

Configuring the JAXB Resource Manager: A mini-tutorial

The following slides provide a brief introduction to the XML schema defining the JAXB Configurable Resource Manager, with a particular focus on the “control” component.

We first describe briefly the structure of the schema, along with how the resource manager “environment” is wired in the XML document.

We then offer a small step-by-step example in which we add new functionality to an existing resource manager; this will serve to illustrate a representative cross-section of the configurable features in the definition schema.

The Resource Manager Schema (XSD)

The Resource Manager XML Schema

(*resource_manager_type.xsd*) is comprised of three principal sections:

<site-data>: used to provide fixed or default URLs for the control and monitor connections (optional).

<control-data>: pertains to the “control” component which furnishes the job -control actions (submission, cancellation, status update, etc.); usually the part most often reconfigured or customized according to system, scheduler or the special needs of an application community.

<monitor-data>: settings necessary for configuring the *LLview* monitoring client (adjustments will be made less often to this part of the definition).

We will concentrate here on the second section of the schema.

The (Job) Control Component

A definition instance supports
only one of these two pairs

control-type	
e property	[0..*] property-type
e attribute	[0..*] attribute-type
e managed-files	[0..*] managed-files-type
e script	[0..1] script-type
e start-up-command	[0..*] command-type
e submit-interactive	command-type
e submit-interactive-debug	[0..1] command-type
e submit-batch	command-type
e submit-batch-debug	[0..1] command-type
e get-job-status	[0..1] command-type
e terminate-job	[0..1] command-type
e suspend-job	[0..1] command-type
e resume-job	[0..1] command-type
e hold-job	[0..1] command-type
e release-job	[0..1] command-type
e shut-down-command	[0..*] command-type
e button-action	[0..*] command-type
e launch-tab	[0..1] launch-tab-type

properties = arbitrary variables internal to the client

attributes = “external” properties (such as those defined by the scheduler)

managed-files = local files, either pre-existent or written out from the definition itself, meant to be staged in conjunction with a job submission

script = for scheduler (batch) systems, the “batch script” (if any) to be staged in conjunction with a job submission

start-up-command(s) = arbitrary (remote) commands run when the resource manager is started

shut-down-command(s) = arbitrary (remote) commands run when the resource manager is stopped

submit-, terminate-, suspend-, resume-, hold-, release- = commands for controlling job submission

get-job-status = on-demand check of status of job

button-action = special command run via a Launch-Tab push-button

launch-tab = section describing the parts and disposition of the UI wizard used to configure and launch a job

A detailed guide to the XSD will be given in the Eclipse PTP developer documentation included as Help pages in the 5.0.1 release and also at <http://wiki.eclipse.org/PTP>.

The Resource Manager “Environment”

The *properties* and *attributes* defined in the XML, along with several preset properties provided internally, constitute the “environment” in which the JAXB Resource Manager runs.

When configuring the Resource Manager, this environment can be referenced in one of two ways.

- Some XML elements have attributes which take the name of the resource manager *attribute* or *property*.

Ex. 1: parser adds entry to the *value* field (a List) of the *queues* property:

```
<target ref="queues">  
    ... <add field="value"> ...
```

Ex. 2: combo sets the *value* field of *destination* to the selected item:

```
<widget type="combo" style="SWT.BORDER" readOnly="true"  
       saveValueTo="destination">
```

- A string value for the *property*’s or *attribute*’s fields can be obtained using the Eclipse variable resolution syntax. The namespace for the JAXB Resource Manager resolver is *ptp_rm*. The part preceding ‘#’ indicates the name of the *property* or *attribute*, that following, the field:

Ex.: tooltip on widget references that field of *destination* attribute:

```
<tooltip>${ptp_rm:destination#tooltip}</tooltip>
```

“Wiring” the Resource Manager Definition

A particularly powerful aspect of the Resource Manager’s configurability derives from the ability, on the basis of this “environment”, to make one part of the XML definition refer to another. We will call this “wiring” the definition.

On the next two slides, we present and explain an example of this procedure, based on the setting and reading of variables related to the choice of the scheduler queue.

On the slide following this discussion, we will then note two special procedures necessary for capturing certain values present only after you select “Run” and launch the job.

Wiring the XML: Example (queue)



“Wiring” the Resource Manager Definition

1. There are two main “variables” for handling the queue name: *destination*, which points to the actual queue chosen, and *queues*, which provides a list of available queues.
2. The *control.queue.name* is an internal property used under the covers to convey information to the monitor. The <link-value-to> element means that this property’s value is set to that of *destination*, or in the case that *destination*’s value is undefined, to any default value given to *control.queue.name* (here none).
3. The script element has a line whose only argument references the value of *destination*; if this is empty, the argument will resolve to “#PBS –q”, and since it is indicated that this should be considered equivalent to undefined, the argument will be eliminated.
4. The *tooltip* string is given to the label widget pointing to the combo list where the queue can be selected.
5. The combo list widget itself takes its preset values (notice that it is “readOnly”, so the user cannot type in a value here but is constrained to the provided choices) from the list of available *queues*.
6. The selection made via the combo widget will become the value of the *destination* attribute.
7. The start-up command runs “qstat –Q –f”, then uses a regular expression to parse the standard output and to accumulate the matching values as entries in a List which it sets as the value of the *queues* property.

“Wiring” the Resource Manager Definition

SPECIAL NOTES: **visible**, **@jobId**, <managed-file> target paths

The **visible** attribute on a *property* or *attribute* is a way of indicating to the Resource Manager that the user will not be directly changing its value via the Launch Tab interface. Certain widgets (such as the attribute table or tree) check this to see if the *property* or *attribute* should be included automatically in its list.

@jobId: This is a special property name designating the runtime id for a job instance. In the lifecycle of the run/launch call, this value begins as an internally generated unique id which then is swapped for the id returned by the scheduler.

The @jobId, along with the target paths for <managed-file> elements, are not known at configuration time (i.e., before the user hits “Run”). While the former is made visible to the parsers and the returned status object of the submit command, neither is available for reference in other managed files or in the <script> element, because these elements are generated prior to the actual submission.

If the <script> needs to refer to the **@jobId**, it must do so via the variable made available by the particular scheduler it is written for. An example of *how to reference the target paths of other <managed-file>s inside the script* is included in the illustration which follows.

Customizing the Resource Manager Definition: An Illustration

In the following illustration, we will take a pre-existing Resource Manager definition and modify it by adding functionality to support a particular application scenario.

Let us suppose we wish to tailor the definition for use with a simulation code which requires an input file; we would like the user to be able to choose this file before launching the job.

Let us further say that the file could be chosen either from a predetermined location on the remote host, or could be selected from a file edited locally.

In the former case, we would simply need to select a (remote) path, whereas in the latter, we would actually need to stage the file over before executing the call to submit the job.

Customizing the Resource Manager Definition: An Illustration

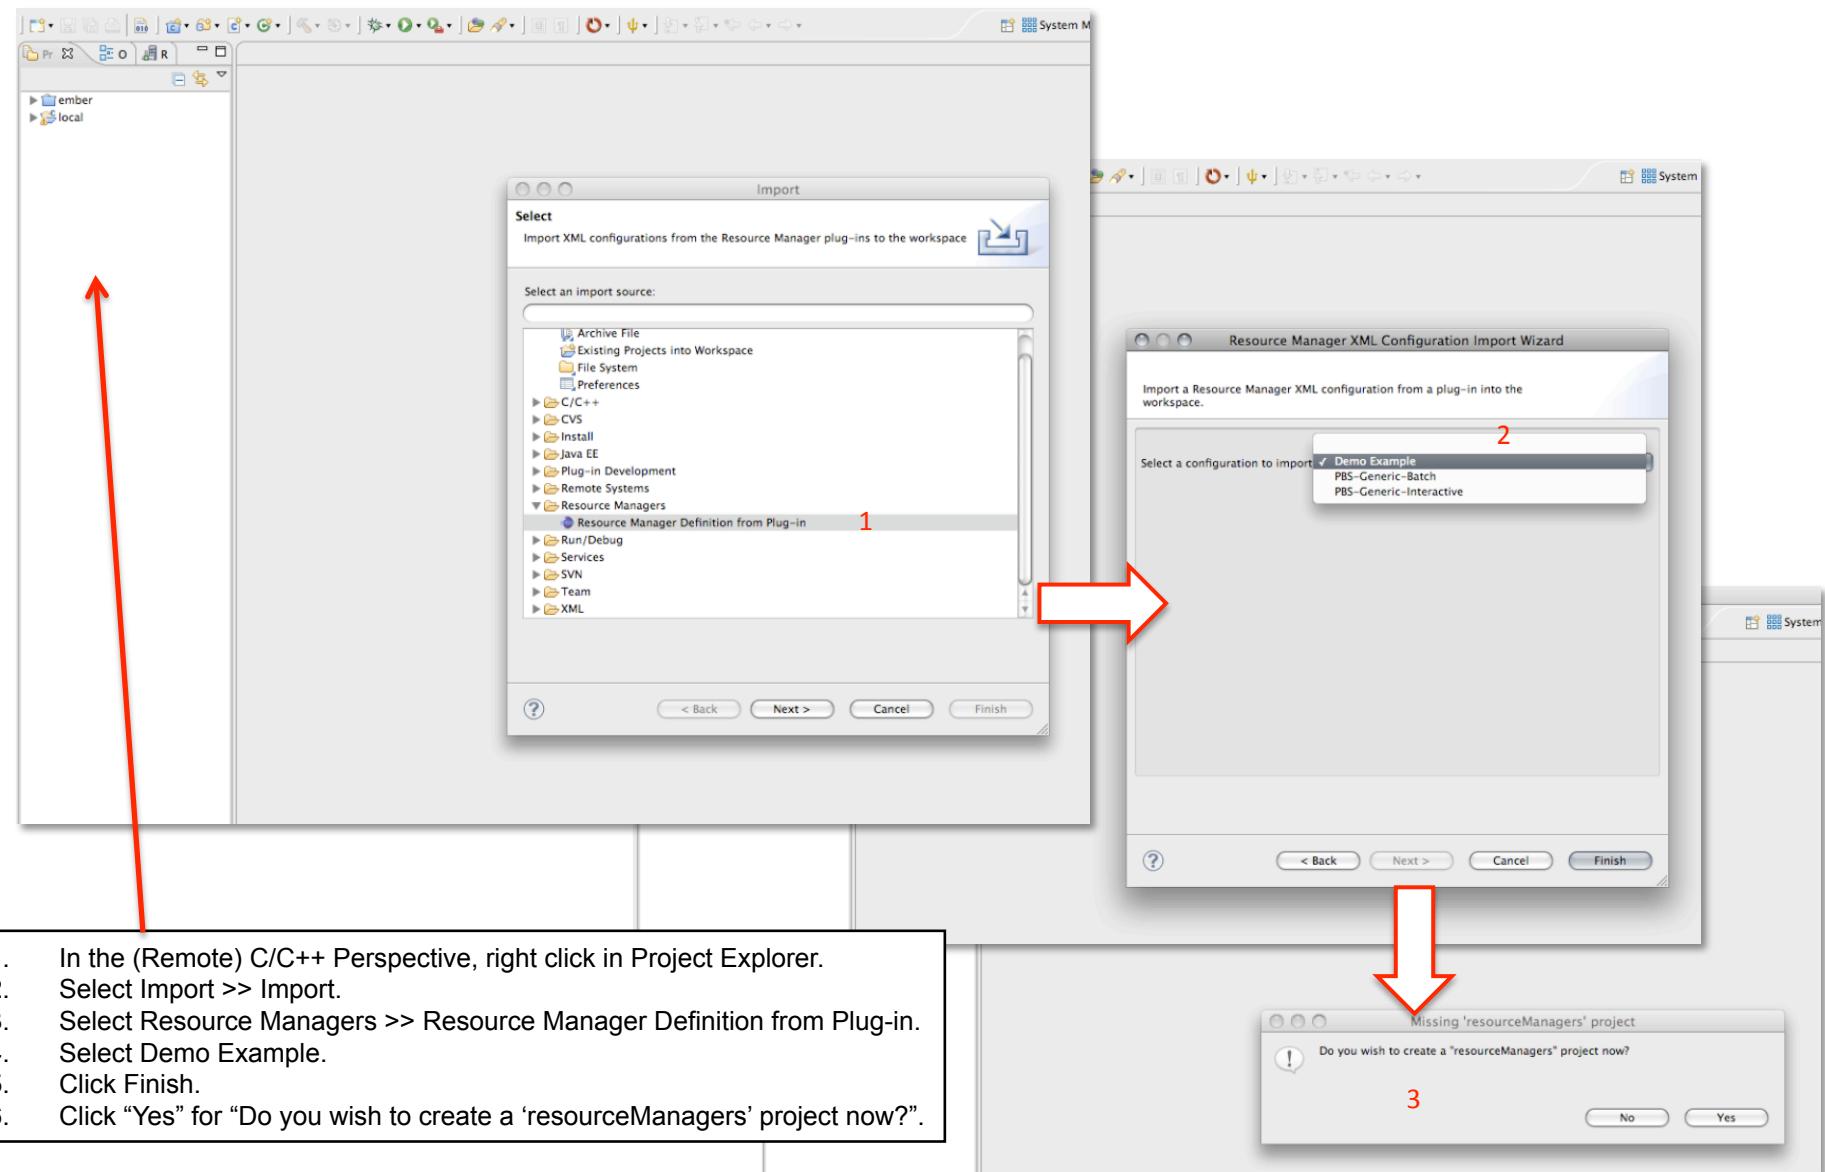
What we will add to the existing definition:

1. Properties for handling the path(s).
2. Two widgets:
 - a. A combo list populated from the contents of a remote directory;
 - b. A browse text + button for selecting a local file.
3. A managed file for staging in case of 2b being selected.
4. A start-up command and parser which will populate the items of 2a.
5. The necessary context for conveying the path to the batch script.
6. The additional reference to the path on the execution line of the batch script.

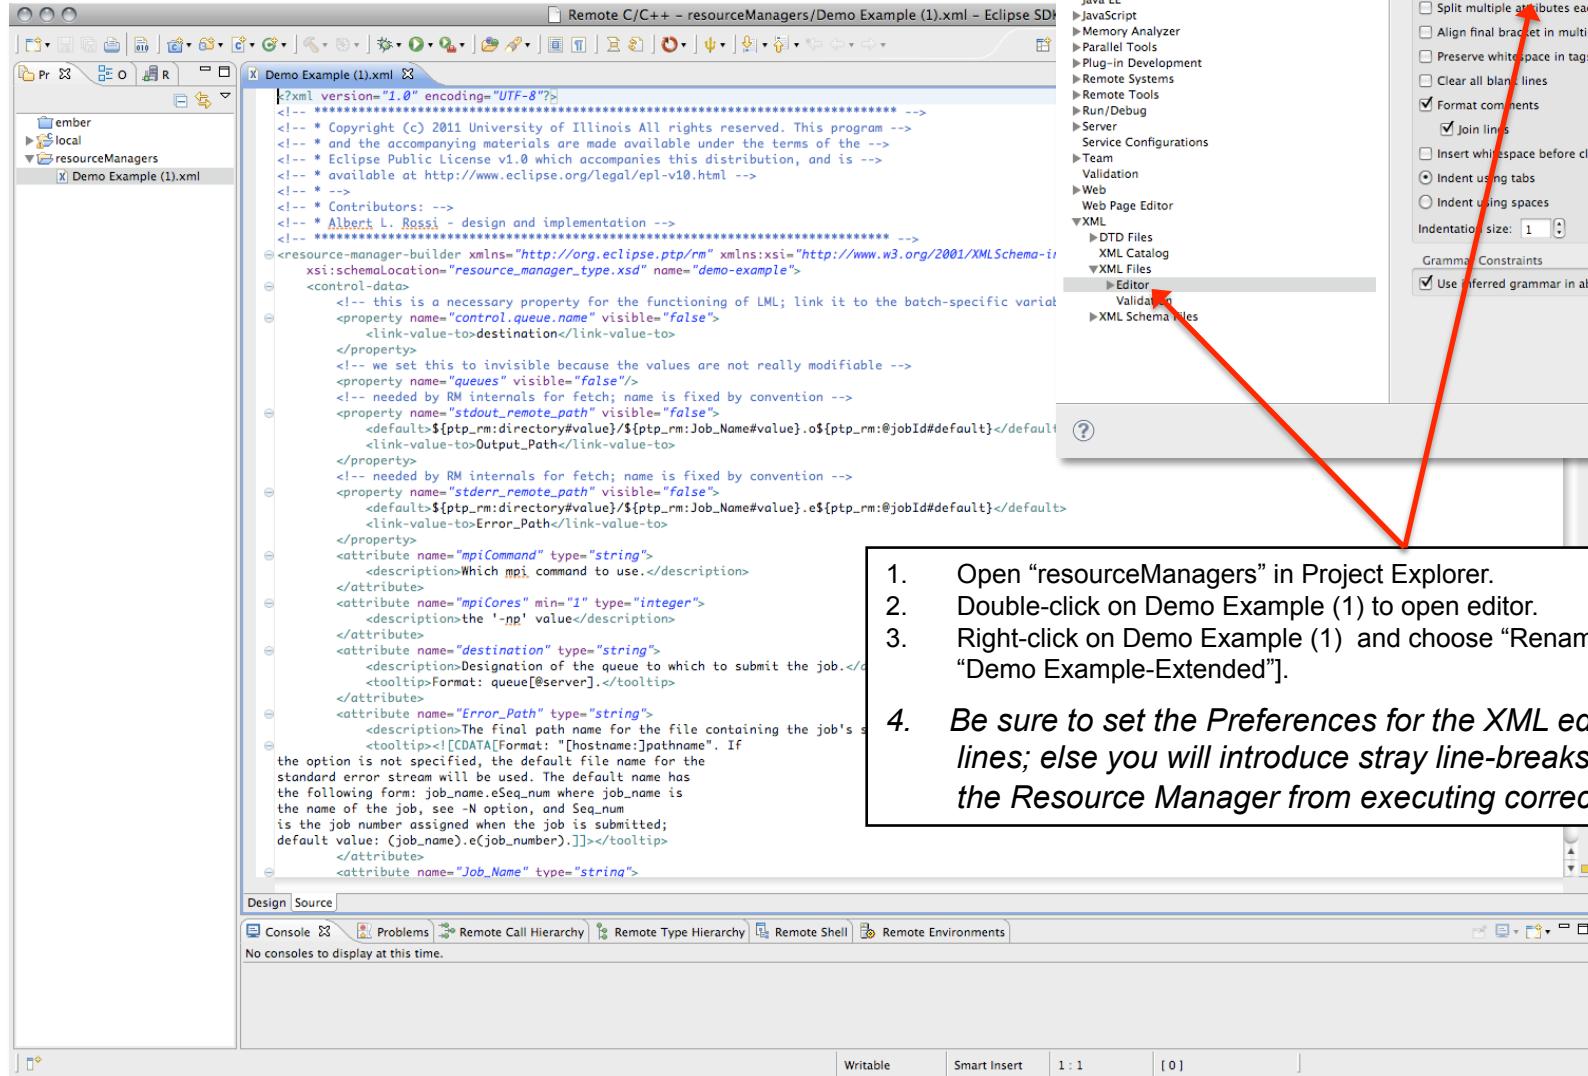
First, we will need to import the provided definition file into our workspace for editing.

(Note: “Demo Example” is not provided in the Indigo release; you will need to select one of the PBS-generic configurations for modification.)

