



HPCC System Administrator's Guide

Boca Raton Documentation Team

HPCC System Administrator's Guide

Boca Raton Documentation Team

Copyright © 2016 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.

2016 Version 6.2.0-1

Introducing HPCC Systems® Administraton	4
Introduction	4
Architectural Overview	5
Hardware and Software Requirements	11
Network Switch	11
Load Balancer	12
System Sizings	13
Nodes-Software	14
Hardware and Component Sizing	16
Thor Hardware	17
Roxie Hardware Configurations	18
Dali and Sasha Hardware Configurations	19
Other HPCC Components	20
Routine Maintenance	21
Data Handling	22
Back Up Data	22
Log Files	25
Preflight	27
Preflight System Servers	28
Preflight Target Clusters	31
Preflight Thor	34
Preflight the Roxie Cluster	37
System Configuration and Management	39
Running the Configuration Manager	42
Environment.conf	48
Configuring HPCC for Authentication	50
User Security Maintenance	64
Workunits and Active Directory	102
Using Cassandra for Workunit Storage	103
Redefining nodes in a Thor Cluster	110
Best Practices	111
Cluster Redundancy	111
High Availability	113
Best Practice Considerations	115
Capacity Planning	117
Sample Sizings	120
System Resources	122
HPCC Resources	122
Additional Resources	123

Introducing HPCC Systems® Administraton

Introduction

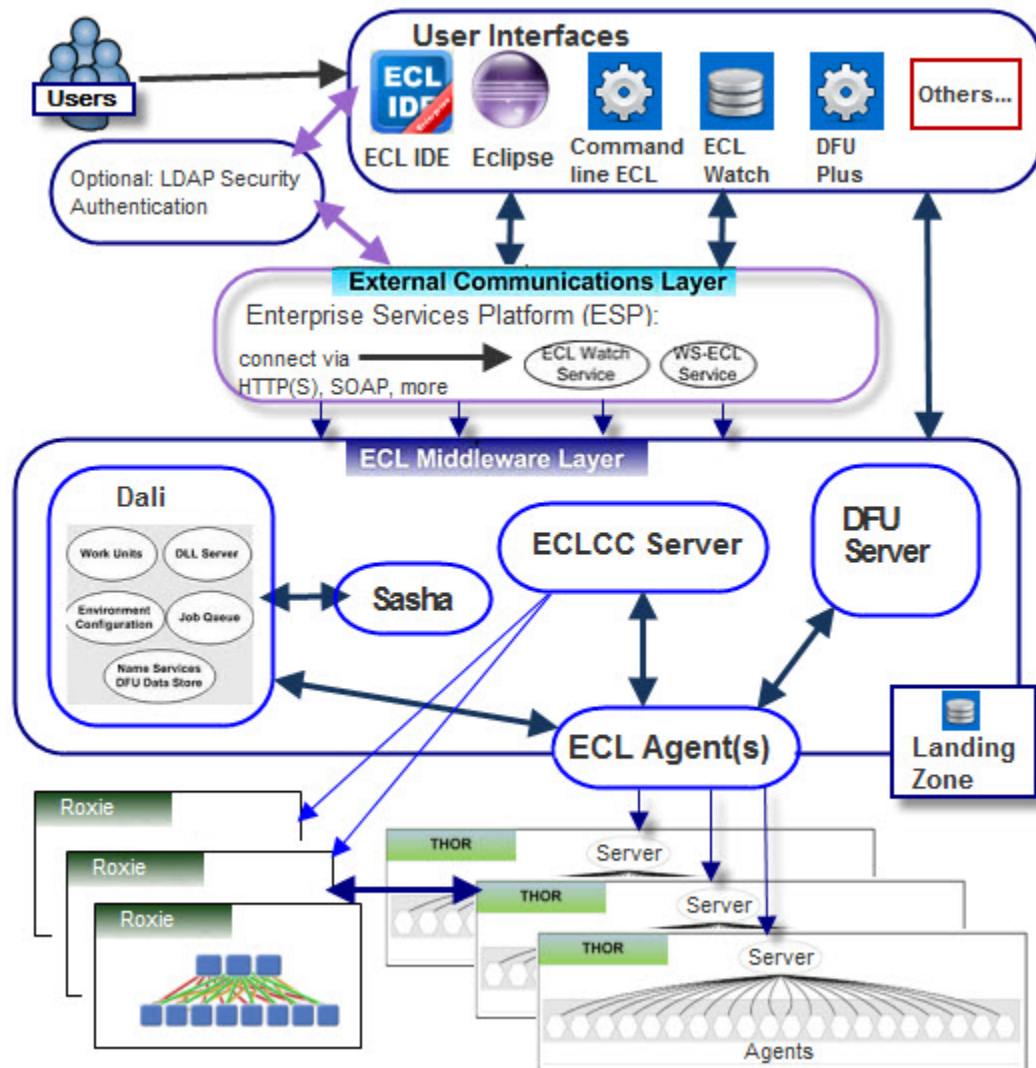
HPCC (High Performance Computing Cluster) is a massive parallel-processing computing platform that solves Big Data problems.

HPCC stores and processes large quantities of data, processing billions of records per second using massive parallel processing technology. Large amounts of data across disparate data sources can be accessed, analyzed, and manipulated in fractions of seconds. HPCC functions as both a processing and a distributed data storage environment, capable of analyzing terabytes of information.

Architectural Overview

An HPCC Systems® Platform consists of the following components: Thor, Roxie, ESP Server, Dali, Sasha, DFU Server, and ECLCC Server. LDAP security is optionally available.

Figure 1. HPCC Architectural Diagram

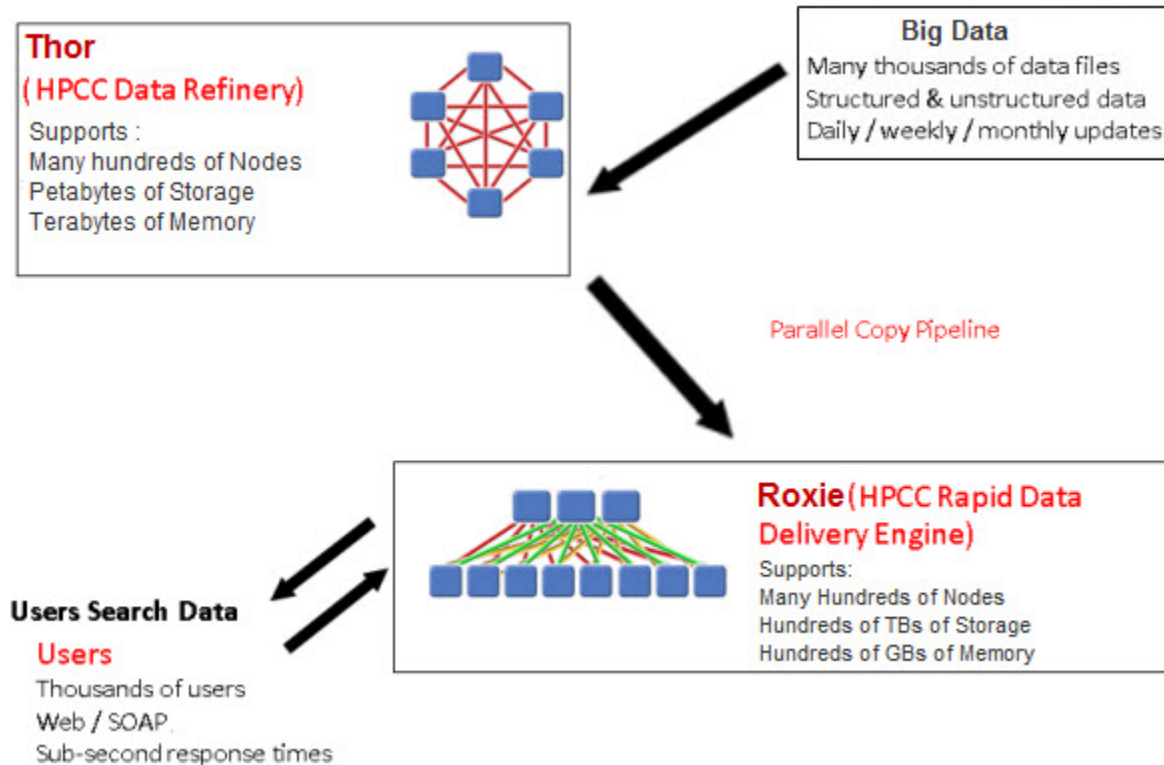


Data loading is controlled through the Distributed File Utility (DFU) server.

Data typically arrives on the landing zone (for example, by FTP). File movement (across components) is initiated by DFU. Data is copied from the landing zone and is distributed (sprayed) to the Data Refinery (Thor) by the ECL code. Data can be further processed via ETL (Extract, Transform, and Load process) in the refinery.

A single physical file is distributed into multiple physical files across the nodes of a cluster. The aggregate of the physical files creates one logical file that is addressed by the ECL code.

Figure 2. Data Processing



The data retrieval process (despraying) places the file back on the landing zone.

Clusters

HPCC environment contains clusters which you define and use according to your needs. The types of clusters used in HPCC:

Thor

Data Refinery (Thor) – Used to process every one of billions of records in order to create billions of "improved" records. ECL Agent (hThor) is also used to process simple jobs that would be an inefficient use of the Thor cluster.

Roxie

Rapid Data Delivery Engine (Roxie) – Used to search quickly for a particular record or set of records.

Queries are compiled and published, usually in ECL Watch. Data moves in parallel from Thor nodes to the receiving Roxie nodes. Parallel bandwidth utilization improves the speed of putting new data into play.

ECL Agent

The ECL Agent's primary function is to send the job to execute on the appropriate cluster. The ECL Agent can act as a single-node cluster. That is called spawning an hThor cluster. hThor is used to process simple jobs that would otherwise be an inefficient use of Thor. For simple tasks, the ECL Agent will make a determination and perform the execution itself by acting as an hThor cluster.

Figure 3. Clusters



System Servers

The System Servers are integral middleware components of an HPCC system. They are used to control workflow and intercomponent communication.

Dali

Dali is also known as the system data store. It manages workunit records, logical file directory, and shared object services.

It maintains the message queues that drive job execution and scheduling. It also enforces the all LDAP security restrictions.

Sasha

The Sasha server is a companion “housekeeping” server to the Dali server. It works independently of all other components. Sasha’s main function is to reduce the stress on the Dali server. Whenever possible, Sasha reduces the resource utilization on Dali.

Sasha archives workunits (including DFU Workunits) which are stored in a series of folders.

Sasha also performs routine housekeeping such as removing cached workunits and DFU recovery files.

DFU Server

DFU server controls the spraying and despraying operations used to move data in and out of Thor.

DFU services are available from:

- Standard libraries in ECL code.
- Client interfaces: Eclipse, ECL Playground, ECL IDE, and the ECL command line interface.
- DFU Plus command line interface.

ECLCC Server

ECLCC Server is the compiler that translates ECL code. When you submit ECL code, the ECLCC Server generates optimized C++ which is then compiled and executed. ECLCC Server controls the whole compilation process.

When you submit workunits for execution on Thor, they are first converted to executable code by the ECLCC Server.

When you submit a workunit to Roxie, code is compiled and later published to the Roxie cluster, where it is available to execute multiple times.

ECLCC Server is also used when the ECL IDE requests a syntax check.

ECLCC Server uses a queue to convert workunits one at a time, however you can have ECLCC Servers deployed in the system to increase throughput and they will automatically load balance as required.

ECL Agent

ECL Agent (hThor) is a single node process for executing simple ECL Queries.

ECL Agent is an execution engine that processes workunits by sending them to the appropriate cluster. ECL Agent processes are spawned on-demand when you submit a workunit.

ESP Server

ESP (Enterprise Service Platform) Server is the inter-component communication server. ESP Server is a framework that allows multiple services to be “plugged in” to provide various types of functionality to client applications via multiple protocols.

Examples of services that are plugged into ESP include:

- **WsECL:** Interface to published queries on a Roxie, Thor, or hThor cluster.
- **ECL Watch:** A web-based query execution, monitoring, and file management interface. It can be accessed via the ECL IDE or a web browser. See *Using ECL Watch*.

The ESP Server supports both XML and JSON Formats.

LDAP

You can incorporate a Lightweight Directory Access Protocol (LDAP) server to work with Dali to enforce the security restrictions for file scopes, workunit scopes, and feature access.

When LDAP is configured, you need to authenticate when accessing ECL Watch, WsECL, ECL IDE, or any other client tools. Those credentials are then used to authenticate any requests from those tools.

Client Interfaces

The following Client Interfaces are available to interact with the HPCC Platform.

Eclipse

With the ECL plug-in for Eclipse, you can use the Eclipse IDE to create and execute queries into your data on an HPCC platform using Enterprise Control Language (ECL). Eclipse is open-source, and multi-platform and it can be used to interface with your data and workunits on HPCC. The ECL plug-in for Eclipse is also open-source.

ECL IDE

ECL IDE is a full-featured GUI providing access to your ECL code for ECL development. ECL IDE uses various ESP services via SOAP.

The ECL IDE provides access to ECL Definitions to build your queries. These definitions are created by coding an expression that defines how some calculation or record set derivation is to be done. Once defined, they can be used in succeeding ECL definitions.

ECL Watch

ECL Watch is a web-based query execution, monitoring, and file management interface. It can be accessed via ECL IDE, Eclipse, or a web browser. ECL Watch allows you to see information about and manipulate workunits. It also allows you monitor cluster activity and perform other administrative tasks.

Using ECL Watch you can:

- Browse through previously submitted workunits (WU). You can see a visual representation (graphs) of the data flow within the WU, complete with statistics which are updated as the job progresses.
- Search through files and see information including record counts and layouts or sample records.

- See the status of all system servers.
- View log files.
- Add users or groups and modify permissions.

See the *Using ECL Watch* Manual for more details.

Command Line Tools

Command line tools: **ECL**, **DFU Plus**, and **ECL Plus** provide command line access to functionality provided by the ECL Watch web pages. They work by communicating with the corresponding ESP service via SOAP.

See the *Client Tools* Manual for more details.

Hardware and Software Requirements

This chapter describes some of the hardware and software requirements in order to run the HPCC System. HPCC is designed to run on commodity hardware, which makes building and maintaining large scale (petabytes) clusters economically feasible. When planning your cluster hardware, you will need to balance a number of considerations specific to your needs.

This section provides some insight into the hardware and infrastructure that HPCC works well on. This is not an exclusive comprehensive set of instructions, nor a mandate on what hardware you must have. Consider this as a guide to use when looking to implement or scale your HPCC system. These suggestions should be taken into consideration for your specific enterprise needs.

Network Switch

The network switch is a significant component of the HPCC System.

Switch requirements

- Sufficient number of ports to allow all nodes to be connected directly to it;
- IGMP v.2 support
- IGMP snooping support

Ideally your HPCC system will perform better when each node is connected directly into a single switch. You should be able to provide a port for each node on a single switch to optimize system performance. Your switch size should correspond to the size of your system. You would want to ensure that the switch you use has enough capacity for each node to be plugged into it's own port.

Switch additional recommended features

- Gigabit speed
- Non-blocking/Non-oversubscribed backplane
- Low latency (under 35usec)
- Layer 3 switching
- Managed and monitored (SNMP is a plus)
- Port channel (port bundling) support

Generally, higher-end, higher throughput switches are also going to provide better performance. For larger systems, a high-capacity managed switch that can be configured and tuned for HPCC efficiency is the best choice.

Load Balancer

A load balancer distributes network traffic across a number of servers. Each Roxie Node is capable of receiving requests and returning results. Therefore, a load balancer distributes the load in an efficient manner to get the best performance and avoid a potential bottleneck.

Load Balancer Requirements

Minimum requirements

- Throughput: 1 Gigabit
- Ethernet ports: 2
- Balancing Strategy: Round Robin

Standard requirements

- Throughput: 8Gbps
- Gigabit Ethernet ports: 4
- Balancing Strategy: Flexible (F5 iRules or equivalent)

Recommended capabilities

- Ability to provide cyclic load rotation (not load balancing).
 - Ability to forward SOAP/HTTP traffic
 - Ability to provide triangulation/n-path routing (traffic incoming through the load balancer to the node, replies sent out the via the switch).
 - Ability to treat a cluster of nodes as a single entity (for load balancing clusters not nodes)
- or
- Ability to stack or tier the load balancers for multiple levels if not.

System Sizings

This section provides some guidance in determining the sizing requirements for an initial installation of HPCC. The following are some suggested configuration guides that can be helpful when planning your system.

Minimum Suggested Hardware

HPCC was designed to run on common commodity hardware, and could function on even lesser hardware. The following list is the suggested minimum hardware specifications. At the very minimum you should consider the following hardware components for your HPCC system. These guidelines were put together based on real world usage of mission critical (uptime) with high volume data.

Thor slave	Processor	4 x 64-bit Intel Processor per
	RAM	8GB per daemon
	Storage	RAID - 200MB/sec Sequential Read/Write per node
	Network	1 Gb/sec bandwidth
Roxie	Processor	4 x 64-bit Intel Processor
	RAM	12GB per Roxie
	Storage	400 IOPS & 2 Volumes per (RAID optional)
	Network	1 Gb/sec bandwidth
Dali	Processor	4 x 64-bit Intel Processor each
	RAM	24GB per Dali
	Storage	RAID 1, 5, 6, 10 Volume 200GB
Other	Processor	4 x 64-bit Intel Processor
	RAM	12GB
	Storage	RAID 1, 5, 6, 10 Volume 200GB
	Network	1 Gb/sec bandwidth

Nodes-Software

All nodes must have the identical operating systems. We recommend all nodes have identical BIOS settings, and packages installed. This significantly reduces variables when troubleshooting. It is easier to manage a system where all nodes are identical, but this is not required.

Operating System Requirements

Binary installation packages are available for many Linux Operating systems. HPCC System platform requirements are readily available on the HPCC Portal.

<http://hpccsystems.com/permlink/requirements>

Dependencies

Installing HPCC on your system depends on having required component packages installed on the system. The required dependencies can vary depending on your platform. In some cases the dependencies are included in the installation packages. In other instances the installation may fail, and the package management utility will prompt you for the required packages. Installation of these packages can vary depending on your platform. For details of the specific installation commands for obtaining and installing these packages, see the commands specific to your Operating System.

Note: For CentOS installations, the Fedora EPEL repository is required.

SSH Keys

The HPCC components use ssh keys to authenticate each other. This is required for communication between nodes. A script to generate keys has been provided. You should run that script and distribute the public and private keys to all nodes after you have installed the packages on all nodes, but before you configure a multi-node HPCC.

- As root (or sudo as shown below), generate a new key using this command:

```
sudo /opt/HPCCSystems/sbin/keygen.sh
```

- Distribute the keys to all nodes. From the **/home/hpcc/.ssh** directory, copy these three files to the same directory (**/home/hpcc/.ssh**) on each node:
 - **id_rsa**
 - **id_rsa.pub**
 - **authorized_keys**

Make sure that files retain permissions when they are distributed. These keys need to be owned by the user "**hpcc**".

Hardware and Component Sizing

This section provides some insight as to what sort of hardware and infrastructure optimally HPCC works well on. This is not an exclusive comprehensive set of instructions, nor a mandate on what hardware you must have. Consider this as a guide to use when looking to implement or scale your HPCC system. These suggestions should be taken into consideration for your specific enterprise needs.

HPCC is designed to run on commodity hardware, which makes building and maintaining large scale (petabytes) clusters economically feasible. When planning your cluster hardware, you will need to balance a number of considerations, including fail-over domains and potential performance issues. Hardware planning should include distributing HPCC across multiple physical hosts, such as a cluster. Generally, one type of best practice is to run HPCC processes of a particular type, for example Thor, Roxie, or Dali, on a host configured specifically for that type of process.

Thor Hardware

Thor slave nodes require a proper balance of CPU, RAM, network, and disk I/O in order to operate most efficiently. A single Thor slave node works optimally when allocated 4 CPU cores, 8GB RAM, 1Gb/sec network and 200MB/sec sequential read/write disk I/O.

Hardware architecture can provide higher value within a single physical server. In such cases you can use multi-slave to configure your larger physical servers to run multiple Thor slave nodes per physical server.

It is important to note that HPCC by nature is a parallel processing system and all Thor slave nodes will be exercising at precisely the same time. So when allocating more than one HPCC Thor slave per physical machine assure that each slave meets the recommended requirements.

For instance, 1 physical server with 48 cores, 96GB RAM, 10Gb/sec network and 2GB/sec sequential I/O would be capable of running ten (10) HPCC Thor slaves at optimal efficiency. The order of optimization for resource usage in a Thor slave node is disk I/O 60%, network 30%, and CPU 10%. Any increase in sequential I/O will have the most impact on speed, followed by improvements in network, followed by improvements in CPU.

Network architecture is also an important consideration. HPCC Thor nodes work optimally in a streamlined network architecture between all Thor slave processes.

RAID is recommended and all RAID levels suitable for sequential read/write operations and high availability are acceptable. For example, RAID1, RAID10, RAID5 (preferred), and RAID6.

Roxie Hardware Configurations

HPCC Roxie processes require a proper, yet different (from Thor) balance of CPU, RAM, network, and disk I/O in order to ensure efficient operations. A single HPCC Roxie node works optimally when allocated 6 or more CPU cores, 24GB RAM, 1Gb/sec network backbone, and 400/sec 4k random read IOPS.

Each HPCC Roxie node is presented two hard drives, each capable of 200/sec 4k random seek IOPS. Hard drive recommendations for Roxie efficiency are 15K SAS, or SSD. A good rule of thumb is the more random read IOPS the better and faster your Roxie will perform.

Running multiple HPCC Roxie nodes on a single physical server is not recommended, except in the cases of virtualization or containers.

Configure your system to balance the size of your Thor and Roxie clusters. The number of Roxie nodes should never exceed the number of Thor nodes. In addition, the number of Thor nodes should be evenly divisible by the number of Roxie nodes. This ensures an efficient distribution of file parts from Thor to Roxie.

Dali and Sasha Hardware Configurations

HPCC Dali processes store cluster metadata in RAM. For optimal efficiency, provide at least 48GB of RAM, 6 or more CPU cores, 1Gb/sec network interface and a high availability disk for a single HPCC Dali. HPCC's Dali processes are one of the few active/passive components. Using standard “swinging disk” clustering is recommended for a high availability setup. For a single HPCC Dali process, any suitable High Availability (HA) RAID level is fine.

Sasha does not store any data. Sasha reads data from Dali then processes it. Sasha does store archived workunits (WUs) on a disk. Allocating a larger disk for Sasha reduces the amount of housekeeping needed. Since Sasha assists Dali by performing housekeeping, it works best when on its own node. You should avoid putting Sasha and Dali on the same node.

Other HPCC Components

ECL Agent, ECLCC Server, DFU Server, the Thor master, and ECL Watch are administrative processes which are used for supporting components of the main clusters.

For maximum efficiency you should provide 24GB RAM, 6+ CPU cores, 1Gb/sec network and high availability disk(s). These components can be made highly available in an active/active fashion.

Routine Maintenance

In order to ensure that your HPCC system keeps running optimally, some care and maintenance is required. The following sections address routine maintenance tasks for your HPCC system.

Data Handling

When you start working with your HPCC system, you will want to have some data on the system to process. Data gets transferred to and the HPCC system by a process called a spray. Likewise to get data out from an HPCC system it must be desprayed.

As HPCC is a computer cluster the data gets deployed out over the nodes that make up the cluster. A *spray* or import is the relocation of a data file from one location (such as a Landing Zone) to a cluster. The term spray was adopted due to the nature of the file movement – the file is partitioned across all nodes within a cluster.

A *despray* or export is the relocation of a data file from a Data Refinery cluster to a single machine location (such as a Landing Zone). The term despray was adopted due to the nature of the file movement – the file is reassembled from its parts on all nodes in the cluster and placed in a single file on the destination.

A *Landing Zone* (or drop zone) is a physical storage location defined in your system's environment. There can be one or more of these locations defined. A daemon (dfilesrv) must be running on that server to enable file sprays and desprays. You can spray or despray some files to your landing zone through ECL Watch. To upload large files, you will need a tool that supports the secure copy protocol, something like a WinSCP.

For more information about HPCC data handling see the *HPCC Data Handling* and the *HPCC Data Tutorial* documents.

Back Up Data

An integral part of routine maintenance is the back up of essential data. Devise a back up strategy to meet the needs of your organization. This section is not meant to replace your current back up strategy, instead this section supplements it by outlining special considerations for HPCC Systems®.

Back Up Considerations

You probably already have some sort of a back up strategy in place, by adding HPCC Systems® into your operating environment there are some additional considerations to be aware of. The following sections discuss back up considerations for the individual HPCC system components.

Dali

Dali can be configured to create its own back up, ideally you would want that back up kept on a different server or node. You can specify the Dali back up folder location using the Configuration Manager. You may want to keep multiple copies that back up, to be able to restore to a certain point in time. For example, you may want to do daily snapshots, or weekly.

You may want to keep back up copies at a system level using traditional back up methods.

Sasha

Sasha itself generates no original data but archives workunits to disks. Be aware that Sasha can create quite a bit of archive data. Once the workunits are archived they are no longer available in the Dali data store. The archives can still be retrieved, but that archive now becomes the only copy of these workunits.

If you need high availability for these archived workunits, you should back them up at a system level using traditional back up methods.

DFU Server

DFU Server has no data. DFU workunits are stored in Dali until they are archived by Sasha.

ECLCC Server

ECLCC Server stores no data. ECL workunits are stored in Dali and archived by Sasha.

ECL Agent

ECL Agent stores no data.

ECL Scheduler

ECL Scheduler stores no data. ECL Workunits are stored in Dali.

ESP Server

ESP Server stores no data. If you are using SSL certificates, public and private keys they should be backed up using traditional methods.

Thor

Thor, the data refinery, as one of the critical components of HPCC Systems® needs to be backed up. Back up Thor by configuring replication and setting up a nightly back up cron task. Back up Thor on demand before and/or after any node swap or drive swap if you do not have a RAID configured.

A very important part of administering Thor is to check the logs to ensure the previous back ups completed successfully.

