

HPCC Systems™

Installing & Running the HPCC Platform

Boca Raton Documentation Team

Installing & Running the HPCC Platform

Boca Raton Documentation Team

Copyright © 2012 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. 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.

February 2012 Version 3.6.2.1

Welcome	4
Hardware and Software Requirements	5
Network Switch	5
Load Balancer	8
Nodes-Hardware	10
Nodes-Software	11
User Workstation Requirements	13
HPCC Installation and Startup	14
Initial Setup-Single Node	16
Configuring a Multi-Node System	27
Starting and Stopping	33
Configuring HPCC to use LDAP Authentication	35
User Security Maintenance	42
More Examples	58
ECL Example: Anagram1	58
Roxie Example: Anagram2	61
Next Steps	75
Appendix	76
Example Scripts	76
Uninstalling the HPCC Platform	80
Helper Applications	81
hpcc-init	82
Running the ECL IDE under WINE	84

Welcome

These instructions will guide you through installing and running the HPCC¹ Community Edition on a single node to start and then optionally, expand it to a larger cluster of nodes.

The HPCC Thor technology is designed to effectively process, analyze, and find links and associations within high volumes of complex data. This can detect non-obvious relationships, scale to support petabytes of data, and is significantly faster than competing technologies while requiring less hardware and resources.

The HPCC Roxie technology - also known as the Rapid Data Delivery Engine or RDDE - uses a combination of technologies and techniques that produce extremely fast throughput for queries on indexed data.

This translates into better quality answers in less time so that organizations can cope with massive data and efficiently turn information into knowledge.



We suggest reading this document in its entirety before beginning. The entire process can take an hour or two, depending on your download speed.

¹High Performance Computing Cluster (HPCC) is a massively parallel processing computing platform that solves Big Data problems. See <http://hpccsystems.com/Why-HPCC/How-it-works> for more details.

Hardware and Software Requirements

The following section describes the various hardware and software required in order to run the HPCC.

Network Switch

A significant component of HPCC is the infrastructure it runs on, specifically the switch.

Switch requirements

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

Small: For a very small test system, almost any gigabit switch will suffice. These are inexpensive and readily available in six to 20-port models.

Figure 1. 1 GigE 8-port Switch



Medium: For medium sized (10-48 node) systems, we recommend using a Force10 s25, s50, s55, or s60 switch

Figure 2. Force10 S55 48-port Network Switch



Large: For large (48-350 node) system, the Force10 c150 or c300 are good choices.

Figure 3. Force 10 c150



Very Large: For very large (more than 300 nodes) system, the Force10 e600 or e1200 are good choices.

Figure 4. Force 10 e600 and e1200



Switch additional recommended features

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

Load Balancer

In order to take full advantage of a Roxie cluster, a load balancer is required. 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.

We recommend the Web Accelerator product line from F5 Networks. See <http://www.f5.com/pdf/products/big-ip-webaccelerator-ds.pdf> for more information.

Figure 5. F5 Load Balancers



Load Balancer Requirements

Minimum requirements

- Throughput: 1Gbps 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.

Nodes-Hardware

The HPCC can run as a single node system or a multi node system.

These hardware recommendations are intended for a multi-node production system. A test system can use less stringent specifications. Also, while it is easier to manage a system where all nodes are identical, this is not required. However, it is important to note that your system will only run as fast as its slowest node.

Node mandatory requirements

- Pentium 4 or newer CPU
- 64-bit
- 4GB RAM
- Two Hard Drives (with sufficient free space to handle the size of the data you plan to process)
- 1 GigE network interface

Node recommended specifications

- Nehalem Core i7 CPU
- 64-bit
- 4 GB RAM (or more)
- 1 GigE network interface
- PXE boot support in BIOS

PXE boot support is recommended so you can manage OS, packages, and other settings when you have a large system

- Optionally IPMI and KVM over IP support

For Roxie nodes:

- Two 10K RPM (or faster) SAS Hard Drives

Typically, drive speed is the priority for Roxie nodes

For Thor nodes:

- Two 7200K RPM (or faster) SATA Hard Drives (Thor)
- Optionally 3 or more hard drives can be configured in a RAID 5 container for increased performance and availability

Typically, drive capacity is the priority for Thor nodes

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

You will need one of the following:

- 64-bit LINUX CentOS 5.x / Red Hat Enterprise Linux
- 64-bit Ubuntu 10.04 LTS, 11.04, or 11.10
- 64-bit OpenSuSe 11.3 or 11.4
- 64-bit Debian 6.x (Squeeze)

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.

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**".

User Workstation Requirements

- Running the HPCC platform requires communication from your user workstation with a browser to the HPCC. You will use it to access ECL Watch—a Web-based interface to your HPCC system. ECL Watch enables you to examine and manage many aspects of the HPCC and allows you to see information about jobs you run, data files, and system metrics.

Use one of the supported web browsers with Javascript enabled.

- Internet Explorer® 7 (or later)
- Firefox™ 3.0 (or later.)
- Google Chrome 10 (or later)

If browser security is set to **High**, you should add ECLWatch as a Trusted Site to allow Javascript execution.

- Install the ECL IDE

The ECL IDE (Integrated Development Environment) is the tool used to create queries into your data and ECL files with which to build your queries.

From the ECLWatch web page, download the Windows install set. If the link is not visible, either follow the link to the HPCC System's portal or install the Optional Packages.

You can reach this page using the following URL:

<http://nnn.nnn.nnn.nnn:8010>, where nnn.nnn.nnn.nnn is your node's IP address.

The ECL IDE was designed to run on Windows machines. See the appendix for instructions on running on Linux workstations using Wine.

- Microsoft VS 2008 C++ compiler (either Express or Professional edition). This is needed if you are running Windows and want to compile queries locally. This allows you to compile and run ECL code on your Windows workstation.
- GCC. This is needed if you are running under Linux and want to compile queries locally on a standalone Linux machine, (although it may already be available to you since it usually comes with the operating system).

HPCC Installation and Startup

Follow these steps to install the packages and start components in a single-node configuration to begin. Once it is successfully installed, you will use the Configuration Manager to customize or expand your system.

Configuration Manager is the utility with which we configure the HPCC platform. It is run on your Linux Server and you access its interface using a browser.

Figure 6. System Overview: Thor

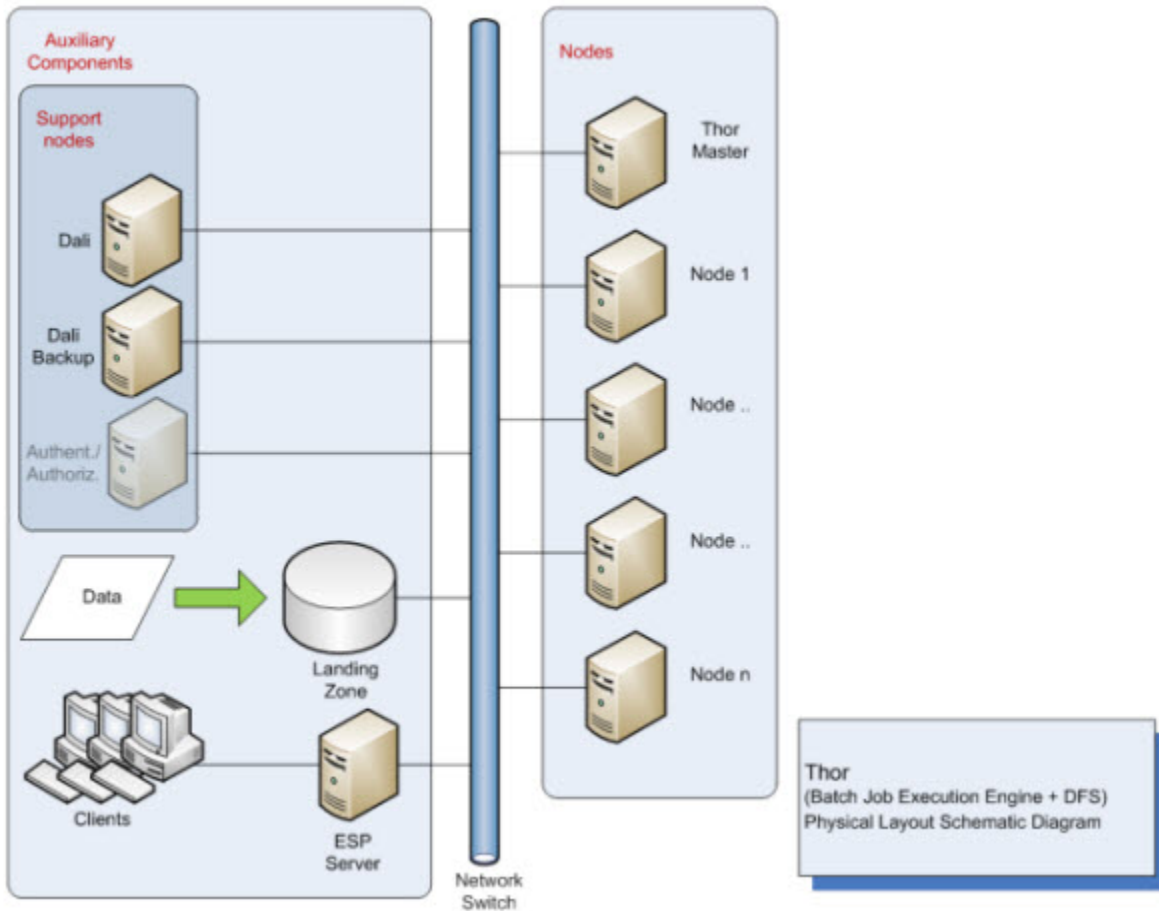



Figure 7. System Overview: Roxie



Initial Setup-Single Node

This section covers installing the HPCC on a single node. This will enable the HPCC system to operate successfully; however, the real strength of the HPCC is when it is run in a multi-node environment and can leverage the ability to perform operations using Massively Parallel Processing (MPP).

In addition, on a production system, you would dedicate one or more nodes to each server process. See the *Using Configuration Manager* manual for more details.

Installing the Package

The installation and package that you download is different depending on the operating system you plan to use. The installation packages will fail to install if their dependencies are missing from the target system.

Packages are available from the HPCC Systems website: <http://hpccsystems.com/download/free-community-edition>

To install the package, follow the appropriate installation instructions:

Centos/Red Hat/SuSe

Install RPM with the -Uvh switch.

This is the upgrade command and will perform an automatic upgrade if a previous version is installed or it will just install fresh if no other version has been installed.

```
sudo rpm -Uvh <rpm file name>
```

NOTE : For ANY version of SuSe you must set a password for the hpcc user on all nodes. One way to do this is to issue the following command:

```
sudo passwd hpcc
```

Be sure to set the password on ALL nodes.

Ubuntu/Debian

For Ubuntu installations a Debian package is provided. To install the package, use:

```
sudo dpkg -i <deb filename>
```

Installing the Optional Packages

There are three optional packages you can install which will enable resources to be available from your ESP server's ECL Watch page. These resources include:

- Documentation
- ECL IDE
- Graph Control

We recommend installing these packages on your ESP Server node.

You can find these packages at <http://hpccsystems.lexisnexis.com/download/free-community-edition>

To install the package, follow the appropriate installation instructions:

Centos/Red Hat/SuSe

Install RPM with the -Uvh switch.

This is the upgrade command and will perform an automatic upgrade if a previous version is installed or it will just install fresh if no other version has been installed.

```
sudo rpm -Uvh <rpm file name>
```

Ubuntu/Debian

For Ubuntu installations a Debian package is provided. To install the package, use:

```
sudo dpkg -i <deb filename>
```

Initial Startup

1. Start the system using the default configuration.

Centos/Red Hat/SuSe

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

Ubuntu

```
sudo service hpcc-init start
```

Debian 6 (Squeeze)

```
sudo /etc/init.d/hpcc-init start
```



```
root@node219008:~# sudo /sbin/service hpcc-init start
Starting mydali... [ OK ]
Starting mydafilesrv... [ OK ]
Starting mydfuserver... [ OK ]
Starting myeclagent... [ OK ]
Starting myeclccserver... [ OK ]
Starting myesp... [ OK ]
Starting myroxie... [ OK ]
Starting mysasha... [ OK ]
Starting mythor... [ OK ]
[root@node219008 ~]#
```



There are log files for each component in directories below **/var/log/HPCCSystems** (default location). If any component fails to start, these logs can help in troubleshooting.

Running an ECL Query on your Single-Node System

The single node system is running, and you can now create and run some ECL¹ code using either ECL IDE the command line ECL compiler, or the ECLPlus tool.

Install the ECL IDE and HPCC Client Tools

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



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

2. From ECL Watch page, click on the **Resources** link in the menu on the left side.

Figure 8. ECL Watch Resource Page

The screenshot shows the EclWatch web interface. On the left is a navigation menu for 'HPCC Systems' with categories like Clusters, ECL Workunits, Query Sets, Topology, DFU Workunits, DFU Files, Roxie Queries, and Resources. The 'Resources' link is circled in red. The main content area is titled 'EclWatch' and contains a download instruction: 'Click a link below to download a version from your installation. You may visit <http://hpccsystems.com/download> for other versions.' Below this are three tables:

- graphcontrol:** A table with columns Name, Version, and Description. It lists 'Graph Control Installer' (Version 20110523, Description: Windows installer for Graph Control).
- documentation:** A table with columns Name, Version, and Description. It lists various resources like 'HPCC Data Tutorial', 'ECL Language Reference', 'ECL Programmers Guide', 'Six Degrees of Kevin Bacon Example', 'HPCC Client Tools', 'HPCC Data Handling', 'Original Person File Zip', 'Original Person Compressed Tar file', 'HPCC Configuration Manager', and 'Rapid Data Delivery Engine Reference'.
- clienttools:** A table with columns Name, Version, and Description. It lists 'ECL IDE Installer' (Version ver-6_0_1_5, Description: Windows installer for ECL IDE), which is circled in red.

¹Enterprise Control Language (ECL) is a declarative, data centric programming language used to manage all aspects of the massive data joins, sorts, and builds that truly differentiate HPCC (High Performance Computing Cluster) from other technologies in its ability to provide flexible data analysis on a massive scale.

Note: The order of items on the page may vary depending on the order in which optional packages are installed.

If this page is not visible, either follow the link to the HPCC System's portal or install the Optional Packages,

3. Click on the ECL IDE Installer link. When prompted save this file to your PC and then run it—do not run directly from your browser. This is a Windows installer for a Windows application.
4. You can close your browser, if desired.
5. Install the ECL IDE, following the prompts in the installer. Once the ECL IDE is installed successfully, you can proceed.

Running a basic ECL program

Now that the package is installed on your Linux node and ECL IDE is installed on your Windows workstation, you can run your first ECL program. ECL programs may be run locally or remotely. For larger ECL jobs, you will want to target a remote cluster of machines, which may not be running the same operating system as the machine you are working on.

