

# CrushSync Documentation

by Ben Spink <ben@crushsync.com> (4/10/2009)



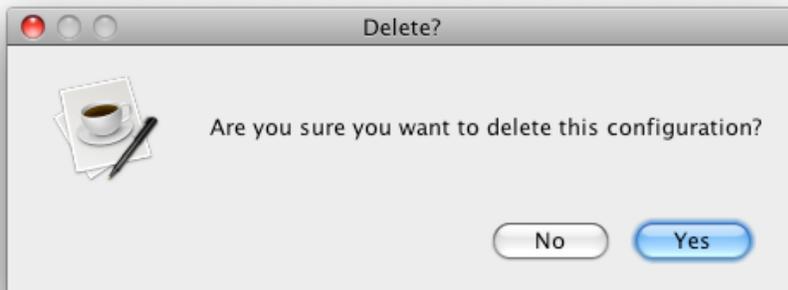
This documentation was written as of CrushSync 1.2. I will try and keep it up to date as CrushSync evolves. If you see any discrepancies in what I am describing in the documentation versus the actual application, please let me know.

## Setup Tab

The top area of this tab contains the syncs that you have configured in CrushSync. You can create new syncs, duplicate existing ones, delete them, run them, or preview them. Additionally you have options on the right to create a sync group, as well as alter the order of the syncs.



The top left button is the New Sync button. It allows you to create a new sync. The next button is the duplicate button and will make a copy of the currently selected sync in the list below. The next button is the delete button and will delete the currently selected sync.



The next button is the green play button which will run the currently selected sync. Once you click this button, you may wish to change to the Monitor tab to see the sync's progress. The most common button looks like the play button, but also has the word "Preview" on it. This will run the sync in preview mode, and display a list of all items that will be affected along with the action that is going to be applied to the item. You can either cancel this preview window, or click the "Run Now" button on the preview window to run the sync as its listed.

Action	src	dest	src modified	dest modified	path
copy	source	destination	02/28/09 12:46:57 PM	04/10/09 12:49:12 AM	.DS_Store
copy	source	destination	02/10/05 07:22:27 AM		b.jar

Buttons: Cancel, Run Now

On the right side of the sync listing is the group control buttons. The plus button will create a new group while the up and down buttons will move the selected sync around in the list. To add a sync to a group, drag and drop the sync in the list. If you drop it when its over a group, it will be added to the group.



A group can only have its name set, and be scheduled to run at different times. It does not have other options. If a group is run, each sync inside of the group is run sequentially, with a 30 second pause after each sync is finished. This allows you to have an easier organization and control the order of when syncs will run.

## Type

There are three different synchronization modes that CrushSync can do.

**1 : Type**   **2 : Source/Destination**   **3 : Schedule**   **Options**

**Two Way Synchronization**  
Two way synchronizations work by copying files between the two locations in both directions. A snapshot is taken every time a synchronization finishes. The snapshot is used for the next synchronization to determine what has changed, and what should be synchronized. Files are only deleted if they existed in a snapshot, and are now no longer present.

**Merged Synchronization**  
Merged synchronizations work by copying the files from the source location to the destination. It is very similar to the copy synchronization except it will not delete files from the destination. It will over write a file if it exists already but it will not delete files that exist only in the destination.

**Copy Synchronization**  
Copy synchronizations works by copying the files from the source location to the destination. Any file in the destination that is not in the source is deleted. As a result, **Caution** should be used when specifying a destination. Choosing a folder that has other important files in it that are not in the source folder, will result in those items being deleted. Its best to choose an empty folder for the destination.

### Two Way Synchronization

Two way syncs work by getting a list of files from the source and a list of files from the destination. Any files that are missing from either direction are then copied to the other side where the file is missing. Any files who's modified date doesn't match the other side will have the newer file copied as well.

Once the very first synchronization has been run, it then saves a snapshot of what the source looks like, and what the destination looks like. This is done so that the next time you run the synchronization, CrushSync can determine what files have been deleted since the last time the sync was run. If a file exist in the previous snapshot for either the source or destination, and no longer exists in the real listing of files, then that item must have been deleted. So the same action is taken on the opposite side to keep things synchronized.

### Merged Synchronization

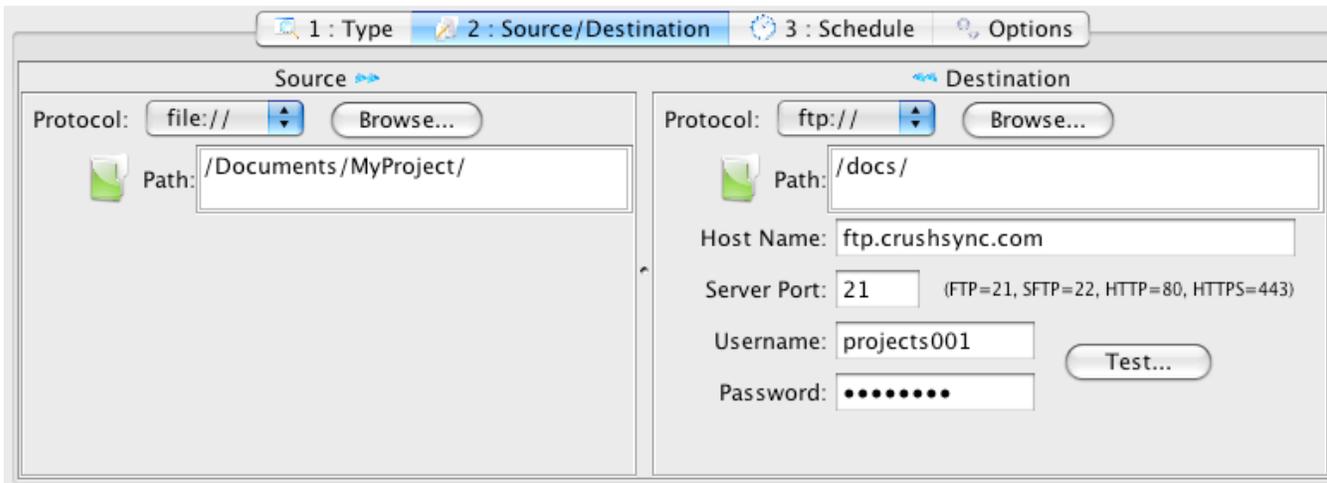
A merged sync is basically a one way synchronization. Its the most common type of sync used for most situations. A merged sync only copies files from the source to the destination, it never deletes files. It will not copy a file if the file already exists and has the same modification date. It will copy the newer file if there are two files with different modified dates.

