

Using ECL Watch

Boca Raton Documentation Team



Using ECL Watch

Boca Raton Documentation Team

Copyright © 2020 HPCC Systems®. All rights reserved

We welcome your comments and feedback about this document via email to <docfeedback@hpccsystems.com>

Please include **Documentation Feedback** in the subject line and reference the document name, page numbers, and current Version Number in the text of the message.

LexisNexis and the Knowledge Burst logo are registered trademarks of Reed Elsevier Properties Inc., used under license.

HPCC Systems® is a registered trademark of LexisNexis Risk Data Management Inc.

Other products, logos, and services may be trademarks or registered trademarks of their respective companies.

All names and example data used in this manual are fictitious. Any similarity to actual persons, living or dead, is purely coincidental.

2020 Version 7.10.22-1

Introducing ECL Watch	4
Common Elements	5
ECL Watch Home Page	14
Activity	15
ECL Event Scheduler	24
Interface in ECL Watch	25
ECL Workunits	28
ECL Workunits Page	29
Using the ECL Playground	51
Files	58
Files	59
DFU Workunits Page	68
Landing Zones	75
Spray/Despray	79
XRef	92
Queries	95
Queries Page in ECL Watch	95
Operations	111
Topology	112
Disk Usage	117
Operations: Target Clusters	118
Dynamic ESDL	119
Preflight System Servers	125
Preflight Thor	129
Users Permissions	132
User Administration	132
Plugins	152
Ganglia in ECL Watch	153
Nagios in ECL Watch	155
Resources	157
A. HPCC Systems Session Management	158
Session Management	158

Introducing ECL Watch

ECL Watch is a service that runs on the Enterprise Services Platform (ESP), a middleware component on the HPCC Systems platform.

ECL Watch provides an interface to the HPCC Systems platform and allows you to view information and interrogate nodes to confirm all expected processes are running. It is a plugin that is useful for Systems Administrators to check processes, examine topology, and view logs. It is useful to ECL Programmers to monitor the status of jobs and files, and other pertinent information. This provides a simple view into the system and a means to perform Workunit and data files maintenance.

The ECL Watch interface is a browser based set of pages where you can access and interface with the HPCC Systems platform. To Run **ECL Watch** using your browser, go to the node on your system running **ESP/ECL Watch** on port 8010. For example, <http://nnn.nnn.nnn.nnn:8010>, where nnn.nnn.nnn.nnn is your ESP/ECL Watch node's IP address. This will launch ECL Watch.

ECL Watch is organized by categories based on several system functions. These functions or actions are accessible through the links along the top of the main home page. However there are several ECL Watch items with common functionality universally across pages of ECL Watch.

Common Elements

There are common elements that are useful anywhere within ECL Watch.

User Login

When you open ECL Watch a log in window displays.

Figure 1. Login



You login with the credentials supplied by your Systems Administrator. After a configurable period of inactivity your ECL Watch session locks and you need to enter your credentials to unlock and resume your session. The session remains active for as long as there is regular keyboard or mouse interaction.

Additionally, there is a Logout menu item available you can select to close your session when you complete your work.

For more information and frequently asked questions about session management, please refer to Appendix A.

Open in New Page

The **Open in New Page** link allows you to open a new window directly to the item which you are viewing.

Figure 2. Open in New Page



Opening a new page of the current view is useful in several ways:

- **Open in New Page** is a shareable link. Right-click on **Open in New Page** and select *Copy link address* from the context menu. You can now save, share, or send that link and the recipient can open a window directly to the (new) page.
- New pages maintain filter states. If you have set a filter and share that link, the new page has the same filter set.
- The new page is live. (If that page has auto-refresh enabled) The recipient doesn't need to reset the page or refresh to see current state.

Maximize Pane



The **Maximize/Restore** icon appears throughout ECL Watch. Press the button to maximize the corresponding window pane. Press the button again to restore the pane to its original size.

List widgets

Figure 3. List Widgets



ECL Watch commonly displays items in the forms of a list. Lists of workunits, lists of logical files, lists of Roxie queries, etc. The following elements in ECL Watch work on any list widget:

- Download as CSV (Workunits, DFU Workunits, Logical files, and Roxie queries)
- Column check box. Select all rows by selecting the check box at top (Use Shift+Click to select a range)
- Sort by column header. Click on the column header to sort that column according to that column's contents.
- Resize Column Widths by dragging the borders.
- Filters that allow you to set criteria for the items to display in the list.

List Filters

Filter options on a list allow you to use specified criteria to set what you display in the list. Press the **Filter** action button to display the filter options. When you specify any Filter options, the Filter action button changes and then displays **Filter Set** and the filter indicator image lights up. When you then copy an **Open in New Page** link with the filter set, the filter condition persists. Filter conditions limit the results displayed until you clear the filter.

The Mine Button

The **Mine** button appears along the top of every list page. Check the Mine box to filter for only items that you own.

Figure 4. Mine button



Uncheck the Mine button to display all jobs.

Banner Items

The banner area of ECL Watch appears at the top of all pages.

Home Button

The HPCC Systems home button is also the main Activity menu link. More than just a decoration the home ECL Watch button actually opens up the Activity page, discussed in more detail in the ensuing chapter.

The icons to the left of the home button are links to different areas of ECL Watch. Each of the link areas are discussed in subsequent chapters.

Monitor indicator light

The Monitor indicator light is used to provide an at a glance snapshot of your overall system health. If you have monitoring and reporting configured (see *HPCC Systems Monitoring and Reporting*) then this light will display a color based on the the overall (configurable) system health. The light is black if there is no monitoring or reporting configured for your system.

Download as CSV

With most lists in ECL Watch you have the ability to download selected rows as a CSV formatted file. You can then open the CSV file in any spreadsheet application which supports CSV format.

Global Search

The global search box can be found on the navigation bar at the top of the ECL Watch page.

Figure 5. Global Search box



You can search DFU Workunits, ECL Workunits, Logical Files, and Queries using the global search box. The global search excludes ECL code by default. To search for ECL code within workunits, use the **ecl:** keyword. For example, **ecl:SOAPCALL**. To limit or filter your search results you can use keywords as displayed in the empty search box.

- file:** Preface the search string with *file:* to search Logical Files.
- wuid:.** Preface the search string with *wuid:* to search only Workunit ids.
- ecl:** Preface the search string with *ecl:* to search ECL code in the workunits.
- dfu:** Preface the search string with *dfu:* to search only DFU workunits.
- query:** Preface the search string with *query:* to search only published queries.

The global search box also supports using wildcards.

Examples of using the global search:

Enter *W201510** into the search box, and it will return all of the workunits from October 2015.

Enter *file:keys* into the search box, and it will return all of the logical files that contain "keys".

Figure 6. Global Search Example



Advanced Menu

There is a section at the top right on the navigation bar with some useful information and features. This section shows you who you are logged in as (if your system has authentication enabled). The Advanced menu is located on the right hand side of the navigation bar. There are several items that you can access from the advanced menu.

Figure 7. Advanced menu



There may be a number displayed on the menu link. A number displayed next to it indicates how many errors and warnings have been generated during your session. Click on the advanced menu to display a list of features.

Access the Advanced menu

You access the advanced menu items from the advanced menu link at the top right corner of ECL Watch.

The **Set Banner** link allows you to set a custom banner message at the top of your browser window when you open ECL Watch. Check the **Enable** box to enable a banner. You can use the banner to deliver a short message about the environment to users. Customize the appearance of the message banner with the other controls. Banner settings persist until the ESP Server restarts.

The **Set Toolbar** link allows you to customize the toolbar at the top of the ECL Watch page. Check the **Enable Environment Text** box to display the **Name of Environment** text at the top of the page and in the browser tab. Labeling browser tabs is helpful when working with multiple environments. Toolbar settings persist through restarts of the ESP Server.

The **Error/Warning(s)** link displays a tab showing you Errors, Warnings, and Information messages. You can filter this page by checking the boxes at the bottom of the tab. A copy facility is also provided.

The **Release Notes** link opens a new browser tab to the HPCC Systems release notes page where you can find more release specific information about the contents of each version of HPCC.

The **Documentation** link opens a new browser tab to the HPCC Systems documentation page, where you can view and download the HPCC Systems platform documentation.

The **Downloads** link opens a new browser tab to the HPCC Systems downloads page, where you can find and download the HPCC Systems platform, client tools, and plugins.

The **Additional Resources** link opens a submenu that provides links to areas on the HPCC Systems web site where you can find additional resources, such as the **Red Book**, **User Forums**, **Issue Reporting**, and the **Transition Guide**.

The **Configuration** link opens for viewing the XML version of the configuration file in use on your system.

The **About** link opens a dialog to display information about the version of the HPCC Systems platform installed on your server.

Logged In As

The **Logged In As:** links at the top of the ECL Watch page display information about the current user logged onto a cluster configured for authentication.

Figure 8. Logged In As



1. Click on your username link and the User Details window opens.

Figure 9. User Details window

The image shows a 'User Details' window with a title bar containing a close button. Below the title bar is a 'Save' button. The main content area displays the username 'FranklinX' in a large font. Below this, there are several fields: 'Username:' with the value 'FranklinX', 'Employee ID:' with the value '99999', 'First Name:' with the value 'Franklin', and 'Last Name:' with the value 'Xavier'. There are three empty text input fields for 'Old Password:', 'New Password:', and 'Confirm Password:'. The 'Password Expiration:' field is set to 'Never'. The window is overlaid on a dark, textured background that appears to be a terminal or command prompt window with some text visible at the bottom.

2. From The User Details page, you can Confirm the User Name that you are logged in as.

You can change your password.

Note that Administrator rights are needed to manage users and permissions.

Ensure you are using an account with Administrator rights if you intend to manage users or permissions.

3. Verify the password expiration date, or if password is set to expire.

Change Password

If authentication is enabled on your HPCC Systems platform, you can change your password, right from the User Details window.

1. Click on your username link under the **LOGGED IN AS:**

The User Details window opens. There are fields where you can change your password.

2. Enter your Old Password.

3. Enter your desired new password.

Make sure it meets whatever criteria your system may have for passwords.

4. Confirm your new password.

Make sure that it matches the password you entered in the previous field.

5. Press the Save button. It is in the upper left portion of the window.

Lock

Click the **Lock** link next to the username link to lock your session. You would then need to enter your credentials to unlock your session.

ECL Watch Home Page

Click on the **ECL Watch** home page link in the navigation bar at the top of the ECL Watch page to find the **Activity**, **Scheduler**, and **Search Results** links. You can access the respective pages from the links or tabs along the top of the ECL Watch home page.

Figure 10. ECL Watch home page link



There are several icons used throughout ECL Watch. The following table describes most of the icons that you will encounter.

Icon	Definition
	System cluster
	System cluster paused
	System cluster not found
	Workunit Compiled, Completed
	Workunit Running, Compiling, Debug Running
	Workunit Failed, Aborted
	Workunit Blocked, Scheduled, Wait, Uploading Files, Debug Paused, Paused
	Workunit Archived
	Workunit Aborting
	Workunit Submitted

	Workunit Deleted.
	Workunit Unknown State

Activity

The Activity tab displays activity on all clusters in the environment. The Activity page provides access to Cluster Job Queue administration tasks such as: monitoring progress, setting priority, moving a job up or down in the queue, pausing a job, aborting a job, and pausing or resuming a queue.

When you access the ECL Watch URL, the ECL Watch Home Page displays the Activity tab. To access the Activity tab from any other page in ECL Watch click on the **ECL Watch** image at the top of any page, as shown above.


Cluster Disk Usage

The Cluster Activity page displays graphs along the top showing cluster disk usage. Clicking on each of the images displays more information about the individual disk usage activity.

Cluster Activity

Information about your systems' clusters, and any activity on those clusters can be accessed from the ECL Watch Activity tab. Select the **Activity** tab link below the **ECL Watch** image in the navigation sub-menu. This displays the Cluster Activity tab.

Figure 11. Clusters Activity



All the system clusters display. If there is any activity on a cluster there is an icon next to the cluster to indicate some activity. In the above example the icon to expand the Thor cluster indicates there is some activity on the Thor cluster. Click the icon to expand the cluster to see the activity on that cluster.

Cluster Action Buttons

Figure 12. Cluster Action Buttons



Check the box next to a cluster to enable the Cluster Action buttons. The Cluster Action buttons will then allow you to perform the following actions on the selected cluster's job queue.

Pause

Pauses the cluster's job queue. The currently running job will complete, but no other jobs will execute until the queue is resumed.

Resume

Resumes a paused job queue. Any waiting jobs will resume execution in order.

Clear

Removes all workunits from the job queue. The removed workunits' state is then set to aborted. Any workunits that were waiting in the queue can be resubmitted manually later, if desired.

Cluster Workunit Activity

Figure 13. Workunit Activity Buttons



Workunit Action Buttons

Check the box next to a Workunit to select it and enable the Workunit Action buttons.

The enabled Workunit Action buttons will then allow you to perform the following actions on the selected Workunit(s).

Open

Opens a (workunit) tab for the selected workunit.

Pause

Press the **Pause** button to complete the current subgraph and then put the job into a paused state.

Pause Now

Press the **Pause Now** button to interrupt the current subgraph (abort it) and put the job into a paused state.

Resume

Resumes processing of any paused job.

Abort

Aborts a running job. An aborted job cannot be resumed.

High, Normal, Low

Figure 14. Priority




You can raise or lower the priority of the job in the queue. Select the job, then press one of the appropriate buttons, **High** or **Low**. Typically all jobs are Normal priority by default.

Press the **High** button to raise the priority of the select job to High. Press the **Low** button to lower the priority of the processing job.

Top, Bottom, Up, Down

Figure 15. Queue Position



You can change the position of a job in the queue using the Top, Bottom, Up, and Down buttons.

Select the workunit to move, that will enable the action buttons.

Press the **Top** button to move the select job to the top of the processing queue. Press the **Up** button to move the job up one position in the queue. Press the **Down** button to move the job down one position in the queue. Press the **Bottom** button to move the job down to the bottom of the queue.

Cluster Information

You can access more information about your Thor clusters from the main Activity tab.

Select the target cluster from the main **Activity** tab, by checking the box next to it.

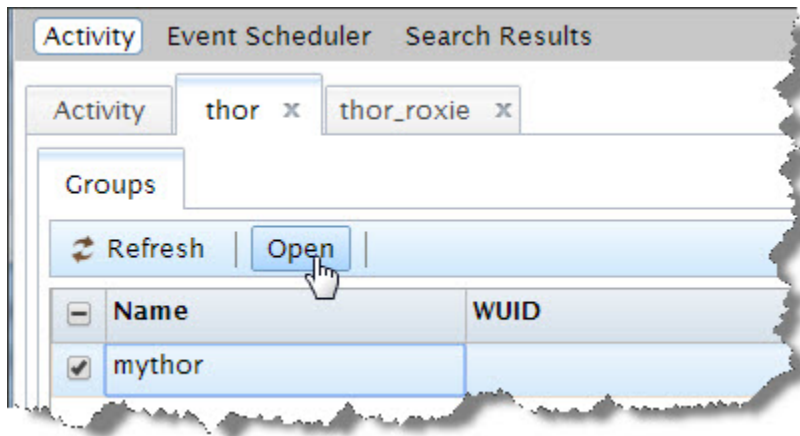
Figure 16. Open Cluster



This enables the **Open** action button. Press the Open action button to open a new tab for that cluster.


The cluster tab displays the groups on that cluster. Check the box next to the cluster group, and press the open button.

Figure 17. Open Cluster Groups



Open up the cluster group tab (for example, mythor), and select the group from the Groups tab to see the activity information on that cluster group.

Figure 18. Cluster Activity tabs



On that cluster group tab, you can access the information about that cluster. There are three tabs on that cluster group tab. The **Summary** tab provides a snapshot of that group.

Alternatively, you can click the link on the cluster name to examine.

The Cluster Usage Tab

The **Usage** tab provides access to a usage graph. The usage tab provides information about the cluster usage.

Figure 19. Usage Graph



To display the usage graph, you can enter some values in the in the fields displayed on the initial usage tab. Optionally, you can just accept the default, the values for the last 30 days. Then press the **Get Usage Graph** button to display the graph.

The Graph shows the cluster usage over time. More information about the cluster usage is specified using a color code. The column on the right breaks down the overall usage as well as the percentage during standard business hours and non-business hours. You can change these values using the fields above the chart, then press the **Get Usage Graph** button again.

Cluster Log File

The **Log File** tab is where you can view that cluster group's log.

Figure 20. Cluster Log File



There are several log display options that you can use to filter the log file. You can filter by rows, pages, or by time. Use the filter options on the Log File tab to filter the displayed log. You can also download the log file to view offline.

Auto Refresh

Figure 21. Auto Refresh



The **Activity** page displays active ECL or DFU workunits either running or in the queue on your cluster. To refresh the list, press the **Refresh** button. Auto Refresh shows the list in real-time, but this feature is disabled by default.

To enable Auto Refresh, toggle the Auto Refresh button.

In an environment with a large number of active users, Auto Refresh could impact performance of your ESP server.

ECL Event Scheduler

The Event Scheduler page provides an interface to the ECL Scheduler. The ECL Scheduler interface allows you to see a list of scheduled workunits. It can also trigger an event. An Event is a case-insensitive string constant naming the event to trap.

ECL Scheduling provides a means of automating processes within ECL code or to chain processes together to work in sequence. For example, you can write ECL code that watches a landing zone for the arrival of a file, and when it arrives, sprays it to Thor, processes it, builds an index, and then adds it to a superfile.

ECL Scheduling is event-based. The ECL Scheduler monitors a Schedule list containing registered Workunits and Events and executes any Workunits associated with an Event when that Event is triggered.

Your ECL Code can execute when an Event is triggered, or can trigger an Event. If you submit code containing a **WHEN** clause, the Event and Workunit registers in the Schedule list. When that Event triggers, the Workunit compiles and executes. When the Workunit completes, ECL Scheduler removes it from the Schedule list.

For example, if you submit a Workunit using **WHEN('Event1','MyEvent', COUNT(2))** in the appropriate place, it will execute twice (the value of **COUNT**) before the ECL Scheduler removes it from the Schedule list and the Workunit is marked as completed.

For more details about both **WHEN** or **NOTIFY** or any ECL Language functions or keywords please see the ECL Language reference. A copy of which can be found online at <http://hpccsystems.com/download/docs/learning-ecl> on the HPCC Systems® web site.

Interface in ECL Watch

To access the ECL Scheduler interface in ECL Watch, click on the **Event Scheduler** link in the navigation sub-menu. The Scheduler interface displays and you can see the scheduled workunits, if any.

The list of scheduled workunits has two significant columns, the **EventName** and the **EventText**.