In this section we will use the **command line interface** to the compiler to compile and run ECL code locally.

The ECL compiler (eclcc) installs on to the eclcc server node when a package is installed. This should be in your path, so you can run it from anywhere on the server. It is also installed on a Windows machine when you install the ECL IDE. To compile and run on Windows, you also need the Visual Studio 2008 C++ compiler (see *User Workstation Requirements* for details).

1. Create a file called hello.ecl and type in the following text (including the quotes):

```
output('Hello world');
```

You can either use your favorite editor, or you can use the command line by typing the following

```
echo "output('Hello world');" > hello.ecl
```

2. Compile your program using eclcc by typing the following command:

```
eclcc hello.ecl
```

3. An executable file is created which you can run as follows:

```
# on a Linux machine:  
./a.out  
# on a Windows machine:  
a.out
```

This generates the output "Hello world" (excluding quotes), to the std output, your terminal window in this example. You can redirect or pipe the output to a file or program if you choose. This verifies that the compiler is working properly.

Running remotely using ECLPlus

The **ECLPlus** application accepts command line parameters to send directly to an ECL execution engine. You can use eclplus to control the creation and execution of larger ECL jobs which target a remote system. To compile jobs on a remote system, eclcc is used to create an archive of the ECL code to be compiled, and eclplus is used to submit it to a queue for compilation by the remote compiler server (eclccserver).

To submit a job to the queue using ECL Plus, make sure the HPCC has been started and use the following syntax:

```
eclplus @hello.ecl cluster=hthor server=<IP Address of the ESP node>:8010
```

or

```
eclplus @hello.ecl cluster=hthor server=.
```

Where "." indicates the IP of the current box.

The workunit² ID, status, and result are shown at the command line.

²A Workunit is a record of a task submitted to an HPCC. It contains an identifier--workunit ID, the ECL code, results, and other information about the job.

View the full details of the workunit using the ECL Watch interface for your HPCC at this location <http://nnn.nnn.nnn.nnn:8010>, where nnn.nnn.nnn.nnn is the IP of your node. Either search for the workunit using the workunit ID or select ECL Workunits/Browse and find your workunit in the list provided.

Setting up an **ecplus.ini** file makes running a workunit a little easier when you want to use the same settings every time you submit a workunit in this way. See the *HPCC Client Tools* manual for details.

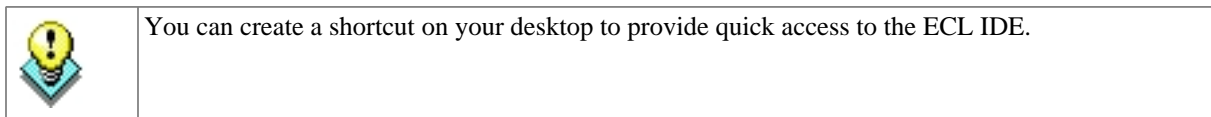
If your ECL is more complex than a single source file, you can use the eclcc compiler locally to create an archive to be sent to the eclccServer:

```
eclcc hello.ecl -E | eclplus @ cluster=thor server=<IP Address of the ESP>:8010
```

The cluster parameter must name a valid target cluster name as listed in your environment's topology section.

Running a basic ECL program from the ECL IDE

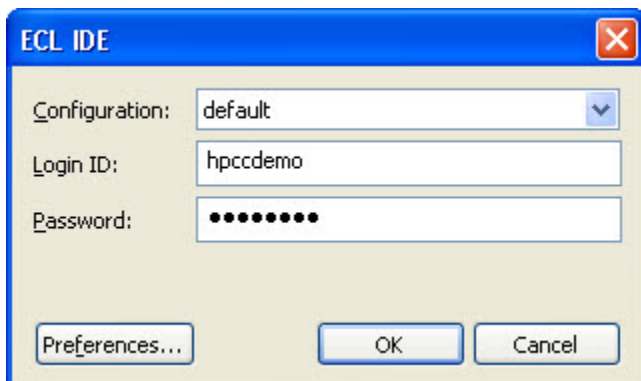
1. Open the ECL IDE on your Windows workstation, from your start menu. (Start >> All Programs >> HPCCSystems >> ECL IDE).



2. Enter the **Login ID** and **Password** provided in the Login dialog.

Login ID	hpccdemo
Password	hpccdemo

Figure 9. Login Window



3. Open a new **Builder Window** (CTRL+N) and write the following code:

```
OUTPUT('Hello World');
```

This could also be written as:


```
'Hello World';
```

In the second program listing, the OUTPUT keyword is omitted. This is possible because the language is declarative and the OUTPUT action is implied.

4. Select **thor** as your target cluster.

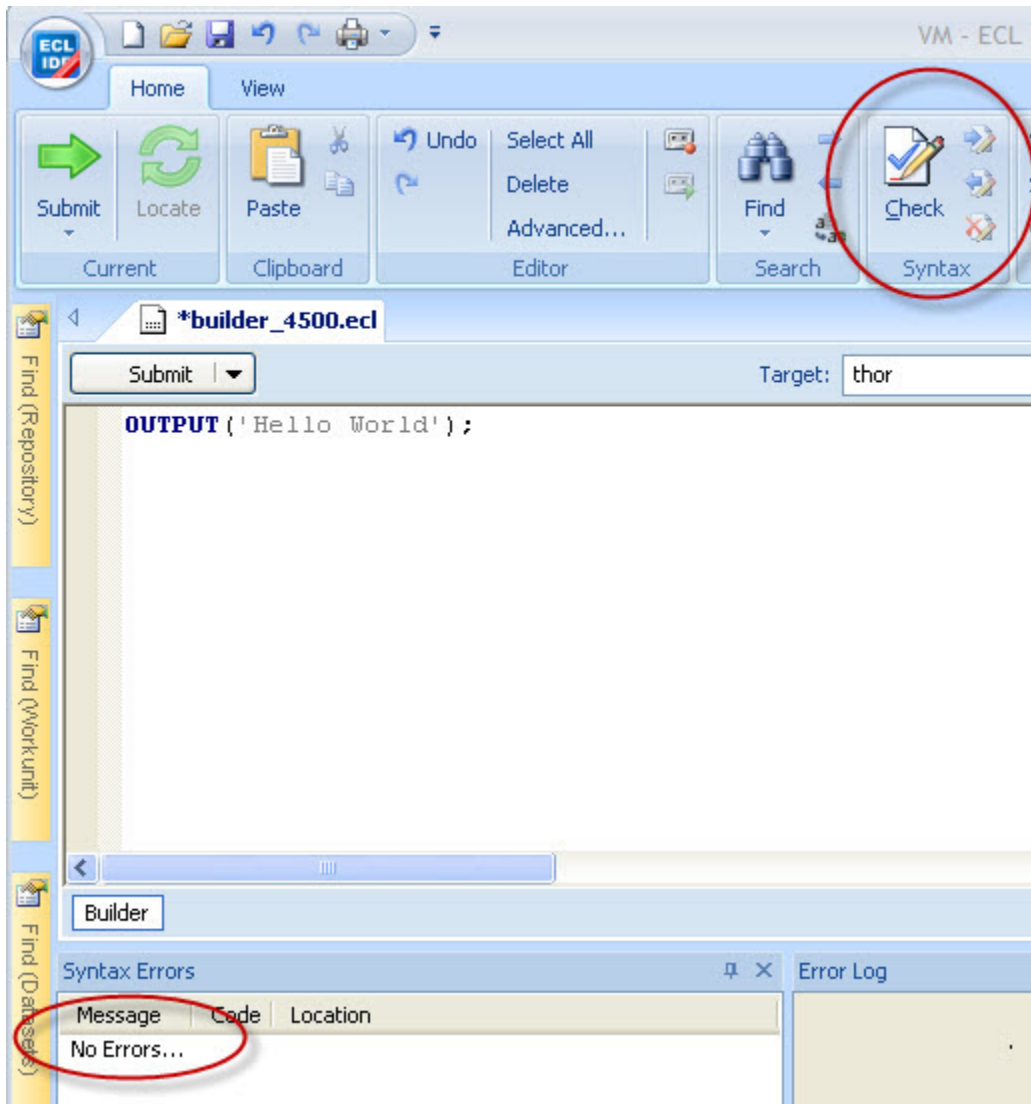
Thor is the Data Refinery component of your HPCC. It is a disk based massively parallel computer cluster, optimized for sorting, manipulating, and transforming massive data.

Figure 10. Select target



5. Press the syntax check button on the main toolbar (or press F7).

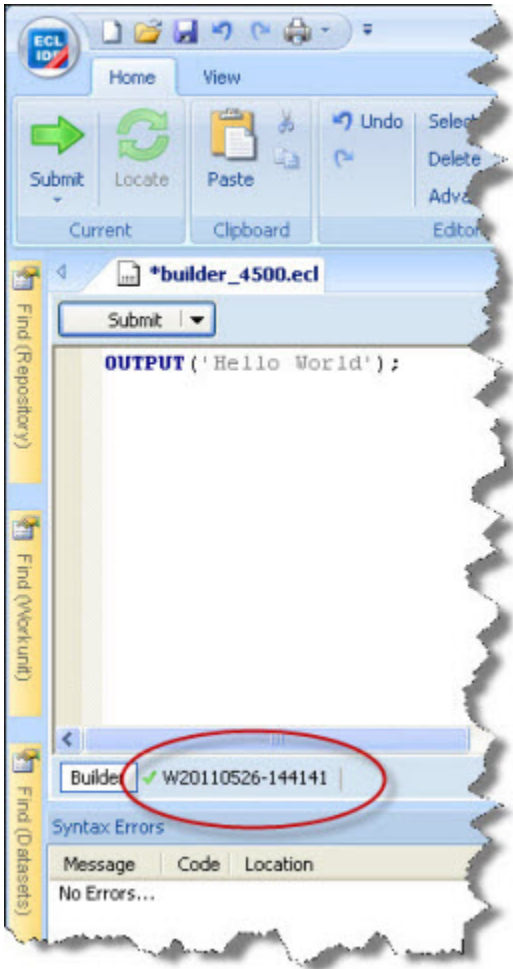
Figure 11. Syntax Check



A successful syntax check displays the "No Errors" message.

6. Press the **Go** button (or press ctrl+enter).

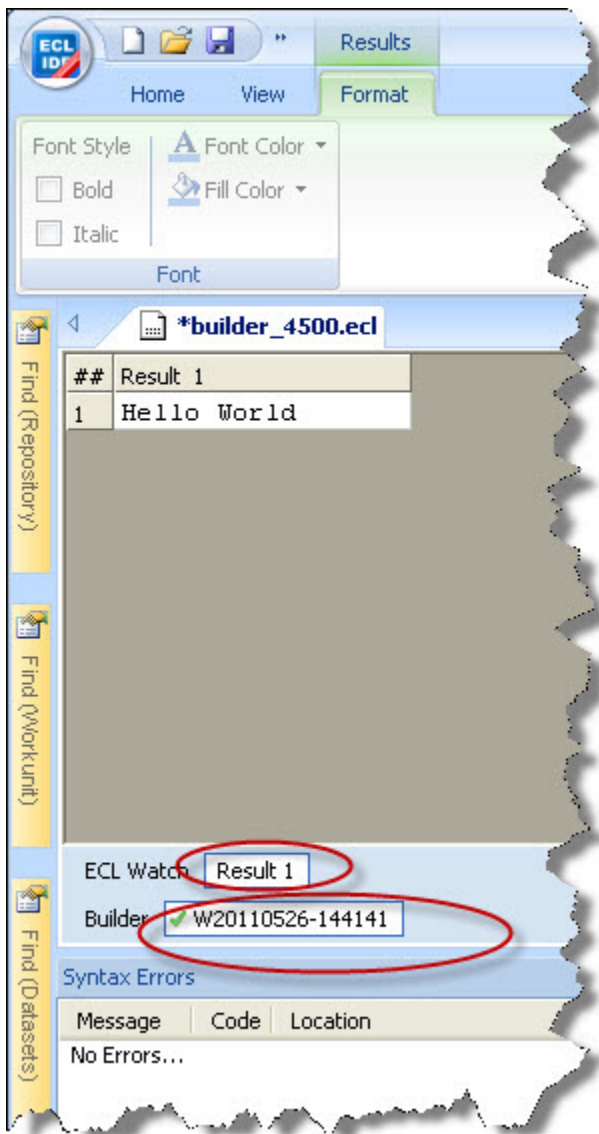
Figure 12. Completed job



The green check mark indicates successful completion.

7. Click on the workunit number tab and then on the Result 1 tab to see the output.

Figure 13. Completed job output



Configuring a Multi-Node System

While the single-node system is fully-functional, it does not take advantage of the true power of an HPCC—the ability to perform operations using Massively Parallel Processing (MPP). This section provides the steps to expand your single-node system into a multi-node system using the Configuration Manager Wizard.