Import the XML into your Project Workspace



Open/Rename the XML



1. Open “resourceManagers” in Project Explorer.
2. Double-click on Demo Example (1) to open editor.
3. Right-click on Demo Example (1) and choose “Rename” [here we use “Demo Example-Extended”].
4. **Be sure to set the Preferences for the XML editor to format long lines; else you will introduce stray line-breaks which will stop the Resource Manager from executing correctly.**

Add the Resource Manager

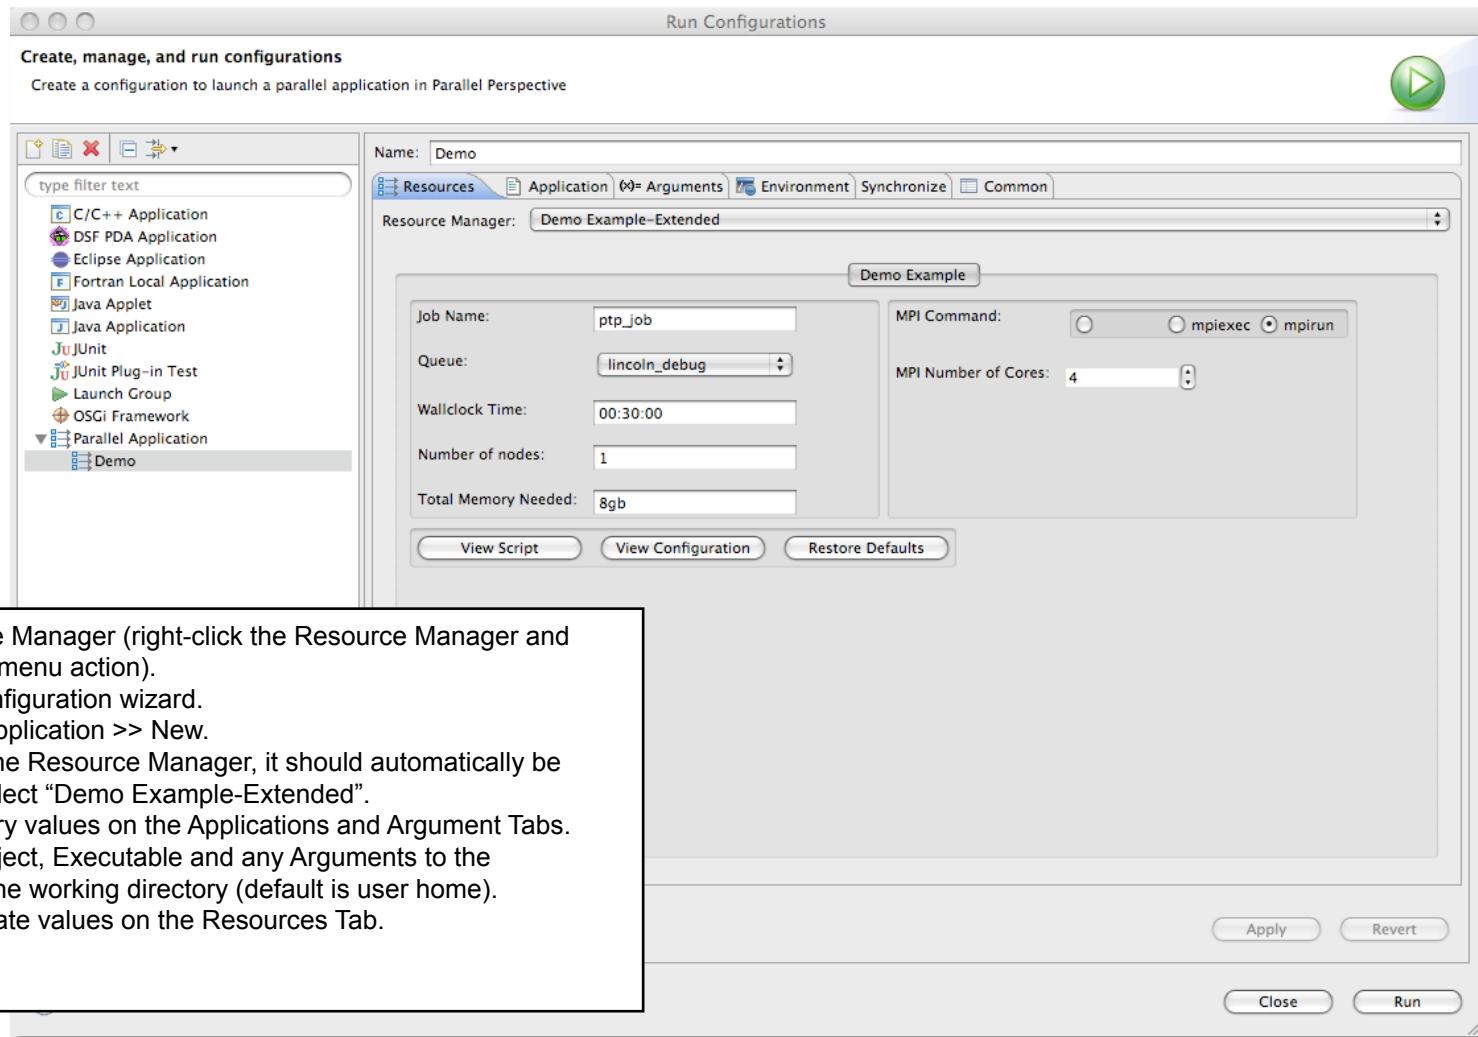
1. Go to System Monitoring perspective.
2. Right click in Resource Managers.
3. Select "Add Resource Manager" from context menu.
4. In dialog "Choose Resource Manger Type", select Demo Example-Extended.

Choose or create a remote connection

1. From "Preferences" Menu, select Parallel Tools >> Resource Managers.
2. Select "Configurable Resource Manager".
3. Change default setting (unchecked) for "Always reload"; this allows you to see changes made to the XML each time you restart the Resource Manager (otherwise, the file is cached and reused for the life of the ResourceManager).
4. Click "Apply", then "OK".

Launch Tab Set-Up

Proceed with the usual Launch Tab configuration.



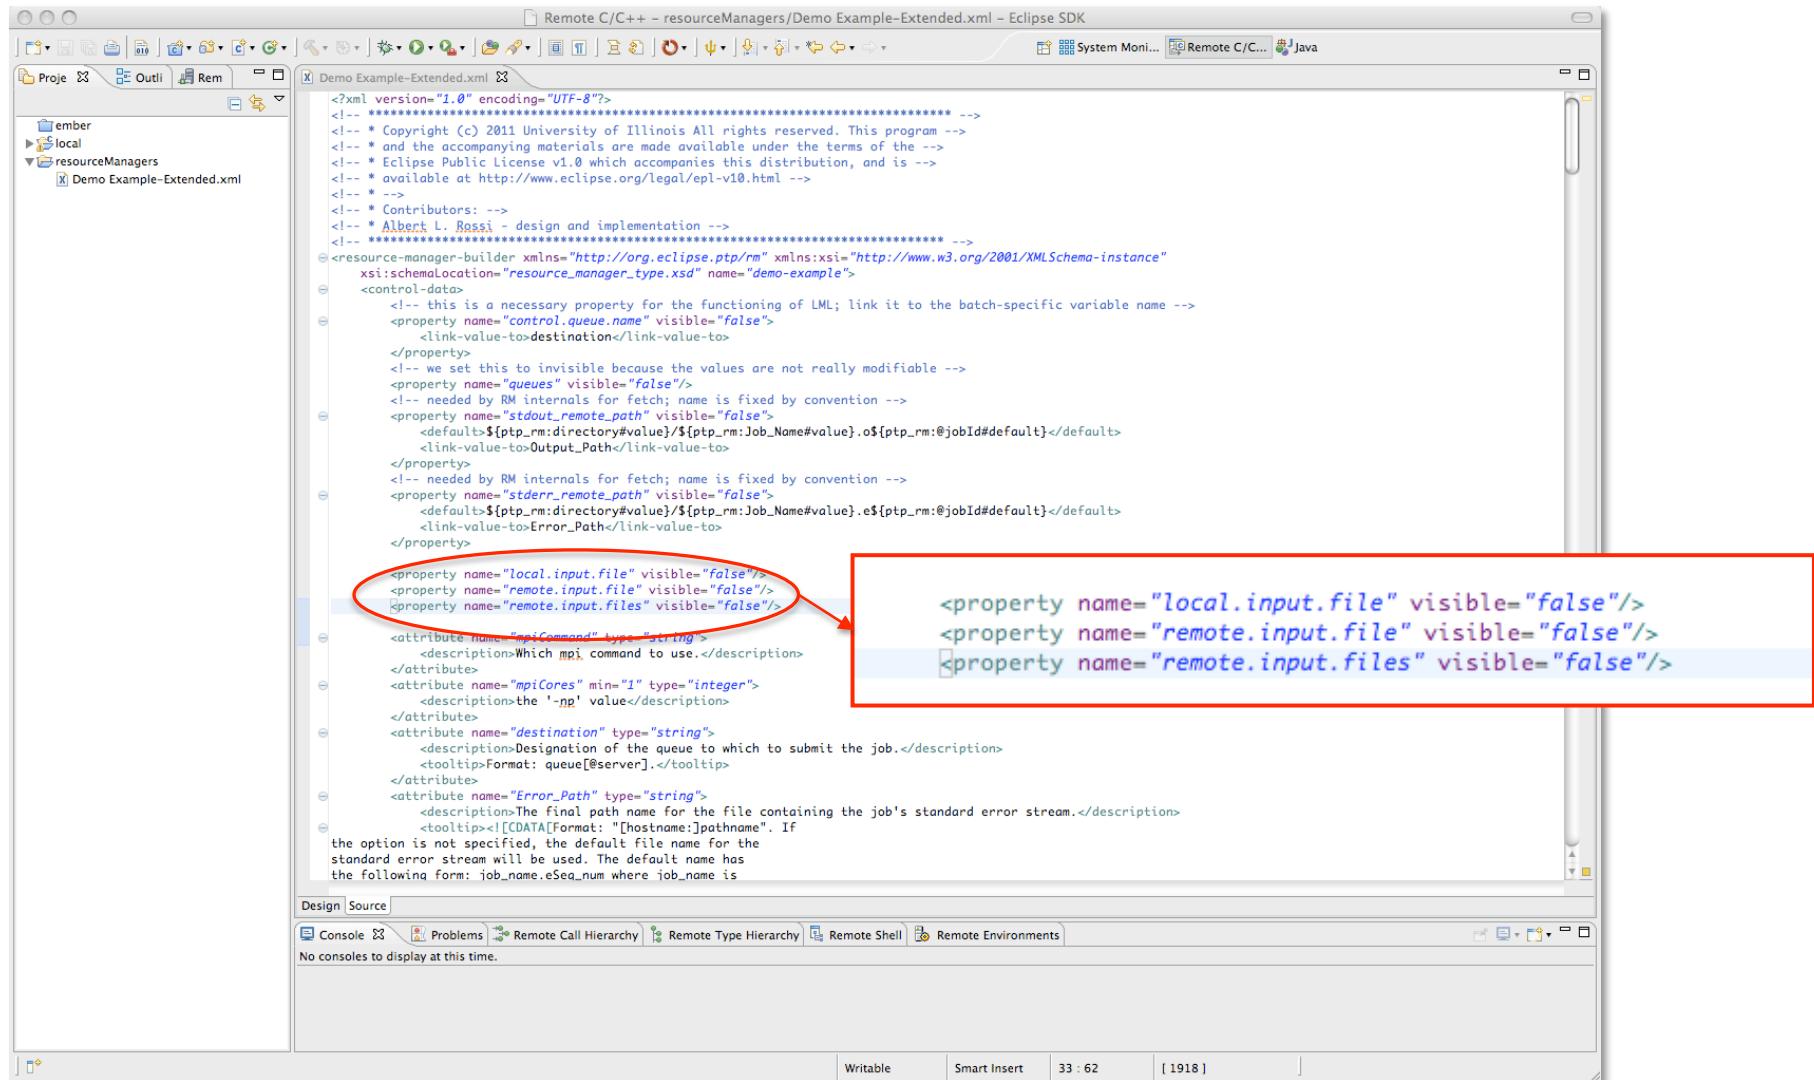
1. Start the Resource Manager (right-click the Resource Manager and select the context menu action).
2. Open the Run Configuration wizard.
3. Choose Parallel Application >> New.
4. If you have only one Resource Manager, it should automatically be selected; if not, select "Demo Example-Extended".
5. Fill in the necessary values on the Applications and Argument Tabs. These include Project, Executable and any Arguments to the Executable, plus the working directory (default is user home).
6. Fill in the appropriate values on the Resources Tab.
7. Click Apply.

