

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

May 1, 2012

The Eclipse Parallel Tools Platform

Toward an Integrated Development Environment
for Improved Software Engineering on Crays



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Agenda

1. What is the Eclipse Parallel Tools Platform (PTP)
2. Tour of features available in Eclipse/PTP
 - Features added to support Blue Waters
3. How could Cray-PTP integration be improved

What is the Eclipse Parallel Tools Platform?

- **Eclipse**

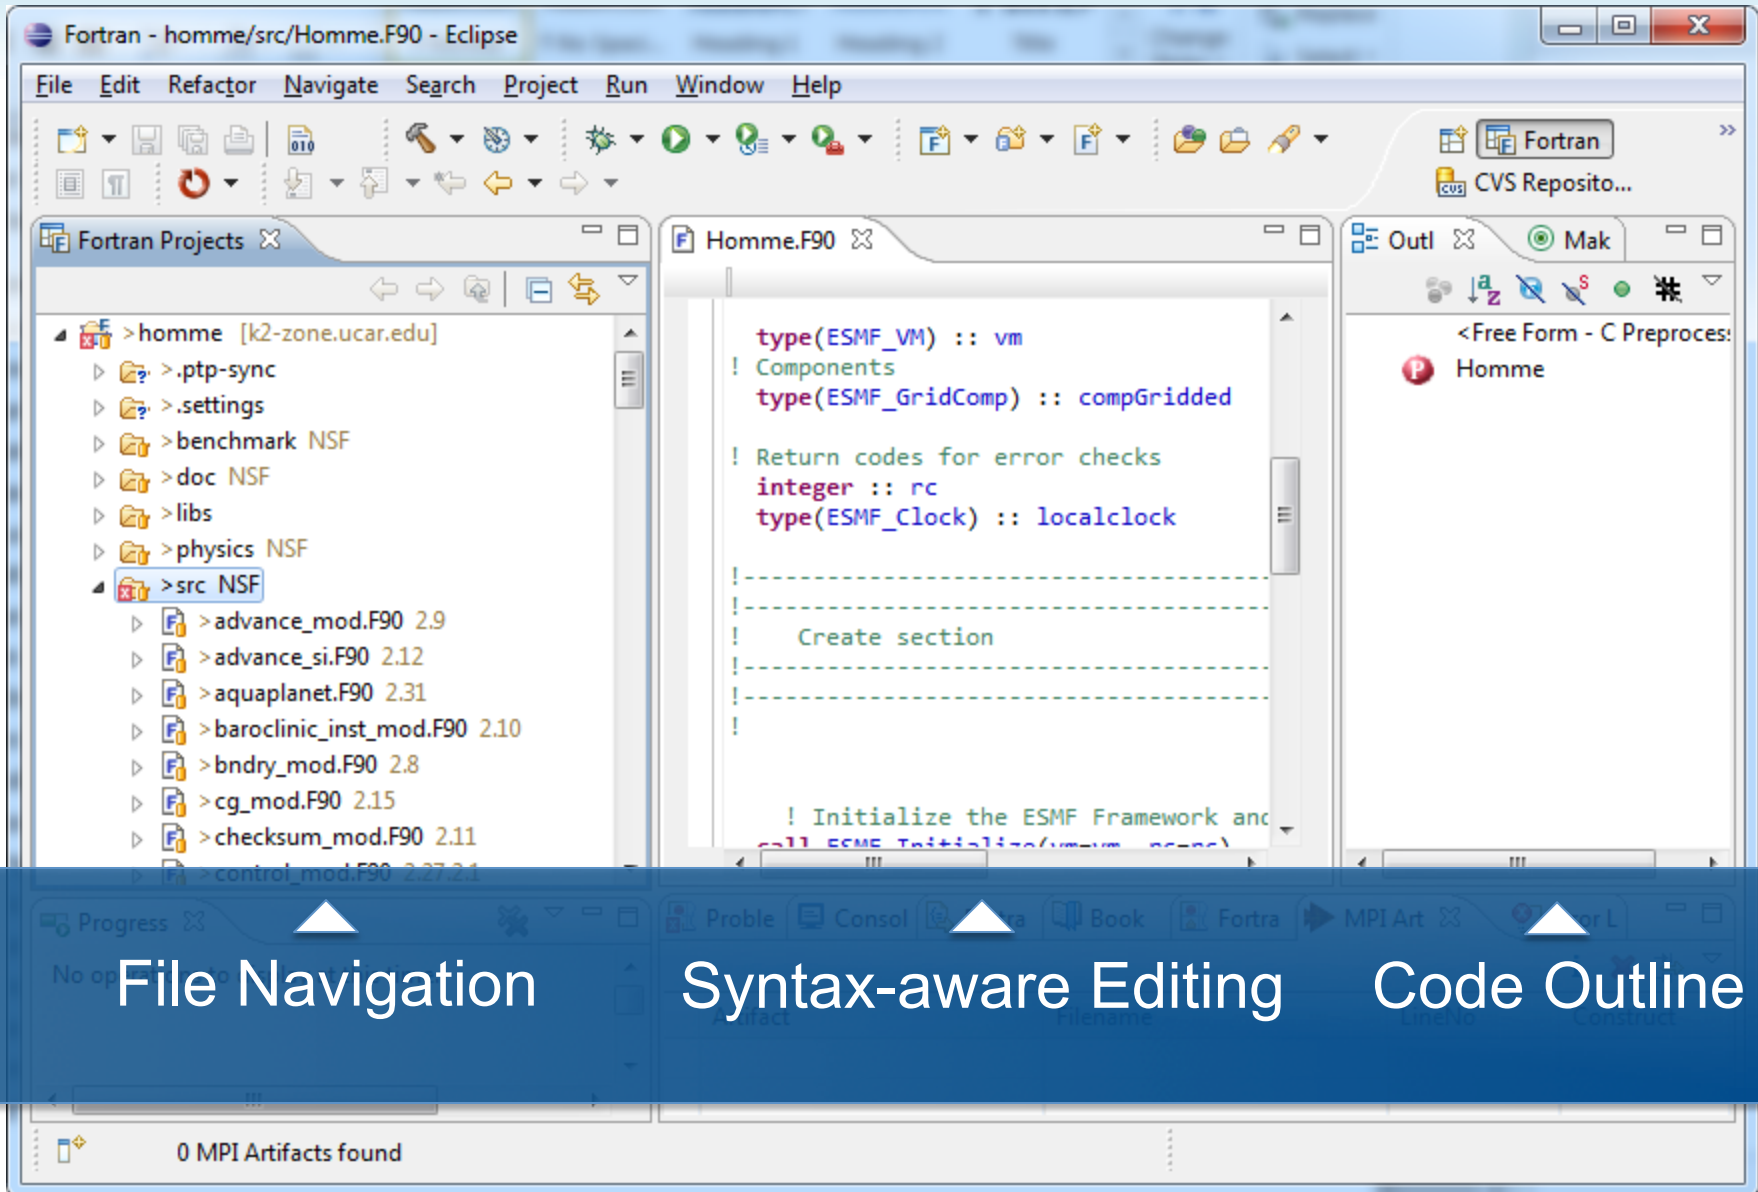
- Multi-platform integrated development environment
- Extremely popular as a Java IDE
- Excellent support for C/C++, UPC, Fortran, Python
- Extensible via third-party plug-ins
 - Actually, *everything* is a plug-in!
 - Java, C/C++, Fortran support are all plugged in
 - CVS, Subversion, Git support are plugged in
 - What about plug-ins to support HPC?

