close the configuration file.
When you start Media Server it connects to the shared database.

Event Handlers

Some of the actions that you can send to Media Server are asynchronous. Asynchronous actions do not run immediately, but are added to a queue. This means that the person or application that sends the action does not receive an immediate response. However, you can configure Media Server to call an event handler when an asynchronous action starts, finishes, or encounters an error.

You might use an event handler to:

- Return data about an event back to the application that sent the action.
- Write event data to a text file, to log any errors that occur.

Media Server can call an event handler for the following events.

OnStart The `OnStart` event is called when Media Server starts processing an asynchronous action from the queue.

OnFinish The `OnFinish` event is called when Media Server successfully finishes processing an asynchronous action.

OnError The `OnError` event is called when an asynchronous action fails and cannot continue.

Configure an Event Handler

Use the following procedure to configure an event handler.

To configure an event handler

1. Open the Media Server configuration file in a text editor.
2. Use the `OnStart`, `OnFinish`, or `OnError` parameter to specify the name of a section in the configuration file that contains event handler settings for the corresponding event.
 - To run an event handler for all asynchronous actions, set these parameters in the `[Actions]` section. For example:

```
[Actions]
OnStart=NormalEvents
OnFinish=NormalEvents
OnError=ErrorEvents
```

- To run an event handler for a specific action, set these parameters in the `[ActionName]` section, where `ActionName` is the name of the action. The following example runs an event handler when the `Process` action starts and finishes successfully:

```
[Process]
OnStart=NormalEvents
OnFinish=NormalEvents
```

3. Create a new section in the configuration file to contain the settings for your event handler. You must name the section using the name you specified with the `OnStart`, `OnFinish`, or `OnError` parameter.
4. In the new section, set the following parameters.

`LibraryName` (Required) The name of the library to use to handle the event.

- To write event data to a text file, set this parameter to `TextFileHandler`, and then set the `FilePath` parameter to specify the path of the file.
- To send event data to a URL, set this parameter to `HttpHandler`, and then use the HTTP event handler parameters to specify the URL, proxy server settings, credentials, and so on.

For example:

```
[NormalEvents]
LibraryName=TextFileHandler
FilePath=./events.txt

[ErrorEvents]
LibraryName=HttpHandler
URL=http://handlers:8080/10-proxy/callback.htm?
```

5. Save and close the configuration file. You must restart Media Server for your changes to take effect.

The following XML is an example of the ACI response sent to the event handler on the OnFinish event:

```
<action>
  <status>Finished</status>
  <queued_time>2009-Jun-08 17:57:29</queued_time>
  <time_in_queue>0</time_in_queue>
  <process_start_time>2012-Jun-08 17:57:29</process_start_time>
  <time_processing>86</time_processing>
  <token>MTkyLjE2OC45NS4yNDoxMTAwMDpJTkdFU1Q6LTEyOTc5NjgxODY=</token>
</action>
```

Use XSL Templates to Transform Action Responses

You can transform the action responses returned by Media Server using XSL templates. You must write your own XSL templates and save them with either an `.xsl` or `.tmpl` file extension. For more information about XSL, see <http://www.w3.org/TR/xslt>.

After creating the templates, you must configure Media Server to use them, and then apply them to the relevant actions.

To enable XSL transformations

1. Ensure that the `autnxslt` library is located in the same directory as your configuration file. If the library is not included in your installation, you can obtain it from HPE Support.
2. Open the Media Server configuration file in a text editor.
3. In the `[Server]` section, set the `XSLTemplates` parameter to `True`.

Caution: If `XSLTemplates` is set to `True` and the `autnxslt` library is not present in the same directory as the configuration file, the server will not start.

4. In the `[Paths]` section, set the `TemplateDirectory` parameter to the path to the directory that contains your XSL templates.
5. Save and close the configuration file.
6. Restart Media Server for your changes to take effect.

To apply a template to action output

- Add the following parameters to the action.

<code>Template</code>	The name of the template to use to transform the action output. Exclude the folder path and file extension.
<code>ForceTemplateRefresh</code>	(Optional) If you modified the template after the server started, set this parameter to <code>True</code> to force the ACI server to reload the template from disk rather than from the cache.

For example:

```
action=QueueInfo&QueueName=Process
      &QueueAction=GetStatus
```

```
&Token=MTkyLjE2OC45NS4yNDox  
&Template=Form
```

In this example, Media Server applies the XSL template `Form.tmp1` to the response from a `QueueInfo` action.

Note: If the action returns an error response, Media Server does not apply the XSL template.

Part II: Process Video

This section describes how to configure the operations that Media Server can perform.

- ["Introduction"](#)
- ["Ingest Video"](#)
- ["Analyze Video"](#)
- ["Event Stream Processing"](#)
- ["Transform Data" on page 150](#)
- ["Encode Video"](#)
- ["Output Data"](#)

Chapter 6: Introduction

This section describes how to create a configuration and use that configuration to start processing video.

- [Configuration Overview](#) 72
- [Create a Configuration](#) 76
- [Example Configuration](#) 80
- [Example Configuration - Advanced](#) 81
- [Start Processing Video](#) 84
- [Verify Media Server is Processing Video](#) 85
- [Stop Processing Video](#) 85

Configuration Overview

Before you begin processing a file or stream, you must create a configuration that instructs Media Server how to process the video.

You can configure Media Server by modifying the Media Server configuration file. However, you might want to process videos from different sources, or run different analysis tasks for different video clips. Therefore HPE recommends that you do not configure ingestion, analysis, encoding, event stream processing, or output in the Media Server configuration file. Instead, pass these settings to Media Server in the `process` action, when you start processing.

To pass a configuration to Media Server through the `process` action, you can:

- Create a configuration and save the file in the directory specified by the `ConfigDirectory` parameter, in the `[Paths]` section of the Media Server configuration file. When you run the `process` action to start processing, use the `ConfigName` action parameter to specify the name of the configuration file to use.
- Base-64 encode the configuration and send it with the `process` action when you start processing. Use the `Config` action parameter to include the base-64 encoded data.

A configuration that you specify in the `process` action overrides the `[Ingest]`, `[Analysis]`, `[Transform]`, `[Encoding]`, `[EventProcessing]`, and `[Output]` sections of the Media Server configuration file.

Tasks

To create a configuration for processing video, you must define *tasks*, which are individual operations that Media Server can perform.

- **Ingest tasks** bring video into Media Server so that it can be processed. Ingestion splits the data contained in a video file or stream into images (video frames), audio, and metadata. A configuration must contain exactly one ingest task.

- **Encoding tasks** make a copy of ingested video, or encode it into different formats. A configuration can contain any number of encoding tasks.
- **Analysis tasks** run analysis on ingested video, and produce metadata that describes the content of the video. A configuration can contain any number of analysis tasks.
- **Event stream processing (ESP) tasks** introduce additional custom logic into video analysis. For example, you can filter or deduplicate the information produced during analysis. A configuration can contain any number of ESP tasks.
- **Transform tasks** transform the data produced by other tasks, but do not modify its schema. A configuration can contain any number of transformation tasks.
- **Output tasks** send the data produced by Media Server to other systems. A configuration can contain any number of output tasks.

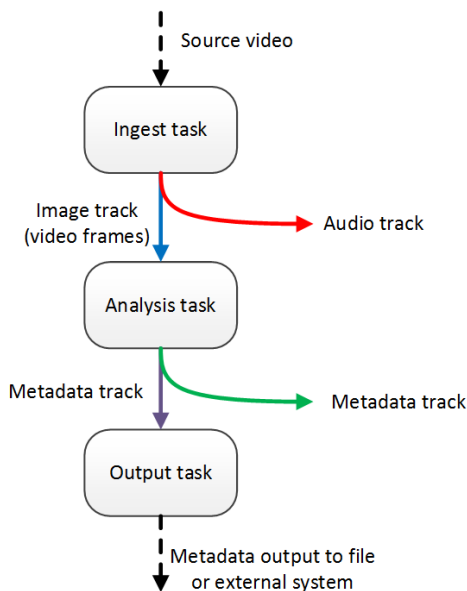
In a complete configuration, tasks that you define are connected. The output of an ingest task becomes the input for an analysis or encoding task. The output of an analysis task might become the input for an encoding, ESP, or output task.

Tracks

Information is passed between tasks in the form of *tracks*.

Tasks usually produce multiple tracks. For example, an ingest task can produce an image track that contains frames from the ingested video, and an audio track that contains the audio packages.

In the following example, an ingest task takes the source video and produces an image track and an audio track. The image track is used as the input for an analysis task. This could be object detection, OCR, or another type of analysis that operates on images. In this example the audio track is not used. The analysis task produces some metadata tracks which contain information about the video content. One of the tracks is used as the input for an output task.



When you configure a task, you might need to specify the track(s) to use as the input for the task:

Task type	Default input tracks
Ingest	Ingest tasks do not accept input tracks. Specify the video source in the process action, when you start processing video.
Encoding	Encoding tasks automatically use the first image and audio tracks produced by your ingest task, so you only need to specify the input for an encoding task if you want to encode different data, for example: <ul style="list-style-type: none"> • Some video sources contain more than one audio stream, to supply audio in multiple languages. You might want to encode an audio stream other than the default. • You might want to encode data produced by another task. For example, you might want to encode the keyframes identified during keyframe analysis, instead of all ingested frames.
Analysis	Most analysis tasks automatically analyze the first image or audio track produced by your ingest task, so in most cases you do not need to specify an input track. However, some analysis operations require additional data. For example, face recognition requires the metadata produced by face detection, so when you configure face recognition you must specify an input track.
ESP	You must always specify the input track(s) to use.
Transform	You must always specify the input track to use.
Output	Output tasks automatically use the default output tracks produced by your analysis tasks, but you can specify the input track(s) so that the output tasks use different tracks.

Records

Tracks contain *records*. A record is a collection of metadata that contains information about the video content. For example, the LibAv ingest engine produces a record for each ingested frame.

A record can represent data extracted from a single frame, or span multiple frames. Each record that is output from Media Server contains timestamps that describe the start time, duration, and end time of the record.

A record also contains metadata that describes the video content.

The following record is an example record produced by the OCR analysis engine. It has a unique ID to identify the event. Multiple records can relate to the same event (see ["Analysis Task Output Tracks" on the next page](#)) so you can use the ID to identify related records. The metadata also describes the text that was extracted from the video, the position of the text in the video frame, and the confidence score for the OCR process.

```
<record>
  <timestamp>
    <startTime iso8601="2015-09-22T14:30:35.531005Z">1442932235531005</startTime>
    <duration iso8601="PT00H01M16.820000S">76820000</duration>
    <peakTime iso8601="2015-09-22T14:30:35.531005Z">1442932235531005</peakTime>
```

```
<endTime iso8601="2015-09-22T14:31:52.351005Z">1442932312351005</endTime>
</timestamp>
<trackname>OCR.Result</trackname>
<OCRResult>
  <id>bcba9983-8dee-450e-9d67-3234e1a0c17a</id>
  <text>FOCUS ON ECONOMY</text>
  <region>
    <left>114</left>
    <top>476</top>
    <width>194</width>
    <height>19</height>
  </region>
  <confidence>100</confidence>
</OCRResult>
</record>
```

Analysis Task Output Tracks

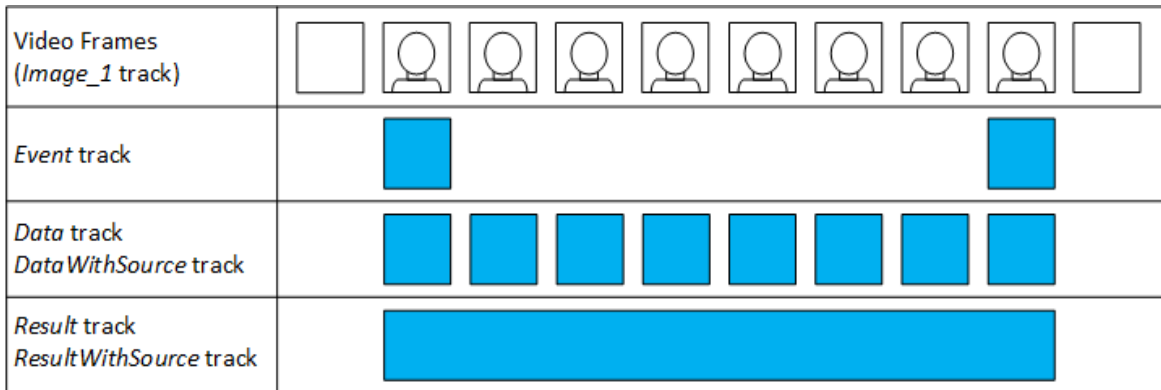
The events that occur in video usually span multiple frames. For example, a person, object, or logo might appear on screen and remain there for seconds or minutes. Media Server analysis tasks run analysis on single frames, but also follow events across multiple frames.

Analysis tasks can produce the following tracks:

- *Data* tracks contain a record for every frame of an event. For example, the data track from a face detection task contains one or more records for every frame in which a face is detected (Media Server creates multiple records when a frame contains more than one face). If a person remains in the scene for several seconds, this track could contain hundreds of records that represent the same face.
- *DataWithSource* tracks are similar to data tracks because they contain one or more records for every frame of an event. However, each record also includes the image analyzed by the task to produce the record.
- *Result* tracks contain a single record for each event. The result track summarizes the results of analysis for each event. For example, the result track from a face detection task contains a single record for each detected face. Each record has a start time, a duration, and an end time, and can span many frames. Result tracks are, by default, used as input tracks for output tasks that you configure.
- *ResultWithSource* tracks are similar to result tracks because they contain a single record for each event. However, each record also includes the source image that produced the best result. For example, when you run face detection, the frame with the highest confidence score is added to the record.
- *Event* tracks contain records that describe the beginning or end of events in the video. For example, the event track for a face detection task contains a record when a face appears in the video, and another record when the face disappears from the scene. Event tracks are, by default, used as input tracks for output tasks that you configure.

The following diagram shows how Media Server creates records (represented by blue squares) when a person appears in a video.

- The Face Detection task creates records in the Event track when the person appears and disappears from the scene.
- The task creates one record in the Data and DataWithSource tracks for each analyzed frame. If there are multiple people in the scene, Media Server creates multiple records in the Data and DataWithSource tracks for each frame.
- The task creates a single record in the Result and ResultWithSource tracks for each event (in this case a detected face). This record spans the event and summarizes the analysis results.



For information about the tracks that are produced by Media Server tasks, and the data contained in each track, refer to the *Media Server Reference*.

Create a Configuration

To create a configuration, you need to consider:

- the video source, because this determines how to ingest the video.
- the analysis tasks that you want to run, for example face detection or speech-to-text.
- the way in which data will flow between tasks. For example, you cannot run face recognition without first running face detection. You must configure the input of the face recognition task to be the output from the face detection task.
- how to output data.

A Media Server configuration can include [Ingest], [Analysis], [Transform], [Encoding], [EventProcessing], and [Output] sections. Each section contains a list of tasks that you want to run:

```
[Ingest]
IngestEngine=LibAvIngest

[Analysis]
AnalysisEngine0=FaceDetect
AnalysisEngine1=FaceRecognize
...

[Transform]
```

```
TransformEngine0=ResizeFaceImagesForEncoding  
...
```

```
[EventProcessing]  
EventProcessingEngine0=Deduplicate  
...
```

```
[Encoding]  
EncodingEngine0=Image  
...
```

```
[Output]  
OutputEngine0=IdolServer  
...
```

These tasks are configured in other sections of the configuration file. The previous example defines an ingest task named `LibAvIngest` which you would configure in a section named `[LibAvIngest]`.

Every task configuration section contains the `Type` parameter to specify the engine to use to complete the task, and contains any settings that the engine requires to complete the task. For example:

```
[LibAvIngest]  
Type=libav
```

Ingestion

Ingestion decodes video so that Media Server can process it. An ingest engine extracts the video and audio from its container and decodes the streams so that they can be analyzed and transcoded.

Your configuration must include exactly one ingest task. For example:

```
[Ingest]  
IngestEngine0=LibAvIngest  
  
[LibAvIngest]  
Type=libav
```

This example specifies a single ingest task named `LibAvIngest`. The engine used to complete the task is the `libav` ingest engine. Notice that no source file or stream is specified in the configuration. You provide the path of a file or the URL of a stream to Media Server in the `Process` action when you start processing.

Ingest engines produce one or more image and audio tracks:

- Each image track is named `Image_n`, where `n` is a unique number. Tracks are numbered from 1.
- Each audio track is named `Audio_Lang_n`, where `Lang` is the language and `n` is a unique number. Tracks are numbered from 1. If the language is unknown, each track is named `Audio__n` (note the double underscore), where `n` is a unique number.

For example, the engine might ingest a TV broadcast and produce one image track, `Image_1`, and three audio tracks: `Audio_French_1`, `Audio_English_2`, and `Audio_German_3`.

Analysis

The [Analysis] section of the configuration lists the analysis tasks that you want to run. A configuration can contain any number of analysis tasks. For example, you can run face detection and object detection at the same time.

The following example includes a single analysis task named OCRtext. The task uses the OCR analysis engine:

```
[Analysis]
AnalysisEngine0=OCRtext
```

```
[OCRtext]
Type=ocr
Input=Image_1
```

An analysis engine accepts input of a particular type. For example, the OCR engine requires an image track, and the SpeakerID engine requires an audio track. The analysis engine only processes records from the track specified by the Input configuration parameter. In this example, the OCR engine processes the Image_1 track produced by the ingest engine. If you do not specify an input track, the engine uses the first track of the required type produced by the ingest engine. The engine would automatically attempt to process the Image_1 track.

To find out what type of input is required by an analysis engine, send action=getstatus to Media Server. The response includes information about the input and output tracks for the analysis engines.

The input could be an output track from another engine, provided the track type is compatible. In the following example, the input of the face recognition task is the ResultWithSource output track of the face detection task:

```
[Analysis]
AnalysisEngine0=FaceDetect
AnalysisEngine1=FaceRecognize

[FaceDetect]
Type=FaceDetect
...

[FaceRecognize]
Type=FaceRecognize
Input=FaceDetect.ResultWithSource
...
```

To be used as the input for a task, a track must provide at least the required record types. For information about the output tracks produced by Media Server engines, use action=getstatus or refer to the *Media Server Reference*.

Transform

The [Transform] section of the configuration lists the transformation tasks that you want to run. A configuration can contain any number of transformation tasks.

A transformation task requires a single input, transforms the data in some way, and produces a single output with the same schema as the input. For example, you can use a transformation task to resize keyframes extracted by keyframe analysis, before sending them to the image encoder and writing them to disk.

The following example includes a single transformation task named `ScaleKeyframes`. The task uses the `Scale` transformation engine:

```
[Transform]
TransformEngine0=ScaleKeyframes
```

```
[ScaleKeyframes]
Type=Scale
Input=Keyframe.ResultWithSource
ImageSize=300,0
```

All transformation engines produce a single output track with the name `TaskName.Output`, where `TaskName` is the name of the transformation task.

Encoding

The `[Encoding]` section lists the names of tasks that produce encoded copies of the video during the session. A configuration can contain any number of encoding tasks. For example:

```
[Encoding]
EncodingEngine0=MyRollingBuffer

[MyRollingBuffer]
Type=rollingbuffer
// rolling buffer configuration
```

This example specifies a single encoding task named `MyRollingBuffer`.

An encoding engine accepts image and/or audio tracks produced by an ingest or analysis engine. For example, you can:

- make a copy of the ingested video.
- make a copy of the ingested video at a different resolution or bitrate to the source.
- encode the output of an analysis engine - for example use the Image Encoder to write the keyframes identified by keyframe analysis to disk.

Output

The `[Output]` section lists the tasks that output information produced by Media Server. Usually a configuration contains at least one output task. For example:

```
[Output]
OutputEngine0=IDOL
OutputEngine1=Vertica

[IDOL]
Type=idol
```

```
// idol engine options  
  
[Vertica]  
...
```

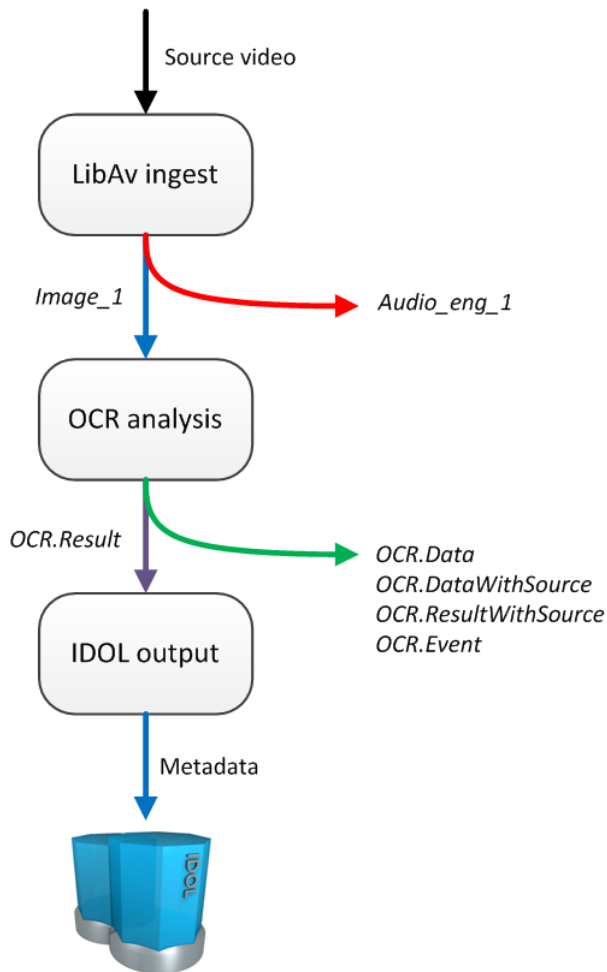
This example specifies two output tasks, named IDOL and Vertica.

All output tasks can include the configuration parameter `Input`. This specifies a comma-separated list of tracks that you want to output to an external system. If you do not specify a list of tracks, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to the *Media Server Reference*.

Output engines do not produce records, and therefore do not have output tracks.

Example Configuration

The following diagram shows a simple configuration for performing OCR on a video stream and sending the results to IDOL Server. The setup includes an ingest engine, an analysis engine, and an output engine. These are usually the minimum required for any configuration.



1. Media Server receives a Process action, which starts a new session. The Libav ingest engine receives video from the source specified in the action. The Libav engine produces an image track and audio track.
2. The image track is used as the input for the OCR analysis engine.
3. The audio track cannot be used by the OCR analysis engine so is discarded.
4. The OCR engine produces several output tracks. The Result track contains the OCR results and is used as the input for the IDOL output task.
5. The other tracks produced by the OCR engine are not required and so are discarded.
6. The IDOL output engine indexes the data produced by Media Server into IDOL.

A Media Server configuration that matches this process is shown below.

- The OCRtext task includes the parameter `Input=Image_1`, which specifies that the input for the OCR task is the `Image_1` track from the ingest engine. Setting the `Input` parameter is optional, because by default Media Server processes the first available track of the correct type.
- The IDOLOutput task includes the parameter `Input=OCRtext.Result`. This specifies that the input for the IDOL output task is the `Result` track from the OCRtext task.

```
[Ingest]
IngestEngine=LibAvIngest
```

```
[LibAvIngest]
Type=libav
```

```
[Analysis]
AnalysisEngine0=OCRtext
```

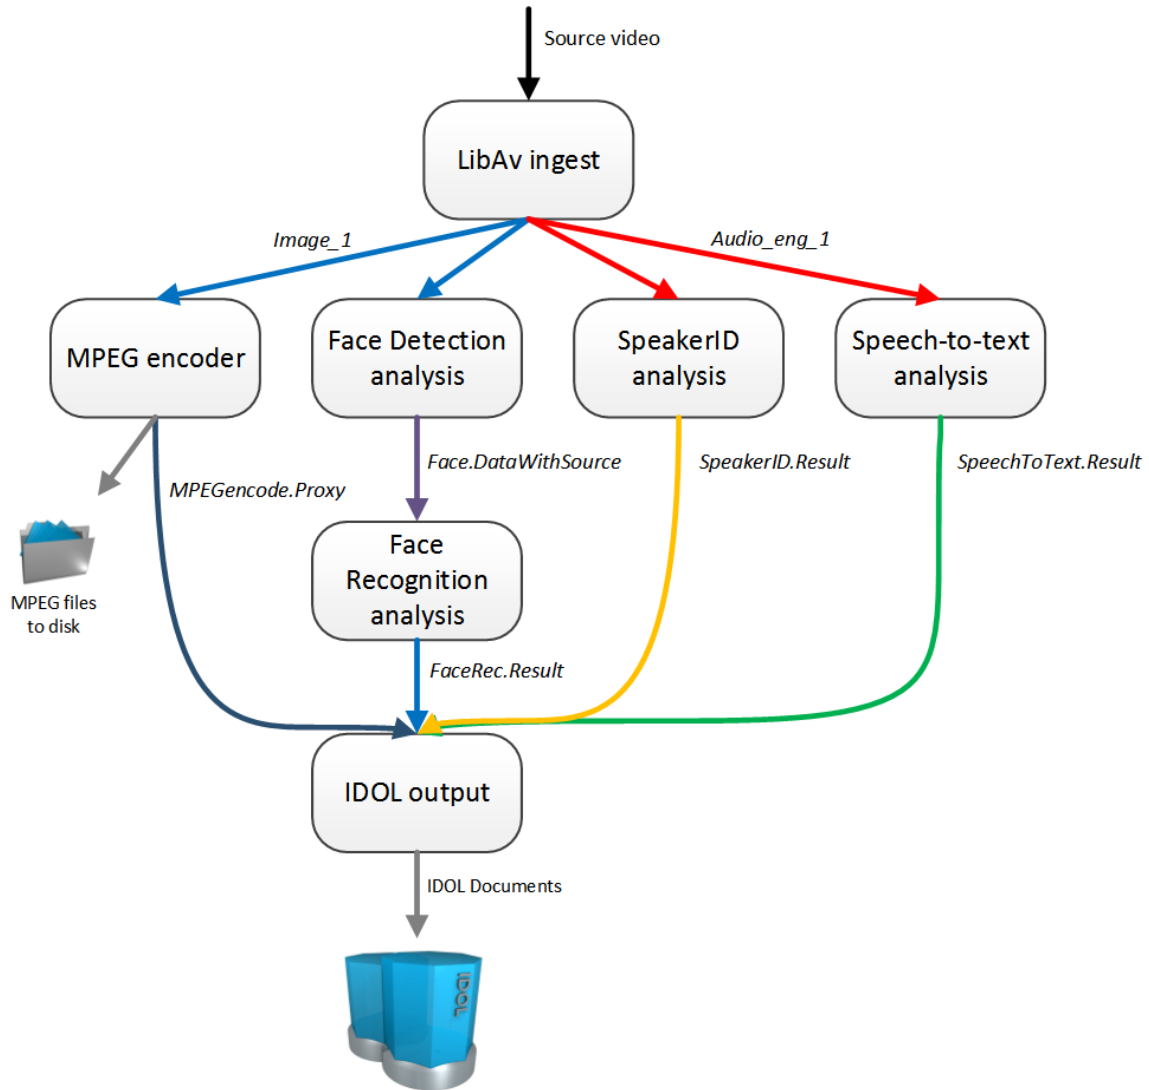
```
[OCRtext]
Type=ocr
Input=Image_1
```

```
[Output]
OutputEngine0=IDOLOutput
```

```
[IDOLOutput]
Input=OCRtext.Result
Type=idol
Mode=time
OutputInterval=60
XslTemplate=./toIDX.xsl
```

Example Configuration - Advanced

The following diagram shows a more complex configuration, which includes an encoding task and multiple analysis tasks.



