STRONGBOX
DATA SOLUTIONS

StrongBox Appliance Description

Version 2.0 | June 2017

CONTENTS

Understanding the StrongBox Appliance

The Strongbox® is a self-contained converged file data storage system that is utilized to store data as a NAS for normal file operations as well as being an archive data for long term retention at a cost effect manner for businesses. Data volumes and unstructured data repositories are growing beyond the ability to manage them as well as guarantying the files are backed up and available in case of a disaster. Disaster comes in many forms and it is not always a natural disaster as it could easily be a corrupted file system or a double disk failure.

With the ever growing amounts of data, organizations needs a cost effective solution to store and keep data that is relevant as well as needed for compliance for many years. The perceptible issue for each and every organization is the ability to economically and effectively store the data and be able to retrieve when it is needed in a straightforward fashion. The cost can be seen in a multitude of different views including the process or workflow needed for the data and how it is processed. What maybe a cost conscious decision for some companies may not be for another organization. While the amount of storage and the cost of disk maybe a number one requirement for some customers, workflow maybe a higher decision based on the cost to the business.

The Strongbox is a self-contained converged file data storage system that actively manages the preservation of the data based upon policies that has been set by the company or organization. All aspects of preservation are handled by the system including management of the attached tape library. Unlike traditional backup systems that require administration of a tape library, there is extremely limited interaction between administrator and the inherent tasks being performed by the StrongBox system with the library. The library is managed day to day by the StrongBox storage system and the only active tasks for the administrator would be the initial load of the tape media in to the library as well as removing export copies or adding tapes back in that maybe needed but were removed. Tapes will be formatted by the system when they are needed for sue by the system. This process is an autonomous management of the tape library and eases the burden of management of media.

The cost effective manner that the strongbox utilizes frees up the full time employee to manage the important day to day tasks that keep a business running. The StrongBox decreases the amount of data on high cost storage to the truly active data that needs to reside there and removing the static un-accessed data that is simply taking up space and costing the business money. By removing the un-accessed data from the expensive tier 1 storage, the backup of data nightly is now being reduced and handled in a cost effective manner that is key for preservation requirements.

The ability to determine the amount of copies that a piece of data will have significantly limit the amount of user created copies as well as copies residing on backup tapes and snapshots and will reduce storage footprint and OPEX costs associated with the data. This converged technology will also reduce the cost of the horizontal scale out that happens intrinsically with backup

infrastructures and reduce the complexity and administration of the company's backup schema.

Overview of StrongBox

StrongBox is a network-attached storage (NAS) appliance that provides fast, cost-effective, long-term archiving of unstructured data and fixed content. StrongBox relies on disk for fast file storage and retrieval, and it relies on physical tape storage for cost-effective, long-term storage. You can leverage the capacity savings of tape storage without the complexities of the tape interface and management. It can store and archive data that is growing and unstructured, such as video tapes and movies, legal and medical records, telephone conversations, or cloud applications, to meet long-term data preservation needs or compliance requirements. StrongBox presents itself as a network-attached storage (NAS) target.

A typical use case for StrongBox includes media content from a large broadcasting company. Once aired, the broadcast content may not need to be accessed for a long time, but it needs to be securely archived for retention requirements or other production purposes. StrongBox provides the ideal destination for this type of fixed content, while enabling a simple, file-based way to archive or access data. When a user needs to access a file, they can simply click on the file from the workstation. The complexity of tape is alleviated as the end user experiences seamless file access.

StrongBox leverages the Disk/Tape File System (DTFS) to expose its file system as network shares, allowing end users to mount a StrongBox shares from any Common Internet File System (CIFS) or Network File System (NFSv3) client. An end user can also enable the S3 interface provided by StrongBox for S3-compatible applications to access an object share via the standard RESTful S3 interface.

After the shares are mounted on the client workstations or accessed using S3 clients, end users or applications can simply copy files to the share or retrieve files from the share as needed. Files are initially stored on the StrongBox appliance's disk for fast file storage and retrieval, and then they can be transferred to tape in the attached tape library for cost-effective long-term capacity storage.

Files written to StrongBox are migrated to tape using Linear Tape File System (LTFS) technology. LTFS enables tape to be mounted as a file system independent of the underlying operating system. This relieves tape of its age-long dependency on backup applications and other proprietary software to store and retrieve data. Because StrongBox is a file-based online archive system, data is continuously available for reading and retrieval in an open, non-proprietary and fully portable format.

