

DiskBoss

DATA MANAGEMENT



File Synchronization


Version 14.1

Jan 2024

www.diskboss.com
info@flexense.com
Flexense Ltd.

1 DiskBoss Overview

DiskBoss is an automated, policy-based data management solution allowing one to analyze disks, directories and network shares, classify and categorize files, search and cleanup duplicate files, perform automated file management operations according to user-defined rules and policies, synchronize disks, directories and network shares, compare directories and files, perform bulk file delete and secure data wiping operations, detect unauthorized changes in files and directories, etc.



All disk space analysis and file management operations are integrated into a centralized and easy-to-use GUI application allowing one pre-configure analysis and file management operations as user-defined commands and execute any required command in a single mouse click using the DiskBoss GUI application or direct desktop shortcuts.

- Disk Space Analysis
- File Classification and Organizing
- Duplicate Files Search and Cleanup
- Bulk File Delete and Secure Data Wiping
- Automated, Policy-Based File Management
- Real-Time Disk Change Monitoring
- High-Speed File Synchronization
- Secure File Transfer Operations
- File Integrity Monitoring

DiskBoss allows one to generate various types of pie charts and save HTML, PDF, Excel, text, CSV and XML reports for all types of disk space analysis, file classification and file search operations. The user is provided with the ability to categorize and filter analysis and file classification results and perform file management operations on categories of files.

IT administrators are provided with extensive SQL database integration capabilities allowing one to submit disk space analysis, file classification, duplicate files search and disk change monitoring reports into an SQL database. Reports from multiple servers and NAS storage devices may be submitted to a centralized SQL database allowing one to display charts showing the used disk space, file categories and duplicate files per user or per server and providing an in-depth visibility into how the disk space is used, what types of files are stored and how much space is wasted on duplicate files across the entire enterprise.

In addition to the DiskBoss GUI application, IT and storage administrators are provided with the DiskBoss command line utility, which can be used to execute all types of analysis and file management operations from batch files and shell scripts. The command line utility provides an extensive set of command line options allowing one to execute various types of disk space analysis, file synchronization, data migration and bulk file delete operations pre-configured for user-custom needs and hardware configurations.

```

c:\Program Files\DiskBoss Ultimate\bin>diskboss -classify -dir c:\

DiskBoss Ultimate v9.0.18 (Mar  8 2018)

Operation Successfully Completed

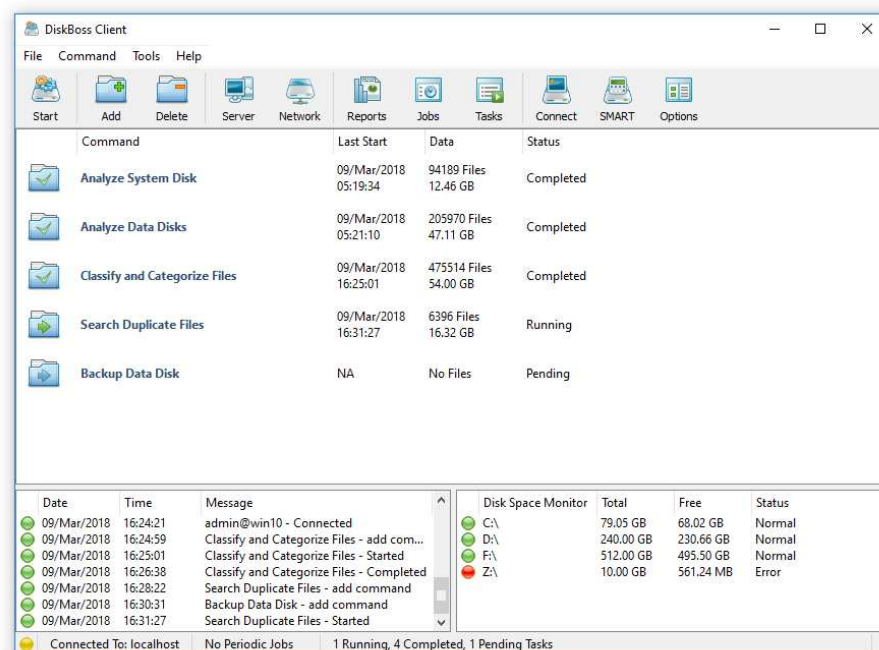
Date           - 2018/03/08
Time           - 18:04:30
Total Files    - 127539
Total Space    - 13.96 GB
Performance   - 14653 Files/Sec
Process Time   - 8.70 Sec
Host Name      - Windows10-64
User Name      - alex-s
Input Dirs     - c:\

Programs, Extensions and Script Files - 23267 Files, 9.43 GB
Windows Files                          - 43840 Files, 1.85 GB
Miscellaneous Files                    - 40489 Files, 1.64 GB
Archive, Backup and Disk Image Files   - 2540 Files, 286.23 MB
Unknown Files                          - 5321 Files, 258.63 MB
Internet Files                         - 4709 Files, 153.99 MB
Game Files                             - 18 Files, 118.05 MB
Development Files                      - 1099 Files, 103.72 MB
Images, Pictures and Graphic Files      - 5140 Files, 57.13 MB
Documents, Books and Help Files         - 960 Files, 35.46 MB
Movies, Clips and Video Files           - 144 Files, 33.95 MB
Music and Audio Files                   - 7 Files, 13.21 MB
Temporary Files                         - 5 Files, 304.63 KB

c:\Program Files\DiskBoss Ultimate\bin>