1. Media Server receives a Process action, which starts a new session. The Libav ingest engine receives video from the source specified in the action. The Libav engine produces an image track and an audio track.
2. The image track is used as the input for the MPEG encoder and Face Detection analysis tasks.
3. The audio track is used as the input for the Speaker ID and Speech-to-text analysis tasks.
4. The MPEG encoder writes a copy of the video to disk as a set of MPEG files. It also produces proxy information that is available to the output engine.
5. The face detection engine produces a track that contains information about detected faces. This is used as the input for the face recognition analysis engine.
6. The output tracks from the analysis engines are used as the input for the IDOL Server output engine. The analysis engines all produce multiple tracks, some of which are not used.
7. The IDOL output engine transforms the records produced by the MPEG encoder and analysis engines into documents and indexes the information into IDOL Server.

A Media Server configuration that matches this process is shown below.

```
[Ingest]
IngestEngine=LibAvIngest

[LibAvIngest]
Type=libav

[Encoding]
EncodingEngine0=MPEGencode

[MPEGencode]
Type=mpeg
OutputPath=\\server\folder\file.mpg
URLBase=https://www.myserver.com/folder/
Segment=TRUE

[Analysis]
AnalysisEngine0=Face
AnalysisEngine1=FaceRec
AnalysisEngine2=SpeakerID
AnalysisEngine3=SpeechToText

[Face]
Type=FaceDetect

[FaceRec]
Type=FaceRecognize
Input=Face.DataWithSource

[SpeakerID]
Type=SpeakerID
AstPath=speakers.ast
SpeakerIDServers=10.0.0.2:13000

[SpeechToText]
Type=SpeechToText
Language=ENUK
Mode=relative
ModeValue=0.90
SpeechToTextServers=10.0.0.2:13000

[Output]
OutputEngine0=IDOLOutput

[IDOLOutput]
Type=IDOL
```

```
Input=FaceRec.Result,SpeakerID.Result,SpeechToText.Result,MPEGencode.Proxy
Mode=time
OutputInterval=60
XslTemplate=./toIDX.xsl
SavePostXML=true
XMLOutputPath=./output/idol/
```

Start Processing Video

To start processing video, send the `process` action to Media Server.

Media Server adds your request to the `process` action queue. When the request reaches the top of the queue and a thread is available, Media Server starts to ingest the video and perform the configured analysis, encoding, event stream processing, and output tasks.

To start processing video

- Send the `process` action to Media Server. Set the following parameters:

Source	The video source to process. <ul style="list-style-type: none">• To ingest video from a file or stream, specify either the path to a file or a URL.• To ingest video from a DirectShow source, such as a video capture card, specify the name of the video device.• To ingest video from Milestone XProtect, specify the camera GUID.
Persist	Specifies whether the action restarts in the event that processing stops for any reason. For example, if you are processing video from an RTSP camera which becomes unreachable and <code>persist=true</code> , Media Server will wait for the stream to become available again. Persistent actions only stop when you stop them using the <code>QueueInfo</code> action with <code>QueueAction=stop</code> , or when Media Server finishes processing the video.
Config	(Optional) A base-64 encoded configuration to use to override the default configuration file.
ConfigName	(Optional) The name of a configuration file to use to override the Media Server configuration file. The file must be stored in the directory specified by the <code>ConfigDirectory</code> parameter in the <code>[Paths]</code> section of the configuration file.

For example,

```
http://localhost:14000/action=Process&Source=.\video\broadcast.mpeg
&ConfigName=broadcast
```

This action is asynchronous, so Media Server returns a token for the request. You can use the `QueueInfo` action, with `QueueName=Process` to retrieve more information about the request.

Verify Media Server is Processing Video

You can run the action `GetLatestRecord` to verify that Media Server is processing video. This action returns the latest records that have been added to the tracks you specify in the action.

For example, send the following action:

```
http://localhost:14000/action=GetLatestRecord
                                &Token=...
                                &Tracks=Keyframe.Result,OCR.Result
```

- where the `token` parameter specifies the asynchronous action token returned by the `process` action,
- and `tracks` specifies the tracks to retrieve the latest records from. You can set this parameter to an asterisk (*) to retrieve the latest records from all tracks.

Stop Processing Video

When you process a file, the `process` action finishes when Media Server reaches the end of the file. However, you might be processing a video stream that does not end. In this case, you might want to stop processing.

To stop processing video

- Use the `QueueInfo` action, with the following parameters:

<code>QueueName</code>	The name of the queue. Processing is initiated using the <code>process</code> action, so set <code>QueueAction=process</code> .
<code>QueueAction</code>	The action to perform on the queue. To stop processing, set <code>QueueAction=stop</code> .
<code>Token</code>	(Optional) The token for the request that you want to stop. If you don't specify a token, Media Server stops the task that is currently processing.

For example,

```
http://localhost:14000/action=queueinfo
                                &queueName=process
                                &queueAction=stop
                                &token=MTAuMi4xMDQuODI6MTQwMDA6UFJlPQ0VTUzoxMTgzMzYzMjQz
```

Media Server returns a response stating whether the request was successfully completed.

Chapter 7: Ingest Video

Media Server ingests video through an ingest engine. An ingest engine receives video input and produces demuxed and decoded streams ready for processing by other engines.

This section describes how to configure Media Server to ingest video.

- [Supported Codecs and Formats](#)86
- [Choose the Rate of Ingestion](#)87
- [Ingest a Video File or Stream](#) 88
- [Ingest Video from a DirectShow Device](#) 89
- [Ingest Video from Milestone XProtect](#)89

Supported Codecs and Formats

This page lists the audio and video codecs, and the file formats, supported by Media Server. To ingest files that use these codecs and formats, use a LibAv ingest task (see "[Ingest a Video File or Stream](#)" on page 88).

Other codecs and file formats might work, but they are not supported.

Video Codecs

- MPEG1
- MPEG2
- MPEG4 part 2
- MPEG4 part 10 (Advanced Video Coding) (h264)
- MPEGH part 2 (High Efficiency Video Coding) (h265)
- Windows media 7
- Windows media 8

Audio Codecs

- MPEG audio layer 1
- MPEG audio layer 2
- MPEG audio layer 3
- MPEG 4 audio (Advanced Audio Coding) (AAC)
- PCM (wav)
- Windows media audio 1
- Windows media audio 2
- AC3 (ATSC A/52)

File Formats

- MPEG packet stream (for example .mpg)
- MPEG2 transport stream (for example .ts)
- MPEG4 (for example .mp4)
- WAVE (Waveform Audio) (.wav)
- asf (Windows media) (.asf, .wmv)
- raw aac (.aac)
- raw ac3 (.ac3)

Choose the Rate of Ingestion

When Media Server ingests video the decoded frames are pushed to the analysis engines. Some analysis operations are resource-intensive and cannot process frames as quickly as they can be ingested. This means that you have to decide whether to prioritize processing speed by dropping some of the frames, or pause ingestion until the analysis engine has processed the preceding frames.

The rate of ingestion has no effect on the speed at which video is analyzed. The speed of analysis is determined by the resources available to Media Server.

The rate at which Media Server ingests video is controlled by the `IngestRate` configuration parameter. You can set this parameter to one of the following values:

- `1` - This is the default value. Media Server ingests video at normal (playback) speed. Assuming there are no external limitations such as network bandwidth, Media Server ingests one minute of video every minute. If an analysis task cannot process frames at this speed, frames are dropped.
- `0` - The rate at which video is ingested is determined by the slowest analysis task that you have configured. Media Server ingests video without dropping frames. Analysis tasks such as keyframe analysis are very fast, and can analyze thousands of frames per second. Others, such as optical character recognition, are slower. At this rate, the amount of time required to ingest a video file depends on the analysis tasks that you have configured and the resources available to Media Server.

There are various scenarios when you might want to consider the rate of ingestion:

- To process a live stream, you must use the default value. You cannot change the rate at which a stream is ingested.
- To prioritize speed and prevent resource-intensive analysis operations resulting in long ingest times, set `IngestRate=1`.
- If it is critical that all frames are analyzed, set `IngestRate=0` to that ensure no frames are dropped. This is the only mode that can produce deterministic results (results that are the same every time a video is processed). The other modes are non-deterministic because different frames might be dropped each time the video is analyzed.
- If you have configured only fast analysis operations and want to ingest files as fast as possible, set `IngestRate=0`. Be aware that ingestion will be slower than normal speed if you have configured resource-intensive analysis operations, because in this mode Media Server is not permitted to drop frames.

Ingest a Video File or Stream

Media Server can ingest video files and streams.

To configure a task to ingest a video file or stream

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Ingest]` section, set the `IngestEngine` parameter to the name of a section in the configuration file that will contain the ingestion settings. For example:

```
[Ingest]
IngestEngine=VideoIn
```

3. Create a new section in the configuration file to contain the ingestion settings, and set the following parameters:

<code>Type</code>	The ingest engine to use. Set this parameter to <code>libav</code> .
<code>IngestTime</code>	(Optional) The ingest engine produces output tracks that contain timestamped records. You can use this parameter to control the timestamp values, which is useful if you are ingesting video that was originally broadcast at an earlier time. Specify the time in epoch seconds.
<code>MaximumDuration</code>	(Optional) The maximum amount of video and audio to ingest, before stopping ingestion. If you do not set this parameter, there is no limit on the amount of video and audio that is ingested. You can set this parameter to ingest part of a video file, or a fixed amount of a continuous stream. If the maximum duration is reached, Media Server stops ingestion but does finish processing any frames that have already been ingested.

For example:

```
[VideoIn]
Type=libav
```

Tip: There is no configuration parameter to specify the source. You must specify the path of the source file or the IP address of the stream that you want to ingest as the value of the `Source` parameter in the `Process` action, when you start processing.

For more information about the parameters you can set to configure ingestion, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Ingest Video from a DirectShow Device

Media Server can ingest video directly from a DirectShow device, such as a video capture card or camera.

Note: Ingestion from a DirectShow device is supported only on Windows.

To ingest video from a DirectShow device

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Ingest]` section, set the `IngestEngine` parameter to the name of a section in the configuration file that will contain the ingestion settings. For example:

```
[Ingest]
IngestEngine=CaptureCard
```

3. Create a new section in the configuration file to contain the ingestion settings, and set the following parameters:

Type	The ingest engine to use. Set this parameter to <code>libav</code> .
Format	The video format. Set this parameter to <code>dshow</code> .
IngestTime	(Optional) The ingest engine produces output tracks that contain timestamped records. You can use this parameter to control the timestamp values, which is useful if you are ingesting video that was originally broadcast at an earlier time. Specify the time in epoch seconds.

For example:

```
[CaptureCard]
Type=libav
Format=dshow
```

Tip: There is no configuration parameter to specify the source. You must specify the name of the DirectShow device that you want to use as the value of the `Source` parameter in the Process action, when you start processing.

For more information about the parameters you can set to configure ingestion, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Ingest Video from Milestone XProtect

Media Server can ingest live video from Milestone XProtect Enterprise and Milestone XProtect Corporate.

Media Server connects to the Milestone XProtect recording server to obtain video. The recording server requires user authentication, so you must specify a user name and password in the Media Server configuration. HPE recommends that you encrypt passwords before entering them into the configuration file. For information on how to do this, see ["Encrypt Passwords" on page 52](#).

To ingest video from Milestone XProtect

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (mediaserver.cfg).
2. In the [Ingest] section, create an ingest task by setting the IngestEngine parameter.

```
[Ingest]
IngestEngine=XProtect
```

3. Create a configuration section to contain the task settings. Type the task name inside square brackets, and then set the following parameters:

Type	The ingest engine to use. Set this parameter to <code>milestone</code> .
RecorderHost	The host name or IP address of the XProtect recording server.
RecorderPort	The port of the XProtect recording server.

For example:

```
[XProtect]
Type=milestone
RecorderHost=XProtect
RecorderPort=7563
```

4. Configure how the ingest engine should authenticate with the Milestone XProtect system.

Note: To ingest video from Milestone XProtect Corporate, you must use NTLM authentication against the Milestone XProtect authentication server. To ingest video from Milestone XProtect Enterprise, you can use any of the following options.

- To use NTLM authentication against a Milestone authentication server, set the following parameters:

SOAPAuthentication	Set this parameter to <code>true</code> . The engine authenticates against the XProtect authentication server.
NTLMUsername	The NT user name to use to retrieve video.
NTLMPassword	The NT password to use to retrieve video.
AuthenticationHost	The host name or IP address of the XProtect authentication server.
AuthenticationPort	The port of the XProtect authentication server.

- To use Milestone credentials (Basic authentication) against a Milestone authentication server, set the following parameters:

<code>SOAPAuthentication</code>	Set this parameter to <code>true</code> . The engine authenticates against the XProtect authentication server.
<code>BasicUsername</code>	The user name to use to retrieve video.
<code>BasicPassword</code>	The password to use to retrieve video.
<code>AuthenticationHost</code>	The host name or IP address of the XProtect authentication server.
<code>AuthenticationPort</code>	The port of the XProtect authentication server.

- To use Milestone credentials (Basic authentication) against the Milestone recording server, set the following parameters:

<code>SOAPAuthentication</code>	Set this parameter to <code>false</code> . The engine sends credentials to the recording server.
<code>BasicUsername</code>	The user name to use to retrieve video.
<code>BasicPassword</code>	The password to use to retrieve video.

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

You can now start ingesting video. When you use the `process` action, you must set the `source` parameter to a Milestone camera GUID. For information about how to find the GUID, refer to the Milestone documentation. Unless you are using SOAP authentication, you can specify the camera ID instead. For more information about how to start processing, see ["Start Processing Video" on page 84](#).

Chapter 8: Analyze Video

This section describes how to set up analysis tasks to analyze video.

- [Analyze Speech](#) 92
- [Detect QR Codes](#) 95
- [Detect, Recognize, and Analyze Faces](#) 96
- [Extract Keyframes](#) 106
- [Recognize Objects](#) 107
- [Classify Objects](#) 118
- [Perform Optical Character Recognition \(OCR\)](#) 125
- [Perform Color Analysis](#) 126
- [Analyze Vehicles](#) 127
- [Scene Analysis](#) 133

Analyze Speech

Media Server can use an IDOL Speech Server to analyze the speech in video.

- **Speech-to-text** extracts speech and converts it into text.
- **Speaker identification** identifies the people who speak in the video.

When the audio or video source contains narration or dialogue, running speech analysis and indexing the resulting metadata into IDOL Server means that IDOL can:

- search all of the video that you have processed to find clips where a speaker talks about a specific topic.
- provide parametric search to filter the results by speaker.
- categorize video clips.
- cluster together video clips that contain related concepts.

To analyze speech, you must have an IDOL Speech Server. For more information about using IDOL Speech Server, refer to the IDOL Speech Server documentation.

Transcribe Speech

To run speech-to-text

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the [Analysis] section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

[Analysis]

AnalysisEngine0=TranscribeSpeech

3. Create a new section to contain the settings for the task and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>speechtotext</code> .
Input	(Optional) The audio track to analyze. If you do not specify an input track, Media Server processes the first track of the correct type produced by the ingest engine.
SpeechToTextServers	The host name and ACI port of an IDOL Speech Server. Separate the host name and port with a colon (for example, <code>speechserver:15000</code>). You can specify multiple IDOL Speech Servers by using a comma-separated list. Media Server can connect to only one IDOL Speech Server at a time, but you can provide multiple servers for failover. Tip: You can specify a default IDOL Speech Server to use for all speech-to-text tasks by setting the <code>SpeechToTextServers</code> parameter in the [Resources] section of the Media Server configuration file.
Language	The language pack to use. For a list of available language packs, refer to the <i>IDOL Speech Server Administration Guide</i> .
Mode	The mode for speech-to-text transcription. Media Server supports the following modes: <ul style="list-style-type: none">• Fixed• Relative These modes are described in the <i>IDOL Speech Server Reference</i> .
ModeValue	The mode value for speech-to-text transcription. For more information about this parameter, refer to the <i>IDOL Speech Server Reference</i> .
FilterMusic	(Optional) Specifies whether to ignore speech-to-text results for audio segments that Speech Server identifies as music or noise. To filter these results from the output, set this parameter to <code>true</code> .
SampleFrequency	(Optional) The sample frequency of the audio to send to the IDOL Speech Server for analysis, in samples per second (Hz). IDOL Speech Server language packs are dependent on the audio sample rate, and accept audio at either 8000Hz or 16000Hz.

For example:

```
[TranscribeSpeech]
```

```
Type=speechtotext
```

```
SpeechToTextServers=speechserver:13000
```

```
Language=ENUK
```

```
Mode=relative
```

```
ModeValue=0.8  
FilterMusic=TRUE
```

For more information about the parameters that you can use to configure speech-to-text, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Identify Speakers

Note: Before IDOL Speech Server can identify speakers, you must train the Speech Server. Without training, Speech Server can divide the audio into different speakers and identify the gender of each speaker. For information about training IDOL Speech Server, refer to the *IDOL Speech Server Administration Guide*.

To identify speakers in video

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the [Analysis] section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]  
AnalysisEngine0=IDSpeakers
```

3. Create a new section to contain the settings for the task, and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>speakerid</code> .
Input	(Optional) The audio track to process. If you do not specify an input track, Media Server processes the first track of the correct type produced by the ingest engine.
SpeakerIDServers	The host name and ACI port of an IDOL Speech Server. Separate the host name and port with a colon (for example, <code>speechserver:13000</code>). You can specify multiple IDOL Speech Servers by using a comma-separated list. Media Server can connect to only one IDOL Speech Server at a time, but you can provide multiple servers for failover.

Tip: You can specify a default IDOL Speech Server to use for all speaker identification tasks by setting the `SpeakerIdServers` parameter in the [Resources] section of the Media Server configuration file.

AstPath	(Optional) The path to the speaker classifier (.ast) file to use for speaker identification. You must create this file. Specify the path relative to the directory defined by the <code>SpeakerIDDir</code> parameter in the IDOL Speech Server configuration file. If you do not set this parameter Speech Server
---------	--

cannot identify speakers, but can divide the audio into different speakers and detect the gender of each speaker.

SampleFrequency (Optional) The sample frequency of the audio to send to the IDOL Speech Server for analysis, in samples per second (Hz). IDOL Speech Server accepts audio at either 8000Hz or 16000Hz.

For example:

```
[IDspeakers]  
Type=speakerid  
SpeakerIDServers=speechserver:13000  
AstPath=speakers.ast
```

For more information about the parameters that you can use to configure this task, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Detect QR Codes

Media Server can detect and read QR codes in video. The Barcode analysis engine produces a track that contains details of detected QR codes.

The Barcode engine automatically merges duplicate records to produce a single record for each detected appearance of a QR code. For records to be identical, the QR code contents must match and the QR code must appear in a nearly-identical location. For example, if a QR code appears in the same location for five consecutive frames, the engine merges the results for each frame into a single record with a timestamp duration that covers the five frames. If the QR code then moves to a different location on the screen, the engine counts this as a new detection and produces a separate record. If the QR code disappears then reappears, the engine reports the two detections separately.

To detect QR codes in video

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]  
AnalysisEngine0=QRcode
```

3. Create a new configuration section to contain the task settings and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>barcode</code> .
Input	(Optional) The image track to process. If you do not specify an input track, Media Server processes the first track of the correct type that is produced by the ingest engine.

- Region** (Optional) Search for QR codes in a region of the image, instead of the entire image.
- RegionUnit** (Optional) By default, the **Region** parameter specifies the size and position of a region using percentages of the frame dimensions. Set the **RegionUnit** parameter to **pixel** if you want to use pixels instead.

For example:

```
[QRcode]  
Type=barcode
```

For more information about the parameters that you can use to configure this task, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Detect, Recognize, and Analyze Faces

Media Server can detect faces that appear in video. You can then run several types of analysis:

- recognize faces that appear in video by comparing the detected faces to faces stored in your databases.
- obtain demographic information for detected faces, for example an estimated age, gender, and ethnicity.
- obtain information about the facial expression or whether the person is wearing spectacles.

You can run face detection and face analysis operations out-of-the-box. To run face recognition, you must train Media Server using images of people that you want to identify. For information about how to train Media Server, see "[Train Media Server to Recognize Faces](#)" on the next page.

The Face Detection engine requires an image input track. The image track can be produced by an ingest engine such as the LibAv ingest engine. The Face Recognition, Face State, and Face Demographics engines require an input track that contains face detection data, which you must generate by running face detection.

Detect Faces

To detect faces in video, follow these steps.

To detect faces in video

1. Create a new configuration or open an existing configuration to send to Media Server with the **process** action. Alternatively, you can modify the Media Server configuration file (**mediaserver.cfg**).
2. In the **[Analysis]** section, add a new analysis task by setting the **AnalysisEngineN** parameter. You can give the task any name, for example:

```
[Analysis]  
AnalysisEngine0=FaceDetect
```


3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>FaceDetect</code> .
Input	(Optional) The image track to process. If you do not set this parameter, Media Server processes the first track of the correct type.
Region	(Optional) The region that you want to search for faces.
RegionUnit	(Optional) The units used to specify the region. By default, the <code>Region</code> is specified in pixels. Set the <code>RegionUnit</code> parameter to <code>percent</code> if you want to specify percentages of the image dimensions instead.
FaceDirection	(Optional) By default, Media Server only detects faces looking directly at the camera. Set this parameter if you want to detect faces looking in other directions.
ColorAnalysis	(Optional) A Boolean value that specifies whether to perform color analysis on detected faces. This can reduce the number of false detections.
MinSize	(Optional) The minimum expected size of faces in the video (in pixels unless you set the <code>SizeUnit</code> parameter).
SizeUnit	(Optional) By default, the <code>MinSize</code> parameter specifies the minimum face size in pixels. Set the <code>SizeUnit</code> parameter if you want to specify the size as a percentage of the smallest image side.

For example:

```
[FaceDetect]
Type=facedetect
FaceDirection=Any
MinSize=200
```

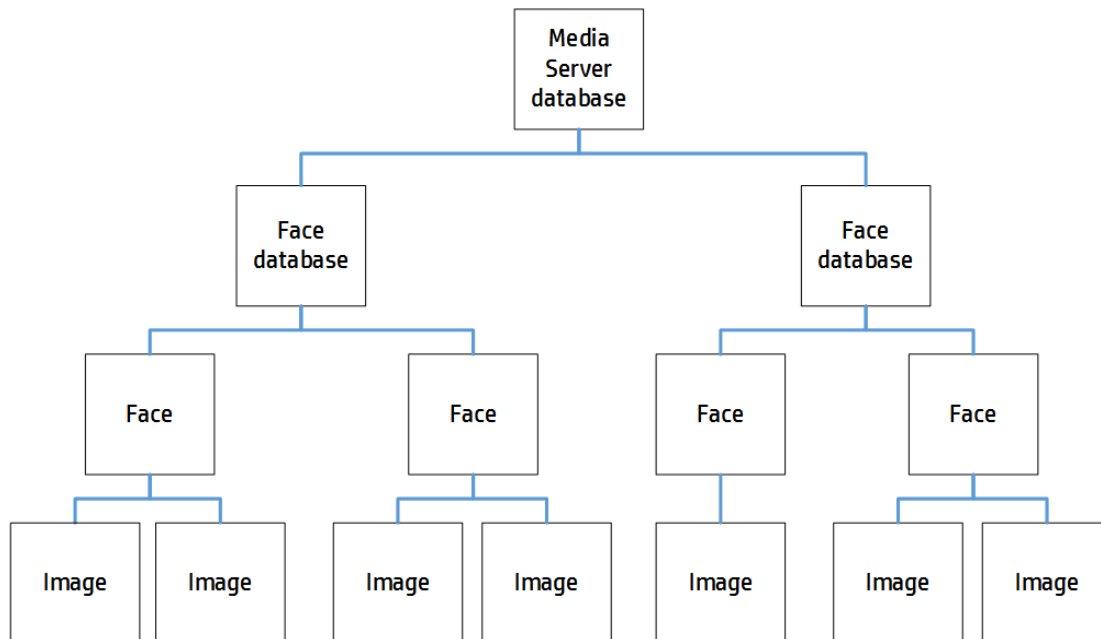
For more information about the parameters that you can use to configure Face Detection, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Train Media Server to Recognize Faces

To run face recognition you must train Media Server by providing images of people that you want to identify. When you run face recognition, Media Server compares the faces it detects in the video to the faces that you added during training.

The following diagram shows how Media Server stores information you provide during training.