### Copy Synchronization

A copy sync can be potentially very dangerous if used incorrectly. Caution should be used! A copy sync works just like a merged sync, except it also will delete any file it finds on the destination that does not exist on the source. So if you accidentally pointed a copy sync to a location that was the wrong place, and it didn't have the same files as the source, each and every item would be deleted. You could delete all your files on the destination by configuring this incorrectly. The purpose of the copy sync is to ensure that the destination only contains an exact copy of what the source has. That way if you delete something on the source, the copy sync will then remove it from the destination too.

## Source / Destination

Each sync has a source and destination that you configure. The source or destination can point to local files on your hard drive, or a network server, as well as files located on a remote internet server using the FTP, FTPS, SFTP, HTTP, or HTTPS protocols. This is part of the real power of CrushSync. It can work with files on a remote server just like they are local files on our machine. Further, when connecting to a CrushFTP based server, there are many more enhancements that CrushSync can do as CrushFTP implements many different advanced options.



Both the left and right side allow you to drag and drop a file from your local Finder / Explorer windows and drop it on CrushSync. This may be easier than using the browse functionality. The path for a source and destination should always start and end with a "/" (forward slash.)

#### FILE://

The file protocol is simply referring to local files on your hard drive, or a connected server. If dealing with a file server, remember that the file server needs to be connected and available when CrushSync runs the sync. There are some tricks you can do with this is you see below under the Options tab.

(On windows, the file paths will have forward leaning slashes instead of back slashes. Either is acceptable, CrushSync will just convert them to the standard URL notation using forward slashes.)

#### Internet Protocols

You configure the server's IP, username, and password of the server you are connecting to. You can click the "Test..." button to verify your settings. If you are unsure about the path, start by setting it to a "/". Then click the browse button to get a folder list that you can double click on to move around on the server to locate the source or destination where you want the sync to transfer files to or from.



#### FTP://

The ftp protocol is the most common server protocol used for transferring files. CrushSync should work with most FTP servers that support some basic standard FTP commands, in general this means the "MDTM" command. If syncing with a CrushFTP server, CrushSync can ask CrushFTP for a file listing of all files, and their modified dates in one quick command. When working with other servers, CrushSync will go through each directory one at a time building a the file listing of everything. This takes more time, but it will still work.

When working with a CrushFTP server, CrushSync also uses compression when transferring files and directory listings allowing for faster than normal speeds.

#### FTPS://

The FTPS protocol is normal FTP, but uses encryption. Other than the encryption, its the exact same as FTP. Many FTP servers support FTPS including CrushFTP.

#### SFTP://

The SFTP protocol is another secure protocol. Although the acronym is similar to FTP, the protocol is not based on FTP in any way. SFTP works well when dealing with firewalls as at times, firewalls are not configured properly for FTP traffic. SFTP however suffers from speed issues due to the nature of the protocol. So if you could potentially get speeds faster than roughly two to three megabyte per second, then you may want to consider FTPS. SFTPs max speed is roughly two to three megabyte per second.

#### HTTP://

The HTTP protocol is using WebDAV traffic to transmit files, and get file listings. So technically, the HTTP server must have WebDAV enabled in order for CrushSync to be able to deal with it. Specifically this works well when dealing with a CrushFTP HTTP WebDAV server. HTTP is very good when dealing with firewalls as it only uses a single port. So in that respect is better than FTP. It has the same performance as FTP. The one

issue HTTP can suffer from is when dealing with many thousands of files. Say 50,000 small files. HTTP is using XML messages, and unfortunately sometimes these messages can be bigger than the actual file being transferred. HTTP is faster than normal FTP. However, the CrushFTP optimized FTP or FTPS methods are much faster than all other protocols. The HTTP protocol will likely only work with a CrushFTP HTTP server as most, if not all HTTP WebDAV server don't have the ability to set modified dates on uploaded files in a standard way.

### HTTPS://

The HTTPS protocol is identical to the HTTP protocol except it uses encryption. Other than this, it is identical.

## Schedules

Schedules allow for unattended syncs to run in the background. You can configure syncs to run as quickly as every minute. You can set the schedule to run hourly, daily, weekly, or monthly as well. The daily, weekly, and monthly schedules allow you to set the different hours and minutes of the day that the sync should run as well. So you can have a sync daily sync that runs multiple times per hour for instance.

