

Table of Contents

Table of Contents	1
Job Manager for remote execution of QuantumATK scripts	2
A single remote machine	3
Settings	4
Environment	7
Resources	8
Notifications	9
Diagnostics	9
Save and test the new machine	10
Custom job settings	14
Debugging	15
Adding several remote machines	15
Rename and export	16
A machine for threaded jobs	17



[Docs](#) » [Tutorials](#) » [Introduction to using NanoLab - the graphical user interface of QuantumATK](#) »
Job Manager for remote execution of QuantumATK scripts

Job Manager for remote execution of QuantumATK scripts

Version: 2017.0

Downloads & Links

- [PDF version](#)
- [PDF version for QuantumATK 2016](#)
- [mpi_check.py](#)
- [silicon.py](#)

In this tutorial you will learn how to use the Job Manager for execution of QuantumATK jobs on remote computing clusters. In particular, you will learn how to:

- add a remote machine to the Machine Manager;
- use custom Machine Settings for individual jobs;
- add several different machines and import/export machine settings.

Important

You will set up a remote machine for running jobs in parallel using MPI, as well as with threading. We strongly recommend you go through the tutorial [Job Manager for local execution of QuantumATK scripts](#) before continuing with this one.

Since ATK 2017 Intel's `mpiexec.hydra` is provided on both Windows and Linux versions - **this is the recommended way to run QuantumATK in parallel**. The `mpiexec.hydra` binary shipped with QuantumATK is located in the folder `libexec/mpiexec.hydra` present in your installation folder.

Note

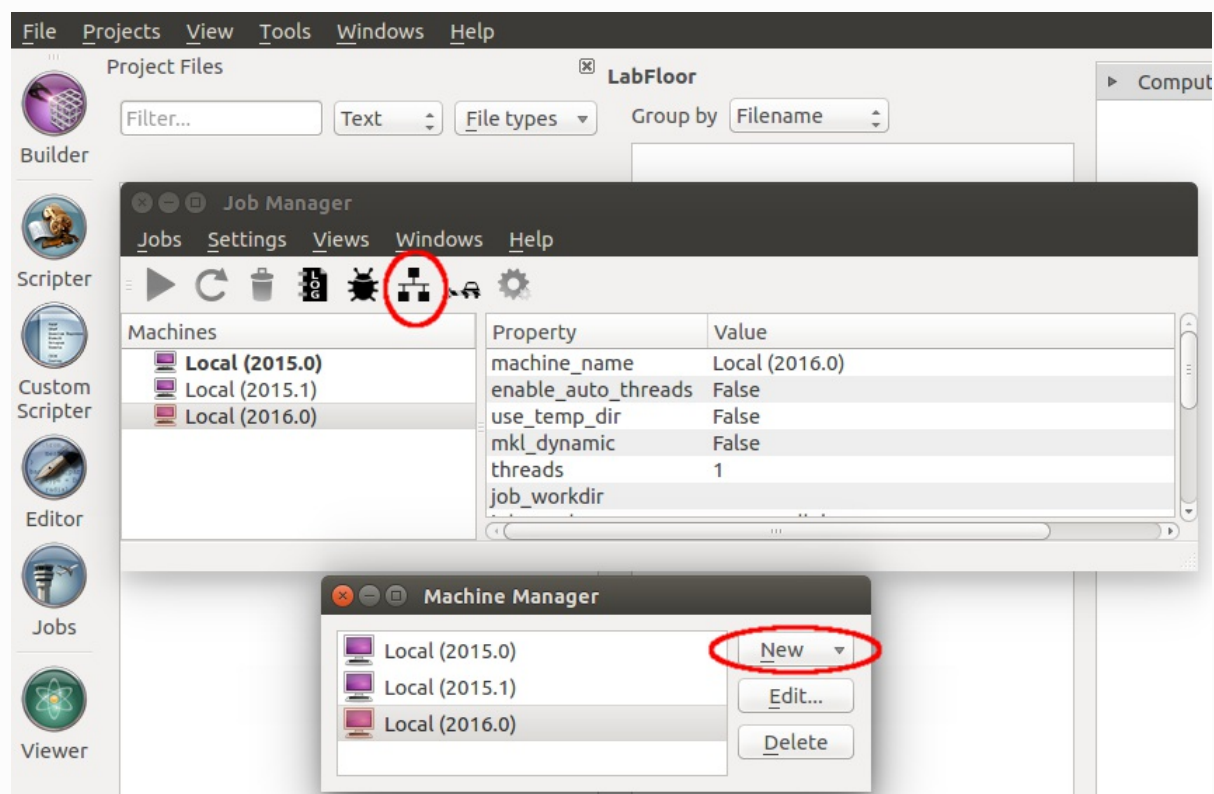
There are two essential requirements when using the Job Manager for executing and managing QuantumATK jobs on a remote cluster. You need:

1. QuantumATK installed on the local machine **and** on the cluster;
2. an SSH connection from your local machine to the cluster.

Please refer to the tutorial [SSH keys](#) if you need help setting up the SSH connection.

A single remote machine

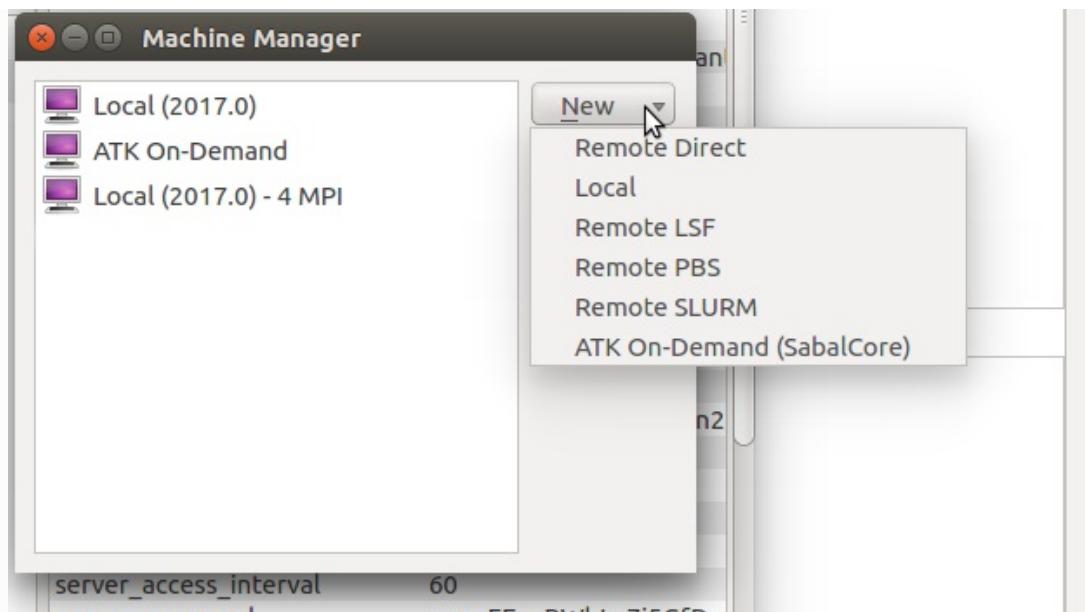
Open the  **Machine Manager**, and click *New* in order to add a new machine.



Note

Please note that the Machine Manager options may differ between versions of QuantumATK.


The menu that appears has six options, *Local*, *Remote PBS*, *Remote LSF*, *Remote SLURM*, *Remote Direct*, and *ATK On-Demand*.



Choose the type of remote machine that matches the job scheduling on your remote Linux cluster, and start setting up the connection. For more information about the QuantumATK On-Demand option [click here](#). Here, you will choose the *Remote PBS* machine to set up a machine with a PBS job scheduling system. The option *Remote Direct* is appropriate for clusters that have no queue system.