The "Media Server database" can represent either the internal Media Server database or a database on an external database server. For information on setting up this database, see ["Set up a Database to Store Training Data"](#) on page 26.

You can organize faces into databases. These face databases are for organizational purposes. For example, you might have a database for politicians, a database for actors and a database for musicians. If you want to recognize only actors, you can then run face recognition using only that database.

A database can contain any number of faces, each representing a single person. Each face has an identifier, which must be unique. You can associate metadata with a face, for example the person's name or other information.

To each face you must add one or more training images. HPE recommends adding multiple training images for best performance. For information about choosing suitable training images, see ["Select Images for Training"](#) below.

Select Images for Training

Add one or more training images to each face that you want to recognize.

HPE recommends using images that are representative of how you will use face recognition. For example, if you intend to use face recognition in a situation that is not controlled for facial expression, add several images to each face showing different facial expressions.

It is better to use fewer high-quality images, rather than many low-quality images, because low-quality images can reduce accuracy.

A good training image for a face:

- must contain only one face.
- must be rotated correctly (the face cannot be upside down).

- contains a face that is wider than approximately 150 pixels; the face should constitute a large proportion of the image.
- contains a front view of the face, so that both eyes are visible.
- is not blurry.
- has even illumination, because dark shadows or bright highlights on the face reduce recognition accuracy.
- has not been converted from another file format to JPEG. Converting images to JPEG introduces compression artifacts. If you have JPEG images, avoid making many rounds of edits. Each time that a new version of the image is saved as a JPEG, the image quality degrades.

Create a Database to Contain Faces

To create a database to contain faces, use the following procedure.

To create a database to contain faces

- Use the `CreateFaceDatabase` action with the `database` parameter.

`database` The name of the new database (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=CreateFaceDatabase  
                              -F database=Faces
```

Add a Face to a Database

To add a face to a database of faces, you can:

- **Perform training in separate steps.** Create a new (empty) face, then add training images, and then train Media Server, using separate actions. You can also add metadata to the face but this is an optional step. You might want to perform training in this way if you are building a front end application that responds to user input.
- **Perform training in a single action.** You might use this approach when writing a script to train Media Server from a collection of images and metadata.

Add a Face to a Database (Using Separate Steps)

This section describes how to create a new (empty) face, then add training images, and then train Media Server, using separate actions. You can also add metadata to the face but this is an optional step. You might want to perform training in this way if you are building a front end application that responds to user input.

Alternatively, you can train Media Server to recognize a face by sending a single action. To do this, see ["Add a Face to a Database \(Using a Single Action\)" on page 101](#).

To add a face to a database (using separate steps)

1. Add a face using the `NewFace` action. Set the following parameters:

<code>database</code>	The name of the database to add the face to. The database must already exist.
<code>identifier</code>	(Optional) A unique identifier for the face (maximum 254 bytes). If you do not set this parameter, Media Server generates an identifier automatically. You can use the person's name as the identifier, but you might have several people with the same name and all identifiers within a database must be unique. Alternatively, allow Media Server to generate an identifier and add the person's name to a metadata field (see step 3, below).

For example:

```
curl http://localhost:14000 -F action=NewFace  
                             -F database=Faces
```

Media Server adds the face and returns the identifier.

2. Add one or more training images to the face using the `AddFaceImages` action. Set the following parameters:

<code>database</code>	The name of the database that contains the face.
<code>identifier</code>	The identifier for the face, returned by the <code>NewFace</code> action.
<code>imagedata</code>	The training images to add. Files must be uploaded as multipart/form-data. For more information about sending data to Media Server, see "Send Data by Using a POST Method" on page 61 .
<code>imagelabels</code>	A comma-separated list of labels. One label is associated with each image. (maximum 254 bytes for each label). The number of labels must match the number of images. If you do not set this parameter, Media Server generates labels automatically.

For example:

```
curl http://localhost:14000 -F action=AddFaceImages  
                             -F database=Faces  
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62  
                             -F imagedata=@face1_smile.jpg,face1_neutral.jpg  
                             -F imagelabels=image1,image2
```

3. (Optional) Add metadata to the face using the `AddFaceMetadata` action. You can add any number of key-value pairs. Set the following parameters:

<code>database</code>	The name of the database that contains the face.
<code>identifier</code>	The identifier for the face, returned by the <code>NewFace</code> action.
<code>key</code>	The key to add (maximum 254 bytes).
<code>value</code>	The value to add (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=AddFaceMetadata
                             -F database=Faces
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62
                             -F key=firstname
                             -F value=John
```

4. Complete the training for the face using the `BuildFace` action. Set the following parameters:

<code>database</code>	The name of the database that contains the face.
<code>identifier</code>	The identifier for the face, returned by the <code>NewFace</code> action.

For example:

```
curl http://localhost:14000 -F action=BuildFace
                             -F database=Faces
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62
```

Add a Face to a Database (Using a Single Action)

You can train Media Server to recognize a face by sending a single action (`TrainFace`).

Running this action is equivalent to running the following actions in the following order:

- `NewFace`
- `AddFaceImages`
- `AddFaceMetadata` (optional)
- `BuildFace`

The `TrainFace` action is atomic, so that any interruption to the server does not leave the database in an inconsistent state.

Alternatively, you can train Media Server by sending these actions individually. You might want to do this if you are building a front end application that trains Media Server in response to user input. For more information about how to do this, see ["Add a Face to a Database \(Using Separate Steps\)" on page 99](#).

To add a face to a database (using a single action)

- Add a face using the `TrainFace` action. Set the following parameters:

<code>database</code>	The name of the database to add the face to. The database must already exist.
<code>identifier</code>	(Optional) A unique identifier for the face (maximum 254 bytes). If you do not set this parameter, Media Server generates an identifier automatically.
<code>imagedata</code>	The training images to add. Files must be uploaded as multipart/form-data. For more information about sending data to Media Server, see "Send Data by Using a POST Method" on page 61 .
<code>imagelabels</code>	(Optional) A comma-separated list of labels. One label is associated with each image (maximum 254 bytes for each label). The number of labels must match the number of images. If you do not set this parameter, Media Server generates labels automatically.
<code>metadata</code>	(Optional) A comma-separated list of metadata key-value pairs to add to the face. Separate keys from values using a colon (:). To include a comma or colon in a key name or value, you must enclose the key name or value in quotation marks (") and escape any quotation marks that occur within the string with a backslash (\).

For example:

```
curl http://localhost:14000 -F action=TrainFace
                             -F database=Faces
                             -F imagedata=@face1_smile.jpg,face1_neutral.jpg
                             -F imagelabels=image1,image2
                             -F metadata=lastname:Smith,fullname:"John Smith, Jr"
```

Media Server adds the face to the database and returns the identifier.

List the Faces in a Database

To list the faces that you have added to a database, and check whether training was successful, use the following procedure.

To list the faces in a database

1. (Optional) First list the databases that have been created to store faces. Use the action `ListFaceDatabases`:

```
http://localhost:14000/action=ListFaceDatabases
```

Media Server returns a list of databases that you have created to store faces.

2. List the faces that exist in one of the databases. Use the action `ListFaces`, for example:

```
http://localhost:14000/action=ListFaces&database=faces
                                &metadata=true
                                &imagestatus=true
```

Media Server returns a list of faces in the specified database, and the number of training images associated with each face.

If you set the action parameter `metadata` to `true`, Media Server returns the metadata you have added to a face.

If you set the action parameter `imagestatus` to `true`, Media Server returns the status of each training image associated with each face.

- A status of `trained` indicates that training was successful.
- If images have a status of `untrained`, run training for the face using the action `BuildFace`, or run training for all faces that have incomplete training using the action `BuildAllFaces`.
- If images have a status of `failed`, Media Server could not use the image for training. For example, if Media Server does not detect a face in an image, it cannot be used as a training image. Remove the failed images using the action `RemoveFaceImages`.

Update or Remove Faces and Databases

To update or remove a face use the following actions:

- To complete training for all faces that exist in the Media Server database but have incomplete training, use the action `BuildAllFaces`. To confirm that training was successful, use the action `ListFaces` (see "[List the Faces in a Database](#)" on the previous page).
- To add additional training images to a face, use the action `AddFaceImages`. Media Server does not use the new images for face recognition until you run training using the action `BuildFace` or `BuildAllFaces`. To confirm that training was successful, use the action `ListFaces` (see "[List the Faces in a Database](#)" on the previous page).
- To remove a training image from a face, for example if Media Server cannot detect a face in the image, use the action `RemoveFaceImages`.
- To change the label of an image that you added to a face, use the action `RelabelFaceImage`.
- To move an image of a face from one face to another, for example if you add an image to the wrong face, use the action `MoveFaceImage`.
- To add, remove, or update custom metadata for a face, use the actions `AddFaceMetadata`, `RemoveFaceMetadata`, and `UpdateFaceMetadata`.
- To change the identifier of a face you added to a face database, use the action `RenameFace`.
- To remove a face from a database and discard all associated images and metadata, use the action `RemoveFace`.

To update or remove face databases, use the following actions:

- To rename a face database, use the action `RenameFaceDatabase`.
- To delete a face database and all of the information that it contains, use the action `RemoveFaceDatabase`.

For more information about the actions that you can use to train Media Server, refer to the *Media Server Reference*.

IDOL Admin

You can train Media Server using IDOL Admin. To open the IDOL Admin interface, send the following action to Media Server's ACI port:

<http://localhost:14000/action=admin>

For more information about using IDOL Admin refer to the *IDOL Admin User Guide*.

Recognize Faces

This section describes how to create an analysis task to recognize faces that appear in video.

To run face recognition, you must first detect faces by setting up a face detection task. For information about how to do this, see ["Detect Faces" on page 96](#).

To recognize faces in video

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the [Analysis] section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=FaceDetect
AnalysisEngine1=FaceRecognize
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>FaceRecognize</code> .
Input	The track that contains detected faces that you want to recognize. Set this parameter to the <code>ResultWithSource</code> output track from your face detection task. For example, if your face detection task is named <code>FaceDetect</code> , set this parameter to <code>FaceDetect.ResultWithSource</code> .
RecognitionThreshold	(Optional) The minimum confidence score required to recognize a face.
MaxRecognitionResults	(Optional) The maximum number of results to return, if the face matches more than one entry in the training database(s).
Database	(Optional) The database to use for recognizing the detected faces. By default, Media Server uses all available data. Database names are case-sensitive.

For example:

```
[FaceRecognize]
Type=FaceRecognize
Input=FaceDetect.ResultWithSource
RecognitionThreshold=60
MaxRecognitionResults=1
```

For a complete list of parameters that you can use to configure a face recognition task, and more information about these parameters, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Obtain Demographic Information

You can use the Face Demographics analysis engine to obtain demographic information for detected faces. This section describes how to configure a face demographics task.

To obtain demographic information for detected faces

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=FaceDetect
AnalysisEngine1=FaceDemographics
```

3. Create a new configuration section to contain the task settings and set the following parameters:

Type The analysis engine to use. Set this parameter to `Demographics`.

Input The track that contains detected faces that you want to analyze. Set this parameter to the `ResultWithSource` output track from your face detection task. For example, if your face detection task is named `FaceDetect`, set this parameter to `FaceDetect.ResultWithSource`.

For example:

```
[FaceDemographics]
Type=demographics
Input=FaceDetect.ResultWithSource
```

For more information about these parameters, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Analyze Facial Expression

You can use the Face State analysis engine to obtain information about facial expression for detected faces.

Note: If the person has only one eye open (for example, if they are winking), Media Server reports that their eyes are open.

Note: To detect whether a person is wearing spectacles, the person must be looking directly at the camera.

To analyze facial expression

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=FaceDetect
AnalysisEngine1=FaceState
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type The analysis engine to use. Set this parameter to `FaceState`.

Input The track that contains detected faces that you want to analyze. Set this parameter to the `ResultWithSource` output track from your face detection task. For example, if your face detection task is named `FaceDetect`, set this parameter to `FaceDetect.ResultWithSource`.

For example:

```
[FaceState]
Type=facestate
Input=FaceDetect.ResultWithSource
```

For more information about these parameters, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Extract Keyframes

Media Server can identify the keyframes in a video. A *keyframe* is the first frame following a significant scene change. Keyframes are often used as preview images for video clips.

To extract keyframes from video

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=Keyframes
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type The analysis engine to use. Set this parameter to `keyframe`.

Input	(Optional) The name of the image track to process.
KeyAtLeastSec	(Optional) The maximum time interval in seconds between keyframes.
KeyAtMostSec	(Optional) The minimum time interval in seconds between keyframes.
Sensitivity	(Optional) The sensitivity of the engine to scene changes.

For example:

```
[Keyframes]  
Type=keyframe  
Sensitivity=0.6  
KeyAtLeastSec=120
```

For more information about the parameters that you can use to configure this task, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Recognize Objects

Object recognition detects the appearance of specific objects in video; for example, company logos or product packaging. Media Server reports the frames that an object was detected in and its location within the frames.

The Object engine automatically merges duplicate records to produce a single record for each appearance of a recognized object. For records to be identical, the object's unique name must match and the object must appear in a nearly-identical location. For example, if a recognized object appears in the same location for five consecutive frames, the engine merges the results for each frame into a single record with a timestamp duration that covers the five frames. If the object then moves to a different location on the screen, the engine counts this as a new detection and produces a separate record. If the object disappears then reappears, the engine reports the two detections separately.

Introduction

You can train Media Server to detect objects, such as logos, that appear in video. These objects can be two-dimensional or three-dimensional.

2-D object with 2-D similarity transform

For example, a company logo. Media Server can still detect the logo when it is subject to 2-D similarity transformations such as rotation and scaling.



2-D object with perspective transform

For example, a video containing a sign displaying a company logo.

Perspective transformations result from viewing the object from any angle other than directly perpendicular to its plane. This can result in skewing of the image. Perspective transformations often occur in photographs of objects when the camera axis is not directly aligned with the object.

Media Server can only detect 2-D objects that appear on flat, rigid objects. For example, if you attempt to detect a logo printed on a flag or item of clothing, detection may not be reliable.



3-D object

For example, a video of product packaging.



If you supply sufficient training images, Media Server can detect a 3-D object from any angle.











2D Object Detection

Media Server can detect 2-D objects in video, even when they are subject to similarity or perspective transformation. However, Media Server expects the object to appear on a flat, rigid surface.

The following table provides some examples of transformations that are supported and are not supported.

Transformation	Example	Supported?
Original object		✓
Scaling		✓

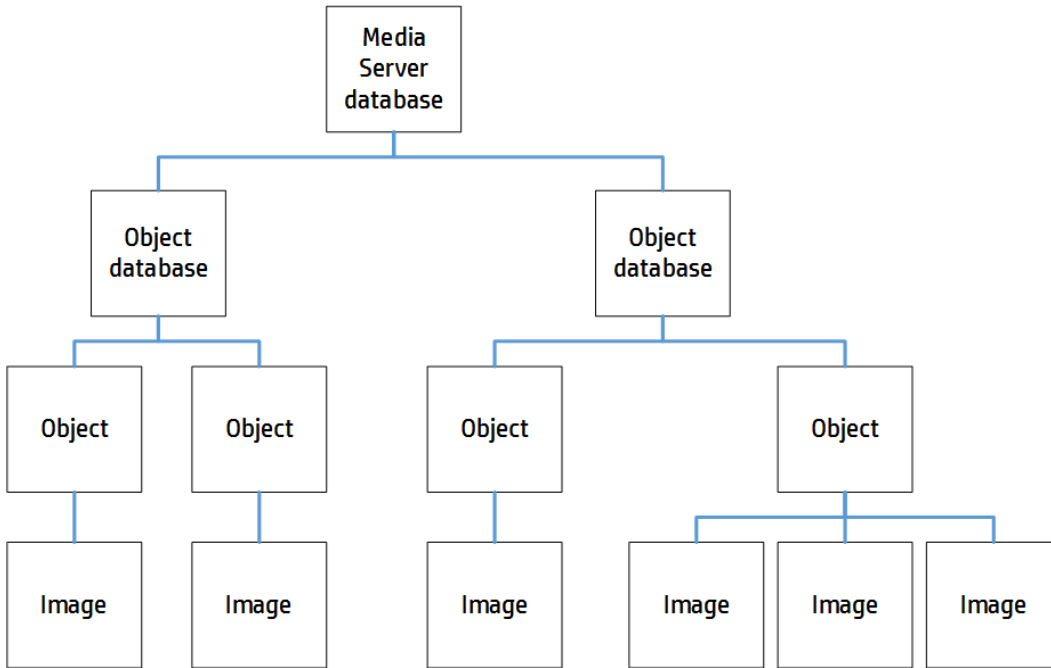
Rotation		
Perspective distortion or skew (horizontal or vertical only)		
Perspective distortion or skew (both horizontal and vertical)		
2-D object on a surface that is not flat		

Tip: HPE Media Server does not support the detection of 2-D objects on surfaces that are not flat. However, acceptable accuracy is achievable if you train Media Server using images of the object on the curved surface, rather than images of the object on a flat surface (so that the training is representative of the images you want to analyze).

Train Media Server to Recognize Objects

You must train Media Server by providing images of objects that you want to recognize. When you run Object Detection, Media Server uses the training data to recognize objects in the video.

The following diagram shows how Media Server stores information you provide during training.



The "Media Server database" represents the Media Server datastore file or a database on an external database server. For information on setting up this database, see ["Set up a Database to Store Training Data" on page 26](#).

You can organize objects into databases. These object databases are for organizational purposes. For example, you might have a database for company logos, and a database for your company's products. If you want to recognize only one category of objects in your images, you can use only that database.

A database can contain any number of objects, both 2D and 3D. You can associate training options and custom metadata with each object.

To each object you must add a training image (multiple training images for a 3-D object). For information about choosing suitable training images, see ["Select Images for Training" below](#).

Select Images for Training

The following table describes the training requirements for 2-D and 3-D objects:

Object type	Requirements
2-D	<p>One or more training images.</p> <p>Media Server creates a separate model for each training image. For example, if a company has many different variations of its logo, you only need to add one 2-D object to your object database. Add all of the training images for the logo to the same object, even if some variations of the logo do not look like the others.</p>
3-D	<p>Multiple training images depicting the same 3-D object from all angles.</p> <p>Provide images of the object from all angles that you expect it to appear in target images. HPE recommends that the images are taken in one</p>

	<p>continuous process. One way to obtain training images is by recording a video of the object from all angles and extracting images from the video.</p> <p>As a rough estimate, between 20 and 40 images of an object are usually sufficient to provide accurate detection in a small image.</p> <p>Media Server uses all of the training images to create a single model.</p>
--	---

A good training image for an object:

- contains only the object, with as little background as possible. The object should be the dominant feature in the image but there should be at least 16 pixels of background surrounding the object. If any parts of an object touch the image edges, the features corresponding to those parts of the image are likely to be lost.
- includes transparency information, if you intend to detect the object against many different backgrounds (for example, superimposed on TV footage). Ideally all parts of the image that are not part of the object are transparent (by having their alpha channel set appropriately).

Note: Only some image formats support transparency information, for example .PNG and .TIFF. Media Server does not support transparency information in .GIF files.

- has not been converted from another file format to JPEG. Converting images to JPEG introduces compression artifacts. If you have JPEG images, avoid making many rounds of edits. Each time that a new version of the image is saved as a JPEG, the image quality degrades.

Create a Database to Contain Objects

To create a database to contain objects, use the following procedure.

To create a database to contain objects

- Use the `CreateObjectDatabase` action, with the `database` parameter:

`database` The name of the new database (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=CreateObjectDatabase  
-F database=CompanyLogos
```

Add an Object to a Database

To add an object to database of objects, you can:

- **Perform training in separate steps.** Create a new (empty) object, then add training images, and then train Media Server, using separate actions. You can also add metadata to the object and configure training options, but these are optional steps. You might want to perform training in this way if you are building a front end application that responds to user input.
- **Perform training in a single action.** You might use this approach when writing a script to train Media Server from a collection of images and metadata.

Add an Object to a Database (Using Separate Steps)

This section describes how to create a new (empty) object, then add training images, and then train Media Server, using separate actions. You can also add metadata to the object, and configure training options, but these are optional steps. You might want to perform training in this way if you are building a front end application that responds to user input.

Alternatively, you can train Media Server to recognize an object by sending a single action. To do this, see ["Add an Object to a Database \(Using a Single Action\)" on the next page.](#)

To add an object to a database (using separate steps)

1. Add an object using the `NewObject` action. Set the following parameters:

<code>database</code>	The name of the database to add the object to. The database must already exist.
<code>identifier</code>	(Optional) A unique identifier for the object (maximum 254 bytes). If you do not set this parameter, Media Server generates an identifier automatically.

For example:

```
curl http://localhost:14000 -F action=NewObject  
                             -F database=CompanyLogos
```

Media Server adds an object to the database and returns the identifier.

2. Add one or more training images to the object using the `AddObjectImages` action. Set the following parameters:

<code>database</code>	The name of the database that contains the object.
<code>identifier</code>	The identifier for the object, returned by the <code>NewObject</code> action.
<code>imagedata</code>	The training images to add. Files must be uploaded as multipart/form-data. For more information about sending data to Media Server, see "Send Data by Using a POST Method" on page 61.
<code>imageLabels</code>	A comma-separated list of labels. One label is associated with each image. (maximum 254 bytes for each label). The number of labels must match the number of images. If you do not set this parameter, Media Server generates labels automatically.

For example:

```
curl http://localhost:14000 -F action=AddObjectImages  
                             -F database=CompanyLogos  
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62  
                             -F imagedata=@logo1.png
```

3. (Optional) Configure the way that Media Server is trained by setting training options for the object. To do this use the `SetObjectTrainingOption` action with the following parameters:

<code>database</code>	The name of the database that contains the object.
<code>identifier</code>	The identifier for the object, returned by the <code>NewObject</code> action.

key	The name of the training option that you want to change. For more information about training options, see "Object Training Options" on the next page .
value	The new value for the training option.

For example:

```
curl http://localhost:14000 -F action=SetObjectTrainingOption
                             -F database=CompanyLogos
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62
                             -F key=useColor
                             -F value=true
```

4. (Optional) Add metadata to the object using the AddObjectMetadata action. You can add any number of key-value pairs. Set the following parameters:

database	The name of the database that contains the object.
identifier	The identifier for the object, returned by the NewObject action.
key	The key to add (maximum 254 bytes).
value	The value to add (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=AddObjectMetadata
                             -F database=CompanyLogos
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62
                             -F key=CompanyName
                             -F value=HewlettPackard
```

5. Complete the training for the object using the BuildObject action. Set the following parameters:

database	The name of the database that contains the object.
identifier	The identifier for the object, returned by the NewObject action.

For example:

```
curl http://localhost:14000 -F action=BuildObject
                             -F database=CompanyLogos
                             -F identifier=6600dc0f9dd72d0cb55589e8f1d28b62
```

Add an Object to a Database (Using a Single Action)

You can train Media Server to recognize an object by sending a single action (TrainObject).

Running this action is equivalent to running the following actions in the following order:

- NewObject
- AddObjectImages
- SetObjectTrainingOption (optional)
- AddObjectMetadata (optional)
- BuildObject

The `TrainObject` action is atomic, so that any interruption to the server does not leave the database in an inconsistent state.

Alternatively, you can train Media Server by sending these actions individually. You might want to do this if you are building a front end application that trains Media Server in response to user input. For more information about how to do this, see ["Add an Object to a Database \(Using Separate Steps\)" on page 112](#).

To add an object to a database (using a single action)

- Add an object using the `TrainObject` action. Set the following parameters:

<code>database</code>	The name of the database to add the object to. The database must already exist.
<code>identifier</code>	(Optional) A unique identifier for the object (maximum 254 bytes). If you do not set this parameter, Media Server generates an identifier automatically.
<code>imagedata</code>	The training images to add. Files must be uploaded as multipart/form-data. For more information about sending data to Media Server, see "Send Data by Using a POST Method" on page 61 .
<code>imagelabels</code>	(Optional) A comma-separated list of labels. One label is associated with each image. (maximum 254 bytes for each label). The number of labels must match the number of images. If you do not set this parameter, Media Server generates labels automatically.
<code>metadata</code>	(Optional) A comma-separated list of metadata key-value pairs to add to the object. Separate keys from values using a colon (:). To include a comma or colon in a key name or value, you must enclose the key name or value in quotation marks (") and escape any quotation marks that occur within the string with a backslash (\).
<code>trainingoptions</code>	(Optional) A comma-separated list of training options to apply to the object. Separate training options from their values using a colon (:).

For example:

```
curl http://localhost:14000 -F action=TrainObject
-F database=CompanyLogos
-F imagedata=@logo.png
-F metadata=CompanyName:HPE, "another:value": "1,000"
-F trainingoptions=useColor:true,3D:false
```

Media Server adds the object to the database and returns the identifier.

Object Training Options

After you create a new object in an object database, you can set training options for the object. Training options configure how Media Server uses the training data that you supply to create a model of the object.

You can set the following training options:

Training Option	Description	Acceptable values	Default value
3D	Specifies whether the object is a three-dimensional object.	true, false	false
boundaryFeatures	Specifies whether there are important features at the edges of the training images. For example if you supply a training image of a rectangular flag, and the edge of the image represents the edge of the flag, set this option to true.	true, false	false
useColor	Specifies whether Media Server adds color information from the training images into the models that it creates for detection. For example, if you are detecting company logos but one company often prints its logo in different colors, set this option to false.	true, false	false

For an example that shows how to add a new object and set training options, see ["Add an Object to a Database" on page 111](#).

List the Objects in a Database

To list the objects that you have added to a database, and check whether training was successful, use the following procedure.

To list the objects in a database

1. (Optional) First list the databases that have been created to store objects. Use the action `ListObjectDatabases`:

```
http://localhost:14000/action=ListObjectDatabases
```

Media Server returns a list of databases that you have created.

2. List the objects that exist in one of the databases. Use the action `ListObjects`, for example:

```
http://localhost:14000/action=ListObjects&database=logos
&metadata=true
&trainingoptions=true
&imagestatus=true
```

Media Server returns a list of objects in the specified database, and the number of training images associated with each object.

If you set the action parameter `metadata` to `true`, Media Server returns the metadata you have added to the object.

If you set the action parameter `trainingoptions` to `true`, Media Server returns the training options you have set for the object.

If you set the action parameter `imagestatus` to `true`, Media Server returns the status of each training image associated with each object.

- A status of `trained` indicates that training was successful.
- If images have a status of `untrained`, run training for the object using the action `BuildObject`.
- If images have a status of `failed`, Media Server could not use the image for training. Remove the failed images using the action `RemoveObjectImages`.

Update or Remove Objects and Databases

To update or remove an object use the following actions:

- To add additional training images to an object, use the action `AddObjectImages`. Media Server does not use the new images for object detection until you run the action `BuildObject`. To confirm that training was successful, use the action `ListObjects` (see ["List the Objects in a Database" on the previous page](#)).
- To remove a training image from an object, use the action `RemoveObjectImages`.
- To modify the training options for an object, use the actions `SetObjectTrainingOption` and `UnsetObjectTrainingOption`.
- To change the label of an image that you added to an object, use the action `RelabelObjectImage`.
- To move an image of an object from one object to another, for example if you add an image to the wrong object, use the action `MoveObjectImage`.
- To add, remove, or update custom metadata for an object, use the actions `AddObjectMetadata`, `RemoveObjectMetadata`, and `UpdateObjectMetadata`.
- To change the identifier of an object you added to an object database, use the action `RenameObject`.
- To remove an object from a database and discard all associated images and metadata, use the action `RemoveObject`.

To update or remove object databases, use the following actions:

- To rename an object database, use the action `RenameObjectDatabase`.
- To delete an object database and all of the information that it contains, use the action `RemoveObjectDatabase`.

For more information about the actions that you can use to train Media Server, refer to the *Media Server Reference*.

IDOL Admin

You can train Media Server using IDOL Admin. To open the IDOL Admin interface, send the following action to Media Server's ACI port:

```
http://localhost:14000/action=admin
```

For more information about using IDOL Admin refer to the *IDOL Admin User Guide*.

Recognize Objects

To recognize objects in video, configure an object analysis task by following these steps.

To recognize objects in video

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=Object
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>object</code> .
Input	(Optional) The image track to process.
Database	(Optional) The database to use for recognizing objects. By default, Media Server searches for objects from all object databases. Database names are case-sensitive.
ColorAnalysis	(Optional) A Boolean value that specifies whether to check the color of detected objects in order to reduce false identification. Set this parameter if the objects are primarily distinguished by color; for example, some flags differ from each other by color only.
ObjectEnvironment	(Optional) Some objects, such as logos, might be partially or completely transparent, meaning that their appearance changes depending on the background on which they are superimposed. In such cases, set this parameter to specify the type of background in target images. Note: For the <code>ObjectEnvironment</code> parameter to have an effect, the image of the object used for training the database must contain transparent pixels.
Occlusion	(Optional) By default, Media Server assumes that an object might be partially hidden in target images, for example by overlapping objects. If the object is never obscured in target images, set this parameter to <code>FALSE</code> to reduce false positives.
Region	(Optional) Search for objects in a region of the image, instead of the entire image.

For example:

```
[Object]
Type=Object
Database=CompanyLogos
```

For more details about the parameters that you can use to customize object recognition, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

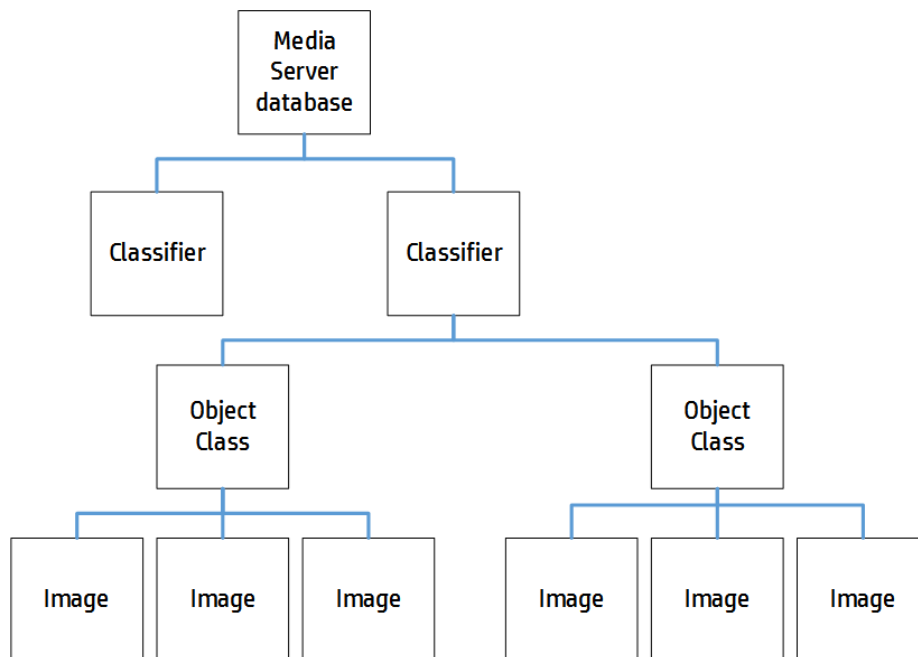
Classify Objects

Object classification classifies objects that are identified in a video. For example, you could create a classifier named `vehicles` and train Media Server to classify each vehicle detected by Automatic Number Plate Recognition or Intelligent Scene Analysis as a car, truck, or motorcycle.

Train Media Server to Classify Images

To classify objects, you must train Media Server by providing images that are representative of your chosen object classes.

The following diagram shows how Media Server stores information you provide during training.



The "Media Server database" represents the Media Server datastore file or a database on an external database server. For information about how to set up this database, see ["Set up a Database to Store Training Data" on page 26](#).

A classifier contains the object classes that you want to use for a classification task. When you run object classification, Media Server classifies your images into the classes in your chosen classifier. For example, to classify vehicles into "cars", "trucks", "motorcycles", and so on, you would create a classifier named "vehicles" and create object classes named "car", "truck", and "motorcycle".

A classifier must contain at least two object classes. To determine whether an image contains a car or doesn't contain a car, you would need to train a "car" object class and also a "not car" object class.

To each object class you must add training images. Usually around 100 training images per object class is sufficient to obtain good performance. For information about choosing suitable training images, see ["Training Requirements" on page 120](#).

Classifier Types

Media Server supports several types of object classifier.

HPE recommends trying each type of classifier to determine which produces the best results for your purposes. When you test your classifier, ensure that the images you use for testing are different to your training images.

Bayesian

The default type of classifier is the Bayesian classifier. Use a Bayesian classifier when you have a small number of classes and many training images for each class.

The Bayesian classifier is better than the Maxvote classifier at distinguishing between classes that have a high degree of similarity, and requires less time to train than the CNN classifier.

When you run classification, the Bayesian classifier outputs a confidence score for each class. These scores can be compared across classifiers, and you can set a threshold to discard results below a specified confidence level.

This is the only type of classifier that supports prior probabilities.

Maxvote

Use a Maxvote classifier when you have many classes with fewer training images per class. This type of classifier requires fewer training images and can be trained faster than either the Bayesian or CNN classifier.

When you run classification, the Maxvote classifier outputs a number of votes for each class in the classifier. This means that you cannot compare scores across classifiers.

CNN

The Convolutional Neural Network (CNN) classifier usually produces the most accurate results, but can require a significant amount of time to train.

The more time you allow Media Server to train the classifier, the greater the accuracy. Before you train a CNN classifier, you can choose how many training iterations to run. The time required to train the classifier is proportional to the number of training iterations and the number of training images. Increasing the number of iterations always improves the training and results in better accuracy, but each additional iteration that you add has a smaller effect.

For classifiers that have four or five dissimilar classes with around 100 training images per class, approximately 500 iterations produces reasonable results. This number of iterations with this number of training images requires approximately five hours to complete. HPE recommends a larger number of iterations for classifiers that contain many similar classes. For extremely complex classifiers that have hundreds of classes, you might run 200,000 training iterations. Be aware that running this number of training iterations with large numbers of training images is likely to take weeks.

To find the optimum number of iterations, HPE recommends that you start with a small number of iterations. Double the number of iterations each time you train, until classification accuracy is acceptable.

When you run classification, the CNN classifier outputs a confidence score for each class. These scores can be compared across classifiers, and you can set a threshold to discard results below a specified confidence level.

Training Requirements

The performance of classification is generally better if:

- the classifier contains only a few object classes (but it must contain at least two classes).
- the object classes are dissimilar. For example, when training a 'field' class and a 'beach' class, the presence of clouds in the sky in both sets of training images might cause confusion between the classes.
- the object classes are trained with many images. Usually around 100 images are sufficient to train an object class. If the images in a class are very similar, fewer images might be sufficient.
- the training images are representative of the variation typically found within the object class. For example, to train a "dog" class, use images of dogs of different sizes, breeds, colors, and from different viewpoints.
- the training images contain little background or clutter around the object in the image.
- the longest dimension (width or height) of the training image is at least 500 pixels - smaller images might result in reduced accuracy.

Tip: High-resolution images where the object covers a small proportion of the image make poor training images. If you have a large image showing the object and it can be cropped such that its longest dimension still exceeds 500 pixels, HPE recommends cropping the image. If you crop an image, leave a gap around the object of at least 16 pixels.

Create a Classifier

A classifier contains object classes. When you run object classification, Media Server uses a single classifier and classifies images into the classes in that classifier.

To create a classifier

1. Use the `CreateClassifier` action with the `classifier` parameter.

`classifier` The name of the new classifier (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=CreateClassifier  
                              -F classifier=vehicles
```

2. (Optional) Set training options on the classifier (such as selecting the [type of classifier](#)) using the `SetClassifierTrainingOption` action:

`classifier` The name of the classifier that you created in the previous step.

`key` The name of the training option to set.

`value` The value for the training option.

For example:


```
curl http://localhost:14000 -F action=SetClassifierTrainingOption  
-F classifier=vehicles  
-F key=classifierType  
-F value=maxvote
```

Classifier Training Options

After you create a new classifier, you can set training options for it.

Set or modify training options using the actions `SetClassifierTrainingOption` and `UnsetClassifierTrainingOption`. When you change a training option all training associated with the classifier becomes invalid and you must retrain the classifier using the `BuildClassifier` action. For more information about these actions, refer to the *Media Server Reference*.

You can set the following training options:

Training Option	Description	Default value
<code>classifierType</code>	The type of classifier that you want to train: <ul style="list-style-type: none">• Bayesian• Maxvote• CNN - Convolutional Neural Network. For information about each type of classifier, see " Classifier Types " on page 119	Bayesian
<code>iterations</code>	The number of iterations to perform when training a neural network classifier. For information about how to set this training option, see " Classifier Types " on page 119.	500

For an example that shows how to add a new classifier and set training options, see "[Create a Classifier](#)" on the previous page.

Add Classes to a Classifier

To add object classes to a classifier, complete the following procedure.

To add object classes to a classifier

1. Add each new class using the action `CreateClass`. Set the following parameters:

<code>classifier</code>	The name of the classifier to add the object class to. The classifier must already exist. To create a classifier, see "Create a Classifier" on page 120 .
<code>identifier</code>	(Optional) A unique identifier for the object class (maximum 254 bytes). If you do not set this parameter, Media Server generates an identifier automatically.
<code>imagedata</code>	The training images to add. Files must be uploaded as multipart/form-data. For more information about sending data to Media Server, see "Send Data by Using a POST Method" on page 61 .
<code>imagelabels</code>	(Optional) A comma-separated list of labels. One label is associated with each image. (maximum 254 bytes for each label). The number of labels must match the number of images. If you do not set this parameter, Media Server generates labels automatically.

For example:

```
curl http://localhost:14000 -F action=createclass
                             -F classifier=vehicles
                             -F identifier=cars
                             -F imagedata=@car1.jpg,car2.jpg,...
```

Media Server adds the new class.

2. (Optional) Add metadata to the class using the `AddClassMetadata` action. You can add any number of key-value pairs. Set the following parameters:

<code>classifier</code>	The name of the classifier that contains the class.
<code>identifier</code>	The identifier for the object class, as returned by the <code>CreateClass</code> action.
<code>key</code>	The key to add (maximum 254 bytes).
<code>value</code>	The value to add (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=AddClassMetadata
                             -F classifier=vehicles
                             -F identifier=cars
                             -F key=notes
                             -F value=motor%20vehicles
```

3. Complete the training for the classifier by running the action `BuildClassifier`:

```
curl http://localhost:14000 -F action=BuildClassifier
                             -F classifier=vehicles
```

This action is asynchronous, so Media Server returns a token. You can use the token with the `QueueInfo` action to retrieve the response.

List Classifiers and Object Classes

To list the classifiers that you have created, and check their status (whether they need training), use the following procedure.

To list classifiers

- Run the action `ListClassifiers`:

```
http://localhost:14000/action=ListClassifiers&trainingoptions=TRUE
```

Media Server returns a list of classifiers that you have created. For example:

```
<autnresponse>  
  <action>LISTCLASSIFIERS</action>  
  <response>SUCCESS</response>  
  <responsedata>  
    <classifier>  
      <classifier>vehicles</classifier>  
      <state>STALE</state>  
      <numclasses>3</numclasses>  
      <trainingoptions>  
        <trainingoption>  
          <key>classifiertype</key>  
          <value>bayesian</value>  
        </trainingoption>  
      </trainingoptions>  
    </classifier>  
  </responsedata>  
</autnresponse>
```

The `state` element specifies whether the classifier needs training. If this element contains the value "stale", train the classifier using the action `BuildClassifier`. The other possible states are `training`, `trained`, and `failed`.

The `numclasses` element specifies the number of object classes in the classifier.

If you set the `trainingoptions` parameter to `true`, the response also includes training options that you have set for the classifier.

To list the object classes in a classifier

- Run the action `ListClasses`:

```
http://localhost:14000/action=ListClasses&classifier=vehicles  
                                &metadata=true  
                                &imageLabels=true
```

Media Server returns a list of object classes in the specified classifier.

If you set the action parameter `imageLabels` to `true`, Media Server returns the labels of images associated with each object class.

If you set the action parameter `metadata` to `true`, Media Server returns the metadata you have added to each object class.

Update or Remove Object Classes and Classifiers

To update or remove an object class use the following actions:

- To add additional training images to an object class, use the action `AddClassImages`. After adding images to an object class you must retrain the classifier using the action `BuildClassifier`. To confirm that training was successful, use the action `ListClassifiers` (see "[List Classifiers and Object Classes](#)" on the previous page).
- To remove training images from an object class, use the action `RemoveClassImages`. After removing images from an object class you must retrain the classifier using the action `BuildClassifier`. To confirm that training was successful, use the action `ListClassifiers` (see "[List Classifiers and Object Classes](#)" on the previous page).
- To modify the training options for a classifier, use the actions `SetClassifierTrainingOption` and `UnsetClassifierTrainingOption`.
- To change the label of an image that you added to an object class, use the action `RelabelClassImage`.
- To move an image from one object class to another, for example if you add an image to the wrong object class, use the action `MoveClassImage`.
- To add, remove, or update custom metadata for an object class, use the actions `AddClassMetadata`, `RemoveClassMetadata`, and `UpdateClassMetadata`.
- To rename an object class, use the action `RenameClass`.
- To remove an object class from a classifier and discard all associated images and metadata, use the action `RemoveClass`.

To update or remove a classifier, use the following actions:

- To rename a classifier, use the action `RenameClassifier`.
- To delete a classifier and all of the information that it contains, use the action `RemoveClassifier`.

For more information about the actions that you can use to train Media Server, refer to the *Media Server Reference*.

Import a Classifier

HPE may provide classifiers that you can use with Media Server to classify video. These classifiers are pre-trained, and to use one you only need to import the training data into your Media Server database.

For a list of the available pre-trained classifiers, see "[Pre-Trained Classifiers](#)" on page 195.

To import a classifier

1. Go to <https://downloads.autonomy.com/productDownloads.jsp> and download the classifier to your Media Server machine.
2. Import the training data by running the action `ImportClassifier`, for example if you downloaded the file to the folder `pretrained/objectclass`, in the Media Server installation folder:

```
/action=ImportClassifier&classifier=imagenet  
&classifierpath=./pretrained/objectclass/imagenet.dat
```

where,

- | | |
|-----------------------------|---|
| <code>classifier</code> | The name to give to the imported classifier. |
| <code>classifierpath</code> | The path of the classifier data file on disk |
| <code>classifierdata</code> | The classifier data (you can set this as an alternative to <code>classifierpath</code>). |

Media Server imports the data. You can now run classification.

Classify Objects

To classify objects

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=ObjectClassRecog
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

- | | |
|-------------------------|--|
| <code>Type</code> | The analysis engine to use. Set this parameter to <code>objectclass</code> . |
| <code>Input</code> | (Optional) The image track to process. If you do not specify an input track, Media Server processes the first track of the correct type that is produced by the ingest engine. |
| <code>Classifier</code> | The classifier to use for object classification. Media Server categorizes objects into the classes in this classifier. |

For example:

```
[ObjectClassRecog]
Type=objectclass
Classifier=Vehicles
```

For more information about these parameters, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Perform Optical Character Recognition (OCR)

OCR converts text in a video frame into a text file. Media Server's OCR:

- produces output in several different formats.
- provides options to increase the accuracy of the OCR process for different types of text, such as different fonts.

- searches video frames for text-like regions, and only performs OCR on those regions.
- supports many languages, font types, and character sets.

To perform OCR

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the [Analysis] section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=OCR
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>OCR</code> .
Input	(Optional) The image track to process. If you do not specify an input track, Media Server processes the first track of the correct type that is produced by the ingest engine.
Languages	The language of the text. Setting this parameter narrows the character and word choices for the OCR process, which increases speed and accuracy. For a list of supported languages, see "Supported Languages" on page 194 .
CharacterTypes	(Optional) If the document uses a particular type of characters only, such as all uppercase, or all lowercase, you can specify the type in the <code>CharacterTypes</code> parameter. This can improve accuracy.
HollowText	(Optional) Set this parameter if you are processing subtitles that consist of white characters with black outlines.
Region	(Optional) Restrict OCR to a region of the image, instead of the entire image.

For example:

```
[OCR]
Type=ocr
Languages=en
```

For more information about the parameters you can use to customize OCR, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Perform Color Analysis

Media Server can identify the dominant colors in video frames.

The Color Cluster analysis engine determines the dominant colors in a video frame by clustering similar colors. The operation returns the color at the center of each cluster as a value in the selected color

space (for example, RGB). It also returns the proportion of the pixels in the frame that belong to each cluster.

To configure color analysis

1. Open the Media Server configuration file with a text editor (or create a new configuration file).
2. In the [Analysis] section, add a new analysis task using the AnalysisEngine parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=ColorClusterTask
```

3. Create a new configuration section to contain the task settings and set the following parameters:

Type	The analysis engine to use. Set this to <code>ColorCluster</code> .
Input	(Optional) The name of the track that contains the images to process. If you do not specify an input track, Media Server processes the first track of the correct type.
ColorSpace	(Optional) The color space in which the results of analysis are provided.

For example:

```
[ColorClusterTask]
Type=colorcluster
Input=keyframe.ResultWithSource
ColorSpace=RGB
```

For more information about the configuration parameters that you can set to customize the task, refer to the *Media Server Reference*.

4. Save and close the configuration file.

Analyze Vehicles

Media Server can detect and analyze vehicles in video. These features are suitable for many different applications. You can use Media Server to monitor car parks, provide a drive-off deterrent at petrol filling stations, and detect stolen vehicles.

Media Server can:

- **Detect and read the number plates of vehicles in the scene.** Number plate recognition has many applications; you can detect stolen and uninsured vehicles, and monitor the length of stay for vehicles in car parks.
- **Identify the manufacturer and model of a vehicle identified during number plate recognition.** By comparing the model detected by Media Server to a database, you can use this feature to identify vehicles that have false number plates.
- **Identify the color of a vehicle identified during vehicle model detection.**

Requirements for ANPR

For reliable number plate detection and recognition, ensure that your system meets the following requirements.

Camera	<p>HPE recommends a separate camera for each lane of traffic. If you use a single camera for several lanes of traffic, high-definition recording is required (1080p or better).</p> <p>Manually focus the camera on the middle of the region where number plates are read, and use a fast shutter speed to avoid blurring.</p>
Image contrast	<p>The contrast must be sufficient for the number plates to be human-readable. If you are reading retroreflective number plates, the best results are obtained with infra-red (IR) illumination.</p>
Number plate size	<p>The number plates must be human-readable and characters in the video must not be less than 10 pixels high.</p>
Number plate rotation	<p>Position the camera above the traffic, so that number plates appear to be horizontal. The closer the plates are to horizontal, the better the performance.</p>
Video frame rate	<p>Media Server improves accuracy by reading number plates across multiple frames. For moving traffic Media Server requires 25 or 30 frames per second. Do not reduce the frame rate of the captured video.</p>

Detect and Read Number Plates

Media Server can detect and read number plates that appear in video.

To detect and read number plates in video

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=ANPR
```

3. Create a new section in the configuration file to contain the task settings and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>numberplate</code> .
Input	(Optional) The image track to process. If you do not specify an input track,

	Media Server processes the first track of the correct type that is produced by the ingest engine.
Location	The location in which you are reading number plates. There are significant differences in the number plates used by countries and states, so this parameter is necessary to read the plates successfully.
Region	(Optional) The region of interest (ROI) to monitor for number plates (left, top, width, height). If you do not specify a region, Media Server detects and reads number plates anywhere in the scene.
RegionUnit	(Optional) The units used to specify the position and size of the region of interest (<code>pixel</code> or <code>percent</code>).
Sensitivity	(Optional) The confidence level required to detect a number plate.
Integration	(Optional) The method to use to combine the number plates read from successive frames into a single result (<code>none</code> , <code>moving</code> , or <code>static</code>).
MinValidScore	(Optional) The average character score required for a number plate to be recognized.

For example:

```
[ANPR]
Type=numberplate
Location=uk
RegionUnit=percent
Region=20,10,60,90
Sensitivity=14
Integration=moving
```

For more information about the parameters that customize number plate recognition, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Train Media Server to Recognize Vehicles

Before using vehicle model identification, you must train Media Server by providing images of the vehicle models that you want to recognize. Vehicle identification performs best when vehicles are moving towards the camera, when the front of the vehicle is visible in the video.

Vehicles must be added to an object database as two-dimensional objects.

Obtain Training Images

You can obtain training images by running vehicle model identification on a sample video. Output the result images to the image encoder. For example, you could use the following configuration:

```
[Ingest]
IngestEngine0=libav
```

```
[libav]
Type=libav

[Analysis]
AnalysisEngine0=ANPR
AnalysisEngine1=VehicleModel

[ANPR]
Type=numberplate
Location=uk

[VehicleModel]
Type=vehiclemodel
Input=ANPR.DataWithSource

[Encoding]
EncodingEngine0=VehicleImages

[VehicleImages]
Type=imageencoder
ImageInput=VehicleModel.ResultWithSource
OutputPath=./training/image-%segment.sequence%.jpg
Crop=TRUE
```

The vehicle model analysis engine requires as input either the `DataWithSource` or `ResultWithSource` track from a number plate analysis task.

Usually, the vehicle model analysis task would include the `Database` parameter, which specifies the database to use for recognizing vehicles. If you are obtaining training images, you can run vehicle identification without setting this parameter.

When you output images using the image encoder, set `Crop=TRUE` so that the images are cropped to the region identified by the vehicle model analysis engine.

Train Media Server

Vehicle identification is trained in the same way as object detection (see "[Train Media Server to Recognize Objects](#)" on page 109). Vehicles must be trained as two-dimensional objects.

Tip: The following procedure explains how to train Media Server by sending actions to Media Server's ACI port. However, you can also train Media Server using IDOL Admin. For information about opening IDOL Admin, see "[Access IDOL Admin](#)" on page 59. For information about using IDOL Admin, refer to the *IDOL Admin User Guide*.

To train Media Server to recognize vehicles

1. Create a new object database. Use the `CreateObjectDatabase` action, with the database parameter:

database The name of the new database (maximum 254 bytes).

For example:

```
curl http://localhost:14000 -F action=CreateObjectDatabase
                             -F database=vehicles
```

2. Add each vehicle to the database using the TrainObject action. Use the following parameters:

database	The name of the database to add the object to. The database must already exist.
identifier	A unique identifier for the object (maximum 254 bytes). If you do not set this parameter, Media Server generates an identifier automatically.
imagedata	The training images to add. Files must be uploaded as multipart/form-data. For more information about sending data to Media Server, see "Send Data by Using a POST Method" on page 61 .
imagelabels	A comma-separated list of labels to associate with the images that you are adding (maximum 254 bytes for each label). The number of labels must match the number of images provided. If you do not set this parameter, Media Server generates labels automatically.
trainingoptions	A comma-separated list of training options to apply to the object. Separate training options from their values using a colon (:). For information about the training options that you can set, see "Object Training Options" on page 114 .

For example:

```
curl http://localhost:14000 -F action=TrainObject
                             -F database=vehicles
                             -F identifier=ford_focus
                             -F imagedata=@image-1.jpg,image-27.jpg,image-33.jpg
                             -F imagelabels=2013focus,2005focus,2002focus
                             -F trainingoptions=useColor:true,3D:false
```

3. Verify that the training was successful.

- The TrainObject action is asynchronous. To obtain the response to these actions, use the QueueInfo action. For example:

```
curl
"http://localhost:14000/a=QueueInfo&queuename=trainobject&queueaction=getstatus"
```

- To check the status of the vehicles database, and verify that training was successful, use the ListObjects action. For example:

```
curl "http://localhost:14000/a=ListObjects&database=vehicles"
```

Identify Vehicle Models

To identify vehicles in video, use the following procedure.

To identify vehicles

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=ANPR
AnalysisEngine1=VehicleModel
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The analysis engine to use. Set this parameter to <code>vehiclemodel</code> .
Input	The track to process. This track must be the <code>DataWithSource</code> or <code>ResultWithSource</code> track from a number plate analysis task.
Database	The name of the object database to use to identify vehicles. For information about creating this database, see " Train Media Server to Recognize Vehicles " on page 129.
Perspective	(Optional) If the video frames show the vehicle from a different perspective to that used in the training images, set this parameter to <code>TRUE</code> so that Media Server compensates accordingly.
ColorRegion	(Optional) The region to output to the <code>ColorRegionWithSource</code> output track, so that you can configure an analysis task to identify the color of the vehicle . If you don't specify a region, Media Server uses a default region.

For example:

```
[VehicleModel]
Type=vehiclemodel
Input=ANPR.DataWithSource
Database=vehicles
Perspective=FALSE
```

For more information about the parameters that you can use, including the values that they accept, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Identify Vehicle Colors

Media Server can detect the color of vehicles identified by vehicle model detection.

To detect the color of vehicles

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).

2. In the [Analysis] section, add a new analysis task by setting the AnalysisEngineN parameter. You can give the task any name, for example:

```
[Analysis]
AnalysisEngine0=ANPR
AnalysisEngine1=VehicleModel
AnalysisEngine2=VehicleColor
```

3. Create a new section in the configuration file to contain the task settings and set the following parameters:

Type	The analysis engine to use. Set this parameter to ColorClusterRegion .
Input	The image track to process. This track must be the ColorRegionWithSource output track from a vehicle model analysis task. This track contains a region identified by the vehicle model task, which should contain a sample of the vehicle's color. You can customize the selection of this region by setting the parameter ColorRegion in your vehicle model analysis task .
ColorThreshold	The analysis task discards colors that do not make up a significant proportion of the region (as specified by this parameter).
ColorSpace	The color space in which the results are provided (RGB, YCbCr, HSL, HSV, CIELAB).

For example:

```
[VehicleColor]
Type=ColorClusterRegion
Input=VehicleModel.ColorRegionWithSource
ColorThreshold=20
ColorSpace=HSV
```

For more information about the parameters that customize color analysis, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Scene Analysis

Scene analysis detects important events that occur in video. You can use scene analysis to monitor video streamed from CCTV cameras, to assist you with the detection of potential threats, illegal actions, or alert you to situations where help is required. The Scene Analysis engine can be trained to detect whatever you require.

Typical examples of objects and events you might want to detect in CCTV footage include:

- Abandoned bags or vehicles.
- Zone breaches and trip wire events, for example a person entering a restricted area.
- A vehicle breaking traffic laws or running a red light.
- Pedestrian or vehicular traffic congestion.

Train the Scene Analysis Engine

To run scene analysis, you must create a training configuration that describes the events that you want to detect. To create the training configuration, use the Scene Analysis Training Utility.

You can use the Training Utility to:

- Define one or more regions of interest in the scene.
- Mask parts of the scene that you do not want to monitor.
- Define the size, shape, orientation, velocity, and color of the objects that you want to detect, and the permitted variations in all of these properties.
- Define the position of traffic lights in the scene, so that Media Server can read the lights.
- Display the video being analyzed by Media Server so that you can confirm objects are being tracked correctly.
- Set up filters to reduce the number of false alarms. For example, you might want Media Server to generate alarms only for objects that are larger than a particular size.
- Review the alarms that have been generated using your training configuration, and classify those alarms as suspicious (true alarms) or not suspicious (false alarms). The Training Utility can then suggest improvements to the configuration to minimize the number of false alarms and the number of missed alarms.

While you are training Media Server, you can use the Training Utility to start and stop ingestion and analysis (`process` actions). When you start processing through the Training Utility, Media Server makes video frames and alarm data available to the Training Utility. The Training Utility needs this data during the training process.

After you have finished training, send your finished configuration to Media Server. To run Scene Analysis in a production environment, do not start processing through the Training Utility. Instead, [create a new configuration containing a scene analysis task](#), and start the `process` action as described in ["Start Processing Video" on page 84](#).

For more information about using the Scene Analysis Training Utility, refer to the documentation for the Training Utility.

Run Scene Analysis

While you are training Media Server, you can use the Training Utility to start and stop ingestion and analysis (`process` actions). However, to run scene analysis in a production environment, do not start processing through the Training Utility. Instead, use the following procedure to create a configuration that contains a scene analysis task, and start the `process` action as described in ["Start Processing Video" on page 84](#).

To detect important events in video

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Analysis]` section, add a new analysis task by setting the `AnalysisEngineN` parameter.

You can give the task any name, for example:

```
[Analysis]  
AnalysisEngine0=ISAS
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The analysis engine to use. Set this parameter to SceneAnalysis .
TrainingCfgName	The name of the training configuration file to use to detect important events. This is the name of the file that you created using the Scene Analysis Training Utility.

For example:

```
[ISAS]  
Type=SceneAnalysis  
TrainingCfgName=redlights
```

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Chapter 9: Event Stream Processing

This section describes event stream processing in Media Server.

- [Introduction to Event Stream Processing](#) 136
- [Filter a Track](#) 137
- [Deduplicate Records in a Track](#) 138
- [Combine Tracks](#) 141
- [Identify Time-Related Events in Two Tracks–And Engine](#) 142
- [Identify Time-Related Events in Two Tracks–AndThen Engine](#) 144
- [Identify Isolated Events–AndNot Engine](#) 145
- [Identify Isolated Events–AndNotThen Engine](#) 147
- [Write a Lua Script for an ESP Engine](#) 148

Introduction to Event Stream Processing

Event Stream Processing (ESP) applies rules to the output of other tasks, including other ESP tasks.

One use of ESP is to identify meaningful events among a large number of records. A simple example is detecting the occurrence of a particular word in an audio stream. A more complex example is detecting when two events occur within a specific time interval; for example, a number plate is detected shortly after a traffic light changes to red.

The tracks passed to ESP engines are not modified. An ESP task produces a new output track that contains records which meet the specified conditions.

ESP engine	Description	Example
Filter	Filters a track, producing an output track that contains only records that meet specified conditions.	Filter speech-to-text results for a news channel to produce an output track that only contains records for the word "weather".
Deduplicate	Identifies duplicate records in a track, and produces a new track that has the duplicate records removed.	Face recognition can produce a record for each frame that a recognized person appears in. The Deduplicate engine can filter the output to remove any identical records produced within a specific number of seconds.
Or	Combines two or more tracks into a single track. The "Or" ESP engine creates an output track that contains the records from all of the input tracks.	Combine tracks from OCR, speech-to-text, and face recognition engines to produce an output track containing information about text, speech, and faces detected in the video.

And	Compares two tracks to identify events in the second track that occur within a specific time interval of events in the first track (before or after the first track's event). The engine produces a track containing the identified record pairs.	Take two OCR tracks from the same news broadcast, one track filtered for the string "election" and the other filtered for the string "results", and extract the records that fulfill the condition that "results" occurs within three seconds of "election" (either before or after).
AndThen	Compares two tracks to identify events in the second track that occur within a specific time interval after (or at the same time as) events in the first track. The engine produces a track containing the identified record pairs.	Take two OCR tracks from the same news broadcast, one track filtered for the string "breaking news" and the other filtered for the string "Syria", and extract the records that fulfill the condition that "Syria" occurs up to three seconds after "breaking news".
AndNot	Compares two tracks and produces a track containing records from the first track for which there is no event in the second track within a specified time interval (before or after the first track's event).	Produce a track containing records for all appearances of a logo that are not accompanied by the company name in the ten seconds preceding or following the logo's appearance.
AndNotThen	Compares two tracks and produces a track containing records from the first track for events that are not followed within a specified time interval by events in the second track.	Produce a track containing records for all appearances of a logo that are not followed within ten seconds by the company name.

The ESP engines support configuration parameters that allow you to customize the operations for your data. Some engines also allow you to run scripts written in the Lua scripting language. For more information about Lua scripts, see ["Write a Lua Script for an ESP Engine" on page 148](#).

ESP engines can accept any track. They also accept the output of other ESP engines.

Filter a Track

You can filter a track to extract records that match particular conditions. For example, you can:

- extract records from OCR analysis results that match or contain a specified string
- extract records from a speech-to-text task that match or contain a specified string
- extract records from a face detection task that describe faces that appear in a particular region of the frame
- extract records from an object detection task that match a specific object

To filter a track

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).

2. In the [EventProcessing] section, add a new task by setting the EventProcessingEngineN parameter. You can give the task any name, for example:

```
[EventProcessing]
EventProcessingEngine0=Weather
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type The ESP engine to use. Set this parameter to `filter`.

Input The output track, produced by another Media Server task, that you want to filter.

4. Set one of the following parameters to specify how the input track is filtered:

RequiredString A string that a record must match to be included in the output track. (The input track must contain text data).

RequiredSubString A string that a record must contain to be included in the output track. (The input track must contain text data).

LuaScript The name of a Lua script that defines conditions that a record must meet in order to be included in the output track from the ESP engine. For more information, see ["Write a Lua Script for an ESP Engine" on page 148](#).

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following configuration produces a new track called `Weather.Output`. This track only contains records that include the word "weather".

```
[EventProcessing]
EventProcessingEngine0=Weather
```

```
[Weather]
Type=filter
Input=speechtotext.result
RequiredSubString=weather
```

Deduplicate Records in a Track

Deduplication identifies duplicate records in a track, and produces a new track that has the duplicate records removed.

The engine identifies two identical records that occur within a specific time interval and discards the second record. For example, face recognition can produce a record for each frame that a person is recognized in. The Deduplicate ESP engine can remove duplicate records so that the track contains a single record for each recognized person.

You can specify the conditions that make two records identical. There are several options:

- any records are considered identical; Media Server discards any record that occurs within the minimum time interval of the first record

- use the default equivalence conditions for the track. Each type of track has its own default equivalence conditions; for example, OCR records are considered equivalent if the text is identical. The table lists the equivalence conditions for each output track.

Analysis engine	Output tracks	Equivalence conditions
Barcode	Data, DataWithSource, Result, ResultWithSource	Text field must be identical.
Face Detection	Data, DataWithSource, Result, ResultWithSource	Rectangle field must be identical.
Face Demographics	Result, ResultWithSource	All custom fields must be identical.
Face Recognition	Result, ResultWithSource	Uniquename field must be identical.
Face State	Result, ResultWithSource	All custom fields must be identical.
Numberplate	Data, DataWithSource, Result, ResultWithSource	Text field must be identical.
	PlateRegion	Polygon field must be identical.
Object	Data, DataWithSource, Result, ResultWithSource	Uniquename field must be identical.
ObjectClass	Result, ResultWithSource	Uniquename field must be identical.
OCR	Data, DataWithSource, Result, ResultWithSource	Text field must be identical.
SceneAnalysis	Data, DataWithSource, Result, ResultWithSource	All custom fields must be identical.
SpeakerID	Result	Text field must be identical.
SpeechToText	Result	Text field must be identical.

- specify equivalence conditions using a Lua script. For example, you might want to declare two Face records identical if they contain the same person, even if the location of the face in the frame is different. For more information about Lua scripts, see ["Write a Lua Script for an ESP Engine" on page 148](#).

To deduplicate a track

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (mediaserver.cfg).
2. In the [EventProcessing] section, add a new task by setting the EventProcessingEngineN

parameter. You can give the task any name, for example:

```
[EventProcessing]
EventProcessingEngine0=Deduplicate
```

3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The ESP engine to use. Set this parameter to <code>deduplicate</code> .
Input	The name of the track to deduplicate. This must be an output track produced by another task.
MinTimeInterval	The minimum time in milliseconds between records. The engine only discards duplicate records that occur within this time interval.
PredicateType	(Optional) The conditions to use to determine whether two records are considered identical. You can set one of: <ul style="list-style-type: none">• <code>always</code>. Any records are considered identical.• <code>default</code>. Use the default equivalence conditions for the track type.• <code>lua</code>. Use the conditions defined in a Lua script specified in the <code>LuaScript</code> parameter.
LuaScript	(Optional) The name of a Lua script that determines whether two records are considered identical. For more information, see "Write a Lua Script for an ESP Engine" on page 148 . If this parameter is set, Media Server uses the Lua script to determine whether records are identical, regardless of the <code>PredicateType</code> parameter setting.

For more information about these parameters, including the values that they accept, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following example deduplicates the output track from an OCR task by discarding all identical records that occur within 1 second after a record. The records are judged to be identical based on the default equivalence conditions for the OCR track (the text is identical).

```
[EventProcessing]
EventProcessingEngine0=DeduplicateOCR

[DeduplicateOCR]
Type=deduplicate
Input=myocr.data
MinTimeInterval=1000
PredicateType=default
```

Combine Tracks

You can combine two or more tracks into a single track. The "Or" ESP engine creates an output track that contains the records from all of the input tracks. Each record in the resulting track includes the name of the track that it originated from.

For example, you could combine output tracks from speech-to-text and face recognition. The records in the resulting track would contain a transcript of speech in the video, or details of recognized faces.

Tip: This engine combines tracks, not records.

To combine tracks

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[EventProcessing]` section, add a new ESP task by setting the `EventProcessingEngineN` parameter. You can give the task any name, for example:

```
[EventProcessing]
EventProcessingEngine0=Combine
```
3. Create a new configuration section to contain the task settings, and set the following parameters:

Type	The ESP engine to use. Set this parameter to <code>or</code> .
Input N	The names of the tracks that you want to output as a single track. Specify two or more tracks that are produced by other tasks.
4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following configuration combines output tracks from a Barcode task and an OCR task:

```
[EventProcessing]
EventProcessingEngine0=Combine
```

```
[Combine]
Type=or
Input0=mybarcode.result
Input1=myocr.result
```

This task produces a new track, named `Combine.Output`, that contains all of the records from `mybarcode.result` and `myocr.result`.

Identify Time-Related Events in Two Tracks–And Engine

The *And* ESP engine compares two tracks to identify events in the second track that occur within a specific time interval of events in the first track (before and/or after the first track's event). The engine produces a track containing the identified record pairs.

Note: To detect events in the second track that occur only after (or at the same time as) events in the first track, use the *AndThen* engine. For more information, see ["Identify Time-Related Events in Two Tracks–AndThen Engine" on page 144](#).

For example, you might want to identify all the times that a specific person appears in conjunction with a specific product. There are several methods to perform this task:

- Set up the analysis engines so they produce output tracks containing only relevant events; in this case, you would configure the Face engine to detect appearances of the specified person only, and the Object engine to detect appearances of the specified product only. You could then send the output tracks from the analysis tasks to the And engine to produce a track containing only record pairs that occurred within the specified time interval.
- Filter the output tracks from the analysis engines, before sending the filtered tracks to the And engine. In this case, you do not need to configure the Face and Object engines to detect only specific faces and objects; these would be extracted by Filter engines before being sent to the And engine. As with the previous method, the And engine produces a track containing only record pairs that occurred within the specified time interval.
- Send the unfiltered output tracks from the analysis engines to an And engine that uses a Lua script to determine which events it should consider. Each time the engine detects record pairs that occur within the specified time interval of each other, the engine runs the function in the Lua script to determine if the record pairs should be included in the And output track. For more information on writing a Lua script, see ["Write a Lua Script for an ESP Engine" on page 148](#).

To identify time-linked events in two tracks

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[EventProcessing]` section, add a new task by setting the `EventProcessingEngineN` parameter. You can give the task any name, for example:

```
[EventProcessing]
EventProcessingEngine0=AndEvents
```

3. Create a new configuration section for the task, and set the following parameters:

Type	The ESP engine to use. Set this parameter to <code>and</code> .
Input0	The first input track. This track must be an output track produced by another task.

Input1	The second input track. This track must be an output track produced by another task.
MaxTimeInterval	The maximum time in milliseconds from the record in the first track to the record in the second track, for the two records to be considered a pair. (The engine compares the timestamp values.)
MinTimeInterval	(Optional) The minimum time in milliseconds from the record in the first track to the record in the second track, for the two records to be considered a pair. (The engine compares the timestamp values.) The default value is the negative of the MaxTimeInterval value, meaning that the event in the second track can occur before the event in the first track (up to the specified number of milliseconds).
LuaScript	(Optional) The name of a Lua script that defines conditions that a record pair must meet in order to be included in the output track. For information about writing the script, see "Write a Lua Script for an ESP Engine" on page 148 .

For more details about these parameters, including the values that they accept, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following example produces an output track named `BreakingNewsSyria.Output`. This track contains speech-to-text records that contain the word "Syria", and OCR records that match the string "Breaking News". However, the records are only included when they occur within two seconds (2000 milliseconds) of the other record type.

ESP filter tasks are used to filter the OCR and speech-to-text results, before those results are passed to the "and" ESP task.

```
[EventProcessing]
EventProcessingEngine0=BreakingNews
EventProcessingEngine1=Syria
EventProcessingEngine2=BreakingNewsSyria
```

```
[BreakingNews]
Type=filter
Input=ocr.result
RequiredString=Breaking News
```

```
[Syria]
Type=filter
Input=speechtotext.result
RequiredSubString=Syria
```

```
[BreakingNewsSyria]
Type=and
Input0=BreakingNews.output
```

```
Input1=Syria.output  
MaxTimeInterval=2000
```

Identify Time-Related Events in Two Tracks–AndThen Engine

The *AndThen* ESP engine compares two tracks to identify events in the second track that occur within a specific time interval *after* (or at the same time as) events in the first track. The engine produces a track containing the identified record pairs.

Note: The *AndThen* engine enforces the order of events in the two tracks: events in the second track must occur after events in the first track. To detect events in two tracks that occur within a specified time interval, when the event order does not matter, use an *And* task. For more information, see ["Identify Time-Related Events in Two Tracks–And Engine" on page 142](#).

To identify time-linked events in two tracks

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[EventProcessing]` section, add a new task by setting the `EventProcessingEngineN` parameter. You can give the task any name, for example:

```
[EventProcessing]  
EventProcessingEngine0=AndThen
```

3. Create a new configuration section for the task and set the following parameters:

Type	The ESP engine to use. Set this parameter to <code>andthen</code> .
Input0	The first input track. This track must be an output track produced by another task.
Input1	The second input track. This track must be an output track produced by another task.

Note: The events to detect in the `Input1` track must occur after, or at the same time as, the events to detect in the `Input0` track.

MaxTimeInterval	The maximum time in milliseconds from the record in the first track to the record in the second track. (The engine compares the timestamp values.)
MinTimeInterval	(Optional) The minimum time in milliseconds from the record in the first track to the record in the second track. (The engine compares the timestamp values.) The default value is 0 (zero), meaning that the event in the second track must occur after (or at the same time as) the event in the first track.
LuaScript	(Optional) The name of a Lua script that defines conditions that a record pair

must meet in order to be included in the output track. For information about writing the script, see ["Write a Lua Script for an ESP Engine" on page 148](#).

For more information about these parameters and the values they accept, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following example produces an output track (`RedLightBreach.Output`) containing records produced when a number plate is detected up to five seconds (5000 milliseconds) after traffic lights turn red. The engine takes tracks produced by `SceneAnalysis` and `NumberPlate` engines. The `SceneAnalysis` output track contains a record for every time the traffic lights turned red. The `Numberplate` output track contains a record for every number plate detected in the video.

```
[EventProcessing]
EventProcessingEngine0=RedLightBreach

[RedLightBreach]
Type=andthen
Input0=sceneanalysis.ResultWithSource
Input1=numberplate.ResultWithSource
MaxTimeInterval=5000
```

Identify Isolated Events–AndNot Engine

The *AndNot* engine compares two tracks and produces a track containing records from the first track for which there is no event in the second track within a specified time interval (before and/or after the first track's event).

For example, you can identify all the occasions when a company logo appears in a video without mention of the company name in the speech:

1. An object detection task detects appearances of the logo.
2. A speech-to-text task produces a transcript of the speech.
3. An ESP filter task extracts records from the speech-to-text output that contain the company name and outputs them to a new track.
4. An ESP task (using the *AndNot* engine) uses the output tracks from the object detection and filter tasks. It compares events in both tracks and identifies isolated events in the first track (the object track). The engine produces an output track containing records from the object track for when a logo is detected but is not followed within the specified time interval by the company name in the filtered `SpeechToText` track.

Note: The *AndNot* engine does not enforce an order of events in the two tracks: the first track's event is considered isolated only if an event in the second track does not occur either before or after it. If you want to consider only events in the second track occurring after (or at the same time as) events in the first track, use the *AndNotThen* engine (see ["Identify Isolated Events–AndNotThen Engine" on page 147](#)).

To identify isolated events

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[EventProcessing]` section, add a new task by setting the `EventProcessingEngineN` parameter. You can give the task any name, for example:

```
[EventProcessing]
EventProcessingEngine0=IsolatedEvents
```

3. Create a new configuration section for the task and set the following parameters:

Type	The ESP engine to use. Set this parameter to <code>andNot</code> .
Input0	The first input track. This track must be an output track produced by another task.
Input1	The second input track. This track must be an output track produced by another task.
MaxTimeInterval	The maximum time in milliseconds from the record in the first track to the record in the second track. (The engine compares the timestamp values.) If an event in the second track occurs within this time interval from the event in the first track, the engine discards the event in the first track.
MinTimeInterval	(Optional) The minimum time in milliseconds from the record in the first track to the record in the second track, for the two records to be considered a pair. (The engine compares the timestamp values.) The default value is the negative of the <code>MaxTimeInterval</code> value, meaning that the event in the second track can occur before the event in the first track (up to the specified number of milliseconds).
LuaScript	(Optional) The name of a Lua script that defines conditions for a discarding a record from the first track. For information about writing the script, see "Write a Lua Script for an ESP Engine" on page 148 .

For more details about the parameters, including the values that they accept, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following example produces an output track called `LogoWithoutCompanyName.Output`. This track contains records from an object detection task that indicate when the company logo was detected. However, the track only contains the appearances when the company name was not mentioned in the audio within five seconds.

```
[EventProcessing]
EventProcessingEngine0=FilterAudio
EventProcessingEngine1=LogoWithoutCompanyName
```

```
[FilterAudio]
```

```
Type=filter
...

[LogoWithoutCompanyName]
Type=andnot
Input0=DetectCompanyLogo.Result
Input1=FilterAudio.output
MaxTimeInterval=5000
```

Identify Isolated Events–AndNotThen Engine

The *AndNotThen* engine compares two tracks and produces a track containing records from the first track for events that are not followed within a specified time interval by events in the second track.

Note: The *AndNotThen* engine enforces the order of events in the two tracks: the first track's event is considered isolated only if an event in the second track does not occur after (or at the same time as) it. If you want to consider events in the second track occurring both before or after events in the first track, use the *AndNot* engine (see "[Identify Isolated Events–AndNot Engine](#)" on page 145).

To identify isolated events

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[EventProcessing]` section, add a new task by setting the `EventProcessingEngineN` parameter. You can give the task any name, for example:

```
[EventProcessing]
EventProcessingEngine0=MyAndNotThen
```

3. Create a configuration section for the task and set the following parameters:

Type	The ESP engine to use. Set this parameter to <code>andNotThen</code> .
Input0	The first input track. This track must be an output track produced by another task.
Input1	The second input track. This track must be an output track produced by another task.
MaxTimeInterval	The maximum time in milliseconds from the record in the first track to the record in the second track. (The engine compares the timestamp values.) If an event in the second track occurs within this time interval from the event in the first track, the engine discards the event in the first track.
MinTimeInterval	(Optional) The minimum time in milliseconds from the record in the first track to the record in the second track. (The engine compares the timestamp values.)
LuaScript	(Optional) The name of a Lua script that defines conditions for a discarding

a record from the first track. For information about writing the script, see ["Write a Lua Script for an ESP Engine" below](#).

For more information about these parameters, including the values that they accept, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following example produces an output track containing records produced when a number plate is not detected within thirty seconds (30,000 milliseconds) of a barrier being raised. The AndNot engine takes tracks produced by SceneAnalysis and NumberPlate engines. The SceneAnalysis output track contains a record for every time a barrier is raised. The Numberplate output track contains a record for every number plate detected in the video. The AndNot engine output track contains all records from the SceneAnalysis output track that are not followed within thirty seconds by an event in the NumberPlate track.

```
[EventProcessing]
EventProcessingEngine0=Barrier

[Barrier]
Type=andnotthen
Input0=sceneanalysis.Result
Input1=numberplate.Result
MaxTimeInterval=30000
```

Write a Lua Script for an ESP Engine

All of the ESP engines, except for the "Or" engine, can run a Lua script to determine whether to include a record in the task's output track. Writing a Lua script allows you to specify more complex rules than you can specify using configuration parameters. For example, a Lua script might specify where in an image recognized text must appear. If text appears within this region, then depending on the engine type, the engine includes or excludes this record from its output track.

The Lua script must define a function with the name `pred`. This function takes one or two parameters (depending on the engine type) and must return `true` or `false`. Each parameter is a record object: this is a representation of the record being processed and has the same structure as the record XML.

The `pred` function is called once for each record (or pair of records) that the engine has to consider. The engine's response to the function depends on the engine type.

ESP Engine	Number of record parameters	Engine response if the function returns true
And	2	The record pair is included in the engine output track.
AndNot	2	The record in the first track is discarded. (Records from the second track are never included in the output anyway.)
AndNotThen	2	The record in the first track is discarded. (Records from the

		second track are never included in the output anyway.)
AndThen	2	The record pair is included in the engine output track.
Deduplicate	2	The second record is discarded.
Filter	1	The record is included in the engine output track.

When the `pred` function takes two parameters, each individual record may feature in many different pairs, so might be processed by the `pred` function any number of times. For example, for the `AndNot` or `AndNotThen` engine, a record in the first track might be passed to the function several times, being paired with different records from the second track. The record will only appear in the output track if `pred` returns `false` every time.

The ESP engine cannot modify the record objects passed to the `pred` function. Any effects of the function other than the return value are ignored.

To run a Lua script from an ESP engine, add the `LuaScript` parameter to the task configuration section and set it to the path of the script. For example, `LuaScript=./scripts/breakingnews.lua`.

Example Script

The following is an example script for the `Filter` ESP engine. The script filters records based on where text, read by an OCR task, appears in the image. The function returns `true` if text appears in the region between `x=100` and `x=300`, and the record is included in the output track. If the text region is outside these coordinates, the record is discarded.