Key Features

The following features differentiate StrongBox from other data-vault applications:

1. **Data Performance and Policy-based Management**

    a. Built-in policy-based management

        i. The Delayed Action policy determines the period of time after which policies are evaluated for the share, including Library and Export Copies policies, which instruct StrongBox to write files to tape. This is referred to as the "delayed action period". The Delayed Action policy value must be between 10 minutes and one year (365 days).
        ii. The file exclusion policy allows specific files to reside only on disk for fast retrieval and application specific modifications. This policy will also prevent archival of these files to tape.
        iii. The compression policy compresses data to maximize storage capacity (on tape).
        iv. The Dual-copy, Export, and Tape Regeneration policies determine how and when the StrongBox will create copies of tape for local access, data protection, and distribution as needed by the organization.

    b. Performance optimization and tuning

        i. The read cache retains whole files that are smaller in size than what the read cache setting is configured. Files that are larger than the setting are truncated after being moved to disk. A portion of the file remains on disk for initial reads until the tape is loaded, which prevents application timeouts. The stub or truncated size of the file remaining on disk is a configurable setting.
        ii. The file cache retains the most recently accessed files, retrieved from a tape, on disk for period of time to enhance reread performance.
        iii. The file retention policy enables a StrongBox to retain data in file (disk) cache for repeated access.

    c. The Application Programming Interface (API) enables a way for a program to be coded to access to data stored on the appliance, including prefetching files back to disk cache in advance and the ability to provide unique identifier (UID) referencing for any and all file assets associated with a file set.

2. Data Protection and Security

a. The dual-copy (library) policy stores two copies of data with the primary copy stored on one tape and the second copy created on a different tape within the tape library for data protection and redundancy.
b. The export policy enables writing additional copies of the data files to tape for transport, remote storage, or for sharing and distribution of the files among a remote workgroup.
c. The S3 integration enables cloud-based tape storage, allowing end users or applications to communicate through StrongBox's implementation of the standard, RESTful S3 interface.
d. Supports WORM-like storage behavior to prevent deletions, which is necessary for long-term data preservation and regulatory compliance.
e. The StrongBox can also optionally authenticate CIFS shares through Active Directory integration or local system user authentication allowing for an efficient mechanism to safeguard data with standard authentication methodologies.
f. An aliased IP enables a Strongbox to isolate access to network shares by tying an IP address to a share for complete data confidentiality.
g. Tape regeneration provides the ability to create a "tape copy" after initial execution of policies may have occurred. An exact copy of the tape is created based on "live" files (those that have not been deleted from a share). This feature provides further data integrity, data protection, efficiency, and portability when managing tapes in the library.

3. Data Integrity

a. "Fingerprinting" generates unique fingerprints by using SHA512 hash for every file, and then StrongBox verifies the hash on a read request to denote discrepancies between data received and recalled from the archive.
b. No-deletion policy prevents end users or applications from deleting data on shares, thereby meeting data preservation rules.

4. Reliability

a. Automatic database snapshots, database change logs, and configuration backup to help rebuild the system from an appliance failure.

5. Monitoring and Reporting

a. Provides real-time monitoring of tape drives, media, and library.
b. Provides alerts for system, drive, and media errors.
c. Provides Tape Archive Optimization reporting, such as simultaneous drive usage and drive I/O performance.
d. Provides the ability to compare data points; and to identify problem areas utilizing the dynamic, real-time reporting to dynamically view the data collected from the library, drives, and tapes.
e. The StrongBox allows a user to export the Database CSV reporting metrics to a comma-separated value (CSV) file for easy review and manipulation for detailed reports.

6. Scalability

a. Provides seamless capacity expansion by adding disk volumes, tape slots, drives, and media to meet growing archiving needs of businesses. No-cost performance tuning with the addition of tape drives and disk volumes.
b. The feature of "Vaulting" is also available and, if licensed, allows capacity to be excluded from managed capacity calculations if the data is stored on tapes that are not present in the library.

7. Replication

a. Replication enables a source StrongBox system to replicate files (copied to shares) to another StrongBox appliance. If desired, you can configure a StrongBox to replicate data to another system (only) and eliminate the use of tape at the first StrongBox applaince, thereby enabling StrongBox to act as a gateway.

8. High Availability

a. Redundant, hot swappable power supplies.
b. RAID 6+ spare for all data storage (on T20 and T30 appliances).
c. Expedited backup and redundancy of files achieved through seed loading of shares.

DTFS — Moving Data to and from Disk

The DTFS presents a file system from which shares are exposed to the network. DTFS creates a mapping from the presented file system to the file system(s) it uses for back-end storage. Files that are written to the StrongBox appliance can be written to tape after a defined period of time; this is referred to as the “delayed action period" and is configured by the Delayed Action policy.

A hash code is generated for every file when written to tape. When a file is retrieved from DTFS, the hash code is verified. If verification fails, an error is returned to the client and an alert is generated. This hash code verification detects any change to the file content after delivery to the system. This might occur if, for instance, a tape was removed from the system and the content of its files were changed, and then the tape was re-inserted into the system.

DTFS handles the migration of files from disk to tape using LTFS for tape access. When files are migrated to tape, they are migrated as non-fragmented files; individual files do not span multiple tapes. DTFS maintains mapping information to correlate disk files with their corresponding tape files. Once a file has been migrated to tape, file data is removed from disk except for a small amount of data that remains in stub files (as specified by the Read Cache policy for the share). DTFS maintains accurate file size information on migrated files to provide information upon request.

To streamline reads and writes, DTFS performs the following automatically:

- Optimizes disk-to-tape operations to ensure that the tape drive remains busy and to avoid “shoeshining” (when the transfer rate is low enough that it allows the data buffer to empty, causing the drive mechanism to stop, rewind, and wait for additional data)
- Tracks available space on all tapes and uses this information when selecting the destination for a file migration. The tape with the most available space is selected over a tape with less space
- Minimizes tape movements when selecting a destination tape for a file migration
- Distributes work to the available tape drives appropriately
- Prioritizes tape read operations over tape write operations.