```

Finally, IT professionals and enterprise customers are provided with DiskBoss Server – a server-based product version, which runs in the background as a service and is capable of performing all type of disk space analysis and file management operations in a fully automatic and unattended mode according to a user-specified schedule.




DiskBoss Server can be controlled locally or through the network using the DiskBoss client GUI application or the command line utility. DiskBoss Server provides the ability to pre-configure various types of disk space analysis and/or policy-based file management operations, schedule periodic jobs, save analysis reports into a number of different formats, export analysis results to an SQL database, periodically synchronize disks, directories and network shares and monitor critical disks and directories for unauthorized changes.

2 DiskBoss File Synchronization

DiskBoss provides advanced file synchronization capabilities allowing one to synchronize files between local disks, directories, network shares, NAS storage devices and enterprise storage systems. The DiskBoss file synchronization engine is capable of synchronizing files using multiple one-way and two-way file synchronization modes, provides performance tuning options, periodic and real-time file synchronization capabilities and allows one to synchronize specific file types and categories.



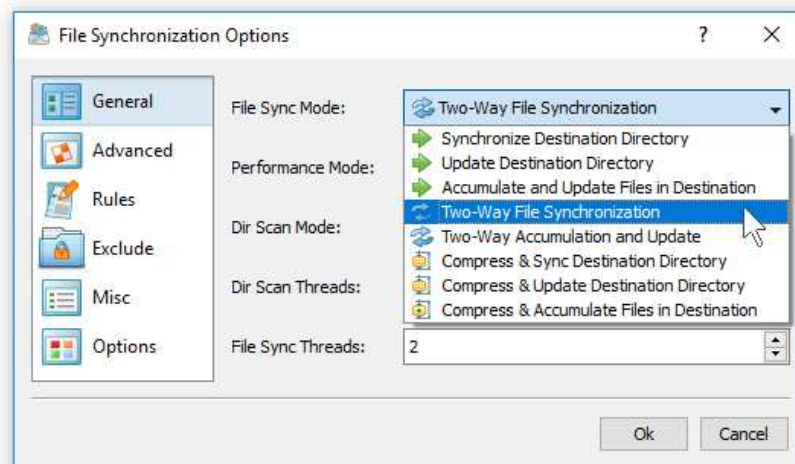
In order to synchronize a pair of directories, press the 'Sync' button located on the main toolbar, enter the source and destination directories and press the 'Preview' button. DiskBoss will scan the specified source and destination directories, compare files between the source and destination directories and show a list of file synchronization actions that should be performed.



On the file synchronization preview dialog, review sync actions, select/unselect actions that should be performed and press the 'Sync' button to execute the selected file synchronization actions. Optionally, press the 'Options' button to control advanced file synchronization options or press the 'Export' button to save file synchronization actions to an HTML, XML, text or Excel CSV report. The DiskBoss file synchronization engine is especially optimized for modern hardware and storage devices and is capable of effectively utilizing RAID storage devices, multi-core/multi-CPU systems and Gigabit networks delivering maximum possible performance.

3 Selecting File Synchronization Mode

Depending on specific needs, it may be required to perform different types of file synchronization operations. For example, sometimes users may be interested in keeping all files in the destination directory even when some of them were deleted from the source directory. Another option is to perform two-way accumulation and update keeping both directories completely identical. DiskBoss provides eight different one-way and two-way file synchronization modes allowing one to perform different types of file synchronization operations according to user-specific needs and requirements.

