

**MATLAB® Parallel Server™**  
System Administrator's Guide



**MATLAB®**

R2021b



## How to Contact MathWorks



Latest news: [www.mathworks.com](http://www.mathworks.com)  
Sales and services: [www.mathworks.com/sales\\_and\\_services](http://www.mathworks.com/sales_and_services)  
User community: [www.mathworks.com/matlabcentral](http://www.mathworks.com/matlabcentral)  
Technical support: [www.mathworks.com/support/contact\\_us](http://www.mathworks.com/support/contact_us)



Phone: 508-647-7000



The MathWorks, Inc.  
1 Apple Hill Drive  
Natick, MA 01760-2098

*MATLAB® Parallel Server™ System Administrator's Guide*

© COPYRIGHT 2005–2021 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

FEDERAL ACQUISITION: This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

### Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See [www.mathworks.com/trademarks](http://www.mathworks.com/trademarks) for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

### Patents

MathWorks products are protected by one or more U.S. patents. Please see [www.mathworks.com/patents](http://www.mathworks.com/patents) for more information.

## Revision History

November 2005	Online only	New for Version 2.0 (Release 14SP3+)
December 2005	Online only	Revised for Version 2.0 (Release 14SP3+)
March 2006	Online only	Revised for Version 2.0.1 (Release 2006a)
September 2006	Online only	Revised for Version 3.0 (Release 2006b)
March 2007	Online only	Revised for Version 3.1 (Release 2007a)
September 2007	Online only	Revised for Version 3.2 (Release 2007b)
March 2008	Online only	Revised for Version 3.3 (Release 2008a)
October 2008	Online only	Revised for Version 4.0 (Release 2008b)
March 2009	Online only	Revised for Version 4.1 (Release 2009a)
September 2009	Online only	Revised for Version 4.2 (Release 2009b)
March 2010	Online only	Revised for Version 4.3 (Release 2010a)
September 2010	Online only	Revised for Version 5.0 (Release 2010b)
April 2011	Online only	Revised for Version 5.1 (Release 2011a)
September 2011	Online only	Revised for Version 5.2 (Release 2011b)
March 2012	Online only	Revised for Version 6.0 (Release 2012a)
September 2012	Online only	Revised for Version 6.1 (Release 2012b)
March 2013	Online only	Revised for Version 6.2 (Release 2013a)
September 2013	Online only	Revised for Version 6.3 (Release 2013b)
March 2014	Online only	Revised for Version 6.4 (Release 2014a)
October 2014	Online only	Revised for Version 6.5 (Release 2014b)
March 2015	Online only	Revised for Version 6.6 (Release 2015a)
September 2015	Online only	Revised for Version 6.7 (Release 2015b)
March 2016	Online only	Revised for Version 6.8 (Release 2016a)
September 2016	Online only	Revised for Version 6.9 (Release 2016b)
March 2017	Online only	Revised for Version 6.10 (Release 2017a)
September 2017	Online only	Revised for Version 6.11 (Release 2017b)
March 2018	Online only	Revised for Version 6.12 (Release 2018a)
September 2018	Online only	Revised for Version 6.13 (Release 2018b)
March 2019	Online only	Revised for Version 7.0 (Release 2019a)
September 2019	Online only	Revised for Version 7.1 (Release 2019b)
March 2020	Online only	Revised for Version 7.2 (Release 2020a)
September 2020	Online only	Revised for Version 7.3 (Release 2020b)
March 2021	Online only	Revised for Version 7.4 (Release 2021a)
September 2021	Online only	Revised for Version 7.5 (Release 2021b)



## Introduction

### 1

<b>MATLAB Parallel Server Product Description</b> .....	<b>1-2</b>
<b>Product Overview</b> .....	<b>1-3</b>
Parallel Computing Concepts .....	<b>1-3</b>
Determining Product Installation and Versions .....	<b>1-3</b>

## Network Administration

### 2

<b>Requirements and Ports for MATLAB Parallel Server</b> .....	<b>2-2</b>
Port Configuration .....	<b>2-2</b>
Fully Qualified Domain Names .....	<b>2-2</b>
Security Considerations .....	<b>2-2</b>
<b>Use Different MPI Builds on UNIX Systems</b> .....	<b>2-4</b>
Build MPI .....	<b>2-4</b>
Use Your MPI Build .....	<b>2-4</b>
<b>Shut Down a Job Manager Cluster</b> .....	<b>2-6</b>
Linux and Macintosh Operating Systems .....	<b>2-6</b>
Microsoft Windows Operating Systems .....	<b>2-7</b>
<b>Customize Startup Parameters</b> .....	<b>2-9</b>
Define Script Defaults .....	<b>2-9</b>
Override Script Defaults .....	<b>2-10</b>
Increase Heap Memory .....	<b>2-11</b>
<b>Access Service Record Files</b> .....	<b>2-12</b>
Locate Log Files .....	<b>2-12</b>
Locate Checkpoint Folders .....	<b>2-12</b>
<b>Set MATLAB Job Scheduler Cluster Security</b> .....	<b>2-14</b>
Set the Security Level .....	<b>2-14</b>
Local, MATLAB Job Scheduler, and Network Passwords .....	<b>2-15</b>
Authorize Users for Job and Task Access .....	<b>2-16</b>
Set Secure Communication .....	<b>2-16</b>
<b>Troubleshoot Common Problems</b> .....	<b>2-18</b>
License Errors .....	<b>2-18</b>
Memory Errors on UNIX Operating Systems .....	<b>2-19</b>
Run Server Processes on Windows Network Installation .....	<b>2-19</b>

Required Ports .....	2-19
Ephemeral TCP Ports with Job Manager .....	2-20
Host Communications Problems .....	2-20
Verify Network Communications for Cluster Discovery .....	2-21

## Product Installation

### 3

<b>Choose Solution to Install and Configure MATLAB Parallel Server .....</b>	<b>3-2</b>
Activate the License for, Install, and Configure MATLAB Parallel Server on the Cluster .....	3-2
Configure a MATLAB Client for a User to Submit Jobs to the Cluster .....	3-3
Upgrade MATLAB Parallel Server .....	3-4
<b>Install and Configure MATLAB Parallel Server for Slurm .....</b>	<b>3-6</b>
Activate Your MATLAB Parallel Server License .....	3-6
Get the Installation Files .....	3-6
Install License Manager .....	3-7
Install Software on Compute Nodes .....	3-8
Install Software on Local Desktop .....	3-8
Configure Client Machine .....	3-9
Validate the Cluster Profile .....	3-10
Run Parallel Code .....	3-11
<b>Install and Configure MATLAB Parallel Server for Third-Party Schedulers .....</b>	<b>3-12</b>
Activate Your MATLAB Parallel Server License .....	3-12
Get the Installation Files .....	3-12
Install License Manager .....	3-13
Install Software on Compute Nodes .....	3-14
Install Software on Local Desktop .....	3-14
Configure Your Cluster .....	3-14
<b>Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager .....</b>	<b>3-16</b>
Activate Your MATLAB Parallel Server License .....	3-16
Get the Installation Files .....	3-16
Install Software on the Head Node .....	3-17
Install Software on Compute Nodes .....	3-18
Configure the MATLAB Job Scheduler .....	3-18
Connect the MATLAB Client to the MATLAB Parallel Server Cluster .....	3-20
<b>Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Online Licensing .....</b>	<b>3-23</b>
Check License Type and Users .....	3-23
Get the Installation Files .....	3-24
Install Software on All Nodes .....	3-25
Configure the MATLAB Job Scheduler with Admin Center .....	3-25
Connect the MATLAB Client to the MATLAB Parallel Server Cluster .....	3-27
<b>Configure Advanced Options for MATLAB Job Scheduler Integration ..</b>	<b>3-29</b>
Run Multiple MATLAB Parallel Server Versions .....	3-29

Set Up Windows Cluster Hosts .....	3-30
Open Required Ports on Server .....	3-30
Stop mjs Services of Old Installation .....	3-31
Set the MATLAB Job Scheduler Security Level .....	3-32
Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line) .....	3-32
Install the mjs Service to Start Automatically at Boot Time (UNIX) .....	3-35
Validate Installation with MATLAB Job Scheduler .....	3-37
<b>Configure for HPC Pack .....</b>	<b>3-38</b>
Configure Cluster for Microsoft HPC Pack .....	3-38
Configure Client Computer for HPC Pack .....	3-38
Validate Installation Using Microsoft HPC Pack .....	3-39
<b>Configure for Slurm, PBS Pro, OpenPBS, LSF, TORQUE .....</b>	<b>3-42</b>
Create Cluster Profile and Validate Installation .....	3-42
Configure LSF Scheduler on Windows Cluster .....	3-44
Configure Windows Firewalls on Server .....	3-45
<b>Configure Using the Generic Scheduler Interface .....</b>	<b>3-46</b>
Interface with Third-Party Schedulers .....	3-46
Create a Generic Cluster Profile .....	3-47
Special Configurations .....	3-52
<b>Distribute a Generic Cluster Profile and Plugin Scripts .....</b>	<b>3-54</b>
Decide How Users Access the Plugin Scripts .....	3-54
Shared PluginScriptsLocation Folder .....	3-54
Distribute Copies of the PluginScriptsLocation Folder .....	3-55
Further Considerations for Clusters with Shared File Systems .....	3-55
<b>Customize Behavior of Sample Plugin Scripts .....</b>	<b>3-56</b>
Set Additional Properties .....	3-59
<b>Configure a Hadoop Cluster .....</b>	<b>3-62</b>
Cluster Configuration .....	3-62
Client Configuration .....	3-62
Kerberos Authentication .....	3-62
Hadoop Version Support .....	3-63

## Admin Center

# 4

<b>Start Admin Center .....</b>	<b>4-2</b>
<b>Set Up Resources .....</b>	<b>4-3</b>
Add Hosts .....	4-3
Start mjs Service .....	4-4
Start a MATLAB Job Scheduler .....	4-5
Start Workers .....	4-6
Stop, Destroy, Resume, Restart Processes .....	4-7
Move a Worker .....	4-7
Update the Display .....	4-8

<b>Test Connectivity</b> .....	<b>4-9</b>
<b>Export and Import Sessions</b> .....	<b>4-12</b>
<b>Prepare for Cluster Profiles</b> .....	<b>4-13</b>

## Control Scripts

### 5

## Reference Architectures

### 6

<b>Run MATLAB Parallel Server on Amazon Web Services</b> .....	<b>6-2</b>
Requirements .....	<b>6-2</b>
Run from GitHub .....	<b>6-2</b>
<b>Run MATLAB Parallel Server on AWS Batch</b> .....	<b>6-4</b>
Requirements .....	<b>6-4</b>
Run from GitHub .....	<b>6-4</b>
<b>Run MATLAB Parallel Server on Microsoft Azure Using Reference Architecture</b> .....	<b>6-5</b>
Requirements .....	<b>6-5</b>
Run from GitHub .....	<b>6-5</b>
<b>Run MATLAB Parallel Server from Microsoft Azure Marketplace</b> .....	<b>6-7</b>
Requirements .....	<b>6-7</b>
Deploy MATLAB Parallel Server Resources in Azure .....	<b>6-8</b>
Delete MATLAB Parallel Server Resources from Azure .....	<b>6-13</b>
Architecture and Resources for MATLAB Parallel Server in Azure .....	<b>6-14</b>
Marketplace vs. Reference Architecture .....	<b>6-15</b>
Support .....	<b>6-15</b>
<b>Configure MATLAB Parallel Server Licensing for Cloud Platforms</b> .....	<b>6-17</b>
Before You Begin .....	<b>6-17</b>
Select Licensing Configuration Option .....	<b>6-17</b>
Change License Manager Type .....	<b>6-18</b>
Add End Users to License .....	<b>6-18</b>
Disable Network License Manager for MATLAB Parallel Server .....	<b>6-19</b>



# Introduction

---

- “MATLAB Parallel Server Product Description” on page 1-2
- “Product Overview” on page 1-3

## **MATLAB Parallel Server Product Description**

### **Perform MATLAB and Simulink computations on clusters and clouds**

MATLAB Parallel Server lets you scale MATLAB programs and Simulink® simulations to clusters and clouds. You can prototype your programs and simulations on the desktop and then run them on clusters and clouds without recoding. MATLAB Parallel Server supports batch jobs, interactive parallel computations, and distributed computations with large matrices.

All cluster-side licensing is handled by MATLAB Parallel Server. Your desktop license profile is dynamically enabled on the cluster, so you do not need to supply MATLAB licenses for the cluster. The licensing model includes features to support unlimited scaling.

MATLAB Parallel Server runs your programs and simulations as scheduled applications on your cluster. You can use the MATLAB optimized scheduler provided with MATLAB Parallel Server or your own scheduler. A plugin framework enables direct communication with popular cluster scheduler submission clients.

## Product Overview

### In this section...

“Parallel Computing Concepts” on page 1-3

“Determining Product Installation and Versions” on page 1-3


## Parallel Computing Concepts

A *job* is some large operation that you need to perform in your MATLAB session. A job is broken down into segments called *tasks*. You decide how best to divide your job into tasks. You could divide your job into identical tasks, but tasks do not have to be identical.

The MATLAB session in which the job and its tasks are defined is called the *client* session. Often, this is on the machine where you program MATLAB. The client uses Parallel Computing Toolbox software to perform the definition of jobs and tasks. The MATLAB Parallel Server product performs the execution of your job by evaluating each of its tasks and returning the result to your client session.

Parallel Computing Toolbox software allows you to run a cluster of MATLAB workers on your local machine in addition to your MATLAB client session. MATLAB Parallel Server software allows you to run as many MATLAB workers on a remote cluster of computers as your licensing allows.

The MATLAB Job Scheduler is the part of the server software that coordinates the execution of jobs and the evaluation of their tasks. The MATLAB Job Scheduler distributes the tasks for evaluation to the server's individual MATLAB sessions called *workers*. Use of the MATLAB Job Scheduler is optional; the distribution of tasks to workers can also be performed by a third-party scheduler, such as Windows HPC Server (including CCS), a Platform LSF<sup>®</sup> scheduler, or a PBS Pro<sup>®</sup> scheduler.



### Basic Parallel Computing Configuration

## Determining Product Installation and Versions

To determine if Parallel Computing Toolbox software is installed on your system, type this command at the MATLAB prompt:

`ver`

When you enter this command, MATLAB displays information about the version of MATLAB you are running, including a list of all toolboxes installed on your system and their version numbers.

You can run the `ver` command as part of a task in a distributed or parallel application to determine what version of MATLAB Parallel Server software is installed on a worker machine. Note that the toolbox and server software must be the same version.

# Network Administration

---

This chapter provides information useful for network administration of Parallel Computing Toolbox software and MATLAB Parallel Server software.

- “Requirements and Ports for MATLAB Parallel Server” on page 2-2
- “Use Different MPI Builds on UNIX Systems” on page 2-4
- “Shut Down a Job Manager Cluster” on page 2-6
- “Customize Startup Parameters” on page 2-9
- “Access Service Record Files” on page 2-12
- “Set MATLAB Job Scheduler Cluster Security” on page 2-14
- “Troubleshoot Common Problems” on page 2-18

## Requirements and Ports for MATLAB Parallel Server

### In this section...

“Port Configuration” on page 2-2  
 “Fully Qualified Domain Names” on page 2-2  
 “Security Considerations” on page 2-2

This section discusses the requirements and port configuration for your network to support parallel computing.

### Port Configuration

Before you can use MATLAB Parallel Server, you must configure certain required ports. For details, see “Required Ports” on page 2-19. If you need more help during configuration, see this information from MathWorks® Support Team on MATLAB Answers: MATLAB Job Scheduler, or Third-Party Scheduler.

### Fully Qualified Domain Names

MATLAB Parallel Server software and Parallel Computing Toolbox software support both short host names and fully qualified domain names. The default usage is short host names. Check the following considerations depending on your scheduler type:

Scheduler	Consideration
MATLAB Job Scheduler	<ul style="list-style-type: none"> <li>If your network requires fully qualified host names, you can use the <code>mjs_def</code> file to identify the worker nodes by their full names. See “Customize Startup Parameters” on page 2-9.</li> <li>To set the host name used for a MATLAB client session, see the <code>pctconfig</code> reference page.</li> </ul>
Third-Party scheduler	<ul style="list-style-type: none"> <li>To set the host name used for a MATLAB client session, see the <code>pctconfig</code> reference page.</li> </ul>

### Security Considerations

Check the following table for security considerations when using MATLAB Parallel Server:

Scheduler	Security Consideration
MATLAB Job Scheduler	<ul style="list-style-type: none"> <li>• MATLAB workers run as whatever user the administrator starts the node's mjs service under. By default, the mjs service starts as <code>root</code> on UNIX® operating systems, and as <code>LocalSystem</code> on Microsoft® Windows® operating systems. Because MATLAB provides system calls, users can submit jobs that execute shell commands. If you want to run tasks as the user that submitted the job, use security level 3. For more information, see “Set the Security Level” on page 2-14.</li> <li>• By default, the mjs service does not enforce any access control or authentication. Anyone with local or remote access to the mjs services can start and stop their workers and job managers, and query for their status. For authentication and access control options, see “Set MATLAB Job Scheduler Cluster Security” on page 2-14.</li> <li>• The job manager does not restrict access to the cluster, nor to job and task data. For information on security options, see “Set MATLAB Job Scheduler Cluster Security” on page 2-14. Using a third-party scheduler instead of the MathWorks job manager could allow you to take advantage of the security measures it provides.</li> <li>• The parallel computing processes must all be on the same side of a firewall, or you must take measures to enable them to communicate with each other through the firewall. Workers running tasks of the same communicating job cannot be firewalled off from each other, because their MPI-based communication will not work.</li> <li>• If certain ports are restricted, you can specify the ports used for parallel computing. See “Define Script Defaults” on page 2-9.</li> <li>• If your organization is a member of the Internet Multicast Backbone (MBone), make sure that your parallel computing cluster is isolated from MBone access if you are using multicast for parallel computing. Isolation is generally the default condition. If you have any questions about MBone membership, contact your network administrator.</li> </ul>
Third-Party scheduler	<ul style="list-style-type: none"> <li>• The parallel computing processes must all be on the same side of a firewall, or you must take measures to enable them to communicate with each other through the firewall. Workers running tasks of the same communicating job cannot be firewalled off from each other, because their MPI-based communication will not work.</li> <li>• If users will schedule communicating jobs on the cluster, each cluster host participating in communicating jobs requires unchallenged SSH. For example, to enable unchallenged SSH the cluster admin can set up host based authentication/host based authorization for approved cluster users.</li> </ul>

## See Also

“Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16 | “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12

## Use Different MPI Builds on UNIX Systems

### In this section...

“Build MPI” on page 2-4

“Use Your MPI Build” on page 2-4

### Build MPI

On Linux® and Macintosh operating systems, you can use an MPI build that differs from the one provided with Parallel Computing Toolbox. This topic outlines the steps for creating an MPI build for use with the generic scheduler interface. If you already have an alternative MPI build, proceed to “Use Your MPI Build” on page 2-4.

- 1 Unpack the MPI sources into the target file system on your machine. For example, suppose you have downloaded `mpich2-distro.tgz` and want to unpack it into `/opt` for building:

```
# cd /opt
# mkdir mpich2 && cd mpich2
# tar zxvf path/to/mpich2-distro.tgz
# cd mpich2-1.4.1p1
```

- 2 Build your MPI using the `enable-shared` option (this is vital, as you must build a shared library MPI, binary compatible with `MPICH2-1.4.1p1` for R2013b to R2018b, or `MPICH3.2.1` for R2019a and later). For example, the following commands build an MPI with the `nemesis` channel device and the `gforker` launcher.

```
#!/configure -prefix=/opt/mpich2/mpich2-1.4.1p1 \
--enable-shared --with-device=ch3:nemesis \
--with-pm=gforker 2>&1 | tee log
# make 2>&1 | tee -a log
# make install 2>&1 | tee -a log
```

### Use Your MPI Build

When your MPI build is ready, this stage highlights the steps to use it with a generic scheduler. To get your cluster working with a different MPI build, follow these steps.

- 1 Test your build by running the `mpiexec` executable. The build should be ready to test if its `bin/mpiexec` and `lib/libmpich.so` are available in the MPI installation location.

Following the example in “Build MPI” on page 2-4, `/opt/mpich2/mpich2-1.4.1p1/bin/mpiexec` and `/opt/mpich2/mpich2-1.4.1p1/lib/libmpich.so` are ready to use, so you can test the build with:

```
$ /opt/mpich2/mpich2-1.4.1p1/bin/mpiexec -n 4 hostname
```

- 2 Create an `mpiLibConf` function to direct Parallel Computing Toolbox to use your new MPI. Write your `mpiLibConf.m` to return the appropriate information for your build. For example:

```
function [primary, extras] = mpiLibConf
primary = '/opt/mpich2/mpich2-1.4.1p1/lib/libmpich.so';
extras = {};
```

The primary path *must* be valid on the cluster; and your `mpiLibConf.m` file must be higher on the cluster workers’ path than `matlabroot/toolbox/parallel/mpi`. (Sending



`mpiLibConf.m` as an attached file for this purpose does not work. You can get the `mpiLibConf.m` function on the worker path by either moving the file into a folder on the path, or by having the scheduler use `cd` in its command so that it starts the MATLAB worker from within the folder that contains the function.)

- 3** Determine necessary daemons and command-line options.
  - Determine all necessary daemons (often something like `mpdboot` or `smpd`). The `gforker` build example in this section uses an MPI that needs no services or daemons running on the cluster, but it can use only the local machine.
  - Determine the correct command-line options to pass to `mpiexec`.
- 4** To set up your cluster to use your new MPI build, modify your communicating job wrapper script to pick up the correct `mpiexec`. Additionally, there might be a stage in the wrapper script where the MPI process manager daemons are launched.

The communicating job wrapper script must:

- Determine which nodes are allocated by the scheduler.
- Start required daemon processes. For example, for the MPD process manager this means calling "`mpdboot -f <nodefile>`".
- Define which `mpiexec` executable to use for starting workers.
- Stop the daemon processes. For example, for the MPD process manager this means calling "`mpdallexit`".

For examples of communicating job wrapper scripts, see “Sample Plugin Scripts” (Parallel Computing Toolbox).

## Shut Down a Job Manager Cluster

### In this section...

“Linux and Macintosh Operating Systems” on page 2-6

“Microsoft Windows Operating Systems” on page 2-7

If you are done using the job manager and its workers, you might want to shut down the server software processes so that they are not consuming network resources. You do not need to be at the computer running the processes that you are shutting down. You can run these commands from any machine with network access to the processes. The following sections explain shutting down the processes for different platforms.

### Linux and Macintosh Operating Systems

Enter the commands of this section at the prompt in a shell.

#### Stopping the Job Manager and Workers

- 1 To shut down the job manager, enter the commands

```
cd matlabroot/toolbox/parallel/bin
```

(Enter the following command on a single line.)

```
stopjobmanager -remotehost <job manager hostname> -name  
<MyJobManager> -v
```

If you have more than one job manager running, stop each of them individually by host and name.

For a list of all options to the script, type

```
stopjobmanager -help
```

- 2 For each MATLAB worker you want to shut down, enter the commands

```
cd matlabroot/toolbox/parallel/bin  
stopworker -remotehost <worker hostname> -v
```

If you have more than one worker session running, you can stop each of them individually by host and name.

```
stopworker -name worker1 -remotehost <worker hostname>  
stopworker -name worker2 -remotehost <worker hostname>
```

For a list of all options to the script, type

```
stopworker -help
```

#### Stop and Uninstall the mjs Daemon

Normally, you configure the mjs daemon to start at system boot time and continue running until the machine shuts down. However, if you plan to uninstall the MATLAB Parallel Server product from a machine, you might want to uninstall the mjs daemon also, because you no longer need it.

---

**Note** You must have root privileges to stop or uninstall the mjs daemon.

---

- 1 Use the following command to stop the mjs daemon:

```
/etc/init.d/mjs stop
```

- 2 Remove the installed link to prevent the daemon from starting up again at system reboot:

```
cd /etc/init.d/
rm mjs
```

### Stop the Daemon Manually

If you used the alternative manual startup of the mjs daemon, use the following commands to stop it manually:

```
cd matlabroot/toolbox/parallel/bin
mjs stop
```

## Microsoft Windows Operating Systems

### Stop the Job Manager and Workers

Enter the commands of this section at the prompt in a DOS command window.

- 1 To shut down the job manager, enter the commands

```
cd matlabroot\toolbox\parallel\bin
```

(Enter the following command on a single line.)

```
stopjobmanager -remotehost <job manager hostname> -name
<MyJobManager> -v
```

If you have more than one job manager running, stop each of them individually by host and name.

For a list of all options to the script, type

```
stopjobmanager -help
```

- 2 For each MATLAB worker you want to shut down, enter the commands

```
cd matlabroot\toolbox\parallel\bin
stopworker -remotehost <worker hostname> -name <worker name> -v
```

If you have more than one worker session running, you can stop each of them individually by host and name.

```
stopworker -remotehost <worker hostname> -name <worker1 name>
stopworker -remotehost <worker hostname> -name <worker2 name>
```

For a list of all options to the script, type

```
stopworker -help
```

### **Stop and Uninstall the mjs Service**

Normally, you configure the mjs service to start at system boot time and continue running until the machine shuts down. If you need to stop the mjs service while leaving the machine on, enter the following commands at a DOS command prompt:

```
cd matlabroot\toolbox\parallel\bin  
mjs stop
```

If you plan to uninstall the MATLAB Parallel Server product from a machine, you might want to uninstall the mjs service also, because you no longer need it.

You do not need to stop the service before uninstalling it.

To uninstall the mjs service, enter the following commands at a DOS command prompt:

```
cd matlabroot\toolbox\parallel\bin  
mjs uninstall
```

## Customize Startup Parameters

### In this section...

“Define Script Defaults” on page 2-9

“Override Script Defaults” on page 2-10

“Increase Heap Memory” on page 2-11

The MATLAB Parallel Server scripts run using several default parameters. You can customize the scripts, as described in this section.

### Define Script Defaults

The scripts for the server services require values for several parameters. These parameters set the process name, the user name, log file location, ports, etc. Some of these can be set using flags on the command lines, but the full set of user-configurable parameters are in the `mjs_def` file.

---

**Note** The startup script flags take precedence over the settings in the `mjs_def` file.

---

The default parameters used by the server service scripts are defined in the file:

- `matlabroot\toolbox\parallel\bin\mjs_def.bat` (on Microsoft Windows operating systems)
- `matlabroot/toolbox/parallel/bin/mjs_def.sh` (on Linux or Macintosh operating systems)

To set the default parameters, edit this file before installing or starting the `mjs` service.

The `mjs_def` file is self-documented, and includes explanations of all its parameters.

---

**Note** If you want to run more than one job manager on the same machine, they must all have unique names. Specify the names using flags with the startup commands.

---

### Set the User

By default, the job manager and worker services run as the user who starts them. You can run the services as a different user with the following settings in the `mjs_def` file.

Parameter	Description
MJSUSER	Set this parameter to run the <code>mjs</code> services as a user different from the user who starts the service. On a UNIX operating system, set the value before starting the service; on a Windows operating system, set it before installing the service.
MJSPASS	On a Windows operating system, set this parameter to specify the password for the user identified in the <code>MJSUSER</code> parameter; otherwise, the system prompts you for the password when the service is installed.

On UNIX operating systems, `MJSUSER` requires that the current machine has the `sudo` utility installed, and that the current user be allowed to use `sudo` to execute commands as the user

identified by MJSUSER. For further information, refer to your system documentation on the `sudo` and `sudoers` utilities (for example, `man sudo` and `man sudoers`).

The MJSUSER is granted these permissions on Windows systems:

Privilege	Purpose	Local Security Settings Policy
SeServiceLogonRight	Required to log on using the service logon type.	Log on as a service
SeAssignPrimaryTokenPrivilege	Required to start a process under a different user account.	Replace a process level token
SeIncreaseQuotaPrivilege	Required to start a process under a different user account.	Adjust memory quotas for a process

To modify or remove these privileges,

- 1 Select the Windows menu **Start > Settings > Control Panel**.
- 2 Double-click **Administrative Tools**, then **Local Security Policy**.
- 3 In the tree, select **Local Policies > User Rights Assignment**.

The table above indicates which policies are affected by MJSUSER. Double-click any of the listed policies in the Local Security Settings GUI to alter its setting or remove a user from that policy.

## Override Script Defaults

### Specify an Alternative Defaults File

The default parameters used by the `mjs` service, job managers, and workers are defined in the file:

- `matlabroot\toolbox\parallel\bin\mjs_def.bat` (on Windows operating systems)
- `matlabroot/toolbox/parallel/bin/mjs_def.sh` (on Linux or Macintosh operating systems)

Before installing and starting the `mjs` service, you can edit this file to set the default parameters with values you require.

Alternatively, you can make a copy of this file, modify the copy, and specify that this copy be used for the default parameters.

