

Index Server Upgrade Guide

Environment

Messaging Architects Netmail Archive 5.1.0 HP5
Messaging Architects M+Archive 2010.1.2 HP4

If you are upgrading from M+Archive to Netmail Archive 5.1, you will also need to upgrade to the new Exalead CloudView (Index Server) architecture. In order to do so, you must regroup as many existing build groups as possible onto the same CloudView servers. This guide includes useful information related to the new CloudView architecture and describes the steps you need to follow to move build groups between machines.

In this guide:

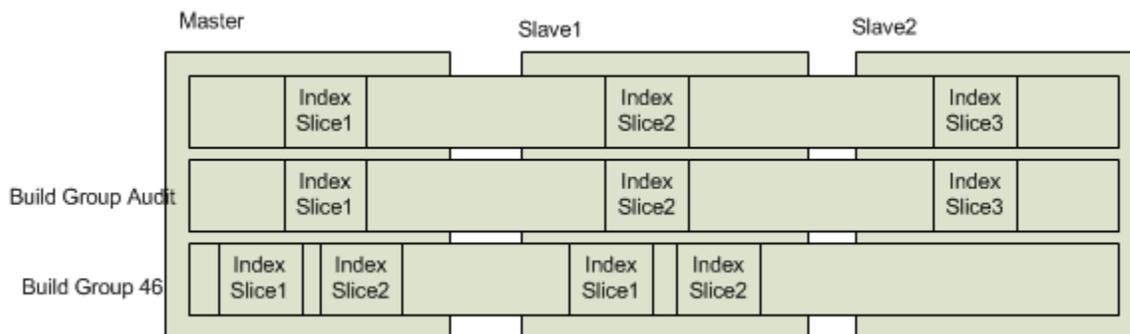
- Environment
- CloudView Infrastructure Changes
 - Old Architecture
 - New Architecture
 - Significance of the Architectural Changes
- Configuration Changes Overview
- Upgrading the Exalead Server (Single Node - Properly Sized System)
- Upgrading the Exalead Server (Single Node - Undersized)
 - Prepare the New Server
 - Install the Slave Server
 - Configure the New Slave Server
 - Move the Indexes
- Upgrading the Exalead Server (Multiple Nodes - Properly Sized System)
 - Update the Configuration
 - Move the Indexes
- Upgrading the Exalead Server (Multiple Nodes - Undersized)
 - Configuring an Existing Build Group
 - Adding Index Slices for 4.6 Conversion
- Adding One Node to the Exalead Cluster
 - Extend the Cluster
 - Install the Slave
- Post-Upgrade Verification Tasks

CloudView Infrastructure Changes

Old Architecture

The architecture of the Exalead version used prior to Netmail Archive 5.1 has the following characteristics:

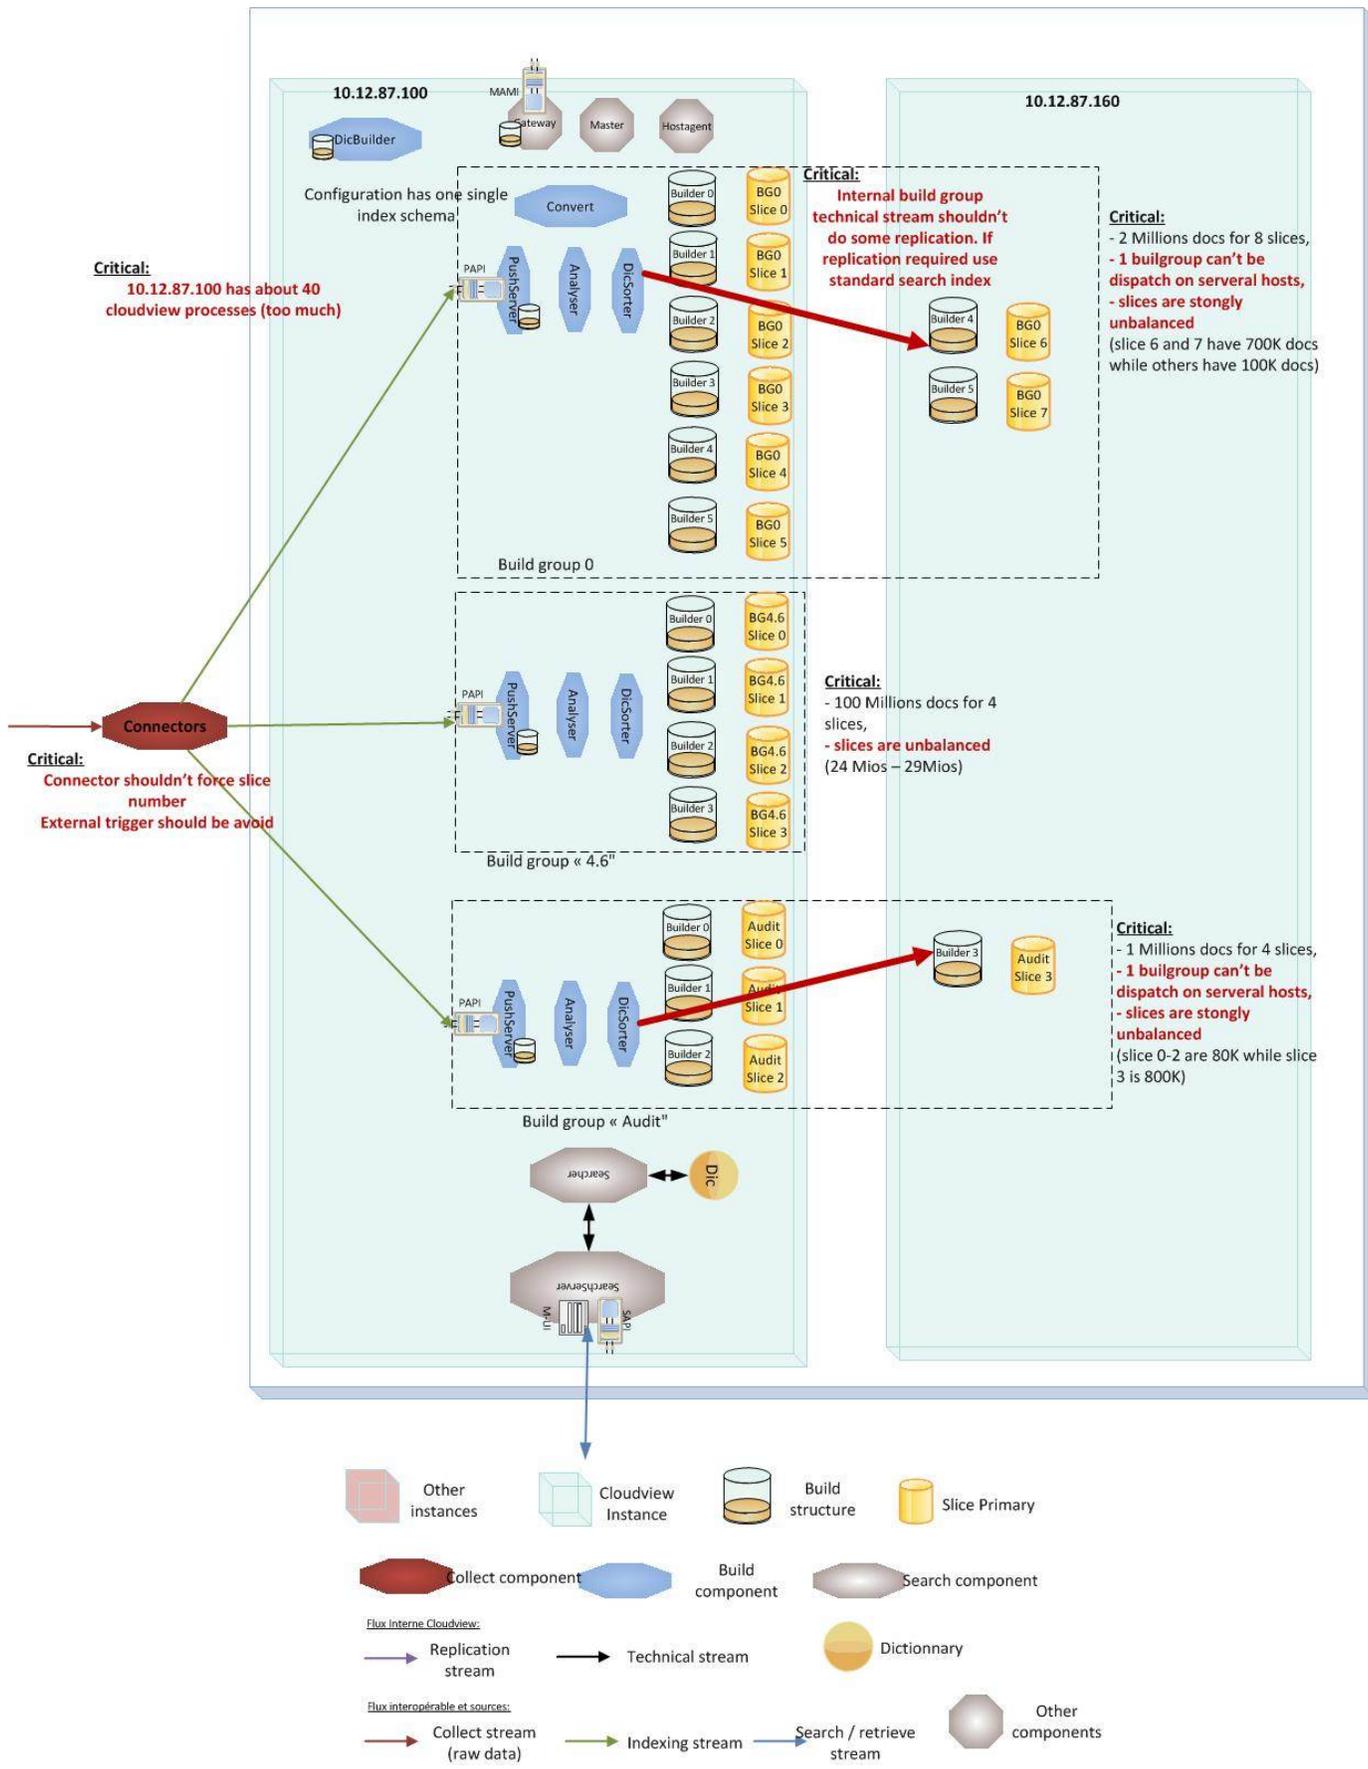
- A cluster is extended by adding a machine to the cluster and Index slices to the build group
- Internal Load Balancing pushes on the least busy machines/slices
- One push port is assigned per build group
- Distributed Analysis against all the servers
- One DIH exists per document type (we know where the documents are)



The following issues arise from this architecture:

- CloudView is not optimized to have build groups spread over the network, thus creating communication protocol resilience problems.