LTFS — Moving Data to and from Tape

The primary goal of StrongBox is to move data to tape. StrongBox leverages LTFS technology to move data from disk to tape while preserving original file formats, which enables the StrongBox to quickly access data and provide independence from software and hardware dependencies and application constraints. LTFS was developed by IBM and HP to address tape archive requirements. Files can be written directly to tape by StrongBox and read independently of operating systems or third-party applications. Tapes written using LTFS can be used independently of any external database or storage system, allowing direct access to file content data and file metadata. Data can then be presented in a standard file-system view, which means files can be accessed as they would on other forms of storage media, such as disk or removable flash drives. By using LTFS and keeping all files in the formats in which they were written, StrongBox offers a non-proprietary archiving data storage system.

If you want to use the tape outside of StrongBox, you may need to decode the tape or copy the contents of the tape and decode the paths to restore the reserved characters. This is because StrongBox encodes reserved characters used in file and directory paths when files are moved from disk to tape. These characters are encoded:

- The ~ (tilde) character
- The : (colon) character
- All control characters, which are defined by the locale and include all non-printable characters

These characters are replaced with ~ followed by the hexadecimal value of the character. This is similar to the approach used for URL encoding, except that ~ is used rather than % as the escape character. See http://en.wikipedia.org/wiki/Percent-encoding for more information.

> ***Note*** *Under normal operation, StrongBox tapes are not accessible to end users, and users have access to files stored on disk only (through shares).*

The following are key LTFS benefits leveraged by StrongBox:

- Use of tape media as a disk — When an LTFS tape is mounted, the files and directories stored on tape appear to the client in the same way as those in a disk directory listing. StrongBox leverages this capability to provide faster access to data and enables you to simply drag-and-drop files to and from tape.
- Compatibility across customer environments — Tape media written using LTFS is self-describing, which means that data recovery from tape is independent of hardware or software platforms.
- Data mobility — Content can be easily shared to increase data mobility. LTFS provides the ability to share data across platforms, as with a USB drive or memory stick.
- Cost — Using tape for data storage delivers cost savings on expensive disk storage as well as data center space and power usage. For example, you can store data on tape at a lower cost-location rather than using data center space for archive disk arrays. If you choose to archive data for long periods of time (5+ years) and depending on the frequency of changes and the type of data, the volume of archive data can exceed the amount of production data by a factor of 10 times or more. Therefore, storing all archive data on disk can be unnecessarily expensive; tape provides a low-cost alternative.

Overview of How StrongBox Moves Data to and from Shares

When a file or directory is copied to a share, the following happens in the background:

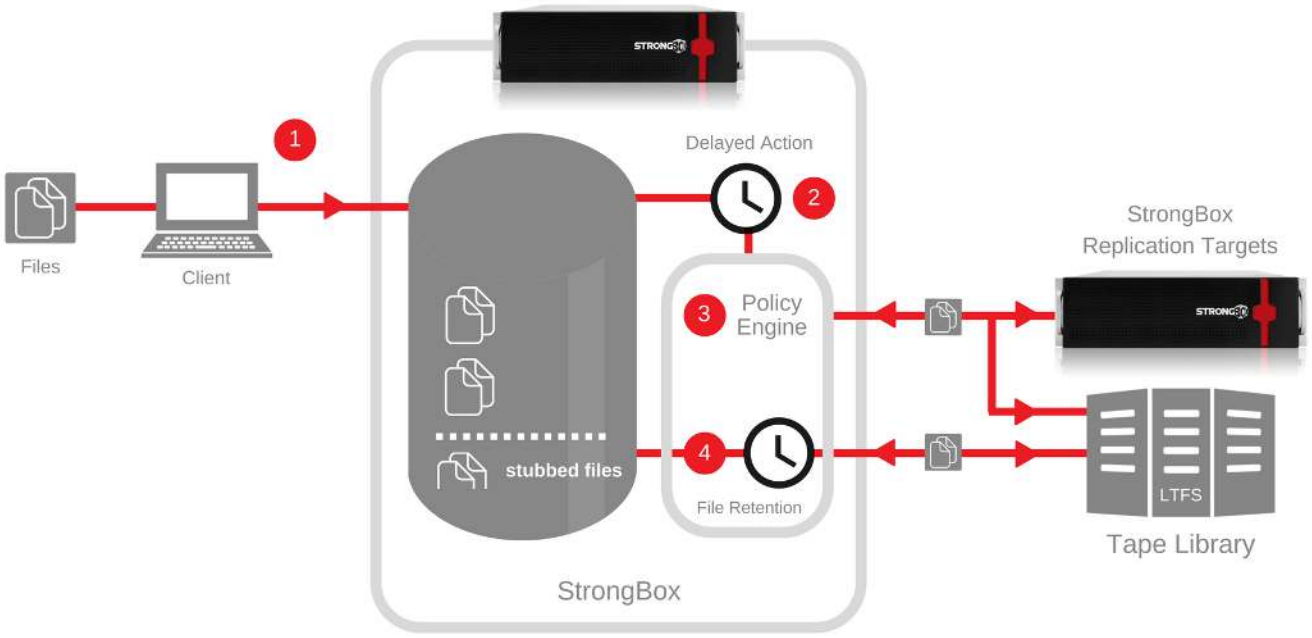

1. Immediately after a file or directory is copied to a share, it is written to disk on the StrongBox appliance. Within five minutes of the last file modification, the delayed action period begins. This is on a file by file basis for the delayed action policy period to start.
2. After the delayed action period expires (as set by the Delayed Action policy for the share), the StrongBox policy engine runs to evaluate all other policies set for the share.

