Eclipse PTP Support for UT TACC Stampede Progress Report

Brian Watt, briandwatt@gmail.com
Doug James, UT TACC, djames@tacc.utexas.edu

September 13, 2013

Agenda

- UT TACC Stampede
- SLURM sbatch Support
- Stampede Target System Configuration
- Module Support
- Inject Commands into Batch Script
- System Monitoring View
- Acknowledgements

UT TACC Stampede - 1

- Dell PowerEdge C8220 Cluster with Intel Xeon Phi Many Integrated Core (MIC) CoProcessors
  - Cluster contains 6,400+ Dell Zeus C8220 Nodes
  - Typical node consists of
    - Two Xeon Intel 8-Core 64-bit E5-processors w/ 32GB – for a total of 2.2 PF
    - One or two 61-core Xeon Phi MIC coprocessor w/ 8GB – for an additional 7+ PF
  - Specialty nodes – 1TB large mem, GPU, etc.
- CentOS 6.3 OS
- SLURM 2.4 w/ mods
- Intel Compilers & Libs

UT TACC Stampede - 2

- Different ways to run on MIC coprocessor
  - Native – run on MIC (serial, MPI, OpenMP)
  - Offload – from host offload work to MIC
  - Symmetric – MPI across one or more hosts and MICs
- Initial Eclipse PTP support for MIC provides
  - Native and Offload
  - Future – Symmetric (dual executable launch)

SLURM sbatch Support

- Based upon slurm-generic-batch XML file in org.eclipse.ptp.rm.jaxb.configs
- Noticed that it did NOT contain all sbatch command line arguments. Why? So added
  - For example, acctg_freq, clusters, comment, constraint, contiguous, cores_per_socket, cpu_bind, cpus_per_task, distribution, exclude, exclusive, export_file, extra_node_info, gres, hint, input, licenses, mem, mem_per_cpu, mem_bind, mincpus, network, nodefile, nodelist, no_kill, no_requeue, ntasks, ntasks_per_core, ntasks_per_node, ntasks_per_socket, overcommit, propagate, qos, requeue, share, sockets_per_node, switches, threads_per_core, time_min, tmp, uid, wait_all_nodes, wckey
- Wrote Bug 416962 - Update slurm-generic-batch with all document sbatch command line arguments
  - There are multiple SLURM XML files for ALPS, BGP, BGQ, and generic - only change generic one
  - A future enhancement – UT TACC not interested in just generic SLURM

Stampede Environment

- Fast moving infrastructure
  - Intel software stack – Use Intel Eclipse Plugins
  - Support both mvapich2 and Intel MPI libraries
  - Propagating inherited environment to MIC
  - Symmetric MPI – ibrun.symm -c <host exec> -m <MIC exec>
- Custom SLURM batch XML file – edu.utexas.tacc.stampede.batch in org.eclipse.ptp.rm.jaxb.configs
  - Customize which sbatch arguments needed
  - No srun instead replaced with ibrun
  - Explicitly set environment variables - MIC_OMP_NUM_THREADS, MIC_PPN, and OMP_NUM_THREADS
  - Add Module support (LMOD) as GUI control

Stampede Target System Configuration

- Wrote Bug 412925 - UT Ranger decommissioned - remove edu.utexas.tacc.ranger.sge.batch.xml – DONE, checked into master
- Wrote Bug 412926 - Add UT Stampede - add edu.utexas.tacc.stampede.slurm.batch.xml – Work in progress, checked into master
- After adding all missing sbatch command line arguments, reviewed with Doug James and other TACC personnel
- Upon review needed to adjust the generic SLURM batch XML file
  - Some definitely worked and were primary - 10
  - Some definitely worked and were secondary - 7
  - Some might work and were secondary - 12
  - Some were not supported and should be removed! - 27
  - Some were discouraged/conflict and should be removed! - 5
  - Some were uncertain and should be removed! – 9
- Can NOT use generic SLURM batch, but has to reorganized and simplified it

Basic Settings

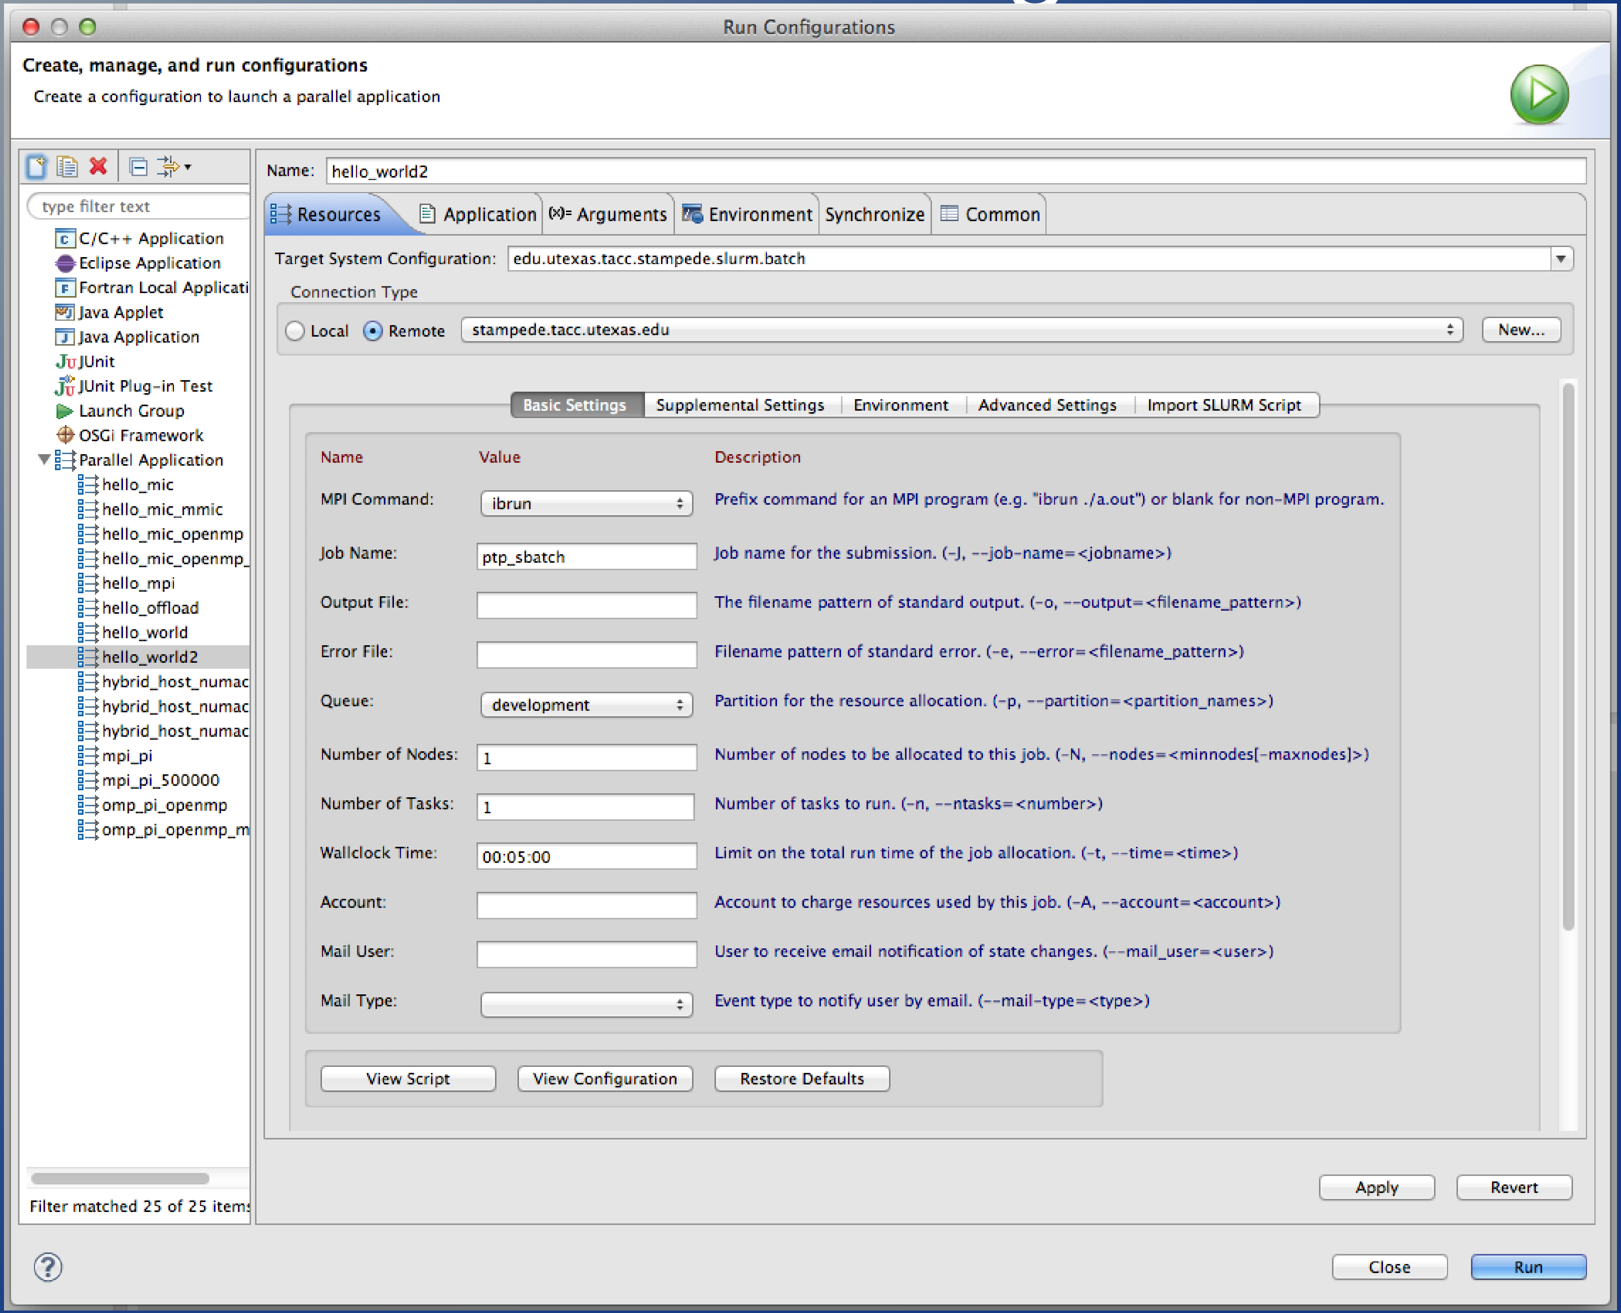

Create, manage, and run configurations

Create a configuration to launch a parallel application

Name: hello_world2

Target System Configuration: edu.utexas.tacc.stampede.slurm.batch

Connection Type

○ Local ● Remote stampede.tacc.utexas.edu

Name	Value	Description
MPI Command:	ibrun	Prefix command for an MPI program (e.g. "ibrun ./a.out") or blank for non-MPI program.
Job Name:	ptp_sbatch	Job name for the submission. (-J, --job-name=<jobname>)
Output File:		The filename pattern of standard output. (-o, --output=<filename_pattern>)
Error File:		Filename pattern of standard error. (-e, --error=<filename_pattern>)
Queue:	development	Partition for the resource allocation. (-p, --partition=<partition_names>)
Number of Nodes:	1	Number of nodes to be allocated to this job. (-N, --nodes=<minnodes[-maxnodes]>)
Number of Tasks:	1	Number of tasks to run. (-n, --ntasks=<number>)
Wallclock Time:	00:05:00	Limit on the total run time of the job allocation. (-t, --time=<time>)
Account:		Account to charge resources used by this job. (-A, --account=<account>)
Mail User:		User to receive email notification of state changes. (--mail_user=<user>)
Mail Type:		Event type to notify user by email. (--mail-type=<type>)