The **Machine Settings** widget pops up. It has five main tabs:

- Settings (*remote connection*),
- Environment (*software on the remote cluster*),
- Resources (*allocated computing resources and time*),
- Notifications (*job progress updates*),
- Diagnostics (*check the current setting*).



Hint

Most options have default values (which must be checked). Red fields are mandatory, but have no default.

Settings

Connection to the remote cluster.

Machine Settings

Settings
Environment
Resources
Notifications
Diagnostics

Machine name* Remote PBS

Connection

Hostname*

Port 22

Username*

Use ssh key ☒

Private key path /home/user/.ssh/id_rsa ... +

Key password Passwordless key

Queuing System

Node type names

Queue names

Path to PBS binaries use PATH environment variable

* Required fields

Export... Import... Preview Script Cancel OK

You need to specify the following fields:

Machine name
 Hostname
 Username
 Use ssh key

The following options vary from cluster to cluster.

Private key path
 The directory containing your private SSH key.

Node type names
 Optional, but in this example you have access to node type "orange".

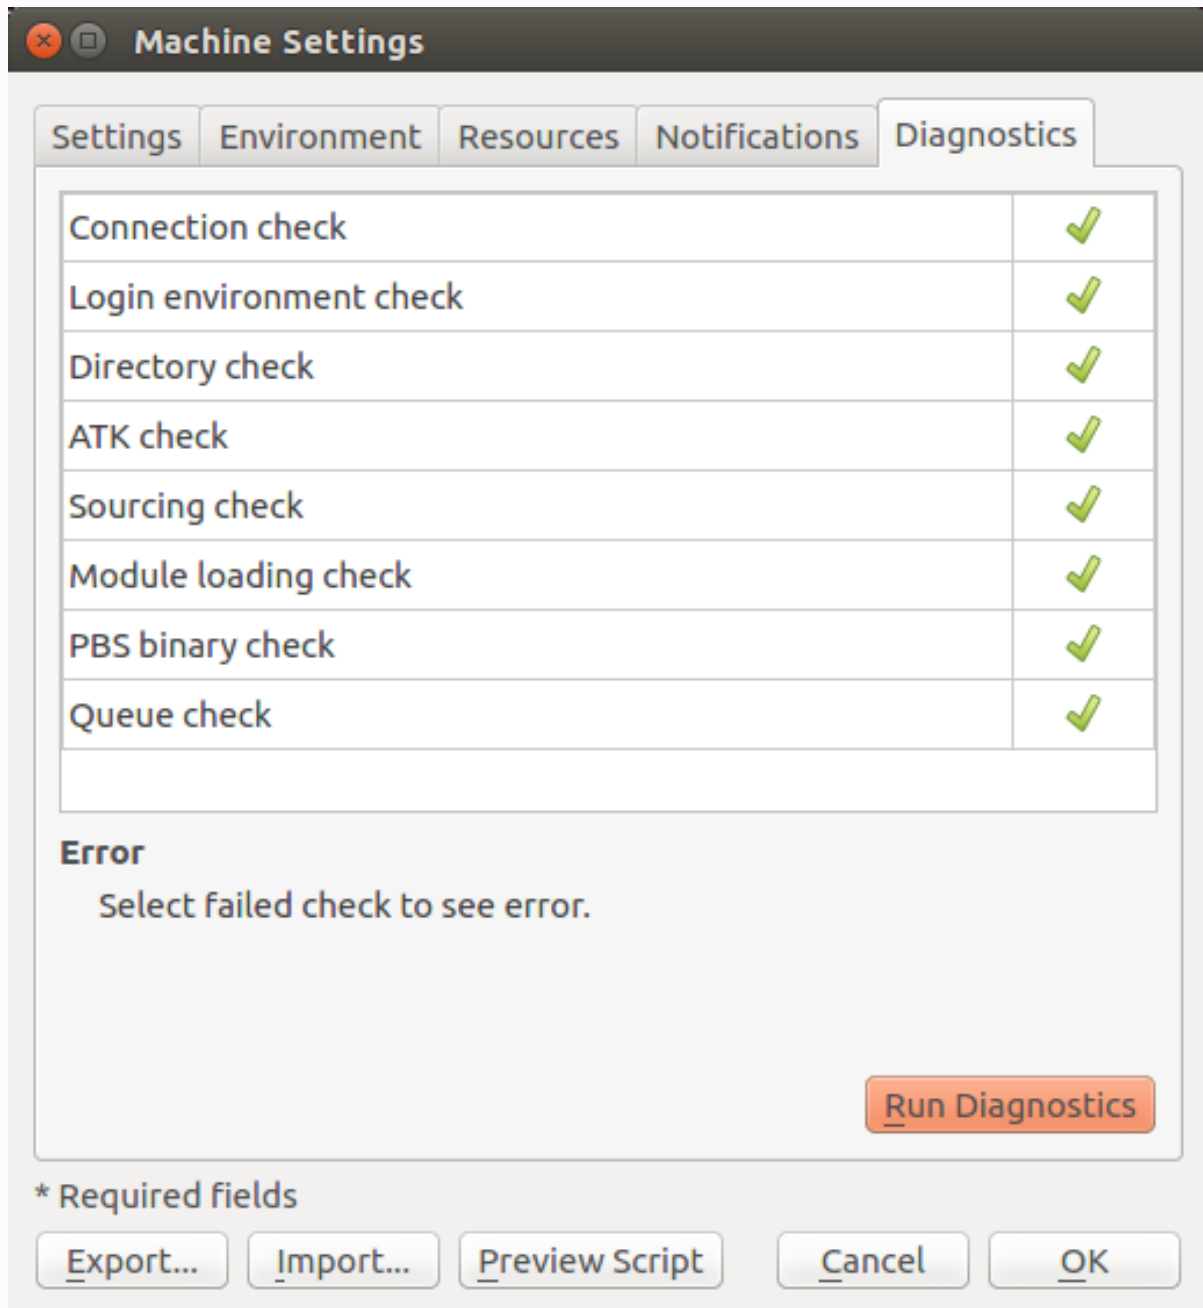
Queue names
 Optional, could for example be "long".

Path to PBS binaries

The directory containing the **qsub** and other PBS executables must be specified. Log on to the cluster and use the `which` command to locate it (e.g. `/usr/local/torque-4.2.8/bin`):

```
$ which qsub
/usr/local/torque-4.2.8/bin/qsub
```

Once all settings are added, you can check if they are correct by navigating to the Diagnostics tag and clicking the “Run Diagnostics” button.



Machine Settings

Settings Environment Resources Notifications **Diagnostics**

Connection check	✓
Login environment check	✓
Directory check	✓
ATK check	✓
Sourcing check	✓
Module loading check	✓
PBS binary check	✓
Queue check	✓

Error
Select failed check to see error.

Run Diagnostics

* Required fields

Export... **Import...** **Preview Script** **Cancel** **OK**

Tip

The Diagnostics tab checks if the options in the **Settings** and **Environment** tabs allow your local computer to connect to the remote cluster and execute the commands needed for job submission and management.