Figure 22. ECL Scheduler Interface



The EventName is created when scheduling a workunit. The EventText is an accompanying sub event.

You can trigger an event by entering the EventName and Event Text in the entry boxes and then pressing the **PushEvent** button. This is the same as triggering an event using NOTIFY.

Scheduler Workunit List

You can search scheduled workunits by cluster or event name. To filter by cluster or event name, click on the **Filter** Action button. The Filter sub-menu displays. Fill in values for the filter criteria, Eventname or Cluster, then press the **Apply** button. When you specify any Filter options, the Filter Action button displays *Filter Set*.

Figure 23. Workunits in the Scheduler Interface



You can sort the workunits by clicking on the column header.

To view the workunit details, click on the workunit ID (WUID) link for the workunit.

You can modify scheduled workunits from the workunit details page in ECL Watch. Select the workunit details page, then press the **Reschedule** button to reschedule a descheduled workunit. Press the **Deschedule** button to stop a selected scheduled workunit from running. You can also access the Reschedule and Deschedule options from the context menu when you right click on a workunit.

If you are using a WHEN clause and it contains a COUNT number, when rescheduled the workunit will continue the COUNT from the point where it stopped and resumes the remaining COUNT. Once a workunit completes the COUNT, there is no reschedule option.

Pushing Events

The Event Scheduler allows you to trigger or "push" an event to help manage and test your scheduled jobs.

1. Press the **PushEvent** action button.

The Push Event dialog opens.

2. Enter the EventName:

The EventName is a case-insensitive string constant naming the event to trap.

See Also: EVENT

3. Enter the EventText:

The EventText is case-insensitive string constant naming the specific type of event to trap. It may contain * and ? to wildcard-match.

See Also: EVENT


4. Press the **Apply** button

This is the equivalent of

```
NOTIFY( EVENT( EventName, EventText ) );
```

See Also: NOTIFY, EVENT

Figure 24. PushEvent



ECL Workunits

ECL Watch provides information about ECL jobs and workunits. Links to pages that contain information about ECL workunits appear in the navigation bar along top of the ECL Watch page. There you will find links to ECL Workunits, simply labelled as Workunits, and you will also find the link to the ECL Playground. Not only can you find information about workunits there, but you can perform operations on workunits.

ECL Workunits Page

The ECL Workunits page contains a list of all the ECL workunits on your system. It provides access to more details about the workunits. You can also perform actions on the selected workunit using the Workunit Action buttons.

To access the workunits page click on the **ECL** icon, then click the **Workunits** link from the navigation sub-menu.

Figure 25. ECL Files


The page displays the ECL workunits on your system. Choose the **Workunits** Navigation tab to display the workunits.

Figure 26. ECL Watch Browse Workunits



To further examine a workunit or to perform some action on it, you must select it. You can select the workunit by checking the check box. You can also open a work unit by double-clicking on the workunit.

Figure 27. Select ECL Workunit



When you select a workunit, the Action buttons are enabled. You can also use the Context Menu when you right-click on a workunit to perform an action on it.

You can select multiple workunits by checking the check box next to each workunit. You can also click-and-drag over the workunit check boxes to select multiple workunits. When you select multiple workunits, each workunit will open its own tab.

The Action buttons allow you to perform actions on the selected workunits.

- Press the **Open** button to open the selected workunit(s).
- Press the **Delete** button to delete selected workunit(s).
- Press the **Set to Failed** button to set the state of the selected workunit(s) to failed.
- Press the **Abort** button to stop a selected workunit that is running and abort the job.
- Press the **Protect** button to lock the selected workunit(s). This prevents it from archiving by the Sasha server.
- Press the **Unprotect** button to unlock the selected locked workunit(s).
- Press the **Reschedule** button to reschedule selected workunit(s) which have been descheduled.
- Press the **Deschedule** button to stop a selected scheduled workunit from running.
- Press the **Filter** button to display additional filter options. These options filter the displayed list of workunits.

Filter Options

You can filter the workunits displayed on the Workunits tab by clicking on the **Filter** Action button. The Filter sub-menu displays. Fill in values to specify the filter criteria, then press the **Apply** button.

Figure 28. The Filter sub-menu

The screenshot shows the 'Filter' sub-menu in ECL Watch. The menu is titled 'Filter' and contains various input fields for filtering workunits. The fields are: Archived Only (checkbox), WUID (text field with 'W20130222-171723'), Owner (text field with 'jsmi*'), Job Name (text field with 'log_analysis_1*'), Cluster (dropdown), State (dropdown), ECL (text field with ':=dataset*'), Logical File (text field with '*::somefile*'), Logical File Type (dropdown), From Date (text field with '7/28/2013'), To Date (text field with '7/28/2013'), Last N Days (text field with '2'), and time range selectors for 7:30 AM and 7:30 PM. There are 'Clear' and 'Apply' buttons at the bottom.

The workunit filter options allow you to filter workunits using the specified criteria. Workunits can be filtered by:

- **Archived Only** - when checked, this filter will search only archived workunits.
- **WUID** - filter workunits for specific workunit ID (wuid).
- **Owner** - filter workunits for specific owners. Supports wildcards.
- **Job Name** - filter workunits by job name. Supports wildcards.
- **Cluster** - filter workunits by cluster. Select the cluster from the drop list.
- **State** - filter workunits by State. Select the state from the drop list.
- **ECL** - filter workunits by specific ECL. For example, `:=dataset`. Supports wildcards.
- **Logical File** - filter workunits by Logical File name, or some portion of it. Supports wildcards.


- **Logical File Type** - filter workunits by Logical File type. Select the Logical file type from the drop list.
- **From date** - filter workunits from a specific date and/or time. Select the date and time from the drop list.
- **To date** - filter workunits up to a specific date and/or time. Select the date and time from the drop list.
- **Last N Days** - filter workunits for a certain number (N) of days. Enter a number in this field.

Some filter fields support wild card filtering. Wildcards can substitute for one or more characters when filtering data in the filter. The wild card characters are *, and ?. Where * can substituted for all possible characters, and a ? can be substituted for any single character.

Note: Filter criteria are not case sensitive.

When you specify any Filter options, the Filter Action button displays **Filter Set**.


Figure 29. Filter Set



Sorting Columns

You can sort a column by clicking on the column heading. Click once for ascending, click again to toggle to descending. The direction of the arrow indicates the sort order.

Figure 30. Sort by column




The screenshot shows the ECL Workunits interface. At the top, there is a 'Workunits' tab and a toolbar with buttons: Refresh, Open, Delete, Set To Failed, Abort, Protect, Unprotect, and Res. Below the toolbar is a table with columns: Wuid, Owner, and Job Name. The 'Owner' column is highlighted with a red circle and a mouse cursor, indicating it is the selected column for sorting. The table contains 10 rows of workunit data.

Wuid	Owner	Job Name
W20130425-145253	earl	CleanActor
W20130425-144924	earl	File_OriginalPerson
W20130424-122215	earl	LayoutActors
W20130423-093425	earl	CleanActor
W20130422-162036	earl	CleanActor
W20130422-162024	franklin	Layout_People
W20130422-162014	franklin	CleanActor
W20130422-122527	sort	ActorsInMovies

Context Menu

You can right-click on a workunit to get a context menu of actions, including filter options. These are the same set of actions that you could perform from the Action buttons.

Figure 31. Context menu



Workunit Details

The Workunit Details page provides more information about the workunit. You can see more specific information about the selected workunit by selecting the various Workunit Details tabs.

Figure 32. Workunit Details



The Workunit Details Action Buttons allow you to perform actions on the selected workunit. Press the corresponding Action button to allow you to perform the following actions.


- **Refresh** - Redisplays the workunit details.
- **Save** - Saves any changes to the workunit.
- **Delete** - Deletes the workunit.
- **Restore** - Restores an archived workunit. (Workunits are archived by Sasha).
- **Reschedule** - Reschedules a workunit which has been descheduled.
- **Deschedule** - Stops the scheduled workunit from running.
- **Set To Failed** - Changes the workunit state to failed.
- **Abort** - Stops a running workunit and aborts the job.
- **Recover** - Resubmits the workunit without resetting the workflow. This resumes processing from the same point in the process where it ended previously.
- **Resubmit** - Resubmits the workunit. This resets the workflow and starts it over from the beginning.
- **Clone** - Creates a new copy of the workunit. The new workunit is now owned by the user who cloned it.
- **Publish** - Publishes the workunit as a published query.
- **Z.A.P.** - Packages up workunit and system information into a Zip file that can be shared. This is useful for troubleshooting and bug reporting.

- **Slave Logs** - Download the logs for the specified Thor cluster. This is useful for troubleshooting any Thor issues.

Publish Action Button

Click on the Publish action button to publish a query.

Figure 33. Publish Menu



The screenshot shows the ECL Workunits interface with a menu bar at the top containing 'Hule', 'Set To Failed', 'Abort', 'Recover', 'Resubmit', 'Clone', 'Publish', and 'Z.A.'. The 'Publish' button is circled in red. Below the menu bar is a form for publishing a job. The form contains the following fields and options:

- Job Name:** A text box containing 'FindNearestWNAMES'.
- Remote Dali:** An empty text box.
- Source Process:** An empty text box.
- Comment:** An empty text box.
- Priority:** A dropdown menu with 'None' selected.
- Allow Foreign Files:** A checkbox that is checked.
- Update Super Files:** A checkbox that is checked.
- Submit:** A button at the bottom right of the form.

Fill in the values for the Job Name, the Remote Dali, Source Process, Priority (optional), and you can add a comment. After you fill in the required values you can press the **Submit** button to submit your query.


Options for ECL Publish

- **Job Name:** The name of the job to publish. Auto-populated if you choose to publish from the workunit details page.
- **Remote Dali:** The IP or hostname of the DALI to be used to resolve remote files. (optional)
- **Source Process:** Process cluster from which to copy files. (optional)
- **Comment:** If desired, add a comment. The comment displays on the published queries details page.
- **Priority:** Sets the priority for the query. Values can be LOW, HIGH, SLA, or NONE. NONE will clear current setting.
- **Allow Foreign Files:** Check the box to allow the use of foreign files in a Roxie query. If a Roxie query references foreign files and this is not enabled, publish will fail.
- **Update Super Files:** Use when a query uses foreign superfiles or a remote Dali. When such a query is published the superfiles are copied from the remote Dali. If superfiles already exist locally, then the current definition is overwritten only when this box is checked. If it is not checked, the current definition will not change.

Z.A.P. Utility

The Zipped Analysis Package (Z.A.P.) button is a utility for collecting system information and encapsulating it into a shareable package. It is a useful tool for reporting errors, inconsistencies, or other unexpected behavior. When there is such an occurrence, this utility packages up information to send for further analysis.

Figure 34. Z.A.P. Button



To use the Z.A.P. utility, press the Z.A.P. button on the workunit details page from the appropriate workunit. The button opens the Zipped Analysis Package dialog.

Figure 35. The Zipped Analysis Package Dialog

The screenshot shows a dialog box titled "Zipped Analysis Package" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- File Name:** An empty text input field.
- WUID:** A text input field containing "W20190325-153705".
- ESP Build Version:** A text input field containing "internal_7.2.1-closedown0[re".
- ESP Network Address:** A text input field containing "192.168.114.4".
- Thor Network Address:** An empty text input field.
- Description:** An empty text input field.
- History:** An empty text input field.
- Timings:** An empty text input field.
- Password to open ZAP (optional):** An empty text input field.
- Include slave logs:** A checkbox that is currently unchecked.
- Send Email:** A checkbox that is currently unchecked.
- Email Address (To):** A text input field containing "See Configuration Manager.".
- Email Address (From):** A text input field containing "See Configuration Manager.".
- Email Subject:** An empty text input field.
- Email Body:** An empty text input field.

At the bottom of the dialog are two buttons: "Cancel" on the left and "Apply" on the right.

Some of the required fields are populated. Fill in the corresponding values under Description, History, and Timings fields. Optionally, you can password protect the ZAP package and choose to include slave logs. Slave logs are not included by default. If there are slave logs, the option to include them is available. You must check that field when available to include the slave logs.

If your data contains sensitive information, such as personally identifiable information (PII), save the ZAP package, sanitize the data, then email it manually. If appropriate to share your data, you can take advantage of the Send Email field.

The Send Email field is only available if email is configured for the ESP service in the Configuration Manager. If available, check the Send Email box to email the ZAP report. Only an Administrator can configure the email. The (To) Email Address is also set by the Administrator and can only be changed in the configuration. The (From) Email

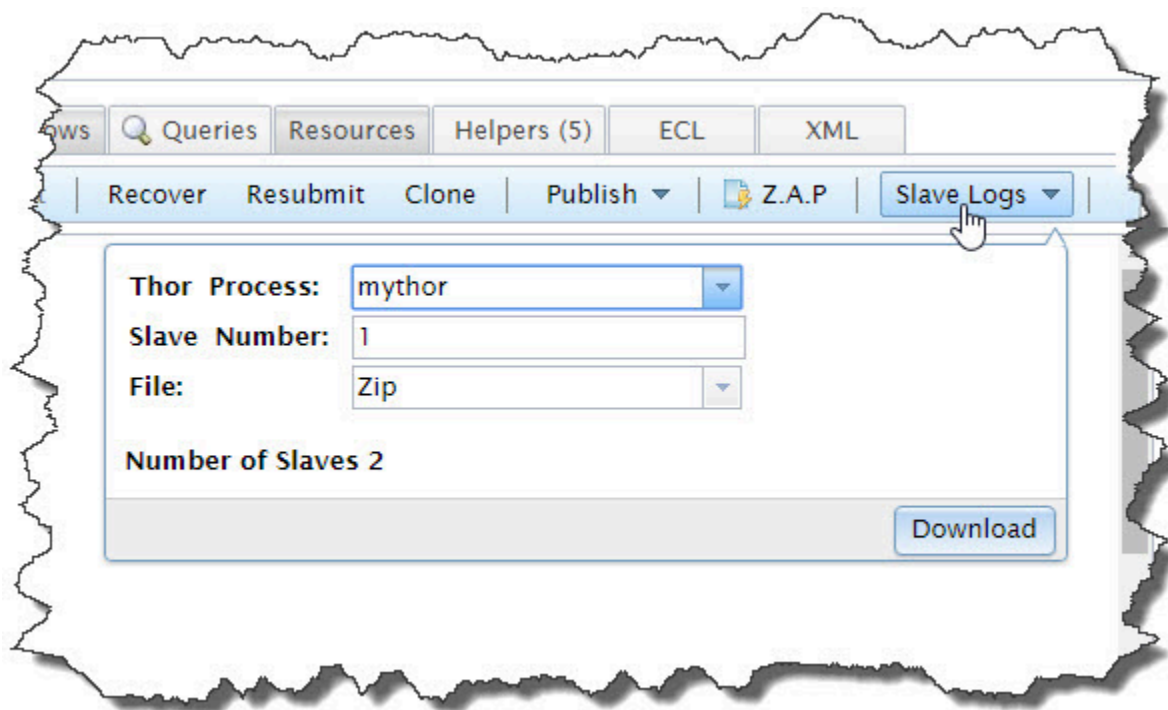
Address can be set in the Configuration Manager, but can be changed if desired. The Email Subject is required, but the Email Body is optional.

Press the **Apply** button when all the dialog fields are completed. At that point if you checked the Send Email box, the Z.A.P. report gets sent. If email is not configured, the Z.A.P. utility generates a zip file with all the appropriate information for troubleshooting.

You can find the generated zip file in your browser's designated download directory. You can now manually send this file to the person handling your support request, or you can upload the file into the issue tracking system. Remember, you should only use the email feature if appropriate to share your data.

Slave Logs

Figure 36. The Slave Logs Dialog



The **Slave Logs** action button opens a dialog where you can choose to download the logs for a specified Thor cluster. You can select the cluster, a specific Slave number, and the log file format (plain or compressed). This is useful for troubleshooting any Thor issues.

Outputs tab


Click on the Outputs tab to see all results.

Figure 37. Outputs



Click on the **Result #** link to open a tab and display the results.

Figure 38. Results Tab




You can filter the result set. Press the **Filter** action button to further filter down the results.

Press the **Download** action buttons to download the output files. The output files are available in 3 formats.

- **GZIP**
- **ZIP**
- **XLS:** Download the output in an Excel spreadsheet format.


Visualizations

You can see visual representations of select workunits. Visualizations are accessible from the workunit details page. On the workunit details page select the **Outputs** tab. Check the box to select the result and press the **Visualize** action button.



The Visualize tab provides a number of chart types you can generate from your result.

Figure 39. Visualization




You can view different visualization types by clicking on the drop list on the Visualize tab. Click on the **Mappings** drop menu, to change the parameters.

Inputs tab

Click on the **Inputs** tab to see the workunit's input files. The input files are the source files that make up the workunit.

Figure 40. Inputs




The input files are displayed as a link. You can double-click the link to open a tab for each input file. You can select more than one input file and press the open button.

The **Usage** header indicates how many times a file was used in the workunit.

Timers tab

Click on the **Timers** tab to see the workunit timings.

Figure 41. Timers



Timers depict everything that happened with the workunit and for how long. Timers include graphs which also provide more information as to the processing of the workunit. The graphical heat map indicates by a darker color where more time was needed, while the lighter the color indicates that portion took less time.

Graphs tab

Click on the **Graphs** tab to see the graphs produced by the workunit.

Figure 42. GRAPHS



The Graphs tab shows a list of each graph and the heat map. Double click on the heat map to go to the corresponding portion of the graph.

Timers tab

Click on the **Timers** tab to see the workunit timings.

Figure 43. Timers



Timers depict everything that happened with the workunit and for how long. Timers include graphs which also provide more information as to the processing. The Heat map indicates by a darker color where more time was needed, while the lighter the color indicates that portion took less time.

Usage: how many times a file was used in the work unit.

Stats

On the workunit Timers tab is another tab for Stats. The Stats tab is a visual representation of all the workunit timings.

Figure 44. Stats



Select the timer values from the drop list on the Stats tab to view the various charts and graphs.

Workflows tab

The workflows tab only exists if you have an attribute scheduled. There are multiple workflows when your code contains more than one WHEN statement.

Counts: How many "events" are scheduled to happen.

Remain: How many "events" remain to occur.

Queries tab

Queries will only appear in the work unit details if the work unit is a published QUERY. Displays what queries were published from that work unit. Same as the queries tab: suspend, unsuspend, activate or deactivate. delete tabs with out any data.

Helpers tab

The Helpers tab displays several helpful elements of a workunit, which might include: the submitted ECL code, the Workunit XML, the archived query, the DLL (SO), the generated C++ file(s), the Thor log, the Thor slave logs, the ECL Agent log, and the Compiler log. Log files only show the portion of the log that pertains to the selected workunit.