Supplemental Settings

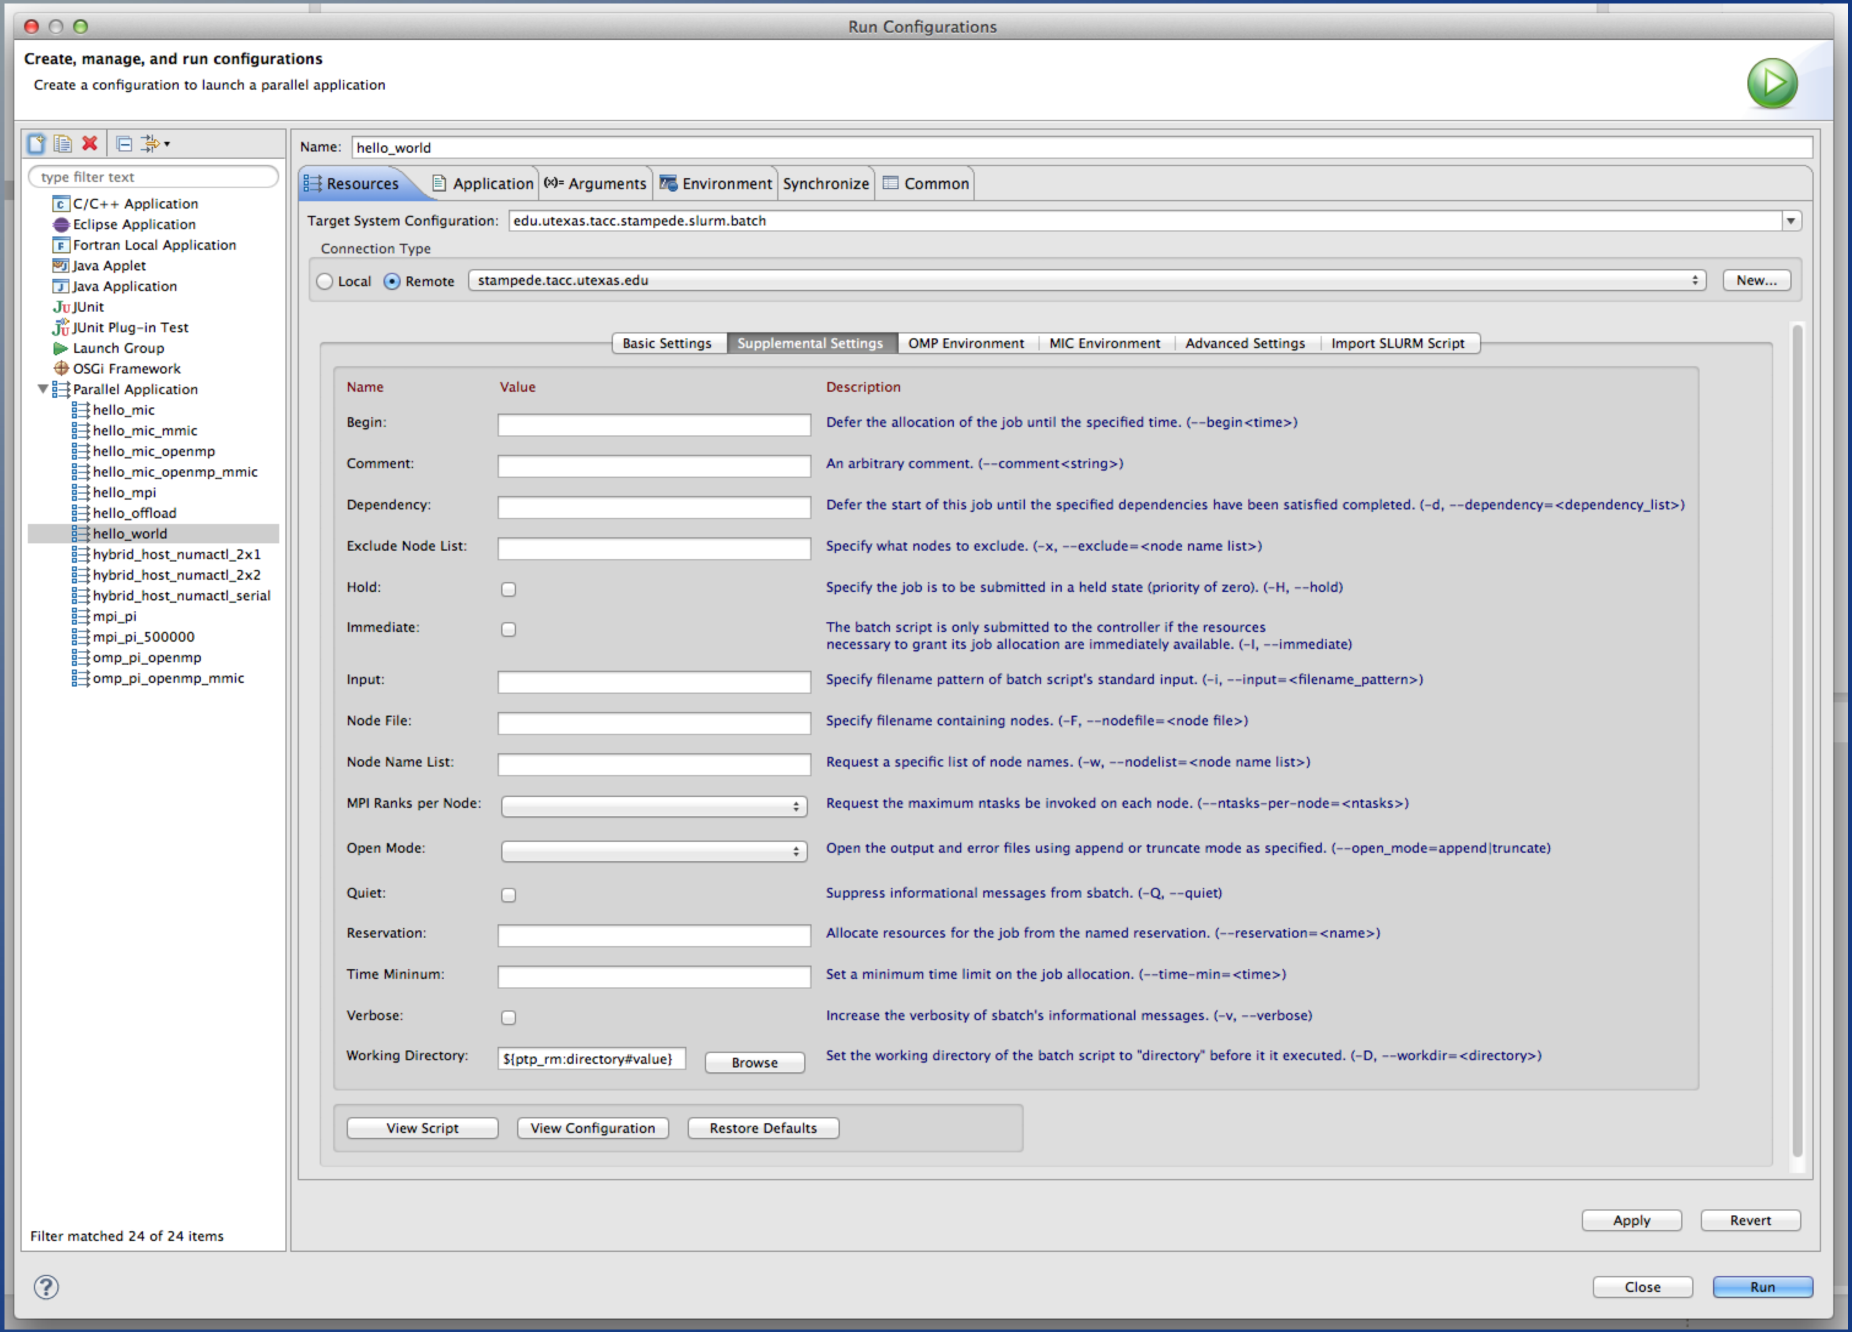

Run Configurations

Create, manage, and run configurations

Create a configuration to launch a parallel application

type filter text

- C/C++ Application
- Eclipse Application
- Fortran Local Application
- Java Applet
- Java Application
- JUnit
- JUnit Plug-in Test
- Launch Group
- OSGi Framework
- Parallel Application
  - hello_mic
  - hello_mic_mmic
  - hello_mic_openmp
  - hello_mic_openmp_mmic
  - hello_mpi
  - hello_offload
  - hello_world
  - hybrid_host_numactl_2x1
  - hybrid_host_numactl_2x2
  - hybrid_host_numactl_serial
  - mpi_pi
  - mpi_pi_500000
  - omp_pi_openmp
  - omp_pi_openmp_mmic

Filter matched 24 of 24 items

Name: hello_world

Target System Configuration: edu.utexas.tacc.stampede.slurm.batch

Connection Type

○ Local ● Remote stampede.tacc.utexas.edu New...

Name	Value	Description
Begin:		Defer the allocation of the job until the specified time. (--begin<time>)
Comment:		An arbitrary comment. (--comment<string>)
Dependency:		Defer the start of this job until the specified dependencies have been satisfied completed. (-d, --dependency=<dependency_list>)
Exclude Node List:		Specify what nodes to exclude. (-x, --exclude=<node name list>)
Hold:	□	Specify the job is to be submitted in a held state (priority of zero). (-H, --hold)
Immediate:	□	The batch script is only submitted to the controller if the resources necessary to grant its job allocation are immediately available. (-I, --immediate)
Input:		Specify filename pattern of batch script's standard input. (-i, --input=<filename_pattern>)
Node File:		Specify filename containing nodes. (-F, --nodefile=<node file>)
Node Name List:		Request a specific list of node names. (-w, --nodelist=<node name list>)
MPI Ranks per Node:		Request the maximum ntasks be invoked on each node. (--ntasks-per-node=<ntasks>)
Open Mode:		Open the output and error files using append or truncate mode as specified. (--open_mode=append|truncate)
Quiet:	□	Suppress informational messages from sbatch. (-Q, --quiet)
Reservation:		Allocate resources for the job from the named reservation. (--reservation=<name>)
Time Mininum:		Set a minimum time limit on the job allocation. (--time-min=<time>)
Verbose:	□	Increase the verbosity of sbatch's informational messages. (-v, --verbose)
Working Directory:	${ptp_rm:directory#value} Browse	Set the working directory of the batch script to "directory" before it it executed. (-D, --workdir=<directory>)