The screenshot shows the 'Schedule' configuration window for a sync named 'test'. The window has four tabs: '1 : Type', '2 : Source/Destination', '3 : Schedule' (which is active), and 'Options'. In the 'Schedule' tab, there is a list of schedules with 'test' selected. To the right of the list are 'Add New Schedule' and 'Delete Schedule' buttons. Below the list, there are radio buttons for different frequency options: 'Minutely Every 1 minutes', 'Hourly Every 1 hours', 'Daily Every 1 days' (which is selected), 'Weekly Every 1 week(s) on:', and 'Monthly Every 1 months(s) on:'. A 'Scheduled to next run at:' field shows '04/10/2009 08:05:28 AM' and a 'Test' button. There is also a field for 'Hour and Minutes (HH:mm):' with the value '8:05,9:00,13:23,22:45'. Below these are checkboxes for 'Email Sync Log' (unchecked) and 'Email Sync Failures' (checked). At the bottom, there are input fields for 'To:' (ben@crushsync.com), 'From:' (ben@crushsync.com), 'CC:', 'BCC:', and 'Subject:' (Sync Failure).

Schedules also allow for notifications since you won't be there watching the sync run. You can configure an email alert to be sent to you anytime that a sync fails for instance. You can also set the sync to include a copy of the log when it sends the email.

## Options

There are many complex options that you can setup CrushSync to do. When in doubt, leave the options at their defaults. The defaults have been configured to settings that should work in nearly all situations. Configuring options incorrectly can make your sync not run at all potentially.

### Filters

Filters allow very specific actions to be taken on whether or not to include a folder or file as part of a synchronization.

Filters Comparison Preview Performance Archive Scripts

Filters to Apply

test filter

Add

Remove

Move Up

Move D...

Filter what items?  Files  Folders

Do what with them?  Include  Exclude

Look at what part?  Item's Path  Item's Name

Criteria: \*.jpg  RegEx

If you add a filter that is of the "Include" type, then no files will be processed unless it matches the filter you have set. You can have multiple include filters allowing for various filenames as long as the filename matches at least one of the "Include" filters.

If you have only exclude filters, then all files will be accepted unless they match one of the exclude filters in which case it will be skipped. You can have multiple "Exclude" filters as well.

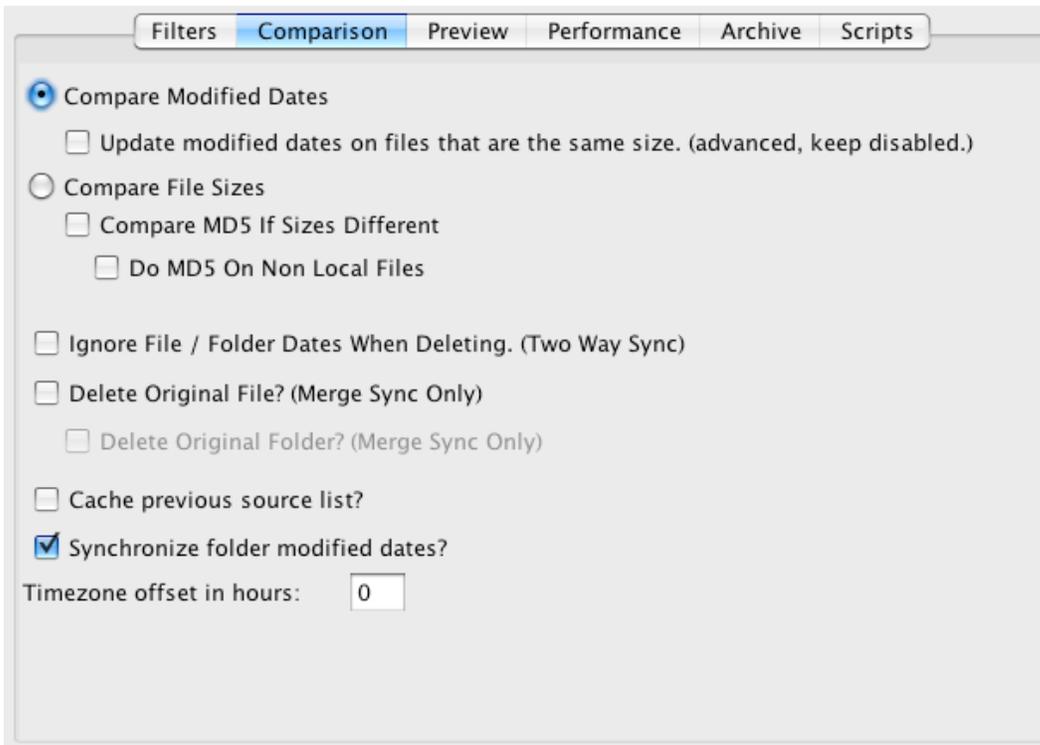
The Files / Folders option allows you to identify the type of item the filter can apply to. The Path / Name option allows you to specify what the filter applies to. Either the path to the file or the actual items name.

The criteria accepts wild cards of "?" and "\*". The question mark means any single character, while the asterisk means any continuous set of characters.

For example, "\*.jpg" applied to "Item's Name" will match on all items that end with a period and JPG. Filters are case insensitive. Another example "\*.??g" would match files such as "test.jpg", "test.png", "test.frg" but not "test.pg" or "test.ppag" as those two do not have a period, followed by any two characters, and then a letter G.

RegEx is an advanced option. If you know what RegEx is, use it...if you don't, then you really don't need it.

## Comparison



#### Modified Dates:

There are several comparison options that can affect the logic CrushSync uses when comparing files. The default, and most common comparison method is by looking at the modified date of the file. The newest file is the one always considered to be the master file.