ECL Tab

Shows the ECL code for that workunit. It is the same thing as the Helpers ECL link.

XML tab

The workunits XML record as stored in Dali.

Using the ECL Playground

ECL Playground is a tool hosted on an ESP server. A page runs in your browser, allowing you to access and execute self-contained ECL code on your HPCC Systems platform without the use of any other tools. The ECL Playground then shows you the results and the graph in your browser. The view is very similar to what the ECL IDE displays.

Accessing ECL Playground

ECL Playground is installed with the HPCC Systems platform. You can access it through the ECL Watch page.

1. In your browser, go to the **ECL Watch** URL. For example, `http://nnn.nnn.nnn.nnn:8010`, where `nnn.nnn.nnn.nnn` is your ESP server node's IP address.



Your IP address could be different from the ones provided in the example images. Please use the IP address of **your** node.

2. From ECL Watch, click on the **ECL** icon, then click the **Playground** link from the navigation sub-menu.

Figure 45. ECL Playground link




The ECL Playground displays.

Introducing the ECL Playground

The ECL Playground page is a work area where you can see and run self-contained ECL code. You can see the code, submit it, and see the results. You can even change the code and resubmit it to instantly see the new results right in your browser. This is an ideal tool for the user who is not an ECL programming expert who wants to change some of the ECL code and see the results.

Figure 46. The ECL Playground



The ECL Playground page is divided into areas. The top portion contains the *Editor* area and the *Graph Viewer*. The Sample code drop list is at the top right. The bottom portion of the page displays the results.

The ECL Playground comes with a set of ready to run sample ECL code. The drop list contains code samples. Select any one of these samples and it loads in the editor.

Figure 47. Sample drop list



The selected code displays in the *Editor* area. You can then submit it as-is, or modify and submit. The results display at the bottom portion of the page.

Running ECL Code

To run the selected sample code, choose a target cluster from the drop list, then press the **Submit** button.

A successful run displays the word **completed** as the status and the results display in the results viewer. You can also view the graph in the upper right.

Figure 48. Success

The screenshot shows the ECL Watch Playground interface. At the top, there's a blue header with the ECL Watch logo and navigation icons. Below it, a tab bar shows 'Workunits' and 'Playground'. The main area is titled 'ECL Playground' and includes a 'Sample' dropdown set to 'JOIN_dupes'. On the left, a code editor displays ECL code for creating datasets and performing a transform. On the right, a graph view shows a data flow diagram. A red box labeled 'Graph' points to this view. Below the code editor, a 'Submit' button is circled in red. To its right is a 'Target' dropdown set to 'hthor'. Further right, a 'completed' status is circled in red. At the bottom, a table displays the results of the job, showing columns for 'fred' and 'sue'. The table has 5 rows of data. Below the table, it indicates '1 - 50 of 100 results' and includes pagination controls.

```
1 set1 := [1,2,3,4,5,6,7,8,9,10];
2 set2 := [10,20,30,40,50,60,70,80,90,100];
3
4 r1 := {integer1 fred};
5 r2 := {integer1 fred, integer1 sue};
6 ds1 := dataset(set1, r1);
7
8 ds2 := dataset(set2, r1);
9
10 r2 XF(ds1 L, ds2 R) := transform
11   self.fred := L.fred;
12   self.sue := R.fred;
13 end;
```


##	fred	sue
1	1	10
2	1	20
3	1	30
4	1	40
5	1	50

1 - 50 of 100 results

A completed job generates a graph. You can examine the graph in greater detail by double-clicking the graph to zoom in. You can also zoom in with the mouse wheel. A double-click on a blank area of the graph will zoom out. You can use the scroll bars on the border of the graph to navigate or you can drag the graph with your mouse.

Selecting a node in the graph highlights the relevant section of the code in the Editor. This is helpful in troubleshooting or modifying code since it shows you the code that corresponds to a node in the graph.

Figure 49. Error



The status area displays the job status. If a job fails, errors display in the result viewer and the code is highlighted in red in the *Editor*. If there are warnings they are displayed in yellow.

Analyze the results

When running ECL Code that has multiple results, each result is on a separate tab. Select a tab to see that set of results. You can also change number of results displayed or page through the results with the links at the bottom.

Figure 50. Multiple results

The screenshot shows the ECL Playground interface. The top section displays a code editor with the following ECL code:

```
1 SomeFile := DATASET([{'A'},{'B'},{'C'},{'D'},{'E'},  
2 {'F'},{'G'},{'H'},{'I'},{'J'},  
3 {'K'},{'L'},{'M'},{'N'},{'O'},  
4 {'P'},{'Q'},{'R'},{'S'},{'T'},  
5 {'U'},{'V'},{'W'},{'X'},{'Y'}],  
6 {STRING1 Letter});  
7  
8 Set1 := ENTH(SomeFile,2,10,1);  
9 Set2 := ENTH(SomeFile,2,10,2);  
10 Set3 := ENTH(SomeFile,2,10,3);  
11 Set4 := ENTH(SomeFile,2,10,4);  
12 Set5 := ENTH(SomeFile,2,10,5);  
13
```

Below the code editor, there is a "Submit" button and a "Target" dropdown set to "hthor". The status "completed" is shown on the right. The results section shows a table with columns "##" and "letter". The table contains 5 rows of data. Below the table, there is a pagination bar showing "1 - 5 of 5 results". To the right of the pagination bar, there are navigation controls including "Result 2", "Result 3", "Result 4", and "Result 5". A red box labeled "Results Navigation" points to these controls. Another red box highlights the pagination controls, which include "1", "2", "3", "4", "5", and a "50" dropdown.

ECL from a Workunit

You can access ECL code from inside a Workunit Details page in ECL Watch.




1. Select **Workunits** from the ECL Watch  menu.

Figure 51. Browse Workunits



2. Click on a workunit hyperlink to open the Workunit Details page.
3. Click on the **ECL** tab to view the workunit's ECL code.

Figure 52. ECL link



Files

This chapter contains sections dealing with HPCC Systems platform Files, found on the **Files** link in ECL Watch.

In an HPCC Systems platform, data files are partitioned across nodes. The file parts, referenced using Logical File-names, are stored in the Distributed File Utility. This allows the collection of file parts to be referenced as a single entity.


Files

The **Files** page contains features relating to the process of getting data files on to your cluster, managing these files and also the workunits associated with them. Click on the **Files** icon for access to the Files features. You can also perform actions on selected files and superfiles using the Workunit Action buttons.

Logical Files Page

To access the Files page click on the **Files** icon, then click the **Logical Files** link from the navigation sub-menu.

Figure 53. Logical Files




You can browse or search for logical files from this page using the Filter drop menu.

Note: Filter criteria are not case sensitive.

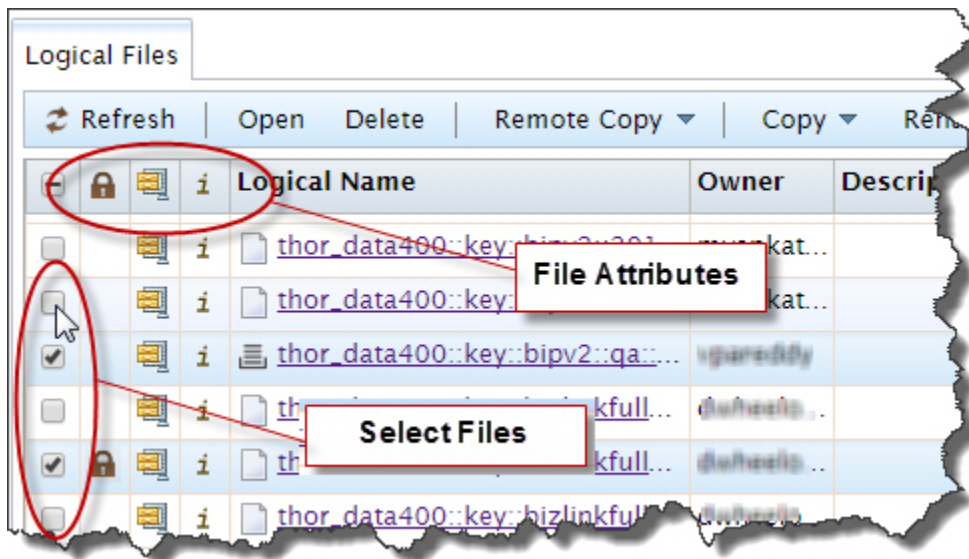
If there are more than 100,000 logical files present you may see a warning message.

Figure 54. Logical Files Page



To see details for a particular file, or to perform some action on it you must select it. You can select a file or files by checking the check box. You can also click and drag your mouse over a group of check boxes, to select multiple files.

Figure 55. Select Files



There are three columns with icons to indicate some file attributes. There is also an icon that appears next to the Logical Name that indicates what type of file it is.

	Locked File
	Compressed File
	Key File (index)
	Logical file
	Superfile

You can also sort a column by clicking on the column heading. Click once for ascending, click again to toggle to descending. The arrow shows the sort order.

Once you have selected a file or files, the Action Buttons are enabled. You can perform actions on selected files.

- Press the **Open** button to open the Logical Files Details page(s).
- Press the **Delete** button to delete file(s).
- Press the **Remote Copy** button to open the dialog where you can copy files from a foreign HPCC Systems platform. You will need permission to access the foreign Dali server.
- Press the **Copy** button to copy a file. You can modify some of the copy options from the drop menu.
- Press the **Rename** button to rename a logical file. You can modify some rename options from the drop menu.
- Press the **Add To Superfile** button to create and add file(s) to a superfile.
- Press the **Despray** button to despray the file. You can modify some despray options from the drop menu.
- Press the **Filter** button to display additional filter options. Use these options to filter the list.

- Press the Tree image button (to the right of the Filter button) to view files by scope in a tree view.

You can press the **Open** button to open a tab with the details for each selected file.

Remote Copy

Press the **Remote Copy** button to open the dialog where you can copy files to or from a foreign dali.

Figure 56. Remote Copy Dialog

The Remote Copy Dialog box is shown with the following fields and options:

- Source:**
 - Dali: 10.239.219.2
 - User ID: EmilyKate
 - Password:
 - Logical Name: tutorial::ek::originalperson
- Target:**
 - Group: mythor
 - Logical Name: tutorial::ek::originalperson
- Options:**
 - Overwrite: ☐
 - Compress: ☐
 - Replicate: ☒
 - No Split: ☐
 - Wrap: ☐
 - Retain Superfile Structure: ☐

A **Submit** button is located at the bottom right of the dialog.

Fill in the values for the Source file, and the Target destination, check any appropriate options then press the **Submit** button.

Copy File

Press the **Copy** button to display the copy drop menu, with additional file copy options.


- Check the **Overwrite** box to overwrite files of the same name.
- Check the **Compress** box to compress the file copy.
- Check the **Retain Superfile Structure** box for the copy to retain the structure of a Superfile. If you are copying a superfile that contains INDEXes, you must enable this option.

- Check the **Replicate** box to create backup copies of all file parts.
- Check the **No Split** box to prevent splitting up the file copy into parts.
- Check the **Wrap** box to keep the number of parts the same and wrap if the target cluster is smaller than the original.
- Check the **Preserve Compression** box to retain the file compression.
- Check the **Expire in (days)** box to enter a number of days before automatically removing the file. If omitted, the default is -1 (never expires).

Sorting Columns

You can sort a column by clicking on the column heading. Click once for ascending, click again to toggle to descending. The direction of the arrow indicates the sort order.

Figure 57. Logical File sort by column



Logical Files Filter Options

You can filter the logical files displayed on the Logical Files tab by clicking on the **Filter** Action button. The Filter sub-menu displays. Fill in values to specify the filter criteria, then press the **Apply** button.

Figure 58. Logical Files Filter sub-menu



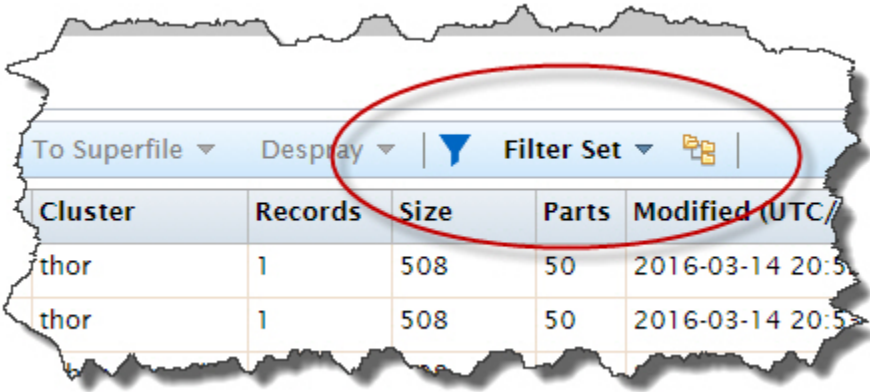
The logical file filter options allow you to filter files using the specified criteria. Logical files can be filtered by:

- **Name** - filter files by name. Supports wildcards.
- **Description** - filter files by description. Supports wildcards.
- **Owner** - filter files by owners. Supports wildcards.
- **Index** - include only Index files if checked.
- **Cluster** - filter files by cluster. Select the cluster from the drop list.
- **From Sizes** - filter files from a specific size.
- **To Sizes** - filter files up to a specific size.
- **File Type** - filter files by type.
- **From date** - filter files from a specific date and/or time. Select the date and time from the drop list.
- **To date** - filter files up to a specific date and/or time. Select the date and time from the drop list.

Note: Filter criteria are not case sensitive.

When you specify any Filter options, the Filter Action button displays *Filter Set*.

Figure 59. Logical Files Filter Set



The screenshot shows a web-based interface for ECL Watch. At the top, there are three dropdown menus: 'To Superfile', 'Display', and 'Filter Set'. The 'Filter Set' dropdown is highlighted with a red circle. Below these menus is a table with five columns: 'Cluster', 'Records', 'Size', 'Parts', and 'Modified (UTC/...)'. The table contains two rows of data, both with 'thor' in the 'Cluster' column, '1' in 'Records', and '508' in 'Size'. The 'Parts' column shows '50' and the 'Modified' column shows '2016-03-14 20:5...'. The table is styled with a light blue header and a white body with orange borders.

Cluster	Records	Size	Parts	Modified (UTC/...
thor	1	508	50	2016-03-14 20:5...
thor	1	508	50	2016-03-14 20:5...


Logical Files Details

The Logical Files Detail page shows specific details for the file selected, including: Workunit ID, Owner, Super Owner, Cluster Name, Description, Job Name, Protected state, Content Type, Key Type, File Size, Format, Compressed state and information, Modified Date, Expire in (days) value, Directory and Path Mask, Record Size, Record Count, Replication state, File Parts, and Skew information.

Check the Protected box to protect the file from deletion or expiration.

The Key Type value only displays if the file is an INDEX. Key Type can have three possible values: Distributed, Local, or Partitioned.

Figure 60. Logical Files Detail Page



The Logical File details summary appears in the main File Details portion of the files page. You can view other file details using the **File Details Tabs** at the top portion of the Page.

- Select the **Summary** tab to view a summary of the file details.
- Select the **Contents** tab to view file contents.
- Select the **ECL** tab to view the ECL code.
- Select the **DEF** tab to view the ECL definitions.
- Select the **XML** tab to view the XML representation of the logical file.
- Select the **Superfiles** tab (when enabled) to display the superfile information.

- Select the **File Parts** tab to view information about the various file parts.
- Select the **Queries** tab to see which queries use which logical files.
- Select the **Graphs** tab (when enabled) to display any graphs associated with the file.
- Select the **Workunit** tab to view the corresponding workunit details. Note that the workunit tab shows you the same information that you would see if you selected it through the workunit link.
- Select the **History** tab to display a list of DFU Actions taken such as copy, remote copy, spray, etc. for the file.

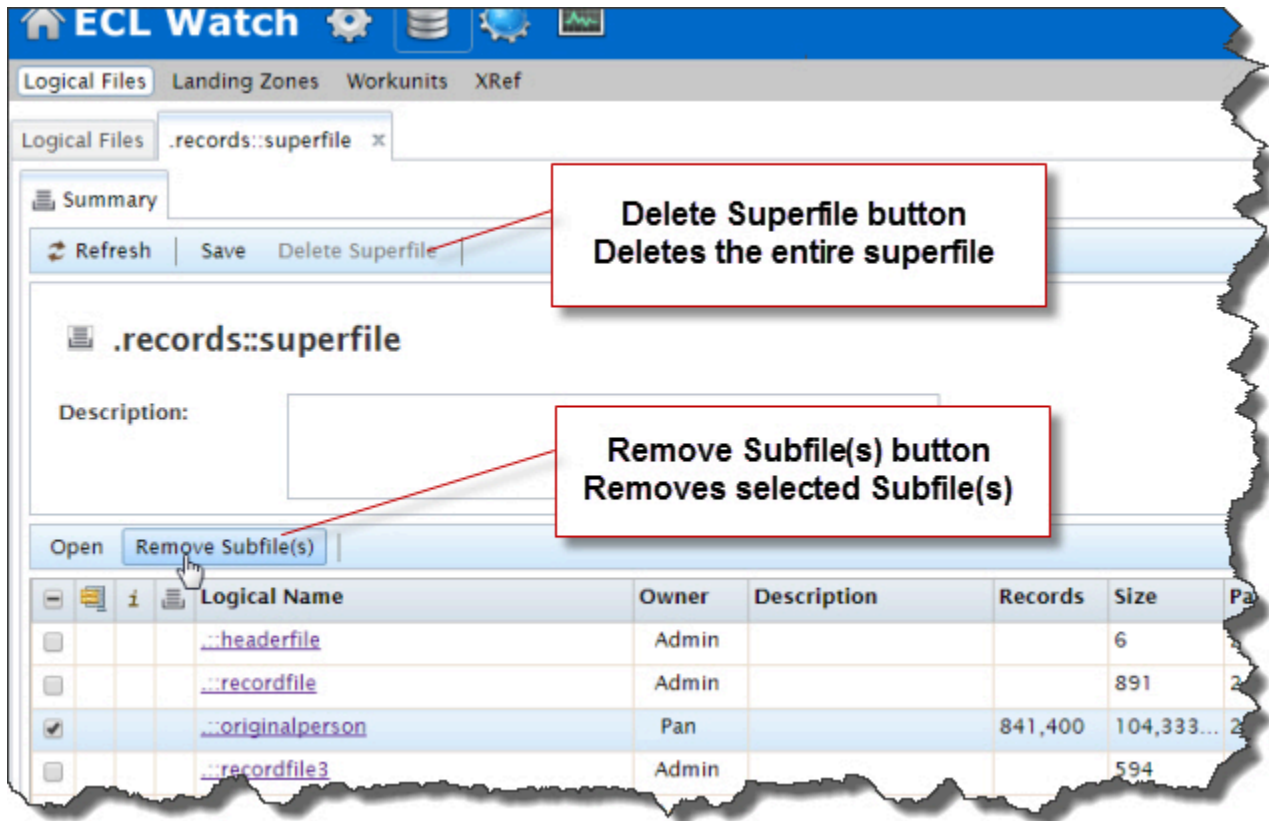
On the File Details Summary tab, you can perform some actions on the selected file.

