


Parallel Computing Toolbox™ Release Notes



MATLAB®

How to Contact MathWorks



Latest news: www.mathworks.com
Sales and services: www.mathworks.com/sales_and_services
User community: www.mathworks.com/matlabcentral
Technical support: www.mathworks.com/support/contact_us



Phone: 508-647-7000



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

Parallel Computing Toolbox™ Release Notes

© COPYRIGHT 2006–2018 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.

Parfeval callbacks: New afterAll and afterEach methods for parallel futures	1-2
Slurm support: Slurm is a fully supported scheduler	1-2
Support for NVIDIA Volta: Update to CUDA 9, support for Volta class GPUs	1-2
Improved file mirroring: Performance of file mirroring for generic scheduler integration	1-2
Parfor performance improvements: More efficient broadcast variables in parfor for non-local clusters	1-2
GPU Array Support: Use enhanced gpuArray functions	1-3
Distributed Array Support: Use enhanced distributed array functions	1-3
Parameter Sweep Example: Use a DataQueue to monitor results during computations on a parallel pool	1-3
Discontinued Support for GPU Devices of Compute Capability less than 3.0	1-3
Backwards Compatibility for MATLAB Job Scheduler: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use previous releases of Parallel Computing Toolbox	1-4
Upgrade Parallel Computing Products Together	1-4
Functionality Being Removed or Changed	1-5

Improved Parallel Language Performance: Execute parallel language constructs with reduced overhead	2-2
Tall Array Support: Use tall arrays with Windows client access to Linux Spark clusters	2-2
Improved Parallel Pool Robustness: Run pools without Message Passing Interface (MPI) by default, making pools resilient to workers crashing	2-2
Improved MATLAB Integration with Third-Party Schedulers: Use the Generic Profile Wizard for easier installation and setup of MATLAB Distributed Computing Server	2-2
Cloud Storage: Work with data in Microsoft Azure Blob Storage	2-2
Copy Client Environment to Workers on Any Cluster: Specify which environment variables on your client machine your workers should automatically inherit	2-3
Client Path Sharing: Add user-added-entries on the client's path to the workers' paths	2-3
GPU Array Support: Use enhanced gpuArray functions	2-3
Distributed Array Support: Use enhanced distributed array functions	2-3
Enhanced Support for Microsoft Windows HPC Pack	2-4
Discontinued Support for GPU Devices of Compute Capability less than 3.0	2-4
Backwards Compatibility for MATLAB Job Scheduler: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use previous releases of Parallel Computing Toolbox	2-4
Upgrade Parallel Computing Products Together	2-5

Access to Intermediate Results and Updates in Parallel Computations: Poll for messages or data from different workers during parallel workflows	3-2
Tall Array Support: Use parallel execution environments for tall timetables and enhanced tall array functions	3-2
More Responsive Job Monitor: Automatic updates for new, submitted, or deleted jobs or tasks	3-3
Re-create Jobs: Easily rerun all failed or cancelled tasks	3-3
GPU Array Support: Use enhanced gpuArray functions	3-3
Distributed Array Support: Use enhanced distributed array functions	3-4
Simplified Integration for Third-Party Cluster Schedulers: Updates to generic scheduler integration allow folder-based configuration and eliminate the need to specify function handles	3-4
MATLAB String Support: Functions in Parallel Computing Toolbox accept MATLAB strings as input	3-4
Support for Spark 2.x enabled Hadoop clusters	3-4
Upgraded CUDA Toolkit version	3-5
Discontinued Support for GPU Devices of Compute Capability less than 3.0	3-5
Backwards Compatibility for MATLAB Job Scheduler: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use previous releases of Parallel Computing Toolbox	3-5

Upgrade Parallel Computing Products Together	3-5
Properties of generic cluster objects have changed	3-6
Functionality Being Removed or Changed	3-7

R2016b

Parallel Support for Tall Arrays: Process big data with tall arrays in parallel on your desktop, MATLAB Distributed Computing Server, and Spark clusters	4-2
Support for GPU Arrays: Use enhanced gpuArray functions, including new sparse iterative solver bicg	4-2
Parallel Menu Enhancement: Use the new menu items in the Parallel Menu to configure and manage cloud based resources	4-2
New Data Types in Distributed Arrays: Use enhanced functions for creating distributed arrays of: datetime; duration; calendarDuration; string; categorical; and table	4-2
Loading Distributed Arrays: Load distributed arrays in parallel using datastore	4-4
Cluster Profile Validation: Choose which validation stages run and the number of MATLAB workers to use	4-4
Backwards Compatibility: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use the previous release of Parallel Computing Toolbox	4-4
datetime Support for Timestamps: Use built-in datetime objects in MATLAB to access timestamp information for jobs and tasks	4-5

Data Transfer Measurement: Use ticBytes and tocBytes to measure the data transfer between MATLAB workers in a parallel pool	4-5
Multithreaded Workers: Use multiple computational threads on your MATLAB workers	4-5
Support for Distributed Arrays: Use enhanced distributed array functions, including sparse input to iterative solvers (gmres and lsqr)	4-6
Increased Data Transfer Limits: Send messages larger than 2 GB using labSend, labSendReceive and labBroadcast	4-6
Upgrade Parallel Computing Products Together	4-6
Functionality Being Removed or Changed	4-7

R2016a

GPU Support for Sparse Matrices: Use enhanced gpuArray functions for sparse matrices on GPUs	5-2
Support for Distributed Arrays: Use enhanced distributed array functions including sparse input to direct (mldivide) and iterative solvers (cgs and pcg)	5-2
GPU-Accelerated Deep Learning: Use Neural Network Toolbox to train deep convolutional neural networks with GPU-enabled acceleration for image classification tasks	5-3
GPU-enabled MATLAB Functions: Accelerate applications using GPU-enabled MATLAB functions for linear equations, descriptive statistics and set operations	5-3
Upgraded CUDA Toolkit version	5-3

Parallel-Enabled Gradient Estimation: Accelerate more nonlinear solvers in the Optimization Toolbox with parallel finite difference estimation of gradients and Jacobians	5-4
Hadoop Kerberos Support: Improved support for Hadoop in a Kerberos authenticated environment	5-4
Transfer unlimited data between client and workers, and attached files up to 4GB in total, in any job using a MATLAB Job Scheduler cluster	5-4
matlabpool function removed	5-4
Upgrade parallel computing products together	5-5

R2015b

Discontinued support for parallel computing products on 32-bit Windows operating systems	6-2
More than 90 GPU-enabled functions in Statistics and Machine Learning Toolbox, including probability distribution, descriptive statistics, and hypothesis testing	6-2
Additional GPU-enabled MATLAB functions, including support for sparse matrices	6-2
Additional GPU-enabled MATLAB functions	6-2
Sparse arrays with GPU-enabled functions	6-3
mexcuda function for easier compilation of MEX-files containing CUDA code	6-3
Upgraded CUDA Toolkit version	6-4
Scheduler integration scripts for SLURM	6-4
parallel.pool.Constant function to create constant data on parallel pool workers, accessible within parallel language constructs such as parfor and parfeval	6-4

Improved performance of mapreduce on Hadoop 2 clusters	6-4
Enhanced and additional MATLAB functions for distributed arrays	6-5
Warnings property for tasks	6-5
More consistent transparency enforcement	6-5
Upgrade parallel computing products together	6-6

R2015a

Support for mapreduce function on any cluster that supports parallel pools	7-2
Sparse arrays with GPU-enabled functions	7-2
Additional GPU-enabled MATLAB functions	7-2
pagefun support for mrdivide and inv functions on GPUs	7-3
Enhancements to GPU-enabled linear algebra functions	7-3
Parallel data reads from a datastore with MATLAB partition function	7-3
Using DNS for cluster discovery	7-3
MS-MPI support for local and MJS clusters	7-4
Ports and sockets in mdce_def file	7-4
Improved profiler accuracy for GPU code	7-5
Upgraded CUDA Toolkit version	7-5

Discontinued support for GPU devices on 32-bit Windows computers	7-5
Discontinued support for parallel computing products on 32-bit Windows computers	7-5
matlabpool function removed	7-5

R2014b

Parallelization of mapreduce on local workers	8-2
Additional GPU-enabled MATLAB functions, including accumarray, histc, cummax, and cummin	8-2
pagefun support for mldivide on GPUs	8-3
Additional MATLAB functions for distributed arrays, including fft2, fftn, ifft2, ifftn, cummax, cummin, and diff	8-3
Data Analysis on Hadoop clusters using mapreduce	8-4
Discover Clusters Supports Microsoft Windows HPC Server for Multiple Releases	8-4
Upgraded CUDA Toolkit Version	8-4
Discontinued Support for GPU Devices of Compute Capability 1.3	8-4
Discontinued Support for GPU Devices on 32-Bit Windows Computers	8-5

Number of local workers no longer limited to 12	9-2
Additional GPU-enabled MATLAB functions: interp3, interpn, besselj, bessely	9-2
Additional GPU enabled Image Processing Toolbox functions: bwdist, imreconstruct, iradon, radon	9-2
Enhancements to GPU-enabled MATLAB functions: filter (IIR filters); pagefun (additional functions supported); interp1, interp2, conv2, reshape (performance improvements)	9-2
Duplication of an existing job, containing some or all of its tasks	9-3
More MATLAB functions enhanced for distributed arrays ...	9-3
GPU MEX Support for Updated MEX	9-4
Old Programming Interface Removed	9-4
matlabpool Function Being Removed	9-5
Removed Support for parallel.cluster.Mpiexec	9-5

parpool: New command-line interface (replaces matlabpool), desktop indicator, and preferences for easier interaction with a parallel pool of MATLAB workers	10-2
Parallel Pool	10-2
New Desktop Pool Indicator	10-3
New Parallel Preferences	10-4