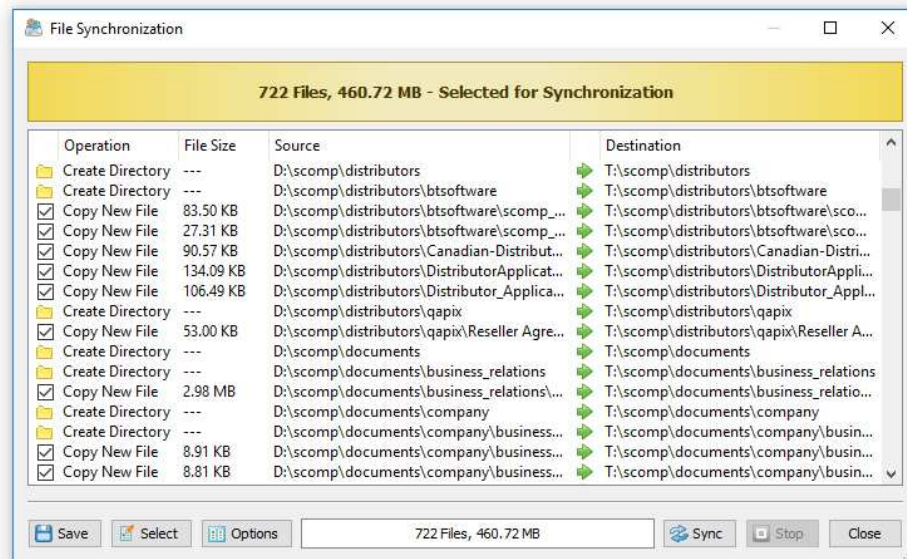


In order to change the file synchronization mode for a file synchronization command, open the file synchronization command dialog, select the 'General' tab and select one of the following file synchronization modes:

- **Synchronize Destination Directory** (one-way) - All changes made in the source directory will be propagated to the destination directory. Files deleted from the source directory will be deleted from the destination directory. All files changed or deleted in the destination directory will be restored using files from the source directory. No changes will be made in the source directory. After the file synchronization process is complete both locations will be identical.
- **Update Destination Directory** (one-way) - Newly created and modified source files will be copied to the destination directory. Files deleted from the source directory will be deleted from the destination directory. Files deleted from the destination directory will be restored using files from the source directory. Newly created and modified destination files will be kept in place. No changes will be made in the source directory.
- **Accumulate and Update Files in Destination** (one-way) - Newly created and modified source files will be copied to the destination directory. Files deleted from the destination directory will be restored using files from the source directory. Newly created and modified destination files will be kept in place. No files will be deleted from the destination directory. No changes will be made in the source directory.
- **Two-Way File Synchronization** - Newly created and modified files will be synchronized in both directions. Files deleted from the source directory will be deleted from the destination directory. Files deleted from the destination directory will be restored from the source directory. Changes may be made in both locations. After the file synchronization process is complete both locations will be identical.
- **Two-Way Accumulation and Update** - Newly created and modified files will be updated in both directions. Files deleted from one location will be restored from the second location. Changes may be made in both locations. After the file synchronization process is complete both locations will be identical.

4 Synchronizing Files with Preview

File synchronization with preview is very useful providing the user with a clear picture about what files will be synchronized. By default, DiskBoss synchronizes directories with preview, but the user is provided with the ability to disable the preview mode and synchronize all the required files on-the-fly without showing the preview dialog.



5 Synchronizing Files Without Preview

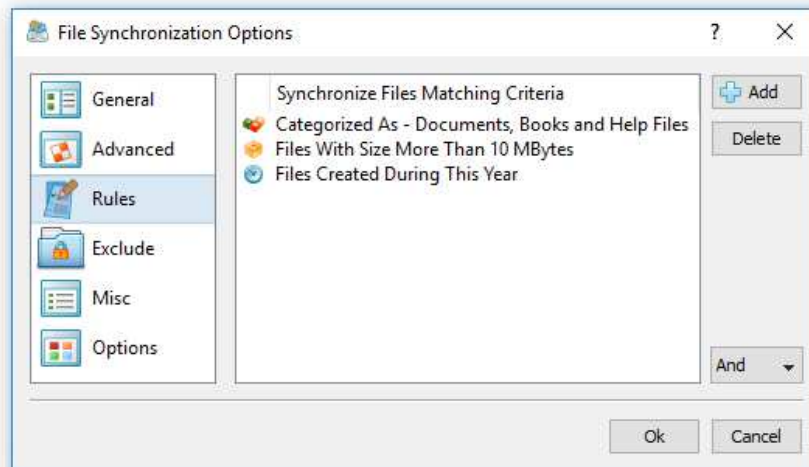
Sometimes, when the user needs to synchronize large directories or disks containing many thousands of files, file synchronization with preview may be ineffective or completely inappropriate especially due to the fact that none will have the time to review lists of file synchronization actions containing huge amounts of items. Moreover, unattended file synchronization operations performed periodically at specific time intervals or triggered by a disk change monitor, should be performed fully automatically without requiring any user intervention.