You now have the basic tab which we will proceed to modify.

This Resource Manager provides basic PBS batch settings, and on start-up looks for the available queues. The actions enabled are for job submission, getting job status and job cancellation.

1. Properties

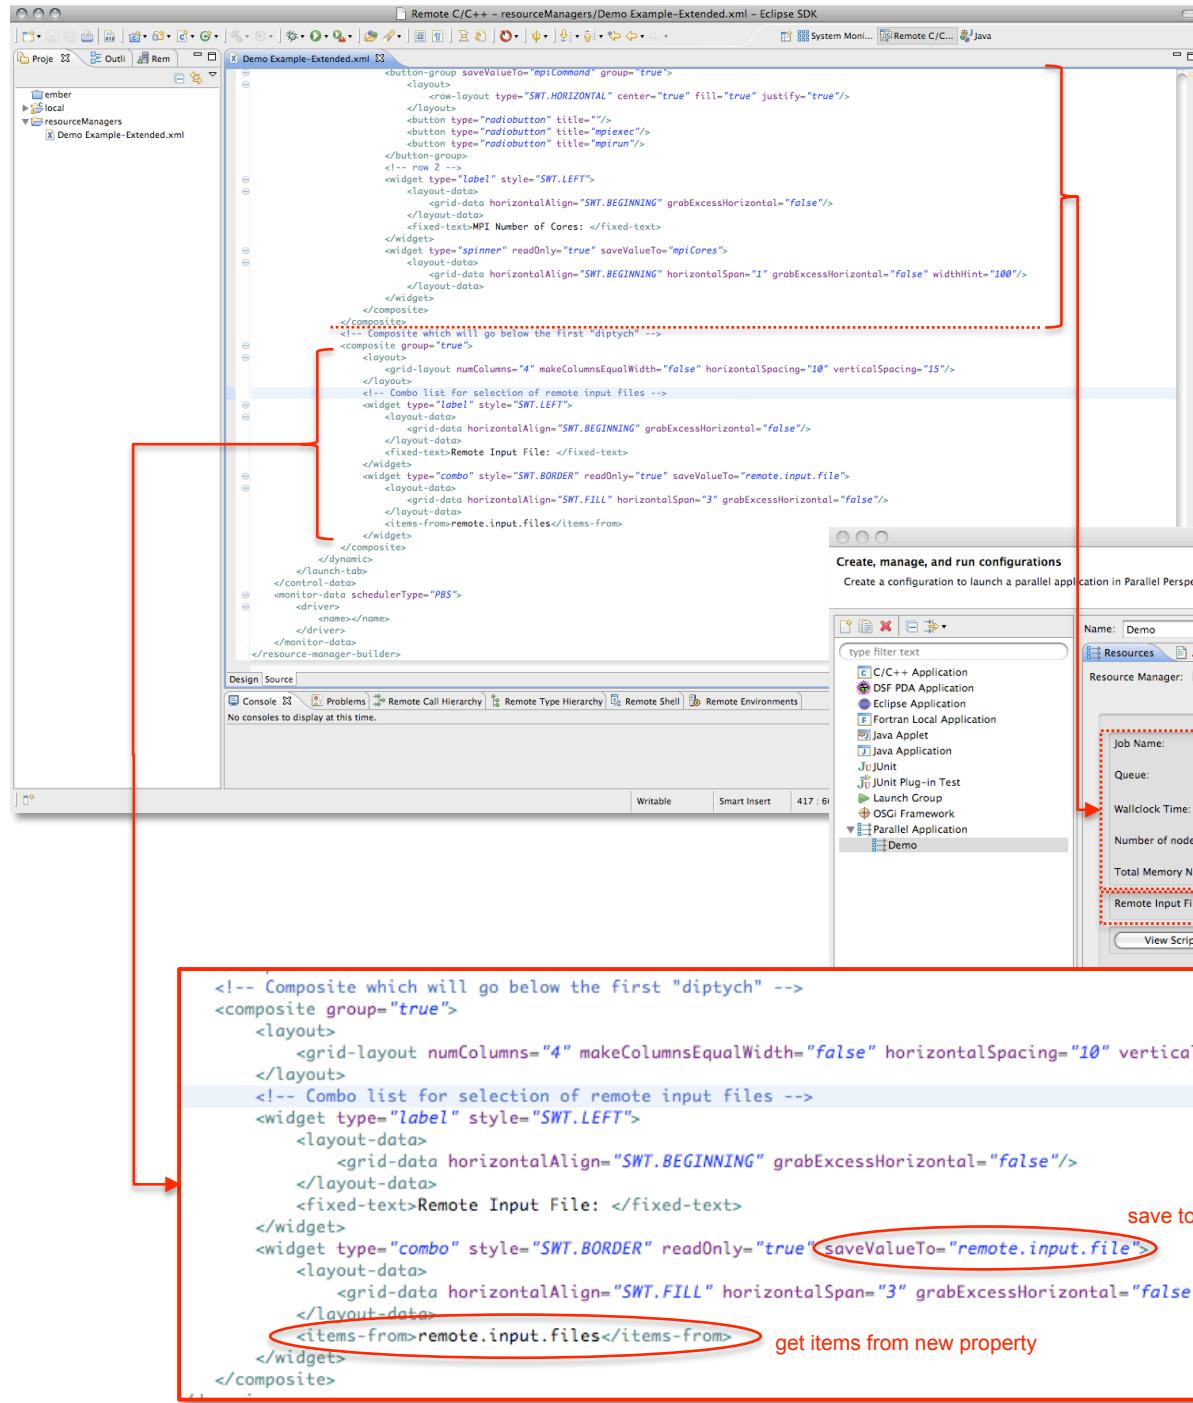
Add three properties, two to reference the path choice, and a third to hold a list of remote paths.