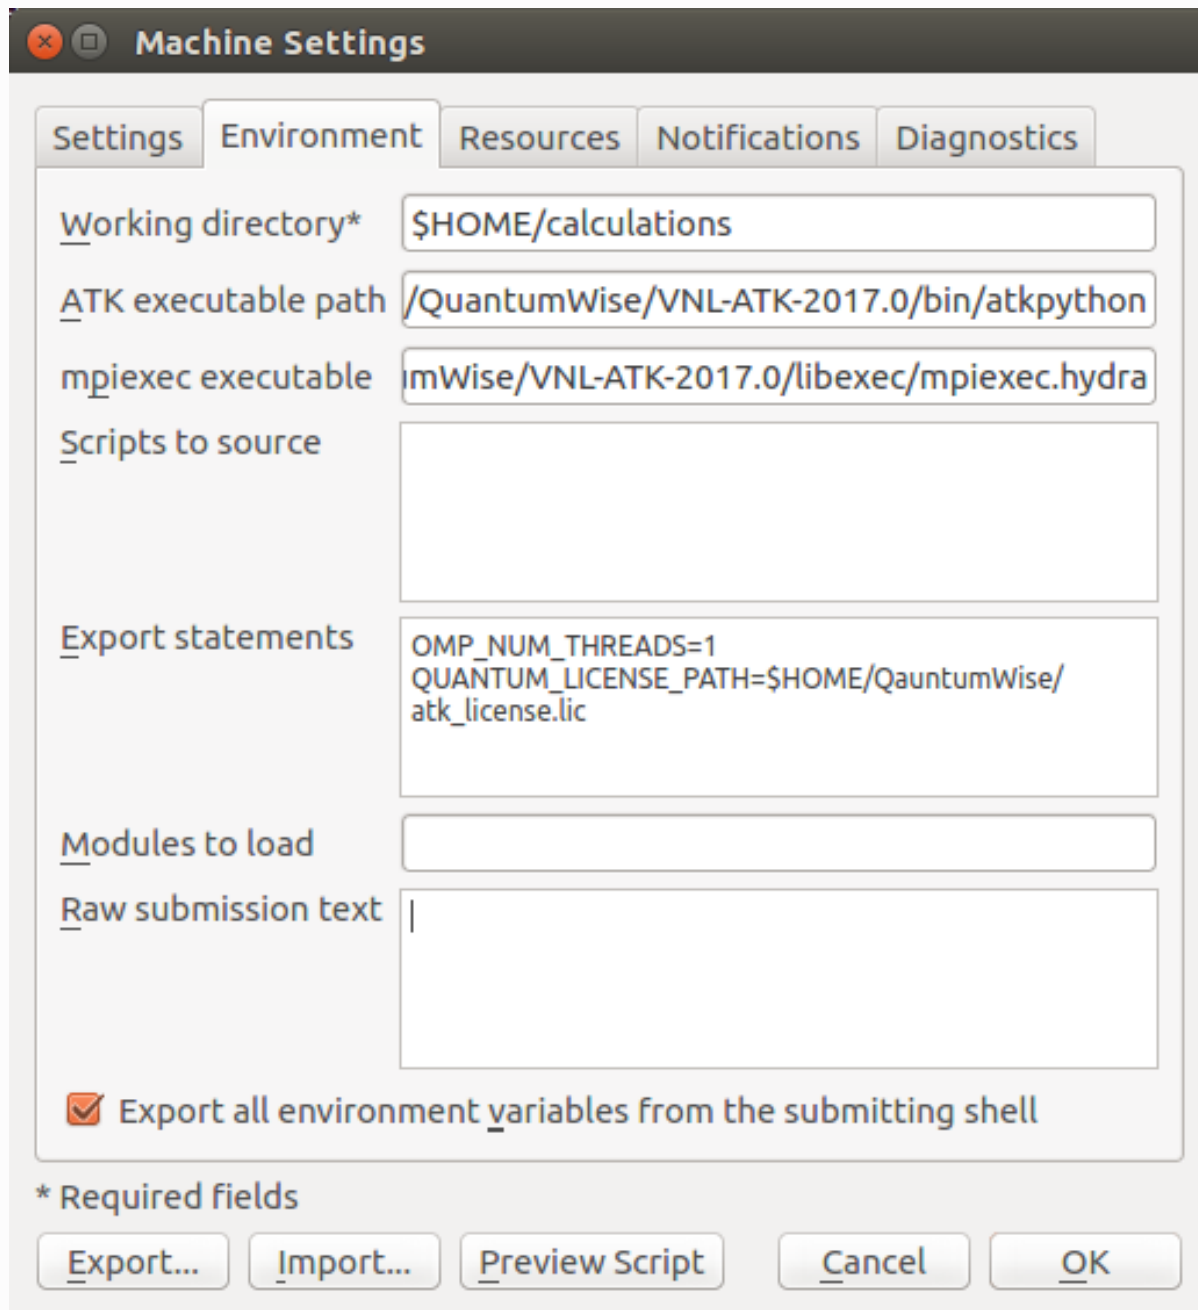
If some field is not specified, the diagnostics will check the default setting. If the connection to the

cluster works well with that default, or is at least not disrupted, it will be marked as OK.

You therefore need to run an actual test job to make absolutely sure that all settings are indeed OK.

Environment

Computing environment on the cluster.



The image shows a 'Machine Settings' dialog box with the 'Environment' tab selected. The dialog has five tabs: 'Settings', 'Environment', 'Resources', 'Notifications', and 'Diagnostics'. The 'Environment' tab contains several fields for configuring the computing environment on a remote cluster. The fields are: 'Working directory*' with the value '\$HOME/calculations'; 'ATK executable path' with the value '/QuantumWise/VNL-ATK-2017.0/bin/atkpython'; 'mpiexec executable' with the value 'mWise/VNL-ATK-2017.0/libexec/mpiexec.hydra'; 'Scripts to source' (an empty text area); 'Export statements' with the text 'OMP_NUM_THREADS=1' and 'QUANTUM_LICENSE_PATH=\$HOME/QauntumWise/atk_license.lic'; 'Modules to load' (an empty text area); and 'Raw submission text' (an empty text area). At the bottom of the tab, there is a checked checkbox labeled 'Export all environment variables from the submitting shell'. Below the tabs, there is a section labeled '* Required fields' with four buttons: 'Export...', 'Import...', 'Preview Script', 'Cancel', and 'OK'.

Machine Settings

Settings Environment Resources Notifications Diagnostics

Working directory* \$HOME/calculations

ATK executable path /QuantumWise/VNL-ATK-2017.0/bin/atkpython

mpiexec executable mWise/VNL-ATK-2017.0/libexec/mpiexec.hydra

Scripts to source

Export statements OMP_NUM_THREADS=1
QUANTUM_LICENSE_PATH=\$HOME/QauntumWise/
atk_license.lic