There is a checkbox to update modified dates on files whose modified date doesn't match, but the file size does. This option should only be one time, and only on the first time when syncing files that you know are the same, but their dates may be off due to a mismatch on system clocks or such. Using this at any other time is a bad idea, and could make you think files are up to date when in fact they are not. The perfect scenario for when this option is appropriate is as follows: You have a very large set of files on an office computer. You copy the files to a portable drive, and move them to your server when you copy them off. Then you run CrushSync and point it at this new location where you have copied the files. Due to the process of moving the files across multiple systems, the time zone, or possibly the modified dates on the files might be out of sync. Using this option the first time will put things back in sync, then this option should be turned off, and left off.

#### File Sizes:

A file size comparison can also be done. The larger of the two files is the file that is always considered to be the master file and is copied over the other file. This is not appropriate in many situations. The MD5 option will check the contents of files who are the same size, and verify that they are in fact identical. If they are not found to be identical, then the newer of the two files that are the same size is considered to be the master file.

#### More Comparison Options:

The "Two Way" sync is very powerful, but you may not care about the folder modified dates when its determining what files to delete. Some operating systems don't handle updating folder modified dates very consistently, and as such a folder may always appear to have changed when it has not. Enabling this option allows you to allow a folder to be deleted even though the OS says its date changed.

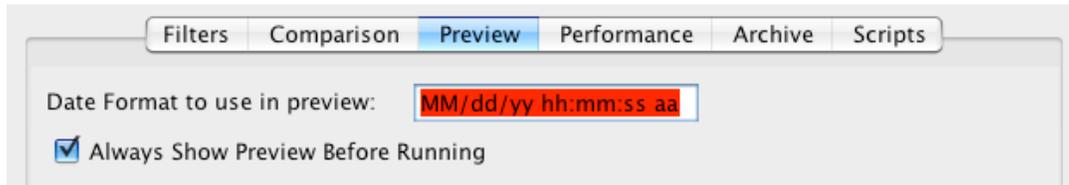
The "Delete Original File" option allows a "Merge Sync" to essentially become a "Move Sync". Anytime a file is found that is missing on the destinations, the file is first copied tot he destination, then deleted from the source. This allows you to make a sort of drop folder where you place files, and they are then transfered to the other location for you automatically. You can also choose if you want to keep the folder structure intact and only delete files after transferring them, or also remove the folders once the files have been transferred.

The "Cache previous source list" option allows for faster syncs when dealing with a slow remote server. What it does is after the first sync, and each sync there after, it keeps a copy of what files it had the last time. Then the next time the sync is run, rather than comparing what the destination has with the source, the source is just compared with the last known version of the source. This type of customizations can have some great performance benefits. Its especially beneficial when dealing with syncs for the purpose of backing files up...where you know you are the only one putting files there, so you can always trust that you know what files were last transferred. It also allows for the ability to use the "Archive" feature of CrushSync. Essentially this means CrushSync will upload a new sip file each time the sync is run containing only the few files that changed since the last time the sync was run. You could for instance then have a daily history of zip files containing only the files that had changed for that particular day. A great archiving tool.

The "Sync folder modified dates" option allows you to disable the updating of folder dates. Many times folder dates really don't matter, or are not even accurate as tot he real status of when files have been modified in those folders. Turning this option off allows you to skip this step during synchronization and get a decrease in the amount of time a sync takes to run.

The "timezone" setting allows you to specify a number of hours to add or subtract to the date when setting the modified time of a file. This is particularly useful when working across time zones and various servers.

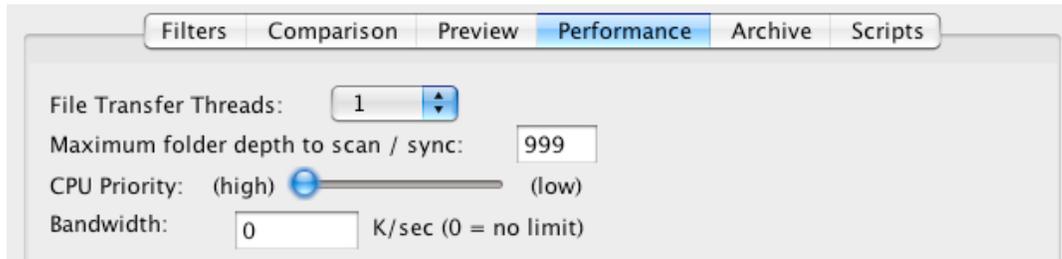
## Preview



The screenshot shows the 'Preview' tab selected in a menu bar with options: Filters, Comparison, Preview, Performance, Archive, and Scripts. Below the menu, there is a text input field for 'Date Format to use in preview:' containing the text 'MM/dd/yy hh:mm:ss aa'. Below that is a checked checkbox labeled 'Always Show Preview Before Running'.

This tab allows you to set the date format used when displaying the preview list of what files are going to be affected by a sync via the "Preview" button. You can also enable the "Always Show Preview Before Running" option for more sensitive syncs that you want to force a preview on. This option is only ignored when the sync is scheduled instead of being manually run.

## Performance



The screenshot shows the 'Performance' tab selected in a menu bar with options: Filters, Comparison, Preview, Performance, Archive, and Scripts. Below the menu, there are four settings: 'File Transfer Threads:' with a dropdown menu set to '1'; 'Maximum folder depth to scan / sync:' with a text input field set to '999'; 'CPU Priority:' with a slider between '(high)' and '(low)' positioned towards the high end; and 'Bandwidth:' with a text input field set to '0' followed by the text 'K/sec (0 = no limit)'.

