

Release Notes for Parallel Computing Toolbox™

How to Contact MathWorks



www.mathworks.com Web
comp.soft-sys.matlab Newsgroup
www.mathworks.com/contact_TS.html Technical Support



suggest@mathworks.com Product enhancement suggestions
bugs@mathworks.com Bug reports
doc@mathworks.com Documentation error reports
service@mathworks.com Order status, license renewals, passcodes
info@mathworks.com Sales, pricing, and general information



508-647-7000 (Phone)



508-647-7001 (Fax)



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

For contact information about worldwide offices, see the MathWorks Web site.

Release Notes for Parallel Computing Toolbox™

© COPYRIGHT 2006–2012 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 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 for more information.

R2012b

More MATLAB functions enabled for GPUs, including convn, cov, and normest	2
GPU-enabled functions in Neural Network Toolbox, Phased Array System Toolbox, and Signal Processing Toolbox	4
Performance improvements to GPU-enabled MATLAB functions and random number generation	5
Automatic detection and selection of specific GPUs on a cluster node when multiple GPUs are available on the node	6
More MATLAB functions enabled for distributed arrays, including sparse constructor, bsxfun, and repmat	7
Detection of MATLAB Distributed Computing Server clusters that are available for connection from user desktops through Profile Manager	8
gpuArray Class Name	9
Diary Output Now Available During Running Task	10
Upgrade Parallel Computing Products Together	11

R2012a

New Programming Interface	14
Cluster Profiles	22
Enhanced GPU Support	25
Enhanced Distributed Array Support	30
Random Number Generation on Workers	31

R2011b

New Job Monitor	34
Run Scripts as Batch Jobs from the Current Folder Browser	35

Number of Local Workers Increased to Twelve	36
Enhanced GPU Support	37
Enhanced Distributed Array Support	40
Conversion of Error and Warning Message Identifiers ...	41
Task Error Properties Updated	42

R2011a

Deployment of Local Workers	44
New Desktop Indicator for MATLAB Pool Status	45
Enhanced GPU Support	46
Distributed Array Support	48
Enhanced parfor Support	49
Enhanced Support for Microsoft Windows HPC Server ...	50
Enhanced Admin Center Support	51
New Remote Cluster Access Object	52

R2010b

GPU Computing	54
Job Manager Security and Secure Communications	55
Generic Scheduler Interface Enhancements	56
batch Now Able to Run Functions	58
batch and matlabpool Accept Scheduler Object	59
Enhanced Functions for Distributed Arrays	60
Support for Microsoft Windows HPC Server 2008 R2	62
User Permissions for MDCEUSER on Microsoft Windows	63

R2010a

New Save and Load Abilities for Distributed Arrays	66
Enhanced Functions for Distributed Arrays	67
Importing Configurations Programmatically	68
Enhanced 2-D Block-Cyclic Array Distribution	69
New Remote Startup of mdce Process	70

Obtaining mdce Process Version	71
Demo Updates	72
taskFinish File for MATLAB Pool	73

R2009b

New Distributed Arrays	76
Renamed codistributor Functions	77
Enhancements to Admin Center	79
Adding or Updating File Dependencies in an Open MATLAB Pool	80
Updated globalIndices Function	81
Support for Job Templates and Description Files with HPC Server 2008	82
HPC Challenge Benchmarks	83
pctconfig Enhanced to Support Range of Ports	84
Random Number Generator on Client Versus Workers ...	85

R2009a

Number of Local Workers Increased to Eight	88
Admin Center Allows Controlling of Cluster Resources ...	89
Support for Microsoft Windows HPC Server 2008 (CCS v2)	90
New Benchmarking Demos	91
Pre-R2008b Distributed Array Syntax Now Generates Error	92
LSF Support on Mac OS X 10.5.x	93

R2008b

MATLAB Compiler Product Support for Parallel Computing Toolbox Applications	96
spm� Construct	97
Composite Objects	98
Configuration Validation	99

Rerunning Failed Tasks	100
Enhanced Job Control with Generic Scheduler Interface ..	101
Changed Function Names for Codistributed Arrays	102
Determining if a MATLAB Pool is Open	103

R2008a

Renamed Functions for Product Name Changes	106
New batch Function	107
Enhanced Job Creation Functions	108
Increased Data Size Transfers	109
Changed Function Names for Distributed Arrays	110
Support for PBS Pro and TORQUE Schedulers	111
findResource Now Sets Properties According to Configuration	112
parfor Syntax Has Single Usage	113
dfeval Now Destroys Its Job When Finished	114

R2007b

New Parallel for-Loops (parfor-Loops)	116
Configurations Manager and Dialogs	117
Default Configuration	118
Parallel Profiler	119
MDCE Script for Red Hat Removed	120

R2007a

Local Scheduler and Workers	122
New pmode Interface	123
New Default Scheduler for pmode	124
Vectorized Task Creation	125
Additional Submit and Decode Scripts	126
Jobs Property of Job Manager Sorts Jobs by ID	127
New Object Display Format	128
Enhanced MATLAB Functions	129

darray Function Replaces distributor Function	130
rand Seeding Unique for Each Task or Lab	131
Single-Threaded Computations on Workers	132

R2006b

Support for Windows Compute Cluster Server (CCS)	134
Windows 64 Support	135
Parallel Job Enhancements	136
Distributed Arrays	137
Interactive Parallel Mode (pmode)	138
Moved MDCE Control Scripts	139
rand Seeding Unique for Each Task or Lab	141
Task ID Property Now Same as labindex	142

R2012b

Version: 6.1
New Features: Yes
Bug Fixes: Yes

More MATLAB functions enabled for GPUs, including `convn`, `cov`, and `normest`

- “gpuArray Support” on page 2
- “MATLAB Code on the GPU” on page 2

gpuArray Support

The following functions are enhanced to support `gpuArray` data, or are expanded in their support:

<code>bitget</code>	<code>cov</code>	<code>normest</code>
<code>bitset</code>	<code>issparse</code>	<code>pow2</code>
<code>cond</code>	<code>mpower</code>	<code>var</code>
<code>convn</code>	<code>nnz</code>	

The following functions are not methods of the `gpuArray` class, but they now work with `gpuArray` data:

<code>blkdiag</code>	<code>ismatrix</code>	<code>isvector</code>
<code>cross</code>	<code>isrow</code>	<code>std</code>
<code>iscolumn</code>		

For complete lists of functions that support `gpuArray`, see “Built-In Functions That Support `gpuArray`”.

MATLAB Code on the GPU

`bsxfun` now supports the same subset of the language on a GPU that `arrayfun` does.

GPU support is extended to include the following MATLAB[®] code in functions called by `arrayfun` and `bsxfun` to run on the GPU:

bitget
bitset
pow2

GPU-enabled functions in Neural Network Toolbox, Phased Array System Toolbox, and Signal Processing Toolbox

A number of other toolboxes now support enhanced functionality for some of their functions to perform computations on a GPU. For specific information about these other toolboxes, see their respective release notes. Parallel Computing Toolbox™ is required to access this functionality.

Performance improvements to GPU-enabled MATLAB functions and random number generation

The performance of some MATLAB functions and random number generation on GPU devices is improved in this release.

You now have a choice of three random generators on the GPU: the combined multiplicative recursive MRG32K3A, the Philox4x32-10, and the Threefry4x64-20. For information on these generators and how to select them, see “Control the Random Stream for gpuArray”. For information about generating random numbers on a GPU, and a comparison between GPU and CPU generation, see “Control Random Number Streams”.

Automatic detection and selection of specific GPUs on a cluster node when multiple GPUs are available on the node

When multiple workers run on a single compute node with multiple GPU devices, the devices are automatically divided up among the workers. If there are more workers than GPU devices on the node, multiple workers share the same GPU device. If you put a GPU device in 'exclusive' mode, only one worker uses that device. As in previous releases, you can change the device used by any particular worker with the `gpuDevice` function.

More MATLAB functions enabled for distributed arrays, including sparse constructor, bsxfun, and repmat

The following functions now support distributed arrays with all forms of codistributor (1-D and 2DBC), or are enhanced in their support for this release:

atan2d	mrdivide	struct2cell
bsxfun	repmat	typecast
cell2struct	sparse	

Note the following enhancements for some of these functions:

- `cell2struct`, `struct2cell`, and `typecast` now support 2DBC in addition to 1-D distribution.
- `mrdivide` is now fully supported, and is no longer limited to accepting only scalars for its second argument.
- `sparse` is now fully supported for all distribution types.

This release also offers improved performance of fft functions for long vectors as distributed arrays.

Detection of MATLAB Distributed Computing Server clusters that are available for connection from user desktops through Profile Manager

You can let MATLAB discover clusters for you. Use either of the following techniques to discover those clusters which are available for you to use:

- On the **Home** tab in the **Environment** section, click **Parallel > Discover Clusters**.
- In the Cluster Profile Manager, click **Discover Clusters**.

For more information, see “Discover Clusters”.

gpuArray Class Name

Compatibility Considerations: Yes

The object type formerly known as a GPUArray has been renamed to gpuArray. The corresponding class name has been changed from parallel.gpu.GPUArray to the shorter gpuArray. The name of the function gpuArray remains unchanged.

Compatibility Considerations

You cannot load gpuArray objects from files that were saved in previous versions.

Code that uses the old class name must be updated to use the shorter new name. For example, the functions for directly generating gpuArrays on the GPU:

Previous version form	New version form
<code>parallel.gpu.GPUArray.rand</code> <code>parallel.gpu.GPUArray.ones</code>	<code>gpuArray.rand</code> <code>gpuArray.ones</code>
etc.	etc.