Using the Configuration Manager Wizard

This section details reconfiguring a system to use multiple nodes. Before you start this section, you must have already downloaded the correct packages for your distro from the HPCC Systems website: <http://hpccsystems.com/download/free-community-edition>

1. If it is running, stop the HPCC system, using this command:

Centos/Red Hat/SuSe


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

Ubuntu

```
sudo service hpcc-init stop
```

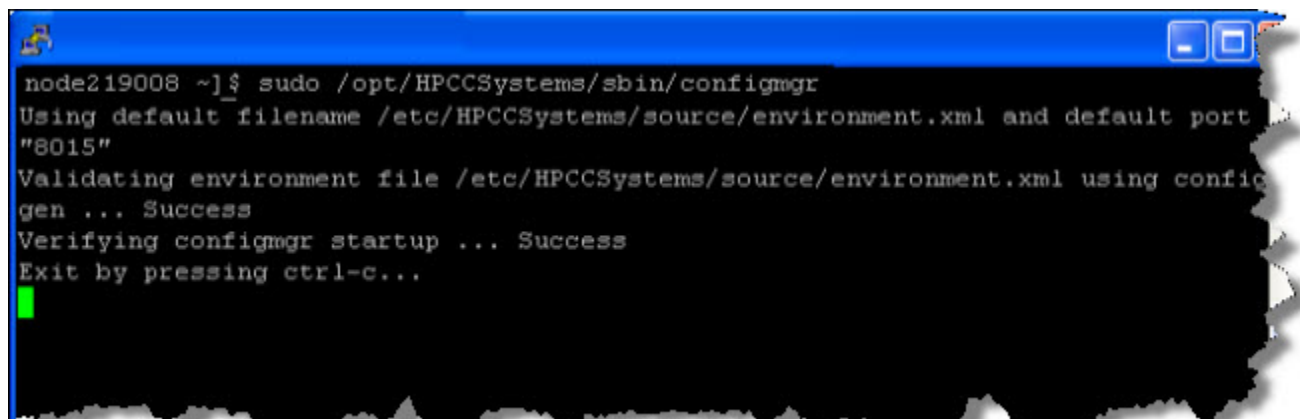
Debian 6 (Squeeze)

```
sudo /etc/init.d/hpcc-init stop
```

	You can use this command to confirm HPCC processes are stopped (on Centos/Red Hat/SuSe):
	<pre>sudo /sbin/service hpcc-init status</pre>
	For Ubuntu
	<pre>sudo service hpcc-init status</pre>
	For Debian 6 (Squeeze)
	<pre>sudo /etc/init.d/hpcc-init status</pre>

2. Start the Configuration Manager service.

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

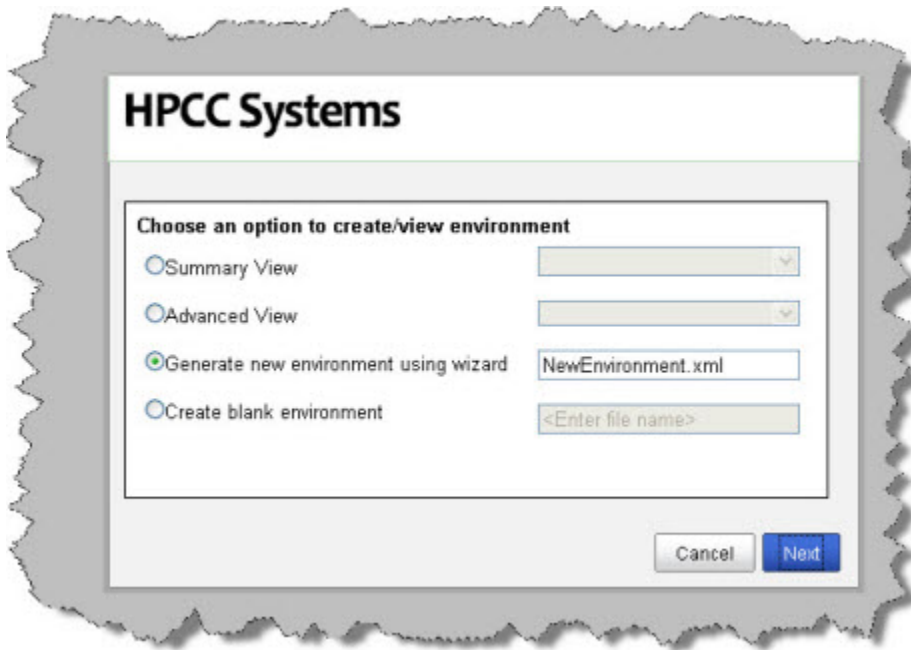


```
node219008 ~]$ sudo /opt/HPCCSystems/sbin/configmgr
Using default filename /etc/HPCCSystems/source/environment.xml and default port
"8015"
Validating environment file /etc/HPCCSystems/source/environment.xml using config
gen ... Success
Verifying configmgr startup ... Success
Exit by pressing ctrl-c...
```

3. Leave this window open. You can minimize it, if desired.
4. Using a Web browser, go to the Configuration Manager's interface:

`http://<node ip >:8015`

5. The Configuration Manager startup wizard displays. To use the wizard, select the Generate new environment using wizard button.



6. Provide a name for the environment file.

This will then be the name of the configuration xml. For example, we will name this *NewEnvironment.xml*.

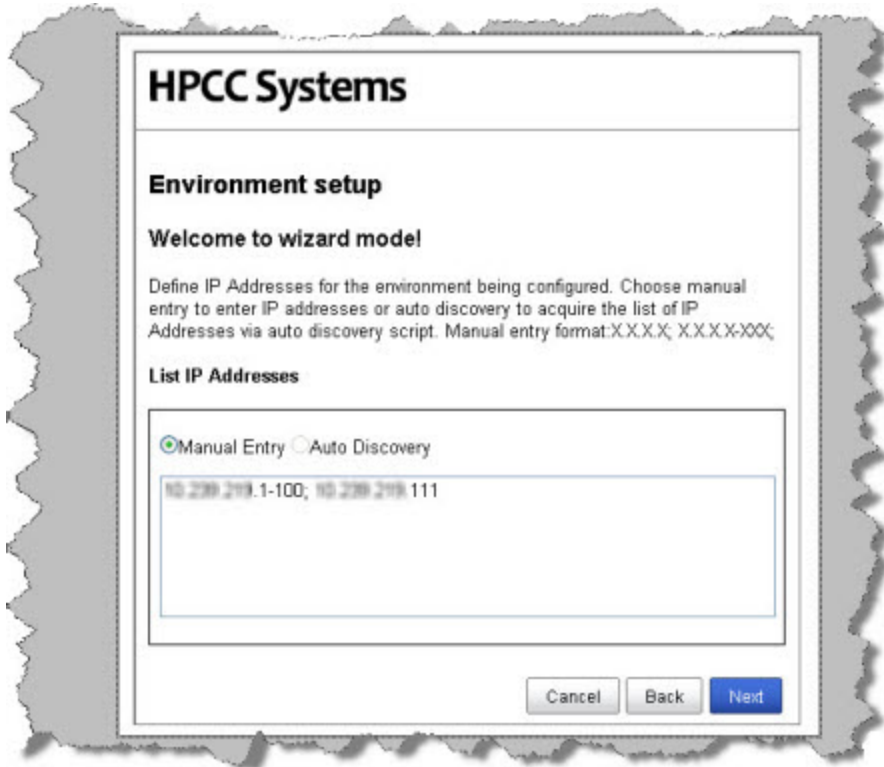
7. Press the **Next** button.

Next you will need to define the IP addresses that your system will use.

8. Enter the all the IP addresses you want to use in this HPCC.

The IP addresses do not need to be contiguous. In the image below, we specified the IP addresses nn.nnn.nnn.1-100 and nn.nnn.nnn.111. These are separated with a semi-colon.

You can specify a range of IPs using a hyphen (for example, NNN.NNN.NNN.1-100). IP Addresses can be specified individually using semi-colon delimiters.



9. Press the **Next** button.

Alternatively, you could find the IP addresses using Auto Discovery by selecting the Auto Discovery button. (only available in Enterprise Edition)

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

10. 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	<input type="text" value="1"/>
Number of nodes for Roxie cluster	<input type="text" value="0"/>
Number of slave nodes for Thor cluster (A master node will be automatically added to the cluster)	<input type="text" value="1"/>
Number of Thor slaves per node (default 1)	<input type="text" value="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.
- 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)

11. Press the Next button

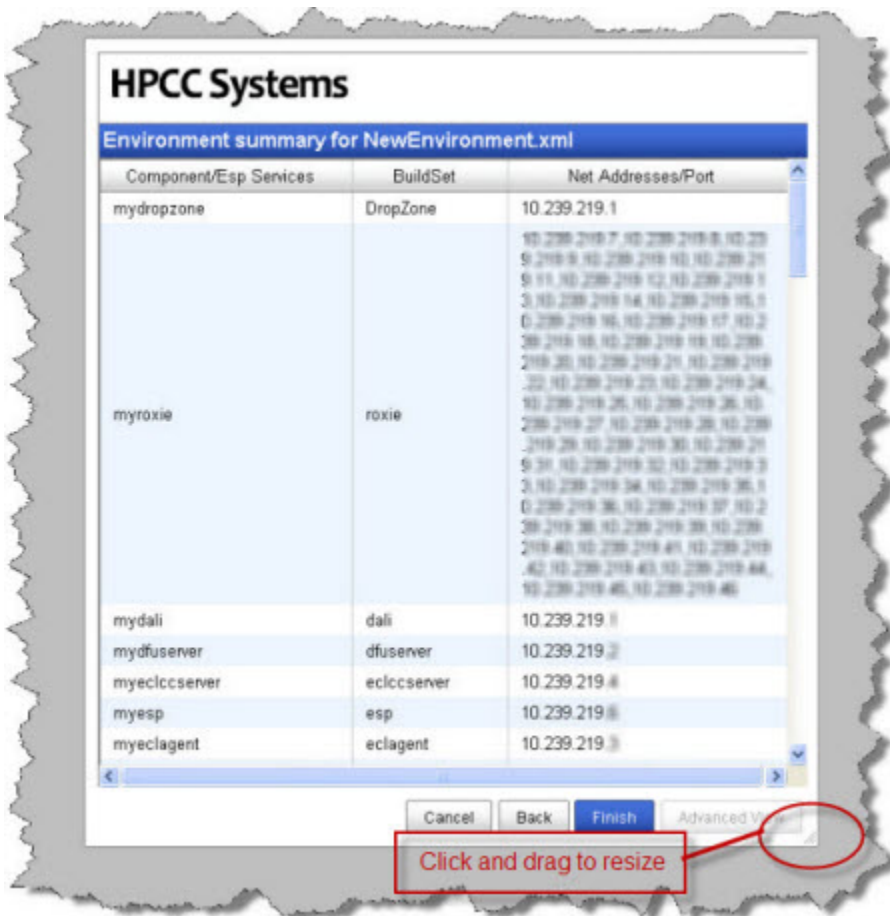
The Environment Summary displays.

Installing & Running the HPPC Platform HPPC Installation and Startup

12. Click on **Finish** to accept these values. This saves the file.



Keep in mind, that your HPPC 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.



You can resize the Environment Summary by clicking and dragging the lower right corner.

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



At this point the system has created a file named NewEnvironment.xml in the `/etc/HPPCSysyems/source` directory

14. Stop the Configuration Manager in the terminal where you started it by pressing CTRL-C.



Be sure system is stopped before attempting to move the environment.xml file.

15. Copy the NewEnvironment.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
```



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**. The environment.xml file **MUST** be owned by the **hpcc** user.

16. If you have added new machines to the cluster, you need to copy and install the HPCC package onto all nodes, and generate and clone the SSH keys. This can be done using the install-cluster.sh script which is provided with HPCC. Use the following command:

```
/opt/HPCCSystems/sbin/install-cluster.sh -k <package-file-name>
```

Where <package-file-name> is the name of the package file that you want to install on every node - this will be in the form hpccsystems-platform-xxx-n.n.nmmn.rpm (or .deb) depending on the version and distro. More details including other options that may be used with this command are included in the appendix.

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

You may want to create a script to push out 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 -u hpcc /opt/HPCCSystems/sbin/hpcc-push.sh <sourcefile> <destinationfile>
```

See the appendix for more information on using this script.

18. Restart the HPCC system on **every** node. The following command starts the HPCC system on an individual node:

Centos/Red Hat/SuSe

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

Ubuntu

```
sudo service hpcc-init start
```

Debian 6 (Squeeze)

```
sudo /etc/init.d/hpcc-init start
```

You may want to create a script to push this command out to every node. A sample script is provided with HPCC. Use the following command to start HPCC on all nodes:

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

This script can also be used to stop HPCC on all nodes and to stop and start individual components on all nodes. See the appendix for more details.

Starting and Stopping

Start, Stop, Restart the System

Once you have your system environment established, the **init** system can be used to start, stop, or restart components.

The following commands can be used:

To start the system:

Centos/Red Hat/SuSe

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

Ubuntu

```
sudo service hpcc-init start
```

Debian 6 (Squeeze)

```
sudo /etc/init.d/hpcc-init start
```

To stop the system:

Centos/Red Hat/SuSe

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

Ubuntu

```
sudo service hpcc-init stop
```

Debian 6 (Squeeze)

```
sudo /etc/init.d/hpcc-init stop
```



You can use a script to start or stop multiple nodes in the system. See *Example Scripts* in the Appendix section for samples.

Start or Stop Single Components

To start or stop a single component, you can use the **-c** flag in the init system as follows.

Centos/Red Hat/SuSe

```
sudo /sbin/service hpcc-init -c <component name> <command>
```

Ubuntu

```
sudo service hpcc-init -c <component name> <command>
```

Debian 6 (Squeeze)

```
sudo /etc/init.d/hpcc-init -c <component name> <command>
```



To stop dafilesrv (a helper application), you must use this command: `sudo /sbin/service dafilesrv stop`. See Helper Applications for details.

Start or Stop Configuration Manager

Configure the system as desired using Configuration Manager.