   You can continue to edit files or rename directories after the delayed action period expires, though for directories, you cannot update any attributes such as timestamps, permissions, Windows ACLs, or properties stored as an extended attribute.

3. Based on share or top-level directory policy settings, the following occurs:
   - If File Exclusion is set, the file remains on disk but is never written to tape.
   - If Library Copies policy or Export Copies policy is set for the share, the file is written to tape (up to four, based on the settings for both policies).

     It is important to understand that the StrongBox appliance writes to tape only if 37GB+ of data is in the share OR if one hour has elapsed since the write chain was created. (A write chain is the list of files to be written to tape.) For example, say:

     - the delay action policy is set to five minutes
     - a 50GB file was written to the share at 1:00 PM
     - a 1MB file was written to the share at 1:01 PM

     Example:

     At 1:05 PM, the delay action timer runs and writes the 50GB file to tape. The 1MB file is not written to tape at this time because it has been in the share for four minutes, not five. The delay action timer runs again at 1:10 PM. Because the 1MB file is less than 37GB, a write chain is created with the 1MB file and it is scheduled to be written to tape at 2:10 PM (one hour after the write chain was created).

     - If Compression policy is set for the share, the file is compressed.
     - If Replication policy is set, the file is replicated to the enabled replication target.
     - If File Retention policy is set, the file retention period begins after all writes complete.
     - If Deletion Prevention policy is set, permissions on the file are changed (on StrongBox disk) so that end users cannot delete the file from the share.

4. After the file retention period expires (if set by File Retention policy), the file is truncated to the size set by Read Cache policy. On Windows clients, the file icon changes from this icon  to this icon  to indicate that the file is now offline (no visual indication is given on Mac and Linux clients).

If Deletion Prevention policy is set on the share, the end user cannot delete the file; either an attempt to delete the file generates an error or the file appears to be deleted but is shown again when the file browser is refreshed.

When an end user or application reads a file on a share, the following occurs:

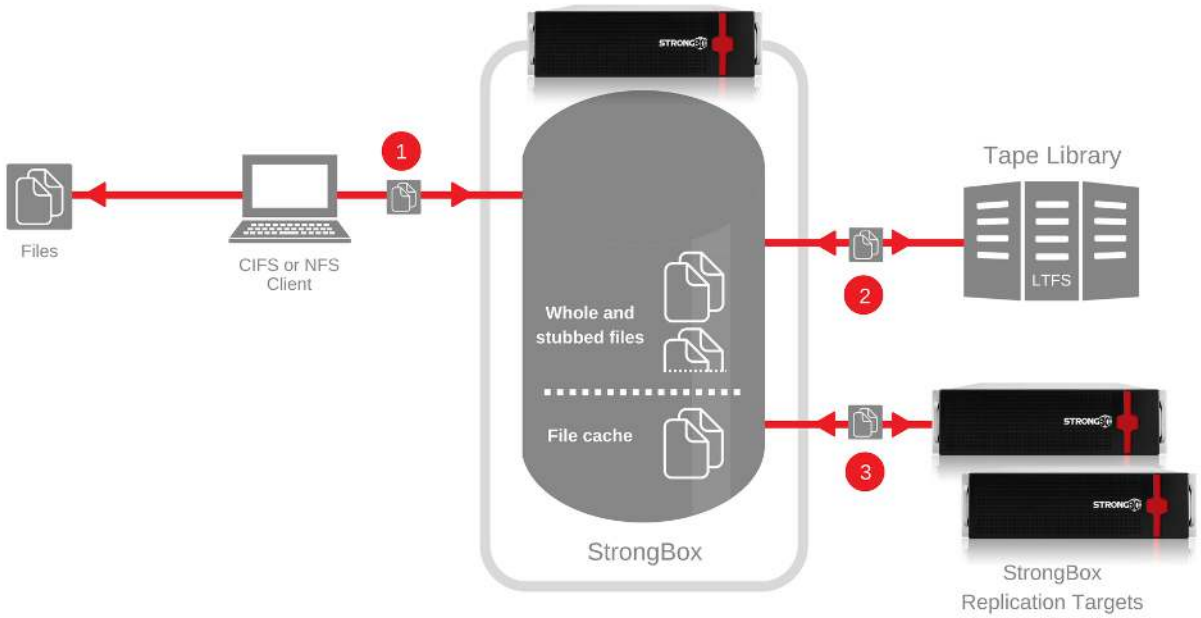

1. Immediately after the end user (or application) opens a file to read or edit it, the following happens:
   - If the file is not stubbed, the file is wholly on disk, and all file data is sent to the client at full speed.
   - If the file is stubbed, StrongBox begins to "trickle" file data from the read cache to the client (according to the Trickle Read Delay setting).