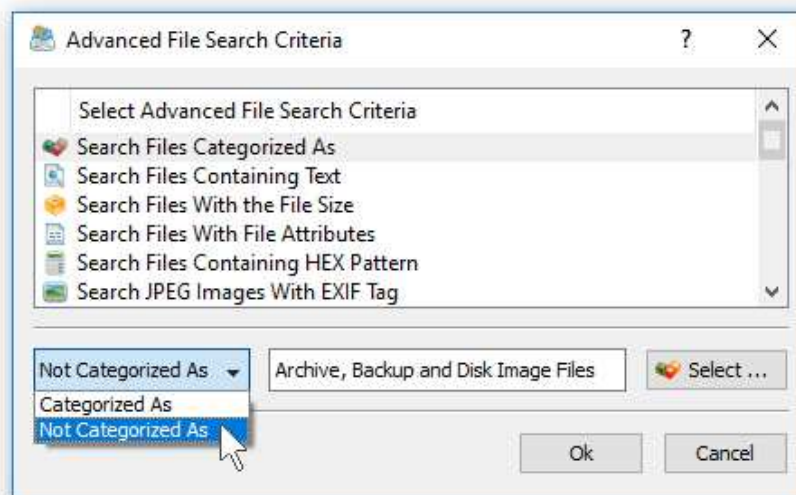
In order to disable the preview dialog for a file synchronization operation, open the operation dialog and set the process mode combo box to 'Sync Without Preview'. The stream file synchronization mode is especially designed for file systems containing huge amounts of files and is capable of synchronizing large directories, disks and NAS storage devices very fast and effectively. In addition, the stream file synchronization mode is ideal for automated, periodic file synchronization operations performed in the background without any user intervention.

6 Synchronizing Specific File Types of Files

DiskBoss Ultimate and DiskBoss Server provide power computer users and IT administrators with the ability to synchronize specific file types or file categories using one or more flexible file matching rules. For example, the user may specify to synchronize all types of documents and digital images with the file size more than 10 MB that were created during this year.



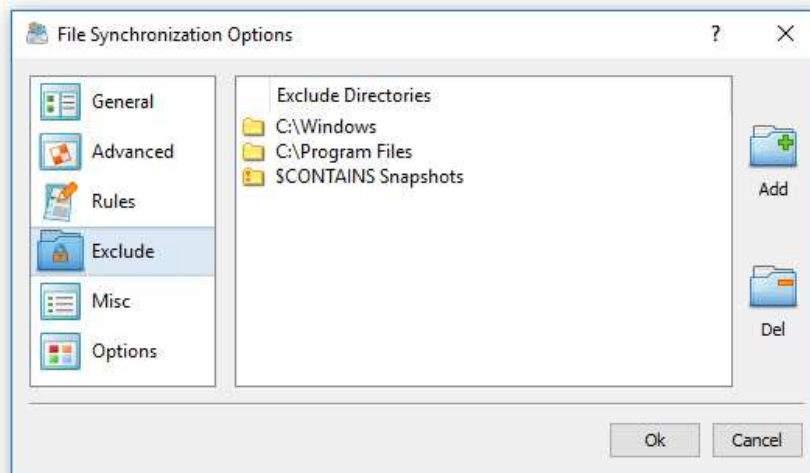
In order to add one or more file matching rules, open the file synchronization command dialog, press the 'Options' button, select the 'Rules' tab and press the 'Add' button. On the 'Rules' dialog select an appropriate rule type and specify all the required parameters. During file synchronization, DiskBoss will scan the entered source and destination directories and apply the specified file matching rules to all the existing files. Files not matching the specified rules will be just skipped from the file synchronization process.



Another option is to exclude files from the file synchronization process using one or more negative file matching rules. For example, in order to exclude all types of backup and archive files from the file synchronization process, add the 'Category' rule, select the 'Backup and Archive Files' category and set the operator to 'Not Categorized'.

7 Exclude Directories From File Synchronization

Sometimes, it may be required to exclude one or more subdirectories from the file synchronization process. For example, if you need to synchronize two directories excluding one or two special subdirectories, you may specify the top level directories as the source and destination directories and add the subdirectories that should be skipped to the exclude list.



In order to add one or more directories to the exclude list, open the file synchronization command dialog, press the 'Options' button, select the 'Exclude' tab and press the 'Add' button. Keep in mind that exclude directories are case sensitive and should be specified with the same case as stored on the disk. All files and subdirectories located in the specified exclude directories will be excluded from the file synchronization process. Select an exclude directory and press the 'Delete' button, to remove the selected directory from the exclude list.

In addition, advanced users are provided with a number of exclude directories macro commands allowing one to exclude multiple directories using a single macro command.


