Avid Media | Index

Configuration Guide

Version 2.7

Contents

Using This Guide

This guide is intended for users of Avid MediaCentral Platform Services (MCS) with the Media | Index feature. This guide describes the configuration procedures for Media Index, such as configuring Media Index for single-server and cluster installations.

For information on installing MCS, see the *MediaCentral Platform Services Installation and Configuration Guide*. For administrative information for MediaCentral | UX, see the *Avid MediaCentral | UX Administration Guide*.

Symbols and Conventions

Avid documentation uses the following symbols and conventions:

Symbol or Convention	Meaning or Action
	A note provides important related information, reminders, recommendations, and strong suggestions.
	A caution means that a specific action you take could cause harm to your computer or cause you to lose data.
	A warning describes an action that could cause you physical harm. Follow the guidelines in this document or on the unit itself when handling electrical equipment.
>	This symbol indicates menu commands (and subcommands) in the order you select them. For example, File > Import means to open the File menu and then select the Import command.
▸	This symbol indicates a single-step procedure. Multiple arrows in a list indicate that you perform one of the actions listed.
(Windows), (Windows only), (Macintosh), or (Macintosh only)	This text indicates that the information applies only to the specified operating system, either Windows or Macintosh OS X.
Bold font	Bold font is primarily used in task instructions to identify user interface items and keyboard sequences.
Italic font	Italic font is used to emphasize certain words and to indicate variables.
`Courier Bold font`	Courier Bold font identifies text that you type.
Ctrl+key or mouse action	Press and hold the first key while you press the last key or perform the mouse action. For example, Command+Option+C or Ctrl+drag.

If You Need Help

If you are having trouble using your Avid product:

1. Retry the action, carefully following the instructions given for that task in this guide. It is especially important to check each step of your workflow.

2. Check the latest information that might have become available after the documentation was published. You should always check online for the most up-to-date release notes or ReadMe because the online version is updated whenever new information becomes available. To view these online versions, select ReadMe from the Help menu, or visit the Knowledge Base at www.avid.com/support.

3. Check the documentation that came with your Avid application or your hardware for maintenance or hardware-related issues.

4. Visit the online Knowledge Base at www.avid.com/support. Online services are available 24 hours per day, 7 days per week. Search this online Knowledge Base to find answers, to view error messages, to access troubleshooting tips, to download updates, and to read or join online message-board discussions.

Avid Training Services

Avid makes lifelong learning, career advancement, and personal development easy and convenient. Avid understands that the knowledge you need to differentiate yourself is always changing, and Avid continually updates course content and offers new training delivery methods that accommodate your pressured and competitive work environment.

For information on courses/schedules, training centers, certifications, courseware, and books, please visit www.avid.com/support and follow the Training links, or call Avid Sales at 800-949-AVID (800-949-2843).

1 Media | Index Configuration Overview

This document describes the tasks required to configure a Media Index system.

Media Index comprises the following components:

- Elasticsearch® service, which manages data storage and runs the search engine.
- Elasticsearch tribe service, the federated client that enables multi-zone search functionality.
- Data Import API, a bus-based service for metadata indexing, search, configuration, status, permissions, and updates for Media Index.
- Search API, a bus-based service that provides the search API
- Media Index Configuration API, a bus-based service that provides the API for the Media Index configuration
- Media Index Status API, a bus-based service that provides the API for the Media Index
- Media Index Permissions API, a bus-based service that provides per asset permissions for the search service
- Media Index Feed API, a bus-based service that provides an updated assets feed for external services
- Media Index Thesaurus API, a bus-based service that provides API for importing structured metadata (dictionaries)

Elasticsearch is an open source, distributed, real-time search and analytics engine. Elasticsearch stores information about assets and makes it accessible across multiple sites.

For more information on Elasticsearch operation and term definitions, visit the Elasticsearch website: https://www.elastic.co/guide/en/elasticsearch/guide/master/distributed-cluster.html

Prerequisites for Media | Index Configuration

Before you configure Media Index, you need to install and configure the following MediaCentral components:

- MediaCentral Platform Services (MCS) — Media Index is configured after you have installed and set up your MCS system, either for a single server or for a cluster configuration. For more information about installing and configuring MCS, see the *MediaCentral Platform Services Installation and Configuration Guide*.
- Multi-zone components — If you use MCS in a multi-zone environment, you must first configure the multi-zone components. For information on configuring a multi-zone environment, see the *MediaCentral Platform Services Installation and Configuration Guide*.

MediaCentral servers configured with Media Index require a minimum of 128 GB of RAM in each server. Prior to configuring Media Index, verify that your server meets these requirements.

For more information on hardware requirements, see the *MediaCentral Platform Services Hardware Guide*.

Media Index is installed with MediaCentral Platform Services by default, but the Media Index services are not started by default. Once configured, Media Index services start automatically when the server is powered-on.

Media Index can be configured in both standalone and multi-zone environments. Multi-zone configurations federate separate MCS installations, which can be co-located in a single facility or geographically separated — for example, one zone could be located in Toronto and another in Munich.

If configuring Media Index with an Interplay Production system, review the following:

- The Production Engine Bus Connector (PEBCo) service requires Interplay Engine v3.1 or later.

- If you are installing or updating your configuration with Media Index v2.3.x, you must install the Interplay Access v3.3 client. You do not need to update your Interplay Production Engine to v3.3, but you must have the Interplay Administration tool that installs with Interplay Access v3.3 to enable sync events for Media Index.

- If you are installing or updating your configuration with Media Index v2.5.x, you must install the Interplay Access v3.5 or later client. You do not need to update your Interplay Production Engine to v3.5, but you must have the Interplay Administration tool that installs with Interplay Access v3.5 or later to enable sync events for Media Index.

Checklist for Configuring Media | Index

The following table provides a checklist for configuring Media Index components. The checklist also includes chapter and page references to complete each step.

Task	Section Reference
❑ Review the prerequisites.	"Prerequisites for Media	Index Configuration" on page 8.
❑ Configure the Media Index and Elasticsearch services on the MediaCentral server or servers.	"Configuring Media	Index for a Single Server" on page 11 or "Configuring Media	Index for a Cluster" on page 16.
❑ Validate your configuration.	"Validating the Media	Index Configuration" on page 14 or "Validating the Media	Index Configuration" on page 23.
❑ If you have a multi-zone setup, configure Media Index for multiple zones.	"Setting Up Multi-Zone Search" on page 26.
❑ If you are integrating with an Interplay Production system, configure the Production Engine Bus Connector.	"Installing the Production Engine Bus Connector" on page 33.
❑ If you are integrating with an Interplay Production system that includes Production Services such as Interplay Delivery, configure additional provider services used with Media Index workflows.	"Installing and Configuring Production Services for MediaCentral Delivery" on page 41.

Task	Section Reference
❑ If you are integrating with an Avid iNEWS system, configure the iNEWS servers for Media Index.	"Configuring iNEWS for Media	Index" on page 53.
❑ If you are integrating with an Interplay MAM system, configure the MAM servers for Media Index.	"Setting Up Interplay	MAM Sync Service" on page 59. For information on installing and configuring Interplay MAM, see the *Interplay	MAM Installation Manual* and the *Interplay	MAM Sync Service Administrator User's Guide*.

2 Configuring Media | Index for a Single Server

When you set up Media Index, you enable Elasticsearch and several other services related to Media Index. Elasticsearch and Elasticsearch Tribe are used to index, store and search assets while the Media Index services add specific Avid logic to Elasticsearch.

Elasticsearch indexes are split into "shards" and "replicas". Each shard contains indexed data and each replica contains a copy of that data. Replicas are used to spread the search load and provide redundancy. Single server installations default to one shard and zero replicas. This is the appropriate number of shards and replicas for single server configuration and in most cases the numbers should not be altered.

The configuration of shards and replicas on a single server is fairly simple, but the deployment becomes more complex in clustered systems. If you are interested in learning more about Elasticsearch shards and replicas, see of this guide.

Preparing for the Media | Index Configuration

When you install MediaCentral Platform Services, Media Indexes services are stopped, except for avid-acs-media-index-configuration. Before you begin configuring Media Index, you must start this configuration service if it has stopped.

You can also change the default method of naming Elasticsearch nodes. The Media Index API uses the hostname of the Elasticsearch node as the display name of the node. If you disable this default behavior, Elasticsearch automatically generates node names. You can also supply a different value to use for assigning node names.

To check the status of the Media Index configuration service:

1. Open a Linux terminal window on your MCS server and type the following command:

   `service avid-acs-media-index-configuration status`

   The result must report that the service has started:

   `avid-acs-media-index-configuration is running`

2. If the configuration service is not running, type the following command:

   `service avid-acs-media-index-configuration start`

To change the default Elasticsearch node name method:

1. Open the configuration file for editing by typing the following command:

   `vi /etc/sysconfig/avid-acs-media-index-configuration`

2. Using the text editor, locate the following line:

   `export ELASTICSEARCH_NODE_NAME=hostname`

3. Delete the `hostname` value and type the value you want to use to assign to node.

📖 *If the node name variable is not set, Elasticsearch automatically generates node names when your node starts. If it is set to* `hostname`*, Elasticsearch uses the hostname of the node where the service is running.*

4. Press the Escape key to return to command mode, then type **`:wq`** and press Return to write and quit the text editor.

5. Restart the configuration service to enable the configuration changes:

```
service avid-acs-media-index-configuration restart
```

Configuring Media | Index

You configure Media Index in the Media Central | UX System Settings. You must provide the following properties in your settings:

- Cluster Name — The "Cluster Name" is simply an identifier for the Media Index configuration on this system. In a standalone configuration (non-multi-zone), Avid recommends using the local hostname or site location (e.g. NewYork). In a multi-zone configuration, Avid recommends using the same value that you used for "Zone Name" during the multi-zone configuration process. Each zone in a multi-zone configuration requires a unique Media Index Cluster Name.

- Elasticsearch Cluster Hosts — Specify the host name or IP address of the MediaCentral server running the Elasticsearch service.

- Elasticsearch TCP Port — Specify the TCP port of the Elasticsearch data node, used for binary connection protocol and for inter-node communications. The default value is 9300. In most cases you, you do not need to modify this property.

- Elasticsearch HTTP Port — Specify the port of the Elasticsearch data node HTTP API. The default value is 9200. In most cases you, you do not need to modify this property.

- Elasticsearch Tribe TCP Port — Specify the TCP port of the Elasticsearch tribe node, used for inter-node communications. The default value is 9305. In most cases you, you do not need to modify this property.

- Elasticsearch Tribe HTTP Port — Specify the port of the Elasticsearch tribe node HTTP API. The default value is 9201. In most cases you, you do not need to modify this property

- Default Shards — Specify default shards count. Avid recommends using a default value of 1 shard.

- Default Replicas — Specify default replica count. The default number of replicas is 0, which is the required replica count for a single-server configuration.

- Supported languages — Lists the languages supported by Media Index.

Media Index supports multiple languages for your search. However, enabling all languages can slow down searches. You should enable only those languages your users require. If you work with a multi-zone configuration, you can enable different languages for each zone, and these languages are then available for all indexes within the specified zone. If possible, your language selection should be the same across all zones, because a search using one language returns results only from indexes that you have configured for this language.

The following table provides an example of which languages are used in searches within a multi-zone configuration with different languages enabled for different zones.

Enabled Languages	Language Selected for Search	Languages in Results
Zone 1: English, French	English	Zone 1: English
Zone 2: English, Korean		Zone 2: English
Zone 1: English, French	French	Zone 1: French
Zone 2: English, Korean		Zone 2: [no results]
Zone 1: English, French, Korean	French	Zone 1: French
Zone 2: English, French, Korean		Zone 2: French

⚠ **If you change your language settings after you have configured Media Index, you must reset Media Index. This deletes all indexes, which you then need to re-index.**

To configure Media Index:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Settings.

   The Media Index Settings display in the Details pane.

4. Specify the required settings for the following properties:

   - Cluster Name (hostname or IP address of the local Media Central server)

📄 *Hostnames can include only letters, numbers, dashes, and underscore symbols.*

- Elasticsearch Cluster Hosts (hostname or IP address of the local Media Central node)
- Elasticsearch TCP Port
- Elasticsearch HTTP Port
- Elasticsearch Tribe TCP Port
- Elasticsearch Tribe HTTP Port
- Default Shards
- Default Replicas

5. Click Apply. The system will display a message on the screen indicating "Update in progress...".

   After a few moments, a second message should appear indicating that Index has been configured successfully:

Update status. ⊠
Update was successful.
Close
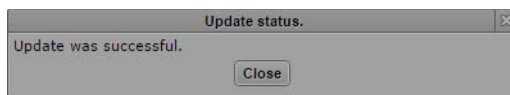
6. Click Close.

Validating the Media | Index Configuration

You can validate your setup by checking the status of the Media Index services and ensuring that they are running. Using the Cluster State pane in MediaCentral UX, you can view your node with core services running and see that the system data index is available.

To check the current status of Media Index:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Cluster State.

   The Media Index Cluster State display in the Details pane.
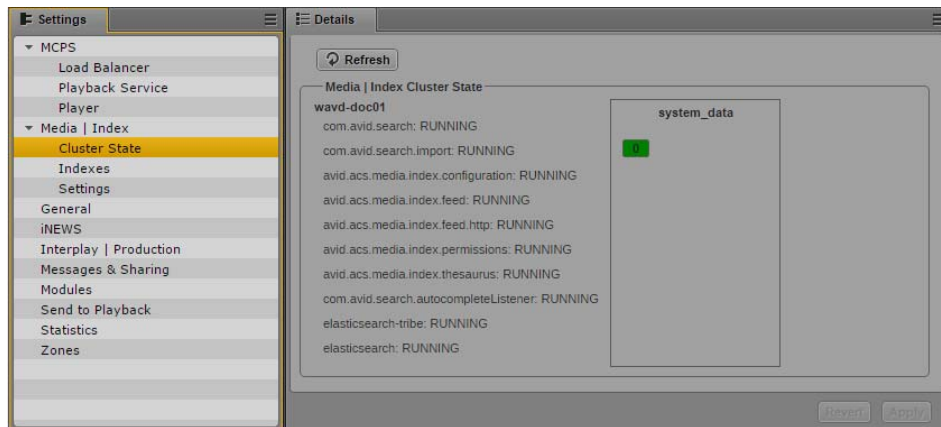
Media Index Cluster State pane, with the system_data index and one index available

The core services for your configuration must display a status of "Running," and the system_data must display the shard count in green.

Later, once you have configured an Interplay Production, MAM or iNEWS system, additional information on the indexes will appear in this window. The following image shows an index for an Interplay Production system:
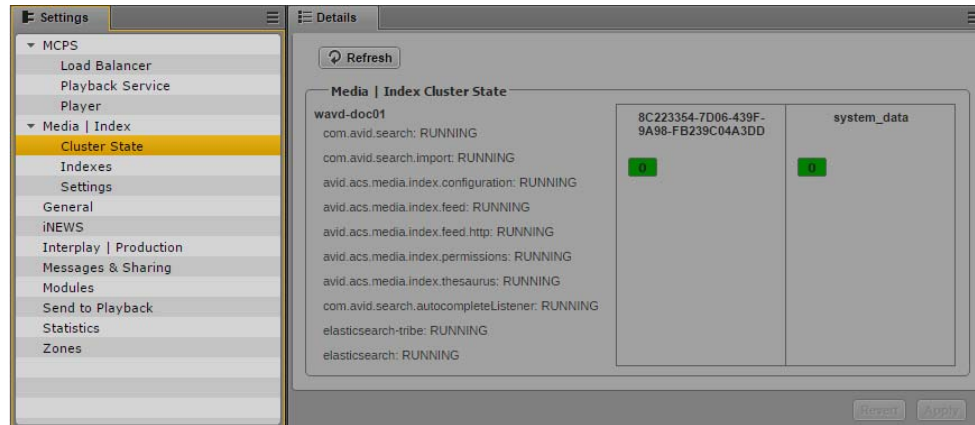
Once configured, the string of characters listed in the index will match the System ID that appears in the System Settings > Media Index > Indexes pane.

Continuing the Installation

Depending upon your workflow, proceed to one or more of the following sections as applicable:

- If you are installing Media Index in a multi-zone configuration, see "Setting Up Multi-Zone Search" on page 26.

- If you are integrating with and Interplay Production or Interplay Archive system, you need to configure and start the Production Engine Bus Connector (PEBCo) service. For information on PEBCo, see "Installing the Production Engine Bus Connector" on page 33.

- If you are integrating with an Avid iNEWS system, you must configure the iNEWS servers to synchronize data with the Media Index services. For information, see "Configuring iNEWS for Media | Index" on page 53.

- If you are integrating with an Interplay MAM system, you must configure the MAM servers to communicate with Media Index. For more information, see "Setting Up Interplay | MAM Sync Service" on page 59.

3 Configuring Media | Index for a Cluster

When you set up Media Index, you enable Elasticsearch and several other services related to Media Index. Every node in an MCS cluster runs an Elasticsearch instance as well as an Elasticsearch-tribe instance. Elasticsearch and Elasticsearch Tribe are used to index, store and search assets while the Media Index services add specific Avid logic to Elasticsearch. Tribe instances allow you to perform read operations against the nodes in multiple zones.

Although the MediaCentral Corosync cluster includes resources that relate to Media Index, the Corosync cluster and the Elasticsearch cluster do not share configuration details and use different mechanisms for the clustering.

You can configure Elasticsearch in a federated way so that a tribe instance of one Elasticsearch zone can be connected to other zones (one or more Elasticsearch data instances). For MediaCentral multi-zone configurations, this functionality is required The tribe node configuration process connects multiple systems to each other so that searches from each zone can find assets from all other remotely connected Elasticsearch zones.

When you create an index, Elasticsearch creates a number of "primary shards" and "replicas," each shard contains indexed data and each replica contains a copy of that data. You use shards and replicas to spread the search load and provide redundancy.

Shard Count and Usage

Elasticsearch shards are available in two varieties: primaries and replicas. The key difference is that primary shards allow for both read and write while replica shards are read-only.

During the configuration of Media Index, you configure the number of (primary) shards and replicas through the Media Central UX System Settings. In a single-server configuration, a single shard is configured with zero replicas. Since there is only one server, there is little benefit to having a replica of the shard on the same system.

In cluster configurations, Avid recommends setting the shard count to one and the replica count to one fewer than the number of cluster nodes to accommodate fail-over safety and load balancing. For example, if your configuration includes a cluster with five nodes, set the replica count to four. This allows for one node to maintain the primary shard with a replica of all your indexed data on the remaining nodes. This horizontal scaling allows the load to be distributed among the nodes which increases efficiency and enables redundancy.

Configuring a Single Primary Shard

Consider a MediaCentral cluster consisting of three nodes, as shown below:
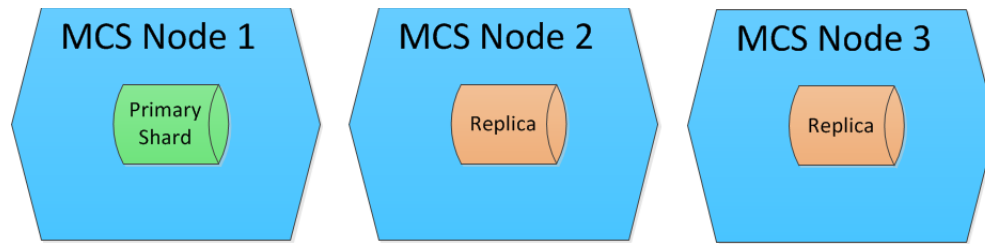
In this configuration, the MediaCentral UX System Settings were configured for:

- 1 Default (primary) Shard
- 2 Index Replicas

Elasticsearch intelligently locates the shards across the three nodes. In this example the primary shard happens to be located on Node 1, but that might not always the case. A single primary shard, coupled with replicas, provides redundancy and flexibility for your index.

In most situations, a single primary shard provides enough resources to host your entire index. MediaCentral systems that are indexing very large databases, such as large Interplay MAM deployments, can increase the shard count to decrease the amount of data per shard. In such configurations, an MCS cluster is used to spread the increased load of multiple primary shards across the cluster nodes.

Configuring Multiple Primary Shards

In the following example, the System Settings in the same MediaCentral cluster have been configured for 3 primary shards and 2replica shards:
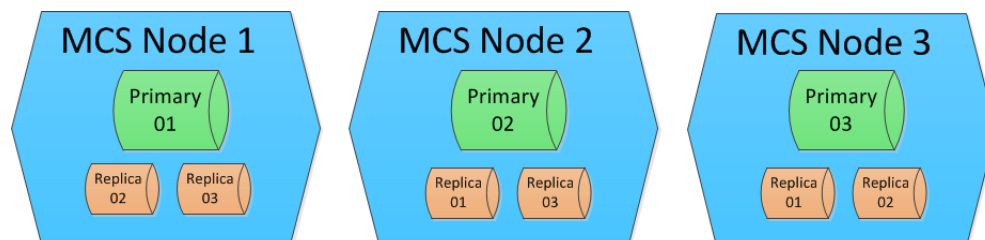
As previously stated, this configuration spreads the load across multiple systems which increases database efficiency and enables additional redundancy. Notice that the primary shard and its replica are never co-located on the same server. The number of primary shards in your configuration affects the amount of data that can be stored in your index, although the type of data and metadata affects the size of the index as well.

It is important to plan your configuration early in the process. Once configured, users can add replicas to the configuration, but primary shards cannot be added without reconfiguring and resynchronizing the indexes which can be a time consuming process. For example, if you start with two server nodes but know that at some point you will scale your configuration up to six nodes, you

might want to configure Media Index initially for six shards. As the additional nodes are added to the configuration, Elasticsearch relocates the primary shards to the new nodes so that the shards are evenly distributed across the six servers (one shard on each node).

Increasing the amount of shards without increasing the number of nodes on the system adds load to the system resources. For systems with spinning disks, there should ideally be nor more than two shards per node. Only increase the number of primary shards if you know that additional nodes will be added in the near future or if directed to do so by Avid.

Explaining Failover

If a primary node is lost, it must have a replica for the system to continue to operate normally. If available the former read-only replica shard is automatically promoted to primary, providing uninterrupted read/write access to the Elasticsearch database. If a replica is not available, Indexed searches will fail until the node containing the missing sharded is repaired. Configuring multiple shards and replicas gives your index more stability in case of a failure or if you need to take one or more nodes offline for maintenance.

Example 1: In this illustration, Node 1 has failed. Elasticsearch dynamically promoted the replica of the shard on Node 2 to primary. Since Node 2 and Node 3 both contain a replica of Node 1, Elasticsearch decides which replica becomes the primary based on an internal algorithm.
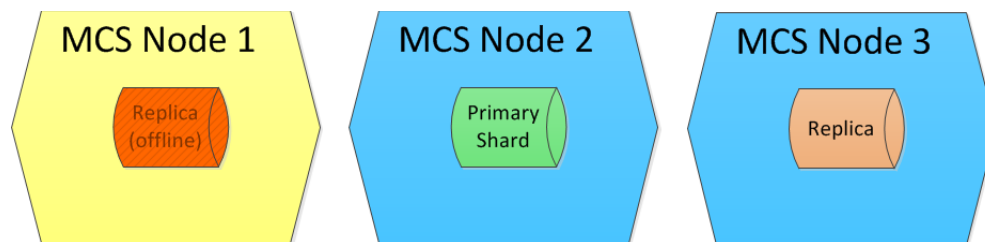
Note that in this configuration, the system is still fully redundant. The primary shard has a replica on Node 3. Since replicas load-balance read requests, it is possible that users might see a small reduction in search performance.

Example 2: Similar to the previous illustration, Node 2 was taken offline which resulted in Elasticsearch dynamically promoting the replica shard on Node 3 to primary.
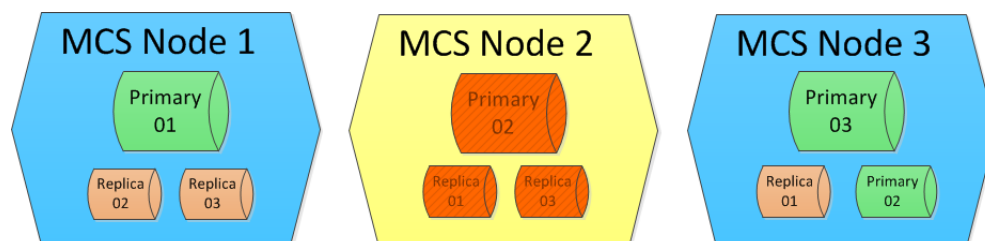
Note that in this configuration, the system is still fully redundant. Each primary shard has a replica on another node. Node 3 now hosts two primary shards.

