



ATTIX⁵
BACKUP PROFESSIONAL
SERVER EDITION

WHITE PAPER v5.1

ATTIX⁵ **BACKUP PROFESSIONAL**

SERVER EDITION V5.1

This document is intended to give a more technical overview to the Attix5 Backup Professional Server Edition Backup Client. For further operational and installation details please consult the Backup Professional Server Edition User Guide.

INSTALLATION

Server Edition (SE) is a JAVA client that is installed as an OS service on the server to be backed up. The installation file is approximately 15MB, depending on the options selected in the Deployment Wizard, and includes the Java Runtime Environment (JRE) from Sun Microsystems. Installation and configuration should take no longer than 15 minutes and does not require an OS reboot. While idle the service uses about 3 MB of memory, and increases this to approximately 25 MB during the backup or recovery process, depending on the amount of data selected. These processes are typically scheduled but can also be initiated manually via a runner in the system tray (Windows version) and via the GUI (Windows, Mac and Linux versions). Server Edition Remote Management from the Storage Platform Console and the Command Line Interface enable you to remotely manage any Server Edition in your organisation. Full installation and operating procedures are available in the User Manual.

DATA SELECTION

A JAVA GUI is loaded in order to select data for backup. Single files, folders, or entire directory structures can be selected for backup either manually or by using customisable inclusion and exclusion filters. These can be based on file or folder name, file types (e.g. .doc) or file modification dates (e.g. no files older than 1/1/02). Files and folders can also be specified in an Exclusions list; these files or folders are to be excluded, no matter where they are located on the drive/volume. Each backup initiates a new data scan to update the backupset with any additions or deletions. The nature of the backup process does require disk workspace equivalent to roughly 50% of the backupset size. See details of this process under **BACKUP PROCESS** below.

APPLICATION PLUG-INS

SE uses Application Program Interface (API) based plug-ins to backup data from common databases and applications. These can be installed together with SE or at a later date and multiple plug-ins can be installed on a single server. Once configured, the plug-in makes an API call to the application to do a full data backup to a specific folder on the server. SE then backs up this folder as part of its normal backup procedure. This is done without shutting down the application or interfering with user access. A VSS Database plug-in is also available for all VSS supported applications that negate the need for the data dump.

OPEN FILE MANAGEMENT

Files locked by the OS or an application are typically skipped during standard backup processes. MS Volume Shadow Copy Services (VSS) that is supported from Windows Server 2003 ensures that locked files are correctly handled during the backup process. Open File Manager from St Bernard Software is also supported for older MS operating systems.

COMMUNICATIONS

SE makes use of the SSL (Secure Socket Layer) protocol standard for secure data transmission (1024 bit RSA key exchange, 128 bit RCA stream cipher and SHA-1 integrity checking). All communications use SSL on port 443/8443.

SCHEDULING

Backups can be scheduled (once a day, multiple times daily down to per minute schedule, once a week, etc.) and can also be initiated via the GUI, system tray icon, command prompt or the Storage Platform Console.

THE BACKUP PROCESS

The Backup Process

Initial Backup: The data selected for backup is compressed in a backup file. A separate index file is also created detailing files and their hashes. The maximum size of this backup file can be limited, in which case the Backup Client will compress until the limit is reached and then flag any further files for backup at the next schedule time. Thus on the first backup only the first gigabyte of data (for example) will be backed up. The next time (which could be an hour or a day later) the first gigabyte will be patched and the next gigabyte will be added, and so on until all files have been backed up and only patches need to be sent. This process can automatically be cycled in Server Edition to speed up and complete the process.

Patching: Subsequent backups use a proprietary patching technique to reduce the size of the backup. This works as follows:

1. SE does a scan of the files and folders and filters selected for backup.
2. New files are compressed and form the new backup file.
3. Deleted or removed files are added to a deletion list.
4. Modified files (we use the modification date and do not reset the archive flag) are first checked against the local cache to see if a previous version of that file is in the cache. If so then the files are compared on a binary level (and thus the process is file type independent) and any changes are stored in a file patch and added to the backup file. Thus the greater the modification date setting of the cache the greater the chances of a previous version of a file being present to allow for the patching process.