DiskBoss provides the following exclude directories macro commands:

- **\$BEGINS <Text String>** - this macro command excludes all directories beginning with the specified text string.
- **\$CONTAINS <Text String>** - this macro command excludes all directories containing the specified text string.
- **\$ENDS <Text String>** - this macro command excludes all directories ending with the specified text string.
- **\$REGEX <Regular Expression>** - this macro command excludes directories matching the specified regular expression.
- **\$DIRLIST <File Name>** - this macro command excludes all directories listed in the specified text file.
- **\$FILELIST <File Name>** - this macro command excludes all files listed in the specified text file.

For example, the exclude macro command '\$CONTAINS Temporary Files' will exclude all directories with 'Temporary Files' in the full directory path and the exclude macro command '\$REGEX \.(TMP|TEMP)\$' will exclude directories ending with '.TMP' or '.TEMP'.


8 Advanced File Synchronization Options

DiskBoss provides a large number of advanced files synchronization options allowing one to tune the performance of file synchronization operations, control the file copy recovery and fault tolerance options, select which file meta data to copy, etc. In order to access the performance tuning options, open the file sync command dialog, press the '**Options**' button and select the '**General**' tab.



- **Performance Mode** - the performance mode allows one to select one of the following modes: Full Speed, Medium Speed and Low Speed. The performance mode provides the ability to intentionally slow down the performance of the file sync operation in order to minimize the potential impact on the running production applications.
- **Dir Scan Mode** - the directory scanning mode allows one to select one of the following directory scanning modes: Sequential or Parallel. In the Sequential directory scanning mode, DiskBoss scans directories one-after-one using a single processing thread. In the Parallel directory scanning mode, DiskBoss scans directories in parallel using a number of processing threads according to the 'Dir Scan Threads' option, which significantly improves the performance of the file synchronization operation when synchronizing large numbers of files via a slow, high-latency network.
- **Dir Scan Threads** - this option specifies the number of parallel directory scanning threads to use in the **Parallel** directory scanning mode.
- **File Sync Threads** - this option specifies the number of parallel file synchronization threads to use during the file synchronization process. Multi-Stream file synchronization significantly improves the file synchronization performance when synchronizing a large number of files between multiple disks, NAS storage devices or enterprise storage systems.

The '**Advanced**' file synchronization options tab provides a number of different file copy modes optimized for different hardware configurations and user-specific needs:




- **Active File Copy Mode** - in this file copy mode DiskBoss automatically switches between different file copy modes depending on the size and the type of each file. Small files are transferred using a regular, buffered file copy mode, while large files are transferred using a non-buffered file copy mode and optionally, when the hardware platform supports it, allows to enable the write-through I/O mode.
- **Regular, Buffered File Copy Mode** - in this file copy mode DiskBoss performs regular, buffered file copy operations according to the specified file I/O block size and memory alignment.
- **Fault-Tolerant File Copy Mode** - in this file copy mode DiskBoss creates a temporary file for each file that should be copied and only after a successful file copy operation replaces the original file with the new one. If a network failure will occur during a file copy operation, original files will remain in place without any changes.
- **Operating System Native File Copy Mode** - this file copy mode is optimized for modern operating systems such as Windows Server 2012 and newer. This file copy mode supports SMB direct, direct RDMA file transfers, multi-channel and write-through file I/O operations. In order to take full advantage of all performance optimizations when synchronizing files via the network, Windows Server 2012 or newer should be installed on both sides.

In addition, the '**Advanced**' file synchronization options tab provides a number of performance, fault tolerance and recovery options:

- **Verification Mode** - the file copy verification mode specifies how to verify files copied to the destination directory. By default, verification of copied files is disabled and in order to enable verification of copied files the user needs to select an appropriate verification mode and then enable verification on the 'Options' tab.
- **Memory Buffer** - this option specifies the size of the file copy memory buffer used for each file copy thread.
- **Memory Alignment** - this option specifies the alignment of the file copy memory buffer, which is critical for fast, non-buffered file copy operations. By default, the memory alignment is set to 4096 bytes, which is appropriate for modern file systems and storage devices.
- **Copy Retry Count** - this option specifies how many times to retry a failed file copy operation. Sometimes, a file may be temporarily locked by an application and this option allows one to retry a failed file copy operation after a user specified delay.
- **Copy Retry Delay** - this option specifies the time delay (in milliseconds) between failed file copy retries.