Automatic start of a parallel pool when executing code that uses parfor or spmd	10-4
Option to start a parallel pool without using MPI	10-5
More GPU-enabled MATLAB functions (e.g., interp2, pagefun) and Image Processing Toolbox functions (e.g., bwmorph, edge, imresize, and medfilt2)	10-5
More MATLAB functions enabled for distributed arrays: permute, ipermute, and sortrows	10-6
Enhancements to MATLAB functions enabled for GPUs, including ones, zeros	10-7
gputimeit Function to Time GPU Computations	10-7
New GPU Random Number Generator NormalTransform Option: Box-Muller	10-8
Upgraded MPICH2 Version	10-8
Discontinued Support for GPU Devices of Compute Capability 1.3	10-9
Discontinued Support for parallel.cluster.Mpiexec	10-9

R2013a

GPU-enabled functions in Image Processing Toolbox and Phased Array System Toolbox	11-2
More MATLAB functions enabled for use with GPUs, including interp1 and ismember	11-2
Enhancements to MATLAB functions enabled for GPUs, including arrayfun, svd, and mldivide ()	11-2

Ability to launch CUDA code and manipulate data contained in GPU arrays from MEX-functions	11-3
Automatic detection and transfer of files required for execution in both batch and interactive workflows	11-3
More MATLAB functions enabled for distributed arrays	11-4

R2012b

More MATLAB functions enabled for GPUs, including convn, cov, and normest	12-2
gpuArray Support	12-2
MATLAB Code on the GPU	12-2
GPU-enabled functions in Neural Network Toolbox, Phased Array System Toolbox, and Signal Processing Toolbox	12-2
Performance improvements to GPU-enabled MATLAB functions and random number generation	12-3
Automatic detection and selection of specific GPUs on a cluster node when multiple GPUs are available on the node	12-3
More MATLAB functions enabled for distributed arrays, including sparse constructor, bsxfun, and repmat	12-3
Detection of MATLAB Distributed Computing Server clusters that are available for connection from user desktops through Profile Manager	12-4
gpuArray Class Name	12-4
Diary Output Now Available During Running Task	12-5

New Programming Interface	13-2
General Concepts and Phrases	13-2
Objects	13-2
Functions and Methods	13-4
Properties	13-4
Getting Help	13-5
Enhanced Example Scripts	13-8
Cluster Profiles	13-8
New Cluster Profile Manager	13-8
Programming with Profiles	13-9
Profiles in Compiled Applications	13-10
Enhanced GPU Support	13-10
GPUArray Support	13-10
Reset or Deselect GPU Device	13-11
Asynchronous GPU Calculations and Wait	13-12
Verify GPUArray or CUDAKernel Exists on the Device	13-12
MATLAB Code on the GPU	13-13
Set CUDA Kernel Constant Memory	13-14
Latest NVIDIA CUDA Device Driver	13-14
Enhanced Distributed Array Support	13-14
Newly Supported Functions	13-14
Random Number Generation on Workers	13-14

New Job Monitor	14-2
Run Scripts as Batch Jobs from the Current Folder	
Browser	14-2
Number of Local Workers Increased to Twelve	14-3

Enhanced GPU Support	14-3
Latest NVIDIA CUDA Device Driver	14-3
Deployment of GPU Applications	14-3
Random Number Generation	14-3
GPUArray Support	14-4
MATLAB Code on the GPU	14-5
Enhanced Distributed Array Support	14-5
Newly Supported Functions	14-5
Conversion of Error and Warning Message Identifiers	14-6
Task Error Properties Updated	14-6

R2011a

Deployment of Local Workers	15-2
New Desktop Indicator for MATLAB Pool Status	15-2
Enhanced GPU Support	15-2
Static Methods to Create GPUArray	15-2
GPUArray Support	15-2
GPUArray Indexing	15-3
MATLAB Code on the GPU	15-3
NVIDIA CUDA Driver 3.2 Support	15-3
Distributed Array Support	15-4
Newly Supported Functions	15-4
Enhanced mtimes Support	15-4
Enhanced parfor Support	15-4
Nested for-Loops Inside parfor	15-4
Enhanced Support for Microsoft Windows HPC Server	15-5
Support for 32-Bit Clients	15-5
Search for Cluster Head Nodes Using Active Directory	15-5
Enhanced Admin Center Support	15-5

New Remote Cluster Access Object	15-5
---	-------------

R2010b

GPU Computing	16-2
Job Manager Security and Secure Communications	16-2
Generic Scheduler Interface Enhancements	16-2
Decode Functions Provided with Product	16-2
Enhanced Example Scripts	16-3
batch Now Able to Run Functions	16-4
batch and matlabpool Accept Scheduler Object	16-4
Enhanced Functions for Distributed Arrays	16-4
qr Supports Distributed Arrays	16-4
mldivide Enhancements	16-4
chol Supports 'lower' option	16-4
eig and svd Return Distributed Array	16-5
transpose and ctranspose Support 2dbc	16-5
Inf and NaN Support Multiple Formats	16-5
Support for Microsoft Windows HPC Server 2008 R2	16-5
User Permissions for MDCEUSER on Microsoft Windows ...	16-6

R2018a

Version: 6.12

New Features

Bug Fixes

Compatibility Considerations

Parfeval callbacks: New afterAll and afterEach methods for parallel futures

You can use the new methods `afterAll` and `afterEach` to specify functions to automatically execute when parallel futures complete. You create parallel futures using the `parfeval` function. A useful application for `afterEach` and `afterAll` is to update user interfaces such as plots and apps during parallel computations using `parfeval`. For an example, see “Update a User Interface Asynchronously Using `afterEach` and `afterAll`”.

Slurm support: Slurm is a fully supported scheduler

Slurm is now a fully supported scheduler. You can create Slurm profiles directly in the Profile Manager without needing templates and scripts from File Exchange. For more information, see “Configure for Slurm, PBS Pro, Platform LSF, TORQUE” (MATLAB Distributed Computing Server).

Support for NVIDIA Volta: Update to CUDA 9, support for Volta class GPUs

The parallel computing products are now using CUDA® Toolkit version 9.0. To compile CUDA code for `CUDAkernel` or CUDA MEX-files, you must use toolkit version 9.0.

This update improves support for Volta class GPUs.

Improved file mirroring: Performance of file mirroring for generic scheduler integration

File mirroring for the generic scheduler integration has been improved through the use of zip archives.

Parfor performance improvements: More efficient broadcast variables in parfor for non-local clusters

If you use broadcast variables in `parfor` loops with non-local clusters, now performance is improved. Broadcast variables in `parfor` loops are sent only once to the cluster, instead of once per worker.

GPU Array Support: Use enhanced gpuArray functions

You can now use the following enhanced gpuArray functions:

- `rescale`
- new syntaxes of `sort` and `sortrow`

For more information, see “Run Built-In Functions on a GPU”.

Distributed Array Support: Use enhanced distributed array functions

You can now use the following enhanced distributed array functions:

- `rescale`
- new syntaxes of `sort` and `sortrow`
- `bicgstabl`

For more information, see “Using MATLAB Functions on Distributed Arrays”.

Parameter Sweep Example: Use a DataQueue to monitor results during computations on a parallel pool

The new example *Plot during Parameter Sweep with parfor* shows how to perform a parameter sweep in parallel and plot progress during parallel computations. For details, see “Plot during Parameter Sweep with parfor”.

Discontinued Support for GPU Devices of Compute Capability less than 3.0

In R2018a, support for GPU devices of compute capability less than 3.0 is removed. A minimum compute capability of 3.0 is required for GPU computing in MATLAB®. For more information on using GPUs, see “GPU Computing in MATLAB”.

Compatibility Considerations

In R2018a, any use of `gpuDevice` to select a GPU with compute capability less than 3.0 generates an error. Support is removed for devices with compute capability less than 3.0.

Backwards Compatibility for MATLAB Job Scheduler: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use previous releases of Parallel Computing Toolbox

You can upgrade your MATLAB Job Scheduler (MJS) clusters to the latest release and still connect to it using previous releases (back to R2016a) of Parallel Computing Toolbox on your MATLAB desktop client. For the supported older releases, you must maintain an installation of the same release of MATLAB Distributed Computing Server™ for each release of MATLAB you want to use. The latest MATLAB Job Scheduler routes your jobs accordingly. You can configure MATLAB Job Scheduler with the location of these installations in the `mdce_def` file. For more information, see “Install Products and Choose Cluster Configuration” (MATLAB Distributed Computing Server).

Upgrade Parallel Computing Products Together

As with every new release, if you upgrade Parallel Computing Toolbox you must upgrade MATLAB Distributed Computing Server at the same time so that they interact properly with each other. Note that you do not need to do this if you are taking advantage of MATLAB Job Scheduler’s (MJS) backwards compatibility.

The R2018a 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. This requirement applies only if you are not taking advantage of MATLAB Job Scheduler (MJS) backwards compatibility.

Jobs created in one version of Parallel Computing Toolbox software do 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`.

Functionality Being Removed or Changed

Functionality	Result	Use Instead	Compatibility Considerations
Support for GPU devices of compute capability less than 3.0 is removed.	Errors	Use GPUs with a minimum compute capability of 3.0.	Support is removed for GPU devices of compute capability 2.x. Use GPUs with a minimum compute capability 3.0. Use the function <code>gpuDevice</code> to query the compute capabilities of your GPUs. For more information, see "GPU Computing in MATLAB".
<code>parallel.gpu.rng</code> is renamed to <code>gpurng</code>	Still runs	<code>gpurng</code>	Replace all instances of <code>parallel.gpu.rng</code> with <code>gpurng</code>

R2017b

Version: 6.11

New Features

Bug Fixes

Compatibility Considerations

Improved Parallel Language Performance: Execute parallel language constructs with reduced overhead

All parallel language constructs, including `parfor`, run with reduced overhead, particularly constructs with short duration. The scheduling of `parfor`-loop intervals has been changed to utilize the cluster hardware more efficiently.

Tall Array Support: Use tall arrays with Windows client access to Linux Spark 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.