On Linux or Macintosh operating systems, enter the command

```
mjs start -mjsdef my_mjs_def.sh
```

On Windows operating systems, enter the command

```
mjs install -mjsdef my_mjs_def.bat
mjs start -mjsdef my_mjs_def.bat
```

If you specify a new `mjs_def` file instead of the default file for the service on one computer, the new file is not automatically used by the `mjs` service on other computers. If you want to use the same alternative file for all your `mjs` services, you must specify it for each `mjs` service you install or start.

For more information, see “Define Script Defaults” on page 2-9.

---

**Note** The startup script flags take precedence over the settings in the `mjs_def` file.

---

### Start in a Clean State

When a job manager or worker starts up, it normally resumes its session from the past. This way, a job queue is not destroyed or lost if the job manager machine crashes or if the job manager is inadvertently shut down. To start up a job manager or worker from a clean state, with all history deleted, use the `-clean` flag on the start command:

```
startjobmanager -clean -name MyJobManager
startworker -clean -jobmanager MyJobManager
```

### Increase Heap Memory

An mjs service can use up to 4000 workers across a collection of nodes. When you scale up the number of workers or tasks in your cluster, you must increase the heap memory available to the job manager. To do so, set the `JOB_MANAGER_MAXIMUM_MEMORY` parameter in the `mjs_def` file based on the following recommendations.

- Use a minimum value of 1000m (1000 MiB).
- Use 1000m for every 1000 workers.
- Add 1000m for every 100,000 tasks expected to be queued at peak load.

For example, for a cluster with 4000 workers and a peak queue size of 200,000 tasks, set `JOB_MANAGER_MAXIMUM_MEMORY` to 6000m.

## Access Service Record Files

### In this section...

“Locate Log Files” on page 2-12

“Locate Checkpoint Folders” on page 2-12

The MATLAB Parallel Server services generate various record files in the normal course of their operations. The mjs service, job manager, and worker sessions all generate such files. This section describes the types of information stored by the services.

### Locate Log Files

Log files for each service contain entries for the service’s operations. These might be of particular interest to the network administrator in cases when problems arise.

Operating System	File Location
Windows	<p>The default location of the log files is &lt;TEMP&gt;\MJS\Log, where &lt;TEMP&gt; is the value of the system TEMP variable. For example, if TEMP is set to C:\TEMP, the log files are placed in C:\TEMP\MJS\Log.</p> <p>You can set alternative locations for the log files by modifying the LOGBASE setting in the mjs_def.bat file before starting the mjs service.</p>
Linux and Macintosh	<p>The default location of the log files is /var/log/mjs/.</p> <p>You can set alternative locations for the log files by modifying the LOGBASE setting in the mjs_def.sh file before starting the mjs service.</p>

### Locate Checkpoint Folders

Checkpoint folders contain information related to persistence data, which the server services use to create continuity from one instance of a session to another. For example, if you stop and restart a job manager, the new session continues the old session, using all the same data.

A primary feature offered by the checkpoint folders is in crash recovery. This allows server services to automatically resume their sessions after a system goes down and comes back up, minimizing the loss of data. However, if a MATLAB worker goes down during the evaluation of a task, that task is neither reevaluated nor reassigned to another worker. In this case, a finished job may not have a complete set of output data, because data from any unfinished tasks might be missing.

---

**Note** If a job manager crashes and restarts, its workers can take up to 2 minutes to reregister with it.

---



<b>Platform</b>	<b>File Location</b>
Windows	<p>The default location of the checkpoint folders is &lt;TEMP&gt;\MJS\Checkpoint, where &lt;TEMP&gt; is the value of the system TEMP variable. For example, if TEMP is set to C:\TEMP, the checkpoint folders are placed in C:\TEMP\MJS\Checkpoint.</p> <p>You can set alternative locations for the checkpoint folders by modifying the CHECKPOINTBASE setting in the mjs_def.bat file before starting the mjs service.</p>
Linux and Macintosh	<p>The checkpoint folders are placed by default in /var/lib/mjs/.</p> <p>You can set alternative locations for the checkpoint folder by modifying the CHECKPOINTBASE setting in the mjs_def.sh file before starting the mjs service.</p>

## Set MATLAB Job Scheduler Cluster Security

### In this section...

“Set the Security Level” on page 2-14

“Local, MATLAB Job Scheduler, and Network Passwords” on page 2-15

“Authorize Users for Job and Task Access” on page 2-16

“Set Secure Communication” on page 2-16

### Set the Security Level

You set the MATLAB Job Scheduler security level with the `SECURITY_LEVEL` parameter in the `mjs_def` file before starting the `mjs` service on your cluster nodes. The `mjs_def` file indicates what values are allowed, and briefly describes each security level.

The following table describes the available security levels for accessing a MATLAB Job Scheduler and its jobs.

Security Level	Description	User Requirements
0	<p>No security.</p> <ul style="list-style-type: none"> <li>Any user can access any job.</li> <li>Tasks run as the user who started the <code>mjs</code> process on the worker machines (typically <code>root</code> or <code>Local System</code>).</li> <li>This is the default, and is the behavior in all releases prior to R2010b.</li> </ul>	<ul style="list-style-type: none"> <li>Jobs are associated with the default user name of the programmer, but no protection is provided.</li> </ul>
1	<p>Jobs are identified with the submitting user.</p> <ul style="list-style-type: none"> <li>Any user can access any job; a dialog warns if the accessed job belongs to another user.</li> <li>Tasks run as the user who started the <code>mjs</code> process on the worker machines (typically <code>root</code> or <code>Local System</code>).</li> </ul>	<ul style="list-style-type: none"> <li>A dialog requires you to establish a user name when you first access the job manager.</li> <li>Your MATLAB Job Scheduler user name does not have to match your system/network user name.</li> <li>No passwords are used.</li> </ul>
2	<p>Job manager MATLAB Job Scheduler password protection on jobs.</p> <ul style="list-style-type: none"> <li>Jobs and tasks are identified with the submitting user, and are password protected. Other unauthorized users cannot access your jobs.</li> <li>Tasks run as the user who started the <code>mjs</code> process on the worker machines (typically <code>root</code> or <code>Local System</code>).</li> </ul>	<ul style="list-style-type: none"> <li>When you start the MATLAB Job Scheduler, it prompts you to provide a new password for that job manager's admin account, which can be used for accessing all users' jobs and tasks.</li> <li>A dialog box requires you to establish a user name and password when you first access the MATLAB Job Scheduler from the MATLAB client.</li> <li>Your MATLAB Job Scheduler user name and password do not have to match your system/network user name and password.</li> </ul>

Security Level	Description	User Requirements
3	<p>In addition to the security of level 2, tasks run as the submitting user on worker machines.</p> <ul style="list-style-type: none"> <li>Jobs and tasks are identified with the submitting user, and are password protected. Other unauthorized users cannot access your jobs.</li> <li>Tasks run as the user who submitted the job.</li> </ul>	<ul style="list-style-type: none"> <li>The MATLAB Job Scheduler must use secure communication with the workers (set in the <code>mjs_def</code> file).</li> <li>When you start the MATLAB Job Scheduler, it prompts you to provide a new password for that job manager's admin account, which can be used for accessing all users' jobs and tasks.</li> <li>A dialog box requires you to establish a user name and password when you first access the MATLAB Job Scheduler from the MATLAB client.</li> <li>Your job manager MATLAB Job Scheduler user name and password must be the same as your system/network user name and password, because the worker must log you in to run the task as you.</li> <li>All users that tasks run as, require read and write permissions to the <code>CHECKPOINTBASE</code> folder and all its subfolders.</li> <li>On UNIX systems, the <code>mjs</code> process on the cluster nodes must be started by the root user.</li> <li>On Windows systems, the submitting user must have "Allow log on locally" permission enabled on every worker machine to successfully run jobs on the cluster. If this permission is disabled, all submitted jobs will fail when started by the cluster.</li> </ul> <p>To enable this permission, change the <code>SeInteractiveLogonRight</code> constant for the user in the User Rights Assignment security policy settings for each machine in the cluster.</p>

The job manager and the workers should run at the same security level. A worker running at too low a security level will fail to register with the job manager, because the job manager does not trust it.

## Local, MATLAB Job Scheduler, and Network Passwords

For any security above level 0, when you start the MATLAB Job Scheduler (for example, with the `startjobmanager` command), a cluster user account named `admin` is created for this cluster, and you are prompted to provide a password for this new account. The cluster `admin` account has all the necessary permissions for accessing the cluster and all its jobs.

For any security level, the MATLAB Job Scheduler identifies every job with the user who submits the job. Therefore, whenever you access the MATLAB Job Scheduler or a job, the MATLAB Job Scheduler must be aware of who you are.

At security level 0, the MATLAB Job Scheduler and job objects' `UserName` property is set to the login name of the person who creates the job; this setting can be changed at any time. For all higher security levels, the first access to the MATLAB Job Scheduler causes a dialog box to open which asks for your username; if the security level is 2 or 3, you must also provide a password. The username and password you provide for the MATLAB Job Scheduler needs to match your network username and password *only* if you are using security level 3; otherwise, you can create a new username and password unique for the MATLAB Job Scheduler. For your convenience, you can choose how long to save your username and password on the local computer, so that you do not need to enter them every time you access your job.

For information about changing a password and logging out of a MATLAB Job Scheduler, see `changePassword` and `logout`.

## Authorize Users for Job and Task Access

This example shows how to authorize users to access your job on a MATLAB Job Scheduler cluster with security level 2 or 3. For security level 2 and 3, when you create and submit a job to a MATLAB Job Scheduler cluster, jobs and tasks are identified with the submitting user. These jobs and tasks are password protected, and therefore other unauthorized users cannot access your jobs.

Use `parcluster` to create a cluster object using the cluster profile 'MyMJSCluster'. Replace 'MyMJSCluster' with the name of your cluster profile. Then, use `batch` to create and submit a job on the cluster.

```
c = parcluster('MyMJSCluster');  
j = batch(c,@rand,1,{2});
```

You can use the `AuthorizedUsers` of a job to authorize user access to a job and the job's and tasks. Each user that you specify must have previously interacted with MATLAB Job Scheduler cluster. To authorize access to a job for users "user1" and "user2", set the `AuthorizedUsers` of the job to ["user1", "user2"].

```
j.AuthorizedUsers = ["user1", "user2"];
```

## Set Secure Communication

To establish secure encrypted communication between MATLAB Job Scheduler, client and workers, set

- `USE_SECURE_COMMUNICATION = true`
- `ALL_SERVER_SOCKETS_IN_CLUSTER = true` (default)

in the `mjs_def` file. Secure encrypted communication is provided via `SSLSocket` using TLSv1.2 only.

---

**Note** If `ALL_SERVER_SOCKETS_IN_CLUSTER = false` in the `mjs_def` file, then secure encrypted communication is established between MATLAB Job Scheduler and workers only.

---

You must also provide a value for the `SHARED_SECRET_FILE` parameter in the `mjs_def` file, identifying where the file can be found from the MATLAB Job Scheduler perspective. To create this file, run either script:

- `matlabroot/toolbox/parallel/bin/createSharedSecret` (UNIX)
- `matlabroot\toolbox\parallel\bin\createSharedSecret.bat` (Windows)

The secret file establishes trust between the processes on different machines.

- In a shared file system, all the nodes can point to the same secret file, and they can even all share the same `mjs_def` file.
- In a nonshared file system, create a secret file with the provided script, then copy the file to each node and make sure each node's `mjs_def` file indicates where its particular secret file is located.

---

**Note** Secure communication is required when using MATLAB Job Scheduler security level 3.

---

## Troubleshoot Common Problems

### In this section...

“License Errors” on page 2-18

“Memory Errors on UNIX Operating Systems” on page 2-19

“Run Server Processes on Windows Network Installation” on page 2-19

“Required Ports” on page 2-19

“Ephemeral TCP Ports with Job Manager” on page 2-20

“Host Communications Problems” on page 2-20

“Verify Network Communications for Cluster Discovery” on page 2-21

This section offers advice on solving problems you might encounter with MATLAB Parallel Server software.

### License Errors

When starting a MATLAB worker, a licensing problem might result in the message

```
License checkout failed. No such FEATURE exists.
License Manager Error -5
```

There are many reasons why you might receive this error:

- This message usually indicates that you are trying to use a product for which you are not licensed. Look at your `license.dat` file located within your MATLAB installation to see if you are licensed to use this product.
- If you are licensed for this product, this error may be the result of having extra carriage returns or tabs in your license file. To avoid this, ensure that each line begins with either `#`, `SERVER`, `DAEMON`, or `INCREMENT`.

After fixing your `license.dat` file, restart the network license manager and MATLAB should work properly.

- This error may also be the result of an incorrect system date. If your system date is before the date that your license was made, you will get this error.
- If you receive this error when starting a worker with MATLAB Parallel Server software:
  - You may be calling the `startworker` command from an installation that does not have access to a worker license. For example, starting a worker from a client installation of the Parallel Computing Toolbox product causes the following error:

```
The mjs service on the host hostname
returned the following error:
```

```
Problem starting the MATLAB worker.
```

```
The cause of this problem is:
```

```
=====
Most likely, the MATLAB worker failed to start due to a
licensing problem, or MATLAB crashed during startup. Check
```

```

the worker log file
/tmp/mjs_user/node_node_worker_05-11-01_16-52-03_953.log
for more detailed information. The mjs log file
/tmp/mjs_user/mjs-service.log
may also contain some additional information.

```

=====

In the worker log files, you see the following information:

```

License checkout failed.
License Manager Error -15
MATLAB is unable to connect to the license server.
Check that the license manager has been started, and that the
MATLAB client machine can communicate with the license server.

```

```

Troubleshoot this issue by visiting:
https://www.mathworks.com/support/lme/R2009a/15

```

```

Diagnostic Information:
Feature: MATLAB_Distrib_Comp_Engine
License path: /apps/matlab/etc/license.dat
FLEXnet Licensing error: -15,570. System Error: 115

```

- If you installed only the Parallel Computing Toolbox product, and you are attempting to run a worker on the same machine, you will receive this error because the MATLAB Parallel Server product is not installed, and therefore the worker cannot obtain a license.

## Memory Errors on UNIX Operating Systems

If the number of processes created by the server services on a machine running a UNIX operating system (Linux or Macintosh) exceeds the operating system limits, the services fail and generate an out-of-memory error. It is recommended that you adjust your system limits. For more information, see “Recommended System Limits for Macintosh and Linux” (Parallel Computing Toolbox).

## Run Server Processes on Windows Network Installation

Many networks are configured not to allow LocalSystem to have access to UNC or mapped network shares. In this case, run the mjs process under a different user with rights to log on as a service. See “Set the User” on page 2-9.

## Required Ports

### With Job Manager

#### BASE\_PORT

The `mjs_def` file specifies and describes the ports required by the job manager and all workers. See the following file in the MATLAB installation used for each cluster process:

- `matlabroot/toolbox/parallel/bin/mjs_def.sh` (on UNIX operating systems)
- `matlabroot\toolbox\parallel\bin\mjs_def.bat` (on Windows operating systems)

### Communicating Jobs

On worker machines running a UNIX operating system, the number of ports required by MPICH for the running of communicating jobs ranges from `BASE_PORT + 1000` to `BASE_PORT + 2000`.

### With Third-Party Scheduler

Before the worker processes start, you can control the range of ports used by the workers for communicating jobs by defining the environment variable `MPICH_PORT_RANGE` with the value `minport:maxport`.

### Client Ports

With the `pctconfig` function, you specify the ports used by the client. If the default ports cannot be used, this function allows you to configure ports separately for communication with the job scheduler and communication with a parallel pool.

## Ephemeral TCP Ports with Job Manager

If you use the job manager on a cluster of nodes running Windows operating systems, you must make sure that a large number of ephemeral TCP ports are available on the job manager machine. By default, the maximum valid ephemeral TCP port number on a Windows operating system is 5000, but transfers of large data sets might fail if this setting is not increased. In particular, if your cluster has 32 or more workers, you should increase the maximum valid ephemeral TCP port number using the following procedure:

- 1 Start the Registry Editor.
- 2 Locate the following subkey in the registry, and click **Parameters**:  
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters`
- 3 On the Registry Editor window, select **Edit > New > DWORD Value**.
- 4 In the list of entries on the right, change the new value name to `MaxUserPort` and press **Enter**.
- 5 Right-click on the `MaxUserPort` entry name and select **Modify**.
- 6 In the Edit DWORD Value dialog, enter `65534` in the **Value data** field. Select `Decimal` for the **Base** value. Click **OK**.

This parameter controls the maximum port number that is used when a program requests any available user port from the system. Typically, ephemeral (short-lived) ports are allocated between the values of 1024 and 5000 inclusive. This action allows allocation for port numbers up to 65534.

- 7 Quit the Registry Editor.
- 8 Reboot your machine.

## Host Communications Problems

If a worker is not able to make a connection with its MATLAB Job Scheduler, or if a client session cannot validate a profile that uses that scheduler, this might indicate communications problems between nodes.

### With Command-Line Interface

First, be sure that the machines in question agree on their IP resolutions. The IP address for a particular host should be the same for itself as it is from the perspective of another host. For



example, if a process on `hostB` cannot connect to one on `hostA`, find out the `hostA` IP address for itself, then see what the IP address for `hostA` is from `hostB`. They should be the same.

If the machines can identify each other, the `nodestatus` command can be useful for diagnosing problems between their processes. Use the function to determine what MATLAB Parallel Server processes are running on the local host, and which are accessible from remote hosts. If a worker on `hostA` cannot register with its job manager on `hostB`, run `nodestatus` on both hosts to see what each can see on `hostB`.

On `hostB`, execute:

```
nodestatus -remotehost hostB
```

Then on `hostA`, run exactly the same command:

```
nodestatus -remotehost hostB
```

The results should be the same, showing the same listing of job managers and workers.

If the output indicates problems, run the command again with a higher information level to receive more detailed information:

```
nodestatus -remotehost hostB -infolevel 3
```

### With Admin Center GUI

You can diagnose some communications problems using Admin Center.

If you cannot successfully add hosts to the listing by specifying host name, you can use their IP addresses instead (see “Add Hosts” on page 4-3). If you suspect any communications problems, in the Admin Center GUI click **Test Connectivity** (see “Test Connectivity” on page 4-9). This testing verifies that the nodes can identify each other and allow their processes to communicate with each other.

## Verify Network Communications for Cluster Discovery

If you want to use the discover cluster capabilities in Parallel Computing Toolbox, your network must be configured with at least one of the following:

- “DNS SRV Record” on page 2-21
- “Multicast” on page 2-22

### DNS SRV Record

When you use DNS for cluster discovery you require a DNS SRV record for each domain. You can have multiple DNS SRV records for multiple MATLAB Job Schedulers. Use the following general form for each DNS SRV record.

```
_mdcs._tcp.<domain> <TTL> IN SRV <priority> <weight> <port> <hostname>.
```

Construct a DNS SRV record for a MATLAB Job Scheduler server using the following parts.

- `<domain>` is the domain name (like `company.com` or `university.edu`) that the client machine searches.

- `<TTL>` indicates how long (in seconds) the DNS record can be cached. 3600 is recommended.
- `IN SRV` is required as shown, indicating that this is a service record.
- `<priority>` and `<weight>` indicate priority and weight values. If you create multiple DNS SRV records, you can specify their priority with these fields. A value of 0 is recommended for each. The lower `<priority>` is, the *higher* priority the host has. When two records have the same `<priority>`, the record with the highest `<weight>` is used first. Use the `<weight>` value to specify server preference.
- `<port>` is the port on which you connect to the MATLAB Job Scheduler server. The default port is 27350. If you change port for the MATLAB Job Scheduler server, change `<port>` accordingly.
- `<hostname>` is the fully qualified domain name for the host serving the MATLAB Job Scheduler. The machine `mjs-1` on the domain `company.com` has a fully qualified domain name `mjs-1.company.com`.

A valid DNS SRV record for the `company.com` network running a MATLAB Job Scheduler on machine `mjs-1` might look like this:

```
_mdcs._tcp.company.com 3600 IN SRV 0 0 27350 mjs-1.company.com.
```

---

**Note** If multiple domains are required to locate the cluster, use a DNS SRV record for each domain. If the network accessed by users via VPN has different DNS SRV records to your internal network, ensure that a DNS SRV record exists for each domain.

---

Use the standard procedure for your DNS system to create appropriate DNS SRV records. You can use standard utilities such as the `nslookup` command to verify that your network is configured with the necessary DNS SRV records. To examine MATLAB Job Scheduler DNS SRV records for the `company.com` domain, use the following command.

```
nslookup -type=SRV _mdcs._tcp.company.com
```

### Multicast

To use multicast, it is required on the head node running the MATLAB Job Scheduler and on the client system.

Multicast, unlike TCP/IP or UDP, is a subscription-based protocol where a number of machines on a network indicate to the network their interest in particular packets originating somewhere on that network. By contrast, both UDP and TCP packets are always bound for a single machine, usually indicated by its IP address.

The main tools for investigating this type of packet are:

- `tcpdump` for UNIX operating systems
- `winpcap` and `ethereal` for Microsoft Windows operating systems
- A Java<sup>®</sup> class included with the parallel computing products.

The Java class is called `com.mathworks.toolbox.distcomp.mjs.test.MulticastTester`. Both its static main method and its constructor take two input arguments: the multicast group to join and the port number to use.

This Java class has a number of simple methods to attempt to join a specified multicast group. Once the class has successfully joined the group, it has methods to send messages to the group, listen for

messages from the group, and display what it receives. You can use this class both from a command-line call to Java software and inside MATLAB.

From a shell prompt (assuming that `java` is on your path), type

```
java -cp distcomp.jar com.mathworks.toolbox.distcomp.mjs.test.MulticastTester
```

You should see an output something like this:

```
0 : host1name : 0
1 : host2name : 0
```

The following example shows how to use the Java class inside MATLAB.

Start MATLAB on two machines (e.g., `host1name` and `host2name`) for which you want to test multicast. In each MATLAB session, enter the following commands:

```
m = com.mathworks.toolbox.distcomp.mjs.test.MulticastTester('239.1.1.1', 9999);
m.startSendingThread;
m.startListeningThread;
```

These instructions cause each MATLAB session to issue a stream of multicast test packets, and to listen for test packets. If multicast is working between the machines, you see a stream of lines like the following:

```
0 : host1name : 0
1 : host2name : 0
2 : host2name : 1
3 : host2name : 2
```

The number on the left in each character vector is the line number for the received packet. The text in the center is the host from which the packet is received. The number on the right is the packet number sent by the sending host. It is normal for a host to report a test packet from itself.

If either machine does not receive a stream of test packets, or if the remote host is not included in either stream, then multicast communication is not operating properly.

To terminate the test stream, execute the following in both MATLAB sessions:

```
m.stopSendingThread;
m.stopListeningThread;
```



# Product Installation

---

## Choose Solution to Install and Configure MATLAB Parallel Server


Consult the following tables for help with the most common tasks while setting up MATLAB Parallel Server. The tables help you to choose when you want to:

In this section...
“Activate the License for, Install, and Configure MATLAB Parallel Server on the Cluster” on page 3-2
“Configure a MATLAB Client for a User to Submit Jobs to the Cluster” on page 3-3
“Upgrade MATLAB Parallel Server” on page 3-4

### Activate the License for, Install, and Configure MATLAB Parallel Server on the Cluster

Consult the following table if you are getting started and you want to install and configure MATLAB Parallel Server for your cluster.

Scenario	Solution	
You have a Slurm scheduler in your cluster.	Follow “Install and Configure MATLAB Parallel Server for Slurm” on page 3-6.	
You have an existing scheduler in your cluster, such as <ul style="list-style-type: none"> <li>• LSF®</li> <li>• PBS</li> <li>• Torque</li> <li>• Grid Engine</li> <li>• HPC Pack</li> <li>• Hadoop®</li> </ul>	Follow “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12.	
You do not have an existing scheduler in your cluster.	You want to manage licensing of your cluster users on-premise.	Set up the MATLAB Job Scheduler, which comes with MATLAB Parallel Server. Follow “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16.
	You want to use MathWorks online licensing for your cluster users.	Set up the MATLAB Job Scheduler, which comes with MATLAB Parallel Server. Follow “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Online Licensing” on page 3-23.



### Configure a MATLAB Client for a User to Submit Jobs to the Cluster

Consult the following table if you already have a cluster configured and you want to configure a new MATLAB client to submit jobs to the cluster.

Scenario	Solution
<p>The cluster uses a third-party scheduler and you meet all of these conditions:</p> <ul style="list-style-type: none"> <li>• The scheduler has out-of-the-box support in MATLAB. The schedulers with out-of-the-box support are Slurm, LSF, PBS Pro, OpenPBS, Grid Engine, Torque, HPC Pack.</li> <li>• There is a shared file system between the client machine and the cluster nodes.</li> <li>• The client machine has the third-party scheduler submission tools installed. You can check this by executing the relevant commands in a command line. For example, the Slurm submission script is <code>sbatch</code>.</li> </ul>	<p>See “Install Software on Local Desktop” on page 3-14. Then,</p> <ul style="list-style-type: none"> <li>• Slurm, LSF, PBS Pro, OpenPBS, Grid Engine, Torque - If the client is a Windows machine, start from “Configure Windows Firewalls on Server” on page 3-45. Otherwise, start from “Create Cluster Profile and Validate Installation” on page 3-42.</li> <li>• HPC Server - Start from “Configure Client Computer for HPC Pack” on page 3-38.</li> </ul>
<p>The cluster is a Hadoop cluster.</p>	<p>Start from “Client Configuration” on page 3-62.</p>
<p>The cluster uses a third-party scheduler.</p>	<p>Follow “Configure Using the Generic Scheduler Interface” on page 3-46.</p>
<p>The cluster uses the MATLAB Job Scheduler.</p>	<ul style="list-style-type: none"> <li>• If the cluster is using MathWorks online licensing, ask your cluster administrator to add the new user. For more information, see “Add Licensed End Users” on page 3-23. Then, follow “Connect the MATLAB Client to the MATLAB Parallel Server Cluster” on page 3-27.</li> <li>• Otherwise, follow “Connect the MATLAB Client to the MATLAB Parallel Server Cluster” on page 3-20.</li> </ul>

**Note** You can export a cluster profile and import it in another client machine. For more information, see “Import and Export Cluster Profiles” (Parallel Computing Toolbox). This is useful when you are configuring another client machine with the same configuration, to avoid recreating the cluster profile manually.

## Upgrade MATLAB Parallel Server

Consult the following table if you want to configure a new release of MATLAB Parallel Server.

Scenario	Solution
<p>The cluster uses a third-party scheduler.</p>	<p>Follow “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12 to</p> <ul style="list-style-type: none"> <li>• Activate and install the new version of MATLAB Parallel Server on the cluster.</li> <li>• Configure the cluster and client machines.</li> </ul>



<b>Scenario</b>	<b>Solution</b>
The cluster uses the MATLAB Job Scheduler.	<ol style="list-style-type: none"><li data-bbox="865 300 1471 394"><b>1</b> Stop the mjs service of the older release. For instructions, see “Stop mjs Services of Old Installation” on page 3-31.</li><li data-bbox="865 405 1471 716"><b>2</b> Install the new version of MATLAB Parallel Server on the cluster and configure the MATLAB Job Scheduler. For more information, see “Run Multiple MATLAB Parallel Server Versions” on page 3-29. For help with activation, installation, and configuration, consult the table in “Activate the License for, Install, and Configure MATLAB Parallel Server on the Cluster” on page 3-2.</li></ol>

## Install and Configure MATLAB Parallel Server for Slurm

If you have a cluster with Slurm, follow these instructions to integrate MATLAB with your scheduler using MATLAB Parallel Server. If you do not have an existing scheduler in your cluster, see: “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16.

These instructions guide you through the following tasks:

In this section...
“Activate Your MATLAB Parallel Server License” on page 3-6
“Get the Installation Files” on page 3-6
“Install License Manager” on page 3-7
“Install Software on Compute Nodes” on page 3-8
“Install Software on Local Desktop” on page 3-8
“Configure Client Machine” on page 3-9
“Validate the Cluster Profile” on page 3-10
“Run Parallel Code” on page 3-11

After you integrate MATLAB with Slurm, you can access workers in your cluster from a desktop MATLAB client session with Parallel Computing Toolbox. Workers are MATLAB computational engines that typically correspond to a core.

The setup in these steps uses the network license manager.

### Activate Your MATLAB Parallel Server License

To install MATLAB Parallel Server, you must activate your license. To activate your MATLAB Parallel Server license:

- 1 Navigate to <https://www.mathworks.com/licensecenter>.
- 2 Log into the Administrator’s MathWorks Account.
- 3 Select your MATLAB Parallel Server license, and click the **Install and Activate** tab.
- 4 At the rightmost side, under **RELATED TASKS**, select **Activate to Retrieve License File**.
- 5 Fill in the requested information. This information must refer to the machine that hosts the license manager. In these instructions, it is the head node. For more information, see “Install License Manager” on page 3-13.
- 6 After filling in the information, download or email the License File and copy the File Installation Key. These are used later in the process.

---

**Note** Activation is not necessary for trials. Contact your sales representative to obtain the License File and the File Installation Key.