Backupnode

Backupnode is a tool that is packaged with HPCC. Backupnode allows you to back up Thor nodes on demand or in a script. You can also use backupnode regularly in a crontab. You would always want to run it on the Thor master of that cluster.

The following example is one suggested way for invoking backupnode manually.

```
/bin/su - hpcc -c "/opt/HPCCSystems/bin/start_backupnode thor" &
```

The command line parameter must match the name of your Thor cluster. In your production environment, it is likely that you would provide descriptive names for your Thor clusters.

For example, if your Thor cluster is named thor400_7s, you would call start_backupnode thor400_7s.

```
/bin/su - hpcc -c "/opt/HPCCSystems/bin/start_backupnode thor400_7s" &
```

To run backupnode regularly you could use cron. For example, you may want a crontab entry (to back up thor400_7s) set to run at 1am daily:

```
0 1 * * * /bin/su - hpcc -c "/opt/HPCCSystems/bin/start_backupnode thor400_7s" &
```

Backupnode writes out its activity to a log file. That log can be found at:

/var/log/HPCCSystems/backupnode/MM_DD_YYYY_HH_MM_SS.log

The (MM) Month, (DD) Day, (YYYY) 4-digit Year, (HH) Hour, (MM) Minutes, and (SS) Seconds of the back up comprising the log file name.

The main log file exists on the Thor master node. It shows what nodes it is run on and if it finished. You can find other backupnode logs on each of the Thor nodes showing what files, if any, it needed to restore.

It is important to check the logs to ensure the previous back ups completed successfully. The following entry is from the backupnode log showing that back up completed successfully:

```
00000028 2014-02-19 12:01:08 26457 26457 "Completed in 0m 0s with 0 errors"
00000029 2014-02-19 12:01:08 26457 26457 "backupnode finished"
```

Roxie

Roxie data is protected by three forms of redundancy:

- **Original Source Data File Retention:** When a query is published, the data is typically copied from a remote site, either a Thor or a Roxie. The Thor data can serve as back up, provided it is not removed or altered on Thor. Thor data is typically retained for a period of time sufficient to serve as a back up copy.
- **Peer-Node Redundancy:** Each Slave node typically has one or more peer nodes within its cluster. Each peer stores a copy of data files it will read.
- **Sibling Cluster Redundancy:** Although not required, Roxie may run multiple identically-configured Roxie clusters. When two clusters are deployed for Production each node has an identical twin in terms of queries and/or data stored on the node in the other cluster. This configuration provides multiple redundant copies of data files. With three sibling Roxie clusters that have peer node redundancy, there are always six copies of each file part at any given time; eliminating the need to use traditional back up procedures for Roxie data files.

Landing Zone

The Landing Zone is used to host incoming and outgoing files. This should be treated similarly to an FTP server. Use traditional system level back ups.

Misc

Back up of any additional component add-ons, your environment files (environment.xml), or other custom configurations should be done according to traditional back up methods.

Log Files

You can review system messages and see any error messages as they are reported and captured in log files. Log files can help you in understanding what is occurring on the system and useful in troubleshooting.

Component Logs

There are log files for each component in directories below `/var/log/HPCCSystems` (default location). You can optionally configure the system to write the logs in a different directory. You should know where the log files are, and refer to the logs first when troubleshooting any issues.

There are log files which record activity among the various components. You can find the log files in subdirectories named corresponding to the components that they track. For example, the Thor logs would be found in a directory named `mythor`, the sasha log would be in the `mysasha` directory, the esp log in the `myesp` directory.

In each of the component subdirectories, there are several log files. Most of the log files use a logical naming convention that includes the component name, the date, and time in the name of the log file. There is also usually a link for the component with a simple name, such as `esp.log` which is a short cut to the latest current log file for that component.

Understanding the log files, and what is normally reported in the log files, helps in troubleshooting the HPCC system.

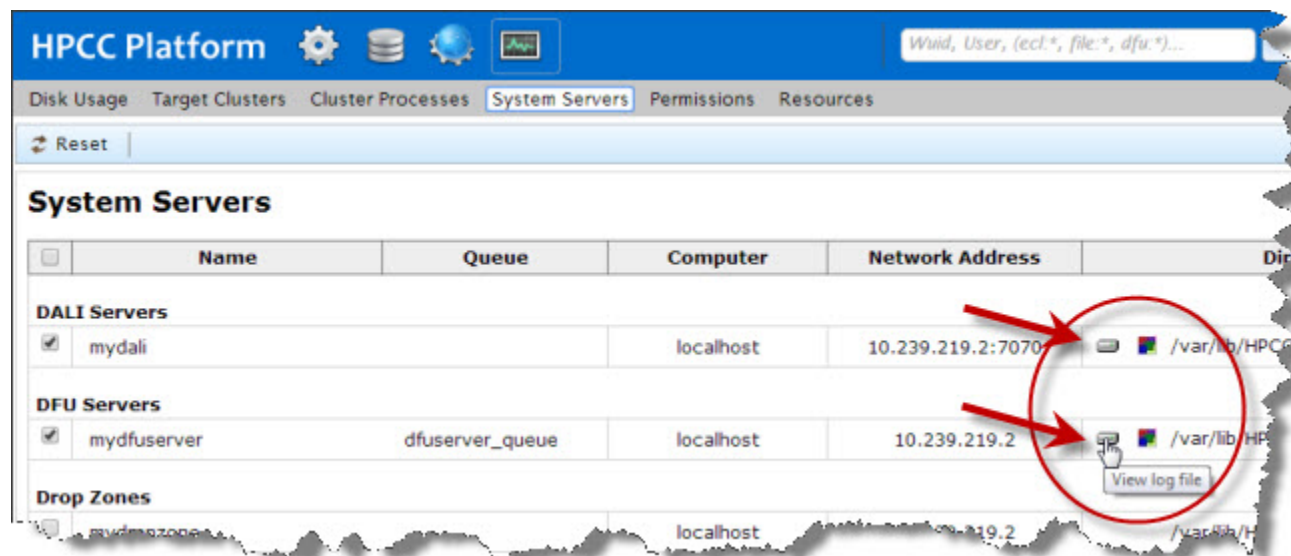
As part of routine maintenance you may want to back up, archive, and remove the older log files.

Accessing Log Files

You can access and view the log files directly by going to the component log directory from a command prompt or a terminal application. You can also view the component log files through ECL Watch.

To view logs on ECL Watch, click on the **Operations** icon, then click on the **System Servers** link. That opens the System Servers page in ECL Watch. There are several HPCC system components listed on that page. In the **Directory** column for each component there is a computer drive icon. Click the icon in the row for the component log you wish to view.

Figure 4. Logs in ECL Watch



You can also view log files from the other links under the Operations icon in ECL Watch.

1. Click on the **Target Clusters** link to open the tab with links to your system's clusters.
2. Click on the computer drive icon (circled in red in the above figure), in the row of the cluster and node of the component log you wish to view.

To view cluster process logs:

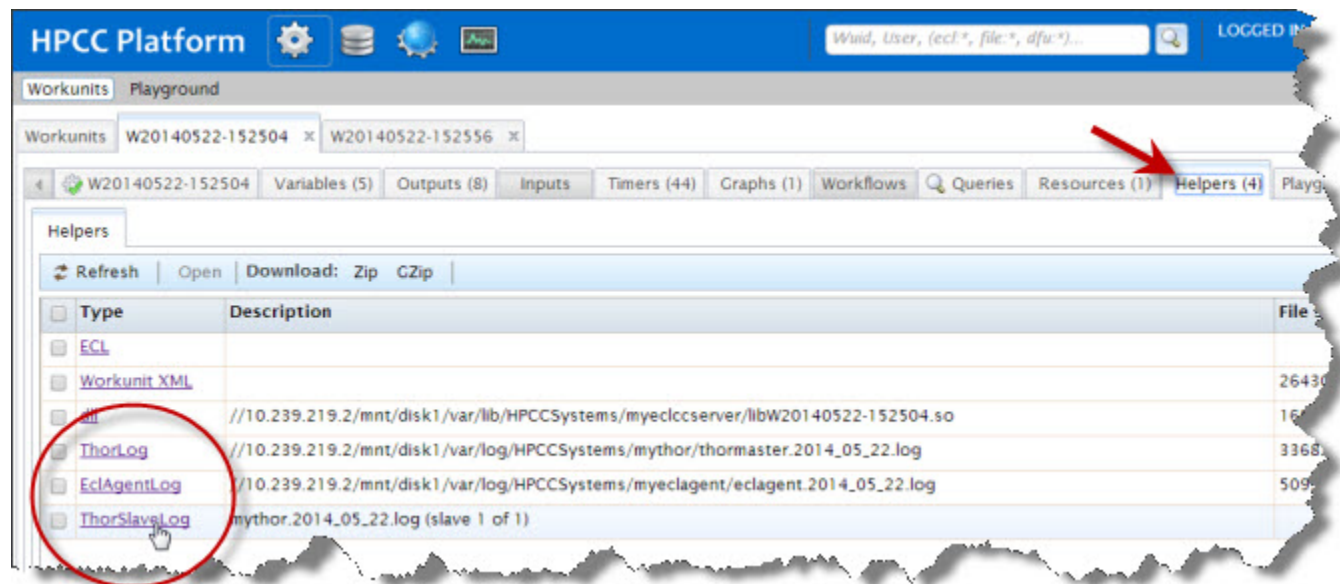
1. Click on the **Cluster Processes** link to open the tab with links to your system's clusters processes.
2. Click on the cluster process you wish to view more information about.

For example, click on the **myroxie** link. You will then see a page of all that components nodes. You will see computer drive icon, in the row of each node. Click that icon to see the logs for the cluster process for that node.

Log files in ECL Workunits

You can also access the Thor or ECL Agent log files from the ECL Workunits. (not available for Roxie workunits) In ECL Watch when examining the Workunit details, you will see a **Helpers** tab. Click on the Helpers tab to display the relevant log files for that particular workunit.

Figure 5. Logs in ECL Watch Workunits



Preflight

The first step in certifying that the platform is installed and configured properly is to run a preflight check on the components. This ensures that all machines are operating and have the proper executables running. This also confirms there is adequate disk space, available memory, and acceptable available CPU % values.

- 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)




Note: That your IP address could be different from the ones provided in these figures. Please use the IP address provided by your installation.

Preflight System Servers

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

Figure 6. System Servers link



A screen similar to the following displays.

Figure 7. System Servers page

System Servers				
<input type="checkbox"/>	Name	Queue	Computer	Network
Dali Servers				
<input checked="" type="checkbox"/>	mydali		localhost	192.168.1.1
DFU Servers				
<input checked="" type="checkbox"/>	mydfuserver	dfuserver_queue	localhost	192.168.1.1
Drop Zones				
<input type="checkbox"/>	mydropzone		localhost	192.168.1.1
ECL Agents				
<input checked="" type="checkbox"/>	myeclagent		localhost	192.168.1.1

2. Press the **Submit** button at the bottom of this page to start preflight.

Figure 8. Submit

☒ Get storage information
☒ Local File Systems Only
☒ Get software information
☒ Show processes using filter
 Additional processes to filter:
☒ Auto Refresh every 5 mins

EXPECTED RESULTS:

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

Figure 9. System Component Information

Machine Information

<input checked="" type="checkbox"/>	Location	Component	Condition	State	Up Time	Processes Down	
<input checked="" type="checkbox"/>	10.239.219.3 /var/lib/HPCCSystems/myesp	Esp [myesp]	Normal	Ready	09:38	-	60%
<input checked="" type="checkbox"/>	10.239.219.3 /var/lib/HPCCSystems/myeclscheduler	Ecl Scheduler [myeclscheduler]	Normal	Ready	11:32	-	60%
<input checked="" type="checkbox"/>	10.239.219.3 /var/lib/HPCCSystems/myeclagent	Agent Exec [myeclagent]	Normal	Ready	11:35	-	60%
<input checked="" type="checkbox"/>	10.239.219.3 /var/lib/HPCCSystems/myeclccserver	Ecl CC Server [myeclccserver]	Normal	Ready	11:33	-	60%
<input checked="" type="checkbox"/>	10.239.219.4 /var/lib/HPCCSystems/mysasha	Sasha Server [mysasha]	Normal	Ready	11:51	-	60%
<input checked="" type="checkbox"/>	10.239.219.4 /var/lib/HPCCSystems/mydali	Dali Server [mydali]	Normal	Ready	11:54	-	60%
<input checked="" type="checkbox"/>	10.239.219.5 /var/lib/HPCCSystems/mydfuserver	Dfu Server [mydfuserver]	Normal	Ready	11:29	-	60%

☒ Select All / None
 Fetched: 11/11/11 14:13:09
 Action:

☒ Get processor information Warn if CPU usage is over 95%
☐ Get memory information Warn if memory usage is over 95%

This screen displays information on several system components. This information indicates whether several components are actually running appropriately. The resulting page shows useful information about each component. The component name, the 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 failed components, they are highlighted in orange, indicating they are not ready.

Figure 10. Failed Component

Machine Information

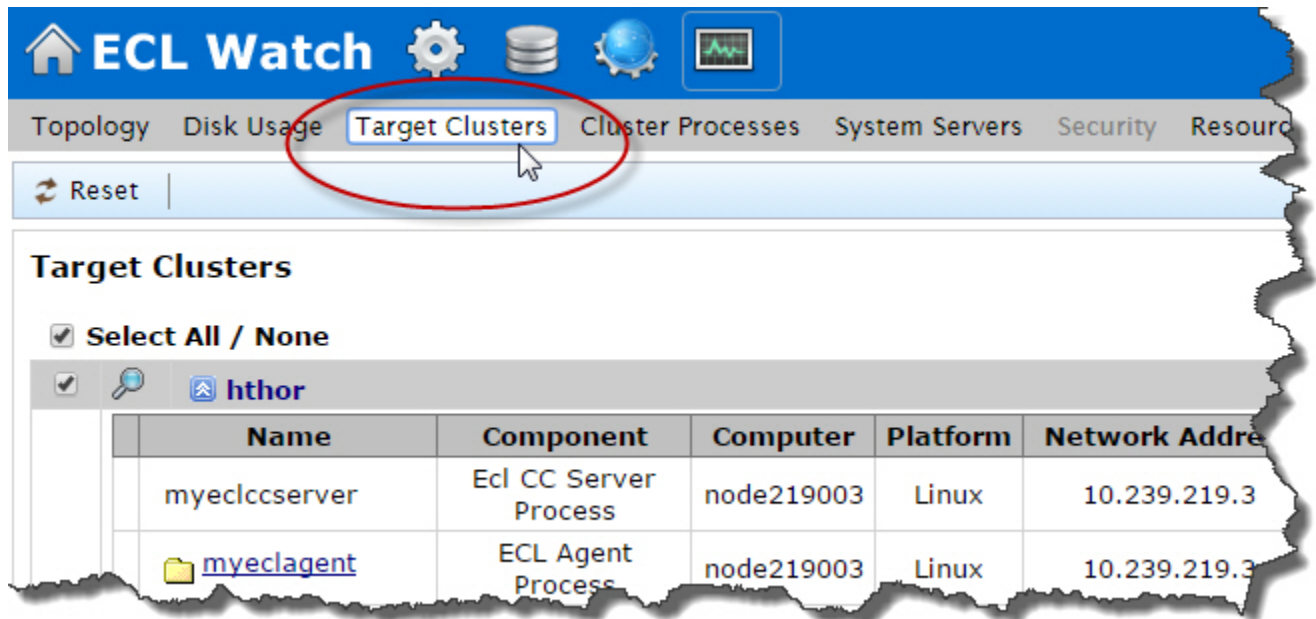
<input checked="" type="checkbox"/>	Location	Component	Condition	State	Up Time	Processes Down	/	/mnt/disk1	Physical Memory
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/mydali	Dali Server [mydali]	Normal	Ready	76 day(s) 00:50:02	-	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/mydfuserver	Dfu Server [mydfuserver]	Warning	Unknown		mydfuserver	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/myeclagent	Ecl Agent [myeclagent]	Normal	Ready	-	-	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/myeclagent	Agent Exec [myeclagent]	Normal	Ready	76 day(s) 00:50:00	-	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/myeclccserver	Ecl CC Server [myeclccserver]	Normal	Ready	76 day(s) 00:49:59	-	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/myeclscheduler	Ecl Scheduler [myeclscheduler]	Normal	Ready	76 day(s) 00:49:57	-	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/myesp	Esp [myesp]	Normal	Ready	76 day(s) 00:48:10	-	43%	97%	96%
<input checked="" type="checkbox"/>	10.239.219.2 /var/lib/HPCCSystems/mysasha	Sasha Server [mysasha]	Normal	Ready	76 day(s) 00:49:54	-	43%	97%	96%

Preflight Target Clusters

Use the Target Clusters link to preflight all your clusters at once. To preflight clusters individually, see the following *Preflight Thor* or *Preflight Roxie* sections.

1. Click on the **Operations** icon then click on the **Target Clusters** link.

Figure 11. Target Clusters Link

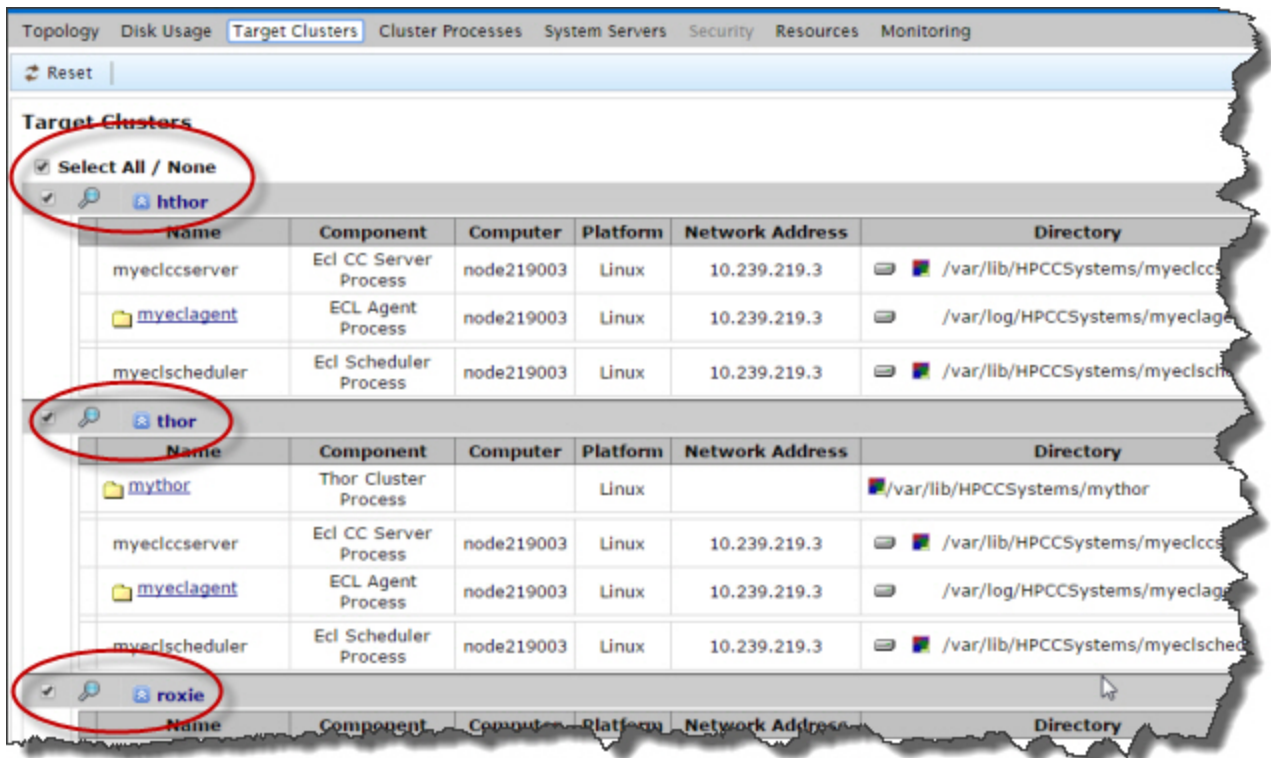


This displays all of your system Clusters.

2. Click to check on the **Select All / None** box, to select all of the clusters.

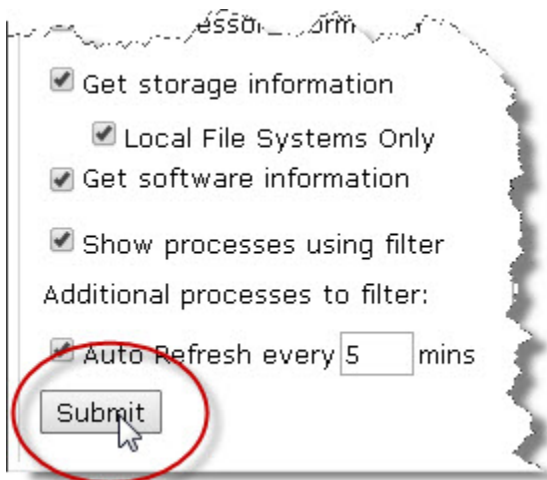
Optionally, you can just check the box(es) next to the cluster(s) you want to check.

Figure 12. Select Target Clusters



3. Press the **Submit** button at the bottom of the page to start preflight.

Figure 13. Submit



EXPECTED RESULTS:

After pressing **Submit**, a screen similar to the following should display.

Figure 14. Target Cluster Information

Location	Component	Condition	State	Roxie State	Up Time	Processes Down	/	/mnt/disk1	Physical Memory	Swap	CPU Load
10.239.219.3 /var/lib/HPCCSystems/myeclccserver	Ecl CC Server [myeclccserver]	Normal	Ready		62 day(s) 02:59:36	-	28%	99%	96%	100%	0 %
10.239.219.3 /var/lib/HPCCSystems/myeclagent	Agent Exec [myeclagent]	Normal	Ready		62 day(s) 02:59:37	-	28%	99%	96%	100%	0 %
10.239.219.3 /var/lib/HPCCSystems/myeclscheduler	Ecl Scheduler [myeclscheduler]	Normal	Ready		62 day(s) 02:59:35	-	28%	99%	96%	100%	0 %
10.239.219.5 /var/lib/HPCCSystems/mythor	Thor Slave [mythor, 1]	Normal	Ready		07:50:21	-	28%	99%	97%	100%	0 %
10.239.219.4 /var/lib/HPCCSystems/mythor	Thor Slave [mythor, 2]	Normal	Ready		07:50:21	-	28%	99%	97%	100%	0 %
10.239.219.3 /var/lib/HPCCSystems/mythor	Thor Master	Normal	Ready		07:50:21	-	28%	99%	96%	100%	0 %
10.239.219.3 /var/lib/HPCCSystems/myeclccserver	Ecl CC Server [myeclccserver]	Normal	Ready		62 day(s) 02:59:36	-	28%	99%	96%	100%	0 %
10.239.219.3 /var/lib/HPCCSystems/myeclagent	Agent Exec [myeclagent]	Normal	Ready		62 day(s) 02:59:37	-	28%	99%	96%	100%	0 %

This screen displays information on your system's component nodes. This information can help to indicate if everything is operating normally or can help to point out any potential concerns.

If there are any failed components or notable alerts, they are highlighted in orange. These alerts usually require some attention.

Figure 15. Failed Component

Location	Component	Slave Number	Condition	State	Up Time	Processes Down	/	/mnt/disk1	Physical Memory	CPU Load
10.239.219.6 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	3	Warning	Unknown		mythor...	16%	95%	90%	
10.239.219.5 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	2	Normal	Ready	04:32	-	52%	99%	97%	
10.239.219.4 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	1	Normal	Ready	04:32	-	52%	99%	96%	
10.239.219.3 /var/lib/HPCCSystems/mythor	Thor Master		Normal	Ready	04:32	-	51%	99%	97%	


☒ Select All / None
 Fetched: 11/08/12 11:28:41
 Action: Machine Information
☒ Get processor information Warn if CPU usage is over 95 %
☒ Get storage information Warn if available memory is under 5 %
☒ Local File Systems Only
☒ Get software information Warn if available disk space is under 5 %
☒ Show processes using filter
 Additional processes to filter:
☐ Auto Refresh every 0 mins.

If you have any orange fields you should examine the specified component further. It is indicative of some kind of problem or abnormality.

Preflight Thor


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

Figure 16. Cluster Processes Link



2. Click on the **mythor** link.

Figure 17. mythor link



3. Check the **Select All** checkbox (if necessary).
4. Press the **Submit** button to start preflight.

Figure 18. Submit

☒ Get storage information
☒ Local File Systems Only
☒ Get software information
☒ Show processes using filter
 Additional processes to filter:
☒ Auto Refresh every 5 mins

EXPECTED RESULTS:

After pressing Submit, a screen similar to the following should display.


Figure 19. ESP mythor system component information

Thor Cluster 'mythor'

<input checked="" type="checkbox"/>	Location	Component	Slave Number	Condition	State	Up Time	Processes Down	/	/mnt/disk
<input checked="" type="checkbox"/>	10.239.219.4 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	2	Normal	Ready	03:17:11	-	51%	99%
<input checked="" type="checkbox"/>	10.239.219.5 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	1	Normal	Ready	03:17:11	-	51%	99%
<input checked="" type="checkbox"/>	10.239.219.3 /var/lib/HPCCSystems/mythor	Thor Master		Normal	Ready	03:17:11	-	51%	99%

☒ Select All / None
 Fetched: 06/13/14 11:56:33
 Action:

This screen displays information on your Thor cluster. This information can help to indicate if everything is operating normally or can help to point out any potential concerns.



If your system has more than 1 Thor cluster, repeat these steps for each cluster.

If there are any failed nodes or notable alerts, they are highlighted in orange. The orange alerts usually require some additional attention.