Modules to load

Raw submission text

☒ Export all environment variables from the submitting shell

* Required fields

Export... Import... Preview Script Cancel OK

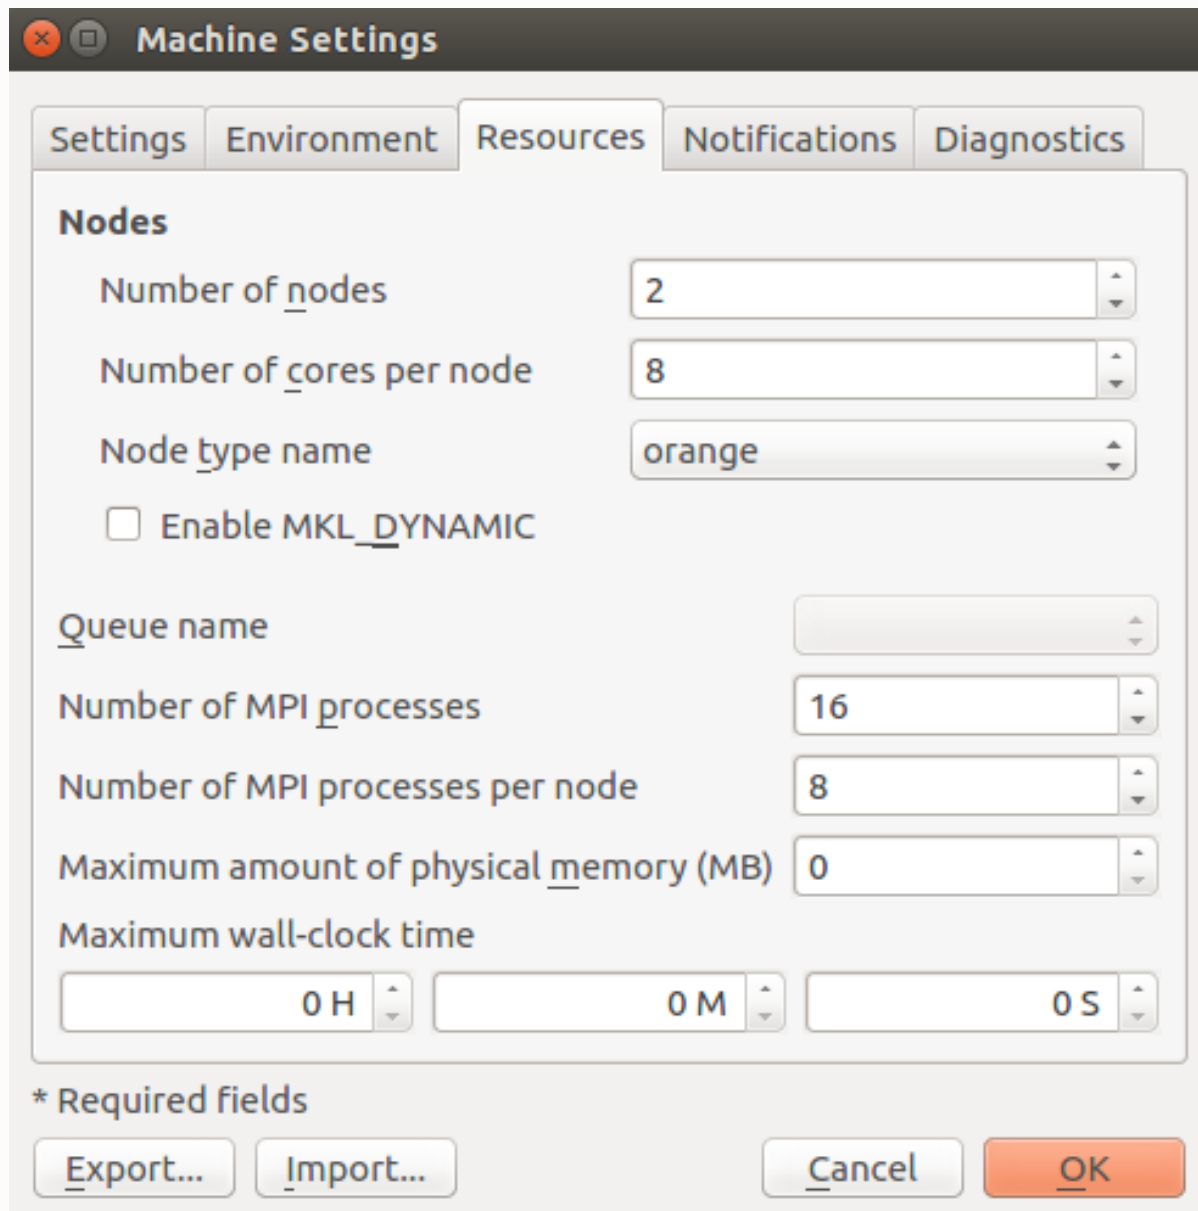
This tab concerns the environment (directories, executables, modules, etc.) on the *remote cluster*, not on your local computer. `$HOME` is therefore your home directory on the cluster, and QuantumATK must of course be installed on the cluster.

Since ATK 2017 Intel's `mpiexec.hydra` is provided on both Windows and Linux versions - **this is the recommended way to run QuantumATK in parallel**. The `mpiexec.hydra` binary shipped with QuantumATK is located in the folder `libexec/mpiexec.hydra` present in your installation folder.

Any scripts that should be sourced in order to get the environment working must be listed. The same goes for required export statements (may be needed to correctly set the `SNPSLMD_LICENSE_FILE`) and cluster modules that should be loaded.

Resources

Computing resources requested at job submission.



The image shows a 'Machine Settings' dialog box with the 'Resources' tab selected. The dialog has five tabs: 'Settings', 'Environment', 'Resources', 'Notifications', and 'Diagnostics'. Under the 'Resources' tab, there is a section titled 'Nodes' with the following settings: 'Number of nodes' set to 2, 'Number of cores per node' set to 8, 'Node type name' set to 'orange', and an unchecked checkbox for 'Enable MKL_DYNAMIC'. Below this is a 'Queue name' field. Further down are 'Number of MPI processes' set to 16, 'Number of MPI processes per node' set to 8, 'Maximum amount of physical memory (MB)' set to 0, and 'Maximum wall-clock time' set to 0 H, 0 M, and 0 S. At the bottom, there is a section for '* Required fields' with 'Export...' and 'Import...' buttons, and a 'Cancel' button next to an 'OK' button.

Setting	Value
Number of nodes	2
Number of cores per node	8
Node type name	orange
Enable MKL_DYNAMIC	<input type="checkbox"/>
Queue name	
Number of MPI processes	16
Number of MPI processes per node	8
Maximum amount of physical memory (MB)	0
Maximum wall-clock time	0 H 0 M 0 S

* Required fields

Export... Import... Cancel OK

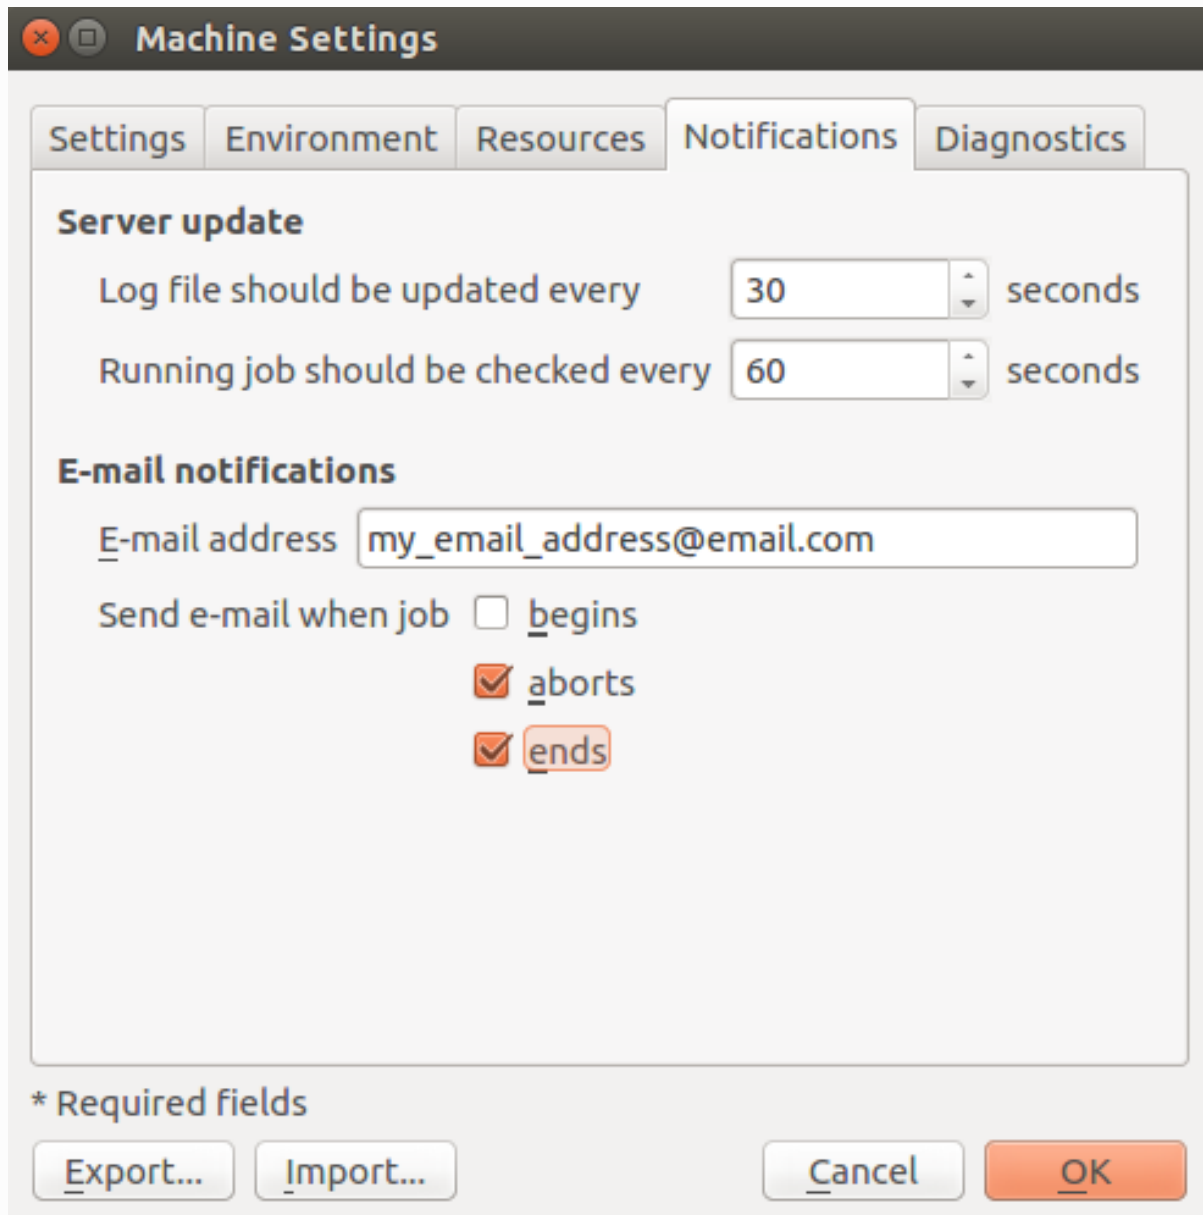
This tab specifies the default computing resources (nodes, cores, queue, time, etc.) requested at job submission. Note that `MKL_DYNAMIC` is disabled by default, which means that MKL is not allowed to dynamically decrease the number of threads at runtime.