- Press the **Refresh** button to refresh the file details.
- Press the **Save** button to save any changes you make to the file details.
- Press the **Delete** button to delete the file.
- Press the **Copy** button to copy a file. You can also modify some file attributes from the drop menu.
- Press the **Rename** button to provide a name and rename the file.
- Press the **Despray** button to despray the file. You can also modify some despray options from the drop menu.

Superfiles

A superfile is a managed list of subfiles (Logical Files) treated as a single logical entity. When a file is a superfile, the **Summary** tab displays the superfile details, such as each subfile. Select a superfile from the Logical Files list, then press the Open action button. This displays the superfile details page.

Figure 61. Superfile Details page




On the superfile details page you can:

- See the list of subfiles in the superfile. Click on any subfile hyperlink to see details for that subfile.
- See the details of the subfiles.
- Press the **Save** button to save any changes to the superfile.
- Press the **Delete Superfile** button to delete the entire superfile.
- Press the **Remove Subfile(s)** button to remove any selected subfile from the superfile.

DFU Workunits Page

The DFU Workunits page contains a list of all the DFU workunits on your system. It provides access to more details about the workunits. You can also perform actions on the selected workunit using the Workunit Action buttons.

Figure 62. Files Link



To access the DFU workunits page, click on the **Files** link on the navigation pane, then click the **Workunits** link from the navigation sub-menu. This action opens the DFU Workunits page. The page displays the DFU Workunits on your system.

Figure 63. Browse DFU Workunits



To further examine a workunit or to perform some action on it you must select it. You can select the workunit by checking the check box. You can also double-click on the workunit to select and open it in a new tab.

Figure 64. Select DFU Workunit



You can select multiple workunits by checking the check box next to each workunit. You can also click-and-drag over a group of workunit check boxes to select multiple workunits.

The enabled Action buttons now allow you to perform some actions on the selected workunits.

- Press the **Refresh** button to refresh the list.
- Press the **Open** button to open the workunit(s) details tab.
- Press the **Delete** button to delete selected workunit(s).
- Press **Set to Failed** button to set the workunit(s) state to failed.
- Press the **Protect** button to lock the workunit(s). This prevents it from archiving by the Sasha server.
- Press the **Unprotect** button to unlock the selected protected workunit(s).
- Press the **Filter** button to display additional filter/search options.

Note: Filter criteria are not case sensitive.


When you select a workunit you can then press the **Open** action button to view the workunit details. You can also double-click on a particular workunit to open the details tab.

When you select and then open multiple units, they will each open their own tab.

DFU Workunits Filter Options

You can filter the workunits displayed on the Workunits tab by clicking on the **Filter** Action button. The Filter sub-menu displays. Fill in values to specify the filter criteria, then press the **Apply** button.

Figure 65. The DFU Workunit Filter




The DFU Workunit filter options allow you to filter workunits using the specified criteria. Workunits can be filtered by:

- **Archived Only** - when checked, this filter will search only archived workunits.
- **Owner** - filter workunits for specific owners. Supports wildcards.
- **Job Name** - filter workunits by job name. Supports wildcards.
- **Cluster** - filter workunits by cluster. Select the cluster from the drop list.
- **State** - filter workunits by state. Select the state from the drop list.

Note: Filter criteria are not case sensitive.

When you specify any Filter options, the Filter Action button displays *Filter Set*.


Figure 66. DFU Filter Set



Sorting Columns

You can sort a column by clicking on the column heading. Click once for ascending, click again to toggle to descending. The direction of the arrow indicates the sort order.


Figure 67. DFU Workunit Sort by column



DFU Workunit Details Page

When you open the selected workunit(s) you will see the workunit details. The Workunit Details tab provides information about a workunit. You can see more information about workunit details by selecting the various Workunit Details tabs. You can also perform actions on the selected workunits using the Workunit Action buttons.

Figure 68. Workunit Details



Additional Workunit details are located in the Workunit Details section of the page. Job name, queue, command, time, completion percentage, along with specific process messages display here.


You can use the Workunit Action buttons on the Active Workunit tab to perform actions on the selected workunit. Press the appropriate Workunit Action button to perform the following actions.

- Press the **Refresh** button to refresh the workunit details.
- Press the **Save** button to save the workunit.
- Press the **Delete** button to delete the workunit.
- Press the **Abort** button to abort a running workunit.
- Press the **Resubmit** button to resubmit the workunit (not yet implemented).
- Press the **Modify** button to modify the workunit (not yet implemented).

XML Tab

The XML Tab on the workunit details page allows you to see the XML representation of the workunit.

Figure 69. Workunit Detail XML tab



Source Tab

The Source Tab on the workunit details page allows you to view the source file(s) of the DFU workunit.


Figure 70. Workunit Detail Source tab



Target Tab

The Target Tab on the workunit details page allows you to view the target of the DFU workunit.

Figure 71. Workunit Detail Target tab



Landing Zones

To access the Landing Zones page click on the **Files** icon, then click the **Landing Zones** link from the navigation sub-menu. The Landing Zone link displays the Landing Zones page. The Landing Zone Page shows you each landing zone you have configured for your cluster and its contents.

Figure 72. Landing Zone Page



Click on the arrow next to a drop zone container, server, or folder to expand. The files on the drop zone display. You can choose to upload, download, or delete any files on the drop zone using the landing zone action buttons. You can also spray files to a cluster from this page.

Upload files

You can upload files to your landing zone from the Landing Zone page.

Figure 73. Landing Zone Upload




1. Press the **Upload** action button.



The upload utility in ECL Watch is limited by the browser's file size limitation. This is typically 4 GB. For production systems, we recommend a secure copy protocol (scp) utility.

2. Choose the file from the window that displays.
3. Verify the file, folder, and dropzone information are correct in the **File Uploader** dialog that displays.

Figure 74. Info Dialog



4. Press the **Start** button to begin the upload.

The Upload action button displays the progress as the file uploads.

Download files

You can download files from your landing zone to your computer.

1. From the Landing Zone page, select a file (or files) to download by checking the box next to it.
2. Press the **Download** button to download the file.

The file will download to your browser's download directory as specified in your browser settings.

Delete files

You can delete files from your landing zone.

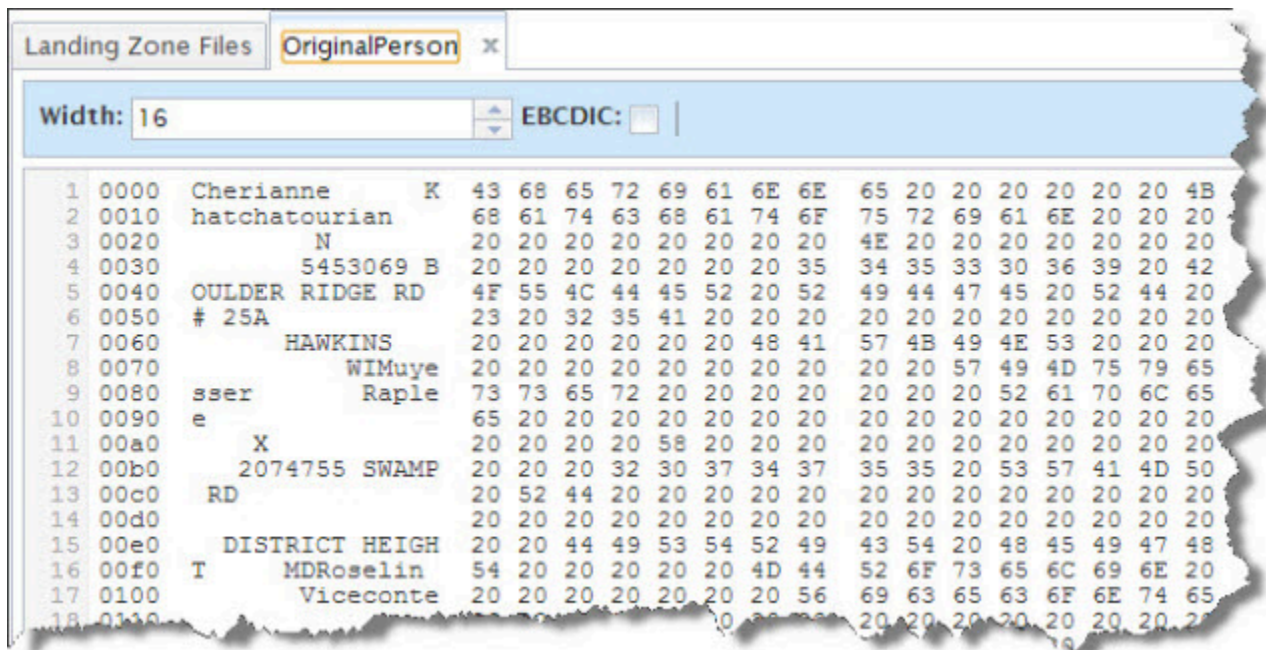
1. From the Landing Zone page, select a file (or files) to delete by checking the box next to it.
2. Press the **Delete** action button to delete the file from your landing zone.

Hex Preview

The Hex Preview shows the contents of a file on the landing zone in hexadecimal form. If the file is large, then only the first 32k display. Hex preview is designed for fixed length files, although it can also work for delimited files but may be limited in that regards.

1. Select a file by checking the box next to it.
2. Press the **Hex Preview** action button to display the selected file(s) in a hex format.

Figure 75. Hex Preview



You can adjust the width of the view on the hex preview page using the spinbox controls on the **Width** box.

If you have an EBCDIC file check the box next to **EBCDIC:** for it to display properly.


Spray/Despray

This section details how to Spray and Despray a data file to your cluster using ECL Watch. The spray function is integrated into the Landing Zone page as detailed in the Upload Files section .

In order to spray a file to your cluster you must first upload the file to your landing zone. The file upload steps are detailed in the preceding section.

With the file successfully uploaded to the landing zone you can choose the file to spray from the Landing Zone page. Once selected the Spray buttons become enabled.

Figure 76. Landing Zone Page



Spray Data to a Cluster

With the **Spray:** Action buttons enabled, you press the appropriate button for the Spray you wish to perform. Fill in the appropriate values when prompted to complete the spray.

Spray Files

One way you can spray files to your clusters is from the **Landing Zone** page in ECL Watch.

1. Select the file from your drop zone by checking the box next to it.
2. Select the appropriate drop menu option for the type of spray you want.

For example, to spray a delimited file, select the **Delimited** action button.

Figure 77. Landing Zone Spray

The screenshot shows the 'Landing Zone Spray' configuration window in ECL Watch. The window has a tabbed interface at the top with options: 'Fixed', 'Delimited', 'XML', 'JSON', 'Variable', and 'BLOB'. The 'Delimited' tab is selected. Below the tabs, there are several input fields and checkboxes. The 'Target' section includes a 'Group' dropdown set to 'mythor', a 'Queue' dropdown set to 'dfuserver_queue', and a 'Target Scope' text field containing 'some:prefix'. Below these is a 'Target Name' list box containing 'actors.list'. The 'Options' section includes a 'Format' dropdown set to 'ASCII', a 'Max Record Length' text field set to '8192', a 'Separators' text field containing '\', an 'Omit Separator' checkbox, an 'Escape' text field, a 'Line Terminators' text field containing '\n,\r\n', a 'Quote' text field containing ' ', an 'Overwrite' checkbox, a 'Replicate' checkbox (checked), a 'No Split' checkbox, a 'Compress' checkbox, a 'Fail if No Source File' checkbox, a 'Record Structure Present' checkbox, a 'Quoted Terminator' checkbox, and an 'Expire in (days)' text field. A 'Spray' button is located at the bottom right of the window.

3. Fill in the values as appropriate for the spray.
4. Press the **Spray** button to spray the file(s).

Spray multiple files

You can choose to spray multiple files with the multi-file spray feature. This is useful for spraying a number of files of the same type using the same spray options.

Fixed (length) files can have different record lengths and XML files can have different row tags which must be specified individually for each file. To specify these differences select the files you want to spray and the spray type. You will then see the files listed. Enter the record length or row tag information for each file if using the Fixed or XML spray types, then check all other applicable options and Spray.

Spray Fixed

- Click on the **Files** icon, then click the **Landing Zones** link from the navigation sub-menu.
- Click on the arrow next to your dropzone to expand the list.

The files on your drop zone display.

- Check the checkboxes for the file(s) you want to spray, then press the Spray: **Fixed** action button.

The **Spray Fixed** dialog displays.

- Fill in relevant details:

Target	
Group	Select the name of cluster to spray to. You can only select a cluster in your environment.
Queue	Select the queue for the spray.
Target Scope	The prefix for the logical file
Target Name	The logical filename to create. This is pre-filled with the name of the source file on the landing zone, but can be changed.
Record Length	The size of each record.
Options:	
Overwrite	Check this box to overwrite files of the same name.
Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).
This option is only available on systems where replication has been enabled.	
Compress	Check this box to compress the files.
Expire in (days)	An integer value indicating the number of days before automatically removing the file. If omitted, the default is -1 (never expires).
No Split	Check this box to prevent splitting file parts to multiple target parts.
Fail if no source file	Check this box to allow the spray to fail if no source file is found.

- Press the **Spray** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each spray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Spray Delimited

- Click on the **Files** icon, then click the **Landing Zones** link from the navigation sub-menu.
- Click on the arrow next to your dropzone to expand the list.

The files on your drop zone display.

- Check the checkboxes for the file(s) you want to spray, then press the Spray: **Delimited** action button.

The **Spray Delimited** page displays.

- Fill in relevant details:

Target	
Group	Select the name of cluster to spray to. You can only select a cluster in your environment.
Queue	Select the queue for the spray.
Target Scope	The prefix for the logical file
Target Name	The logical filename to create. This is pre-filled with the name of the source file on the landing zone, but can be changed.

Options:

Format	Select the format from the droplist
Max Record Length	The length of longest record in the file.
Separators	The character(s) used as a separator in the source file.
Omit Separator	Check this box to omit the separator.
Escape	A null-terminated string containing the CSV escape characters.
Line Terminators	The character(s) used as a line terminators in the source file.
Quote	The character used as a quote in the source file.
Overwrite	Check this box to overwrite files of the same name.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Fail if no source file	Check this box to allow the spray to fail if no source file is found.
Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).

This option is only available on systems where replication has been enabled.

Compress	Check this box to compress the files.
Expire in (days)	An integer value indicating the number of days before automatically removing the file. If omitted, the default is -1 (never expires).
Quoted Terminator	Check this box to indicate that the terminator character can be included in a quoted field. If unchecked, it allows quicker partitioning of the file (avoiding a complete file scan).

Record Structure Present Flag indicating whether to derive the record structure from the header of the file.

- Press the **Spray** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each spray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Spray XML

- Click on the **Files** icon, then click the **Landing Zones** link from the navigation sub-menu.
- Click on the arrow next to your dropzone to expand the list.

The files on your drop zone display.

- Check the checkboxes for the file(s) you want to spray, then press the Spray: **XML** action button.

The **Spray XML** dialog displays.

- Fill in relevant details:

Target	
Group	Select the name of cluster to spray to. You can only select a cluster in your environment.
Queue	Select the queue for the spray.
Target Scope	The prefix for the logical file
Target Name	The logical filename to create. This is pre-filled with the name of the source file on the landing zone, but can be changed.
Row Tag	The tag name of the row delimiter. Required.
Options:	
Format	Select the format from the droplist
Max Record Length	The length of longest record in the file.
Overwrite	Check this box to overwrite files of the same name.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Expire in (days)	An integer value indicating the number of days before automatically removing the file. If omitted, the default is -1 (never expires).
Fail if no source file	Check this box to allow the spray to fail if no source file is found.
Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).
This option is only available on systems where replication has been enabled.	
Compress	Check this box to compress the files.

- Press the **Spray** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each spray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Spray JSON

- Click on the **Files** icon, then click the **Landing Zones** link from the navigation sub-menu.
- Click on the arrow next to your dropzone to expand the list.

The files on your drop zone display.

- Check the checkboxes for the file(s) you want to spray, then press the Spray: **JSON** action button.

The dialog displays.

- Fill in relevant details:

Target	
Group	Select the name of cluster to spray to. You can only select a cluster in your environment.
Queue	Select the queue for the spray.
Target Scope	The prefix for the logical file
Target Name	The logical filename to create. This is pre-filled with the name of the source file on the landing zone, but can be changed.
Row Path	The path specifier to the JSON content. The default takes the root level content as an array of objects to be treated as rows.

Options:	
Format	Select the format from the droplist
Max Record Length	The length of longest record in the file.
Overwrite	Check this box to overwrite files of the same name.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Fail if no source file	Check this box to allow the spray to fail if no source file is found.
Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).
This option is only available on systems where replication has been enabled.	
Expire in (days)	An integer value indicating the number of days before automatically removing the file. If omitted, the default is -1 (never expires).
Compress	Check this box to compress the files.

- Press the **Spray** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each spray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Spray Variable

- Click on the **Files** icon, then click the **Landing Zones** link on the navigation sub-menu.
- Click on the arrow next to your dropzone to expand the list.

The files on your drop zone display.

- Check the checkboxes for the file(s) you want to spray, then press the Spray: **Variable** action button.

The Spray **Variable** dialog displays.

- Fill in relevant details:

Target	
Group	Select the name of cluster to spray to. You can only select a cluster in your environment.
Queue	Select the queue for the spray.
Target Scope	The prefix for the logical file
Target Name	The logical filename to create. This is pre-filled with the name of the source file on the landing zone, but can be changed.

Options:

Source Type	Select the source type from the drop list. Values: recfmv, recfmvb, Variable, or Variable Big-endian.
Overwrite	Check this box to overwrite files of the same name.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Fail if no source file	Check this box to allow the spray to fail if no source file is found.
Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).

This option is only available on systems where replication has been enabled.

Expire in (days)	An integer value indicating the number of days before automatically removing the file. If omitted, the default is -1 (never expires).
Compress	Check this box to compress the files.

- Press the **Spray** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each spray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Spray Blob

- Click on the **Files** icon, then click the **Landing Zones** link on the navigation sub-menu.
- Click on the arrow next to your dropzone to expand the list.

The files on your drop zone display.

- Check the checkboxes for the file(s) you want to spray, then press the Spray: **BLOB** action button.

The Spray **BLOB** dialog displays.