Figure 20. Failed Component

Thor Cluster 'mythor'

<input checked="" type="checkbox"/>	Location	Component	Slave Number	Condition	State	Up Time	Processes Down	/	/mnt/disk1	Physic Memor
<input checked="" type="checkbox"/>	10.239.219.6 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	3	Warning	Unknown		mythor...	16%	95%	90%
<input checked="" type="checkbox"/>	10.239.219.5 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	2	Normal	Ready	04:32	-	52%	99%	97%
<input checked="" type="checkbox"/>	10.239.219.4 /var/lib/HPCCSystems/mythor	Thor Slave [mythor]	1	Normal	Ready	04:32	-	52%	99%	96%
<input checked="" type="checkbox"/>	10.239.219.3 /var/lib/HPCCSystems/mythor	Thor Master		Normal	Ready	04:32	-	51%	99%	97%

☒ Select All / None
Fetched: 11/08/12 11:28:41

Action: Machine Information ▾

☒ Get processor information Warn if CPU usage is over %

☒ Get storage information Warn if available memory is under % ▾

☒ Local File Systems Only

☒ Get software information Warn if available disk space is under % ▾

☒ Show processes using filter


Additional processes to filter:

☐ Auto Refresh every mins.

Preflight the Roxie Cluster

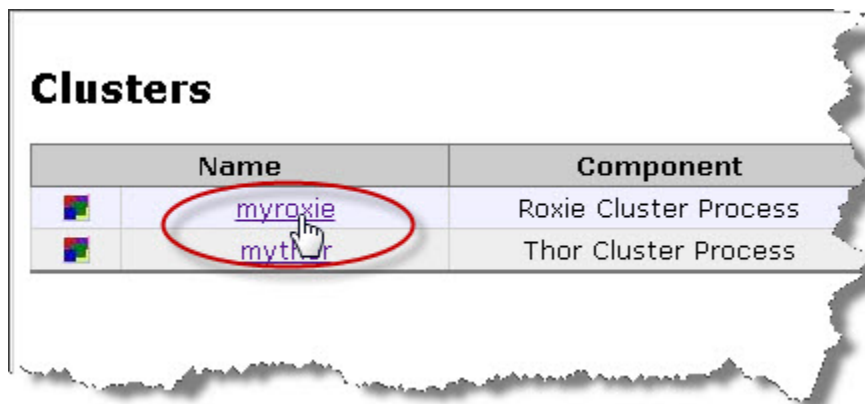
1. Click on the **Operations** icon then click on the **Cluster Processes** link.

Figure 21. Cluster Processes Link



2. Click on the **myroxie** link.

Figure 22. myroxie link



3. Press the **Submit** button to start preflight.

EXPECTED RESULTS

After pressing Submit, a screen similar to the following should display.

Figure 23. Roxie system information

Roxie Cluster 'myroxie'

<input checked="" type="checkbox"/>	Location	Component	Condition	State	Up Time	Processes Down	/
<input checked="" type="checkbox"/>	10.239.219.5 /var/lib/HPCCSystems/myroxie	Roxie Server	Normal	Ready	6 day(s) 23:27:08	-	51%
<input checked="" type="checkbox"/>	10.239.219.4 /var/lib/HPCCSystems/myroxie	Roxie Server	Normal	Ready	6 day(s) 23:27:10	-	51%

☒ **Select All / None**
Fetches: 06/13/14 12:09:27

Action: Machine Information ▾

☒ Get processor information Warn if CPU usage is over %

☒ Get storage information Warn if available memory is under % ▾

☒ Local File Systems Only

☒ Get software information Warn if available disk space is under % ▾

☒ Show processes using filter

Additional processes to filter:

☐ Auto Refresh every mins

This indicates whether the Roxie nodes are running, and some additional information about them.

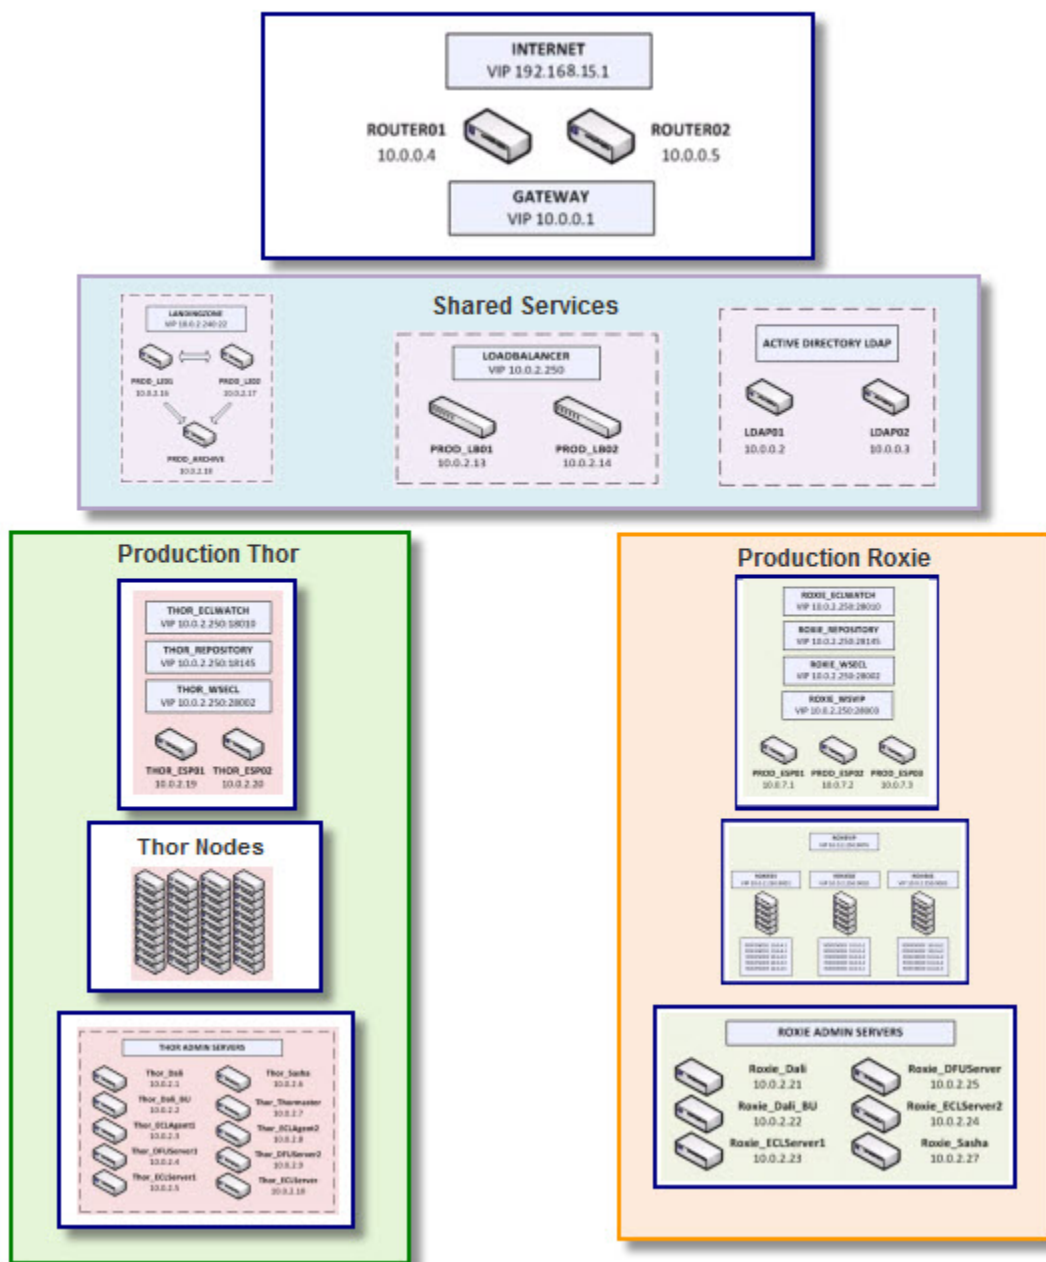


If your system has more than 1 Roxie cluster, repeat these steps for each cluster.

System Configuration and Management

The HPCC system requires configuration. The Configuration Manager tool (configmgr) included with the system software is a valuable piece of setting up your HPCC system. The Configuration Manager is a graphical tool provided that can be used to configure your system. Configuration Manager has a wizard that you can run which will easily generate an environment file to get you configured, up and running quickly. There is an advanced option available through Configuration Manager which allows for a more specific configuration, while still using the graphical interface. If desired you can edit the environment files using any xml or text editor however the file structure must remain valid.

Figure 24. Sample Production Configuration



Configuration Manager is the utility with which we configure the HPCC platform. The HPCC platform's configuration is stored in an XML file named **environment.xml**. Once you generate an environment (xml) file, it gets saved into a source directory (default is **/etc/HPCCSystems/source**). You then need to stop the system to copy it into the active HPCC directory, then distribute it into place on to each node and restart the HPCC system. At no time during configuration do you work on the live environment file.

When you install the HPCC system package, a default single-node environment.xml file is generated. After that, you can use the Configuration Manager to modify it and/or create a different environment file to configure components, or add nodes. There is a Configuration Manager wizard to help create an environment file. Give any environment file you create a descriptive name that would indicate what it is for in the source. For example, you might create an environment without a Roxie, you could call that file *environmentNoRoxie.xml*.

You would then copy the new configuration file you generate from the source directory to the **/etc/HPCCSystems** directory. Rename the file to `environment.xml`, and restart the system in order to reconfigure your system.

Configuration Manager also offers an **Advanced View** which allows more granularity for you to add instances of components or change the default settings of components for more advanced users. Even if you plan to use the Advanced View, it is a good idea to start with a wizard generated configuration file and use Advanced View to edit it.

More information and specific details for each Configuration Manager component and attributes of those components is detailed in *Using Configuration Manager*.

Running the Configuration Manager


This section will guide you through configuring an HPCC environment using the Configuration Manager.

The HPCC package should already be installed on ALL nodes.

You can use any tool or shell script you choose.

1. SSH to a node in your environment and login as a user with sudo privileges. We would suggest that it would be the first node, and that it is a support node, however that is up to your discretion.
2. Start the Configuration Manager service on the node (again we would suggest that it should be on a support node, and further that you use the same node to start the Configuration Manager every time, but this is also entirely up to you).

```
sudo /opt/HPCCSystems/sbin/configmgr
```



3. Using a Web browser, go to the Configuration Manager's interface:


```
http://<ip of installed system>:8015
```

The Configuration Manager startup wizard displays.

There are different ways to configure your HPCC system. You can use the **Generate environment wizard** and use that environment or experienced users can then use the **Advanced View** for more specific customization. There is also the option of using **Create blank environment** to generate an empty environment that you could then go in and add only the components you would want.

Environment Wizard

1. To use the wizard select the **Generate new environment using wizard** button.



2. Provide a name for the environment file.


This will then be the name of the configuration XML file. For example, we will name our environment *NewEnvironment* and this will produce a configuration XML file named *NewEnvironment.xml* that we will use.

3. Press the Next button.

Next you will need to define the IP addresses that your HPCC system will be using.

4. Enter the IP addresses.

IP Addresses can be specified individually using semi-colon delimiters. You can also specify a range of IPs using a hyphen (for example, nnn.nnn.nnn.x-y). In the image below, we specified the IP addresses 10.239.219.1 through 10.239.219.100 using the range syntax, and also a single IP 10.239.219.111.



5. Press the Next button.

Now you will define how many nodes to use for the Roxie and Thor clusters.

6. Enter the appropriate values as indicated.

HPCC Systems

Environment setup

Enter number of nodes for Roxie and Thor clusters. No Roxie/Thor cluster will be generated for zero (0) number of nodes.

Number of support nodes	0
Number of nodes for Roxie cluster	0
Number of slave nodes for Thor cluster (A Thor Master will be added to the cluster and assigned to a support node)	1
Number of Thor slaves per node (default 1)	1
Enable Roxie on demand	<input checked="" type="checkbox"/>


Cancel Back Next

Number of support nodes:	Specify the number of nodes to use for support components. The default is 1.
Number of nodes for Roxie cluster:	Specify the number of nodes to use for your Roxie cluster. Enter zero (0) if you do not want a Roxie cluster.
Number of slave nodes for Thor cluster	Specify the number of slave nodes to use in your Thor cluster. A Thor master node will be added automatically. Enter zero (0) if you do not want any Thor slaves.
Number of Thor slaves per node (default 1)	Specify the number of Thor slave processes to instantiate on each slave node. Enter zero (0) if you do not want a Thor cluster.
Enable Roxie on demand	Specify whether or not to allow queries to be run immediately on Roxie. This must be enabled to run the debugger. (Default is true)


7. Press the **Next** button

The wizard displays the configuration parameters.


8. Press the **Finish** button to accept these values or press the **Advanced View** button to edit in advanced mode.



You will now be notified that you have completed the wizard.



At this point, you have created a file named NewEnvironment.xml in the `/etc/HPCCSystems/source` directory



Keep in mind, that your HPCC configuration may be different depending on your needs. For example, you may not need a Roxie or you may need several smaller Roxie clusters. In addition, in a production [Thor] system, you would ensure that Thor and Roxie nodes are dedicated and have no other processes running on them. This document is intended to show you how to use the configuration tools. Capacity planning and system design is covered in a training module.

Distribute the Configuration

1. Stop the HPCC system.

If it is running stop the HPCC system (on every node), using a command such as this:

```
sudo /sbin/service hpcc-init stop
```

Note: You may have a multi-node system and a custom script such as the one illustrated in Appendix of the [Installing and Running the HPCC Platform](#) document to start and stop your system. If that is the case please use the appropriate command for stopping your system on every node.



Be sure HPCC is stopped before attempting to copy the environment.xml file.

2. Back up the original environment.xml file.

```
# For example
sudo -u hpcc cp /etc/HPCCSystems/environment.xml /etc/HPCCSystems/source/environment-date.xml
```

Note: The live environment.xml file is located in your **/etc/HPCCSystems/** directory. Configuration Manager works on files in **/etc/HPCCSystems/source** directory. You must copy from this location to make an environment.xml file active.

You can also choose to give the environment file a more descriptive name, to help differentiate any differences.

Having environment files under source control is a good way to archive your environment settings.

3. Copy the new .xml file from the source directory to the /etc/HPCCSystems and rename the file to *environment.xml*

```
# for example
sudo -u hpcc cp /etc/HPCCSystems/source/NewEnvironment.xml /etc/HPCCSystems/environment.xml
```

4. Copy the **/etc/HPCCSystems/environment.xml** to the **/etc/HPCCSystems/** on to *every* node.

You may want to use a script to push out the XML file to all nodes. See the *Example Scripts* section in the Appendix of the [Installing and Running the HPCC Platform](#) document. You can use the scripts as a model to create your own script to copy the environment.xml file out to all your nodes.

5. Restart the HPCC platform on all nodes.

Environment.conf

Another component of HPCC system configuration is the environment.conf file. Environment.conf contains some global definitions that the configuration manager uses to configure the HPCC system. In most cases, the defaults are sufficient.



WARNING: These settings are essential to proper system operation. Only expert level HPCC administrators should attempt to change any aspects of this file.

By default the environment.conf file is located:

```
/etc/HPCCSystems
```

Environment.conf is required upon startup of HPCC. The environment.conf is where the HPCC environment file is defined.

```
/opt/HPCCSystems/environment.xml
```

This is also where the working path is defined.

```
path=/opt/HPCCSystems
```

The working path is used by several aspects of the application, changing this could cause needless complications. By default the application installs there, and sets many resources to that as well.

The default environment.conf:

```
## HPCC Systems default environment configuration file

[DEFAULT SETTINGS]
configs=/etc/HPCCSystems
path=/opt/HPCCSystems
classpath=/opt/HPCCSystems/classes
runtime=/var/lib/HPCCSystems
lock=/var/lock/HPCCSystems
# Supported logging fields: AUD,CLS,DET,MID,TIM,DAT,PID,TID,NOD,JOB,USE,SES,
#                          COD,MLT,MCT,NNT,COM,QUO,PFX,ALL,STD
logfields=TIM+DAT+MLT+MID+PID+TID+COD+QUO+PFX
pid=/var/run/HPCCSystems
log=/var/log/HPCCSystems
user=hpcc
group=hpcc
home=/Users
environment=environment.xml
sourcedir=/etc/HPCCSystems/source
blockname=HPCCSystems
interface=*
# enable epoll method for notification events (true/false)
use_epoll=true
```

Path considerations

Most of the directories are defined as absolute paths:

```
configs=/etc/HPCCSystems
path=/opt/HPCCSystems
classpath=/opt/HPCCSystems/classes
runtime=/var/lib/HPCCSystems
```



```
lock=/var/lock/HPCCSystems
```

HPCC will not run properly without the proper paths, and in some cases needs the absolute path. If a process or component can't find a path you will get an error message such as the following:

```
"There are no components configured to run on the node..."
```

If the path changes from HPCCSystems, it does NOT change in the environment.xml file. Any changes would require manually modifying the environment.xml file.

The log file, *hpcc-init.log* is written to the HPCCSystems path.

Other Environment.conf items

Some other items used by or referred to in environment.conf.

Use_epoll It is an event mechanism to achieve better performance in more demanding applications where number of watched file descriptors is large.

Logfields Categories available to be logged. These consist of Time(TIM), Date(DAT), Process ID (PID), Thread ID (TID), etc.

Interface In the default environment.conf there is a value for interface. The default value for that is:

```
interface=*
```

The default value of * assigns the interface to an open ip address, in any order. Specifying an interface, such as Eth0, will assign the specified node as the primary.

Remote Access over TLS

Configuring your system for remote file access over Transport Layer Security (TLS) requires modifying the **dafilesrv** setting in the *environment.conf* file.

To do this either uncomment (if they are already there), or add the following lines to the *environment.conf* file. Then set the values as appropriate for your system.

```
#enable SSL for dafilesrv remote file access
dfsUseSSL=true
dfsSSLCertFile=/certfilepath/certfile
dfsSSLPrivateKeyFile=/keyfilepath/keyfile
```

Set the **dfsUseSSL=true** and set the value for the paths to point to the certificate and key file paths on your system. Then deploy the *environment.conf* file (and cert/key files) to all nodes as appropriate.

When dafilesrv is enabled for TLS (port 7600), it can still connect over a non-TLS connection (port 7100) to allow legacy clients to work.

Configuring HPCC for Authentication

This section details the steps to configure your HPCC platform to use authentication. There are currently a few ways to use authentication with your HPCC system: simple htpasswd authentication, LDAP, or another plug-in security method.

The htpasswd authentication method is basic password authentication. It only grants or denies access to a user, based upon MD5 encrypted password authentication.

LDAP authentication offers more features and options. LDAP can not only authenticate users, but adds granularity to the authentication. LDAP allows you to control grouped access to features, functions, and files.

You should consider your system needs and decide which of these methods is appropriate for your environment.



When implementing any form of authentication, we strongly recommend that you enable your ESP server to use HTTPS (SSL) and set ALL service bindings to only use HTTPS. This ensures that credentials are passed over the network using SSL encryption. See *Configuring ESP Server to use HTTPS (SSL)* for details.

You should not attempt this until you have already deployed, configured, and certified the environment you will use.

Using htpasswd authentication

htpasswd provides basic password authentication to the entire system. This section contains the information to install and implement htpasswd authentication.

Connect to Configuration Manager

In order to change the configuration for HPCC components, connect to the Configuration Manager.

1. Stop all HPCC Components, if they are running.
2. Verify that they are stopped. You can use a single command, such as :

```
sudo /opt/HPCCSystems/sbin/hpcc-run.sh -a hpcc-init status
```

3. Start Configuration Manager.

```
sudo /opt/HPCCSystems/sbin/configmgr
```

4. Connect your web browser to the Configuration Manager web interface.

(using the url of `http://<configmgr_IP_Address>:8015`, where `<configmgr_IP_Address>` is the IP address of the node running Configuration Manager)

5. Select the **Advanced View** radio button.
6. Use the drop list to select the XML configuration file.

Note: Configuration Manager **never** works on the active configuration file. After you finish editing you will have to copy the `environment.xml` to the active location and push it out to all nodes.

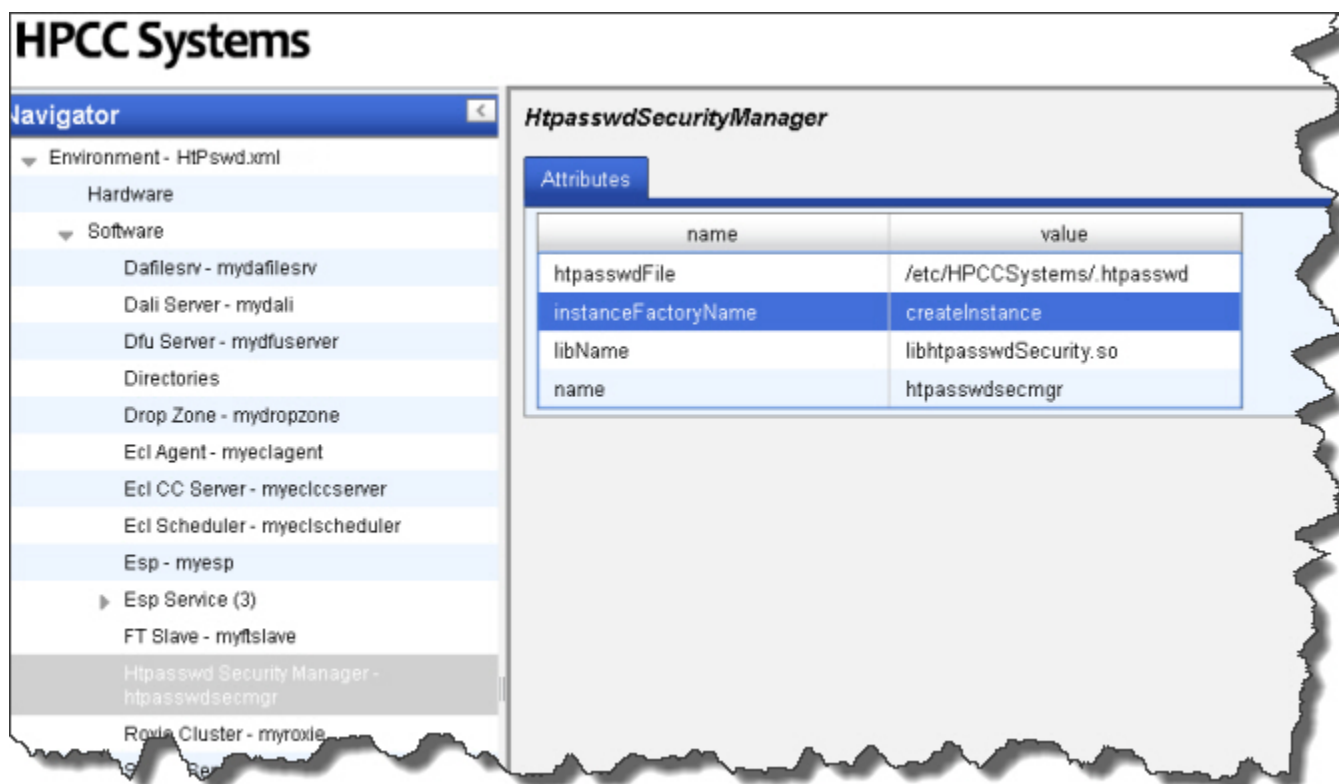
7. Check the **Write Access** box.

Default access is read-only. Many options are only available when write-access is enabled.

Enabling htpasswd authentication in HPCC

8. Create an instance of the **Security Manager** Plug-in:
 - a. Right-click on Navigator Pane on the left side.
 - b. Select **New Components**
 - c. Select the **htpasswdsecmgr** component
9. Configure the htpasswd plug-in

Figure 25. Security Mgr Configuration page




- a. Enter the location of the Htpasswd file containing the username and password on the Linux file system for the value of **htpasswdFile**
 - b. **InstanceFactoryName** is the name of the security manager factory function, implemented in the security library. The default is "createInstance". For implementing Htpasswd, leave the default.
 - c. Provide a library name value for **libName**. For htpasswd, use [libhtpasswdSecurity.so](#)
 - d. Provide an instance **name** for the name value. For example, [htpasswdsecmgr](#).
10. Select **Esp - myesp** in the Navigator panel on the left hand side.

Note: If you have more than one ESP Server, you would only use one of them for authentication.

11. Associate the Security Manager Plug-in with the ESP binding(s)

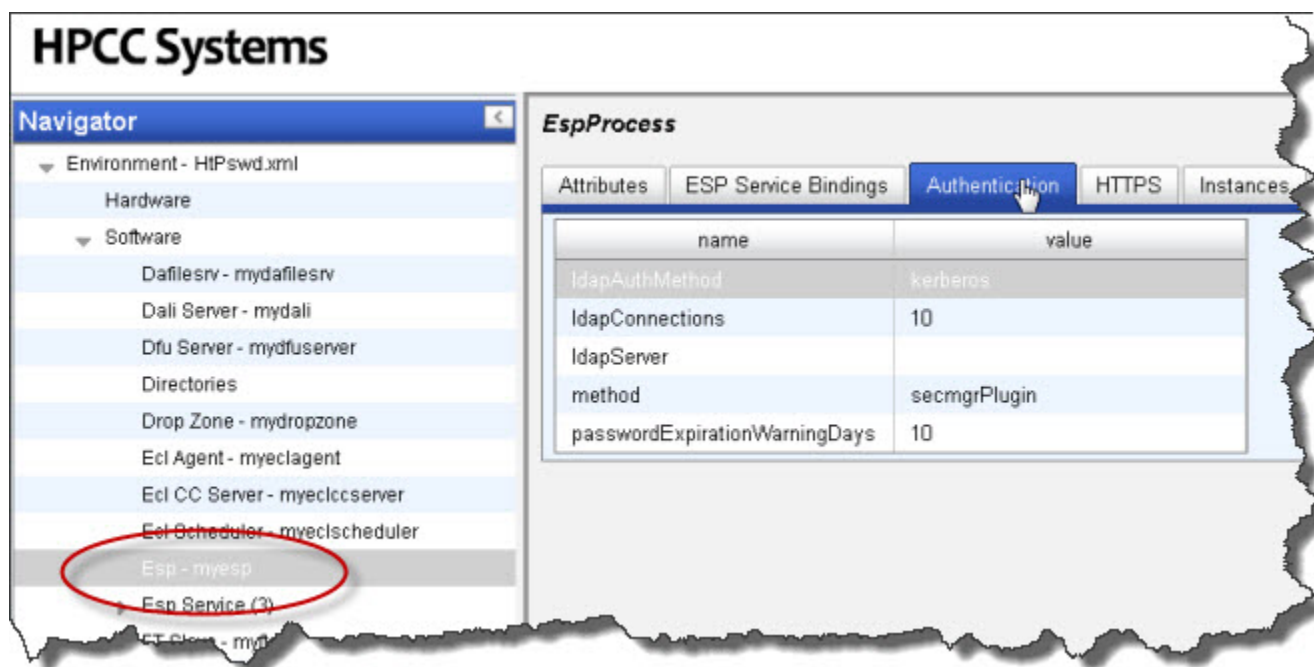
- Click on the target **Esp** in the Navigator Pane on the left side.
- Select the **ESP Service bindings** tab
- On the target binding(s) select the appropriate securityMgrPlugin instance from the drop list.