---

### Get the Installation Files


To save time and eliminate the need for the installer-based download process for each computer in your cluster, download the installation files prior to installation. Doing so facilitates installation in a

large number of machines. If you have access to an Administrator's account for your license, you can use the installer to download files without installing them. If not, contact the administrator of your license to obtain a copy of the installation files. For more information, see "Download Products Without Installing". When using the installer to download the files, choose the following options:

- Select the appropriate operating system for the cluster nodes.
- Select all products for download. MATLAB Parallel Server cannot run jobs requiring products that are not installed.

There are three server-side components of MATLAB Parallel Server:

- 1 The license manager, which hosts the MATLAB Parallel Server license used by each worker. For more information, see "Install License Manager" on page 3-13.
- 2 Your third-party job scheduler, which runs on the head node and manages jobs on your cluster. You integrate this scheduler with MATLAB Parallel Server. For more information, see "Install License Manager" on page 3-13 and "Configure Your Cluster" on page 3-14.
- 3 MATLAB Parallel Server, which runs on the compute nodes. For more information, see "Install Software on Compute Nodes" on page 3-14.



## Install License Manager

Choose a computer node to host the license manager. For the installation, use the offline installer from the previous step. For more information on the offline installation procedure, see "Install Products Using File Installation Key".

- 1 Start the MATLAB installer from the installation files acquired in "Get the Installation Files" on page 3-12.
- 2 Select **Use a File Installation Key**, and continue through the prompts.
- 3 In the product selection, select only the license manager.
- 4 In the License File step, browse to your `license.lic` file (obtained from "Activate Your MATLAB Parallel Server License" on page 3-12).
- 5 Start the license manager.

## Install Software on Compute Nodes

This procedure is similar to “Install License Manager” on page 3-13. For more information on the offline installation procedure, see “Install Products Using File Installation Key”.

- 1 Start the MATLAB installer from the installation files acquired in “Get the Installation Files” on page 3-12.
- 2 Select **Use a File Installation Key** and continue through the prompts.
- 3 Select all products. However, do not install the license manager. Alternatively, to save space, install only the products that the users of the cluster are licensed for.
- 4 Use the `license.dat` file from the head node. You can obtain this file from the `matlabroot/etc` folder, where `matlabroot` is the MATLAB installation folder.

For best performance, install locally on each node. However, you can also install in a network share location.

---

**Note** Install noninteractively (silently) instead if you want to

- Install the software on a machine without graphical user interface (GUI).
- Save the time that it takes to configure the installer for each compute node.

For more information, see “Install Noninteractively”.

---

You can test the installation and licensing by running the following command in a command-line interface. `matlabroot` is the MATLAB installation folder. `filename` is the location to write the log file to, for example, a temporary location. You must have write permissions in this location.

```
matlabroot/toolbox/parallel/bin/checkLicensing -logfile filename
```

---

**Note** You do not need to start the `mjs` service when you want to configure MATLAB Parallel Server for Slurm.

---

## Install Software on Local Desktop

To use MATLAB Parallel Server, you must use a local desktop running MATLAB and Parallel Computing Toolbox. Install the MathWorks products for which you are licensed, including Parallel Computing Toolbox, on the local desktops from which you want to submit jobs to the cluster. For help with this step, see .

Any MATLAB cluster workers that you start use dynamic licensing: they can use all the functionality you are licensed for in the MATLAB client, while checking out only MATLAB Parallel Server licenses in the cluster.

## Configure Client Machine

Next step	More information
Use out-of-the-box support if you meet all of these conditions. <ul style="list-style-type: none"> <li>• The client machine uses Linux.</li> <li>• There is a shared file system between the client machine and the cluster nodes.</li> <li>• The client machine has the third-party scheduler submission tools installed. You can check this by executing the relevant commands in a command line. For example, the submission script is <code>sbatch</code>.</li> </ul>	Follow “Create a Cluster Profile Using Out-of-the-Box Support” on page 3-9.
Use the generic scheduler interface.	Follow “Create a Cluster Profile Using the Generic Scheduler Interface” on page 3-10.

### Create a Cluster Profile Using Out-of-the-Box Support

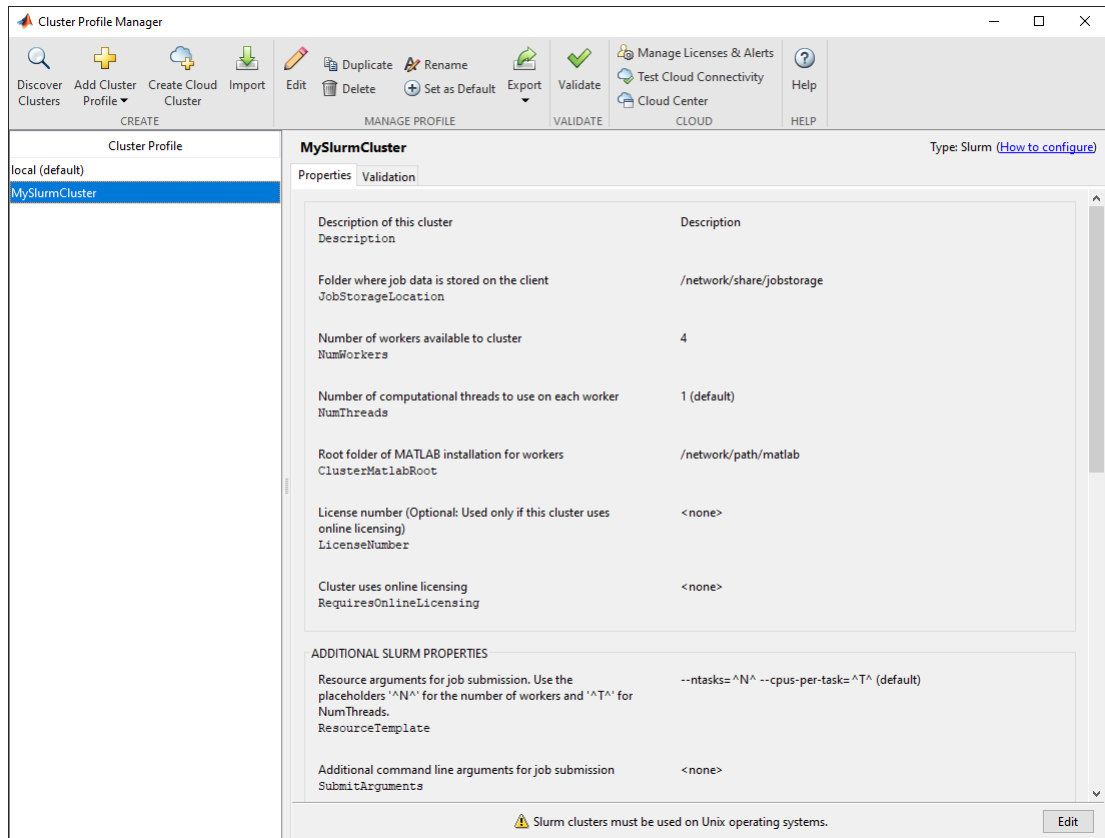
- 1 Start the Cluster Profile Manager. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- 2 Create a new Slurm profile. In the Cluster Profile Manager, select **Add Cluster Profile > Slurm**.
- 3 To give it a different name, select the new profile and click **Rename**.
- 4 In the **Properties** tab, provide settings for the following fields.
  - a (Optional) Set the **Description** field.
  - b Set the **JobStorageLocation** to the location where you want job and task data to be stored. This location must be accessible to all the worker machines.

---

**Note** Sharing `JobStorageLocation` with parallel computing products running different versions is not recommended; each version on your cluster should have its own `JobStorageLocation`.

---

- c Set the **NumWorkers** field to the number of workers you want to use with this profile, within the limitation of your licensing.
- d Set the **ClusterMatlabRoot** to the installation location of the MATLAB to be executed by the worker machines.
- e (Optional) Set the **SubmitArguments** to include any additional command arguments required by your particular cluster and scheduler.
- f After filling the fields, the dialog box looks like this:



5 Click **Done** to save your cluster profile.

### Create a Cluster Profile Using the Generic Scheduler Interface

Download the Slurm plugin scripts to interface with Slurm from the MATLAB client. On the **Home** tab, in the **Environment** area, click **Add-Ons**. In the Add-On Explorer, search for the add-on **Parallel Computing Toolbox plugin for MATLAB Parallel Server with Slurm** and install it. Alternatively, you can download the add-on from here: [Slurm](#). To open it, double click it or drag it and drop it in MATLAB.

For more information on the generic scheduler interface, see “Configure Using the Generic Scheduler Interface” on page 3-46.


### Validate the Cluster Profile

In this step you verify your cluster profile, and thereby your installation. You can specify the number of workers to use when validating your profile. If you do not specify the number of workers in the **Validation** tab, then the validation will attempt to use as many workers as the value specified by the **NumWorkers** property on the **Properties** tab. You can specify a smaller number of workers to validate your configuration without occupying the whole cluster.

- 1 If it is not already open, start the Cluster Profile Manager. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- 2 Select your cluster profile in the listing.
- 3 Click the **Validation** tab.

- 4 Use the checkboxes to choose all tests, or a subset of the validation stages, and specify the number of workers to use when validating your profile.
- 5 Click **Validate**.

The Validation Results tab shows the output. The following figure shows the results of a profile that passed all validation tests.




---

**Note** If your validation does not pass, contact the MathWorks install support team.

---

If your validation passed, you now have a valid profile that you can use in other parallel applications. You can make any modifications to your profile appropriate for your applications, such as `NumWorkersRange`, `AttachedFiles`, `AdditionalPaths`, etc.

To save your profile for other users, select the profile and click **Export**, then save your profile to a file in a convenient location. Later, when running the Cluster Profile Manager, other users can import your profile by clicking **Import**.

## Run Parallel Code

After you complete the configuration, you can submit jobs to the cluster. For examples, see “Running Code on Clusters and Clouds”.

## See Also

“Running Code on Clusters and Clouds”

## Install and Configure MATLAB Parallel Server for Third-Party Schedulers

If you already have a cluster with a scheduler, follow these instructions to integrate MATLAB with your scheduler using MATLAB Parallel Server. If you do not have an existing scheduler in your cluster, see: “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16.

These instructions guide you through the following tasks:

In this section...
“Activate Your MATLAB Parallel Server License” on page 3-12
“Get the Installation Files” on page 3-12
“Install License Manager” on page 3-13
“Install Software on Compute Nodes” on page 3-14
“Install Software on Local Desktop” on page 3-14
“Configure Your Cluster” on page 3-14

After you integrate MATLAB with a scheduler, you can access workers in your cluster from a desktop MATLAB client session with Parallel Computing Toolbox. Workers are MATLAB computational engines that typically correspond to a core.

The setup in these steps uses the network license manager.

### Activate Your MATLAB Parallel Server License

To install MATLAB Parallel Server, you must activate your license. To activate your MATLAB Parallel Server license:

- 1 Navigate to <https://www.mathworks.com/licensecenter>.
- 2 Log into the Administrator’s MathWorks Account.
- 3 Select your MATLAB Parallel Server license, and click the **Install and Activate** tab.
- 4 At the rightmost side, under **RELATED TASKS**, select **Activate to Retrieve License File**.
- 5 Fill in the requested information. This information must refer to the machine that hosts the license manager. In these instructions, it is the head node. For more information, see “Install License Manager” on page 3-13.
- 6 After filling in the information, download or email the License File and copy the File Installation Key. These are used later in the process.

---

**Note** Activation is not necessary for trials. Contact your sales representative to obtain the License File and the File Installation Key.

---

### Get the Installation Files

To save time and eliminate the need for the installer-based download process for each computer in your cluster, download the installation files prior to installation. Doing so facilitates installation in a

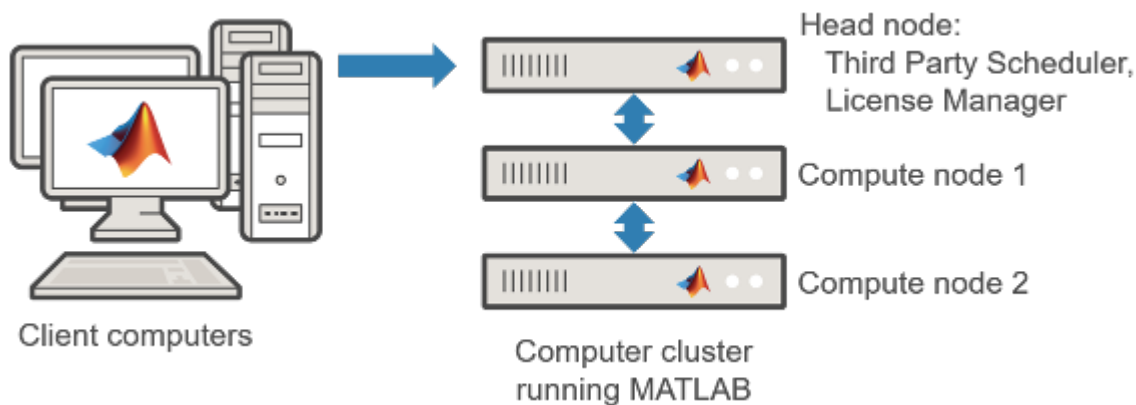


large number of machines. If you have access to an Administrator's account for your license, you can use the installer to download files without installing them. If not, contact the administrator of your license to obtain a copy of the installation files. For more information, see "Download Products Without Installing". When using the installer to download the files, choose the following options:

- Select the appropriate operating system for the cluster nodes.
- Select all products for download. MATLAB Parallel Server cannot run jobs requiring products that are not installed.

There are three server-side components of MATLAB Parallel Server:

- 1 The license manager, which hosts the MATLAB Parallel Server license used by each worker. For more information, see "Install License Manager" on page 3-13.
- 2 Your third-party job scheduler, which runs on the head node and manages jobs on your cluster. You integrate this scheduler with MATLAB Parallel Server. For more information, see "Install License Manager" on page 3-13 and "Configure Your Cluster" on page 3-14.
- 3 MATLAB Parallel Server, which runs on the compute nodes. For more information, see "Install Software on Compute Nodes" on page 3-14.



## Install License Manager

Choose a computer node to host the license manager. For the installation, use the offline installer from the previous step. For more information on the offline installation procedure, see "Install Products Using File Installation Key".

- 1 Start the MATLAB installer from the installation files acquired in "Get the Installation Files" on page 3-12.
- 2 Select **Use a File Installation Key**, and continue through the prompts.
- 3 In the product selection, select only the license manager.
- 4 In the License File step, browse to your `license.lic` file (obtained from "Activate Your MATLAB Parallel Server License" on page 3-12).
- 5 Start the license manager.

## Install Software on Compute Nodes

This procedure is similar to “Install License Manager” on page 3-13. For more information on the offline installation procedure, see “Install Products Using File Installation Key”.

- 1 Start the MATLAB installer from the installation files acquired in “Get the Installation Files” on page 3-12.
- 2 Select **Use a File Installation Key** and continue through the prompts.
- 3 Select all products. However, do not install the license manager. Alternatively, to save space, install only the products that the users of the cluster are licensed for.
- 4 Use the `license.dat` file from the head node. You can obtain this file from the `matlabroot/etc` folder, where `matlabroot` is the MATLAB installation folder.

For best performance, install locally on each node. However, you can also install in a network share location.

---

**Note** Install noninteractively (silently) instead if you want to

- Install the software on a machine without graphical user interface (GUI).
- Save the time that it takes to configure the installer for each compute node.

For more information, see “Install Noninteractively”.

---

You can test the installation and licensing by running the following command in a command-line interface. `matlabroot` is the MATLAB installation folder. `filename` is the location to write the log file to, for example, a temporary location. You must have write permissions in this location.

```
matlabroot/toolbox/parallel/bin/checkLicensing -logfile filename
```

---

**Note** You do not need to start the `mjs` service when you want to configure MATLAB Parallel Server for a third-party scheduler.

---

## Install Software on Local Desktop

To use MATLAB Parallel Server, you must use a local desktop running MATLAB and Parallel Computing Toolbox. Install the MathWorks products for which you are licensed, including Parallel Computing Toolbox, on the local desktops from which you want to submit jobs to the cluster. For help with this step, see .

Any MATLAB cluster workers that you start use dynamic licensing: they can use all the functionality you are licensed for in the MATLAB client, while checking out only MATLAB Parallel Server licenses in the cluster.

## Configure Your Cluster

When the cluster and client installations are complete, you can proceed to configure the products for the job scheduler of your choice. Use this table to choose a suitable guide to complete your configuration.

Next step	More information
<p>Use out-of-the-box support if</p> <ul style="list-style-type: none"> <li>• Your scheduler has out-of-the-box support in MATLAB. These schedulers are Slurm, PBS Pro, OpenPBS, Torque, LSF, and HPC Pack.</li> <li>• There is a shared file system between the client machine and the cluster nodes.</li> <li>• The client machine has the third-party scheduler submission tools installed. You can check this by executing the relevant commands in a command line. For example, the Slurm submission script is <code>sbatch</code>.</li> </ul>	<ul style="list-style-type: none"> <li>• “Configure for HPC Pack” on page 3-38</li> <li>• “Configure for Slurm, PBS Pro, OpenPBS, LSF, TORQUE” on page 3-42</li> </ul>
<p>Use out-of-the-box support if you have a Hadoop cluster.</p>	<ul style="list-style-type: none"> <li>• “Configure a Hadoop Cluster” on page 3-62</li> </ul>
<p>Use the generic scheduler interface.</p>	<ul style="list-style-type: none"> <li>• “Configure Using the Generic Scheduler Interface” on page 3-46</li> </ul>

## See Also

“Running Code on Clusters and Clouds”

## Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager

If you do not have an existing scheduler in your cluster, follow these instructions to integrate the MATLAB Job Scheduler, which is provided with MATLAB Parallel Server. If you already have a cluster with a scheduler, see “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12.

These instructions guide you through the following tasks:

In this section...
“Activate Your MATLAB Parallel Server License” on page 3-16
“Get the Installation Files” on page 3-16
“Install Software on the Head Node” on page 3-17
“Install Software on Compute Nodes” on page 3-18
“Configure the MATLAB Job Scheduler” on page 3-18
“Connect the MATLAB Client to the MATLAB Parallel Server Cluster” on page 3-20

After you integrate MATLAB with a scheduler, you can access workers in your cluster from a desktop MATLAB client session with Parallel Computing Toolbox. Workers are MATLAB computational engines that typically correspond to a core.

The setup in these steps uses the network license manager.

### Activate Your MATLAB Parallel Server License

To install MATLAB Parallel Server, you must activate your license. To activate your MATLAB Parallel Server license:

- 1 Navigate to <https://www.mathworks.com/licensecenter>.
- 2 Log into the Administrator’s MathWorks Account.
- 3 Select your MATLAB Parallel Server license, and click the **Install and Activate** tab.
- 4 At the rightmost side, under **RELATED TASKS**, select **Activate to Retrieve License File**.
- 5 Fill in the requested information. This information must refer to the machine that hosts the license manager. In these instructions, it is the head node. For more information, see “Install Software on the Head Node” on page 3-17.
- 6 After filling in the information, download or email the License File and copy the File Installation Key. These are used later in the process.

---

**Note** Activation is not necessary for trials. Contact your sales representative to obtain the License File and the File Installation Key.

---

### Get the Installation Files

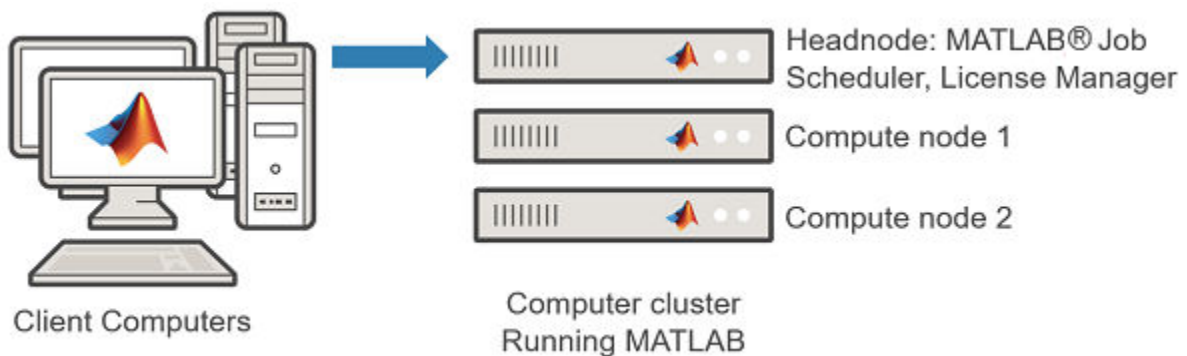
To save time and eliminate the need for the installer-based download process for each computer in your cluster, download the installation files prior to installation. Doing so facilitates installation in a

large number of machines. If you have access to an Administrator's account for your license, you can use the installer to download files without installing them. If not, contact the administrator of your license to obtain a copy of the installation files. For more information, see "Download Products Without Installing". When using the installer to download the files, choose the following options:

- Select the appropriate operating system for the cluster machines.
- Select all products for download. MATLAB Parallel Server cannot run jobs requiring products that are not installed.

There are three server-side components of MATLAB Parallel Server:

- 1 The license manager, which hosts the MATLAB Parallel Server license used by each worker. For more information, see "Install Software on the Head Node" on page 3-17.
- 2 The MATLAB Job Scheduler, which runs on the head node and manages jobs on your cluster. For more information, see "Install Software on the Head Node" on page 3-17.
- 3 MATLAB Parallel Server, which runs on the compute nodes. For more information, see "Install Software on Compute Nodes" on page 3-18.



## Install Software on the Head Node

Use the offline installer from the previous step. For more information on the offline installation procedure, see "Install Products Using File Installation Key".

- 1 Choose a computer to host the license manager and the MATLAB Job Scheduler. This computer is your head node.
- 2 Start the MATLAB installer from the installation files acquired in "Get the Installation Files" on page 3-16.
- 3 Select **Use a File Installation Key**, and continue through the prompts.
- 4 Select all products, including the license manager. Alternatively, to save space, install only the products that the users of the cluster are licensed for, in addition to the license manager.
- 5 In the License File step, browse to your `license.lic` file (obtained from "Activate Your MATLAB Parallel Server License" on page 3-16).
- 6 Start the license manager.

You can test the installation and licensing by running the following command in a command-line interface. `matlabroot` is the MATLAB installation folder. `filename` is the location to write the log file to, for example, a temporary location. You must have write permissions in this location.

```
matlabroot/toolbox/parallel/bin/checkLicensing -logfile filename
```

## Install Software on Compute Nodes

This procedure is similar to “Install Software on the Head Node” on page 3-17. For more information on the offline installation procedure, see “Install Products Using File Installation Key”.

- 1 Start the MATLAB installer from the installation files acquired in “Get the Installation Files” on page 3-16.
- 2 Select **Use a File Installation Key** and continue through the prompts.
- 3 Select all products. However, do not install the license manager. Alternatively, to save space, install only the products that the users of the cluster are licensed for.
- 4 Use the `license.dat` file from the head node. You can obtain this file from the `matlabroot/etc` folder, where `matlabroot` is the MATLAB installation folder.

For best performance, install locally on each node. However, you can also install in a network share location.

---

**Note** Install noninteractively (silently) instead if you want to

- Install the software on a machine without graphical user interface (GUI).
- Save the time that it takes to configure the installer for each compute node.

For more information, see “Install Noninteractively”.

---

You can test the installation and licensing by running the following command in a command-line interface. `matlabroot` is the MATLAB installation folder.

```
matlabroot/toolbox/parallel/bin/checkLicensing
```

## Configure the MATLAB Job Scheduler

The MATLAB Job Scheduler is a scheduler that ships with MATLAB Parallel Server. The MATLAB Job Scheduler is intended primarily for clusters that run only MATLAB jobs. The scheduler interface is a high-level abstraction that lets you submit jobs to your computation resources, so you do not have to deal with differences in operating systems and environments.

The following steps configure the MATLAB Job Scheduler with Admin Center, the graphical interface. If your machine does not provide graphics, use the command-line interface instead. For more information, see “Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line)” on page 3-32.

- 1 On the head node, start Admin Center. Browse to `matlabroot/toolbox/parallel/bin` and execute the file named `admincenter`. `matlabroot` is the MATLAB installation folder.
- 2 Click **Add or Find**, and specify the computers that you are using as your head node and compute nodes.
- 3 Progress through the prompts, and confirm to start the mjs service. If necessary, manually start the mjs service using the command-line interface. For more information, see “Use the Command-Line Interface (Windows)” on page 3-32 or “Use the Command-Line Interface (UNIX)” on page 3-34.

- 4 In the MATLAB Job Scheduler section, click **Start**. Specify a name for your MATLAB Job Scheduler, and select the head node from the dropdown list.
- 5 To add the MATLAB Parallel Server workers, click **Start** in the **Workers** section of the Admin Center.
  - a Select the computers to host the workers.
  - b Select the number of workers per computer.
- 6 To verify your configuration, review worker status in the **Workers** section.
- 7 To troubleshoot issues, click **Test Connectivity** in the **Host** section.
- 8 If you are using UNIX, configure the mjs service to start automatically at start time. For instructions, see “Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line)” on page 3-32.

The head node uses computational resources to run the MATLAB Job Scheduler. If you set up workers on the head node, they compete for resources with the MATLAB Job Scheduler.

---

**Tip** Avoid setting up workers on the head node. If workers use too many system resources, such as memory, processor, network, or local storage, the job manager can become unresponsive.

---

The following screenshot shows a final setup in Admin Center:

The screenshot shows the Admin Center interface with three main sections: Hosts, MATLAB Job Scheduler, and Workers. Each section contains a table of system components and their status.

**Hosts Section:**

	Host			MJS Service		MATLAB Job Sche...	Workers
	Hostname	Reachable	Cores	Status	Up Since	Name	Count
ComputeNodeHostname	yes	6	running	2019-01-03 15:55			4
HeadNodeHostname	yes	6	running	2019-01-03 15:53	myJobManager		0

**MATLAB Job Scheduler Section:**

	Name	Hostname	Status	Up Since	Workers
myJobManager	HeadNodeHostname	running	2019-01-03 16:03	4	

**Workers Section:**

	Worker				MATLAB Job Scheduler		
	Name	Hostname	Status	Up Since	Connection	Name	Hostname
Worker1	ComputeNode...	idle	2019-01-03 16:03	connected	myJobManager	HeadNodeHos...	
Worker2	ComputeNode...	idle	2019-01-03 16:04	connected	myJobManager	HeadNodeHos...	
Worker3	ComputeNode...	idle	2019-01-03 16:04	connected	myJobManager	HeadNodeHos...	
Worker4	ComputeNode...	idle	2019-01-03 16:04	connected	myJobManager	HeadNodeHos...	

At the bottom of the interface, it shows "Last updated: 03/01/19 16:04" and an "Update" button set to "every 2 minutes" with an "Update Now" button.

**Note** If you need more help during the configuration, such as your cluster requires firewall configuration or you want to set up multiple mjs installations, see this more detailed guide: “Configure Advanced Options for MATLAB Job Scheduler Integration” on page 3-29.

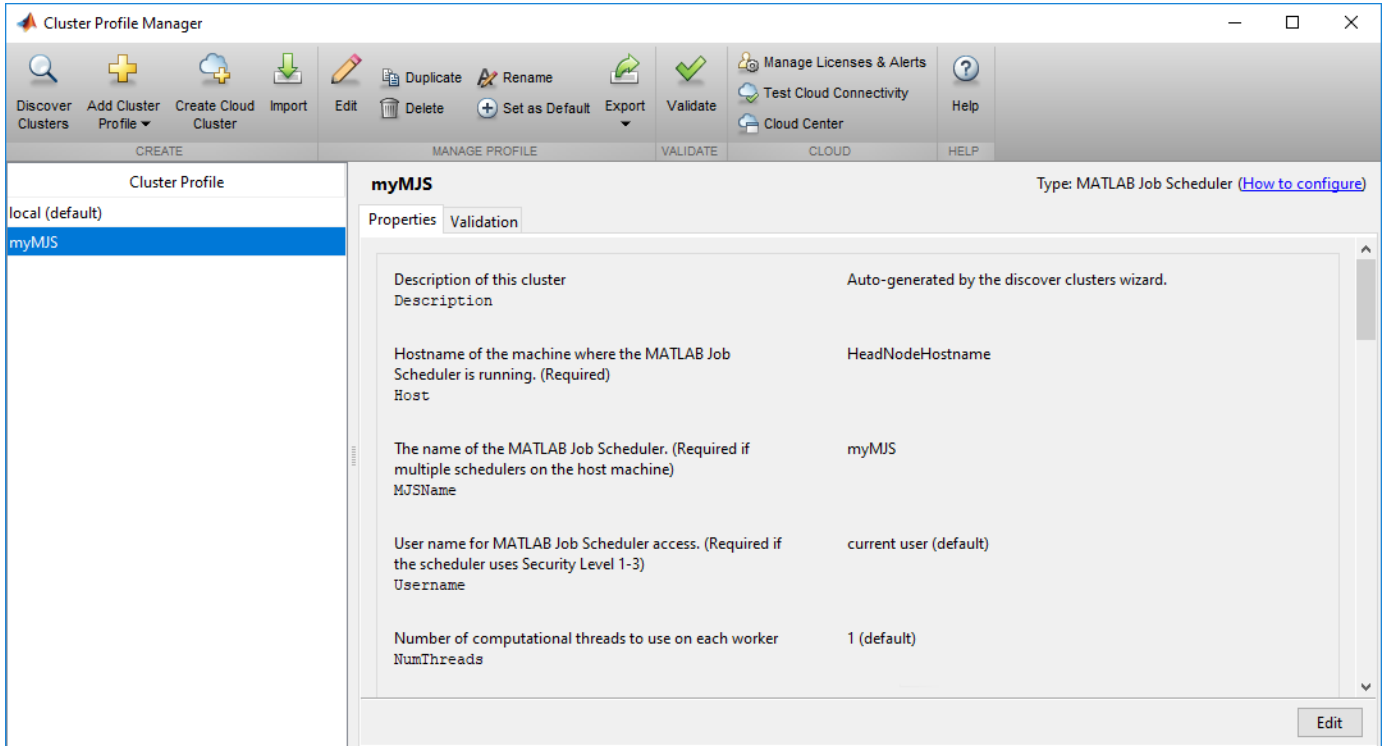
### Connect the MATLAB Client to the MATLAB Parallel Server Cluster

To use MATLAB Parallel Server, you must use a client computer running MATLAB and Parallel Computing Toolbox. In the MATLAB toolstrip, use **Parallel > Discover Clusters** and follow the instructions to automatically discover and set up your cluster. Alternatively, you can configure it manually as follows:

- 1 In MATLAB, on the **Home** tab, select the **Parallel** menu. Select **Create and Manage Clusters**.
- 2 Click **Add Cluster Profile > MATLAB Job Scheduler**.
- 3 To modify the name of the MATLAB Job Scheduler profile, double click the profile name.
- 4 To edit the profile, select it and click **Edit** in the toolbar.