Tip

For a job with only MPI parallelization (no threading), the number of nodes times the number of cores per node should equal the (total) number of MPI processes. In the example above, you have $2 \times 8 = 16$ cores, so you ask for 16 MPI processes and specify that each node should have 8 of those processes (only one MPI process per core).

Notifications

Job progress reports.



The image shows a 'Machine Settings' dialog box with the 'Notifications' tab selected. The dialog has a title bar with a close button and a maximize button. Below the title bar are five tabs: 'Settings', 'Environment', 'Resources', 'Notifications', and 'Diagnostics'. The 'Notifications' tab is active and contains the following settings:

- Server update**
 - Log file should be updated every seconds
 - Running job should be checked every seconds
- E-mail notifications**
 - E-mail address
 - Send e-mail when job
 - ☐ begins
 - ☒ aborts
 - ☒ ends

At the bottom of the dialog, there is a section labeled '* Required fields' with three buttons: 'Export...', 'Import...', and 'Cancel'. To the right of these buttons is an 'OK' button.

The Job Manager will regularly check the job progress on the remote cluster and report it in the log file. You can also receive e-mail notifications from the PBS scheduler when the job starts, finishes, or terminates.

Diagnostics

Use this tab to test the machine settings. Green check-mark indicates that all settings appear to be OK. Red circle indicates some problem that should be fixed in one of the tabs.

Machine Settings

SettingsEnvironmentResourcesNotificationsDiagnostics

Connection check	✓
Login environment check	✓
Directory check	✓
ATK check	✓
Sourcing check	✓
Module loading check	✓
PBS binary check	✓
Queue check	✓

Error

Select failed check to see error.

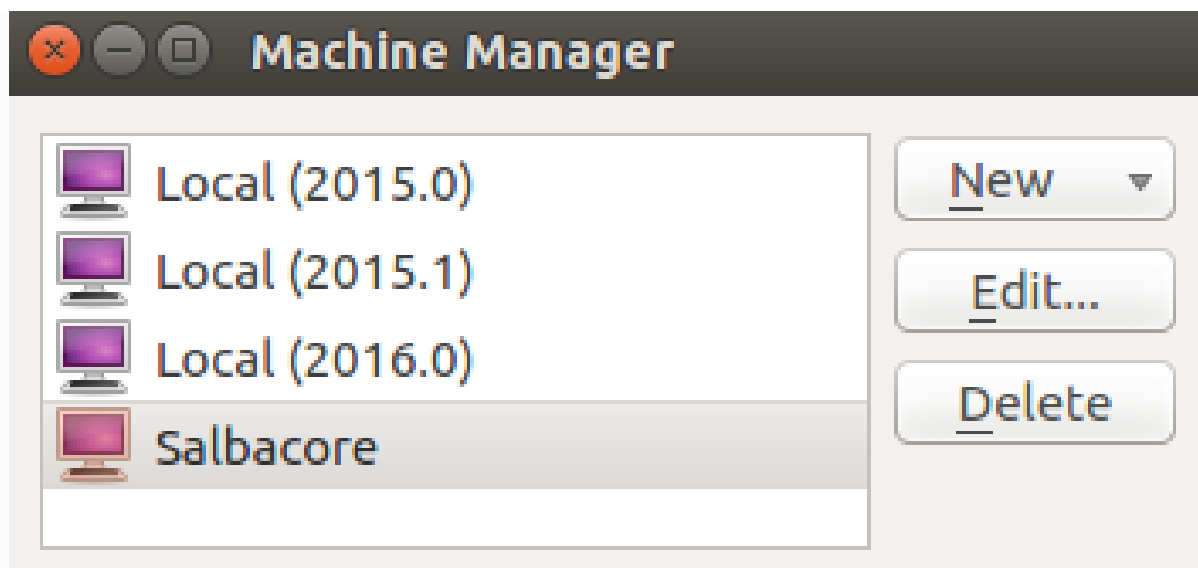
Run Diagnostics

* Required fields


Export...Import...Preview ScriptCancelOK

Save and test the new machine

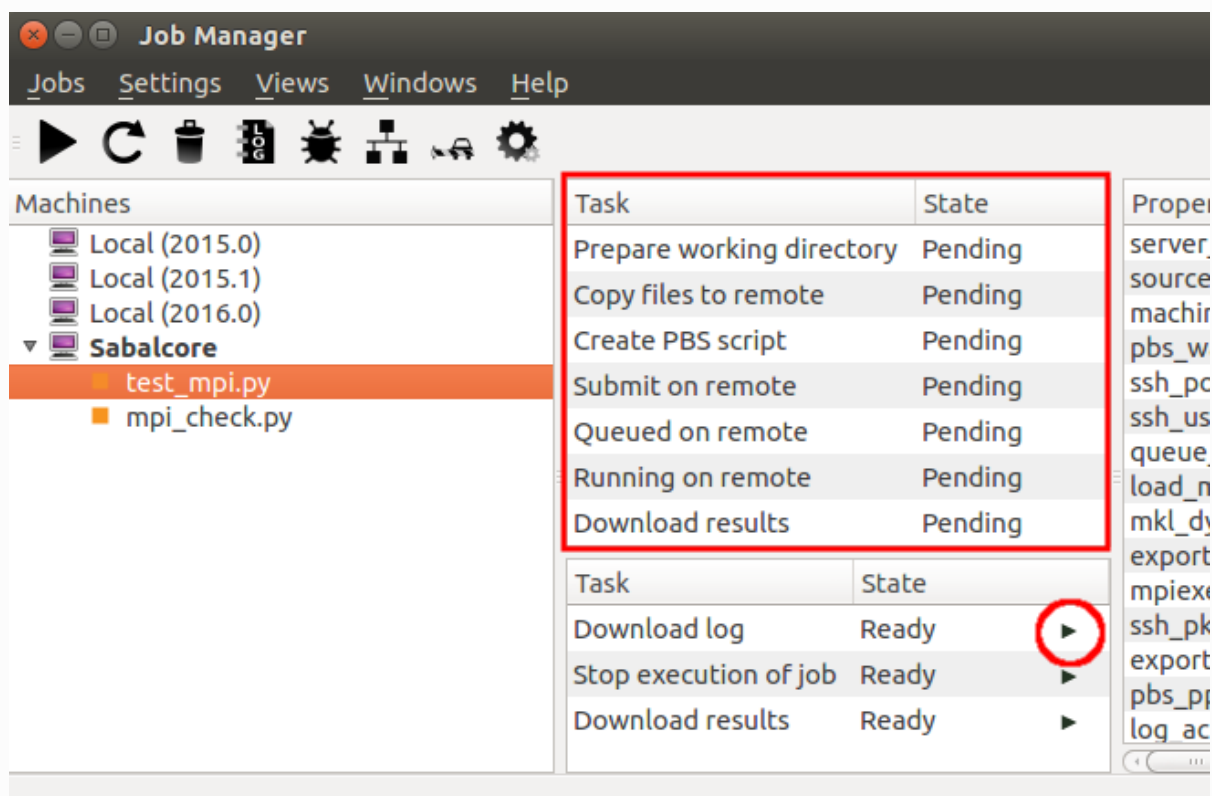
Click *OK* to add the machine to the **Machine Manager**.