Improved Parallel Pool Robustness: Run pools without Message Passing Interface (MPI) by default, making pools resilient to workers crashing

You can now run parallel pools without MPI by default, improving the resilience of your parallel pool to workers crashing.

Improved MATLAB Integration with Third-Party Schedulers: Use the Generic Profile Wizard for easier installation and setup of MATLAB Distributed Computing Server

You can now use the new workflow to automatically create generic profiles using the support packages for third-party schedulers. For more details, see [Distribute a Generic Cluster Profile](#) and [Integration Scripts](#).

Cloud Storage: Work with data in Microsoft Azure Blob Storage

You can now use Microsoft® Azure® Blob Storage to access data in the cloud using `datastore`. For more information on this topic, see [Distribute a Generic Cluster Profile](#) and [Integration Scripts](#).

Copy Client Environment to Workers on Any Cluster: Specify which environment variables on your client machine your workers should automatically inherit

You can now mirror any of your client environment variables to the workers on a cluster. For example, you can include any environment variables required by third-party software. Or you can include cloud service provider credentials in your local environment to give worker processes access to your data stored in the cloud.

You can now set `EnvironmentVariables` in `parpool`, `batch`, `createJob` and in the Cluster Profile Manager. `EnvironmentVariables` is also a common property of the `parallel.Job` class.

Client Path Sharing: Add user-added-entries on the client's path to the workers' paths

You can now automatically add user-added-entries on the client's path to the workers' paths for `batch` and independent jobs. This functionality can be enabled or disabled by specifying the `AutoAddClientPath` option to the `parpool`, `batch` and `createJob` commands. `AutoAddClientPath` is also a common property of the `parallel.Job` class.

GPU Array Support: Use enhanced `gpuArray` functions

You can now use the following enhanced `gpuArray` functions:

- `bounds` for computing `max` and `min` simultaneously
- Sparse iterative solvers `cgs` and `lsqr`
- `spdiags`

For more information on this topic, see [Run Built-In Functions on a GPU](#).

Distributed Array Support: Use enhanced distributed array functions

You can now use the following enhanced distributed array functions:

- Sparse iterative solvers `minres`, `symmlq` and `bicgstab`

- `bounds` for computing `max` and `min` simultaneously
- `eigs`
- `spdiags`

For more information, see [Using MATLAB Functions on Distributed Arrays](#).

Enhanced Support for Microsoft Windows HPC Pack

The parallel computing products now support Microsoft Windows® HPC Pack 2016. For details, see [Configure for HPC Pack](#).

Discontinued Support for GPU Devices of Compute Capability less than 3.0

In a future release, support for GPU devices of compute capability less than 3.0 will be removed. At that time, a minimum compute capability of 3.0 will be required.

Compatibility Considerations

In R2017b, any use of `gpuDevice` to select a GPU with compute capability less than 3.0, generates a warning. The device is still supported in this release, but in a future release support will be completely removed for devices with compute capability less than 3.0.

Backwards Compatibility for MATLAB Job Scheduler: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use previous releases of Parallel Computing Toolbox

You can upgrade your MATLAB Job Scheduler (MJS) clusters to R2017b and still connect to it using the R2017a, R2016b, and R2016a releases of Parallel Computing Toolbox on your MATLAB desktop client. For releases prior to R2016b, you must maintain an installation of the same release of MATLAB Distributed Computing Server for each release of MATLAB you want to use. The latest MATLAB Job Scheduler routes your jobs accordingly. You can configure MATLAB Job Scheduler with the location of these installations in the `mdce_def` file. For more information, see [Install Products and Choose Cluster Configuration](#).

Upgrade Parallel Computing Products Together

The R2017b 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. This requirement applies only if you are not taking advantage of MATLAB Job Scheduler (MJS) backwards compatibility.

Jobs created in one version of Parallel Computing Toolbox software do 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`.

Functionality Being Removed or Changed

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 or higher installed to run MATLAB MapReduce.	Migrate MATLAB MapReduce code that runs on Hadoop 1.x to Hadoop 2.x.

For more information, see [Configure a Hadoop Cluster](#).

R2017a

Version: 6.10

New Features

Bug Fixes

Compatibility Considerations

Access to Intermediate Results and Updates in Parallel Computations: Poll for messages or data from different workers during parallel workflows

You can now transfer intermediate results when you carry out `parfor`, `spmd`, or `parfeval` calculations. Use the `send` and `poll` methods together to transfer and retrieve messages or data from different workers, using a pollable `DataQueue`.

Tall Array Support: Use parallel execution environments for tall timetables and enhanced tall array functions

You can now use the following enhanced tall array functions:

- `timetable` supports:

<code>head</code>	<code>join</code>	<code>stack</code>	<code>unique</code>
<code>height</code>	<code>ndims</code>	<code>standardizeMissing</code>	<code>varfun</code>
<code>isempty</code>	<code>numel</code>	<code>table2array</code>	<code>width</code>
<code>innerjoin</code>	<code>size</code>	<code>table2cell</code>	
<code>ismember</code>	<code>sortrows</code>	<code>tail</code>	
<code>ismissing</code>	<code>splitapply</code>	<code>topkrows</code>	

- `array2table`
- `bsxfun`
- `conv`
- `cumsum`, `cumprod`, `cummax`, `cummin`
- `diff`
- `ismember`
- `issorted`, `issortedrows`
- `movmad`, `movmax`, `movmean`, `movmedian`, `movmin`, `movprod`, `movstd`, `movsum`, `movvar`
- `pie` for categoricals
- `repmat`, `repelem`

-
- `sort`, `sortrows`
 - `table2timetable`, `timetable2table`
 - `varfun`
 - Improved indexing, allowing sorted (ascending or descending) indices in the first dimension

For more information, see [Functions That Support Tall Arrays \(A-Z\)](#).

More Responsive Job Monitor: Automatic updates for new, submitted, or deleted jobs or tasks

You can now use the Job Monitor to get an up-to-date view of your cluster. Verify your changes without manually querying the cluster or forcing an update to the Job Monitor user interface. Examine your latest changes and execute quickly without interrupting your workflow. For more details, see [Job Monitor](#).

Re-create Jobs: Easily rerun all failed or cancelled tasks

If your job failed or was cancelled, you can now use the `recreate` method to return a new job object. Call `submit` on the new job object to easily rerun all failed or cancelled tasks.

GPU Array Support: Use enhanced `gpuArray` functions

You can now use the following enhanced `gpuArray` functions:

- `head`, `tail`
- Cubic support for `interp1` and `interp2`
- `movmean`, `movstd`, `movsum`, `movvar`
- `svds`
- `tril` and `triu` for sparse arrays

For more information on this topic, see [Run Built-In Functions on a GPU](#).

Distributed Array Support: Use enhanced distributed array functions

You can now use the following enhanced distributed array functions:

- `issymmetric`, `ishermitian`
- `qmr`
- `svds`
- `tfqmr`
- `write` for storing a snapshot of a distributed array in a format suitable for `datastore`
- `mldivide`, providing improved performance for triangular and diagonal systems

For more information, see [Using MATLAB Functions on Distributed Arrays](#).

Simplified Integration for Third-Party Cluster Schedulers: Updates to generic scheduler integration allow folder-based configuration and eliminate the need to specify function handles

You no longer need to specify function handles in your generic scheduler profile. Instead, specify only one folder name and some `AdditionalProperties`, which are similar to name-value pairs. For more information, see “[Configure for a Generic Scheduler](#)” (MATLAB Distributed Computing Server).

MATLAB String Support: Functions in Parallel Computing Toolbox accept MATLAB strings as input

You can now use MATLAB strings for functions in Parallel Computing Toolbox. For information on MATLAB strings, see [Create String Arrays](#).

Support for Spark 2.x enabled Hadoop clusters

Tall array integration on a Spark enabled Hadoop cluster, running MATLAB Distributed Computing Server, now supports both versions 1.x and 2.x.

Upgraded CUDA Toolkit version

The parallel computing products are now using CUDA Toolkit version 8.0. To compile CUDA code for CUDAKernel or CUDA MEX-files, you must use toolkit version 8.0.

Discontinued Support for GPU Devices of Compute Capability less than 3.0

In a future release, support for GPU devices of compute capability less than 3.0 will be removed. At that time, a minimum compute capability of 3.0 will be required.

Compatibility Considerations

In R2017a, any use of `gpuDevice` to select a GPU with compute capability less than 3, generates a warning. The device is still supported in this release, but in a future release support will be completely removed for devices with compute capability less than 3.

Backwards Compatibility for MATLAB Job Scheduler: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use previous releases of Parallel Computing Toolbox

You can upgrade your MATLAB Job Scheduler (MJS) clusters to R2017a and still use the R2016b and R2016a releases of Parallel Computing Toolbox on your MATLAB desktop client to connect to it. This situation only applies to the R2016a release onward. You must maintain an installation of the same release of MATLAB Distributed Computing Server for each release of MATLAB you want to use. The latest MATLAB Job Scheduler routes your jobs accordingly. You can configure MATLAB Job Scheduler with the location of these installations in the `mdce_def` file. For more information, see “Install Products and Choose Cluster Configuration” (MATLAB Distributed Computing Server).

Upgrade Parallel Computing Products Together

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 do 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`.

Properties of generic cluster objects have changed

The properties of new generic cluster objects have changed.

Compatibility Considerations

From R2017a onwards, you can no longer use the following old properties:

- `CancelJobFcn`
- `CancelTaskFcn`
- `CommunicatingSubmitFcn`
- `DeleteJobFcn`
- `DeleteTaskFcn`
- `GetJobStateFcn`
- `IndependentSubmitFcn`

Instead, set the new `IntegrationScriptsLocation` property to the folder containing your cluster's integration scripts. You must now use integration scripts with the default function name and signature. Specify any additional input arguments to these scripts using the new `AdditionalProperties` property. For more details, see

- “Configure for a Generic Scheduler” (MATLAB Distributed Computing Server)

-
- “Program Communicating Jobs for a Generic Scheduler”
 - “Integration Scripts for Generic Schedulers”

Functionality Being Removed or Changed

Functionality	Result	Use Instead	Compatibility Considerations
Support for running MATLAB MapReduce on Hadoop 1.x clusters will be removed in a future release.	Warns	Use clusters that have Hadoop 2.x or higher installed to run MATLAB MapReduce.	Migrate MATLAB MapReduce code that runs on Hadoop 1.x to Hadoop 2.x.