Environment

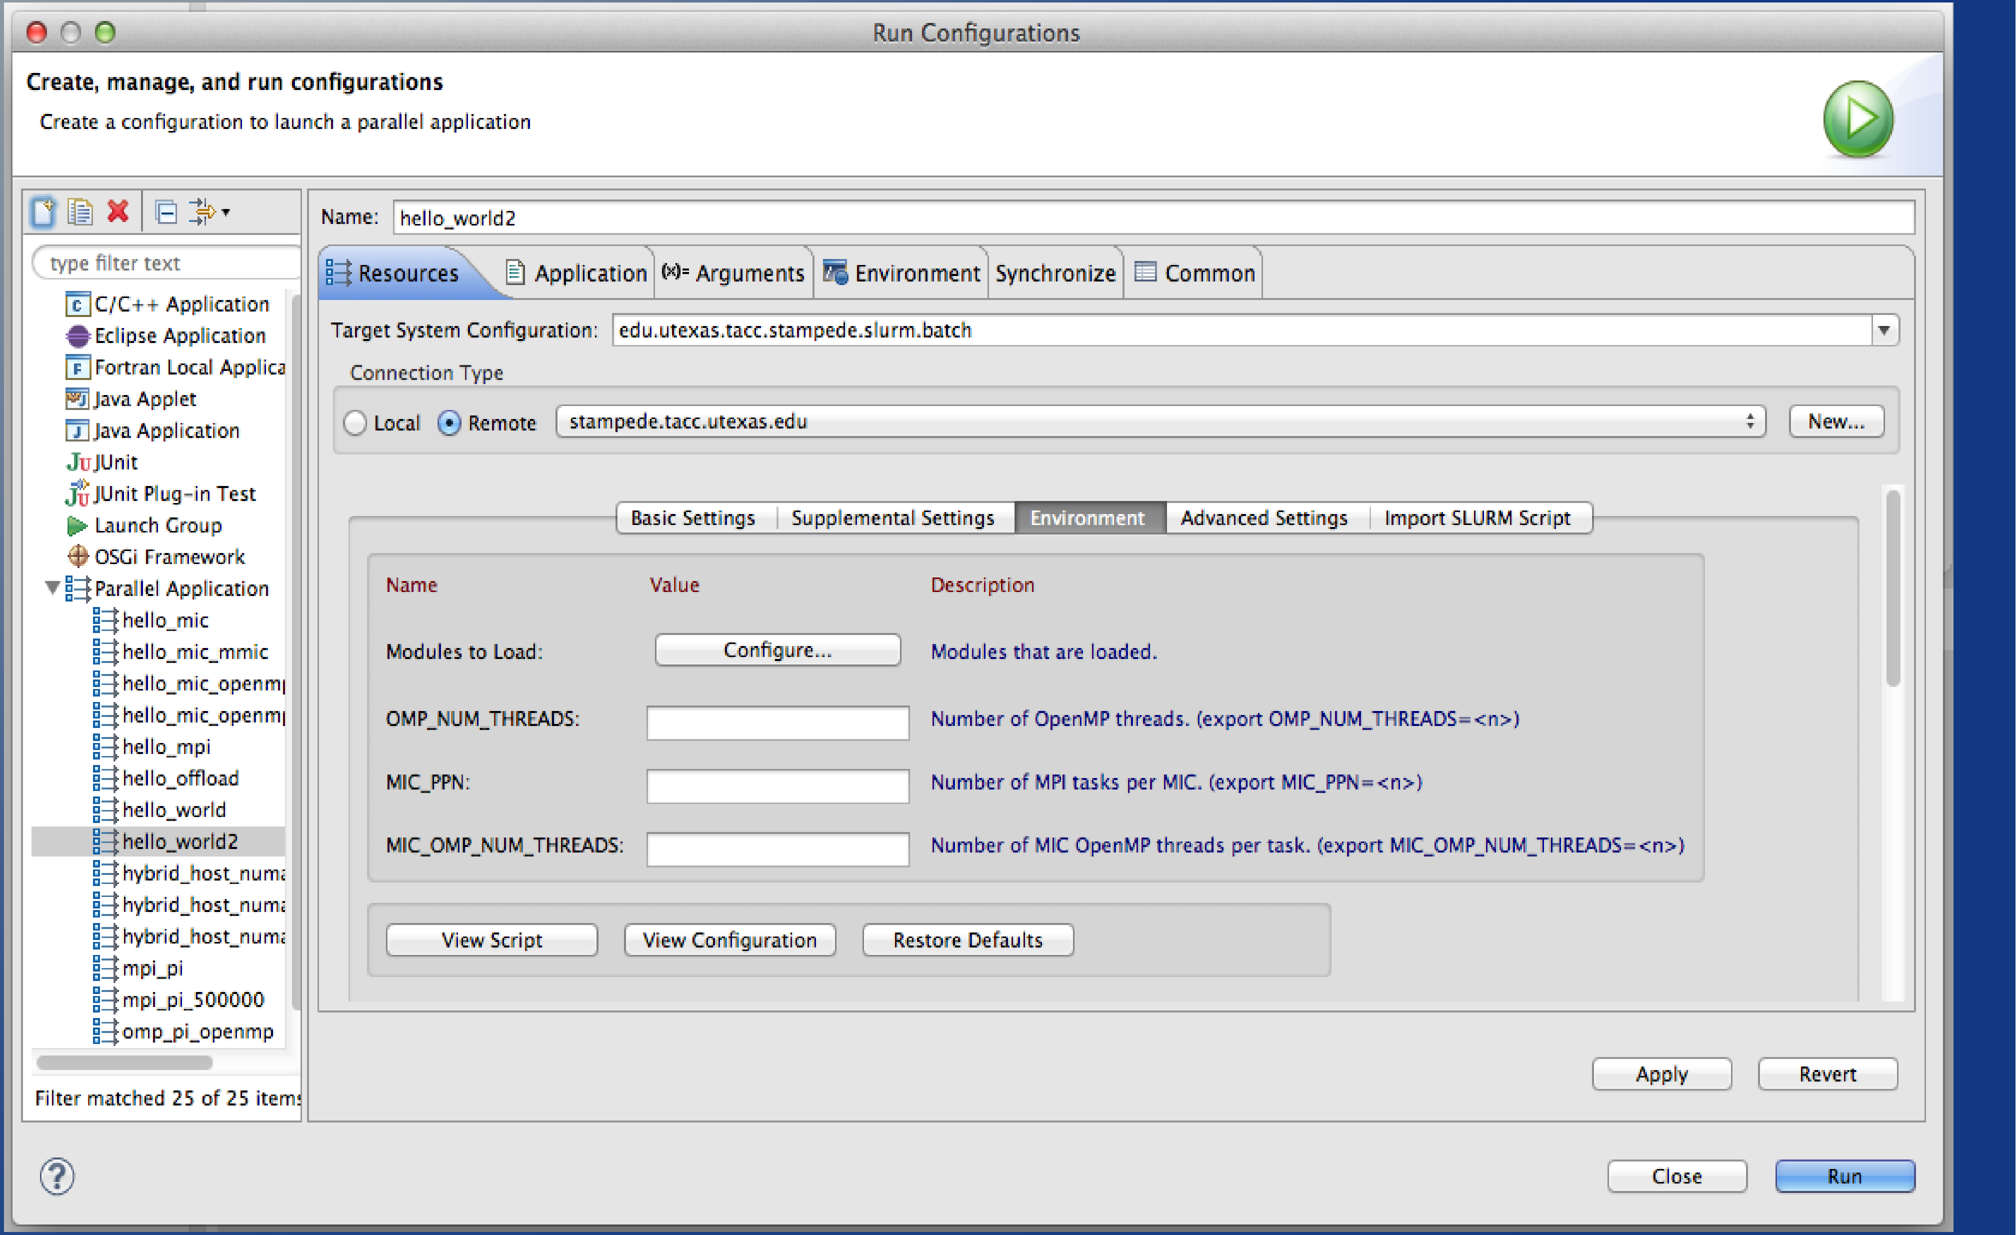

Run Configurations

Create, manage, and run configurations

Create a configuration to launch a parallel application

Name: hello_world2

Target System Configuration: edu.utexas.tacc.stampede.slurm.batch

Connection Type

○ Local ● Remote stampede.tacc.utexas.edu New...

Name	Value	Description
Modules to Load:	Configure...	Modules that are loaded.
OMP_NUM_THREADS:		Number of OpenMP threads. (export OMP_NUM_THREADS=<n>)
MIC_PPN:		Number of MPI tasks per MIC. (export MIC_PPN=<n>)
MIC_OMP_NUM_THREADS:		Number of MIC OpenMP threads per task. (export MIC_OMP_NUM_THREADS=<n>)

Advanced Settings

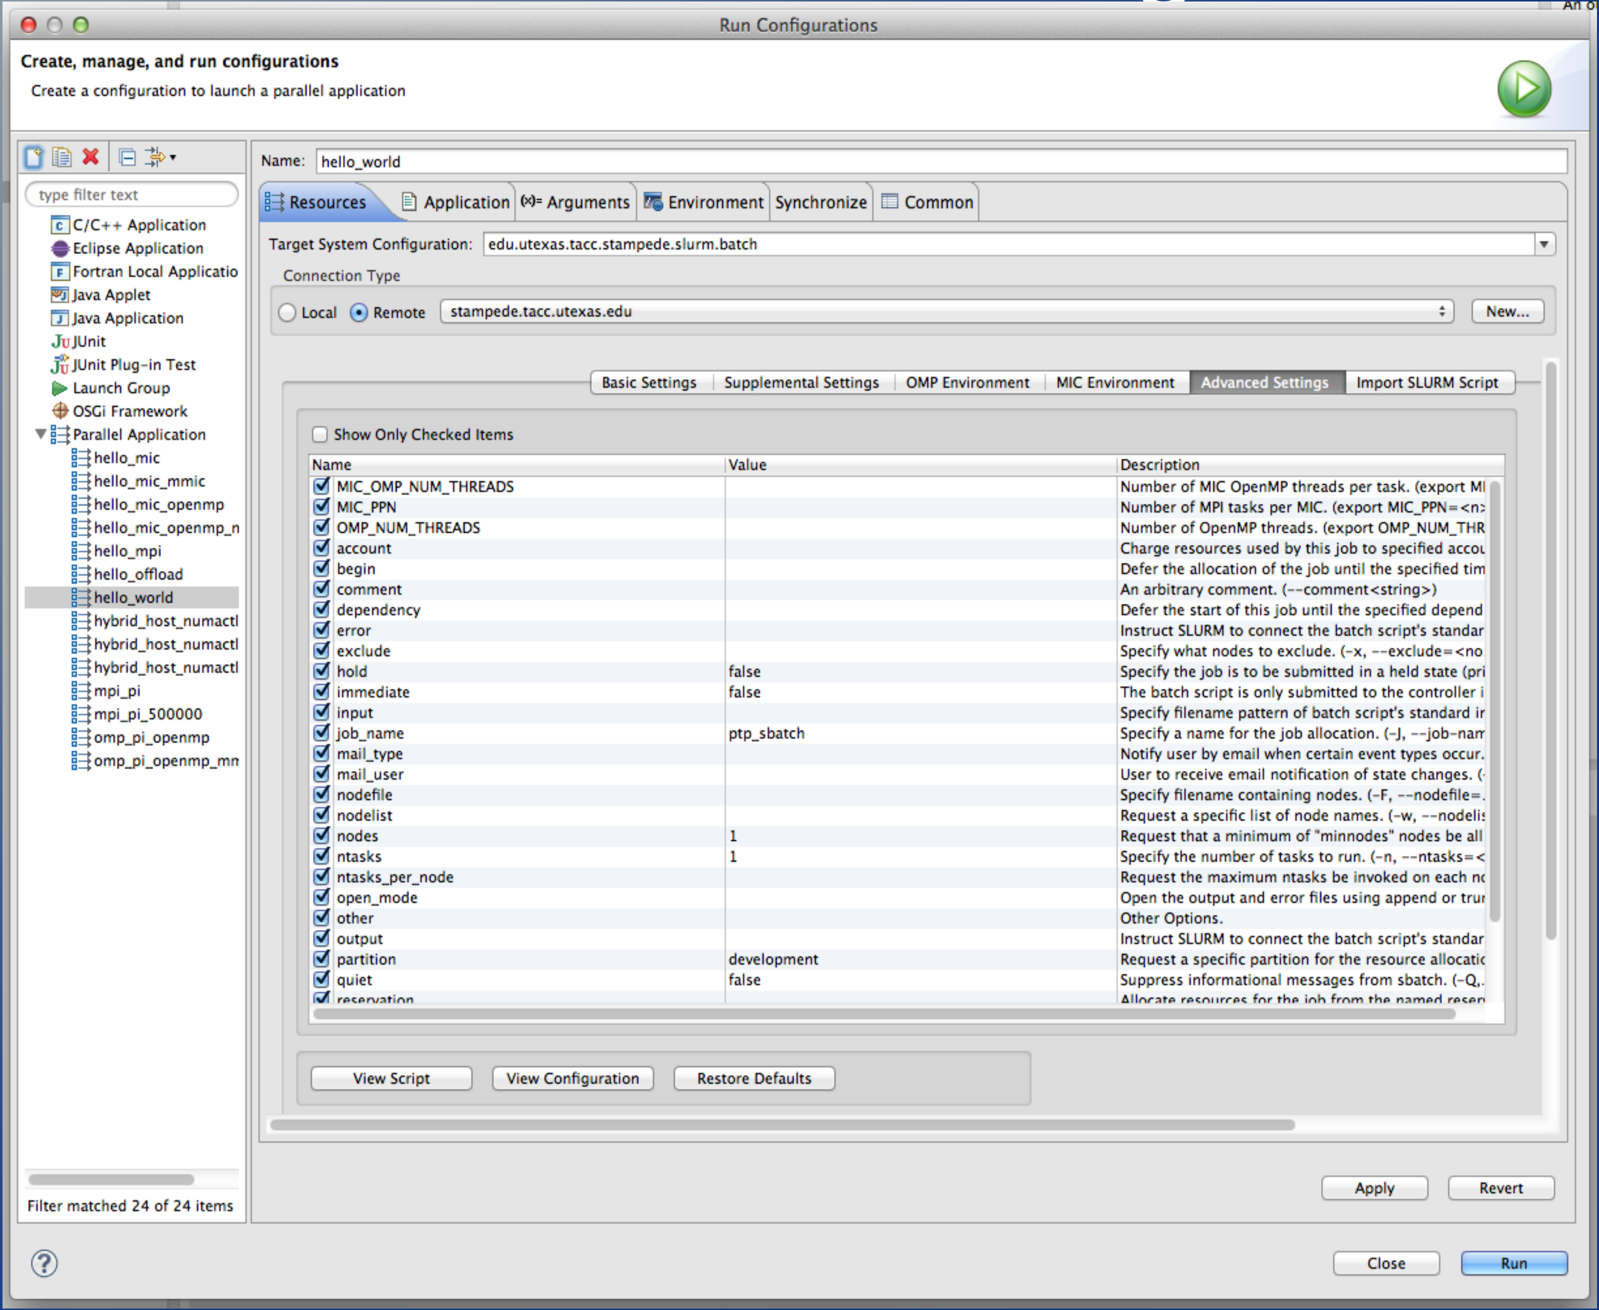

Create, manage, and run configurations

Create a configuration to launch a parallel application

Name: hello_world

Target System Configuration: edu.utexas.tacc.stampede.slurm.batch

Connection Type

○ Local ● Remote stampede.tacc.utexas.edu

□ Show Only Checked Items

Name	Value	Description
☑ MIC_OMP_NUM_THREADS		Number of MIC OpenMP threads per task. (export MI
☑ MIC_PPN		Number of MPI tasks per MIC. (export MIC_PPN=<n:
☑ OMP_NUM_THREADS		Number of OpenMP threads. (export OMP_NUM_THR
☑ account		Charge resources used by this job to specified accou
☑ begin		Defer the allocation of the job until the specified tim
☑ comment		An arbitrary comment. (--comment<string>)
☑ dependency		Defer the start of this job until the specified depend
☑ error		Instruct SLURM to connect the batch script's standar
☑ exclude		Specify what nodes to exclude. (-x, --exclude=<no
☑ hold	false	Specify the job is to be submitted in a held state (pri
☑ immediate	false	The batch script is only submitted to the controller i
☑ input		Specify filename pattern of batch script's standard ir
☑ job_name	ptp_sbatch	Specify a name for the job allocation. (-J, --job-nam
☑ mail_type		Notify user by email when certain event types occur.
☑ mail_user		User to receive email notification of state changes. (
☑ nodefile		Specify filename containing nodes. (-F, --nodefile=.
☑ nodelist		Request a specific list of node names. (-w, --nodelis
☑ nodes	1	Request that a minimum of "minnodes" nodes be all
☑ ntasks	1	Specify the number of tasks to run. (-n, --ntasks=<
☑ ntasks_per_node		Request the maximum ntasks be invoked on each nc
☑ open_mode		Open the output and error files using append or trui
☑ other		Other Options.
☑ output		Instruct SLURM to connect the batch script's standar
☑ partition	development	Request a specific partition for the resource allocatic
☑ quiet	false	Suppress informational messages from sbatch. (-Q,.
☑ reservation		Allocate resources for the job from the named reser