These options affect how fast syncs can operate. The default settings are in general the most appropriate, but you can tweak these as desired. Just be sure to test before and after you make a change.

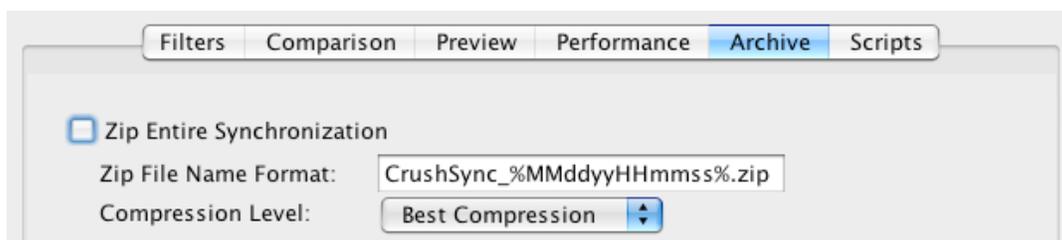
The "File Transfer Threads" options allows you to control the number of simultaneous workers (threads) that are processing files when a sync is running. The sync always uses two threads when preparing to run the sync. One for getting the source list, and one for getting the destination list. Then, once the file transfer starts for transferring the files, multiple threads can be used. The benefit for this can be on machines with multiple processors or cores and transferring files locally. The other benefit can come when doing internet server transfers and dealing with many small files. For example, much of the FTP protocol time is spent preparing to start the transfer, and setting the modified date of a file after the transfer is finished. Using multiple threads allows this to be more efficient and the sync will take less time. However if most of your files are large files, instead of thousands of small files, multiple threads won't really provide any benefit to you.

The "maximum folder depth" allows you to configure how deep CrushSync will scan into folders to find files. There may be special cases where you only want to go one or two levels deep. In general though, leaving it at 999 is the best option.

The "CPU Priority" helps you configure how aggressive CrushSync is when its building the file listings, and transferring files. The slider is very heavily weighted to the left side. Every pixel that you move the slider to the right is a millisecond that CrushSync will pause between operations. So even sliding it 10 pixels to the right could mean a huge performance change in how long a sync takes to run. For instance, 10ms, means at maximum, CrushSync could only build a file listing of 100 files per second. If you had 50,000 files to analyze, that is going to take quite a while to list the files. On the other hand if you leave it at the far left with its current zero value, it can build a listing of a thousand or more files per second. There are other factors that come into play on how fast the listing can be built such as if its a local hard drive, or a remote internet server.

The "Bandwidth" option allows you to set a maximum speed at which transfers will run. This can be useful if the sync is running during the day when other things may be using your internet connection or network as well. Setting a limit will keep CrushSync restricted to the speed you set.

## Archive

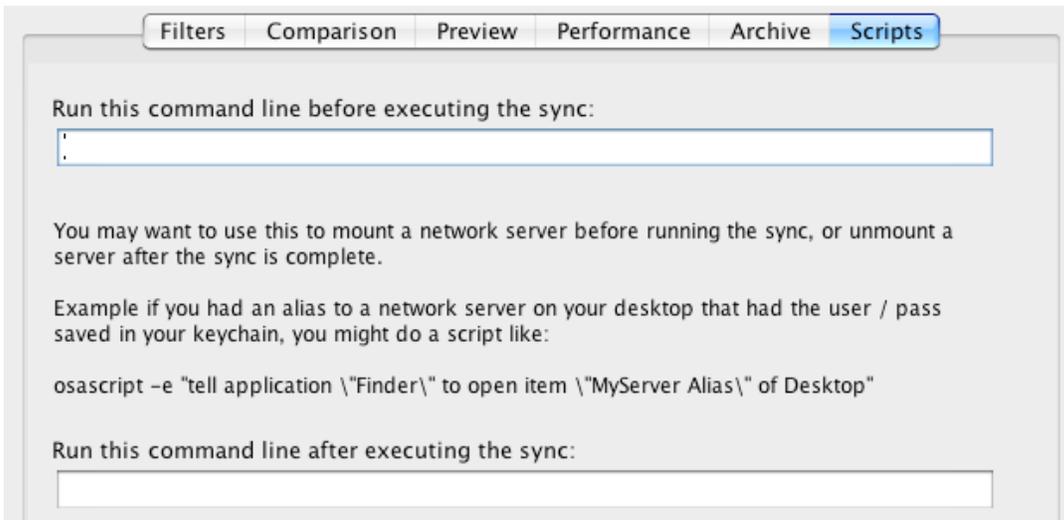


The screenshot shows the 'Archive' tab selected in a menu bar with options: Filters, Comparison, Preview, Performance, Archive, and Scripts. Below the menu, there is an unchecked checkbox labeled 'Zip Entire Synchronization'. Below that is a text input field for 'Zip File Name Format:' containing the text 'CrushSync\_%MMddyHHmms%.zip'. Below that is a dropdown menu for 'Compression Level:' set to 'Best Compression'.

The zip option will take the entire list of source files that need to be transferred, and send them as one single zip file to the server. If the server is a CrushFTP server, this file may then be automatically expanded in CrushFTP and written out to the folder it was intended for. This can allow for a very fast single file transfer of many smaller files. Additionally you may want to use the "cache previous source list" option mentioned above to allow for an archived history of file changes.

You can set the filename and format to the filename of the zip file that is uploaded. You can also specify the compression level used for the zip file. On a local high speed network, you may not want to bother with compression as that may slow down the actual file transfer.