- 5 In the **Host** field, enter the host name of the head node.
- 6 Click **Done**. The following image shows a MATLAB Job Scheduler cluster profile after configuration:



- 7 To make this profile the default, select **Set as Default**.
- 8 **Validate** the cluster profile.

If validation of your cluster is successful, your MATLAB session can now submit jobs to the MATLAB Parallel Server cluster.

---

**Note** If your validation does not pass, contact the MathWorks install support team.

---

Any MATLAB cluster workers that you start use dynamic licensing: they can use all the functionality you are licensed for in the MATLAB client, while checking out only MATLAB Parallel Server licenses in the cluster.

To configure more advanced options for your cluster, see “MATLAB Job Scheduler Cluster Customization”. For example, you can set the security of the cluster in “Set MATLAB Job Scheduler Cluster Security” on page 2-14. After you finish your configuration, try some examples of cluster workflows in “Running Code on Clusters and Clouds”.

## See Also

### Related Examples

- “MATLAB Job Scheduler Cluster Customization”

- “Set MATLAB Job Scheduler Cluster Security” on page 2-14

# Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Online Licensing

If you do not have an existing scheduler in your cluster, follow these instructions to integrate the MATLAB Job Scheduler, which is provided with MATLAB Parallel Server. If you already have a cluster with a scheduler, see “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12.

These instructions guide you through the following tasks:

In this section...
“Check License Type and Users” on page 3-23
“Get the Installation Files” on page 3-24
“Install Software on All Nodes” on page 3-25
“Configure the MATLAB Job Scheduler with Admin Center” on page 3-25
“Connect the MATLAB Client to the MATLAB Parallel Server Cluster” on page 3-27

After you integrate MATLAB with a scheduler, you can access workers in your cluster from a desktop MATLAB client session (requires Parallel Computing Toolbox). Workers are MATLAB computational engines that typically correspond to a core.

The setup in these steps uses online licensing.

## Check License Type and Users

### Check License Type

**Note** If you have a campus-wide license for MATLAB Parallel Server, then you already have access to online licensing and can skip to “Add Licensed End Users” on page 3-23. If you are unsure about your license details, then contact the license administrator.

To install MATLAB Parallel Server using online licensing, you must check your license type.


- 1 In your browser, go to License Center and log in with your MathWorks Administrator Account.
- 2 Select the MATLAB Parallel Server license that you plan to use.
- 3 On the **Install and Activate** tab, look for **License Manager**: followed by the license manager type currently assigned to this license.
  - If the license manager is already the one you want, then you do not need to do anything. Go to “Add Licensed End Users” on page 3-23.
  - To change the license manager, click the pencil icon and follow the onscreen instructions. When you have finished, go to “Add Licensed End Users” on page 3-23.

### Add Licensed End Users

With online licensing, any user of MATLAB Parallel Server must be added as a licensed end user. To add licensed end users, you must be a license administrator. If you are not a license administrator,

provide a list of required licensed end users to the license administrator along with the following steps. If you are a license administrator, follow these steps.

- 1 If you are not already logged in as an administrator, go to License Center and log in with your MathWorks Administrator Account.
- 2 Select your MATLAB Parallel Server license, and then click **Manage Users**.
- 3 Click **Add User** to add a user to the list.



**Add User**

A MathWorks Account will be automatically created for the user, if one does not already exist.

Email Address:

First Name:

Last Name:

Country:

Download Software

- 4 Provide the user's email address, first and last names, and country. Click **Add User**. Note that if the specified email address does not correspond to an existing MathWorks Account, a new account is created for that user.
- 5 Add end users as necessary.

## Get the Installation Files

If you are not a license administrator, you can download the Internet-based installer. For more information, see "Install Products Using Internet Connection". Note that this method performs a full download of the necessary files per compute node that you are setting up.


To save time and eliminate the need to run the installer-based download process on each computer in your cluster, download the installation files prior to installation. Doing so facilitates installation in a large number of machines. If you have access to an Administrator's account for your license, you can use the installer to download files without installing them. If not, contact the administrator of your license to obtain a copy of the installation files. For more information, see "Download Products Without Installing". Alternatively, proceed with the Internet-based installer. When you download the files using the installer, you must:

- Select the operating system for the cluster machines.
- Select all products for download. MATLAB Parallel Server cannot run jobs requiring products that are not installed.

MATLAB Parallel Server has two server-side components:

- The MATLAB Job Scheduler, which runs on the head node and manages jobs on your cluster. For more information, see "Install Software on All Nodes" on page 3-25.

- MATLAB Parallel Server, which runs on the compute nodes. For more information, see “Install Software on All Nodes” on page 3-25.



## Install Software on All Nodes

To install the software on each computer in your cluster, follow these steps:

- 1 Start the MATLAB installer from the installation files downloaded in “Get the Installation Files” on page 3-24.
- 2 Select **Log in with a MathWorks Account** and follow the prompts.
- 3 Select all products that the end users will use, and exclude the license manager.
- 4 After the installation completes, update the `mjs_def` file in `matlabroot/toolbox/parallel/bin`. Uncomment and set:
  - Unix: `USE_ONLINE_LICENSING="true"`
  - Windows: `USE_ONLINE_LICENSING=true`

For best performance, install the software locally on each node. However, you can also install the software in a network share location.

---

**Note** Install noninteractively (silently) instead if you want to

- Install the software on a machine without graphical user interface (GUI).
- Save the time that it takes to configure the installer for each compute node.

For more information, see “Install Noninteractively”.

---

You can test the installation and licensing by running the following command in a command-line interface. `matlabroot` is the MATLAB installation folder. `filename` is the location to write the log file to, for example, a temporary location. You must have write permissions in this location.

```
matlabroot/toolbox/parallel/bin/checkLicensing -logfile filename
```

## Configure the MATLAB Job Scheduler with Admin Center

The MATLAB Job Scheduler is a scheduler that is provided with MATLAB Parallel Server. The MATLAB Job Scheduler is intended primarily for clusters that run only MATLAB jobs. The scheduler interface is a high-level abstraction that enables you to submit jobs to your computation resources, and allows you to avoid dealing with differences in operating systems and environments.

The following steps configure the MATLAB Job Scheduler with Admin Center, the graphical interface. If your machine does not provide graphics, use the command-line interface instead. For more information, see “Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line)” on page 3-32.

- 1** On the head node, start Admin Center. Go to `matlabroot/toolbox/parallel/bin` and execute the file named `admincenter`. `matlabroot` is the MATLAB installation folder.
- 2** Click **Add or Find**, and specify the computers that you are using as your head node and compute nodes.
- 3** Follow the prompts and confirm to start the mjs service. If necessary, manually start the mjs service using the command-line interface. For more information, see “Use the Command-Line Interface (Windows)” on page 3-32 or “Use the Command-Line Interface (UNIX)” on page 3-34.
- 4** In the MATLAB Job Scheduler section, click **Start**. Specify a name for your MATLAB Job Scheduler and select the head node from the dropdown list.
- 5** To add MATLAB Parallel Server workers, click **Start** in the **Workers** section of the Admin Center.
  - a** Select the computers to host the workers.
  - b** Select the number of workers per computer.
- 6** Verify your configuration by checking worker status in the **Workers** section.
- 7** To troubleshoot issues, click **Test Connectivity** in the **Host** section.
- 8** If you use UNIX, configure the mjs service to start automatically at start time. For instructions, see “Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line)” on page 3-32.

The head node uses computational resources to run the MATLAB Job Scheduler. If you set up workers on the head node, they compete for resources with the MATLAB Job Scheduler.

---

**Tip** Avoid setting up workers on the head node. If workers use too many system resources, such as memory, processor, network, or local storage, the job manager can become unresponsive.

---

The following screenshot shows the final setup in Admin Center.

The screenshot shows the MATLAB Admin Center interface with three main sections: Hosts, MATLAB Job Scheduler, and Workers. Each section contains a table of system components and their status.

**Hosts Section:**

Add or Find... Start MJS Service... Stop MJS Service... Test Connectivity...	Host			MJS Service		MATLAB Job Sche...	Workers
	Hostname ^	Reachable	Cores	Status	Up Since	Name	Count
	ComputeNodeHostname	● yes	6	● running	2019-01-03 15:55		4
	HeadNodeHostname	● yes	6	● running	2019-01-03 15:53	myJobManager	0

**MATLAB Job Scheduler Section:**

Start... Stop... Resume	Name ^	Hostname	Status	Up Since	Workers
		myJobManager	HeadNodeHostname	● running	2019-01-03 16:03

**Workers Section:**

Start... Stop... Resume	Worker				MATLAB Job Scheduler		
	Name ^	Hostname	Status	Up Since	Connection	Name	Hostname
	Worker1	ComputeNode...	● idle	2019-01-03 16:03	● connected	myJobManager	HeadNodeHos...
	Worker2	ComputeNode...	● idle	2019-01-03 16:04	● connected	myJobManager	HeadNodeHos...
	Worker3	ComputeNode...	● idle	2019-01-03 16:04	● connected	myJobManager	HeadNodeHos...
	Worker4	ComputeNode...	● idle	2019-01-03 16:04	● connected	myJobManager	HeadNodeHos...

At the bottom of the interface, it shows "Last updated: 03/01/19 16:04" and an "Update" button set to "every 2 minutes" with an "Update Now" button.

**Note** If you need more help during the configuration, such as your cluster requires firewall configuration or you want to set up multiple mjs installations, see this more detailed guide “Configure Advanced Options for MATLAB Job Scheduler Integration” on page 3-29.

## Connect the MATLAB Client to the MATLAB Parallel Server Cluster

To use MATLAB Parallel Server, you must have a client computer running MATLAB and Parallel Computing Toolbox. In the MATLAB toolstrip, select **Parallel > Discover Clusters** and follow the instructions to automatically discover and set up your cluster.

Alternatively, you can configure it manually as follows:

- 1 In MATLAB, on the **Home** tab, select **Parallel > Create and Manage Clusters**.
- 2 Select **Add Cluster Profile > MATLAB Job Scheduler**.

- Create your MATLAB Job Scheduler profile and click **Edit**.
- Update the hostname of the head node.
- Update the license number.
- Click **Done** and select **Set as Default** (optional) .

After you successfully validate your cluster, you can now use your MATLAB session to submit jobs to the MATLAB Parallel Server cluster.

---

**Note** If your validation does not pass, contact the MathWorks install support team.

---

Any MATLAB cluster workers that you start use dynamic licensing: they can use all the functionality you are licensed for in the MATLAB client, while checking out only MATLAB Parallel Server licenses in the cluster.

For information on configuring more advanced options for your cluster, see “MATLAB Job Scheduler Cluster Customization”. For example, you can set the security of the cluster in “Set MATLAB Job Scheduler Cluster Security” on page 2-14. After you finish your configuration, try some examples of cluster workflows in “Running Code on Clusters and Clouds”.

## See Also

### Related Examples

- “MATLAB Job Scheduler Cluster Customization”
- “Set MATLAB Job Scheduler Cluster Security” on page 2-14



## Configure Advanced Options for MATLAB Job Scheduler Integration

### In this section...

“Run Multiple MATLAB Parallel Server Versions” on page 3-29

“Set Up Windows Cluster Hosts” on page 3-30

“Open Required Ports on Server” on page 3-30

“Stop mjs Services of Old Installation” on page 3-31

“Set the MATLAB Job Scheduler Security Level” on page 3-32

“Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line)” on page 3-32

“Install the mjs Service to Start Automatically at Boot Time (UNIX)” on page 3-35

“Validate Installation with MATLAB Job Scheduler” on page 3-37

Follow these instructions to configure advanced options during integration of MATLAB Job Scheduler with your cluster.

**Note** If this is the first time you integrate MATLAB Job Scheduler, see the following for the most common configuration options: “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16.

In the following instructions, *matlabroot* refers to the location of your installed MATLAB Parallel Server software. Where you see this term used in the instructions that follow, substitute the path to your location.

### Run Multiple MATLAB Parallel Server Versions

You can upgrade your MATLAB Job Scheduler clusters and continue to use the R2016a release onwards of Parallel Computing Toolbox on your MATLAB desktop client to connect to it. To take advantage of this backward compatibility feature:

- 1 Install the latest version of MATLAB Parallel Server on your cluster. You must use this version to run MATLAB Job Scheduler on your cluster.
- 2 Install MATLAB Parallel Server for each release that you want to support in the cluster. For example, to use R2016a and R2016b with your cluster, install both the R2016a and R2016b releases of MATLAB Parallel Server.
- 3 Configure MATLAB Job Scheduler with the location of these installations. In the `mjs_def` configuration file, specify the location of each installation of MATLAB Parallel Server in the `MJS_ADDITIONAL_MATLABROOTS` variable. You can find this file in `matlabroot/toolbox/parallel/bin` for Linux (`mjs_def.sh`) and Windows (`mjs_def.bat`). For more information, see `mjs`.

With this configuration, the MATLAB Job Scheduler allows MATLAB clients from the installed releases to submit jobs to the cluster. The MATLAB Job Scheduler dynamically starts the right version of the MATLAB worker to run the job.

## Set Up Windows Cluster Hosts

If this is the first installation of MATLAB Parallel Server on a cluster of Windows machines, you need to configure these hosts for job communications.

---

**Note** If you do not have a Windows cluster, or if you have already installed a previous version of MATLAB Parallel Server on your Windows cluster, you can skip this step.

---

### Configure Windows Firewalls on the Client

If you are using Windows firewalls on your cluster nodes,

- 1 Log in as a user with administrator privileges.
- 2 Execute the following in a DOS command window.

```
matlabroot\toolbox\parallel\bin\addMatlabToWindowsFirewall.bat
```

This command adds MATLAB as an allowed program. If you are using other firewalls, you must configure them for similar accommodation.

### Configure Windows User Access for mjs

The user that mjs runs as requires access to the cluster MATLAB installation location. By default, mjs runs as the user LocalSystem. If your network allows LocalSystem to access the install location, you can skip this step. (If you are not sure of your network configuration and the access provided for LocalSystem, contact the MathWorks install support team.)

---

**Note** If LocalSystem cannot access the install location, you must run mjs as a different user.

---

You can set a different user with these steps:

- 1 With any standard text editor (such as WordPad) open the mjs\_def file found at:  

```
matlabroot\toolbox\parallel\bin\mjs_def.bat
```
- 2 Find the line for setting the MJSUSER parameter, and provide a value in the form domain \username:

```
set MJSUSER=mydomain\myusername
```

- 3 Provide the user password by setting the MJSPASS parameter:

```
set MJSPASS=password
```

- 4 Save the file.

## Open Required Ports on Server

The mjs service uses as many ports as required, starting with BASE\_PORT. By default, BASE\_PORT is 27350.

If you use a machine that runs a total of nJ job managers and nW workers, the mjs service reserves a total of  $6+2*nJ+4*nW$  consecutive ports for its own use. All job managers and workers, even those on

different hosts, that are going to work together must use the same base port. Otherwise the job managers and workers will not be able to contact each other. In addition, MPI communication occurs on ports starting at `BASE_PORT+1000` and use  $2*nw$  consecutive ports.

For example, if you use a machine with 1 job manager and 16 workers, then you need the following ranges of ports to be open:

- 27350 – 27422 for the mjs service.
- 28350 – 28382 for MPI communication.

To connect from MATLAB to a cluster with a non-default `BASE_PORT`, you must append the value of `BASE_PORT` to the 'Host' property in the MATLAB Job Scheduler cluster profile. You must do this in the form `Hostname:BASE_PORT`, for example `myMJSHost:44001`.

## Stop mjs Services of Old Installation

If you have an older version of MATLAB Parallel Server running on your cluster nodes, you should stop the mjs services before starting the services of the new installation.

### Stop mjs on Windows

- 1 Open a DOS command window with the necessary privileges:
  - a If you are using Windows 7 or Windows Vista™, you must run the command window with administrator privileges. Click the Windows menu **Start > (All) Programs > Accessories**; then right-click **Command Window**, and select **Run as Administrator**. This option is available only if you are running User Account Control (UAC).
  - b If you are using Windows XP, open a DOS command window by selecting the Windows menu **Start > Run**, then in the **Open** field, type

```
cmd
```

- 2 In the command window, navigate to the folder of the old installation that contains the control scripts.

```
cd oldmatlabroot\toolbox\parallel\bin
```

- 3 Stop and uninstall the old service and remove its associated files by typing the following command.

```
mjs uninstall -clean
```

In releases before R2019a, the service is called mdce. Type the following commands instead.

```
cd oldmatlabroot\toolbox\distcomp\bin
mdce uninstall -clean
```

---

**Note** Using the `-clean` flag permanently removes all existing job data. Be sure this data is no longer needed before removing it.

---

- 4 Repeat the instructions of this step on all worker nodes.

### Stop mjs on UNIX

- 1 Log in as root. If you cannot log in as root, you must alter the following parameters in the `oldmatlabroot/toolbox/parallel/bin/mjs_def.sh` file to point to a folder for which you

have write privileges: CHECKPOINTBASE, LOGBASE, PIDBASE, and LOCKBASE if applicable. In releases before R2019a, this file is *oldmatlabroot/toolbox/distcomp/bin/mdce\_def.sh* instead.

- 2 On each cluster node**, stop the mjs service and remove its associated files by typing the commands:

```
cd oldmatlabroot/toolbox/parallel/bin
./mjs stop -clean
```

In releases before R2019a, the service is called mdce. Type the following command instead.

```
cd oldmatlabroot/toolbox/distcomp/bin
./mdce stop -clean
```

---

**Note** Using the `-clean` flag permanently removes all existing job data. Be sure this data is no longer needed before removing it.

---

## Set the MATLAB Job Scheduler Security Level

Before starting the mjs service on your cluster nodes, set a security level. For instructions, see “Set the Security Level” on page 2-14. For additional security considerations, see “Set MATLAB Job Scheduler Cluster Security” on page 2-14.

## Start the mjs Service, MATLAB Job Scheduler, and Workers (Command-Line)

You can start MATLAB Job Scheduler using a graphical interface or the command line. For instructions on how to use the graphical interface, see “Configure the MATLAB Job Scheduler” on page 3-18. To use the graphical interface, Admin Center, you must run it on a computer that has direct network connectivity to all the nodes of your cluster. If you cannot run Admin Center on such a computer, you must use the command-line interface. For instructions on how to use the command-line interface, follow the next steps.

### Use the Command-Line Interface (Windows)

#### 1 Start the mjs Service

You must install the mjs service on all nodes (head node and worker nodes). Begin on the head node.

- a** Open a DOS command window with the necessary privileges:
  - i** If you are using Windows or Windows Vista, you must run the command window with administrator privileges. Click the Windows menu **Start > (All) Programs > Accessories**; then right-click **Command Window**, and select **Run as Administrator**. This option is available only if you are running User Account Control (UAC).
  - ii** If you are using Windows XP, open a DOS command window by selecting the Windows menu **Start > Run**, then in the **Open** field, type:

```
cmd
```

- b** In the DOS command window, navigate to the folder with the control scripts:

```
cd matlabroot\toolbox\parallel\bin
```

- c** Install the mjs service by typing the command:

```
mjs install
```

- d** Start the mjs service by typing the command:

```
mjs start
```

- e** Repeat the instructions of this step on all worker nodes.

As an alternative to items 3-5, you can install and start the mjs service on several nodes remotely from one machine by typing:

```
cd matlabroot\toolbox\parallel\bin
remotemjs install -remotehost hostA,hostB,hostC . . .
remotemjs start -remotehost hostA,hostB,hostC . . .
```

where `hostA`, `hostB`, `hostC` refers to a list of your host names. Note that there are no spaces between host names, only a comma. If you need to indicate protocol, platform (such as in a mixed environment), or other information, see the help for `remotemjs` by typing:

```
remotemjs -help
```

Once installed, the mjs service starts running each time the machine reboots. The mjs service continues to run until explicitly stopped or uninstalled, regardless of whether a MATLAB Job Scheduler or worker session is running.

## 2 Start the MATLAB Job Scheduler

To start the MATLAB Job Scheduler, enter the following commands in a DOS command window. You do not have to be at the machine on which the MATLAB Job Scheduler runs, as long as you have access to the MATLAB Parallel Server installation.

- a** In your DOS command window, navigate to the folder with the startup scripts:

```
cd matlabroot\toolbox\parallel\bin
```

- b** Start the MATLAB Job Scheduler, using any unique text you want for the name `<MyMJS>`:

```
startjobmanager -name <MyMJS> -remotehost <MATLAB Job Scheduler host name> -v
```

- c** Verify that the MATLAB Job Scheduler is running on the intended host.

```
nodestatus -remotehost <MATLAB Job Scheduler host name>
```

---

**Note** If you are executing `startjobmanager` on the host where the MATLAB Job Scheduler runs, you do not need to specify the `-remotehost` flag.

If you have more than one MATLAB Job Scheduler on your cluster, each must have a unique name.

---

## 3 Start the Workers

---

**Note** Before you can start a worker on a machine, the mjs service must already be running on that machine. If you are using the network license manager, it must be running on the network.

---

For each node used as a worker, enter the following commands in a DOS command window. You do not have to be at the machines where the MATLAB workers will run, as long as you have access to the MATLAB Parallel Server installation.

- a** Navigate to the folder with the startup scripts:

```
cd matlabroot\toolbox\parallel\bin
```

- b** Start the workers on each node, using the text for <MyMJS> that identifies the name of the MATLAB Job Scheduler you want this worker registered with. Enter this text on a single line:

```
startworker -jobmanagerhost <MATLAB Job Scheduler host name>  
            -jobmanager <MyMJS> -remotehost <worker host name> -v
```

To run more than one worker session on the same node, give each worker a unique name by including the `-name` option on the `startworker` command, and run it for each worker on that node:

```
startworker ... -name <worker1 name>  
startworker ... -name <worker2 name>
```

- c** Verify that the workers are running.

```
nodestatus -remotehost <worker host name>
```

- d** Repeat items 2-3 for all worker nodes.

For more information about `mjs`, MATLAB Job Scheduler, and worker processes, such as how to shut them down or customize them, see “MATLAB Job Scheduler Cluster Customization”.

## Use the Command-Line Interface (UNIX)

### 1 Start the `mjs` Service

On each cluster node, start the `mjs` service by typing the commands:

```
cd matlabroot/toolbox/parallel/bin  
./mjs start
```

Alternatively (on Linux, but not Macintosh), you can start the `mjs` service on several nodes remotely from one machine by typing

```
cd matlabroot/toolbox/parallel/bin  
./remotemjs start -remotehost hostA,hostB,hostC . . .
```

where `hostA,hostB,hostC` refers to a list of your host names. Note that there are no spaces between host names, only a comma. If you need to indicate protocol, platform (such as in a mixed environment), or other information, see the help for `remotemjs` by typing

```
./remotemjs -help
```

### 2 Start the MATLAB Job Scheduler

To start the MATLAB Job Scheduler, enter the following commands. You do not have to be at the machine on which the MATLAB Job Scheduler runs, as long as you have access to the MATLAB Parallel Server installation.

- a** Navigate to the folder with the startup scripts:

```
cd matlabroot/toolbox/parallel/bin
```

- b** Start the MATLAB Job Scheduler, using any unique text you want for the name <MyMJS>. Enter this text on a single line.

```
./startjobmanager -name <MyMJS> -remotehost <MATLAB Job Scheduler host name> -v
```

- c** Verify that the MATLAB Job Scheduler is running on the intended host:

```
./nodestatus -remotehost <MATLAB Job Scheduler host name>
```

---

**Note** If you have more than one MATLAB Job Scheduler on your cluster, each must have a unique name.

---

### 3 Start the Workers

---

**Note** Before you can start a worker on a machine, the mjs service must already be running on that machine. If you are using the network license manager, it must be running on the network.

---

For each computer hosting a MATLAB worker, enter the following commands. You do not have to be at the machines where the MATLAB workers run, as long as you have access to the MATLAB Parallel Server installation.

- a** Navigate to the folder with the startup scripts:

```
cd matlabroot/toolbox/parallel/bin
```

- b** Start the workers on each node, using the text for <MyMJS> that identifies the name of the MATLAB Job Scheduler you want this worker registered with. Enter this text on a single line:

```
./startworker -jobmanagerhost <MATLAB Job Scheduler host name>  
-jobmanager <MyMJS> -remotehost <worker host name> -v
```

To run more than one worker session on the same machine, give each worker a unique name with the `-name` option:

```
./startworker ... -name <worker1>  
./startworker ... -name <worker2>
```

- c** Verify that the workers are running. Repeat this command for each worker node:

```
./nodestatus -remotehost <worker host name>
```

For more information about mjs, MATLAB Job Scheduler, and worker processes, such as how to shut them down or customize them, see “MATLAB Job Scheduler Cluster Customization”.

## Install the mjs Service to Start Automatically at Boot Time (UNIX)

Although this step is not required, it is helpful in case of a system crash. Once configured for this, the mjs service starts running each time the machine reboots. The mjs service continues to run until explicitly stopped, regardless of whether a MATLAB Job Scheduler or worker session is running.

You must have root privileges to do this step.

Choose your platform:

- “Debian, Fedora, SUSE, and Red Hat (non-Fedora) Platforms” on page 3-36

- “Macintosh Platform” on page 3-36

### **Debian, Fedora, SUSE, and Red Hat (non-Fedora) Platforms**

On each cluster node, register the mjs service as a known service and configure it to start automatically at system boot time by following these steps:

- 1** Create the following link, if it does not already exist:

```
ln -s matlabroot/toolbox/parallel/bin/mjs /etc/mjs
```

- 2** Create the following link to the boot script file:

```
ln -s matlabroot/toolbox/parallel/bin/mjs /etc/init.d/mjs
```

- 3** Set the boot script file permissions:

```
chmod 555 /etc/init.d/mjs
```

- 4** Find your default run level. If you have a SysV linux machine, you can determine the default run level by booting your machine and immediately executing the `$runlevel` command. The second number output is the default run level of your system. If your Linux machine does not support SysV, look in `/etc/inittab` for the default run level.

- 5** When you have determined the run level, create a link in the `rc` folder associated with that run level. For example, if the run level is 5, execute one of the following sets of platform-specific commands.

- Debian and Fedora platforms:

```
cd /etc/rc5.d;  
ln -s ../init.d/mjs S99MJS
```

- SUSE platform:

```
cd /etc/init.d/rc5.d;  
ln -s ../mjs S99MJS
```

- Red Hat platform (non-Fedora):

```
cd /etc/rc.d/rc5.d;  
ln -s ../../init.d/mjs S99MJS
```

### **Macintosh Platform**

On each cluster node, register the mjs service as a known service with `launchd`, and configure it to start automatically at system boot time by following these steps:

- 1** Navigate to the toolbox folder and stop the running mjs service:

```
cd matlabroot/toolbox/parallel/bin  
sudo ./mjs stop
```

- 2** Create the following link if it does not already exist:

```
sudo mkdir -p /usr/local/sbin/  
sudo ln -s matlabroot/toolbox/parallel/bin/mjs /usr/local/sbin/mjs
```

- 3** Copy the `launchd.plist` file for mjs to `/Library/LaunchDaemons`:

```
sudo cp ./util/com.mathworks.mjs.plist /Library/LaunchDaemons
```

- 4** Open the copied `.plist` file in a text editor. Ensure that the leading part of the `StandardOutPath` and `StandardErrorPath` fields match the `LOGBASE` value as defined in the



mjs\_def.sh file. For example, if LOGBASE is /var/log/mjs, then you must define StandardOutPath and StandardErrorPath as follows:

```
<key>StandardOutPath</key>
<string>/var/log/mjs/launchctl.stdout</string>
<key>StandardErrorPath</key>
<string>/var/log/mjs/launchctl.stderr</string>
```

- 5 Restart your machine and observe that mjs is running using nodestatus:

```
cd matlabroot/toolbox/parallel/bin
./nodestatus
```

## Validate Installation with MATLAB Job Scheduler

To verify that your MATLAB Parallel Server products are installed and configured correctly, create a cluster profile and validate it. For instructions, see “Connect the MATLAB Client to the MATLAB Parallel Server Cluster” on page 3-20. You can specify the number of workers to use when validating your profile, to avoid occupying the whole cluster. If your validation does not pass, contact the MathWorks Install Support Team, or see “Troubleshoot Common Problems” on page 2-18.