2. If the file is stubbed, StrongBox attempts to locate the file on tape. If a copy of the file is found on tape, StrongBox reads file data from tape onto disk (into the file cache) and sends the file data to the client at full speed.

3. If StrongBox cannot locate a tape containing the file, StrongBox attempts to recall the file from the selected replication target. However, on the client, if the read fails, the end user must attempt to read the file again. This is the same with any file system in which the file is not available. The file is then read from the selected target onto disk. After the entire file is on disk, file data is sent to the client at full speed. Depending on the stub size if the file can alleviate time outs.

Understanding the StrongBox Shares

StrongBox shares are an integral part of the StrongBox and the ability to present a workflow or storage location to end user simply and effectively. The share is a standard CiFS(SMB) or NFS share that allows end users or application to access and write data in a manner that is already prevalent in the enterprise today. The unique features of the StrongBox on each share allows for the preservation of data and reduction of data on disk. These features allow of duplication of the data and being able to create and maintain distribution copies of the data as needed by the organization as well as replication and vaulting for disaster recovery purposes.

Managing Directories and Policies

The Manage Directories and Policies page enables you to perform the following:

- View information about directories on StrongBox that is used to store files or objects copied to shares. Each directory is mapped to a share that is exposed to end users.
- Edit a directory to define policies that govern how data copied to the share is stored on disk and tape. If you edit a directory to add a new policy or modify an existing policy, the policy affects data written to the share after you save the changes to the share. The policy is not retroactive and does not affect files or objects that have already been written to the share unless the files are modified after the new policy changes have occurred.

You can also view or edit the share associated with a directory; when you click a directory's share, the Manage Shares page is displayed in the management UI.

To Set or Edit Policies for a Directory

Here are the policies that may be available for each directory:

1. Tape: Library Copies

For each file written to the directory, this policy will define how many copies are stored on tapes that will remain in each enabled library. If two copies are specified in a library, StrongBox writes a copy of each file to two tapes. Tapes used for library copies are intended to remain in the library at all times as to be available for the StrongBox. After StrongBox copies files to tape, you can identify these tapes by viewing the Policy column in the Manage Tapes page. The policy column will designate the type of tape media it is: Library copy, Export Copy, or Distribution media

Note: *This policy affects how used capacity is calculated and shown in the management GUI.*

When setting the Library Copies policy, be aware of the following:

- The number of library and export tape copies are cumulative when configured for the same share. For example, if you set Library Copies to 2 and Export Copies to 1, three tapes are required for the share. You can specify up to four copies in all libraries.
- To set "zero copy policy", set the Library Copies and Export Copies policies to 0 and then enable the Replication policy. The web interface will not allow you to set the Library and Export Copies policies to 0 without enabling replication.

Note: *If the Library Copies or Export Copies policy is set to 1 or 2 but no library is enabled, files will remain in the Delayed Action state and will not be stubbed.*

2. Tape: Export Copies

For each file written to the directory, the Tape export copy setting defines how many copies are stored on tapes that can be exported from each enabled library. If two copies are specified for a library, StrongBox writes a copy of each file to two tapes. Tapes used for export copies can be ejected from the library at any time. After StrongBox copies files to tape, you can identify export tapes by viewing the Policy column on the Manage Tapes page. Note that this policy affects how used capacity is calculated.

When setting the Library Copies policy, be aware of the following:

- The numbers of Library and export tape copies are cumulative when configured for the same share. For example, if you set Library Copies to 2 and Export Copies to 1, three tapes are required for the share. You can specify up to four copies in all libraries.
- To set "zero copy policy", set the Library Copies and Export Copies policies to 0 and then enable the Replication policy. The web interface will not allow you to set the Library and Export Copies policies to 0 without enabling replication.

Note: *If the Library Copies or Export Copies policy is set to 1 or 2 but no library is enabled, files will remain in the Delayed Action state and will not be stubbed.*

3. Tape: Compression

The compression setting defines whether files written to the directory are compressed when they are stored on tape. Compression is performed by the tape drive, not the StrongBox appliance. The StrongBox appliance will (LTFS) format the tapes according to the policy setting. You may want to disable compression, for example, if the content written to the share is already compressed or if regulatory rules disallow altering data.

It is extremely important to be aware of the following when you edit the Compression policy for a share:

- *If the policy change occurs while StrongBox is writing files to tape, writes will continue according to the original policy setting.*
- *Files that are queued (in the delayed action period) will be written to tape according to the new policy setting.*

4. Tape: File Grouping

The File grouping setting determines the name of the directory that is used when files are written to tape.

- If you choose the Top-level Directory Name, each top-level directory (and its contents) on the share is stored on a tape. If files are stored in the share's root directory, those files are stored in a directory that uses the name (UUID) of the directory that is mapped to the share (the share's directory listed on the Manage Shares page).

  Example

  On the share:
  /share_name/file1
  /file2
  /directory1/...
  /directory2/...