Two patching techniques are available in Server Edition, the default Binary Patching and Delta Blocking, which reduces this requirement. The Delta Blocking process is significantly faster than Binary Patching and it does not require a complete copy of a file in the cache to calculate the patch, only a footprint file. The patches are, however, much bigger than Binary patches.

Backup Transfer: An SSL connection is established to the NameServer (NS) which verifies the account and passes the IP address of the FileServer (FS) where that accounts' data is stored. The connection is dropped and a new SSL connection is established to the FS. Files are sent to the FS as soon as they are compressed or patched, where they are stored in a dedicated account directory. File hashes are checked to ensure file integrity. Files are stored using 448-bit Blowfish encryption in CBC mode.

Local SnapShot: Should the initial backup be too large for the available bandwidth, a local backup can be initiated to disk or to a portable FS (or DR Box). The data or DR Box is then physically transported to the remote FS to which the backup is uploaded. Subsequent backups, being incremental and thus smaller in size, will revert to using the available bandwidth. This process is known as a Local SnapShot. The same logical process applies to recovering large amounts of data to the server. Proxy settings can be specified during these processes, if needed.

Local Cache: Once the FS has confirmed the successful transfer of the initial backup, the Backup Client drops the connection and creates a local cache. This is a compressed folder containing copies of all files that are backed up that had been modified in the previous 14 (default setting) days. The size of this folder is dependent on the type and size of the files concerned and also the modification date setting. To reduce the size of the cache the date can be reduced to files changed in the last 7 days etc. or turned off altogether. This will impact on the patching process.

The Restore Process

When a user wants to restore files, the Backup Client connects to the Storage Platform and retrieves the list of all the stored backup dates. The user then selects a date and requests the index file for that date from the SP. The index file contains a list of all the files that were protected with that backup. Once the index file has been transmitted, it is displayed in the client interface where the user can select the file(s) to be restored. The user can also use the Find option to search for certain filenames or file types. The list of files to be restored is then sent to the server. The server goes through a process of finding the right files in previous backups, applying any applicable patches, and then compressing the files after which the requested files are transmitted. Options to restore to the original location, recreate directory structure, restore empty folders, overwrite, and restore folder and file permissions are available. Compression can also be disabled in the Backup Client. This speeds up the restore process on a LAN as the Storage Platform will not compress any patched files before transferring them to the Backup Client.

The Storage Platform Console, introduced in Backup Professional v4.4 also enables you to remotely connect to, and manage any Server Edition Backup Client in your organisation. This includes the ability to restore files and folders to the remote server.

SAMPLE STATISTICS

1. Binary Patching of 1GB/10,000 files (user data – doc, ppt, xls, pdf)

Time to compress: 10 minutes
 Compressed size: 620 MB
 Time to patch: 3 minutes (200 modified files)
 Patch size: 5 MB
 Size of cache: 200 MB

2. Binary Patching of a 20GB Exchange Database

	Process	Time	Result	Percent
Day 1	Compression	1.6 hrs	8.4 GB	42%
Day 2	Patching	3.3 hrs	190 MB	0.95%
Day 3	Patching	3.3 hrs	182 MB	0.89%
Day 4	Patching	1.9 hrs	52 MB	0.26%
Day 5	Patching	4.1 hrs	260 MB	1.60%
Day 6	Patching	2.6 hrs	112 MB	0.56%
Day 7	Patching	3.1 hrs	172 MB	0.86%

Average process time (excl. Day 1): **2.8 hrs**
 Average backup size (excl. Day 1): **161 MB (0.81%)**
 Average local disk cache size: **10.8 GB**

3. Delta Blocking of a 20GB Exchange Database

	Process	Time	Result	Percent
Day 1	Compression	1.2 hrs	10.6 GB	53%
Day 2	Patching	0.8 hrs	2.01 GB	10%
Day 3	Patching	0.5 hrs	0.92 GB	4.60%
Day 4	Patching	0.7 hrs	1.34 GB	6.70%
Day 5	Patching	0.81 hrs	1.69 GB	8.45%
Day 6	Patching	0.74 hrs	1.39 GB	6.95%
Day 7	Patching	0.82 hrs	1.83 GB	9.15%

Average process time (excl. Day 1): **0.72 hrs**
 Average backup size (excl. Day 1): **1.53 GB (7.65%)**
 Average local disk cache size: **16.4 MB**