What is the Eclipse Parallel Tools Platform?

- **Parallel Tools Platform (PTP)**
 - Set of Eclipse plug-ins
 - Adds support for HPC development to Eclipse
 - Write code on your laptop; compile on an HPC resource
 - Submit jobs to a batch scheduler; monitor jobs
 - Debug remote MPI applications (parallel debugger)
 - Get assistance with MPI, OpenMP development
 - Current release: about 50,000 downloads

Supporting Blue Waters

- **Blue Waters:** Cray XE6/XK6 at NCSA
- PTP did not work with Crays “out of the box”
 - Could not submit jobs with appropriate aprun options
 - Could not monitor status of compute nodes
 - Could not set environment modules for build
 - Did not recognize Cray, PGI compilers’ errors messages
 - Did not support OpenACC
 - ...
- Less than 6 months to fix these for PTP 6.0 (!)



Static analyses
built in, available
in real time while
coding

< Static call hierarchy

Call Hierarchy

Calls from main(int, char **) - /proxy/org.eclipse.ptp.debug.sdm_5.0.4.201111121445/src/master

- main(int, char **) : int
 - shortopts : char *
 - longopts : option []
 - opt_type : int
 - find_dbg_backend(char *, dbg_backend **) : int
 - dbg_backends : dbg_backend [] (2 matches)
 - dbg_backend::db_name : char *
 - backend_set_path(dbg_backend *, char *) : void
 - find_proxy(char *, proxy **) : int
 - sdm_init(int, char **) : int
 - sdm_route_get_id() : sdm_id (2 matches)
 - SDM_MASTER : int
 - master(char *, char *, int) : void
 - sdm_route_get_size() : int
 - DbgMasterInit(int, int, char *, proxy_svr_helper_funcs *, proxy_commands *) : int
 - sdm_route_get_id() : sdm_id
 - helper_funcs : proxy_svr_helper_funcs
 - command_tab : proxy_commands
 - DbgMasterCreateSession(int, char *, int) : int
 - DbgGetErrorStr() : char *
 - DbgMasterQuit(int, int, char **) : int
 - DbgMasterProgress() : int (2 matches)
 - DbgMasterIsShutdown() : int

Integrated
OpenACC
documentation
(added for BW)

Documentation
also available for
MPI, OpenMP

The screenshot shows an IDE window with a tab labeled 'test1.f90'. The code in the editor is as follows:

```
15  
16 !$acc parallel loop  
17 ..do i = 1, 1000  
18 ...c(:, :) = (a(:, :) + b(:, :)) / 2.d0  
19 ...a(:, :) = (a(:, :) + c(:, :)) / 2.d0  
20 ...b(:, :) = (b(:, :) + c(:, :)) / 2.d0  
21 ..end do  
22 !$acc end parallel loop
```

Below the code editor, there is a 'Problems' pane and a 'Fortran Declaration' pane. The 'Fortran Declaration' pane is active and displays the following content:

OpenACC™ parallel directive

Delineates a block of code that will be executed on an accelerator device.

<pre>!\$acc parallel [clause [, clause ...]] block !\$acc end parallel</pre>	<pre>#pragma acc parallel [clause [, clause ...]] block</pre>
--	---

Supported clauses are `if`, `async`, `num_gangs`, `num_workers`, `vector_length`, `reduction`, `copy`, `copyin`, `copyout`, `create`, `present`, `present_or_copy`, `present_or_copyin`, `present_or_copyout`, `present_or_create`, `deviceptr`, `private`, `firstprivate`.


```
!$acc ·  
· · do !$acc cache – OpenACC cache directive  
· · · · !$acc data – OpenACC data directive  
· · · · !$acc end data – OpenACC end data directive  
· · en !$acc declare – OpenACC declare directive  
! · · !$acc host_data – OpenACC host_data directive  
· · !p !$acc end host_data – OpenACC end host_data c  
· · !p !$acc kernels – OpenACC kernels directive  
· · !p !$acc end kernels – OpenACC end kernels direct  
· · !p !$acc kernels loop – OpenACC kernels loop direc  
· · pr  
end · · !$acc end kernels loop – OpenACC end kernels l  
!$acc loop – OpenACC loop directive
```