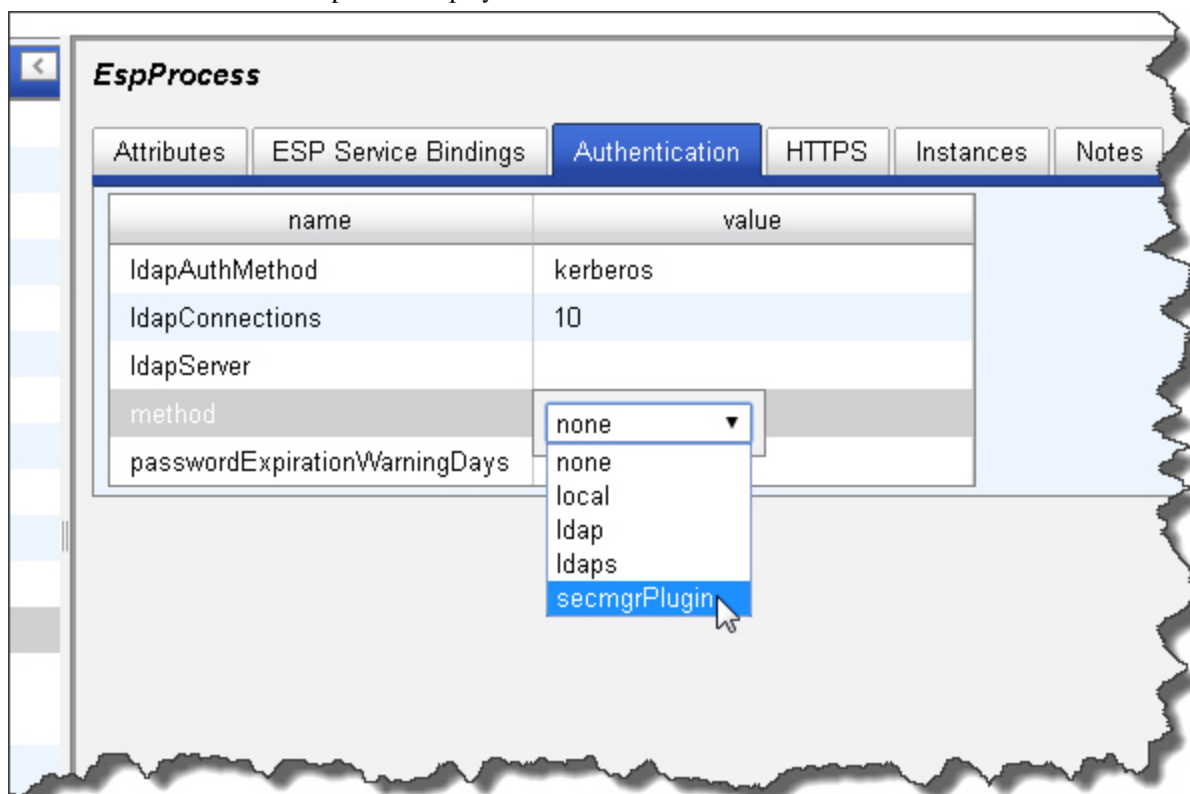
12. Select the security Plugin for each service that requires a security manager.

For example, in the above image, select [httpasswdsecmgr](#) for the smc service. Then, select it for ws_ecl and every other service that you want to use httpassword security.

13. Select the **Authentication** tab



14. Click on the value column drop list to display the choices for **method**.



15. Choose **secmgrPlugin** from the drop list.

16. Click on the disk icon to save.

User administration with htpasswd

Users and passwords are kept in the htpasswd file. The htpasswd file must exist on the ESP Node where you have enabled authentication. HPCC only recognizes MD5 encrypted passwords.

The default location is: **/etc/HPCCSystems/.htpasswd** on the ESP node that has been configured to authenticate, but it is configurable from the Htpasswd Security Manager as outlined above (step 9).

You can use the htpasswd utility to create the .htpasswd file to administer users.

You may already have the htpasswd utility on your system, as it is a part of some Linux distributions. Check your Linux distribution to see if you already have it. If you do not have it you should download the utility for your distribution from The Apache Software Foundation.

For more information about using htpasswd see: <http://httpd.apache.org/docs/2.2/programs/htpasswd.html>.

Using LDAP Authentication

This section contains the information to install and implement LDAP based authentication. LDAP Authentication provides the most options for securing your system, or parts of your system. In addition to these configuration settings you should run the **initldap** utility to create the required default HPCC Admin user on your LDAP server.

If you choose to use LDAP authentication you must enable LDAP security in your HPCC System configuration. With LDAP security enabled on your system you can then choose to enable file scope security. You can choose to use LDAP authentication without enabling file scope security. The following sections describe how to enable LDAP authentication and file scope security for your HPCC system.

Connect to Configuration Manager

In order to change the configuration for HPCC components, connect to the Configuration Manager.

1. Stop all HPCC Components, if they are running.
2. Verify that they are stopped. You can use a single command, such as :

```
sudo /opt/HPCCSystems/sbin/hpcc-run.sh -a hpcc-init status
```

3. Start Configuration Manager.

```
sudo /opt/HPCCSystems/sbin/configmgr
```

4. Connect to the Configuration Manager web interface.

(using the url of `http://<configmgr_IP_Address>:8015`, where `<configmgr_IP_Address>` is the IP address of the node running Configuration Manager)

5. Select the **Advanced View** radio button.
6. Use the drop list to select the XML configuration file.

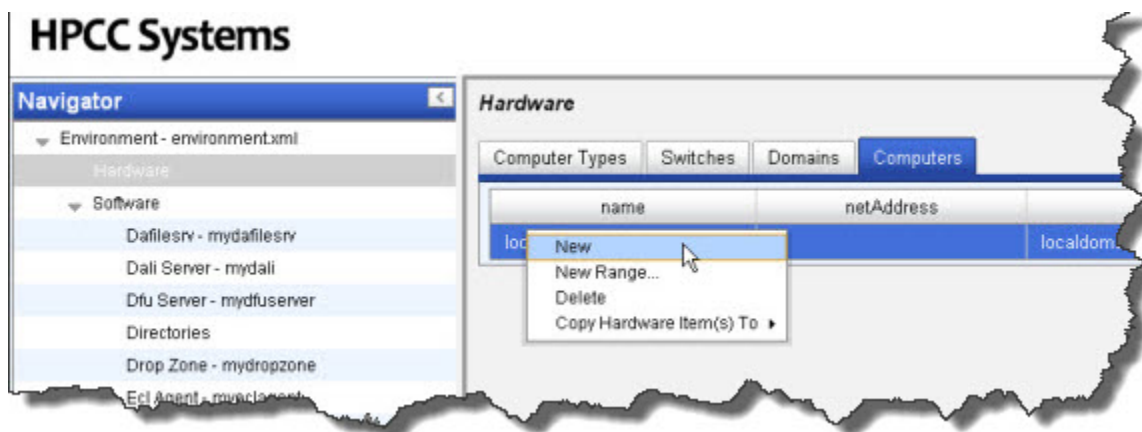
Note: Configuration Manager **never** works on the active configuration file. After you finish editing you will have to copy the `environment.xml` to the active location and push it out to all nodes.

Modifying the configuration

Follow the steps below to modify your configuration.

1. Check the box for **Write Access**.
2. From the **Navigator** pane, select **Hardware**.
3. Select the **Computers** tab from the panel on the right.

4. Right-click on the table below computers and select **New** from the pop up menu.



The **Add New Computers** dialog displays.

5. Fill in the values for the **Computer Attributes**

The 'Add New Computers' dialog box is shown. It has a title bar with 'Add New Computers' and a close button. The dialog is divided into two main sections. The first section, 'Computer Attributes', contains three rows: 'Name Prefix' with a text input field containing 'ldap', 'Domain' with a dropdown menu showing 'localdomain', and 'Type' with a dropdown menu showing 'linuxmachine'. The second section, 'IP address/range', contains a 'Range' checkbox (unchecked), a 'Start IP Address' text input field, and a 'Stop IP Address' text input field. At the bottom of the dialog are 'Ok' and 'Cancel' buttons.

- a. Provide a **Name Prefix**, for example: **ldap**.

This helps you to identify it in the list of computers.

- b. Fill in **Domain** and **Type** with the values of your domain name, as well as the types of machines you are using.

In the example above, **Domain** is **localdomain**, and the **Type** is **linuxmachine**. These should correspond to your domain and type.

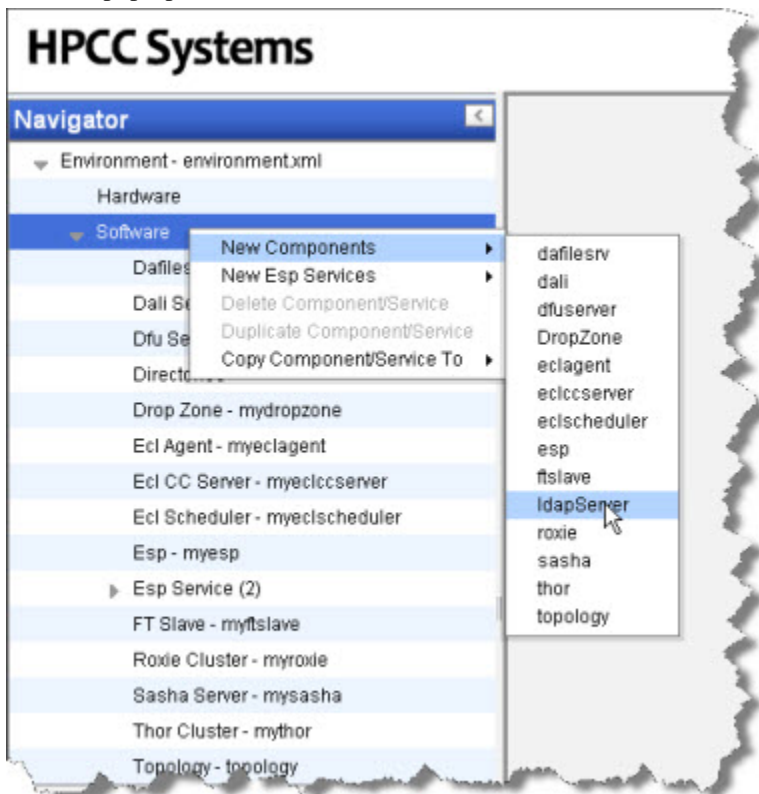
If you need to add a new domain or machine type to your system to be able to define an existing LDAP server, you should set these up first in the other two tabs in the hardware section.

- c. Add the IP address as appropriate for the LDAP server.
- d. Press the **Ok** button.
- e. Click on the disk icon to save.

Adding the ldapServer component

After the LDAP Server node has been added to the Hardware configuration, configure the Software LDAP server definition.

1. Right-click on **Navigator** Pane and choose **New Components** from the pop-up menu, then choose **ldapServer** from the pop-up menu.



Note: The ldapServer component is merely a definition that specifies an existing LDAP server. It does not install one.

2. Fill in the **LDAP Server Process** properties:

- a. On the **Instances** tab, Right-click on the table on the right hand side, choose **Add Instances...**



The **Select computers** dialog appears.

- b. Select the computer to use by checking the box next to it.



This is the computer you added in the **Hardware / Add New Computers** portion earlier.


- c. Press the **Ok** button.

- d. Fill in the **Attributes** tab with the appropriate settings from your existing LDAP Server.

The screenshot shows a configuration window titled "LDAPServerProcess". It has three tabs: "Attributes" (selected), "Instances", and "Notes". Below the tabs is a table with two columns: "name" and "value". The table contains the following entries:

name	value
cacheTimeout	5
description	LDAP server process
filesBasedn	ou=files,ou=ecl
groupsBasedn	ou=groups,ou=ecl
ldapPort	389
ldapSecurePort	636
modulesBasedn	ou=modules,ou=ecl
name	ldapserver
serverType	ActiveDirectory
sudoersBasedn	ou=SUDOers
systemBasedn	cn=Users,ou=ecl
systemCommonName	TheAdmin
systemPassword	*****
systemUser	TheAdmin
usersBasedn	ou=users,ou=ecl
workunitsBasedn	ou=workunits,ou=ecl

- e. Choose the LDAP server type from the serverType attribute drop box.



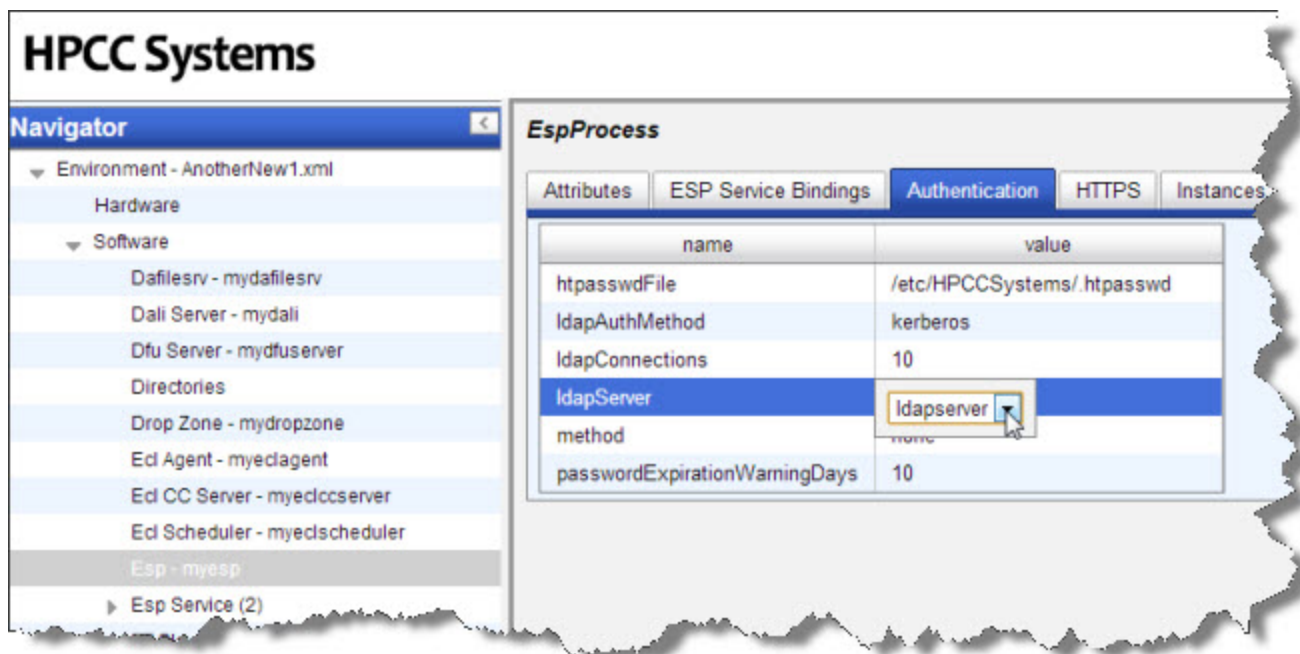
NOTE: Support for OpenLDAP has been deprecated. The option is included only for legacy purposes.

- f. Click on the disk icon to save.

Note: The **cacheTimeout** value is the number of minutes that permissions are cached in ESP. If you change any permissions in LDAP, the new settings will not take effect until ESP and Dali refresh the permissions. This could take as long as the cacheTimeout. Setting this to 0 means no cache, but this has performance overhead so it should not be used in production.

3. In the Navigator pane, click on **ESP – myesp**

4. On the **EspProcess** page on the right hand side, select the **Authentication** tab.




Fill in the appropriate values:

- Change the **ldapAuthMethod** to [kerberos](#).
- Change the **ldapConnections** to the number appropriate for your system (10 is for example only, may not be necessary in your environment).
- Select the **ldapServer** component that you added earlier from the drop list, for example: [Idapserver](#).
- Change the **method** value to [ldap](#).
- Select the ESP Service bindings tab. Verify that your LDAP settings appear in the **resourcesBasedn** and **workunitsBasedn**
- Click on the disk icon to save.

5. To enable the file scope permissions, configure the file scope security for the Dali Server.

In the Navigator pane, click on the **Dali Server – mydali**



Fill in the values as appropriate:

- Select the **LDAP** tab.
- Change the **authMethod** to **simple**
- Set the **checkScopeScans** value to **true**.

Only set this value to true when you want file scope security enabled. Security settings can have three states.

- None, no authentication and no file scope security.
- LDAP security for authentication only, without enabling file scope security.
- LDAP authentication and file scope security enabled.

- Change the LDAP values as appropriate to match the settings in your LDAP server component in configuration manager.

For example, change the **ldapServer** to the value you gave your LDAP Server, in our example it is: *ldapservers*.


Confirm the change when prompted.

The **filesDefaultUser** is an LDAP account used to access files when no user credentials are supplied. This is similar to a guest account, so it should be an account with **very** limited access, if used at all. To disable access without credentials, leave **filesDefaultUser** blank.

The **filesDefaultPassword** is the password for that account.

- Click on the disk icon to save.

6. In the Navigator pane, click on the **Roxie Cluster – myroxie**



- On the **RoxieCluster** page on the right hand side, select the **LDAP** tab.
- Locate the **ldapUser** field and verify that there is a valid HPCC user who is a member of the Authenticated Users group on your LDAP server. For example, the "roxie" user assumes that the "roxie" user is a valid HPCC authenticated user.
- Add the password security for Roxie by adding it to the **ldapPassword** field on the same tab.

In order to run Roxie queries with File Scope security, ensure that a Roxie user is created in the list of authenticated users.

In the following section, *Adding and editing users*, add the *roxie* user and make sure that password is the same as the one entered in Configuration Manager.

Installing the Default Admin user

After enabling your configuration for LDAP security, you must copy your environment file to the /etc/HPCCSystems directory. See the section *Configuring a Multi-Node System* for more info about configuring your system. With the correct environment.xml file in place, you must then run the **initldap** utility that initializes the security components and the default users.

The initldap Utility

The **initldap** utility creates the HPCC Administrator's user account and the HPCC OUs for a newly defined LDAP server. The **initldap** utility extracts these settings from the LDAPServer component(s) in the environment.xml bound to the configured ESPs.

You run the **initldap** utility once you complete your configuration with LDAP components enabled and have distributed your environment.xml file to all nodes.

```
sudo /opt/HPCCSystems/bin/initldap
```

The **initldap** utility prompts you for LDAP Administrator credentials. Enter the appropriate values when prompted.

The following example of `initldap` for a `389DirectoryServer` deployment.

```
Enter the '389DirectoryServer' LDAP Admin User name on '10.123.456.78'...Directory Manager
Enter the LDAP Admin user 'Directory Manager' password...*****

Ready to initialize HPCC LDAP Environment, using the following settings
  LDAP Server      : 10.123.456.78
  LDAP Type        : 389DirectoryServer
  HPCC Admin User  : HPCCAdmin389
Proceed? y/n
```

Using the `addScopes` tool

When a new ESP user account is created, a private “`hpccinternal::<user>`” file scope is also created granting new users full access to that scope and restricting access to other users. This file scope is used to store temporary HPCC files such as spill files and temp files.

If you are enabling LDAP file scope security and already have user accounts, you should run the `addScopes` utility program to create the `hpccinternal::<user>` scope for those existing users.

Users which already have this scope defined are ignored and so it can be used on both new and legacy ESP user accounts safely.

The tool is located in the `/opt/HPCCSystems/bin/` folder and to run it you must pass the location of **`daliconf.xml`**, for example:

```
/opt/HPCCSystems/bin/addScopes /var/lib/HPCCSystems/mydali/daliconf.xml
```

User Security Maintenance

Configuring an HPCC System to use Active Directory or LDAP-based security allows you to set permissions to control access to Features, File Scopes, and Workunit Scopes.

Introduction

HPCC systems[®] maintains security in a number of ways. HPCC Systems[®] can be configured to manage users' security rights by pointing either at Microsoft's Active Directory on a Windows system, or a 389Directory Server on Linux systems.

Using the Permissions interface in ECL Watch, administrators can control access to features in ECL IDE, ECL Watch, ECL Plus, DFU Plus, and the ECL modules within the Attribute Repository. Optionally, you can also implement file and workunit access control by enabling that setting in the Dali server.

Establish permissions by group or by user and define them by association with a particular feature of the HPCC System. Permissions can be defined for each unique combination of group and feature. Permissions are separated into the following categories:


Esp Features for SMC	Controls access to features in ECL Watch and similar features accessed from ECL IDE.
Esp Features for WsEclAccess	Controls access to the WS-ECL web service
Esp Features for EclDirectAccess	Controls access to the ECLDirect web service
File Scopes	Controls access to data files by applying permissions to File scopes
Workunit Scopes	Controls access to Workunits by applying permissions to Workunit scopes
Repository Modules	Controls access to the Attribute Repository and Modules in the repository (legacy)

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 System where permissions may be set. Use this area to view the permissions currently set for any area of the HPCC System, 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 the **Logged In As:** link at the top of the ECL Watch page.



1. Click on the **Logged In As:** link.

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. Inside, there is a 'Save' button at the top left. The user's name 'FranklinX' is displayed prominently. Below it, the following fields are shown: 'Username: FranklinX', 'First Name: Franklin', 'Last Name: Xavier', 'Old Password:' with an empty input field, 'New Password:' with an empty input field, 'Confirm Password:' with an empty input field, and 'Password Expiration: 2016-06-26'. The dialog is set against a background of a torn paper effect.

2. Confirm the User Name that you are logged in as.

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.

You can also change your password here, if desired.

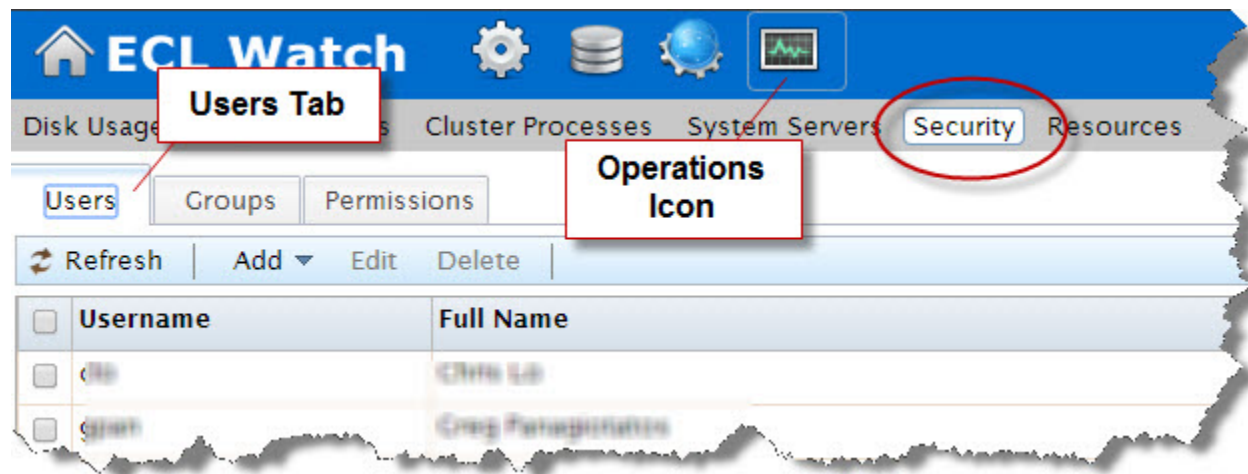
Setting and modifying user permissions

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 System 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:

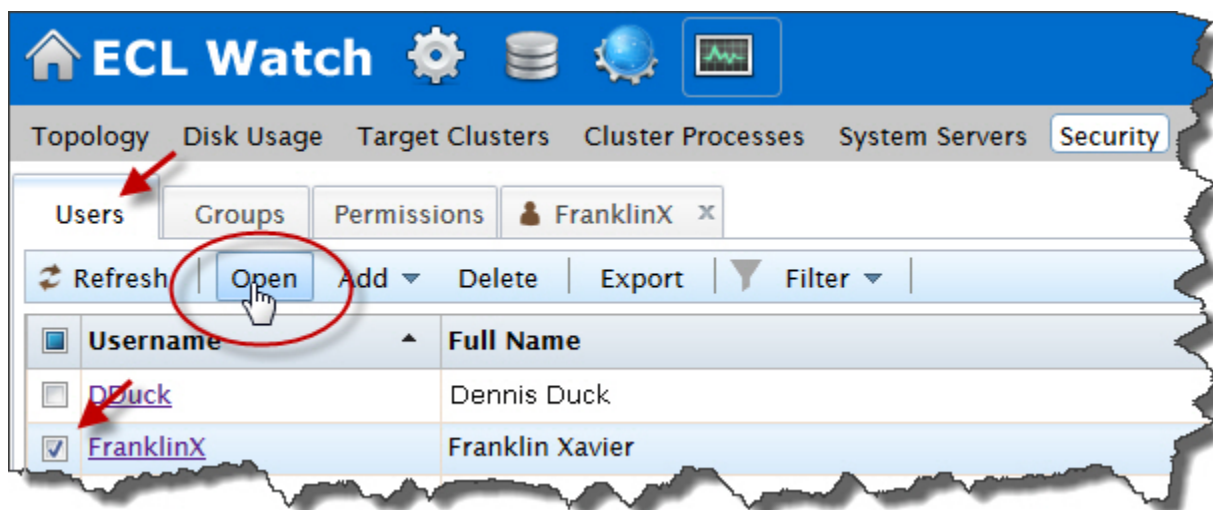
Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

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:

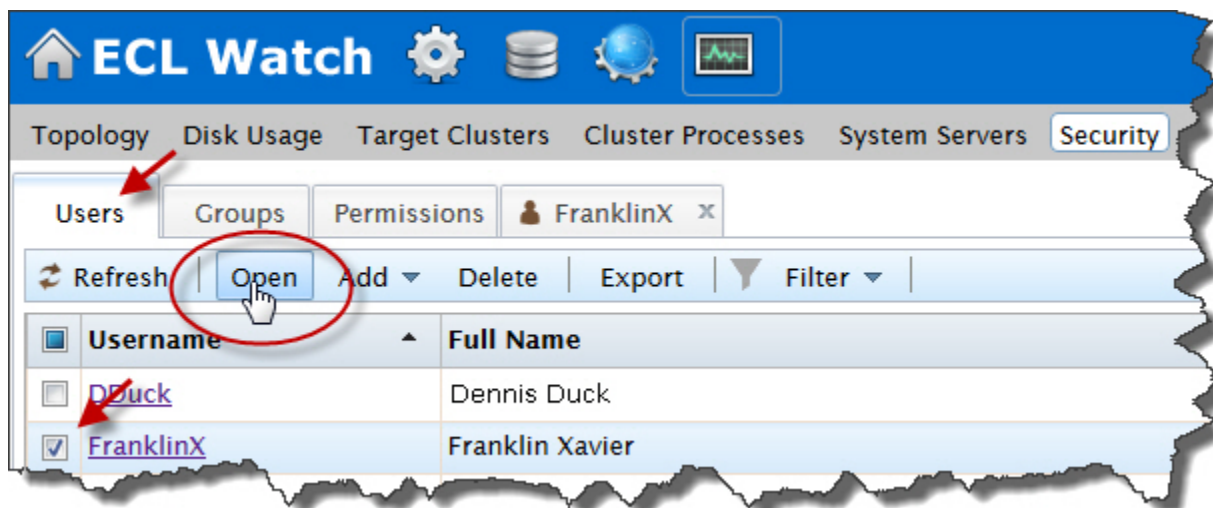
Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

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.