Diary Output Now Available During Running Task

Diary output from tasks (including those of batch jobs) can now be obtained while the task is still running. The diary is text output that would normally be sent to the Command Window. Now this text is appended to the task's `Diary` property as the text is generated, rather than waiting until the task is complete. You can read this property at any time. Diary information is accumulated only if the job's `CaptureDiary` property value is `true`. (**Note:** This feature is not yet available for SOA jobs on HPC Server clusters.)

Upgrade Parallel Computing Products Together

Compatibility Considerations: Yes

This version of Parallel Computing Toolbox software is accompanied by a corresponding new version of MATLAB Distributed Computing Server™ software.

Compatibility Considerations

As with every new release, if you are using both parallel computing products, you must upgrade Parallel Computing Toolbox and MATLAB Distributed Computing Server together. These products must be the same version to interact properly with each other.

Jobs created in one version of Parallel Computing Toolbox software will not run in a different version of MATLAB Distributed Computing Server software, and might not be readable in different versions of the toolbox software. The job data stored in the folder identified by `JobStorageLocation` (formerly `DataLocation`) might not be compatible between different versions of MATLAB Distributed Computing Server. Therefore, `JobStorageLocation` should not be shared by parallel computing products running different versions, and each version on your cluster should have its own `JobStorageLocation`.

R2012a

Version: 6.0
New Features: Yes
Bug Fixes: Yes

New Programming Interface

Compatibility Considerations: Yes

This release provides a new programming interface for accessing clusters, jobs, and tasks.

General Concepts and Phrases

This table maps some of the concepts and phrases from the old interface to the new.

Previous Interface	New Interface in R2012a
MathWorks job manager	MATLAB job scheduler (MJS)
Third-party or local scheduler	Common job scheduler (CJS)
Configuration	Profile
Scheduler	Cluster

The following code examples compare programming the old and new interfaces, showing some of the most common commands and properties. Note that most differences involve creating a cluster object instead of a scheduler object, and some of the property and method names on the job. After the example are tables listing some details of these new objects.

Previous Interface	New Interface
<pre> sched = findResource('scheduler',... 'Configuration','local'); j = createJob(sched,... 'PathDependencies',{'/share/app/'},... 'FileDependencies',{'funa.m','funb.m'}); createTask(j,@myfun,1,{3,4}); submit(j); waitForState(j); results = j.getAllOutputArguments; </pre>	<pre> clust = parcluster('local'); j = createJob(clust,... 'AdditionalPaths',{'/share/app/'},... 'AttachedFiles',{'funa.m','funb.m'}); createTask(j,@myfun,1,{3,4}); submit(j); wait(j); results = j.fetchOutputs; </pre>

Objects

These tables compare objects in the previous and new interfaces.

Previous Scheduler Objects	New Cluster Objects
ccsscheduler	parallel.cluster.HPCServer
genericscheduler	parallel.cluster.Generic
jobmanager	parallel.cluster.MJS
localscheduler	parallel.cluster.Local
lsfscheduler	parallel.cluster.LSF
mpiexec	parallel.cluster.Mpiexec
pbsproscheduler	parallel.cluster.PBSPro
torquescheduler	parallel.cluster.Torque

For information on each of the cluster objects, see the `parallel.Cluster` reference page.

Previous Job Objects	New Job Objects
job (distributed job)	parallel.job.MJSIndependentJob
matlabpooljob	parallel.job.MJSCommunicatingJob where ('Type' = 'Pool')
paralleljob	parallel.job.MJSCommunicatingJob where ('Type' = 'SPMD')
simplejob	parallel.job.CJSIndependentJob
simplematlabpooljob	parallel.job.CJSCommunicatingJob ('Type' = 'Pool')
simpleparalleljob	parallel.job.CJSCommunicatingJob ('Type' = 'SPMD')

For information on each of the job objects, see the `parallel.Job` reference page.

Previous Task Objects	New Task Objects
task	parallel.task.MJSTask
simpletask	parallel.task.CJSTask

For information on each of the task objects, see the `parallel.Task` reference page.

Previous Worker Object	New Worker Objects
worker	parallel.cluster.MJSWorker, parallel.cluster.CJSWorker

For information on each of the worker objects, see the `parallel.Worker` reference page.

Functions and Methods

This table compares some functions and methods of the old interface to those of the new. Many functions do not have a name change in the new interface, and are not listed here. Not all functions are available for all cluster types.

Previous Name	New Name
findResource	parcluster
createJob	createJob (no change)
createParallelJob	createCommunicatingJob (where 'Type' = 'SPMD')
createMatlabPoolJob	createCommunicatingJob (where 'Type' = 'Pool')
destroy	delete
clearLocalPassword	logout
createTask	createTask (no change)
getAllOutputArguments	fetchOutputs
getJobSchedulerData	getJobClusterData
setJobSchedulerData	setJobClusterData

Previous Name	New Name
getCurrentJobmanager	getCurrentCluster
getFileDependencyDir	getAttachedFilesFolder

Properties

In addition to a few new properties on the objects in the new interface, some other properties have new names when using the new interface, according to the following tables.

Previous Scheduler Properties	New Cluster Properties
DataLocation	JobStorageLocation
Configuration	Profile
HostAddress	AllHostAddresses
Computer	ComputerType
ClusterOsType	OperatingSystem
IsUsingSecureCommunication	HasSecureCommunication
SchedulerHostname, MasterName, Hostname	Host
ParallelSubmissionWrapperScript	CommunicatingJobWrapper
ClusterSize	NumWorkers
NumberOfBusyWorkers	NumBusyWorkers
NumberOfIdleWorkers	NumIdleWorkers
SubmitFcn	IndependentSubmitFcn
ParallelSubmitFcn	CommunicatingSubmitFcn
DestroyJobFcn	DeleteJobFcn
DestroyTaskFcn	DeleteTaskFcn

Previous Job Properties	New Job Properties
FileDependencies	AttachedFiles
PathDependencies	AdditionalPaths
MinimumNumberOfWorkers, MaximumNumberOfWorkers	NumWorkersRange

Previous Task Properties	New Task Properties
MaximumNumberOfRetries	MaximumRetries

Getting Help

In addition to these changes, there are some new properties and methods, while some old properties are not used in the new interface. For a list of the methods and properties available for clusters, jobs, and tasks, use the following commands for help on each class type:

```
help parallel.Cluster
help parallel.Job
help parallel.Task
```

```
help parallel.job.CJSIndependentJob
help parallel.job.CJSCommunicatingJob
help parallel.task.CJSTask
```

```
help parallel.job.MJSIndependentJob
help parallel.job.MJSCommunicatingJob
help parallel.task.MJSTask
```

There might be slight changes in the supported format for properties whose names are still the same. To get help on an individual method or property, the general form of the command is:

```
help parallel.obj-type.method-or-property-name
```

You might need to specify the subclass in some cases, for example, where different cluster types use properties differently. The following examples display the help for specified methods and properties of various object types:

```
help parallel.cluster.LSF.JobStorageLocation
help parallel.Job.fetchOutputs
help parallel.job.MJSIndependentJob.FinishedFcn
help parallel.Task.delete
```

Compatibility Considerations

Jobs

This release still supports the old form of interface, however, the old interface and the new interface are not compatible in the same job. For example, a job manager scheduler object that you create with `findResource` requires that you use only the old interface methods and properties, and cannot be used for creating a communicating job (`createCommunicatingJob`). A cluster that was defined with `parcluster` must use the new interface, and cannot be used to create a parallel job (`createParallelJob`).

Graphical interfaces provide access only to the new interface. The Configurations Manager is no longer available and is replaced by the Cluster Profile Manager. Actions that use profiles automatically convert and upgrade your configurations to profiles. Therefore, if you already have a group of configurations, the first time you open the Cluster Profile Manager, it should already be populated with your converted profiles. Furthermore, you can specify cluster profiles when using the old interface in places where configurations are expected.

One job manager (or MJS) can accommodate jobs of both the old and new interface at the same time.

Creating and Finding Jobs and Tasks

In the old interface, to create or find jobs or tasks, you could specify property name and value pairs in structures or cell arrays, and then provide the structure or cell array as an input argument to the function you used. In the new interface, property names and values must be pairs of separate arguments, with the property name as a string expression and

its value of the appropriate type. This applies to the functions `createJob`, `createCommunicatingJob`, `createTask`, `findJob`, and `findTask`.

Batch

Jobs created by the `batch` command use the new interface, even if you specify configurations or properties using the old interface. For example, the following code generates two identical jobs of the new interface, even though job `j2` is defined with an old interface property. The script `randScript` contains just the one line of code, `R = rand(3)`, and the default profile is `local`.

```
j1 = batch('randScript','AdditionalPaths','c:\temp');
j1.wait;
R1 = j1.load('R');
```

or

```
j2 = batch('randScript','PathDependencies','c:\temp');
j2.wait;
R2 = j2.load('R');
```

`whos`

Name	Size	Bytes	Class
R1	1x1	248	struct
R2	1x1	248	struct
j1	1x1	112	parallel.job.CJSIndependentJob
j2	1x1	112	parallel.job.CJSIndependentJob

Communicating Job Wrapper

In the old interface, for a parallel job in an LSF, TORQUE, or PBS Pro scheduler, you would call the scheduler's `setupForParallelExecution` method with the necessary arguments so that the toolbox could automatically set the object's `ClusterOSType` and `ParallelSubmissionWrapperScript` properties, thus determining which wrapper was used. In the new interface, with a communicating job you only have to set the LSF, Torque, or PBSPro cluster object's `OperatingSystem` and `CommunicatingJobWrapper` properties, from which the toolbox calculates which wrapper to use. For more information about these properties and their possible values, in MATLAB type

```
help parallel.cluster.LSF.CommunicatingJobWrapper
```

You can change the LSF to PBSPro or Torque in this command, as appropriate.