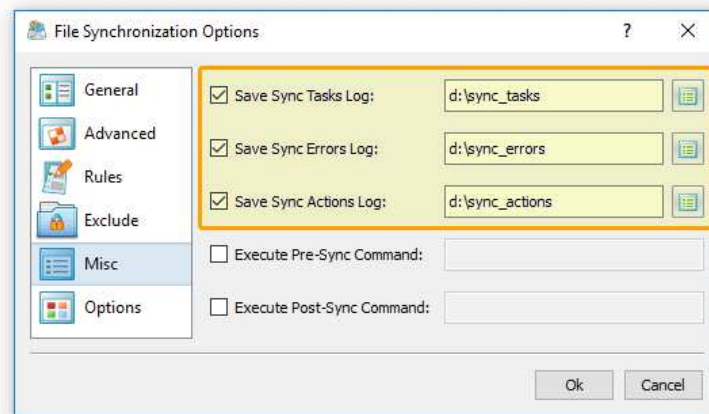
The '**Options**' tab provides the ability to control which file meta data is copied and allows one to enable/disable various security and verification options.



- **Sync File Attributes** - when this option is enabled, DiskBoss will copy file attributes for each copied file. By default, this option is disabled and files are created in the destination directory with normal file attributes.
- **Sync User/Group Security Descriptor** - when this option is enabled, DiskBoss will copy the file ownership information for each copied file. By default, this option is disabled and files are created in the destination directory with the user account running the DiskBoss application. In addition, in order to be able to use this option, the user needs to have administrative permissions. Finally, when synchronizing files to a remote destination directory located on a NAS device, the NAS device should provide support for the NTFS security capabilities.
- **Sync Discretionary Access Control List (DACL)** - when this option is enabled, DiskBoss will copy the DACL security information for each copied file. By default, this option is disabled and files are created in the destination directory with permissions inherited from the parent directory. In addition, in order to be able to use this option, the user needs to have administrative permissions. Finally, when synchronizing files to a remote destination directory located on a NAS device, the NAS device should provide support for the NTFS security capabilities.
- **Sync System Access Control List (DACL)** - when this option is enabled, DiskBoss will copy the SACL security information for each copied file. By default, this option is disabled and files are created in the destination directory with permissions inherited from the parent directory. In addition, in order to be able to use this option, the user needs to have administrative permissions. Finally, when synchronizing files to a remote destination directory located on a NAS device, the NAS device should provide support for the NTFS security capabilities.
- **Enable-Bit Level File Synchronization** - when this option is enabled, DiskBoss will check which data blocks have changed in each new file and write to the destination file only the changed data blocks.
- **Enable File Write Trough** - when this option is enabled and DiskBoss is running on a newer operating system such as Windows Server 2012 or newer, files will be copied using the operating system native RDMA mode.
- **Verify Files After Copy** - when this option is enabled, DiskBoss will verify the integrity of each copied file using the file verification mode selected on the 'Advanced' tab.

9 File Synchronization Logs

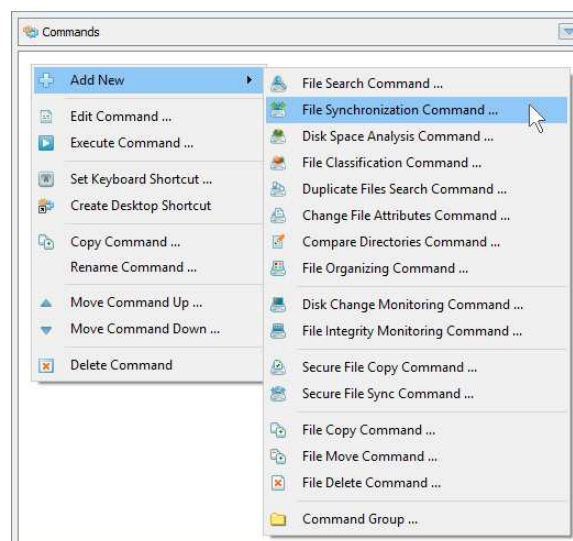
DiskBoss provides the ability to save different types of file synchronization logs allowing one to keep track of performed file synchronization operations and record file synchronization errors. In order to enable file synchronization logs for a file synchronization command, open the file synchronization options dialog, select the 'Misc' tab, enable one or more file synchronization logs and specify directories where to save the log files.



For each log file, the user is provided with the ability to specify the log file format (plain text or Excel CSV), the log file mode and the log files history. By default, DiskBoss will save an individual log file for each file synchronization operation, but the user can switch to consolidated hourly, daily or monthly logs files, which will be kept in the file synchronization logs directory according to the user-specified number of days.