After you create a cluster profile, you can make any modifications appropriate for your applications, such as NumWorkersRange, AttachedFiles, or AdditionalPaths. To save your profile for other users, in the Cluster Profile Manager, select the profile and click **Export**, then save your profile to a file in a convenient location. Later, when running the Cluster Profile Manager, other users can import your profile by clicking **Import**. For more information about cluster profiles, see “Discover Clusters and Use Cluster Profiles” (Parallel Computing Toolbox).

## See Also

### Related Examples

- “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16

### More About

- “MATLAB Job Scheduler Cluster Customization”
- “Troubleshooting in MATLAB Parallel Server”
- “Get Started with MATLAB Parallel Server”

## Configure for HPC Pack

### In this section...

“Configure Cluster for Microsoft HPC Pack” on page 3-38

“Configure Client Computer for HPC Pack” on page 3-38

“Validate Installation Using Microsoft HPC Pack” on page 3-39

### Configure Cluster for Microsoft HPC Pack

Follow these instructions to configure your MATLAB Parallel Server installation to work with Microsoft HPC Pack or Compute Cluster Server (CCS). In the following instructions, *matlabroot* refers to the MATLAB installation location.

Supported versions: MATLAB Compute Cluster Server 2003, Windows HPC Server 2008, Windows HPC Server 2008 R2, Microsoft HPC Pack 2012, Microsoft HPC Pack 2012 R2, and Microsoft HPC Pack 2016.

---

**Note** If you are using an HPC Pack in a network share installation, the network share location must be in the “Intranet” zone. You might need to adjust the Internet Options for your cluster nodes and add the network share location to the list of Intranet sites.

---

- 1 Log in on the cluster head node as a user with administrator privileges.
- 2 Open a command window with administrator privileges and run the following file command

```
matlabroot\toolbox\parallel\bin\MicrosoftHPCServerSetup.bat -cluster
```

If you are using an HPC Pack in a network share installation, this command performs some of the setup required for all machines in the cluster. The location of the MATLAB installation must be the same on every cluster node. If you are not using an HPC Pack in a network share installation, then you must run this command on every cluster node.

---

**Note** If you need to override the script default values, modify the values defined in `MicrosoftHPCServerSetup.xml` before running `MicrosoftHPCServerSetup.bat`. Use the `-def_file` argument to the script when using a `MicrosoftHPCServerSetup.xml` file in a custom location. For example:

```
MicrosoftHPCServerSetup.bat -cluster -def_file <filename>
```

You modify the file only on the node where you actually run the script.

An example of one of the values you might set is for `CLUSTER_NAME`. If you provide a friendly name for the cluster in this parameter, it is recognized by MATLAB’s discover clusters feature and displayed in the resulting cluster list.

---

### Configure Client Computer for HPC Pack

This configuring applies to all versions of HPC Pack.

---

**Note** If you are using HPC Pack in a network share installation, the network share location must be in the “Intranet” zone. You might need to adjust the Internet Options for your cluster nodes and add the network share location to the list of Intranet sites.

---

- 1 Open a command window with administrator privileges and run the following file command

```
matlabroot\toolbox\parallel\bin\MicrosoftHPCServerSetup.bat -client
```

This command performs some of the setup required for a client machine.

---

**Note** If you need to override the default values the script, modify the values defined in `MicrosoftHPCServerSetup.xml` before running `MicrosoftHPCServerSetup.bat`. Use the `-def_file` argument to the script when using a `MicrosoftHPCServerSetup.xml` file in a custom location. For example:

```
MicrosoftHPCServerSetup.bat -client -def_file <filename>
```

---

- 2 To submit jobs or discover the cluster from MATLAB, the Microsoft HPC Pack client utilities must be installed on your MATLAB client machine. If they are not already installed and up to date, ask your system administrator for the correct client utilities to install. The utilities are available from Microsoft download center.

If you have installed multiple versions of the Microsoft HPC Pack client utilities, MATLAB uses the most recent install. To configure MATLAB to use a specific install, set the environment variable 'MATLAB\_HPC\_SERVER\_HOME' to the install location of the client utilities you want to use.

## Validate Installation Using Microsoft HPC Pack

This procedure verifies that your parallel computing products are installed and configured correctly for using Microsoft Windows HPC Pack or Compute Cluster Server (CCS).

### Step 1: Create a Cluster Profile

In this step you create a cluster profile to use in subsequent steps. To create a cluster profile, try discovering your cluster. On the **Home** tab, in the **Environment** area, select **Parallel > Discover Clusters**. For more information, see “Discover Clusters” (Parallel Computing Toolbox). Alternatively, if your cluster is not available for discovery, follow these steps.

- 1 Start the Cluster Profile Manager. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- 2 Create a new profile in the Cluster Profile Manager by selecting **Add Cluster Profile > HPC Server**.
- 3 With the new profile selected in the list, click **Rename** and edit the profile name to be `HPCtest`. Press **Enter**.
- 4 In the Properties tab, provide text for the following fields:
  - a Set the **Description** field to `For testing installation with HPC Server`.
  - b Set the **NumWorkers** field to the number of workers you want to run the validation tests on, within the limitation of your licensing.

- c Set the **Host** field to the name of the host on which your scheduler is running. Depending on your network, this might be a simple host name, or it might have to be a fully qualified domain name.


Note: The following four property settings (**JobStorageLocation**, **ClusterMatlabRoot**, **ClusterVersion**, and **UseSOAJobSubmission**) are optional, and need to be set in the profile here only if you did not run **MicrosoftHPCServerSetup.bat** as described in “Configure Cluster for Microsoft HPC Pack” on page 3-38, or if you want to override the setting established by that script.

- d Set the **JobStorageLocation** to the location where you want job and task data to be stored. This must be accessible to all the worker machines.

**Note** **JobStorageLocation** should not be shared by parallel computing products running different versions; each version on your cluster should have its own **JobStorageLocation**.

- e Set the **ClusterMatlabRoot** to the installation location of the MATLAB to be executed by the worker machines, as determined in Chapter 1 of the installation instructions.
- f Set the **ClusterVersion** field to **HPCServer** or **CCS**.
- g If you want to test SOA job submissions on an HPC Server cluster, set **UseSOAJobSubmission** to **true**. If you plan to use SOA job submissions with your cluster, you should test this first without SOA submission, then later return and test it with SOA job submission. The default value is determined at runtime based on your scheduler.

So far, the dialog box should look like the following figure:




- 5 Click **Done** to save your cluster profile.

## Step 2: Validate the Configuration

In this step you validate your cluster profile, and thereby your installation. You can specify the number of workers to use when validating your profile. If you do not specify the number of workers in the **Validation** tab, then the validation will attempt to use as many workers as the value specified by the NumWorkers property on the **Properties** tab. You can specify a smaller number of workers to validate your configuration without occupying the whole cluster.

- 1 If it is not already open, start the Cluster Profile Manager from the MATLAB desktop by selecting on the **Home** tab in the **Environment** area **Parallel > Create and Manage Clusters**.
- 2 Select your cluster profile in the listing.
- 3 Click **Validation** tab.
- 4 Use the checkboxes to choose all tests, or a subset of the validation stages, and specify the number of workers to use when validating your profile.
- 5 Click **Validate**.

The Validation Results tab shows the output. The following figure shows the results of a profile that passed all validation tests.




---

**Note** If your validation does not pass, contact the MathWorks install support team.

---

If your validation passed, you now have a valid profile that you can use in other parallel applications. You can make any modifications to your profile appropriate for your applications, such as NumWorkersRange, AttachedFiles, AdditionalPaths, etc.

To save your profile for other users, select the profile and click **Export**, then save your profile to a file in a convenient location. Later, when running the Cluster Profile Manager, other users can import your profile by clicking **Import**.

## Configure for Slurm, PBS Pro, OpenPBS, LSF, TORQUE

Follow these instructions to configure your MATLAB Parallel Server installation to work with Slurm, PBS Pro, OpenPBS, LSF, and TORQUE using built-in cluster types.

You can create a cluster profile using either a built-in cluster type or `Generic`. As a best practice, use built-in cluster types where possible.

You must use a `Generic` cluster profile when:

- You connect to a cluster that does not have a built-in cluster type
- The MATLAB client and the cluster nodes do not have a shared file system
- The MATLAB client machine is unable to directly submit jobs to the third-party scheduler, or
- You need to fully customize how parallel jobs are submitted to the cluster

To configure a cluster using the `Generic` cluster type, see “Configure Using the Generic Scheduler Interface” on page 3-46.

### Create Cluster Profile and Validate Installation

This procedure verifies that the parallel computing products are installed and configured correctly on your cluster.

#### Step 1: Create a Cluster Profile

In this step you create a cluster profile to use in subsequent steps.

- 1 Start the Cluster Profile Manager. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- 2 Create a new profile in the Cluster Profile Manager by selecting **Add Cluster Profile > LSF** (or **Slurm, PBS Pro** or **Torque**, as appropriate).
- 3 With the new profile selected in the list, click **Rename** and edit the profile name to be `InstallTest`. Press **Enter**.
- 4 In the Properties tab, provide settings for the following fields:
  - a Set the **Description** field to `For testing installation`.
  - b Set the **JobStorageLocation** to the location where you want job and task data to be stored (accessible to all the worker machines if you have a shared file system).

---

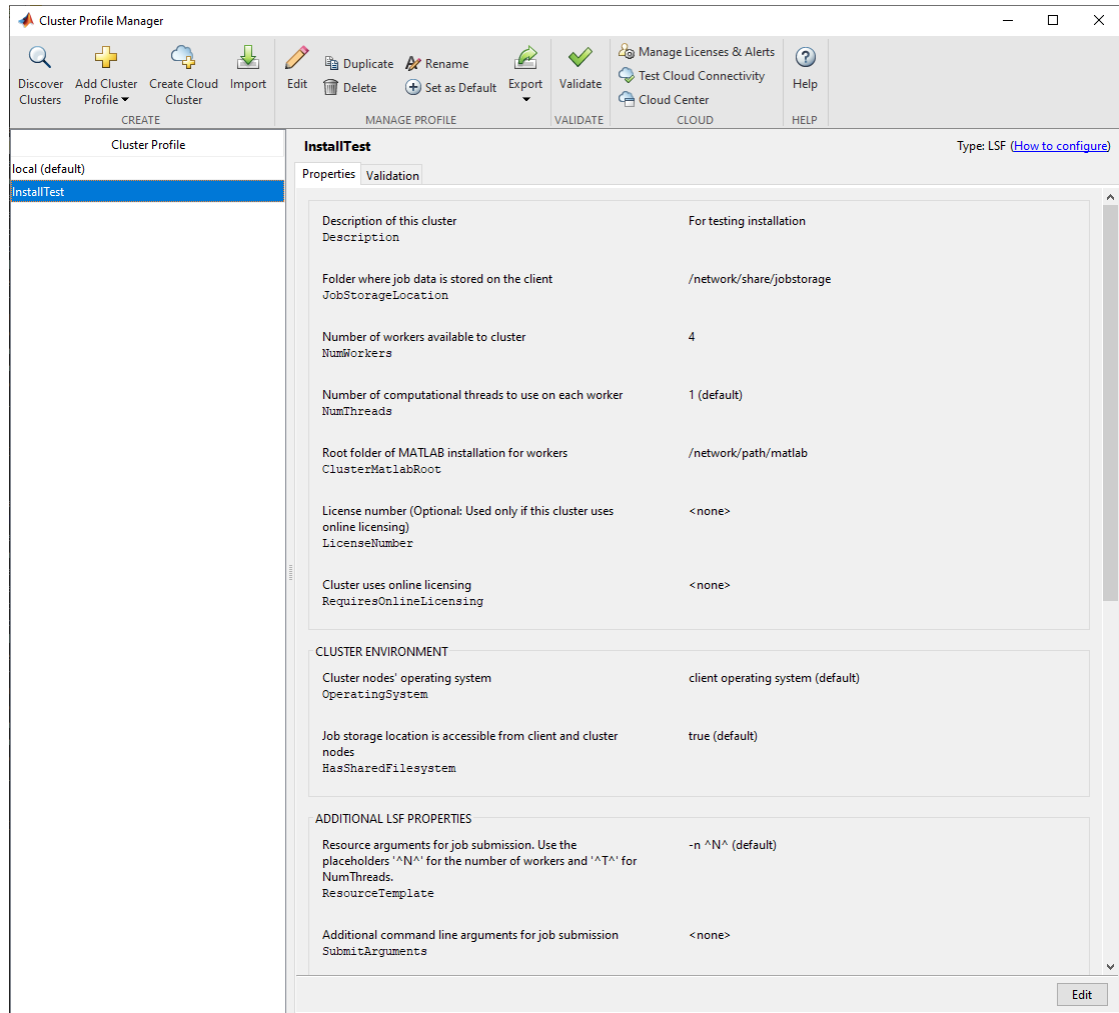
**Note** `JobStorageLocation` should not be shared by parallel computing products running different versions; each version on your cluster should have its own `JobStorageLocation`.

---

- c Set the **NumWorkers** field to the number of workers you want to run the validation tests on, within the limitation of your licensing.
- d Set the **ClusterMatlabRoot** to the installation location of the MATLAB to be executed by the worker machines.
- e Set the **SubmitArguments** to include any additional command arguments required by your particular cluster and scheduler.
- f If you are using LSF, set the **OperatingSystem** to the operating system of your worker machines.

- 9 Set **HasSharedFilesystem** to indicate if client and workers can share the same data location.

The dialog box should look something like this, or slightly different for Slurm, PBS Pro, OpenPBS, or TORQUE schedulers.



- 5 Click **Done** to save your cluster profile.


## Step 2: Validate the Cluster Profile

In this step you verify your cluster profile, and thereby your installation. You can specify the number of workers to use when validating your profile. If you do not specify the number of workers in the **Validation** tab, then the validation will attempt to use as many workers as the value specified by the **NumWorkers** property on the **Properties** tab. You can specify a smaller number of workers to validate your configuration without occupying the whole cluster.

- 1 If it is not already open, start the Cluster Profile Manager from the MATLAB desktop. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- 2 Select your cluster profile in the listing.
- 3 Click **Validation** tab.

- 4 Use the checkboxes to choose all tests, or a subset of the validation stages, and specify the number of workers to use when validating your profile.
- 5 Click **Validate**.

The Validation Results tab shows the output. The following figure shows the results of a profile that passed all validation tests.




---

**Note** If your validation does not pass, contact the MathWorks install support team.

---

If your validation passed, you now have a valid profile that you can use in other parallel applications. You can make any modifications to your profile appropriate for your applications, such as `NumWorkersRange`, `AttachedFiles`, `AdditionalPaths`, etc.

To save your profile for other users, select the profile and click **Export**, then save your profile to a file in a convenient location. Later, when running the Cluster Profile Manager, other users can import your profile by clicking **Import**.

## Configure LSF Scheduler on Windows Cluster

If your cluster is already set up to use `mpiexec` and `smpd`, you can use Parallel Computing Toolbox software with your existing configuration if you are using a compatible MPI implementation library (as defined in `matlabroot\toolbox\parallel\mpi\mpiLibConf.m`). However, if you do not have `mpiexec` on your cluster and you want to use it, you can use the `mpiexec` software shipped with the parallel computing products.

For further information about `mpiexec` and `smpd`, see the MPICH home page. For user's guides and installation instructions on that page, select **Documentation > User Docs**.

In the following instructions, `matlabroot` refers to the MATLAB installation location.

To use `mpiexec` to distribute a job, the `smpd` service must be running on all nodes that will be used for running MATLAB workers.



---

**Note** The `smpd` executable does not support running from a mapped drive. Use either a local installation, or the full UNC path name to the executable. Microsoft Windows Vista does not support the `smpd` executable on network share installations, so with Vista the installation must be local.

---

- 1 Log in as a user with administrator privileges.
- 2 Start `smpd` by typing in a DOS command window:

```
matlabroot\bin\win64\smpd -install
```

This command installs the service and starts it. As long as the service remains installed, it will start each time the node boots.

- 3 If this is a worker machine and you did not run the installer on it to install MATLAB Parallel Server software (for example, if you are running MATLAB Parallel Server software from a shared installation), execute the following command in a DOS command window.

```
matlabroot\bin\matlab.bat -install_vcrt
```

This command installs the Microsoft run-time libraries needed for running jobs with your scheduler.

- 4 If you are using Windows firewalls on your cluster nodes, execute the following in a DOS command window.

```
matlabroot\toolbox\parallel\bin\addMatlabToWindowsFirewall.bat
```

This command adds MATLAB as an allowed program. If you are using other firewalls, you must configure them to make similar accommodation.

- 5 Log in as the user who will be submitting jobs for execution on this node.
- 6 Register this user to use `mpiexec` by typing:

```
matlabroot\bin\win64\mpiexec -register
```

- 7 Repeat steps 5-6 for all users who will run jobs on this machine.
- 8 Repeat all these steps on all Windows nodes in your cluster.

## Configure Windows Firewalls on Server

If you are using Windows firewalls on your cluster nodes,

- 1 Log in as a user with administrative privileges.
- 2 Execute the following in a DOS command window.

```
matlabroot\toolbox\parallel\bin\addMatlabToWindowsFirewall.bat
```

This command adds MATLAB as an allowed program. If you are using other firewalls, you must configure them for similar accommodation.

## Configure Using the Generic Scheduler Interface

The generic scheduler interface provides flexibility to configure the interaction of the MATLAB client, MATLAB workers, and a third-party scheduler. Use the generic scheduler interface when you want complete customization for interfacing MATLAB with your scheduler setup.

You can create a cluster profile using either a built-in cluster type or `Generic`. As a best practice, use built-in cluster types where possible.

You must use a `Generic` cluster profile when:

- You connect to a cluster that does not have a built-in cluster type
- The MATLAB client and the cluster nodes do not have a shared file system
- The MATLAB client machine is unable to directly submit jobs to the third-party scheduler, or
- You need to fully customize how parallel jobs are submitted to the cluster

To configure a cluster using a built-in cluster type, see “Configure for Slurm, PBS Pro, OpenPBS, LSF, TORQUE” on page 3-42, “Configure a Hadoop Cluster” on page 3-62, or “Configure for HPC Pack” on page 3-38.

### Interface with Third-Party Schedulers

The generic scheduler interface provides a means of getting tasks from your Parallel Computing Toolbox client session to your scheduler and cluster nodes. To achieve this, you must provide your MATLAB client with a set of plugin scripts. The scripts contain instructions specific to your cluster infrastructure, such as how to communicate with the job scheduler, and how to transfer job and task data to cluster nodes.

#### Support Scripts

To help you to interface with your scheduler, MathWorks provides add-ons, or plugins, for the following third-party schedulers, which you can download and install from the Add-On Manager.

- Parallel Computing Toolbox plugin for MATLAB Parallel Server with Slurm
- Parallel Computing Toolbox plugin for MATLAB Parallel Server with IBM Platform LSF
- Parallel Computing Toolbox plugin for MATLAB Parallel Server with Grid Engine
- Parallel Computing Toolbox plugin for MATLAB Parallel Server with PBS
- Parallel Computing Toolbox plugin for MATLAB Parallel Server with HTCondor

On the **Home** tab, in the **Environment** area, click **Add-Ons**. In the Add-On Explorer, search for the add-on and install it. The installer includes a wizard to guide you through creating a cluster profile for your cluster configuration.

Alternatively, you can download the plugins from these locations:

- SLURM
- LSF
- Grid Engine family
- PBS family

- HTCondor

To install a plugin, double click it or drag it and drop it in MATLAB.

Each installer creates scripts for three possible submission modes:

- Shared - When the client can submit directly to the scheduler, and the client and the cluster machines have a shared file system.
- Remote - When the client and cluster machines have a shared file system, but the client machine cannot submit directly to the scheduler, such as when the client utilities of the scheduler are not installed. In this case, a remote host submits commands to the scheduler using the `ssh` protocol.
- Nonshared - When the client and cluster machines do not have a shared file system. This mode uses the `ssh` protocol to submit commands to the scheduler using a remote host, and it uses the `sftp` protocol to copy job and task files to the cluster file system.

Each submission mode has its own subfolder within the installation folder. These subfolders contain a `README` file that provides specific instructions on how to use the scripts. Before using the scripts, decide which submission mode describes your network setup.

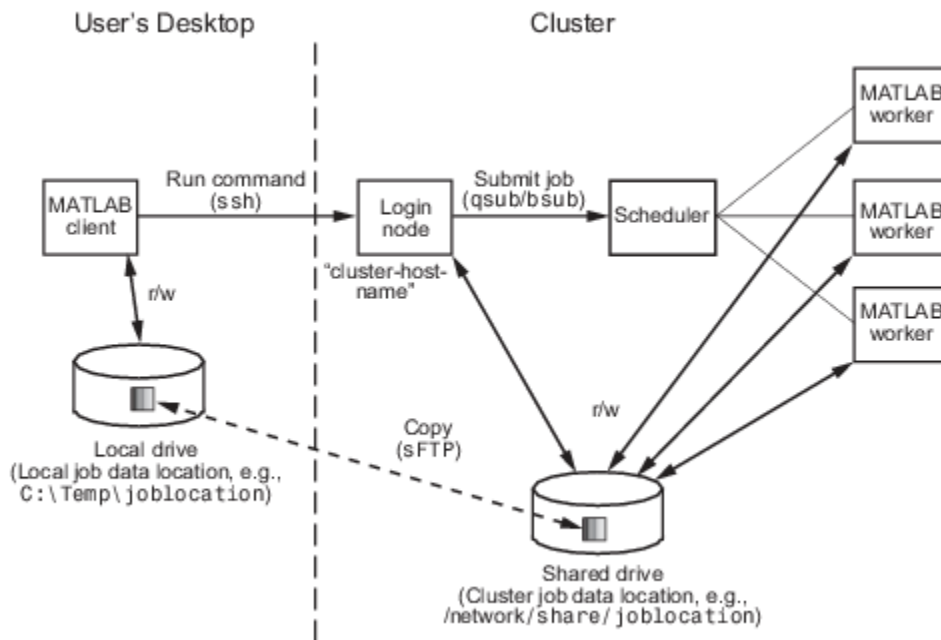
If you want to customize the behavior of the plugin scripts, you can set additional properties, such as `AdditionalSubmitArgs`. For more information, see “Customize Behavior of Sample Plugin Scripts” on page 3-56.

If your scheduler or cluster configuration is not supported by one of the support packages, it is recommended that you modify the scripts of one of these packages. For more information on how to write a set of plugin scripts for generic schedulers, see “Plugin Scripts for Generic Schedulers” (Parallel Computing Toolbox).

## Create a Generic Cluster Profile

### Sample Setup for LSF

This example shows how to set up your cluster profile to use the generic scheduler interface. It shows the set up of an LSF scheduler in a network without a shared file system between the client and the cluster machines. The following diagram illustrates the cluster setup:



In this type of configuration, job data is copied from the client host running a Windows operating system to a host on the cluster (cluster login node) running a UNIX® operating system. From the cluster login node, the LSF `bsub` command submits the job to the scheduler. When the job finishes, its output is copied back to the client host.

### Requirements

The setup must meet the following conditions:

- The client node and cluster login node must support `ssh` and `sftp`.
- The cluster login node must be able to call the `bsub` command to submit a job to an LSF scheduler. You can find more about this in the `README` file in the `nonshared` subfolder within the installation folder.

### Run the LSF Installer

- 1 Download the installer for LSF from [here](#).
- 2 Run the installer by opening the file from within your MATLAB client.

The installer downloads the plugin scripts to the `nonshared` subfolder within the installation folder. The installer wizard guides you through the steps to create and validate a cluster profile.

### Validate Cluster Profile and Installation

You can specify the number of workers to use when validating your profile. If you do not specify the number of workers in the **Validation** tab, then the validation process attempts to use as many workers as the value specified by the **NumWorkers** property on the **Properties** tab. You can specify a smaller number of workers to validate your configuration without occupying the whole cluster.

- 1 Start the Cluster Profile Manager from the MATLAB desktop. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.

- 2 Select your cluster profile in the listing.
- 3 Click the **Validation** tab.
- 4 Use the checkboxes to choose all tests, or a subset of the validation stages, and specify the number of workers to use when validating your profile.
- 5 Click **Validate**.

The **Validation** results tab shows the output. The following figure shows the results of a profile that passed all validation tests.

Stage	Status	Description
<input checked="" type="checkbox"/> Cluster connection test (parcluster)	Passed	
<input checked="" type="checkbox"/> Job test (createJob)	Passed	
<input checked="" type="checkbox"/> SPMD job test (createCommunicatingJob)	Passed	Job ran with 8 workers.
<input checked="" type="checkbox"/> Pool job test (createCommunicatingJob)	Passed	Job ran with 8 workers.
<input checked="" type="checkbox"/> Parallel pool test (parpool)	Passed	Parallel pool ran with 8 workers.

Number of workers to use: Use default

STAGE DETAILS

Stage started at 16:31:18. Completed in 0 min 11 sec.

**Note** If your validation fails any stage, contact the MathWorks install support team.

If your validation passes, you have a valid profile that you can use in other parallel applications. You can make any modifications to your profile that are appropriate for your applications, such as **NumWorkersRange**, **AttachedFiles**, or **AdditionalPaths**.

To save your profile for other users, select the profile, and click **Export**. Then save your profile to a file in a convenient location. When running the Cluster Profile Manager, other users can import your profile by clicking **Import**.

To learn how to distribute a generic cluster profile and plugin scripts for others to use, see “Distribute a Generic Cluster Profile and Plugin Scripts” on page 3-54.

## Manually Configure a Cluster Profile

If you want to modify an existing generic cluster profile, you can configure the profile manually. If you are creating the profile for the first time and you are using one of the provided support scripts, use the installer wizard instead: “Run the LSF Installer” on page 3-48. The following steps reproduce manually the configuration performed by the installer. You can modify any of these options depending on your setup.

- 1 Start a MATLAB session on the client host.
- 2 Start the Cluster Profile Manager from the MATLAB desktop. On the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- 3 Create a new profile in the Cluster Profile Manager by selecting **Add > Custom > Generic**.
- 4 With the new profile selected in the list, select **Rename** and change the profile name to `InstallTest`. Press **Enter**.
- 5 In the **Properties** tab, select **Edit** and provide settings for the following fields:
  - a Set the **Description** field to *For testing installation*.
  - b Set the **JobStorageLocation** to the location where you want job and task data to be stored on the client machine (not the cluster location), for example, `C:\Temp\joblocation`.

You must not share **JobStorageLocation** among parallel computing products running different versions. Each version on your cluster must have its own **JobStorageLocation**.
  - c Set **NumWorkers** to the number of workers for which you want to test your installation.
  - d Set **NumThreads** to the number of threads to use on each worker.
  - e Set **ClusterMatlabRoot** to the installation location of MATLAB to run on the worker machines.
  - f If the cluster uses online licensing, set **RequiresOnlineLicensing** to true.
  - g If you set **RequiresOnlineLicensing** to true, enter your **LicenseNumber**.
  - h Set **OperatingSystem** to the operating system of your cluster worker machines.
  - i Set **HasSharedFilesystem** to false. This setting indicates that the client node and worker nodes cannot share the same data location.
  - j Set the **PluginScriptsLocation** to the location of your plugin scripts. In this example, the location is the nonshared subfolder within the LSF installation folder. As part of using the example scripts in nonshared submission mode, set the properties in steps k and l.
  - k In the **AdditionalProperties** table, select **Add**. Specify a new property with name `ClusterHost`, value `cluster-host-name`, and type `String`.
  - l In the **AdditionalProperties** table, select **Add**. Specify a new property with name `RemoteJobStorageLocation`, value `/network/share/joblocation`, and type `String`.
- 6 Click **Done** to save your cluster profile changes. The dialog box looks as follows:

**InstallTest**
Type: Generic ([How to configure](#))

Properties
Validation

Description of this cluster Description	For testing installation
Folder where job data is stored on the client JobStorageLocation	C:\Temp\joblocation
Number of workers available to cluster NumWorkers	4
Number of computational threads to use on each worker NumThreads	1
Root folder of MATLAB installation for workers ClusterMatlabRoot	/network/installs/MATLAB/R2019a
License number (Optional: Used only if this cluster uses online licensing) LicenseNumber	< none >
Cluster uses online licensing RequiresOnlineLicensing	< none >

**CLUSTER ENVIRONMENT**

Cluster nodes' operating system OperatingSystem	unix
Job storage location is accessible from client and cluster nodes HasSharedFilesystem	false

**SCHEDULER INTEGRATION**

Folder containing scheduler integration scripts IntegrationScriptsLocation	C:\ProgramData\MATLAB\SupportPackages\R2019a\paralle\sl...									
Additional properties for integration scripts AdditionalProperties	<table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 33%;">Name</th> <th style="width: 33%;">Value</th> <th style="width: 33%;">Type</th> </tr> </thead> <tbody> <tr> <td>ClusterHost</td> <td>cluster-host-name</td> <td>String</td> </tr> <tr> <td>RemoteJobStorageLoc...</td> <td>/network/share/jobloc...</td> <td>String</td> </tr> </tbody> </table>	Name	Value	Type	ClusterHost	cluster-host-name	String	RemoteJobStorageLoc...	/network/share/jobloc...	String
Name	Value	Type								
ClusterHost	cluster-host-name	String								
RemoteJobStorageLoc...	/network/share/jobloc...	String								

Edit

To check that the profile works, perform a validation following the steps in “Validate Cluster Profile and Installation” on page 3-48.