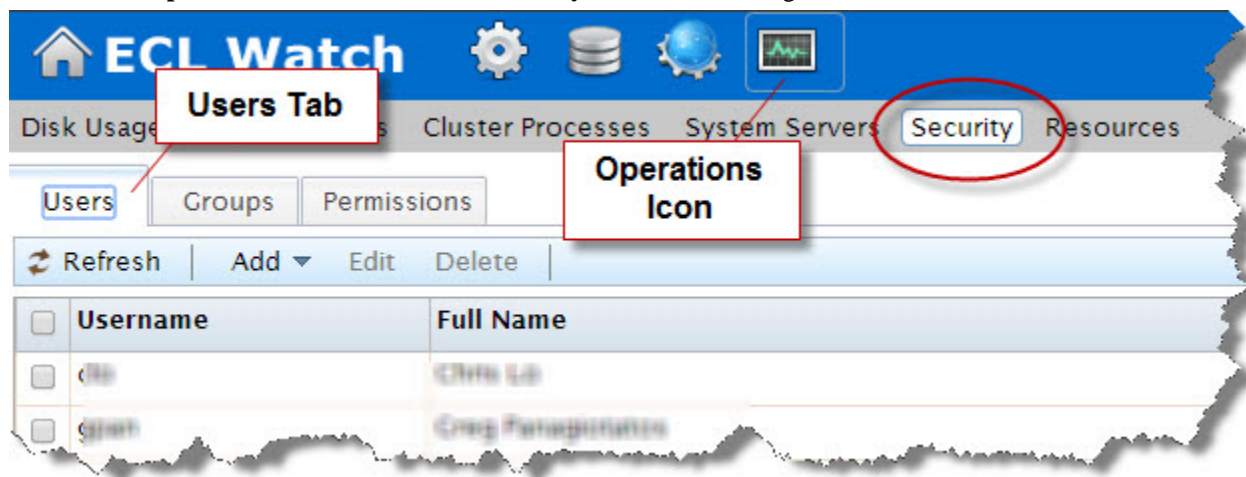
To add the user to the group, check the box next to the desired group.

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

To promote a user to an Administrator

To modify a users credentials you must have Administrator level access. To promote a user to an HPCC Administrator, add the user to the **Administrators** group.

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

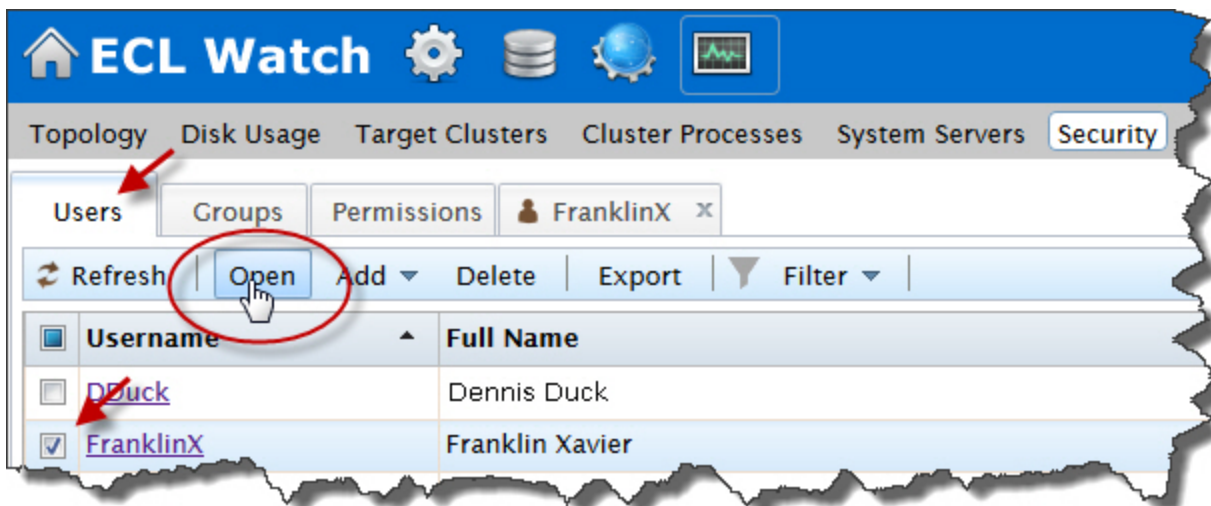


1. Click on the **Users** tab.

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.



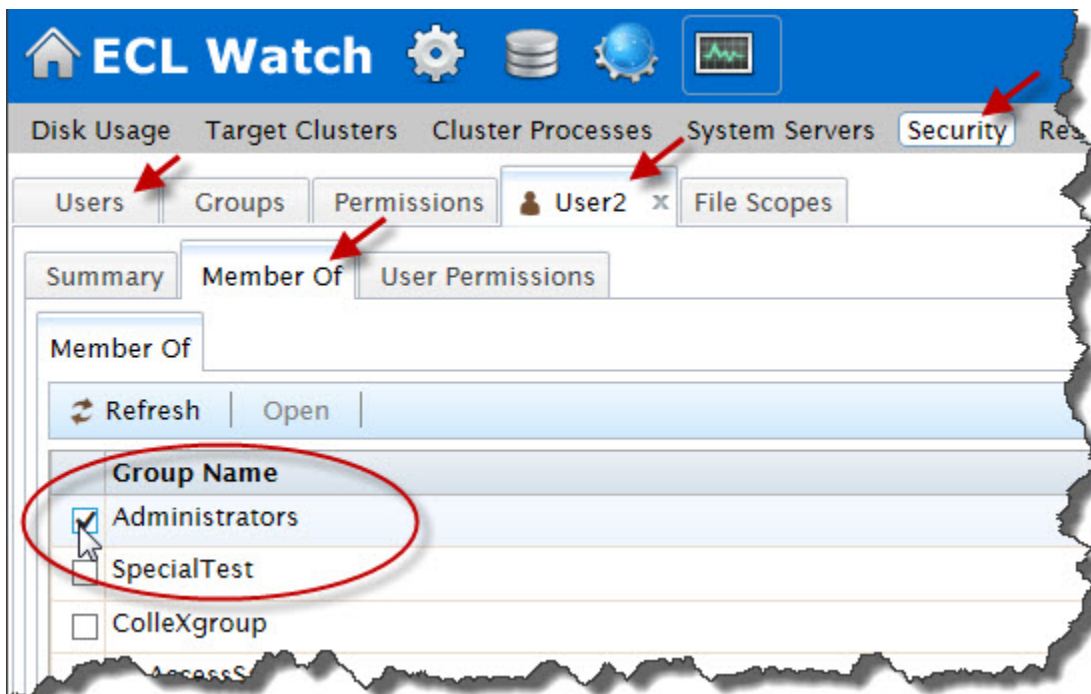
A tab opens for each user selected. On each user's tab there are several sub-tabs.

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 Administrator group could vary. For example, in Active Directory, it is "Administrators", in OpenLDAP it is "Directory Administrators".

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.

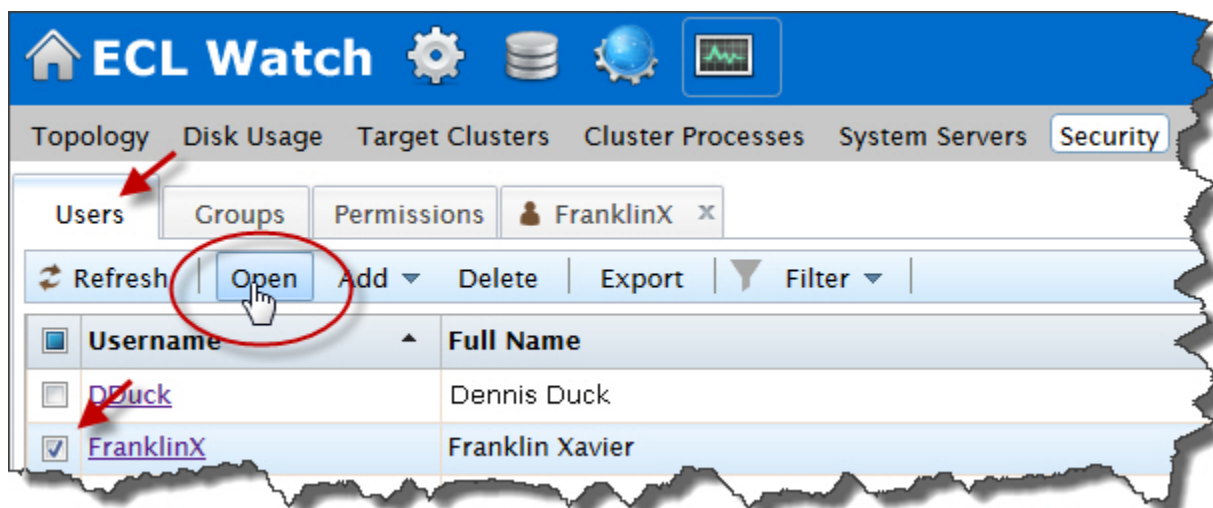
Click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

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 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.

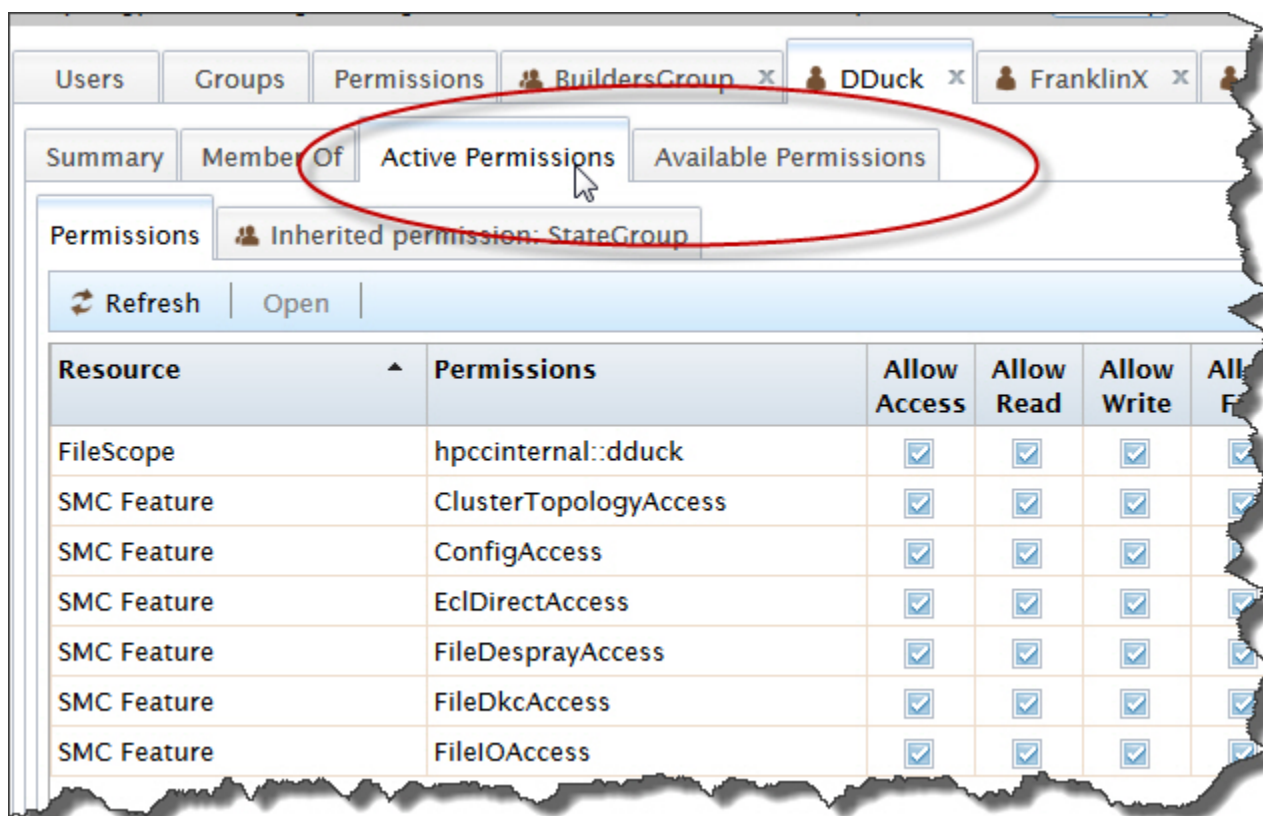
The users display in a list.

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.

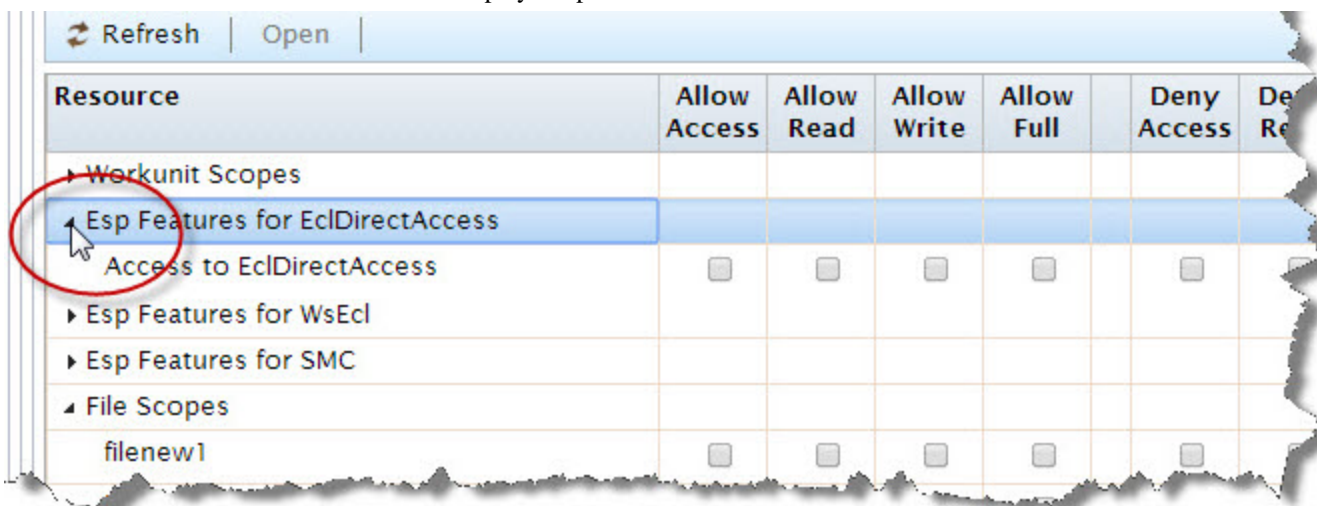


Click on the **Active Permissions** sub-tab to view the user's current permissions.

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.

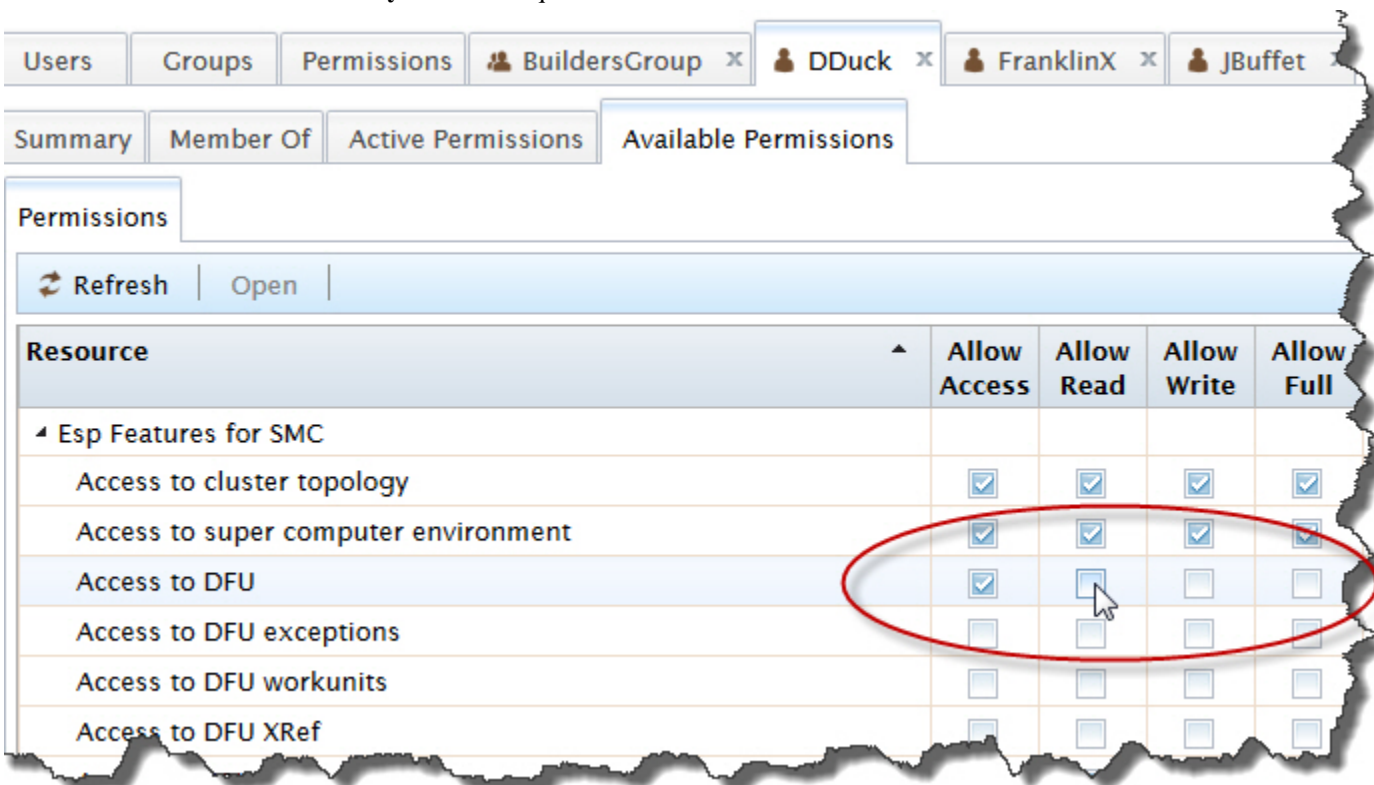
5. Click on the arrow next to the resource to display the permissions that can be set for that resource.



Resource	Allow Access	Allow Read	Allow Write	Allow Full	Deny Access	Deny Read	Deny Write	Deny Full
Workunit Scopes								
↳ Esp Features for EclDirectAccess								
Access to EclDirectAccess	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>			
↳ Esp Features for WsEcl								
↳ Esp Features for SMC								
File Scopes								
filenew1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>			

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.



Resource	Allow Access	Allow Read	Allow Write	Allow Full
↳ 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 checked="" type="checkbox"/>
Access to DFU	<input checked="" 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. 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 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.

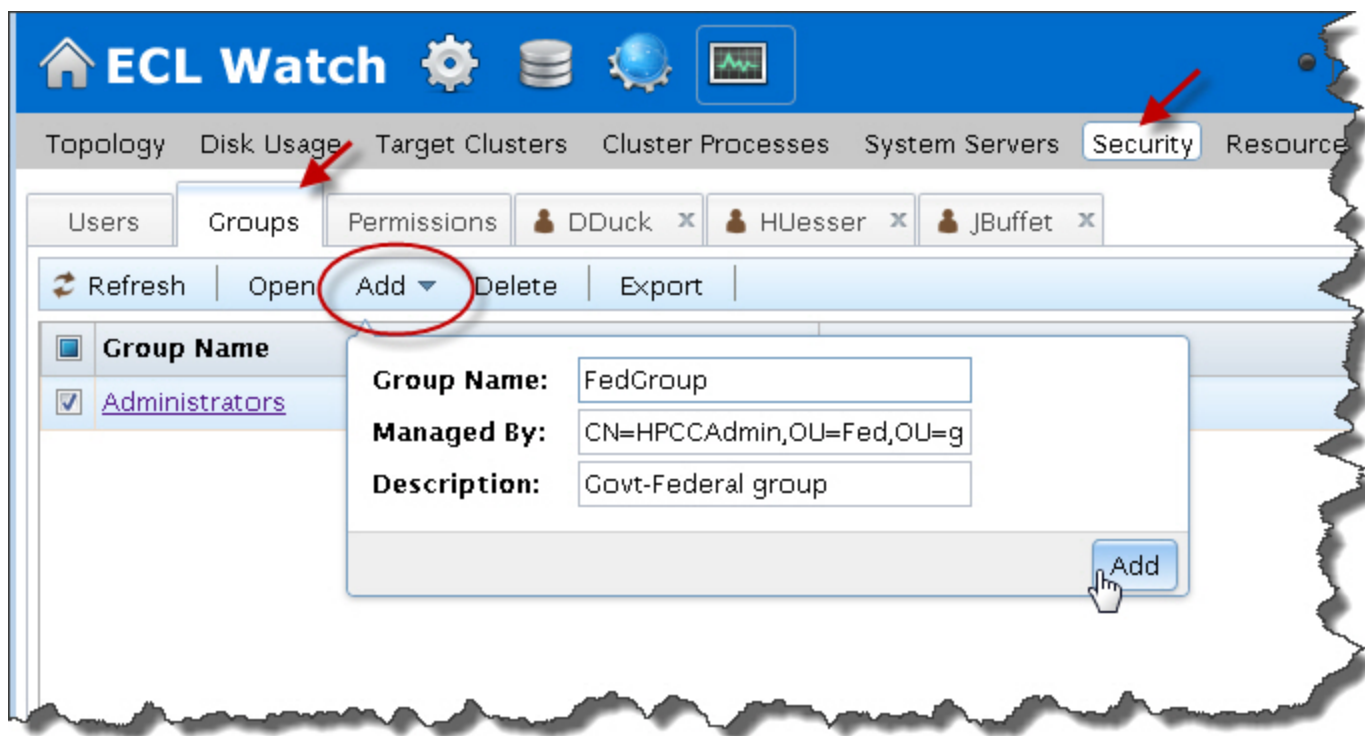


To modify groups, click on the **Operations** icon, then click the **Security** link from the navigation sub-menu. Click on the **Groups** tab.

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.

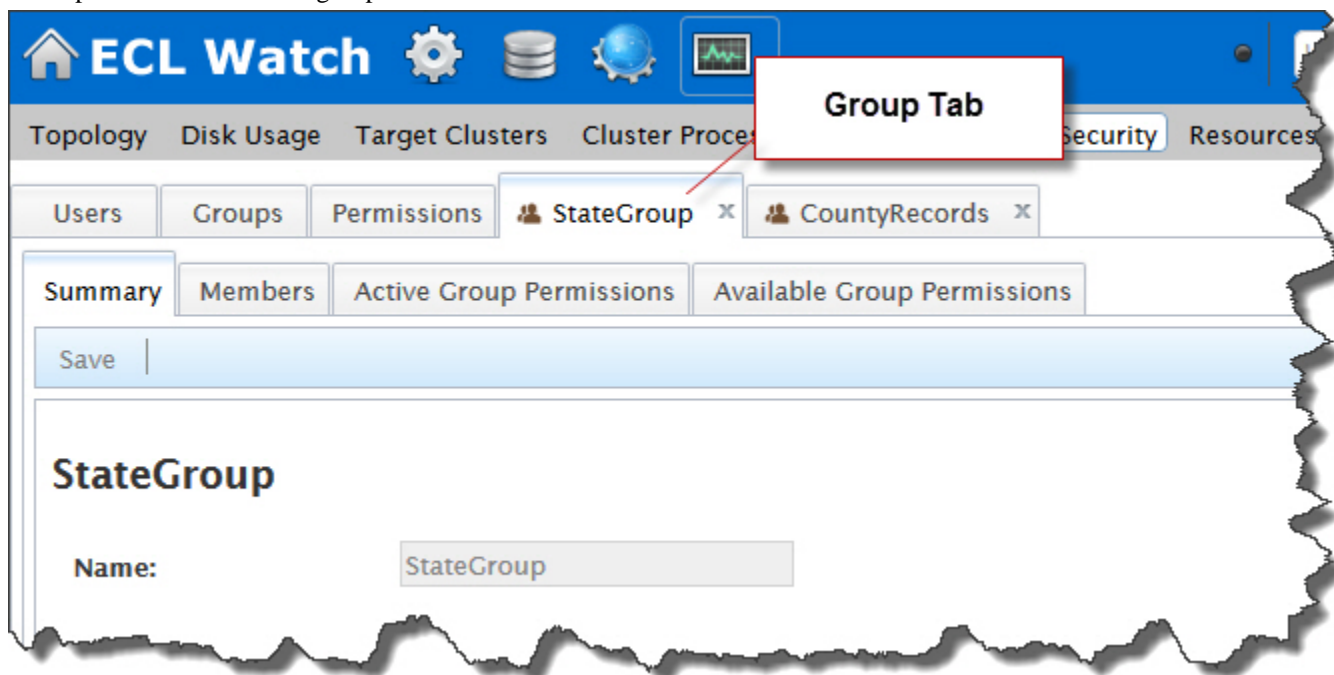
To set permissions for a group:

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.