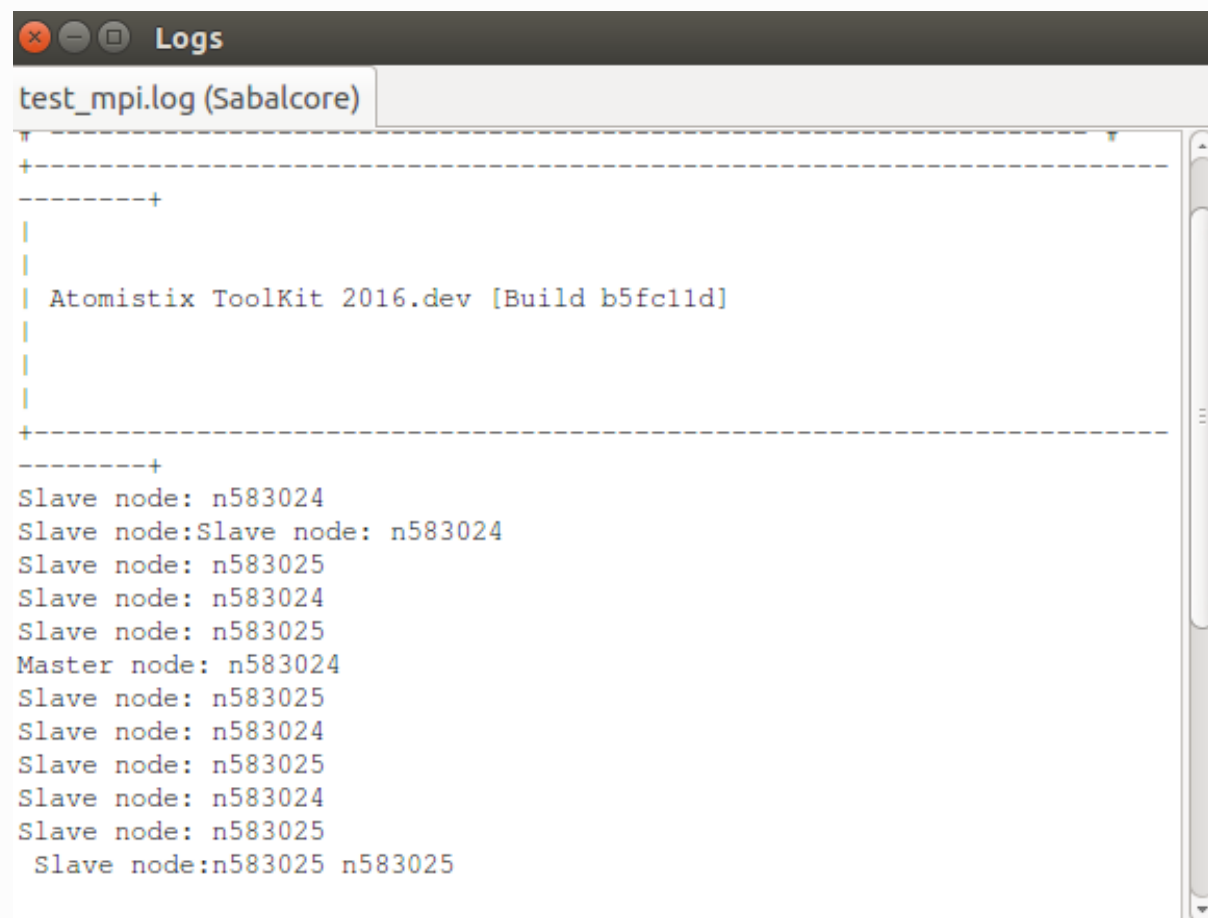
Next, you should run the test scripts [mpi_check.py](#) and [test_mpi.py](#) to test the machine settings. In the QuantumATK main window, drag and drop a script onto the Job Manager. Choose the newly added machine (in this example "Salbacore"), and click *OK*.



Click to submit the job and watch how the **Task State** changes from *Pending* to *Finished*. For a longer job, you can click *Download log* to regularly retrieve the log file during the remote job execution.



Once a job has finished you should check the log file to inspect the job output to see if the expected number of processes were used. Click the  icon to open the log. In the two examples below, 16 cores were used in total on two different nodes.



```
test_mpi.log (Sabalcore)
-----+
|
| Atomistix ToolKit 2016.dev [Build b5fc11d]
|
|
+-----+
Slave node: n583024
Slave node:Slave node: n583024
Slave node: n583025
Slave node: n583024
Slave node: n583025
Master node: n583024
Slave node: n583025
Slave node: n583024
Slave node: n583025
Slave node: n583024
Slave node: n583025
Slave node:n583025 n583025
```

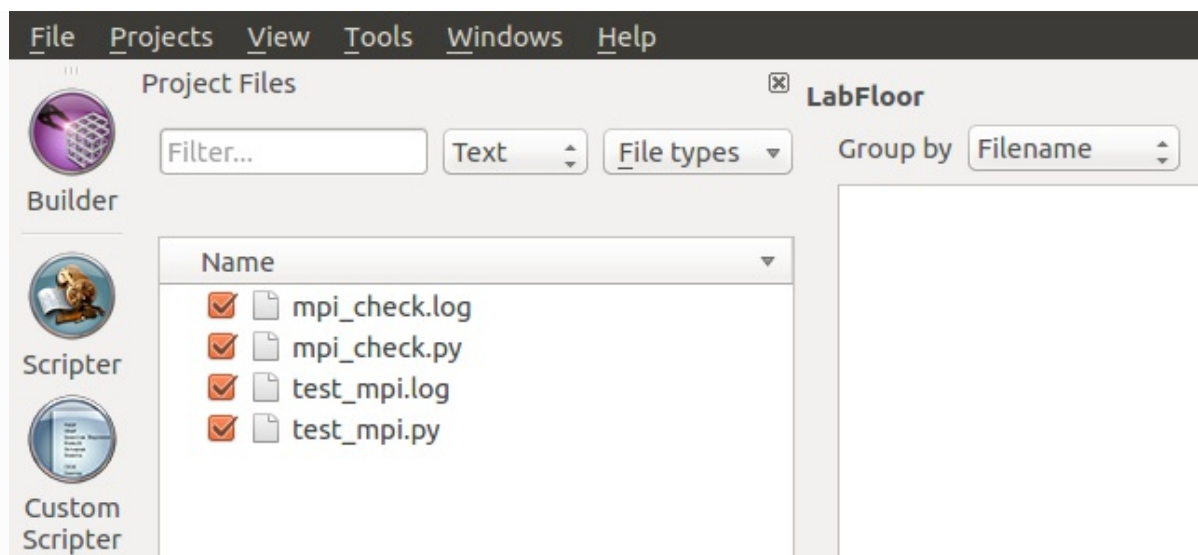
```
Logs
mpi_check.log (Sabalcore)
+-----+
|
| Atomistix ToolKit 2016.dev [Build b5fc11d]
|
+-----+
MPI process 3 of 16 reporting.MPI process 1 of 16 reporting.
MPI process 4 of 16 reporting.
MPI process 8 of 16 reporting.MPI process 14 of 16 reporting.
MPI process 11 of 16 reporting.
MPI process 12 of 16 reporting.
MPI process 5 of 16 reporting.
MPI process 13 of 16 reporting.
MPI process 6 of 16 reporting.
MPI process 7 of 16 reporting.

MPI process 15 of 16 reporting.


MPI process 0 of 16 reporting.
MPI process 2 of 16 reporting.
MPI process 10 of 16 reporting.
MPI process 9 of 16 reporting.

Timing:                                Total      Per Step      %
-----
Loading Modules + MPI    :      6.36 s      6.36 s      95.02% |
=====|
-----
Total                    :      6.70 s
```