```
function pred(rec)
    return rec.OCRResult.region.left > 100 and rec.OCRResult.region.right < 300
end
```

Chapter 10: Transform Data

This section describes how to transform the data produced by Media Server, so that you can customize how it is encoded or output to external systems.

- [Resize Images](#)150
- [Change the Format of Images](#) 151

Resize Images

Many processing tasks produce images. When Media Server ingests video, it decodes the video into individual frames. Analysis tasks can also produce images, for example keyframe extraction, face detection, and number plate recognition produce images of keyframes, faces, and number plates, respectively.

You might want to resize these images before encoding them. For example, you might extract keyframes for use in a web application, so that your users can navigate through a video file. In this case you might prefer to have thumbnail-size images rather than the high-definition video frames.

You can use a `scale` transformation task to duplicate an existing track and resize the images in the new track (the input track is not modified). The new output track is named `taskName.Output`, where `taskName` is the name of the transformation task.

The scale transformation engine also scales metadata that refers to the position of an object in the video frame, so that the positions remain correct after the image has been scaled. For example, when you resize images of faces, the metadata that describes the bounding box surrounding the face is also scaled.

To resize images

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Transform]` section, add a new transformation task by setting the `TransformEngineN` parameter. You can give the task any name, for example:

```
[Transform]
TransformEngine0=ScaleKeyframes
```

3. Create a new configuration section to contain the task settings and set the following parameters:

Type	The transformation engine to use. Set this parameter to <code>Scale</code> .
Input	The name of the image track to process.
ImageSize	The output image size in pixels (width followed by height). If you specify one dimension and set the other to <code>0</code> (zero), Media Server preserves the aspect ratio of the original image.

For example:

```
[ScaleKeyFrames]  
Type=Scale  
Input=Keyframe.ResultWithSource  
ImageSize=300,0
```

For more information about the parameters that you can use to configure this task, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Change the Format of Images

Many processing tasks produce images. When Media Server ingests video, it decodes the video into individual frames. Analysis tasks can also produce images, for example keyframe extraction, face detection, and number plate recognition produce images of keyframes, faces, and number plates, respectively.

You might want to change the format of these images before sending them to an output task (output tasks can output base-64 encoded image data). The `ImageFormat` transformation task transforms images in a track into another format. The new output track is named `taskName.Output`, where `taskName` is the name of the transformation task (the input track is not modified).

Tip: You can use the `ImageFormat` transformation engine to change the format of images before sending them to an output engine. When writing image files to disk through the image encoder, use the image encoder `Format` configuration parameter to specify the image format.

To change the format of images

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Transform]` section, add a new transformation task by setting the `TransformEngineN` parameter. You can give the task any name, for example:

```
[Transform]  
TransformEngine0=KeyframesFormat
```

3. Create a new configuration section to contain the task settings and set the following parameters:

Type	The transformation engine to use. Set this parameter to <code>ImageFormat</code> .
Input	The name of the image track to process.
ImageSize	The output format for the images in the new track.

For example:

```
[KeyframesFormat]  
Type=ImageFormat
```

```
Input=Keyframe.ResultWithSource  
Format=jpg
```

For more information about the parameters that you can use to configure this task, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Chapter 11: Encode Video

This section describes how to encode video using Media Server.

- [Introduction](#) 153
- [Encode Video to MPEG Files or a UDP Stream](#) 153
- [Encode Images](#) 154
- [Store Video in a Rolling Buffer](#) 156

Introduction

Media Server can encode the video it ingests and write the video to files, to a stream, or to a rolling buffer.

There are several reasons for encoding video:

- If you are ingesting video from a stream, you can encode the video for playback at a later time. If your analysis tasks detect interesting events in the video, you might want to review those events.
- You can choose the size and bit rate of the encoded video. If you are analyzing sources such as high-definition television broadcasts, you might want to reduce the size and bit rate of the video for storage and streaming to users when they search for clips.

Encoding video does not affect the video frames used for analysis. Analysis always uses the source video.

Media Server provides several audio and video profiles that contain settings for encoding. For example, there are profiles for high-quality output, which you can use if you want to prioritize quality over disk space. By default, the files are stored in the `profiles` folder in the Media Server installation directory. You should not need to modify the profiles. When you configure encoding, you can select the profiles that you want to use.

Encode Video to MPEG Files or a UDP Stream

To encode video to an MPEG file or UDP stream, follow these steps.

To encode video to a file or stream

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Encoding]` section, add a new encoding task by setting the `EncodingEngineN` parameter. You can give the task any name, for example:

```
[Encoding]
EncodingEngine0=MyEncodingTask
```

3. Below the `[Encoding]` section, create a configuration section for the engine by typing the task

name inside square brackets. For example:

```
[MyEncodingTask]
```

4. In the new section, set the following parameters:

Type	The encoding engine type. Set this parameter to <code>mpeg</code> .
VideoSize	(Optional) The size of the encoded video, in pixels, width followed by height. If you are encoding video in MPEG-2 format and you want to use the same dimensions as the source, set this parameter to <code>copy</code> .

To write the encoded video to an MPEG file, set the following parameters:

OutputPath	The path of the encoded output file(s). You can use macros to create output paths based on the information contained in the encoded records.
UrlBase	The base URL that will be used to access the encoded files. This is used when Media Server generates proxies. You can use macros to create the URL base from information contained in the encoded records.
Segment	(Optional) Specifies whether to split the output file into segments.
SegmentDuration	(Optional) The maximum duration of a segment, in seconds.

To generate a live UDP stream of the content that Media Server is ingesting, set the following parameters:

OutputURL	The UDP output stream. For example, <code>udp://239.255.1.123:4321</code> .
Format	The container format. For example, <code>mpegts</code> for MPEG transport stream.

For example:

```
[MyEncodingTask]
Type=mpeg
OutputPath=./mp4/%year%/%month%/%day%/%timestamp%_%segment.sequence%.mp4
UrlBase=http://www.mysite.com/video/clips/mp4/%year%/%month%/%day%
```

For more information about the configuration parameters and macros that you can use, refer to the *Media Server Reference*.

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Encode Images

The Image Encoder engine can save image records that Media Server produces as image files.

To encode images

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Encoding]` section, add a new encoding task by setting the `EncodingEngineN` parameter.

You can give the task any name, for example:

```
[Encoding]  
EncodingEngine0=KeyframeImages
```

3. Below the [Encoding] section, create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[KeyframeImages]
```

4. In the new section, set the following parameters:

Type	The encoding engine to use. Set this parameter to ImageEncoder .
ImageInput	The name of the track that contains the images to encode.
OutputPath	The path and file name for the output files. If the directory does not exist, Media Server creates it. You can use macros to create output paths based on the information contained in a record.
UrlBase	The URL that will be used to access the encoded files.
ImageSize	(Optional) The output image size in pixels, width followed by height. If you do not set this parameter, Media Server uses the original image size. If you specify only one dimension, Media Server calculates the other, maintaining the original aspect ratio. For example, to specify a width of 300 pixels and have Media Server calculate the appropriate height, set this parameter to ImageSize=300x0 . Setting this parameter only modifies the size of the encoded image. The image encoder does not scale any metadata that describes the position of an object in the image. To scale images and position metadata, consider using a scale transformation task .
FrameRateMax	(Optional) The maximum number of images to encode per second of video.

For example:

```
[KeyframeImages]  
Type=ImageEncoder  
ImageInput=myKeyframeEngine.ResultWithSource  
  
OutputPath=c:\output\keyframes\%record.starttime.year%-%record.starttime.month%  
-%record.starttime.day%\%record.starttime.timestamp%.jpg  
  
UrlBase=http://www.mysite.com/keyframes/%record.starttime.year%-%record.startti  
me.month%-%record.starttime.day%/  
ImageSize=192x108
```

For more information about the parameters that you can use to configure the image encoder, refer to the *Media Server Reference*.

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Store Video in a Rolling Buffer

Media Server can save a copy of the video it ingests to a rolling buffer. A *rolling buffer* is a fixed amount of storage where the oldest content is discarded to make space for the latest.

Media Server can write video directly into the rolling buffer, producing an evidential copy that has not been modified in any way. You might require an evidential copy for legal reasons. Media Server can also encode video before writing it to the buffer, creating a non-evidential copy that is optimized for storage and playback.

Note: Evidential mode is not supported when the source video is in MJPEG format.

A rolling buffer is useful for both broadcast monitoring and surveillance applications. For example, in a surveillance application you could configure the rolling buffer with sufficient storage to contain all the video captured by a camera in the last 7 days. If an event occurs you can play the content from the rolling buffer and view the video around the time of the event. If you needed to document the event, you would have 7 days to extract images and video from the rolling buffer before the video is overwritten.

Rolling buffers are configured in a separate configuration file (not in the Media Server configuration file). The rolling buffer configuration file contains settings such as the paths to the storage locations on disk, and the amount of storage to allocate to each rolling buffer.

You can configure as many rolling buffers as you need. For example, you could choose to save an evidential copy of the ingested video to one rolling buffer and a smaller compressed copy to another. You might also want to set up multiple rolling buffers if you ingest video from separate cameras or channels. For example, if you ingest 12 hours of video from one channel and then 48 hours from another you can use multiple rolling buffers to store the last 12 hours from each channel.

Video is served from the rolling buffer using the HTTP Live Streaming (HLS) protocol. An HLS-compliant media player can request an HLS playlist from Media Server or the HPE MMAP REST endpoint. The playlist points to video segments in the rolling buffer, and these segments are served by an external component such as a Web server or HPE MMAP. Media Server does not include a Web server.

Calculate Storage Requirements

When setting up a rolling buffer, consider what you intend to store and determine the amount of storage you need. It is important to allocate sufficient storage space because as soon as the buffer is full, Media Server overwrites the oldest video to make space for the latest. To avoid losing important data, ensure you have sufficient capacity.

The amount of storage you should allocate to a rolling buffer depends on:

- the amount of time for which you want to keep video before it is overwritten.
- the bitrate of the video you encode to the rolling buffer.

For example, if you intend to store video for 1 week, and you encode the video at 4 megabits per second, you would need approximately 350GB of disk space.

After setting up your rolling buffer, HPE recommends that you check that the buffer is storing as much video as you expected.

Set Up Rolling Buffers

The settings for your rolling buffers are stored in a separate configuration file. This is so that you can share the rolling buffer configuration between multiple Media Servers. For example, you might have one Media Server writing video to a rolling buffer and another generating playlists.

A default rolling buffer configuration file is included with Media Server in the `/encoding/rollingBuffer` folder.

In `rollingBuffer.cfg`, you can control the following settings:

- the location of the root folder for each rolling buffer.
- the maximum number and size of the files allocated to each rolling buffer. The number of files multiplied by their size gives the total amount of storage for a rolling buffer. Media Server saves video across multiple files because this is beneficial to disk performance. You can set a default value for all rolling buffers and override it for individual rolling buffers.
- the prefixes added to URLs used in playlists.

Note: Apart from the settings in the `[Defaults]` section, do not edit the rolling buffer configuration file in a text editor. To add or modify rolling buffers, use the ACI actions provided by Media Server.

To set the path to the rolling buffer configuration file

1. Open the Media Server configuration file and find the `[Paths]` section.
2. Ensure that the path to the rolling buffer configuration file is set, for example:

```
[Paths]
RollingBufferConfigPath=./encoding/rollingBuffer/rollingBuffer.cfg
```

Relative paths must be relative to the Media Server working directory. If you share a rolling buffer configuration file between multiple Media Servers, specify a UNC path.

3. Save and close the configuration file.
4. If you made any changes to the configuration file, restart Media Server.

To configure default settings for rolling buffers

1. Open the rolling buffer configuration file.
2. In the `[Defaults]` section, set the following parameters:

<code>RootPath</code>	The path of the directory to store the rolling buffer files in. Relative paths must be relative to the rolling buffer configuration file.
<code>MaxFiles</code>	The maximum number of files for each rolling buffer.
<code>MaxFileSizeMB</code>	The maximum size in MB for each file.
<code>MediaSegmentTemplate</code>	(Optional) A template to use to construct the URL of every media segment in a playlist.
<code>VariantSegmentTemplate</code>	(Optional) A template to use to construct the URL of every <code>GetPlaylist</code> action that is produced by Media Server.

For example:

```
[Defaults]
RootPath=C:\VideoRecording\
MaxFiles=10
MaxFileSizeMB=100
```

For more information about these configuration parameters, refer to the *Media Server Reference*.

To add a new rolling buffer

- Use the ACI action `AddStream`, for example:

```
/action=AddStream&name=NewsChannel&maxfiles=10&maxfilesize=100
```

where the `name` parameter is required and specifies the name of the rolling buffer. The other parameters are optional and override the default rolling buffer settings.

To remove a rolling buffer

Caution: This deletes all video that has been stored in the rolling buffer.

- Use the ACI action `RemoveStream`, for example:

```
/action=RemoveStream&name=NewsChannel
```

where the `name` parameter is required and specifies the name of the rolling buffer to remove.

For more information about the actions that you can use to configure rolling buffers, refer to the *Media Server Reference*.

Pre-Allocate Storage for a Rolling Buffer

Media Server must allocate storage for a rolling buffer before it can write encoded video to the buffer. Allocating storage is not instantaneous so to ensure that Media Server can start recording from a stream immediately, pre-allocate storage before you start a session.

To pre-allocate storage in the rolling buffer

- Send the `PreAllocateStorage` action to Media Server:

```
http://localhost:14000/action=PreAllocateStorage
```

For more information about this action, refer to the *Media Server Reference*.

Write Video to a Rolling Buffer

To write ingested video to a rolling buffer, follow these steps.

To write video to a rolling buffer

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).

2. In the [Encoding] section, add a new encoding task and a configuration section for the task:

```
[Encoding]
EncodingEngine0=RollingBufferEvidential
```

```
[RollingBufferEvidential]
```

3. In the new section, set the following parameters:

Type The type of engine to use. Set this parameter to **RollingBuffer**.

Stream The name of the rolling buffer to write video to. You must have created the rolling buffer (see "[Set Up Rolling Buffers](#)" on page 157).

4. Decide whether to store an evidential copy of the video or encode the video so that it is optimized for storage and playback:

- To store an evidential copy of the video, set the following parameter:

EvidentialMode To write an evidential copy of the video to the rolling buffer, set **EvidentialMode=true**. Media Server ignores any encoding settings defined for this task, but you must also set **AudioInput=none** and **ImageInput=none**.

- To store an encoded (non-evidential) copy of the video, set the following parameters:

AudioProfile The audio encoding profile to use.

VideoProfile The video encoding profile to use.

For example:

```
[Encoding]
EncodingEngine0=RollingBufferEvidential
EncodingEngine1=RollingBufferCompressed
```

```
[RollingBufferEvidential]
Type=RollingBuffer
Stream=evidential
EvidentialMode=true
ImageInput=none
AudioInput=none
```

```
[RollingBufferCompressed]
Type=RollingBuffer
Stream=compressed
AudioProfile=mpeg4audio
VideoProfile=mpeg4video_h264_sd
```

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

View the Rolling Buffer Contents

To retrieve a list of all the rolling buffers that have been configured use the following procedure.

To list your rolling buffers

- Send the `GetStreams` action to Media Server. For example:

```
http://localhost:14000/action=GetStreams
```

Media Server returns a list of all of your rolling buffers.

You can also retrieve information about the video stored in a rolling buffer. For example, you might record video for one hour each morning and two hours each afternoon. To get information about what has been stored in a rolling buffer, use the following procedure.

To return a list of the content in a rolling buffer

- Send the `GetRecordedRanges` action to Media Server.

You can filter the results by setting the following optional parameters:

<code>Stream</code>	The rolling buffer to list of recorded content for. If this parameter is not set, Media Server returns results for all rolling buffers.
<code>StartTime</code>	Only show content available after this time. Specify the start time in epoch milliseconds or ISO-8601 format.
<code>Duration</code>	The length of time after the start time to return a list of available content for. Specify the duration in milliseconds.

For example, to retrieve a list of content available in the `BBCNews` rolling buffer for 24 hours from 16:27:54 on 21 February 2014:

```
http://localhost:14000/action=GetRecordedRanges&Stream=BBCNews
                                     &StartTime=1393000074243
                                     &Duration=86400000
```

Retrieve an HLS Playlist

Media Server generates HTTP Live Streaming (HLS) playlists that can be used by a media player to request content from a rolling buffer. To play video from a rolling buffer, your system must meet the following requirements:

- You must configure a Web server or use the HPE MMAP REST endpoint to serve video segments from your file system to the media player.
- The media player that you use must be HLS-compliant. By default, Media Server generates HLS version 4 playlists, but you can also configure Media Server to generate HLS version 1 playlists.

To retrieve a playlist

- Send the `GetPlaylist` action to Media Server. Set the following parameters:

<code>Stream</code>	The name of the rolling buffer to request video from.
<code>StartTime</code>	(Set this parameter or <code>Offset</code>) The start time of the playlist in ISO 8601 format or epoch milliseconds.
<code>Offset</code>	(Set this parameter or <code>StartTime</code>) Specifies the start time of the playlist by calculating an offset from the current time. For example, if you specify an offset of one hour, the start time for the playlist is one hour ago. Specify the offset in milliseconds.
<code>Duration</code>	(Optional) The length of time after the start time that the playlist covers. Specify the duration in milliseconds.
<code>HLSVersion</code>	(Optional) By default, Media Server generates HLS version 4 playlists. To obtain an HLS version 1 playlist set this parameter to 1.

For example, to retrieve a playlist that contains five minutes of content from the `BBCNews` rolling buffer, starting from 16:27:54 on 21 February 2014:

```
http://localhost:14000/action=GetPlaylist&Stream=BBCNews
&StartTime=2014-02-21T16:27:54Z
&Duration=300000
```

To retrieve a playlist that contains content from the `BBCNews` rolling buffer, starting from 16:27:54 on 21 February 2014, with no end time, use the following action. If you play the content at normal speed and there is no break in recording, the media player could continue playing content forever:

```
http://localhost:14000/action=GetPlaylist&Stream=BBCNews
&StartTime=2014-02-21T16:27:54Z
```

Media Server returns the playlist. If you open the playlist with an HLS-compliant media player, the player will play the video from the rolling buffer.

Create a Clip from a Rolling Buffer

Use the following procedure to retrieve a section of video from the rolling buffer.

To create a clip from a rolling buffer

- Send the `CreateClip` action to Media Server. Set the following parameters:

<code>Stream</code>	The name of the rolling buffer to create the clip from.
<code>StartTime</code>	The start time of the clip in epoch-milliseconds or ISO-8601 format.
<code>Duration</code>	The duration of the clip in milliseconds.
<code>Path</code>	The path to save the clip to. The directory must be on a file system that Media Server can access.

For example,

```
http://localhost:14000/action=CreateClip&Stream=BBCNews
&StartTime=1393000074243
&Duration=300000
&Path=./temp/News1.ts
```

This action instructs Media Server to create a five minute clip from the rolling buffer `BBCNews`, beginning from Fri, 21 Feb 2014 16:27:54 GMT, and to save the clip as the `News1.ts` file in the `temp` directory.

Create an Image from a Rolling Buffer

To obtain a single video frame from a rolling buffer, use the following procedure.

To create an image from a rolling buffer

- Send the `CreateImage` action to Media Server. Set the following parameters:

`Stream` The name of the rolling buffer to create the image from.

`Time` The time that the desired frame occurs in the rolling buffer. Specify the time in epoch-milliseconds or ISO-8601 format.

For example,

```
http://localhost:14000/action=CreateImage&Stream=BBCNews
&Time=1393000074243
```

Use Multiple Media Servers

It is possible to configure multiple Media Servers to use the same rolling buffer configuration file. You can have a maximum of one Media Server writing video into a rolling buffer, but other Media Servers can read content from the rolling buffer.

If you configure multiple Media Servers to use the same rolling buffer configuration file, you must:

- grant read access to the rolling buffer configuration file to all of the Media Servers.
- grant write access to the rolling buffer configuration file to only one Media Server. If you need to change the settings for your rolling buffers, you must send actions to this Media Server.
- allow only one Media Server to write video into each rolling buffer. The `Stream` parameter in a rolling buffer encoding task specifies which rolling buffer to write video to.

Chapter 12: Output Data

Media Server can send the metadata it produces to many systems including IDOL Server, CFS, HPE Broadcast Monitoring, Vertica and Milestone XProtect. Media Server can also write metadata to XML files on disk.

This section describes how to configure Media Server to output data.

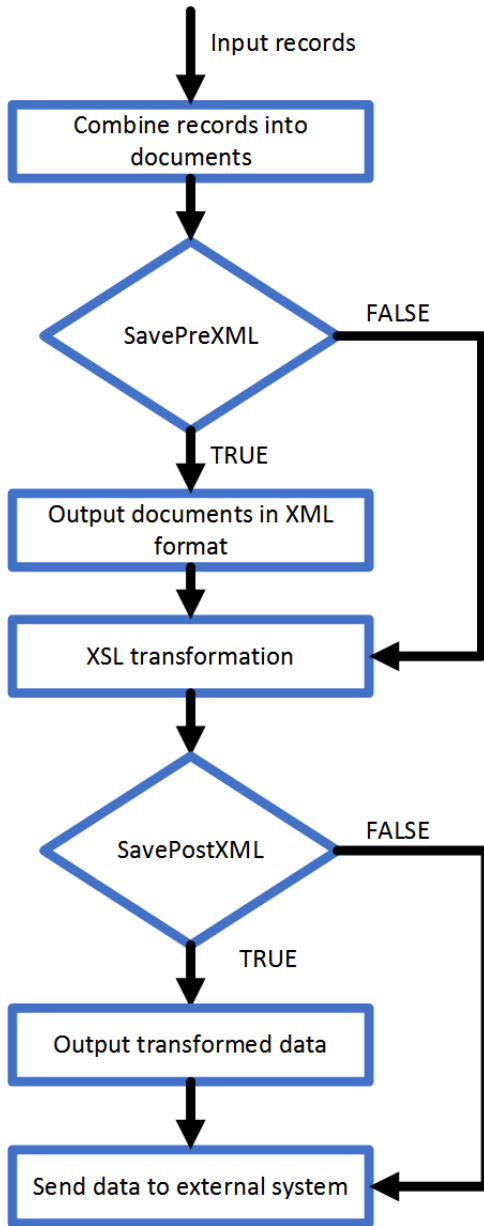
- [Introduction](#) 163
- [Choose How to Output Data](#)165
- [Output Data to Files](#) 172
- [Output Data to the Process Action Response](#) 174
- [Send Information over HTTP](#) 175
- [Index Documents into IDOL Server](#)176
- [Insert Data into a Vertica Database](#) 178
- [Insert Records into a Database](#) 180
- [Send Documents to Connector Framework Server](#) 183
- [Index Records into Broadcast Monitoring](#) 185
- [Send Data to Milestone XProtect](#) 187

Introduction

Media Server output tasks export the metadata produced by Media Server. This can include information about the ingested video, information about copies of the video that you create through [encoding](#), and any information extracted from the video during [analysis](#).

Tip: Output tasks do not output video. For information about saving video, see "[Encode Video](#)" on [page 153](#)

The following diagram shows the steps that occur when you configure Media Server to output data.



Select Input Records

An output task receives records that are produced by your ingest, analysis, encoding, and ESP tasks. You can choose the information to output by setting the `Input` configuration parameter in the output task. If you do not set this parameter, the output task receives records from all tracks that are considered by default as 'output' tracks. For information about whether a track is considered by default to be an 'output' track, refer to *Media Server Reference*.

Combine Records into Documents

An output task receives individual records, but you might want to combine the information from many records and index that information as a single document. For example, if you are processing a news broadcast, you might want to index a document for each news story, rather than a document for each word spoken in the audio and a document for each recognized face. To do this, the Media Server must combine records representing recognized faces, speech-to-text results, recognized objects, and so on.

Most output engines have several indexing modes so that you can configure how to create documents. For more information about these indexing modes, see ["Choose How to Output Data" below](#).

XSL Transformation

The output task performs an XSL transformation to convert the combined records into a format that is suitable for the destination repository.

Media Server is supplied with XSL templates to transform data into IDOL documents, Broadcast Monitoring assets, and other formats. You can modify the default templates to suit your needs. Set the `XSLTemplate` configuration parameter in the output task to specify the path of the XSL template to use.

If you want to customize the XSL template, set the configuration parameters `SavePreXML=TRUE`, `SavePostXML=TRUE`, and `XMLOutputPath` so that Media Server writes documents to disk before and after performing the transformation. Viewing the data before and after transformation might help you optimize your XSL template.

Send the Data to the External System

After performing the XSL transformation, Media Server sends the data to its destination.

Choose How to Output Data

Media Server can output data in many formats, including:

- XML
- documents that you can send to CFS or index into IDOL Server
- assets for indexing into HPE Broadcast Monitoring
- rows of data suitable for inserting into a database

Media Server analysis engines produce records, each of which describes something that happens in a video. A record might describe a recognized face, a scene change, or a word spoken in the audio. When you index data into some systems, such as IDOL Server or Broadcast Monitoring, you might want to combine the data from many records to produce documents that represent video segments or clips.

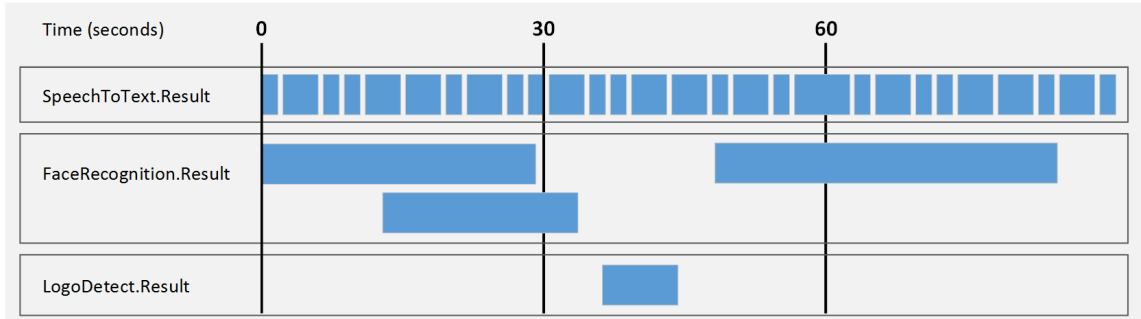
For example, the speech-to-text analysis engine produces a record for each word in the audio. If you are indexing data into a SQL database, you could insert a separate row into a database table for each word. If you are indexing data into IDOL Server, you might prefer to combine records so that each document contains the transcription from an entire news story or interview.