1. If the system is running, stop the HPCC system, using this command on **every** node:

Centos/Red Hat/SuSe

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

Ubuntu

```
sudo service hpc-init stop
```

Debian 6 (Squeeze)

```
sudo /etc/init.d/hpc-init stop
```

2. Start the Configuration Manager service on one node (usually the first node is considered the head node and is used for this task, but this is 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
```

Configuring HPCC to use LDAP Authentication

This section details the steps to connect your HPCC platform to an existing LDAP Server to enable user security.



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

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 -u hpcc /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.

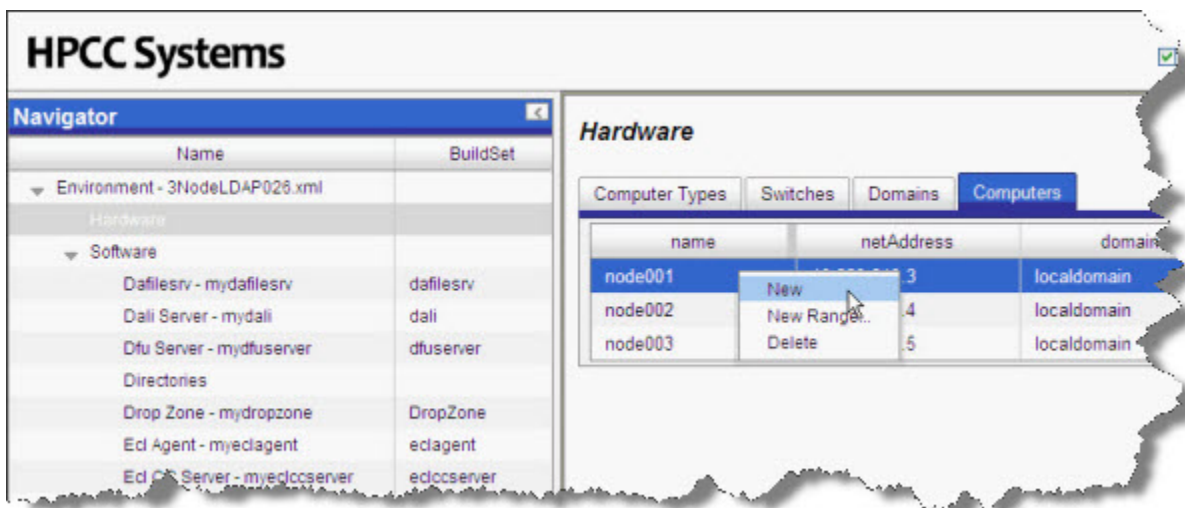
Note: Configuration Manager **never** works on the active configuration file. After you finish editing you will have to copy the `environment.xml` to the active location and push it out to all nodes.

Modifying the configuration

Follow the steps below to modify your configuration.

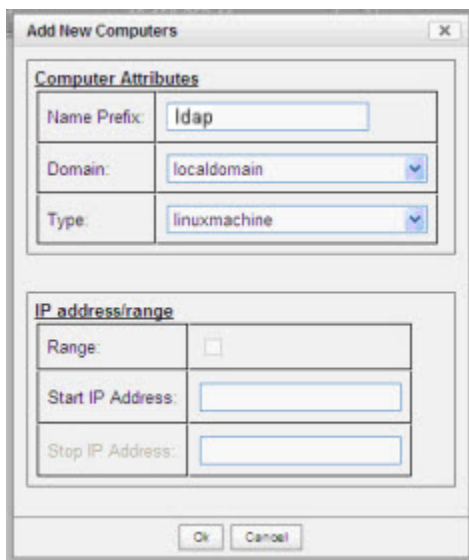
1. Check the box for **Write Access**.
2. From the **Navigator** pane, select **Hardware**.
3. Select the **Computers** tab from the panel on the right.

4. Rt-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**



- 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.

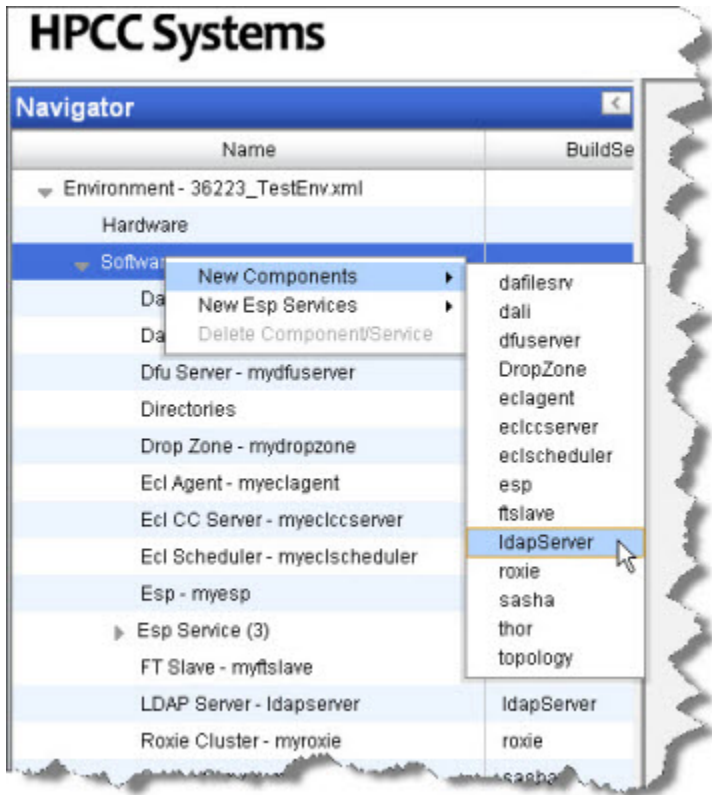
- c. Add the IP address as appropriate for the LDAP server.

- d. Press the **Ok** button.
- e. Click on the disk icon to save.

Adding the IdapServer component

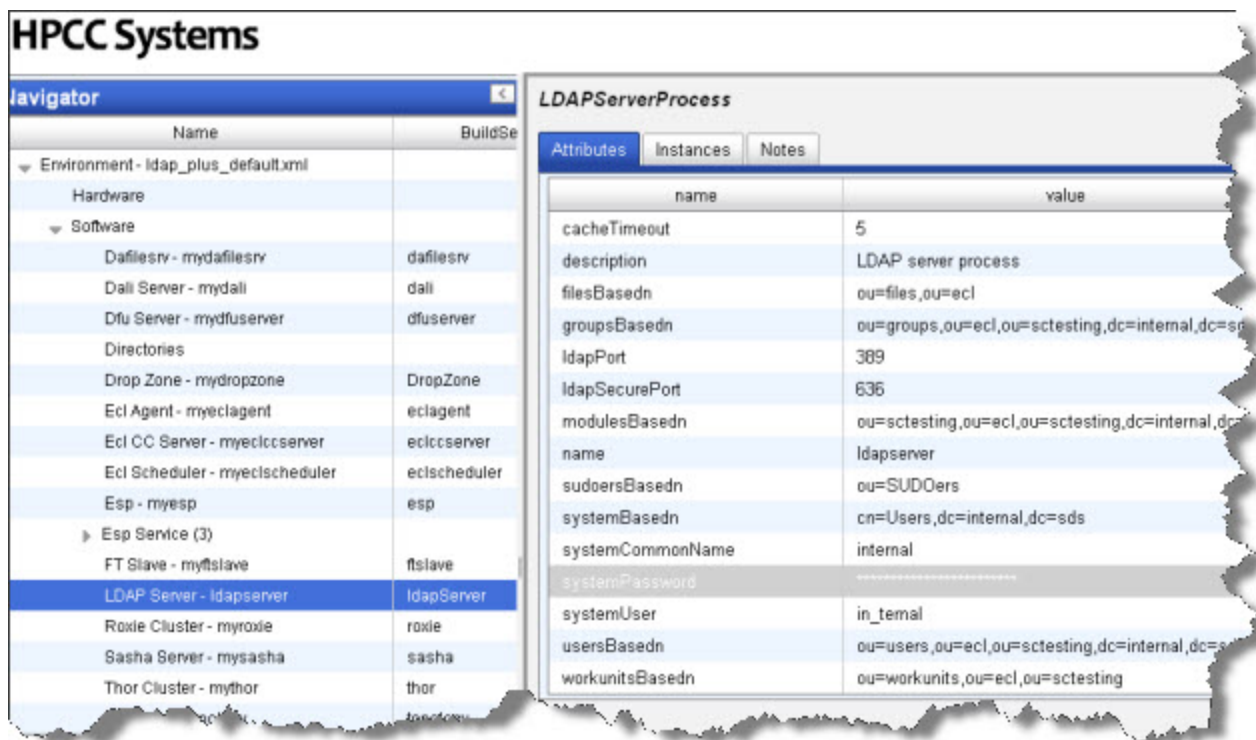
After the LDAP Server node has been added to the Hardware configuration, configure the Software LDAP server definition.

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



Note: The IdapServer 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, Rt-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.

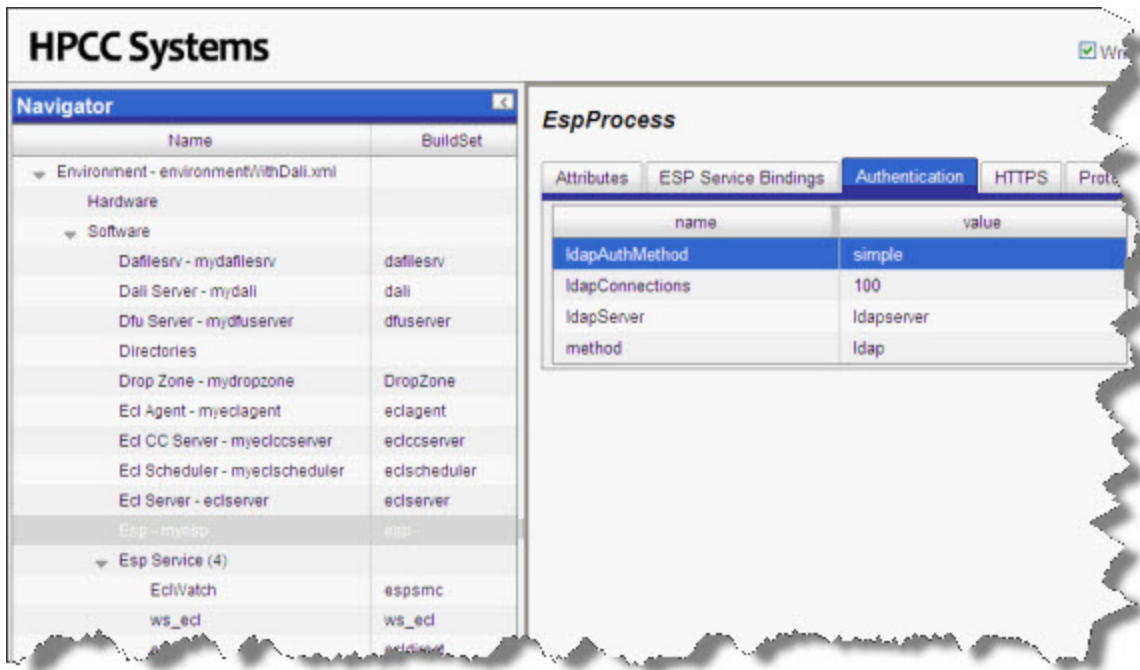
d. Fill in the **Attributes** tab with the appropriate settings from your existing LDAP Server.

e. 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**

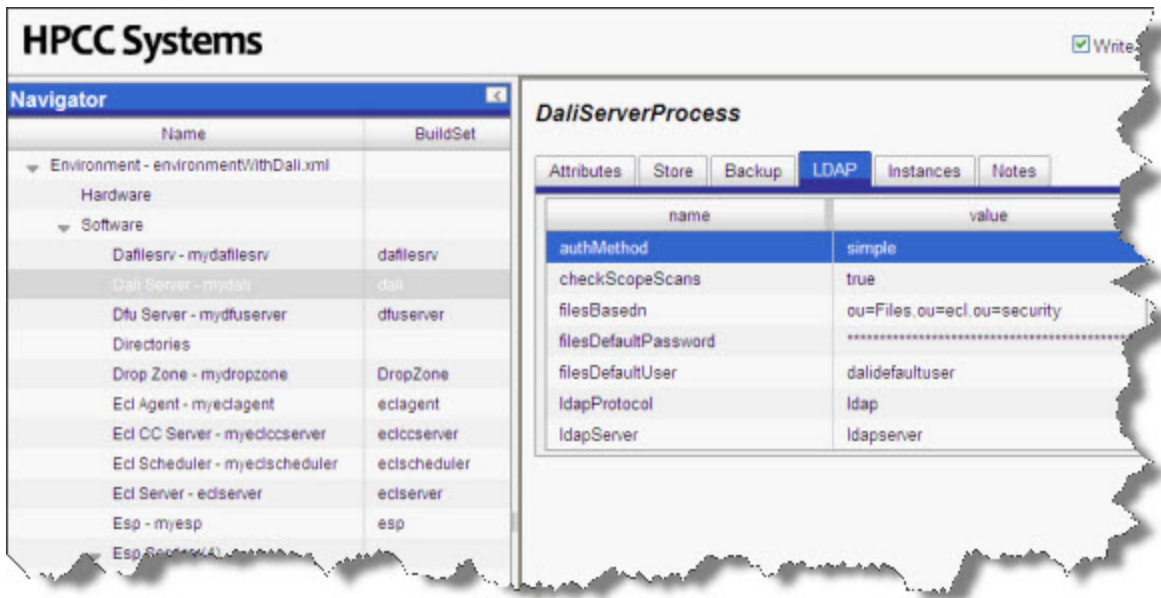
4. On the **EspProcess** page on the right hand side, select the **Authentication** tab.



Fill in the appropriate values:

- Change the **ldapAuthMethod** to [simple](#).
- Change the **ldapConnections** to the number appropriate for your system (100 is for example only, may not be necessary in your environment).
- Change **ldapServer** value to the name you gave your ldapServer, for example: [ldapsrv](#).
- Change the **method** value to [ldap](#).
- For each ESP Service binding, edit the **resourcesBasedn** and **workunitsBasedn** to match your LDAP server settings.
- Click on the disk icon to save.

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



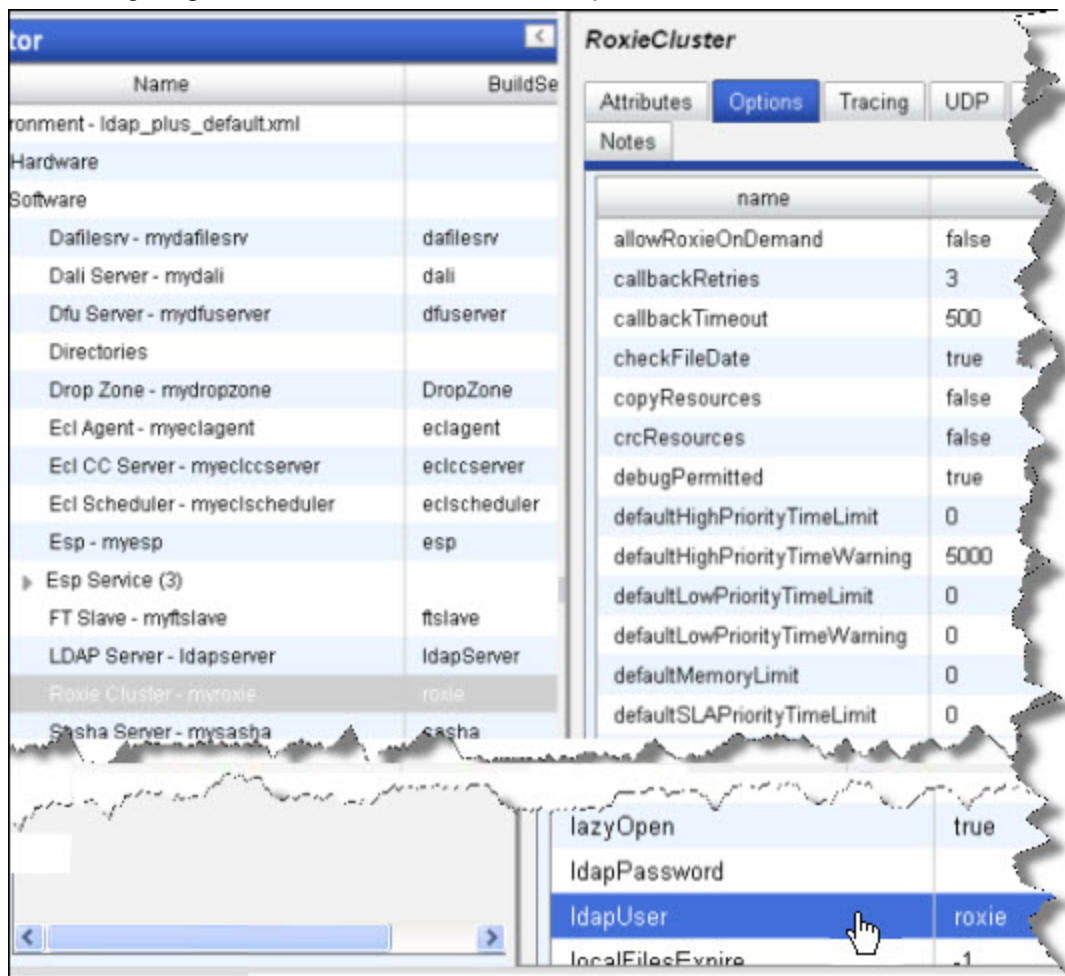
Fill in the values as appropriate:

- Select the **LDAP** tab.
- Change the **authMethod** to **simple**
- Change the LDAP values as appropriate to match the settings in your LDAP server.

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

- 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 **Options** tab.
- Scroll down to the **ldapUser** field and verify that there is a "roxie" user.
- You can add 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 the roxie user is created in the list of authenticated users.
In the following section, *Adding and editing users*, add "roxie" as a user and make sure the password is the same as the one entered in Configuration Manager.

User Security Maintenance

Configuring an HPCC System to use LDAP security will give you greater control over users and the security of your HPCC system.

Introduction

HPCC systems maintain 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 OpenLDAP 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. Additional "file access control" can be implemented over data files by configuring the Dali server to point to the Active Directory/LDAP server. This is what is known as enabling file security.

Permissions are established by group or by user and are defined as they are associated with a particular feature of the HPCC System. Only one set of 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 ECLWatch and similar features accessed from ECL IDE.
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
Esp Features for WsEclAccess	Controls access to the WS-ECL web service
Repository Modules	Controls access to the Attribute Repository and Modules in the repository
Esp Features for EclDirectAccess	Controls access to the ECLDirect web service

Access to the permission settings for each of these areas is available using the **Users/Permissions** area in ECL Watch.

Security Administration using ECL Watch

Administrator rights are needed to manager 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**). For example: **http://10.150.51.27:8010/**.

Security administration is controlled using the **Users/Permissions** area of ECL Watch. 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 member 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

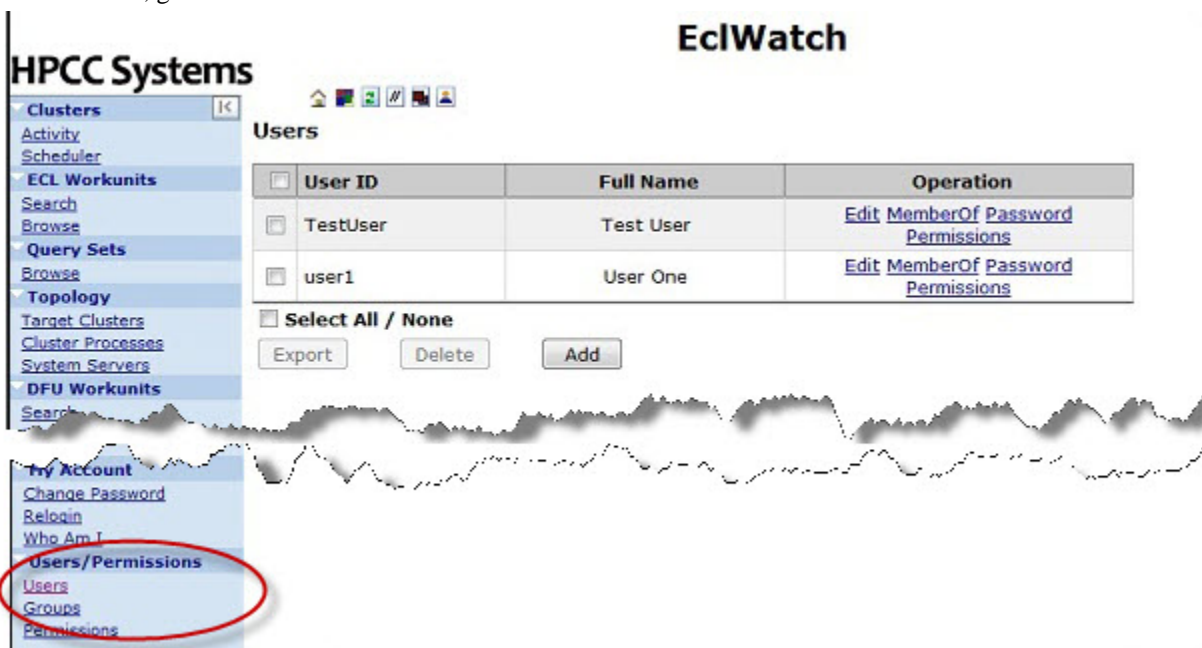
Setting and modifying user permissions

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
- 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

In ECL Watch, go to the **Users/Permissions** menu item and click **Users**:



All current users are identified in the list by their UserID and full name.

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

1. Press the **Add** button.

The **Add User** window is displayed.

2. Enter a **UserID**.

This is the login name for using 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.

5. Press **Submit**.

Confirmation of the request is shown.

Once added, the user is displayed in the list and you can modify the user's details and set permissions as required.

To modify a user's personal details:

1. Click the **Edit** link.

The **User Info Edit** window is displayed.

2. Change the **First Name** and **Last Name** as required.

Note: The **User Name** cannot be changed.

3. Press **Submit**.

Confirmation of the request is shown.

To add the user to a group:

1. Click on the **Member of** link.

The list of groups the user is already associated with is displayed.

2. To add the user to a group press **ADD**.

The list of available groups is displayed.

3. Click the checkbox to the left of the group(s) you want to add the user to and click **ADD**.

4. Click OK to confirm.

Confirmation of the request is shown.

To delete the user from a group:

1. Click on the **Member of** link.

The list of groups the user is already associated with is displayed.

2. Click the checkbox to the left of the group(s) you want to remove the user from and click **DELETE**.

3. Click OK to confirm.

Confirmation of the request is shown.

To change a user's password:

1. Click on the **Password** link.

The **Reset Password** window is displayed.

2. Complete the **New Password** and **Retype New Password** fields as required.

Press **Clear** to empty these fields and start again.

3. Click **Submit**.

Confirmation of the request is shown.

To delete a user from the list of authenticated users:

1. Check the checkbox to the left of the user(s) you want to remove.

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

2. Click **Delete**.

Confirmation of the request is shown.

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 and any permissions you set here overwrite those set in any group to which the user belongs.

To set new permissions for an individual user:

1. Click the **Users** menu item in ECL Watch
2. Locate the user in the list of authenticated users and click on the **Permissions** link in the **Operations** column.

The list of permissions currently set for this user are displayed and the groups from which the user has inherited permissions are also listed.

HPCC Systems **EclWatch**

Permissions of user1

Workunit Scopes

Inherited Permissions from Group: Authenticated Users (Changes inside this section will be applied to the whole group.)

Resource	Permission	allow				deny				Operation
		access	read	write	full	access	read	write	full	
SMC Feature	ClusterTopology1Access	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	ClusterTopologyAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	ConfigAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	DfuAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	DfuExceptions	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	DfuWorkunitsAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update

Workunit Scopes

3. A drop down list showing the 6 feature areas of the HPCC is provided. Select the feature area you are interested in and press **Add**.

The **Add Permissions for Authenticated Users** page is displayed.

4. There may be more than one resource setting available, select the one you require from the **Select Resource** drop down list provided.
5. Check the boxes that **allow** and **deny** access as required for the user.
6. Click **Add**.

Confirmation of the request is shown.

To modify permissions for an individual user:

1. Click the **Users** menu item in ECL Watch.
2. Locate the user in the list of authenticated and click on the **Permissions** link in the **Operations** column.
3. Locate the feature you want to modify.
4. Check the box(es) that **allow** and **deny** access as required for the feature and press **Update** or press **Delete** if you want to remove access to that feature.
5. Press OK to confirm.

Confirmation of the request is shown.

Setting and modifying group permissions

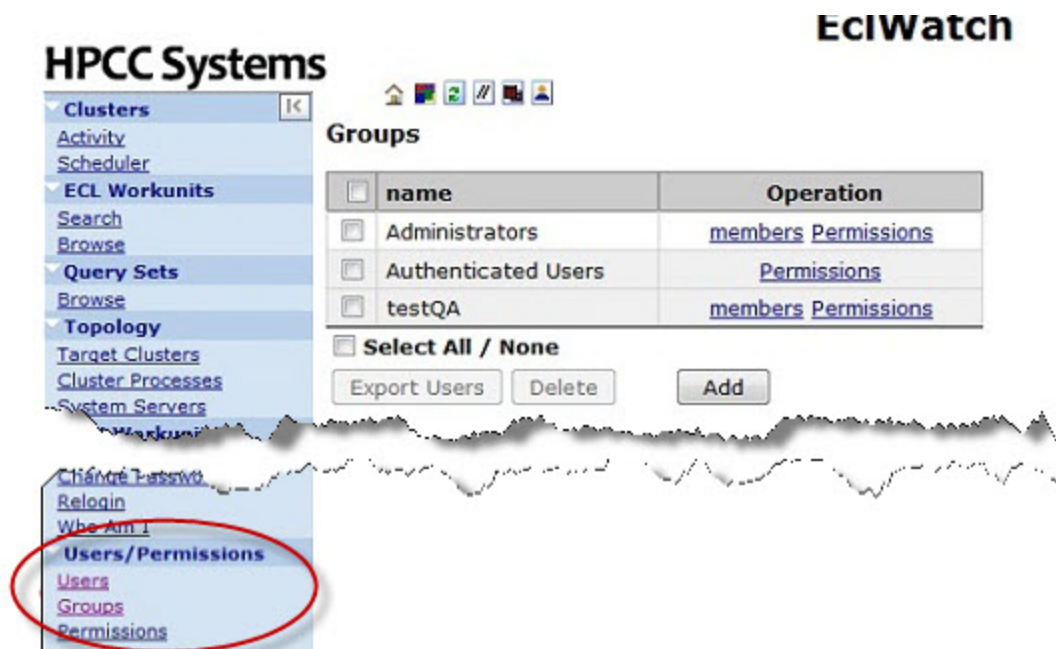
Setting up groups ensures that all users with the same permission needs have the same permission settings. You can give users the access they require to the feature areas of HPCC they need. There is no limit to the number of groups you can create. You can create as many groups as you need to control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

Use the **Group** menu item to:

- Add a new group
- Delete a group
- Add members to a groups
- 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. Since individual permission settings take precedence over the group settings, you can safely change the individual settings for a user without affecting the rest of the group(s) to which they belong.



In ECL Watch, go to the **Users/Permissions** menu item and click **Groups**:

To add a new group:

1. Press the **Add** button.

The **Add Group** window is displayed.

2. Enter a **Group Name**.
3. Press **Submit**.

Confirmation of the request is shown. **Permissions** may now be set for this new group.

To delete a group:

1. Locate the group in the list and check the checkbox to the left.
2. Press the **Delete** button.
3. Press **OK** to confirm.

Confirmation of the request is shown.

To add new members to a group:

1. Locate the group in the list and click on the **Members** link in the **Operations** column.

All current members of the group are listed.

2. Press **Add**.

All authenticated users are listed.

3. Check the box(es) to the left for all users you want to add to the group and press **Add**.
4. Press **OK** to confirm.

Confirmation of the request is shown.

To delete members from a group:

1. Locate the group in the list and click on the **Members** link in the **Operations** column.
2. Check the box(es) for all users you want to delete from the group and press **Delete**.
3. Press **OK** to confirm.

Confirmation of the request is shown.

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 controls.

If you intend to restrict permissions for some users, you must remove **Authenticated Users** from the sections you wish to limit. 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.