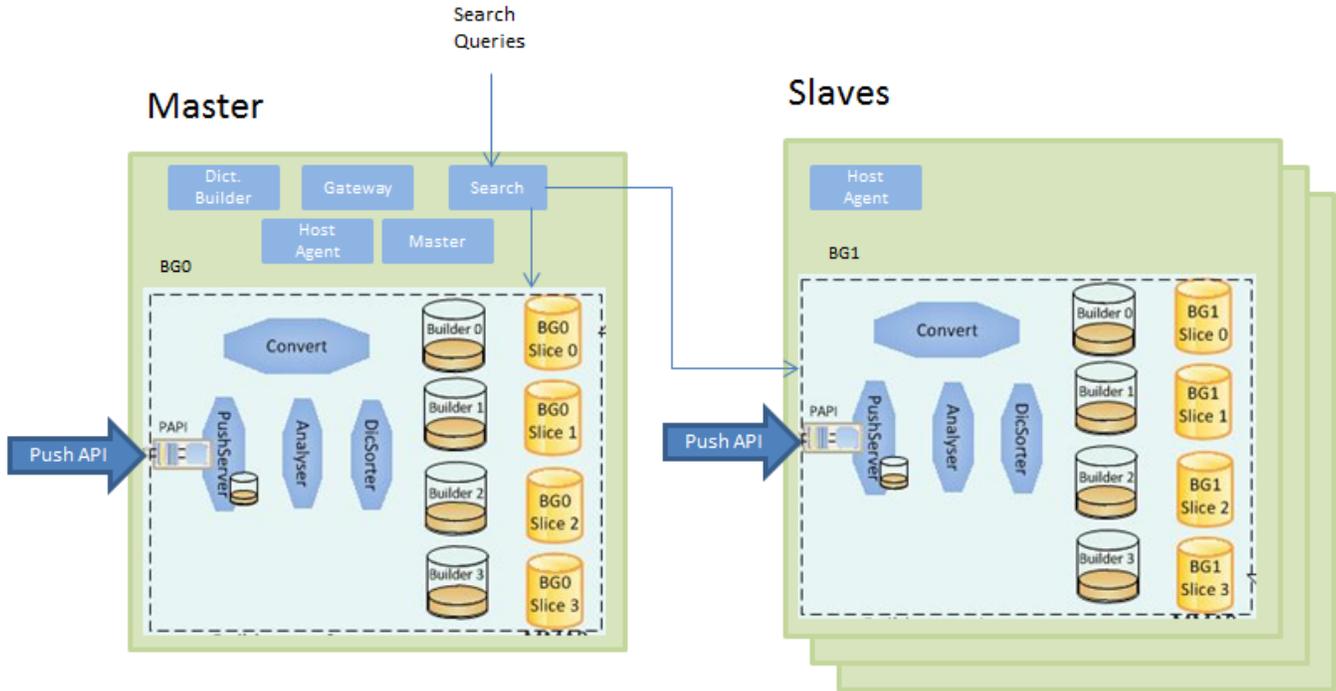
- CloudView does not like an “unbalanced” number of documents among the slices of a single build group. This causes search speed to decrease.
- Having multiple Analyzers per build group is not supported.
- Having multiple build groups per machine forces too many “Exa” processes to run in parallel, which is not recommended. Exalead recommends a single build group per machine.
- See sample deployment architecture below (Weitz).



New Architecture

The new Exalead CloudView architecture used in Netmail Archive 5.1 has the following improvements over the old architecture:

- A single build group exists per machine. This minimizes the number of processes running concurrently and prevents processing across the network.
- A single Build group type exists. All documents (audits, messages/attachments, 4.6 index records) are stored inside the same build group (e.g., BG0, BG1,...).
- Slice load balancing between slices is given back to CloudView.
- Growth is supported through the addition of a new server and a new build group to the cluster.
- Indexing load balancing (documents) is still supported by Netmail Archive based on the number of documents per machine.
- Index analysis trigger configuration is based on a queue size condition in order to regulate the amount of memory usage during analysis and minimize the import time (1GB).