- Fill in relevant details:

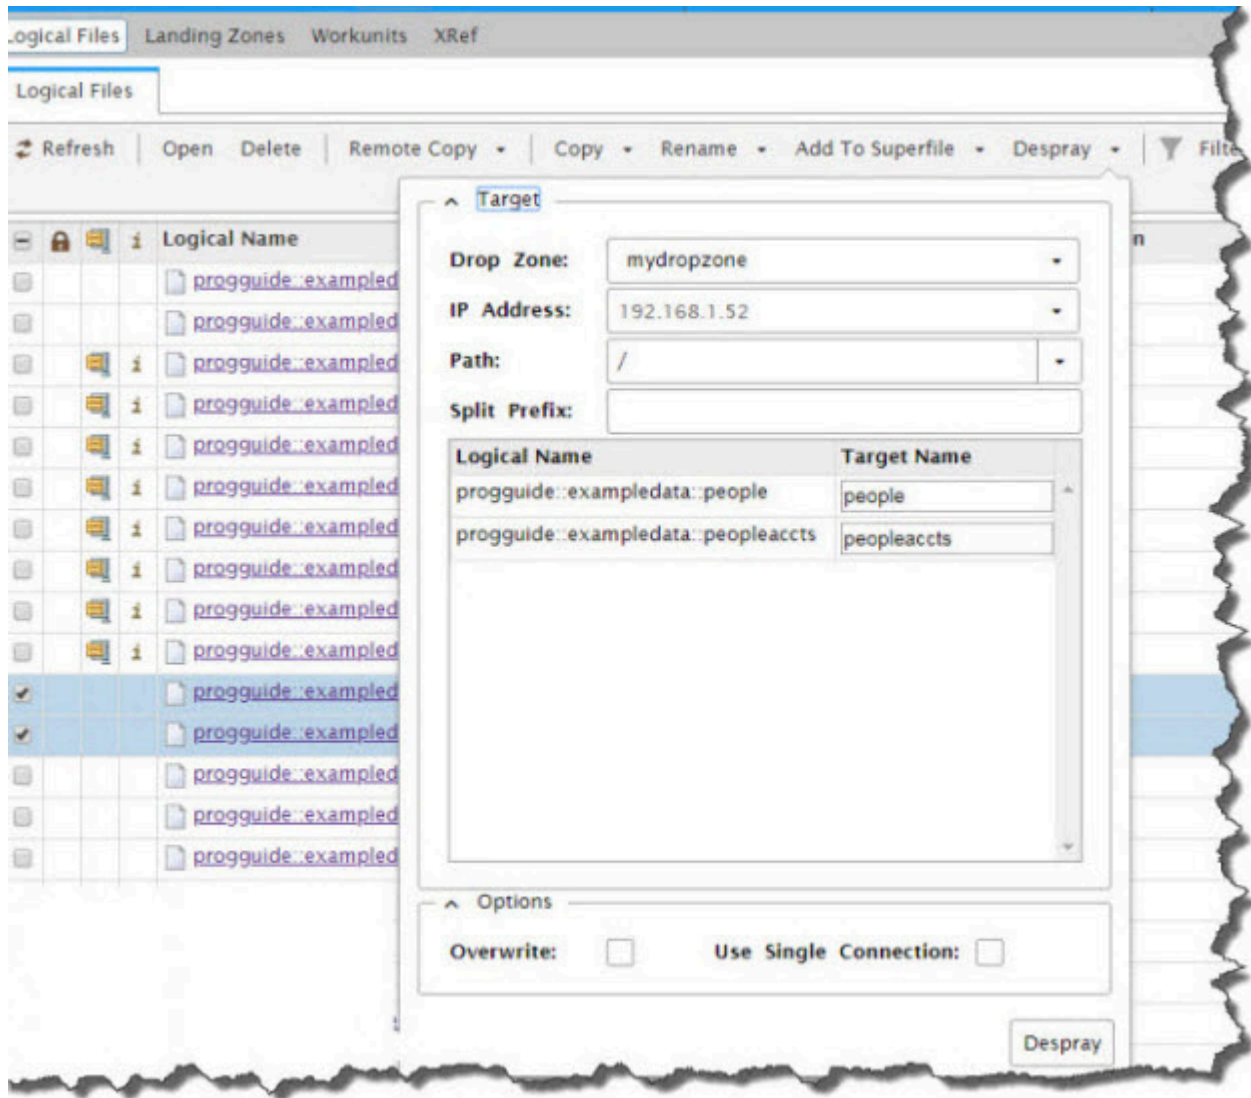
Target	
Group	Select the name of cluster to spray to. You can only select a cluster in your environment.
Queue	Select the queue for the spray.
Target Name	The logical target name to create. Required. You must provide a target name.
Source Path	The path to the file. This is pre-filled with the name of the selected source file(s) on the landing zone, but can be changed. Supports wildcards.
Options:	
Blob Prefix	The prefix for the file.
Overwrite	Check this box to overwrite files of the same name.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).
This option is only available on systems where replication has been enabled.	
Compress	Check this box to compress the files.
Expire in (days)	An integer value indicating the number of days before automatically removing the file. If omitted, the default is -1 (never expires).
Fail if no source file	Check this box to allow the spray to fail if no source file is found.

- Press the **Spray** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each spray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Desprays

- Locate the file(s) to despray in the list of files, then Press the the **Despray** action button.



- Provide **Destination** information.

Drop Zone

Use the drop list to select the machine to despray to. The items in the list are landing zones defined in the system's configuration.

IP Address

This is prefilled based upon the selected machine.

Path

Provide the complete file path of the destination including file name and extension.

Split Prefix

Prefix

Overwrite

Check this box to overwrite a file with the same name if it exists.

Use Single Connection

Check this box to use a single network connection to despray.

- Press the **Despray** button.

A DFU Workunit tab for each job opens. You can see the progress of each despray operation on the tab. If a job fails, information related to the cause of the failure also displays.

Copy

- Click on the **Files** icon, then click the Logical Files button on the navigation bar.
- Select the file(s) to copy in the list of files, then click on the **Copy** action button.
- Fill in **Destination** and **Options** information.

Target:

Group	Select the name of cluster to copy to. You can only select a cluster in your environment.
Target Name	The logical filename to create. This is pre-filled with the name of the source file on the landing zone, but can be changed.

Options:

Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).
	This option is only available on systems where replication has been enabled.
Wrap	Check this box to keep the number of parts the same and wrap if the target cluster is smaller than the original.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Overwrite	Check this box to overwrite files of the same name.
Compress	Check this box to compress the files.
Retain Superfile Structure	Check this box to retain the superfile structure.
Preserve Compression	Check this box to preserve the compression of the original file when copying

- Press the **Copy** button.

A **DFU Workunit** tab displays for each job. You can see the progress of each copy operation on the tab. If a job fails, information related to the cause of the failure also displays.

Remote Copy

Remote Copy allows you to copy data from a cluster outside your environment to one in your environment.

- Click on the **Files** icon, then click the Logical Files button on the navigation bar.

- Click on the **Remote Copy** link

The **Copy File** page displays.

- Fill in **Source**, **Destination**, and **Options** information.

Source:

Dali	The Dali Server in the remote environment
User ID	The Username to use to authenticate on the Remote environment (if needed)
Password	The password to use to authenticate on the Remote environment (if needed)
Logical File	The logical filename in the remote environment.

Destination:

Group	Select the name of cluster to copy to. You can only select a cluster in your environment.
Logical Name	The logical name for the copied file.

Options:

Replicate	Check this box to create backup copies of all file parts in the backup directory (by convention on the secondary drive of the node following in the cluster).
	This option is only available on systems where replication has been enabled.
Wrap	Check this box to keep the number of parts the same and wrap if the target cluster is smaller than the original.
Overwrite	Check this box to overwrite files of the same name.
Compress	Check this box to compress the files.
No Split	Check this box to prevent splitting file parts to multiple target parts.
Retain Superfile Structure	Check this box to retain the superfile structure.

- Press the **Submit** button.

A **DFU Workunit** tab displays. You can see the progress of the copy operation on the tab. If a job fails, information related to the cause of the failure also displays.

- Press the **Refresh** button periodically until the status of your request indicates it is **Finished** or click on the **View Progress** hyperlink to see a progress indicator.

XRef

Under the **Files** Icon on the navigation sub-menu there is a link for XRef. This link will take you to the XRef page. On the XRef page you can run the XREF utility.

Figure 78. XRef page



The XRef Utility

The XREF utility provides the ability to find "orphaned", "lost", and "found" files. These "orphan" are files which are found on the nodes but not registered in the Distributed File System.

Found File A found file has file parts on disk that are not referenced in the Dali server. All the file parts are accounted for so they can be added back to the Dali server. They can also be deleted from the cluster, if required.

Orphan File An orphan file is a file without a Dali entry and without a complete set of physical files. Since some of the physical files are missing it is an incomplete file and it cannot be used to add a logical file entry back into the Dali server. These orphan file parts do not have a reference in the Dali server.


Lost File A logical file that is missing at least one file part on both the primary and replicated locations in storage. The logical file is still referenced in the Dali server. Deleting the file removes the reference from the Dali server and any remaining parts on disk.



On a large system, we suggest limiting the number of users who can Generate XREF reports by setting DfuXrefAccess access to FULL for only those users.

To generate a list:

- Press the **Generate** button.

	Sasha Server typically runs Xref at the times scheduled when deployed.
---	--

To view results:

- Click on one of the tabs (Found Files, Orphan Files, Lost Files, Directories, or Errors/Warnings).

XREF with multiple Thor clusters

XREF runs on the primary Thor cluster. For a configuration with more than one Thor cluster on the same physical node group, ECL Watch only displays the primary Thor and not any other Thor that shares the same node group. This is the Thor cluster whose name matches the *nodeGroup* in the configuration.

To allow XREF to run in an environment with multiple Thor clusters, set the *nodeGroup* to the same value for all Thor clusters.

Working with XREF results:

After XRef completes you can see a list of available reports. Click on one of the tabs to see the results page of each type.

Figure 79. XRef Errors




Errors/Warnings for: mythor

Refresh | Open

File	Message	Status
certification::full_test_distributed	Recent file ignored	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::superfile::people2 exists	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::superfile::people5 exists	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::superfile::people4 exists	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::superfile::people1 exists	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::superfile::people6 exists	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::superfile::people3 exists	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::exampledata::peopleacc...	Warning
/var/lib/HPCCSystems/hpcc-data/thor...	Orphans ignored as progguide::exampledata::xml_ucc e...	Warning

The Orphan and Lost Files pages list any Orphan or Lost files. Orphan and Lost files are difficult to recover, unless you have copies of missing parts needed to reconstruct the file (for example, if a missing part is on a hard drive that was replaced.)


Figure 80. XRef: Delete Orphan Files



Typically, these files will need to be deleted. To Delete, check the boxes next to files you want to delete, then press the **Delete** button.

Found files can and usually are reattached.

Figure 81. XRef: Attach Found Files




Check the boxes next to files you want to reattach, then press the **Attach** button.

Queries

The link for the Published Queries (icon) header provides more information and details about queries on available targets.

Click on the Published Queries (icon) hyperlink to display the published queries and package maps on that cluster.

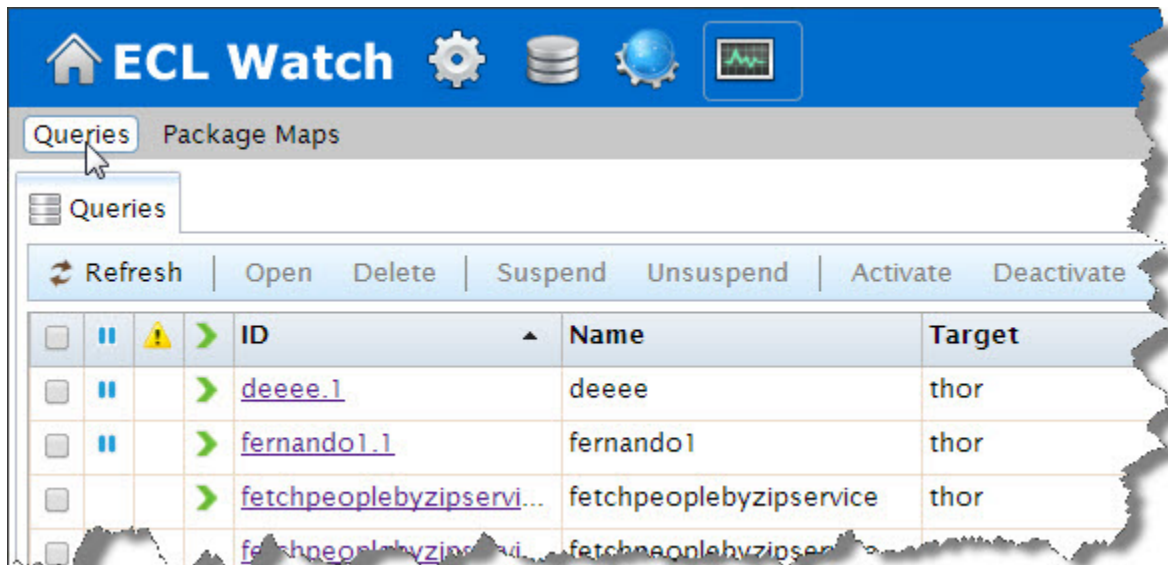
Figure 82. Query Link



Queries Page in ECL Watch




The Queries page lists published queries for each target cluster. On this page you can see the published queries . You can also perform some actions on the selected queries.

Figure 83. Browse Query Sets



The Queries interface provides some information at a glance, there are three columns on the left side of each listed query. These three columns provide information about these queries.

Using ECL Watch Queries

	Indicates a paused query
	Indicates an activated query
	Indicates a query suspended by the system

The queries page also provides other information at a glance:

- the query ID
- the query name
- the target
- the workunit id (WUID)
- the dll
- Published by

All the above available at a glance on the main queries page, with further actions that can be performed from the action buttons along the top of the tab. You can sort a column by clicking on the column heading. Click once for ascending, click again to toggle to descending. The direction of the arrow indicates the sort order.

To see the details page for a particular query, or to perform some action on it you must select it. You can select a query or queries by checking the check box. You can also open a particular query by double clicking on it.

Queries Tab

When you select the Published Queries hyperlink you open the Queries tab. This tab displays published queries on the system. The Action buttons allow you to perform operations on the published queries selected.

Figure 84. Published Query Action buttons



Open	Opens the selected query (or queries).
Delete	Deletes the selected query (or queries).
Suspend	Suspends the selected active query (or queries).
Unsuspend	Unsuspects the selected suspended query (or queries).
Activate	Activates the selected query (or queries). This assigns a query to the active alias with the same name as the query.
Deactivate	Deactivates the selected active query (or queries) by removing the active query alias from the given queryset.

Filter

Allows you to filter the queries for the criteria you enter. When the Filter is applied the action button displays **Filter Set**. This icon indicates that the published queries displayed are filtered.

The screenshot shows a web interface with a 'Filter Set' dialog box. The dialog has a title bar with 'Filter Set' and a dropdown arrow. Below the title bar are several input fields: 'ID:' with the value 'som?q*ry.1', 'Name:' with 'My?Su?erQ*ry', 'Published By:' with 'Published By', 'WUID:' with 'W2016*', 'Cluster:' with a dropdown arrow, 'Logical File:' with 'some::logical::name', 'Libraries Used:' with an empty field, 'Suspended:' with a dropdown arrow set to 'All', and 'Active:' with a dropdown arrow set to 'All'. At the bottom of the dialog are 'Clear' and 'Apply' buttons. A red circle highlights the 'Filter Set' button in the top navigation bar. The background shows a list of queries with columns like 'Target', 'Name', 'WUID', 'Cluster', 'Logical File', 'Libraries Used', 'Suspended', and 'Active'.

You can filter for several query attributes. You can filter by:

- ID
- Name
- Published by
- WUID
- Cluster
- Logical File Name
- Libraries Used
- Suspended queries.
- Active queries.

The Filter also supports wild cards.

Options

Provides the option to search/display queries on a single node or all nodes. Using this option can improve performance if you have a large multi-node cluster.

Recreate Query Recompiles a query into a new workunit and republishes the new workunit. This is useful when upgrading to a new ECL compiler and you want to recompile a query from the exact same source. The ECL archive must be available within the workunit of the query.


Query Details

To examine the Query Details page, you select and open the query or queries. This opens a tab containing the query details. From the query details page you can get more information about the specific query. You can also perform some actions on that query. There are several tabs with additional information about the selected query.

Query Summary Tab

The default query tab opened when you select a query is the Summary tab. The summary tab shows you some detail information about the query.

Figure 85. Query detail page



There are a few actions that you can perform on the query from this tab. Press the action buttons for the desired activity for the selected query.

Refresh Refreshes the information displayed for the selected query.

Save Saves the selected query (or queries).


Delete Deletes the selected query (or queries).

Errors Tab

For each selected query there is an Errors tab. The Errors tab displays any errors that may have been encountered during the compiling and publishing of that query. If there aren't any errors the errors tab will be blank. If there are

errors, you can further examine any specific error by checking the box and selecting it, and then press the open action button. You could also just double click on the selected error.


Figure 86. Query Error



Graphs Tab

The graphs tab provides access to graphical interpretations of the query. This can be helpful in visualizing how the query ran. The graphs tab displays a list of any graphs generated by the selected query, along with some additional information like timing. To display a specific graph, you must select it, and choose to open it, or you can double click on listed graph.

Figure 87. Graphs list



Opening a graph will open a new tab showing the selected graph(s).

Figure 88. Graphs



When you open a graph the visible area splits into three smaller sections each displaying some relevant component of the query graph. Notice the myriad of graph controls, and tabs in the border area of each tab. Manipulate these controls to view different aspects of the graphs.

The Advanced action button on the main graph control area, provides access to even more advanced graphing options.

Logical Files Tab

The Published queries details page provides a link to the queries Logical Files tab. The Logical Files tab shows all logical files that are used by the query. To view the logical file details for any file listed, select one or more files by checking the checkbox and press the Open action button. Tabs for each file selected opens where you can view and make changes to the file(s) without the need to go back to the logical files page.

Figure 89. Queries:Logical Files Tab




The above image shows the list of Logical files on the Logical Files tab. To view more detail about a logical file listed here, check the box next to the file, and then press the **Open** action button. You can also just double click on the logical file you want to view.

Using ECL Watch Queries

Once open, you can select any of the tabs to see Summary, Contents, ECL, DEF, XML, File Parts, Queries, or the Workunit.