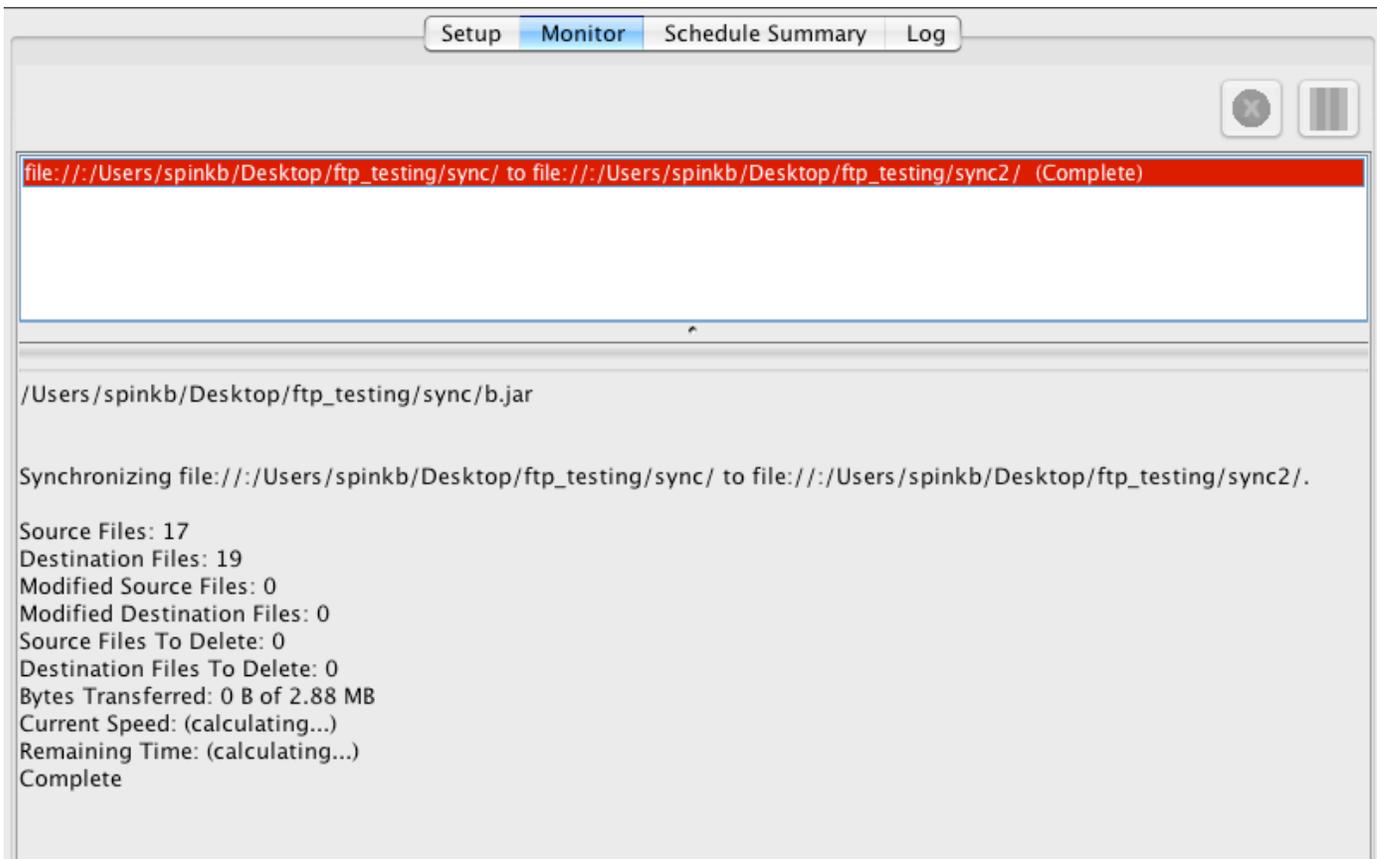
## Scripts



Scripts allow you to execute a command line task before and after a sync is run. This can be very useful in the case when backing up to a network server which may or may not be already connected. Running a script before the sync allows you the chance to connect to the network server so that it is available to the sync when it runs. Additionally, you may want to run some post operations as well, and there is the option to do that too.

## Monitor Tab

This tab shows the current syncs that are running and their progress. If there is only one sync in the list, it will automatically be selected for you. Below the list will be a progress bar that shows the completion percentage of the running sync. The progress bar is measuring the bytes transferred compared to the total bytes to transfer. Depending on the number of threads you have configured for the sync, you may see several lines indicating the actual file each thread is working on.



The current speed, and calculated time remaining will show an estimate of how much longer CrushSync thinks the synchronizations is going to take. Once complete, the status will indicate the sync is complete. It will then remain in the list for another 30 seconds for you to view it before it will disappear. Then the log tab is the only place to look at the status and history of a sync.

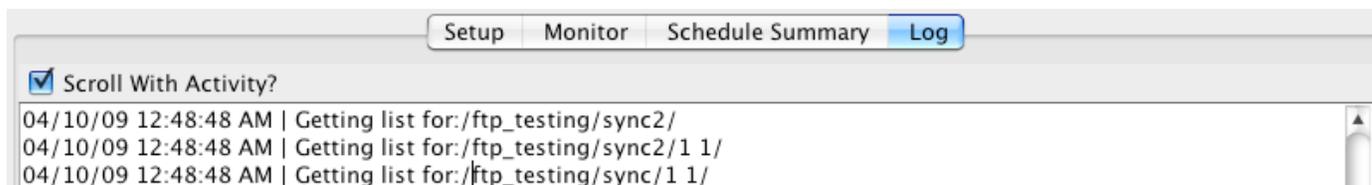
## Schedule Summary Tab

This tab gives a quick over view of all the syncs you have, and their schedules. You can see at a glance when the next scheduled time is that a particular sync will be running.