Significance of the Architectural Changes

- **One DIH per Build Group: We no longer know where a document is indexed**
 - Index process deletes documents on all build groups before indexing a record.
 - Documents are pushed onto the least busy machine/build group.
- **Build Group names are no longer meaningful. Some Exceptions:**
 - Checkpoints are always stored on BGAudits if it exists; otherwise BG0 is used.
 - The BG_46 name is still used to detect if a build group contains converted documents (due to performance reasons).
- **If you are upgrading from the older version of Exalead to CloudView, you can have more than one build group per machine.**
- **If you are upgrading from the older version of Exalead to CloudView, BG_46, BGAudits and BG0 will be populated with "mismatched" document types.**
- **The ExaleadClusterManager tool is still required to extend the cluster and evaluate the number of index slices.**
- **The new conversion process is similar, except that it no longer relies on the BG_46 name to detect empty slices.**
 - Conversion allows for the parallel conversion of multiple servers because each build group has its own DIH.

Configuration Changes Overview

Single Node Configurations

Netmail 5.1 upgrades

- Properly sized system:
 - Three build groups: BG0, BGADT, BG4.6
 - No build group configuration changes

- Update the Index builder trigger configuration (size condition)
- Downside: Multiple "Exa" processes running
- Undersized system:
 - Three build groups: BG0, BGADT, BG4.6
 - Move BG0 to a new server
 - Update the Index builder trigger configuration (size condition)
 - Upside: New documents will get indexed against BG0 on the new server: two BGs on Svr1, one BG on Svr2

M+Archive 2010.1.x upgrades

- Properly sized system
 - Migrate/convert ALL v4.6 slices to BG0
 - New documents will be pushed against BG0
 - Upside: Only one build group per server
- Undersized system
 - Migrate/convert a subset of instances to BG0 on Svr1 until close to maximum capacity
 - Migrate/convert the remaining instance to BG1 on Svr2
 - Upside: Only one build group per server; added capacity.

Multi-Node Configurations

Netmail 5.1 upgrades

- Properly sized systems
 - Three build groups: BG0, BGADT, BG4.6
 - Consolidate build groupw per server by moving slices: BG4.6 + BGADT on Svr1, BG0 on Svr2
 - Update the Index builder trigger configuration (size condition)
 - Downside: Multiple "Exa" processes running on Svr1
- Undersized systems
 - Three build groups: BG0, BGADT, BG4.6
 - Consolidate build groups per server by moving slices: BG4.6 on Svr1, BGADT on Svr2 and BG0 on the new server
 - Update the Index builder trigger configuration (size condition)
 - Upside: New documents will get indexed against BG0 on the new server

M+Archive 2010.1 upgrades

- Properly sized systems
 - One build group per server
 - Migrate/convert v4.6 slices to build groups on Svr1, Svr2, etc. until close to maximum capacity
 - Migrate/convert the remaining v4.6 slices to the build group running on the last server
 - New documents will be pushed against the last & smallest build group
 - Upside: Only one build group per server
- Undersized systems
 - One build group per server
 - Add more server(s)
 - Migrate/convert v4.6 slices to build groups on Svr1, Svr2, etc. until close to maximum capacity
 - Migrate/convert the remaining v4.6 slices to the build group running on the last server
 - New documents will be pushed against the last and smallest build group
 - Upside: Only one build group per server

Upgrading the Exalead Server (Single Node - Properly Sized System)

Simply stop the Netmail Indexing service and run the Exalead installer from the SP1 package. It will detect that this is an upgrade scenario. Although no Exalead components have been updated for the 5.1 release, the ExaleadClusterManager tool needs to be updated for future use.

Upgrading the Exalead Server (Single Node - Undersized)

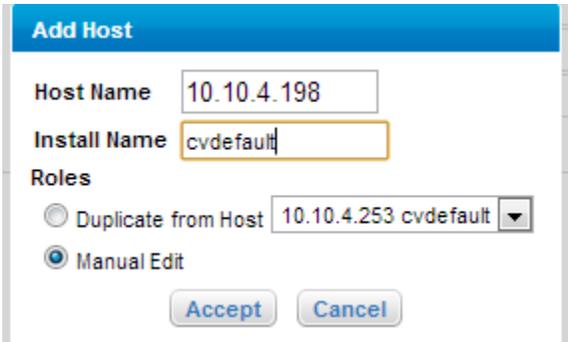
Stop the Netmail Indexing service and run the Exalead installer from the SP1 package. It will detect that this is an upgrade scenario. Although no Exalead components have been updated for the 5.1 release, the ExaleadClusterManager tool needs to be updated for future use.

Prepare the New Server

All the configuration changes are done using the CloudView Roles editor console. Follow these steps to prepare your server:

1. From the CloudView Roles editor console, click **Add Host**. A dialog box pops up, requesting the following new host parameters:
 - **Host Name:** Enter the IP address of the new host.
 - **Install Name:** Enter **cvdefault**.
 - **Roles:** Select **Manual Edit**.

Click **Accept**.