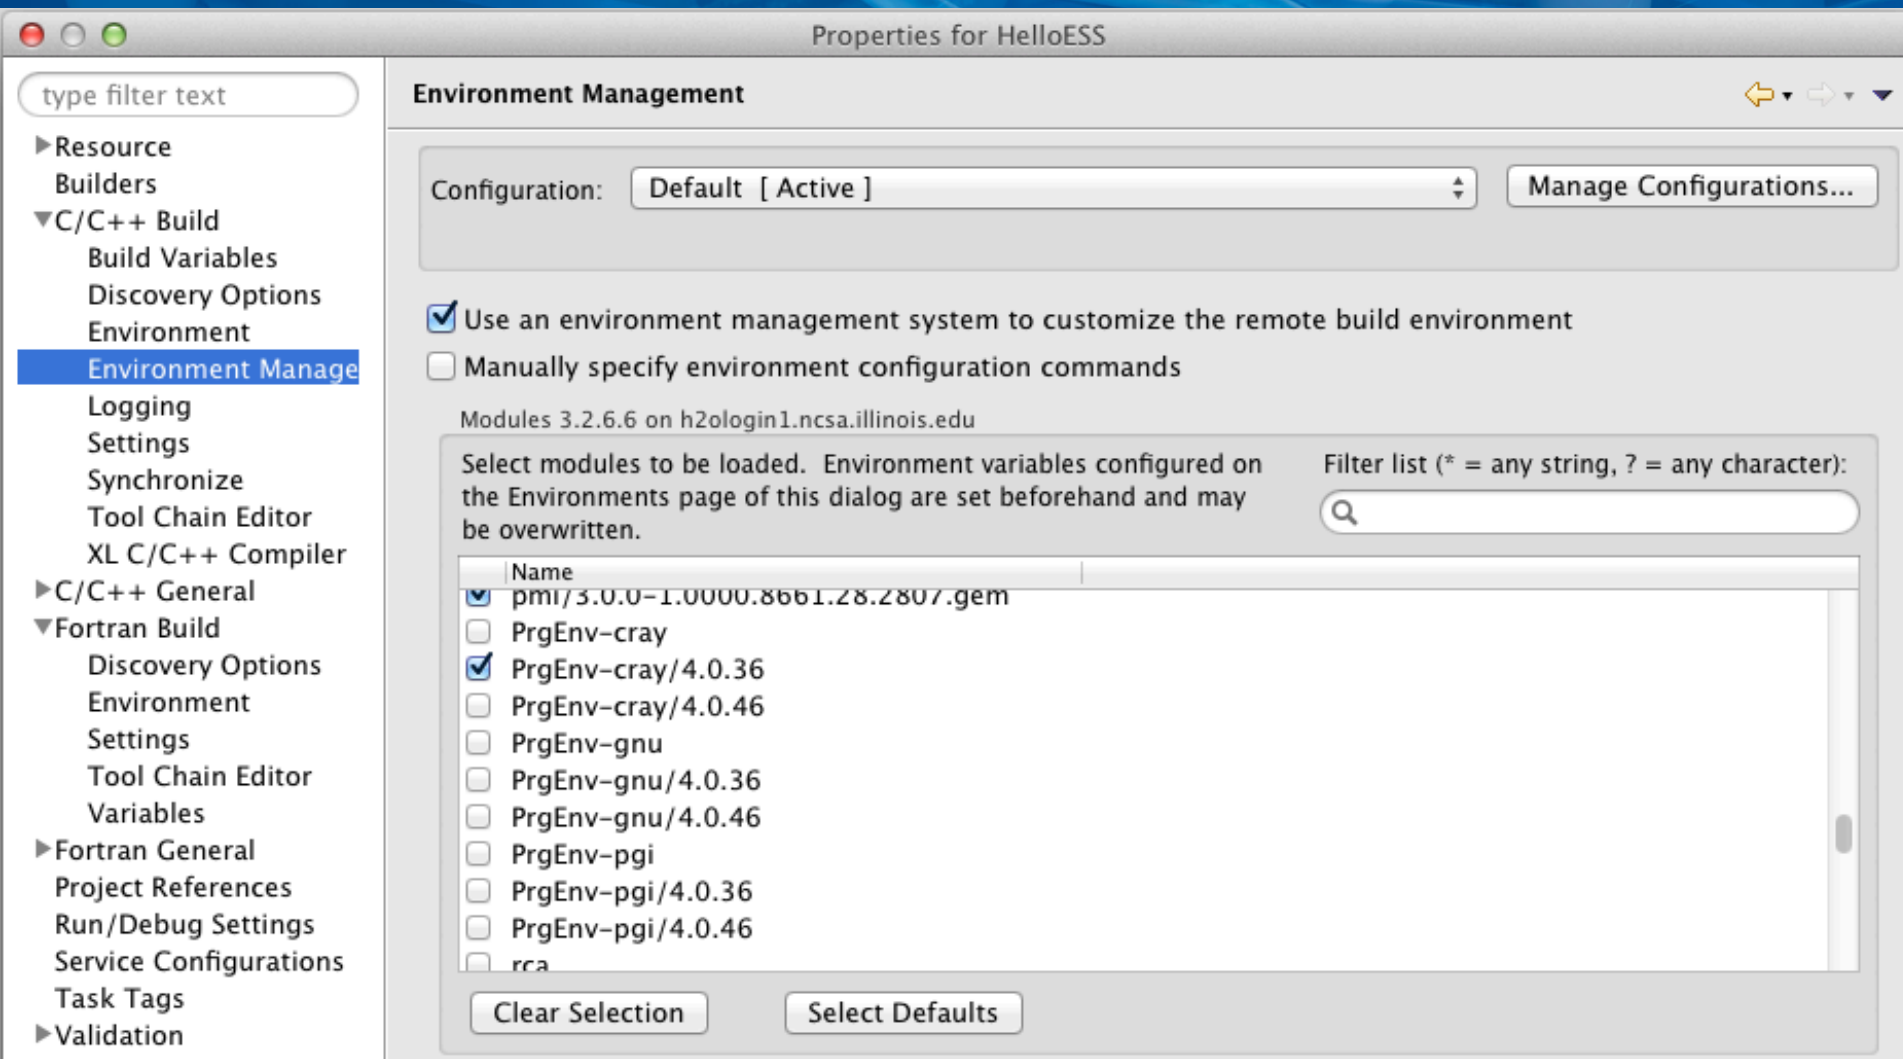
Code completion for OpenACC directives (added for BW)