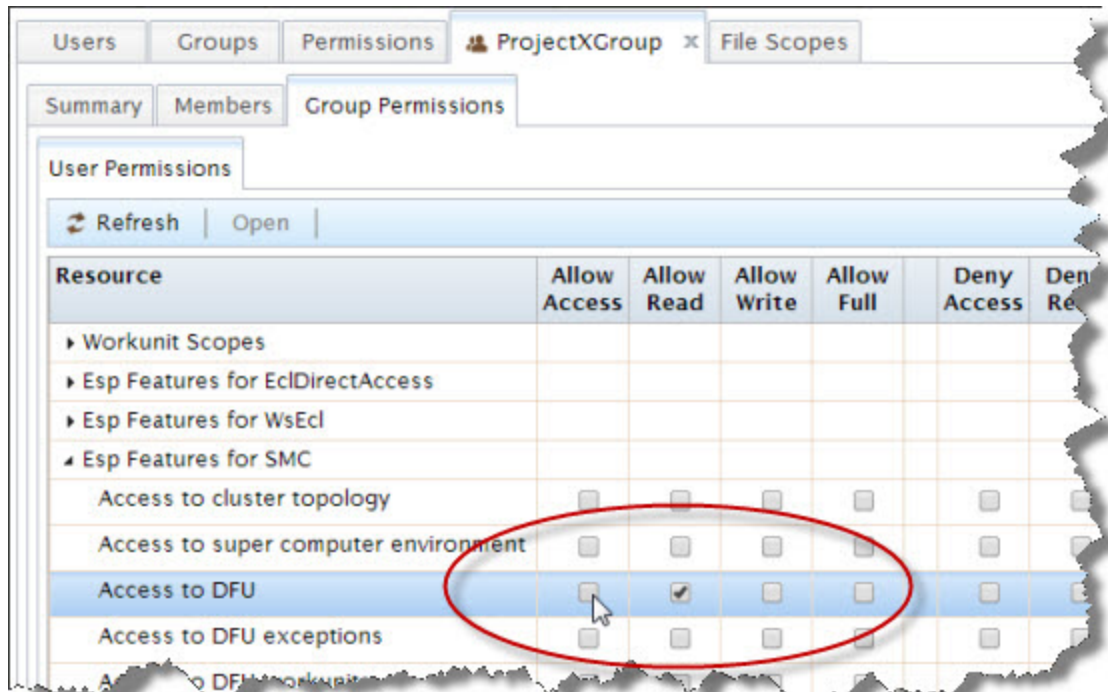


The group tab displays the sub-tabs: **Summary**, **Members**, **Active Group Permissions** and **Available Group Permissions**.

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.

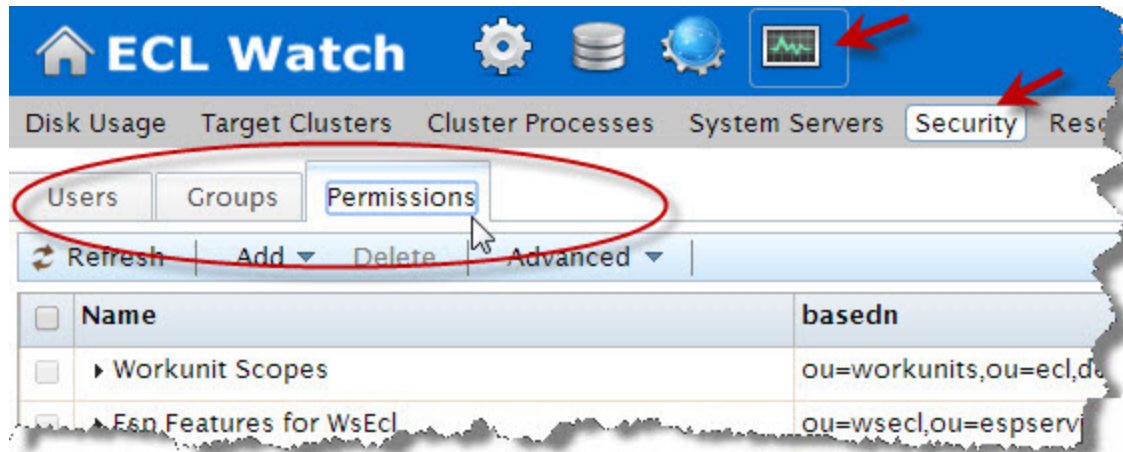


NOTE: Use caution when setting any explicit **deny** permission setting. The most restrictive permission always applies.

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.

Feature level access control

Access to the feature permissions is available through ECL Watch. In order to modify feature permissions you must have Administrator level access. To access the feature permissions click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.




To use the feature level controls, apply the feature resource from the **Available Permissions** tab to the **Active Permissions** for users and groups. Using the feature level controls allow you to:

- View the features and permissions for any resource
- Edit the permissions for any feature
- Update the permissions for users and groups for a specific resource

Feature resources

There are several features for which you can set up access control in HPCC. Access to features of the HPCC system is controlled by via the **ESP Features for SMC** category.



The available features are listed under the **Permissions** tab. You can view and gain access to the feature controls from here. However, the feature controls must be applied to users, or to groups. If you click on the feature name link, a tab opens that displays the users and groups where those feature permissions are applied.

ECL Watch feature permission settings that are not listed are not relevant and should not be used.

Apply permissions for a feature resource:

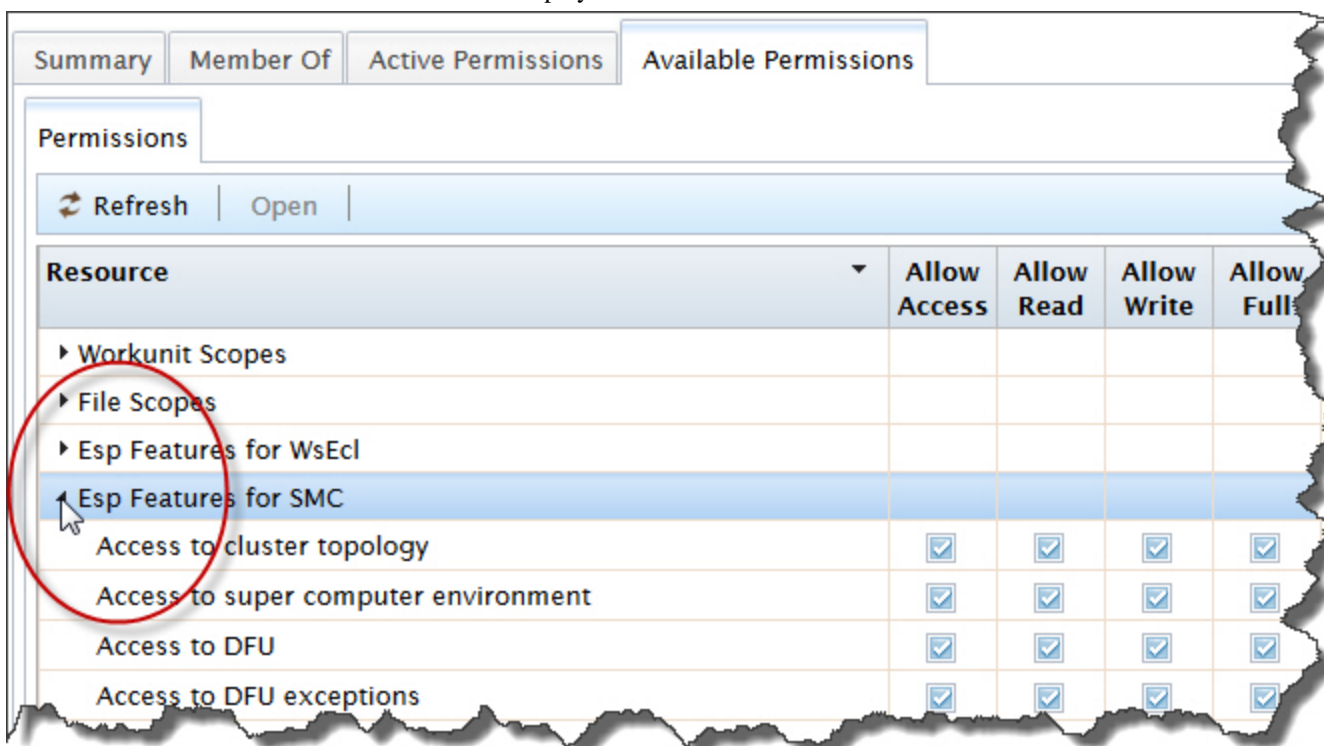
To use the feature permissions, you must apply them to a user or group(s). To access the feature permissions click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Identify the user(s) or group(s) which you want to modify the feature permissions.

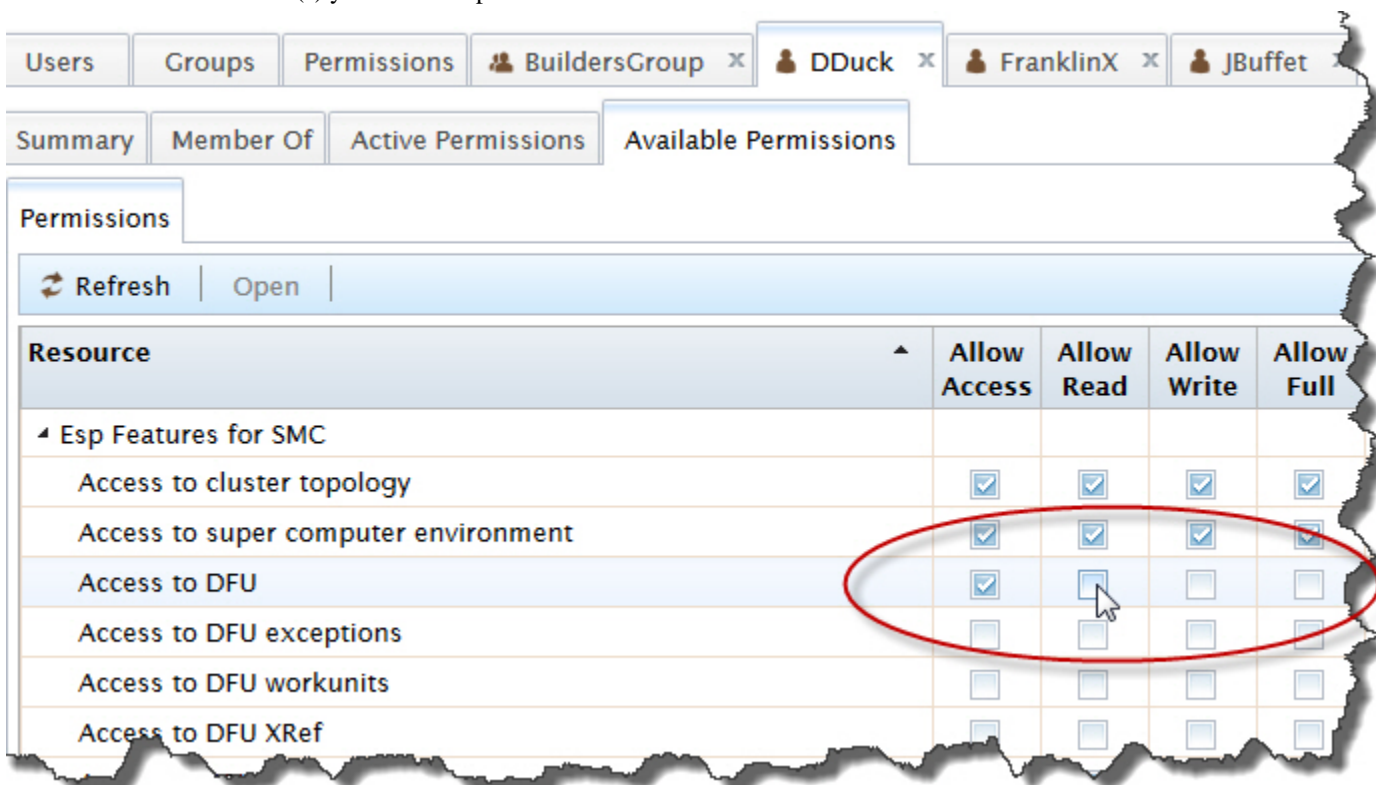
Select the appropriate tab. (Users or Groups)

2. Check the checkbox(es) next to the user(s) or group(s) to modify.
3. Press the **Open** action button. A tab for each user or group selected opens.
4. Click the **Available Permissions** sub-tab.

5. Click on the arrow to the left of the resource to display the features of that resource.



6. Locate the feature resource(s) you want to update.



7. Click the checkbox(es) in the **allow** and **deny** columns as appropriate.

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

Note: You must follow this process for each user or group(s) separately.

SMC Feature Permissions

The following table describes the level of access required to be able to use these HPCC ECL Watch features.


Name	Description	Access
ClusterTopologyAccess	Access to Cluster Topology	Read
	Access to log files.	Full
DfuAccess	Access to DFU Logical Files	Read
	Delete Files, add to superfiles, and remove from superfiles	Write
DfuExceptions	Access to DFU Exceptions	Read
DfuWorkunitsAccess	Access to View DFU Workunits	Read
	Access to Create, Delete, Update, Submit, and Abort DFU Workunits	Write
DfuXrefAccess	Access to DFU XREF	Read
	Clean directory	Write
	Make changes and generate XREF Reports	Full
EclDirectAccess	Access to ECL direct service.	Full
ESDLConfigAccess	ESDL Config Access	Read
	Publish ESDL definition and ESDL binding, configure ESDL binding method.	Write
	Delete ESDL definitions, delete ESDL bindings.	Full
FileDesprayAccess	Allows a user to despray logical files.	Write
FileIOAccess	Access to read files in Drop zone	Read
	Access to write to files in Drop zone	Write
FileScopeAccess	Allows access to query, set, modify, and delete File Scope Permissions	Full
FileSprayAccess	Access to Spraying and Copying	Read
	Rename, spray, copy, and replicate files	Write
	Delete from Drop zone	Full
MachineInfoAccess	Access to machine/Preflight Information	Read
MetricsAccess	Access to SNMP Metrics Information (Roxie Metrics)	Read
OthersWorkunitsAccess	Access to View Other User's Workunits	Read
	Access to Modify or Resubmit User's Workunits	Write
	Access to Delete Other User's Workunits	Full
OwnWorkunitsAccess	Access to View Own Workunit	Read
	Access to Create or Modify Own Workunit	Write
	Access to Delete Own Workunits	Full
RoxieControlAccess	Access to Roxie control commands	Read
SmcAccess	Access to ECL Watch (SMC Service)	Read
ThorQueueAccess	Access to Thor Job Queue Control	Full
WsEclAccess	Access to WS ECL service	Full

Some Feature Permissions Notes

- SMCAccess is required to be able to successfully login to ECL Watch.
- ThorQueueAccess allows you to manipulate the queue by promoting/demoting queued workunits according to priority.
- ThorQueueAccess also allows you to pause or clear the Thor queue. You can also view Thor usage statistics.
- Depending on the level of access the user has, they can view, modify, and delete their own, or others workunits. This is OwnWorkunitsAccess, and OthersWorkunitsAccess respectively.
- DfuWorkunitsAccess permissions allow users to view and/or manipulate DFU Workunits.
- Users need permission to see files on the dropzone and also to put files there. They need further permissions to be able to spray and copy files from the dropzone to their cluster and also to despray files from the cluster back to the dropzone.

DFU Xref

XREF is used for monitoring files on the cluster(s). Reports generated show where housekeeping is required on the cluster(s) and users require additional permission to use this feature.

	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.
---	--

Users/Permissions

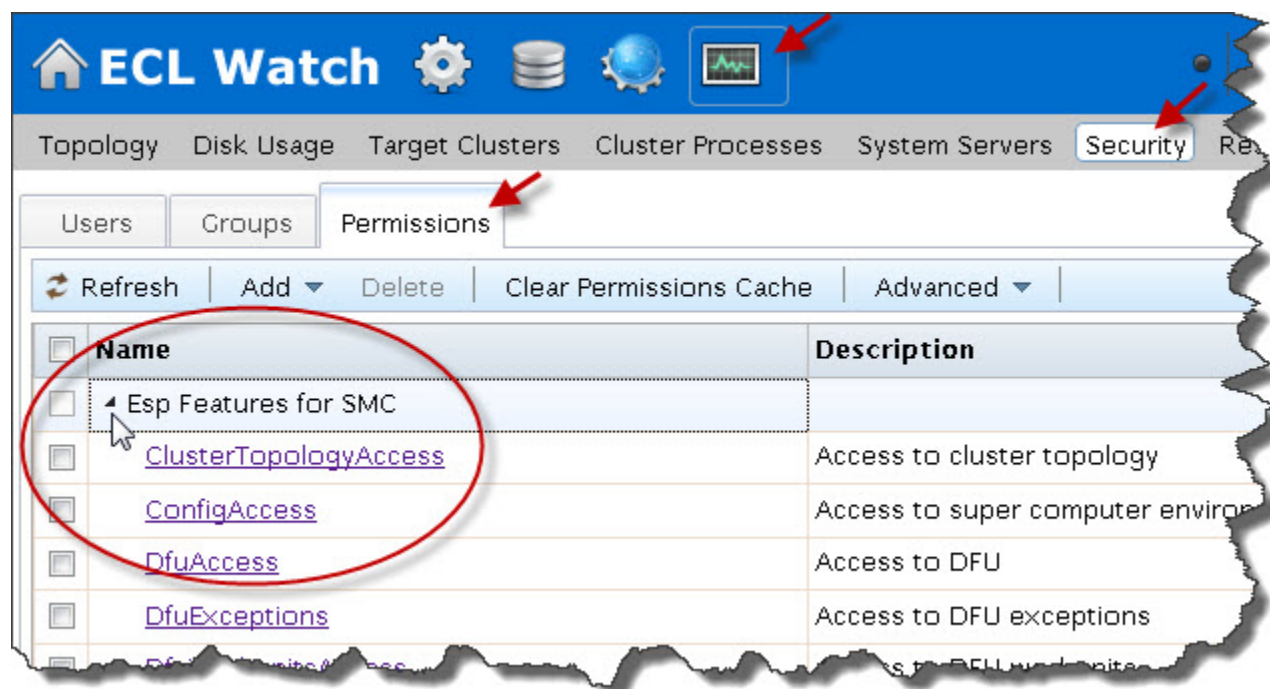
To be able to view the **Users/Permissions** area in ECL Watch, a user must be a member of the Administrators (or similarly named) group with the appropriate permissions on the LDAP or Active Directory server.

File Access Control

The HPCC's LDAP **Dali Server** technology provides the ability to set secure access permissions to data file folders (or file scopes). This is controlled by the use of file scope resources.

An OU called **Files** is automatically created when the Dali server starts. To secure data folders, create a file scope for that folder and apply rights to each scope.

Figure 26. File Scopes Permissions



For example, below **Files** there is a unit (OU) representing the cluster, such as **thor** (or the name that you set up for your cluster). Furthering the example, below that could be a unit named **collectionx** which contains two units, **publicdata** and **securedata**. The **publicdata** folder has rights granted to a large group of users and the **securedata** folder has limited access granted. This allows you to prevent unauthorized users from any access to files in the **securedata** folder.

The structure described above corresponds to this logical structure:

collectionx::securedata

Which corresponds to this physical structure:

/var/lib/HPCCSystems/hpcc-data/thor/collectionx/securedata

All HPCC components and tools respect LDAP file access security. The following exceptions are assumed to be system level or for administrative users:

- Network file access using UNC's, Terminal Services, or SSH.
- Administrative utilities

Attempting to access a file in a folder for which access is not granted will result in one of the following errors:

DFS Exception: 4 Create access denied for scope <filepath>

or

DFS Exception: 3 Lookup access denied for scope <filepath>

(where <filepath> is the full logical file scope path)

Creating file scopes


To apply permissions to a file scope, you must first create the file scope(s).

To create file scope(s) click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Click the **Permissions** tab.


The feature resources display.

2. Click on the arrow to the left of the **File Scopes** resource to display the file scopes.




3. Press the **Add** button.

4. Choose **File Scopes** from the drop list.



5. Enter the exact name of the scope you want to add in the **Name** field.



Enter a short description in the **Description** field.

6. Press the **Add** button.

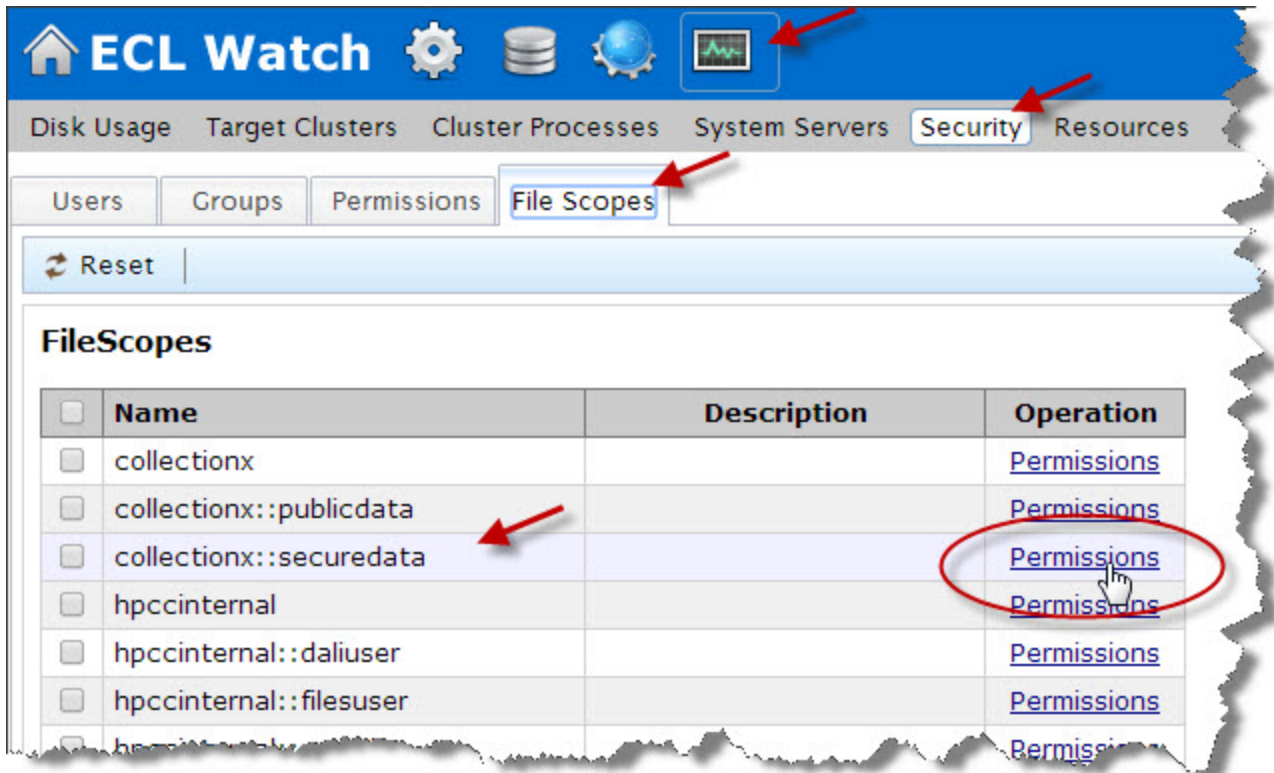
The new scope displays in the list.

Setting permissions for file scopes

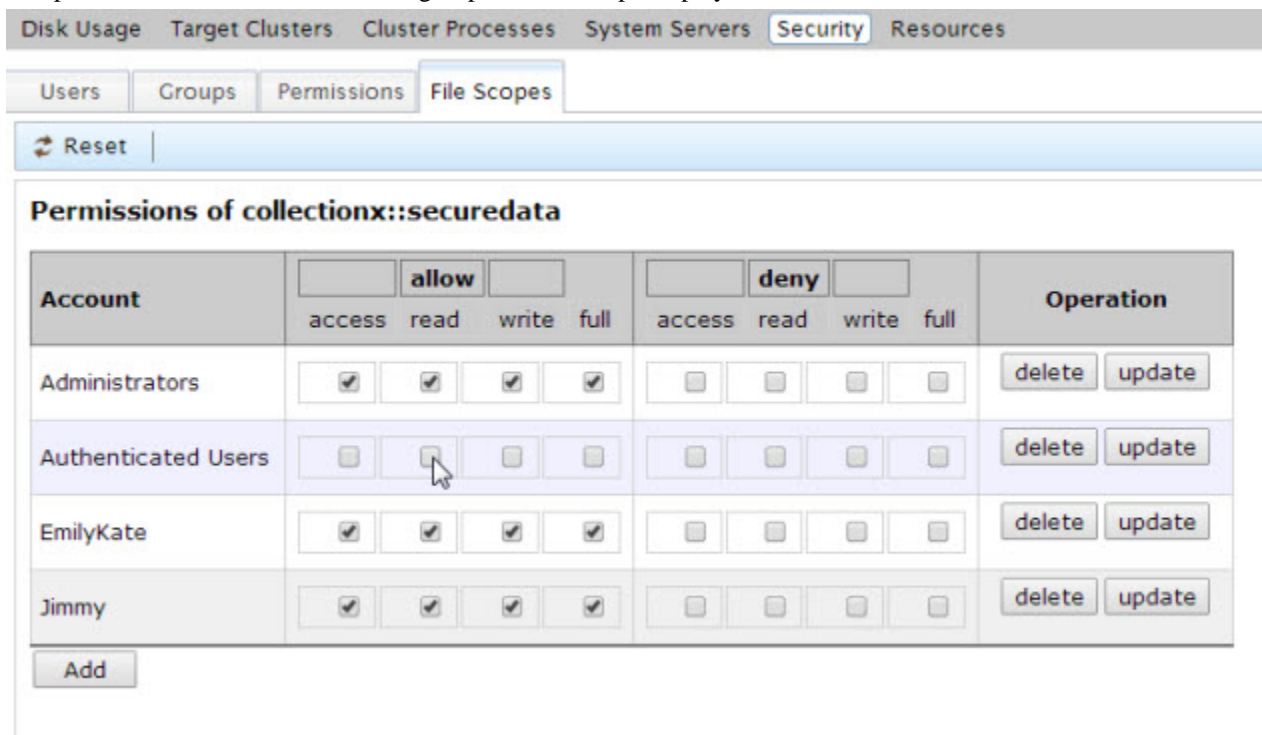
You must apply permissions for file scopes to users or group(s). If you want to apply the scope to a new group, create the group(s) as required.