  On tape1:
  /4adc1f83-ab70-4b43/file1
  /file2

  On tape2:
  /directory1/...

  On tape3:
  /directory2/...

Note: *If you select this option, you cannot edit this policy setting (you cannot select Share UUID later) after files are written to tapes for this share. In addition, if you select Top-level Directory Name, you cannot move the share's content between directories once files are copied to tape.*

- If you choose **Share UUID**, files are stored in a directory that uses the name (UUID) of the directory that is mapped to the share (the share's directory listed on the Manage Shares page).

Example

On the share:	On tape1:
/share_name/file1	/4adc1f83-ab70-4b43/file1
/file2	/file2
/directory1/...	/directory1/...
/directory2/...	/directory2/...

If multiple copies of a tape are created based on the Library Copies and Export Copies policies, the directory structure determined by this policy is used on those tapes as well.

It is extremely important to be aware of the following when you edit the Compression policy for a share:

- If the policy change occurs while StrongBox is writing files to tape, writes will continue according to the original policy setting.
- Files that are queued (in the delayed action period) will be written to tape according to the new policy setting.

5. Files: Directory Pre-fetch on Read

This policy determines whether all files in a directory, on a share will be read from tape and restored to disk if a single file in the directory is read/recalled. (This is not recursive; files in subdirectories are not read.) When this policy is set to **off**, only the requested file is read from tape.

Note: If a file was replicated instead of stored on tape, it is pre-fetched from the replication target if this policy is set.

6. Deletion Prevention

This policy determines whether files can be deleted (regardless of permissions set on the share). When set to **off**, end users can delete files from the share.

Note: *When an end user deletes a file, it is deleted immediately from the share. Then, the file is deleted from the StrongBox disk and from local tape, and its uncompressed file size is returned to the free capacity space. Due to how LTFS manages data, when the file is deleted from the tape, it is deleted from the tape index but the space cannot be reclaimed or reused and the data is not over-written or removed.*

Note: *The file is deleted from exported tapes and ejected tapes when they reenter the library. However, the file is not deleted from replicated tapes.*

7. Files: File Exclusion

The File Exclusion policy specifies the files to exclude from being copied to tape; if the specified files are written to the share. The files are stored on disk but never written to tape. For each file type you want to exclude, type a wildcard pattern that represents the file(s) to exclude. In general, it is a best practice to include a file pattern for any file that does not contain user data and does not need to be archived.

To save tape space, you can exclude temporary files by adding these patterns to the policy (type them as shown below due to the patterns being case-sensitive):

- For all client types (to exclude temporary files created by applications):
  - *.tmp
- For Mac clients (to exclude temporary files and files that stores custom attributes, used by Finder):
  - TemporaryItems
  - .TemporaryItems
  - ._.TemporaryItems
- For Windows clients (to exclude the Windows file to governs folder viewing):
  - desktop.ini
- For UNIX clients (to exclude swap files created by UNIX):
  - *.swp

You can test the file exclusion by typing a file name in the second text field and clicking Test. This is especially useful for complex wildcard patterns. After all files have been written to the share, you can verify that the specified files were excluded by reviewing file and tape reports.

8. Files: File Retention

The file retention policy determines when a file on disk is truncated (Stubbed). If a file is larger than the Read Cache policy setting, the file on disk is always truncated to a stub file (size set by the Read Cache policy) and the file is moved to tape. The File Retention value must be between 0 minutes and 10 years.

The file retention period begins after all files in the share have been written to all tapes (Library and Export Copies, across all libraries). After the file retention period expires, the file is truncated.

***Note**: If you choose to leave files on disk for a long period of time, more disk space is consumed over time. Also, note that you cannot alter when files are stubbed after the delayed action period has expired; therefore, be careful when setting this policy value.*

9. Delayed Action

Determines the period of time after which policies are evaluated for the share, including Library and Export Copies policies, which instruct StrongBox to write files to tape. This is referred to as the "delayed action period". The Delayed Action policy value must be between 10 minutes and one year (365 days), though setting this value below one hour is not recommended.

The delayed action timer runs every five minutes, and it checks each file's last modification time. For directories, it checks the creation time. If file modification time or directory creation time is older than the specified delayed action period, policies are then evaluated. You can continue to edit files or rename directories after the delayed action period expires, though for directories, you cannot update any attributes such as timestamps, permissions, Windows ACLs, or properties stored as an extended attribute. (Please refer to the See Overview of how StrongBox moves data to and from shares for more information about how policies are evaluated and how when files are written to tape.)

If StrongBox is ingesting data (writing files from shares to disk) while it is also writing files to tape, different areas of the disk are used simultaneously for reads and writes. This may cause performance issues. Set the Delayed Action policy to allow enough time to ingest data before writing to tape, which depends on the size and number of files being written. You may need to reset this value several times to tune it based on observed write speeds.

10. Read Cache

The read cache specifies the maximum size that each file can grow on disk; stated differently, this determines the size of each stub file. After a file is migrated from disk to tape, file data is removed from disk and only a portion of the file remains on disk. This remaining file is referred to as a "stub file". When a client requests a file, data is delivered from the read cache (stub file) if StrongBox must load the tape that contains the full file. Stub files enable StrongBox to minimize application timeouts and respond immediately to file requests. You can specify up to 1,024GB (1TB), and the default stub file size is 4MB.