Figure 90. Queries:Logical Files:Contents Tab

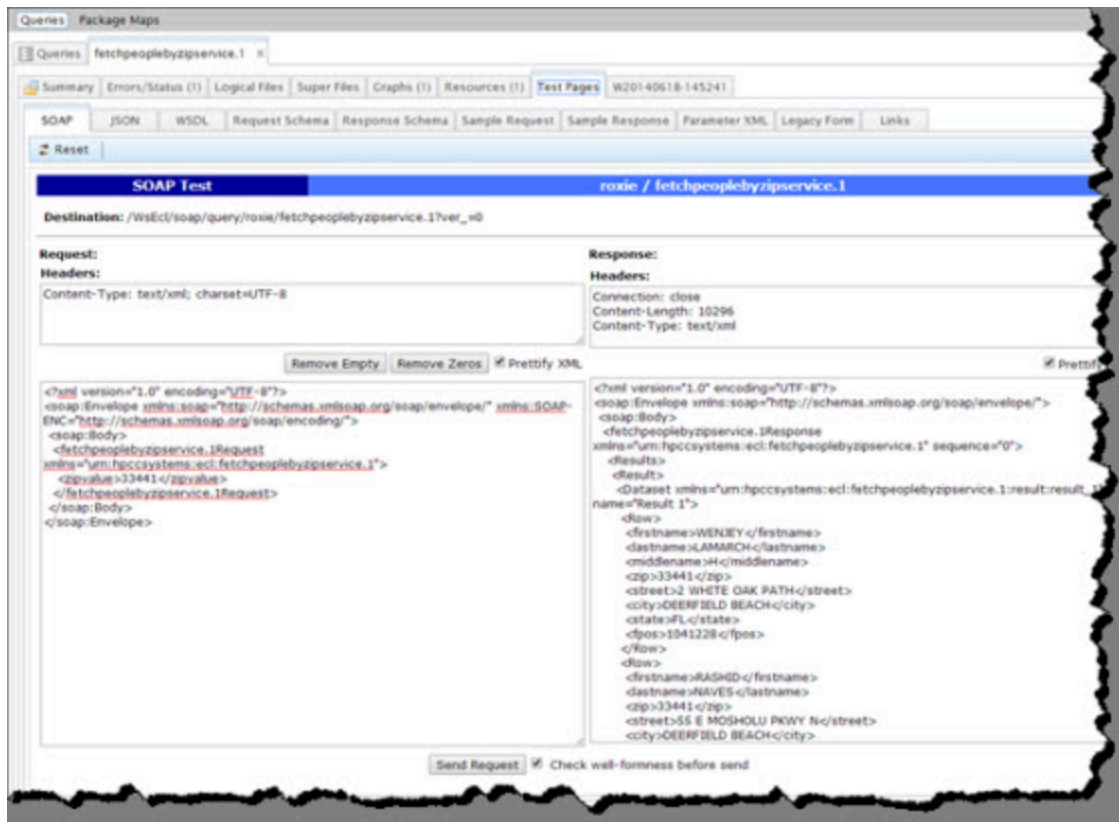


#	firstname	lastname	middlename	zip	street
1	CHERIANNE	KHATCHATOURIAN	N	545...	69 BOULDER RIDGE RD # 25A
2	MUYESSER	RAFLEE	X	207...	55 SWAMP RD
3	ROSELIN	VICECONTE		978...	107 HILL TER
4	INDA	PROVINES		729...	290 W MOUNT PLEASANT AVE
5	INDERDEEP	LAURENCE	D	323...	44 PROSPECT PL
6	CHRYSTINE	MANGIAPANE		800...	1806 1ST AVE APT 8F
7	ADELENE	STOCK	R	199...	1117 FARM RD
8	MENDY	RUFENBLANCHETTE		296...	3 W 83RD ST APT 4C
9	LANNIE	AMERANTES	I	253...	200 W 20TH ST APT 909
10	TARE	GONYEAU	T	799...	6 CANDLE CT
11	FINNEY	ARISTILDE	P	312...	222 1ST AVE APT 2B
12	OREOLUWA	MARTHALER		042...	176 CLAREMONT GDNS
13	SURGE	ABBOTTREPP	D	440...	22 LE PARC CT
14	DAVE	MCNURY		884...	510 COOPER RD

Test Pages

The Test Pages tab provides a number of resources you can use to test your query including SOAP/JSON/WSDL and the legacy WS-ECL form, as well as other tabs showing useful information or sample details about the query.

Figure 91. Test Pages tab




Information available from the Test pages tab.

- **SOAP** This tab provides an interactive interface to submit the query (with input data) and see the response in XML format.
- **JSON** This tab provides an interactive interface to submit the query (with input data) and see the response in JSON format.
- **WSDL** This tab provides a WSDL definition describing the functionality offered by the query (web service).
- **Request Schema** This tab provides a schema in XSD format describing a request for the query (web service).
- **Response Schema** This tab provides a schema in XSD format describing a response from the query (web service).
- **Sample Request** This tab provides a sample request for the query (web service) in XML Format.
- **Sample Response** This tab provides a sample response from the query (web service) in XML Format.
- **Parameter XML** This tab provides Parameterized XML representation of the query interface.
- **Legacy Form** This tab provides a form that can be used to submit a query and get a response. This is similar to the WsECL form.
- **Links** Provides a list of useful links such as: the Form, a sample REST URL, sample request, sample response, parameter XML, SOAP POST, WSDL, XSD, and the result schema.

The Workunits link

The Published queries details page provides a link to the workunits, page. This tab is a shortcut that takes you to the same workunits tab you can get to through the ECL workunits menu.

Figure 92. Queries Workunit




This is the same as the **ECL Workunits** page from the navigation sub-menu Workunits link. You can perform the same operations here. Notice that there are some other familiar tabs here as well, for example the Graphs tab, both from the Queries details page, and from the workunit tab nested here.

Package Maps

A package map provides a reference to the contents of a superkey used in queries that overrides the original definition. Package map file mappings can be organized into a collection of files defining some subsets of queries or to organize by various groupings such as functions, files, developers, etc. These subsets are called **parts**. For more information about Package Maps see the *Roxie Reference* guide.

From the Queries icon link, you can access the Package Maps page. Press the **Package Maps** button on the navigation sub-menu bar, to access the Package Maps on your cluster.

Figure 93. Package Maps



The package maps page displays all the package maps loaded on your cluster. You can Add, Activate, Deactivate, Delete, or Open a package map. To examine a package map, select a package map from the list.

To update the package maps you are using, you would either edit the package map file or add a new one and then activate it. You could later delete the old one.

Package Map Actions

You can perform actions on your package maps from the Package Maps tab in ECL Watch.

Package Map Open

To examine a package map, select the package map and press the **Open** action button. This opens a new tab where you can access additional tabs with the package information, the XML, and validate the package map.

Package Map Add

To Add a package map to the target cluster:

1. Select the package map to add by checking the box next to it.
2. Press the **Add** action button and open the Add Package Map dialog.

Figure 94. Add Package Maps

#	Type	File Name	Size
1	TXT	Mastr.txt	581

3. Press the **Select Package File** button and select the package files to add.
4. Choose the **Target** to associate the package map with.
5. Select a **Process Filter** from the drop list. The process filter determines which physical Roxie clusters will actually load the package map.
6. Enter the IP address or hostname of the remote Dali to use for logical file lookups for the **Remote Dali IP Address** field.
7. Check the boxes to Activate or Overwrite as desired.

Activate Package Map

Press the **Activate** button to deactivate the currently active package map and make the selected package map active.

Deactivate Package Map

Press the **Deactivate** button to deactivate the currently active package map.

Package Map Delete

To delete a package map:

1. Select the package map to delete by checking the box next to it.

2. Press the **Delete** action button.
3. Press **OK** when prompted to confirm.

Package Map Parts

You can see more information and perform some action on package map parts. Open the package map to see the package parts tab.

Figure 95. Package Parts



You can examine the individual parts, add parts, or remove parts through this interface in ECLWatch.

Add Part

To add a part to the package map:

1. Select the **Package Parts** tab.
2. Press the **Add** button.
3. Fill in the appropriate information.
4. Press **Apply**.


Remove Part

To remove a part from the package map:

1. Select the **Package Parts** tab.
2. Check the box next to the part to remove.
3. Press the **Remove Part** button.
4. Press **OK** when prompted to confirm.

Get Part

Figure 96. Get Part



Press the **Get Part** button to view the contents of the selected part.

Figure 97. Package Part Contents



Validate Package Map

The **Validate Package Map** tab is used to validate active package maps. The **Validate Package Content** tab is used to validate package map content that is not yet loaded. To validate an active package map:

Figure 98. Validate Package Maps



1. Select the **Validate Package Map** tab
2. Choose the **Target** and **Process** from the drop lists on the **Validate Package Map** tab.
3. Press the **Validate** button to validate the package map.

The result is shown on the **Validate Active Package Map** tab.

You can validate any package map, active, inactive, external or one not even uploaded onto the environment.

To validate an external package map:

1. Go to the **Package Maps** tab.
2. Select the package map to validate.
3. Press the **Open** action button.
4. Select the **Validate** tab.

The Validate Package Content tab allows you to open any package file, or insert any package content into the form and validate it. The content does not have to be published onto the system.

Operations

The Operations link provides access to several components useful for the day-to-day operation of your system, and some system administration access as well.

Figure 99. Operations Menu



There are links to Target Clusters, Cluster Processes, and System Servers. These links open pages with more information about the specific topology for the selected clusters. These pages are helpful in certifying that your system is up and running properly.

There are also some system administration type links such as, Users, Groups, Permissions, and Resources. These links allow you to perform some system administration tasks through ECL watch.

Topology

The Topology page provides a visual tree display with information about your clusters, services, and nodes.

Click on the **Topology** link from the Operations navigation sub-menu to access the topology page.

Figure 100. Topology Page



Targets

The **Targets** Action button displays your clusters by type.

Click on the arrow to the left of the Cluster folder/object to expand. The expanded view displays.

Figure 101. Expanded View




The expanded view displays the objects and nodes in the selected container. Select the node or object to display more information or to access the logs. The Summary, Configuration, and Log tabs on the left side of the page display the relevant information for the selected component.

Services

Press the **Services** Action button to display information on the various services running on your cluster.

Click on the arrow to the left of the service you wish to expand. The expanded view displays.

Figure 102. Expanded Services



The services view provides a service oriented tree view that provides access to the services. Expand the tree, and select the component to view the Summary, Configuration, or Logs tabs for the selected component.

Machines

Press the **Machines** Action button for more information on the various machines or nodes running in your cluster(s).

Figure 103. Machines View




The **Machines** action button opens a node oriented view of the tree. Expand the nodes to see the services on each node. Select the component to view the Summary, Configuration, or Logs tabs for that selected component.

Roxie File Copy Status

In the Topology section for Roxie clusters, the **Roxie Files Copy Status** tab shows the number of files a cluster has left to copy.


Figure 104. File Copy Status



Disk Usage

When you click on the Operations link, it opens the Disk Usage page by default. The Disk Usage page provides information about the available space on your system, and what is using that space.

Figure 105. Disk Usage



Disk Usage					
Refresh Filter					
Grouping	File Counts	Total Size	Largest File	Largest Size	Smallest File
Thor::in::imdb::actors.list	3	1,263,305,...	thor::in::imdb::actors.list	731,640,773	tutorial::...
temp::imdb::actorsinmovies	8	1,542,667,...	temp::imdb::actorsinmovies	1,081,189,...	tutorial::...
nearestneighbor::all::fields	18	4,517,084,...	nearestneighbor::all::fields	2,241,164,...	d::n

You can search or filter the results, using the **Filter** action button.

Operations: Target Clusters

The Target Clusters link from the navigation sub-menu bar, on the Operations page, opens the link to the Target Clusters page. This page provides machine information on the clusters you have set up on your machine.

Figure 106. Target Clusters

The screenshot shows the ECL Watch interface with the 'Target Clusters' tab selected. The 'Machine Information' section displays a table of machine details for two clusters: 'hthor' and 'thor'. Each cluster's table includes columns for Location, Component, Condition, State, Up Time, Processes Down, /, /mnt/disk1, Physical Memory, Swap, CPU Load, and Compute Up Time.

Location	Component	Condition	State	Up Time	Processes Down	/	/mnt/disk1	Physical Memory	Swap	CPU Load	Compute Up Time
10.239.219.3 /var/lib/HPCCSystems/myeclccserver	Ecl CC Server [myeclccserver]	Normal	Ready	3 day(s) 23:48:29	-	51%	99%	98%	100%	0 %	7 days, 22:13
10.239.219.3 /var/lib/HPCCSystems/myeclagent	Ecl Agent [myeclagent]	Normal	Ready	-	-	51%	99%	98%	100%	0 %	7 days, 22:13
10.239.219.3 /var/lib/HPCCSystems/myeclagent	Agent Exec [myeclagent]	Normal	Ready	3 day(s) 23:48:30	-	51%	99%	98%	100%	0 %	7 days, 22:13
10.239.219.3 /var/lib/HPCCSystems/myeclscheduler	Ecl Scheduler [myeclscheduler]	Normal	Ready	3 day(s) 23:48:28	-	51%	99%	98%	100%	0 %	7 days, 22:13


Location	Component	Condition	State	Up Time	Processes Down	/	/mnt/disk1	Physical Memory	Swap	CPU Load	Compute Up Time
10.239.219.5 /var/lib/HPCCSystems/mythor	Thor Slave [mythor, 1]	Normal	Ready	07:36:23	-	51%	99%	98%	100%	0 %	4 days, 3:51
10.239.219.4 /var/lib/HPCCSystems/mythor	Thor Slave [mythor, 2]	Normal	Ready	07:36:23	-	51%	99%	98%	100%	0 %	4 days, 3:51
10.239.219.3 /var/lib/HPCCSystems/mythor	Thor Master	Normal	Ready	07:36:23	-	51%	99%	98%	100%	0 %	7 days, 22:13
10.239.219.3 /var/lib/HPCCSystems/myeclccserver	Ecl CC Server [myeclccserver]	Normal	Ready	3 day(s) 23:48:29	-	51%	99%	98%	100%	0 %	7 days, 22:13

Dynamic ESDL

The Dynamic ESDL tab in ECL Watch displays the available ESP Services. You can explore the DESDL services and ESDL bindings, also known as service configurations.

To access Dynamic ESDL through ECL Watch, click on the **Operations** link, then click on **Dynamic ESDL** from the navigation sub-menu bar.

Figure 107. Dynamic ESDL sub-menu



The **Dynamic ESDL** tab contains a list of all DESDL based ESP Services and their ESDL Binding information. The DESDL-based ESP services available are listed in the navigator pane on the left as children of their parent ESP process.

The **Definitions** tab, also accessible from the **Dynamic ESDL** button's sub-menu, lists all available ESDL definitions and provides a view into any of the definitions. These definitions are used to dynamically define interfaces for existing ESP web services.


For more information about Dynamic ESDL refer to the documentation; *Dynamic ESDL* available from the HPCC Systems portal: <https://hpccsystems.com/training/documentation/learning-ecl/dynamic-esdl>

Using Dynamic ESDL

In order to use this interface choose an ESP service from the service list. Click on the triangle icon next to the ESP process (myesp) to expand and display the DESDL services. Check the box to select the desired DESDL service, then press the Open button.

The selected service's information opens and displays the **Summary** tab.

Figure 108. Dynamic ESDL Summary




Viewing the binding information

If the selected service contains an ESDL binding, you can select the binding tabs to view the binding definition (the service interface) or the configuration information.

Select the **Binding Definition** to display the ESDL definition in XML format.

Figure 109. Dynamic ESDL Binding Definition



Configuring ESDL Bindings

You can select a service and if there is a binding for it you can review, delete, or modify the configuration of that binding.

Select the **Binding Configuration** tab to view or edit the **Methods**. Press **Save** when finished.

Figure 110. Dynamic EDSL Binding Configuration



Add a Binding

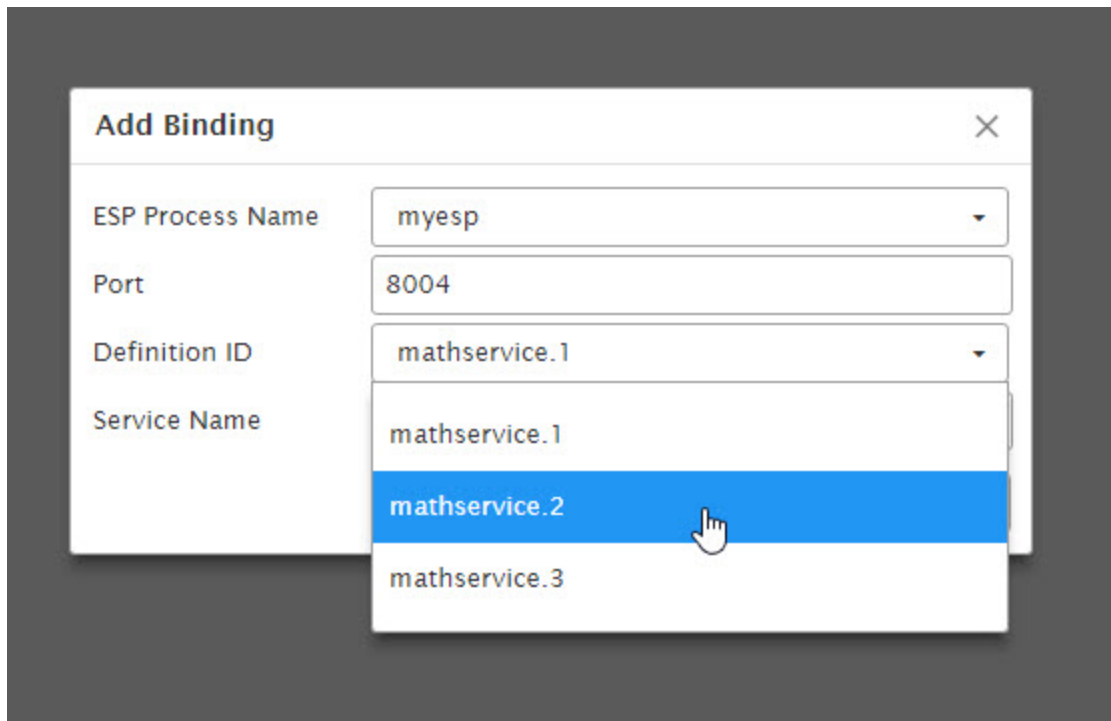
If a configuration does not have a binding, you can add a binding. To add a service binding to an *unconfigured* ESP Service. Select the unconfigured ESP service, then press the enabled **Add Binding** button.

Figure 111. Adding a service binding



This will open a dialog listing the available interfaces. Provide the information requested, an ESP Process Name and a Port. Select a Definition and then provide a Service Name.