Example 3: In this final example, MediaCentral has been configured for 3 primary shards and 1 replica:
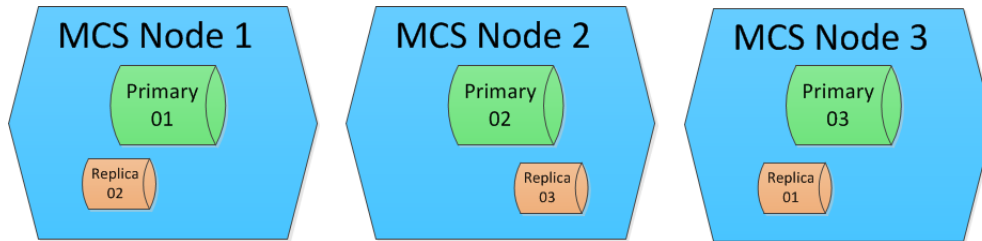
If Node 1 goes offline, Replica 1 on Node 3 becomes a primary shard and the index remains fully functional. However, since Node 1 also hosted the replica for Node 2, Elasticsearch then re-creates the missing replica, as shown in the following illustration:
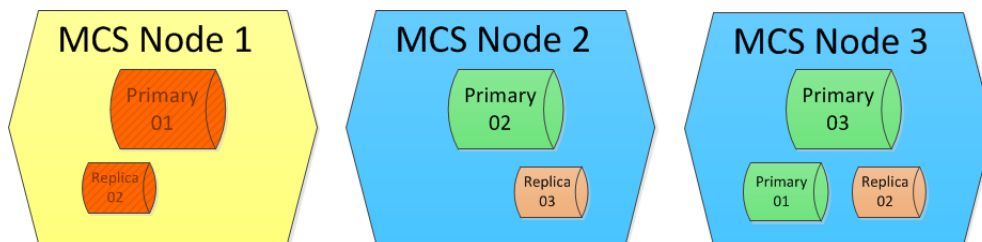
If Node 2 goes offline, the index remains functional because Node 3 contains Shards 1 and 3, as well as a replica of Shard 2. However, if Node 2 goes offline before Elasticsearch re-creates the replica on Node 3, then the index is incomplete and cannot function.

Preparing for the Media | Index Configuration

Prior to configuring Media Index on a MediaCentral cluster, it is good practice to check that the cluster is healthy. The Cluster Resource Monitor, `crm_mon`, is used to verify this information.

To monitor the cluster:

1. Log in to any node in the cluster and open the Cluster Resource Monitor:

   `crm_mon -f`

2. If any failure are listed, use the cleanup command to reset the fail-counts:

   `crm resource cleanup <resource>`

   Where *<resource>* is the name of the Pacemaker resource you want to clean — for example, avid-acs-search.

3. The Cluster Resource Monitor is needed again in the following process. If desired, you can leave the monitor tool running or close it and reopen it again later.

   To close the Cluster Resource Monitor, press CTRL-C on a Windows keyboard or CMD-C on a Mac keyboard.

For more information on using the cluster resource monitor utility, see "Cluster Resource Monitor" in the MediaCentral Platform Services Concepts and Clustering Guide.

Adding Resources to the Cluster

When you install a new MediaCentral system, all Media Index services are stopped except for the "avid-acs-media-index-configuration" service. After issuing the setup command, as detailed in the following process, additional services are started and resources are added to the cluster.

To add resources to the cluster:

1. Log in to any node in the cluster as root and type the following command:

   `/opt/avid/cluster/bin/search-cluster setup`

   The necessary resources are added to the cluster configuration, and the related services are started. The following example shows the new cluster resources:

   ```
   pgpool (lsb:pgpool):    Started wavd-mcs01
   pgpoolchecker  (lsb:pgpoolchecker):    Started wavd-mcs01
   Clone Set: elasticsearchEverywhere [elasticsearch]
       Started: [ wavd-mcs01 wavd-mcs02 wavd-mcs03 ]
   Clone Set: elasticsearchTribeEverywhere [elasticsearchTribe]
       Started: [ wavd-mcs01 wavd-mcs02 wavd-mcs03 ]
   Clone Set: AvidSearchConfigEverywhere [AvidSearchConfig]
       Started: [ wavd-mcs01 wavd-mcs02 wavd-mcs03 ]
   ```

   If the cluster is not detected, the following message is displayed:

   ```
   The cluster services are either down or the cluster is not yet configured.
   Please make sure that the corosync and pacemaker services are running
   correctly and that the cluster is properly configured before rerunning the
   search-cluster script.
   ```

2. If not already running, open the Cluster Resource Monitor to verify that the resources have been added successfully:

   `crm_mon -f`

   When configuring a cluster for Media Index a number of new resources are added to Pacemaker, making it difficult to see the fail-counts at the end of the output. If this is the case, you can modify the command with a "1" which tells the Cluster Resource Monitor to print the output to the screen once and exit:

   `crm_mon -f1`

3. If any failure are listed, use the cleanup command to reset the fail-counts:

   `crm resource cleanup <resource>`

   Where <resource> is the name of the Pacemaker resource you want to clean — for example, avid-acs-search.

4. Press CTRL-C on a Windows keyboard or CMD-C on a Mac keyboard to exit the crm_mon utility.

Configuring Media | Index

You configure Media Index in the Media Central | UX System Settings. You must provide the following properties in your settings:

- Cluster Name — This name has no direct relation to the virtual host name given to the Corosync cluster. The "Cluster Name" is simply an identifier for the Media Index configuration on this system. In a standalone configuration (non-multi-zone), Avid recommends using the Corosync

cluster hostname or site location (e.g. NewYork). In a multi-zone configuration, Avid recommends using the same value that you used for "Zone Name" during the multi-zone configuration process. Each zone in a multi-zone configuration requires a unique Media Index Cluster Name.

⚠ **Avid recommends that you not change the Cluster Name after your initial configuration. If you change the cluster name, Elasticsearch creates a new data folder on the next restart. While the index data is not lost, any indexes that existed before changing the cluster name fail to appear in your configuration and the information in them is not accessible. When you work in a multi-zone set-up, you should consider which cluster names you plan to use in all zones and not change them after the indexing process has started.**

- Elasticsearch Cluster Hosts — Specify the hostname or IP address of each MediaCentral server running the Elasticsearch service. Each must be separated by a comma.

- Elasticsearch TCP Port — Specify the TCP port of the Elasticsearch data node, used for binary connection protocol and for inter-node communications. The default value is 9300. In most cases you, you do not need to modify this property.

- Elasticsearch HTTP Port — Specify the port of the Elasticsearch data node HTTP API. The default value is 9200. In most cases you, you do not need to modify this property.

- Elasticsearch Tribe TCP Port — Specify the TCP port of the Elasticsearch tribe node, used for inter-node communications. The default value is 9305. In most cases you, you do not need to modify this property.

- Elasticsearch Tribe HTTP Port — Specify the port of the Elasticsearch tribe node HTTP API. The default value is 9201. In most cases you, you do not need to modify this property

- Default Shards — Specify the number of default shards. Avid recommends using a default value of 1 shard. This only changes if indexing a very large database such as an Interplay MAM system. If you are considering altering the number of default shards, consult your Avid representative before making this change.

- Default Replicas — Specify default replica count. The default number of replicas is 0, which is the required replica count for a single-server configuration. In an MCS cluster installation, you must set the default number of replicas to a number greater than zero.

⚠ **A replica count greater than zero is necessary in a cluster configuration in order to ensure that your data gets distributed properly and remains safe. For maximum security, you should set the replica count to a value of *<number of MCS cluster nodes>* minus 1. Do not set the replica count to a value greater than this. For example, if your configuration consists of two nodes, then set the replica count to 1.**

- Supported languages — Lists the languages supported by Media Index.

Media Index supports multiple languages for your search. However, enabling all languages can slow down searches. You should enable only those languages your users require. If you work with a multi-zone configuration, you can enable different languages for each zone, and these languages are then available for all indexes within the specified zone. If possible, your language selection should be the same across all zones, because a search using one language returns results only from indexes that you have configured for this language. The following table provides an example of which languages are used in searches within a multi-zone configuration with different languages enabled for different zones.

Enabled Languages	Language Selected for Search	Languages in Results
Zone 1: English, French	English	Zone 1: English
Zone 2: English, Korean		Zone 2: English
Zone 1: English, French	French	Zone 1: French
Zone 2: English, Korean		Zone 2: [no results]
Zone 1: English, French, Korean	French	Zone 1: French
Zone 2: English, French, Korean		Zone 2: French

⚠ **If you change your language settings after you have configured Media Index, you must reset Media Index. This deletes all indexes, which you then need to re-index.**

When you set these values and apply your changes, Media Index writes to several configuration files and restarts cluster services and resources. It also creates the system_data index. This process might take some time, depending on your cluster configuration.

To configure Media Index:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Settings.

   The Media Index Settings display in the Details pane.

4. In the Supported Languages section, select the languages you want to make available for searching. You can select as many languages as needed, but Avid recommends you select only those languages your users are likely to use in a search.

5. Click the Default Language menu, and select the default language for indexed searches.

6. Specify the required settings for the following properties:

   - Cluster Name

   - Elasticsearch Cluster Hosts

   - Elasticsearch TCP Port

   - Elasticsearch HTTP Port

   - Elasticsearch Tribe TCP Port

   - Elasticsearch Tribe HTTP Port

   - Default Shards

   - Default Replicas

7. Click Apply. The system will display a message on the screen indicating "Update in progress...".

   After a few moments, a second message should appear indicating that Index has been configured successfully:

Update status.	☒
Update was successful.	
Close	
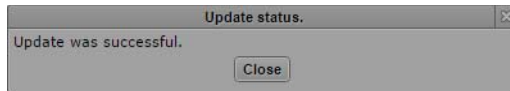
8. Click Close.

Validating the Media | Index Configuration

You can validate your setup by checking the status of the Media Index services and ensuring that they are running. Using the Cluster State pane in MediaCentral UX, you can view all nodes with core services running and see that the system data index is available. You should also use the Cluster Resource Monitor utility, crm_mon, to verify the status and failure counts of the MediaCentral services. Any failures should be reset or investigated.

📄 *For information on using the cluster resource manager, see "Monitoring Services and Resources" in the* MediaCentral Platform Services Installation and Configuration Guide.

To check the current status of Media Index:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Cluster State.

   The Media Index Cluster State displays in the Details pane.
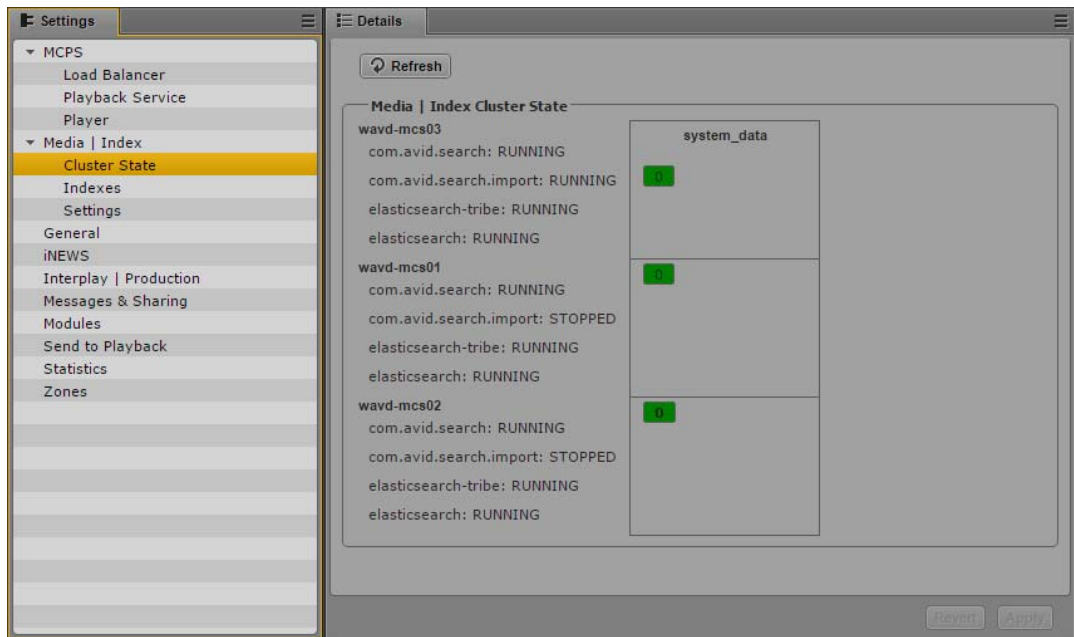
The core services for your configuration must display a status of "Running," and the indexes must display the shard count in green. Note that the avid-acs-search-import service is installed on all cluster nodes, but is only ever active on one node.

To monitor the fail count within the cluster:

1. Log in to any node in the cluster as root and verify the cluster fail counts by typing the following command:

```
crm_mon -f
```

When configuring a cluster for Media Index a number of new resources are added to Pacemaker, making it difficult to see the fail-counts at the end of the output. If this is the case, you can modify the command with a "1" which tells the command to print the output to the screen once and exit:

```
crm_mon -f1
```

Additional "AvidSearch" resources will have been added to the cluster.

2. Use the cleanup command to reset any observed failure counts:

```
crm resource cleanup <resource>
```

<resource> is the name of the resource you want to monitor — for example, avid-acs-search.

For more information on using the cluster resource monitor utility, see "Cluster Resource Monitor" in the MediaCentral Platform Services Concepts and Clustering Guide.

Continuing the Installation

Depending upon your workflow, proceed to one or more of the following sections as applicable:

- If you are installing Media Index in a multi-zone configuration, see "Setting Up Multi-Zone Search" on page 26.

- If you are integrating with and Interplay Production or Interplay Archive system, you need to configure and start the Production Engine Bus Connector (PEBCo) service. For information on PEBCo, see "Installing the Production Engine Bus Connector" on page 33.

- If you are integrating with an Avid iNEWS system, you must configure the iNEWS servers to synchronize data with the Media Index services. For information, see "Configuring iNEWS for Media | Index" on page 53.

- If you are integrating with an Interplay MAM system, you must configure the MAM servers to communicate with Media Index. For more information, see "Setting Up Interplay | MAM Sync Service" on page 59.

Once configured, taking a cluster node configured for Media Index offline requires special procedures. for more information, see "Taking Nodes Offline in a MediaCentral Cluster" on page 80.

4 **Setting Up Multi-Zone Search**

In a multi-zone MCS configuration, each zone runs its own Elasticsearch cluster with its own Elasticsearch Cluster Name. This is true for MCS single-servers as well as MCS servers in a Corosync cluster configuration. Single servers run a Elasticsearch cluster consisting of one node. These different Elasticsearch clusters are then connected to each other using the Elasticsearch tribe service. For example, in a two zone configuration, if you call one cluster "Zone1" and another cluster "Zone2," you configure the Media Index Cluster Name to Zone1 on all nodes in one MCS cluster and to Zone2 on all nodes in the second MCS cluster.

> *Avid recommends that you name each cluster using the same value that you used for "Zone Name" during the multi-zone configuration process.*

You can configure the MCS multi-zones before or after starting the Media Index multi-zone configuration. However, if you complete your Media Index multi-zone configuration without having previously configured your MCS multi-zones, you cannot play back remote assets until the multi-zone configuration process is complete. Avid recommends configuring and testing the multi-zone configuration before configuring Media Index.

If you need to change the Media Index Cluster Name, you must make the change before proceeding with the Media Index multi-zone configuration process detailed in this chapter. For information on changing the Cluster Name, see "Modifying the Media | Index System Settings" on page 84.

> ⚠ **Avid recommends that you not change the Cluster Name after your initial configuration. If you change the cluster name, Elasticsearch creates a new data folder on the next restart. While the index data is not lost, any indexes that existed before changing the cluster name fail to appear in your configuration and the information in them is not accessible. When you work in a multi-zone set-up, you should consider which cluster names you plan to use in all zones and not change them after the indexing process has started.**

For information for configuring MCS for multi-zones, see the *MediaCentral Platform Services Installation and Configuration Guide*. For information on user management and multi-zones in MediaCentral UX, see "Managing Multi-Zone Environments" in the *Avid MediaCentral | UX Administration Guide*.

Setting the Zone Bindings

To join two or more individual zones in a multi-zone setup, you must update the Elasticsearch tribe configuration on all of your tribe nodes and set bindings for all Elasticsearch clusters. If one or more zones share the same Cluster Name, you also have to reconfigure the Cluster Name so that each zone has a unique name. This requires changes in the Elasticsearch and the Elasticsearch tribe.

> *You must update the Elasticsearch configuration file — /etc/elasticsearch-tribe/elasticsearch.yml — on all cluster nodes in each of the zones.*

To update the Elasticsearch tribe configuration:

1. Using the Linux text editor, *vi*, open the Elasticsearch tribe configuration file for editing:

   `vi /etc/elasticsearch-tribe/elasticsearch.yml`

2. Add the following lines in the configuration file to add the binding information that allows the tribe to connect to another Media Index cluster:

   `tribe.<RemoteClusterName>.cluster.name: <RemoteClusterName>`

   `tribe.<RemoteClusterName>.transport.tcp.port: <port>`

   Where the following variables are used:

   - `<RemoteClusterName>` - This value represents the Media Index Cluster Name specified in the MediaCentral UX System Settings of the remote zone.

   - `<port>` - The `transport.tcp.port` number must be unique for each Media Index cluster. Avid recommends using port number 9313 as the starting port. If additional zones are added to your configuration, increment the port number by one for each zone (9314, 9315, and so on). Avid recommends using sequential port numbers for clarity, but this is not required.

   You must use a space between the colon (:) and your custom variable for each of these lines.

3. Add the following line to configure Unicast for tribe node binding:

   ```
   tribe.<RemoteClusterName>.discovery.zen.ping.unicast.hosts:
     - "<remoteMCSNode1>:9300"
     - "<remoteMCSNode2>:9300"
   ...
     - "<remoteMCSNodeN>:9300"
   ```

   Where the following variables are used:

   - `<RemoteClusterName>` - This value represents the Media Index Cluster Name specified in the MediaCentral UX System Settings of the remote zone.

   - `<remoteMCSNode1>` through `<remoteMCSNodeN>` - Each value represents the short host name or IP address of an MCS server in the remote zone. If the remote zone is a single server, only one line is added. If the remote zone is a cluster, a line is added for each node.

   When adding this information to the configuration file, strict adherence to formatting is important. You must add two spaces before the dash "-" and one space after the dash. Do not add a space between the node name, the colon and the port number.

4. Save and exit the vi session:

   Press <ESC> and type: `:wq`

5. If you have a cluster configuration, repeat these steps for each node within the cluster.

6. Once complete, restart all Elasticsearch tribe nodes affected by your configuration changes:

   ▶ For single-server configurations, type the following command:

   `service elasticsearch-tribe restart`

   ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

   `crm resource restart elasticsearchTribeEverywhere`

   The following examples show an environment with three zones, where the first zone is a single server, the second zone is a three node cluster and the third zone is another single server.

Example from the master zone of the multi-zone configuration:

```
[root@wavd-doc01 ~]# cat /etc/elasticsearch-tribe/elasticsearch.yml

discovery.zen.ping.multicast.enabled: false
tribe.local.cluster.name: Master_Zone
tribe.local.transport.tcp.port: 9312
tribe.local.discovery.zen.ping.unicast.hosts:
  - "wavd-doc01:9300"
http.port: 9201
transport.tcp.port: 9305
node.name: wavd-doc01-tribe

tribe.Slave_Zone.cluster.name: Slave_Zone
tribe.Slave_Zone.transport.tcp.port: 9313
tribe.Slave_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-mcs01:9300"
  - "wavd-mcs02:9300"
  - "wavd-mcs03:9300"

tribe.3rd_Zone.cluster.name: 3rd_Zone
tribe.3rd_Zone.transport.tcp.port: 9314
tribe.3rd_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-news:9300"
```

Example from the first slave zone of the multi-zone configuration:

```
[root@wavd-mcs01 ~]# cat /etc/elasticsearch-tribe/elasticsearch.yml

discovery.zen.ping.multicast.enabled: false
tribe.local.cluster.name: Slave_Zone
tribe.local.transport.tcp.port: 9312
tribe.local.discovery.zen.ping.unicast.hosts:
  - "wavd-mcs01:9300"
  - "wavd-mcs02:9300"
  - "wavd-mcs03:9300"
http.port: 9201
transport.tcp.port: 9305
node.name: wavd-mcs01-tribe

tribe.Master_Zone.cluster.name: Master_Zone
tribe.Master_Zone.transport.tcp.port: 9313
tribe.Master_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-doc01:9300"

tribe.3rd_Zone.cluster.name: 3rd_Zone
tribe.3rd_Zone.transport.tcp.port: 9314
tribe.3rd_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-news:9300"
```

Example from the second slave zone of the multi-zone configuration:

```
[root@wavd-news ~]# cat /etc/elasticsearch-tribe/elasticsearch.yml

discovery.zen.ping.multicast.enabled: false
tribe.local.cluster.name: 3rd_Zone
tribe.local.transport.tcp.port: 9312
tribe.local.discovery.zen.ping.unicast.hosts:
  - "wavd-news:9300"
```

```
http.port: 9201
transport.tcp.port: 9305
node.name: wavd-news-tribe

tribe.Master_Zone.cluster.name: Master_Zone
tribe.Master_Zone.transport.tcp.port: 9313
tribe.Master_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-doc01:9300"

tribe.Slave_Zone.cluster.name: Slave_Zone
tribe.Slave_Zone.transport.tcp.port: 9314
tribe.Slave_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-mcs01:9300"
  - "wavd-mcs02:9300"
  - "wavd-mcs03:9300"
```

Validating the Elasticsearch Configuration

When you finish updating the Elasticsearch configuration file, you can validate your setup by checking the Elasticsearch information for each server.

The example in this section is a multi-zone system consisting of the following:

- wavd-doc01: Multi-zone / master zone, single server

- wavd-news: Multi-zone / slave zone, single server

- wavd-mcs01 - wavd-mcs03: Multi-zone / slave zone, 3-node cluster

To validate the Elasticsearch configuration:

1. Open a browser and navigate to http://*<server_name>*:9200/_plugin/head, where *sever_name* is the host name or IP address of the server in one zone in your configuration — for example, the local zone.

   The Elasticsearch head plug-in Web page opens in your browser.

2. Verify that the local zone and the bindings between zones are listed.
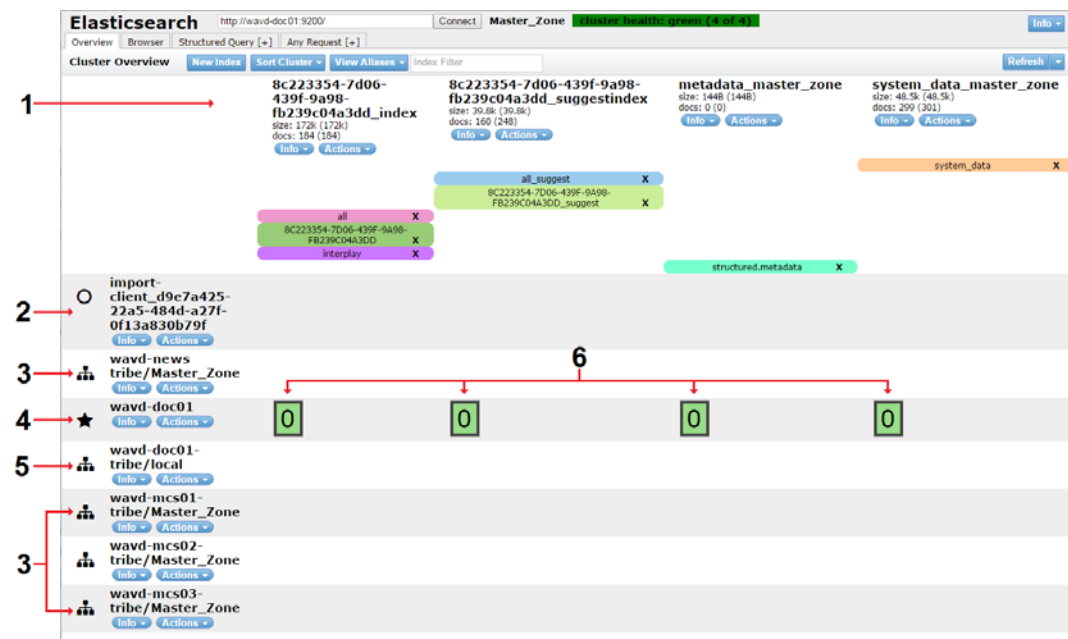
1   Elasticsearch Indexes:

- *<System_ID>*_index - Index that stores assets. The type of index is identified below in the same column. In the example above, this is an index for an Interplay Production system.