Add Host

Host Name

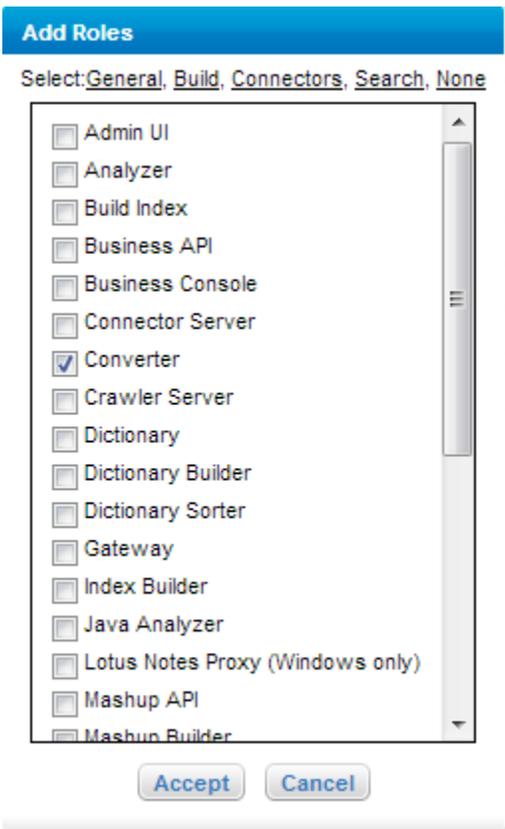
Install Name

Roles

Duplicate from Host ▼

Manual Edit

2. The **Add Roles** dialog box appears. Select **Converter**, and then click **Accept**.



Add Roles

Select: General, Build, Connectors, Search, None

Admin UI

Analyzer

Build Index

Business API

Business Console

Connector Server

Converter

Crawler Server

Dictionary

Dictionary Builder

Dictionary Sorter

Gateway

Index Builder

Java Analyzer

Lotus Notes Proxy (Windows only)

Mashup API

Mashup Builder

3. Edit the new Converter created on the new host:

- Change the **Instance** parameter to **c1**.
- Change the **Last** port value from 10099 to **10199**.

10.10.4.198
Install : cvdefault

▼ Ports

First
Host Agent
Last

[Add Roles](#) [Delete Host](#)

Color role (for visual help only)

▼ Converter new ✕

Instance

Click **Apply Changes** at the top of the screen to save the configuration changes.

4. Restart the Master service.

Install the Slave Server

To install the Slave server, run through the installer and then start the service the service;

Configure the New Slave Server

Follow these steps to configure the new Slave server:

1. For one of the build groups, move the roles from the Master machine to the Slave machine. The biggest build group should go on the more powerful machine. The roles can be moved by dragging and dropping them in the appropriate place in the GUI. Every role having the chosen build group name in parenthesis needs to be moved.

Host	General	Build	Connectors	Search
10.10.4.253 Install : cvdefault ▼ Ports First <input type="text" value="10000"/> Host Agent <input type="text" value="10009"/> Last <input type="text" value="10199"/> Add Roles Delete Host	► Scheduler ✕	► Analyzer (bgAudits) ✕	► Connector Server (exa) ✕	► Searcher (s0) ✕
	► Master ✕	► Dictionary Sorter (bgAudits) ✕	► Connector Server (java) ✕	► Search API ✕
	► Gateway ✕	► Build Index (bgAudits - 0) ✕	► Push Server (bgAudits) ✕	► Search UI ✕
	► Admin UI ✕	► Index Builder (bgAudits - 0) ✕		► Dictionary ✕
	► PerfMonitor ✕	► Converter ✕		
	► ResourceBuilder ✕			
	► DictionaryBuilder ✕			
10.10.4.198 Install : cvdefault ▼ Ports First <input type="text" value="10000"/> Host Agent <input type="text" value="10009"/> Last <input type="text" value="10199"/> Add Roles Delete Host		► Converter ✕	► Push Server (bg0) ✕	
		► Analyzer (bg0) ✕		
		► Dictionary Sorter (bg0) ✕		
		► Build Index (bg0 - 0) ✕		
		► Index Builder (bg0 - 0) ✕		
		► Build Index (bg0 - 1) ✕		
	► Index Builder (bg0 - 1) ✕			

2. Change the push port for the Push Server (bg0) to 10198 (or to base port + 198).

▼ Push Server ✕

Build Group

Papi Port

3. Click **Apply Changes** at the top of the screen to save the configuration changes.

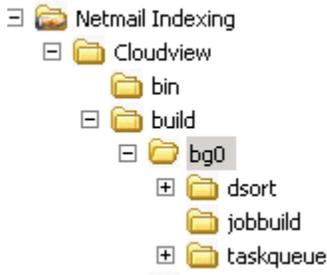
4. Restart all the services.

Move the Indexes

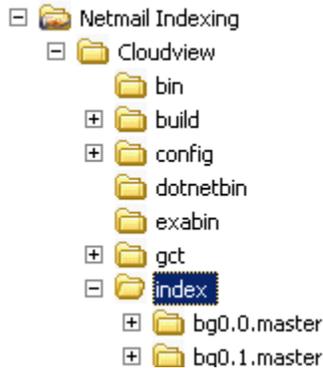
To move the indexes, follow these steps:

1. Make sure all services are stopped on both machines.

2. Move the DIH folder from the Master to the Worker. The DIH is located in the **build** folder of the Index folder. You need to move the **bg** folder associated with the moved build group from the Master to the Slave. In our example, this is the bg0 folder. The already existing folder on the destination machine should be deleted before the move.