Enhanced Example Scripts

An updated set of example scripts is available in the product for using a generic scheduler with the new programming interface. These currently supported scripts are provided in the folder:

```
matlabroot/toolbox/distcomp/examples/integration
```

For more information on these scripts and their updates, see the README file provided in each subfolder, or see Supplied Submit and Decode Functions.

Compatibility Considerations


Scripts that use the old programming interface are provided in the folder *matlabroot/toolbox/distcomp/examples/integration/old*. The scripts that resided in this folder in previous releases are no longer available. The scripts currently in this folder might be removed in future releases.

Cluster Profiles

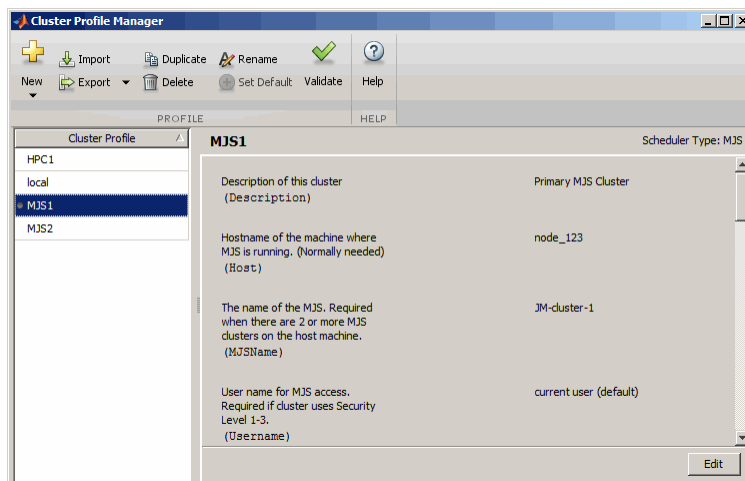
Compatibility Considerations: Yes

New Cluster Profile Manager

Cluster profiles replace parallel configurations for defining settings for clusters and their jobs. You can access profiles and the Cluster Profile Manager from the desktop **Parallel** menu. From this menu you can select or import profiles. To choose a profile as the default, select the desktop menu **Parallel > Select Cluster Profile**. The current default profile is indicated with a bold dot.



The Cluster Profile Manager lets you create, edit, validate, import, and export profiles, among other actions. To open the Cluster Profile Manager, select **Parallel > Manage Cluster Profiles**.



For more information about cluster profiles and the Cluster Profile Manager, see Cluster Profiles.

Programming with Profiles

These commands provide access to profiles and the ability to create cluster objects.

Function	Description
<code>p = parallel.clusterProfiles</code>	List of your profiles
<code>parallel.defaultClusterProfile</code>	Specifies which profile to use by default
<code>c = parcluster('clustername')</code>	Creates cluster object, <code>c</code> , for accessing parallel compute resources
<code>c.saveProfile</code>	Saves changes of cluster property values to its current profile
<code>c.saveAsProfile</code>	Saves cluster properties to a profile of the specified name, and sets that as the current profile for this cluster
<code>parallel.importProfile</code>	Import profiles from specified <code>.settings</code> file
<code>parallel.exportProfile</code>	Export profiles to specified <code>.settings</code> file

Profiles in Compiled Applications

Because compiled applications include the current profiles of the user who compiles, in most cases the application has the profiles it needs. When other profiles are needed, a compiled application can also import profiles that have been previously exported to a `.settings` file. The new `ParallelProfile` key supports exported parallel configuration `.mat` files and exported cluster profile `.settings` files; but this might change in a future release. For more information, see Export Profiles for MATLAB Compiler.

Compatibility Considerations

In past releases, when a compiled application imported a parallel configuration, that configuration would overwrite a configuration of the same if it existed. In this release, imported profiles are renamed if profiles already exist with the same name; so the same profile might have a different name in the compiled application than it does in the exported profile file.

Enhanced GPU Support

Compatibility Considerations: Yes

GPUArray Support

The following functions are added to those that support GPUArray data, or are enhanced in their support for this release:

<code>beta</code>	<code>ifft</code>	<code>mldivide</code>
<code>betaln</code>	<code>ifft2</code>	<code>min</code>
<code>bsxfun</code>	<code>ifftn</code>	<code>mrdivide</code>
<code>circshift</code>	<code>ind2sub</code>	<code>norm</code>
<code>det</code>	<code>int2str</code>	<code>num2str</code>
<code>eig</code>	<code>inv</code>	<code>permute</code>
<code>fft</code>	<code>ipermute</code>	<code>qr</code>
<code>fft2</code>	<code>isequaln</code>	<code>shiftdim</code>
<code>fftn</code>	<code>issorted</code>	<code>sprintf</code>
<code>fprintf</code>	<code>mat2str</code>	<code>sub2ind</code>
<code>full</code>	<code>max</code>	

Note the following enhancements and restrictions to some of these functions:

- GPUArray usage now supports all data types supported by MATLAB, except `int64` and `uint64`.
- For the list of functions that `bsxfun` supports, see the `bsxfun` reference page.
- The full range of syntax is now supported for `fft`, `fft2`, `fftn`, `ifft`, `ifft2`, and `ifftn`.
- `eig` now supports all matrices, symmetric or not.
- `issorted` supports only vectors, not matrices.
- `max` and `min` can now return two output arguments, including an index vector.
- `mldivide` supports complex arrays. It also supports overdetermined matrices (with more rows than columns), with no constraints on the second input.

- `mrdivide` supports underdetermined matrices (more columns than rows) as the second input argument.
- `norm` now supports the form `norm(X,2)`, where `X` is a matrix.

The following functions are not methods of the `GPUArray` class, but they do work with `GPUArray` data:

<code>angle</code>	<code>ifftshift</code>	<code>rank</code>
<code>fliplr</code>	<code>kron</code>	<code>squeeze</code>
<code>flipud</code>	<code>mean</code>	<code>rot90</code>
<code>flipdim</code>	<code>perms</code>	<code>trace</code>
<code>fftshift</code>		

Reset or Deselect GPU Device

Resetting a GPU device clears your `GPUArray` and `CUDAKernel` data from the device. There are two ways to reset a GPU device, while still keeping it as the currently selected device. You can use the `reset` function, or you can use `gpuDevice(idx)` with the current device's index for `idx`. For example, use `reset`:

```
idx = 1
g = gpuDevice(idx)
reset(g)
```

Alternatively, call `gpuDevice` again with the same index argument:

```
idx = 1
g = gpuDevice(idx)
.
.
.
g = gpuDevice(idx) % Resets GPU device, clears data
```

To deselect the current device, use `gpuDevice([])` with an empty argument (as opposed to no argument). This clears the GPU of all arrays and kernels, and invalidates variables in the workspace that point to such data.

Asynchronous GPU Calculations and Wait

All GPU calculations now run asynchronously with MATLAB. That is, when you initiate a calculation on the GPU, MATLAB continues execution while the GPU runs its calculations at the same time. The `wait` command now accommodates a GPU device, so that you can synchronize MATLAB and the GPU. The form of the command is

```
wait(gpudev)
```

where `gpudev` is the object representing the GPU device to wait for. At this command, MATLAB waits until all current calculations complete on the specified device.

Compatibility Considerations

In previous releases, MATLAB and the GPU were synchronous, so that any calls to the GPU had to complete before MATLAB proceeded to the next command. This is no longer the case. Now MATLAB continues while the GPU is running. The `wait` command lets you time GPU code execution.

Verify GPUArray or CUDAKernel Exists on the Device

The new function `existsOnGPU` lets you verify that a `GPUArray` or `CUDAKernel` exists on the current GPU device, and that its data is accessible from MATLAB. It is possible to reset the GPU device, so that a `GPUArray` or `CUDAKernel` object variable still exists in your MATLAB workspace even though it is no longer available on the GPU. For example, you can reset a GPU device using the command `gpuDevice(index)` or `reset(dev)`:

```
index = 1;
g = gpuDevice(index);
R = parallel.gpu.GPUArray.rand(4,4)
    0.5465    0.3000    0.4067    0.6110
    0.9024    0.8965    0.6635    0.7709
    0.8632    0.7481    0.9901    0.0420
    0.2307    0.7008    0.7516    0.5059

existsOnGPU(R)
    1
```

```
reset(g); % Resets GPU device
existsOnGPU(R)
    0
```

```
R % View GPUArray contents
Data no longer exists on the GPU.
```

Any attempt to use the data from R generates an error.

MATLAB Code on the GPU

GPU support is extended to include the following MATLAB code in functions called by `arrayfun` to run on the GPU:

<code>and</code>	<code>intmin</code>	<code>power</code>
<code>beta</code>	<code>ldivide</code>	<code>rdivide</code>
<code>betaln</code>	<code>le</code>	<code>realmax</code>
<code>eq</code>	<code>lt</code>	<code>realmin</code>
<code>ge</code>	<code>minus</code>	<code>times</code>
<code>gt</code>	<code>ne</code>	<code>uint8</code>
<code>int8</code>	<code>not</code>	<code>uint16</code>
<code>int16</code>	<code>or</code>	
<code>intmax</code>	<code>plus</code>	

Note the following enhancements and restrictions to some of these functions:

- `plus`, `minus`, `ldivide`, `rdivide`, `power`, `times`, and other arithmetic, comparison or logical operator functions were previously supported only when called with their operator symbol. Now they are supported in their functional form, so can be used as direct argument inputs to `arrayfun` and `bsxfun`.
- `arrayfun` and `bsxfun` on the GPU now support the integer data types `int8`, `uint8`, `int16`, and `uint16`.