- *<System_ID>*_suggestindex - Index that stores "suggest" data. Search suggestions appear below the Search text field when users begin typing.

- metadata_*<zone_name>* - This index stores thesaurus data. This index is also identified by its alias, "structured.metadata".

- system_data_*<zone_name>* - This is a system index that stores information about the indexes registered with Media Index. Permission information and other data is also stored in this index. This index is also identified by its alias, "system_data".

2	The import client listens for updates from the PEBco service. There is one instance of this service in each zone. In cluster configurations, the import client may or may not reside on the same server as the master.
3	Each remote zone tribe is bound to the local cluster. This this example, two remote zones are shown. Note that the tribe name reflects the local zone.
4	The local zone's data node is listed.
5	The local zone tribe is bound to the local cluster.
6	The numbered green boxes represent Elasticsearch shards. The local node in this example (wavd-doc01) has been configured for one primary shard and no replicas. If the system had been configured for two primary shards, the image would show two green boxes (numbered 0 and 1). Shard numbering begins with zero and increments for each additional shard. This remains true for replica shards as well. Configurations that include primary and replica shards show the primary shard outlined in bold.

3. Open a browser and navigate to http://*<server_name>*:9200/_plugin/head, where *sever_name* is the host name or IP address of the server in another zone in your configuration — for example, the remote zone.

   The Elasticsearch head plug-in Web page opens in your browser.

4. Verify that this local zone (the remote zone in your multi-zone configuration) and the bindings between zones are listed.
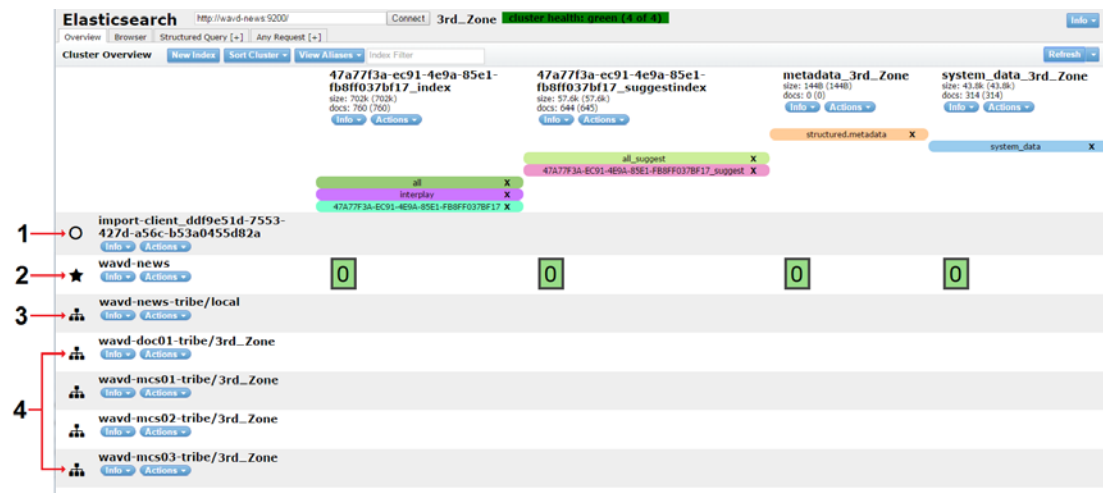
1	The import client listens for updates from the PEBco service. There is one instance of this service in each zone. In cluster configurations, the import client may or may not reside on the same server as the master.
2	The local zone data node is listed.
3	The local zone tribe is bound to the local cluster.
4	Each remote zone tribe is bound to the local cluster. This this example, two remote zones are shown. Note that the tribe name reflects the local zone.

5. Open a browser and navigate to http://<*server_name*>:9201/_plugin/head, where *sever_name* is the host name or IP address of a server in a zone in your configuration — for example, the local zone — and the port number is configured for the tribe node of the zone.

    The Elasticsearch head plug-in Web page for the tribe node opens in your browser.

6. Verify that the tribe node lists the all indexes in your configuration as well as the local and remote data nodes.
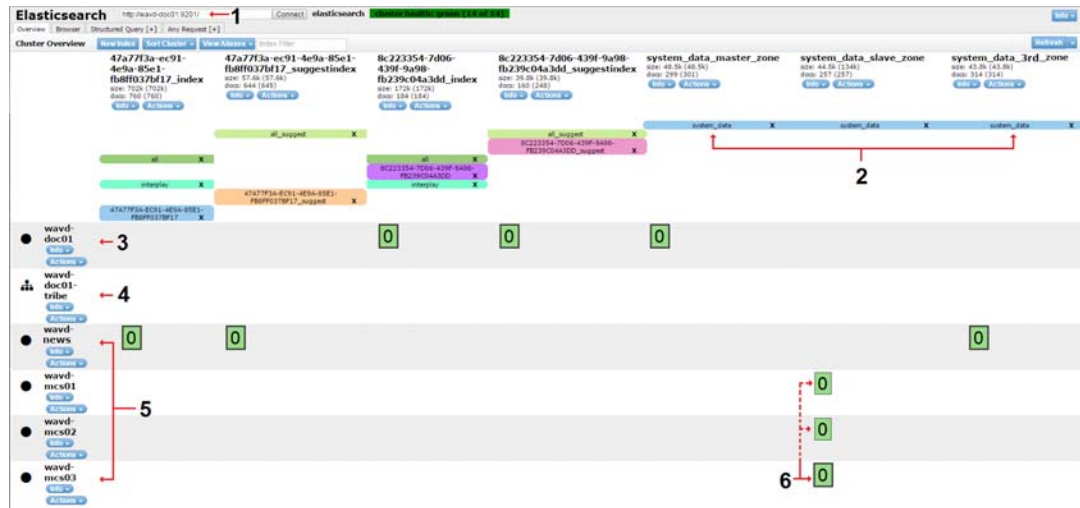
Some data not present do to sizing limitations within this document.

1	Shows the connected tribe node.
2	Indexes from all nodes are listed.
3	The local zone data node is listed.
4	The local zone tribe node is listed.
5	The remote zone data node is listed.
6	In zones with multiple shards, the primary shard is outlined in bold.

5 Installing the Production Engine Bus Connector

The Production Engine Bus Connector (PEBCo) queries the change table of the Interplay Production Engine and delivers that data to the Avid Common Services (ACS) bus where the import service takes these messages and sends the data to the Elasticsearch index. A timestamp is saved on the Interplay Production Engine, and the PEBCo service queries the Interplay Production Engine regularly and uses the timestamp to determine which assets have been already indexed. After each indexing operation, the timestamp is updated.

If you have a cluster configuration, you can use the search-cluster script to set up the necessary search services on a node and configure the PEBCo services (see "Creating a PEBCo Instance on an MCS Cluster" on page 35).

Media Index v2.5 and later supports searching Interplay Archive databases. If you want to use Media Index with Interplay Archive, create a PEBCo instance and configure the bus service for Interplay Production, and then repeat the procedure for Interplay Archive. You must have a separate PEBCo instance for each database.

📄 *For more information on configuring the ACS bus service for Interplay Archive, see the* Interplay | Engine and Interplay | Archive Engine Administration Guide *and the* Avid MediaCentral Platform Services ReadMe.

Configuring PEBco involves the following steps:

- Enabling Sync Events in the Interplay Administrator
- Creating a PEBCo Instance on a Single MCS Server
- Creating a PEBCo Instance on an MCS Cluster
- Configuring the ACS Bus Service in Interplay Administrator
- Configuring and Using the PEBco Service

Enabling Sync Events in the Interplay Administrator

Before you set up your PEBCo instance, you need to enable sync events. Syncing events allows changes in the search properties specified in Interplay Production and Interplay Archive to be pushed to Media Index.

- If your configuration includes Interplay Production v3.3 or later, the Interplay Administrator tool is used to enable sync events.
- If your configuration includes Media Index v2.2.x and Interplay Production v3.2 or earlier, the InterplayEngineUpdateActivation utility on the MCS installation USB key is used to update your registry setting to allow sync events.

To enable the sync events (Media Index v2.3 or later):

1. Start Interplay Administrator and log on to the Interplay Production or Interplay Archive database that you want to be indexed.

2. In the Server section of the Interplay Administrator, click the Server Settings icon.

   The Server Settings view opens.
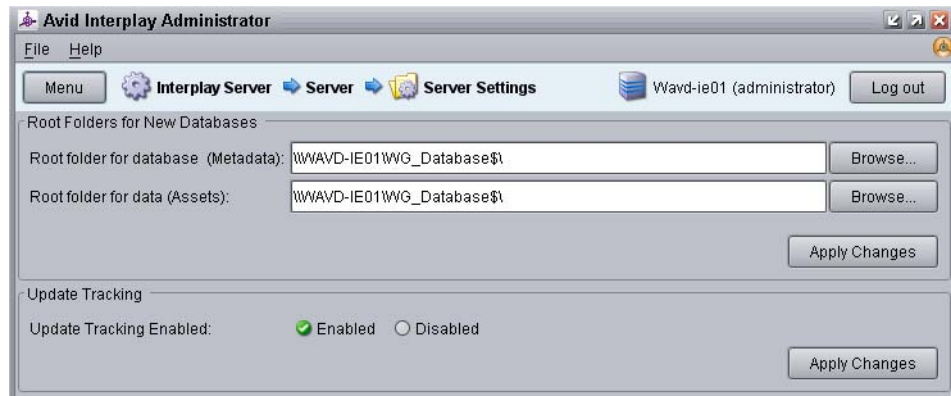
3. In the Update Tracking section, select Enabled to enable sync events.

📄 *You do not need to restart the Interplay Production or Archive Engine.*

4. Click Apply Changes to save the settings change.

5. If you have additional Interplay Production or Interplay Archive Engines, repeat steps 1 through 4 to enable syncing on those systems.

To enable the sync events (Media Index v2.2.x):

1. On the server running the Interplay Production Engine, insert the MCS installation USB key.

2. Navigate to the root directory of the USB key, and then double-click the following utility: InterplayEngineUpdateActivation.exe.

3. Restart the Interplay Production Engine by doing one of the following:

   ▶ For single-server configurations: in the Interplay Administrator tool, click the Restart Server view, and then click Restart.

   ▶ For cluster configurations: in the Cluster Administrator tool, open the Avid Workgroup Server resource group, select Avid Workgroup Engine Monitor, and then change the state to "Offline." Once the engine is offline, you can then right-click the Avid Workgroup Server group (not the resource) and select Bring Online.

Creating a PEBCo Instance on a Single MCS Server

The Production Engine Bus Connector (PEBco) is automatically installed during the MCS installation. Although before the service can be used, it must first be configured.

To create a PEBCo instance:

1. Type the following command to create a PEBCo instance on the MCS server:

   pam-agent-ctrl add <instance_name>

   <instance_name> can be any meaningful name. Avid recommends that you use the MCS host name or Corosync cluster virtual host name so the PEBCo instance can be easily identified.

   A single server will return no information after issuing this command.

2. Verify that the PEBCo instance has been created through the pam-agent-ctrl script:

   pam-agent-ctrl list

   The screen displays a list of all available PEBCo (pam-agent) instances on this machine. The following is a sample output from a single server:

   ```
   [root@wavd-doc01 ~]# pam-agent-ctrl list
   wavd-doc01
   ```

3. Type the following command to start the PEBCo instance:

   pam-agent-ctrl start <instance_name>

   The screen displays the status of all available PEBCo instances on this machine. The following is a sample output:

   ```
   [root@wavd-doc01 ~]# pam-agent-ctrl start wavd-doc01
   Starting pam-agent-service.wavd-doc01:                    [  OK  ]
   ```

4. To verify that the PEBCo instance is running, type the following command:

   pam-agent-ctrl status

   The command should return a confirmation that the service is running, similar to the following:

   ```
   [root@wavd-doc01 ~]# pam-agent-ctrl status
   Instance 'wavd-doc01'                    (pid  85540) is running...
   ```

5. If your configuration includes an Interplay Archive Engine, a separate PEBco instance must be created. Repeat steps 1-4 to create an additional instance.

Creating a PEBCo Instance on an MCS Cluster

The Production Engine Bus Connector (PEBco) is automatically installed during the MCS installation. Although before the service can be used, it must first be configured.

To create a PEBCo instance:

1. Type the following command on any cluster node to create a PEBCo instance on the MCS cluster.

   pam-agent-ctrl cluster add <instance_name>

   <instance_name> can be any meaningful name. Avid recommends that you use the MCS host name or Corosync cluster virtual host name so the PEBCo instance can be easily identified.

> *In a cluster, the addition of the* `cluster` *switch to this command adds the PEBco instance to all nodes simultaneously.*

A cluster configuration should return a response similar to the following:

```
[root@wavd-mcs01 ~]# pam-agent-ctrl cluster add wavd-mcs

Add resource: AvidPamAgent-wavd-mcs, Broker Host: 192.168.10.50
```

2. Verify that the PEBCo instance has been created through the `pam-agent-ctrl` script:

   `pam-agent-ctrl cluster list`

   The screen displays a list of all available PEBCo (pam-agent) instances on this machine. The following is a sample output from a cluster configuration:

```
[root@wavd-mcs01 ~]# pam-agent-ctrl cluster list

Cluster resources:
----------------------------------
 AvidPamAgent-wavd-mcs  (ocf::avid:pam-agent):  Started
----------------------------------

Note: 'AvidPamAgent-' is the common prefix for pam agent cluster resources.
```

3. Use the Cluster Resource Monitor to verify that a new resource has been added to Pacemaker:

   `crm_mon -f1`

   A new resource similar to the following should be found near the end of the output:

```
 AvidPamAgent-wavd-mcs  (ocf::avid:pam-agent):  Started wavd-mcs01
```

4. Type the following command to start the PEBCo instance:

   `crm resource start AvidPamAgent-<instance_name>`

5. To verify that the PEBCo instance is running, type the following command:

   `pam-agent-ctrl status`

   The screen displays the status of all available PEBCo instances on this machine. The following is a sample output:

```
[root@wavd-mcs01 ~]# pam-agent-ctrl status

Instance 'wavd-mcs'                    (pid  117058) is running...
```

6. If your configuration includes an Interplay Archive Engine, a separate PEBco instance must be created. Repeat steps 1-5 to create an additional instance.

Configuring the ACS Bus Service in Interplay Administrator

The PEBCo configuration requires a series of bus calls to the MCS server. For the configuration to be successful, the Interplay Production and /or Interplay Archive Engines must have the proper bus URL so that the Interplay Administrator can configure the service correctly.

To configure the ICS bus service:

1. Start Interplay Administrator and log on to the Interplay Production database that you want to have indexed.

2. In the Site Settings section of the Interplay Administrator, click the MediaCentral Platform Services Settings icon.

📄 *In versions of Interplay Production prior to v3.6, the setting is located in the Server Hostname Settings section of the tool.*

3. In the ACS Bus Service Settings section, type the name of the Bus AMQP (Advanced Message Queuing Protocol) URL.
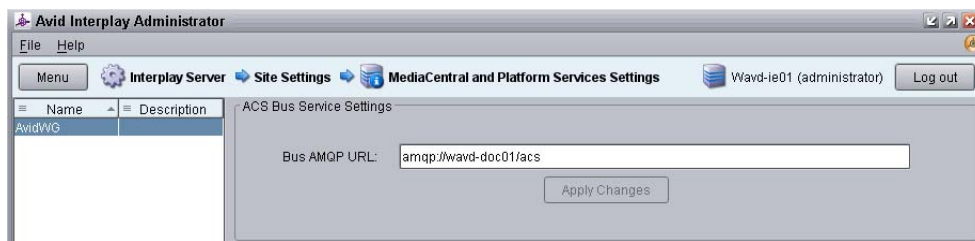
The bus URL can take the following form:

- For single-server configurations: `amqp://<MCS-server-IP-or-hostname>/acs`

- For cluster configurations: `amqp://<virtual-cluster-IP-or-hostname>/acs`

If you do not configure this setting correctly, Interplay Administrator cannot communicate with the bus on the MCS server and no data displays in the Production Engine Bus Connector settings since the needed services cannot be queried. Instead, an error message appears: `ACS Bus error: Error connecting to the bus (Make sure ACS Bus URL is correctly configured).`

4. Click the Apply Changes button.

5. (Optional) If your configuration includes Interplay Archive, log out of the Interplay Production database and repeat the procedure to configure the bus service for the Archive system.

Configuring and Using the PEBco Service

Every PEBCo instance is bound to no more than one Interplay Production or Interplay Archive database.

You configure the PEBCo service in the Production Engine Bus Connector view in the Interplay Administrator. This view has two tabs:

• The Manage/Status tab provides the following:

- Fields that let you assign an instance of the PEBCo service to your Interplay Production or Interplay Archive database.

- Buttons that let you start, stop, and resync the PEBCo service.

- Information about the status of the instance which is assigned to the selected database.

Before you can configure the PEBCo service, you must configure the ACS Bus URL in the Interplay Administrator.

• The Canonical Data tab lets you select the Properties used in the indexing process and that users can select when performing an indexed search in MediaCentral UX. Some properties are part of the standard data model. These properties are already selected in the Canonical Data tab, and you cannot deselect them.

Before you can start the first indexing operation, you must specify the Canonical Data set in the Canonical Data tab.

⚠️ **When you change the database properties, a re-index operation occurs to adapt the index structure to your change. This runs in the background but does not add data from these new properties to the existing data within the index. Only newly indexed assets contain the data from this property. To have this property available for all assets, you need to use the Resync button on the Manage/Status tab. This resync might take a lot of time for bigger databases, so make changes to your database properties carefully.**

Indexed search maps Interplay Production properties to Media Index criterion fields. The following table describes how Media Index categorizes some common Interplay Production properties:

Interplay Production Property	Media Index Field
Name	Name
Created By	Creator
Creation Date	Created
Modified Date	Modified
Changed By	Modifier
Start	Start
End	End
Duration	Duration
Comments	Description
Mime Type	Type

Due to differences in the properties and fields used by Interplay Production and Media Index, some duplication appears in the criteria available for indexed searches with both Interplay Production and Media Index fields appearing. If you add the wrong fields as criteria, your search might not return any usable results.

If your configuration includes multiple systems that use the same custom field ID, once you mark a field ID as hidden in one system, the field will be not searchable on any of the systems. This might be an issue when fields are removed in a multi-zone setup with multiple Interplay Production systems. If you add a field and then remove it for one Interplay Production system, then it is hidden for all other systems as well. Also, if the field IDs (keys) are the same in different systems, this could cause an issue. Avid recommends that you keep the selected fields consistent across all systems of the same type (Interplay Production, Interplay Archive, iNEWS, MAM). If this is not possible, then you need to re-index the index of the affected systems.

To select database properties that can be used in an indexed search:

1. Click the Canonical Data tab.

2. Select the properties that you want to be available when using the MediaCentral UX Search pane.

📑 *You must select at property set before you can start the indexing process*

▸ Use individual checkboxes, the Select All button, or the Clear All button. You should select just the set of data that users will search for and not rely on the Select All button.

> ▶ Click the Sync with Property Layout button to match these selections with those in the Interplay Administrator Property Layout view.

📄 *Media Index base properties and common properties are selected by default, and you cannot deselect these properties.*

3. Click Apply to save your selections and to make them available in MediaCentral UX.

To configure an instance of the PEBCo service:

1. Start Interplay Administrator and log on to the Interplay database that you want to be indexed.

2. In the Site Settings section of the Interplay Administrator, click the Production Engine Bus Connector icon.

📄 *The Production Engine Bus Connector view is not available on Macintosh clients.*

3. Click the Manage Status tab.

4. In the Assign PEBCo Service section, do the following:

   a. Click the Instance menu and select the PEBCo instance name you want to use.

   b. Enter a user name and password for access to the Interplay Production or Archive Engine.

      As a best practice, create a separate user and password for this purpose (in the User Management settings) that you can easily identify in case troubleshooting is necessary. Avid recommends creating an Interplay user called "pebco" with administrative privileges. This can make it easier to track down issues in log files for debugging issues.

📄 *The PEBCo settings respect privileges set for users logging on to Interplay Production and Interplay Archive. Users can only view assets in search results for which they have permissions.*

5. Click Assign.

   The Interplay Administrator should reply with the following message:

   Service instance assigned successfully

   If you need to remove the service from the selected Interplay Production or Archive Engine, click UnAssign.

To manage an instance of the PEBCo service:

▶ In the Manage section of the Manage/Status tab, click one of the following:

   - Start to start the service.

   - Stop to stop the service.

   - Resync to update the whole database content independent from any existing saved timestamp. Resyncing the database might take a long time to complete and invalidates the existing content of your index. During this process, the search only returns assets that are already processed within this run.

To update the status and statistics for the PEBCo service:

▶ In the Status section, click the Refresh button.

   The information includes requests sent from the PEBCo to the import service for indexing the data.

6 Installing and Configuring Production Services for MediaCentral Delivery

The Production Automation service and the Interplay Consolidate service are used for the MediaCentral Delivery feature. The following topics describe installation and configuration of these services:

- Understanding Production Services for MediaCentral Delivery
- Check List for Installing and Configuring Automation and Consolidate Services
- Prerequisites for Installing and Configuring Automation and Consolidate Services
- Installing the Automation and Consolidate Services
- Registering the Automation and Consolidate Services
- Registering a Provider with the Production Services Engine
- Creating a Consolidate Profile

MediaCentral Delivery requires an Interplay Delivery provider. For information about MediaCentral Delivery, see the *Avid MediaCentral | UX User's Guide*. For information about Interplay Delivery, see the *Interplay | Production Services Setup and User's Guide*.

Understanding Production Services for MediaCentral Delivery

A Production Services Engine and three Interplay Production Services are needed for delivering media and assets between Interplay Production workgroups or MediaCentral zones:

- Interplay Production Services Automation: This service handles the logic of chaining a consolidate job and a delivery job. It triggers a consolidate job and monitors its status. If the consolidate job is successful, it triggers a delivery job.
- Interplay Consolidate: This service uses the In and Out marks specified for a loaded clip or subclip to create a new media and asset. This service is used if a user selects "Deliver from Mark In to Mark Out" in either the Deliver To dialog box or the Deliver To Me dialog box.
- Interplay Delivery: This service executes the delivery job.

The following illustration shows the Production Services components and a Delivery workflow.
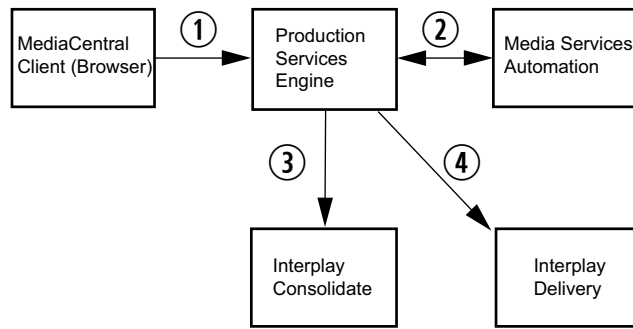
1. The user initiates a Delivery request.

2. The Production Services Engine forwards the request to the Production Services Automation service, which determines the next step.

3. If the request includes In to Out marks, the Automation service tells the Production Services Engine to send the request to the Consolidate service. The Consolidate service creates a new asset and new media.

4. The Automation Service tells the Production Services Engine to initiate a Delivery job.

Check List for Installing and Configuring Automation and Consolidate Services

The following table provides a check list of steps for installing and configuring the Interplay Production Services that support MediaCentral UX Delivery. The check list also provides references where to find more information about each step.

Production Services Check List

Task	Section Reference
❏ Check the prerequisites	"Prerequisites for Installing and Configuring Automation and Consolidate Services" on page 43.
❏ Install the Production Services Automation service, the Interplay Consolidate service, and Interplay Access.	"Installing the Automation and Consolidate Services" on page 44.
❏ Verify that each service is running.	"Installing the Automation and Consolidate Services" on page 44.
❏ Register each service.	"Registering the Automation and Consolidate Services" on page 44.
❏ Configure each service to communicate with the Production Services Engine.	"Registering a Provider with the Production Services Engine" on page 46
❏ Verify that each service is connected.	"Registering a Provider with the Production Services Engine" on page 46

Production Services Check List(Continued)

Task	Section Reference
❏ Connect the Avid shared storage client to the System Director and mount the required workspaces.	Avid shared storage client documentation.
❏ Create at least one Interplay Consolidate profile. No profile is needed for Production Services Automation.	"Creating a Consolidate Profile" on page 50

Prerequisites for Installing and Configuring Automation and Consolidate Services

Where to install: Production Services Automation and Interplay Consolidate can be installed on any server currently configured in an Interplay Production workgroup except for the servers hosting the following components:

• Interplay Engine

• Interplay Archive Engine

Avid recommends installing Production Services Automation on the same server as the Production Services Engine. In a high-workload facility, Avid recommends installing Interplay Consolidate on a separate server. Do not install Interplay Consolidate on the same server as Interplay Transcode.