The job outputs have of course also appeared on the LabFloor:




Custom job settings

As explained in the tutorial [Job Manager for local execution of QuantumATK scripts](#), you can customize many of the job settings before submitting a job. Use the script [silicon.py](#) as an example QuantumATK script. Download the script, send it to the **Job Manager**, and choose the newly created remote machine. Then click the  **Job Settings** plug-in.


You can now customize the *Resources* and *Notifications* tabs, and thereby submit the job with settings different from the default ones you specified above. For example, you can change the number of requested cores and cluster queue and/or maximum wall-clock time. You can also change the notifications settings.

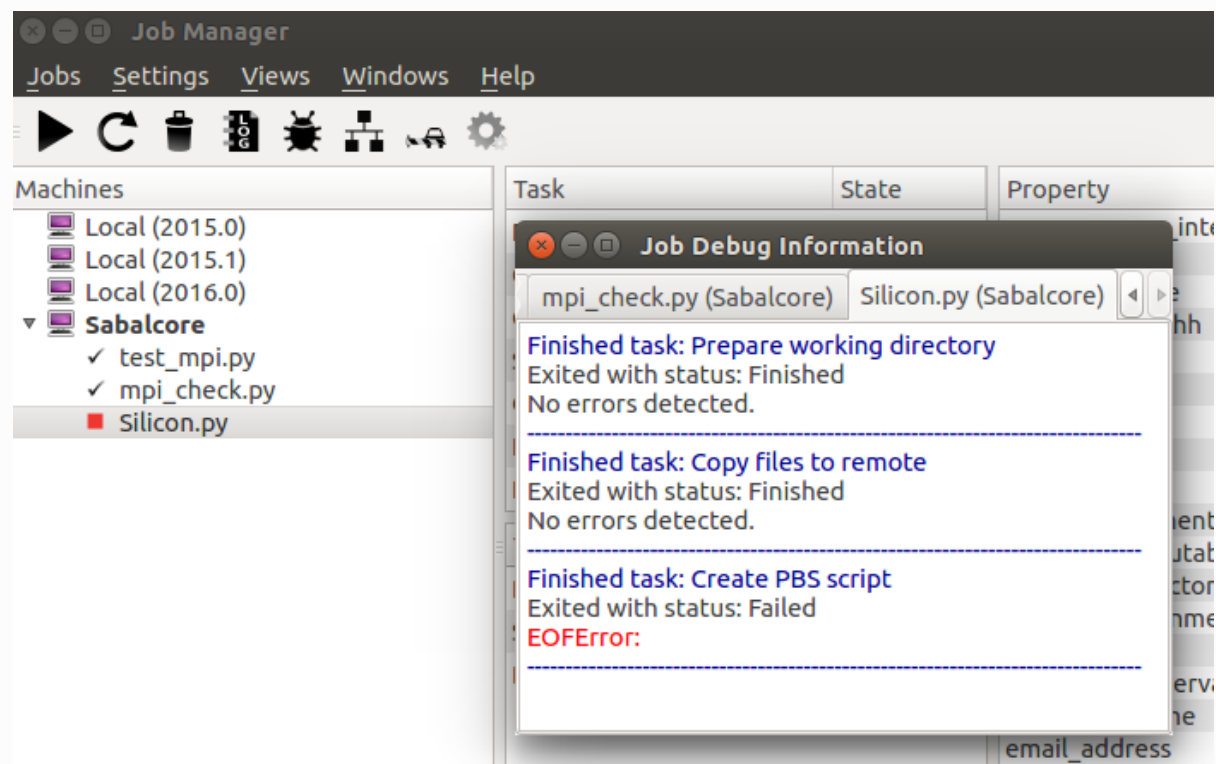
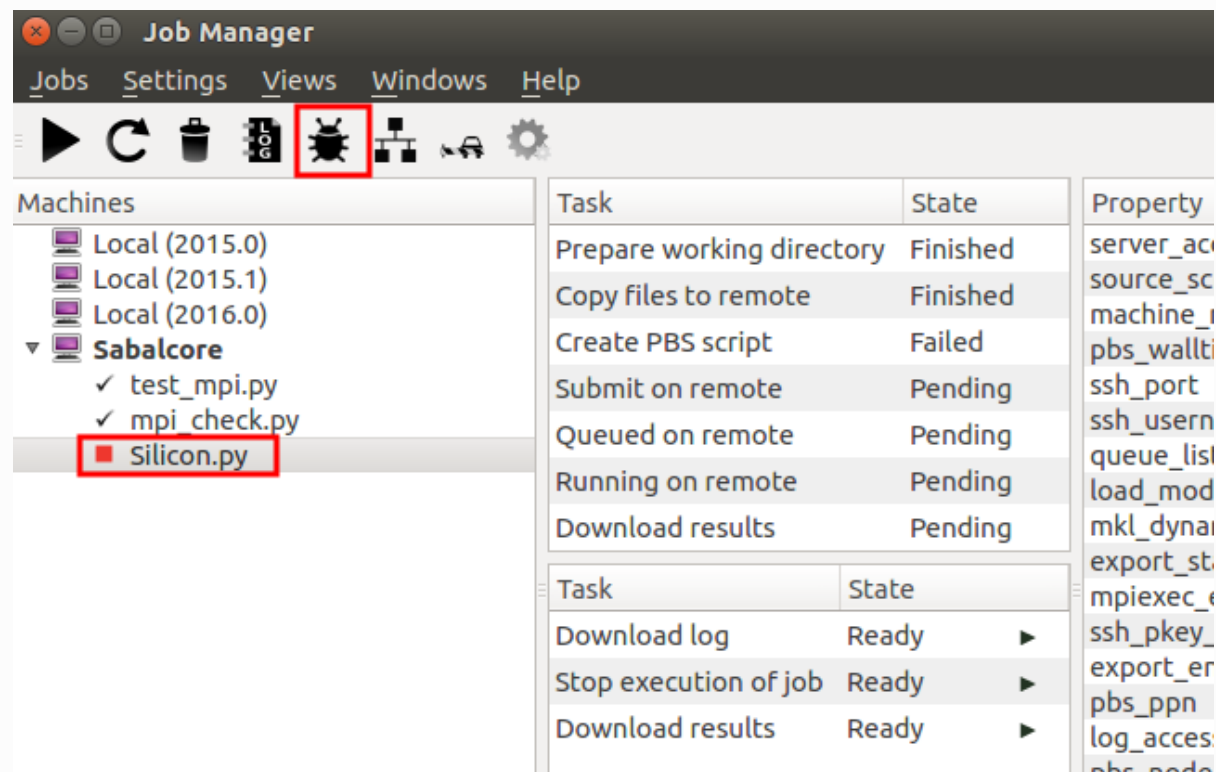
Set up the job settings as you like, e.g. 4 cores on a single node and 4 MPI processes. Then click *OK* and submit the calculation to the remote cluster.