For more information, see “Configure a Hadoop Cluster” (MATLAB Distributed Computing Server).

R2016b

Version: 6.9

New Features

Bug Fixes

Compatibility Considerations

Parallel Support for Tall Arrays: Process big data with tall arrays in parallel on your desktop, MATLAB Distributed Computing Server, and Spark clusters

Speed up your tall array workflows with Parallel Computing Toolbox. You can use Parallel Computing Toolbox to evaluate tall array expressions in parallel using a parallel pool on your desktop. You can also use Parallel Computing Toolbox to scale up tall-array processing by connecting to a parallel pool running on a MATLAB Distributed Computing Server cluster, or to a Spark enabled Hadoop cluster running MATLAB Distributed Computing Server. For more information, see

- [Big Data Workflow Using Tall Arrays and Datastores](#)
- [Use Tall Arrays on a Parallel Pool](#)
- [Use Tall Arrays on a Spark Enabled Hadoop Cluster](#)

Support for GPU Arrays: Use enhanced gpuArray functions, including new sparse iterative solver bicg

- [New sparse iterative solver bicg](#)
- [Support for tolerance set functions: ismembertol, uniquetol](#)
- [Create sparse arrays using sprandsym](#)

For more information on this topic, see [Run Built-In Functions on a GPU](#).

Parallel Menu Enhancement: Use the new menu items in the Parallel Menu to configure and manage cloud based resources

Open the Cloud Center web application, and view MATLAB Distributed Computing Server license usage. For more information, see [Use Parallel Menu and Cluster Profiles](#).

New Data Types in Distributed Arrays: Use enhanced functions for creating distributed arrays of: datetime; duration; calendarDuration; string; categorical; and table

Distributed calendarDuration Arrays:

calendarDuration	calquarters	cellstr	time
caldays	calweeks	datevec	
calmonths	calyears	iscalendarduration	

Distributed categorical Arrays:

categorical	iscategorical	isundefined	setcats
addcats	iscategory	removecats	
categories	isordinal	renamecats	
countcats	isprotected	reordercats	

Distributed datetime Arrays:

datetime	exceltime	isweekend	string
between	hms	juliandate	timeofday
cellstr	hour	minute	tzoffset
datenum	isbetween	month	week
dateshift	isdatetime	posixtime	year
datevec	isdst	quarter	ymd
day	isnat	second	yyyymmdd

Distributed duration Arrays:

duration	datevec	hours	minutes
cellstr	days	isduration	seconds
datenum	hms	milliseconds	years

Distributed string Arrays:

string	erase	insertBefore	replaceBetween
cellstr	eraseBetween	ismissing	reverse
compose	extractAfter	isstring	startsWith
contains	extractBefore	lower	strip

count	extractBetween	pad	strlength
endsWith	insertAfter	replace	upper

Distributed tables:

table	istable	table2cell
head	standardizeMissing	tail
ismissing	table2array	

For more information, see [Using MATLAB Functions on Distributed Arrays](#).

Loading Distributed Arrays: Load distributed arrays in parallel using datastore

Create distributed arrays more easily using `datastore`, and eliminate the need to create distributed arrays with `codistributed`. For more information, see [Load Distributed Arrays in Parallel Using datastore](#).

Cluster Profile Validation: Choose which validation stages run and the number of MATLAB workers to use

In previous releases, validating your cluster profile ran all validation stages and used a fixed number of workers determined from your profile. You can now choose to run a subset of the validation stages and specify the number of workers to use when validating your profile. For more information on the detailed validation steps in your cluster profile, see [Validate Cluster Profiles](#).

Backwards Compatibility: Upgrade MATLAB Job Scheduler (MJS) clusters, and continue to use the previous release of Parallel Computing Toolbox

You can upgrade your MATLAB Job Scheduler (MJS) clusters to R2016b and still use the R2016a release of Parallel Computing Toolbox on your MATLAB desktop client to connect to it. This situation only applies to the R2016a release onward. You must maintain an installation of the same release of MATLAB Distributed Computing Server for each release of MATLAB you want to use. The latest MATLAB Job Scheduler will route your

jobs accordingly. You can configure MATLAB Job Scheduler with the location of these installations in the `mdce_def` file. For more information, see [Install Products and Choose Cluster Configuration](#).

datetime Support for Timestamps: Use built-in datetime objects in MATLAB to access timestamp information for jobs and tasks

You can now use the following properties for MATLAB jobs:

```
CreateDateTime  
SubmitDateTime  
StartDateTime  
FinishDateTime
```

You can use the following properties for MATLAB tasks:

```
CreateDateTime  
StartDateTime  
FinishDateTime
```

For more information, see [parallel.Job](#) and [parallel.Task](#)

Data Transfer Measurement: Use ticBytes and tocBytes to measure the data transfer between MATLAB workers in a parallel pool

You can now measure how much data needs to be passed around to carry out `parfor`, `spmd` or `parfeval`. Use `ticBytes` and `tocBytes` to optimize your code and pass around less data.

Multithreaded Workers: Use multiple computational threads on your MATLAB workers

MATLAB workers used to run in single-threaded mode. Now you can control the number of computational threads so that workers can run in multithreaded mode and use all the

cores on your cluster. This enables you to increase the number of computational threads, `NumThreads`, on each worker, without increasing the number of workers, `NumWorkers`. If you have more cores available, increase `NumThreads` to take full advantage of the built-in parallelism provided by the multithreaded nature of many of the underlying MATLAB libraries. For more information, see [Create and Modify Cluster Profiles](#).

Support for Distributed Arrays: Use enhanced distributed array functions, including sparse input to iterative solvers (gmres and lsqr)

- Sparse support for `gmres` and `lsqr`
- Support for `isdiag`, `istril`, `istriu`, `isbanded`, `bandwidth`, and `mrdivide` for both sparse and dense arrays
- N-dimensional convolutions with `convn`
- `find` now supports the 2-D block-cyclic (2dbc) distribution scheme

For more information, see [Using MATLAB Functions on Distributed Arrays](#).

Increased Data Transfer Limits: Send messages larger than 2 GB using labSend, labSendReceive and labBroadcast

In previous releases, it was not possible to use `labSend`, `labSendReceive`, or `labBroadcast` to send messages larger than 2GB when the message was not a real dense numeric array (such as a complex array, a sparse array, a cell array or a `struct`). In R2016b, this limitation has been removed.

Upgrade Parallel Computing Products Together

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

Compatibility Considerations

If you are running MATLAB Job Scheduler on your cluster, then you can upgrade MATLAB® Distributed Computing Server without upgrading Parallel Computing Toolbox.

However, you cannot upgrade Parallel Computing Toolbox without upgrading MATLAB Distributed Computing Server.

If you are not running MATLAB Job Scheduler, then you must upgrade Parallel Computing Toolbox and MATLAB Distributed Computing Server together. 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`.

Functionality Being Removed or Changed

Functionality	Result	Use Instead	Compatibility Considerations
Support for running MATLAB MapReduce on Hadoop 1.x clusters will be removed in a future release.	Still runs	Use clusters that have Hadoop 2.x or higher installed to run MATLAB MapReduce.	Migrate MATLAB MapReduce code that runs on Hadoop 1.x to Hadoop 2.x

For more information, see [Configure a Hadoop Cluster](#).

R2016a

Version: 6.8

New Features

Bug Fixes

Compatibility Considerations

GPU Support for Sparse Matrices: Use enhanced gpuArray functions for sparse matrices on GPUs

- Support for 5-argument form of sparse constructor
- Support for sparse inputs to bicgstab and pcg

You can create a sparse gpuArray either by calling sparse with a gpuArray input, or by calling gpuArray with a sparse input.

For more information on this topic, see Sparse Arrays on a GPU.

Support for Distributed Arrays: Use enhanced distributed array functions including sparse input to direct (mldivide) and iterative solvers (cgs and pcg)

The following functions are new in supporting distributed arrays:

```
cgs
conv
conv2
expint
ischar
pcg
superiorfloat
```

For more details, see MATLAB Functions on Distributed and Codistributed Arrays.

In addition, the following features are new in providing enhanced functionality for distributed arrays:

- Sparse input to mldivide (direct system solve)
- Sparse input to cgs and pcg (iterative system solve)
- Single argument syntax for sprand and sprandn

GPU-Accelerated Deep Learning: Use Neural Network Toolbox to train deep convolutional neural networks with GPU-enabled acceleration for image classification tasks

For specific information, see the Neural Network Toolbox™ release notes. To access this functionality, Parallel Computing Toolbox is required.

GPU-enabled MATLAB Functions: Accelerate applications using GPU-enabled MATLAB functions for linear equations, descriptive statistics and set operations

The following functions are new in their support for `gpuArrays`:

```
bicg  
detrend  
discretize  
expint  
pcg  
spconvert  
sprand  
sprandn
```

New support is available for the 'rows' option in the following Set Operations:

```
intersect  
ismember  
setdiff  
setxor  
union  
unique
```

For more details, see [Run Built-In Functions on a GPU](#).

Upgraded CUDA Toolkit version

The parallel computing products are now using CUDA Toolkit version 7.5. To compile CUDA code for `CUDAKernel` or CUDA MEX-files, you must use toolkit version 7.5.

Parallel-Enabled Gradient Estimation: Accelerate more nonlinear solvers in the Optimization Toolbox with parallel finite difference estimation of gradients and Jacobians

Hadoop Kerberos Support: Improved support for Hadoop in a Kerberos authenticated environment

In R2016a, enhanced support for the recommended setup for the Cloudera distribution of Hadoop has been provided.