Sync	Schedule	Next Run
fullrights sync -> sync2	No Schedules Configured	
Sync Big	No Schedules Configured	
Real Syncs	No Schedules Configured	
Backup Documents	No Schedules Configured	
CrushFTP Source Code - Production!	No Schedules Configured	
Work Stuff	No Schedules Configured	
Keep Project Synced on Server		
test		Fri Apr 10 08:05:00 CDT 2009

## Log Tab

This tab shows a live log of the activity CrushSync is performing. Changing the debug level in the preferences will affect the amount of information being logged. If the "Scroll With Activity" checkbox is selected, the log will scroll as new lines are added.



## Preferences

### Settings

The "Start CrushSync GUI When I Login" option is available for OS X users. Enabling this option will add CrushSync to your system preferences, accounts, login items tab. This is different than the Startup Service / Daemon option...it will start even before your user has logged into the computer.

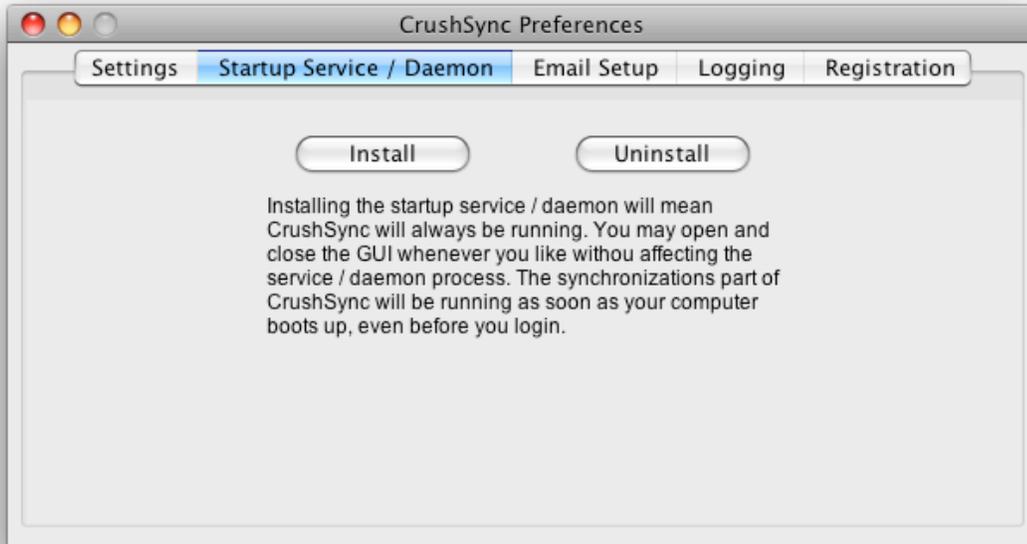


The "Check for Updates at Startup" option connects to the CrushSync website and gets the "version.html" file to see if your copy of CrushSync is the current version. If it is not, a window is displayed containing the version.html page showing you the new changes that are available.

### Startup Service / Daemon

The install button will install a link to start CrushSync by the OS even before a user has logged into the computer. This also means the GUI to CrushSync doesn't need to be open, and CrushSync is still running in the background processing any scheduled syncs that you may have configured. Daemon mode is the recommended way to run CrushSync. Then, when you open the GUI, it just connects to the daemon to show

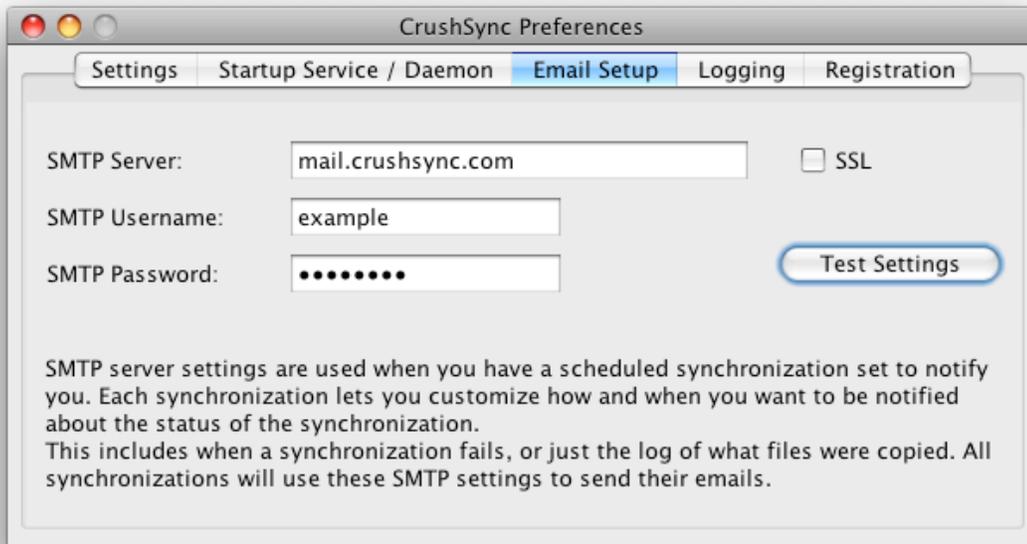
you what the daemon is doing and allow you to control it. When daemon mode is enabled, the title bar of CrushSync will indicate that it is connected to the daemon.