10 Pre-Configured File Synchronization Commands

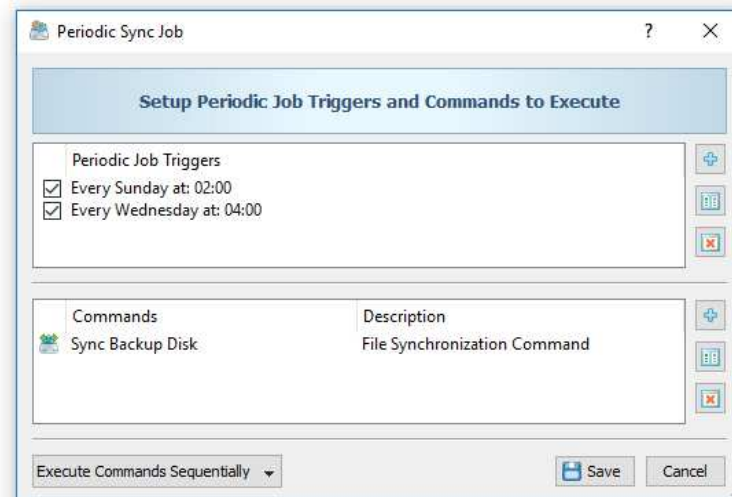
DiskBoss provides the user with the ability to pre-configure custom file synchronization operations as user-defined commands and execute such commands in a single mouse click using the DiskBoss GUI application or direct desktop shortcuts.



User-defined commands may be managed and executed through the commands dialog or the commands tool pane. In order to add a new command through the commands pane, press the right mouse button over the pane and select the 'Add New - File Synchronization Command' menu item. In order to execute a previously saved command, just click on the command item in the commands tool pane or create a direct desktop shortcut on the desktop.

11 Periodic File Synchronization

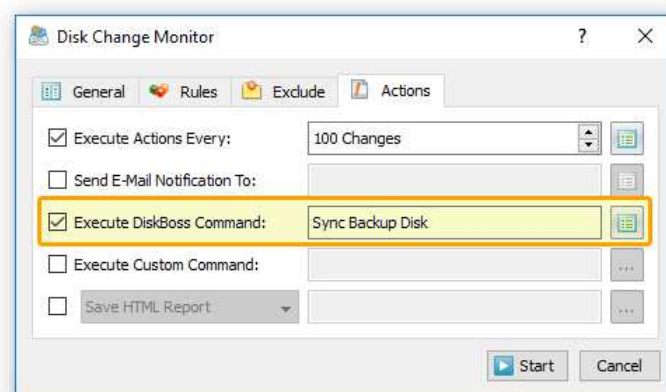
Sometimes, it may be required to execute a file synchronization operation periodically at specific time intervals. The first option to do that is to configure a periodic job in the DiskBoss GUI application and select the required file synchronization command to be executed periodically. Another option is to use DiskBoss Server, which runs as a service in the background and may be used to periodically execute file synchronization operations in a fully automatic, unattended mode.



In order to add a new periodic job, select the 'Tools – Manage Periodic Jobs' menu item and press the 'Add' button. On the periodic job dialog, select the file synchronization command that should be executed and specify the time period.

12 Real-Time File Synchronization

DiskBoss Ultimate and DiskBoss Server provide the ability to monitor one or more disks or directories and automatically execute a file synchronization operation after a user-specified number of changes is detected.



In order to execute a file synchronization command using the DiskBoss real-time disk change monitor, create a new disk change monitoring command, specify one or more disks or directories that should be monitored, select the 'Actions' tab, specify the number of disk changes that should trigger the file synchronization command and selected the name of the file synchronization command that should be executed once the disk change monitor reaches the specified number of changes.

13 Using the DiskBoss Command Line Utility

In addition to the GUI application, DiskBoss Ultimate and DiskBoss Server provide a command line utility allowing one to execute file synchronization commands from batch files and shell scripts. The command line utility is located in the '**<ProductDir>/bin**' directory.

Command Line Syntax:

diskboss -sync -source <Source Directory> -dest <Dest Directory>

This command synchronizes files between directories, local disks or network shares.

diskboss -execute <File Synchronization Command>

This command executes the specified user-defined file synchronization command pre-configured using the DiskBoss GUI application or imported from an XML file.

Parameters:

-source <Source Directory>

This parameter specifies the source directory for file synchronization. In order to ensure proper parsing of command line arguments, directories and file names containing space characters should be double quoted.

-dest <Destination Directory>

This parameter specifies the destination directory for file synchronization. In order to ensure proper parsing of command line arguments, directories and file names containing space characters should be double quoted.

Options:

-sync_mode <SD | SUD | SAD | STW | STA>

SD - Sets the sync-destination file synchronization mode (Default)

All changes made in the source directory will be propagated to destination. Files deleted from the source directory will be deleted from destination. All files changed or deleted in destination will be restored from source.

SUD - Sets the update-destination file synchronization mode

Newly created and modified source files will be copied to destination. Files deleted from the source directory will be deleted from destination. Files deleted from the destination directory will be restored from source.