Hardware requirements: Production Services Automation and Interplay Consolidate share the same hardware requirements:

• Minimum 512 MB RAM

• Approximately 380 MB of hard drive space.

• Port 1099 for communication with the Production Services Engine.

📑 *A license is not required for these services.*

Software requirements: Production Services Automation and Interplay Consolidate share the same hardware requirements:

• Windows 7 or Windows Server 2012 Standard R2

• Avid Service Framework (ASF)

• Avid shared storage client software

Production Services Engine: Production Services Automation and Interplay Consolidate require network access to an installed and configured Production Services Engine. For more information about installing and configuring a Production Services Engine, see the *Interplay | Production Services Setup and Installation Guide*.

Interplay Delivery Service: The Delivery service is configured to work with the Production Services Engine and provides the profile for transferring the assets. For more information about Interplay Delivery, see the *Interplay | Production Services Setup and Installation Guide*.

⚠ **When you create a Delivery profile for MediaCentral, use a network path to an Avid shared storage share for the Temporary Workspace on Sender parameter. For example, use \\avidstorage_server\Delivery_temp. This storage space is used for temporary media files created by the Interplay Consolidate service. If you use a local folder, the large amount of media created can cause system problems for the Delivery server. The files will automatically be deleted when the Delivery job is completed.**

Installing the Automation and Consolidate Services

You install both services from the Interplay Production Servers installer.

📝 *Starting with Interplay Consolidate v3.6, separate installation of Interplay Access is not required.*

To install the Production Services Automation service:

1. Select the following from the Interplay Server Installer Main Menu:

   Servers > Production Services > Install Media | Index Support

   The Install Media | Index support window opens.

2. Click Production Services Automation.

   The Production Services Automation installer window opens.

3. Click Next.

4. Accept the license agreement and click Next.

5. Follow the prompts to install the service.

   The service is started automatically by the installer after the installation is complete.

To install the Interplay Consolidate service:

1. Select the following from the Interplay Server Installer Main Menu:

   Servers > Production Services > Install Media | Index Support

   The Install Media | Index support window opens.

2. Click Consolidate.

   The Interplay Consolidate installer window opens.

3. Install the Sentinel USB Driver (if prompted).

4. Install the PACE Interlok Driver (if prompted)

5. Accept the license agreement and click Next.

6. Follow the prompts to install the service.

   The service is started automatically by the installer after the installation is complete.

Registering the Automation and Consolidate Services

To register the Automation and Consolidate services, you need to open the Production Services and Transfer Status tool and install the service description. You only need to install a service description once, even if you configure multiple providers.

📝 *Registration of other Production Services is automatic.*

To install service descriptions:

1. Start the Production Services Engine and log in as administrator.

2. Click the Admin Tool button.

   The Interplay Production Services and Transfer Status tool opens.

3. Click the Services tab.

   The Services page displays the currently configured services.

| Production Services Login/Logout | Jobs | Providers | Services | Users | Profiles |

Name	Version	Manufacturer	Description
Avid Interplay Archive Service	1.3.3	Avid Technology, Inc.	Avid Interplay Archive Service
Avid Interplay Copy Service	3.1	Avid Technology, Inc.	Avid Interplay Copy Service
Avid Interplay Delivery Service	3.1	Avid Technology, Inc.	Avid Interplay Delivery Service

Help   Install/Upgrade   Delete   Details

4. Click Install/Upgrade.

   The Install/Upgrade Service dialog box opens.

Install/Upgrade Service

Package [   ] Browse...

Install/Upgrade   Cancel
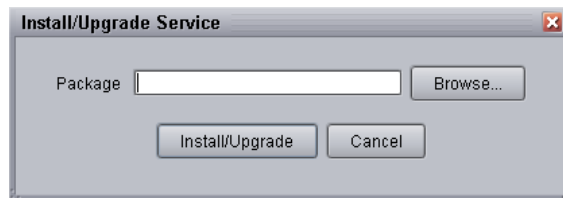
5. Click the Browse button and navigate to the folder containing the service package (.zip file). Make sure you have access to the folder.

   The following are the default locations of the service packages:

   - C:\Program Files\Avid\Interplay Consolidate\TranscodeService\ConsolidateService.zip

   - C:\Program Files\Avid\Interplay Production Services Automation\MediaServicesAutomation.zip

6. In the folder, select the *service*.zip file.

7. Click Save.

   The path to the file appears in the Install/Upgrade Service dialog box.

Install/Upgrade Service

Package [codeService\ConsolidateService.zip] Browse...

Install/Upgrade   Cancel
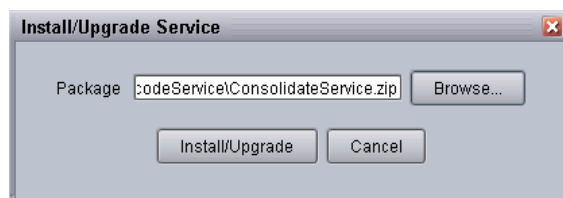
8. Click Install/Upgrade.

The service and its description appear on the Services page. The following illustration shows both services after they are registered.

Registering a Provider with the Production Services Engine

A service that is installed and registered is referred to as a *provider*. Starting with Interplay Production Services v3.1, some service providers operate as Windows services, and others operate as Windows applications.

• Interplay Production Services Automation operates as a Windows service. See "Registering a Production Services Automation Provider" on page 46

• Interplay Consolidate operates as a Windows application. See "Registering and Connecting a Consolidate Provider" on page 48

Registering a Production Services Automation Provider

Production Services Automation runs as a Windows service. Information about the service is included in an .ini file.

• If you are performing a new installation, the installed .ini file uses "localhost" as the name of the associated Production Services Engine. This is appropriate if you are installing the service on the same system as the Production Services Engine. If you are installing the service on a different system, you need to change the default hostname by editing the .ini file.

• If you are upgrading from a previous version, the existing .ini file is copied to the appropriate new folder. If the associated Production Services Engine remains the same, no further configuration is needed.

To register a Production Services Automation provider with the Production Services Engine:

1. Locate the .ini file and open it in a text editor, such as Windows Notepad.

The following table lists the services with the folder that holds the .ini file and the log files.

Service Name	Folder for .ini File and Log Files
Interplay	Production Services Automation	C:\ProgramData\Avid\Production Services Automation\DMSCMService.ini

2. Substitute the hostname of the Production Services Engine.

   Following is an example of an Automation service .ini file with the default hostname "localhost":

   ```
   #Mon May 05 11:34:54 EDT 2014
   @5%?MaxJobs=1
   @2%tMedia_Services_Engine_Hostname=localhost
   @1%tProvider_Name=a51-WG6_Consolidate_0123
   ```

   Following is the same .ini file, edited to use the hostname "a51-MSE."

   ```
   #Mon May 05 11:34:54 EDT 2014
   @5%?MaxJobs=1
   @2%tMedia_Services_Engine_Hostname=a51-MSE
   @1%tProvider_Name=a51-WG6_Consolidate_0123
   ```

   The provider name is automatically created by adding a number string to the local hostname. You can change this name if you want.

3. A default name for the provider is automatically created by the provider software using the following syntax:

   <hostname>_<servicename>_<unique#>

   The provider name in the previous example is a51-WG6_Consolidate_0123

   Accept the default name or type a new name. This name appears in the Production Services and Transfer Status tool.

4. Save and close the .ini file.

   The service automatically reads the .ini file and tries to connect to the Production Services Engine. If the service fails to connect, it continues to try to connect every thirty seconds.

To verify that the Production Services Automation service provider is running and connected:

1. Start the Production Services Engine.

2. Click the Admin Tool button.

   The Interplay Production Services and Transfer Status tool opens.

3. Click the Providers tab.

   The system displays the status of the Production Services providers. In the following illustration, all services are connected, indicated by a check mark in the Status column. Production Services Automation is circled.

You can also view the services in the Services tab of the Computer Management window.
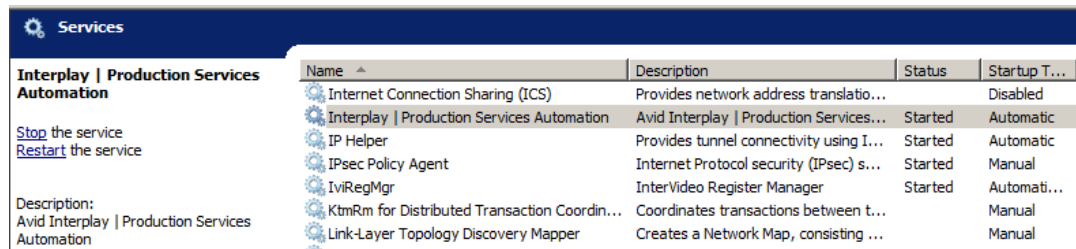
Registering and Connecting a Consolidate Provider

Interplay Consolidate runs as a Windows application. You need to register the provider by connecting to the Production Services Engine.

If you try to connect to the Production Services Engine before the latest service is registered, the Status line in the Transcode Service dialog box reads:
Error From Broker! UNKNOWN_SERVICE.

To connect the Consolidate provider to the Production Services Engine:

1. Click Start and select Programs > Avid > Avid Interplay Consolidate.

   The Consolidate Service dialog box opens.
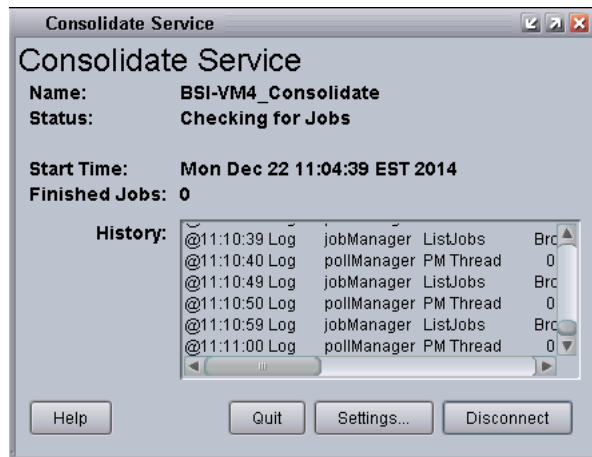
2. Click Settings.

   The Provider Settings dialog box opens.

3. Do the following:

   a. Provider Name — A default name for the provider is automatically supplied. Accept the default name or type a new name. In this example, the name is Consolidate_01.

   b. Production Services Engine Host Name — Type the name of the system running the Production Services Engine application.

   c. Automatically Connect — Select Yes to automatically connect the provider to the Production Services Engine when the application starts. To prevent automatic connection, select No.

4. Click OK.

5. Click Connect in the Service window.

   The Consolidate Service dialog box now shows that the service is connected and shows the name of the provider you connected.
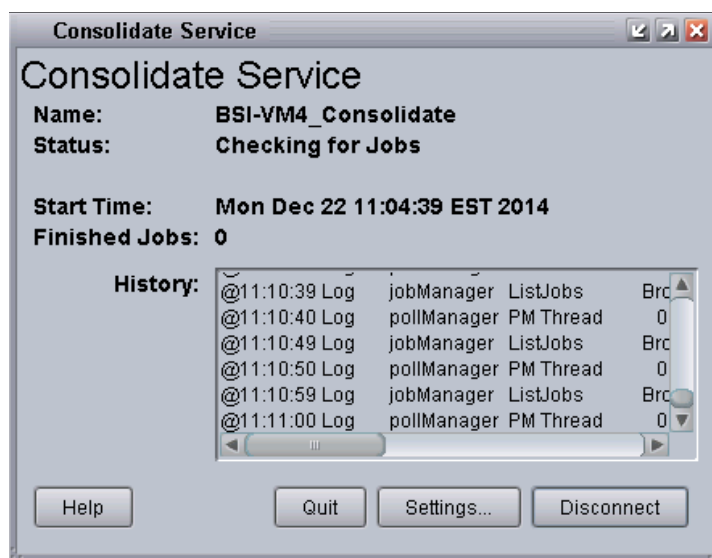
The Provider page in the Production Services and Transfer Status tool now shows that the Consolidate service is connected, indicated by a check mark in the Status column.
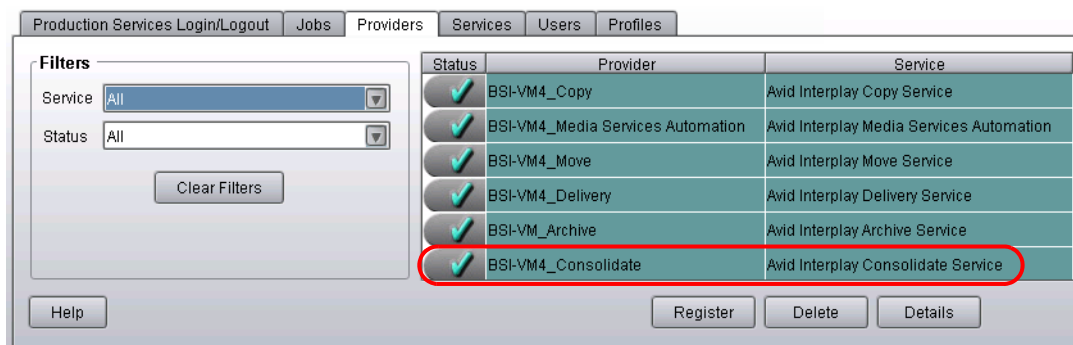
Creating a Consolidate Profile

If users in your MediaCentral zone are going to perform MediaCentral Delivery operations that use "Deliver from Mark In to Mark Out," an administrator must create at least one Consolidate profile. This default profile must be named "default_for_delivery." The target video quality for this profile is displayed in the Deliver or Deliver To dialog box, and is used if the user selects "Deliver from Mark In to Mark Out." and the Delivery profile does not include a specific Target Video Quality.

The following illustration shows a profile with DV 25 420 as the TargetVideoQuality.
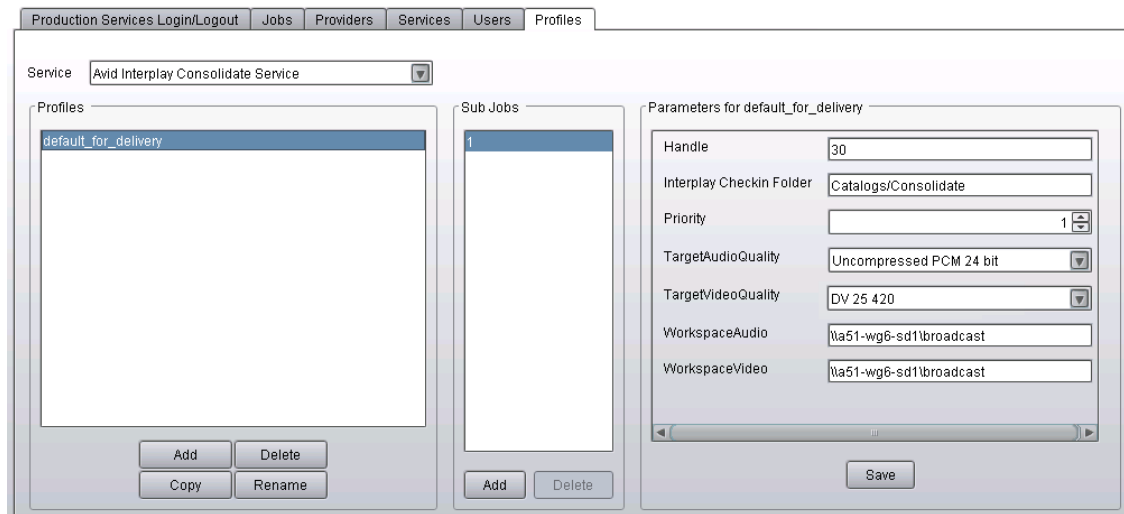
The following illustration shows the Deliver To dialog box, with Deliver from Mark In to Mark Out selected, and the Target Video Quality as DV 25 420. The Target Video Quality is taken from the Consolidate profile. This resolution will be used for the new media, because the Target Video Quality specified in the Delivery profile is "highest" rather than a specific resolution.
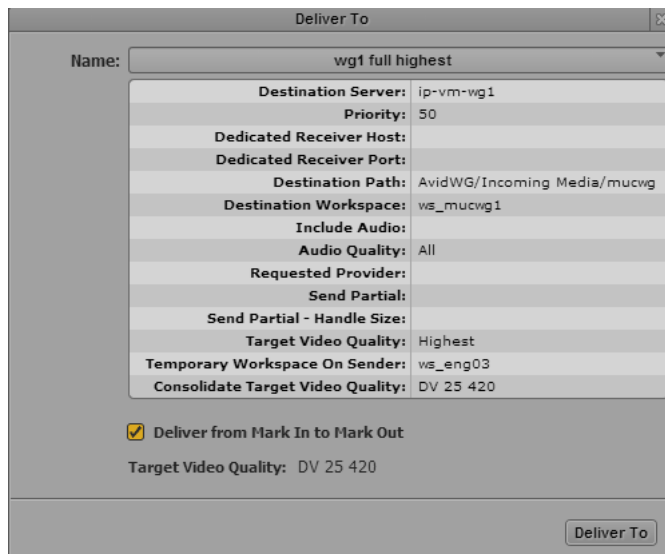
Other options in the Consolidate profile, such as Handle and Interplay Checkin Folder, are also used for Deliver from Mark In to Mark Out operations. However, they are not shown in the Deliver To or Deliver to Me dialog box.

To create a Consolidate profile:

1. Start the Production Services Engine.

2. Click the Admin Tool button.

   The Interplay Production Services and Transfer Status tool opens.

3. Click the Profiles tab.

4. In the Service menu, select Avid Interplay Consolidate Service.

5. Click Add in the Profiles area.

   The Add Profile dialog box opens.
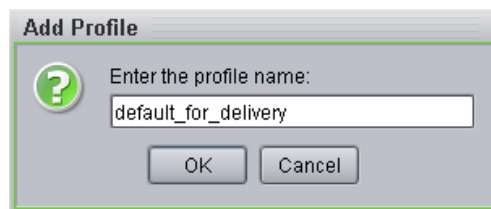
6. Type "default_for_delivery" in the Add Profile dialog box.

7. Click OK.

   The name appears in the Profiles list and an empty template appears in the Parameters area.

8. In the Parameters area, define the values you want in the profile.

   The following table describes the parameters for a Consolidate profile.

Consolidate Profile Settings

Parameter	Description
Handle	The number of additional frames you want to include on both ends of a subclip or a cut within a sequence. If the handle length exceeds either end of the master clip, then the handle is truncated to the amount of source media available.
Interplay Checkin Folder	The folder in the database that will contain the consolidated asset (metadata). For example, Catalogs/Consolidated.
Priority	A value allows you to assign job priorities to different profiles. Priority numbers range from 1 (highest priority) through 100 (lowest priority). The default priority number assigned to each job is 50.

Consolidate Profile Settings(Continued)

Parameter	Description
TargetAudioQuality	Select an audio resolution for the consolidate operation: • None • MPEG1 Layer 2: Digital audio compressed to the MPEG1 Layer 2 specification at 96 Kb/s. • Uncompressed PCM: 16-bit 48-kHz digital audio • Uncompressed PCM 24-bit: 24-bit 48-kHz digital audio
TargetVideoQuality	Select a video resolution for the consolidate operation. Select None if you do not want the transcoded clip to include a video track.
WorkspaceAudio	The name of the Avid shared-storage workspace that will hold the new audio media files. Use the following syntax: *hostname**workspace_name*
WorkspaceVideo	The name of the Avid shared-storage workspace that will hold the new video media files.

9. Click Save in the Parameters area.

   The Save Profile dialog box opens.

10. Click Yes to save your changes.

For information about using MediaCentral Delivery, see the *Avid MediaCentral | UX User's Guide*.

7 Configuring iNEWS for Media | Index

When integrating iNEWS with Media Index, the iNEWS system relies on the sync agent in the Mediaindex utility program to push stories to Media Index. It must be configured to control what data gets indexed. To specify which Media Index the iNEWS system should use also requires configuration within iNEWS.

This integration first appeared with the release of iNEWS v5.2, for which the ftsindex utility program was used to drive indexing for both Media Index and Fast Text Search (FTS), a feature of iNEWS. Beginning with iNEWS v5.3, however, indexing for Media Index was decoupled from that of FTS. This chapter has been updated to provide set up for iNEWS v5.3 and later. For information on how to set up integration to work with the earlier 5.2 version of iNEWS, refer to the *iNEWS & Media | Index Integration Guide* appended to the *iNEWS v5.2 ReadMe*.

The following table summarizes the differences between the Media Index configuration for iNEWS v5.2 and iNEWS v5.3:

iNEWS v5.2.x	iNEWS v5.3 and later
The ftsindex utility program pushes to both FTS server and the BUS (for Media Index).	The ftsindex utility program pushes to FTS server only; The Mediaindex utility program pushes to the BUS (for Media Index).
Indexes remain separate, as defined by separate tokens in /site/dict/words: W_BINDFTSI for the FTS server port and W_BINDBUS for the MediaCentral system's bus server port.	Indexes remain separate, as defined by separate tokens in /site/dict/words: W_BINDFTS for the FTS server port and W_BINDBUS for the MediaCentral system's bus server port.
Queue for index requests is shared, as defined by token in /site/dict/queues: SYSTEM.INDEX is the queue as defined by the Q_INDEX token.	Queue for index requests are separate, as defined by tokens in /site/dict/queues: SYSTEM.INDEX is still the queue for handling FTS index requests, as defined by Q_INDEX, but now SYSTEM.MEDIA-INDEX is the queue for handling index requests for Media Index, as defined by the Q_MEDIA_INDEX token.
The server's device number, notify value, and mailbox number are the same for both FTS Index and Media Index — in this case 120 — and must match the mailbox assigned to the SYSTEM.INDEX queue.	The server's device number, notify value, and mailbox number are the same — in this case 120 — and must match the mailbox assigned to the SYSTEM.MEDIA-INDEX queue. But for Media Index, the server's device number, notify value, and mailbox are changed — in this case 122 — and must match the mailbox assigned to the SYSTEM.MEDIA-INDEX queue.

Logging was enabled for both indexes by creating a file in /site/env/ftsindex with the following on separate lines:
 	
`DEBUG=7`
`DEBUGFILE=/tmp/ftsindex`
	Logging for Media Index is enabled by creating a file in /site/env/mediaindex with the following on separate lines:
	
	`DEBUG=7`
	`DEBUGFILE=/tmp/mediaindex`

iNEWS v5.2.x	iNEWS v5.3 and later

Config file (/site/config) needed a line in the host section and a line in the SEEK & FTS SERVERS section for ftsindex:
 	
`host abc b <section>`
 	
`servers 105 ; keyword = 2`
 	
`servers 113:115 ; seek = 5`
 	
`servers 120:121 ; fts = 2`
 	
`servers 200:207 ; action, distribution = 20`
 	
`; SEEK & FTS SERVERS`
 	
`;`
 	
`server 113 seek 110 - ; seek server`
 	
`server 114 seek 110 - ; seek server`
 	
`server 115 seek 110 - ; seek server`
 	
`server 120 ftsindex 120 - ; Fast Text Indexing`
 	
`server 121 ftsseek 121 - ; Fast Text Search`
	Config file (/site/config) needs a separate line in both sections for Media Index:
	
	`host abc b <section>`
	
	`servers 105 ; keyword = 2`
	
	`servers 113:115 ; seek = 5`
	
	`servers 120:121 ; fts = 2`
	
	`servers 122 ; media-index = 1`
	
	`servers 200:207 ; action, distribution = 20`
	
	`; SEEK & FTS SERVERS`
	
	`;`
	
	`server 113 seek 110 - ; seek server`
	
	`server 114 seek 110 - ; seek server`
	
	`server 115 seek 110 - ; seek server`
	
	`server 120 ftsindex 120 - ; Fast Text Indexing`
	
	`server 121 ftsseek 121 - ; Fast Text Search`
	
	`server 122 mediaindex 122 - ; Media-Index`

⚠ **In a multi-zone configuration, all iNEWS servers must belong to the same iNEWS Community.**

The following steps are required to configure Media Index on the iNEWS Server:

- Manually Installing the log4cpp File
- Editing iNEWS Site Files
- (Optional) Enabling Logging
- Assigning the Media | Index Attribute
- Activating the Mediaindex Program

📄 *These procedures should only be performed by qualified Avid personnel.*

This chapter only includes the steps related to Media Index integration. Procedures for configuring Fast Text Search are not represented. For more information about FTS, see the "Fast Text Search Servers" section in the *Avid iNEWS Setup and Configuration Guide*.