## Special Configurations

Depending on your cluster architecture, you might need to perform additional tasks before you connect to your generic scheduler.

### Custom MPI builds

You can use an MPI build that differs from the one provided with Parallel Computing Toolbox. For more information about using this option with the generic scheduler interface, see “Use Different MPI Builds on UNIX Systems” on page 2-4.

### Run Communicating Jobs with the Grid Engine Family

The sample scripts for Grid Engine family rely on the presence of a `matlab` parallel environment. Parallel environments (PE) are programming environments designed for parallel computing in clusters. To run communicating jobs with MATLAB Parallel Server and a Grid Engine family cluster, you must establish a `matlab` parallel environment.

#### Create the Parallel Environment

The following steps create the parallel environment, and then make it runnable on all queues. As a best practice, perform these steps on the head node of your cluster. Some steps require administrator access.

- 1 Download and run the installer for Grid Engine from Grid Engine family.
- 2 Navigate to the location of the relevant plugin scripts for your submission mode in the installation folder.
- 3 Modify the contents of `matlabpe.template` to use the number of slots you want and the correct location of the `startmatlabpe.sh` and `stopmatlabpe.sh` files. These files can exist in a shared location accessible by all hosts, or you can copy them to the same location on each host. You can also change other values or add additional values to `matlabpe.template` to suit your cluster. For more information, refer to the `sgpe` documentation provided with your scheduler.
- 4 Add the `matlab` parallel environment, using a shell command such as:

```
qconf -Ap matlabpe.template
```

- 5 Make the `matlab` parallel environment runnable on all queues:

```
qconf -mq all.q
```

This command brings up a text editor for you to make changes. Search for the line `pe_list`, and add `matlab`.

- 6 Ensure you can submit a trivial job to the PE:

```
$ echo "hostname" | qsub -pe matlab 1
```

- 7 Use `qstat` to check that the job runs correctly, and check that the output file contains the name of the host that ran the job. The default file name for the output file is `~/STDIN.o###`, where `###` is the Grid Engine job number.



---

**Note** If you change the name of the parallel environment to something other than `matlab`, also change the submit functions.

---

### **Configure Firewalls on Windows Cluster**

If you are using Windows firewalls on your cluster nodes, you can add MATLAB as an allowed program.

In the following instructions, *matlabroot* refers to the MATLAB installation location.

- 1** Log in as a user with administrative privileges.
- 2** Execute the following script in a DOS command window:

```
matlabroot\toolbox\parallel\bin\addMatlabToWindowsFirewall.bat
```

If you are using other firewalls, you must configure these separately to add MATLAB as an allowed program.

### **See Also**

#### **Related Examples**

- “Plugin Scripts for Generic Schedulers” (Parallel Computing Toolbox)

## Distribute a Generic Cluster Profile and Plugin Scripts

You can distribute a Generic cluster profile and plugin scripts for others to use. If necessary, you can create a Generic cluster profile and plugin scripts as follows:

- 1 Install the appropriate support package for your third-party scheduler (see “Support Scripts” on page 3-46).
- 2 Use the Generic Profile Wizard to create a Generic cluster profile with the default MATLAB plugin scripts.

### Decide How Users Access the Plugin Scripts

The **PluginScriptsLocation** property of your Generic cluster profile specifies the folder containing the plugin scripts that your cluster profile uses to submit MATLAB jobs to the cluster. Other users must have access to these plugin scripts or a copy of them in order to submit jobs to the cluster. As the one distributing a Generic cluster profile and plugin scripts, you must decide how other users will access these scripts.

- If you prefer to put the plugin scripts in a read-only shared location, follow the steps in “Shared PluginScriptsLocation Folder” on page 3-54. This option simplifies subsequent steps and allows any changes you make to the plugin scripts to take effect immediately for all users.
- If you prefer to give other users their own copy of your plugin scripts, follow the steps in “Distribute Copies of the PluginScriptsLocation Folder” on page 3-55.

### Shared PluginScriptsLocation Folder

If other users have read access to the **PluginScriptsLocation** folder specified in your cluster profile, they need only a copy of your profile to submit MATLAB jobs to the cluster.

---

**Note** If you have moved your PluginScriptsLocation folder to a shared location, remember to update the **PluginScriptsLocation** property of your cluster profile before continuing.

---

- To distribute a copy of your profile, you must:
  - 1 Open MATLAB and navigate to **Home > Parallel > Create and Manage Clusters** to open the Cluster Profile Manager.
  - 2 Select your profile in the list and click **Export**.
  - 3 Choose a name for the `.settings` file, which contains your profile, and click **Save**.
  - 4 Send a copy of the `.settings` file to other users.
- To import your profile, other users must:
  - 1 Save the `.settings` file to a location of their choice.
  - 2 Open MATLAB and navigate to **Home > Parallel > Create and Manage Clusters**.
  - 3 Open the Cluster Profile Manager and click **Import**.
  - 4 Select the `.settings` file and click **Open**. A copy of the profile appears in their profile list.
  - 5 Check that the profile works by selecting the profile in the Cluster Profile Manager and clicking **Validate**.

## Distribute Copies of the PluginScriptsLocation Folder

If you cannot or prefer not to put your plugin scripts in a shared location, you can give users a copy of your plugin scripts in addition to your profile.

---

**Note** If you make changes to your plugin scripts, you will have to distribute copies of your updated plugin scripts for the changes to take effect for other users.

---

- To distribute a copy of your profile and plugin scripts, you must:
  - 1 Open MATLAB and navigate to **Home > Parallel > Create and Manage Clusters** to open the Cluster Profile Manager.
  - 2 Select your profile in the list and click **Export**.
  - 3 Choose a name for the `.settings` file, which contains the exported profile, and click **Save**.
  - 4 Send other users a copy of
    - The `.settings` file.
    - The **PluginScriptsLocation** folder and all files therein.
- To import your profile and plugin scripts, other users must:
  - 1 Save all files to a location of their choice.
  - 2 Open MATLAB and navigate to **Home > Parallel > Create and Manage Clusters**.
  - 3 Open the Cluster Profile Manager and click **Import**.
  - 4 Select the `.settings` file and click **Open**. A copy of the profile appears in their profile list.
  - 5 Select the profile in the Cluster Profile Manager and click **Edit**. Scroll down to the Scheduler Plugin section of the profile and change the **PluginScriptsLocation** property to point to their copy of the **PluginScriptsLocation** folder.
  - 6 Check that the profile works by selecting the profile in the Cluster Profile Manager and clicking **Validate**.

## Further Considerations for Clusters with Shared File Systems

If you are distributing a Generic cluster profile with the **HasSharedFileSystem** property set to true (e.g. shared or remote submission modes, see “Interface with Third-Party Schedulers” on page 3-46), the cluster machines must have read and write access to the folder specified in the **JobStorageLocation** property of the profile. Remind other users receiving the profile that they must either:

- 1 Set the **JobStorageLocation** property of their profile to a shared location, preferably one that is unique to their user name and MATLAB version.
- 2 Only create or submit jobs to the cluster when the current working folder is a shared location.

## See Also

“Interface with Third-Party Schedulers” on page 3-46 | “Support Scripts” on page 3-46

## **Customize Behavior of Sample Plugin Scripts**

When using the generic scheduler interface, you can modify the behavior of the plugin scripts by setting additional properties for a generic cluster profile or object using `AdditionalProperties`. For more information on the generic scheduler interface, see “Configure Using the Generic Scheduler Interface” on page 3-46.

The sample plugin scripts allow you to set the following properties.

**Properties for All Schedulers**

<b>Property</b>	<b>Description</b>	<b>Type</b>
<code>AdditionalSubmitArgs</code>	Additional scheduler arguments for job submission. The sample plugin scripts add the value of this property to the scheduler submission string.	String
<code>AuthenticationMode</code>	Option to indicate how you are authenticated when you connect to the cluster, specified as one of the following: <ul style="list-style-type: none"> <li>• "Password" - the client prompts you for your SSH password. Your user name is specified by the <code>Username</code> additional property.</li> <li>• "IdentityFile" - the client uses the identity file specified by the <code>IdentityFile</code> additional property.</li> <li>• "Agent" - the client interfaces with an SSH agent running on the client machine. If the client is running Windows, only the Pageant SSH agent is supported.</li> </ul>	String
<code>ClusterHost</code>	Host name of the cluster machine that has the scheduler utilities to submit jobs. This property is available in remote and nonshared submission modes only.  The cluster machine must run Linux.	String
<code>IdentityFile</code>	Location on the client machine of the SSH identity file that identifies you in <code>ClusterHost</code> .	String
<code>IdentityFileHasPassphrase</code>	Set this property to <code>true</code> if <code>IdentityFile</code> requires a passphrase.	Logical
<code>RemoteJobStorageLocation</code>	Location to store job files on the cluster. This property is available in nonshared submission mode only.	String

<b>Property</b>	<b>Description</b>	<b>Type</b>
UseIdentityFile	Option to use an identity file. Set this property to <code>true</code> if you want to use an SSH identity file to log in to <code>ClusterHost</code> . If you set this property, then also set <code>IdentityFile</code> and <code>IdentityFileHasPassphrase</code> .	Logical
Username	User name to log in to <code>ClusterHost</code> with.	String
UseUniqueSubfolders	Option to use unique subfolders. Set this property to <code>true</code> if you want MATLAB to store job files under different subfolders based on the user name and MATLAB version. Doing so helps to prevent conflicts between jobs submitted from different users and MATLAB versions.	Logical

**Properties for PBS Only**

Property	Description
ProcsPerNode	<p>Number of processors per node, specified as a finite positive integer scalar.</p> <p>When you submit a job to the cluster, the number of cores per node that MATLAB requests is guaranteed to be less than or equal to ProcsPerNode. Set ProcsPerNode equal to the maximum number of processors you want MATLAB to request from each cluster node.</p> <p>MATLAB requests the smallest number of cores per node required to run the job.</p> <ul style="list-style-type: none"> <li>• If the NumThreads property of the cluster is less than or equal to ProcsPerNode, MATLAB requests NumThreads processors per worker, then maximizes the number of workers per node. For example if NumThreads is 16 and ProcsPerNode is 5, MATLAB requests 15 cores, the smallest multiple of 5 that is less than NumThreads.</li> <li>• If NumThreads of the cluster is greater than ProcsPerNode, MATLAB requests ProcsPerNode processors per node.</li> </ul> <p>When NumThreads is greater than ProcsPerNode, you might encounter performance issues. As a best practice, set NumThreads less than or equal to ProcsPerNode. For more information, see “Edit Number of Workers and Cluster Settings” (Parallel Computing Toolbox).</p>

When you create a generic cluster profile for the first time with the Generic Profile Wizard, the installer guides you through setting some of these properties. For more information on the installer, see “Interface with Third-Party Schedulers” on page 3-46.

You can also define your own properties to customize the behavior of the plugin scripts. To do so, you must modify the plugin scripts. For more information, see “Adding User Customization” (Parallel Computing Toolbox).

**Set Additional Properties**

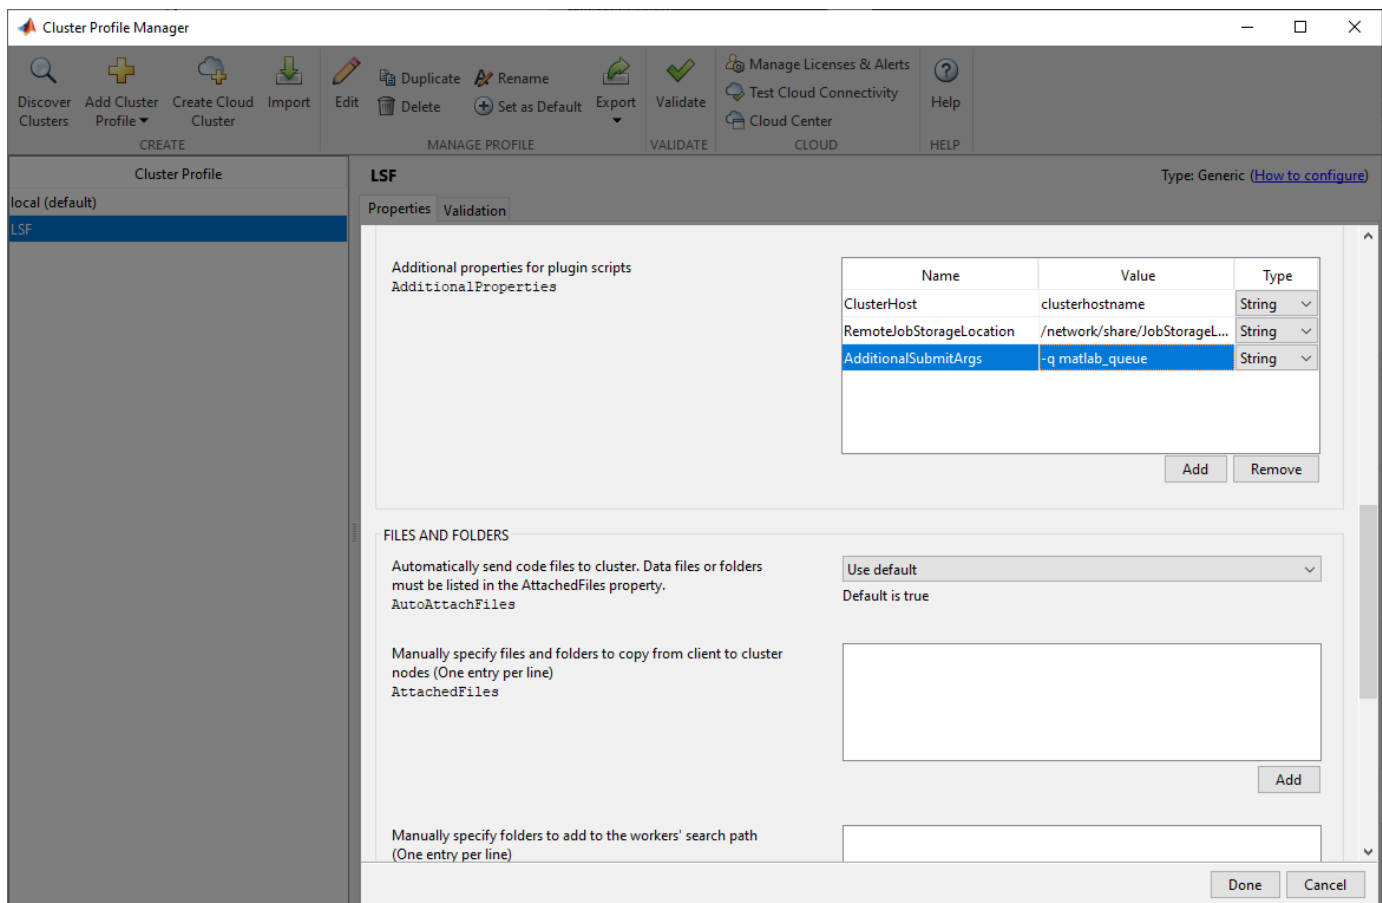
You can set additional properties in the cluster profile, with the Cluster Profile Manager, or programmatically.

**Set Properties in Cluster Profile**

When you set the additional properties in the cluster profile, the properties apply every time you use the cluster.

- In the MATLAB toolstrip, on the **Home** tab, in the **Environment** area, select **Parallel > Create and Manage Clusters**.
- In the Cluster Profile Manager, click on the generic profile that you want to modify.
- Click **Edit** at the bottom-right.
- Go to the **AdditionalProperties** table.
- To add a new property, click **Add**. To modify an existing property, click on the property.

The following image shows an example of setting `AdditionalSubmitArgs` for an LSF cluster profile.



### Set Properties Programmatically

You can set additional properties programmatically by accessing the `AdditionalProperties` of a generic cluster object. Note that this action does not update the cluster profile and the properties only apply to that particular cluster object. The following is an example of how to set `AdditionalSubmitArgs` for an LSF cluster.

First, create a cluster object by using the `parcluster` function. In the following code, change `MyLSFCluster` to the name of your cluster profile.

```
c = parcluster("MyLSFCluster");
```

Next, set `AdditionalSubmitArgs` so that the plugin scripts use a different job queue.



```
c.AdditionalProperties.AdditionalSubmitArgs = '-q matlab_queue';
```

With this change, MATLAB passes the additional arguments to the scheduler when you submit a job. For example, submit a batch job.

```
job = batch(c, "myScript");
```

## See Also

### Related Examples

- “Configure Using the Generic Scheduler Interface” on page 3-46
- “Plugin Scripts for Generic Schedulers” (Parallel Computing Toolbox)

## Configure a Hadoop Cluster

### In this section...

“Cluster Configuration” on page 3-62

“Client Configuration” on page 3-62

“Kerberos Authentication” on page 3-62

“Hadoop Version Support” on page 3-63

Parallel MATLAB code that contains `tall` (MATLAB) arrays and `mapreduce` (MATLAB) functions can be submitted to the Hadoop cluster from suitably configured MATLAB clients.

To configure the client to run MATLAB code on the cluster, you must already be able to submit to the cluster from the intended client machine. The client machine must have a Hadoop installation that can access the cluster outside of MATLAB.

Many Hadoop distributions do not support direct access of Linux based clusters from Windows clients. Users of Windows clients typically need to set up a Linux gateway node that can be accessed from the Windows client via SSH or VNC. The cluster can then be accessed from this gateway node.

### Cluster Configuration

- 1 Integrate MATLAB Parallel Server with your cluster infrastructure. For instructions, see “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12.
- 2 If your cluster requires Kerberos authentication, ensure your MATLAB Parallel Server installations have been configured correctly. For instructions, see “Kerberos Authentication” on page 3-62.

### Client Configuration

- 1 Ensure your client can access the Hadoop cluster outside MATLAB.
- 2 Ensure your client MATLAB installation has been configured for Kerberos authentication if your cluster requires it. For instructions, see “Kerberos Authentication” on page 3-62.

To access the cluster from within MATLAB, set up a `parallel.cluster.Hadoop` (Parallel Computing Toolbox) object using the following statements.

```
setenv('HADOOP_HOME', '/path/to/hadoop/install')
cluster = parallel.cluster.Hadoop;
```

Use `mapreducer` (MATLAB) to specify `mapreduce` to run on the Hadoop cluster object.

For examples of how to run parallel MATLAB code on your Hadoop cluster, see “Run mapreduce on a Hadoop Cluster” (Parallel Computing Toolbox) and “Use Tall Arrays on a Spark Enabled Hadoop Cluster” (Parallel Computing Toolbox).

### Kerberos Authentication

If the cluster uses Kerberos authentication that requires the Oracle® Java Cryptography Extension, you must configure all installations of MATLAB and MATLAB Parallel Server. If you are using

Hortonworks® or Cloudera® distributions, it is likely that you need to complete these configuration steps.

The configuration instructions are the same for client and worker MATLAB installations.

Starting in R2018b, configure your MATLAB installation by enabling the appropriate security policy in the Java installation.

- 1 In the MATLAB Editor, open the file `${MATLAB_ROOT}/sys/java/jre/${ARCH}/jre/lib/security/java.security`.

- 2 Change the line

```
#crypto.policy=unlimited
```

to

```
crypto.policy=unlimited
```

For previous releases, you must download additional security files from Oracle.

- 1 Download the Oracle Java Cryptography Extension zip file from the Oracle Java SE page.
- 2 Unzip the downloaded zip file into a temporary folder.
- 3 Replace the files `local_policy.jar` and `US_export_policy.jar` in the folder `${MATLABROOT}/sys/java/jre/${ARCH}/jre/lib/security` with the downloaded versions.

## Hadoop Version Support

- MATLAB mapreduce is supported on Hadoop 2.x clusters. Note that support for Hadoop 1.x clusters has been removed.
- MATLAB tall arrays are supported on Spark™ enabled Hadoop 2.x clusters. You can use tall arrays on Spark enabled Hadoop clusters supporting all architectures for the client, while supporting Linux and Mac architectures for the cluster. This includes cross-platform support.

Functionality	Result	Use Instead	Compatibility Considerations
Support for running MATLAB mapreduce on Hadoop 1.x clusters has been removed.	Errors	Use clusters that have Hadoop 2.x installed to run MATLAB mapreduce.	Migrate MATLAB mapreduce code that runs on Hadoop 1.x to Hadoop 2.x.

## See Also

`parallel.cluster.Hadoop`

## Related Examples

- “Install and Configure MATLAB Parallel Server for Third-Party Schedulers” on page 3-12
- “Use Tall Arrays on a Spark Enabled Hadoop Cluster” (Parallel Computing Toolbox)
- “Run mapreduce on a Hadoop Cluster” (Parallel Computing Toolbox)
- “Read and Analyze Hadoop Sequence File” (MATLAB)



# Admin Center

---

- “Start Admin Center” on page 4-2
- “Set Up Resources” on page 4-3
- “Test Connectivity” on page 4-9
- “Export and Import Sessions” on page 4-12
- “Prepare for Cluster Profiles” on page 4-13


## Start Admin Center

Admin Center is a graphical user interface with which you can control and monitor the MATLAB Parallel Server processes of a MATLAB Job Scheduler cluster. Admin center does not support any common job scheduler (CJS) clusters or third-party schedulers.

You start Admin Center outside a MATLAB session by executing the following:

- `matlabroot/toolbox/parallel/bin/admincenter` (on UNIX operating systems)
- `matlabroot\toolbox\parallel\bin\admincenter.bat` (on Microsoft Windows operating systems)

The first time you start Admin Center, you see the following welcome dialog box.



A new session of Admin Center has no cluster hosts listed, so the usual first step is to identify the hosts you want to include in your listing. To do this, click **Add or Find**. Further information continues in the next section, “Set Up Resources” on page 4-3.

If you start Admin Center again on the same host, your previous session for that machine is loaded; and unless the update rate is set to *never*, Admin Center performs an update immediately for the listed hosts and processes. To clear this information and start a new session, select the pull-down **File** > **New Session**.

## See Also

### More About

- “Set Up Resources” on page 4-3
- “Export and Import Sessions” on page 4-12

## Set Up Resources

### In this section...

“Add Hosts” on page 4-3

“Start mjs Service” on page 4-4

“Start a MATLAB Job Scheduler” on page 4-5

“Start Workers” on page 4-6

“Stop, Destroy, Resume, Restart Processes” on page 4-7

“Move a Worker” on page 4-7

“Update the Display” on page 4-8

### Add Hosts

To specify the hosts you want listed in Admin Center, click **Add or Find** in the Welcome dialog box, or if this is not a new session, click **Add or Find** in the Hosts module.

In the Add or Find Hosts dialog box, identify the hosts you want to add to the listing, by one of the following methods:

- Select **Enter Hostnames** and provide short host names, fully qualified domain names, or individual IP addresses for the hosts.
- Select **Enter IP Range** and provide the range of IP addresses for your hosts.

If one of the hosts you have specified is running a MATLAB Job Scheduler, Admin Center automatically finds and lists all the hosts running workers registered with that MATLAB Job Scheduler. Similarly, if you specify a host that is running a worker, Admin Center finds and lists the host running that worker’s MATLAB Job Scheduler, and then also all hosts running other workers under that MATLAB Job Scheduler.

The screenshot shows the Admin Center window with three main sections:

- Hosts:** A table with columns: Hostname, Reachable, Cores, Status, Up Since, Name, and Count. It lists two nodes: node1 (6 Cores, running since 2018-11-08 10:54) and node2 (2 Cores, running since 2018-11-08 14:31).
- MATLAB Job Scheduler:** A table with columns: Name, Hostname, Status, Up Since, and Workers. It is currently empty.
- Workers:** A table with columns: Name, Hostname, Status, Up Since, Connection, Name, and Hostname. It is currently empty.

At the bottom, there is a status bar showing "Last updated: 08/11/18 15:45" and an "Update" button set to "every 2 minutes" with an "Update Now" button.

## Start mjs Service

A host must be running the mjs service if a MATLAB Job Scheduler or worker is to run on that host. Normally, you set this up with Admin Center or command-line scripts during the installation of MATLAB Parallel Server on your cluster, as described in the installation instructions available at “Install and Configure MATLAB Parallel Server for MATLAB Job Scheduler and Network License Manager” on page 3-16.

If you want to add or remove hosts to your cluster, Admin Center allows you to start and stop the mjs service on those hosts. To start the mjs service on a group of hosts with the same platform, select all those hosts in the Hosts module, and click **Start MJS Service** in the left column of the panel.

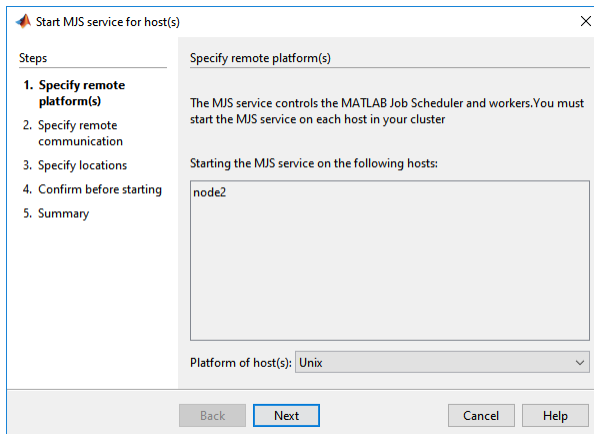
Alternative methods for starting mjs include selecting the pull-down **Hosts > Start MJS Service**, or right-clicking a listed host and selecting, **Start MJS Service**.

A dialog box leads you through the procedure of starting the mjs service on the selected hosts. There are five steps to the procedure in which you provide or confirm information for the service:

- 1 Specify remote platform — Windows or UNIX. You can start mjs on multiple hosts at the same time, but they all must be the same platform. If you have a mixed platform cluster, run the mjs startup separately for each type of platform.
- 2 Specify remote communication — Choose the protocol for communication with the hosts.
- 3 Specify locations — Specify the location of the MATLAB installation and the mjs\_def file for the hosts.
- 4 Confirm before starting — Review information before proceeding.
- 5 Summary — Status about the startup attempt.

The dialog box looks like this for the first step:






At each step, you can click **Help** to read detailed information about that step.

## Start a MATLAB Job Scheduler


To start a MATLAB Job Scheduler, click **Start** in the MATLAB Job Scheduler module.

In the New MATLAB Job Scheduler dialog box, provide a name for the MATLAB Job Scheduler, and select a host to run it on.



Alternative methods for starting a MATLAB Job Scheduler include selecting the pull-down **Scheduler > Start**, or right-clicking a listed host and selecting, **Start Scheduler**.

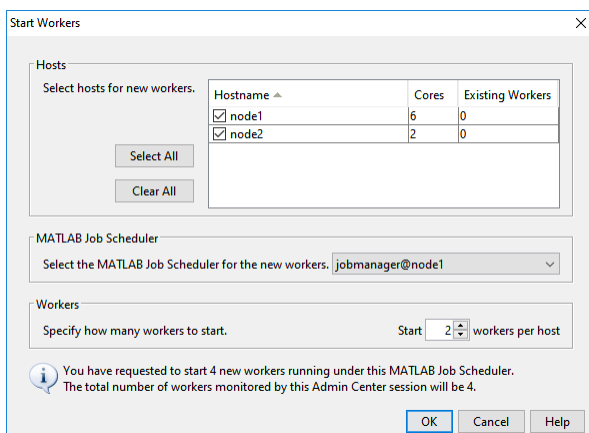
With a MATLAB Job Scheduler running on your cluster, Admin Center might look like the following figure, with the MATLAB Job Scheduler listed in the MATLAB Job Scheduler module, as well as being listed by name in the Hosts module in the line for the host on which it is running.



## Start Workers

To start MATLAB workers, click **Start** in the Workers module.

In the Start Workers dialog box, specify the numbers of workers to start on each host, and select the hosts to run them. From the list, select the MATLAB Job Scheduler for these workers. Click **OK** to start the workers. Admin center automatically provides names for the workers, based on the hosts running them.



Alternative methods for starting workers include selecting the pull-down **Workers > Start**, or right-clicking a listed host or MATLAB Job Scheduler and selecting **Start Workers**.

With workers running on your cluster, Admin Center might look like the following figure, which shows the workers listed in the Workers module. Also, the number of workers running under the MATLAB Job Scheduler is listed in the MATLAB Job Scheduler module, and the number of workers for each MATLAB Job Scheduler is listed in the Hosts module.