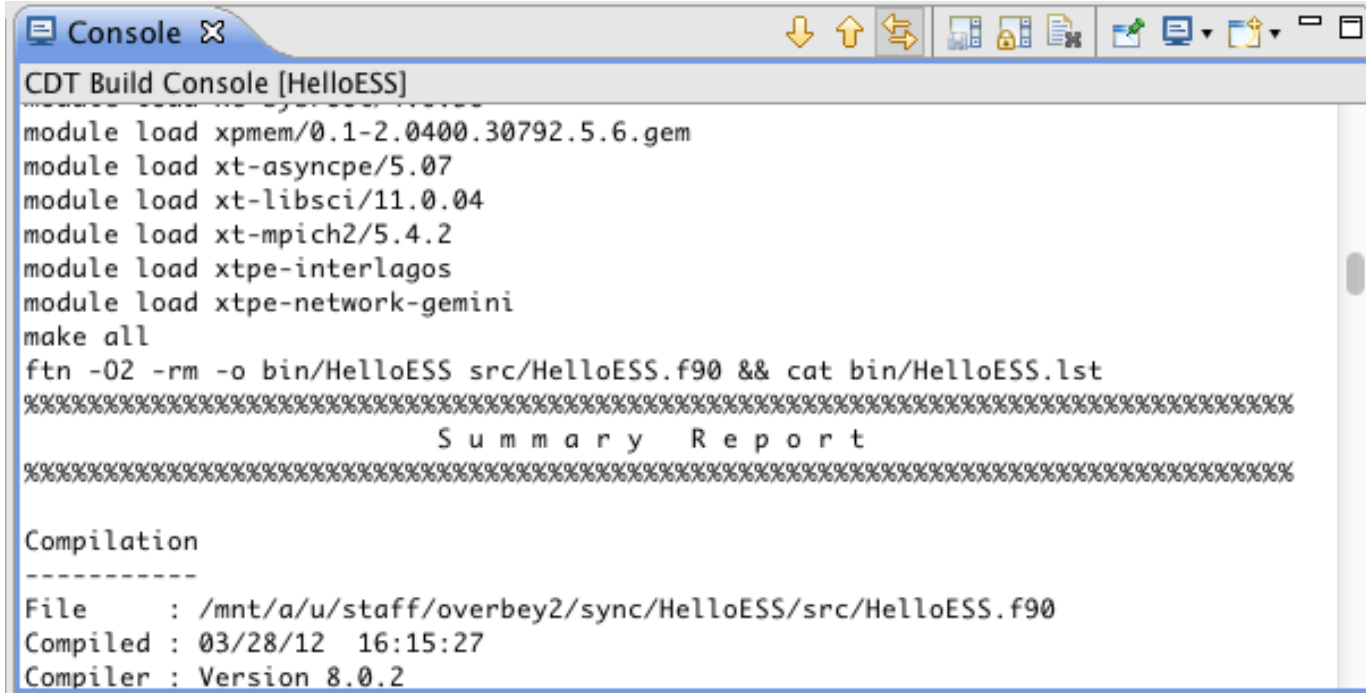
- Source code editing
- Code search/navigation
- Static analysis