The code in your function can now call any functions defined in your function file or on the search path. You are no longer restricted to calling only those supported functions listed in the table of Supported MATLAB Code.

Set CUDA Kernel Constant Memory

The new `setConstantMemory` method on the `CUDAKernel` object lets you set kernel constant memory from MATLAB. For more information, see the `setConstantMemory` reference page.

Latest NVIDIA CUDA Device Driver

This version of Parallel Computing Toolbox GPU functionality supports only the latest NVIDIA CUDA device driver.

Compatibility Considerations

Earlier versions of the toolbox supported earlier versions of CUDA device driver. Always make sure you have the latest CUDA device driver.

Enhanced Distributed Array Support

Newly Supported Functions

The following functions now support distributed arrays with all forms of codistributor (1-D and 2DBC), or are enhanced in their support for this release:

`isequaln`

Random Number Generation on Workers

Compatibility Considerations: Yes

MATLAB worker sessions now generate random number values using the combined multiplicative recursive generator (mrg32k3a) by default.

Compatibility Considerations


In past releases, MATLAB workers used the same default generator as a MATLAB client session. This is no longer the case.

R2011b


Version: 5.2
New Features: Yes
Bug Fixes: Yes

New Job Monitor

The Job Monitor is a tool that lets you track the jobs you have submitted to a cluster. It displays the jobs for the scheduler determined by your selection of a parallel configuration. Open the Job Monitor from the MATLAB desktop by selecting **Parallel > Job Monitor**.




Right-click a job in the list to select a command from the context menu for that job:



For more information about the Job Monitor and its capabilities, see [Job Monitor](#).

Run Scripts as Batch Jobs from the Current Folder Browser

From the Current Folder browser, you can run a MATLAB script as a batch job by browsing to the file's folder, right-clicking the file, and selecting **Run Script as Batch Job**. The batch job runs on the cluster identified by the current default parallel configuration. The following figure shows the menu option to run the script from the file `script1.m`:



Number of Local Workers Increased to Twelve

You can now run up to 12 local workers on your MATLAB client machine. If you do not specify the number of local workers in a command or configuration, the default number of local workers is determined by the value of the local scheduler's `ClusterSize` property, which by default equals the number of computational cores on the client machine.

Enhanced GPU Support

Compatibility Considerations: Yes

Latest NVIDIA CUDA Device Driver

This version of Parallel Computing Toolbox GPU functionality supports only the latest NVIDIA CUDA device driver.

Compatibility Considerations

Earlier versions of the toolbox supported earlier versions of CUDA device driver. Always make sure you have the latest CUDA device driver.

Deployment of GPU Applications

MATLAB Compiler™ generated standalone executables and components now support applications that use the GPU.

Random Number Generation

You can now directly create arrays of random numbers on the GPU using these new static methods for GPUArray objects:

```
parallel.gpu.GPUArray.rand  
parallel.gpu.GPUArray.randi  
parallel.gpu.GPUArray.randn
```

The following functions set the GPU random number generator seed and stream:

```
parallel.gpu.rng  
parallel.gpu.RandStream
```

Also, `arrayfun` called with GPUArray data now supports `rand`, `randi`, and `randn`. For more information about using `arrayfun` to generate random matrices on the GPU, see [Generating Random Numbers on the GPU](#).

GPUArray Support

The following functions now support GPUArray data:

<code>chol</code>	<code>isnan</code>	<code>norm</code>
<code>diff</code>	<code>lu</code>	<code>not</code>
<code>eig</code>	<code>max</code>	<code>repmat</code>
<code>find</code>	<code>min</code>	<code>sort</code>
<code>isfinite</code>	<code>mldivide</code>	<code>svd</code>
<code>isinf</code>		

`mldivide` supports complex arrays. It also supports overdetermined matrices (with more rows than columns) when the second input argument is a column vector (has only one column).

`eig` supports only symmetric matrices.

`max` and `min` return only one output argument; they do not return an index vector.

The following functions are not methods of the `GPUArray` class, but they do work with `GPUArray` data:

<code>angle</code>	<code>flipdim</code>	<code>mean</code>
<code>beta</code>	<code>fftshift</code>	<code>perms</code>
<code>betaln</code>	<code>ifftshift</code>	<code>squeeze</code>
<code>fliplr</code>	<code>kron</code>	<code>rot90</code>
<code>flipud</code>		

The default display of `GPUArray` variables now shows the array contents. In previous releases, the display showed some of the object properties, but not the contents. For example, the new enhanced display looks like this:

```
M = gpuArray(magic(3))
M =
     8     1     6
     3     5     7
     4     9     2
```

To see that `M` is a `GPUArray`, use the `whos` or `class` function.

MATLAB Code on the GPU

GPU support is extended to include the following MATLAB code in functions called by `arrayfun` to run on the GPU:

```
rand  
randi  
randn  
xor
```

Also, the handle passed to `arrayfun` can reference a simple function, a subfunction, a nested function, or an anonymous function. The function passed to `arrayfun` can call any number of its subfunctions. The only restriction is that when running on the GPU, nested and anonymous functions do not have access to variables in the parent function workspace. For more information on function structure and relationships, see [Types of Functions](#).

Enhanced Distributed Array Support

Newly Supported Functions

The following functions are enhanced to support distributed arrays, supporting all forms of codistributor (1-D and 2DBC):

```
inv  
meshgrid  
ndgrid  
sort
```

The following functions can now directly construct codistributed arrays:

```
codistributed.linspace(m, n, ..., codist)  
codistributed.logspace(m, n, ..., codist)
```


Conversion of Error and Warning Message Identifiers

Compatibility Considerations: Yes

For R2011b, error and warning message identifiers have changed in Parallel Computing Toolbox.

Compatibility Considerations

If you have scripts or functions that use message identifiers that changed, you must update the code to use the new identifiers. Typically, message identifiers are used to turn off specific warning messages, or in code that uses a try/catch statement and performs an action based on a specific error identifier.

For example, the 'distcomp:old:ID' identifier has changed to 'parallel:similar:ID'. If your code checks for 'distcomp:old:ID', you must update it to check for 'parallel:similar:ID' instead.

To determine the identifier for a warning, run the following command just after you see the warning:

```
[MSG,MSGID] = lastwarn;
```

This command saves the message identifier to the variable MSGID.

To determine the identifier for an error, run the following commands just after you see the error:

```
exception = MException.last;  
MSGID = exception.identifier;
```

Tip Warning messages indicate a potential issue with your code. While you can turn off a warning, a suggested alternative is to change your code so that it runs warning-free.

Task Error Properties Updated

Compatibility Considerations: Yes

If there is an error during task execution, the task `Error` property now contains the non-empty `MException` object that is thrown. If there was no error during task execution, the `Error` property is empty.

The identifier and message properties of this object are now the same as the task's `ErrorIdentifier` and `ErrorMessage` properties, respectively. For more information about these properties, see the `Error`, `ErrorIdentifier`, and `ErrorMessage` reference pages.

Compatibility Considerations

In past releases, when there was no error, the `Error` property contained an `MException` object with empty data fields, generated by `MException(' ', ' ')`. Now to determine if a task is error-free, you can query the `Error` property itself to see if it is empty:

```
didTaskError = ~isempty(t.Error)
```

where `t` is the task object you are examining.

R2011a

Version: 5.1
New Features: Yes
Bug Fixes: Yes

Deployment of Local Workers

MATLAB Compiler generated standalone executables and libraries from parallel applications can now launch up to eight local workers without requiring MATLAB Distributed Computing Server software.

New Desktop Indicator for MATLAB Pool Status

When you first open a MATLAB pool from your desktop session, an indicator appears in the lower-right corner of the desktop to show that this desktop session is connected to an open pool. The number indicates how many workers are in the pool.



When you close the pool, the indicator remains displayed and shows a value of 0.

Enhanced GPU Support

Compatibility Considerations: Yes

Static Methods to Create GPUArray

The following new static methods directly create GPUArray objects:

```
parallel.gpu.GPUArray.linspace  
parallel.gpu.GPUArray.logspace
```

GPUArray Support

The following functions are enhanced to support GPUArray data:

```
cat  
colon  
conv  
conv2  
cumsum  
cumprod  
eps  
filter  
filter2  
horzcat  
meshgrid  
ndgrid  
plot  
subsasgn  
subsindex  
subsref  
vertcat
```

and all the plotting related functions.

GPUArray Indexing

Because GPUArray now supports `subsasgn` and `subsref`, you can index into a GPUArray for assigning and reading individual elements.

For example, create a GPUArray and assign the value of an element:

```
n = 1000;  
D = parallel.gpu.GPUArray.eye(n);  
D(1,n) = pi
```

Create a GPUArray and read the value of an element back into the MATLAB workspace:

```
m = 500;  
D = parallel.gpu.GPUArray.eye(m);  
one = gather(D(m,m))
```

MATLAB Code on the GPU

GPU support is extended to include the following MATLAB code in functions called by `arrayfun` to run on the GPU:

```
&, |, ~, &&, ||,  
while, if, else, elseif, for, return, break, continue, eps
```

You can now call `eps` with string inputs, so your MATLAB code running on the GPU can include `eps('single')` and `eps('double')`.

NVIDIA CUDA Driver 3.2 Support

This version of Parallel Computing Toolbox GPU functionality supports only NVIDIA CUDA device driver 3.2.

Compatibility Considerations

Earlier versions of the toolbox supported earlier versions of CUDA device driver. If you have an older driver, you must upgrade to CUDA device driver version 3.2.