The screenshot shows the Admin Center interface with three main sections: Hosts, MATLAB Job Scheduler, and Workers. Each section contains a table of system components.

Hosts							
Add or Find...	Host			MJS Service		MATLAB Job Scheduler	Workers
Start MJS Service...	Hostname	Reachable	Cores	Status	Up Since	Name	Count
Stop MJS Service...	node1	yes	6	running	2018-11-08 10:54	jobmanager	2
Test Connectivity...	node2	yes	2	running	2018-11-08 10:59		2

MATLAB Job Scheduler				
Start...	Name	Hostname	Status	Workers
Stop...	jobmanager	node1	running	4
Resume				

Workers							
Start...	Worker				MATLAB Job Scheduler		
Stop...	Name	Hostname	Status	Up Since	Connection	Name	Hostname
Resume	worker1	node1	idle	2018-11-08 13:50	connected	jobmanager	node1
	worker2	node1	idle	2018-11-08 13:50	connected	jobmanager	node1
	worker3	node2	idle	2018-11-08 13:51	connected	jobmanager	node1
	worker4	node2	idle	2018-11-08 13:52	connected	jobmanager	node1

At the bottom of the interface, it says "Last updated: 08/11/18 13:52" and "Update every 2 minutes" with an "Update Now" button.

To get more information on any host, MATLAB Job Scheduler, or worker listed in Admin Center, right-click its name in the display and select **Properties**. Alternatively, you can find the **Properties** option under the **Hosts**, **Scheduler**, and **Workers** drop-down menus.

## Stop, Destroy, Resume, Restart Processes

You can **Stop** or **Destroy** the mjs service, MATLAB Job Schedulers, and workers. The primary difference is that stopping a process shuts it down but retains its data; destroying a process shuts it down and clears its data. Use **Start MJS Service** to have mjs continue with existing data. Use **Resume** to have a MATLAB Job Scheduler or worker continue with its existing data. When you use **Restart**, a dialog box requires you to confirm your intention of starting a new process while keeping or discarding data.

## Move a Worker

To move a worker from one host to another, you must completely shut it down, than start a new worker on the desired host:

- 1 Right-click the worker in the Workers module list.
- 2 Select **Destroy**. This shuts down the worker process and removes all its data.

- 3 If the old worker host is not running any other MATLAB Parallel Server processes (mjs service, MATLAB Job Scheduler, or workers), you might want to remove it from the Admin Center listing.
- 4 If necessary, add the new host to the Admin Center host listing.
- 5 In the Workers module, click **Start**. Select the desired host in the Start Workers dialog box, along with the appropriate number and MATLAB Job Scheduler name.

Use a similar process to move a MATLAB Job Scheduler from one host to another. Note, however, that all workers registered with the MATLAB Job Scheduler must be destroyed and started again, registering them with the new instance of the MATLAB Job Scheduler.

### Update the Display

Admin Center updates its data automatically at regular intervals. To set the update rate, select an option from the **Update** list. Click **Update Now** to immediately update the display data.

### See Also

#### More About

- “Start Admin Center” on page 4-2
- “Test Connectivity” on page 4-9

## Test Connectivity


Admin Center lets you test communications between your MATLAB Job Scheduler node, worker nodes, and the node where Admin Center is running.

The tests are divided into four categories:


- **Client** — Verifies that the node running Admin Center is properly configured so that further cluster testing can proceed.
- **Client to Nodes** — Verifies that the node running Admin Center can identify and communicate with the other nodes in the cluster.
- **Nodes to Nodes** — Verifies that the other nodes in the cluster can identify each other, and that each node allows its mjs service to communicate with the mjs service on the other cluster nodes.
- **Nodes to Client** — Verifies that other cluster nodes can identify and communicate with the node running Admin Center.

First click **Test Connectivity** to open the Connectivity Testing dialog box. By default, the dialog box displays the results of the last test. To run new tests and update the display, click **Run**.

During test execution, Admin Center displays this progress dialog box.



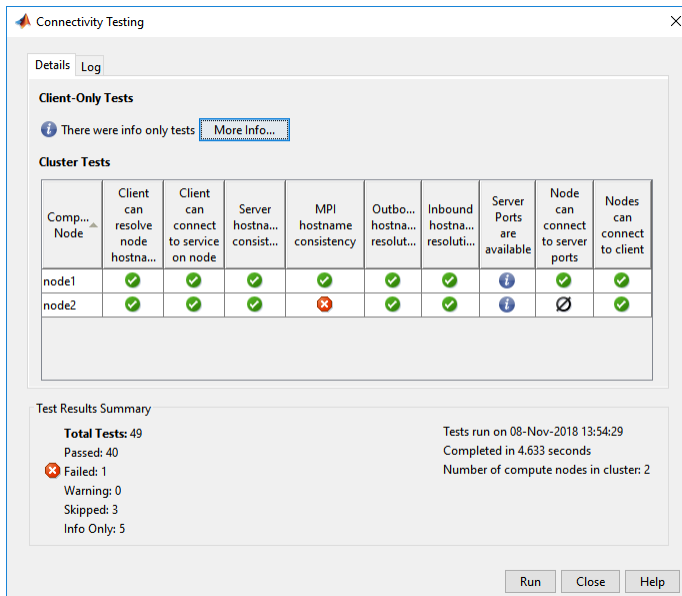
When the tests are complete, the Running Tests dialog box automatically closes, and Admin Center displays the test results in the Connectivity Testing dialog box.



The possible test result symbols are described in the following table.

Test Result	Description
	Test passed.
	Test passed, extra information is available.
	Test passed, but generated a warning.
	Test failed.
	Test was skipped, possibly because prerequisite tests did not pass.

Test that include failures or other results might look like the following figure.



Double-click any of the symbols in the test results to drill down for more detail. Use the **Log** tab to see the raw data from the tests.

The results of the tests that run on only the client are displayed in the lower-left corner of the dialog box. To drill into client-only test results, click **More Info**.

## See Also

### More About

- “Start Admin Center” on page 4-2
- “Export and Import Sessions” on page 4-12

## Export and Import Sessions

By default, Admin Center saves the cluster definition, process status, and test results, so the next time the same user runs Admin Center on the same machine, that saved information is available and displayed by default. You can export session data so that a different user or a different host can access it, by selecting the pull-down **File > Export Session**. Browse to the location where you want to store the session data and provide a name for the file. Admin Center applies the extension `.mjs` to the file name.

You can import that saved session data into a subsequent session of Admin Center by selecting the pull-down **File > Import Session**. The imported data includes cluster definition and test results.

---

**Note** When importing a session file, Admin Center automatically sets its update rate to never (i.e., disabled), so that you can statically examine a cluster setup from the time the session was saved for evaluation or diagnostic purposes.

---

### See Also

#### More About

- “Start Admin Center” on page 4-2
- “Prepare for Cluster Profiles” on page 4-13



## Prepare for Cluster Profiles

Admin Center does not create cluster profiles, but the information displayed in Admin Center is of vital importance when you create your cluster profiles — information such as MATLAB Job Scheduler name, MATLAB Job Scheduler host, and number of workers. For more information about creating and using profiles, see “Discover Clusters and Use Cluster Profiles” (Parallel Computing Toolbox).

### See Also

#### More About

- “Start Admin Center” on page 4-2
- “Discover Clusters and Use Cluster Profiles” (Parallel Computing Toolbox)
- “Benchmark Your Cluster with the HPC Challenge” (Parallel Computing Toolbox)
- “Specify Your Parallel Preferences” (Parallel Computing Toolbox)



# Control Scripts

---

## **admincenter**

Start Admin Center GUI

### **Syntax**

```
admincenter
```

### **Description**

`admincenter` opens the MATLAB Parallel Server Admin Center. When setting up or using a MATLAB job scheduler cluster, Admin Center allows you to establish and verify your cluster, and to diagnose possible problems.

For details about using Admin Center, see:

- “Start Admin Center” on page 4-2
- “Set Up Resources” on page 4-3
- “Test Connectivity” on page 4-9

### **See Also**

`mjs` | `nodestatus` | `remotemjs`

# createSharedSecret

Create shared secret for secure communication

## Syntax

```
createSharedSecret  
createSharedSecret -file <filename>
```

## Description

`createSharedSecret` creates a shared secret file used for secure communication between job managers and workers. The file is named `secret` in the current folder.

`createSharedSecret -file <filename>` create a shared secret file as the given filename.

Before passing sensitive data from one service to another (e.g., between job manager and workers), these services need to establish a trust relationship using a shared secret. This script creates a file that serves as a shared secret between the services. Each service is trusted that has access to that secret file.

Create the secret file only once per cluster on one machine, then copy it into the location specified by `SHARED_SECRET_FILE` in the `mjs_def` file on each machine before starting any job managers or workers. In a shared file system, all nodes can point to the same file. Shared secrets can be reused in subsequent sessions.

## Examples

Create a shared secret file in a central location for all the nodes of the cluster:

```
cd matlabInstallDir/toolbox/parallel/bin  
createSharedSecret -file /share/secret
```

Then make sure that the nodes' shared or copied `mjs_def` files set the parameter `SHARED_SECRET_FILE` to `/share/secret` before starting the `mjs` service on each.

## See Also

`mjs`

## mjs

Install, start, stop, or uninstall mjs service

### Syntax

```
mjs install
mjs uninstall
mjs start
mjs stop
mjs console
mjs restart
mjs ... -mjsdef <mjs_defaults_file>
mjs ... -clean
mjs status
mjs -version
mjs -useonlinelicensing
```

### Description

The mjs service ensures that all other processes are running and that it is possible to communicate with them. Once the mjs service is running, you can use the `nodestatus` command to obtain information about the mjs service and all the processes it maintains.

The mjs executable resides in the folder `matlabroot\toolbox\parallel\bin` (Windows operating system) or `matlabroot/toolbox/parallel/bin` (UNIX operating system). Enter the following commands at a DOS or UNIX command-line prompt, respectively.

`mjs install` installs the mjs service in the Microsoft Windows Service Control Manager. This causes the service to automatically start when the Windows operating system boots up. The service must be installed before it is started.

`mjs uninstall` uninstalls the mjs service from the Windows Service Control Manager. Note that if you wish to install mjs service as a different user, you must first uninstall the service and then reinstall as the new user.

`mjs start` starts the mjs service. This creates the required logging and checkpointing directories, and then starts the service as specified in the mjs defaults file.

`mjs stop` stops running the mjs service. This automatically stops all job managers and workers on the computer, but leaves their checkpoint information intact so that they will start again when the mjs service is started again.

`mjs console` starts the mjs service as a process in the current terminal or command window rather than as a service running in the background.

`mjs restart` performs the equivalent of `mjs stop` followed by `mjs start`. This command is available only on UNIX and Macintosh operating systems.

`mjs ... -mjsdef <mjs_defaults_file>` uses the specified alternative mjs defaults file instead of the one found in `matlabroot/toolbox/parallel/bin`. This file is provided for Linux

(`mjs_def.sh`) and Windows (`mjs_def.bat`). For information on backwards compatibility, see “Run Multiple MATLAB Parallel Server Versions” on page 3-29.

`mjs ... -clean` performs a complete cleanup of all service checkpoint and log files before installing or starting the service, or after stopping or uninstalling it. This deletes all information about any job managers or workers this service has ever maintained.

`mjs status` reports the status of the mjs service, indicating whether it is running and with what PID. Use `nodestatus` to obtain more detailed information about the mjs service. The `mjs status` command is available only on UNIX and Macintosh operating systems.

`mjs -version` prints version information of the mjs process to standard output, then exits.

`mjs -useonlinelicensing` ensures that workers use online licensing. Unless you specify `-useonlinelicensing`, mjs uses the network license manager.

### **See Also**

`nodestatus` | `startjobmanager` | `startworker` | `stopjobmanager` | `stopworker`

## nodestatus

Status of mjs processes running on node

### Syntax

```
nodestatus
nodestatus -flags
```

### Description

`nodestatus` displays the status of the mjs service and the processes which it maintains. The mjs service must already be running on the specified computer.

The `nodestatus` executable resides in the folder `matlabroot\toolbox\parallel\bin` (Windows operating system) or `matlabroot/toolbox/parallel/bin` (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

`nodestatus -flags` accepts the following input flags. Multiple flags can be used together on the same command.

Flag	Operation
<code>-remotehost &lt;hostname&gt;</code>	Displays the status of the mjs service and the processes it maintains on the specified host. The default value is the local host.
<code>-infolevel &lt;level&gt;</code>	Specifies how much status information to report, using a level of 1-3. 1 means only the basic information, 3 means all information available. The default value is 1.
<code>-baseport &lt;port_number&gt;</code>	Specifies the base port that the mjs service on the remote host is using. You need to specify this only if the value of <code>BASE_PORT</code> in the local <code>mjs_def</code> file does not match the base port being used by the mjs service on the remote host.
<code>-v</code>	Verbose mode displays the progress of the command execution.
<code>-json</code>	View the output in JavaScript Object Notation (JSON) format. Output in <code>json</code> format is easy to parse.

### Examples

Display basic information about the mjs processes on the local host.

```
nodestatus
```

Display detailed information about the status of the mjs processes on host `node27`.

```
nodestatus -remotehost node27 -infolevel 2
```



**See Also**

mjs | startjobmanager | startworker | stopjobmanager | stopworker

## remotecopy

(To be removed) Copy file or folder to or from one or more remote hosts using transport protocol

---

**Note** remotecopy will be removed in a future release. Use `scp` or `sftp` instead. For more information, see “Compatibility Considerations”.

---

### Syntax

```
remotecopy <flags> <protocol options>
```

### Description

`remotecopy <flags> <protocol options>` copies a file or folder to or from one or more remote hosts by using a transport protocol (such as `ssh`). Copying from multiple hosts creates a separate file per host, appending the hostname to the specified filename.

The following table describes the supported flags and options. You can combined multiple flags in the same command, preceding each flag by a dash (-).

Flags and Options	Operation
<code>-local &lt;file-or-foldername&gt;</code>	Specify the name of the file or folder on the local host.
<code>-remote &lt;file-or-foldername&gt;</code>	Specify the name of the file or folder on the remote host.
<code>-from</code>	Specify to copy from the remote hosts to the local host. You must use either the <code>-from</code> flag, or the <code>-to</code> flag.
<code>-to</code>	Specify to copy to the remote hosts from the local host. You must use either the <code>-from</code> flag, or the <code>-to</code> flag.
<code>-remotehost host1[,host2[,...]]</code>	Specify the names of the hosts where you want to copy to or from. Separate the host names by commas without any white spaces. This is a mandatory argument.
<code>-remoteplatform { unix   windows }</code>	Specify the platform of the remote hosts. This option is required only if different from the local platform.
<code>-quiet</code>	Prevent <code>remotecopy</code> from prompting for missing information. The command fails if all required information is not specified.
<code>-help</code>	Print the help information for this command.

Flags and Options	Operation
<code>-protocol &lt;type&gt;</code>	<p>Force the usage of a particular protocol type. Specifying a protocol type with all its required parameters also avoids interactive prompting and allows for use in scripts.</p> <p>The supported protocol types are <code>scp</code> and <code>sftp</code>.</p> <p>To get more information about one particular protocol type, enter</p> <pre>remotecopy -protocol &lt;type&gt; -help</pre> <p>For example:</p> <pre>remotecopy -protocol sftp -help</pre>
<code>&lt;protocol options&gt;</code>	Specify particular options for the protocol type being used.

**Note** The file permissions on the copy might not be the same as the permissions on the original file.

## Examples

Copy the local file `mjs_def.sh` to two other machines. (Enter this command on a single line.)

```
remotecopy -local mjs_def.sh -to
  -remote /matlab/toolbox/parallel/bin -remotehost hostA,hostB
```

Retrieve folders of the same name from two hosts to the local machine. (Enter command on a single line.)

```
remotecopy -local C:\temp\log -from -remote C:\temp\mjs\log
  -remotehost winHost1,winHost2
```

## Compatibility Considerations

### remotecopy will be removed

*Warns starting in R2021b*

remotecopy will be removed in a future release.

- Previously you used `remotecopy` and `-protocol scp` to copy files to and from a remote host using the `scp` protocol.

As a best practice use `scp` instead.

Not Recommended	Recommended
<code>remotecopy -remotehost host1 -local /my/file/path -remote host2:/remote/file/path</code>	<code>scp /my/file/path host2:/remote/file/path</code>
<code>remotecopy -remotehost host1,host2 -local /my/file/path -remote host2:/remote/file/path</code>	<code>scp /my/file/path host2:/remote/file/path</code>
<code>remotecopy -remotehost host1 -local /my/file/path -remote host2:/remote/file/path</code>	<code>scp /my/file/path host2:/remote/file/path</code>

- Previously you used `remotecopy` and `-protocol sftp` to copy files to and from a remote host using the `sftp` protocol.

As a best practice use `sftp` instead.

Not Recommended	Recommended
<code>remotecopy -remotehost host1 -local /my/file/path</code>	<code>sftp /my/file/path host1:/remote/file/path</code>
<code>remotecopy -remotehost host1,host2 -local /my/file/path</code>	<code>sftp /my/file/path host1:/remote/file/path -protocol sftp /my/file/path host2:/remote/file/path</code>
<code>remotecopy -remotehost host1 -local /my/file/path</code>	<code>sftp host1:/remote/file/path /my/file/path</code>

**remotecopy function will no longer support the rcp protocol in a future release**

*Warns starting in R2018b*

The `rcp` protocol is insecure. The `remotecopy` function will no longer support `rcp` as an option to the `-protocol` flag in a future release. Instead, specify any other of the supported protocols, such as `scp`.

**See Also**

`mjs`

# remotemjs

(To be removed) Execute mjs command on one or more remote hosts by transport protocol

---

**Note** remotemjs will be removed in a future release. Use ssh instead. For more information, see “Compatibility Considerations”.

---

## Syntax

```
remotemjs <mjs options> <flags> <protocol options>
```

## Description

remotemjs <mjs options> <flags> <protocol options> allows you to execute the mjs service on one or more remote hosts.

For a description of the mjs service, see the mjs reference page.

The following table describes the supported flags and options. You can combined multiple flags in the same command, preceding each flag by a dash (-).

Flags and Options	Operation
<mjs options>	Options and arguments of the mjs command, such as start, stop, etc. See the mjs reference page for a full list.
-matlabroot <installfoldername>	The MATLAB installation folder on the remote hosts, required only if the remote installation folder differs from the one on the local machine.
-remotehost host1[,host2[,...]]	The names of the hosts where you want to run the mjs command. Separate the host names by commas without any white spaces. This is a mandatory argument.
-remoteplatform { unix   windows }	The platform of the remote hosts. This option is required only if different from the local platform.
-quiet	Prevent mjs from prompting the user for missing information. The command fails if all required information is not specified.
-help	Print help information.

Flags and Options	Operation
-protocol <type>	<p>Force the usage of a particular protocol type. Specifying a protocol type with all its required parameters also avoids interactive prompting and allows for use in scripts.</p> <p>The supported protocol types are ssh and winc.</p> <p>To get more information about one particular protocol type, enter</p> <pre>remotemjs -protocol &lt;type&gt; -help</pre> <p>For example:</p> <pre>remotemjs -protocol winc -help</pre> <p>Using the winc protocol requires that you log in as a user with admin privileges on the remote host.</p>
<protocol options>	Specify particular options for the protocol type being used.

---

**Note** If you are using OpenSSH on a Microsoft Windows operating system, you can encounter a problem when using backslashes in path names for your command options. In most cases, you can work around this problem by using forward slashes instead. For example, to specify the file C:\temp\mjs\_def.bat, you should identify it as C:/temp/mjs\_def.bat.

---

## Examples

Start mjs on three remote machines of the same platform as the client:

```
remotemjs start -remotehost hostA,hostB,hostC
```

Start mjs in a clean state on two UNIX operating system machines from a Windows operating system machine, using the ssh protocol. Enter the following command on a single line:

```
remotemjs start -clean -matlabroot /usr/local/matlab
  -remotehost unixHost1,unixHost2 -remoteplatform UNIX
  -protocol ssh
```

## Compatibility Considerations

### remotemjs will be removed

*Warns starting in R2021b*

remotemjs will be removed in a future release.

Previously you used remotemjs to run MATLAB Job Scheduler commands on a remote host using -protocol ssh or -protocol winc.

As a best practice use ssh instead.

Not Recommended	Recommended
remotemjs <mjs options> -matlabroot <installfoldername> -installfolder <host>/toolbox/parallel/bin/mjs	remotemjs <mjs options> -matlabroot <installfoldername> -installfolder <host>/toolbox/parallel/bin/mjs
remotemjs <mjs options> -matlabroot <installfoldername> -installfolder <host1>/toolbox/parallel/bin/mjs	remotemjs <mjs options> -matlabroot <installfoldername> -installfolder <host1>/toolbox/parallel/bin/mjs ssh host2 <installfoldername>/toolbox/parallel/bin/mjs

### remotemjs function will no longer support the rsh protocol in a future release

Warns starting in R2018b

The rsh protocol is insecure. The remotemjs function will no longer support rsh as an option to the -protocol flag in a future release. Instead, specify any other of the supported protocols, such as ssh.

### See Also

mjs

## pausejobmanager

Pause job manager process

### Syntax

```
pausejobmanager
pausejobmanager -flags
```

### Description

pausejobmanager pauses a job manager that is running under the mjs service.

The pausejobmanager executable resides in the folder `matlabroot\toolbox\parallel\bin` (Windows operating system) or `matlabroot/toolbox/parallel/bin` (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

pausejobmanager -flags accepts the following input flags. Multiple flags can be used together on the same command.

Flag	Operation
-name <job_manager_name>	Specifies the name of the job manager to pause. The default is the value of DEFAULT_JOB_MANAGER_NAME parameter the mjs_def file.
-remotehost <hostname>	Specifies the name of the host where you want to pause the job manager. The default value is the local host.
-baseport <port_number>	Specifies the base port that the mjs service on the remote host is using. You need to specify this only if the value of BASE_PORT in the local mjs_def file does not match the base port being used by the mjs service on the remote host.
-v	Verbose mode displays the progress of the command execution.

### Examples

Pause the job manager MyJobManager on the local host.

```
pausejobmanager -name MyJobManager
```

Pause the job manager MyJobManager on the host JMHost.

```
pausejobmanager -name MyJobManager -remotehost JMHost
```

### See Also

mjs | nodestatus | startjobmanager | stopjobmanager | resumejobmanager



# resumejobmanager

Resume job manager process

## Syntax

```
resumejobmanager
resumejobmanager -flags
```

## Description

resumejobmanager resumes a job manager that is running under the mjs service.

The resumejobmanager executable resides in the folder *matlabroot\toolbox\parallel\bin* (Windows operating system) or *matlabroot/toolbox/parallel/bin* (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

resumejobmanager -flags accepts the following input flags. Multiple flags can be used together on the same command.

Flag	Operation
-name <job_manager_name>	Specifies the name of the job manager to resume. The default is the value of DEFAULT_JOB_MANAGER_NAME parameter the mjs_def file.
-remotehost <hostname>	Specifies the name of the host where you want to resume the job manager. The default value is the local host.
-baseport <port_number>	Specifies the base port that the mjs service on the remote host is using. You need to specify this only if the value of BASE_PORT in the local mjs_def file does not match the base port being used by the mjs service on the remote host.
-v	Verbose mode displays the progress of the command execution.

## Examples

Resume the job manager MyJobManager on the local host.

```
resumejobmanager -name MyJobManager
```

Resume the job manager MyJobManager on the host JMHost.

```
resumejobmanager -name MyJobManager -remotehost JMHost
```

## See Also

mjs | nodestatus | startjobmanager | stopjobmanager | pausejobmanager

## startjobmanager

Start job manager process

### Syntax

```
startjobmanager
startjobmanager -flags
```

### Description

`startjobmanager` starts a job manager process and the associated job manager lookup process under the `mjs` service, which maintains them after that. The job manager handles the storage of jobs and the distribution of tasks contained in jobs to MATLAB workers that are registered with it. The `mjs` service must already be running on the specified computer.

The `startjobmanager` executable resides in the folder `matlabroot\toolbox\parallel\bin` (Windows operating system) or `matlabroot/toolbox/parallel/bin` (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

`startjobmanager -flags` accepts the following input flags. Multiple flags can be used together on the same command.

Flag	Operation
-name <job_manager_name>	Specifies the name of the job manager. This identifies the job manager to MATLAB worker sessions and MATLAB clients. The default is the value of the <code>DEFAULT_JOB_MANAGER_NAME</code> parameter in the <code>mjs_def</code> file.
-remotehost <hostname>	Specifies the name of the host where you want to start the job manager and the job manager lookup process. If omitted, they start on the local host.
-clean	Deletes all checkpoint information stored on disk from previous instances of this job manager before starting. This cleans the job manager so that it initializes with no existing jobs or tasks.
-baseport <port_number>	Specifies the base port that the <code>mjs</code> service on the remote host is using. You need to specify this only if the value of <code>BASE_PORT</code> in the local <code>mjs_def</code> file does not match the base port being used by the <code>mjs</code> service on the remote host.
-certificate <path_to_certificate_file>	Specifies the path to the certificate file to use to connect to the job manager. You must use this flag when the <code>mjs_def</code> file of the job manager sets <code>REQUIRE_CLIENT_CERTIFICATE</code> to <code>true</code> .
-useMSMPI	Use Microsoft MPI (MS-MPI) for clusters on Windows platforms.

Flag	Operation
-v	Verbose mode displays the progress of the command execution.

## Examples

Start the job manager MyJobManager on the local host.

```
startjobmanager -name MyJobManager
```

Start the job manager MyJobManager on the host JMHost.

```
startjobmanager -name MyJobManager -remotehost JMHost
```

## See Also

mjs | nodestatus | pausejobmanager | resumejobmanager | startworker | stopjobmanager | stopworker

## startworker

Start MATLAB worker session

### Syntax

```
startworker
startworker -flags
```

### Description

`startworker` starts a MATLAB worker process under the `mjs` service, which maintains it after that. The worker registers with the specified job manager, from which it will get tasks for evaluation. The `mjs` service must already be running on the specified computer.

The `startworker` executable resides in the folder `matlabroot\toolbox\parallel\bin` (Windows operating system) or `matlabroot/toolbox/parallel/bin` (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

`startworker -flags` accepts the following input flags. Multiple flags can be used together on the same command, except where noted.

Flag	Operation
<code>-name &lt;worker_name&gt;</code>	Specifies the name of the MATLAB worker. The default is the value of the <code>DEFAULT_WORKER_NAME</code> parameter in the <code>mjs_def</code> file.
<code>-remotehost &lt;hostname&gt;</code>	Specifies the name of the computer where you want to start the MATLAB worker. If omitted, the worker is started on the local computer.
<code>-jobmanager &lt;job_manager_name&gt;</code>	Specifies the name of the job manager this MATLAB worker will receive tasks from. The default is the value of the <code>DEFAULT_JOB_MANAGER_NAME</code> parameter in the <code>mjs_def</code> file.
<code>-jobmanagerhost &lt;job_manager_hostname&gt;</code>	Specifies the host on which the job manager is running. The worker contacts the job manager lookup process on that host to register with the job manager.  This overrides the setting of <code>JOB_MANAGER_HOST</code> in the <code>mjs_def</code> file on the worker computer. You must specify the job manager host by one of these means.
<code>-clean</code>	Deletes all checkpoint information associated with this worker name before starting.

Flag	Operation
-baseport <port_number>	Specifies the base port that the mjs service on the remote host is using. You only need to specify this if the value of BASE_PORT in the local mjs_def file does not match the base port being used by the mjs service on the remote host.
-v	Verbose mode displays the progress of the command execution.

## Examples

Start a worker on the local host, using the default worker name, registering with the job manager MyJobManager on the host JMHost.

```
startworker -jobmanager MyJobManager -jobmanagerhost JMHost
```

Start a worker on the host WorkerHost, using the default worker name, and registering with the job manager MyJobManager on the host JMHost. (The following command should be entered on a single line.)

```
startworker -jobmanager MyJobManager -jobmanagerhost JMHost
            -remotehost WorkerHost
```

Start two workers, named worker1 and worker2, on the host WorkerHost, registering with the job manager MyJobManager that is running on the host JMHost. Note that to start two workers on the same computer, you must give them different names. (Each of the two commands below should be entered on a single line.)

```
startworker -name worker1 -remotehost WorkerHost
            -jobmanager MyJobManager -jobmanagerhost JMHost
startworker -name worker2 -remotehost WorkerHost
            -jobmanager MyJobManager -jobmanagerhost JMHost
```

## See Also

mjs | nodestatus | startjobmanager | stopjobmanager | stopworker

## stopjobmanager

Stop job manager process

### Syntax

```
stopjobmanager
stopjobmanager -flags
```

### Description

stopjobmanager stops a job manager that is running under the mjs service.

The stopjobmanager executable resides in the folder *matlabroot\toolbox\parallel\bin* (Windows operating system) or *matlabroot/toolbox/parallel/bin* (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

stopjobmanager -flags accepts the following input flags. Multiple flags can be used together on the same command.