📄 *When you finish configuring iNEWS for Media Index, you must restart the Mediaindex program. For more information, see "Activating the Mediaindex Program" on page 58.*

Editing iNEWS Site Files

The following site files must be edited on the iNEWS servers to enable Media Index functionality:

• /site/dict/words

• /site/dict/queues

• /site/config

The Mediaindex program must be pointed to the message broker by adding a definition in the /site/dict/words file:

• W_BINDBUS – Identifies the binding address for the MediaCentral platform's bus server port, which is necessary for central indexing as part of iNEWS integration with Media Index.

If it does not already exist, you might also need to create the queue that holds indexing requests, and to point the Mediaindex program to it by adding a definition in the /site/dict/queues file:

• Q_MEDIA_INDEX – Identifies the SYSTEM.MEDIA-INDEX queue to hold indexing requests necessary for central indexing as part of iNEWS integration with Media Index.

Before editing any system file, Avid recommends you make a backup copy of the file.

To edit the Words dictionary file:

1. From the PuTTYCS application, select the PuTTY Filter created for sending server commands to *all* servers, such as iNEWS Consoles.

2. Use the line editor, *ed*, to modify the /site/dict/words file:

   e**d /site/dict/words**

📄 *If you do not know how to use ed to modify lines in the file, please see "The Line Editor, ed" in the Avid iNEWS Setup and Configuration Guide.*

3. Set the binding address for the MediaCentral system's bus server port by adding the definition for W_BINDBUS. The format is: *<server>:<tcp port number>*. The *<server>* is the hostname of the Avid Common Services (ACS) bus, and *<port>* is its dedicated bus server port number. The default is busserver:5672. You can use a hostname or IP address in place of *<server>*. The port number should be set to 5672.

   For example: **W_BINDBUS /busserver:5672** or **W_BINDBUS /10.105.242.34:5672**

To edit the Queues dictionary file:

1. From the PuTTYCS application, select the PuTTY Filter created for sending server commands to all servers, such as iNEWS Consoles.

2. Use the line editor, *ed*, to modify the /site/dict/queues file:

   ed /site/dict/queues

📄 *If you do not know how to use* ed *to modify lines in the file, please see "The Line Editor,* ed*" in the* Avid iNEWS Setup and Configuration Guide*.

3. Name the queue for index requests by adding the definition for Q_MEDIA_INDEX — for example, **Q_MEDIA_INDEX /system.media-index**

📄 *If the queue SYSTEM.MEDIA-INDEX does not exist, create it. For information on how to create a queue, see the "Creating a New Queue" section in the* Avid iNEWS Setup and Configuration Guide.

To edit the iNEWS configuration file:

1. From the PuTTYCS application, select the PuTTY Filter created for sending server commands to all servers, such as iNEWS Consoles.

⚠️ **Always back up the /site/config file before making any changes.**

2. Use the line editor, *ed*, to modify the configuration file:

   ed /site/config

   After you press Enter, the editor responds by displaying a numerical value indicating the file size expressed as the number of characters, including spaces and returns.

📖 *You can check the iNEWS configuration file (/site/config) for Media Index servers, their assigned device numbers, and whether they are configured or not by using the list c command at the console, in the following format:* list c mediaindex.

   a. Ensure there is a line for the Mediaindex program and its device number in the host section of the /site/config file.

   b. Also ensure there is a line for the Mediaindex program and its device number in the SEEK & FTS SERVERS section of the /site/config file.

   For instance, add lines as shown in the following example in **bold**:

   ```
   host abc b <section>
   servers 105 ; keyword = 2
   servers 113:115 ; seek = 5
   servers 120:121 ; fts = 2
   ```
   servers 122 ; media-index = 1
   ```
   servers 200:207 ; action, distribution = 20


   ; SEEK & FTS SERVERS
   ;
   server 113 seek 110 - ; seek server
   server 114 seek 110 - ; seek server
   server 115 seek 110 - ; seek server
   server 120 ftsindex 120 - ; Fast Text Indexing
   server 121 ftsseek 121 - ; Fast Text Search
   ```
   server 122 mediaindex 122 - ; Media-Index

   Typically a server's device number, notify value, and mailbox number are the same. Make sure the notify value for Mediaindex that appears in the configuration file matches the mailbox assigned to SYSTEM.MEDIA-INDEX. If not, change it.

📖 *SYSTEM.MEDIA-INDEX is the queue used to hold indexing requests for Media Index, as defined in the token Q_MEDIA_INDEX in /site/dict/queues. If this queue does not exist, create it.*

   For more information about mailbox numbers, see "Assigning a Mailbox to a Queue" in the *Avid iNEWS Setup and Configuration Guide*. For more information on editing the configuration file, see "Configuration File" in the *Avid iNEWS Setup and Configuration Guide*.

For more information on the Q_MEDIA_INDEX token in /site/dict/queues, see Appendix C "Standard Dictionaries," in the *Avid iNEWS Setup and Configuration Guide*.

3. Reconfigure the system.

Manually Installing the log4cpp File

The log4cpp library package for RHEL 6.5 is required for iNEWS to work with Media Index. It is necessary for the sync agent in the Mediaindex program to function with the Avid Common Services (ACS) bus.

📄 *Prior to completing this process, your iNEWS servers must already be running RHEL v6.5.*

To install the log4cpp library package:

1. Copy the RHEL v6.5 ISO to a DVD as an image.

   For more information, see the RHEL website.

2. Insert the DVD into the optical drive on the iNEWS server, and mount it:

   mount /dev/cdrom /mnt

3. Navigate to the Packages folder:

   cd /mnt/Packages

4. Install the log4cpp package on the server:

   rpm -ivh log4cpp-1.0-13.e16.x86_64.rpm

   The system reports the installation progress:

   ```
   warning: log4cpp-1.0-13.e16.x86_64.rpm: Header V3 RSA/SHA256 Signature, key
   ID fd431d51: NOKEY

   Preparing...          #################################################
   [100%]

    1:log4cpp             #################################################
   [100%]
   ```

5. Change directory to / and unmount the DVD drive:

   cd /

   umount /mnt

(Optional) Enabling Logging

Should you need to troubleshoot search integration, before restarting Media Index, enable logging.

To enable logging:

1. Create the file /site/env/mediaindex.

2. In the file, on separate lines, type the following:
   DEBUG=7
   DEBUGFILE=/tmp/mediaindex

3. Save and close the file.

Assigning the Media | Index Attribute

Define the queues or folders in the iNEWS directory that will be monitored by Media Index by applying the Media Index attribute (+mi).

To apply the Media Index attribute to a queue or folder:

▶ From the console, type the dbtraits command in the following format:
   dbtraits <*queue or directory*> +mi

   For example: **dbtraits SHOW.RUNDOWN +mi**

📄 *Database traits may also be assigned to queues or directories via the Directory/Queue Properties dialog box at an iNEWS Workstation. For more information, see "Viewing Database Traits" and "Changing Database Traits" in the Avid iNEWS Setup and Configuration Guide.*

To find out which directories or queues have the Media Index attribute:

▶ From the console, type the list command in the following format:
   list flags=m d

   The system will display a line that shows the database traits flags (SRPlo-LIsUGQSXWFiTMm) and some column headers. Each line underneath the header line represents either a queue or directory, as indicated by the Q or D to start each line of text. A dash appears when the attribute is not assigned to the queue or directory. The lowercase "m" is the last flag, coming before the sortfield column, and indicates the directories or queues with the Media Index attribute assigned to them.

```
 SRPlo-LIsUGQSXWFiTMm sortfield     purge dis mbox directory

D-------Is----X-Fi--m TITLE         P1    D7  -    SHOWS

Q-R-----I-----XW-iT-m TITLE         P3.0  D1  -    WIRES.ALL
```

Activating the Mediaindex Program

After completing all other steps in this chapter, you need to activate the Mediaindex program, which acts as the sync agent for Media Index. To do this, you need to know the device number for the program, which will vary by site.

To get the device number for Mediaindex, use the following command:

▶ Type: **list c mediaindex**.

   The system will display information similar to this example, in which case the device number is 122, located in the NOTIFY column:

```
DEV         DEVICE_TYPE    COMPUTER  NOTIFY  OPTIONS  DEVNAME

S122        mediaindex        B        122
```

To activate Mediaindex:

▶ Use the restart console command in the following format:
   restart <d*evice number*>

   For example: select the appropriate server at the console, and type: **restart 122**

📄 *(Optional) Use the stop command to stop the program, such as* **stop 122**.

8 Setting Up Interplay | MAM Sync Service

The following topics provide basic information to help you get started using Avid Interplay | MAM Sync Service Administrator:

- About Interplay | MAM Sync Service Administrator
- Opening Sync Service Administrator
- Understanding the Sync Service Administrator Layout
- Quitting Sync Service Administrator

For more information on using the Interplay MAM Sync Service Administrator, see the *Interplay | MAM Sync Service Administrator User's Guide*.

About Interplay | MAM Sync Service Administrator

This quick start guide is designed to familiarize you with the most important functions of Avid Interplay Media Asset Manager Sync Service Administrator, an Interplay MAM tool that provides features to monitor metadata synchronization between Interplay MAM Data Manager and Media | Index.

Basic Features

Starting with MAM v5.0, Interplay MAM systems can be accessed through MediaCentral | UX. MediaCentral UX provides the capability to search in the databases of different connected systems by using the "Central Index". This requires data synchronization between MediaCentral and Interplay MAM.

On MediaCentral's side, Media Index allows users of MediaCentral UX to search using the Central Index, which provides both data storage and a query engine. The index receives its data from the original data sources — generally, the databases of multiple asset management systems like Interplay MAM — and then pushes the data to the service that performs the index. This allows the index to continuously sync to the database. Indexed search in MediaCentral UX queries a central index synced with multiple databases and finds anything stored in the index.

On Interplay MAM's side, the synchronization of data between the MAM system database and Central Index is done by the SyncCentralIndex Service. It gets data from the MAM Data Manager's DM_SYNC table and pushes it to Central Index. The SyncCentralIndex Service detects data changes for media and EDL objects, reads information from the Data Manager database (for metadata, all attributes; for strata, changed segments), and sends update and delete requests to Central Index. Before SyncCentralIndex Service can synchronize the data, it passes all Interplay MAM attribute names and types to Central Index and converts MAM attributes to Central Index custom attributes. This is called the initialization of the data schema.

Note the following limitations:

- All Attributes marked as "searchable" in the MAM data model are always sent to Central Index during synchronization; attributes marked as "not searchable" are not sent to Central Index and are not available for searching or in search-result lists.

- Multi-value (MV) attributes and multi-value compound (MVC) attributes cannot be used for searching in Central Index and cannot be displayed in search-result lists.

- For timecoded metadata (MAM stratum segments) only searching of text fields is currently supported; non-text metadata fields in segments are not sent to Central Index and are not available for searching or in search-result lists.

- When using legal list attributes in Interplay MAM CBA rules, MAM allows building rules using ">=", "<" for legal list values, since these values are integer values in MAM. On the Central Index side, legal list values are treated as "ID-string," and only comparisons like "=" or "!=" are supported.

The SyncCentralIndex Service has a GUI called "Sync Service Administrator" that can be used by administrators to monitor the metadata synchronization between Interplay MAM and Media Index. The following illustration shows the services that are involved in data synchronization for Central Index on Interplay MAM's and MediaCentral's side.
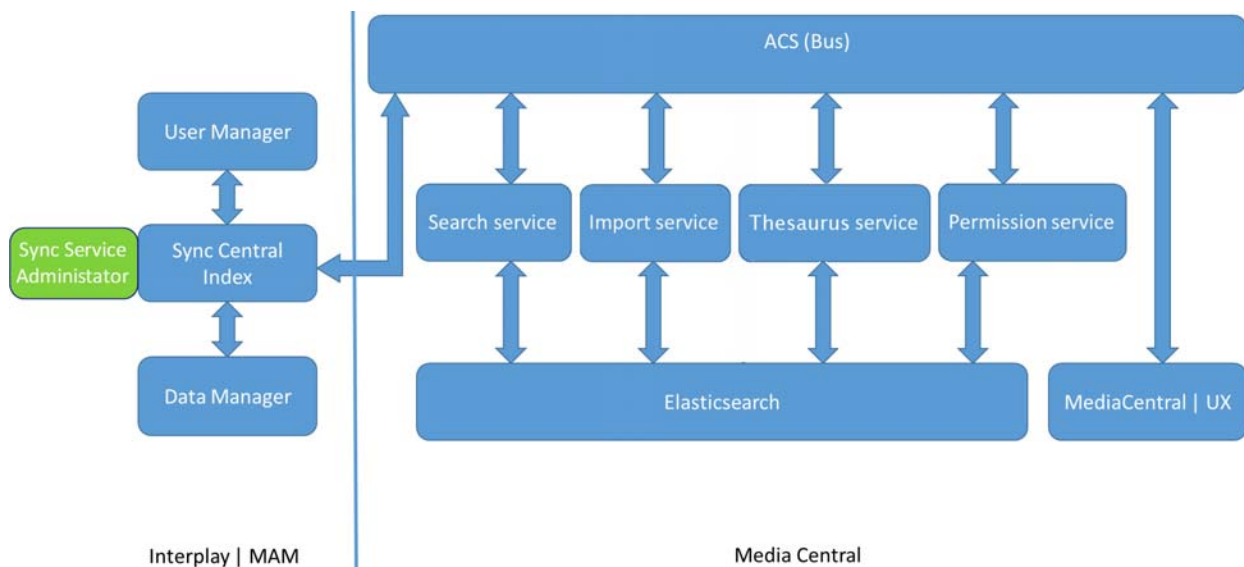
Prerequisites for Using Sync Service Administrator

To use Sync Service Administrator you must be a member of an administrator group that has the following group rights:

- Administration/Platform_Administration

As an administrator you can review and assign these group rights in the User Manager of the Interplay MAM Administrator. For additional information, see the *"Interplay | MAM User Manager User's Guide"* that is shipped with the software.

Sync Service Administrator Configuration Options

Some Sync Service Administrator features are subject to configuration. The default MAM installation comes with a full set of default configuration settings for Sync Service Administrator. As an administrator you can review and adapt these configuration settings in the following MAM System Administrator profiles:

Profile	Section	Affected Sync Service Administrator Views
SyncCentralIndex	Settings	Monitor, Traffic, Errors

For additional information, see the *"Interplay | MAM Configuration Settings Reference"* that is shipped with the software.

Basic Checks

To use Central Index for Interplay MAM in MediaCentral UX, check and adjust the following settings. These settings are configured during the installation of the MAM.Core.MediaCentralConnection package and should already display correct values.

To check the configuration:

1. Sign in to Interplay MAM Control Center as an administrator.

2. Start Interplay MAM Administrator from within MAM Control Center.

3. Open the System Administrator from within Interplay MAM Administrator.

4. In the Configuration tab, select Global > Bus and ensure that the value refers to the correct host where the Bus is running.

If you change the setting, you need to restart the pool MAM_2000_Backend2 in MAM Control Center's Service Controller.

5. Select DataManagerWS > CentralIndex and ensure that the value of the key MarkForSynchronization is set to "true".

If you change the setting, you need to refresh the configuration of the DataMangerWS or restart the pool MAM_1000_Backend1 in MAM Control Center's Service Controller.

6. Select SyncCentralIndex > Settings and ensure that the value of the following keys are set to "true".

   - EnableDataPropagation

   - EnableCBAAssignmentsPropagationOnStart (if needed)

   - SendLegalListsOnStartup (if needed)

   - SendThesauriOnStartup (if needed)

If you change a setting, you need to refresh the configuration of the SyncCentralIndex or restart the pool MAM_2000_Backend2 in MAM Control Center's Service Controller.

7. Select SyncCentralIndex > Toggles and check if the value of the key EnableNewThesauriIds has the required value:

   - True: Enables the new ID format for thesauri and legal lists. Recommended for new installations of Interplay MAM v5.4/MediaCental v2.5 or higher, or updating an installation with a small amount of data. Setting the value to "true" requires rebuilding the index for all assets.

   - False: This is the default value; allows the use of MAM v5.4/MediaCentral v2.5 without the re-indexing requirement.

If you change the setting, you need to refresh the configuration of the SyncCentralIndex or restart the pool MAM_2000_Backend2 in MAM Control Center's Service Controller.

To check required services:

1. Sign in to Interplay MAM Control Center as an administrator.

2. Open the Service Controller from within MAM Control Center.

3. Ensure that the following pools and services are running:

   - Pool MAM_1000_Backend1

   - Pool MAM_2000_Backend2

   - Service DataManagerWS

   - Service SyncCentralIndex

Opening Sync Service Administrator

Sync Service Administrator is started from within Interplay MAM Control Center.

To start Interplay MAM Sync Service Administrator:

1. Open an Internet browser and enter the URL https://<*hostname*>:9911/ControlCenter/web/

   The Interplay MAM Control Center login dialog opens.

2. Type your user name and password, and click Login.

   Interplay MAM Control Center opens.

3. Click MAM Administrator on the Links pane.

   The Interplay MAM Administrator login dialog opens.

4. Type your user name and password, and click Login.

   Interplay MAM Administrator opens.

5. Click the Sync Service Administrator link on the Overview page of Interplay MAM Administrator.

   Sync Service Administrator opens in a new tab.

Understanding the Sync Service Administrator Layout

Once you have logged in, Sync Service Administrator opens and displays the Monitor view. Sync Service Administrator includes the following areas and views:

Area	Function
Title bar	The Show mapping info link provides information on attribute mapping between Interplay MAM and Central Index.
Monitor	Monitors recent synchronization activities, displays details for each activity, and provides controls to manually trigger synchronization on different levels. For additional information, see "The Monitor View" on page 64.
Traffic	Monitors recently synchronized objects and displays details for each object. For additional information, see "The Traffic View" on page 65
Statistics	Provides statistics on how many MAM objects of different classes were synchronized in Central Index. For additional information, see "The Statistics View" on page 66.
MAM/CI Compare	Provides statistics on the number of objects and EDLs per object class in MAM and Central Index, and shows the difference between these numbers. For additional information, see "The MAM / CI Compare View" on page 67.
Sync Queue	Provides statistics on the create, update, and delete entries in the DM_SYNC table. For additional information, see "The Sync Queue View" on page 67.
Errors	Displays information on unsynchronized entries and provides controls to reset or purge erroneous synchronization entries. For additional information, see "The Errors View" on page 68.
Send JSON	Allows the sending of JSON statements to and receives responses from Interplay Central services for debugging and troubleshooting purposes. For security reasons, the Send JSON tab is hidden by default. As an administrator you can show the tab by setting the value of the key EnableMessagingDiagnostics in the Settings section of the SyncCentralIndex configuration profile to "true."

The Monitor View

In most cases, your first experience with Sync Service Administrator will be in Monitor view, where the basic Sync Service Administrator features are close at hand. The Monitor view is arranged in six sections that provide information and controls to set up basic Sync Service Administrator functions. The following illustration and table describe the layout of the Monitor view.

	Item	Function
1	Refresh controls	Sync Service Administrator updates the display of recent synchronization activities every five seconds if the Auto refresh check box is checked. If you uncheck the Auto refresh check box you can update the display manually by clicking the Refresh button.
2	Recent events list	Displays a configurable number of synchronization activities (default is 3,000). For each entry the following information is shown: Time stamp (data and time of synchronization), consecutive number, and message. The following colors indicate the synchronization status: • Red: indicates synchronization errors • Orange: indicates warnings • Black: indicates schema, CBA, legal list, and thesaurus synchronization activities; bold style indicates index rebuild activities

Item	Function
3 Details	Displays detailed information about the entry selected in the Recent events list.
4 Water mark indicator	Indicates the load of the synchronization queue as a traffic light. Sync Index Administrator checks the number of high-priority entries in the DM_Sync table and shows a traffic light depending on water mark settings. You can update the water mark indicator manually by clicking the Refresh button.
5 Synchronization controls	Provides controls to manually trigger synchronization on different levels:

- Synchronization of an individual object using its DMGUID
- Synchronization of all objects of a selected object class
- Synchronization of all objects
- Synchronization of all objects (with lower priority)
- Synchronization of CBA rules to user assignment
- Synchronization of CBA rules
- Synchronization of thesauri
- Synchronization of legal lists
- Rebuild of the entire index

All synchronization options — except for synchronizing an individual object by DMGUID — require that you type a reason in the Synchronization reason text box to activate the Synchronize button.

The Traffic View

The Traffic view displays information about recently synchronized objects; it provides information about recent asset create, asset update, and asset delete activities, and error, warning, and group messages.

	Item	Function
1	Refresh controls	Sync Service Administrator updates the display of recently synchronized objects every five seconds if the Auto refresh check box is checked. If you uncheck the Auto refresh check box you can update the display manually by clicking the Refresh button.
2	Synchronized objects list	Displays a configurable number of recently synchronized objects (default is 3,000). For each object the following information is shown: Time stamp (data and time of synchronization), consecutive number, and a title consisting of an operation type (Created, Updated, or Deleted) and the DMGUID of the object.

The following colors indicate the synchronization status:

- Red: indicates synchronization errors

- Orange: indicates failed synchronizations

- Blue: indicates successfully created or updated objects

- Light blue: indicates deleted objects

- Gray: indicates bulk synchronization activities

	Item	Function
3	Details	Displays detailed information about the object selected in the Recently synchronized objects list.

The Statistics View

The Statistics view provides statistics on how many MAM objects of different classes were synchronized in Central Index. The information is displayed in an overview table, as shown in the following illustration. For each object class the number of successful and failed synchronizations is shown for the current hour, current day, and current week. Additionally, the total number of synchronization activities is given.

The Statistics view is for information purposes only, and does not provide controls to apply actions on the statistics.

Object class	Current hour		Current day		Current week		Total	
	Successful	Failed	Successful	Failed	Successful	Failed	Successful	Failed
TEST	0	0	57	0	57	0	57	0
VIDEO	69	0	138	0	138	0	138	0
IMAGE	0	0	6	0	6	0	6	0
AUDIO	0	0	12	0	12	0	12	0
DOCUMENT	0	0	9	0	9	0	9	0
TEST_CEN	0	0	8	0	8	0	8	0
EDL	0	0	18	0	18	0	18	0
TEST_ED1	0	0	8	0	8	0	8	0
Total for all object classes							256	0

The MAM / CI Compare View

The MAM / CI Compare view provides statistics on the number of objects in MAM and Central Index per object class, and differences between them. Information is provided for object (asset) and EDL (sequence) classes.

The information is displayed in an overview table, as shown in the following illustration. It shows the number of objects for each object and EDL class in MAM and Central Index. Additionally, the number of differing entries is given.

The MAM / CI Compare view is for information purposes only, and does not refresh automatically. It provides a control to update the statistics.

Object class	MAM Count	CI count	Difference
AUDIO	15	12	3
EDL	67	52	15
IMAGE	17	6	11
TRAILER	1	1	0
VIDEO	44	37	7
Total for all object classes	144	108	36

The Sync Queue View

The Sync Queue view provides statistics on the number of entries in the synchronization queue. The information is displayed in an overview table, as shown in the following illustration. For each object class, it shows the number and priority of synchronization requests, and the create, update, and delete entries in the DM_SYNC table. Additionally, the total number of entries in the synchronization table is given.

The Sync Queue view is for information purposes only, and does not refresh automatically. It provides a control to update the statistics.

Object class	Operation			Total per object class
	Create	Update	Delete	
EDL	19	0	0	19
TEST_ED1	4	0	0	4
AUDIO	1014	0	0	1014
DOCUMENT	9	0	0	9
IMAGE	6	0	0	6
TEST	58	0	0	58
TEST_CEN	4	0	0	4
VIDEO	71	0	0	71
Total for all object classes				1185

The Errors View

The Errors view provides statistics on synchronization errors, and controls to reset or purge erroneous synchronization entries in Central Index.The following illustration and table describe the layout of the Errors view.

	Function
1	Refreshes the entries in the Object class table.
2	Displays the number of synchronization errors for each affected object class. Provides check boxes to select individual or all object classes to trigger a retry or purge operation.
3	Provides controls to handle synchronization errors: • Clear error flag (retry): The ERROR_COUNT column value of the selected object class(es) will be reset to 0 and the synchronization can be triggered anew. • Remove Synchronization entries with errors (purge): All rows with an ERROR_COUNT value greater or equal than the configured error count threshold (profile SyncCentralIndex > section Settings > key EntryErrorLimit) will be deleted from the DM_SYNC table. These entries will not be synchronized again.
4	Triggers the selected synchronization operation (retry or purge) for the selected object class(es).