The screenshot shows the Eclipse IDE interface with the XML editor open. The XML file is named "Demo Example-Extended.xml". The code is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- **** Copyright (c) 2011 University of Illinois All rights reserved. This program -->
<!-- and the accompanying materials are made available under the terms of the -->
<!-- Eclipse Public License v1.0 which accompanies this distribution, and is -->
<!-- available at http://www.eclipse.org/legal/epl-v10.html -->
<!-- * -->
<!-- * Contributors: -->
<!-- * Albert L. Rossi - design and implementation -->
<!-- **** -->
<resource-manager-builder xmlns="http://org.eclipse.ptp/rm" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="resource_manager_type.xsd" name="demo-example">
    <control-data>
        <!-- this is a necessary property for the functioning of LML; link it to the batch-specific variable name -->
        <property name="control.queue.name" visible="false">
            <link-value-to>destination</link-value-to>
        </property>
        <!-- we set this to invisible because the values are not really modifiable -->
        <property name="queues" visible="false"/>
        <!-- needed by RM internals for fetch; name is fixed by convention -->
        <property name="stdout_remote_path" visible="false">
            <default>${ptp_rm:directory#value}/${ptp_rm:Job_Name#value}.o${ptp_rm:@jobId#default}</default>
            <link-value-to>Output_Path</link-value-to>
        </property>
        <!-- needed by RM internals for fetch; name is fixed by convention -->
        <property name="stderr_remote_path" visible="false">
            <default>${ptp_rm:directory#value}/${ptp_rm:Job_Name#value}.e${ptp_rm:@jobId#default}</default>
            <link-value-to>Error_Path</link-value-to>
        </property>
        <!-- local input file -->
        <property name="local.input.file" visible="false"/>
        <!-- remote input file -->
        <property name="remote.input.file" visible="false"/>
        <!-- remote input files -->
        <property name="remote.input.files" visible="false"/>
    </control-data>
    <attribute name="mpiCommand" type="string">
        <description>Which mpi command to use.</description>
    </attribute>
    <attribute name="mpiCores" min="1" type="integer">
        <description>The '-np' value</description>
    </attribute>
    <attribute name="destination" type="string">
        <description>Designation of the queue to which to submit the job.</description>
        <tooltip>Format: queue[@server]</tooltip>
    </attribute>
    <attribute name="Error_Path" type="string">
        <description>The final path name for the file containing the job's standard error stream.</description>
        <tooltip><![CDATA[Format: "[hostname:]pathname". If the option is not specified, the default file name for the standard error stream will be used. The default name has the following form: job_name.e$e_a_num where job_name is -->
```

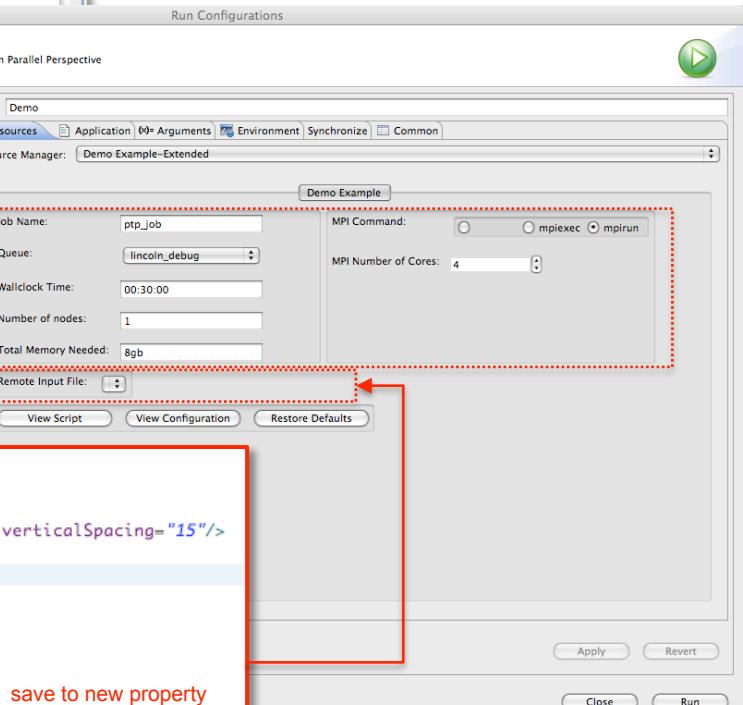
A red box highlights the last three lines of the code, and a red circle highlights the first three lines of the code.

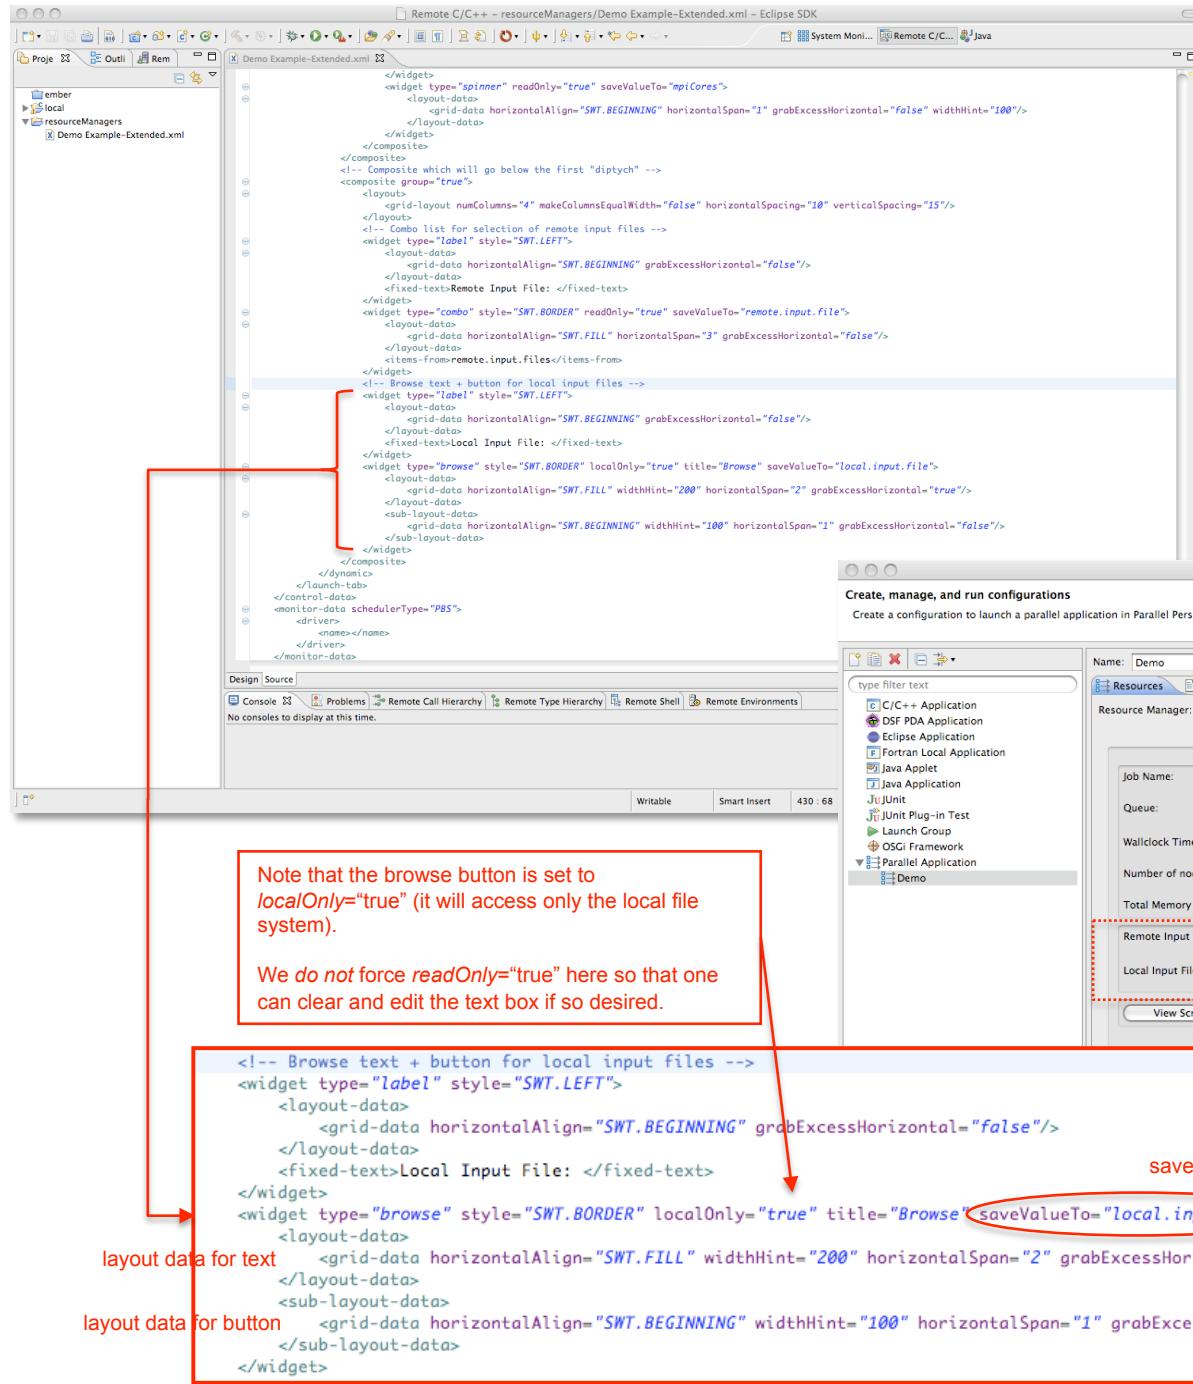


2a. Combo Box

Add group with 4 columns below main composite of two panels.

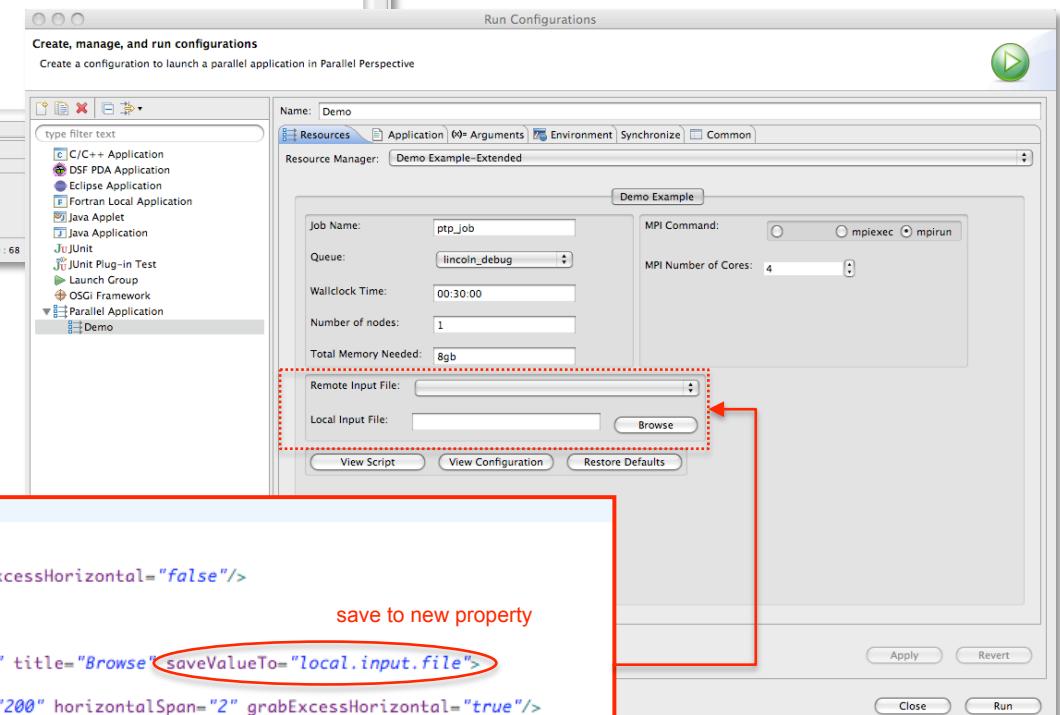
Add label + 3-column combo.





2b. Browse Text + Button

Add label + this paired widget (text gets 2 columns) after the combo box inside the new composite/group.



3. Managed File

The screenshot shows the Eclipse C/C++ IDE interface with the 'Remote C/C++ - resourceManagers/Demo Example-Extended.xml - Eclipse SDK' window active. The left sidebar shows a project structure with 'ember', 'local', and 'resourceManagers' containing 'Demo Example-Extended.xml'. The main editor area displays XML code. A red oval highlights the 'managed-files' section, which contains a 'file-staging-location' set to '.eclipsesettings' and a 'file name' set to 'menu.input'. A red arrow points from this section down to a red box in the 'Source' tab below. The 'Source' tab shows the same XML code with the 'file-staging-location' and 'file name' entries highlighted. A red circle surrounds the 'path' element under the 'file' node, and a red box encloses the entire 'managed-files' section. A tooltip 'references new property' is visible near the circled path element.