Transfer unlimited data between client and workers, and attached files up to 4GB in total, in any job using a MATLAB Job Scheduler cluster

In previous releases, the data transfer limit for a MATLAB Job Scheduler cluster was 2GB. In R2016a, this limit has been removed.

matlabpool function removed

`matlabpool` function has been removed. Use `parpool` instead.

Compatibility Considerations

Functionality	What Happens When You Use This Functionality?	Use This Instead	Compatibility Considerations
matlabpool	Errors	parpool	<p>In R2016a, matlabpool generates the following error:</p> <p>Undefined function or variable</p> <p>In R2015a and R2015b, matlabpool generates the following error:</p> <p>Error using matlabpool. matlabp</p> <p>Use parpool instead.</p>

Upgrade parallel computing products together

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`.

R2015b

Version: 6.7

New Features

Bug Fixes

Compatibility Considerations

Discontinued support for parallel computing products on 32-bit Windows operating systems

This release of MATLAB products no longer supports 32-bit Parallel Computing Toolbox and MATLAB Distributed Computing Server on Windows operating systems.

Compatibility Considerations

You can no longer install the parallel computing products on 32-bit Windows operating systems. If you must use Windows operating systems for the parallel computing products, upgrade to 64-bit MATLAB products on a 64-bit operating system.

More than 90 GPU-enabled functions in Statistics and Machine Learning Toolbox, including probability distribution, descriptive statistics, and hypothesis testing

Statistics and Machine Learning Toolbox™ offers enhanced functionality for some of its functions to perform computations on a GPU. For specific information, see the Statistics and Machine Learning Toolbox release notes. Parallel Computing Toolbox is required to access this functionality.

Additional GPU-enabled MATLAB functions, including support for sparse matrices

- “Additional GPU-enabled MATLAB functions” on page 6-2
- “Sparse arrays with GPU-enabled functions” on page 6-3

Additional GPU-enabled MATLAB functions

The following functions are new in their support for gpuArrays:

assert	isbanded	setxor
bandwidth	rad2deg	sortrows
deg2rad	randperm	union
expm	rectint	unique
gmres	repelem	
intersect	setdiff	

In addition to the functions above, the following functions are enhanced in their `gpuArray` support:

- `arrayfun` now supports code that includes the functions `isfloat`, `isinteger`, `islogical`, `isnumeric`, `isreal`, and `issparse`.
- `pagefun` support for `@mldivide` and `@mrdivide` is no longer limited to 32-by-32 pages.
- `cov`, `max`, `mean`, `median`, `min`, `std`, `sum`, and `var` now support the `'omitnan'` option.
- `accumarray` now supports the logical sparse argument, for generating a sparse `gpuArray` output.

For a list of MATLAB functions that support `gpuArray`, see [Run Built-In Functions on a GPU](#).

Sparse arrays with GPU-enabled functions

The following functions are new in their support for sparse `gpuArrays`:

<code>abs</code>	<code>isequal</code>	<code>round</code>
<code>angle</code>	<code>isequaln</code>	<code>sign</code>
<code>ceil</code>	<code>loglp</code>	<code>spfun</code>
<code>deg2rad</code>	<code>minus</code>	<code>sqrt</code>
<code>expm1</code>	<code>nextpow2</code>	<code>sum</code>
<code>fix</code>	<code>plus</code>	<code>uminus</code>
<code>floor</code>	<code>rad2deg</code>	<code>uplus</code>
<code>gmres</code>	<code>realsqrt</code>	

You can create a sparse `gpuArray` either by calling `sparse` with a `gpuArray` input, or by calling `gpuArray` with a sparse input.

For more information on this topic, see [Sparse Arrays on a GPU](#).

mexcuda function for easier compilation of MEX-files containing CUDA code

A new `mexcuda` function allows you to compile MEX-functions for GPU computation.

For more information, see the `mexcuda` function reference page. See also [Run CUDA or PTX Code on GPU](#).

Upgraded CUDA Toolkit version

The parallel computing products are now using CUDA Toolkit version 7.0. To compile CUDA code for CUDAKernel or CUDA MEX-files, you must use toolkit version 7.0.

Scheduler integration scripts for SLURM

This release offers a new set of scripts containing submit and decode functions to support Simple Linux® Utility for Resource Management (SLURM), using the generic scheduler interface. The pertinent code files are in the folder:

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

where *matlabroot* is your installation location.

The `slurm` folder contains a README file of instructions, and folders for `shared`, `nonshared`, and `remoteSubmission` network configurations.

For more information, view the files in the appropriate folders. See also Program Independent Jobs for a Generic Scheduler and Program Communicating Jobs for a Generic Scheduler.

parallel.pool.Constant function to create constant data on parallel pool workers, accessible within parallel language constructs such as parfor and parfeval

A new `parallel.pool.Constant` function allows you to define a constant whose value can be accessed by multiple `parfor`-loops or other parallel language constructs (e.g., `spmd` or `parfeval`) without the need to transfer the data multiple times.

For more information and examples, see `parallel.pool.Constant`.

Improved performance of mapreduce on Hadoop 2 clusters

The performance of `mapreduce` running on a Hadoop 2.x cluster with MATLAB Distributed Computing Server is improved in this release for large input data.

Enhanced and additional MATLAB functions for distributed arrays

The following functions are new in supporting distributed arrays, with all forms of codistributor (1-D and 2DBC):

bicg	rad2deg
deg2rad	rectint
polyarea	repelem
polyint	

In addition to the new functions above, the following functions are enhanced in their distributed array support:

- `cov`, `max`, `mean`, `median`, `min`, `std`, `sum`, and `var` now support the `'omitnan'` option.

For a list of MATLAB functions that support distributed arrays, see [MATLAB Functions on Distributed and Codistributed Arrays](#).

Warnings property for tasks

There is a new `Warnings` property for `parallel.Task` objects. This property contains warning information that was issued during execution of the task. The data type is a structure array with the fields `message`, `identifier`, and `stack`. This information is part of the job display in the Command Window and Job Monitor.

More consistent transparency enforcement

For more consistent behavior and results, transparency violations in `parfor`-loops and `spmd` blocks are now being enforced more strictly.

Compatibility Considerations

It is possible that `parfor` or `spmd` code that did not error in previous releases will now generate a transparency violation error. For more information, see [Transparency](#).

Anonymous functions that were saved in a previous release and that reference nested functions will not execute in this release. They will load, but when executed will issue an “Undefined function handle” error.

Upgrade parallel computing products together

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`.

R2015a

Version: 6.6

New Features

Bug Fixes

Compatibility Considerations

Support for mapreduce function on any cluster that supports parallel pools

You can now run parallel mapreduce on any cluster that supports a parallel pool. For more information, see [Run mapreduce on a Parallel Pool](#).

Sparse arrays with GPU-enabled functions

This release supports sparse arrays on a GPU. You can create a sparse `gpuArray` either by calling `sparse` with a `gpuArray` input, or by calling `gpuArray` with a sparse input. The following functions support sparse `gpuArrays`.

<code>classUnderlying</code>	<code>isfloat</code>	<code>nnz</code>
<code>conj</code>	<code>isinteger</code>	<code>numel</code>
<code>ctranspose</code>	<code>islogical</code>	<code>nzmax</code>
<code>end</code>	<code>isnumeric</code>	<code>real</code>
<code>find</code>	<code>isreal</code>	<code>size</code>
<code>full</code>	<code>issparse</code>	<code>sparse</code>
<code>gpuArray.speye</code>	<code>length</code>	<code>spones</code>
<code>imag</code>	<code>mtimes</code>	<code>transpose</code>
<code>isaUnderlying</code>	<code>ndims</code>	
<code>isempty</code>	<code>nonzeros</code>	

Note the following for some of these functions:

- `gpuArray.speye` is a static constructor method.
- `sparse` supports only single-argument syntax.
- `mtimes` does not support the case of full-matrix times a sparse-matrix.

For more information on this topic, see [Sparse Arrays on a GPU](#).

A new C function, `mxGPUISSparse`, is available for the MEX interface, to query whether a `gpuArray` is sparse or not. However, even though the MEX interface can query properties of a sparse `gpuArray`, its functions cannot access sparse `gpuArray` elements.

Additional GPU-enabled MATLAB functions

The following functions are new in their support for `gpuArrays`:

<code>cdf2rdf</code>	<code>istril</code>	<code>polyarea</code>
<code>gpuArray.freqspace</code>	<code>istriu</code>	<code>polyder</code>
<code>histcounts</code>	<code>legendre</code>	<code>polyfit</code>
<code>idivide</code>	<code>nonzeros</code>	<code>polyint</code>
<code>inpolygon</code>	<code>nthroot</code>	<code>polyval</code>
<code>isdiag</code>	<code>pinv</code>	<code>polyvalm</code>
<code>ishermitian</code>	<code>planerot</code>	
<code>issymmetric</code>	<code>poly</code>	

Note the following for some of these functions:

- `gpuArray.freqspace` is a static constructor method.

For a list of MATLAB functions that support `gpuArray`, see [Run Built-In Functions on a GPU](#).

pagefun support for mrdivide and inv functions on GPUs

For `gpuArray` inputs, `pagefun` is enhanced to support:

- `@inv`
- `@mrdivide` (for square matrix divisors of sizes up to 32-by-32)

Enhancements to GPU-enabled linear algebra functions

Many of the linear algebra functions that support `gpuArrays` are enhanced for improved performance. Among those functions that can exhibit improved performance are `svd`, `null`, `eig` (for nonsymmetric input), and `mtimes` (for inner products).

Parallel data reads from a datastore with MATLAB partition function

The `partition` function can perform a parallel read and partition of a Datastore. For more information, see `partition` and `numpartitions`. See also [Partition a Datastore in Parallel](#).