Quitting Sync Service Administrator

When you are finished using Sync Service Administrator quit the application.

To quit the application:

▶ Click the X in the Sync Service Administrator tab.

Sync Service Administrator quits, but both Interplay MAM Control Center and Interplay MAM Administrator remain open. You can resume working in both applications or quit.

A **Upgrading Media | Index**

The procedures for upgrading your MediaCentral system configured with Media Index depend on which version of MediaCentral Platform Services (MCS) you are upgrading from. For example, if you upgrade from v2.0 to v2.3, you must reset the Media Index indexes before the upgrade begins. If you upgrade from v2.2 to v2.3, you only need to migrate the Media Index schema data.

Upgrading Media Index from v2.2 or Higher

See the *Avid MediaCentral Platform Services Upgrade Guide* for complete upgrade details.

Upgrading Media Index from v2.0.x or v2.1.x

See the *Avid MediaCentral Platform Services Upgrade Guide* for upgrade details. At the appropriate point in the upgrade process, the Upgrade Guide will direct you to this section of the Media Index Guide.

Media Index v2.0 installations only support indexing Interplay Production databases. Media Index v2.1 introduced support for Avid iNEWS systems and Media Index v2.2 introduced support for MAM. As this chapter relates to v2.0 and v2.1 Media Index installations, it does not apply to configurations that include Interplay MAM.

If you are upgrading a multi-zone configuration, upgrade slave zones first. If you cannot bring all zones offline at once, upgrade the slave zone(s) before the master zone. An upgraded slave zone no longer has access to the centralized user management service (UMS), which is owned by the master zone. Users in the upgraded slave zone can continue to log in using the local, read-only copy of the UMS database. However, new users cannot be added through the slave zone.

This process consists of three steps:

- Resetting Media Index
- Editing the Elasticsearch-tribe Configuration File (if applicable)
- Validating the Media Index Changes

Resetting Media Index

Before upgrading your Media Index deployment, you must reset Media Index. In a multi-zone environment, you must reset the Media Index in each zone. During the Media Index reset procedure, all indexed data is deleted. If you are indexing an Interplay Production database, you must use the Production Engine Bus Connector (PEBCo) service to resync and restore the Media Index database after the upgrade is complete. For information on using PEBCo, see "Installing the Production Engine Bus Connector" on page 33.

To reset Media Index to prepare for an upgrade:

1. If you are indexing an Interplay Production database, you must first determine the name of the PEBCo instance running on the MediaCentral server. Enter the following command on either a single MediaCentral server or any node in a cluster configuration:

   `/opt/avid/bin/pam-agent-ctrl list`

   The server returns the name of your PEBCo instance.

2. If you are indexing an Interplay Production database, stop the PEBCo service by typing the following command:

   ▶ For single-server configurations:

      `/opt/avid/bin/pam-agent-ctrl stop <instance>`

   ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

      `crm resource stop AvidPamAgent-<instance>`

3. From the MediaCentral server, or any node in a cluster configuration, enter the following command to list the systems (Interplay Production or Avid iNEWS) indexed by Elasticsearch:

   `/opt/avid/bin/avid-acs-elastic-index-setup -m info`

   The terminal window displays the system information for your Media Index configuration.

   The following is a sample output:
   ```
   Systems found: 1
   ************************
   System id : BC2F27F5-DBC5-4611-869A-F7A43690EE57
   System type : interplay
   System display name : WAVD-IE01
   Number of assets : 18880
   ************************
   ```

   Take note of the "System id" for all listed systems as this information is required in the next step.

4. Type the following command for each system indexed by Elasticsearch:

   `/opt/avid/bin/avid-acs-elastic-index-setup -m deleteSystem --systemId <id_of_the_system>`

   If you have multiple indexed systems repeat this command for each system.

5. Type the following command to remove the System Data index:

   `/opt/avid/bin/avid-acs-elastic-index-setup -m removeSystemData`

6. Use the `avid-acs-media-index-restart.sh` script to stop all Media Index services:

   ▶ For single-server installations:

      `avid-acs-media-index-restart.sh stop`

   ▶ For cluster installations:

      `avid-acs-media-index-restart.sh crm stop`

📖 *For more information on the* `avid-acs-media-index-restart.sh` *script, see *

Editing the Elasticsearch-tribe Configuration File

In multi-zone configurations, you also must edit the Elasticsearch-tribe configuration file in all zones to remove information about the zone or zones that you are upgrading. For example, in a three-zone configuration where you start by upgrading Zone_1, you must remove Zone_1's information from the elasticsearch-tribe configuration file on Zone_2 and Zone_3. You must also remove the information regarding Zone_2 and Zone_3 from Zone_1's configuration file.

If there are other zones that you are not upgrading, you must still edit the configuration file. However, any zones remaining at the older version can continue to have configuration information that exposes the zone to other old zones.

To remove the zone information from the configuration file:

1. Starting with the zone you wish to upgrade, use the Linux text editor, vi, to edit the `elasticsearch.yml` configuration file:

   vi /etc/elasticsearch-tribe/elasticsearch.yml

2. Delete the zone binding information of the non-local zones (as highlighted in the example below):

```
[root@wavd-doc01 ~]# cat /etc/elasticsearch-tribe/elasticsearch.yml

discovery.zen.ping.multicast.enabled: false
tribe.local.cluster.name: Zone_1
tribe.local.transport.tcp.port: 9312
tribe.local.discovery.zen.ping.unicast.hosts:
  - "zone1-mcs:9300"
http.port: 9201
transport.tcp.port: 9305
node.name: zone1-tribe

tribe.Zone_2.cluster.name: Zone_2
tribe.Zone_2.transport.tcp.port: 9313
tribe.Zone_2.discovery.zen.ping.unicast.hosts:
  - "zone2-mcs01:9300"
  - "zone2-mcs02:9300"
  - "zone2-mcs03:9300"

tribe.Zone_3.cluster.name: Zone_3
tribe.Zone_3.transport.tcp.port: 9314
tribe.Zone_3.discovery.zen.ping.unicast.hosts:
  - "zone3-mcs:9300"
```

3. Press the Escape key to return to command mode, then type **:wq** and press Return to write and quit the text editor.

4. If you have a cluster configuration, repeat these steps for each node within the cluster.

5. Restart the Elasticsearch tribe service on this zone by typing one of the following commands:

   ▶ For single-server configurations:

   service elasticsearch-tribe restart

   ▶ For cluster configurations:

   crm resource restart elasticsearchTribeEverywhere

6. In one of the remaining zones, use the Linux text editor, vi, to edit the `elasticsearch.yml` configuration file:

   `vi /etc/elasticsearch-tribe/elasticsearch.yml`

7. Delete the binding information for the zone that is being upgraded.

   The example below highlights the information for Zone_1 which is the zone being upgraded:

   ```
   [root@wavd-news ~]# vi /etc/elasticsearch-tribe/elasticsearch.yml

   discovery.zen.ping.multicast.enabled: false
   tribe.local.cluster.name: Zone_3
   tribe.local.transport.tcp.port: 9312
   tribe.local.discovery.zen.ping.unicast.hosts:
     - "zone3-mcs:9300"
   http.port: 9201
   transport.tcp.port: 9305
   node.name: zone3-tribe

   tribe.Zone_1.cluster.name: Zone_1
   tribe.Zone_1.transport.tcp.port: 9313
   tribe.Zone_1.discovery.zen.ping.unicast.hosts:
     - "zone1-mcs:9300"

   tribe.Zone_2.cluster.name: Zone_2
   tribe.Zone_2.transport.tcp.port: 9314
   tribe.Zone_2.discovery.zen.ping.unicast.hosts:
     - "zone2-mcs01:9300"
     - "zone2-mcs02:9300"
     - "zone2-mcs03:9300"
   ```

8. Press the Escape key to return to command mode, then type **`:wq`** and press Return to write changes and quit the text editor.

9. If you have a cluster configuration, repeat these steps for each node within the cluster.

10. Restart the Elasticsearch tribe service on this zone by typing one of the following commands:

    ▶ For single-server configurations:

       `service elasticsearch-tribe restart`

    ▶ For cluster configurations:

       `crm resource restart elasticsearchTribeEverywhere`

11. Repeat steps 6 - 10 on all remaining zones.

Validating the Media Index Changes

Once the steps in the previous two processes have been completed, use the Elasticsearch plugin page to verify that the indexes are not listed.

To validate your changes:

1. Open a browser and navigate to http://<*server_name*>:9200/_plugin/head, where *sever_name* is the host name or virtual cluster name of the system that is being upgraded.

   The Elasticsearch head plug-in Web page opens in your browser.
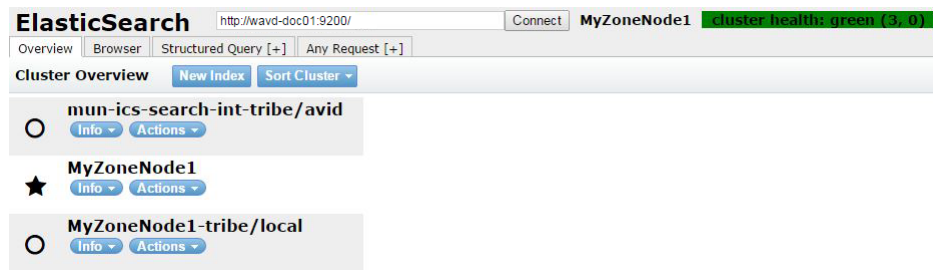
2. Check to make sure the system_data index and index information are not listed.

   When Media Index is reset for your upgrade process, you can install the latest version of Media Central Platform Services. For more information, see one of the following topics:

   Once the MediaCentral upgrade is complete, you must reconfigure your Media Index configuration. See the following sections for more information:

   - For installing Media Index in a single-server configuration, see "Configuring Media | Index for a Single Server" on page 11.

   - For installing Media Index in a cluster configuration, see "Configuring Media | Index for a Cluster" on page 16.

   - For configuring Media Index in a multi-zone environment, see "Setting Up Multi-Zone Search" on page 26.

Media | Index Backup Locations

When upgrading the MediaCentral Platform Services software, users are required to create a backup of all data using the USB key. However, indexes are not backed up by the system backup script. This typically does not cause an issue since users generally do not need to back up their indexes. For cluster configurations, the following notes also apply:

- In a cluster, the index should be configured with data redundancy to avoid data loss.

- In the case of a complete rebuild or of complete data loss within a cluster, Avid recommends that you do not use a backup but re-index the data instead. This is because the backup will no longer be in sync with what the index agents are expecting.

With the release of Elasticsearch v1.6, backups are only allowed on shared file locations, the backup in Media Index v2.4 and higher requires manual configuration of the backup locations. Setting this up correctly requires a file server to be mounted on all machines. Customers should contact Avid Customer Care or Avid Professional Services for instructions and support from Avid Engineering in setting this up.

B Working with Media Index

This section covers the additional commands and utilities that are used with systems configured for Media Index. The following topics are covered:

Additionally, a process for removing a system configured with Media Index in a MediaCentral multi-zone environment is included in this chapter:

Media Index Log Locations

Logs for Elasticsearch are located in the Linux `/var/log/` directory. Additional logs related to Avid Media Index services are located at `/var/log/avid/`.

For detailed information on each of the logs, see the "Best Practices, Troubleshooting and System Logs" chapter of the *MediaCentral Platform Services Concepts and Clustering Guide*.

Updating Localization Strings

Localizations for base, common, and time-based fields are stored in a directory located at `/opt/avid/etc/avid-search-import/resources/` and are automatically loaded into the system data index when it is created.

When would I want to use this procedure? Examples?

To reload localization strings, do the following:

▶ Type the following command:

/opt/avid/bin/avid-acs-elastic-index-setup -m updateLocalization

Displaying Information about Initialized Systems

You can use the `info` command to display the following data for all available systems within all configured Elasticsearch clusters:

- System id

- System type

- System display name

- Number of assets

To display system information, do the following:

▸ Type the following command:

/opt/avid/bin/avid-acs-elastic-index-setup -m info

The following is an example of Media Index configured for an Avid Interplay Production system on a single zone:

```
[root@wavd-doc01 ~]# avid-acs-elastic-index-setup -m info

Fetching system information...

Systems found: 1

*****************************************
System id : 8C213466-7A16-424F-9B72-EB239C04A3BB
System type : interplay
System display name : WAVD-IE01
Number of assets : 138
*****************************************
```

Checking the Current Index Status

The Indexes pane in the MediaCentral UX System Settings can be used to view the status of the indexes in your configuration.

The Indexes pane displays only the indexes in the local zone. To display indexes of other zones, connect directly to the MediaCentral UX of that zone.

To check the current status of your indexes:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Indexes.

   The Media Index Indexes status displays in the Details pane.

All indexes for your configuration display with their status and details listed in the pane.

76

Using the Media | Index Restart Command Line Tool

When Media Index is configured on a MediaCentral system, multiple new services are enabled. In the case of a cluster, many of these resources are added as resources to the Corosync / Pacemaker cluster. In an effort to simplify working with Media Index, Avid includes a utility that allows you to start, stop, or restart many of the services (or resources) in a single command. Additionally, a "status" option is available to provide the user with information on the affected services.

There is however one Media Index service that is not included in the utility. The `avid-acs-media-index-configuration` service is excluded and must be stopped, started, or restarted manually.

Type one of the following commands to interact with the Media Index services:

For single-server configurations:

/opt/avid/bin/avid-acs-media-index-restart.sh *<service>* *<command>*

For cluster configurations:

/opt/avid/bin/avid-acs-media-index-restart.sh *<service>* crm *<command>*

Although the word "restart" is included in the utility name, functionality is not limited to only restarting the services.

The command line tool takes the following parameters:

- ***[service name]***— (optional) If desired, you can interact with one or more services by specifying the service name. Each service must be separated by a space. If you do not specify a service, all Media Index services except for avid-acs-media-index-configuration are affected. The following services or resources can be specified:

Single Server Service Name	Cluster Resource Name
avid-acs-autocomplete	AvidSearchAutocomplete (AvidSearchAutocompleteEverywhere)
avid-acs-media-index-status-provider	AvidSearchIndexStatus (AvidSearchIndexStatusEverywhere)
avid-acs-media-index-feed	AvidSearchIndexFeed (AvidSearchIndexFeedEverywhere)
avid-acs-media-index-permission	AvidSearchPermission (AvidSearchPermissionEverywhere)
avid-acs-media-index-thesaurus	AvidSearchThesaurus (AvidSearchThesaurusEverywhere)
avid-acs-search	AvidSearch (AvidSearchEverywhere)
avid-acs-search-import	AvidSearchImport
elasticsearch	elasticsearch (elasticsearchEverywhere)
elasticsearch-tribe	elasticsearchTribe (elasticsearchTribeEverywhere)

- **crm** — If working in a cluster configuration, the crm flag is required. For single-server configurations, this option is not required.

- **[command]** — Specify one of the following commands to interact with the services:

  - **start** — start services

  - **stop** — stop services

  - **restart** — restart services

  - **status** — get the status

To get command line tool help information, run the command tool without parameters.

The following table includes examples of using the command line tool:

Description	Command
Restart one or more specific services in a single-server configuration	The following example restarts only the elasticsearch and elasticsearch-tribe services:
Alternatively, you can interact with a single service with the standard Linux commands such as: service elasticsearch-tribe restart	**/opt/avid/bin/avid-acs-media-index-restart.sh elasticsearch elasticsearch-tribe restart**
Stop all services in a single-server configuration	**/opt/avid/bin/avid-acs-media-index-restart.sh stop**
Start specific services in a cluster configuration	**/opt/avid/bin/avid-acs-media-index-restart.sh start**
Restart all services in a cluster configuration	**/opt/avid/bin/avid-acs-media-index-restart.sh crm restart**

Since the avid-acs-media-index-configuration service is not included in the restart script, the service must be started, stopped or restarted separately from the other services.

Type one of the following commands to interact with the avid-acs-media-index-configuration service:

- To interact with the service in a single-server configuration:

  service avid-acs-media-index-configuration <command>

- To interact with the service on a single node in a cluster configuration:

  crm resource <command> AvidSearchConfig

- To interact with the service on all nodes in a cluster configuration simultaneously:

  crm resource <command> AvidSearchConfigEverywhere

Command options are **start**, **stop**, and **restart.**

Restarting the Search and PEBCo Services

If you find that the indexed search has stopped working in your MediaCentral environment, you can restart the search service on any node that has a problem and then resync your index.

To check the status of the search service:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Cluster State.

   The Media Index Cluster State displays in the Details pane. You can identify the node displaying problems by checking the status of the search services.

To restart the search service and to resync the PEBCo service:

1. Log in to the MCS node as the Linux root user.

2. Restart the indexed search service by typing one of the following commands:

   ▶ For single-server systems: `service elasticsearch restart`

   ▶ For clustered configurations: `crm resource restart elasticsearchEverywhere`

3. Identify the PEBCo instance name by typing the following command:

   `pam-agent-ctrl list`

4. Restart the PAM agent resource by typing one of the following commands:

   ▶ For single-server systems: `service <PEBCo_instance_name> restart`

   ▶ For clustered configurations: `crm resource restart <PEBCo_instance_name>`

5. Start Interplay Administrator and log on to the Interplay Production database that you want to have indexed.

6. In the Site Settings section of the Interplay Administrator, click the Production Engine Bus Connector icon.

The Production Engine Bus Connector view is not available on Macintosh clients.

7. Click the Manage Status tab.

8. In the Manage section, click Resync to update the whole database content. Resyncing the database might take a long time to complete and invalidates the existing content of your index. During this process, the search only returns assets that are already processed within this run.

Taking Nodes Offline in a MediaCentral Cluster

Depending on the number of primary shards and replicas configured on your system, taking a node offline without following the recommended procedures might temporarily disable Indexed search functionality. If you need to take a node offline, review the following and complete the process that is appropriate for your situation:

• Manually stop Media Index services before taking the node offline using the `avid-acs-media-index-restart.sh` script. This script is cluster-aware and can stop services without introducing fail-counts to the Corosync cluster. This is the best solution when you need to temporarily remove a node from the Media Index cluster.

If you need to take the node completely offline, use the `crm node standby <node>` command to stop the remaining cluster resources.

If you only need to take the Media Index cluster offline, use the `avid-acs-media-index-restart.sh` script to restart the index services when ready.

For more information on the `avid-acs-media-index-restart.sh` script, see "Using the Media | Index Restart Command Line Tool" on page 77.

For more information on the `crm node standby` command, see the *MediaCentral Platform Services Concepts and Clustering Guide*.

• Increase the replica count before removing a node to allow for more nodes to be taken offline. This causes Elasticsearch to copy data on startup, which can consume processor resources. This can be a good solution if you remove multiple nodes. However, increasing the number of replicas puts more load on the existing nodes, both in terms of memory needed for the index and for storage space.

- Decrease number of shards. This option is only appropriate if a node will be taken offline permanently. You can change the replica count at any time, but changing the primary shard count requires you to reset the index, reconfigure it with the new shards count, and then re-index the data. This can be a time consuming process, especially for configurations with large indexes.

📄 *As a reminder, the Media Index clustering functionality and the Corosync Clustering functionality are independent from each other.*

Removing Zones from a Multi-Zone Configuration

If you want to remove a zone from your MediaCentral multi-zone configuration, you must remove the bindings for that zone in the Elasticsearch tribe configuration file on all other zones.

This process only removes the affected zone from the Media Index multi-zone configuration, it does not disable Media Index on any zone.

📄 *The example configuration used in this section consists of three zones: Master_Zone, Slave_Zone, and 3rd_Zone. This process will assume the removal of the zone named "3rd_Zone"*

To remove the zone information from the configuration file:

1. Starting with the zone you wish to remove, use the Linux text editor, vi, to edit the `elasticsearch.yml` configuration file:

   vi /etc/elasticsearch-tribe/elasticsearch.yml

2. Delete the zone binding information of the non-local zones (as highlighted in the example below):

```
[root@wavd-news ~]# vi /etc/elasticsearch-tribe/elasticsearch.yml

discovery.zen.ping.multicast.enabled: false
tribe.local.cluster.name: 3rd_Zone
tribe.local.transport.tcp.port: 9312
tribe.local.discovery.zen.ping.unicast.hosts:
  - "wavd-news:9300"
http.port: 9201
transport.tcp.port: 9305
node.name: wavd-news-tribe

tribe.Master_Zone.cluster.name: Master_Zone
tribe.Master_Zone.transport.tcp.port: 9313
tribe.Master_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-doc01:9300"

tribe.Slave_Zone.cluster.name: Slave_Zone
tribe.Slave_Zone.transport.tcp.port: 9314
tribe.Slave_Zone.discovery.zen.ping.unicast.hosts:
  - "wavd-mcs01:9300"
  - "wavd-mcs02:9300"
  - "wavd-mcs03:9300"
```

3. Press the Escape key to return to command mode, then type **:wq** and press Return to write and quit the text editor.

4. If you have a cluster configuration, repeat these steps for each node within the cluster.

5. Restart the Elasticsearch tribe service on this zone by typing one of the following commands:

   ‣ For single-server configurations:

      `service elasticsearch-tribe restart`

   ‣ For cluster configurations:

      `crm resource restart elasticsearchTribeEverywhere`

6. In one of the remaining zones, use the Linux text editor, vi, to edit the `elasticsearch.yml` configuration file:

   `vi /etc/elasticsearch-tribe/elasticsearch.yml`

7. Delete the binding information for the zone that is being removed from the configuration.

   The example below highlights the information to be removed from zone name "Master_Zone":

   ```
   [root@wavd-doc01 ~]# cat /etc/elasticsearch-tribe/elasticsearch.yml

   discovery.zen.ping.multicast.enabled: false
   tribe.local.cluster.name: Master_Zone
   tribe.local.transport.tcp.port: 9312
   tribe.local.discovery.zen.ping.unicast.hosts:
     - "wavd-doc01:9300"
   http.port: 9201
   transport.tcp.port: 9305
   node.name: wavd-doc01-tribe

   tribe.Slave_Zone.cluster.name: Slave_Zone
   tribe.Slave_Zone.transport.tcp.port: 9313
   tribe.Slave_Zone.discovery.zen.ping.unicast.hosts:
     - "wavd-mcs01:9300"
     - "wavd-mcs02:9300"
     - "wavd-mcs03:9300"

   tribe.3rd_Zone.cluster.name: 3rd_Zone
   tribe.3rd_Zone.transport.tcp.port: 9314
   tribe.3rd_Zone.discovery.zen.ping.unicast.hosts:
     - "wavd-news:9300"
   ```

8. Press the Escape key to return to command mode, then type **`:wq`** and press Return to write changes and quit the text editor.

9. If you have a cluster configuration, repeat these steps for each node within the cluster.

10. Restart the Elasticsearch tribe service on this zone by typing one of the following commands:

    ‣ For single-server configurations:

       `service elasticsearch-tribe restart`

    ‣ For cluster configurations:

       `crm resource restart elasticsearchTribeEverywhere`

11. Repeat steps 6 - 10 on all remaining zones.

12. If you intend to remove this zone from the MediaCentral multi-zone configuration, review the steps for "Dismantling a Multi-Zone Environment" in the *MediaCentral Platform Services Installation and Configuration Guide*.

13. If you are running MediaCentral Platform Services v2.6 or higher, you must also remove the zone from the sharded Mongo configuration.

Disabling Media Index

Media Index components are automatically installed with the MediaCentral Platform Services software installation and cannot be removed. However, if you have configured Media Index on a MediaCentral single server or cluster and need to remove Media Index from the environment, the following processes details the required steps to delete the indexes and remove the Media Index configuration.

This process is destructive as it deletes the indexes from the server! Indexes can only be recovered through a reindex which can take a long time for large databases.

To completely disable Media Index:

1. Delete all Indexes from MCUX System Settings > Media Index > Indexes.

2. Click Reset in MCUX System Settings > Media Index > Settings.

3. If the Media Index is configured with Interplay Production, complete the following in the Interplay Administrator:

   a. Under Server > Server Settings, disable "Update Tracking"

   b. Under Site Settings > Production Engine Bus Connector, Stop the PEBco service and Unassign the PEBCo service

4. If the Media Index is configured with Interplay Production, use Interplay Access to remove the Bus AMQP URL The Interplay Administrator tool allows you to enter a new URL or alter the existing URL, but you cannot delete the URL. The deletion process must be completed using Interplay Access.

   a. If you are logged into the Interplay Administrator, log out at this time.

   b. Launch Interplay Access and log into the application as a user with administrator-level access.

   c. Right-click on the Interplay Production database (AvidWG) at the top of the tree and select Advanced > Get / Set Property...

   d. Type ACS_BUS_URL in the Name field and click the Get button.

   e. Delete the URL in the Value text box and click the Apply button.

   f. Click OK to close the Get / Set Property window and Close Interplay Access.

   g. Launch the Interplay Administrator tool and verify that the Bus AMQP URL has been removed from the MediaCentral and Platform Services Settings.