Import SLURM Script

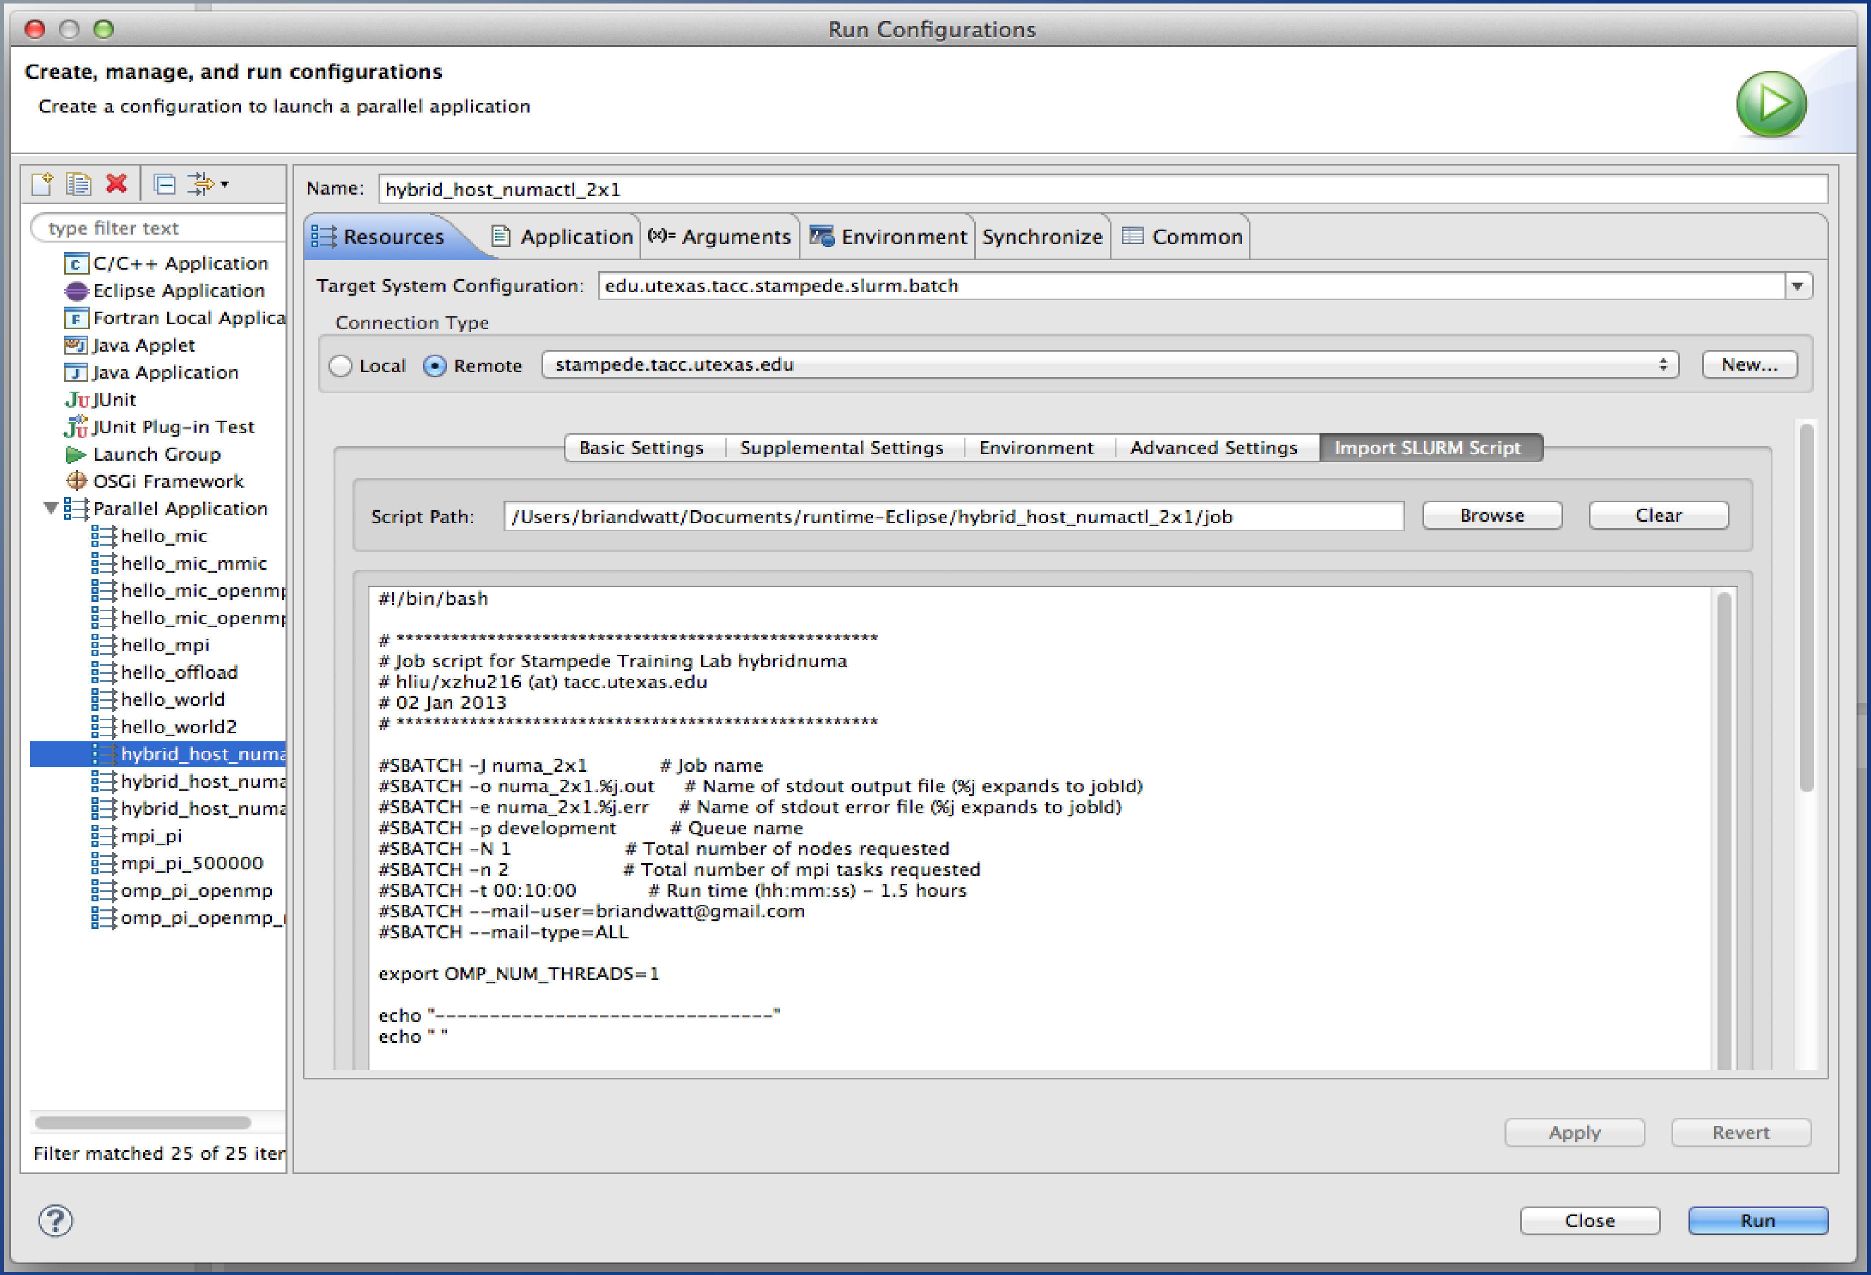

Create, manage, and run configurations

Create a configuration to launch a parallel application

Name: hybrid_host_numactl_2x1

Target System Configuration: edu.utexas.tacc.stampede.slurm.batch

Connection Type

○ Local ● Remote stampede.tacc.utexas.edu

Script Path: /Users/briandwatt/Documents/runtime-Eclipse/hybrid_host_numactl_2x1/job

```
#!/bin/bash

# ****************************************************
# Job script for Stampede Training Lab hybridnuma
# hliu/xzhu216 (at) tacc.utexas.edu
# 02 Jan 2013
# ****************************************************

#SBATCH -J numa_2x1           # Job name
#SBATCH -o numa_2x1.%j.out     # Name of stdout output file (%j expands to jobId)
#SBATCH -e numa_2x1.%j.err     # Name of stdout error file (%j expands to jobId)
#SBATCH -p development         # Queue name
#SBATCH -N 1                   # Total number of nodes requested
#SBATCH -n 2                   # Total number of mpi tasks requested
#SBATCH -t 00:10:00            # Run time (hh:mm:ss) - 1.5 hours
#SBATCH --mail-user=briandwatt@gmail.com
#SBATCH --mail-type=ALL

export OMP_NUM_THREADS=1

echo "--------------------------------"
echo " "
```

Module Support

- Presently generates in script – fetches once at start
  - `module purge > /dev/null 2>&1`
  - `module load <module>` - repeated
- Stampede uses LMOD environmental module system
  - Supports hierarchical modules
  - Supports named saved environments
- Future (for LMOD on Stampede only) – refetches for each change because restore affects unloads affects loads
  - Nothing or `module restore <name>` or `module reset`
  - `module unload <module>` - repeated
  - `module load <module>` - repeated
  - `module list` - optional