Flag	Operation
-name <job_manager_name>	Specifies the name of the job manager to stop. The default is the value of DEFAULT_JOB_MANAGER_NAME parameter the mjs_def file.
-remotehost <hostname>	Specifies the name of the host where you want to stop the job manager and the associated job manager lookup process. The default value is the local host.
-clean	Deletes all checkpoint information stored on disk for the current instance of this job manager after stopping it. This cleans the job manager of all its job and task data.
-baseport <port_number>	Specifies the base port that the mjs service on the remote host is using. You need to specify this only if the value of BASE_PORT in the local mjs_def file does not match the base port being used by the mjs service on the remote host.
-v	Verbose mode displays the progress of the command execution.

### Examples

Stop the job manager MyJobManager on the local host.

```
stopjobmanager -name MyJobManager
```

Stop the job manager MyJobManager on the host JMHost.

```
stopjobmanager -name MyJobManager -remotehost JMHost
```

**See Also**

mjs | nodestatus | startjobmanager | startworker | stopworker

## stopworker

Stop MATLAB worker session

### Syntax

```
stopworker
stopworker -flags
```

### Description

stopworker stops a MATLAB worker process that is running under the mjs service.

The stopworker executable resides in the folder *matlabroot\toolbox\parallel\bin* (Windows operating system) or *matlabroot/toolbox/parallel/bin* (UNIX operating system). Enter the following command syntax at a DOS or UNIX command-line prompt, respectively.

stopworker -flags accepts the following input flags. Multiple flags can be used together on the same command.

Flag	Operation
-name <worker_name>	Specifies the name of the MATLAB worker to stop. The default is the value of the DEFAULT_WORKER_NAME parameter in the mjs_def file.
-remotehost <hostname>	Specifies the name of the host where you want to stop the MATLAB worker. The default value is the local host.
-clean	Deletes all checkpoint information associated with this worker name after stopping it.
-baseport <port_number>	Specifies the base port that the mjs service on the remote host is using. You need to specify this only if the value of BASE_PORT in the local mjs_def file does not match the base port being used by the mjs service on the remote host.
-v	Verbose mode displays the progress of the command execution.

### Examples

Stop the worker with the default name on the local host.

```
stopworker
```

Stop the worker with the default name, running on the computer WorkerHost.

```
stopworker -remotehost WorkerHost
```

Stop the workers named worker1 and worker2, running on the computer WorkerHost.



```
stopworker -name worker1 -remotehost WorkerHost  
stopworker -name worker2 -remotehost WorkerHost
```

**See Also**

mjs | nodestatus | startjobmanager | startworker | stopjobmanager



# Reference Architectures

---

- “Run MATLAB Parallel Server on Amazon Web Services” on page 6-2
- “Run MATLAB Parallel Server on AWS Batch” on page 6-4
- “Run MATLAB Parallel Server on Microsoft Azure Using Reference Architecture” on page 6-5
- “Run MATLAB Parallel Server from Microsoft Azure Marketplace” on page 6-7
- “Configure MATLAB Parallel Server Licensing for Cloud Platforms” on page 6-17

## Run MATLAB Parallel Server on Amazon Web Services

In this section...
“Requirements” on page 6-2
“Run from GitHub” on page 6-2

Use a customizable reference architecture to run MATLAB Parallel Server on Linux virtual machines in Amazon Web Services (AWS®).

Use this reference architecture when you want to launch a cluster of MATLAB workers in a specific region, combine your MATLAB resources with your existing cloud resources, or automate deployment.

For a simpler but less customizable method of launching a MATLAB Parallel Server cluster in AWS, try MathWorks Cloud Center.

### Requirements

To use this reference architecture, you need:

- An AWS account. You are responsible for the costs of all AWS services.
- An SSH key pair for your AWS account in your chosen region. For more information, see Amazon EC2 Key Pairs.
- Parallel Computing Toolbox installed on a client machine. The version must match the MATLAB version of the reference architecture.
- A MATLAB Parallel Server license that meets the following conditions:
  - Current on Software Maintenance Service (SMS).
  - Linked to a MathWorks Account.
  - Configured for cloud use. Individual and Campus-Wide licenses are already configured. For other license types, contact your administrator. You can identify your license type and administrator by viewing your MathWorks Account. Administrators can consult “Configure MATLAB Parallel Server Licensing for Cloud Platforms” on page 6-17.

### Run from GitHub

To launch MATLAB in AWS, use the reference architecture templates provided in the following GitHub repository:

- MATLAB Parallel Server on AWS (GitHub)

You can run the template directly from the link in the GitHub repository. Choose your MATLAB release, and then click the **Launch Stack** icon next to the region in which you want to deploy your resources.

## **See Also**

### **Related Examples**

- “Run MATLAB Parallel Server on Microsoft Azure Using Reference Architecture” on page 6-5

### **External Websites**

- MATLAB Parallel Server on AWS (GitHub)

## Run MATLAB Parallel Server on AWS Batch

Deploy MATLAB Parallel Server in Amazon Web Services (AWS), using AWS Batch, with your Amazon Web Services account.

Use this reference architecture when you want to batch process MATLAB jobs using AWS Batch. AWS Batch is a service that dynamically provisions, manages, monitors, and terminates Amazon EC2 instances based on the volume and resource requirements of the submitted jobs.

For a simpler but less customizable method of launching a MATLAB Parallel Server cluster in AWS, try MathWorks Cloud Center.

### Requirements

To use this reference architecture, you need:

- An AWS account. You are responsible for the costs of all AWS services.
- Parallel Computing Toolbox installed on a client machine. The version must match the MATLAB version of the reference architecture.
- A MATLAB Parallel Server license that meets the following conditions:
  - Current on Software Maintenance Service (SMS).
  - Linked to a MathWorks Account.
  - Configured for cloud use. Individual and Campus-Wide licenses are already configured. For other license types, contact your administrator. You can identify your license type and administrator by viewing your MathWorks Account. Administrators can consult “Configure MATLAB Parallel Server Licensing for Cloud Platforms” on page 6-17.

### Run from GitHub

To launch MATLAB Parallel Server in AWS using AWS Batch, use the reference architecture templates provided in the following GitHub repository:

- MATLAB Parallel Server with AWS Batch (GitHub)

You can run the template directly from the link in the GitHub repository. Choose your MATLAB release, and then click the **Launch Stack** icon next to the region in which you want to deploy your resources.

### See Also

#### Related Examples

- “Run MATLAB Parallel Server on Amazon Web Services” on page 6-2

#### External Websites

- MATLAB Parallel Server with AWS Batch (GitHub)

# Run MATLAB Parallel Server on Microsoft Azure Using Reference Architecture

## In this section...

“Requirements” on page 6-5

“Run from GitHub” on page 6-5

Use a customizable reference architecture to run MATLAB Parallel Server on Windows virtual machines in Microsoft Azure.

Use this reference architecture when you want to launch a cluster of MATLAB workers in Microsoft Azure, combine your MATLAB resources with your existing cloud resources, or automate deployment.

For a simpler but less customizable method of launching a MATLAB Parallel Server cluster in Microsoft Azure, see “Run MATLAB Parallel Server from Microsoft Azure Marketplace” on page 6-7.

## Requirements

To use this reference architecture, you need:

- A Microsoft Azure account. You are responsible for the costs of all Microsoft Azure services.
- Parallel Computing Toolbox installed on a client machine. The version must match the MATLAB version of the reference architecture.
- A MATLAB Parallel Server license that meets the following conditions:
  - Current on Software Maintenance Service (SMS).
  - Linked to a MathWorks Account.
  - Configured for cloud use. Individual and Campus-Wide licenses are already configured. For other license types, contact your administrator. You can identify your license type and administrator by viewing your MathWorks Account. Administrators can consult “Configure MATLAB Parallel Server Licensing for Cloud Platforms” on page 6-17.

## Run from GitHub

To launch MATLAB in Azure, use the reference architecture templates provided in the following GitHub repository:

- MATLAB Parallel Server on Microsoft Azure (GitHub)

You can run the template directly from the link in the GitHub repository. Choose your MATLAB release, and then click the **Deploy to Azure** icon.

## See Also

### Related Examples

- “Run MATLAB Parallel Server on Amazon Web Services” on page 6-2

- “Run MATLAB Parallel Server from Microsoft Azure Marketplace” on page 6-7

### **External Websites**

- MATLAB Parallel Server on Microsoft Azure (GitHub)



## Run MATLAB Parallel Server from Microsoft Azure Marketplace

This topic describes how to use the Azure Marketplace to quickly deploy and run MATLAB Parallel Server in Azure. The elasticity of the cloud infrastructure combined with MATLAB Parallel Server enables you to leverage greater computing resources and keep your calculations close to your data.

Anyone with a valid MATLAB license that includes MATLAB Parallel Server and an Azure account can run MATLAB Parallel Server in the cloud. This offering is called *MATLAB Parallel Server (BYOL)* because you "Bring Your Own License."

### Requirements

To complete these instructions, you need:

- A MATLAB License on page 6-7 that includes MATLAB Parallel Server
- A Microsoft Azure Account
- A working knowledge of Azure Resource Manager

### Licensing

By default, the MATLAB Parallel Server (BYOL) offering in the Azure Marketplace uses online licensing. The following table summarizes which licenses support running MATLAB Parallel Server on the cloud using online licensing.

License Type	Cloud Availability
Individual (a license that only you use that is in your name; not a Home or Student license)	Your license is already configured for MATLAB Parallel Server in the cloud.
Campus-Wide License (a license that you use that belongs to your academic institution)	Your license is already configured for MATLAB Parallel Server in the cloud.
Home and Student (that only you use and is in your name)	These license types are not eligible to use MATLAB Parallel Server in the cloud.
All other license types, including Concurrent and Network Named User	Contact your license administrator.*  The license administrator might need to make some changes to the license to enable you to run MATLAB Parallel Server in the cloud.  <b>Administrators:</b> For more information on this process, see "Configure MATLAB Parallel Server Licensing for Cloud Platforms" on page 6-17.
*Not sure who the license administrator is? Sign in to your MathWorks Account, click the license you are using, then click the tab marked "Contact Administrators."	

To use a license that is managed using a network license manager, you can specify the port and host name or IP address of the network license manager when you create the cloud resources in Azure. You can either specify the location of an existing license manager, or deploy a new license manager in the cloud. You must make sure that the license manager can communicate with your cloud resources. To deploy a Network Manager in Azure using an Azure Marketplace software plan developed by MathWorks, follow the steps at Run Network License Manager from Azure Marketplace.

## Deploy MATLAB Parallel Server Resources in Azure

### Configure and Deploy Template from Marketplace

Use the following steps to configure and deploy MATLAB Parallel Server (BYOL) resources in a resource group.

- 1 Navigate to <https://portal.azure.com> and log in to your Azure account.
- 2 From the Portal, click **Create a Resource**.
- 3 Search for “matlab parallel server” in the marketplace and select the **MATLAB Parallel Server (BYOL)** offering.
- 4 Click **Create** on the offering page to begin setup.



- 5 The setup process uses a Resource Manager template to help you configure the virtual machine (VM) and networking settings. To make setup easy, many of the fields come prepopulated with suitable values. The following table describes how to set the various options in each menu of the template. Click **Next** at the end of each step to proceed to the next menu.

Menu	Option	Description
<b>Basics</b>	Subscription	Select an Azure subscription to use.
	Resource group	You can select an existing resource group from the drop-down menu, or click <b>Create New</b> to create a new resource group. If you select an existing resource group, then it must not have any currently deployed resources.
	Region	Select a region from the drop-down list.
<b>Cluster Settings</b>	Cluster name	Select a name for the cluster. The default name is myCluster.
	Num worker nodes	The number of Azure instances to start for the workers to run on.

Menu	Option	Description
	Num workers per node	The number of MATLAB workers to start on each instance. Specify one worker for every two vCPUs to assign one worker per physical core. For example, a Standard_D64s_v3 instance has 64 vCPUs, so can support 32 MATLAB workers. For details on the number of vCPUs for each instance type, see <a href="https://docs.microsoft.com/en-us/azure/virtual-machines/sizes">https://docs.microsoft.com/en-us/azure/virtual-machines/sizes</a> .
	Headnode instance type	The Azure instance type to use for the head node, which runs the job manager. No workers are started on this node, so this node can be a smaller instance type than the worker nodes. For a list of instance types and sizes, see <a href="https://docs.microsoft.com/en-us/azure/virtual-machines/sizes">https://docs.microsoft.com/en-us/azure/virtual-machines/sizes</a> .
	Worker instance type	The Azure instance type to use for the workers. For a list of instance types and sizes, see <a href="https://docs.microsoft.com/en-us/azure/virtual-machines/sizes">https://docs.microsoft.com/en-us/azure/virtual-machines/sizes</a> .
	Database volume size	The size of the volume in gigabytes used to store the database files. If you set this value to 0, a separate volume is not created and the root volume is used for the database.
	Client IP address	The IP address range that can be used to access the cluster from MATLAB. The range must be a valid IP CIDR range of the form x.x.x.x/x. Use a value of the form x.x.x.x/32 to restrict access to only your computer.

Menu	Option	Description
	Password	Choose the admin password for the user "matlab". This password is required when you log into any instance (headnode or worker) using remote desktop protocol. Your password must have at least 12 characters and meet the Azure password requirements. For information on the password requirements, see What are the password requirements when creating a VM?
	Confirm password	Retype your chosen password.
	Network license manager port@server	If you are using a network license manager, enter the port and hostname or IP address. Make sure that the network license manager can communicate with your Azure resources. To deploy a network license manager in Azure, see Run Network License Manager from Azure Marketplace.
<b>Networking</b>	Virtual network	<p>The <b>Virtual network</b> field is prepopulated with a new virtual network resource named vnet01.</p> <ul style="list-style-type: none"> <li>You can configure the name, address space, or subnets of the new virtual network resource by clicking <b>Create new</b>.</li> <li>You can select an existing virtual network resource in the drop-down menu. If you choose to use an existing resource, then the template does not create any new virtual network resources.</li> </ul>

Menu	Option	Description
	Subnets	<p>The <b>Subnet</b> field is prepopulated with a new subnet named <code>subnet - 1</code> and an associated subnet address.</p> <ul style="list-style-type: none"><li>• You can select a different subnet of the selected virtual network in the drop-down menu.</li><li>• If you selected an existing virtual network rather than creating a new one, then you can click <b>Manage subnet configuration</b> to configure the subnets for the network.</li></ul>

Menu	Option	Description
<b>Review and Create</b>	—	<p>When you advance to the <b>Review and Create</b> menu, Azure automatically runs some final validation checks on the information entered on previous screens. If Azure finds any errors, then you must fix them before proceeding.</p> <p>After the validation completes successfully, review the MathWorks terms of use and privacy policy. You are responsible for all associated costs once you deploy MATLAB Parallel Server in Azure.</p> <p>Once you are satisfied with the values you have entered, click <b>Create</b> to finalize setup and begin deploying the selected resources, or click <b>Download a template for automation</b> to get a copy of the completed template.</p> <ul style="list-style-type: none"> <li>• The deployment can take several minutes. You get a notification in the Azure Portal once the deployment is complete.</li> <li>• When the deployment is complete, follow the instructions in “Connect to your Cluster from MATLAB” on page 6-12 to use the cluster you set up.</li> </ul>

### Connect to your Cluster from MATLAB

Once you have successfully configured and deployed the MATLAB Parallel Server (BYOL) resources, use the following steps to connect to the cluster you set up.

- 1 From the Azure Portal, navigate to the resource group with the resources you deployed.
- 2 Select the Storage Account ending with storage.
- 3 Select the Files container type.
- 4 Select the File Share named "shared."

- 5 Download the `.settings` file.
- 6 Open MATLAB on your client machine.
- 7 In the Parallel drop-down menu in the MATLAB toolstrip, select Create and Manage Clusters....
- 8 Click Import.
- 9 Select the downloaded profile and click open.
- 10 Click Set as Default.
- 11 (Optional) Validate your cluster by clicking the Validate button.

After you set the cloud cluster as the default, the next time you run a parallel language command (such as `parfor`, `spmd`, `parfeval`, or `batch`), MATLAB connects to the cluster. The first time you connect, you are prompted for your MathWorks account login. The first time you run a task on a worker, it can take several minutes for the worker MATLAB to start. This delay is due to provisioning the instance disk. This is a one-time operation, and subsequent tasks begin much faster.

Your cluster is now ready to use. The cluster remains running after you close MATLAB on the client machine.

---

**Caution** Use the profile and client IP address range to control access to your cloud resources. Anyone with the `.settings` file can connect to your resources from a machine within the specified IP address range and run jobs on it.

---

### Port Requirements for Accessing MATLAB Parallel Server

To access a MATLAB Parallel Server cluster from your client MATLAB, your client machine must be able to communicate on specific ports. Make sure that the network firewall allows the following outgoing connections.

Required Ports	Description
TCP 27350 to 27358 + 4*N	Ports 27350 to 27358 + 4*N, where N is the maximum number of workers on a single node
TCP 443	HTTPS access to (at least) <code>*.mathworks</code> and <code>*.microsoft.com</code>
TCP 3389	Remote Desktop Protocol for access to cluster nodes

### Delete MATLAB Parallel Server Resources from Azure

You can remove the resource group and all associated resources when you are done with them to help save on costs.

- 1 Log in to the Azure Portal.
- 2 Select the resource group containing the MATLAB Parallel Server (BYOL) resources you deployed.
- 3 Click the **Delete resource group** icon to delete all resources deployed in the group.
- 4 You are prompted to confirm the deletion.

The deletion is final, and files do not persist between deployments. You must create the resources to make them available again.

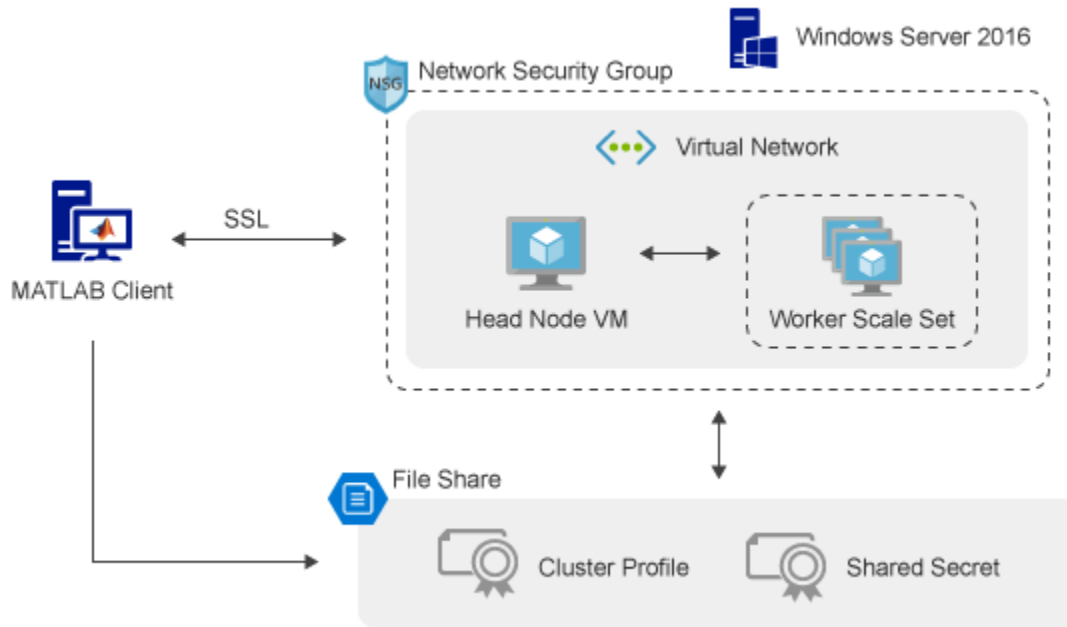
## Architecture and Resources for MATLAB Parallel Server in Azure

Deploying MATLAB Parallel Server in Azure creates several resources in your resource group. The deployment sets up a single Azure VM for the headnode, an Azure virtual machine scaling set (VMSS) for the workers, a network interface with a public IP address to connect to the VM, a network security group that controls network traffic, and a virtual network for communication between resources. The following table summarizes the resources that are created.

Resource Name	Default Resource Name in Azure	Description
Virtual machine	myclus-headnode	A compute instance for the cluster headnode. The MATLAB install is part of the VM image and the job database is stored either locally on the root volume. Optionally, a separate data disk can be used. Communication between clients and the headnode is secured using SSL.
Virtual machine scaling set	mycl-vmss	A scaling set for worker instances to be launched into. The scaling features are not currently used. The scaling set is configured to attach an extension to each instance that is configured at deployment time as a worker node of the cluster. Communication between clients and workers is secured using SSL.
Network interface	myclus-headnodenic	Enables the headnode and workers to communicate with the internet, Azure, and on-premises resources.
Network security group	mycluster-nsg	Allows or denies traffic to and from sources and destinations.
Virtual network	vnet01	Enables resources to communicate with each other. If you select an existing virtual network resource, then the template does not create this resource.

The architecture of these resources is summarized in the following diagram.





## VM Software

A preconfigured Windows VM is provided to make deployment easy. The VM contains the following software:

- MATLAB, Simulink, toolboxes, and support for GPUs.

The license used to run MATLAB determines which products and toolboxes are available for you to use. However, all products are installed on the VM, so the `ver` command returns a list of all products. To add new products to your license, see [Products and Services](#).

- Add-Ons: Deep Learning Toolbox™ Model for AlexNet Network, Deep Learning Toolbox Model for GoogLeNet Network, and Deep Learning Toolbox Model for ResNet-50 Network.

## Marketplace vs. Reference Architecture

The Azure Marketplace offering runs only the latest release of MATLAB Parallel Server. If you want to run older releases of MATLAB Parallel Server, or customize the templates and automation scripts more extensively, then you can use the MATLAB Parallel Server Reference Architecture for Microsoft Azure instead.

## Support

If you encounter an unexpected problem, search MATLAB Answers for solutions first. Most likely, other people have run into the same problem and resolved it already.

If the problem persists, or there are not any relevant posts on MATLAB Answers for the issue, contact Technical Support.

## **See Also**

## **See Also**

## **External Websites**

- [Use MATLAB in the Cloud](#)
- [MATLAB Parallel Server Reference Architecture for Microsoft Azure](#)

# Configure MATLAB Parallel Server Licensing for Cloud Platforms

## Before You Begin

This document shows administrators how to set up a license for use with the MATLAB Parallel Server cloud reference architectures. If you are an end user, contact your administrator.

To use the default MATLAB Parallel Server installation provided as part of a MathWorks reference architecture (Azure or AWS), you must be using release R2018a or later. The installation is configured to use online licensing by default. This license manager might differ from your current license configuration. Follow these instructions for setting up your license to use this installation.

If you have questions about any part of this license configuration process, contact MathWorks Cloud Support.

## License Configuration Requirements

To set up your license to use MATLAB Parallel Server running in the cloud:

- You must be an administrator of a MATLAB Parallel Server license. If you do not have a server license, you can request a free trial or contact your MathWorks sales representative. If you are not sure that you are an administrator, or if you are not sure if you are associated with a MATLAB Parallel Server license, sign in to your MathWorks Account. Click the MATLAB Parallel Server license. If you see a “Manage Users” tab, you are an administrator.
- The MATLAB Parallel Server license you are using must be current on Software Maintenance Service (SMS).
- All licensed end users must be associated with the server license. (Instructions are included in this procedure.)
- The end user’s MATLAB license must include Parallel Computing Toolbox . The toolbox is required so that jobs can be submitted to the cluster running the server. End users must be using the same release as the cluster.

## Select Licensing Configuration Option

From the following scenarios, select a procedure for setting up the license manager with online licensing. Not all procedures are required for all scenarios.

If you have a license setup scenario that is not listed here, Contact Support.

Scenario	Procedure
<ul style="list-style-type: none"> <li>• You have a license for MATLAB Parallel Server, and</li> <li>• The license is already configured for online licensing.</li> </ul>	“Add End Users to License” on page 6-18
You want to get a new on-demand license for MATLAB Parallel Server.	<ol style="list-style-type: none"> <li>1 Get an On-Demand License for MATLAB Parallel Server</li> <li>2 “Add End Users to License” on page 6-18</li> </ol>

Scenario	Procedure
<ul style="list-style-type: none"> <li>You just received a trial or an annual or perpetual license for MATLAB Parallel Server that has not yet been activated, and</li> <li>The license or trial is configured for the network license manager.</li> </ul>	<ol style="list-style-type: none"> <li>“Change License Manager Type” on page 6-18</li> <li>“Add End Users to License” on page 6-18</li> </ol>
<ul style="list-style-type: none"> <li>You have an active annual or perpetual license for MATLAB Parallel Server, and</li> <li>The license is configured for the network license manager, and</li> <li>You are currently actively using that license in a cluster, and</li> <li>You want to use online licensing</li> </ul>	<ol style="list-style-type: none"> <li>“Change License Manager Type” on page 6-18</li> <li>“Add End Users to License” on page 6-18</li> <li>“Disable Network License Manager for MATLAB Parallel Server” on page 6-19</li> </ol>
<ul style="list-style-type: none"> <li>You have an active annual or perpetual license for MATLAB Parallel Server, and</li> <li>The license is configured for the network license manager, and</li> <li>You are currently actively using that license in a cluster, and</li> <li>You want to use the network license manager running on the cloud.</li> </ul>	<ol style="list-style-type: none"> <li>Select one of the following: <ul style="list-style-type: none"> <li>Network License Manager for MATLAB on AWS</li> <li>Network License Manager for MATLAB on Azure</li> </ul> </li> <li>“Add End Users to License” on page 6-18</li> </ol>
<ul style="list-style-type: none"> <li>You are currently using the network license manager, and</li> <li>You want to switch to the network license manager running on the cloud.</li> </ul>	Select one of the following: <ul style="list-style-type: none"> <li>Network License Manager for MATLAB on AWS</li> <li>Network License Manager for MATLAB on Azure</li> </ul>

## Change License Manager Type

- 1 Sign in to your MathWorks Account.
- 2 In your MathWorks Account, click the MATLAB Parallel Server license that you plan to use in the cloud.
- 3 Go to **Install and Activate**.
- 4 Under **Related Tasks**, click **Change License Management Type**.
- 5 In the drop-down list, select **Online licensing**.

## Add End Users to License

For online licensing to know which users have permission to check out a license, you must define a list of allowed users for that license. Allowed users must have a MathWorks Account associated with the on-demand license.

You can associate end users yourself or you can generate emails so that end users can associate themselves. Both actions are performed in License Center.

- 1 Sign in to the MathWorks Account that is associated with the license you want to use.
- 2 In your MathWorks Account, click on the MATLAB Parallel Server license.
- 3 Go to **Manage Users**.
  - Associate end users manually:
    - a Click **Add User**.
    - b Add an end user. Repeat this step for as many end users as you need to associate.
  - Invite end users to associate themselves:
    - a Click **Invite End Users to Add Themselves**.
    - b Enter email addresses in the template.

When you are finished, go to “Disable Network License Manager for MATLAB Parallel Server” on page 6-19.

## Disable Network License Manager for MATLAB Parallel Server

To change your installed network license manager to online licensing, you must shut down the service and either delete the license (if the license manager was only for MATLAB Parallel Server) or modify the license (if license manager is managing other products in addition to the server).

### If Network License Manager Manages Only MATLAB Parallel Server Cluster

Disable the network license manager, and delete the license.

Platform	Procedure
Windows	<ol style="list-style-type: none"> <li>1 Start your current <code>LMtools.exe</code>.</li> <li>2 Navigate to the <b>Start/Stop/Reread</b> tab and choose <b>Stop Server</b>.</li> <li>3 Navigate to the <b>Config Services</b> tab and choose <b>Remove Service</b>.</li> <li>4 Close the current <code>LMtools.exe</code>.</li> </ol>
Linux	<ol style="list-style-type: none"> <li>1 Run <code>lmdown</code>.</li> <li>2 Delete the existing network license manager by deleting the entire folder where you previously placed the license manager binaries.</li> </ol>

### If Network License Manager Manages Multiple Products

Modify the license to remove MATLAB Parallel Server.

- 1 Shut down the network license manager.

Platform	Procedure
Windows	In your current <code>LMtools.exe</code> , navigate to the <b>Start/Stop/Reread</b> tab and choose <b>Stop Server</b> .

Platform	Procedure
Linux	Run <code>lmdown</code> .

- 2 Open the MathWorks license file for editing:

```
matlabroot/etc/license.dat
```

The exact name of your license file and its location might vary. If you have trouble finding the license file, see [Where are the license files for MATLAB located?](#)

- 3 In the license file, locate and remove the entry that begins with the text:

```
INCREMENT MATLAB_Distrib_Comp_Engine ...
```

The entry might span more than one line in the file. Remove the entire entry and save the file.

- 4 Restart the network license manager for the license file change to take effect.

Network license manager utilities, such as `lmstat` - can verify that only the proper products are licensed. To use the `lmstat` function, open Terminal and run the following commands:

```
cd $MATLABROOT/etc/glnxa64
./lmutil lmstat -a
```

## See Also

### External Websites

- [MATLAB Parallel Server on Amazon Web Services \(GitHub\)](#)
- [MATLAB Parallel Server on Microsoft Azure \(GitHub\)](#)