The job log is automatically retrieved when the job finishes. For long jobs, however, you need to click "Download log" if you want to see the log while the job is running.



Debugging

If an error occurs during the execution of the job, this will be indicated by a red square in the queue, as shown below. You can then click the  **Debug logs** icon to open the job debugs information window, which will show you details about the error.



Adding several remote machines

Several remote machines can be added to the **Machine Manager**. You can of course add new machines from scratch, but you can also export/import the settings of an existing machine and use those as a template for new machines.


Tip


The import/export functionality is very convenient for sharing machine settings within a group of researchers.

In the following, you will rename and export the settings of the newly created machine, and then add one more remote machine suitable for threaded calculations.

Rename and export

In the **Job Manager**, remove the jobs that are in the queue of the newly added machine (named “Sabalcore” in our example).

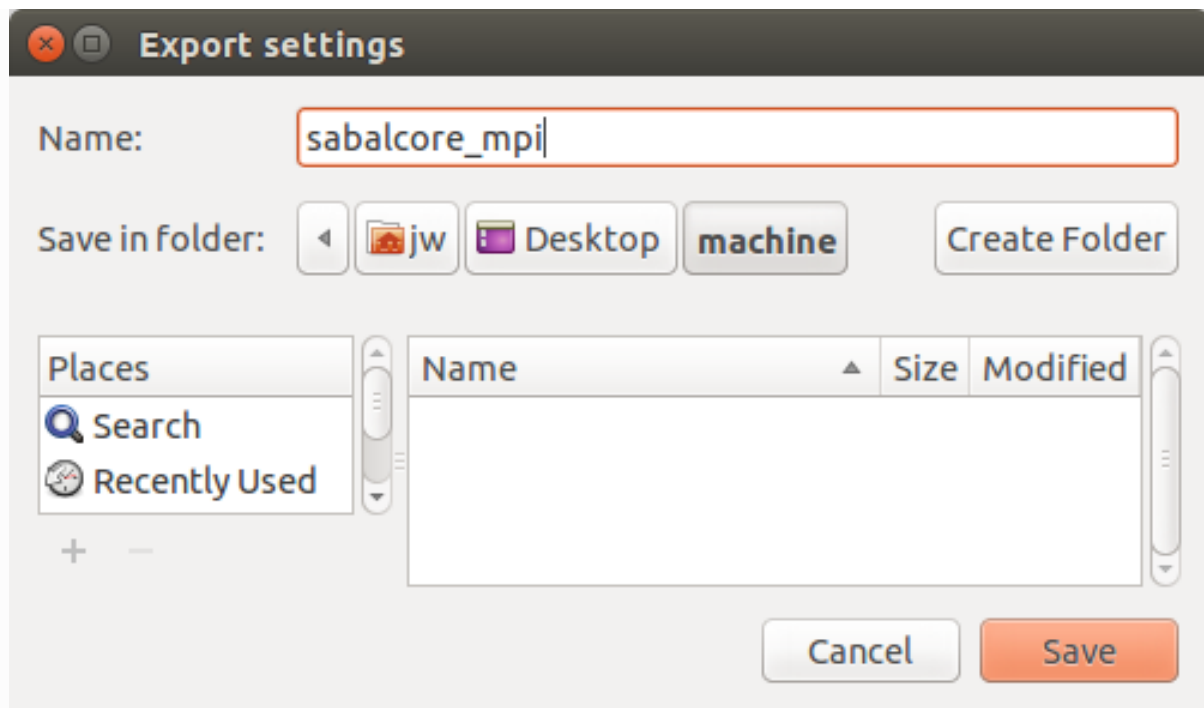
Click  **Machine Settings** and choose to *Edit* the machine.



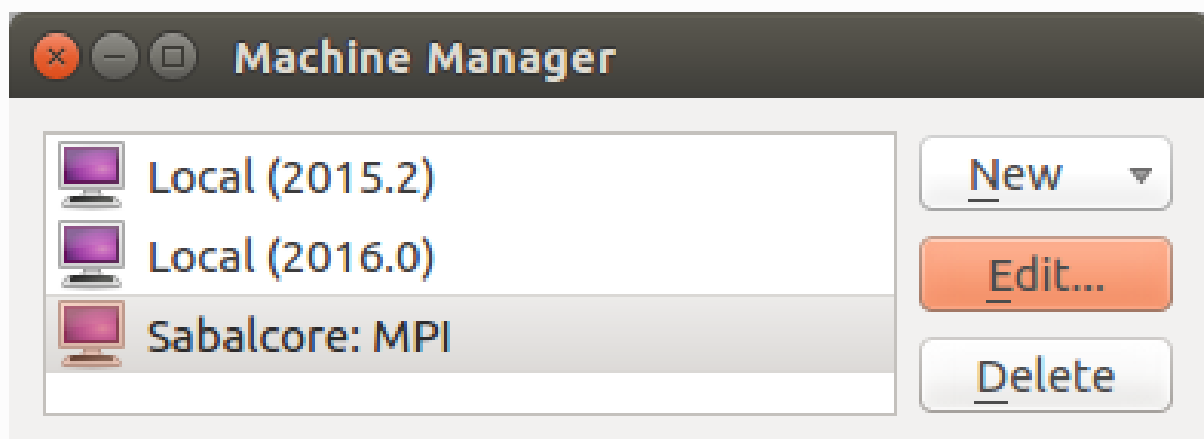
Note

You always need to empty the queue of a machine before you can edit its settings.

This machine is already set up for MPI parallelization. Therefore, append “: MPI” to the machine name, click *Export*, and save the settings in a file.



Click *OK* to accept the renaming and return to the Machine Manager window.



A machine for threaded jobs

Next, add a new remote machine with default settings for “Remote PBS”. Then click *Import* and import the settings file you just saved.

You can now modify the settings to create a machine with default settings that are suitable for a threaded calculation:

- In the Environment tab, remove the `OMP_NUM_THREADS=1` export statement.
- In the Resources tab, use only a single MPI process but enable dynamic scheduling of MKL threads.

Machine Settings

Settings

Environment

Resources

Notifications

Diagnostics

Working directory* \$HOME/calculations

ATK executable path \$HOME/QuantumWise/VNL-ATK-2016.0/bin/atkpython

mpiexec executable mpiexec

Scripts to source ~/.bashrc

Export statements QUANTUM_LICENSE_PATH=\$HOME/QuantumWise/atk_license.lic

Modules to load mpich3

☒ Export all environment variables from the submitting shell

* Required fields

Export...

Import...

Cancel

OK

Machine Settings

Settings Environment Resources Notifications Diagnostics

Nodes

Number of nodes 1

Number of cores per node 8

Node type name orange

☒ Enable MKL DYNAMIC

Queue name

Number of MPI processes 1

Number of MPI processes per node 1

Maximum amount of physical memory (MB) 0

Maximum wall-clock time





0 H 0 M 0 S

* Required fields

Export... Import... Cancel OK

Give the new machine a reasonable name, e.g. "Salbacore: Threading", and click *OK* to add the machine to the Machine Manager. This machine should be convenient for submitting ATK-ForceField calculations.

Machine Manager

 Local (2015.2)	<div>New ▼</div> <div>Edit...</div> <div>Delete</div>
 Local (2016.0)	
 Sabalcore: MPI	
 Sabalcore: Threading	

[← Previous](#)

[Next →](#)

© Copyright 2022 Synopsys, Inc. All Rights Reserved.