To set new permissions for a group:

1. Click the **Groups** menu item in ECL Watch.
2. Locate the group in the list and click on the **Permissions** link in the **Operations** column.

The list of permissions currently set for this group are displayed.

The screenshot shows the HPCC Systems EclWatch interface. On the left is a navigation menu with categories like Clusters, ECL Workunits, Query Sets, DFU Files, Roxie Queries, Resources, My Account, and Users/Permissions. The main area displays 'Permissions of tcstQA' as a table with columns for Resource, Permission, allow (access, read, write, full), deny (access, read, write, full), and Operation (delete, update). Below the table is a 'Workunit Scopes' dropdown and an 'Add' button.

Resource	Permission	allow				deny				Operation
		access	read	write	full	access	read	write	full	
SMC Feature	ClusterTopology1Access	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	ClusterTopologyAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	ConfigAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	DfuAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	DfuExceptions	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update
SMC Feature	DfuWorkunitsAccess	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	delete update

3. A drop down list showing the 6 feature areas of the HPCC is shown. Select the feature area you want and press **Add**.

The **Add Permissions for <GroupName>** page is displayed.

4. There may be more than one resource setting available, select the **Resource** you require from the drop down list provided.

5. Check the boxes that **allow** and **deny** access as required for the group.

6. Click **Add**.

To modify permissions for a group:

1. Click the **Groups** menu item in ECL Watch.

2. Locate the group in the list of authenticated and click on the **Permissions** link in the **Operations** column.

3. Locate the feature you want to modify.

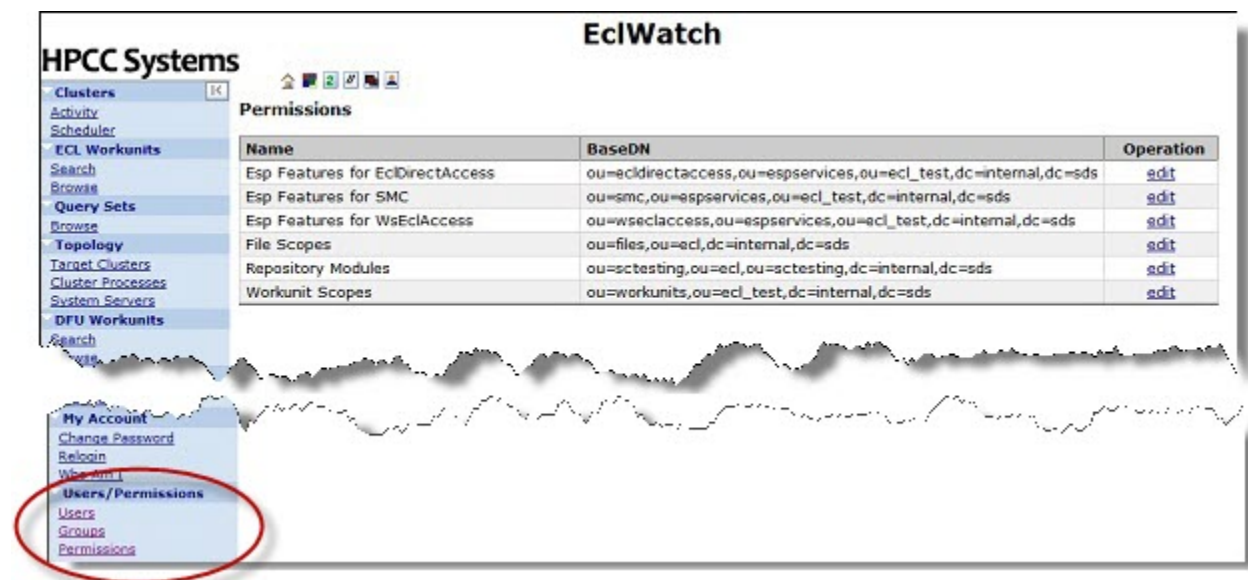
4. Check the box(es) that **allow** and **deny** access as required for the feature and press **Update** or press **Delete** if you want to remove access to that feature.

5. Press OK to confirm.

Confirmation of the request is shown.

Setting and modifying feature permissions

Access to the feature permissions area in ECL Watch is via the **User/Permissions** area using the **Permissions** menu item. All six feature areas are listed.



Use the **Permissions** menu item to:

- Edit the permissions for any feature
- Add new resources to a feature
- Set the permissions for users and groups for a specific resource

- Update the permissions for users and groups for a specific resource
- Delete a resource

Adding and editing feature permissions

Each feature contains a list of resources which are used to control access to an HPCC feature or folders containing files or workunits. The main HPCC feature permission settings are controlled using the **ESP Features for SMC** setting. When new features are added to the HPCC System, the release notes inform you that new permissions may be set. This is also true for **ESP Features for ECLDirectAccess** and **Esp Features for WsEclAccess**. Generally, all the permissions you require to control access to these features are already included.

However, to control access to file or workunit scopes, you must add the location as a resource before you can set permissions.

To add a scope:

1. Click the **Permissions** menu item in ECL Watch, locate the feature you want and click **Edit**.

The resources for that feature are listed.

2. Press the **Add** button.
3. Enter the exact name of the scope you want to add (for example a file or workunit scope) in the **Name** field and also a short **Description**.
4. Click **Submit**.

Confirmation of your request is shown.

5. Go back to the features page using the link provided.
6. Locate your new scope in the list and click **Permissions**.

The **Administrator** and **Authenticated Users** groups are shown showing the default permission settings which you can update as appropriate.

7. To add more users and groups and set permissions for this scope, click **Add**.

The **Add Permissions** window is displayed.

8. Select a **User** or **Group** from the dropdown lists provided and check the checkboxes for **allow** and **deny** as appropriate.
9. Click **Add**. Confirmation of your request is shown.

Note: This description shows how to add a file or workunit scope, However if you do need to add a new resource any other feature area, the process is the same.

To edit the permissions for a feature resource:

1. Click the **Permissions** menu item in ECL Watch, locate the feature you want and click **Edit**. The resources for that feature are listed.
2. Locate the resource you want to update and click **Permissions**. The permissions that are currently set for this resource are listed, including individual users and groups.

3. Locate the user or group you want. Click the checkboxes in the **allow** and **deny** columns as appropriate.
4. Click Update.
5. Click OK to confirm. Confirmation of your request is shown.

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

To delete a resource from a feature list:

1. Click the **Permissions** menu item in ECL Watch, locate the feature you want and click **Edit**. The resources for that feature are listed.
2. Locate the resource you want to remove and check the checkbox to the left.
3. Click **Delete**.
4. Click OK to confirm. Confirmation of your request is shown.

To delete the resource permission settings for a user or group:

1. Click the **Permissions** menu item in ECL Watch, locate the feature you want and click **Edit**.
The resources for that feature are listed.
2. Locate the resource you want to remove and click **Permissions**.
The users and groups are displayed.
3. Locate the user or group in the list and click **Delete**.
4. Click OK to confirm.

Confirmation of your request is shown.

ECL Watch Feature Permissions

Access to features of the HPCC system is controlled by via the **ESP Features for SMC** category. These features are listed as **Resources** when setting permissions using ECL Watch.

The following sections show the level of access required to be able to use ECL Watch features:

Login

SMCAccess is required by all users to be able to successfully login to ECL Watch.

LDAP Path	Description	Access
SmcAccess	Root Access to SMC Service	Read

Clusters

Users may be given access to the thor queue which can be manipulated by promoting/demoting queued workunits according to priority. The thor queue can also be paused or cleared and users can view thor usage statistics.

From this page, users can also click on workunit IDs to view details about the workunit. Depending on the level of access given, they can view, modify and delete their own, or others workunits.

LDAP Path	Description	Access
ThorQueueAccess	Access to Thor Job Queue Control	Full
OwnWorkunitsAccess	Access to View Own Workunit	Read
	Access to Create or Modify Own Workunit	Write
	Access to Delete Own Workunits	Full
OtherWorkunitsAccess	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

ECL Workunits

Workunits can also be viewed using this feature of ECL Watch. The contents of the workunits list reflects whether a user has the permission to view their own and others workunits.

LDAP Path	Description	Access
OwnWorkunitsAccess	Access to View Own Workunit	Read
	Access to Create or Modify Own Workunit	Write
	Access to Delete Own Workunits	Full
OtherWorkunitsAccess	Access to View Other User's Workunits	Read

Installing & Running the HPCC Platform
HPCC Installation and Startup

LDAP Path	Description	Access
	Access to Modify or Resubmit User's Workunits	Write
	Access to Delete Other User's Workunits	Full

Topology

This section shows details about the clusters and other HPCC System components. Preflight provides diagnostic information including disc space, CPU usage and access to logs as well as the ability to swap faulty nodes out of the cluster.

LDAP Path	Description	Access
ClusterTopologyAccess	Access to Cluster Topology	Read
	Set Machine Status	Write
	Swap Node	Full
MachineInfoAccess	Access to machine/Preflight Information	Read
MetricsAccess	Access to SNMP Metrics Information (Roxie Metrics)	Read
ExecuteAccess	Access to Remote Execution in ECL Watch	Full

DFU Workunits

A user must have permission to view DFU Workunits and requires other permissions to be able to manipulate them.

LDAP Path	Description	Access
DfuWorkunitsAccess	Access to View DFU Workunits	Read
	Access to Create, Delete, Update, Submit, and Abort DFU Workunits	Write

DFU Files

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.

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.

LDAP Path	Description	Access
DfuAccess	Access to DFU Logical Files	Read
	Delete Files, add to superfiles	Write
DfuExceptions	Access to DFU Exceptions	Read
DfuWorkunitsAccess	Access to View DFU Workunits	Read

Installing & Running the HPCC Platform
HPCC Installation and Startup

LDAP Path	Description	Access
	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
FileDesprayAccess	Access to De-Spraying Files	Write
FileSprayAccess	Access to Spraying and Copying	Read
	Rename files	Write
	Delete from Drop zone	Full
FileIO	Access to read files in Drop zone	Read
	Access to write to files in Drop zone	Write

Roxie Queries

Additional permission is required to view roxie queries in ECL Watch.

LDAP Path	Description	Access
RoxieQueryAccess	Access to Roxie Queries	Read

Users/Permissions

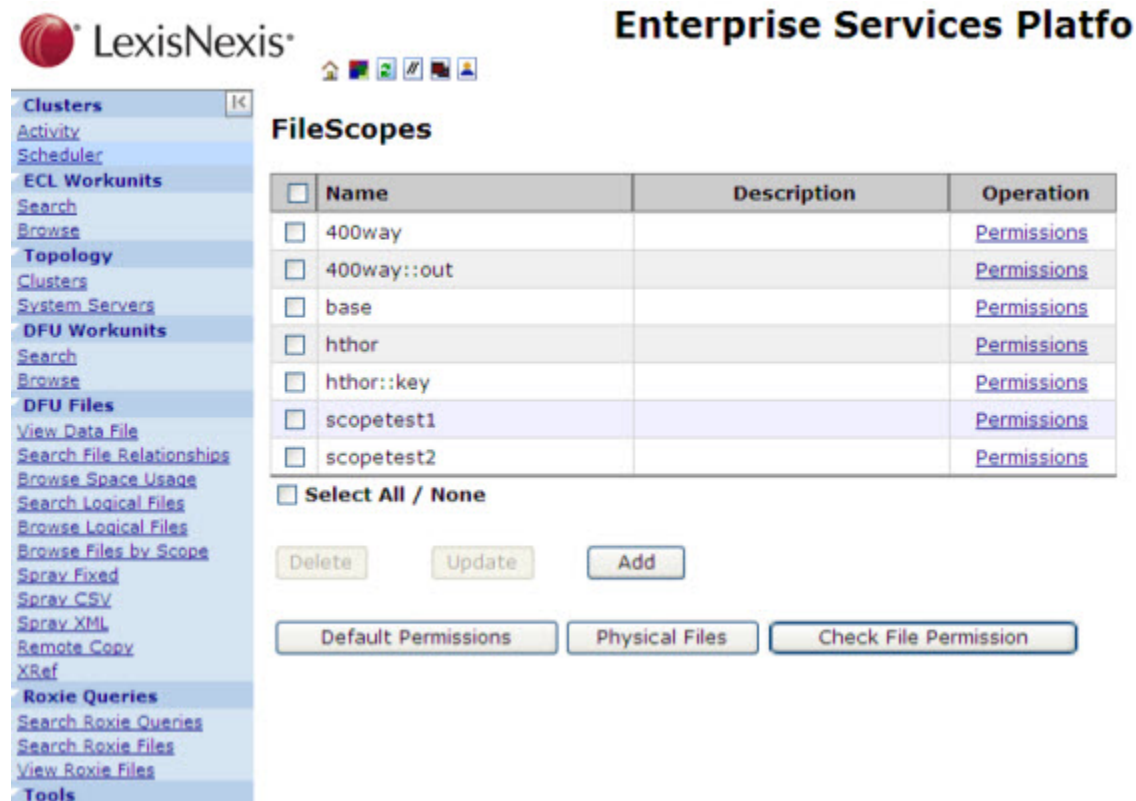
To be able to view the **Users/Permissions** area in ECL Watch, a user must be setup in LDAP as an Administrator.

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 organizational units (OUs) in LDAP.

An OU called **Files** is automatically created when the Dali server starts. To secure data folders, create OUs for each folder level and apply rights to each folder's OU.

Figure 14. File Scopes Page



Files is the top level for file scope security in the Dali Server. Below it there is an OU representing the cluster, for example thor (or the name that was set up for your cluster). Below that is an OU named **specialdata** which contains two OUs, **public** and **secure**. Typically, the **public** folder has rights granted to a large group of users and the **secure** folder has limited access granted. This allows you to prevent unauthorized users from any access to files in the **secure** folder.

The structure in LDAP corresponds to this logical structure in DFU:

thor::specialdata::secure

Which corresponds to this physical structure:

var/lib/HPCCSystems/hpcc-data/thor/specialdata/secure

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

- Remote Execution in ECL Watch
- Network file access using UNC's or Terminal Services
- Administrative utilities such as TreeView

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)

File scope feature permissions

There are some extra features that are available for the **File Scopes** feature in the **Permissions** area of ECL Watch.

- Any file scope in the list may be reset to have the default permission settings for your system.
- Access to the physical file may also be granted separately.
- The **Check File Permissions** button allows you to quickly view the permissions that have been set for any scope listed.

Permissions settings for file scopes may be set as follows:

Description	Access
Read files in that scope	Read
Create/modify a file in that scope	Write
Delete a file in that scope	Full

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 now have a scope value.

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

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 still be denied access to see one of John Doe's WUs.

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');
```