To set the file scope permissions click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.

1. Select the **File Scopes** tab.
2. Choose the scope to modify. Click the **Permissions** link for that scope.



3. The permissions defined for users and groups for that scope display.



4. Check (or clear) the checkbox(es) in the **allow** and **deny** columns as appropriate for the users or groups displayed.
5. To add users or groups to the scope, press the **Add** button.

The Add Permission dialog displays.

6. Select the user or the group to add from the drop list(s).

Add Permission for collectionx::securedata

Reset

Users Groups Permissions File Scopes

Select user: none

Or group: none

allow: access read write full

deny: access read write full

Add

Add user or group permission drop list

Once a user or group is selected, the Add button and the allow and deny checkboxes are active

7. Check the boxes for allow and deny as appropriate to set the permissions for this scope.

Users Groups Permissions File Scopes

Reset

Add Permission for collectionx::securedata

Select user: guser

Or group: none

allow: access read write full

deny: access read write full

Add

8. Press the **Add** button.

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

File scope features

Below the List of File Scopes, there are buttons that allow you to:

- Reset **Default Permissions** to selected file(s)


This allows you to quickly remove any added permission settings for a file and reset to the default access.

- Allow or Deny Access to physical files on Landing Zone

This provides a way to grant or deny access to the top level file scope. By default, only administrators have access to this scope.

- Check File Permissions for a user or group

This provides a way to check a user or group's access to a logical file.



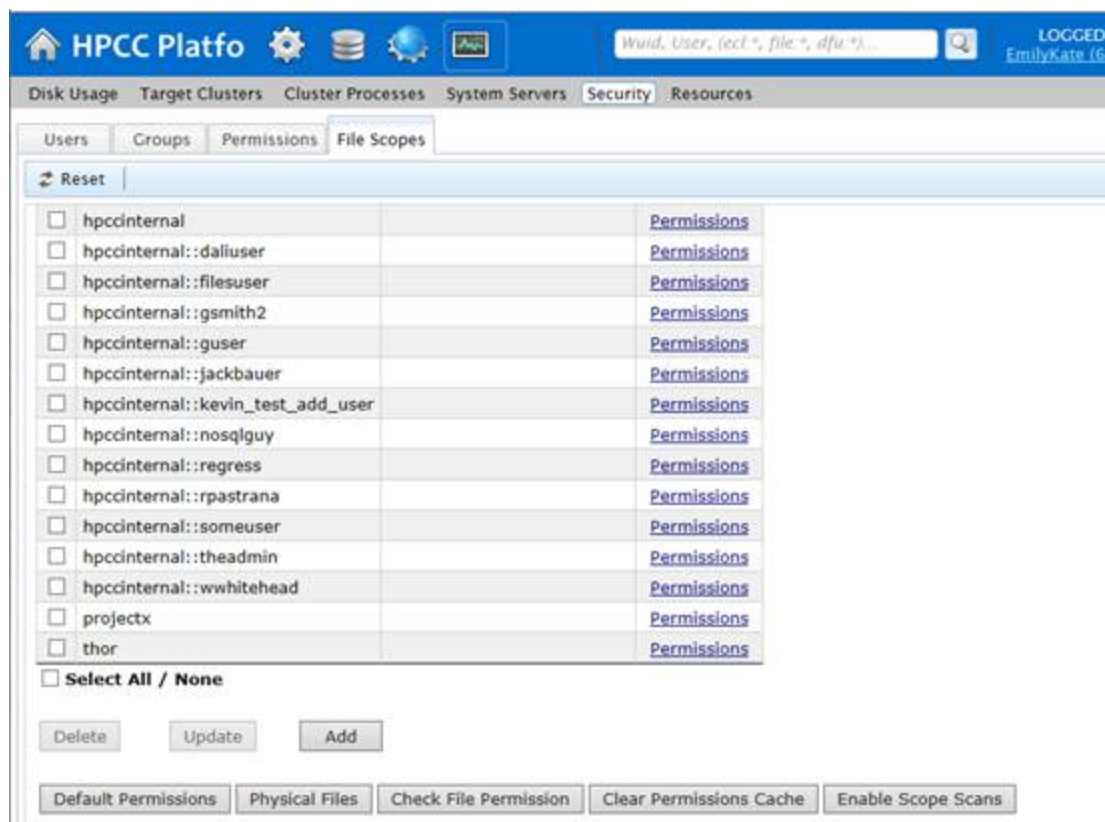
- Clear the Permissions Cache

This clears the permissions cache and allowing any new permission settings to take effect immediately.

- Enable/Disable Scope Scans

This provides a means to enable or disable Scope scans. Enable scope scans to check permissions for users to access scopes. This will impact performance. Disable scope scans ignores any scope permissions and removes all access control, but improves performance. Disabling access control is not recommended.

Changing this setting through ECL Watch, as described here, is only a temporary override. When Dali restarts this setting will revert to what is defined in the configuration environment.xml.



Workunit Access Control

There are 2 aspects of workunit (WU) security:

- Feature Authentication for workunits allows you to set permissions to control whether users can view their own WUs and/or other users' WUs.
- Workunit Scope security provides the ability to set permissions for individual WU scopes. All new workunits have a scope value.

Both methods are valid to use (either separately or together), and the strictest restriction always applies.

In other words, if someone is granted permission to see WUs in the scope *johndoe* but is denied permission to see other users' WUs in the Feature Authentication permissions, this user would be denied access to see the WUs in the *johndoe* scope.

Conversely, if the user is allowed access to see other people's WUs but is denied access to the *johndoe* WU scope, this user will be able to see other WUs in that scope.

Note: If you do not have access to a WU, you will never be able to view it or even know of its existence.

By default, a submitted WU has a scope of the user's ID. For example, a WU JohnDoe submits has *scope=johndoe* in the WU. This value in a WU allows ESP and its services to use LDAP to check for permissions and enforce those permissions.

You can override the default scope using ECL Code:

```
#workunit('scope', 'MyScopeValue');
```

Securing workunit scopes

ESP (on startup) automatically creates an LDAP OU called **Workunits** (unless it already exists). If this OU is automatically created, the OU is made with full permissions granted to all authenticated users. All WU scopes are below the *workunits* OU either implicitly or explicitly.

If a specific scope OU does not exist in LDAP (e.g., the scope johndoe used in earlier example), then the parent OU's permissions are used. In other words, the scope of *johndoe* is implicitly under the *workunits* OU even though it might not be explicitly listed in the LDAP structure and therefore it would use the permissions granted for the parent, *workunits*.

Workunits feature permissions

Using the **Workunit Scopes** feature in the **Permissions** area of ECL Watch the permissions for any scope can be reset to the default permissions settings for your system. Permission settings for Workunit Scopes may be set as follows:

Description	Access
View WUs in that scope	Read
Create/modify a WU in that scope	Write
Delete a WU in that scope	Full

Adding workunit scopes

To add workunit scope permissions click on the **Operations** icon, then click the **Security** link from the navigation sub-menu.


1. Click the **Permissions** tab.

The feature resources display.


2. Click on the arrow to the left of the **Workunit Scopes** resource to display the file scopes.



3. Press the **Add** button.
4. Choose **Workunit Scopes** from the drop list.



5. Enter the exact name of the scope you want to add in the **Name** field.



Enter a short description in the **Description** field.

6. Press the **Add** button.

The new scope displays in the list.


Set permissions to the scope.

You apply the workunit scopes to a group. If you want to apply the scope to a new group, create the group(s) as required.

1. Go to the **Groups** tab.
2. Select a group to apply the scope to by checking the box next to the group name.

Press the **Open** action button. You can select multiple groups, a tab opens for each group.

3. Select the **Group Permissions** tab of that group. (if multiple groups selected, you must repeat for each group)
4. Click on the arrow to the left of the Workunit Scopes to display the available scopes.



The Workunit scopes display. Check the boxes as appropriate to set the permissions for this scope.

5. To set permissions in this scope for another group, open and go to that groups tab.
6. To set permissions in this scope for a user, select the tab.
7. Select the user and press the Edit action button.

A new tab for that user opens.

8. On that tab, click on the **User Permissions** sub-tab.
9. Locate the new scope listed under the appropriate Resource.

Set the access permissions as appropriate for that user.

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

Permission Caching

When you change a permission in ECL Watch, the settings are cached in the ESP server and stored in the Dali server. The information in the cache is updated at a configurable interval. This value can be set in the Configuration Manager under the LDAP Server settings Attributes tab. The default cacheTimeout is 5 minutes.

When you want a permission change to take effect immediately, you can clear the cache and force Dali to update the permission settings by pressing the **Clear Permissions Cache** button. This action transfers the settings when you press the button. Use this feature judiciously as overall system performance is affected temporarily while the LDAP settings in the Dali System Data Store repopulate.

Workunits and Active Directory

The performance of your system can vary depending on how some components interact. One area which could impact performance is the relationship with users, groups, and Active Directory. If possible, having a separate Active Directory specific to HPCC could be a good policy. There have been a few instances where just one Active Directory servicing many, diverse applications has been less than optimal.

HPCC makes setting up your Active Directory OU's relatively easy. ESP creates all the OU's for you when it starts up, based on the settings you defined in Configuration Manager. You can then start Dali/ESP and use ECLWatch to add or modify users or groups.

You can assign permissions to each user individually, however it is more manageable to assign these permissions to groups, and then add users to these groups as appropriate. Create a group for developers and power users (people with full read/write/delete access), and another group for users that only have only read access and perhaps another group that has both read and write access. Add any other groups as appropriate for your environment. Now you can assign users to their appropriate group(s).

Active Directory, and LDAP Commonality

There are components that are common to both Active Directory and LDAP. There are a few relevant terms, that may need some further explanation.

filesBasedn	Deals with restricting access to files. Also referred to as “file scoping”.
groupsBasedn	Controls the groups associated with the environment. For example, administrators, developers, ws_ecl only, etc.
modulesBasedn	Specific to systems using a legacy central repository and controls access to specific modules. Any module you create in the application will create an entry in Eclwatch>>User/Permissions>>Repository Modules
sudoersBasedn	Deprecated.
workunitsBasedn	Controls access to workunits.

Using Cassandra for Workunit Storage

Beginning with version 6.0.0, you can configure your system to use Apache Cassandra™ database for Workunit storage. Apache's Cassandra NoSQL data store is a software framework that supports data-intensive distributed applications under a free license.

The default HPCC configuration continues to be the same--Workunits are stored in Dali's data store and Sasha is typically configured to archive workunits to disk to save memory usage.

Switching to Cassandra Workunit storage eliminates the need to archive workunits and allows you to manage your data independently.

This section details the steps to configure your HPCC platform to use Cassandra as a Workunit server.

Cassandra as a Workunit Server

Pros and Cons

Dali	Cassandra
Storing in Dali requires no additional software or work	Must install and configure a Cassandra server
Sasha provides simple housekeeping	Requires its own housekeeping and maintenance
Limited by Dali memory	Unlimited (elastic scalability)
Archived Workunits are harder to access, must be restored before results can be seen	No need to archive so all Workunits can remain "Live"
Single location	Distributed, scalable, highly available.
Single backup of data supported. Manual Fail over	Inherent fault tolerance with data replication, designed for high availability.

Installation and Configuration Overview

1. Install Cassandra and configure it as you wish. The only requirements to work with an HPCC platform are:
 - a. Create a keyspace to use for Workunit Storage in Cassandra. The default is hpcc.
 - b. If you want security, you should set up Cassandra's authentication and create a superuser for HPCC to use. These credentials must then be entered in the DaliPlugin configuration in Configuration Manager.
 - c. Ensure your Cassandra server can be reached by your HPCC nodes. (listen_address and rpc_address)
 - d. Know your Cassandra settings (IP, Port, keyspace, user credentials, etc.)
2. Use Configuration Manager to configure your HPCC platform to use Cassandra for Workunit storage.
3. Make sure Cassandra is running.

Once you have configured your HPCC Platform to use Cassandra for Workunit storage, you must make sure Cassandra is running before starting your system.

4. Push out your environment.xml and restart your HPCC Platform.
5. If you have existing Workunits, export them from Dali and import them into Cassandra using wutool.
6. If you have existing archived Workunits, and you want them active, import them into Cassandra.
 - a. Once Workunits are stored in a Cassandra Workunit store, Sasha will no longer archive Workunits (despite any settings it may have).
 - b. Any previously archived Workunits will remain in storage on the Sasha server and will continue to be available in the same manner they have been.

Configure your HPCC platform to use Cassandra for Workunit storage

This option is configured by adding a DaliPlugin component to your system.

Connect to Configuration Manager

In order to change the configuration for HPCC components, connect to the Configuration Manager.

1. Stop all HPCC Components, if they are running.
2. Verify that they are stopped. You can use a single command, such as :

```
sudo /opt/HPCCSystems/sbin/hpcc-run.sh -a hpcc-init status
```

3. Start Configuration Manager.

```
sudo /opt/HPCCSystems/sbin/configmgr
```

4. Connect your web browser to the Configuration Manager web interface.

(using the url of `http://<configmgr_IP_Address>:8015`, where `<configmgr_IP_Address>` is the IP address of the node running Configuration Manager)

5. Select the **Advanced View** radio button.
6. Use the drop list to select the XML configuration file.

Note: Configuration Manager **never** works on the active configuration file. After you finish editing you will have to copy the `environment.xml` to the active location and push it out to all nodes.

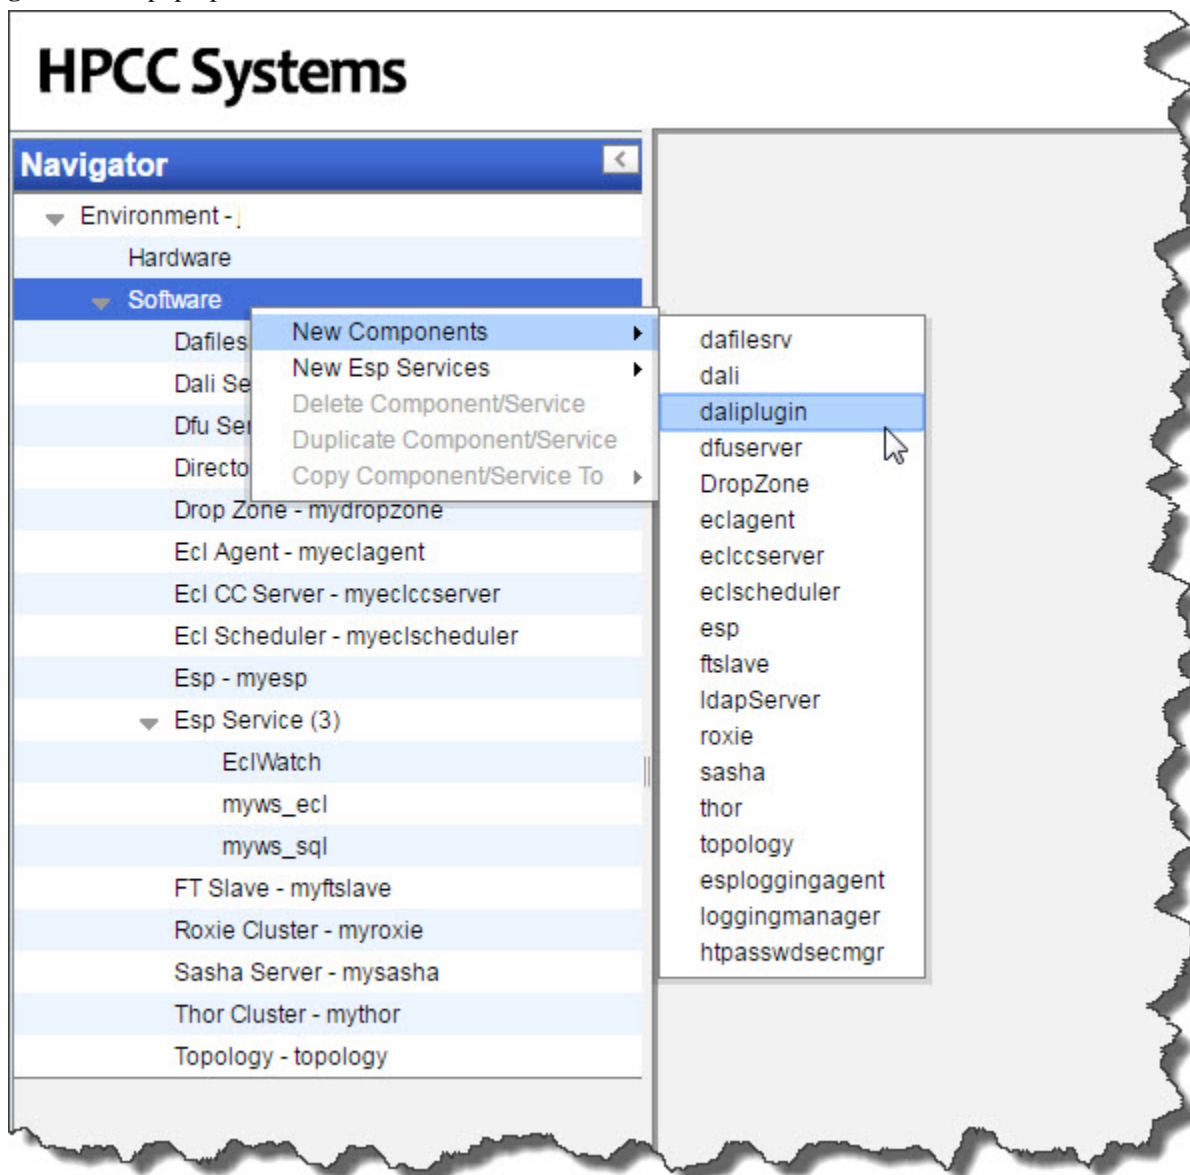
7. Check the **Write Access** box.

Default access is read-only. Many options are only available when write-access is enabled.

Adding the DaliServerPlugin component

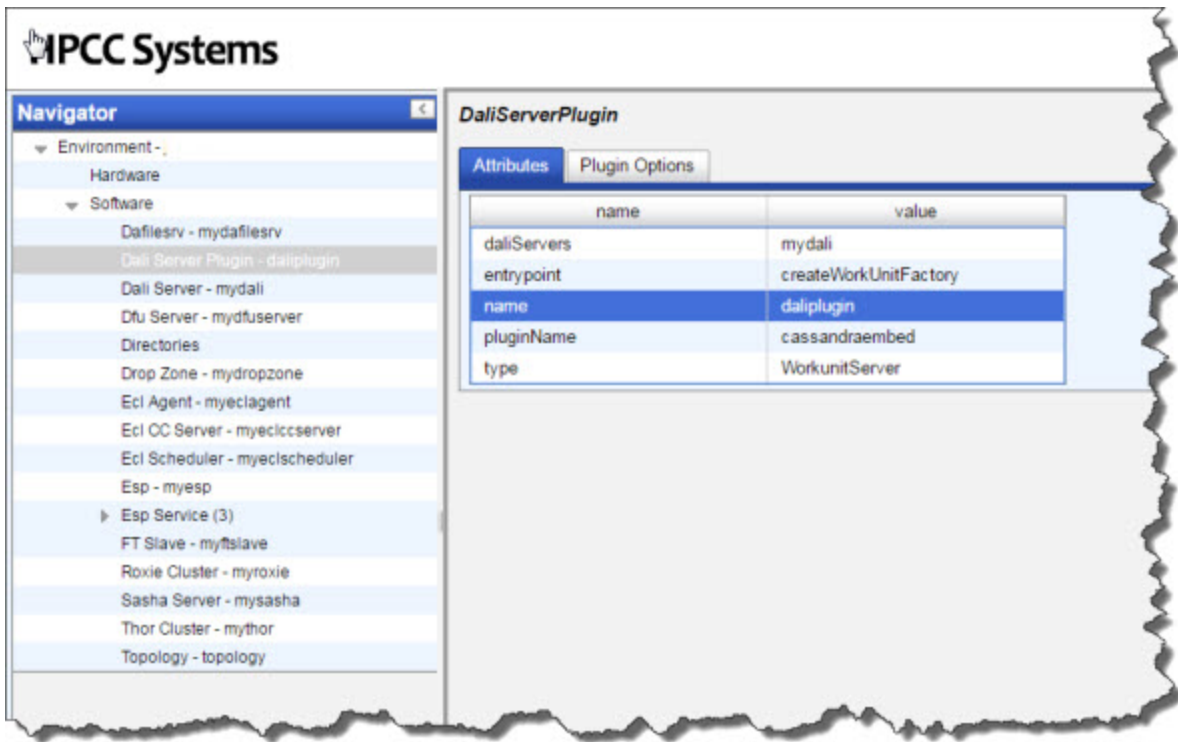
The DaliServerPlugin controls your Cassandra Workunit storage options.

1. Right-click on **Navigator** Pane and choose **New Components** from the pop-up menu, then choose **DaliServerPlugin** from the pop-up menu.

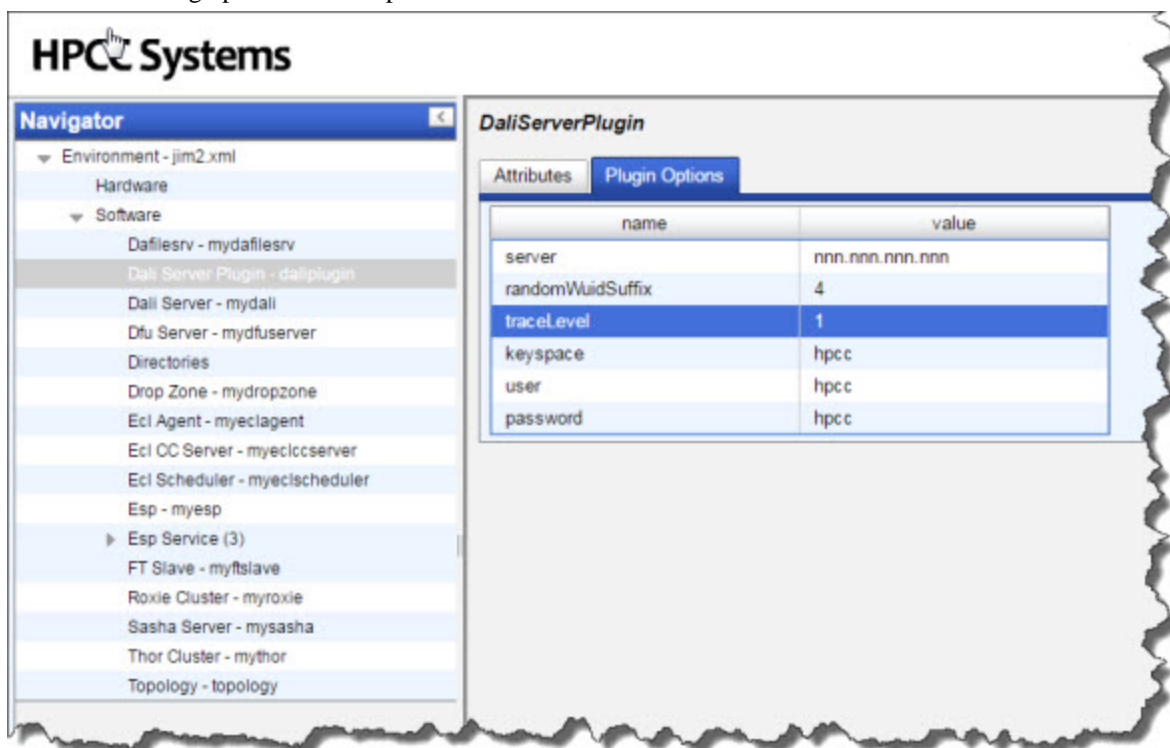


2. Fill in the **DaliServerPlugin** properties:

- a. Fill in the **Attributes** tab with the appropriate settings for your existing Cassandra Server.



- b. Add the following options on the Options tab.



Note: User and password are only needed if your Cassandra server is configured to require credentials.

- c. Click on the disk icon to save.

3. Copy the new xml file from the source directory to the /etc/HPCCSystems and rename the file to environment.xml

```
# for example
sudo cp /etc/HPCCSystems/source/NewEnvironment.xml /etc/HPCCSystems/environment.xml
```



Make sure that you have sufficient privileges to write file(s) to the destination directory before attempting to copy. If prompted to overwrite the destination file, you should answer **yes**.

4. Copy the /etc/HPCCSystems/environment.xml to /etc/HPCCSystems/ on **every** node.

You may want to create a script to push out the XML file to all nodes. A sample script is provided with HPCC. The following command copies the XML files out to all nodes as required:

```
sudo /opt/HPCCSystems/sbin/hpcc-push.sh <sourcefile> <destinationfile>
```

5. Restart the HPCC system on **every** node.

Using wutool

wutool *action* [WUID=nnn] [DALISERVER=ip | CASSANDRASERVER=ip] [option=value]

The wutool is a command line utility used to maintain your Workunit store. It can be found in /opt/HPCCSystems/bin/ on any server where the platform has been installed. You can use this utility to export your Workunits from a Dali server and then import into a Cassandra data store. You can also use it to import archived workunits on a Sasha server.