When the data set consists of many small files, you may want to increase the read cache so that data is read from disk (instead of StrongBox taking time to mount a tape for a small file). For example, if you have thousands of files that are 3MB, you could increase the read cache size to 4MB so that each time a file is requested (read), StrongBox retrieves it from disk. However, this means more disk space is used, which you can add by connecting external disk arrays.

11. Replication

This policy enables replication for the directory. Any files written to the share are replicated to the selected replication target. You can choose to replicate files in the directory when they are first written to StrongBox, or you can choose to replicate files in the directory after they are first written and any time modifications are made to the files. See Overview of Replication for more information.

If the Library Copies and Export Copies policies are set to zero and Replication policy is enabled, files written to this share are stubbed after they are replicated to the target.

Note: *You must enable a WAN Acceleration license to view and configure this policy.*

Understanding Tape States for a StrongBox Library

You can perform actions based on tape state, though in some cases, no action is required.

STATE	STATE DESCRIPTION	AVAILABLE ACTIONS
CHANGED BARCODE	The tape's barcode has changed.	Eject
CHECKS REQUIRED	In-depth checks are required and will run automatically when the tape is loaded and mounted for read or write. **(You cannot initiate a check, unless you read a file from or write a file to the tape.)** *This is not a "bad" state*. Newly added tapes and tapes discovered during initial startup are placed in this state. For tapes that have been in the library, this state indicates that the last time the tape was mounted, it was usable but StrongBox or the library has been restarted since the last status check. ***If a tape is ejected and then re-introduced to StrongBox, its state remains as Checks Required until the tape is selected for a read or write operation. Then, the tape is loaded into the drive and in-depth checks are run.*** Finally, you cannot copy a tape in this state if the tape is not assigned to a share (see the Share column on the Manage Tapes page).	Eject, Load tape, Mark as bad, Copy, Verify tape media
CLEANING	The tape is used for cleaning only.	Eject
DUPLICATE BARCODE	The tape is labeled with a barcode that StrongBox has seen assigned to another tape.	Eject
FOREIGN STRONGBOX	The tape contains StrongBox metadata that was created on another StrongBox appliance.	Eject, Load tape, Mark as scratch, Verify tape media

FORMAT ERROR	StrongBox could not format the tape (LTFS). This error is normally caused by a write protected tape, a bad tape media, or a generation of tape that is too old to be used.	Eject
GOOD	The tape contains data (written by StrongBox). A good tape can move to Checks Required if the tape has been ejected and then put back into the library. Also, the tape can be exported.	Eject, Copy, Mark as scratch, Verify tape media
INDETERMINATE BARCODE	The barcode cannot be determined from tape vendor or serial number.	Eject
LOAD PENDING	The seed tape will be loaded shortly.	Eject
LOADED	The seed tape was successfully loaded and can no longer be written to (read only).	Eject, Copy, Load tape, Show information, Verify tape media
LOADING	The seed tape is loading now.	--
MOUNT FAILURE	The tape failed to mount in a drive.	Eject
NONSTRONGBOX	The tape is formatted using LTFS but was not written by StrongBox.	Eject, Mark as scratch, Verify tape media
OTHER DATA NONLTFS	The tape contains data but the tape is formatted as something other than LTFS.	Eject, Mark as scratch, Verify tape media

REMOVABLE	The tape failed its check during mount time, and the vendor or serial number on the tape did not match the number in the database.	Eject
REMOVED	The user ejected the tape (using the StrongBox web interface), or it resides in an I/O slot.	--
SCRATCH	The tape is empty and writable. The tape's status is moved to Good after a successful write. The tape's status moves to Checks Required if the tape is ejected and then put back into the library.	Eject, Mark as bad, Verify tape media
TAPE ERRORS	Errors occurred on the tape, such as a hash code failure, and the tape is a read-only. The tape can continue to be used for reads.	Eject, Copy, Verify tape media
UNSUPPORTED TYPE	The tape is not an LTO-5, LTO-6, Jaguar JC, and Jaguar JD tape.	Eject
USER ACTION	The user marked the tape as bad and it is read-only. Or, an LTFS or kernel error occurred when attempting to write to tape.	Eject, Copy, Verify tape media

Browsing Files and Objects

The Browse Files page enables an administrator to:

- view all directories and files copied to network shares
- view all buckets and objects copied to object shares
- view the contents of orphaned directories, which are directories whose shares have been deleted, replicated, or seed-loaded
- view in-depth file or object properties
- view information about each tape to which each file or object was written

To View Files or Objects In a Directory

1. Click **Browse Files** on the System menu, which displays all shares and orphaned directories created on the system. Or, click the name of a share on the Manage Shares page. This displays only the contents of the share.