Distributed Array Support

Compatibility Considerations: Yes

Newly Supported Functions

The following functions are enhanced to support distributed arrays, supporting all forms of codistributor (1-D and 2DBC):

```
arrayfun  
cat  
reshape
```

Enhanced `mtimes` Support

The `mtimes` function now supports distributed arrays that use a 2-D block-cyclic (2DBC) distribution scheme, and distributed arrays that use 1-D distribution with a distribution dimension greater than 2. Previously, `mtimes` supported only 1-D distribution with a distribution dimension of 1 or 2.

The `mtimes` function now returns a distributed array when only one of its inputs is distributed, similar to its behavior for two distributed inputs.

Compatibility Considerations

In previous releases, `mtimes` returned a replicated array when one input was distributed and the other input was replicated. Now it returns a distributed array.

Enhanced parfor Support

Nested for-Loops Inside parfor

You can now create nested for-loops inside a parfor-loop, and you can use both the parfor-loop and for-loop variables directly as indices for the sliced array inside the nested loop. See [Nested Loops](#).

Enhanced Support for Microsoft Windows HPC Server

Support for 32-Bit Clients

The parallel computing products now support Microsoft® Windows® HPC Server on 32-bit Windows clients.

Search for Cluster Head Nodes Using Active Directory

The `findResource` function can search Active Directory to identify your cluster head node. For more information, see the `findResource` reference page.

Enhanced Admin Center Support

You can now start and stop mdce services on remote hosts from Admin Center. For more information, see [Start mdce Service](#).

New Remote Cluster Access Object

New functionality is available that lets you mirror job data from a remote cluster to a local data location. This supports the generic scheduler interface when making remote submissions to a scheduler or when using a nonshared file system. For more information, see the `RemoteClusterAccess` object reference page.

R2010b

Version: 5.0
New Features: Yes
Bug Fixes: Yes

GPU Computing

This release provides the ability to perform calculations on a graphics processing unit (GPU). Features include the ability to:

- Use a GPU array interface with several MATLAB built-in functions so that they automatically execute with single- or double-precision on the GPU — functions including `mldivide`, `mtimes`, `fft`, etc.
- Create kernels from your MATLAB function files for execution on a GPU
- Create kernels from your CU and PTX files for execution on a GPU
- Transfer data to/from a GPU and represent it in MATLAB with GPUArray objects
- Identify and select which one of multiple GPUs to use for code execution

For more information on all of these capabilities and the requirements to use these features, see GPU Computing.

Job Manager Security and Secure Communications

You now have a choice of four security levels when using the job manager as your scheduler. These levels range from no security to user authentication requiring passwords to access jobs on the scheduler.

You also have a choice to use secure communications between the job manager and workers.

For more detailed descriptions of these features and information about setting up job manager security, see [Set MJS Cluster Security](#).

The default setup uses no security, to match the behavior of past releases.

Generic Scheduler Interface Enhancements

Compatibility Considerations: Yes

Decode Functions Provided with Product

Generic scheduler interface decode functions for distributed and parallel jobs are now provided with the product. The two decode functions are named:

```
parallel.cluster.generic.distributedDecodeFcn  
parallel.cluster.generic.parallelDecodeFcn
```

These functions are included on the workers' path. If your submit functions make use of the definitions in these decode functions, you do not have to provide your own decode functions. For example, to use the standard decode function for distributed jobs, in your submit function set `MDCE_DECODE_FUNCTION` to `'parallel.cluster.generic.distributedDecodeFcn'`. For information on using the generic scheduler interface with submit and decode functions, see [Use the Generic Scheduler Interface](#).

Enhanced Example Scripts

This release provides new sets of example scripts for using the generic scheduler interface. As in previous releases, the currently supported scripts are provided in the folder

`matlabroot/toolbox/distcomp/examples/integration`

In this location there is a folder for each type of scheduler:

- `lsf` — Platform LSF®
- `pbs` — PBS
- `sge` — Sun™ Grid Engine
- `ssh` — generic UNIX®-based scripts
- `winmpiexec` — mpiexec on Windows

For the updated scheduler folders (`lsf`, `pbs`, `sge`), subfolders within each specify scripts for different cluster configurations: `shared`, `nonshared`, `remoteSubmission`.

For more information on the scripts and their updates, see the `README` file provided in each folder, or see `Supplied Submit and Decode Functions`.

Compatibility Considerations

For those schedulers types with updated scripts in this release (`lsf`, `pbs`, `sge`), the old versions of the scripts are provided in the folder `matlabroot/toolbox/distcomp/examples/integration/old`. These old scripts might be removed in future releases.

batch Now Able to Run Functions

Batch jobs can now run functions as well as scripts. For more information, see the `batch` reference page.

batch and matlabpool Accept Scheduler Object

The `batch` function and the functional form of `matlabpool` now accept a scheduler object as their first input argument to specify which scheduler to use for allocation of compute resources. For more information, see the `batch` and `matlabpool` reference pages.

Enhanced Functions for Distributed Arrays

Compatibility Considerations: Yes

qr Supports Distributed Arrays

The `qr` function now supports distributed arrays. For restrictions on this functionality, type

```
help distributed/qr
```

mldivide Enhancements

The `mldivide` function (`\`) now supports rectangular distributed arrays. Formerly, only square matrices were supported as distributed arrays.

When operating on a square distributed array, if the second input argument (or right-hand side of the operator) is replicated, `mldivide` now returns a distributed array.

Compatibility Considerations

In previous releases, `mldivide` returned a replicated array when the second (or right-hand side) input was replicated. Now it returns a distributed array.

chol Supports 'lower' option

The `chol` function now supports the 'lower' option when operating on distributed arrays. For information on using `chol` with distributed arrays, type

```
help distributed/chol
```

eig and svd Return Distributed Array

When returning only one output matrix, the `eig` and `svd` functions now return a distributed array when the input is distributed. This behavior is now consistent with outputs when requesting more than one matrix, which returned distributed arrays in previous releases.

Compatibility Considerations

In previous releases, `eig` and `svd` returned a replicated array when you requested a single output. Now they return a distributed array if the output is a single matrix. The behavior when requesting more than one output is not changed.

transpose and ctranspose Support 2dbc

In addition to their original support for 1-D distribution schemes, the functions `ctranspose` and `transpose` now support 2-D block-cyclic ('2dbc') distributed arrays.

Inf and NaN Support Multiple Formats

Distributed and codistributed arrays now support `nan`, `NaN`, `inf` and `Inf` for not-a-number and infinity values with the following functions:

Infinity Value	Not-a-Number
<code>codistributed.Inf</code>	<code>codistributed.NaN</code>
<code>codistributed.inf</code>	<code>codistributed.nan</code>
<code>distributed.Inf</code>	<code>distributed.NaN</code>
<code>distributed.inf</code>	<code>distributed.nan</code>

Support for Microsoft Windows HPC Server 2008 R2

Parallel Computing Toolbox software now supports Microsoft Windows HPC Server 2008 R2. There is no change in interface compared to using HPC Server 2008. Configurations and other toolbox utilities use the same settings to support both HPC Server 2008 and HPC Server 2008 R2.

User Permissions for MDCEUSER on Microsoft Windows

Compatibility Considerations: Yes

The user identified by the MDCEUSER parameter in the mdce_def file is now granted all necessary privileges on a Windows system when you install the mdce process. For information about what permissions are granted and how to reset them, see [Set the User](#).

Compatibility Considerations

In past releases, you were required to set the MDCEUSER permissions manually. Now this is done automatically when installing the mdce process.

R2010a

Version: 4.3
New Features: Yes
Bug Fixes: Yes

New Save and Load Abilities for Distributed Arrays

You now have the ability to save distributed arrays from the client to a single MAT-file. Subsequently, in the client you can load a distributed array from that file and have it automatically distributed to the MATLAB pool workers. The pool size and distribution scheme of the array do not have to be the same when you load the array as they were when you saved it.

You also can now load data directly into distributed arrays, even if the originally saved arrays were not distributed.

For more information, see the `dsave` and `dload` reference pages.

Enhanced Functions for Distributed Arrays

The `svd` function now supports single-precision and complex data in distributed arrays. Other functions enhanced to support single-precision distributed arrays are `chol`, `lu`, `mldivide`, and `eig`.

In addition to their original support for 1-D distribution by columns, the enhanced `tril` and `triu` functions now support arrays with 1-D distribution by rows or when the distribution dimension is greater than 2, and 2-D block-cyclic ('2dbc') distributed arrays.

Importing Configurations Programmatically

A new function allows you to programmatically import parallel configurations. For more information, see the `importParallelConfig` reference page.

Enhanced 2-D Block-Cyclic Array Distribution

2-D block-cyclic ('2dbc') array distribution now supports column orientation of the lab grid. The `codistributor2dbc` function now accepts the value 'col' for its orientation argument, which is reflected in the codistributor object's `Orientation` property. For information on '2dbc' distribution and using lab grids, see 2-Dimensional Distribution.

New Remote Startup of mdce Process

New command-line functionality allows you to remotely start up MATLAB Distributed Computing Server processes on cluster machines from the desktop computer. For more information, see the `remotemdce` reference page.

Obtaining mdce Process Version

An enhancement to the `mdce` command lets you get the command version by executing

```
mdce -version
```

For more information on this command, see the `mdce` reference page.

Demo Updates

Product demos are available in the Demos node under Parallel Computing Toolbox in the help browser.

Benchmarking A\b

This new demo benchmarks Parallel Computing Toolbox performance with the `mldivide` function.

BER Performance of Equalizer Types

Demos of BER Performance of Several Equalizer Types have been removed from Parallel Computing Toolbox, because Communications Toolbox now incorporates support for Parallel Computing Toolbox. See the Communications Toolbox release notes.

taskFinish File for MATLAB Pool

Compatibility Considerations: Yes

The taskFinish file (taskFinish.m) for a MATLAB pool now executes when the pool closes.