Current Module Support GUI

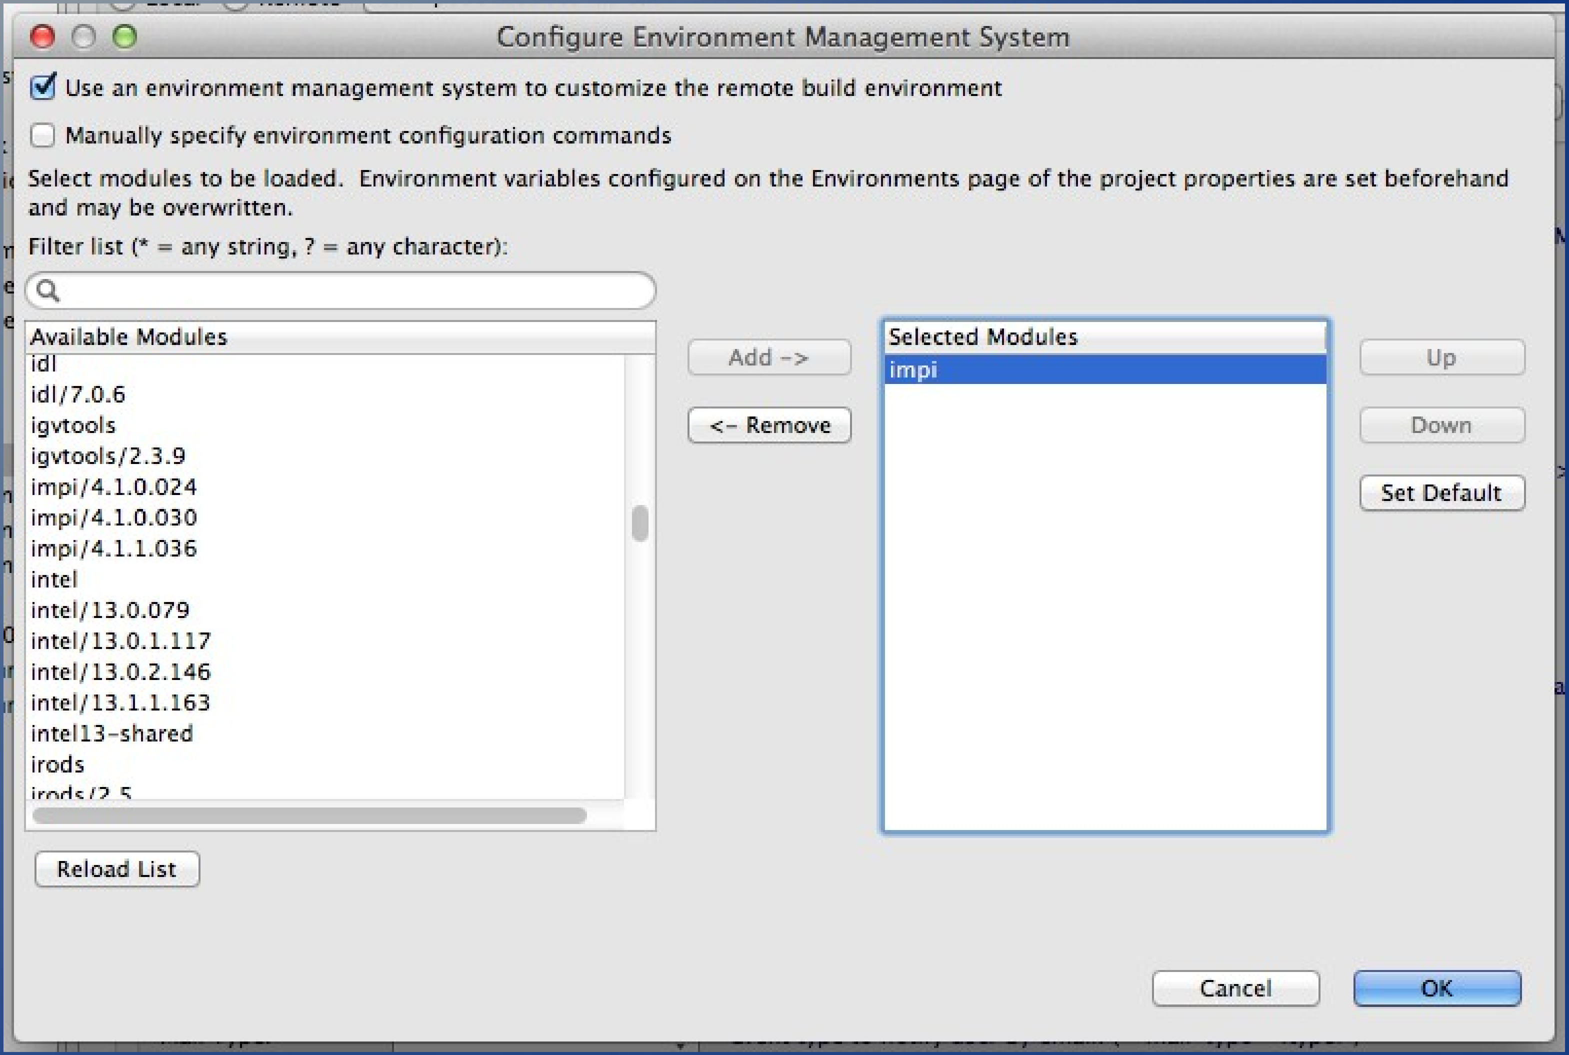

Configure Environment Management System

☑ Use an environment management system to customize the remote build environment

□ Manually specify environment configuration commands

Select modules to be loaded. Environment variables configured on the Environments page of the project properties are set beforehand and may be overwritten.

Filter list (* = any string, ? = any character):

Available Modules

idl
idl/7.0.6
igvtools
igvtools/2.3.9
impi/4.1.0.024
impi/4.1.0.030
impi/4.1.1.036
intel
intel/13.0.079
intel/13.0.1.117
intel/13.0.2.146
intel/13.1.1.163
intel13-shared
irods
irods/2.5

Selected Modules

impi

Future Stampede Module Support GUI

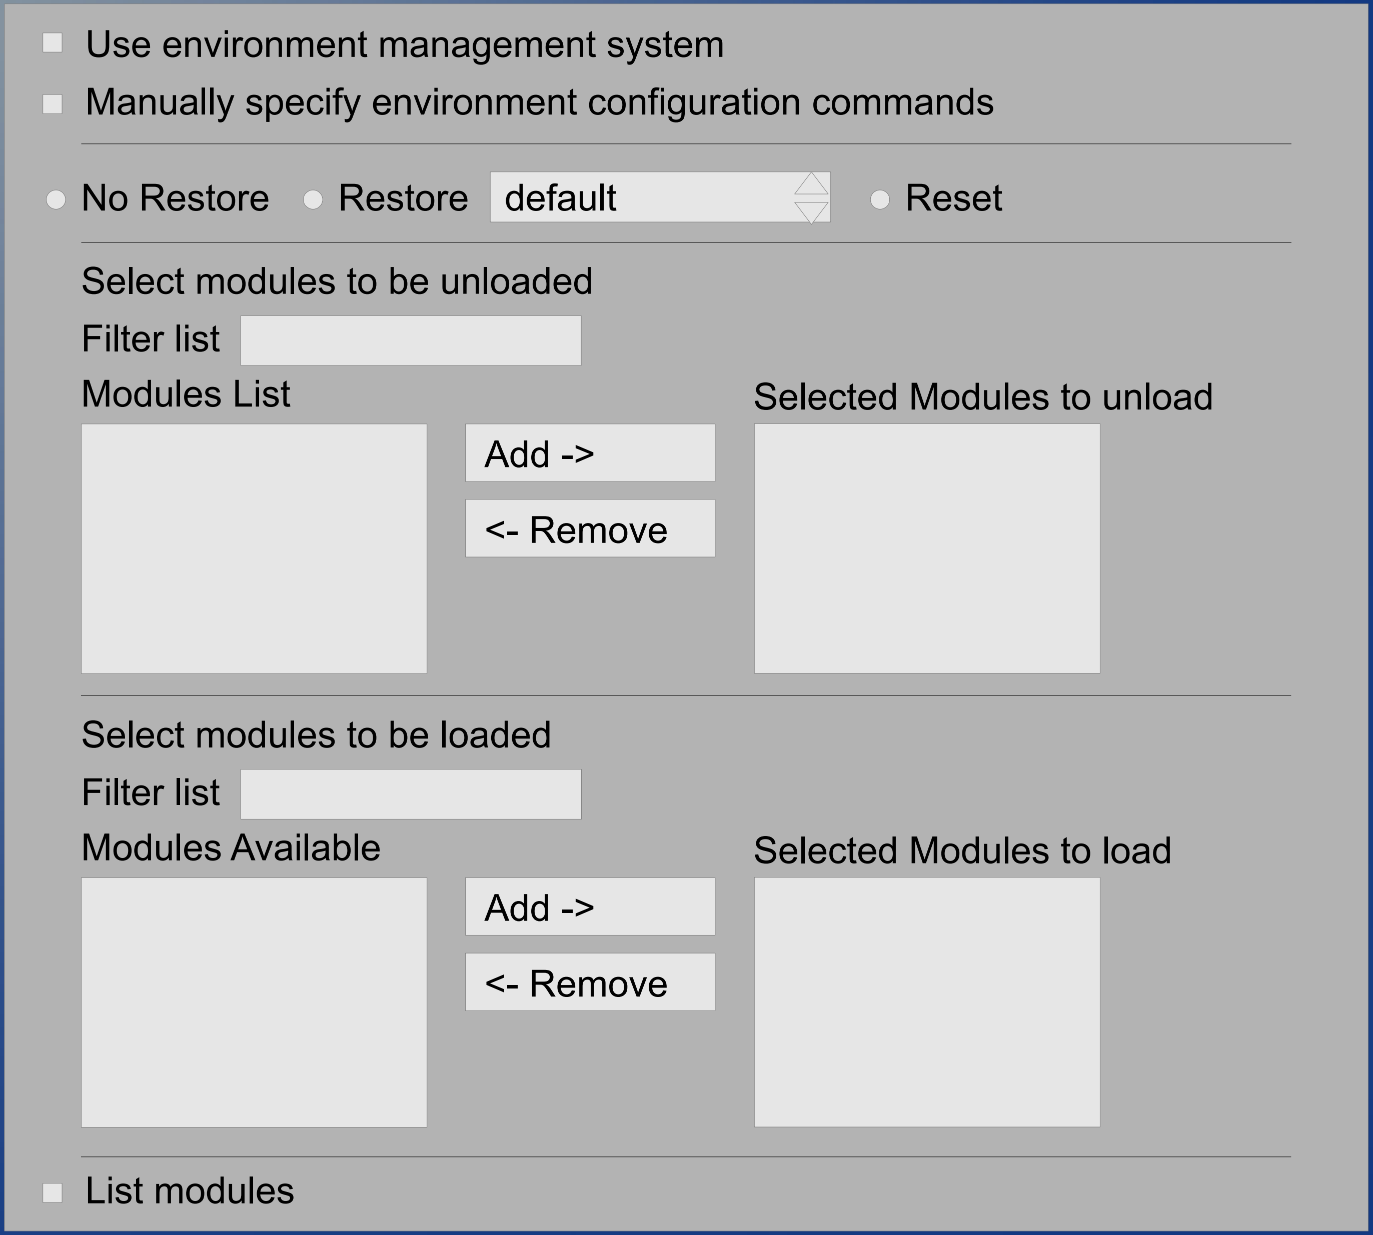

□ Use environment management system

□ Manually specify environment configuration commands

○ No Restore ○ Restore default ○ Reset

Select modules to be unloaded

Filter list

Modules List

Selected Modules to unload

Select modules to be loaded

Filter list

Modules Available

Selected Modules to load

□ List modules

Inject Commands into Batch Script

- Need for 'escape' capability to support more intermediate to advanced users
- Injection of commands into batch script
  - Provides custom processing and/or setup prior to application launch
- Option 1 - User Specified Module Commands (a fudge)
- Option 2 - Propose a new RM 'custom' tab
  - For example, `<inject title=”Additional Lines”>`
  - Provides editor text area where the user enters one or more commands
  - Injects commands after module commands, and before application launch