The following sections describe the indexing modes that you can choose when you configure Media Server to output data.

Single Record Mode

Single record mode creates documents that each represent a single record.

Consider the following records. In single record mode, Media Server creates a separate document for every record. No document contains more than one record.



This mode is suitable when you:

- send data to a database, where each record becomes a row in the database.
- need to index documents that represent discrete records, for example a recognized face or an ANPR result.
- need to use the metadata produced by Media Server in a front end application in real-time. In single record mode, Media Server outputs information about a record as soon as the record has finished. It does not need to hold records so that they can be combined with other related records.

Time Mode

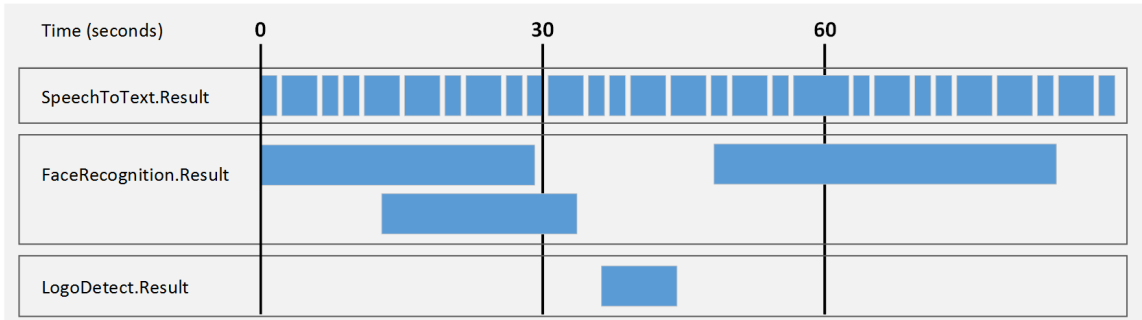
Time mode creates documents that represent a fixed amount of video content. A document contains all of the records that occurred during the time interval, or that overlap the start or end of the interval. You can define the duration of a document, for example 30 seconds. The duration is measured using video time, so you do not need to modify the duration if the **ingest rate** is slower or faster than normal (playback) speed.

Tip: Time-based output does not mean that Media Server outputs data at regular intervals. Media Server cannot output data for a video segment until all records in that segment have finished. As a result, Media Server might produce several documents, which represent consecutive time periods, at the same time.

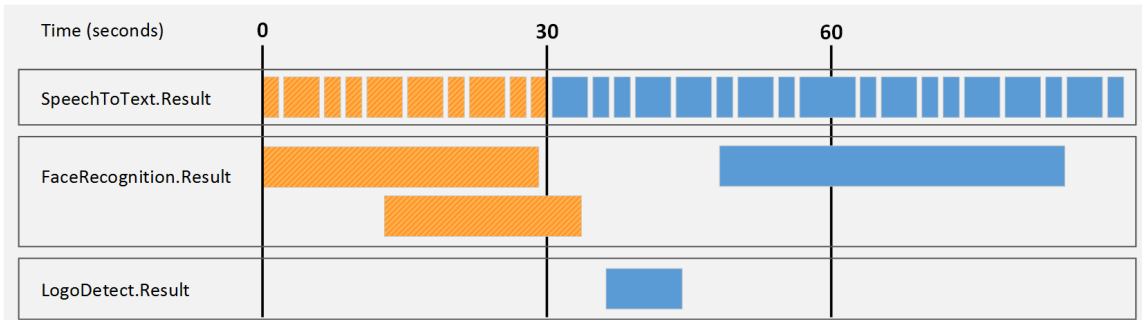
For example, if you run logo detection on a news broadcast, the news logo could be present on screen continuously for an hour. In that case, Media Server does not output any data until the logo disappears. Although Media Server creates documents that each represent 30 seconds of content, all of the documents are output at the end of the hour when the record describing the logo finishes.

In time mode, all records are output to at least one document. If a record spans more than one interval it is output to multiple documents.

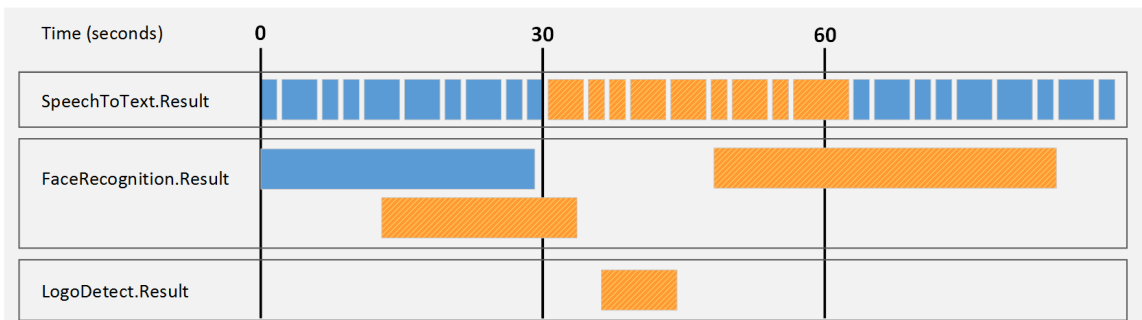
The following diagrams demonstrate how Media Server constructs documents in time mode. Consider the following records:



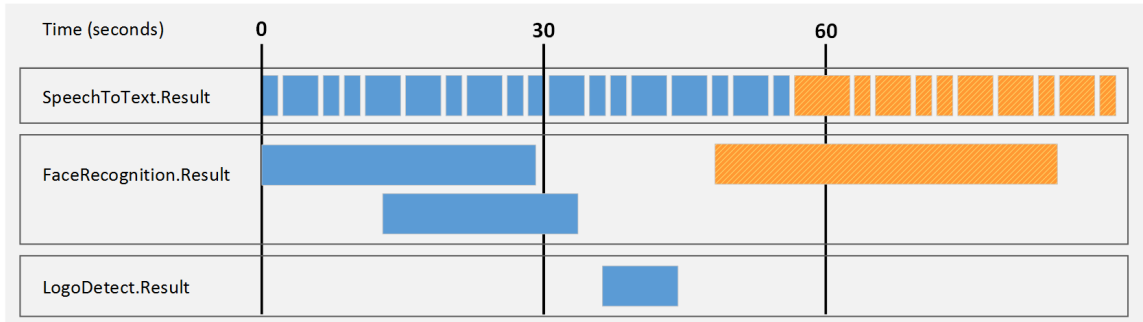
The first document contains the records related to the first interval (in this case, 30 seconds) of video content. Media Server does not output data at the 30-second point because a face recognition record has not ended (the face is still present in the video). Media Server can output data about the first interval as soon as this record ends:



The second document contains the records related to the second interval (in this case, 30 seconds) of video content. Notice that the second face recognition result is output in the second document as well, because it relates to both time intervals.



The third document contains the records related to the third interval (in this case, 30 seconds) of video content:



Time mode is simple to set up but documents might begin and end in the middle of a sentence or segment, rather than at meaningful point in the video.

Event Mode

Event mode helps to create documents that contain information about a single topic, which can provide a better experience for your users when they retrieve the content, and improves the performance of IDOL operations such as categorization.

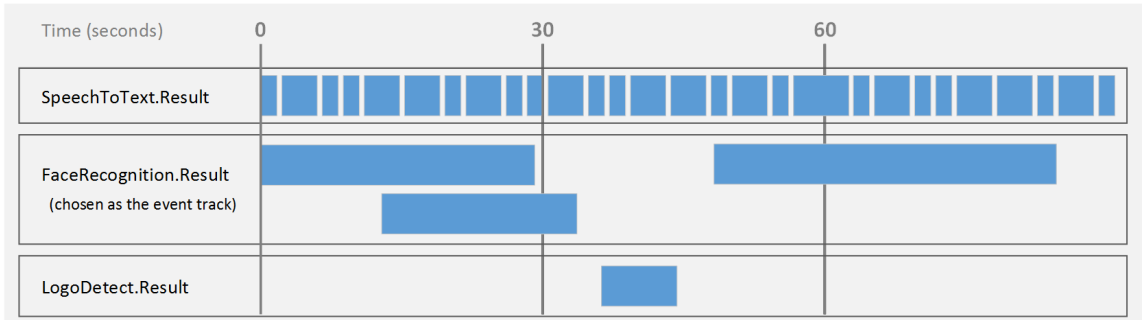
In event mode, Media Server creates a document for each record in an *event track*. The document that is generated contains the event from the event track. Other records are included if they overlap with the event in the event track or if they occurred after the end of the previous event. Each document might represent a different amount of video content.

The event track can be any track that you choose, though often it will be a track that you create using [event stream processing](#). The event track could contain a record whenever there is a scene change, or whenever there is a pause in speech. For example, if you are analyzing news content Media Server can start a new document whenever there is a pause in speech, which could indicate the start of a new story.

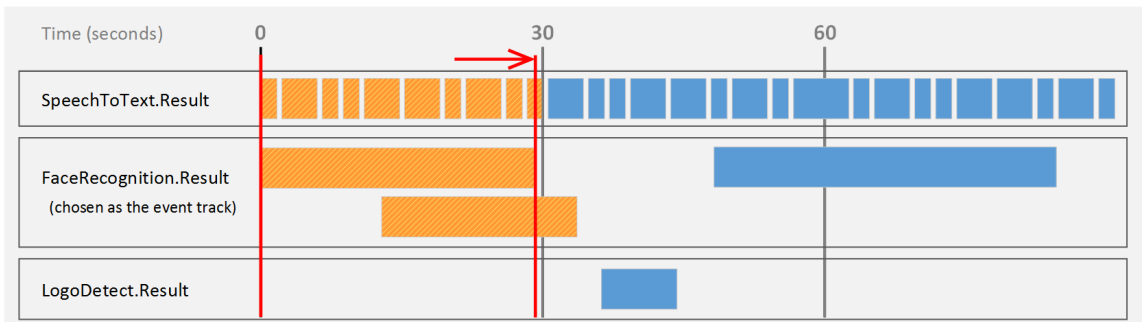
In event mode, all of the records in the selected tracks are output to at least one document. Compare this with ["Bounded Event Mode" on page 170](#), in which some records can be omitted.

Note: Be aware that if you choose an event track that has overlapping records (for example the result track from a face recognition task), the resulting documents might contain more than one record from the event track, and some records will be output to multiple documents.

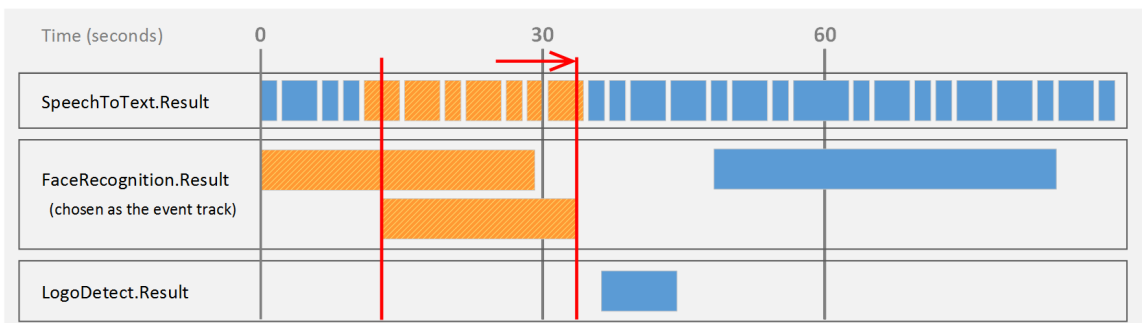
The following diagrams demonstrate how Media Server constructs documents in **event** mode. The FaceRecognition.Result track has been chosen as the event track. Consider the following records:



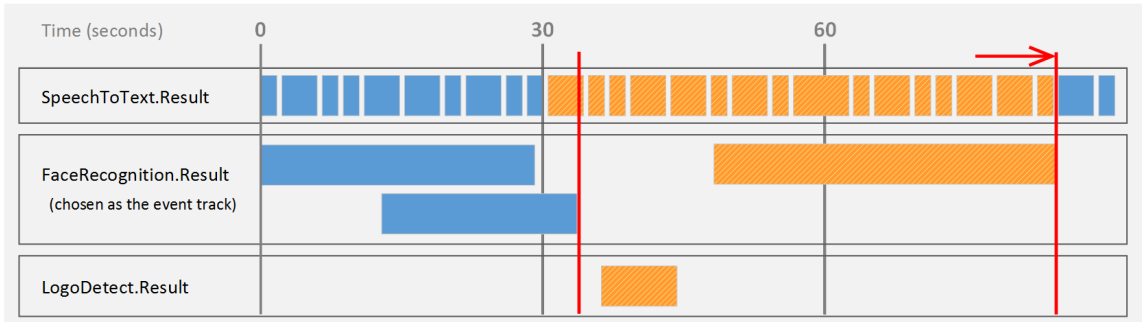
The first document contains the first record in the event track (the first event). Other records are included if they overlap with this event or if they occurred since the end of the previous event. In this case the previous event is the beginning of the video:



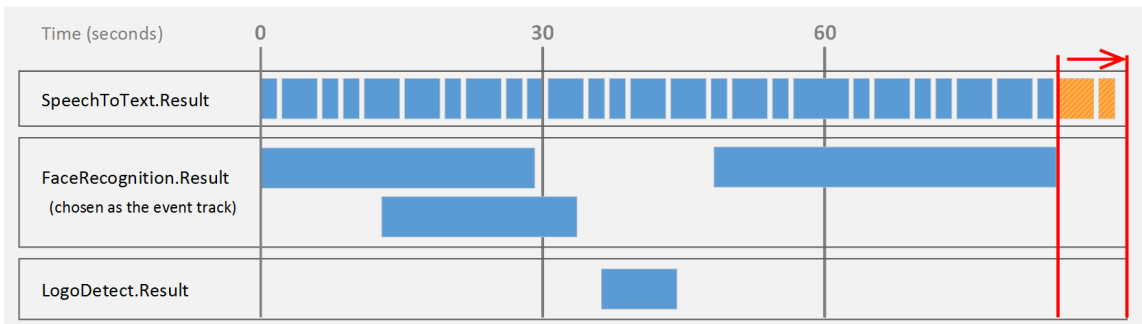
The second document contains the second record in the event track (the second event). Other records are included if they overlap with this event or if they occurred since the end of the previous event:



The third document contains the third record in the event track (the third event). Other records are included if they overlap with this event or if they occurred since the end of the previous event:



The end of the video is considered as an event, so in this case a final document is produced containing the final records in the speech-to-text result track:



Bounded Event Mode

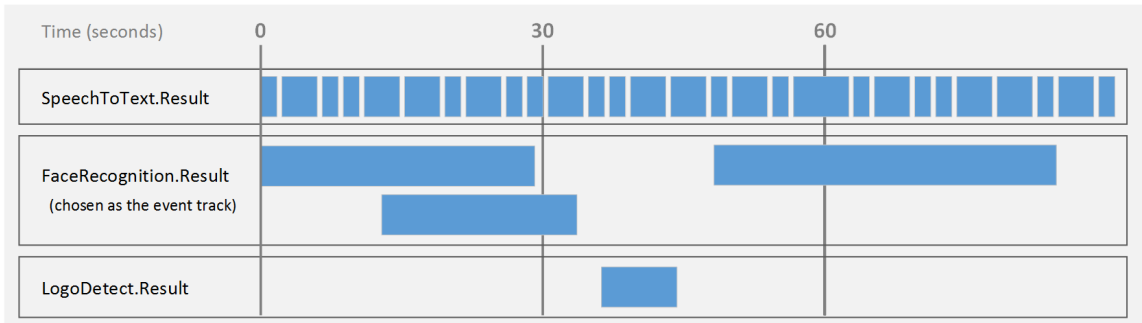
Bounded event mode, like [event mode](#), helps to create documents that contain information about a single topic. Media Server creates a document for each record in an *event track*. The document that is generated contains the event from the event track. However, unlike [event mode](#), other records are included only if they overlap with the event in the event track.

Note: In bounded event mode, it is possible that some records are not output to documents.

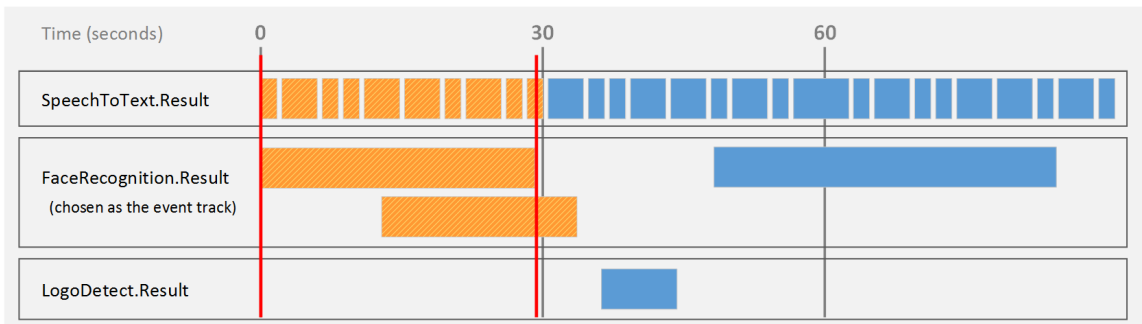
The event track can be any track that you choose, though often it will be a track that you create using [event stream processing](#). You could use speaker identification results as your event track, so that a document is created for each speaker detected in the audio.

Each document might represent a different amount of video content.

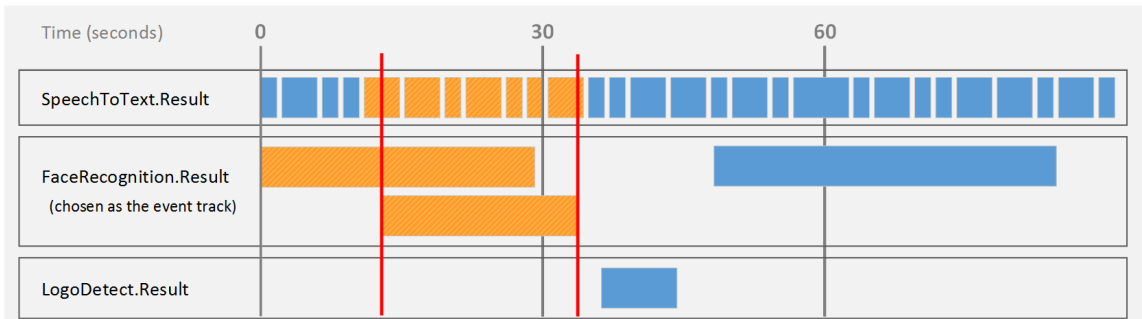
The following diagrams demonstrate how Media Server constructs documents in **bounded event** mode. The `FaceRecognition.Result` track has been chosen as the event track. Consider the following records:



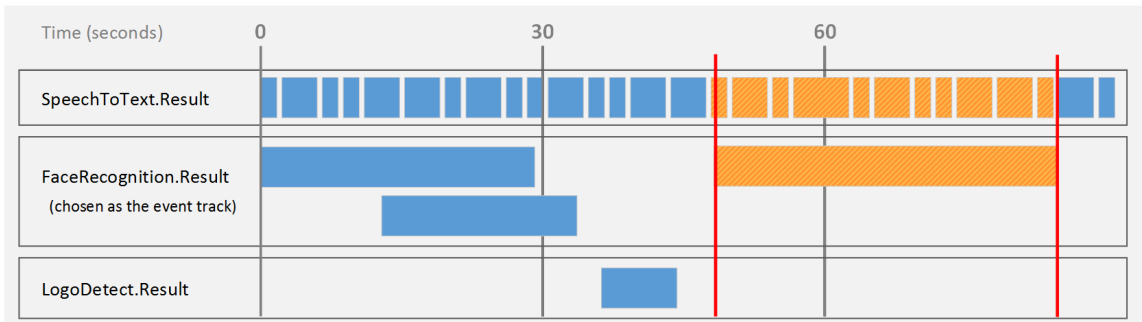
The first document contains the first record that occurs in the event track, and all of the records that occur at the same time:



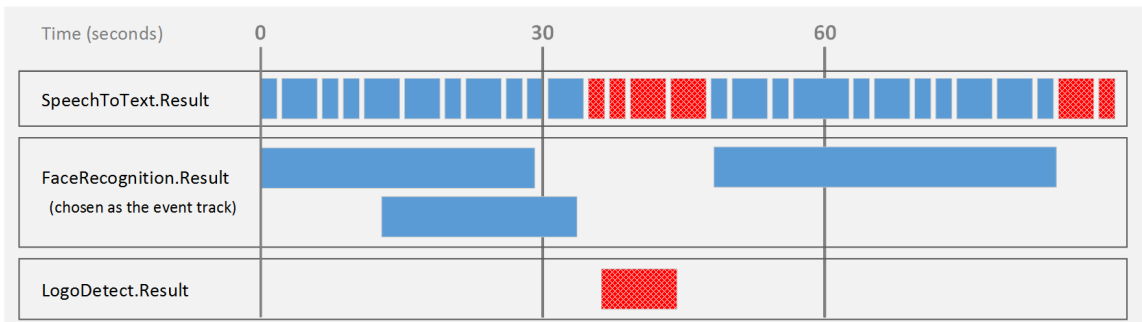
The second document contains the second record that occurs in the event track, and those records that overlap with it. Notice that if records in the event track overlap, the output document can contain more than one record from the event track:



The next document contains the next record from the event track, and the records that overlap with it:



In bounded event mode, the records shown here are not output to any document, because they do not overlap with a record in the event track:



Output Data to Files

Media Server can output records to files, so that you can index the information into any system that accepts data in a format such as XML.

To write records to files, use the XML output engine. The default format is XML but you can configure the engine to apply an XSL transformation to the output, to transform it into another format such as HTML.

To write records to files

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Output]` section, add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=XmlWriter
```

3. Below the `[Output]` section, create a new configuration section by typing the task name inside square brackets. For example:

```
[XmlWriter]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to XML .
Input	(Optional) A comma-separated list of the tracks that you want to output. Specify one or more tracks. If you do not set this parameter, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis and encoding engines in the <i>Media Server Reference</i> .
XmlOutputPath	The path and file name of the XML files to create.
XSLTemplate	(Optional) The path to the XSL template to use to transform the output into the desired format. If you do not set this parameter, the output is not transformed.

For more information about the parameters that you can use to customize an XML output task, refer to the *Media Server Reference*.

5. Configure how to combine records into XML files. For information about how you can combine records, see ["Choose How to Output Data" on page 165](#).
 - To output an XML file for every record, set the following parameters:

Mode	The output mode. Set this parameter to SingleRecord .
------	--
 - To output XML files that represent a fixed amount of video content, set the following parameters:

Mode	The output mode. Set this parameter to Time .
OutputInterval	The amount of video content represented by each output document.
 - To output XML files that represent video segments delimited by events, set the following parameters:

Mode	The output mode. Set this parameter to Event .
EventTrack	The name of the track to use as the event track.
 - To output XML files that contain information about events and any overlapping records, but not other records that occurred since the previous event, set the following parameters:

Mode	The output mode. Set this parameter to Bounded .
EventTrack	The name of the track to use as the event track.
6. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following is an example configuration to output data using the XML output engine.

```
[XmlWriter]  
Type=xml
```

```
XMLOutputPath=./output/html/%segment.type%_
results_%timestamp%_%segment.sequence%.html
XSLTemplate=./output/toHTML.xsl
Mode=Time
OutputInterval=30
```

Output Data to the Process Action Response

Media Server does not output records to the process action response by default.

You can configure Media Server to do this so that another server, such as a Connector Framework Server, can send requests to Media Server for analysis and then retrieve the results from the action response.

Tip: HPE recommends that you do not write results to the action response in cases where Media Server produces a large amount of data. If you output data from tracks that contain a large amount of metadata, the ACI response could become extremely large. This might result in the system running out of resources, or external applications failing to retrieve the response.

To write records to the action response

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the [Output] section, add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=Response
```

3. Create a configuration section for the task and set the following parameters:

Type	The output engine to use. Set this parameter to <code>response</code> .
Input	(Optional) A comma-separated list of the tracks that you want to output. Specify one or more tracks. If you do not set this parameter, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis and encoding engines in the <i>Media Server Reference</i> .

For more information about these parameters, refer to the *Media Server Reference*.

4. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following is an example configuration to output data to the process action response.

```
[Response]
Type=response
Input=OCR.Result,Keyframe.Result
```

Send Information over HTTP

Media Server includes an output engine for sending data to a server through an HTTP POST request.

To send information over HTTP

1. Create a new configuration or open an existing configuration to send to Media Server with the `process` action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. Add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=HTTPpost
```

3. Create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[HTTPpost]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to <code>httppost</code> .
DestinationURL	The URL to send the information to.
XSLTemplate	The path to the XSL template to use to transform records and produce the body of the request.
Input	(Optional) A comma-separated list of the tracks that you want to include in the output. Specify one or more tracks. If you do not set this parameter, the engine includes all tracks that are preset as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis, encoding, and ESP engines in the <i>Media Server Reference</i> .

5. Configure how to combine records into HTTP requests. For information about how you can combine records, see "[Choose How to Output Data](#)" on page 165.

- To send an HTTP POST request for every record, set the following parameter:

Mode	The output mode. Set this parameter to <code>SingleRecord</code> .
------	--

- To send HTTP POST requests that represent a fixed amount of video content, set the following parameters:

Mode	The output mode. Set this parameter to <code>Time</code> .
------	--

OutputInterval	The amount of video content represented by each request.
----------------	--

- To send an HTTP POST request for each video segment, set the following parameters:

Mode	The output mode. Set this parameter to Event .
EventTrack	The name of the track to use as the event track.