In addition, the scope of a workunit can be changed in ECL Watch by opening the WU, editing the scope field and pressing the **Save** button.

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

More Examples

This section contains additional ECL examples you can use on your HPCC cluster. You can run these on a single-node system or a larger multi-node cluster.

ECL Example: Anagram1

This example takes a `STRING` and produces every possible anagram from it. This code is the basis for a second example which evaluates which of these are actual words using a word list data file.


1. Open the ECL IDE (Start >> All Programs >> HPCC Systems >> ECL IDE) and login to your HPCC.
2. Open a new **Builder Window** (CTRL+N) and write the following code:

```
STRING Word := 'FRED' :STORED('Word');
R := RECORD
    STRING SoFar {MAXLENGTH(200)};
    STRING Rest {MAXLENGTH(200)};
END;
Init := DATASET([{'',Word}],R);
R Pluck1(DATASET(R) infile) := FUNCTION
R TakeOne(R le, UNSIGNED1 c) := TRANSFORM
    SELF.SoFar := le.SoFar + le.Rest[c];
    SELF.Rest := le.Rest[..c-1]+le.Rest[c+1..];
// Boundary Conditions handled automatically
END;
RETURN NORMALIZE(infile,LENGTH(LEFT.Rest),TakeOne(LEFT,COUNTER));
END;
L := LOOP(Init,LENGTH(TRIM(Word)),Pluck1(ROWS(LEFT)));
OUTPUT(L);
```

3. Select **thor** as your target cluster.
4. Press the syntax check button on the main toolbar (or press F7)

5. Press the **Submit** button (or press ctrl+enter).


Figure 15. Completed job



The green check mark indicates successful completion.

- Click on the workunit number tab and then on the Result 1 tab to see the output.

Figure 16. Completed job output



Roxie Example: Anagram2

In this example, we will download an open source data file of dictionary words, spray that file to our Thor cluster, then validate our anagrams against that file so that we determine which are valid words. The validation step uses a JOIN of the anagram list to the dictionary file. Using an index and a keyed join would be more efficient, but this serves as a simple example.

Download the word list

We will download the word list from <http://wordlist.sourceforge.net/>

1. Download the *Official 12 Dicts* Package. The files are available in tar.gz or ZIP format.
2. Extract the **2of12.txt** file to a folder on your local machine.

Load the Dictionary File to your Landing Zone

In this step, you will copy the data files to a location from which it can be sprayed to your HPCC cluster. A Landing Zone is a storage location attached to your HPCC. It has a utility running to facilitate file spraying to a cluster.

For smaller data files, maximum of 2GB, you can use the upload/download file utility in ECL Watch. This data file is only ~400 kb.

Next you will distribute (or Spray) the dataset to all the nodes in the HPCC cluster. The power of the HPCC comes from its ability to assign multiple processors to work on different portions of the data file in parallel. Even though the VM Edition only has a single node, the data must be sprayed to the cluster.

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



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

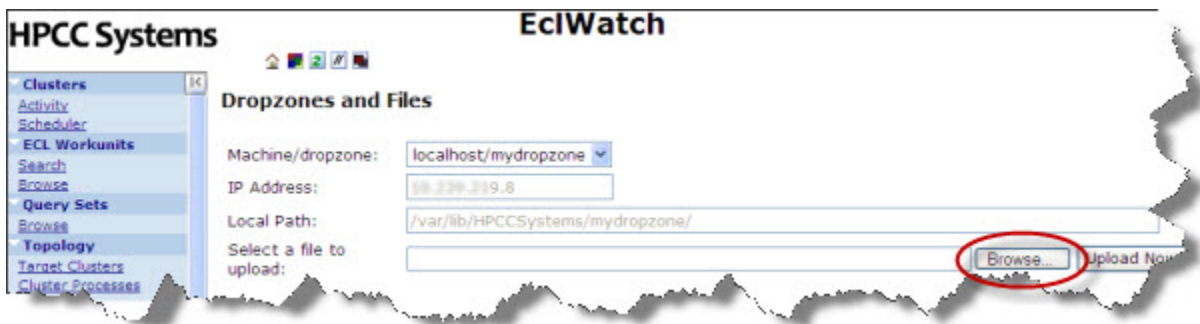
2. From ECL Watch page, click on the **Upload/download File** link in the menu on the left side.