User Specified Module Commands

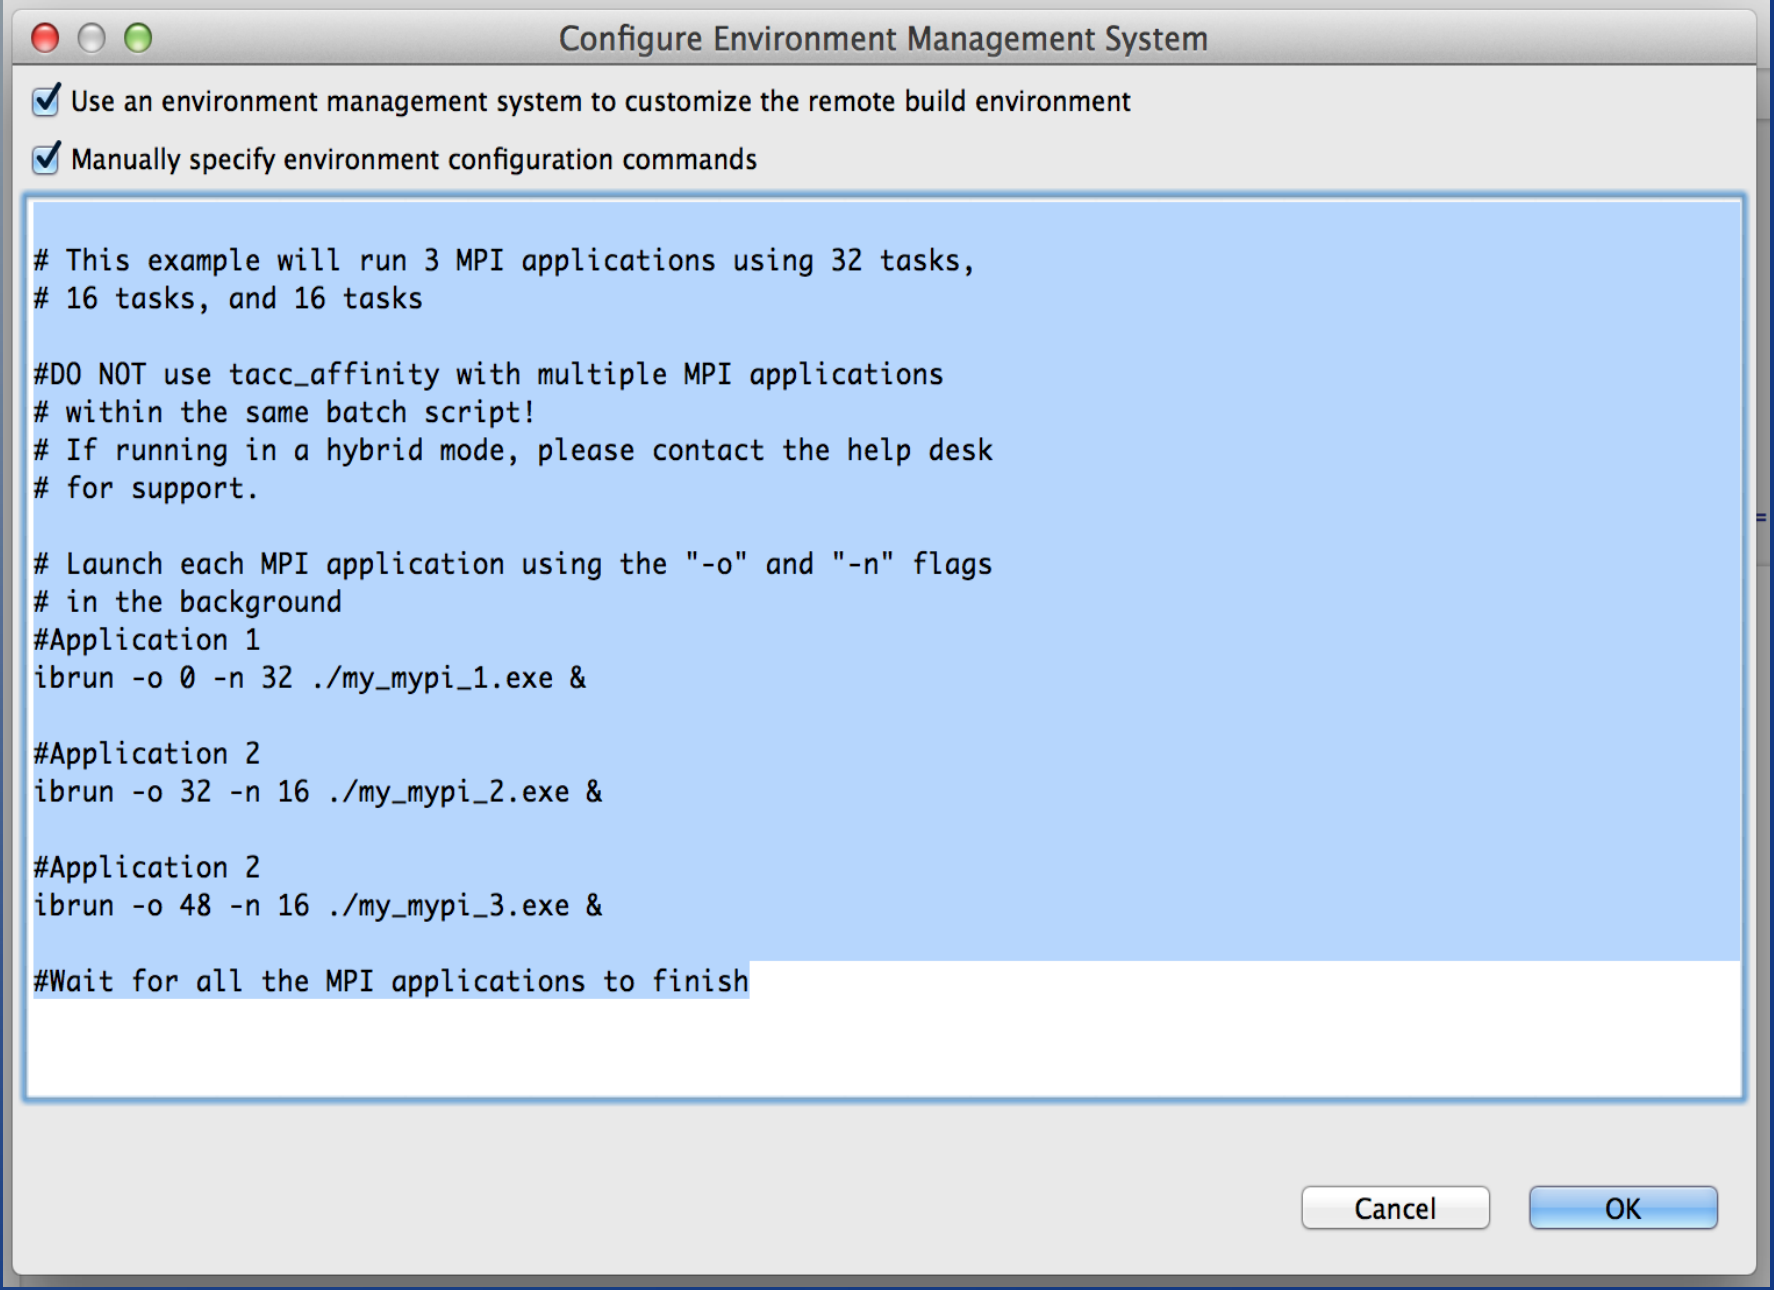

☑ Use an environment management system to customize the remote build environment

☑ Manually specify environment configuration commands

```
# This example will run 3 MPI applications using 32 tasks,
# 16 tasks, and 16 tasks

#DO NOT use tacc_affinity with multiple MPI applications
# within the same batch script!
# If running in a hybrid mode, please contact the help desk
# for support.

# Launch each MPI application using the "-o" and "-n" flags
# in the background
#Application 1
ibrun -o 0 -n 32 ./my_mypi_1.exe &

#Application 2
ibrun -o 32 -n 16 ./my_mypi_2.exe &

#Application 2
ibrun -o 48 -n 16 ./my_mypi_3.exe &

#Wait for all the MPI applications to finish
```

Resulting Batch Script

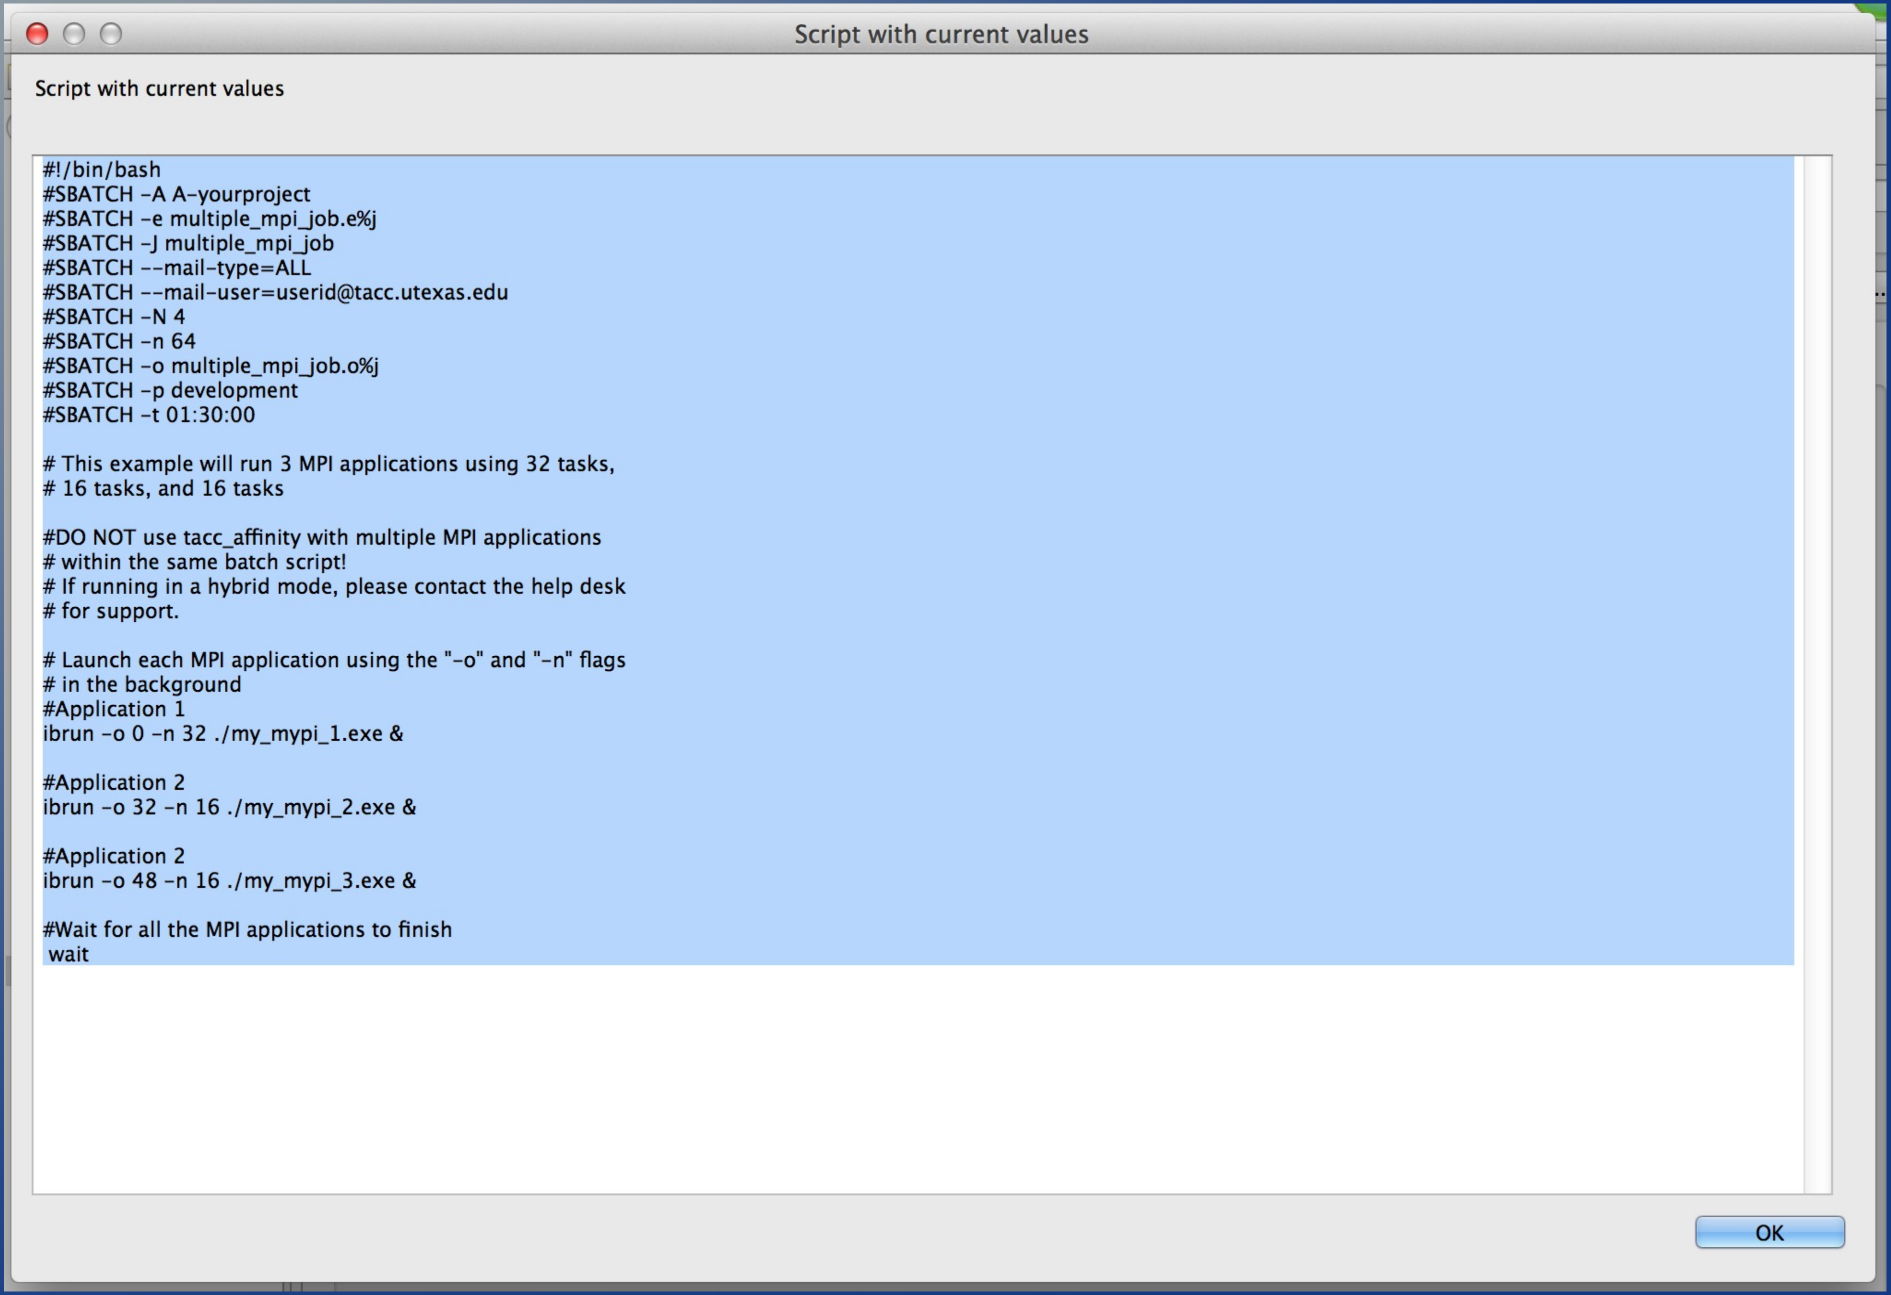

Script with current values

```
#!/bin/bash
#SBATCH -A A-yourproject
#SBATCH -e multiple_mpi_job.e%j
#SBATCH -J multiple_mpi_job
#SBATCH --mail-type=ALL
#SBATCH --mail-user=userid@tacc.utexas.edu
#SBATCH -N 4
#SBATCH -n 64
#SBATCH -o multiple_mpi_job.o%j
#SBATCH -p development
#SBATCH -t 01:30:00

# This example will run 3 MPI applications using 32 tasks,
# 16 tasks, and 16 tasks

#DO NOT use tacc_affinity with multiple MPI applications
# within the same batch script!
# If running in a hybrid mode, please contact the help desk
# for support.

# Launch each MPI application using the "-o" and "-n" flags
# in the background
#Application 1
ibrun -o 0 -n 32 ./my_mypi_1.exe &

#Application 2
ibrun -o 32 -n 16 ./my_mypi_2.exe &

#Application 2
ibrun -o 48 -n 16 ./my_mypi_3.exe &

#Wait for all the MPI applications to finish
 wait
```