SAD - Sets the accumulate-destination file synchronization mode

Newly created and modified source files will be copied to destination. Files deleted from the destination directory will be restored from source. Newly created and modified destination files will be kept in place.

STW - Sets the two-way file synchronization mode

Newly created and modified files will be synchronized in both directions. Files deleted from the source directory will be deleted from destination. Files deleted from the destination directory will be restored from source.

STA - Sets the two-way accumulate file synchronization mode

Newly created and modified files will be updated in both directions. Files deleted in one location will be restored from the second location. After finishing the synchronization process both locations will be identical.

-streams <StreamCount>

This parameter specifies the number of simultaneous file synchronization streams.

-file_delay <delay in milliseconds>

This parameter sets the file delay to the specified value in milliseconds. Use this option when you need to slowly synchronize files to or from a busy production system in order to minimize the potential performance degradation.

-block_delay <delay in milliseconds>

This parameter sets the block delay to the specified value in milliseconds. Use this option when you need to slowly synchronize files to or from a busy production system in order to minimize the potential performance degradation.

-copy_dacl

This parameter instructs to copy files' access control lists (DACL).

-copy_sacl

This parameter instructs to copy files' security control lists (SACL).

-copy_owner

This parameter instructs to copy files' user/group information.

-v


This command shows the product's major version, minor version, revision and build date.

-help


This command shows the command line usage information.

14 File Synchronization Performance

DiskBoss is optimized for the modern hardware architectures and it is capable to sync files using multiple parallel file synchronization streams. When synchronizing large files stored on fast NVMe SSD disks, the performance of the DiskBoss file synchronization operations reaches up to 1,068 MB/sec with a single file synchronization stream, increases up to 1,840 MB/sec with two parallel file synchronization streams and scales up to 3,220 MB/sec with four parallel streams allowing one to synchronize huge amounts of data very fast.




When synchronizing medium-sized files stored on fast NVMe SSD disks, the performance of the file synchronization operations reaches up to 873 MB/sec with a single file synchronization stream, increases up to 1,650 MB/sec with two parallel file synchronization streams and scales up to 2,816 MB/sec. Usually, processing large numbers of small files incurs significant performance degradations, but in the case of NVMe SSD disks, synchronizing small files still results in a very high performance level reaching up to 598 MB/sec with a single file synchronization stream, increasing up to 927 with two parallel file synchronization streams and scaling up to 1,273 MB/sec with four parallel file synchronization streams.




Regular SATA SSD drives are significantly slower than NVMe SSD drives, but still deliver a very respectable performance level. For example, when synchronizing large files stored on regular SATA SSD drives, the performance of the file synchronization operations reaches up to 330 MB/Sec with a single file synchronization stream and scales up to 478 MB/sec when synchronizing large files with 4 parallel file synchronization streams.

Synchronizing files stored on NAS storage devices via a network is more complicated because the user needs to take into account the speed and the latency of the network. If the computer, on which DiskBoss is installed, is connected to the NAS storage device via a high-speed, low-latency network, the performance of the file synchronization operations may reach up to 79 MB/sec with a single file synchronization stream and scale up to 112 MB/sec with four parallel file synchronization streams.




On the other hand, if DiskBoss will need to access network shares via the Internet or via a long-distance, high-latency network, the performance of the file synchronization operations will be relatively slow. One of the options to increase the performance of the file sync operations in such configurations is to set the 'Parallel' directory scanning mode and increase the number of parallel file synchronization streams to 8 or even 16 disregarding how many CPUs are actually installed on the computer.




When synchronizing large files via the 2.5 Gigabit Ethernet, DiskBoss reaches up to 152 MB/Sec with a single file synchronization thread and scales up to 278 MB/Sec with 4 parallel file synchronization threads, which is close to the theoretical maximum of the 2.5 Gigabit Ethernet.

If the user needs to synchronize a large number of small files, the 2.5 Gigabit Ethernet also provides very significant improvements over the regular Gigabit Ethernet network. With a single file synchronization thread, DiskBoss reaches up to 772 Files/Sec and scales up to 2007 Files/Sec, which is a very significant 68% improvement compared to the regular Gigabit Ethernet.



Modern USB flash drives provide plenty of the storage space and are reasonably fast allowing one to store vast amounts of data for backup purposes. When synchronizing large files stored on a high-speed USB flash drive, DiskBoss can reach up to 267 MB/sec with a single file synchronization stream.



But, with two parallel file synchronization streams, the performance of the file synchronization operations degrades to 170 MB/sec and remains slightly degraded with more parallel file synchronization streams showing that USB flash drives are significantly less scalable than NVMe and SATA SSD drives.