Actions	
list <workunits>	List workunits.
dump <workunits>	Dump xml for specified workunits.
delete <workunits>	Delete workunits.
results <workunits>	Dump results from specified workunits.
archive <workunits>	Archive specified Workunits to xml files. The following options are supported: [TO=<directory>] [DEL=1] [DELETERESULTS=1] [INCLUDEFILES=1]
restore <filenames>	Restore from xml files. [INCLUDEFILES=1]
orphans	Delete orphaned information from store
cleanup [days=NN]	Delete workunits older than NN days
validate	Check contents of workunit repository for errors. [FIX=1] will try to repair any issues found.
clear	Delete entire workunit repository (requires entire=1 repository=1)
initialize	Initialize new workunit repository

<workunits> can be specified on the command line or can be specified using a filter owner=XXXX. If omitted, all workunits are selected.

Options	
(These options are valid if CASSANDRASERVER is specified)	
CASSANDRA_KEYSPACE	Cassandra keyspace. default is hpcc.
CASSANDRA_USER	Cassandra username to use (if needed).
CASSANDRA_PASSWORD	Cassandra password to use (if needed).
TRACELEVEL	Trace level (1-10)

Examples:

```
/opt/HPCCSystems/bin/wutool archive DALISERVER=. del=1
```

```
/opt/HPCCSystems/bin/wutool restore *.xml CASSANDRASERVER=192.168.56.120 CASSANDRA_KEYSPACE=hpcc2
```

Redefining nodes in a Thor Cluster

To reconfigure a Thor cluster where you replace existing nodes (with new IP's) or add or remove nodes, you must take an additional step to restructure the group. Dali will not automatically restructure an existing group.

This is because existing published files reference the previous cluster group state by name and therefore changing its structure would invalidate those files and make the physical files inaccessible.

There are a couple of scenarios where you would want to redefine your Thor cluster.

Replacing faulty node(s)

If data files are replicated, replacing a node and forcing the new group to be used by existing files may be desirable. In this scenario, reading an existing file will failover to finding a part on the replicate node, when it tries to find a physical file on the new replacement node.

To force the new group to be used, use the following command:

```
updt dalienv <environment_file> -f
```

In cases where there is no replication, data loss may be unavoidable and forcing the new group may still be the best option.

Resizing the cluster

If you are adding or removing Thor cluster nodes but *all previous nodes remain part of the environment and accessible*, you must **rename** the group that is associated with the Thor cluster (or the Cluster name if there is no group name).

This will ensure all previously existing files, continue to use the old group structure, while new files use the new group structure.


In summary, if the Thor cluster changes it must be updated in the Dali.

Best Practices

This chapter outlines various forms of best practices established by long time HPCC users and administrators running HPCC in a high availability, demanding production environment. While it is not required that you run your environment in this manner, as your specific requirements may vary. This section provides some best practice recommendations established after several years of running HPCC in a demanding, intense, production environment.

Cluster Redundancy

There are several aspects of cluster redundancy that should be considered when setting up your HPCC system.

	Make sure you allocate ample resources to your key components. Dali is RAM intensive. ECL Agent and ECL Server are processor dependent. Thor should have a minimum of 4GB RAM per node.
---	---

Dali

Dali should be run in an active/passive configuration. Active/passive meaning you would have two Dalis running, one primary, or active, and the other passive. In this scenario all actions are run on the active Dali, but duplicated on the passive one. If the active Dali fails, then you can fail over to the passive Dali.

Another suggested best practice is to use standard clustering with a quorum and a takeover VIP (a kind of load balancer). If the primary Dali fails, you move the VIP and data directory over to the passive node and restart the Dali service.

DFU Server

You can run multiple instances of the DFU Server. You can run all instances as active, as opposed to an active/passive configuration. There is no need for a load balancer or VIP. Each instance routinely queries the Dali for workunits. Should one fail, the other(s) will continue to pull new workunits.

ECLCC Server

You can run multiple active instances of the ECLCC Server for redundancy. There is no need for a load balancer or VIP for this either. Each instance will routinely check for workunits. Should one fail, the other(s) will continue to compile.

ESP/ECL Watch/WsECL

To establish redundancy, place the ESP Servers in a VIP. For an active/active design, you must use a load balancer. For active/passive you can use pacemaker/heartbeat. If you run active/active, you should maintain a single client's connection to a single server for the life of a session for ECL Watch (port 8010). Other services, such as WsECL (port 8002) do not require a persistent connection to a single server.

ECL Agent

You can run multiple active instances of the ECL Agent. No need for a load balancer or VIP. Each instance routinely queries for workunits. Should one fail, the other(s) will continue to pull new workunits.

Sasha

Sasha should be run in an active/passive configuration. Active/passive meaning you would have two Sashas configured, one primary (active), and the other standing by.

ECL Scheduler

No need for a load balancer, runs active/active. Each instance routinely queries for workunits. Should one fail, the other(s) will continue to schedule workunits.

Thormaster

Set up Thor in an active/passive configuration. Active/passive meaning you would have two instances running, one primary (active), and the other passive. No load balancer needed. If the active instance fails, then you can fail over to the passive. Failover then uses the VIP (a kind of load balancer) to distribute any incoming requests.

Dropzone

This is just a fileserver that runs the dafiler process. Configure in the same fashion as you would any active/passive file server. One primary, or active, and the other passive. No load balancer needed. If the active instance fails, then you can fail over to the passive.

High Availability

If you require high availability for your HPCC system, there are some additional considerations that you should be aware of. This is not comprehensive list, and it is not meant to be step-by-step instructions for setting up disaster recovery. Instead this section just provides some more information to consider when incorporating HPCC into your disaster recovery plan.

Thor

When designing a Thor cluster for high availability, consider how it actually works -- a Thor cluster accepts jobs from a job queue. If there are two Thor clusters handling the queue, one will continue accepting jobs if the other one fails.

If a single component (thorslave or thormaster) fails, the other will continue to process requests. With replication enabled, it will be able to read data from the back up location of the broken Thor. Other components (such as ECL Server, or ESP) can also have multiple instances. The remaining components, such as Dali, or DFU Server, work in a traditional shared storage high availability fail over model.

The Downside

Costs twice as much initially because you essentially have to have two of everything.

The Upside

Almost 100% of the time you can utilize the additional processing capacity. You can run more jobs, have more space, etc.

Disaster Recovery concerns

The important factor to consider for disaster recovery (DR) is the bandwidth required to replicate your data. Your network administrator should evaluate this aspect carefully.

If you have tens of gigabytes of delta each day then an rsync type replication or some sort of hybrid model should suffice. If you have hundreds of gigabytes to petabytes of deltas, the real limit is your budget.

A best practice is to find where the data is the smallest (at ingestion, after normalization, at Roxie) and replicate from that point and rerun the processing in both locations.

The key to getting disaster recovery right is to know your data flow. For instance, if you are ingesting 20TB of raw data daily, then taking that raw data and rolling it up, scoring it, indexing it, etc. You would be better off replicating an intermediate dataset (that we call base files), rather than replicating the large ingest. If the opposite is occurring (small daily ingest and then blow the data up in size) – you would be better off to ingest the input and then re-run it.

Thor has the ability to do a “Thor copy” which copies data from one cluster to another. You can also do this through ECL code. Additionally, you may decide you don’t want, or need to have a “hot” DR Thor. In that case, the most common minor disasters cause only a relatively brief, less than 1 day disaster. Since Thor is responsible for creating data updates it can take a day or a few to recover. The data just is not quite as fresh but as long as the Roxies are replicated the data is still flowing. In the case of a major disaster such as, a major earthquake, a tidal wave, extended total power loss, multiple fiber cuts, where the systems will be out for a day or more. The likelihood of that occurring may not justify the costs of preventing against it.

Conclusion

Disaster recovery is a calculation. The cost of failure, times the likelihood per year of an event occurring, less than or greater than the cost to prevent against it. Taking all that into consideration can help you to put a sensible DR plan in place.

Roxie

In the case of Roxie, a best practice is to have multiple Roxie clusters and use a proxy to balance. In case of how to keep the data in sync, a pull approach is best. The Roxie automatically pulls the data it needs from the “source” listed in the package file. The data can also be pulled from another Roxie or a Thor. In most cases you would pull to your DR Roxie from the primary Roxie out of the load balancer, but it can also pull from a Thor in the primary location as well.

Middleware

Replication of some components (ECL Agent, ESP/Eclwatch, DFU Server, etc.) are pretty straight forward as they really don't have anything to replicate. Dali is the biggest consideration when it comes to replication. In the case of Dali, you have Sasha as the back up locally. The Dali files can be replicated using rsync. A better approach could be to use a synchronizing device (cluster WAN sync, SAN block replication, etc.), and just put the Dali stores on that and just allow it replicate as designed.

There isn't just a one size fits all approach. Special care, design, and planning are required to make an effective DR strategy that doesn't “over synchronize” across slow WAN links, but still provides you with an acceptable level of redundancy for your business needs.

Best Practice Considerations

There are several other aspects to best practice considerations, and these will change with your system requirements. The following sections are some best practice considerations for some aspects of the HPCC system. Keep in mind that suggested best practices are merely suggested and may not be appropriate for your needs. A thorough review of the considerations highlighted here can be very helpful if your needs align with the stated considerations.

Multiple Thors

You can run multiple Thors on the same physical hardware. Multiple Thors on the same hardware are independent and unaware of each other. The Thors run jobs as they receive them, regardless of what the other(s) is/are doing. The speed of a single job will never be faster with multiple Thors, but the throughput can be. You can run two Thors picking up jobs from two different queues or the same queue.

The downside of running multiple Thors on the same hardware is that the physical memory on the nodes needs to be shared among each of the Thors. This needs to be configured per Thor cluster definition.

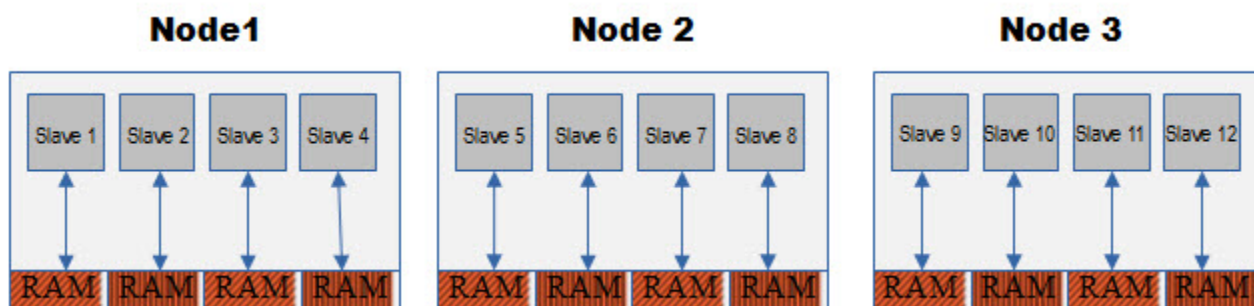
You must not place multiple Thors on hardware which does not have enough CPU cores to support it. You should not have more Thors than number of cores. One good rule is to use a formula where the number of cores divided by two is the maximum number of Thor clusters to use.

Virtual Thor slaves

Beginning in version 6.0.0, Thor clusters can be configured to take full advantage of the resources available per node using Virtual Thor slaves.

In HPCC versions prior to 6.0.0, cluster configurations were typically set to N number of **slavesPerNode** , where N equalled or approached the number of cores per machine.

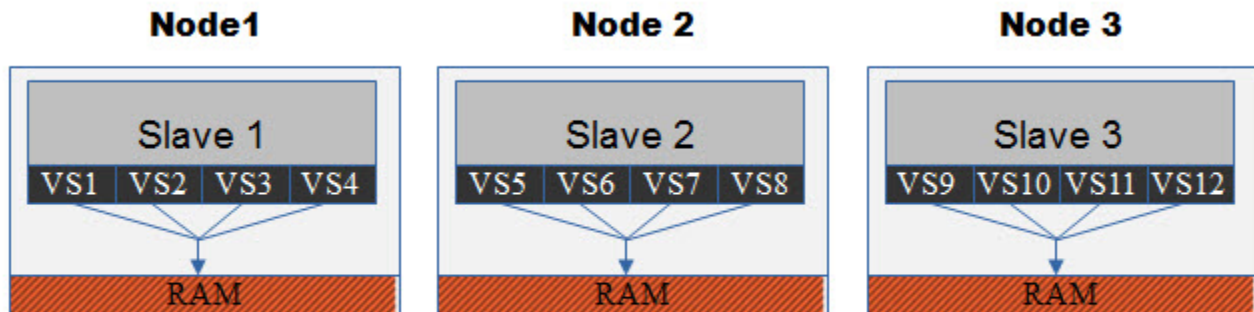
This resulted in N independent slave processes per node, as seen below:



This had several significant disadvantages:

- Each slave process in this configuration has an equal fixed split of the physical memory available to the node.
- Slaves do not share RAM or any other resources.
- Slaves use message passing via the loopback network interface for communication.

Now a new approach is used, allowing virtual slaves to be created with a single slave process, as depicted below:.



* VS = Virtual Slave (also known as a 'channel')

- In this configuration, each physical node has a single Thor slave process.
- Each slave process has N virtual slaves. This is set using a Thor configuration option called **channelsPerSlave**. Under this architecture, slaves within the same process can communicate directly with one another and share resources.

Note: The *slavesPerNode* setting still exists and both may be used in combination if required.

Key advantages:

- Each virtual slave shares cached resources, such as key index pages, etc.
- Slaves can request and share all available RAM.
- Startup and management of the cluster is faster and simpler.
- Allows for future enhancements to bring better management/coordination of CPU cores.

The significance of having access to all available memory becomes very significant for some activities. The clearest example is a SMART or LOOKUP JOIN.

SMART/LOOKUP JOIN example

A LOOKUP JOIN works approximately as follows:

- Streams local slave RHS dataset to all other slaves.
- All slaves gather global RHS into one table.
- A hash table based on the hard key match fields is built.
- Once all slaves are done, the LHS is streamed and matched against the hash table to produce the joined results.

Note: The complete RHS table and hash table must fit into memory; otherwise the join fails with an out of memory error.

SMART JOIN is an evolution of LOOKUP JOIN. If it cannot fit the global RHS into memory, it will HASH PARTITION the RHS and HASH DISTRIBUTE the LHS and perform a LOCAL LOOKUP JOIN.

If it cannot fit the local RHS set into memory on any given node, then it will gather and sort both local datasets and perform a standard JOIN.

The key advantage of LOOKUP JOIN is speed. If the RHS fits into memory, it can perform a very quick gather and streamed JOIN of a large LHS set, without the need of gathering and sorting anything.

The advantages of a virtual slave Thor configuration for ECL code using LOOKUP/SMART JOIN is that in effect it will have N times as much memory before it fails or fails over in the SMART JOIN case.

It is also much quicker; instead of broadcasting the local RHS to N slave process per node; it only has to communicate it to one. That one slave can share the same table and same HT with the other virtual slaves directly.

Key advantages of LOOKUP/SMART JOIN in a Virtual Slave Thor Setup:

- N times as much memory available for RHS. In other words, the RHS can be N times bigger before failing or failing over in SMART JOIN case. (In the illustrated example, the JOIN would have 4 times as much memory available)
- Significantly less communication of row data – equals significantly faster processing for larger RHS sets.

Huge Pages

Linux uses pages as its basic units of memory. Your system may run faster and benefit from huge page support. Huge pages of the appropriate type and size need to be allocated from the operating system. Almost all current Linux systems are set up with Transparent Huge Pages (THP) available by default.

Thor, Roxie, and ECL Agent clusters all have options in the configuration to enable huge page support. The Transparent Huge Pages are enabled for Thor, Roxie, and ECL Agent clusters in the default HPCC environment. Thor clusters can stand to benefit more from huge pages than can Roxie.

You can check the file `/sys/kernel/mm/transparent_hugepage/enabled` to see what your OS setting is. With THP you do not have to explicitly set a size. If your system is not configured to use THP, then you may want to implement Huge Pages.

Setting up Huge Pages

To set up huge page support, consult your OS documentation and determine how to enable huge page support. For example, the administrator can allocate persistent huge pages (for the appropriate OS) on the kernel boot command line by specifying the "hugepages=N" parameter at boot. With huge pages you also need to explicitly allocate the size.

In HPCC, there are three places in the configuration manager to set the attributes to use Huge Pages.

There are attributes in each component, in the ECL Agent attributes, in Roxie attributes, and in Thor attributes. In each component there are two values:

```
heapUseHugePages  
heapUseTransparentHugePages
```

Enable Huge Pages in your operating system, then configure HPCC for the component(s) you wish.

Capacity Planning

Roxie clusters are disk-based High Performance Computing Clusters (HPCC) , typically using indexed files. A cluster is capable of storing and manipulating as much data as its combined hard drive space; however, this does not produce optimal performance.

For maximum performance, you should configure your cluster so slave nodes perform most jobs in memory.

For example, if a query uses three data files with a combined file size of 60 GB, a 40-channel cluster is a good size, while a 60-channel is probably better.

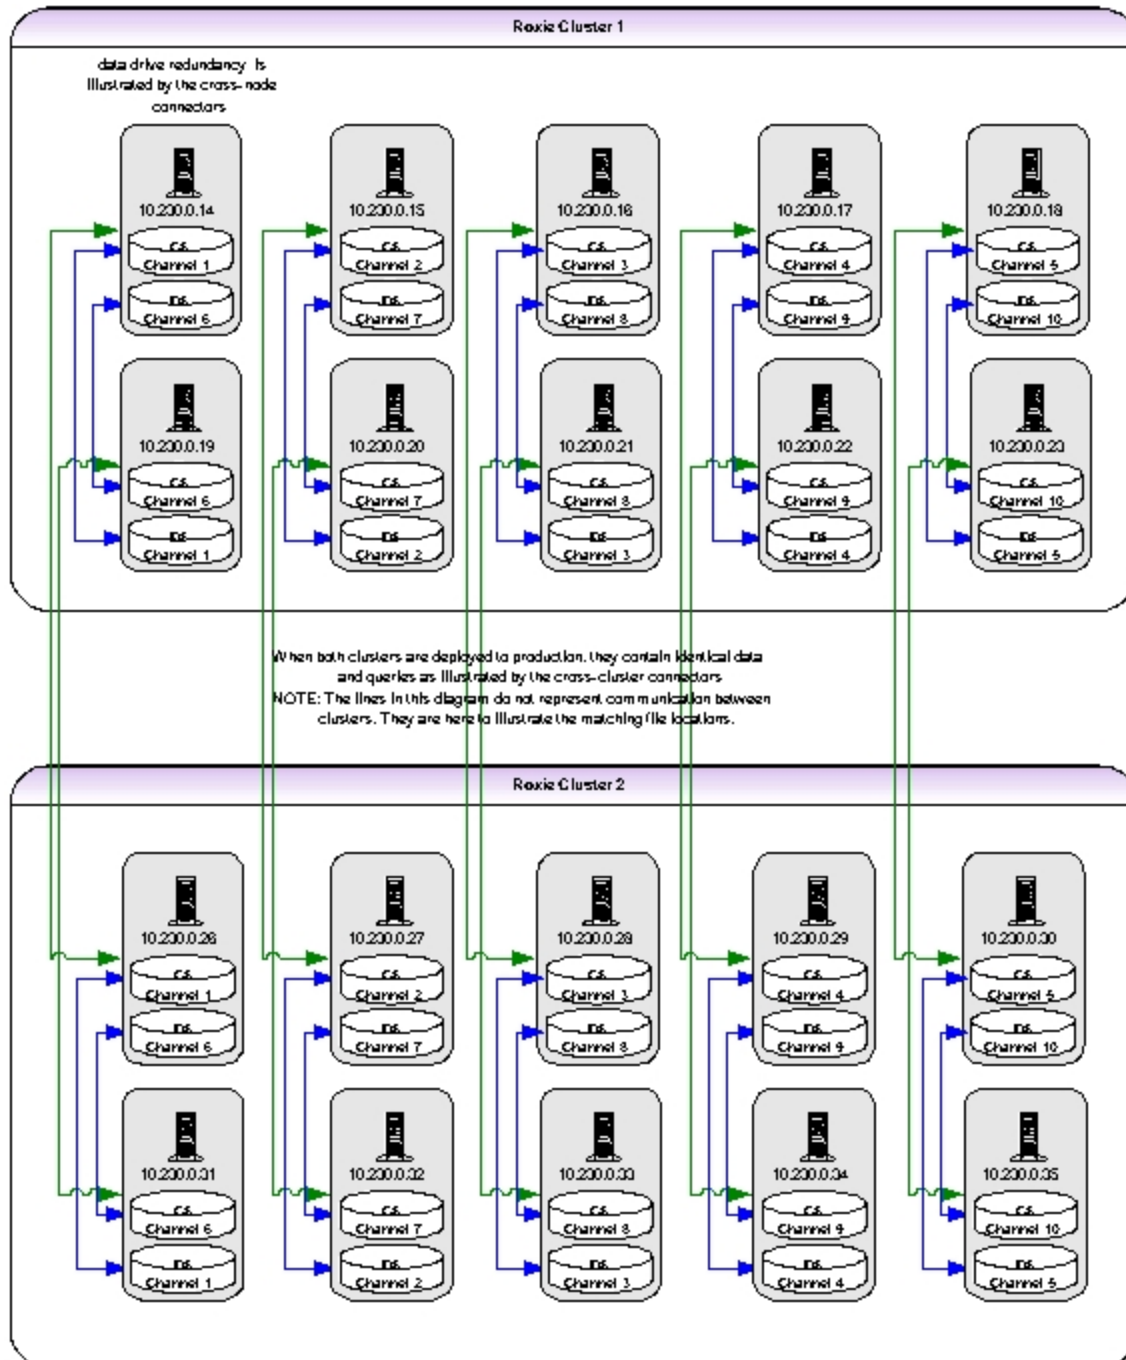
Another consideration is the size of the Thor cluster creating the data files and index files to be loaded. Your target Roxie cluster should be the same size as the Thor on which the data and index files are created or a number evenly divisible by the size of your Roxie cluster. For example, a 100-way Thor to a 20-way Roxie would be acceptable.

This is due to the manner in which data is loaded and processed by Roxie slaves. If data is copied to slave nodes, the file parts are directly copied from source location to target locations. They are NOT split or resized to fit a different sized cluster. Therefore, if you load 50 file parts onto a 40-channel cluster, part one goes to channel one, part two to channel two, etc. Parts 41-50 start at the top again so that part 41 goes to channel 1, and part 42 goes to channel 2, etc. The result is an unevenly distributed workload and would result in reduced performance. A cluster will only perform as fast as its slowest node.

The final consideration is the number of Server processes in a cluster. Each slave must also be a Server, but you can dedicate additional nodes to be only Server processes. This is useful for queries that require processing on the Server after results are returned from slaves. Those Server-intensive queries could be sent only to dedicated Server IP addresses so the load is removed from nodes acting as both Server and slave.

Configuring the Channels

In the illustration below, the nodes are configured using an N+5 scheme to share channels. Channels can be configured in many ways, this is one example.



In this depiction, each enclosure holds five Roxie slave blades (a row of servers in the picture). We will use this example for the rest of this manual.

Sample Sizings

This section illustrates sample system sizings for various work environments. Unlike system requirements, the following samples are suggestions for setting up your system for various operating conditions.

Sample Sizing for High Data volume (Typical)

The most typical scenario for HPCC is utilizing it with a high volume of data. This suggested sample sizing would be appropriate for a site with large volumes of data. A good policy is to set the Thor size to 4 times the source data on your HPCC. Typically, Roxie would be about ¼ the size of Thor. This is because the data is compressed and the system does not hold any transient data in Roxie. Remember that you do not want the number of Roxie nodes to exceed the number of Thor nodes.

High Data Thor sizing considerations

Each Thor node can hold about 2.5 TB of data (MAX), so plan for the number of Thor nodes accordingly for your data.

If possible, SAS drives for both Thor and Roxie as they almost equal to SATA drives now. If not for both, get SAS drives at least for your Roxie cluster.

Thor replicates data and is typically configured for two copies.

High Data Roxie sizing considerations

Roxie keeps most of its data in memory, so you should allocate plenty of memory for Roxie. Calculate the approximate size of your data, and allocate appropriately. You should either increase the number of nodes, or increase the amount of memory.

A good practice is to allocate a Dali for every Roxie cluster.

Roxie should have a mirror. This is useful, when you need to update data. You update the mirror then make that primary and bring the other one down. This is a good practice but not really a necessity except in the case of high availability.

Sample Sizing for Heavy Processing on Low Data Volume

The following section provides some sample sizing for heavy processing with approximately the amount of data indicated.

750 GB of Raw Data

Thor = 3 (slaves) + 2 (management) = 5 Nodes

Roxie = 3 (agents) + 1 (Dali) = 4 Nodes (This will mean that the environment will be down during query deployment)

Spares = 2

Total = 13 nodes

1250 GB of Raw Data

Thor = 6 (slaves) + 2 (management) = 8 Nodes

Roxie = 4 (agents) + 1 (Dali) = 5 Nodes (This will mean that the environment will be down during query deployment)

Spares = 2

Total = 17 nodes

2000 GB of Raw Data

Thor = 8 (slaves) + 3 (management) = 11 Nodes

Roxie = 4 (agents) + 1 (Dali) = 5 Nodes (This will mean that the environment will be down during query deployment)

Spares = 2

Total = 20 nodes

3500 GB of Raw Data

Thor = 12 (slaves) + 5 (management) = 17 Nodes

Roxie = 6 (agents) + 1 (Dali) = 7 Nodes (This will mean that the environment will be down during query deployment)

Spares = 2

Total = 28 nodes

System Resources

There are additional resources available for the HPCC System.

HPCC 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, plug-ins, support, documentation, and more. This is where you can find resources useful for running and maintaining HPCC on the web portal.

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

Additional Resources

Additional help with HPCC and Learning ECL is also available. There are online courses available. Go to :

<https://learn.lexisnexus.com/hpcc>

You may need to register for the site. There are several training videos and other very helpful information.