Using DNS for cluster discovery

In addition to multicast, the discover cluster functionality of Parallel Computing Toolbox can now use DNS to locate MATLAB job scheduler (MJS) clusters. For information about

cluster discovery, see Discover Clusters. For information about configuring and verifying the required DNS SRV record on your network, see DNS SRV Record.

MS-MPI support for local and MJS clusters

On 64-bit Windows platforms, Microsoft MPI (MS-MPI) is now the default MPI implementation for local clusters on the client machine.

For MATLAB job scheduler (MJS) clusters on Windows platforms, you can use MS-MPI by specifying the `-useMSMPI` flag with the `startjobmanager` command.

Ports and sockets in `mdce_def` file

The following parameters are new to the `mdce_def` file for controlling the behavior of MATLAB job scheduler (MJS) clusters.

- `ALL_SERVER_SOCKETS_IN_CLUSTER` — This parameter controls whether all client connections are outbound, or if inbound connections are also allowed.
- `JOBMANAGER_PEERSESSION_MIN_PORT`, `JOBMANAGER_PEERSESSION_MAX_PORT` — These parameters set the range of ports to use when `ALL_SERVER_SOCKETS_IN_CLUSTER = true`.
- `WORKER_PARALLELPPOOL_MIN_PORT`, `WORKER_PARALLELPPOOL_MAX_PORT` — These parameters set the range of ports to use on worker machines for parallel pools.

For more information and default settings for these parameters, see the appropriate `mdce_def` file for your platform:

- `matlabroot\toolbox\distcomp\bin\mdce_def.bat` (Windows)
- `matlabroot/toolbox/distcomp/bin/mdce_def.sh` (UNIX®)

Compatibility Considerations

By default in this release, `ALL_SERVER_SOCKETS_IN_CLUSTER` is `true`, which makes all connections outbound from the client. For pre-R2015a behavior, set its value to `false`, which also initiates a set of inbound connections to the client from the MJS and workers.

Improved profiler accuracy for GPU code

The MATLAB profiler now reports more accurate timings for code running on a GPU. For related information, see [Measure Performance on the GPU](#).

Upgraded CUDA Toolkit version

The parallel computing products are now using CUDA Toolkit version 6.5. To compile CUDA code for CUDA Kernel or CUDA MEX files, you must use toolkit version 6.5.

Discontinued support for GPU devices on 32-bit Windows computers

This release no longer supports GPU devices on 32-bit Windows machines.

Compatibility Considerations

GPU devices on 32-bit Windows machines are not supported in this release. Instead, use GPU devices on 64-bit machines.

Discontinued support for parallel computing products on 32-bit Windows computers

In a future release, support will be removed for Parallel Computing Toolbox and MATLAB Distributed Computing Server on 32-bit Windows machines.

Compatibility Considerations

Parallel Computing Toolbox and MATLAB Distributed Computing Server are still supported on 32-bit Windows machines in this release, but parallel language commands can generate a warning. In a future release, support will be completely removed for these computers, at which time it will not be possible to install the parallel computing products on them.

matlabpool function removed

The `matlabpool` function has been removed.

Compatibility Considerations

Calling `matlabpool` now generates an error. You should instead use `parpool` to create a parallel pool.

R2014b

Version: 6.5

New Features

Bug Fixes

Compatibility Considerations

Parallelization of mapreduce on local workers

If you have Parallel Computing Toolbox installed, and your default cluster profile specifies a local cluster, then execution of `mapreduce` opens a parallel pool and distributes tasks to the pool workers.

Note If your default cluster profile specifies some other cluster, the `mapreduce` function does not use a parallel pool.

For more information, see [Run mapreduce on a Local Cluster](#).

Additional GPU-enabled MATLAB functions, including `accumarray`, `histc`, `cummax`, and `cummin`

The following functions are new in their support for `gpuArrays`:

<code>accumarray</code>	<code>cummax</code>	<code>psi</code>
<code>acosd</code>	<code>cummin</code>	<code>rgb2hsv</code>
<code>acotd</code>	<code>del2</code>	<code>roots</code>
<code>acscd</code>	<code>factorial</code>	<code>secd</code>
<code>asecd</code>	<code>gammainc</code>	<code>sind</code>
<code>asind</code>	<code>gammaincinv</code>	<code>sph2cart</code>
<code>atan2d</code>	<code>gradient</code>	<code>subspace</code>
<code>atand</code>	<code>hankel</code>	<code>superiorfloat</code>
<code>betainc</code>	<code>histc</code>	<code>swapbytes</code>
<code>betaincinv</code>	<code>hsv2rgb</code>	<code>tand</code>
<code>cart2pol</code>	<code>isaUnderlying</code>	<code>toeplitz</code>
<code>cart2sph</code>	<code>median</code>	<code>trapz</code>
<code>compan</code>	<code>mode</code>	<code>typecast</code>
<code>corrcoef</code>	<code>nextpow2</code>	<code>unwrap</code>
<code>cosd</code>	<code>null</code>	<code>vander</code>
<code>cotd</code>	<code>orth</code>	
<code>cscd</code>	<code>pol2cart</code>	

Note the following for some of these functions:

- The first input argument to `gammainc` cannot contain any negative elements.

For a list of MATLAB functions that support `gpuArray`, see [Run Built-In Functions on a GPU](#).

pagefun support for mldivide on GPUs

For gpuArray inputs, pagefun is enhanced to support @mldivide for square matrix divisors of sizes up to 32-by-32.

Additional MATLAB functions for distributed arrays, including fft2, fftn, ifft2, ifftn, cummax, cummin, and diff

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:

besselh	erf	isinteger
besseli	erfc	islogical
besselj	erfcinv	isnumeric
besselk	erfcx	median
bessely	erfinv	mode
beta	fft2	pol2cart
betainc	fftn	psi
betaincinv	gamma	rgb2hsv
betaln	gammainc	sph2cart
cart2pol	gammaincinv	std
cart2sph	gammaln	toeplitz
compan	hankel	trapz
corrcoef	hsv2rgb	unwrap
cov	ifft	vander
cummax	ifft2	var
cummin	ifftn	
diff	isfloat	

Note the following for some of these functions:

- `isfloat`, `isinteger`, `islogical`, and `isnumeric` now return results based on `classUnderlying` of the distributed array.

For a list of MATLAB functions that support distributed arrays, see [MATLAB Functions on Distributed and Codistributed Arrays](#).

Data Analysis on Hadoop clusters using mapreduce

Parallel Computing Toolbox and MATLAB Distributed Computing Server support the use of Hadoop clusters for the execution environment of mapreduce applications. For more information, see:

- Configure a Hadoop Cluster
- Run mapreduce on a Hadoop Cluster

Discover Clusters Supports Microsoft Windows HPC Server for Multiple Releases

The Discover Clusters dialog can now list Microsoft Windows HPC Server clusters for different releases of parallel computing products. For more information about cluster discovery, see Discover Clusters.

For this functionality, the HPC Server client utilities must be installed on the client machine from which you are discovering. See Configure Client Computer for HPC Server.

Upgraded CUDA Toolkit Version

The parallel computing products are now using CUDA Toolkit version 6.0. To compile CUDA code for CUDAKernel or CUDA MEX files, you must use toolkit version 6.0 or earlier.

Discontinued Support for GPU Devices of Compute Capability 1.3

This release no longer supports GPU devices of compute capability 1.3.

Compatibility Considerations

This release supports only GPU devices of compute capability 2.0 or greater.

Discontinued Support for GPU Devices on 32-Bit Windows Computers

In a future release, support for GPU devices on 32-bit Windows machines will be removed.

Compatibility Considerations

GPU devices on 32-bit Windows machines are still supported in this release, but in a future release support will be completely removed for these devices.

R2014a

Version: 6.4

New Features

Bug Fixes

Compatibility Considerations

Number of local workers no longer limited to 12

You can now run a local cluster of more than 12 workers on your client machine. Unless you adjust the cluster profile, the default maximum size for a local cluster is the same as the number of computational cores on the machine.

Additional GPU-enabled MATLAB functions: `interp3`, `interpN`, `besselj`, `bessely`

The following functions are new in their support for `gpuArrays`:

```
besselj  
bessely  
interp3  
interpN
```

For a list of MATLAB functions that support `gpuArray`, see [Run Built-In Functions on a GPU](#).

Additional GPU enabled Image Processing Toolbox functions: `bwdist`, `imreconstruct`, `iradon`, `radon`

Image Processing Toolbox™ offers enhanced functionality for some of its functions to perform computations on a GPU. For specific information, see the [Image Processing Toolbox release notes](#). Parallel Computing Toolbox is required to access this functionality.

Enhancements to GPU-enabled MATLAB functions: `filter` (IIR filters); `pagefun` (additional functions supported); `interp1`, `interp2`, `conv2`, `reshape` (performance improvements)

The following functions are enhanced in their support for `gpuArray` data:

```
conv2  
filter  
interp1  
interp2  
pagefun  
rand  
randi  
randn  
reshape
```

Note the following enhancements for some of these functions:

-
- `filter` now supports IIR filtering.
 - `pagefun` is enhanced to support most element-wise `gpuArray` functions. Also, these functions are supported: `@ctranspose`, `@fliplr`, `@flipud`, `@mtimes`, `@rot90`, `@transpose`.
 - `rand(____, 'like', P)` returns a `gpuArray` of random values of the same underlying class as the `gpuArray P`. This enhancement also applies to `randi`, `randn`.

For a list of MATLAB functions that support `gpuArray`, see [Run Built-In Functions on a GPU](#).

Duplication of an existing job, containing some or all of its tasks

You can now duplicate job objects, allowing you to resubmit jobs that had finished or failed.

The syntax to duplicate a job is

```
newjob = recreate(oldjob)
```

where `oldjob` is an existing job object. The `newjob` object has all the same tasks and settable properties as `oldjob`, but receives a new ID. The old job can be in any state; the new job state is `pending`.