Figure 112. Adding the definition



The screenshot shows a dialog box titled "Add Binding". It contains the following fields:

- ESP Process Name:** A dropdown menu with "myesp" selected.
- Port:** A text input field containing "8004".
- Definition ID:** A dropdown menu with "mathservice.1" selected.
- Service Name:** A list box with three items: "mathservice.1", "mathservice.2" (which is highlighted in blue and has a mouse cursor pointing at it), and "mathservice.3".

Press the **Apply** button to apply the definition.

Delete a Binding

To delete a service binding for a *configured* ESP Service. Expand the Process list as necessary and select the ESP service that contains the binding to delete.

Figure 113. Deleting service binding




Press the **Delete Binding** button. Confirm that you want to delete the binding by pressing OK on the confirmation dialog.

The binding is deleted.

Preflight System Servers

1. Click on the **Operations** icon then click on the **System Servers** link.

Figure 114. System Servers link



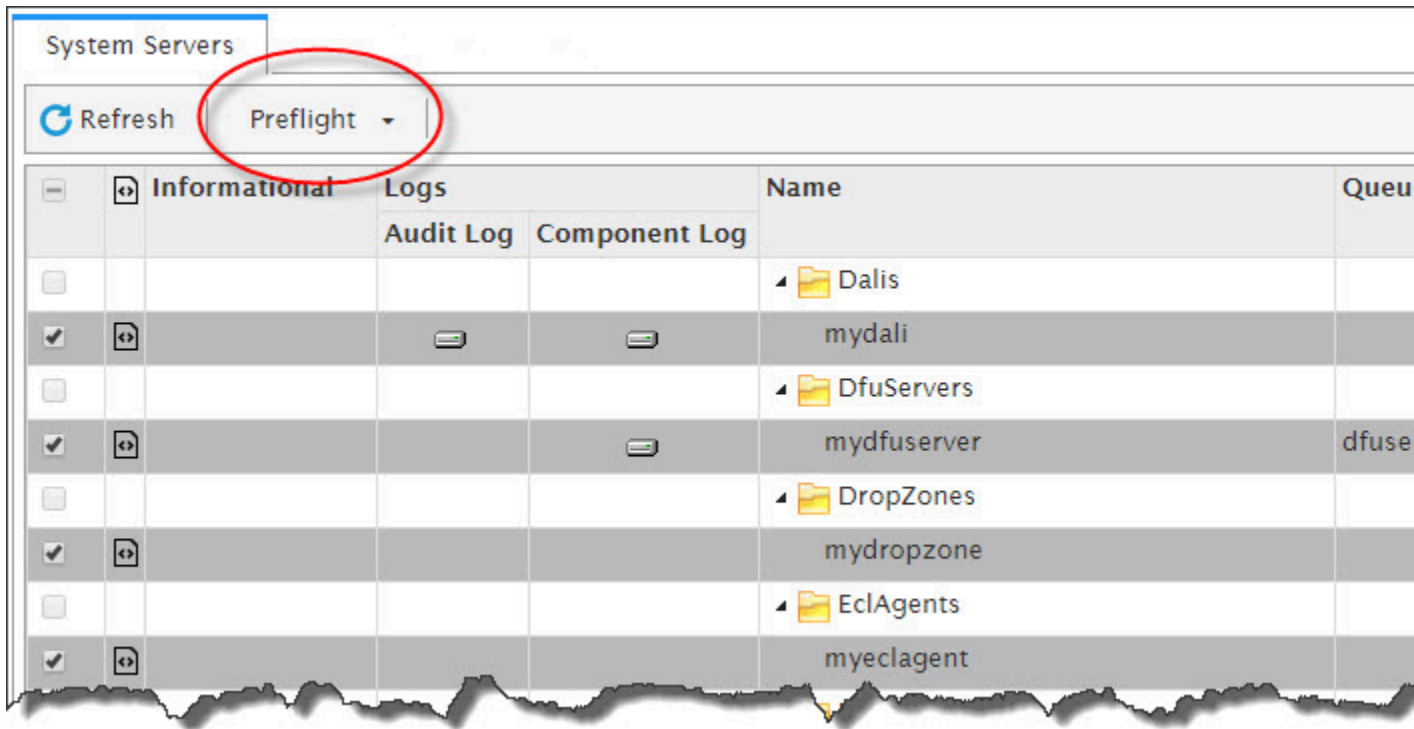
A screen similar to the following displays.

Figure 115. System Servers page



2. Expand the folder for the System Server then check the box next to the desired component(s).

Figure 116. Select System Servers



With the servers selected, the preflight action button activates and you can press it to display the preflight options.

3. Check or uncheck any desired options then Press the **Submit** button to start preflight.

Figure 117. Submit

EXPECTED RESULTS:

After pressing Submit, a screen similar to the following displays.

Figure 118. System Component Information

TopologyDisk UsageTarget ClustersCluster ProcessesSystem ServersSecurityMonitoringDynamic ESOLog Visualization

System ServersMachine Information

Preflight Results

Refresh

Location	Component	Condition	State	Processes Down	Computer Up Time	Physical M
10.176.151.31 /var/lib/HPCCSystems/mydali/	Dali Server[mydali]	Normal	Ready		4 days, 23:06	76%
10.176.151.31 /var/lib/HPCCSystems/mydfuserver/	Dfu Server[mydfuserver]	Normal	Ready		4 days, 23:06	76%
10.176.151.31 /var/lib/HPCCSystems/myeclagent/	Agent Exec[myeclagent]	Normal	Ready		4 days, 23:06	76%
10.176.151.31 /var/lib/HPCCSystems/myeclagent/	Agent Exec[myeclagent]	Normal	Ready		4 days, 23:06	76%


Using ECL Watch Operations

This screen displays information about the selected system components. This information indicates whether the components are actually running appropriately. The resulting page shows useful information about each component. The component name, location, condition, the component state, how long the component has been up and running, the amount of disk usage, memory usage and other information is available at a glance.

If there are any alerts, the component(s) are highlighted, indicating they require further attention.

For example, the following image indicates there is an issue with the DFU Server.

Figure 119. System Server Alert




Location	Component	Condition	State	Processes Down	Computer Up Time	Physical Memory
16.100.100.1 /var/lib/HPCCSystems/mydali/	Dali Server[mydali]	Normal	Ready		5 days, 3:03	81%
16.100.100.1 /var/lib/HPCCSystems/mydfuserver/	Dfu Server[mydfuserver]	Warning	Unk	mydfuserver	5 days, 3:03	81%
16.100.100.1 /var/lib/HPCCSystems/myeclagent/	Agent Exec[myeclagent]	Normal	Ready		5 days, 3:03	81%
16.100.100.1 /var/lib/HPCCSystems/myeclccserver/	Ecl CC Server[myeclccserver]	Normal	Ready		5 days, 3:03	81%
16.100.100.1 /var/lib/HPCCSystems/myeclscheduler/	Ecl Scheduler[myeclscheduler]	Normal	Ready		5 days, 3:03	81%
16.100.100.1 /var/lib/HPCCSystems/myesp/	Esp[myesp]	Normal	Ready		5 days, 3:03	81%
16.100.100.1 /var/lib/HPCCSystems/myftslave/	FT Slave[myftslave]	Normal	Ready		5 days, 3:03	

Preflight Thor


1. Click on the **Operations** icon then click on the **Cluster Processes** link.

Figure 120. Cluster Processes Link



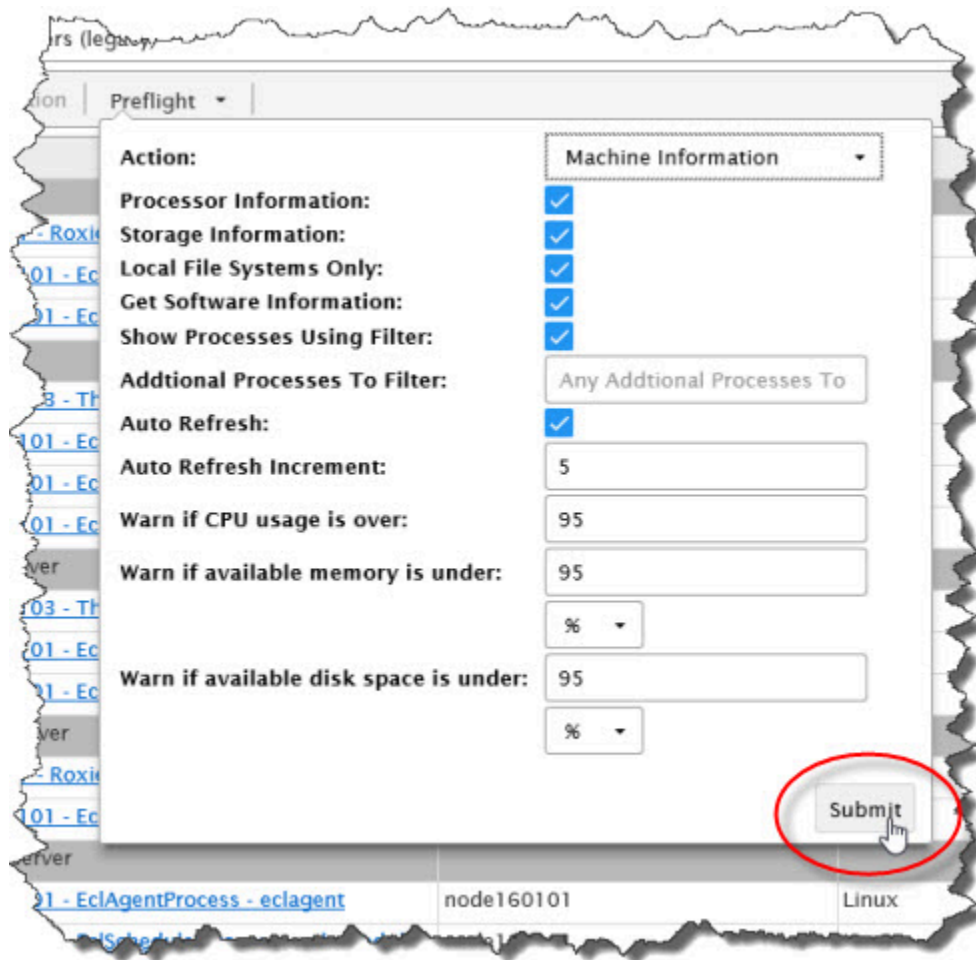
2. Expand the Thor cluster by clicking on the arrow next to the **ThorCluster** link.

Figure 121. Thor Cluster link



3. Check the box next to any individual nodes to examine or check the **Select All** checkbox in the first row.
4. With the systems selected, the preflight action button activates and you can press it to display the preflight options.
5. Select or de-select any desired options, then press the **Submit** button at the bottom to start preflight.

Figure 122. Submit



EXPECTED RESULTS:

After pressing Submit, a screen similar to the following displays.

Figure 123. Cluster Process results

TopologyDisk UsageTarget ClustersCluster ProcessesSystem ServersSecurityMonitoringDynamic ESOLLog Visualization

ClustersMachine Information

Preflight Results

Refresh

Location	Component	Condition	State	Processes Down	Computer Up Time	Physical Mem
10.179.160.1 /var/lib/HPCCSystems/thor_160/	Thor Slave[thor_160]	Normal	Ready		91 days, 2:22	23%
10.179.160.2 /var/lib/HPCCSystems/thor_160/	Thor Slave[thor_160]	Normal	Ready		225 days, 3:54	20%
10.179.160.103 /var/lib/HPCCSystems/thor_160/	Thor Master[thor_160]	Normal	Ready		434 days, 12:06	2%

This displays information on your selected cluster(s). This information can help to indicate if everything is operating normally or can help to point out any potential concerns.

If there are any notable alerts, they are highlighted. The alerts usually require some additional attention.

Users Permissions

User Administration

There are User Administration features available through ECL Watch.

Security Administration using ECL Watch

Administrator rights are needed to manage permissions. Once you have administrator access rights, open ECL Watch in your browser using the following URL:

- <http://nnn.nnn.nnn.nnn:pppp> (where nnn.nnn.nnn.nnn is your ESP Server's IP Address and pppp is the port. The default port is 8010).

Security administration is controlled using the **Security** area of ECL Watch. To access the Security area click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.



There are three areas where permissions may be set:

- **Users.** Shows all the users currently setup. Use this area to add or delete a user, edit a user's details, set/reset a user's password and view the permissions currently assigned to a user.
- **Groups.** Shows all the groups currently setup. Use this area to add or delete a group, view and edit the members of a group, view and edit the permissions that have been set for a group.
- **Permissions.** Shows the features of the HPCC Systems where permissions may be set. Use this area to view the permissions currently set for any area of HPCC Systems, or to add groups and users and set/modify their permission for a specific feature



NOTE: Use caution when setting any explicit **deny** permission setting. The most restrictive permission always applies.

Information about your account

To find out more information about your account, in ECL Watch click on your username link under **Logged In As:** at the top of the ECL Watch page.



- A **User Details** tab with your account information displays.

A screenshot of a 'User Details' dialog box. The dialog has a title bar with 'User Details' and a close button. Below the title bar is a 'Save' button. The main content area displays the following information:
FranklinX
Username: FranklinX
Employee ID: 99999
First Name: Franklin
Last Name: Xavier
Old Password: [text input field]
New Password: [text input field]
Confirm Password: [text input field]
Password Expiration: Never
The dialog box is set against a background that looks like a torn piece of paper with some faint, illegible text visible underneath.

- You can change your password here, if desired.
- You can also verify the password expiration date, if your password is set to expire.


Setting and modifying user permissions

In a security-enabled environment, access to ECL Watch and its features is controlled using a login and password. The **Users** area enables you to control who has access to ECL Watch and the features of your HPCC Systems to which they have access. Permissions can be set for users based on their individual needs and users can also be added to groups which have already been set up. Use the **Users** menu item to:

- Add a new user (**note:** the Username cannot be changed)
- Delete a user
- Add a user to a group
- Change a user's password
- Modify the details/permissions of an individual user

Adding and editing users


To access the user administration sections click on the **Operations** icon, then click the **Security** link from the navigation sub-menu. Click on the **Users** tab to add or edit users.



All current users are identified in the list by their Username and Full Name.

To add a new user to the list of authenticated users:

To add a new user you must have Administrator level access.



1. Press the **Add** button.

The add user dialog displays.

2. Enter a **Username**.

This is the login name to use ECL Watch, ECL IDE, WsECL, etc.

3. Enter the **First Name** and **Last Name** of the user.

This information helps to easily identify the user and is displayed in the **Full Name** field on the main **Users** window.

4. Enter a **Password** for the user and then confirm it in the **Retype Password** field.

NOTE: The password must conform to the policy of your security manager server.

5. Press the **Add** button.

A successful addition opens a new tab where you can verify the new user's information.

6. Press the **Save** button.

Once added, the new user displays in the list and you can modify details and set permissions as required.

To modify a user's details:


Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Users** tab.

The users display in a list.

2. Select the user (or users) to modify. Click on the **Username** link to open the users' details tab.

To select multiple users, check the box next to the Username to select. This enables the Users action buttons. Press the **Open** action button.



A tab opens for each user selected. On each user's tab there are several sub-tabs.

The user's details are on the **Summary** tab.

3. Modify the user's details as required (if more than one user selected, repeat for each user).

Note: The **Username** cannot be changed.

4. Press the **Save** button.

A confirmation message displays.

To add a user to a group:


Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Users** tab.

The users display in a list.

2. Select the user (or users) to modify. Click on the **Username** link to open the users' details tab.

To select multiple users, check the box next to the Username to select. This enables the Users action buttons. Press the **Open** action button.




A tab opens for each user selected. On each user's tab there are several sub-tabs.

The user's details are on the **Summary** tab.

3. Click on the tab for the user to modify (if more than one user selected, repeat for each user).

On the user's tab there are several sub-tabs.



Click on the **Member Of** sub-tab to modify that user's groups.

4. On the **Member Of** tab for that user, a list of the available groups display.


To add the user to the group, check the box next to the desired group.

5. The changes are automatically saved. Close the tab.

To promote a user to an Administrator

To modify a user's credentials you must have Administrator level access. You can designate the HPCC Systems Administrator account to have limited permissions only relating to HPCC Systems elements and not LDAP administrator's rights. To promote a user to an HPCC Systems Administrator, add the user to the configured **Administrators** group.

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.



1. Click on the **Users** tab.

The users display in a list.

2. Select the user (or users) to promote. Click on the **Username** link to open the users' details tab.

To select multiple users, check the box next to the Username to select. This enables the Users action buttons. Press the **Open** action button.



A tab opens for each user selected. On each user's tab there are several sub-tabs.

The user's details are on the **Summary** tab.

3. Click on the tab for the user to modify (if more than one user selected, repeat for each user).

On the user's tab there are several sub-tabs.

Click on the **Member Of** sub-tab.



4. Select **Administrators** by placing a check in box.

NOTE: The name of the default Administrators group could vary. It is a configurable value defined as the value of **adminGroupName** in the configuration. For example, if you set the adminGroupName to "HPCCAdministrators", in the environment then HPCCAdministrators would display in the list.

5. The changes are automatically saved. Close the tab(s).

To delete a user from a group:

To delete a user from a group you must have Administrator level access.

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Users** tab.

The users display in a list.

2. Select the user (or users) to remove. Click on the **Username** link to open the users' details tabs.

To select multiple users, check the box next to the Username to select. This enables the Users action buttons. Press the **Open** action button.



A tab opens for each user selected. On each user's tab there are several sub-tabs.

3. Click on the tab of the user to modify (if multiple users selected, repeat for each user).

On the user's tab there are several sub-tabs.



Click on the **Member Of** sub-tab to modify that user's groups.

4. On the **Member Of** tab for that user, there is a list of the available groups.

There is a check in the box next to each group that user belongs to.

To remove that user from a group, uncheck the box next to the desired group.