Compatibility Considerations

In previous versions of the software, the taskFinish file executed when a MATLAB pool opened. Beginning with this release, it runs when the pool closes.

R2009b

Version: 4.2
New Features: Yes
Bug Fixes: Yes

New Distributed Arrays

A new form of distributed arrays provides direct access from the client to data stored on the workers in a MATLAB pool. Distributed arrays have the same appearance and rules of indexing as regular arrays.

You can distribute an existing array from the client workspace with the command

```
D = distributed(X)
```

where X is an array in the client, and D is a distributed array with its data on the workers in the MATLAB pool. Distributing an array is performed outside an `spm` statement, but a MATLAB pool must be open.

Codistributed arrays that you create on the workers within `spm` statements are accessible on the client as distributed arrays.

The following new functions and methods support distributed arrays.

Function Name	Description
<code>distributed</code>	Distribute existing array from client workspace to workers
<code>distributed.rand</code> , <code>distributed.ones</code> , etc.	Create distributed array consistent with indicated method, constructing on workers only
<code>gather</code>	Transfer data from MATLAB pool workers to client
<code>isdistributed</code>	True for distributed array
<code>C(x,y)</code>	Indexing into distributed array C on client to access data stored as codistributed arrays on workers

Renamed codistributor Functions

Compatibility Considerations: Yes

As part of the general enhancements for distributed arrays, several changes to the codistributed interface appear in this release.

Compatibility Considerations

The following table summarizes the changes in function names relating to codistributed arrays.

Old Function Name	New Function Name
codcolon	codistributed.colon
codistributed(..., 'convert')	codistributed(...)
codistributed(...) without 'convert' option	codistributed.build
codistributed(L, D) using distribution scheme of D to define that of L	codistributed.build(L, getCodistributor(D))
codistributor('1d', ...)	Still available, but can also use codistributor1d
codistributor('2d', ...)	codistributor('2dbc', ...) or codistributor2dbc
codistributor(arrayname)	getCodistributor
defaultLabGrid	codistributor2dbc.defaultLabGrid
defaultPartition	codistributor1d.defaultPartition
isa(X, 'codistributed')	Still available, but can also use iscodistributed(X)
localPart	getLocalPart
redistribute(D) using default distribution scheme	redistribute(D, codistributor())
redistribute(D1, D2) using distribution scheme of D2 to define that of D1	redistribute(D1, getCodistributor(D2))

Some object methods are now properties:

Old Method Name	New Property Name
<code>blockSize(codistObj)</code>	<code>codistObj.BlockSize</code>
<code>defaultBlockSize</code>	<code>codistributor2dbc.defaultBlockSize</code>
<code>distributionDimension(codistObj)</code>	<code>codistObj.Dimension</code>
<code>distributionPartition(codistObj)</code>	<code>codistObj.Partition</code>
<code>labGrid(codistObj)</code>	<code>codistObj.LabGrid</code>

Enhancements to Admin Center

Admin Center has several small enhancements, including more conveniently located menu choices, modified dialog boxes, properties dialog boxes for listed items, etc.

Adding or Updating File Dependencies in an Open MATLAB Pool

Enhancements to the `matlabpool` command let you add or update file dependencies in a running MATLAB pool. The new forms of the command are

```
matlabpool('addfiledependencies', filedepCell)
matlabpool updatefiledependencies
```

where `filedepCell` is a cell array of strings, identical in form to those you use when adding file dependencies to a job or when you open a MATLAB pool. The `updatefiledependencies` option replicates any file dependency changes to all the labs in the pool.

Updated globalIndices Function

Compatibility Considerations: Yes

The `globalIndices` function now requires that you specify the dimension of distribution as its second argument. Because this argument is required, it must precede the optional argument specifying the lab.

Compatibility Considerations

In previous toolbox versions, the `globalIndices` function accepted the lab argument before the dimension argument, and both were optional. Now the dimension argument is required, and it must precede the optional lab argument.

Support for Job Templates and Description Files with HPC Server 2008

Compatibility Considerations: Yes

Using job templates and job description files with Windows HPC Server 2008 lets you specify nodes and other scheduler properties for evaluating your jobs. To support these features, the `ccsscheduler` object has new properties:

- `ClusterVersion` — A string set to 'CCS' or 'HPCServer2008'
- `JobTemplate` — A string set to the name of the job template to use for all jobs
- `JobDescriptionFile` — A string set to the name of the XML file defining a base state for job creation

Compatibility Considerations

CCS is now just one of multiple versions of HPC Server. While 'ccs' is still acceptable as a type of scheduler for the `findResource` function, you can also use 'hpcserver' for this purpose. In the Configurations Manager, the new scheduler type is available by selecting **File > New > hpcserver (ccs)**.

HPC Challenge Benchmarks

Several new MATLAB files are available to demonstrate HPC Challenge benchmark performance. You can find the files in the folder *matlabroot/toolbox/distcomp/examples/benchmark/hpcchallenge*. Each file is self-documented with explanatory comments.

pctconfig Enhanced to Support Range of Ports

Compatibility Considerations: Yes

The `pctconfig` function now lets you specify a range of ports for the Parallel Computing Toolbox client session to use. This range also includes ports used for a `pmode` session.

Compatibility Considerations

You now specify the range of `pctconfig` ports with the `'portrange'` property; you no longer use the `'port'` property. As any client `pmode` session uses those ports of the `'portrange'` setting, you no longer use the `'pmodeport'` property.

Random Number Generator on Client Versus Workers

Compatibility Considerations: Yes

The random number generator of the MATLAB workers now use a slightly different seed from previous releases, so that all the MATLAB workers and the client have separate random number streams.

Compatibility Considerations

In past releases, while all the workers running a job had separate random number streams, the client had the same stream as one of the workers. Now the workers all have unique random number streams different from that of the client.

R2009a

Version: 4.1
New Features: Yes
Bug Fixes: Yes

Number of Local Workers Increased to Eight

Compatibility Considerations: Yes

You can now run up to eight local workers on your MATLAB client machine. If you do not specify the number of local workers in a command or configuration, the default number of local workers is determined by the value of the local scheduler's `ClusterSize` property, which by default is equal to the number of computational cores on the client machine.

Compatibility Considerations

In previous versions, the default number of local workers was four, regardless of the number of cores. If you want to run more local workers than cores (for example, four workers with only one or two cores), you must set the value of `ClusterSize` equal to or greater than the number of workers you need. Then you can specify the increased number of workers in the appropriate command or configuration, or let your `ClusterSize` setting control the default number of workers.

Admin Center Allows Controlling of Cluster Resources

Compatibility Considerations: Yes

When using the MathWorks job manager, the Admin Center GUI now allows you to start, stop, and otherwise control job managers and MATLAB workers on your cluster nodes. For more information about Admin Center, see Admin Center in the MATLAB Distributed Computing Server documentation.

Compatibility Considerations

You can no longer start Admin Center from the MATLAB Desktop **Parallel** pull-down menu. You must start Admin Center from outside MATLAB by executing the following:

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

Support for Microsoft Windows HPC Server 2008 (CCS v2)

The parallel computing products now support Microsoft Windows HPC Server 2008 (CCS v2), including service-oriented architecture (SOA) job submissions. There is no change to the programming interface for `ccs` options, other than the addition of a new CCS scheduler object property, `UseSOAJobSubmission`. For implications to the installation of the MATLAB Distributed Computing Server, see the online installation instructions at <http://www.mathworks.com/distconfig>.

New Benchmarking Demos

New benchmarking demos for Parallel Computing Toolbox can help you understand and evaluate performance of the parallel computing products. You can access these demos in the Help Browser under the **Parallel Computing Toolbox** node: expand the nodes for **Demos** then **Benchmarks**.

Pre-R2008b Distributed Array Syntax Now Generates Error

In R2008b, distributed array syntax was updated for codistributed arrays. In that release, the old form of the syntax still worked, but generated a warning. Now in R2009a, the old forms of the syntax no longer work and generate an error. For a summary of the syntax updates, see “Changed Function Names for Codistributed Arrays” on page 102.

LSF Support on Mac OS X 10.5.x

For availability of Platform LSF support on Macintosh OS X 10.5.x, contact Platform Computing Corporation via their Web site at <http://www.platform.com/Products/platform-lsf/technical-information>. If Platform Computing does not support LSF on Mac OS X 10.5.x, then Parallel Computing Toolbox and MATLAB Distributed Computing Server cannot support this combination.

R2008b

Version: 4.0
New Features: Yes
Bug Fixes: Yes

MATLAB Compiler Product Support for Parallel Computing Toolbox Applications

This release offers the ability to convert Parallel Computing Toolbox applications, using MATLAB Compiler, into executables and shared libraries that can access MATLAB Distributed Computing Server. For information on this update to MATLAB Compiler, see *Applications Created with Parallel Computing Toolbox Can Be Compiled*.

Limitations

- MATLAB Compiler does not support configurations that use the local scheduler or local workers (i.e., workers that run locally on the desktop machine running the MATLAB client session).
- Compiled Parallel Computing Toolbox applications do not support Simulink® software. For a list of other unsupported products, see the Web page http://www.mathworks.com/products/ineligible_programs/.
- When workers are running a task from compiled code, they can execute only compiled code and toolbox code. They cannot execute functions contained in the current directory. Batch and MATLAB pool jobs attempt to change the worker working directory to the client working directory. When noncompiled files in the current directory conflict with compiled versions (for example, files with different extensions), an error is thrown.

spmd Construct

Compatibility Considerations: Yes

A new single program multiple data (spmd) language construct allows enhanced interleaving of serial and parallel programming, with interlab communication.