- Compilation
- Running and debugging
- Performance tuning





Configuring environment modules for build (added for BW)



```
CDT Build Console [HelloESS]
module load xpmem/0.1-2.0400.30792.5.6.gem
module load xt-asyncpe/5.07
module load xt-libsci/11.0.04
module load xt-mpich2/5.4.2
module load xtpe-interlagos
module load xtpe-network-gemini
make all
ftn -O2 -rm -o bin/HelloESS src/HelloESS.f90 && cat bin/HelloESS.lst
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
                        Summary Report
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

Compilation
-----
File      : /mnt/a/u/staff/overbey2/sync/HelloESS/src/HelloESS.f90
Compiled  : 03/28/12 16:15:27
Compiler  : Version 8.0.2
```

Build is performed on remote machine (via SSH)


```

5  ..double precision :: a(SIZE, SIZE), b(SIZE, SIZE), c(SIZE, SIZE)
6
7  ..a(:) = 0.0 ..! This will raise a warning
8
9  ..do i = 1, SIZE
10 ..do j = 1, SIZE
11 ..a(i,j) = i*10.d0+j
12 ..b(i,j) = j*10.d0+i
13 ..end do
14 ..end do
15
16 !$acc parallel loop
17 ..do i = 1, 1000
18 ..c(:,i) = (a(:,i) + b(:,i)) / 2.d0
19 ..a(:,i) = (a(:,i) + c(:,i)) / 2.d0
20 ..b(:,i) = (b(:,i) + c(:,i)) / 2.d0
21 ..end do
22 !$acc end parallel loop
23
24 ..!print *, "Averages:"
25 ..!print *, sum(a(:,i))/(SIZE*SIZE)
26 ..!print *, sum(b(:,i))/(SIZE*SIZE)
27 ..!print *, sum(c(:,i))/(SIZE*SIZE)
28 ..print *, "Minimums:", minval(a(:,i)), minval(b(:,i)), minval(c(:,i))
29

```

After the build, compiler errors, warnings, and loopmark information are shown in the Problems view and source code editor

Problems	
1 error, 1 warning, 17 others	
Description	Resource
▼ ⚠ Warnings (1 item)	
⚠ The number of subscripts is smaller than the number of declared dimensions.	test1.f90
▼ ⓘ Infos (17 items)	
ⓘ A divide was turned into a multiply by a reciprocal	test1.f90
ⓘ A divide was turned into a multiply by a reciprocal	test1.f90
ⓘ A divide was turned into a multiply by a reciprocal	test1.f90
ⓘ A floating point expression involving an induction variable was strength reduced b...	test1.f90
ⓘ A loop nest at line 18 collapsed to a single loop.	test1.f90
ⓘ A loop starting at line 10 was not vectorized because a better candidate was found...	test1.f90
ⓘ A loop starting at line 17 was blocked with block size 512	test1.f90

(Cray, PGI support added for BW)