5. The changes are automatically saved. Close the tab.

To change a user's password:

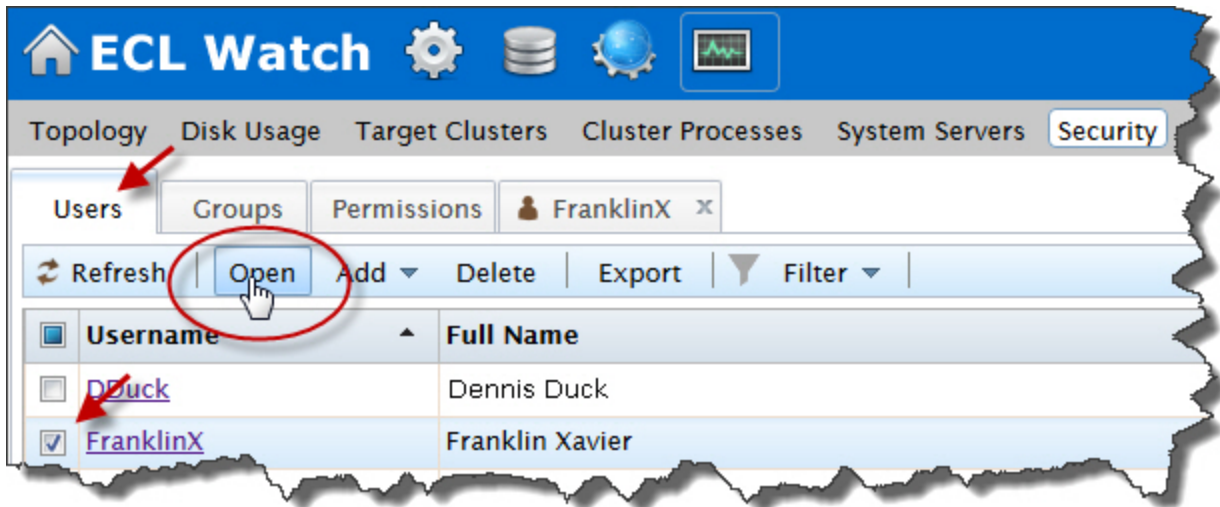
Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Users** tab.

The users display in a list.

2. Select the user (or users) to modify. Click on the **Username** link to open the users' details tab.

To select multiple users, check the box next to the Username to select. This enables the Users action buttons. Press the **Open** action button.



A tab opens for each user selected. On that tab there are several sub-tabs.

The user details are on the **Summary** tab.

3. Select the **Summary** tab.
4. Change the password in the **Password** and **Retype New Password** fields as required on the User details summary tab (if multiple users selected, repeat for each user).

Note: The **Username** cannot be changed.

5. Press the **Save** button.

A confirmation message displays.

To delete a user from the list of authenticated users:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Users** tab.

The users display in a list.

2. Check the box to the left of the user(s) you want to remove.

Note: These users will no longer have access to ECL Watch.

3. Press the **Delete** action button.

Confirmation displays.

Setting permissions for an individual user

There may be occasions when you need to modify the permissions for individual users. For example, users may have individual security needs that are not completely covered in any group or, there may be occasions when a user requires temporary access to an HPCC Systems feature. Permissions set in this area of ECL Watch only affect the user you choose. Most individual permissions you set here overwrite ones set in any group to which the user belongs, except in the case of an explicit deny.

To set permissions for an individual user:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Users** tab.


The users display in a list.

2. Select the user (or users) to modify. Click on the **Username** link to open the users' details tab.

To select multiple users, check the box next to the Username to select. This enables the Users action buttons. Press the **Open** action button.

3. Click on the tab of the username to modify (if multiple users selected, repeat for each user).

On the user's tab there are several sub-tabs.




Click on the **Active Permissions** sub-tab to view the user's current permissions.

4. Click on the **Available Permissions** tab to see all the sets of permissions that are available to apply to that user.

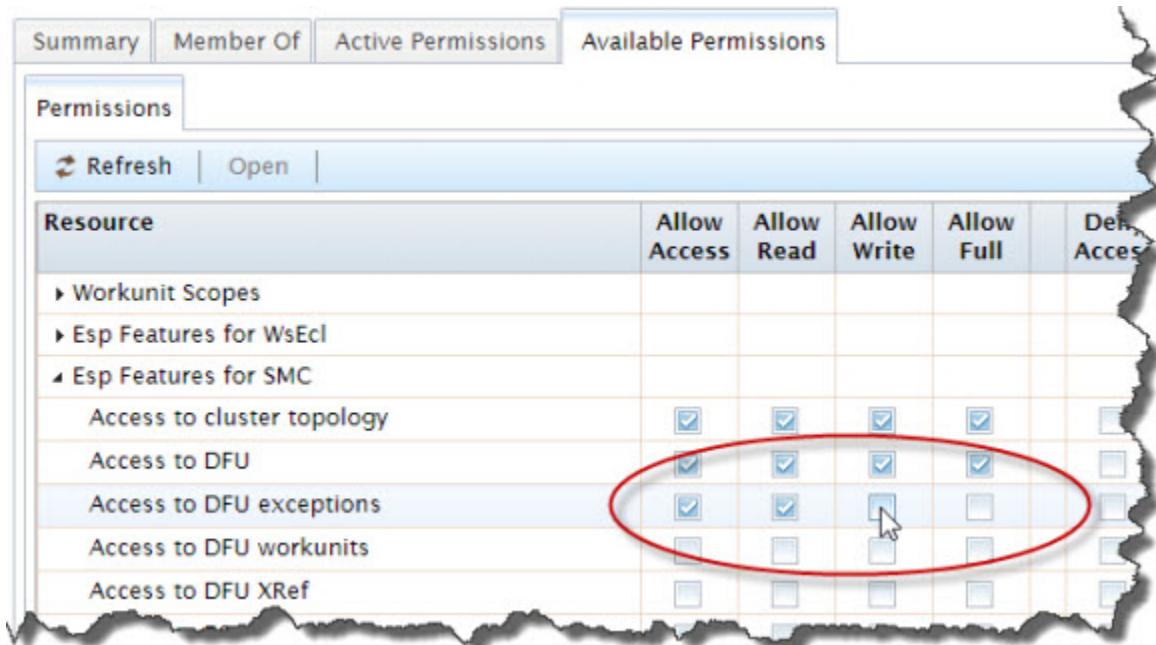
When you select permissions from the Available Permissions tab, they display and can be set in the Active Permissions tab.

5. Click on the arrow next to the resource to display the permissions that can be set for that resource.



The list of permission groups currently set for this user and the ones the user has inherited are also listed. Click the arrow to allow setting the individual resource settings.

6. There may be more than one resource setting available in each group, be sure to set the permissions for each setting as required.
7. Check the boxes that **allow** and **deny** access as required for the user.



NOTE: Use caution when setting any explicit **deny** permission setting. The most restrictive permission always applies.

8. The changes are automatically saved. Close the tab.

Setting and modifying group permissions

Setting up groups ensures that all users with the same permission needs have the same permission settings. You can give users the access they require to the feature areas of HPCC Systems that they need. There is no limit to the number of groups you can create. You can create as many groups as you need to control access for all your users regardless of their tasks.

Use the **Groups** menu item to:

- Add a new group.
- Delete a group.
- Add members to a group.
- Modify the permissions for a group.

Adding and editing groups

When adding or changing the permissions for a group, all members of that group are given those permission settings. So it is important to be sure that you are giving or denying access to features appropriate for the members of that group. If you need to make a change for a single user (or small number of users), it is probably better to make that change for each individual user as illustrated in the previous sections.



To modify groups, click on the **Operations** icon, then click the **Security** link from the navigation sub-menu. Click on the **Groups** tab.

To add a new group:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Groups** tab.
2. Press the **Add** action button.



This opens a dialog where you can enter the name for the group.

3. Enter a **Group Name**.
4. Enter the fully qualified Distinguished Name for the owner of the group **Managed By** field.
5. Enter a description of the group. (optional)
6. Press the **Add** button.

This opens a new tab for the group and several sub tabs

The **Summary** sub-tab displays the group name.

The **Members** tab displays the list of users, check the box next to each user to add to the group.

The **Active Group Permissions** tab displays the permissions applied to the group.

The **Available Group Permissions** tab displays all the available permissions, selecting from the Available Permissions applies them to the Active Group Permissions.

You can set the permissions and add members to this group from the respective sub-tabs on that group tab.

To delete a group:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Groups** tab.
2. Locate the group in the list and check the checkbox next to it.
3. Press the **Delete** action button.

4. Press the **OK** confirmation button.

The group no longer displays in the list.

To add new members to a group:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Groups** tab.
2. Locate the group in the list and check the box next to it.
3. Press **Open** action button.

This opens a new tab for the group.

The sub-tabs display: **Summary**, **Members**, **Active Group Permissions**, and **Available Group Permissions**.

4. Select the **Members** tab.

The members tab displays a list of all users on the system. The users that belong to the selected group have a check in the box next to them.

5. Check the box(es) to the left of the users you want to add to the group.
6. The changes are automatically saved. Close the tab.

To delete members from a group:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click on the **Groups** tab.
2. Locate the group in the list and check the box next to it.
3. Press the **Open** action button.

This opens a new tab for the group.

The Groups tab has several sub-tabs: **Summary**, **Members**, **Active Group Permissions** and **Available Group Permissions**.

4. Select the **Members** tab.

The Members tab displays a list of all users on the system. The users that belong to the selected group have a check in the box next to them.

5. Uncheck the box(es) to the left for all users you want to delete from the group.
6. The changes are automatically saved. Close the tab.

Setting permissions for a group

By default, all users are members of the **Authenticated Users** group. The **Authenticated Users** group has access rights to almost all resources. To set up more restricted controls, you should create specific groups with more restricted permissions.


You can then create groups with only those access rights you wish to grant. This approach allows the most flexibility since a single User ID can have multiple group memberships.

As a best practice, you should use **Allow** instead of **Deny** to control access. Denies should be used only as an exception, when possible. If you wish to deny a user access to some specific control, a good practice would be to create a group for that, place the user(s) in that group, then you can deny access to that group.

Remember the most restrictive control takes precedence. For example, if a user is in a group that has deny permission to file access, and the user is in another group where file access is allowed, that user will still not have file access.


To set permissions for a group:

Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.




1. Click the **Groups** tab.
2. Locate the group in the list and check the box next to it.
3. Press the **Open** action button.

This opens a new tab for the group.



The group tab displays the sub-tabs: **Summary**, **Members**, **Active Group Permissions** and **Available Group Permissions**.

4. Select the **Available Group Permissions** sub-tab. This displays all the available permission resources.
5. Click on the arrow to the left of the **Resource** to expand and expose the permission sets for the resources.



The groups permission resources display.

6. There may be more than one resource setting available in each group, be sure to set the permissions for each setting as required.

7. Check the boxes for **allow** and **deny** as required for the group.

Resource	Allow Access	Allow Read	Allow Write	Allow Full
Workunit Scopes				
Esp Features for WsEcl				
File Scopes				
Esp Features for SMC				
Access to cluster topology	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Access to super computer environment	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Access to DFU	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Access to DFU exceptions	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Access to DFU workunits	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Access to DFU XRef	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>



NOTE: Use caution when setting any explicit **deny** permission setting. The most restrictive permission always applies.

8. There may be more than one resource setting available, select the resource(s) you require from the drop list.

Repeat for each applicable resource.

9. The changes are automatically saved. Close the tab.

Plugins

You can add functionality to ECL Watch by installing plugins. These plugins are designed to integrate into the ECL Watch interface. After you install an approved plugin, the plugin icon displays in the navigation bar at the top of the ECL Watch page to provide access to the plugin(s). Click on the plugin icon to view the plugins page in ECL Watch.


Figure 124. ECL Watch Plugin icon



Ganglia in ECL Watch

With the HPCC Systems® Ganglia-monitoring plugin installed, you can view the Ganglia statistics and graphs through the ECL Watch interface. The default monitoring displays several key statistics, but you can customize and configure the views.


Figure 125. Ganglia in ECL Watch



Customize Monitoring

The default Ganglia page has a tab for Custom Monitoring where you can easily add custom monitoring components.


Figure 126. Ganglia Custom Monitoring



To customize the monitoring page;

1. Select the **Custom Monitoring** tab.
2. Press the **Metrics** button.
3. Use the drop menus to display the various graphing utilities.

Figure 127. Customize the Metrics



Installing Ganglia in ECL Watch

In order to use Ganglia in ECL Watch, you need to have Ganglia installed on your HPCC Systems platform. For details on installing Ganglia for ECL Watch, refer to the *HPCC Systems Monitoring and Reporting* manual.

Nagios in ECL Watch

ECL Watch is set up for monitoring your system with Nagios. ECL Watch has an API that can interface with Nagios and provide Nagios monitoring right in ECL Watch. Nagios escalations can be pointed to any ECL Watch version 5.4 (and later) and are viewable directly in ECL Watch.

By default all ECL Watch services defined in the environment.xml will receive notifications generated using *hpcc-nagios-tools*. You can override that if not desired. The ECL Watch instances need not be in the cluster that is being monitored.

Figure 128. Nagios in ECL Watch




Once you have Nagios configured for your environment, you can see at a glance if there are any alerts. Along the top banner of the ECL Watch window, you will see a small indicator light. The light is darkend (gray) if there is no system data being reported, typically indicative that your system is not yet configured for monitoring.

The light is green when all systems are reporting normal. The light is yellow when there is warning. The light turns red when there is an alert. All the alerts are configurable through the Nagios configuration.

By default ECL Watch monitoring maintains the latest update for 30 minutes. This means that once Nagios stops escalations to ECL Watch any status, including Normal will expire from the list. Nagios escalations notification behavior and frequency is configurable, refer to the Nagios documentation for more information. An empty list could indicate 'no data' or 'no outages', by default no alerts generate when everything is up and running.

To delve further into any warnings or alerts, you can press the indicator light at the top. You can also access the *Primary Monitoring* page by pressing the **Operations** link, then press the **Monitoring** link in the navigation sub-menu.

Figure 129. Nagios in ECL Watch



The screenshot shows the ECL Watch interface with the Monitoring tab selected. A table displays monitoring data. A red oval highlights the first two rows of the table.

Name	Details	URL	IP	Time
HostProcess			Host1:1111	2015-07-07 10:23:10
nagios	Host Is Down	localhost:8010		
EspProcess			10.239.1...	2015-07-07 10:23:10
EspProcess			10.239.1...	2015-07-07 10:23:10
EspProcess			10.239.1...	2015-07-07 10:23:10
EspProcess			10.239.1...	2015-07-07 10:23:10
EspProcess			10.239.1...	2015-07-07 10:23:10

This displays the all the messages and alerts reported to the monitoring system. For more information on a specific message, press the arrow next to the message you want.

Resources

The resources link can be found under the Operations Icon link. The resources link in ECL Watch provides a link to the HPCC Systems® web portal. Visit the HPCC Systems® Web Portal at <http://hpccsystems.com/> for software updates, plugins, support, documentation, and more. This is where you can find resources useful for running and maintaining HPCC Systems on the web portal.

You can also get to the resources link on the HPCC Systems® web portal page, by clicking on the **Additional Resources** link found on the sub-menu of at the top right hand side of navigation bar.

ECL Watch provides a link to the HPCC Systems download page: <http://hpccsystems.com/download>. This is the page where you can download Installation packages, virtual images, source code, documentation, and tutorials.

Appendix A. HPCC Systems Session Management

Session Management

The 7.0 release of HPCC Systems platform introduces the new ESP Session Management security feature. This functions like many banking applications, where after a configurable period of inactivity you are warned with a "You are about to be locked out" pop-up. If no action further is taken, the session is then locked and you would need to enter your credentials to unlock and resume. A session remains active while there is regular user interaction. After a period of inactivity, you are alerted that your session is about to be locked. Sessions are stored in cookies and are shared across tabs and instances of each browser. Activity in any instance will extend the entire session duration. Additionally, a Logout menu option allows you to close your session when you are finished.

FAQ

1. **Q:** Why did we implement this feature?

A: The main motive is to tighten security. Browsers and the IDE left open after hours and over the weekend are a security risk. Additionally, this reduces unnecessary load on ESP since it will not auto refresh inactive ECLWatch sessions.

2. **Q:** How long will an inactive session last?

A: Your administrator can configure this using Configuration Manager. The default setting is two hours of inactivity.

3. **Q:** Does Auto Refresh of active workunits and graphs extend your ESP session?

A: No, only user actions such as typing or mouse clicks extend a session.

4. **Q:** Will I have to login to ECLWatch?

A: Yes, just as you currently do.

5. **Q:** Will I have to login to the ECL IDE?

A: Yes, but you already should be. No perceptible changes here only behind the scenes where you are being authenticated.

6. **Q:** Will I have to login to Configuration Manager?

A: No.

7. **Q:** What credentials should I use to login with?

A: Use your assigned credentials.

8. **Q:** Can I log out of ECL Watch?

A: Yes, there is a link to logout available. You are able to log off, and if you do not your session locks after a configurable period of inactivity.

9. **Q:** Will my sessions get logged off due to inactivity?

A: No. After a configurable period of inactivity your session locks. You then need to unlock to resume your session.

10.Q: How long until my password expires?

A: This depends on your system policies and the configured security manager.

11.Q: Will I be able to log in as a different user?

A: Yes, with our new login screen, you can input previously used IDs or enter a different one. You can have as many user sessions active at any time as permitted by your system's resources.

12.Q: Can I log in concurrently with different credentials?

A: Yes, using different tabs in a single browser, multiple instances of the same browser, or multiple instances of different browsers.

13.Q: Is there an option to stay logged in indefinitely and/or not time out from inactivity?

A: No.

14.Q: Will I lose data if I get automatically logged out?

A: No. You do not get logged out. Your session will get locked. Anything typed into any fields (such as a search box) that has not been submitted or entered could potentially be lost. However, since the session is only locked, it is unlikely that any data will be lost.

15.Q: Will my queued and scheduled workunits run when I am locked out?

A: Yes, the session only applies to ESP/ECL IDE and ESP/ECLWatch communications.

16.Q: Will HPCC Systems command line utilities be affected?

A: Possibly. If you have configured AuthPerSessionOnly then command line utilities will not work. If AuthPerSessionOnly is not enabled then command line utilities will not be effected.

17.Q: Does auto refresh in ECLWatch reset the session expiration timer?

A: No. Only active interactions like mouse clicks and pressing keys extend the timeout. Note that scrolling does not extend the expiration timer.

18.Q: If I am logged in to the same account using multiple tabs in a browser, or multiple instances of the same browser, can I get locked out of one but not the others?

A: No, activity is tracked by your credentials. Activity in one tab or instance extends the session for all.

19.Q: If I am logged in to the same account using different browsers (e.g., Firefox and Chrome), do they share the same session timeout?

A: No. Since each browser has its own cookie store, activity in one does not extend to the other.

20.Q: Can I automatically return to the ECLWatch screen where I was when automatically locked out?

A: Yes. The intent is to lock your session and not completely log you out. Unlocking your session should return you to the same point when your session locked.

21.Q: Will I be able to change an expired password?

A: Yes. You are redirected to a page where you can reset your password.

22.Q: Will access to ECLWatch require SSL/TLS and HTTPS?

A: These secure protocols are already available for your HPCC Systems Administrator to configure. Though not required for session management, hopefully they are currently enabled.

23.Q: Will my programmatic SOAP calls utilizing ESP have any impact?

A: Maybe. If you have configured AuthPerSessionOnly then SOAP calls will not work. If your system is not configured that way, then programmatic SOAP calls continue to operate as they do now.

24.Q: When will I see the Session Management changes?

A: You can configure your system to use Session Management as part of HPCC Systems Version 7.0.