The general form of an spmd statement is:

```
spmd
    <statements>
end
```

The block of code represented by <statements> executes in parallel on workers in the MATLAB pool. Data on the labs is available for access from the client via Composite objects. For more information, see the spmd reference page and Distributing Arrays and Running SPMD.

Compatibility Considerations

Because spmd is a new keyword, it will conflict with any user-defined functions or variables of the same name. If you have any code with functions or variables named spmd, you must rename them.

Composite Objects

Composite objects provide direct access from the client (desktop) program to data that is stored on labs in the MATLAB pool. The data of variables assigned inside an `spmd` block is available via Composites in the client. When a MATLAB pool is open, you can also create Composites directly from the client using the `Composite` function. See also [Distributing Arrays](#) and [Running SPMD](#).

Configuration Validation

The Configurations Manager is enhanced with the capability for validating configurations. Open the Configurations Manager on the MATLAB Desktop by clicking **Parallel > Manage Configurations**. For more information, see [Validate Profiles](#).

Rerunning Failed Tasks

When using a job manager, if a task does not complete due to certain system failures, it can attempt to rerun up to a specified number of times. New properties of a task object to control reruns and access information about rerun attempts are:

- `MaximumNumberOfRetries`
- `AttemptedNumberOfRetries`
- `FailedAttemptInformation`

Enhanced Job Control with Generic Scheduler Interface

The generic scheduler interface now allows you to cancel and destroy jobs and tasks and to investigate the state of a job. The following new properties of the generic scheduler object facilitate these features:

- `GetJobStateFcn`
- `DestroyJobFcn`
- `DestroyTaskFcn`
- `CancelJobFcn`
- `CancelTaskFcn`

New toolbox functions to accommodate this ability are:

- `getJobSchedulerData`
- `setJobSchedulerData`

For more information on this new functionality, see [Manage Jobs with Generic Scheduler](#).

Changed Function Names for Codistributed Arrays

Compatibility Considerations: Yes

What was known in previous releases as distributed arrays are henceforth called *codistributed* arrays. Some functions related to constructing and accessing codistributed arrays have changed names in this release.

Compatibility Considerations

The following table summarizes the changes in function names relating to codistributed arrays. The first three functions behave exactly the same with no change in operation, arguments, etc. The `isa` function takes the argument 'codistributed' in addition to the array in question.

Old Function Name	New Function Name
<code>distributed</code>	<code>codistributed</code>
<code>distributor</code>	<code>codistributor</code>
<code>dcolon</code>	<code>codcolon</code>
<code>isdistributed</code>	<code>isa(X, 'codistributed')</code>

Determining if a MATLAB Pool is Open

The function `matlabpool` now allows you to discover if a pool of workers is already open. The form of the command is:

```
matlabpool size
```

For more information about this option and others, see the `matlabpool` reference page.

R2008a

Version: 3.3
New Features: Yes
Bug Fixes: Yes

Renamed Functions for Product Name Changes

Compatibility Considerations: Yes

As of result of the product name changes, some function names are changing in this release.

Compatibility Considerations

Two function names are changed to correspond to the new product names:

- `dctconfig` has been renamed `pctconfig`.
- `dctRunOnAll` has been renamed `pctRunOnAll`.

New batch Function

The new `batch` function allows you to offload work from the client to one or more workers. The batch submission can run scripts that can include jobs that distribute work to other workers. For more information, see the `batch` reference page, and Getting Started in the Parallel Computing Toolbox User's Guide.

New Matlabpool Job

The batch functionality is implemented using the new `MATLABpool` job feature. A `MATLAB` pool job uses one worker to distribute a job to other workers, thereby freeing the client from the burden of tracking and job's progress and manipulating data. For more information, see the `createMatlabPoolJob` reference page.

Enhanced Job Creation Functions

The `createJob` and `createParallelJob` functions have been enhanced to run without requiring a scheduler object as an argument. This is also true for the new `createMatlabPoolJob` function. When a scheduler is not specified, the function uses the scheduler identified in the applicable parallel configuration. For details, see the reference page for each function.

Increased Data Size Transfers

The default size limitation on data transfers between clients and workers has been significantly increased. In previous releases the default limitation imposed by the JVM memory allocation was approximately 50 MB. The new higher limits are approximately 600 MB on 32-bit systems and 2 GB on 64-bit systems. See [Object Data Size Limitations](#).

Changed Function Names for Distributed Arrays

Compatibility Considerations: Yes

Several functions related to distributed arrays have changed names in this release.

Compatibility Considerations

The following table summarizes the changes in function names relating to distributed arrays.

Old Function Name	New Function Name
darray	distributed, distributor
distribute	distributed
dcolonpartition	defaultPartition
distribdim	distributionDimension
isdarray	isdistributed
labgrid	labGrid
local	localPart
partition	distributionPartition
localspan	globalIndices

Support for PBS Pro and TORQUE Schedulers

Parallel Computing Toolbox software now fully supports PBS Pro® and TORQUE schedulers. These schedulers are integrated into parallel configurations and scheduler-related functions like `findResource`.

Note If you do not have a shared file system between client and cluster machines, or if you cannot submit jobs directly to the scheduler from the client machine, any use of third-party schedulers for parallel jobs (including `pmode`, `matlabpool`, and `parfor`) requires that you use the generic scheduler interface.

findResource Now Sets Properties According to Configuration

Compatibility Considerations: Yes

The `findResource` function now sets the properties on the object it creates according to the configuration identified in the function call.

Compatibility Considerations

In past releases, `findResource` could use a configuration to identify a scheduler, but did not apply the configuration settings to the scheduler object properties. If your code uses separate statements to find an object then set properties, this still works, but is not necessary any more.

parfor Syntax Has Single Usage

Compatibility Considerations: Yes

The `parfor` statement is now recognized only for parallel for-loops, not for loops over a distributed range in parallel jobs.

Compatibility Considerations

In R2007b, the pre-existing form of `parfor` was replaced by `for i = (drange)`, but both forms of syntax were recognized in that release. Now `parfor` has only one context, so `parfor` statements used in parallel jobs in code for versions prior to R2007a must be modified to use `for (drange)`.

Limitations

P-Code Scripts

You can call P-code script files from within a `parfor`-loop, but P-code script cannot contain a `parfor`-loop.

sim Inside parfor-Loops

Running simulations in a `parfor`-loop with the `sim` command at the top level of the loop is not allowed in this release. A `sim` command visible in a `parfor`-loop generates an error, although you can call `sim` inside a function that is called from the loop. Be sure that the various labs running simulations do not have the same working directory, as interference can occur with the simulation data.

dfeval Now Destroys Its Job When Finished

Compatibility Considerations: Yes

When finished performing its distributed evaluation, the `dfeval` function now destroys the job it created.

Compatibility Considerations

If you have any scripts that rely on a job and its data still existing after the completion of `dfeval`, or that destroy the job after `dfeval`, these scripts will no longer work.

R2007b

Version: 3.2
New Features: Yes
Bug Fixes: Yes

New Parallel for-Loops (parfor-Loops)

Compatibility Considerations: Yes

New parallel for-loop (parfor-loop) functionality automatically executes a loop body in parallel on dynamically allocated cluster resources, allowing interleaved serial and parallel code. For details of new parfor functionality, see Parallel for-Loops (parfor) in the Distributed Computing Toolbox™ documentation.

Limitations

P-Code Scripts

You can call P-code script files from within a parfor-loop, but P-code script cannot contain a parfor-loop.

Compatibility Considerations

In past releases, parfor was a different function. The new parfor uses parentheses in defining its range to distinguish it from the old parfor.

New parfor:

```
parfor (ii = 1:N); <body of code>; end;
```

Old parfor:

```
parfor ii = 1:N; <body of code>; end;
```

For this release, the old form of parfor without parentheses is still supported, although it generates a warning. You can read more about the new form of this existing functionality in Using a for-Loop Over a Distributed Range (for-drange). You should update your existing parfor code to use the new form of for-loops over a distributed range (for-drange), thus,

```
for ii = drange(1:N); <body of code>; end;
```

Configurations Manager and Dialogs

Compatibility Considerations: Yes

This release introduces a new graphical user interface for creating and modifying user configurations, and for designating the default configuration used by some toolbox functions. For details about the configurations manager, see Cluster Profiles in the Distributed Computing Toolbox documentation.

Compatibility Considerations

This new feature has no impact on how configurations are used in a program, only on how configurations are created and shared among users. In previous versions of the product, you modified your configurations by editing the file *matlabroot/toolbox/distcomp/user/distcompUserConfig.m*. Now the configuration data is stored as part of your MATLAB software preferences.

The new configurations manager cannot directly import old-style configurations that were defined in the *distcompUserConfig.m* file. However, a utility called `importDistcompUserConfig`, available on the MATLAB Central Web site, allows you to convert and import your existing configurations into the new configurations manager. Visit <http://www.mathworks.com/matlabcentral> and search for `importDistcompUserConfig`.

Default Configuration

This version of the toolbox enables you to select a user configuration to use as the default. Thus, commands such as `pmode` and `matlabpool` will use the default configuration without your having to specify it each time you run the command. You can set the default configuration using the configurations graphical interface, or programmatically with the `defaultParallelConfig` function.

Parallel Profiler

A new parallel profiler graphical user interface generates reports on lab computation and communication times during execution of parallel jobs. For details about this new feature, see [Profiling Parallel Code](#).