Resources
Application
Arguments
Environment
Synchronize
Common

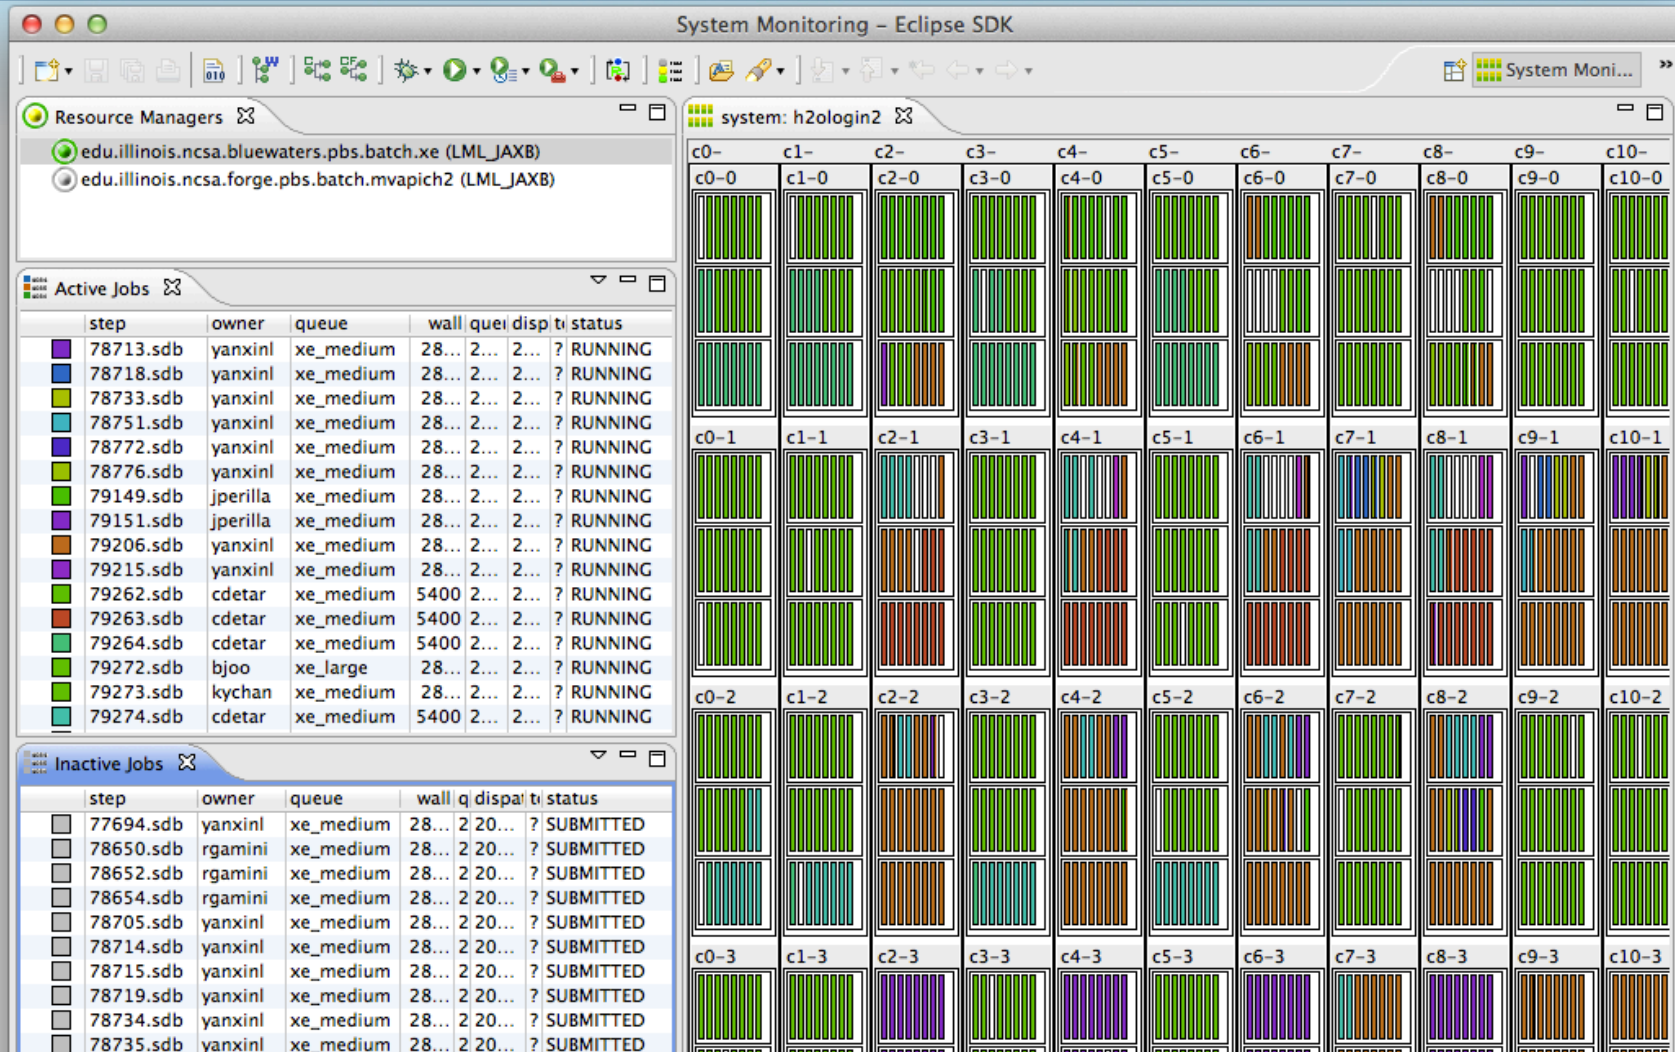
Resource Manager: ESS - Batch (XE)