3. Move the Index folders from the Master to the Slave. The indexes are stored under the **index** folder. The already existing folders on the Slave should be removed. You will have to move as many folders as there are slices defined in the configuration of the build group. The Index folder names that are to be moved start with the build group name (i.e., in our example, it is bg0).



4. Restart all the services, and make sure the document count of each build group is the same after the move.

Upgrading the Exalead Server (Multiple Nodes - Properly Sized System)

Stop Netmail Indexing services on all servers and run the Exalead installer from SP1. It will detect that this is an upgrade. Although no Exalead components have been updated for the 5.1 release, the ExaleadClusterManager tool needs to be updated for future use.

Update the Configuration

Start the Netmail Indexing services and open the Cloudview web console.

The roles must be aggregated so that the build groups are on the same machine. This is done by dragging and dropping the roles from one machine to the other. In our example, all the slices from BGAudit will be moved to the Master and all slices from BG0 moved to the Slave. Here is a sample configuration before the move:

Host	General	Build	Connectors	Search
10.10.4.227 Install : cvdefault ▼ Ports First <input type="text" value="10000"/> Host <input type="text" value="10009"/> Agent <input type="text" value="10009"/> Last <input type="text" value="10199"/> Add Roles Delete Host	<input type="checkbox"/> Scheduler ✘	<input type="checkbox"/> Analyzer (bg0) ✘	<input type="checkbox"/> Push Server (bg0) ✘	<input type="checkbox"/> Searcher (s0) ✘
	<input type="checkbox"/> Master ✘	<input type="checkbox"/> Dictionary Sorter (bg0) ✘	<input type="checkbox"/> Connector Server (exa) ✘	<input type="checkbox"/> Search API ✘
	<input type="checkbox"/> Gateway ✘	<input type="checkbox"/> Analyzer (bgAudits) ✘	<input type="checkbox"/> Connector Server (java) ✘	<input type="checkbox"/> Search UI ✘
	<input type="checkbox"/> Admin UI ✘	<input type="checkbox"/> Dictionary Sorter (bgAudits) ✘	<input type="checkbox"/> Push Server (bgAudits) ✘	<input type="checkbox"/> Dictionary ✘
	<input type="checkbox"/> PerfMonitor ✘	<input type="checkbox"/> Build Index (bg0 - 0) ✘		
	<input type="checkbox"/> ResourceBuilder ✘	<input type="checkbox"/> Index Builder (bg0 - 0) ✘		
	<input type="checkbox"/> DictionaryBuilder ✘	<input type="checkbox"/> Build Index (bgAudits - 0) ✘		
		<input type="checkbox"/> Index Builder (bgAudits - 0) ✘		
		<input type="checkbox"/> Converter ✘		
		<input type="checkbox"/> Build Index (bg0 - 1) ✘		
		<input type="checkbox"/> Index Builder (bg0 - 1) ✘		

10.10.4.198 Install : cvdefault ▼ Ports First <input type="text" value="10000"/> Host <input type="text" value="10009"/> Agent <input type="text" value="10009"/> Last <input type="text" value="10199"/> Add Roles Delete Host	<input type="checkbox"/> Build Index (bg0 - 2) ✘			
	<input type="checkbox"/> Build Index (bgAudits - 1) ✘			
	<input type="checkbox"/> Index Builder (bg0 - 2) ✘			
	<input type="checkbox"/> Index Builder (bgAudits - 1) ✘			
	<input type="checkbox"/> Converter ✘			
	<input type="checkbox"/> Analyzer (bgAudits) ✘			
	<input type="checkbox"/> Analyzer (bg0) ✘			
	<input type="checkbox"/> Build Index (bg0 - 3) ✘			
	<input type="checkbox"/> Index Builder (bg0 - 3) ✘			
	<input type="checkbox"/> Index Builder (bg0 - 3) ✘			

This is the configuration after:

Host	General	Build	Connectors	Search
10.10.4.227 Install : cvdefault ▼ Ports First <input type="text" value="10000"/> Host <input type="text" value="10009"/> Agent <input type="text" value="10009"/> Last <input type="text" value="10199"/> Add Roles Delete Host	<input type="checkbox"/> Scheduler ✘	<input type="checkbox"/> Analyzer (bgAudits) ✘	<input type="checkbox"/> Connector Server (exa) ✘	<input type="checkbox"/> Searcher (s0) ✘
	<input type="checkbox"/> Master ✘	<input type="checkbox"/> Dictionary Sorter (bgAudits) ✘	<input type="checkbox"/> Connector Server (java) ✘	<input type="checkbox"/> Search API ✘
	<input type="checkbox"/> Gateway ✘	<input type="checkbox"/> Build Index (bgAudits - 0) ✘	<input type="checkbox"/> Push Server (bgAudits) ✘	<input type="checkbox"/> Search UI ✘
	<input type="checkbox"/> Admin UI ✘	<input type="checkbox"/> Index Builder (bgAudits - 0) ✘		<input type="checkbox"/> Dictionary ✘
	<input type="checkbox"/> PerfMonitor ✘	<input type="checkbox"/> Converter ✘		
	<input type="checkbox"/> ResourceBuilder ✘	<input type="checkbox"/> Build Index (bgAudits - 1) ✘		
	<input type="checkbox"/> DictionaryBuilder ✘	<input type="checkbox"/> Index Builder (bgAudits - 1) ✘		