MDCE Script for Red Hat Removed

Compatibility Considerations: Yes

The MDCE script `rh_mdce`, specific to Red Hat Linux®, has been removed from `matlabroot/toolbox/distcomp/util/bin`.

Compatibility Considerations

If you make use of this script, you must replace it with its more generic equivalent,

`matlabroot/toolbox/distcomp/bin/mdce`.

R2007a

Version: 3.1
New Features: Yes
Bug Fixes: Yes

Local Scheduler and Workers

A local scheduler allows you to schedule jobs and run up to four workers or labs on a single MATLAB client machine without requiring engine licenses. These workers/labs can run distributed jobs or parallel jobs, including pmode sessions, for all products for which the MATLAB client is licensed. This local scheduler and its workers do not require a job manager or third-party scheduler.

New pmode Interface

Compatibility Considerations: Yes

The interactive parallel mode (pmode) has a new interface. The pmode command input and displays of the lab outputs are provided in a user interface that you can separate from the MATLAB client Command Window.

Compatibility Considerations

In previous versions of Distributed Computing Toolbox, the pmode interface used the MATLAB Command Window, with the pmode input using a different prompt. The output from the labs was intermingled with the MATLAB client output.

New Default Scheduler for pmode

Compatibility Considerations: Yes

If you start pmode without specifying a configuration,

```
pmode start
```

pmode automatically starts a parallel job using the local scheduler with labs running on the client machine. For more information about running pmode, see Interactive Parallel Computation with pmode in the Distributed Computing Toolbox documentation.

Compatibility Considerations

In the previous version of the toolbox, when pmode was started without specifying a configuration, it searched the network for the first available job manager to use as a scheduler.

Vectorized Task Creation

Compatibility Considerations: Yes

The `createTask` function can now create a vector of tasks in a single call when you provide a cell array of cell arrays for input arguments. For full details, see the `createTask` reference page.

Compatibility Considerations

In previous versions of the distributed computing products, if your task function had an input argument that was a cell array of cell arrays, your code will need to be modified to run the same way in this release.

For example, your old code may have been written as follows so that the function `myfun` gets four cell array input arguments:

```
createTask(j, @myfun, 1, {{C1} {C2} {C3} {C4}})
```

In this new version, the same code will produce four tasks. To get the old functionality, you must wrap the four cell arrays in another cell array, so that `createTask` knows to create only one task.

```
createTask(j, @myfun, 1, { {{C1} {C2} {C3} {C4}} })
```

Additional Submit and Decode Scripts

There are several submit and decode functions provided with the toolbox for your use with the generic scheduler interface. These files are in the directory

`matlabroot/toolbox/distcomp/examples/integration`

This version of the toolbox includes new subdirectories for Platform LSF and PBS, to support network configurations in which the client and worker computers do not share a file system. For more information, see Supplied Submit and Decode Functions in the Distributed Computing Toolbox documentation.

Jobs Property of Job Manager Sorts Jobs by ID

Compatibility Considerations: Yes

The `Jobs` property of a job manager object now contains the jobs in the order in which they were created, as indicated by the `ID` property of each job. Similarly, the `findJob` function returns jobs sequenced by their `ID`, unless otherwise specified. This change makes job manager behavior consistent with the behavior of third-party schedulers.

Compatibility Considerations

In previous versions of the distributed computing products, when using a job manager, jobs were arranged in the `Jobs` property or by `findJob` according to the status of the job.

New Object Display Format

Compatibility Considerations: Yes

When you create distributed computing objects (scheduler, job, or task) without a semicolon at the end of the command, the object information is displayed in a new format. This new format is also shown when you use the `display` function to view an object or simply type the object name at the command line.

Compatibility Considerations

With this enhancement, the output format shown when creating an object has changed.

Enhanced MATLAB Functions

Several MATLAB functions have been enhanced to work on distributed arrays:

- `cat`
- `find`
- `horzcat`
- `subsindex`
- `vertcat`

For a complete list of MATLAB functions that are enhanced to work on distributed arrays, see [Using MATLAB Functions on Codistributed Arrays](#) in the Distributed Computing Toolbox documentation.

darray Function Replaces distributor Function

Compatibility Considerations: Yes

The function `darray` now defines how an array is distributed among the labs in a parallel job.

Compatibility Considerations

In the previous version of the toolbox, the `distributor` function was used to define how an array was distributed. In many cases, you can replace a call to `distributor` with a call to `darray`. For example, if you used `distributor` without arguments as an input to an array constructor,

```
rand(m, n, distributor());
```

you can update the code to read,

```
rand(m, n, darray());
```

rand Seeding Unique for Each Task or Lab

Compatibility Considerations: Yes

The random generator seed is now initialized based on the task ID for distributed jobs, or the `labindex` for parallel jobs (including `pmode`). This ensures that the set of random numbers generated for each task or lab within a job is unique, even when you have more than 82 tasks or labs.

Compatibility Considerations

In the previous version of the distributed computing products, the `rand` function could by default generate the same set of numbers for some tasks or labs when these exceeded 82 for a job.

Single-Threaded Computations on Workers

Despite the ability in MATLAB software to perform multithreaded computations on multiple-CPU machines, the workers and labs running distributed and parallel jobs perform only single-threaded computations, so that multiprocessor cluster machines can better accommodate multiple workers or labs.

R2006b

Version: 3.0
New Features: Yes
Bug Fixes: Yes

Support for Windows Compute Cluster Server (CCS)

Distributed Computing Toolbox software and MATLAB Distributed Computing Engine™ software now let you program jobs and run them on a Windows Compute Cluster Server. For information about programming in the toolbox to use Windows Compute Cluster Server (CCS) as your scheduler, see the `findResource` reference page.

Windows 64 Support

The distributed computing products now support Windows 64 (Win64) for both MATLAB client and MATLAB worker machines.

Parallel Job Enhancements

Parallel Jobs Support Any Scheduler

Support for parallel jobs now extends to any type of scheduler. In previous releases, only the MathWorks® job manager and mpiexec scheduler object supported parallel jobs. You can now run parallel jobs on clusters scheduled by a job manager, Windows Compute Cluster Server (CCS), Platform LSF, mpiexec, or using the generic scheduler interface. For programming information, see Program Communicating Jobs.

New labSendReceive Function

The labSendReceive function is introduced in this release. This function performs the same things as both labSend and labReceive, but greatly reduces the risk of deadlock, because the send and receive happen simultaneously rather than by separate statements. For more information, see the labSendReceive reference page.

Improved Error Detection

This release offers improved error detection for miscommunication between labs running parallel jobs. Most notable among the improvements are error detection of mismatched labSend and labReceive statements.

Distributed Arrays

Distributed arrays are partitioned into segments, with each segment residing in the workspace of a different lab, so that each lab has its own array segment to work with. Reducing the size of the array that each lab has to store and process means a more efficient use of memory and faster processing, especially for large data sets. For more information, see [Working with Codistributed Arrays](#).

There are many new and enhanced MATLAB functions to work with distributed arrays in parallel jobs. For a listing of these functions and their reference pages, see [Job Management](#).

parfor: Parallel for-Loops

Parallel for-loops let you run a for-loop across your labs simultaneously. For more information, see [Using a for-Loop Over a Distributed Range \(for-drange\)](#) or the [parfor](#) reference page.

Interactive Parallel Mode (pmode)

The interactive parallel mode (pmode) lets you work interactively with a parallel job running simultaneously on a number of labs. Commands you type at the pmode command line are executed on all labs at the same time. Each lab executes the commands in its own workspace on its own local variables or segments of distributed arrays. For more information, see [Run Parallel Jobs Interactively Using pmode](#).

Moved MDCE Control Scripts

Compatibility Considerations: Yes

To provide greater consistency across all platforms, the MDCE control scripts for Windows have moved and those for UNIX and Macintosh have new names.

Compatibility Considerations

Windows Utilities Moved

In previous versions of the distributed computing products, the MDCE utilities for Windows computers were located in

```
matlabroot\toolbox\distcomp\bin\win32
```

The utilities are now located in

```
matlabroot\toolbox\distcomp\bin
```

The files that have moved are

```
nodestatus  
mdce  
startjobmanager  
stopjobmanager  
startworker  
stopworker  
mdce_def.bat
```

UNIX and Macintosh Utilities Renamed

In previous versions of the distributed computing products, the MDCE utilities for UNIX and Macintosh computers were called by

```
nodestatus.sh  
startjobmanager.sh  
stopjobmanager.sh  
startworker.sh  
stopworker.sh
```

You can now call these with the following commands:

```
nodestatus  
startjobmanager  
stopjobmanager  
startworker  
stopworker
```

Note For UNIX and Macintosh, `mdce` and `mdce_def.sh` have not been moved or renamed.

rand Seeding Unique for Each Task or Lab

Compatibility Considerations: Yes

The random generator seed is now initialized based on the task ID for distributed jobs, or the `labindex` for parallel jobs (including `pmode`). This ensures that the random numbers generated for each task or lab are unique within a job.

Compatibility Considerations

In previous versions of the distributed computing products, the `rand` function would by default generate the same set of numbers on each worker.

Task ID Property Now Same as labindex

Compatibility Considerations: Yes

Although you create only one task for a parallel job, the system copies this task for each worker that runs the job. For example, if a parallel job runs on four workers (labs), the `Tasks` property of the job contains four task objects. The first task in the job's `Tasks` property corresponds to the task run by the lab whose `labindex` is 1, and so on, so that the `ID` property for the task object and `labindex` for the lab that ran that task have the same value. Therefore, the sequence of results returned by the `getAllOutputArguments` function corresponds to the value of `labindex` and to the order of tasks in the job's `Tasks` property.

Compatibility Considerations

In past releases, there was no correlation between `labindex` and the task `ID` property.