5. If changes were made to Interplay Production Services, disable the Configuring Automation and Consolidate Services.

6. If Media Index has been integrated with Avid iNEWS, revert changes on the iNEWS servers.

7. If Media Index has been integrated with Interplay MAM, revert changes on the MAM servers.

C Media | Index Custom Configuration

The default setup described in this document covers most MCS configurations. In some instances, you might need to configure certain property files.

For more information on modifying the application properties file, see "Modifying application.properties" in the MediaCentral Platform Services Installation and Configuration Guide.

This section covers the following custom configuration procedures:

- Modifying the Media | Index System Settings
- Adjusting Settings Through the ACS Monitor
- Setting Elasticsearch Memory Usage
- Configuring the PEBCo for Custom Configurations
- Changing the Default Search Type
- Changing Replica Counts
- Changing the Index Refresh Interval

Modifying the Media | Index System Settings

The Media Index section of the MediaCentral UX System Settings allows you to monitor and modify some settings used by Media Index. You use some of these settings to configure Media Index and others to monitor the systems and the indexes used by Media Index.

Modifying the Indexes Pane

You can use the Indexes pane in the System Settings to add or change the system name that is used to filter some indexed searches, to delete a system from the central index, and to view details about the indexes configured for Media Index.

To use the Media Index settings in the MediaCentral UX Indexes pane:

1. Sign in to MediaCentral UX as an administrator.
2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Indexes.

   The Media Index Indexes display in the Details pane.
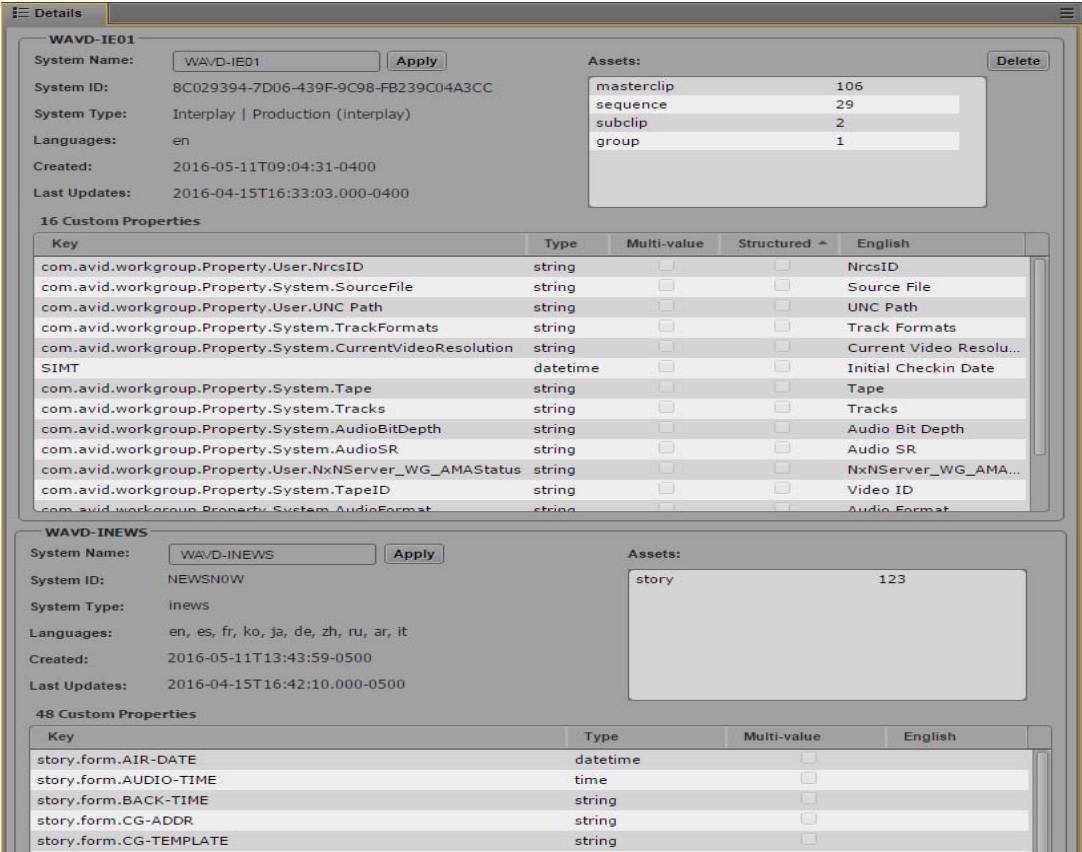
If you have indexed more than one system, such as an Interplay Production system and an iNEWS system, each system will be listed separately in the Details pane.

Area or Component	Description
System Name (optional)	The system name is set by the indexing agent. For Interplay Production, the system name is set by the Production Engine Bus Connector (PEBCo) and uses the Interplay Production Engine's host name or virtual cluster host name. Other indexing agents — for example, the iNEWS sync agent — might not assign a system name by default. You can add or modify the system name so that users can easily identify the index used during an indexed search. The System Name appears as the System filter in indexed searches.
Apply button	Click this button to apply changes made to the System Name field.
Delete button	You can use the Delete button to remove the system and all of its data from the central index. ⚠ **Removing a system removes all data from the index. Users cannot search the system using the indexed search functionality of Media Index.**
Index summary	This area lists some of the details of the index, including the system ID, system type (for example, Interplay Production or iNEWS), and when the index was created.

Area or Component	Description
System ID	The System ID is a unique identifier for an index that is generated by the indexing agent (PEBCO, MAM agent). It uniquely identifies the system in Media Index. Interplay Production systems are identified by a unique string of characters which cannot be altered. iNEWS systems use the system name as the System ID.
Assets summary	This area provides a brief summary of the assets indexed by Media Index.
Custom Properties list	This list includes all properties set for the system that can appear in an indexed search. Search properties are set by administrators for the appropriate system — for example, Interplay Production properties are specified in the Canonical Data tab in the Production Engine Bus Connector view in the Interplay Administrator tool.
	You cannot modify the properties in this pane. For more information, see "Configuring and Using the PEBco Service" on page 37.

4. If you want to modify the name that appears in the System search criterion in indexed searches, enter a new name in the System Name field, and then click Apply.

   For more information on using search criteria to filter searches, see "Advanced Search Filters" in the *Avid MediaCentral | UX User's Guide*.

5. If you want to remove the system from the central index, click Delete.

⚠ **Deleting a system removes the system and all data from the index. You cannot undo this action.**

   The Delete System dialog box opens.

```
┌─────────────────────────────────────────────────┐
│              Delete System                    [×]│
├─────────────────────────────────────────────────┤
│ Type 'DELETE' to delete system. There is no undo for this operation.│
│ ┌─────────────────────────────────────────────┐ │
│ │                                             │ │
│ └─────────────────────────────────────────────┘ │
│           ┌────────┐   ┌────────┐               │
│           │   OK   │   │ Cancel │               │
│           └────────┘   └────────┘               │
└─────────────────────────────────────────────────┘
```
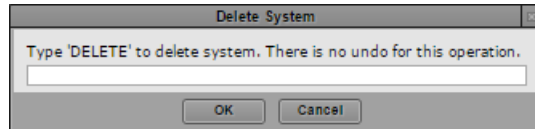
6. Type `DELETE` to delete the system, and then click OK. You can cancel the action by clicking Cancel.

   The system and its data are removed from the central index.

Modifying the Settings Pane

You can use the Settings pane in the System Settings to change the cluster name and to delete all systems and restore Media Index into the initial state.

To change the cluster name:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Settings.

   The Media Index Settings display in the Details pane.

4. If you want to modify the name of the cluster, type the name in the Cluster Name field, and then click Apply.

⚠ **If you change the cluster name, during a restart Elasticsearch creates a new data folder. If you had any indexes before you changed the cluster name, the indexes fail to appear in your configuration and the information in them is not accessible. While the index data is not lost, Avid recommends that you not change the cluster name after your initial configuration. When you work in a multi-zone set-up, you should consider which cluster names you plan to use in all zones and not change them after the indexing process has started.**

To delete all systems and restore Media Index into the initial state:

1. Sign in to MediaCentral UX as an administrator.

2. Select System Settings from the Layout selector.

   The System Settings layout opens.

3. In the Settings pane select Media Index > Settings.

   The Media Index Settings display in the Details pane.

4. In the Reset Media | Index section, click Reset.

⚠ **Deleting all systems and resetting Media Index removes the systems and all data from the index. A reset also removes the Elasticsearch tribe configuration. The old configuration is saved in a backup file next to the new one so that the tribe configuration can be copied and restored.**

⚠ **You cannot undo this action.**

The Reset Media Index dialog box opens.
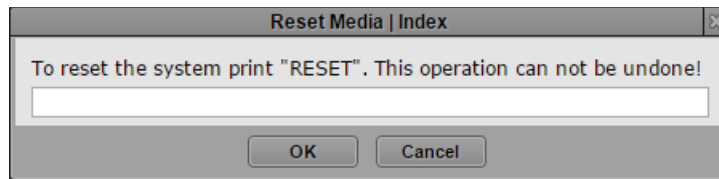
5. Type **RESET** to reset the system, and then click OK. You can cancel the action by clicking Cancel.

   The Media Index system is reset.

6. To restore the Elasticsearch tribe configuration, do the following:

   a. On each node in your configuration, use the Linux *vi* text editor to edit the backup copy of the Elasticsearch tribe configuration file:

      vi /etc/elasticsearch-tribe/elasticsearch.yml_backup<*n*>

   b. Copy the configuration information, and then press the Escape key to return to command mode.

   c. Open the elasticsearch.yml configuration file for editing:

      vi elasticsearch.yml

   d. Paste the information from the backup file into the configuration file.

   e. Press the Escape key to return to command mode, then type **:wq** and press Return to write and quit the text editor.

   f. Restart Elasticsearch tribe service:

      ▶ For single server configurations, type the following command:

      service elasticsearch-tribe restart

      ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

      crm resource restart elasticsearchTribeEverywhere

Adjusting Settings Through the ACS Monitor

Many of the services used to create Media Index are governed by configuration files that contain parameters which can be adjusted to enhance or limit the operation of the system. While these files can be edited manually through Linux, Avid provides a user interface called the Avid ACS Monitor which can be used to adjust many of the same settings more easily.
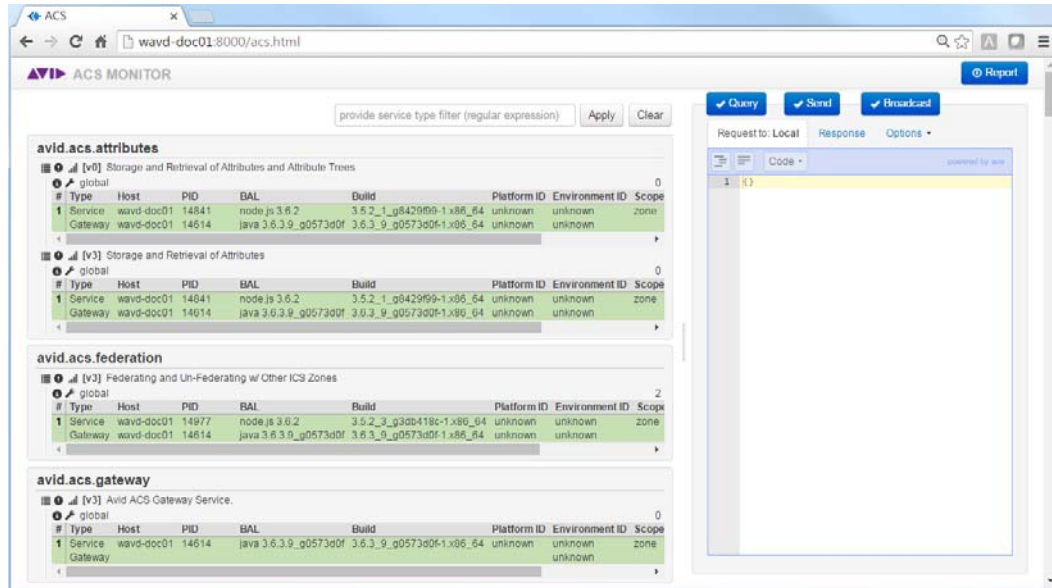
Adjusting settings through the Avid ACS Monitor provides the following advantages:

• Fewer manual adjustments to system files makes the process less prone to error.

• Changes made through direct manipulation of configuration files are often not saved during system upgrades. Changes made through the ACS Monitor are saved through the avid-acs-attributes service and are maintained through all system upgrades.

• In cluster configurations, changes made through the ACS Monitor are pushed to all nodes.

This section includes instructions to complete the following tasks:

• Accessing the Avid ACS Monitor Tool

• Changing the Maximum Number of Markers Returned in a Search

• Disabling Highlighted Markers

• Adjusting Media Index Scalability

• Adjusting Relevance for Search Results

• Disabling Search Suggestions

• Configuring Search Criteria Behavior

Accessing the Avid ACS Monitor Tool

The ACS Monitor is not enabled by default and must be manually started to interact with the tool and adjust system settings.

To adjust the Media Index configuration through the ACS Monitor:

1. Log in to the MCS server as the Linux 'root' user. If you have a clustered configuration, log into the master node.

2. Start the avid-acs-monitor service:

   `service avid-acs-monitor start`

3. Open a web browser and enter the following address to access the Avid ACS Monitor:

   `http://<hostname>:8000`

   Where *<hostname>* is the host name or IP address of your MediaCentral server or cluster master node.

   The Avid ACS Monitor page appears, allowing you to alter the system configuration.

If you do not see the list of attributes on the left side of the page, refresh the browser page.

4. Refer to the processes in this section to alter the system configuration.

5. Once you have completed all system configuration changes, stop the avid-acs-monitor service:

   `service avid-acs-monitor stop`

Changing the Maximum Number of Markers Returned in a Search

When initiating an indexed search to locate text contained in markers, Media Index returns a maximum of five markers in the search results by default. This number can be adjusted through the ACS Monitor.

To change the default number of markers:

1. Follow the instructions for "Accessing the Avid ACS Monitor Tool" on page 90.

2. Enter "avid.acs.media.index.configuration" (without the quotes) in the search field at the top of the page to find the configuration service.

3. Expand the configuration service options by clicking the Service Operations List icon in the upper-left corner of the configuration service field.

4. Locate the setConfigurationByField item in the list of service options. This section includes a number of examples for changes than can be made to the configuration. Click on the link for enableSearchQueryLogger.

   The following text appears in the right pane:

```
1  {
2      "serviceType": "avid.acs.media.index.configuration",
3      "serviceRealm": "global",
4      "serviceVersion": 0,
5      "op": "setConfigurationByField",
6      "paramSet": {
7          "acs_enable_query_logger": true
8      }
9  }
```

If you click on any of the examples in this section, you will notice that they all share the same basic structure. Line 7 is the only line that is different between them. Rather than entering all the text manually, you can alter the contents of this example to update the maximum number of markers.

5. In the "Request to: Local" pane on the right, replace the text on line 7 with the following:

```
"highlight_number_of_fragments": "#"
```

Where "#" represents the maximum number of markers to be returned in a search. The default number is 5.

When entering text, make sure to include the two leading spaces and the quotes as shown in the example above.

6. Click the Query button in the right pane to save the change.

If the query is accepted, a Response value (in milliseconds) will be returned. If you receive any errors, verify that the information you entered does not contain any incorrect characters.

7. Verify your change by selecting the getConfiguration item in the list of service options and click the Query button on the right to poll the service. The updated value appears under the "resultSet" list.

8. Restart the avid-acs-search service to activate the change:

   ▶ For a single server, type the following from the Linux command prompt:

   service avid-acs-search restart

   ▶ For a clustered system, type the following from the Linux command prompt:

   crm resource restart AvidSearchEverywhere

Disabling Highlighted Markers

Text that has been entered in markers through the Logging pane can be found through Indexed searches. When displayed, the text is highlighted in the search results, as indicated below:
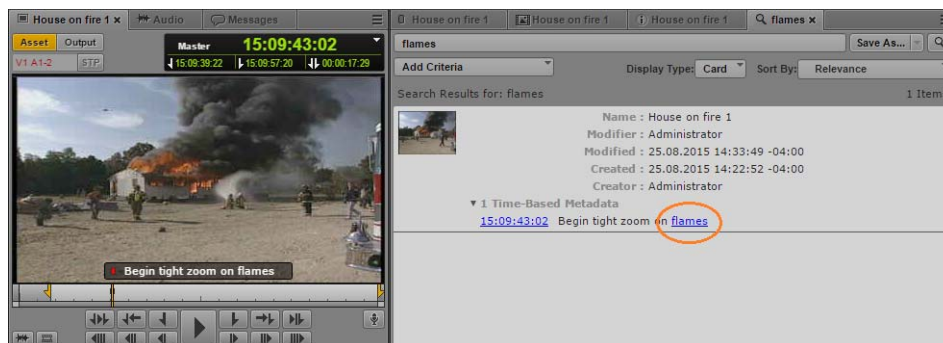
If desired, system administrators can disable this feature for under-performing servers. Disabling the feature reduces load on system resources such as RAM and CPU cycles.

To disable highlighted markers:

1. Enter "avid.acs.media.index.configuration" (without the quotes) in the search field at the top of the page to find the configuration service.

2. Expand the configuration service options by clicking the Service Operations List icon in the upper-left corner of the configuration service field.

3. Locate the setConfigurationByField section and click the link for <u>disableHighlighting</u>.

The following text appears in the right pane:

```
{
  "serviceType": "avid.acs.media.index.configuration",
  "serviceRealm": "global",
  "serviceVersion": 0,
  "op": "setConfigurationByField",
  "paramSet": {
    "acs_disable_highlight": true
  }
}
```

4. No adjustments are required, simply click the Query button in the right pane to save the change.

5. Restart the avid-acs-search service to activate the change:

▶ For a single server, type the following from the Linux command prompt:

service avid-acs-search restart

▶ For a clustered system, type the following from the Linux command prompt:

crm resource restart AvidSearchEverywhere

📄 *Once disabled, highlighted markers can be re-enabled by completing the above process, substituting "enableHighlighting" option in step 3.*

Adjusting Media Index Scalability

You can configure the scalability of the Media Index search service — for example, you can run more than one instance of the search service at the same time.

📄 *Altering the default values of the variables described in this section can negatively affect system performance. Only adjust these values if directed to do so by an Avid representative.*

To adjust Media Index scalability options:

1. Follow the instructions for "Accessing the Avid ACS Monitor Tool" on page 90.

2. Enter "avid.acs.media.index.configuration" (without the quotes) in the search field at the top of the page to find the configuration service.

☰ 3. Expand the configuration service options by clicking the Service Operations List icon in the upper-left corner of the configuration service field.

4. Locate the setConfigurationByField item in the list of service options. This section includes a number of examples for changes than can be made to the configuration. Click on the link for <u>enableSearchQueryLogger</u>.

The following text appears in the right pane:

```
1 ▾ {
2     "serviceType": "avid.acs.media.index.configuration",
3     "serviceRealm": "global",
4     "serviceVersion": 0,
5     "op": "setConfigurationByField",
6 ▾   "paramSet": {
7         "acs_enable_query_logger": true
8     }
9 }
```

If you click on any of the examples in this section, you will notice that they all share the same basic structure. Line 7 is the only line that is different between them. Rather than entering all the text manually, you can alter the contents of this example to alter the Media Index scalability options.

5. In the "Request to: Local" pane on the right, replace the text on line 7 with one of the following values:

Variable	Description
"search_cluster_threads_minimum": "*<value>*"	Sets the minimal number of search service instances. If not set, the variable defaults to the number of CPUs.
"search_cluster_threads_maximum": "*<value>*"	Sets the maximum number of instances. If not set, the variable defaults to the value of the SEARCH_CLUSTER_THREADS_MINIMUM variable multiplied by 2.

When entering text, make sure to include the two leading spaces and the quotes as shown in the example above.

6. Click the Query button in the right pane to save the changes.

If the query is accepted, a Response value (in milliseconds) will be returned. If you receive any errors, verify that the information you entered does not contain any incorrect characters.

7. If you need to adjust both variables, repeat the above steps the configure the second value.

8. Verify your changes by selecting the getConfiguration item in the list of service options and click the Query button on the right to poll the service. The updated values appear under the "resultSet" list.

9. Restart the avid-acs-search service to activate the changes:

▶ For a single server, type the following from the Linux command prompt:

```
service avid-acs-search restart
```

▶ For a clustered system, type the following from the Linux command prompt:

```
crm resource restart AvidSearchEverywhere
```

Adjusting Relevance for Search Results

You can use the indexed search to search by key terms. Adjusting the relevance of certain fields for searches allows you to boost the priority of those terms so that the appropriate assets have a higher relevance in the search results. To adjust the relevance of certain fields, you can modify the index boosting settings in the search attribute service by using the setBoostConfiguration command in the ACS Monitor tool.

To adjust the index boosting settings:

1. Follow the instructions for .

2. Enter "avid.acs.media.index.configuration" (without the quotes) in the search field at the top of the page to find the configuration service.

3. Expand the configuration service options by clicking the Service Operations List icon in the upper-left corner of the configuration service field.

4. Locate the setBoostConfiguration item in the list of service options and click the link for setBoostConfiguration.

The following text appears in the right pane:

```
1  {
2    "serviceType": "avid.acs.media.index.configuration",
3    "serviceRealm": "global",
4    "serviceVersion": 0,
5    "op": "setBoostConfiguration",
6    "paramSet": {
7      "asset": {
8        "systemID": 2
9      },
10     "common": {
11       "globalAssetID": 1,
12       "durationTC": 1,
13       "description": 3
14     },
15     "tbmd": {
16       "id": 3,
17       "category": 1,
18       "creator": 2,
19       "modifier": 4
20     }
21   }
22 }
```

The paramSet lists three different categories or "entities": asset (base), common, and tbmd (time based metadata). Each category is associated with various fields and boost values. For example, the "tbmd" category includes the following default values:

```
"tbmd": {
    "id": 3,
    "category": 1,
    "creator": 2,
    "modifier": 4
```

The default boost value for most categories (displayed or not) is 1. Fields with higher boost values have increased priority when returning search results. For example, the "creator" field has a boost value of 2 which means that it is twice as important than the "category" field. The "id" field has a boost value of 3 indicating that it is three times more important than the "category" field. "modifier" has the highest value and is the most important for time based searches.

Refer to the following tables for lists of fields that can be added to each of the three categories:

asset (base)

id	systemID	type

common

categories	created	creator	description
durationTC	endTC	globalAssetID	location
modified	modifier	name	startTC
thumbnailURL			

94

tbmd			
category	created	creator	duration
id	modified	modifier	startTC
text	thumbnailURL	type	

> *Avid supports making changes to the asset, common, and tbmd categories. Creating additional categories is not supported.*

5. Add or adjust the boost values as desired. In the following example, the "created" parameter has been added and assigned a boost values of 3; giving it a higher value than "category" and "creator".

```
"tbmd": {
      "id": 3,
      "category": 1,
      "creator": 2,
      "created": 3,
      "modifier": 4
```

If you wish to add or adjust multiple fields or boost values, make all adjustments now.

When manually entering text, spacing and formatting is important. Notice that a comma is present after each value except for the final value. If a mistake is made, a red X appears in the configuration indicating an error has been made. In the following example, line 18 is missing the comma at the end of the line:

6. Once you are done editing, click the Query button in the right pane to save the changes.

If the query is accepted, a Response value (in milliseconds) will be returned. If you receive any errors, verify that the information you entered does not contain any incorrect characters.

7. Verify your changes by selecting the getBoostConfiguration item in the list of service options and click the Query button on the right to poll the service. The updated values appear under the "resultSet" list.

8. Restart the avid-acs-search service to activate the changes:

▶ For a single server, type the following from the Linux command prompt:

service avid-acs-search restart

▶ For a clustered system, type the following from the Linux command prompt:

crm resource restart AvidSearchEverywhere

> *If needed, you can restore the original configuration by selecting the eraseBoostConfiguration item in the list of service options and clicking the Query button in the "Request to: Local" pane on the right.*

Disabling Search Suggestions

When entering text in indexed search fields, search suggestions are displayed below the text box to assist users in finding the data they are looking for faster. While this is a useful feature, Media Index configurations that use large indexes might see a slow down in search performance when search suggestions are enabled. If desired, this feature can be disabled through the ACS Monitor tool.