You can also specify which tasks from an existing independent job to include in the new job, based on the task IDs. For example:

```
newjob = recreate(oldjob, 'TaskID', [33:48]);
```

For more information, see the [recreate](#) reference page.

More MATLAB functions enhanced for distributed arrays

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:

eye
ifft
rand
randi
randn

Note the following enhancements for some of these functions:

- `ifft` and `randi` are new in support of distributed and codistributed arrays.
- `rand(____, 'like', D)` returns a distributed or codistributed array of random values of the same underlying class as the distributed or codistributed array `D`. This enhancement also applies to `randi`, `randn`, and `eye`.

For a list of MATLAB functions that support distributed arrays, see [MATLAB Functions on Distributed and Codistributed Arrays](#).

GPU MEX Support for Updated MEX

The GPU MEX support for CUDA code now incorporates the latest MATLAB MEX functionality. See [Streamlined MEX compiler setup and improved troubleshooting](#).

Compatibility Considerations

The latest MEX does not support the `mexopts` files shipped in previous releases. Updated `.xml` files are provided instead. To use the updated MEX functionality and to avoid a warning, replace the `mexopts` file you used in past releases with the appropriate new `.xml` file as described in [Set Up for MEX-File Compilation](#).

Old Programming Interface Removed

The programming interface characterized by distributed jobs and parallel jobs has been removed. This old interface used functions such as `findResource`, `createParallelJob`, `getAllOutputArguments`, `dfeval`, etc.

Compatibility Considerations

The functions of the old programming interface now generate errors. You must migrate your code to the interface described in the R2012a release topic “New Programming Interface” on page 13-2.

matlabpool Function Being Removed

The `matlabpool` function is being removed.

Compatibility Considerations

Calling `matlabpool` continues to work in this release, but now generates a warning. You should instead use `parpool` to create a parallel pool.

Removed Support for `parallel.cluster.Mpiexec`

Support for clusters of type `parallel.cluster.Mpiexec` has been removed.

Compatibility Considerations

Any attempt to use `parallel.cluster.Mpiexec` clusters now generates an error. As an alternative, consider using the generic scheduler interface, `parallel.cluster.Generic`.

R2013b

Version: 6.3

New Features

Bug Fixes

Compatibility Considerations

parpool: New command-line interface (replaces matlabpool), desktop indicator, and preferences for easier interaction with a parallel pool of MATLAB workers

- “Parallel Pool” on page 10-2
- “New Desktop Pool Indicator” on page 10-3
- “New Parallel Preferences” on page 10-4

Parallel Pool

Replaces MATLAB Pool

Parallel pool syntax replaces MATLAB pool syntax for executing parallel language constructs such as `parfor`, `spmnd`, `Composite`, and `distributed`. The pool is represented in MATLAB by a `parallel.Pool` object.

The general workflow for these parallel constructs remains the same: When the pool is available, `parfor`, `spmd`, etc., run the same as before. For example:

MATLAB Pool	Parallel Pool
<pre>matlabpool open 4 parfor ii=1:n X(n) = myFun(n); end matlabpool close</pre>	<pre>p = parpool(4) parfor ii=1:n X(n) = myFun(n); end delete(p)</pre>

Functions for Pool Control

The following functions provide command-line interface for controlling a parallel pool. See the reference pages for more information about each.

Function	Description
<code>parpool</code>	Start a parallel pool
<code>gcp</code>	Get current parallel pool or start pool
<code>delete</code>	Shut down and delete the parallel pool
<code>addAttachedFiles</code>	Attach files to the parallel pool
<code>listAutoAttachedFiles</code>	List files automatically attached to the parallel pool
<code>updateAttachedFiles</code>	Send updates to files already attached to the parallel pool

Asynchronous Function Evaluation on Parallel Pool

You can evaluate functions asynchronously on one or all workers of a parallel pool. Use `parfeval` to evaluate a function on only one worker, or use `parfevalOnAll` to evaluate a function on all workers in the pool.

`parfeval` or `parfevalOnAll` returns an object called a *future*, from which you can get the outputs of the asynchronous function evaluation. You can create an array of futures by calling `parfeval` in a `for`-loop, setting unique parameters for each call in the array.

The following table lists the functions for submitting asynchronous evaluations to a parallel pool, retrieving results, and controlling future objects. See the reference page of each for more details.

Function	Description
<code>parfeval</code>	Evaluate function asynchronously on worker in parallel pool
<code>parfevalOnAll</code>	Evaluate function asynchronously on all workers in parallel pool
<code>fetchOutputs</code>	Retrieve all output arguments from future
<code>fetchNext</code>	Retrieve next available unread future outputs
<code>cancel</code>	Cancel queued or running future
<code>wait</code>	Wait for future to complete

For more information on parallel future objects, including their methods and properties, see the `parallel.Future` reference page.

Compatibility Considerations

This release continues to support MATLAB pool language usage, but this support might discontinue in future releases. You should update your code as soon as possible to use parallel pool syntax instead.

New Desktop Pool Indicator

A new icon at the lower-left corner of the desktop indicates the current pool status.




Icon color and tool tips let you know if the pool is busy or ready, how large it is, and when it might time out. You can click the icon to start a pool, stop a pool, or access your parallel preferences.

New Parallel Preferences

Your MATLAB preferences now include a group of settings for parallel preferences. These settings control general behavior of clusters and parallel pools for your MATLAB session.

You can access your parallel preferences in these ways:

- In the **Environment** section of the **Home** tab, click **Parallel > Parallel Preferences**
- Click the desktop pool indicator icon, and select **Parallel preferences**.
- Type preferences at the command line, and click the Parallel Computing Toolbox node.



Settings in your parallel preferences control the default cluster choice, and the preferred size, automatic opening, and timeout conditions for parallel pools. For more information about these settings, see [Parallel Preferences](#).

Automatic start of a parallel pool when executing code that uses parfor or spmd

You can set your parallel preferences so that a parallel pool automatically starts whenever you execute a language construct that runs on a pool, such as `parfor`, `spmd`, `Composite`, `distributed`, `parfeval`, and `parfevalOnAll`.

Compatibility Considerations

The default preference setting is to automatically start a pool when a parallel language construct requires it. If you want to make sure a pool does not start automatically, you must change your parallel preference setting. You can also work around this by making sure to explicitly start a parallel pool with `parpool` before encountering any code that needs a pool.

By default, a parallel pool will shut down *if idle for 30 minutes*. To prevent this, change the setting in your parallel preferences; or the pool indicator tool tip warns of an impending timeout and provides a link to extend it.

Option to start a parallel pool without using MPI

You now have the option to start a parallel pool on a local or MJS cluster so that the pool does not support running SPMD constructs. This allows the parallel pool to keep running even if one or more workers aborts during `parfor` execution. You explicitly disable SPMD support when starting the parallel pool by setting its 'SpmdEnabled' property `false` in the call to the `parpool` function. For example:

```
p = parpool('SpmdEnabled', false);
```

Compatibility Considerations

Running any code (including MathWorks toolbox code) that uses SPMD constructs, on a parallel pool that was created without SPMD support, will generate errors.

More GPU-enabled MATLAB functions (e.g., `interp2`, `pagefun`) and Image Processing Toolbox functions (e.g., `bwmorph`, `edge`, `imresize`, and `medfilt2`)

Image Processing Toolbox offers enhanced functionality for some of its functions to perform computations on a GPU. For specific information, see the Image Processing Toolbox release notes. Parallel Computing Toolbox is required to access this functionality.

The following functions are new or enhanced in Parallel Computing Toolbox to support `gpuArrays`:

```
flip  
interp2  
pagefun
```

Note the following for some of these functions:

- `pagefun` allows you to iterate over the pages of a `gpuArray`, applying `@mtimes` to each page.

For more information, see the `pagefun` reference page, or type `help pagefun`.

For complete lists of functions that support `gpuArray`, see [Run Built-In Functions on a GPU](#).

More MATLAB functions enabled for distributed arrays: `permute`, `ipermute`, and `sortrows`

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:

<code>ipermute</code>	<code>zeros</code>
<code>permute</code>	<code>ones</code>
<code>sortrows</code>	<code>nan</code>
	<code>inf</code>
<code>cast</code>	<code>true</code>
	<code>false</code>

Note the following enhancements for some of these functions:

- `ipermute`, `permute`, and `sortrows` support distributed arrays for the first time in this release.
- `cast` supports the `'like'` option for distributed arrays, applying the underlying class of one array to another.
- `Z = zeros(___, 'like', P)` returns a distributed array of zeros of the same complexity as distributed array `P`, and same underlying class as `P` if class is not specified in the function call. The same behavior applies to the other similar constructors in the right-hand column of this table.

For more information on any of these functions, type `help distributed.functionname`. For example:

```
help distributed.ipermute
```

For complete lists of MATLAB functions that support distributed arrays, see MATLAB Functions on Distributed and Codistributed Arrays.

Enhancements to MATLAB functions enabled for GPUs, including ones, zeros

The following functions are enhanced in their support for gpuArray data:

zeros	eye
ones	true
nan	false
inf	
	cast

Note the following enhancements for these functions:

- `Z = zeros(___, 'like', P)` returns a gpuArray of zeros of the same complexity as gpuArray P, and same underlying class as P if class is not specified. The same behavior applies to the other constructor functions listed in this table.
- `cast` also supports the 'like' option for gpuArray input, applying the underlying class of one array to another.

For more information on any of these functions, type `help gpuArray.functionname`. For example:

```
help gpuArray.cast
```

For complete lists of MATLAB functions that support gpuArray, see Run Built-In Functions on a GPU.

gputimeit Function to Time GPU Computations

`gputimeit` is a new function to measure the time to run a function on a GPU. It is similar to the MATLAB function `timeit`, but ensures accurate time measurement on the GPU. For more information and examples, see the `gputimeit` reference page.