When you run CrushSync not in daemon mode, CrushSync is launching a temporary background copy of CrushSync which acts as the daemon when the GUI is running. When the GUI closes, this temp copy is closed. CrushSync uses ports 53818, and 53819 for communication with the daemon on the 127.0.0.1 IP.

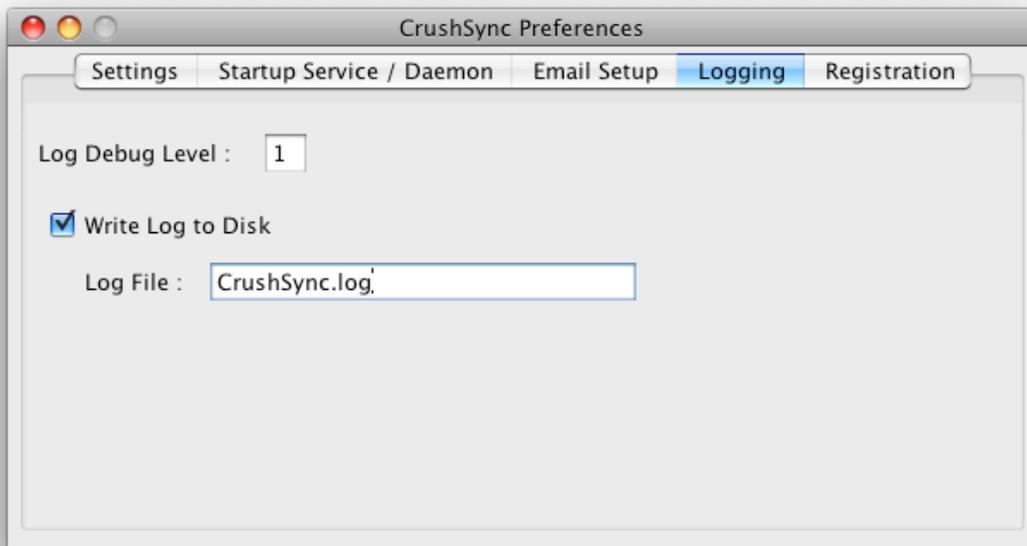
#### Email Setup

This tab allows you to configure your SMTP server's settings for emails that are generated for scheduled syncs. Use the test settings button to verify when your settings are configured correctly.



#### Logging

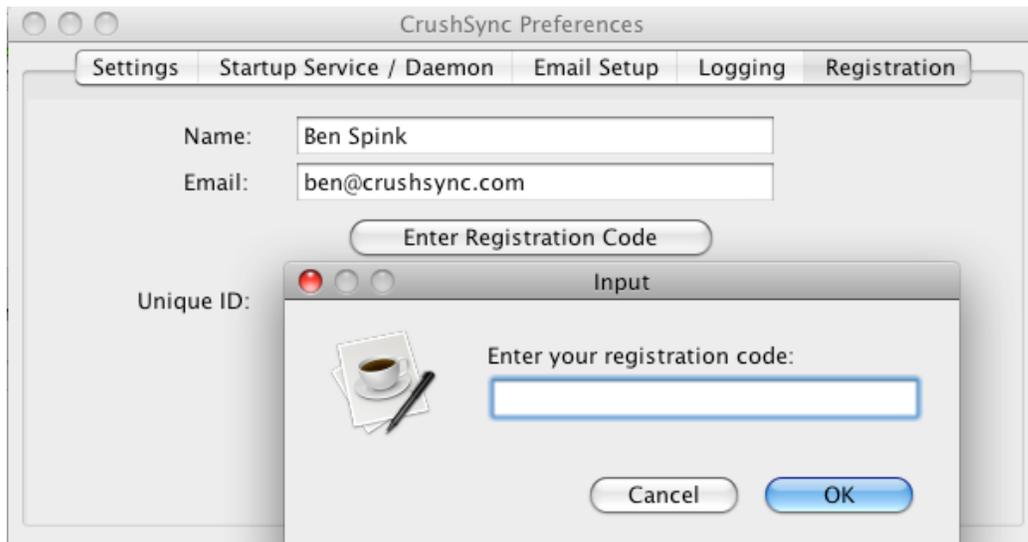
The debug level should generally be set to level 0. If you want to see more detail, try level 1. In general, there is no practical benefit to any level above 2 unless I specifically ask you to try a high level for a debugging purpose. Even then, I generally won't want to see the output of anything higher than level 2.



The "Write Log to Disk" option allows CrushSync to write its log entries to disk as well as the GUI window. If enabled, you can set the log file name to use as well. You could also specify another path to store the log file by putting in a full path to the log file.

### Registration

This tab allows you to enter registration information for CrushSync. Enter your name, and email exactly. Case does not matter on these two items. Then click the button to enter in your registration code. Use control-V or command-V to paste the code to prevent typing errors. CrushSync will reply with a "Thanks!" message when successful.



The unique ID is used when CrushSync is being used in a subscription model. This is where you can purchase several CrushSync licenses, and can enable and disable what machines can use a license. So say you owned ten licenses, if one of the machines you had CrushSync on belonged to a third party customer, you could revoke their right to use CrushSync if their relationship with you changed at a later date. Contact me directly for details if interested in maintaining a subscription model for licenses and potential discounts associated with that.

## Questions?

If you have any questions, don't hesitate to drop me an email. I respond personally to every question.

Ben Spink  
ben@crushsync.com