```
<default>00:30:00</default>
<validator>
    <regex expression="\d\d:\d\d:\d\d" />
    <error-message>format must be hh:mm:ss</error-message>
</validator>
</attribute>

<managed-files>
    <file-staging-location>.eclipsesettings</file-staging-location>
    <file name="menu.input">
        <path>${ptp_rm:local.input.file#value}</path>
    </file>
</managed-files>

<script>
    <line>
        <arg>#!/bin/bash</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -q">#PBS -q ${ptp_rm:destination#value}</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -e">#PBS -e ${ptp_rm:Error_Path#value}</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -N">#PBS -N ${ptp_rm:Job_Name#value}</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -o">#PBS -o ${ptp_rm:Output_Path#value}</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -l mem">#PBS -l
            mem=${ptp_rm:Resource_List.mem#value}</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -l nodes">#PBS -l
            nodes=${ptp_rm:Resource_List.nodes#value}</arg>
    </line>
    <line>
        <arg isUndefinedIfMatches="#PBS -l walltime">#PBS -l
            walltime=${ptp_rm:Resource_List.walltime#value}</arg>
    </line>
    <line>
        <arg>#PBS -V</arg>
    </line>
    <line>
        <arg>MPI_ARGS="--no ${ptp_rm:mpiCores#value}"</arg>
    </line>
</script>
```

<managed-files>
 <file-staging-location>.eclipsesettings</file-staging-location>
 <file name="menu.input">
 <path>\${ptp_rm:local.input.file#value}</path> references new property
 </file>
</managed-files>

- Local path references the new property set from the browse text-widget field.
- Stages the file to ".eclipsesettings" in user home.
- Remote path will be `.eclipsesettings/local.input.file.name`. This is accessed in the environment as: `${ptp_rm:menu.input#value}`.

4. Start-up command

args: "ls -l [remote directory]"

regex parses output for the file name.

The screenshot shows the Eclipse IDE interface with two main windows:

- Left Window (Demo Example-Extended.xml):** Displays an XML configuration file. A red box highlights the `<start-up-command name="get-input">` section. A red arrow points from this section to the `Run Configuration` dialog on the right.
- Right Window (Run Configurations):** Shows a `Parallel Application` configuration named `Demo`. It includes fields for `Job Name` (`ptp_job`), `Queue` (`lincoln_debug`), `Wallclock Time` (`00:30:00`), `Number of nodes` (`1`), and `Total Memory Needed` (`8gb`). The `MPI Command` field has a radio button selected for `mpirun`. A red box highlights the `Remote Input File` dropdown, which lists several files: `menu.input_0`, `menu.input_1`, `menu.input_2`, `menu.input_3`, and `menu.input_4`. A red arrow points from this dropdown to the XML configuration in the left window.

Bottom Left Window (resource-manager-builder/control-data/start-up-command/stdout-parser/target/match/text): Shows the expanded XML code for the `get-input` command. Red annotations explain specific parts:

- `<arg>ls</arg>` and `<arg>-l</arg>` are noted as "input files are in 'input' subdir of the working directory (user home)."
- `<arg>$!{ptp_rm:directory#value}/input</arg>` is noted as "set value on new property".
- A red circle highlights the `<entry valueGroup="1" />` line.
- A dotted line connects this highlighted line to the `valueGroup="1"` entry in the `Remote Input File` dropdown in the Run Configuration dialog.

5. Adjust Environment of Submit Command

The screenshot shows the Eclipse IDE interface with a project named "resourceManagers" containing an XML file "Demo Example-Extended.xml". The code defines a submit batch command with environment settings and a script parser. A red box highlights the first two environment definitions, and a red circle highlights the value attribute of the second one. Another red circle highlights the managed file property reference in the second environment definition.

```
<submit-batch name="submit-batch" waitForId="true">
    <arg>qsub</arg>
    <arg>${ptp_rm:managed_file_for_script#value}</arg>

    <environment name="INPUT">
        <arg isUndefinedIfMatches="-f">-f ${ptp_rm:remote.input.file#value}</arg>
    </environment>
    <environment name="INPUT">
        <arg isUndefinedIfMatches="-f">-f ${ptp_rm:menu.input#value}</arg>
    </environment>

    <stdout-parser delim="\n" all="true" save="1">
        <target ref="@jobId">
            <match>
                <expression>([d]+)([.])[.]+[\s]+.*</expression>
                <append field="name">
                    <entry valueGroup="1" />
                    <entry valueGroup="2" />
                    <entry valueGroup="3" />
                </append>
                <set field="default">
                    <entry valueGroup="1" />
                </set>
                <set field="value">
                    <entry value="SUBMITTED" />
                </set>
            </match>
            <match>
                <expression>([d]+)[.]+</expression>
                <set field="name">
                    <entry valueGroup="0" />
                </set>
                <set field="default">
                    <entry valueGroup="1" />
                </set>
                <set field="value">
                    <entry value="SUBMITTED" />
                </set>
            </match>
        </target>
        <target ref="@jobId">
            <match>
                <expression flags="DOTALL">.*Job not submitted.*</expression>
                <set field="value">
                    <entry value="FAILED" />
                </set>
            </match>
        </target>
    </stdout-parser>
    <stderr-parser delim="\n">
        <target ref="@jobId">
            <match>
                <expression>.*Job not submitted.*</expression>
                <throw message="Job Submit Failed" />
            </match>
        </target>
    </stderr-parser>
</submit-batch>
```

Because the managed files are configured after “Run” is selected, but just before the actual remote submission of the job, the remote target path can be captured and placed in the job’s environment, making it available to the batch script when it becomes active.

Environment definitions follow the command args; the value of an environment variable can be expressed via its “value” attribute, or as embedded <arg>s, as it is here.

references new property

references managed file property (remote path); should overwrite first INPUT definition *only if defined*
(note: a bug in the managed file code, *fixed in release 5.0.1*, was setting this path to the staging directory when the actual file was undefined)

6. Adjust Script Execution Line

<script> is the XML representation of the lines of the batch script. Each line can have an arbitrary number of <arg> elements.

Setting *resolve* to false tells the Resource Manager not to interpret the argument (otherwise it would think the shell variable \${INPUT} was an Eclipse variable).

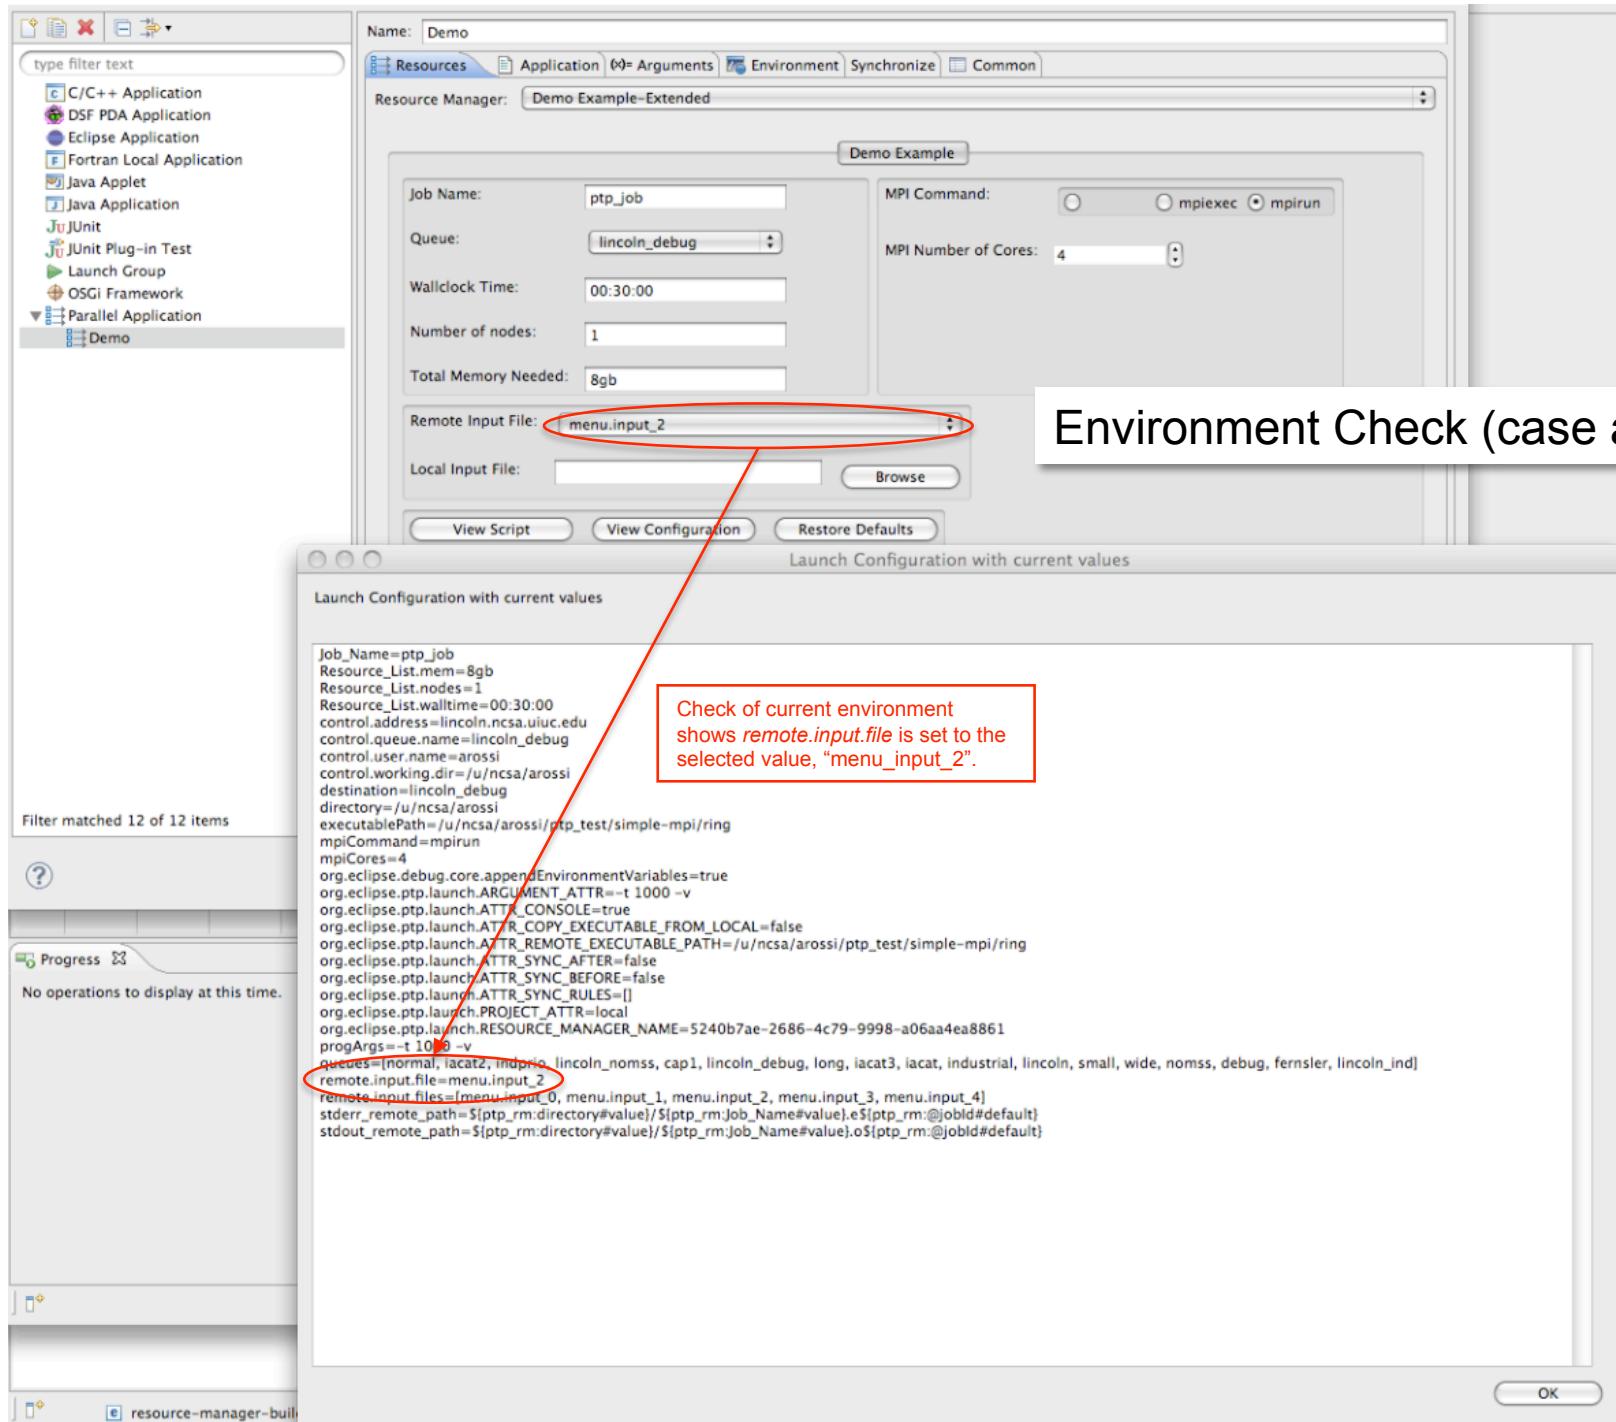
```
<arg resolve="false">if [ -n "${COMMAND}" ] ; then</arg>
<arg resolve="false"> COMMAND=${ptp_rm:mpiCommand#value}</arg>
<arg resolve="false">if [ -n "${COMMAND}" ] ; then</arg>
<arg resolve="false"> COMMAND="$COMMAND ${MPI_ARGS}"</arg>
<arg>${ptp_rm:executablePath#value} ${ptp_rm:progArgs#value}</arg>
<arg resolve="false">${INPUT}</arg>
<arg resolve="false">else</arg>
<arg> COMMAND="${ptp_rm:executablePath#value} ${ptp_rm:progArgs#value}"</arg>
<arg resolve="false">${INPUT}</arg>
<arg resolve="false">fi</arg>
<arg isUndefinedIfMatches="cd > cd > ${ptp_rm:directory#value}"></arg>
<arg resolve="false">${COMMAND}</arg>
</script>
<start-up-command name="get-queues">
<arg>stot</arg>
<arg>-Q</arg>
<arg>-f</arg>
<stdout-parser delim="\n">
<target ref="queues">
<match>
<expr>
<add f="true" />
<e>
</add>
</match>
</target>
</stdout-parser>
</start-up-commands>
```

Add \${INPUT} argument for the input file.