10.10.4.198 Install : cvdefault ▼ Ports First <input type="text" value="10000"/> Host <input type="text" value="10009"/> Agent <input type="text" value="10009"/> Last <input type="text" value="10199"/> Add Roles Delete Host	<input type="checkbox"/> Index Builder (bg0 - 2) ✘	<input type="checkbox"/> Push Server (bg0) ✘		
	<input type="checkbox"/> Converter ✘			
	<input type="checkbox"/> Analyzer (bg0) ✘			
	<input type="checkbox"/> Build Index (bg0 - 3) ✘			
	<input type="checkbox"/> Index Builder (bg0 - 3) ✘			
	<input type="checkbox"/> Dictionary Sorter (bg0) ✘			
	<input type="checkbox"/> Build Index (bg0 - 0) ✘			
	<input type="checkbox"/> Index Builder (bg0 - 0) ✘			
	<input type="checkbox"/> Build Index (bg0 - 1) ✘			
	<input type="checkbox"/> Index Builder (bg0 - 1) ✘			
	<input type="checkbox"/> Build Index (bg0 - 2) ✘			

Change the push port for the Push Server (bg0) to 10198 (or to base port + 198) on the slave server.

▼ Push Server
✘

Build Group

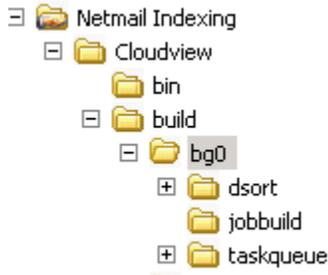
Papi Port

Click **APPLY** to save and effect the changes.

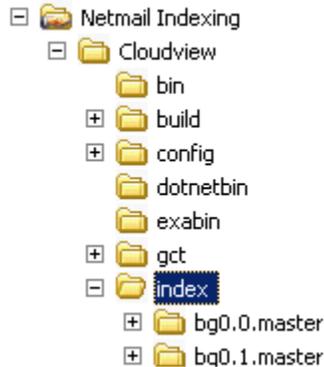
Move the Indexes

To move the indexes, follow these steps:

1. Make sure all services are stopped on both machines.
2. For build groups which have been moved off the Master, move their DIH folder to the server hosting that build group. The DIH is located in the **build** folder of the Index folder. You need to move the **bg** folder associated with the build group. In our example, this is the bg0 folder. The already existing folder on the destination machine should be deleted before the move.



3. Move the Index folders of a build group to the server hosting that build group. The indexes are stored under the **index** folder. The already existing folders on the destination should be removed. You will have to move as many folders as there are slices defined for the build group. The Index folder names that are to be moved start with the build group name.



4. Following the same process, move the Index folders from the Slave to the Master for the build group that is left on the Master server. In our example, we moved all the folders stating with BgAudits from the slave to the master.
5. Restart all the services. Login to the Cloudview console and make sure the document count of each build group is the same after the move.
6. Log into the Netmail Archive web console and verify that the document caps on the index servers is appropriate.

Upgrading the Exalead Server (Multiple Nodes - Undersized)

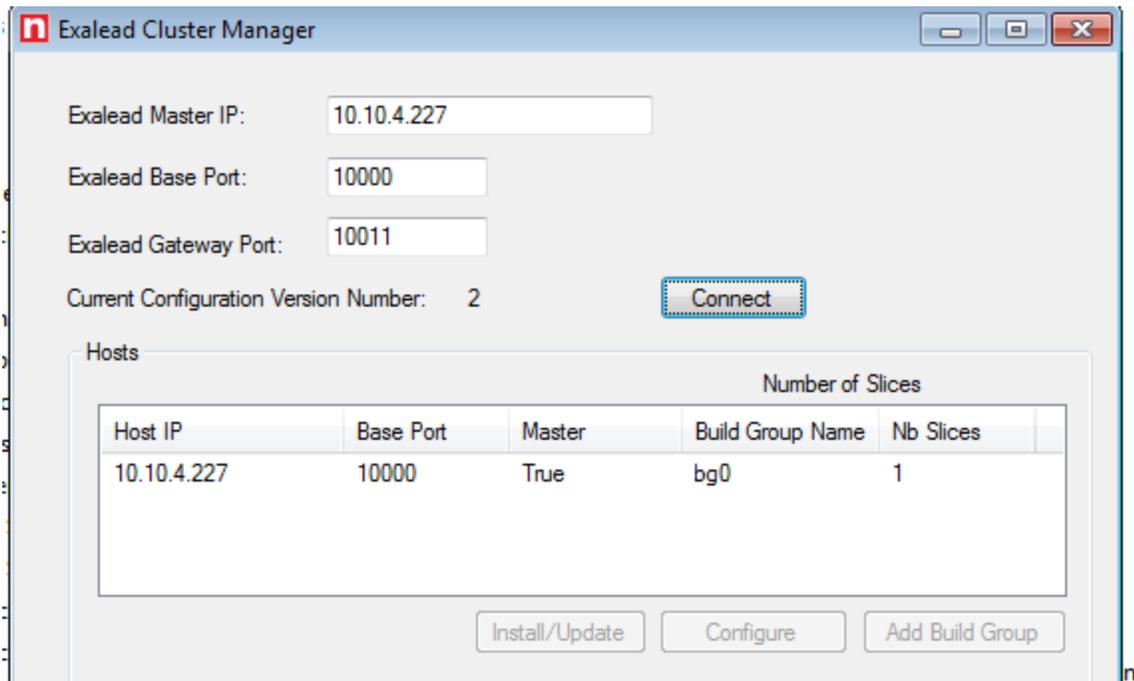
Complete the steps above for a properly sized system.