   The left side of the Browse Files page lists directories; the right side lists directories and files or buckets and objects contained within the selected directory. The following columns are displayed on the right side of the page:

   - **Name** — The name or UUID of the directory. Only orphaned directories are given UUIDs as their names; this is because the mapped share name no longer exists or has been created by replication or seed-loading a tape.
   - **Status** — ***(Detailed on the page 26)*** — The state of the file.
   - **On Disk?** — Whether the file can be read from disk. Values include
     - Yes — The full file resides in the read cache or file cache; note that excluded files can always be read from disk (file state can be Exclusion, Processing Policies, Read Cache, or File Cache)
     - No — The file is stubbed or resides on tape only (file state can be Delayed Action or Stubbed)
   - **Tapes** — The number of tapes to which the file was copied according to policy set on the file's share
   - **Replicated** — Whether the file was replicated to target systems according to Replication policy set on the share.
   - **Size** — The uncompressed size of the file (in bytes).
   - **Date Modified** — The date and time when the file was last accessed for write, even if no writes were actually performed. Be aware that this value is not updated while the file's status is "Delayed Action". The value reflects the creation time of the file until the status changes.

2. To "drill-down" into a share or directory (to view its contents), click the share or directory name on the left side of the page. Repeat this step as necessary to view subdirectories.
3. To view details about a file, click a file name on the right side of the page. The following information is displayed:

   - **Properties:**

     - **Full path** — The full path to the file according to StrongBox. The top level directory name indicates the share name or UUID of the orphaned directory.
     - **Size** — The uncompressed size of the file in bytes, megabytes (MB), gigabytes (GB), or terabytes (TB). As an example, a megabytes equal to 1,048,576, or 1024^2 bytes. In contrast, a megabyte (MB) is equal to 1,000,000, or 1000^2, bytes.
     - **File status** — The state of the file as described above.
     - **Created at** — the date and time when the file was first copied to the share and created on StrongBox.

       Note: *When you copy a file from one location to another (which happens when a file is recalled), the date when the file was created originally becomes the Last modification date and the date when the file was created on the local volume becomes the Created at date.*

     - **Last access** — The date and time when the file was last accessed for read.
     - **Last modification** — The date and time when the file was last accessed for write, even if no writes were actually performed. Be aware that this value is not updated while the file's status is "Delayed Action". The value reflects the creation time of the file until the status changes.
     - **Hash code** — The hash code assigned to the file.
     - **Replicated** — Whether the file was replicated to another StrongBox system.
     - **Available for reading from alternate source (matching hash) on disk** — Whether the file content can be read from a file in a different location (same hash code but different disk location).

   - **Tapes** - For each tape to which the file is copied, the following information is displayed:

     - **Barcode** — The barcode number of the tape.
     - **Vendor** — The manufacturer of the tape.
     - **Serial Number** — The serial number of the tape.
     - **Type** — The type of tape.

The following table depicts the flowchart of the data and what the state of the file means:

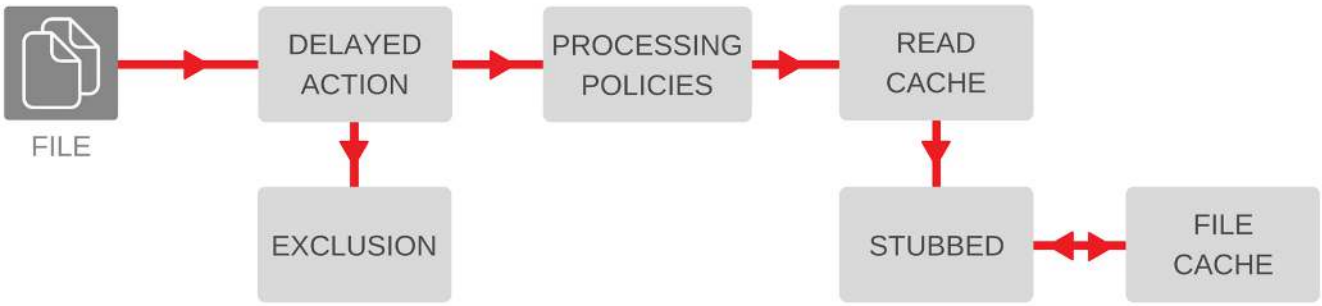

STATE	DESCRIPTION
DELAYED ACTION	The file has been copied to the StrongBox share but the delayed action period has not yet expired (as set by the Delayed Action policy on the share). ***Note:*** If the Library Copies or Export Copies policy is set to 1 or 2 but no library is enabled, files will remain in this state and will not be stubbed.
PROCESSING POLICIES	StrongBox is processing policies applied to the share in which the file resides.
EXCLUSION	The file has been excluded from tape because the file matches a pattern specified by share's File Exclusion policy.
READ CACHE	The entire file resides on disk and was not stubbed; the file was smaller than the share's Read Cache policy setting when it was initially written to StrongBox. ***Note:*** Even if you change the Read Cache policy setting to be smaller than a file in the Read Cache state, the file will not be stubbed. The file is not re-evaluated after it is put into the Read Cache state.
STUBBED	The file has been copied to tape and stubbed according to the Read Cache policy. File stubbing occurs after the retention period expires as specified by the File Retention policy.
FILE CACHE	The file resides in the file cache, which means that it has been read from tape. This file was stubbed and will be restubbed as other files are read from tape and disk space is needed.