```
<lines>
<arg isUndefinedIfMatches="cd > cd > ${ptp_rm:directory#value}"></arg>
</line>
<line>
<arg>#PBS -V</arg>
</line>
<line>
```

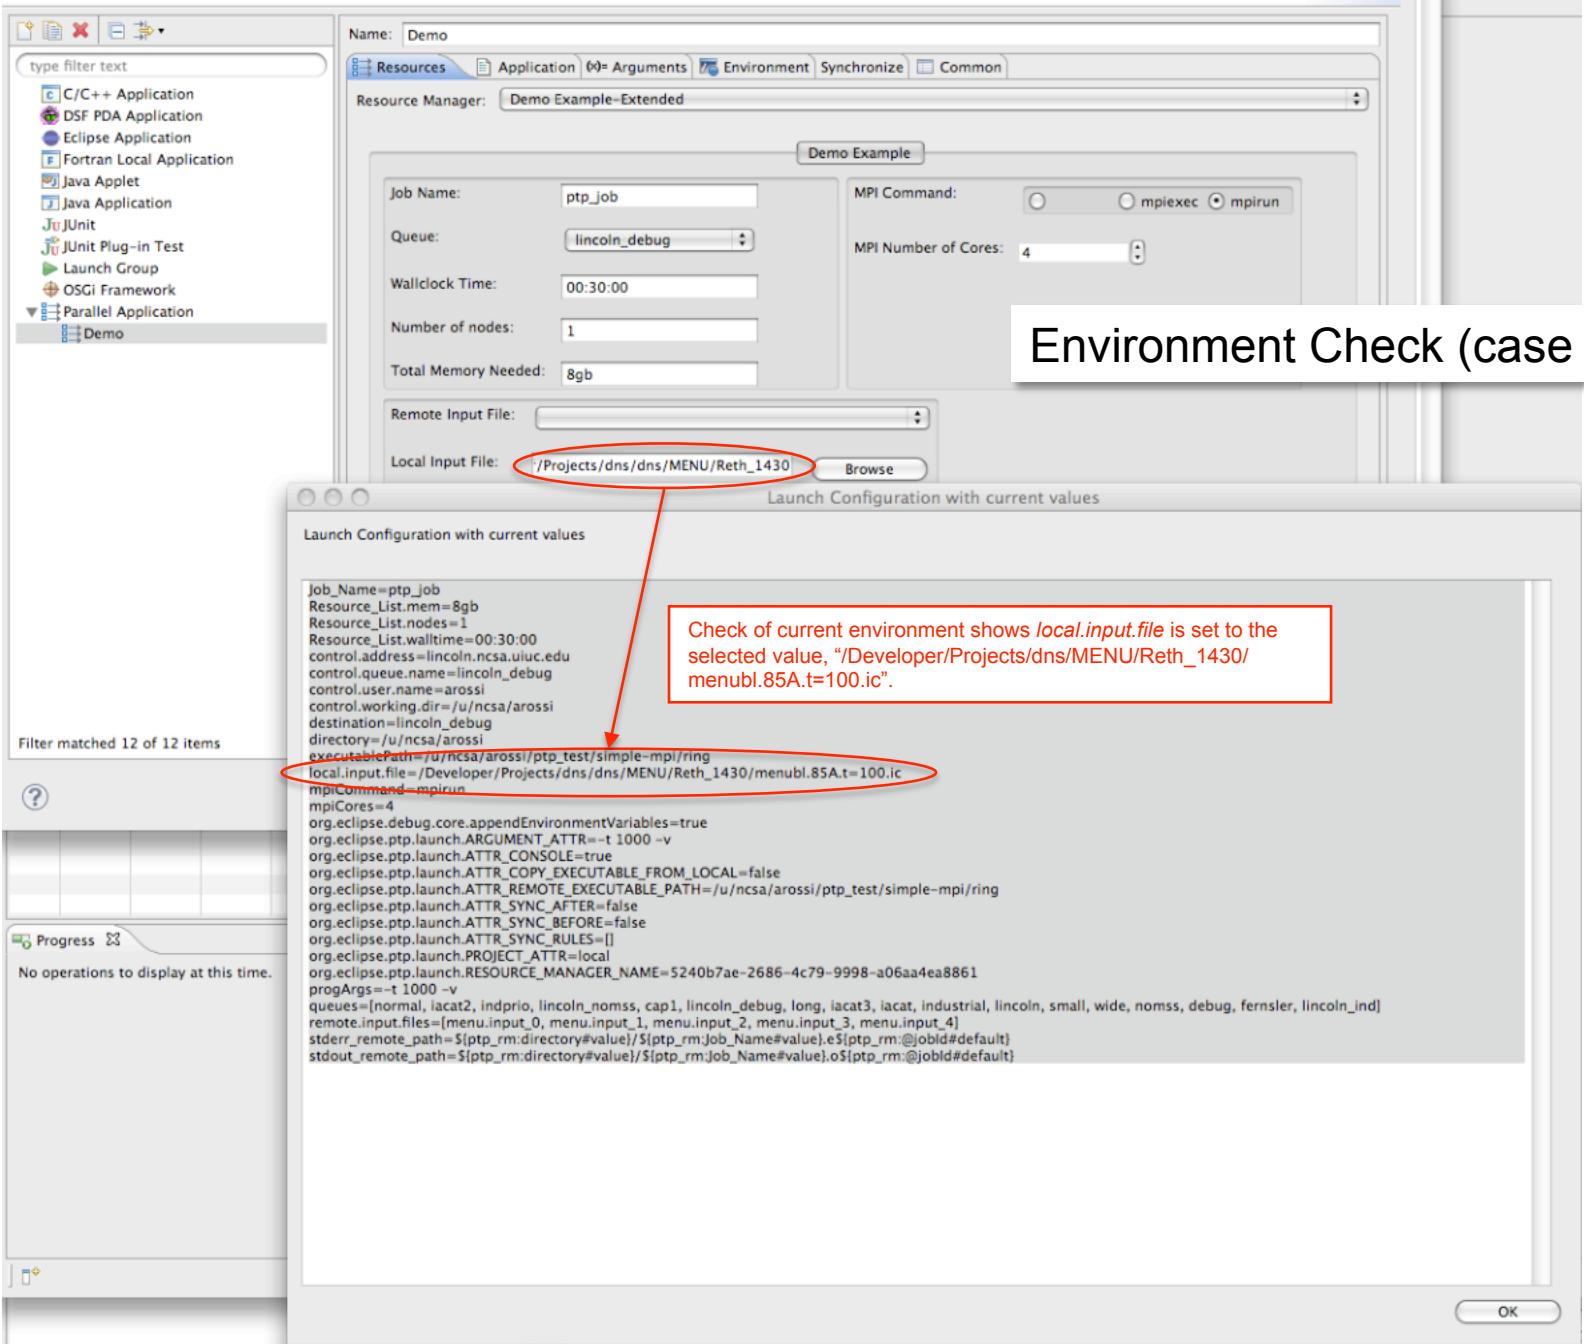
Be sure that your script indicates that the environment be passed on to the child processes.