Basic PBS Settings

Import PBS Script

Name	Value	Description
Total MPI Tasks:	32	Each XE6 node has two AMD Interlagos CPUs for a total of 32 integer cores and 16 floating point units per node. Therefore, the product of the number of MPI tasks per node and the number of OpenMP threads per task must be less than or equal to 32 (or 16 if running in single-stream mode). The number of MPI tasks per node must not exceed the total number of MPI tasks.
MPI Tasks per Node:	32	
OpenMP Threads per Process:		
Run in Dual-Stream Mode:	<input checked="" type="checkbox"/>	XE6 nodes are normally run in "dual-stream mode," where every integer core is allocated one task (i.e., one MPI task or one OpenMP thread). However, this means that every two tasks share a floating point unit. Some floating-point-intensive computations may need to run in "single-stream mode," where every other integer core is idle but every task has exclusive access to a floating point unit.
Job Name:	ptp_job	The name assigned to the job by the qsub or qalter command.
Account:		Account to which to charge this job.
Queue:		Designation of the queue to which to submit the job.
Total Memory Needed:		Maximum amount of memory used by all concurrent processes in the job.
Wallclock Time:	00:30:00	Maximum amount of real time during which the job can be in the running state.
Send E-mail:		Whether e-mail should be sent when the job begins, ends, or fails.
When to Send E-mail:		Whether e-mail should be sent when the job begins, ends, or fails.

Graphical interface for launching a job (customized for BW)



Graphical interface for system monitoring

Enabling Better Software Engineering

- Integrated static analyses, available in real time
- Language-aware code completion
- Language-aware code searching and navigation
- Automated refactoring
- Integrated documentation, available in real time
- Easy-to-use, graphical interfaces for
 - Version control (CVS, Subversion, Git)
 - Issue tracking (Bugzilla, Jira)

Toward Better Cray Support

- PTP parallel debugger does not yet work
 - DDT does not have Eclipse integration, either
- Craypat could be integrated with PTP
 - TAU integrated using PTP's *External Tools Framework (ETFw)*
- Refactorings could be built for OpenACC
- Loopmark information could be better integrated
 - E.g., used to suggest automated refactorings
 - ETFw *Feedback View* is designed to make this possible

Building a Cray-PTP Community

- Try PTP
 - See www.eclipse.org/ptp
 - Tutorials given at Supercomputing, XSEDE
 - Slides online: see wiki.eclipse.org/PTP/tutorials
- Ask questions and give feedback
 - Join the ptp-user mailing list
- Join the open source developer community
 - Join the ptp-dev mailing list
- How could PTP benefit your organization?