System Monitor Display

- Based upon recent presentation: Customizing the PTP Monitoring Layout by Carsten Karbach
    - Custom LML-Layout
    - Define Machine Topology – TBD
    - Setup/Usage
- Update for UT TACC Stampede specifics
- Details/specifics working with Doug James, and consultation with Carsten Karbach
- Review refresh/update performance

Current System Monitor Display

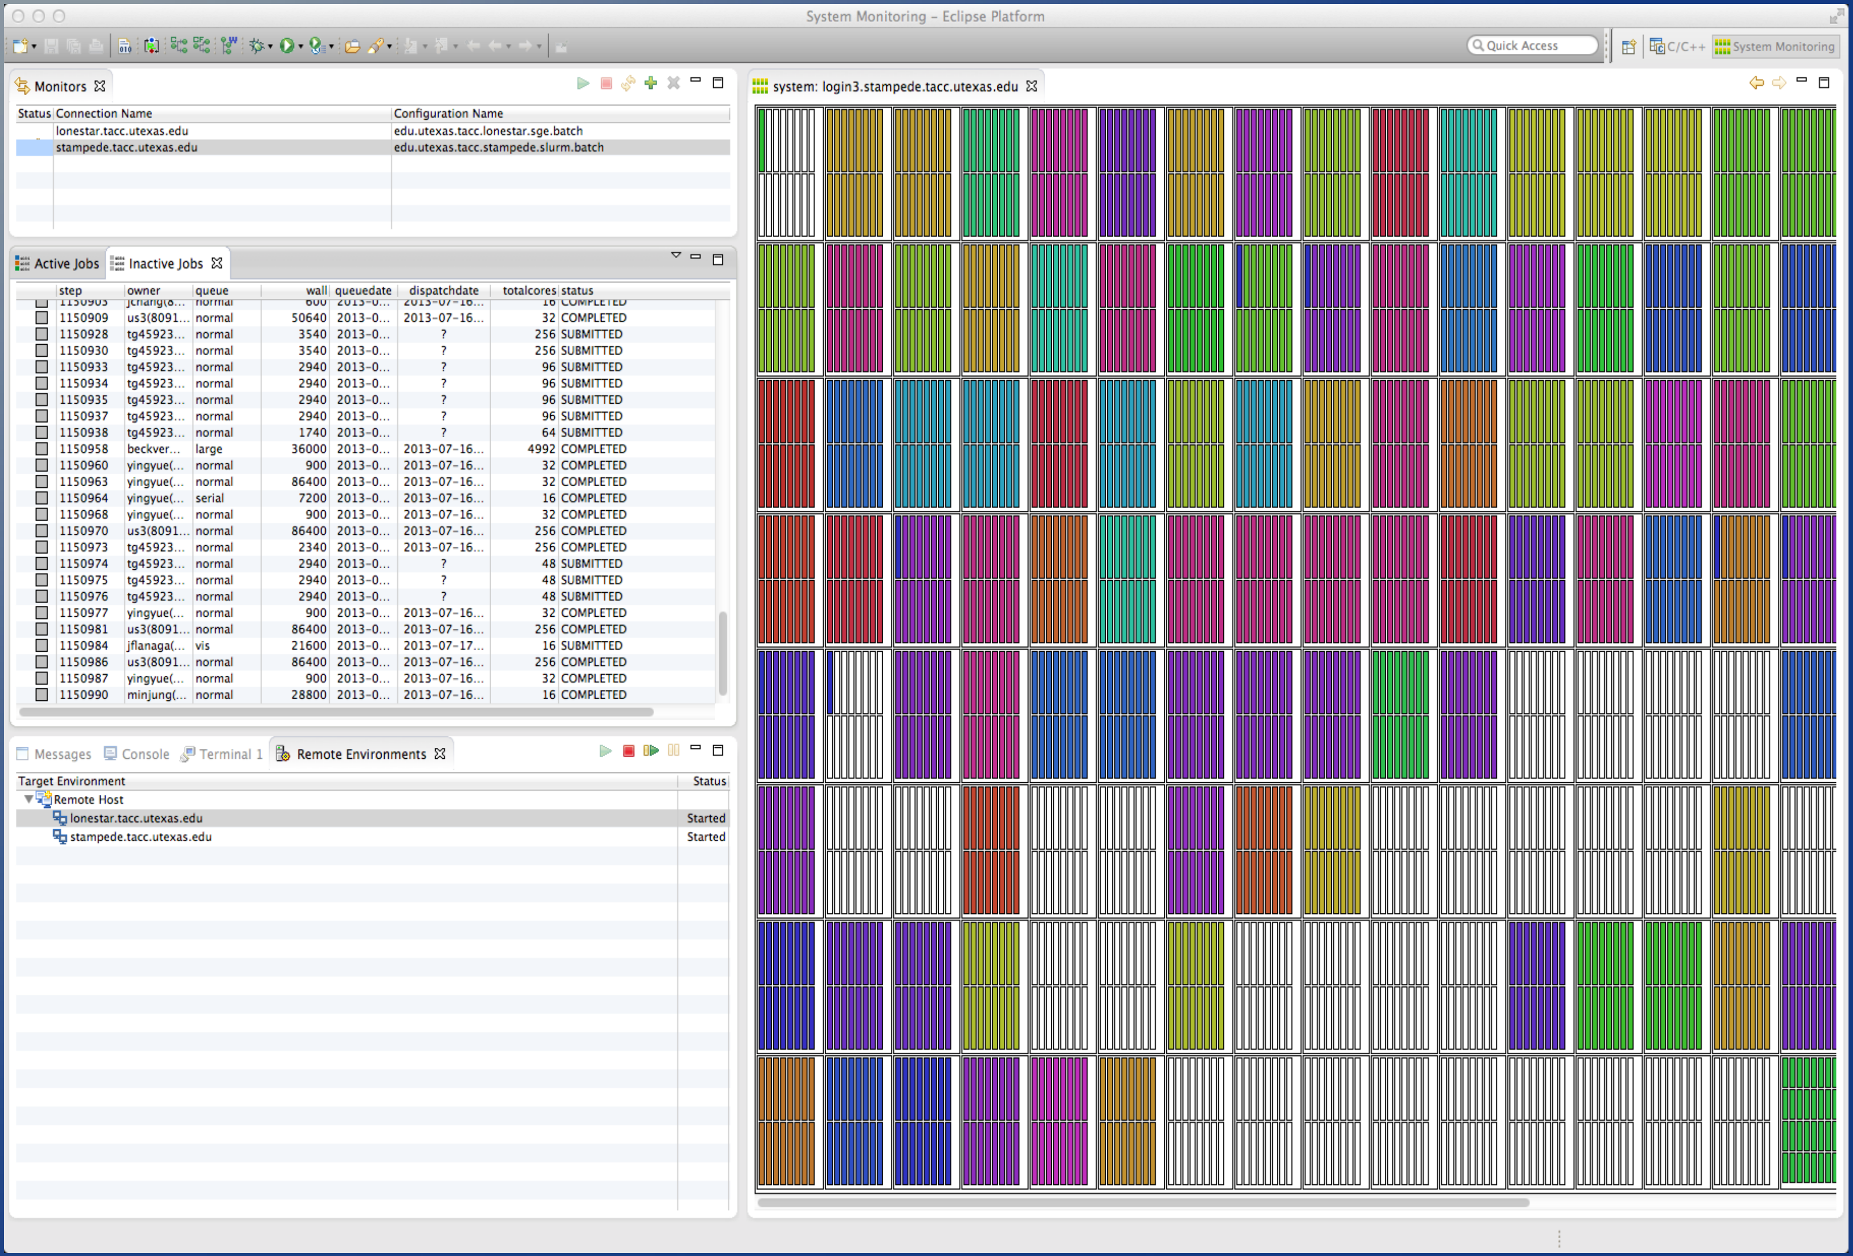

Status	Connection Name	Configuration Name
	lonestar.tacc.utexas.edu	edu.utexas.tacc.lonestar.sge.batch
	stampede.tacc.utexas.edu	edu.utexas.tacc.stampede.slurm.batch

	step	owner	queue	wall	queuedate	dispatchdate	totalcores	status
□	1150903	jchang(8...	normal	600	2013-0...	2013-07-16...	16	COMPLETED
□	1150909	us3(8091...	normal	50640	2013-0...	2013-07-16...	32	COMPLETED
□	1150928	tg45923...	normal	3540	2013-0...	?	256	SUBMITTED
□	1150930	tg45923...	normal	3540	2013-0...	?	256	SUBMITTED
□	1150933	tg45923...	normal	2940	2013-0...	?	96	SUBMITTED
□	1150934	tg45923...	normal	2940	2013-0...	?	96	SUBMITTED
□	1150935	tg45923...	normal	2940	2013-0...	?	96	SUBMITTED
□	1150937	tg45923...	normal	2940	2013-0...	?	96	SUBMITTED
□	1150938	tg45923...	normal	1740	2013-0...	?	64	SUBMITTED
□	1150958	beckver...	large	36000	2013-0...	2013-07-16...	4992	COMPLETED
□	1150960	yingyue(...	normal	900	2013-0...	2013-07-16...	32	COMPLETED
□	1150963	yingyue(...	normal	86400	2013-0...	2013-07-16...	32	COMPLETED
□	1150964	yingyue(...	serial	7200	2013-0...	2013-07-16...	16	COMPLETED
□	1150968	yingyue(...	normal	900	2013-0...	2013-07-16...	32	COMPLETED
□	1150970	us3(8091...	normal	86400	2013-0...	2013-07-16...	256	COMPLETED
□	1150973	tg45923...	normal	2340	2013-0...	2013-07-16...	256	COMPLETED
□	1150974	tg45923...	normal	2940	2013-0...	?	48	SUBMITTED
□	1150975	tg45923...	normal	2940	2013-0...	?	48	SUBMITTED
□	1150976	tg45923...	normal	2940	2013-0...	?	48	SUBMITTED
□	1150977	yingyue(...	normal	900	2013-0...	2013-07-16...	32	COMPLETED
□	1150981	us3(8091...	normal	86400	2013-0...	2013-07-16...	256	COMPLETED
□	1150984	jflanaga(...	vis	21600	2013-0...	2013-07-17...	16	SUBMITTED
□	1150986	us3(8091...	normal	86400	2013-0...	2013-07-16...	256	COMPLETED
□	1150987	yingyue(...	normal	900	2013-0...	2013-07-16...	32	COMPLETED
□	1150990	minjung(...	normal	28800	2013-0...	2013-07-16...	16	COMPLETED

Target Environment	Status
Remote Host	
lonestar.tacc.utexas.edu	Started
stampede.tacc.utexas.edu	Started