```
<line>
<arg resolve="false">if [ -n "${COMMAND}" ] ; then</arg>
</line>
<line>
<arg resolve="false"> COMMAND="${COMMAND} ${MPI_ARGS}"</arg>
<arg>${ptp_rm:executablePath#value} ${ptp_rm:progArgs#value}</arg>
<arg resolve="false">${INPUT}</arg>
</line>
<line>
<arg resolve="false">else</arg>
</line>
<line>
<arg> COMMAND="${ptp_rm:executablePath#value} ${ptp_rm:progArgs#value}"</arg>
<arg resolve="false">${INPUT}</arg>
</line>
<line>
<arg resolve="false">fi</arg>
</line>
```



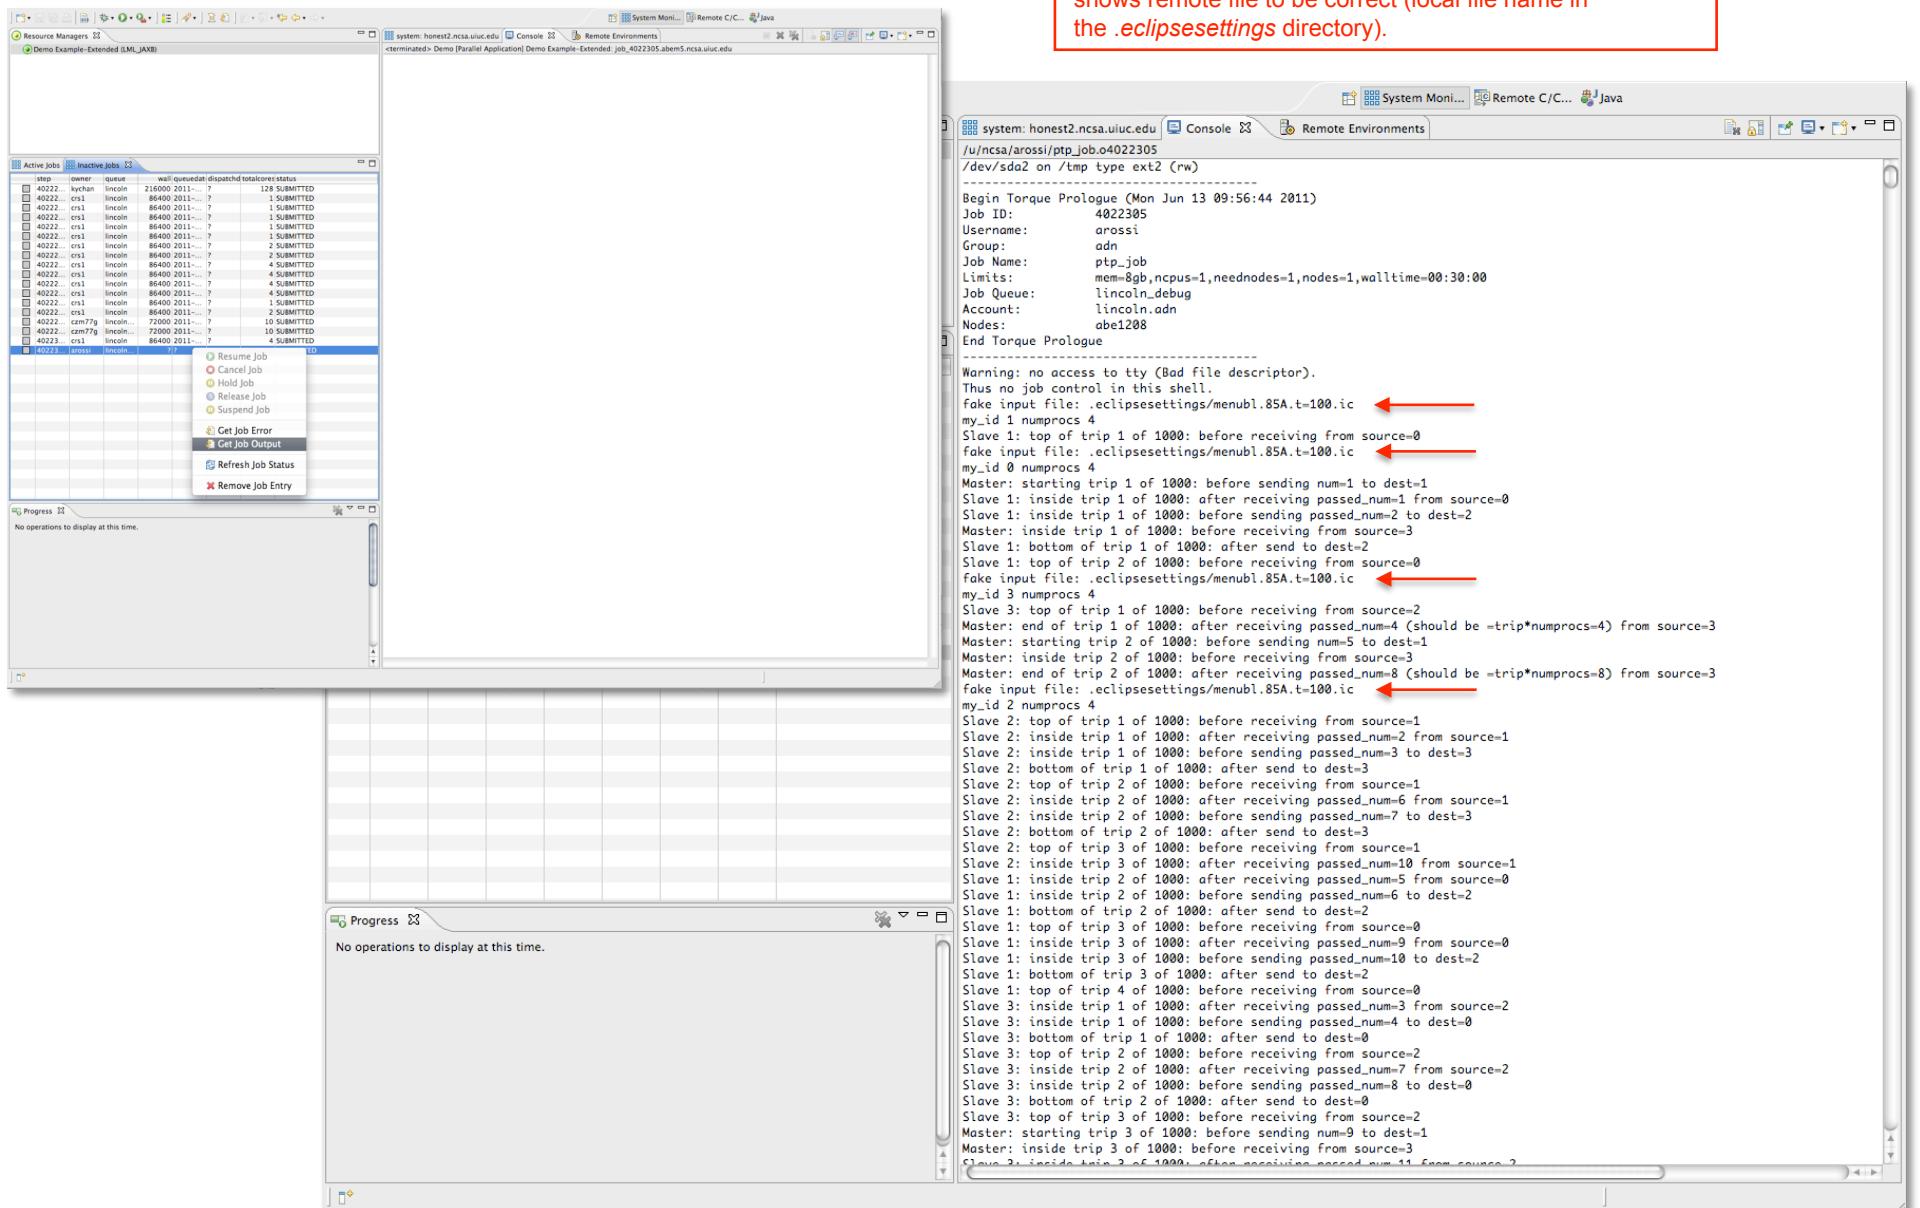
Create, manage, and run configurations

Create a configuration to launch a parallel application in Parallel Perspective



Run Result (case b)

Running from the last setting (`local.input.file`); output of job shows remote file to be correct (local file name in the `.eclipsesettings` directory).



Final Example: Configurable Input

As a final example of some of the possibilities afforded by the Configurable Resource Manager, we show here two screen shots of a dual-panel Launch Tab, the second of which allows the user to configure the values in an input parameter (Fortran “namelist”) file which is staged (as in the preceding example) over as a managed file.

This example required the addition of attributes corresponding to the variables, along with defaults values, specified by the input file, the addition of three composites of text widgets, and the specification of a managed file whose content resembles that of the `<script>` element: a series of `<line>` elements containing resolvable `<arg>` elements.

Configurable Job Input