To disable the search suggestion data:

1. Follow the instructions for "Accessing the Avid ACS Monitor Tool" on page 90.

2. Enter "avid.acs.media.index.configuration" (without the quotes) in the search field at the top of the page to find the configuration service.

3. Expand the configuration service options by clicking the Service Operations List icon in the upper-left corner of the configuration service field.

4. Locate the setConfigurationByField item in the list of service options and click the link for disableSuggestions.

   The following text appears in the right pane:

   ```
   Code ▾
   1   {
   2       "serviceType": "avid.acs.media.index.configuration",
   3       "serviceRealm": "global",
   4       "serviceVersion": 0,
   5       "op": "setConfigurationByField",
   6       "paramSet": {
   7           "acs_enable_suggestions": false
   8       }
   9   }
   ```

5. No adjustments are required, simply click the Query button in the right pane to save the change.

6. Restart the avid-acs-search service to activate the change:

   ▶ For a single server, type the following from the Linux command prompt:

   `service avid-acs-search restart`

   ▶ For a clustered system, type the following from the Linux command prompt:

   `crm resource restart AvidSearchEverywhere`

Configuring Search Criteria Behavior

When you specify search criteria for your search, Media Index returns results that are an exact match for the text field for the criteria used in the search. For example, if you specify "elections" as the text for the Name criterion, your search returns results only with that exact string. By contrast, when you use the search text field in the Search pane, by default Media Index uses text analysis and normalization which allows your search to return results for terms in different forms of the terms, such as conjugated verbs or plural words. For example, if you type "elections" in the Search text field, the natural language search returns results for "elect," "electing," or "election" in addition to the original term.

If desired, you can alter the default configuration for Media Index to return only exact matches in the search results.

To configure the search criteria behavior:

1. Follow the instructions for "Accessing the Avid ACS Monitor Tool" on page 90.

2. Enter "avid.acs.media.index.configuration" (without the quotes) in the search field at the top of the page to find the configuration service.

3. Expand the configuration service options by clicking the Service Operations List icon in the upper-left corner of the configuration service field.

4. Locate the setConfigurationByField item in the list of service options. This section includes a number of examples for changes than can be made to the configuration. Click on the link for enableSearchQueryLogger.

   The following text appears in the right pane:

```
1 ▾ {
2     "serviceType": "avid.acs.media.index.configuration",
3     "serviceRealm": "global",
4     "serviceVersion": 0,
5     "op": "setConfigurationByField",
6 ▾   "paramSet": {
7       "acs_enable_query_logger": true
8     }
9 }
```

   If you click on any of the examples in this section, you will notice that they all share the same basic structure. Line 7 is the only line that is different between them. Rather than entering all the text manually, you can alter the contents of this example to update the search criteria.

5. In the "Request to: Local" pane on the right, replace the text on line 7 with the following:

   ```
   "acs_criteria_filter_settings": "<value>"
   ```

   Where "<*value*>" represents the one of the following:

   ▸ `analyzed` — indexed search uses text analysis and normalization for the Criteria text fields (default)

   ▸ `notAnalyzed` — indexed search uses exact matches for the Criteria text fields

   ▸ `both` — indexed search uses both exact matches and text analysis for the Criteria text fields

   When entering text, make sure to include the two leading spaces and the quotes as shown in the example above.

6. Click the Query button in the right pane to save the changes.

   If the query is accepted, a Response value (in milliseconds) will be returned. If you receive any errors, verify that the information you entered does not contain any incorrect characters.

7. Verify your changes by selecting the getConfiguration item in the list of service options and click the Query button on the right to poll the service. The updated values appear under the "resultSet" list.

8. Restart the avid-acs-search service to activate the changes:

   ▸ For a single server, type the following from the Linux command prompt:

   service avid-acs-search restart

   ▸ For a clustered system, type the following from the Linux command prompt:

   crm resource restart AvidSearchEverywhere

Setting Elasticsearch Memory Usage

The Elasticsearch services require 4 GB of memory, which is set by default for standard configurations. If your configuration makes use of different memory requirements — for testing purposes — you can modify the Elasticsearch configuration files.

To modify the Elasticsearch configuration files:

1. Use the following Linux command to navigate to the directory containing the Elasticsearch configuration file and open it in the Linux text editor, vi:

   `vi /etc/sysconfig/elasticsearch`

2. Uncomment the following lines and edit the values for your configuration:

   `#ES_MIN_MEM=256m`

   `#ES_MAX_MEM=1g`

3. If you are running a cluster configuration, edit the elasticsearch configuration file on each cluster node.

4. Type the following command to restart the Elasticsearch service:

   ▶ For single server configurations, type the following command:

   `service elasticsearch restart`

   ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

   `crm resource restart elasticsearchEverywhere`

5. Type the following command to navigate to the directory containing the Elasticsearch tribe configuration file and open it in the Linux text editor, vi:

   `vi /etc/sysconfig/elasticsearch-tribe`

6. Uncomment the following lines and edit the values for your configuration:

   `#ES_MIN_MEM=256m`

   `#ES_MAX_MEM=1g`

7. Press the Escape key to return to command mode, then type `:wq` and press Return to write and quit the text editor.

8. If you are running a cluster configuration, edit the elasticsearch configuration file on each cluster node.

9. Type the following command to restart the Elasticsearch service:

   ▶ For single server configurations, type the following command:

   `service elasticsearch-tribe restart`

   ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

   `crm resource restart elasticsearchTribeEverywhere`

Configuring the PEBCo for Custom Configurations

For standard installations, you do not need to modify the properties file for the PEBCo. For custom configurations, you can edit the PAM agent services properties file.

To configure the PEBCo:

1. Navigate to the directory containing the PAM agent services properties file:

   /opt/avid/etc/pam-agent-service

2. Using the Linux text editor, vi, open the pamagent-service.properties file for editing:

   vi pamagent-service.properties

3. Configure ACS Bus Access Layer (BAL) URL by editing the following uncommented line:

   bal.amqp.url=amqp://localhost/acs

   Replace localhost with the URL for your ACS bus — for example, amqp://gl-ipix-mcs/acs

4. Type the following command to create a PEBCo instance:

   ▶ For single server configurations, type the following command:

   pam-agent-ctrl add <instance_name>

   ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

   pam-agent-cluster-ctrl add <instance_name>

   <instance_name> can be any meaningful name. By convention, the PEBCo instance name can include the database name and the engine host name (*[db_name]@[engine_host]*) — for example, AvidWG@avid-mpi-kbp.

5. Press the Escape key to return to command mode, then type **:wq** and press Return to write and quit the text editor.

6. Type the following command to start the PEBCo instance:

   ▶ For single server configurations, type the following command:

   pam-agent-ctrl start <instance_name>

   ▶ For cluster configurations, type the following command on one of the nodes in your cluster:

   pam-agent-cluster-ctrl start <instance_name>

Changing the Default Search Type

By default, the Search bar and the Search pane in Media Central UX use the federated search. You can specify the default search type in MediaCentral UX by editing a properties file in your MediaCentral configuration.

If you work with a cluster configuration, you must edit the application properties file on the master and slave nodes. Editing the file on load-balancing nodes is not required.

To specify the default search and search behavior:

1. Log in to the MCS server as the 'root' user. If you have a clustered configuration, log into the master node.

2. Navigate to the directory containing the MediaCentral application properties file:

   /opt/avid/etc/avid/avid-interplay-central/config

3. Using the Linux text editor, vi, open the application.properties file for editing:

   vi application.properties

In some cases, the application.properties file might not exist. Create the file using vi and add the lines indicated in the steps below.

You can examine the contents of the default file in the following directory: /opt/avid/avid-interplay-central/config. *However, do not make your changes in that file, since it is overwritten any time you upgrade MCS. Make your changes in the file you create in the* /opt/avid/etc/avid/avid-interplay-central/config *location as indicated in this step.*

If you use the default file as a model, the one you create should only contain the values you wish to override.

4. Modify the following lines to specify the default search type:

   ▶ To set the default search for all search panes in MediaCentral UX as the indexed search:

   system.client.search.indexed_as_default=true

   By default this value is set to "Federated."

   ▶ To set the default search used in the Search bar as the indexed search:

   system.client.search.quicksearch.default=Indexed

   By default this value is set to "Federated."

5. Press the Escape key to return to command mode, then type **:wq** and press Return to write and quit the text editor.

6. To activate the changes made in the application.properties file, type one of the following commands to restart the avid-interplay-central service (for single-server configurations) or the AvidIPC resource (for clustered configurations):

   ▶ For single-server systems: **service avid-interplay-central restart**

   ▶ For clustered configurations: **crm resource restart AvidIPC**

Changing Replica Counts

You can change the number of replicas for all indexes, for indexes at the specified alias, or for a specific index.

You can view index names and any index alias names by using the Elasticsearch head plug-in Web page, which you can view by opening a browser and navigating to http://*<server_name>*:9200/_plugin/head.
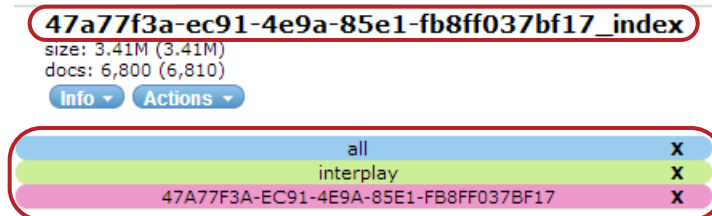
Elasticsearch head plug-in, displaying the index name at the top and index alias names at the bottom

To change the replica count for all indexes, do the following:

▸   Type the following command:

    `avid-acs-elastic-index-setup -m setReplicaCount -r <number_of_replicas>`

    For example, to set the replica count to 5, type the following:

    `avid-acs-elastic-index-setup -m setReplicaCount -r 5`

To change the replica count for all indexes that have a specific alias, do the following:

▸   Type the following command:

    `avid-acs-elastic-index-setup -m setReplicaCount -r <number_of_replicas>  -t <alias_name>`

    For example, to set the replica count to 5 on all indexes using the myAlias alias, type the following: `avid-acs-elastic-index-setup -m setReplicaCount -r 5 -t myAlias`

To change the replica count for a specific index, do the following:

▸   Type the following command:

    `avid-acs-elastic-index-setup -m setReplicaCount -r <number_of_replicas> -t <index_name>`

    For example, to set the replica count to 5 on the index named myIndex, type the following: `avid-acs-elastic-index-setup -m setReplicaCount -r 5 -t myIndex`

Changing the Index Refresh Interval

When you use Media Index to index an asset, you must wait for the next refresh for that asset to appear in a search. Refreshing the index requires some system resources, so by default the index refreshes after regular intervals instead of after each indexing operation. The interval is defined by the index.refresh_interval setting, which you can modify in either the Elasticsearch configuration file (`/etc/elasticsearch/elasticsearch.yml`) or the settings for each index. Because the refresh operation draws on system resources, one way to improve indexing performance is by increasing refresh_interval. The Media Index command line tool provides an option to change the refresh interval per system, and you can use the tool to increase the indexing throughput.

> *By default, all systems have an index refresh interval of 1 second.*

To change the refresh interval for a system:

▶ Type the following command:

```
/opt/avid/bin/avid-acs-elastic-index-setup -m setRefreshInterval
--systemId="<id_of_the_system>" --interval=<n>
```

You can set the interval in seconds (for example, where *<n>* might be 1s, 10s, or 30s), minutes (for example, 1m, 5m, or 10m). You can also type **disable** instead of an interval to completely disable index refreshing.

D Updating Media | Index Property Schemas

Each Media Index index has a schema specific to its source system. You configure the schema in the Production Engine Bus Connector (PEBCo — for more information, see "Configuring and Using the PEBco Service" on page 37). If the list of fields in the PEBCo Canonical Data tab that you configure for indexing changes, then you must migrate the index to this new schema before the system can index new assets. Media Index performs the re-index operation transparently whenever a data source signals a change in the number of fields in its schema. Re-indexing only adjusts the schema, and the process cannot add the data that were not originally indexed. Removing a field from the schema does not affect the index. However, if you add a new field, the index does not have property data for the existing assets. While re-indexing takes care of making the field available for newly indexed assets, you must perform a resync operation to re-send all assets by the data source. This provides the data for already indexed assets.

The re-indexing requires the indexing service to copy all of the existing asset entries to a new copy of the index and then replace the existing index with the new, adjusted copy. Depending on the size of the index this operation might take a long time. While re-indexing is taking place, the index cannot accept any new assets. This affects all indexes and cannot be interrupted.

Alternately, you can delete the index and start indexing from the start, which can take less time. This eliminates the need to re-index, but it causes the search to be partly unavailable until the new indexing completes. While an index is not being updated, it is still readable. Once the indexing finishes, searching switches to the new schema without interruption for the user. Re-indexing, while a costly operation, ensures that the search is always available for the clients.

There is no one best approach for property changes. In general, make sure you apply schema changes only when you can afford the interruption and downtime. Re-indexing allows you to perform a schema change without affecting the users too much but at the cost of limiting updates to the index. Deleting and restarting the indexing process does not affect other indexes, but it takes some time for the index to fill up again.

Property Changes in Production and Schema Updates

The PEBCo offers multiple ways to apply property selection changes to the index. You can make property selections in the Canonical Data tab of the PEBCo settings in the Interplay Administrator at any time. The actual change to the schema is performed when you restart the indexing process, either by pressing the Start button or the Resync button. Start continues the indexing where it left off, while Resync restarts the indexing from the beginning and sends all assets to the index again. Pressing Start or Resync without a change in the property selection does not cause a re-index but only continue indexing or start from the beginning.

See the following topics for examples of the common update cases:

Changing Properties Without Immediately Updating All Assets

This approach adds properties to the index, but only new or updated assets carry the new properties. It performs a re-index so there is no interruption for the user. Index updates are blocked for all indexes until the re-indexing completes.

After pressing Start, you might receive a timeout error after 10 seconds. This is expected, as the re-index operation needs more than 10 seconds to complete. Do not press Start again, but allow the re-indexing to finish. See "Monitor the Re-index Process" on page 106 for information on determining when a re-indexing operation finishes. Once the re-indexing completes, press Start again and indexing resumes normally.

To change properties in a schema without updating all assets:

1. Start Interplay Administrator and log on to the Interplay Production database that you want to be indexed.

2. In the Site Settings section of the Interplay Administrator, click the Production Engine Bus Connector icon.

3. Click the Canonical Data tab.

4. Change the property selection in the Canonical Data tab, and then click Apply. For information on using the Canonical Data tab, see "Configuring and Using the PEBco Service" on page 37.

5. Click the Manage/Status tab.

6. Click Stop to stop the service.

7. Click Start to start the service.

Changing Properties and Immediately Updating All Assets

This procedure follows the steps in the re-indexing with one difference: since indexing starts from the beginning, it takes some time before new changes are sent to the index. Use this method if it is important to get the new properties into the index. If getting new assets immediately into the index is the highest priority, then you can Resync at a later point and go with the re-index method (see "Changing Properties Without Immediately Updating All Assets" on page 104).

To change properties in a schema and update all assets:

1. Start Interplay Administrator and log on to the Interplay Production database that you want to be indexed.

2. In the Site Settings section of the Interplay Administrator, click the Production Engine Bus Connector icon.

3. Click the Canonical Data tab.

4. Change the property selection in the Canonical Data tab, and then click Apply. For information on using the Canonical Data tab, see "Configuring and Using the PEBco Service" on page 37.

5. Click the Manage/Status tab.

6. Click Stop to stop the service.

7. Click Resync to update the whole database content independent from any existing saved timestamp. Resyncing the database might take a long time to complete and will invalidate the existing content of your index. During this process, the search only returns assets that are already processed within this run.

Changing Properties and Restarting With an Empty Index

You can use this method if you can accept that search remains incomplete until the initial indexing finishes. It saves the time needed for re-indexing, but depending on the number of assets in the database it can take quite some time for the initial indexing to finish. This method does not affect the other indexes, so if you are running a multi-system index and want to avoid blocking the indexing process for the other systems then you can consider this approach.

This method is not slower than the resync method. The only difference is in the interruption for the users. Depending on the time the re-indexing needs, it might be significantly faster.

To change schema properties and restart with an empty index:

1. Start Interplay Administrator and log on to the Interplay Production database that you want to be indexed.

2. In the Site Settings section of the Interplay Administrator, click the Production Engine Bus Connector icon.

3. Click the Canonical Data tab.

4. Change the property selection in the Canonical Data tab, and then click Apply

5. Select the Manage/Status tab.

6. Click Stop to stop the service.

7. Sign in to MediaCentral UX as an administrator.

8. Select System Settings from the Layout selector.

    The System Settings layout opens.

9. In the Settings pane select Media | Index > Indexes.

    The Details pane displays the settings for your indexes.

10. To remove the index, click Delete.

⚠️ **Deleting an index removes all data from the index. You cannot undo this action.**

    The Delete System dialog box opens.
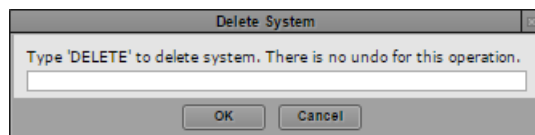
11. Type **DELETE** to delete the index, and then click OK. You can cancel the action by clicking Cancel.

    The index and its data are removed from the central index.

12. In Interplay Administrator, and click Resync to update the whole database content.

Monitor the Re-index Process

You can monitor the re-index process in the Cluster State pane of the Media Index System Settings. During the re-index, the process creates a new, temporary index which gradually fills up with copies of the existing assets. Once the re-index is complete, Media Index deletes the original index and makes the new index active. By monitoring the number of documents in the existing and the new index, you can monitor the progress of the re-index process.

You can check the cluster state by selecting Media Index > Cluster State in the System Settings for MediaCentral UX:

• If there is no active migration, all index labels appear in black.

• If the index has an active migration, its label appears in red with an indication of the migration progress.
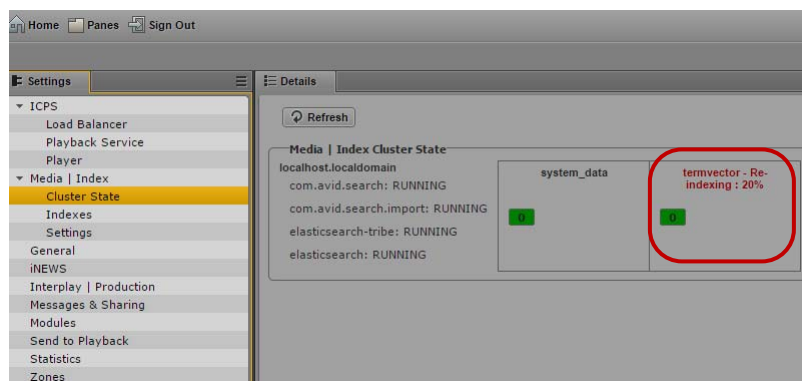
You can also monitor re-indexing using the Elasticsearch Head plug-in. You can access the Head plug-in by opening a browser and navigating to http://*[server_name]*:9200/_plugin/head, where *[server_name]* is any of the nodes in the cluster running Elasticsearch.

Before re-indexing the Head plug-in displays one entry for each index and one system_data entry for the overall metadata store:
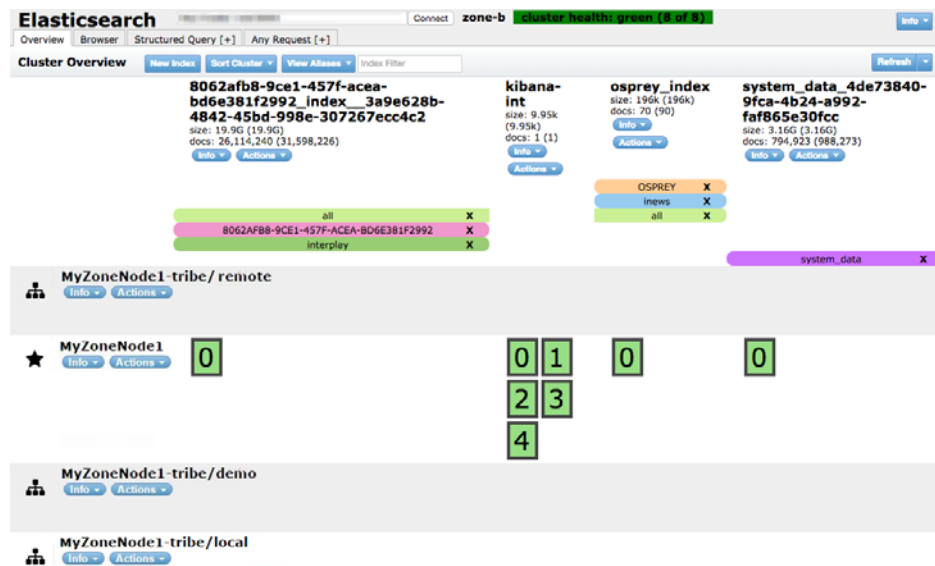
Original index status

When re-index starts, there is one additional column with the same basic name as the original index and a different ID at the end. The new index does not have the color bars that mark it as an active index. While the re-index progresses, the `docs` count slowly increases until it reaches the same number as the original index. Once it reaches the same number, the re-index is complete and the original index disappears with the new index getting the color bars of the original index.
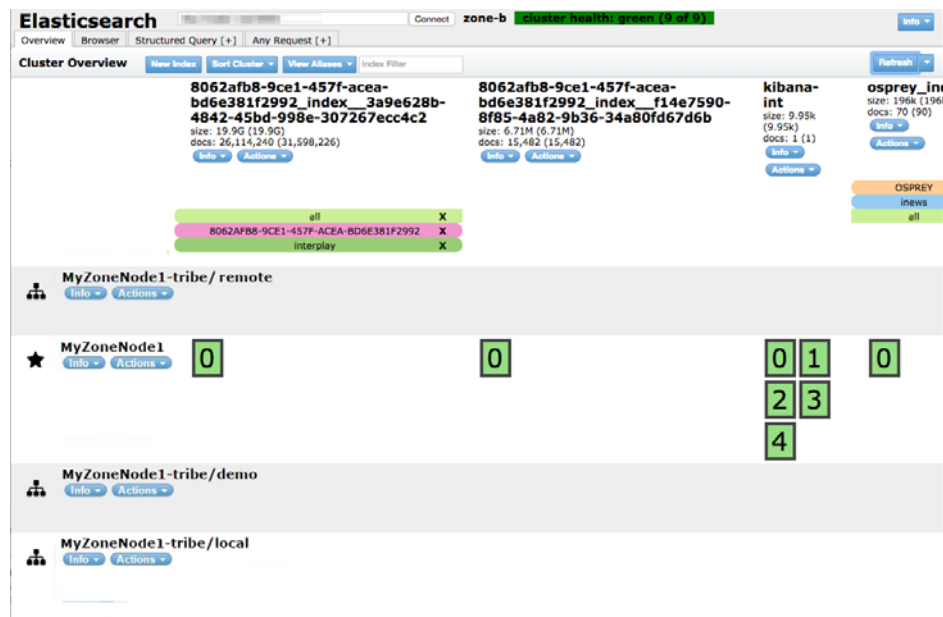
Re-index in progress

You can also monitor the progress and duration of the re-index process by comparing the number of assets listed in the Interplay Administrator tool (in the Object Statistics tab of the Database Information settings) with the assets listed in the MediaCentral UX System Settings (in the Assets list in the Media | Index > Indexes pane).

Database Information

Engine Start Time:	Oct 16, 2015 8:09:48 AM
Name:	AvidWG
Description:	
Database Location (Metadata)	\\Bsi-eng1\WG_Database$\AvidWG_Database\
Data Location (Assets)	\\Bsi-eng1\WG_Database$\AvidWG

| Database Statistics | Deletion Statistics | Command Statistics | **Object Statistics** | Connected User Information |

Total number of database objects:	369259
Database objects per database page:	2.41

≡ Database Object Type	≡ #
masterclip	27847
sequence	7402
subclip	7946
group	188
motioneffect	41
effect	714
renderedeffect	21826
filemob	120601
avid asset (in browse tree)	92529
folder	7255
user	94
other objects	82816

Details

BSI-ENG1

System Name:	BSI-ENG1	Apply
System ID:	9A85D001-0950-4C27-A789-1C5D2EF07A41	
System Type:	Interplay	Production (interplay)	
Languages:	en	
Created:	2015-09-15T10:46:19-0400	
Last Updates:	2069-12-31T20:00:00.000-0500	

Assets: Delete

masterclip	27515
subclip	7910
sequence	6843
effect	714
stereomasterclip	484
group	188
motioneffect	41

0 Custom Properties

Key	Type	Multi-value	Structured	English	

The duration of the re-indexing process depends on a number of variables, such as the number of markers in the assets.