Run *ExaleadClusterManager.exe*. This tool is used to configure the existing build group properly and to add a new node to the cluster.

Configuring an Existing Build Group

Follow these steps to configure an existing build group:

1. Start CloudView. This must be done before running the tool.
2. Open the ExaleadClusterManager tool.
3. Enter IP address and base port of the CloudView Master.
4. Click **Connect**.



5. Select the **bg0** entry and hit click **Configure**.
6. Review the number of cores on the machine to make sure it is correct, and then click **Evaluate**.
7. Click **OK**. A message appears, prompting you to confirm the actions to be taken. Click **Accept**.
8. Click **Apply**. A message appears when the configuration has been updated, and the new configuration number is displayed.
9. Close the tool.
10. Restart all Exalead Services.

Adding Index Slices for 4.6 Conversion

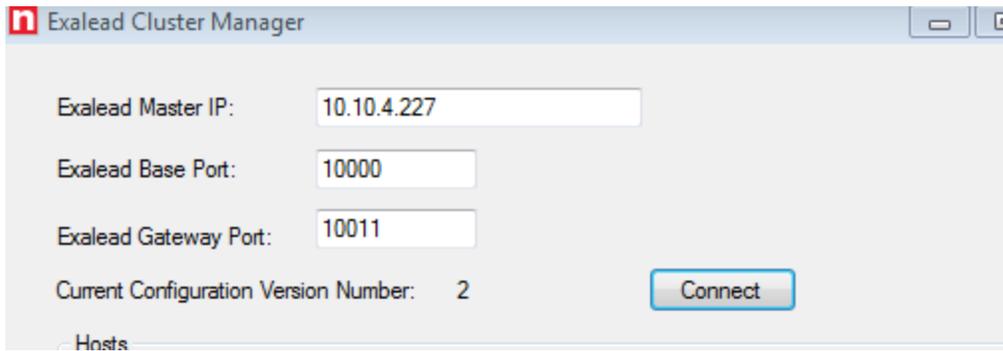
To add Index slices for 4.6 conversion, follow these steps:

1. Start CloudView. This must be done before running the tool.
2. Open the ExaleadClusterManager tool.
3. Enter IP address and base port of the CloudView Master.
4. Click **Connect**.
5. Select the build group, and click **Configure**.
6. Select **Host is used to contain 4.6 converted indexes**, and enter the number of slices to be converted on this machine.
7. Click **OK**. A message appears, prompting you to confirm the actions to be taken. Click **Accept**.
8. Click **Apply**. A message appears when the configuration has been updated, and the new configuration number is displayed.
9. Close the tool.
10. Restart all Exalead Services.

Adding One Node to the Exalead Cluster

Extend the Cluster

To extend the cluster, run the cluster config tool.



Install the Slave

To install the Slave, follow these steps:

1. Install CloudView on the Slave.
2. Restart the Master and Start the slave. Wait for the Slave to complete initialization.
3. Restart the Slave.

Post-Upgrade Verification Tasks

The following settings should be reviewed after the upgrade to make sure they have been applied correctly:

- **One analyzer per build group:** There should be only one analyzer per build group, and the analyzer role should be on the same server as the rest of the build group.
- **One converter per server:** Each machine should have a converter, and only one converter role should be present on each machine. Verify the number is unique compared to the other build group converters.
- **Add a trigger condition for both `ib0_standard` and `ib0_large` index builder policies:**
 1. From the CloudView administration console, select BuildGroups.
 2. Select the **Index Builders** tab.
 3. Select **`ib0_standard`**, and add a condition to trigger indexing:
 - a. Click **Add Condition**.
 - b. Modify the new condition so that it has the following parameters (as shown in the following figure):
 4. Select **`ib0_large`**, and apply the same condition as above.
- Update the **Inactivity Condition** for **`ib0_standard`**, and make sure that the value is set to 7200 seconds.

Size Condition	<input type="text"/>
Trigger indexing job when the "Min. Task Queue size" is reached when there is no pending job. The Task Queue size (before triggering) increases based on the number of pending jobs to a maximum of "Max. Task Queue size" (when the "Jobs threshold" is reached).	
Min. Task Queue size	<input type="text" value="1024"/> MB
Max. Task Queue size	<input type="text" value="2048"/> MB
Jobs threshold	<input type="text" value="5"/>

OR

Inactivity Condition	<input type="text"/>
Trigger after no task has been issued for	<input type="text" value="7200"/> seconds
Trigger after	<input type="text" value="300000"/> tasks have been issued

OR

Once all modification are done, click **Apply**.