Figure 17. Upload/download



Once you click on the Upload/download file link, it will take you to the **Dropzones and Files** page, where you can choose to **Browse** your machine for a file to upload:

Figure 18. Dropzones and Files



3. Press the **Browse** button to browse the files on your local machine, select the file to upload and then press the **Open** button.

The file you selected should appear in the **Select a file to upload:** field. The data file is named: **2of12.txt**.

4. Press on **Upload Now** to complete the file upload.

Spray the Data File to your *Data Refinery (Thor) Cluster*

To use the data file in our HPCC system, we must “spray” it to all the nodes. A *spray* or *import* is the relocation of a data file from one location (such as a Landing Zone) to multiple file parts on nodes in a cluster.

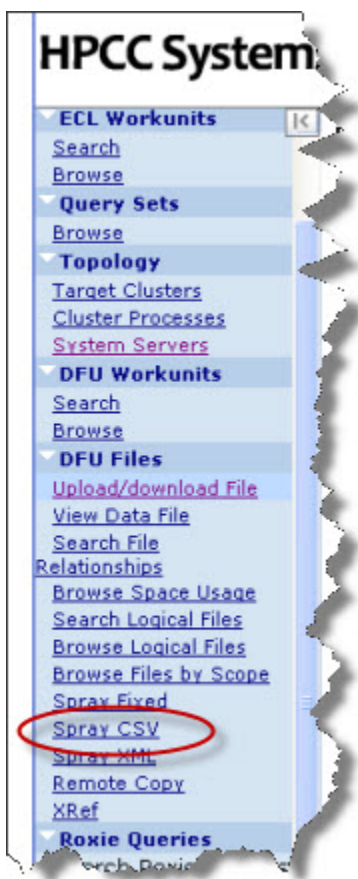
The distributed or sprayed file is given a *logical-file-name* as follows: `~thor::word_list_csv` The system maintains a list of logical files and the corresponding physical file locations of the file parts.

1. Open ECL Watch 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)

2. Click on the Spray CSV hyperlink under the DFU Files menu on the left.

Figure 19. Spray CSV



The **DFU Spray CSV** page displays.

3. Select mydropzone in the Source **Machine/dropzone** drop-list.

The IP Address is automatically filled and the Local Path is partially filled with the default folder on your landing zone. Note: The VM and Community Edition typically only has one landing zone defined.

4. Complete the Local Path to include the complete file name or use the **Choose File** button to select the file from a list of files in the folder. The file is **2of12.txt**.

5. Fill in the rest of the parameters (if they are not filled in already).
 - Max Record Length 8192
 - Separator \,
 - Line Terminator \n,\r\n
 - Quote: '
6. Fill in the Label using the Logical File name desired: thor::word_list_csv
7. Make sure the **Overwrite** and **Replicate** boxes are checked.

Figure 20. Spray the File

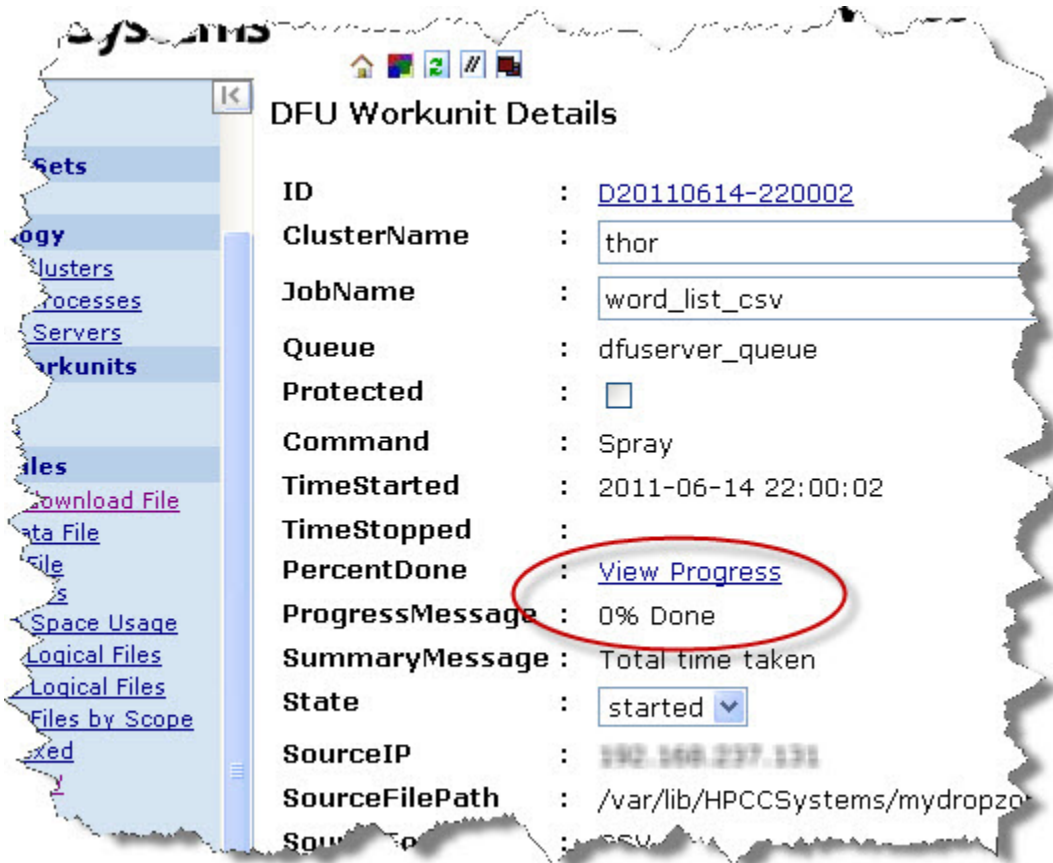
The screenshot shows the HPCC Systems EclWatch interface. The main content area is titled 'Spray CSV'. It contains the following fields and options:

- Source:**
 - Machine/dropzone: localhost/mydropzone
 - IP Address: 192.168.237.131
 - Local Path: /var/lib/HPCCSystems/mydropzone/2of12.txt (with a 'Choose File...' button)
 - Network Path: //192.168.237.131/var/lib/HPCCSystems/mydropzone/2of12.txt
 - Format: ASCII
 - Max Record Length: 8192
 - Separator: \,
 - Line Terminator: \n,\r\n
 - Quote: ''
- Destination:**
 - Group: mythor
 - Label: thor::word_list_csv
 - Mask: thor::word_list_csv_\$\$\$_of_\$\$\$_
 - Prefix: (empty)
- Options:**
 - Overwrite:
 - Replicate:
 - No Split:
 - Compress:

A 'Submit' button is located at the bottom of the form.

8. Press the **Submit** button

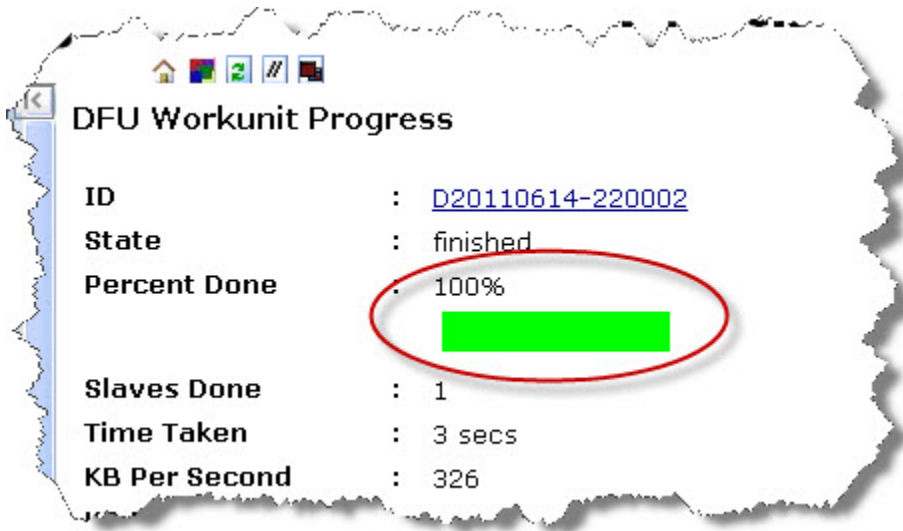
Figure 21. View Progress



9. Click the **View Progress** link

10. The Workunit progress page displays.

Figure 22. Workunit Progress



Run the query on Thor

1. Open a new **Builder Window** (CTRL+N) and write the following code:

```
IMPORT Std;
layout_word_list := record
  string word;
end;
File_Word_List := dataset('~thor::word_list_csv', layout_word_list,
                        CSV(heading(1),separator(','),quote('')));
STRING Word := 'teacher' :STORED('Word');
STRING SortString(STRING input) := FUNCTION
  OneChar := RECORD
    STRING c;
  END;
  OneChar MakeSingle(OneChar L, unsigned pos) := TRANSFORM
    SELF.c := L.c[pos];
  END;
  Split := NORMALIZE(DATASET([input],OneChar), LENGTH(input),
                    MakeSingle(LEFT,COUNTER));
  SortedSplit := SORT(Split, c);
  OneChar Recombine(OneChar L, OneChar R) := TRANSFORM
    SELF.c := L.c+R.c;
  END;
  Recombined := ROLLUP(SortedSplit, Recombine(LEFT, RIGHT),ALL);
  RETURN Recombined[1].c;
END;

STRING CleanedWord := SortString(TRIM(Std.Str.ToUpperCase(Word)));

R := RECORD
  STRING SoFar {MAXLENGTH(200)};
  STRING Rest {MAXLENGTH(200)};
END;
Init := DATASET([{'',CleanedWord}],R);
R Pluck1(DATASET(R) infile) := FUNCTION
  R TakeOne(R le, UNSIGNED1 c) := TRANSFORM
    SELF.SoFar := le.SoFar + le.Rest[c];
```

Installing & Running the HPCC Platform More Examples

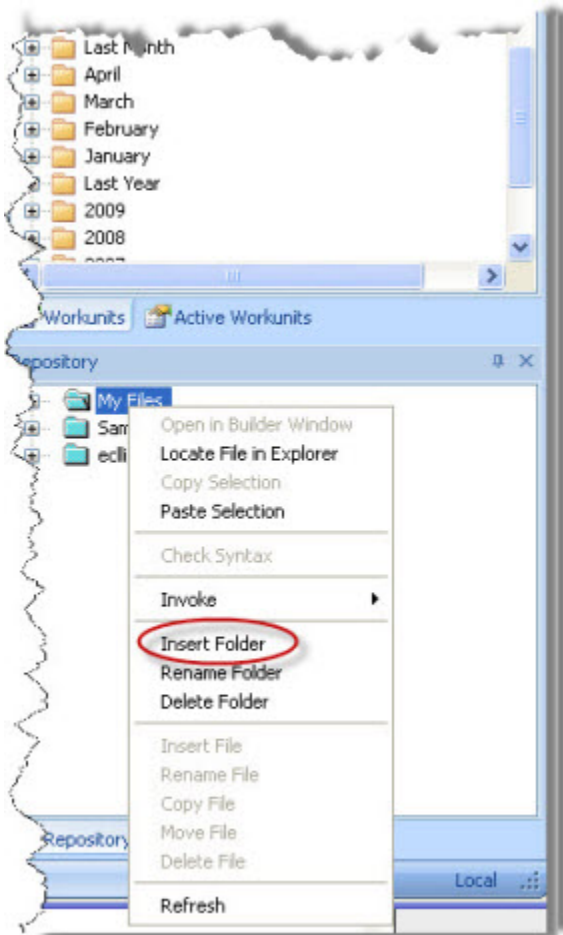
```
SELF.Rest := le.Rest[..c-1]+le.Rest[c+1..];  
// Boundary Conditions  
// handled automatically  
END;  
RETURN DEDUP(NORMALIZE(infile,LENGTH(LEFT.Rest),TakeOne(LEFT,COUNTER)));  
END;  
L := LOOP(Init,LENGTH(CleanedWord),Pluck1(ROWS(LEFT)));  
ValidWords := JOIN(L,File_Word_List,  
LEFT.Sofar=Std.Str.ToUpperCase(RIGHT.Word),TRANSFORM(LEFT));  
OUTPUT(CleanedWord);  
COUNT(ValidWords);  
OUTPUT(ValidWords)
```

2. Select **thor** as your target cluster.
3. Press the syntax check button on the main toolbar (or press F7)
4. Press the **Submit** button.
5. When it completes, select the Workunit tab, then select the Result tab.
6. Examine the result.

Compile and Publish the query to Roxie


1. RT-CLICK on the **My Files** folder in the Repository window, and select **Insert Folder** from the pop-up menu.

Figure 23. Insert Folder



2. Enter **Anagram** for the label, then press the OK button.

Figure 24. Enter Folder Label



3. RT-CLICK on the **Anagram** Folder, and select **Insert File** from the pop-up menu.

4. Enter **ValidateAnagrams** for the label, then press the OK button.

A Builder Window opens.