- To send HTTP POST requests that contain information about events and any overlapping records, but not other records that occurred since the previous event, set the following parameters:

Mode	The output mode. Set this parameter to Bounded .
EventTrack	The name of the track to use as the event track.

6. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following is an example configuration to send documents to a CFS over HTTP.

```
[HTTPpost]
Type=httppost
DestinationURL=http://localhost:7000/action=ingest
XSLTemplate=./toCFS.xsl
Mode=event
OutputInterval=NewsSegment.Output
```

Index Documents into IDOL Server

Media Server's IDOL output engine transforms metadata produced by Media Server into IDOL documents and indexes the documents into an IDOL Server.

The IDOL output engine uses an XSL template to transform records into IDX files. To modify the structure of the IDX file, modify the XSL template. For more information about indexing content into IDOL Server, refer to the *IDOL Server Administration Guide*.

To index information into IDOL Server

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Output]` section, add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=IDOL
```

3. Below the `[Output]` section, create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[IDOL]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to IDOL .
Input	(Optional) A comma-separated list of the tracks that you want to output. Specify one or more tracks. If you do not set this parameter, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis, encoding, and ESP engines in the <i>Media Server Reference</i> .
IdolHost	(Optional) The host name or IP address of the IDOL Server. This overrides the IDOL Server host specified by the <code>IdolServer</code> parameter in the [Resources] section, if it has been set.
IdolPort	(Optional) The ACI port of the IDOL Server (by default, 9000). This overrides the port specified by the <code>IdolServer</code> parameter in the [Resources] section, if it has been set.
IdolDB	(Optional) The IDOL database to index documents into. This overrides the database specified by the <code>IdolServer</code> parameter in the [Resources] section, if it has been set. If you do not set this parameter, documents are indexed into the database specified by their <code>DREDBNAME</code> metadata field. You can modify your XSL template to create this field. If a document does not specify a database it is indexed into the default database specified by the <code>DefaultDatabase</code> parameter, in the [Databases] section of the IDOL Server configuration file.
IDOLParams	(Optional) Additional IDOL index action parameters. The IDOL output engine sends the <code>DREADDDATA</code> action to IDOL Server, which instructs the server to index the data contained in the request. You can send additional parameters with this action. For information about the available index action parameters, refer to the <i>IDOL Server Reference</i> .
XSLTemplate	The path to the XSL template to use to transform records into documents in IDX format. You can modify the default XSL template as required - for example to produce XML rather than IDX files.
SavePostXML	(Optional) Specifies whether to save IDX files produced by the IDOL output engine. If this parameter is set to <code>true</code> , you must also set the <code>XMLOutputPath</code> parameter.
SavePreXML	(Optional) Specifies whether to save records received by the IDOL output engine. If this parameter is set to <code>true</code> , you must also set the <code>XMLOutputPath</code> parameter.
XMLOutputPath	(Optional) The path and file name of the file to save pre-XML and post-XML output to.

5. Configure how to combine records into documents. For information about how you can combine records, see ["Choose How to Output Data" on page 165](#).

- To output a separate document for each record, set the following parameters:

Mode The output mode. Set this parameter to **SingleRecord**.

- To output documents that represent a fixed amount of video content, set the following parameters:

Mode The output mode. Set this parameter to **Time**.

OutputInterval The amount of video content represented by each output document.

- To output documents that represent video segments delimited by events, set the following parameters:

Mode The output mode. Set this parameter to **Event**.

EventTrack The name of the track to use as the event track.

- To output XML files that contain information about events and any overlapping records, but not other records that occurred since the previous event, set the following parameters:

Mode The output mode. Set this parameter to **Bounded**.

EventTrack The name of the track to use as the event track.

6. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following is an example configuration to output data using the IDOL output engine.

```
[IDOLOutput]
Type=IDOL
IdolHost=localhost
IdolPort=9000
IdolDB=BroadcastVideo
Mode=Time
OutputInterval=30
XslTemplate=./toIDX.xsl
SavePreXML=true
SavePostXML=true
XMLOutputPath=./Output/%segment.type%_%segment.sequence%_%segment.timestamp%.xml
```

Insert Data into a Vertica Database

To insert records into a Vertica database, use the Vertica output engine.

The Vertica output engine uses an XSL template to transform the XML produced by Media Server into a format, such as a CSV file, that can be inserted into the database. It then connects to the database using ODBC and inserts the information using a `COPY` query:

```
COPY <table>  
FROM LOCAL '<local_file>  
DELIMITER '<delimiter>  
ENCLOSED BY '<quote>  
ESCAPE AS '<escape>
```

where:

<table> is the Vertica database table to copy data into. This is read from the TrackMapping configuration parameter.

<delimiter>, <quote>, and <escape> are replaced by values from the corresponding configuration parameters.

To insert records into a Vertica database

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (mediaserver.cfg).
2. In the [Output] section, add a new output task by setting the OutputEngineN parameter. You can give the task any name, for example:

```
[Output]  
OutputEngine0=VerticaOutput
```

3. Below the [Output] section, create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[VerticaOutput]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to <code>vertica</code> .
TrackMapping	The tracks that you want to output, mapped to Vertica database tables.
OdbcConnectionString	The ODBC connection string to use to connect to the database. For information about how to connect to a Vertica database, refer to the Vertica documentation.
OdbcDriverManager	(Required only on UNIX platforms) The path of the ODBC driver manager to use.
XMLOutputPath	The path to the directory to use for temporary files and saved output.
XSLTemplate	The XSL template to use to transform records from analysis engines to a format that can be inserted into the database (such as a CSV file).
OutputInterval	(Optional) The interval, in seconds, between inserting batches of records into the database. The default interval is 60 seconds.

For example:

```
[VerticaOutput]
Type=vertica
TrackMapping0=FaceRecog.Result : face_recognition
TrackMapping1=Ocr.Result : ocr
OdbcConnectionString=DSN=mydb
OdbcDriverManager=libodbc.so
XMLOutputPath=./tmp
XSLTemplate=to_vertica.xsl
OutputInterval=120
```

For more information about the parameters that you can set to configure a Vertica output task, refer to the *Media Server Reference*.

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Insert Records into a Database

You can use the ODBC output engine to insert information into a database through ODBC.

The ODBC output engine uses an XSL template to extract values from records. The values are inserted into the database by running the query:

```
INSERT INTO tablename(columns) VALUES values
```

where

- *tablename* is the table name you set using the parameter `TableName`.
- *columns* (optional) are the columns specified by the parameter `Columns`.
- *values* are generated by the XSL template specified by the parameter `XSLTemplate`.

Your XSL template must produce a file of values where:

- each record is on a new line
- the values to insert into the database are tab-separated
- the values are in the order you set in the `Columns` parameter.

Before You Begin

If you are running Media Server on Linux, ODBC connector drivers for your database might be included in the operating system distribution, or be available from a package manager. It is likely that later versions of the connector driver will be available for download directly from the database provider. The later drivers might contain stability and performance improvements. If you experience issues using the ODBC output engine, HPE recommends downloading the latest ODBC connector drivers for your database as the first step in the troubleshooting process.

If you have configured the ODBC output engine to output data into a table or column that has an extended Unicode character (a character that is not included in the ASCII character set) in its name, then you must use a Unicode ODBC driver. These drivers are often identified by a "w" being appended to the driver name.

Configure the Output Task

To configure an output task to insert information into a database, follow these steps.

To insert information into a database

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Output]` section, add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=MyDatabase
```

3. Below the `[Output]` section, create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[MyDatabase]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to ODBC .
Input	(Optional) A comma-separated list of the tracks that you want to output. Specify one or more tracks. If you do not set this parameter, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis, encoding, and ESP engines in the <i>Media Server Reference</i> .
OdbcConnectionString	The ODBC connection string to use to connect to the database.
OdbcDriverManager	The path of the ODBC driver manager shared library. This parameter is not required if you are running Media Server on Windows.
TableName	The name of the table to insert information into.
Columns	A comma-separated list of columns to insert data into.
ColumnTypes	A comma-separated list of data types for the columns specified by the <code>Columns</code> parameter. This parameter must have the same number of values that you set for <code>Columns</code> . You can use the following data types: <ul style="list-style-type: none">• <code>string</code>• <code>int</code>

- bigint
- double
- blob

Note: The ODBC output engine inserts images into the database as binary large objects (BLOBs).

XSLTemplate

The path to the XSL template to use to extract values from records and transform them into a list of values to insert into the database.

5. Configure how to combine records into SQL queries. For more information about how you can combine records, see ["Choose How to Output Data" on page 165](#).

- To create a separate query for each record, set the following parameter:

Mode The output mode. Set this parameter to `SingleRecord`.

- To create a separate query for a fixed amount of video content, set the following parameters:

Mode The output mode. Set this parameter to `Time`.

OutputInterval The amount of video content represented by each query.

- To create a separate query for each video segment, set the following parameters:

Mode The output mode. Set this parameter to `Event`.

EventTrack The name of the track to use as the event track.

- To create a query for each event and any overlapping records, but without including other records that occurred since the previous event, set the following parameters:

Mode The output mode. Set this parameter to `Bounded`.

EventTrack The name of the track to use as the event track.

6. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following is an example configuration to output data using the ODBC output engine.

```
[MyDatabase]
Type=odbc
OdbcConnectionString=DSN=MyDSN
TableName=MyTable
Columns=RecordStartTime,RecordDuration,SpeechToTextResult,Confidence
ColumnTypes=bigint,bigint,string,double
XslTemplate=./toODBC.xsl
```

```
SavePreXML=true  
SavePostXML=true  
XMLOutputPath=./output/odbc/myDatabase_%segment.type%.xml  
Mode=singlerecord
```

Troubleshooting

Information is not inserted into the database.

If information is not inserted into the database, ensure that Media Server can connect to the database. If Media Server cannot send information to the database, the information is saved to an XML file named as follows:

- If XMLOutputPath is defined and contains the %segment.type% macro, the segment type is replaced with failed and the other macros are resolved as usual.
- If XMLOutputPath is defined and does not contain the %segment.type% macro, the prefix failed_ is added to the file name, and other macros are resolved as usual.
- If XMLOutputPath is not defined, the file is saved to ./failed/ODBC/failed_counter, where counter is a number that starts from 0 (zero) and increases each time queries are generated by the output task.

When the connection to the database is reestablished, Media Server continues inserting information, but does not insert the records that were saved to disk. You must insert these into the database manually.

Send Documents to Connector Framework Server

Media Server includes an output engine for indexing documents into IDOL Server. However, if you want to manipulate and enrich documents before they are indexed into IDOL you can send the documents to a Connector Framework Server (CFS) instead.

For example, if you have used speech-to-text analysis to convert the words spoken in a video into text, you might want to run Education on the text to extract the names of people or places and add these to the document metadata.

To send documents to CFS, use the CFS output engine. Media Server includes an XSL template that you can use to transform the metadata produced by Media Server into documents.

For more information about using CFS, refer to the *Connector Framework Server Administration Guide*.

To send documents to CFS

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (mediaserver.cfg).
2. Add a new output task by setting the OutputEngineN parameter. You can give the task any name, for example:

```
[Output]  
OutputEngine0=CFS
```

3. Create a section to configure the task by typing the task name inside square brackets. For example:

[CFS]

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to CFS.
CfsHost	The host name or IP address of your CFS.
CfsPort	The ACI port of your CFS.
XslTemplate	The path to the XSL template to use to transform records into documents. Media Server includes an XSL template for sending documents to CFS. This template is named <code>toCFS.xsl</code> and is in the Media Server installation directory.
Input	(Optional) A comma-separated list of the tracks that you want to output. Specify one or more tracks. If you do not set this parameter, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis, encoding, and ESP engines in the <i>Media Server Reference</i> .

5. Configure how to combine records into documents. For information about how you can combine records, see ["Choose How to Output Data" on page 165](#).

- To output a separate document for each record, set the following parameters:

Mode	The output mode. Set this parameter to <code>SingleRecord</code> .
------	--

- To output documents that represent a fixed amount of video content, set the following parameters:

Mode	The output mode. Set this parameter to <code>Time</code> .
------	--

OutputInterval	The amount of video content represented by each output document.
----------------	--

- To output documents that represent video segments delimited by events, set the following parameters:

Mode	The output mode. Set this parameter to <code>Event</code> .
------	---

EventTrack	The name of the track to use as the event track.
------------	--

- To output documents that contain information about events and any overlapping records, but not other records that occurred since the previous event, set the following parameters:

Mode	The output mode. Set this parameter to <code>Bounded</code> .
------	---

EventTrack	The name of the track to use as the event track.
------------	--

6. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Example

The following is an example configuration to send documents to CFS.

```
[CFS]
Type=CFS
CfsHost=localhost
CfsPort=7000
XSLTemplate=./toCFS.xsl
Mode=time
OutputInterval=20
```

Index Records into Broadcast Monitoring

The Broadcast Monitoring output engine transforms information produced by Media Server and indexes it into Broadcast Monitoring.

Media Server segments data into documents, which each represent a fixed amount of video content. You can specify how many documents to index into each asset.

To index information into Broadcast Monitoring

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Output]` section, add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=BroadcastMonitoring
```

3. Below the `[Output]` section, create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[BroadcastMonitoring]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to <code>BM</code> .
BroadcastMonitoringTrickleURL	The Broadcast Monitoring URL to send data to. Setting this parameter overrides the setting specified in the <code>[Resources]</code> section, if <code>BroadcastMonitoringTrickleURL</code> has been set there.
Input	(Optional) A comma-separated list of the tracks that you want to output. Specify one or more tracks. If you do not set this parameter, the engine outputs all tracks that are configured by default as 'output' tracks. For information about whether a track is considered an 'output' track, refer to documentation for your analysis, encoding, and ESP engines in the <i>Media Server Reference</i> .

OutputInterval	The amount of video content represented by each output document, in seconds.
IntervalsPerAsset	The number of output documents to index into an asset in Broadcast Monitoring. After this number of intervals, Media Server starts a new asset. For example, if you set OutputInterval=30 and IntervalsPerAsset=10, Media Server sends documents to Broadcast Monitoring that represent 30 seconds of video content. After 10 of these documents have been sent to Broadcast Monitoring and the asset represents 5 minutes of video, Media Server starts a new asset.
XSLTemplate	The path to the XSL template to use to transform the output so that it can be indexed into Broadcast Monitoring.
LabelSetBroadcaster	The name of the video broadcaster, to add to the "Broadcaster" field in the Broadcast Monitoring label set.
LabelSetCompleted	The value to use for the "Completed" field in the Broadcast Monitoring label set (true/false).
LabelSetProgram	The program name, to add to the "Program" field in the Broadcast Monitoring label set.
DefaultProxyTrack	The proxy track to output to Broadcast Monitoring as the default proxy track.
ProxyTracks	(Optional) A comma-separated list of additional proxy tracks to output to Broadcast Monitoring.
SavePostXML	(Optional) Specifies whether to save XML files produced by the Broadcast Monitoring engine. If this parameter is set to true, you must also set the XMLOutputPath parameter.
SavePreXML	(Optional) Specifies whether to save records received by the Broadcast Monitoring engine. If this parameter is set to true, you must also set the XMLOutputPath parameter.
XMLOutputPath	(Optional) The path and file name of the file to save pre-XML and post-XML output to.

- (Optional) You can set further optional parameters. For information about the configuration parameters that you can use with a Broadcast Monitoring output task, refer to the *Media Server Reference*.
- Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Send Data to Milestone XProtect

Media Server's Milestone output engine sends data to Milestone XProtect Corporate and XProtect Enterprise surveillance systems.

The Milestone Smart Client displays the events detected by Media Server, by showing information overlaid on the video. For example, detected objects are identified by bounding polygons. The Smart Client also shows metadata produced by Media Server, such as a message, event type, and location.

Before You Begin

To use Media Server with a Milestone XProtect surveillance system, HPE recommends using:

- Milestone XProtect Enterprise 8.1a (96231)
- Milestone XProtect Corporate 2013 R2

Note: Other versions might work but have not been tested.

Configure Media Server

To send data to Milestone XProtect, follow these steps.

To send data to Milestone XProtect

1. Create a new configuration or open an existing configuration to send to Media Server with the process action. Alternatively, you can modify the Media Server configuration file (`mediaserver.cfg`).
2. In the `[Output]` section, add a new output task by setting the `OutputEngineN` parameter. You can give the task any name, for example:

```
[Output]
OutputEngine0=XProtect
```

3. Below the `[Output]` section, create a configuration section for the engine by typing the task name inside square brackets. For example:

```
[XProtect]
```

4. In the new section, set the following parameters:

Type	The output engine to use. Set this parameter to <code>milestoneoutput</code> .
Input	A comma-separated list of the tracks that you want to output. To output information to Milestone, ensure you include the <code>Proxy</code> track generated by the Milestone ingest engine.
ProxyTrack	The <code>Proxy</code> track generated by the Milestone ingest engine.
XSLTemplate	The XSL template to use to transform the output track into a format that can be accepted by the Milestone system.

Host	The host name or IP address of the Milestone XProtect system.
Port	(Optional) The Milestone XProtect server port.
GUID	(Optional) The Milestone GUID of the camera that the events are associated with. This is not required if the video is ingested using the Milestone ingest engine.
Location	(Optional) Location metadata to send with the event.

For example:

```
[XProtect]
Type=milestoneoutput
Input=ANPR.Result,MilestoneIngest.Proxy
ProxyTrack=MilestoneIngest.Proxy
XSLTemplate=./toMilestone.xsl
Host=milestone-server
Port=9090
Location=Cambridge
```

For more information out the configuration parameters that you can use to configure this task, refer to the *Media Server Reference*.

5. Save and close the configuration file. If you modified the Media Server configuration file, you must restart Media Server for your changes to take effect.

Configure Milestone

To configure your Milestone surveillance system to process information sent by Media Server

1. Open the Milestone XProtect management application.
2. Make sure that the Milestone system has *Analytics Events* enabled, and is listening on the same port you specified in the Media Server configuration file.
3. For each type of event that you send to Milestone, add an *Analytic Event* to the Milestone system. The name of the analytic event must match the message produced by Media Server.
 - The default message for intelligent scene analysis events is the iSAS category name.
 - The default message for all ANPR events is ANPR.
 - The default message for recognized faces is the name of the face database.
 - The default message for unrecognized faces is NOT IN DATABASE.

Tip: To modify the message produced by Media Server, modify the `toMilestone.xsl` XSL template. The `<message>` element in the XML sent to Milestone can contain any name, but the names of the analytic events that you create in Milestone must match the messages produced by Media Server.

Note: The names are case-sensitive. For example, if you have a category in your iSAS configuration called "JumpRedLight", create an Analytic Event of the same name.

4. Add *Alarm Definitions* to the Milestone system as necessary. Use the Analytic Events that you created as the *Triggering Events*.
5. In the *Alarm List Configuration (Advanced Configuration > Alarms > Alarm Data Settings)*, select all of the columns. This allows users of the Milestone XProtect Smart Client to view all of the metadata that is provided by Media Server.

For more information about how to configure your Milestone system, refer to the Milestone documentation.

Part III: Distribute Processing

- "Use Multiple Media Servers"

Chapter 13: Use Multiple Media Servers

This section describes how to use multiple Media Servers, so that you can process more than one request simultaneously.

- [Distribute Media Server Operations](#)192

Distribute Media Server Operations

In large systems where you want to process a very large number of videos, you can use multiple Media Server instances. In this case, you can use a Distributed Action Handler (DAH) to distribute actions to each Media Server, to ensure that each Media Server receives a similar number of requests.

Install DAH

You can use the IDOL Server installer to install the DAH. For more information about installing the DAH, refer to the *Distributed Action Handler Administration Guide*.

Appendixes

This section contains the following Appendixes:

- ["Supported Languages" on page 194](#)
- ["Pre-Trained Classifiers" on page 195](#)

Appendix A: Supported Languages

This appendix describes the languages that are supported by Media Server OCR.

Latin Alphabet

Afrikaans (af)	Esperanto (eo)	Italian (it)	Portuguese (pt)
Basque (eu)	Estonian (et)	Irish (ga)	Romanian (ro)
Catalan (ca)	Finnish (fi)	Latin (la)	Slovak (sk)
Croatian (hr)	French (fr)	Latvian (lv)	Slovenian (sl)
Czech (cs)	German (de)	Lithuanian (lt)	Spanish (es)
Danish (da)	Hungarian (hu)	Maltese (mt)	Swedish (sv)
Dutch (nl)	Icelandic (is)	Norwegian (no)	Turkish (tr)
English (en)	Ido (io)	Polish (pl)	Welsh (cy)

Cyrillic Alphabet

Bulgarian (bg)	Serbian (sr)
Macedonian (mk)	Ukrainian (uk)
Russian (ru)	

Other Alphabets

Arabic (ar), Persian (fa), Urdu (ur)
Simplified Chinese (zhs), Traditional Chinese (zht)
Greek (el)
Hebrew (he)
Japanese (ja)
Korean (ko)

Appendix B: Pre-Trained Classifiers

HPE may provide classifiers that you can use with Media Server to classify images.

The following classifiers are currently available:

File name on download center	Description
Imageserver-ObjectClass_ImageNet.dat	A neural net (CNN) classifier that contains training for the Large Scale Visual Recognition Challenge (ILSVRC) classes listed at http://image-net.org/challenges/LSVRC/2012/browse-synsets .

For information about how to import a classifier into your training database, see "Import a Classifier" on [page 124](#)

Glossary

A

ACI (Autonomy Content Infrastructure)

A technology layer that automates operations on unstructured information for cross-enterprise applications. ACI enables an automated and compatible business-to-business, peer-to-peer infrastructure. The ACI allows enterprise applications to understand and process content that exists in unstructured formats, such as email, Web pages, Microsoft Office documents, and IBM Notes.

ACI Server

A server component that runs on the Autonomy Content Infrastructure (ACI).

ACL (access control list)

An ACL is metadata associated with a document that defines which users and groups are permitted to access the document.

action

A request sent to an ACI server.

active directory

A domain controller for the Microsoft Windows operating system, which uses LDAP to authenticate users and computers on a network.

ANPR

Automatic Number Plate Recognition, which reads the number/license plate of a vehicle.

C

Category component

The IDOL Server component that manages categorization and clustering.

Community component

The IDOL Server component that manages users and communities.

connector

An IDOL component (for example File System Connector) that retrieves information from a local or remote repository (for example, a file system, database, or Web site).

Connector Framework Server (CFS)

Connector Framework Server processes the information that is retrieved by connectors. Connector Framework Server uses KeyView to extract document content and metadata from over 1,000 different file types. When the information has been processed, it is sent to an IDOL Server or Distributed Index Handler (DIH).

Content component

The IDOL Server component that manages the data index and performs most of the search and retrieval operations from the index.

D

DAH (Distributed Action Handler)

DAH distributes actions to multiple copies of IDOL Server or a component. It allows you to use failover, load balancing, or distributed content.

DIH (Distributed Index Handler)

DIH allows you to efficiently split and index extremely large quantities of data into multiple copies of IDOL Server or the

Content component. DIH allows you to create a scalable solution that delivers high performance and high availability. It provides a flexible way to batch, route, and categorize the indexing of internal and external content into IDOL Server.

I

IDOL

The Intelligent Data Operating Layer (IDOL) Server, which integrates unstructured, semi-structured and structured information from multiple repositories through an understanding of the content. It delivers a real-time environment in which operations across applications and content are automated.

IDOL Proxy component

An IDOL Server component that accepts incoming actions and distributes them to the appropriate subcomponent. IDOL Proxy also performs some maintenance operations to make sure that the subcomponents are running, and to start and stop them when necessary.

Intellectual Asset Protection System (IAS)

An integrated security solution to protect your data. At the front end, authentication checks that users are allowed to access the system that contains the result data. At the back end, entitlement checking and authentication combine to ensure that query results contain only documents that the user is allowed to see, from repositories that the user has permission to access.

K

KeyView

The IDOL component that extracts data, including text, metadata, and subfiles from over 1,000 different file types.

L

LDAP

Lightweight Directory Access Protocol. Applications can use LDAP to retrieve information from a server. LDAP is used for directory services (such as corporate email and telephone directories) and user authentication. See also: active directory, primary domain controller.

License Server

License Server enables you to license and run multiple IDOL solutions. You must have a License Server on a machine with a known, static IP address.

O

OmniGroupServer (OGS)

A server that manages access permissions for your users. It communicates with your repositories and IDOL Server to apply access permissions to documents.

P

primary domain controller

A server computer in a Microsoft Windows domain that controls various computer resources. See also: active directory, LDAP.

R

record

A single package of metadata in a track. A record produced by an analysis task might describe a recognized face, a word spoken in the audio, or a number plate detected by ANPR. A record can contain a significant amount of information; for example a record describing a number plate includes timestamps describing when the number

plate was detected, the position of the number plate in the video frame, the characters read from the number plate, the confidence score for recognition, and so on.

rolling buffer

A fixed-size storage area on disk where you can save encoded video on a continuous basis. When the rolling buffer is full, the oldest content is discarded to make space for the latest.

S

scene analysis

Scene analysis recognizes suspicious activity in video and produces alarms to alert security personnel. Scene analysis can be trained to recognize many suspicious events, including vehicles driving through red lights, people entering restricted areas, and abandoned bags and vehicles.

T

track

A stream of data produced by a processing task in Media Server. For example, when you ingest video the ingest task produces two tracks: one for video frames and the other for audio packets. Other tasks use these tracks. Analysis tasks read the data and produce tracks that contain analysis results; encoding tasks take the video and audio data to write files to disk. See also record.

V

View

An IDOL component that converts files in a repository to HTML formats for viewing in a Web browser.

W

Wildcard

A character that stands in for any character or group of characters in a query.

X

XML

Extensible Markup Language. XML is a language that defines the different attributes of document content in a format that can be read by humans and machines. In IDOL Server, you can index documents in XML format. IDOL Server also returns action responses in XML format.

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Media Server Administration Guide (Media Server 11.0)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

We appreciate your feedback!