New GPU Random Number Generator NormalTransform Option: Box-Muller

When generating random numbers on a GPU, there is a new option for 'NormalTransform' called 'BoxMuller'. The Box-Muller transform allows faster generation of normally distributed random numbers on the GPU.

This new option is the default 'NormalTransform' setting when using the Philox4x32-10 or Threefry4x64-20 generator. The following commands, therefore, use 'BoxMuller' for 'NormalTransform':

```
parallel.gpu.rng(0, 'Philox4x32-10')
parallel.gpu.rng(0, 'Threefry4x64-20')
```

Note The 'BoxMuller' option is not supported for the CombRecursive (mrg32k3a) generator

Compatibility Considerations

In previous releases, the default 'NormalTransform' setting when using the Philox4x32-10 or Threefry4x64-20 generator on a GPU was 'Inversion'. If you used either of these generators with the default 'NormalTranform' and you want to continue with the same behavior, you must explicitly set the 'NormalTransform' with either of these commands:

```
stream = parallel.gpu.RandStream('Philox4x32-10', 'NormalTransform', 'Inversion')
parallel.gpu.RandStream.setGlobalStream(stream)
```

```
stream = parallel.gpu.RandStream('Threefry4x64-20', 'NormalTransform', 'Inversion')
parallel.gpu.RandStream.setGlobalStream(stream)
```

Upgraded MPICH2 Version

The parallel computing products are now shipping MPICH2 version 1.4.1p1 on all platforms.

Compatibility Considerations

If you use your own MPI builds, you might need to create new builds compatible with this latest version, as described in Use Different MPI Builds on UNIX Systems.

Discontinued Support for GPU Devices of Compute Capability 1.3

In a future release, support for GPU devices of compute capability 1.3 will be removed. At that time, a minimum compute capability of 2.0 will be required.

Compatibility Considerations

In R2013b, any use of `gpuDevice` to select a GPU with compute capability 1.3, generates a warning. The device is still supported in this release, but in a future release support will be completely removed for these 1.3 devices.

Discontinued Support for `parallel.cluster.Mpiexec`

Support for clusters of type `parallel.cluster.Mpiexec` is being discontinued.

Compatibility Considerations

In R2013b, any use of `parallel.cluster.Mpiexec` clusters generates a warning. In a future release, support will be completely removed.

R2013a

Version: 6.2

New Features

Bug Fixes

GPU-enabled functions in Image Processing Toolbox and Phased Array System Toolbox

More toolboxes offer 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.

More MATLAB functions enabled for use with GPUs, including `interp1` and `ismember`

The following functions are enhanced to support `gpuArray` data:

```
interp1          ismember
isfloat          isnumeric
isinteger
```

Note the following for some of these functions:

- `interp1` supports only the linear and nearest interpolation methods.
- `isfloat`, `isinteger`, and `isnumeric` now return results based on `classUnderlying` of the `gpuArray`.
- `ismember` does not support the 'rows' or 'legacy' option for `gpuArray` input.

For complete lists of functions that support `gpuArray`, see [Built-In Functions That Support `gpuArray`](#).

Enhancements to MATLAB functions enabled for GPUs, including `arrayfun`, `svd`, and `mldivide` (`\`)

The following functions are enhanced in their support for `gpuArray` data:

```
arrayfun          mrdivide
bsxfun           svd
mldivide
```

Note the following enhancements for some of these functions:

- `arrayfun` and `bsxfun` support indexing and accessing variables of outer functions from within nested functions.

-
- `arrayfun` supports singleton expansion of all arguments for all operations. For more information, see the `arrayfun` reference page.
 - `mldivide` and `mrdivide` support all rectangular arrays.
 - `svd` can perform economy factorizations.

For complete lists of MATLAB functions that support `gpuArray`, see [Built-In Functions That Support `gpuArray`](#).

Ability to launch CUDA code and manipulate data contained in GPU arrays from MEX-functions

You can now compile MEX-files that contain CUDA code, to create functions that support `gpuArray` input and output. This functionality is supported only on 64-bit platforms (`win64`, `glnxa64`, `maci64`). For more information and examples, see [Execute MEX-Functions Containing CUDA Code](#).

For a list of C functions supporting this capability, see the group of [C Functions in GPU Computing](#).

Automatic detection and transfer of files required for execution in both batch and interactive workflows

Parallel Computing Toolbox can now automatically attach files to a job so that workers have the necessary code files for evaluating tasks. When you set a job object's `AutoAttachFiles` to `true`, an analysis determines what files on the client machine are necessary for the evaluation of your job, and those files are automatically attached to the job and sent to the worker machines.

You can set the `AutoAttachFiles` property in the Cluster Profile Manager, or at the command-line. To get a listing of the files that are automatically attached, use the `listAutoAttachedFiles` method on a job or task object.

For more information, see [Pass Data to and from Worker Sessions](#).

If the `AutoAttachFiles` property in the cluster profile for the MATLAB pool is set to `true`, MATLAB performs an analysis on `spmd` blocks and `parfor`-loops to determine what code files are necessary for their execution, then automatically attaches those files to the MATLAB pool job so that the code is available to the workers.

When you use the MATLAB editor to update files on the client that are attached to a `matlabpool`, those updates are automatically propagated to the workers in the pool.

More MATLAB functions enabled for distributed arrays

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:

<code>cumprod</code>	<code>qr</code>
<code>cumsum</code>	<code>prod</code>
<code>eig</code>	

Note the following enhancements for some of these functions:

- `eig` supports generalized eigenvalues for symmetric matrices.
- `qr` supports column pivoting.
- `cumprod`, `cumsum`, and `prod` now support all integer data types; and `prod` accepts the optional `'native'` or `'double'` argument.

R2012b

Version: 6.1

New Features

Bug Fixes

Compatibility Considerations

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

- “gpuArray Support” on page 12-2
- “MATLAB Code on the GPU” on page 12-2

gpuArray Support

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

bitget	cov	normest
bitset	issparse	pow2
cond	mpower	var
convn	nnz	

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

blkdiag	ismatrix	isvector
cross	isrow	std
iscolumn		

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:

<code>atan2d</code>	<code>mrdivide</code>	<code>struct2cell</code>
<code>bsxfun</code>	<code>repmat</code>	<code>typecast</code>
<code>cell2struct</code>	<code>sparse</code>	

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

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.)

R2012a

Version: 6.0

New Features

Bug Fixes

Compatibility Considerations

New Programming Interface

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

Previous Scheduler Objects	New Cluster Objects
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
<code>findResource</code>	<code>parcluster</code>
<code>createJob</code>	<code>createJob</code> (no change)
<code>createParallelJob</code>	<code>createCommunicatingJob</code> (where 'Type' = 'SPMD')
<code>createMatlabPoolJob</code>	<code>createCommunicatingJob</code> (where 'Type' = 'Pool')
<code>destroy</code>	<code>delete</code>
<code>clearLocalPassword</code>	<code>logout</code>
<code>createTask</code>	<code>createTask</code> (no change)
<code>getAllOutputArguments</code>	<code>fetchOutputs</code>
<code>getJobSchedulerData</code>	<code>getJobClusterData</code>
<code>setJobSchedulerData</code>	<code>setJobClusterData</code>
<code>getCurrentJobmanager</code>	<code>getCurrentCluster</code>
<code>getFileDependencyDir</code>	<code>getAttachedFilesFolder</code>

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
<code>DataLocation</code>	<code>JobStorageLocation</code>
<code>Configuration</code>	<code>Profile</code>

Previous Scheduler Properties	New Cluster Properties
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.

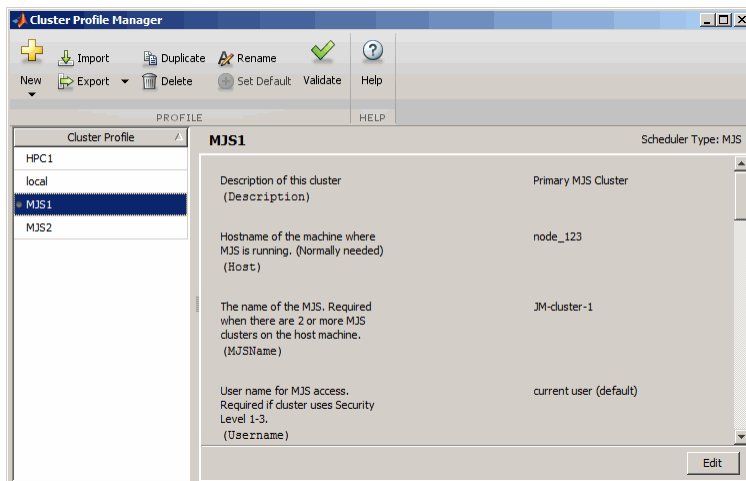
Cluster Profiles

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

Function	Description
<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

GPUArray Support

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

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

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:

angle	ifftshift	rank
fliplr	kron	squeeze
flipud	mean	rot90
flipdim	perms	trace
fftshift		

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

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

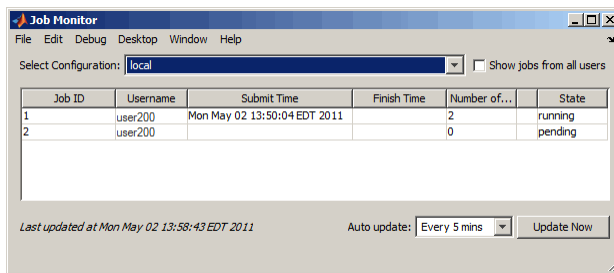
New Features

Bug Fixes


Compatibility Considerations

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

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>betaIn</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

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

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

Bug Fixes

Compatibility Considerations

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

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

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

Bug Fixes

Compatibility Considerations

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

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
- `sgc` — Sun™ Grid Engine
- `ssh` — generic UNIX-based scripts
- `winmpiexec` — mpiexec on Windows

For the updated scheduler folders (`lsf`, `pbs`, `sgc`), 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`, `sgc`), 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

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 `sctranspose` 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

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.