Future System Monitor Display

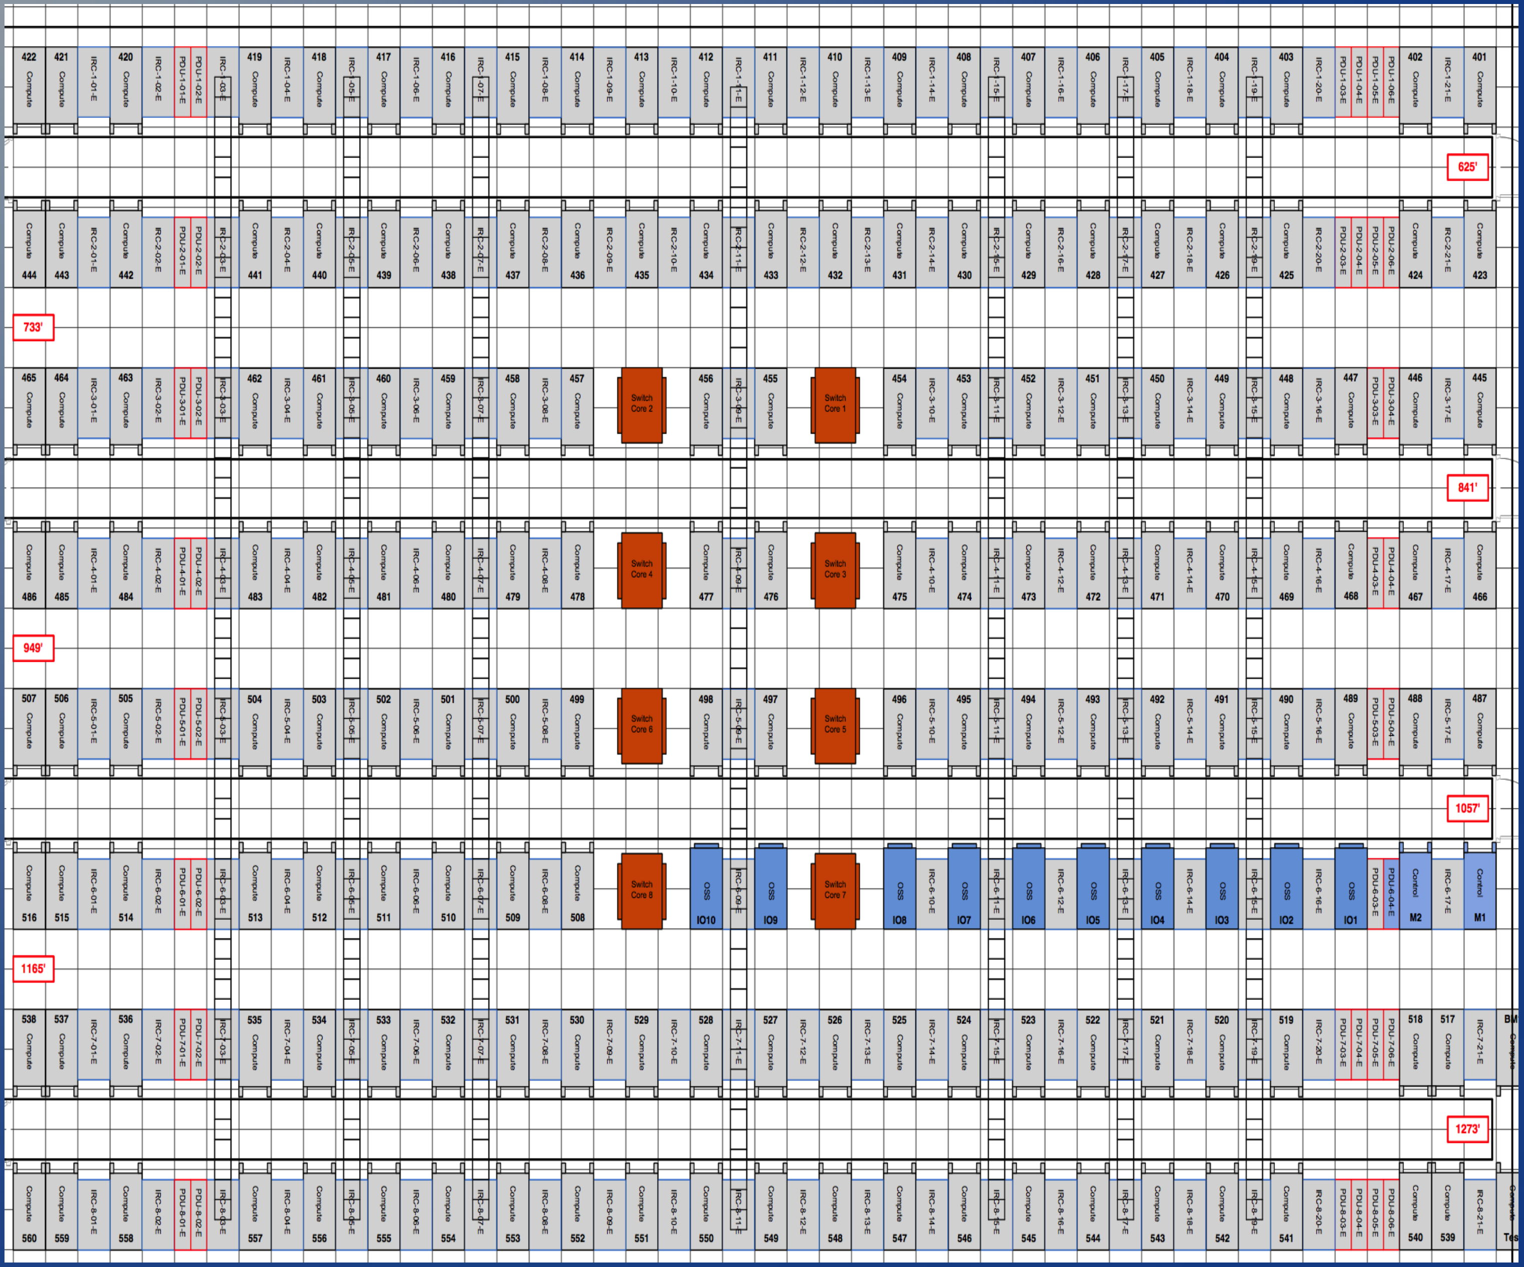
422 Compute 421 Compute IRC-1-01-E 420 Compute IRC-1-02-E PDU-1-01-E PDU-1-02-E IRC-1-03-E 419 Compute IRC-1-04-E 418 Compute IRC-1-05-E 417 Compute IRC-1-06-E 416 Compute IRC-1-07-E 415 Compute IRC-1-08-E 414 Compute IRC-1-09-E 413 Compute IRC-1-10-E 412 Compute IRC-1-11-E 411 Compute IRC-1-12-E 410 Compute IRC-1-13-E 409 Compute IRC-1-14-E 408 Compute IRC-1-15-E 407 Compute IRC-1-16-E 406 Compute IRC-1-17-E 405 Compute IRC-1-18-E 404 Compute IRC-1-19-E 403 Compute IRC-1-20-E PDU-1-03-E PDU-1-04-E PDU-1-05-E PDU-1-06-E 402 Compute IRC-1-21-E 401 Compute
625'
Compute 444 Compute 443 IRC-2-01-E Compute 442 IRC-2-02-E PDU-2-01-E PDU-2-02-E IRC-2-03-E Compute 441 IRC-2-04-E Compute 440 IRC-2-05-E Compute 439 IRC-2-06-E Compute 438 IRC-2-07-E Compute 437 IRC-2-08-E Compute 436 IRC-2-09-E Compute 435 IRC-2-10-E Compute 434 IRC-2-11-E Compute 433 IRC-2-12-E Compute 432 IRC-2-13-E Compute 431 IRC-2-14-E Compute 430 IRC-2-15-E Compute 429 IRC-2-16-E Compute 428 IRC-2-17-E Compute 427 IRC-2-18-E Compute 426 IRC-2-19-E Compute 425 IRC-2-20-E PDU-2-03-E PDU-2-04-E PDU-2-05-E PDU-2-06-E Compute 424 IRC-2-21-E Compute 423
733'
465 Compute 464 Compute IRC-3-01-E 463 Compute IRC-3-02-E PDU-3-01-E PDU-3-02-E IRC-3-03-E 462 Compute IRC-3-04-E 461 Compute IRC-3-05-E 460 Compute IRC-3-06-E 459 Compute IRC-3-07-E 458 Compute IRC-3-08-E 457 Compute Switch Core 2 456 Compute IRC-3-09-E 455 Compute Switch Core 1 454 Compute IRC-3-10-E 453 Compute IRC-3-11-E 452 Compute IRC-3-12-E 451 Compute IRC-3-13-E 450 Compute IRC-3-14-E 449 Compute IRC-3-15-E 448 Compute IRC-3-16-E 447 Compute PDU-3-03-E PDU-3-04-E 446 Compute IRC-3-17-E 445 Compute
841'
Compute 486 Compute 485 IRC-4-01-E Compute 484 IRC-4-02-E PDU-4-01-E PDU-4-02-E IRC-4-03-E Compute 483 IRC-4-04-E Compute 482 IRC-4-05-E Compute 481 IRC-4-06-E Compute 480 IRC-4-07-E Compute 479 IRC-4-08-E Compute 478 Switch Core 4 Compute 477 IRC-4-09-E Compute 476 Switch Core 3 Compute 475 IRC-4-10-E Compute 474 IRC-4-11-E Compute 473 IRC-4-12-E Compute 472 IRC-4-13-E Compute 471 IRC-4-14-E Compute 470 IRC-4-15-E Compute 469 IRC-4-16-E Compute 468 PDU-4-03-E PDU-4-04-E Compute 467 IRC-4-17-E Compute 466
949'
507 Compute 506 Compute IRC-5-01-E 505 Compute IRC-5-02-E PDU-5-01-E PDU-5-02-E IRC-5-03-E 504 Compute IRC-5-04-E 503 Compute IRC-5-05-E 502 Compute IRC-5-06-E 501 Compute IRC-5-07-E 500 Compute IRC-5-08-E 499 Compute Switch Core 6 498 Compute IRC-5-09-E 497 Compute Switch Core 5 496 Compute IRC-5-10-E 495 Compute IRC-5-11-E 494 Compute IRC-5-12-E 493 Compute IRC-5-13-E 492 Compute IRC-5-14-E 491 Compute IRC-5-15-E 490 Compute IRC-5-16-E 489 Compute PDU-5-03-E PDU-5-04-E 488 Compute IRC-5-17-E 487 Compute
1057'
Compute 516 Compute 515 IRC-6-01-E Compute 514 IRC-6-02-E PDU-6-01-E PDU-6-02-E IRC-6-03-E Compute 513 IRC-6-04-E Compute 512 IRC-6-05-E Compute 511 IRC-6-06-E Compute 510 IRC-6-07-E Compute 509 IRC-6-08-E Compute 508 Switch Core 8 OSS IO10 IRC-6-09-E OSS IO9 Switch Core 7 OSS IO8 IRC-6-10-E OSS IO7 IRC-6-11-E OSS IO6 IRC-6-12-E OSS IO5 IRC-6-13-E OSS IO4 IRC-6-14-E OSS IO3 IRC-6-15-E OSS IO2 IRC-6-16-E OSS IO1 PDU-6-03-E PDU-6-04-E Control M2 IRC-6-17-E Control M1
1165'
538 Compute 537 Compute IRC-7-01-E 536 Compute IRC-7-02-E PDU-7-01-E PDU-7-02-E IRC-7-03-E 535 Compute IRC-7-04-E 534 Compute IRC-7-05-E 533 Compute IRC-7-06-E 532 Compute IRC-7-07-E 531 Compute IRC-7-08-E 530 Compute IRC-7-09-E 529 Compute IRC-7-10-E 528 Compute IRC-7-11-E 527 Compute IRC-7-12-E 526 Compute IRC-7-13-E 525 Compute IRC-7-14-E 524 Compute IRC-7-15-E 523 Compute IRC-7-16-E 522 Compute IRC-7-17-E 521 Compute IRC-7-18-E 520 Compute IRC-7-19-E 519 Compute IRC-7-20-E PDU-7-03-E PDU-7-04-E PDU-7-05-E PDU-7-06-E 518 Compute 517 Compute IRC-7-21-E BM Compute
1273'
Compute 560 Compute 559 IRC-8-01-E Compute 558 IRC-8-02-E PDU-8-01-E PDU-8-02-E IRC-8-03-E Compute 557 IRC-8-04-E Compute 556 IRC-8-05-E Compute 555 IRC-8-06-E Compute 554 IRC-8-07-E Compute 553 IRC-8-08-E Compute 552 IRC-8-09-E Compute 551 IRC-8-10-E Compute 550 IRC-8-11-E Compute 549 IRC-8-12-E Compute 548 IRC-8-13-E Compute 547 IRC-8-14-E Compute 546 IRC-8-15-E Compute 545 IRC-8-16-E Compute 544 IRC-8-17-E Compute 543 IRC-8-18-E Compute 542 IRC-8-19-E Compute 541 IRC-8-20-E PDU-8-03-E PDU-8-04-E PDU-8-05-E PDU-8-06-E Compute 540 Compute 539 IRC-8-21-E Compute Tes

Acknowledgements

- Doug James for design, consultation and direction
- Jay Alameda for providing travel and userid/sign-on through XSEDE
  - NSF SI2 grant, NSF OCI-1047956 "SI2-SSI: A Productive and Accessible Development Workbench for HPC Applications Using the Eclipse Parallel Tools Platform"
  - This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575.
  - XSEDE Allocation TG-ASC110002, "A Productive and Accessible Development Workbench for HPC Applications Using the Eclipse Parallel Tools Platform"
- Greg Watson for Eclipse PTP consulting