Figure 25. Builder Window



5. Write the following code (you can copy the code from the other builder window):

```
IMPORT Std;
layout_word_list := record
  string word;
end;
File_Word_List := dataset('~thor::word_list_csv', layout_word_list,
                        CSV(heading(1),separator(','),quote('')));
STRING Word := 'teacher' :STORED('Word');
STRING SortString(STRING input) := FUNCTION
  OneChar := RECORD
    STRING c;
  END;
  OneChar MakeSingle(OneChar L, unsigned pos) := TRANSFORM
    SELF.c := L.c[pos];
  END;
  Split := NORMALIZE(DATASET([input],OneChar), LENGTH(input),
                    MakeSingle(LEFT,COUNTER));
  SortedSplit := SORT(Split, c);
  OneChar Recombine(OneChar L, OneChar R) := TRANSFORM
    SELF.c := L.c+R.c;
  END;
  Recombined := ROLLUP(SortedSplit, Recombine(LEFT, RIGHT),ALL);
  RETURN Recombined[1].c;
END;

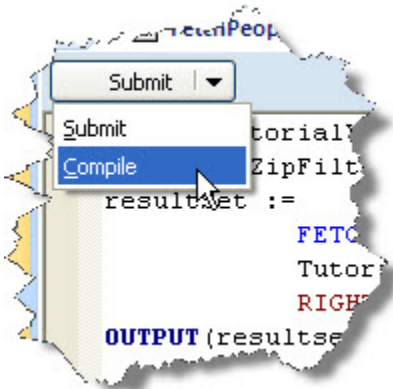
STRING CleanedWord := SortString(TRIM(Std.Str.ToUpperCase(Word)));

R := RECORD
  STRING SoFar {MAXLENGTH(200)};
```

```
STRING Rest {MAXLENGTH(200)};
END;
Init := DATASET(['',CleanedWord],R);
R Pluck1(DATASET(R) infile) := FUNCTION
  R TakeOne(R le, UNSIGNED1 c) := TRANSFORM
    SELF.Sofar := le.Sofar + le.Rest[c];
    SELF.Rest := le.Rest[..c-1]+le.Rest[c+1..];
    // Boundary Conditions
    // handled automatically
  END;
  RETURN DEDUP(NORMALIZE(infile,LENGTH(LEFT.Rest),TakeOne(LEFT,COUNTER)));
END;
L := LOOP(Init,LENGTH(CleanedWord),Pluck1(ROWS(LEFT)));
ValidWords := JOIN(L,File_Word_List,
LEFT.Sofar=Std.Str.ToUpperCase(RIGHT.Word),TRANSFORM(LEFT));
OUTPUT(CleanedWord);
COUNT(ValidWords);
OUTPUT(ValidWords)
```

6. Select **Roxie** as your target cluster.
7. Press the syntax check button on the main toolbar (or press F7)
8. In the Builder window, in the upper left corner the **Submit** button has a drop down arrow next to it. Select the arrow to expose the **Compile** option.


Figure 26. Compile



9. Select **Compile**
10. When it completes, select the Workunit tab, then select the Result tab.

11. When the workunit finishes, it will display a green circle indicating it has compiled.

Figure 27. Compiled



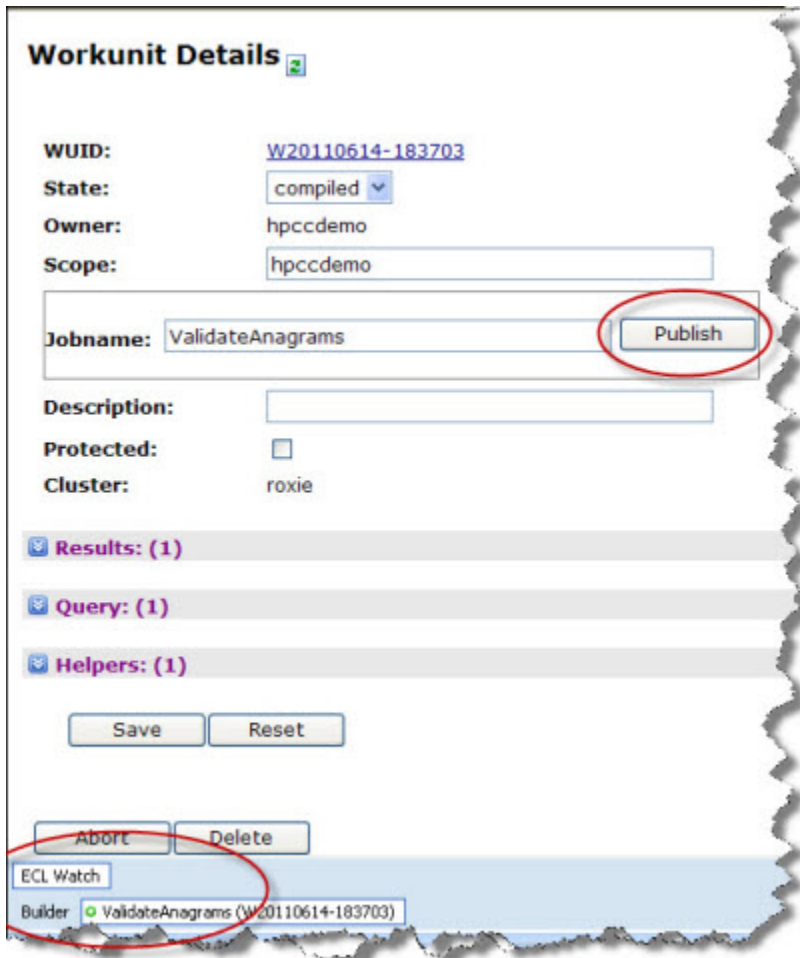
Publish the Roxie query

Next we will publish the query to a Roxie Cluster.

1. Select the workunit tab for the ValidateAnagrams that you just compiled.
2. Select the ECL Watch tab.


3. Press the **Publish** button (you may need to scroll down the main window)

Figure 28. Publish Query



When it successfully publishes, you will see:

Figure 29. Workunit Published



Run the Roxie Query in WsECL

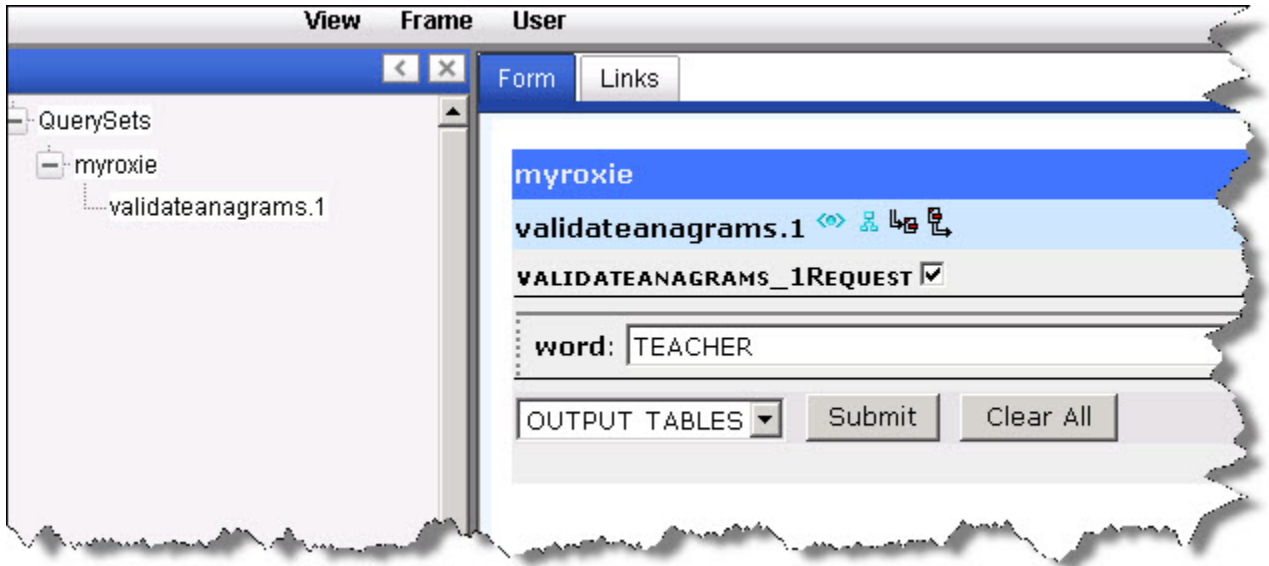
Now that the query is published to a Roxie cluster, we can run it using the WsECL service. WsECL is a web-based interface to queries on an HPCC platform. Use the following URL:

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 8002)

1. Click on the + sign next to **myroxie** to expand the tree.
2. Click on the **ValidateAnagrams.1** hyperlink.

The form for the service displays.

Figure 30. RoxieECL




3. Select Output Tables in the drop-list.

4. Provide a word to make anagrams from (e.g., TEACHER), then press the Submit button.

The results display.

Figure 31. RoxieResults




Next Steps

Available from the menu in the ECL IDE there are several documents which provide details on various aspects of the HPCC.

You can access them from the help menu: Help >> Documentation.

Figure 32. Help Menu



You can also find these from the **Start** menu :

Start >> All Programs >> HPCC Systems >> ECL IDE >> Docs

To familiarize yourself with what your system can do we recommend following the steps in

- The **HPCC Data Tutorial**
- **The Six Degrees of Kevin Bacon** example
- Read **Using Config Manager** to learn how to configure an HPCC platform using Advanced View.
- Use your new skills to process your own massive dataset!

The HPCC Systems Portal is also a valuable resource for more information including:

- Video Tutorials
- Additional examples
- White Papers
- Documentation

Appendix

Example Scripts

For a multi-node configuration, the packages must be installed on each node. You can install each one manually or use scripts to copy and install the packages. On a large system where you have many nodes copying and installing on every node is not practical, therefore we provide some scripts you can use or to serve as examples to give you a start in making your own.

Scripts are installed to the `/opt/HPCCSystems/sbin` directory.



Make sure that you have the sufficient privileges to sudo as an administrator to use the `install-cluster.sh` script. To use the `hpcc-push.sh` or `hpcc-run.sh` scripts, you must sudo as user **hpcc**.

install-cluster.sh

install-cluster.sh [-k] <package-name>

<package-name>	Name of the HPCC package to install. Required
-k	When specified, the script generates and distributes ssh keys to all hosts. Optional.

You can run this script as any user with sufficient permissions to execute it; however, when prompted for username/password, you must provide credentials for a user with sufficient sudo rights to run commands as an administrator on all nodes.

Before you can use this script, you must have already defined and generated an `environment.xml` file (using ConfigMgr's wizard or advanced mode). This script:

- reads the active `environment.xml` file and gathers a list of nodes upon which to act.
- installs the HPCC platform package(s) on all nodes specified.
- pushes out and deploys the environment file (`environment.xml`) to all nodes specified.
- optionally, if you specify the `-k` option it also generates the required ssh keys and deploys them as required to all nodes specified.

Examples:

This example installs the HPCC Platform packages to remaining nodes and pushes out the active environment.xml file to those nodes.:

```
sudo /opt/HPCCSystems/sbin/install-cluster.sh hpccsystems-platform-xxxx-n.n.nnnn
```

(where *n.n.nnnn* is the build number)

This example installs the HPCC Platform packages to all nodes and pushes out the active environment.xml file to those nodes. It also generates ssh keys and pushes them out to all nodes.

```
sudo /opt/HPCCSystems/sbin/install-cluster.sh -k hpccsystems-platform-xxxx-n.n.nnnn
```

(where *n.n.nnnn* is the build number)

hpcc-push.sh

To use this script, the ssh keys need to be properly configured on all nodes, and you must use sudo:

This script "pushes" files from the source filename and path to the destination filename and path for all IP addresses in the active environment.xml.

The IP addresses were defined when editing the environment in ConfigMgr.

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

For example:

```
sudo -u hpcc /opt/HPCCSystems/sbin/hpcc-push.sh \  
    /etc/HPCCSystems/environment.xml /etc/HPCCSystems/environment.xml
```

hpcc-run.sh

hpcc-run.sh [-c component] [-a {hpcc-init|dafilesrv}] {start|stop|restart|status|setup}

- c
- *-c componentname*
Specifies the component upon which to execute the command. If omitted, the default is **all** components on the machine.
 - *-c componenttype*
Specifies the component type upon which to execute the command. If more than one of this type is configured, all will be acted upon. If omitted, the default is **all** components on the machine.
- a
- hpcc-init: [start|stop|restart|status|setup]
 - dafilesrv [start|stop]

To use this script, the ssh keys need to be properly configured on all nodes, and you must sudo as user hpcc:

This script runs a command on all IP addresses in the active environment.xml.

The IP addresses were defined when editing the environment in ConfigMgr. This script supports all the parameters of hpcc-init and dafilesrv.

Example:

This example starts all components on the nodes

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

This example starts all components of the esp type on the nodes

```
sudo -u hpcc /opt/HPCCSystems/sbin/hpcc-run.sh -c esp -a hpcc-init start
```

This example starts all components with a component name myesp on the nodes

```
sudo -u hpcc /opt/HPCCSystems/sbin/hpcc-run.sh -c myesp -a hpcc-init start
```

This example starts the dafilesrv helper application

```
sudo -u hpcc /opt/HPCCSystems/sbin/hpcc-run.sh -a dafilesrv start
```

Uninstalling the HPCC Platform

To uninstall the HPCC platform from your system, issue the following command. If necessary, do so on each node that it is installed on.

```
sudo rpm -e HPCCSystems-platform-xxxx-n.n.nnnn
```

(where *n.n.nnnn* is the build number)

Helper Applications

There is a helper applications that runs on all nodes that you may need to stop or start manually.

Normally, this process is started automatically the first time the hpcc-init service executes.

Enter the following commands to stop or start the helper application:

- dfilesrv

```
sudo /sbin/service dfilesrv stop  
sudo /sbin/service dfilesrv start
```

hpcc-init

sbin/service hpcc-init [*option*] *command*

- | | |
|---------|---|
| option | <ul style="list-style-type: none">• <i>-c componentname, --component=componentname</i>

Specifies the component upon which to execute the command. If omitted, the default is all components on the machine.

<i>-c componenttype, --component=componenttype</i>

Specifies the component type upon which to execute the command. If more than one of this type is configured, all will be acted upon. If omitted, the default is all components on the machine.• <i>--componentlist</i>

Provides a list of all component names on the current node as specified in the environment file.• <i>--typelist</i>

Provides a list of all component types on the current node as specified in the environment file.• <i>-h, --help</i>

Displays a help page |
| command | <ul style="list-style-type: none">• <i>start:</i>

Starts component(s)• <i>stop</i>

Stops component(s)• <i>status</i>

Displays component(s) status• <i>restart</i>

Restarts component(s)• <i>force-reload</i>

Deletes all local configuration files, data files, log files, and then restarts component(s). BE CAREFUL using this command.• <i>setup</i>

Initializes component configuration files but does not start the component(s). |

The **hpcc-init** function is used to start, stop, restart, setup, or check the status of any or all HPCC components.

Examples:

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

sudo /sbin/service hpcc-init -c myeclserver start
sudo /sbin/service hpcc-init --component=myeclserver start

sudo /sbin/service hpcc-init -c esp start
```

Running the ECL IDE under WINE

To run the ECL IDE under WINE in Linux, follow these steps.

1. Install wine1.2 (this corresponds to Wine version 1.1.31) and its dependencies.
2. Download msxml3.msi from Microsoft (Service Pack 7 or later).
<http://support.microsoft.com/kb/308480/en-us>
3. Install msxml3.msi in Wine (DOUBLE-CLICK the msi file and Wine will install it).
4. Open Configure Wine (Applications/Wine/Configure Wine):
5. Select the Libraries tab.
6. In the New override for library droplist, select *msxml3*, then press the add button.
7. Select *msxml3* in the Existing overrides list and press Edit.
8. Select the *Native (Windows)* option and press the OK button.
9. Press the OK button to close the Wine Configuration window.
10. Install the HPCC ECL IDE (DOUBLE-CLICK the setup.msi file and Wine will install it).