

Paragon Drive Backup Enterprise Server Edition

Best Practices for MS Exchange Server

Contents

1	Introduction	4
1.1	About Drive Backup	4
1.2	Backup Concepts	4
1.2.1	File-Level Backup and Volume Imaging	4
1.2.2	Data Consistency	4
1.2.3	Offline and Online Backup	4
1.2.4	Snapshots	5
1.2.5	Copy-On-Write	5
1.2.6	Data Synchronization	6
1.2.7	Write Inactivity Paradigm	7
2	Technology Overview	7
2.1	Paragon HotBackup Description	7
2.1.1	Hotbackup Concepts	7
2.1.2	How it Works	7
2.2	MS VSS Basics	8
2.2.1	VSS Concepts	8
2.2.2	How it Works	9
2.2.3	MS VSS Limitations	9
2.3	How Drive Backup Integrates with MS VSS	9
2.3.1	Enabling Backup via VSS	9
2.3.2	How it Works	10
2.4	Choosing between Online and Offline Backup	10
3	Protecting Microsoft Exchange Server	10
3.1	Microsoft Exchange Server Storage Management Overview	11
3.1.1	Exchange Storage Components	11
3.1.2	Exchange Files	12
3.1.3	Transaction Logs	13
3.1.4	Clean Shutdown State	14
3.1.5	Log Truncation	14
3.1.6	Data Restoration and Rolling Forward	14
3.1.7	Performance	16
3.1.8	Reliability	17
3.1.9	Backup	17
3.2	Backing up Exchange Databases with DBE	17
3.2.1	Recommended Backup Options	18
3.2.2	Recommended Information Store Layouts	18
3.2.3	Unsupported Database Layouts	19
3.3	Restoring Exchange Databases with DBE	20
3.4	Choosing between Hotbackup and VSS Online Backup Options	20
3.5	How to Use DBE for Protecting Exchange Data	21
3.5.1	How to Redistribute Exchange Data	21
3.5.2	How to Backup Exchange Data	24
3.5.3	How to Restore Exchange Data	28
3.5.4	How to backup and restore Exchange databases distributed across multiple volumes	34
4	Appendix	39

4.1 RAID levels39
4.2 Exchange Disaster Recovery and Troubleshooting Resources40

1 Introduction

This paper addresses various aspects of a Microsoft Exchange Server data protection by using Paragon Drive Backup Enterprise (DBE). It describes the concepts, limitations and best practices for Paragon Drive Backup Enterprise to protect "no downtime" operational solutions based on Microsoft Exchange Server. All mentioned recommendations are generic and not specific for a certain Exchange Server configuration.

1.1 About Drive Backup

Paragon Drive Backup Enterprise is a backup tool that implements the best of volume imaging techniques for reliable, fast and convenient data backup and restoration. The program includes end-user tools for building and automating recovery and replication procedures. It implements high-performance algorithms for intelligent data analysis and processing, provides an optimized manipulation for a large set of filesystems that covers all popular filesystems for Windows and Linux platforms and more features.

1.2 Backup Concepts

1.2.1 File-Level Backup and Volume Imaging

There are two concepts about backup subject. A *file-level backup* is oriented to store separate files. A *volume-level backup* or *imaging* is oriented to store whole filesystem of a volume.

A *file-level backup* naturally provides an intuitive and flexible way to select objects to store. File-oriented backup tools allow choosing any combination of both local and networking accessible files. A file-level data restoration allows to selectively restore only damaged files without affecting other ones. However, there are important file-related system objects which are not files and usually cannot be stored, restored and even accessed from a file level.

A *volume imaging* can store any metadata associated with files including distribution information, security data, quotas, extended attributes, named streams, multiple hard and symbolic links and so on. Disk imaging tools generally provide higher backup performance because they do not involve filesystem drivers to the process. In addition, imaging tools can backup offline filesystems and even ones not being supported by a host operating system. A data restoration generally does not require a host operating system to run, so that volume-imaging technique is a perfect choice for system cloning and disaster recovery tools. Disadvantages of volume imaging are that it cannot be applied to remote resources and a general ineffectiveness of backup and restore of selective files within the volume-imaging framework.

1.2.2 Data Consistency

The fundamental requirement to backup is saving of *data consistency*. This means that if applications are stopped, and data are restored, and applications are restarted, they will run smoothly with restored data.

A *data consistency* is conditionally divided into a *physical* and *logical* consistency. A *physical consistency* means storing information about involved data files and file-related attributes in a state that is interpreted by applications as "integral", or "valid", or "repairable on-the-fly". A *logical* consistency means an application-level correctness of stored business data. An automatic correction of minor inconsistency of business data is usually provided by the *transactions mechanism*, which is the basis of modern technologies of information processing.

1.2.3 Offline and Online Backup

As regards to the data consistency concept, offline data are in consistent state (with the only condition that an application or a system was shut down correctly). Offline data archiving is referred to as *offline backup*. Its advantages are ensured data consistency, increased backup speed (due to absence of concurrent data access), resource saving and low impact to a system performance.

The disadvantage is that offline backup is not applicable for "24x7 availability" systems. For continuously working systems *online backup* methods should be applied.

1.2.4 Snapshots

The great challenge to online backup is to provide a coherent state of all open files and databases involved in a backup, under the condition that applications may continue writing to a disk.

The "volume snapshot" concept has met that challenge. A *snapshot* is a point-in-time copy of a volume involved to a backup process. It must be quickly created at a period when applications do not write to disk. Once snapshot has been created, a backup utility copies data from it while applications continue working with an original volume. Modern snapshot technologies reduce *backup window* down to few seconds.

There are hardware and software based snapshot solutions. *Hardware snapshots* such as the *Split-Mirror* technique instantly provide an independent copy of a whole data set; they make no impact to the system performance and can participate in off-host backup solutions. Disadvantages are that hardware snapshots are naturally hardware-dependent, double requirements to storage resources and some time is required in order to re-synchronize a hardware snapshot after it has been used, with the actual disk data.

Instead, *software snapshots* are hardware-independent and resource-saving solutions. The disadvantage is that they cannot make a copy of a whole data set instantly, so that they run and make additional computing load to a host system until a backup routine is ended and a snapshot is destroyed.

Unlike hardware snapshots, types of software snapshots, which are used in practice, are not independent from original data. In case an original volume fails, all of its snapshots fail as well. Software snapshots do not really protect data; they are used as supplementary mechanisms in backup solutions.

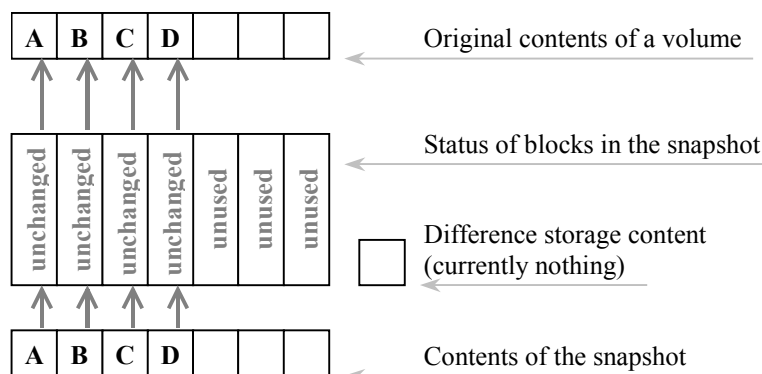
Paragon Drive Backup provides the original online backup technique (named *HotBackup*) and supports any snapshot technologies that can participate in the MS VSS framework.

1.2.5 Copy-On-Write

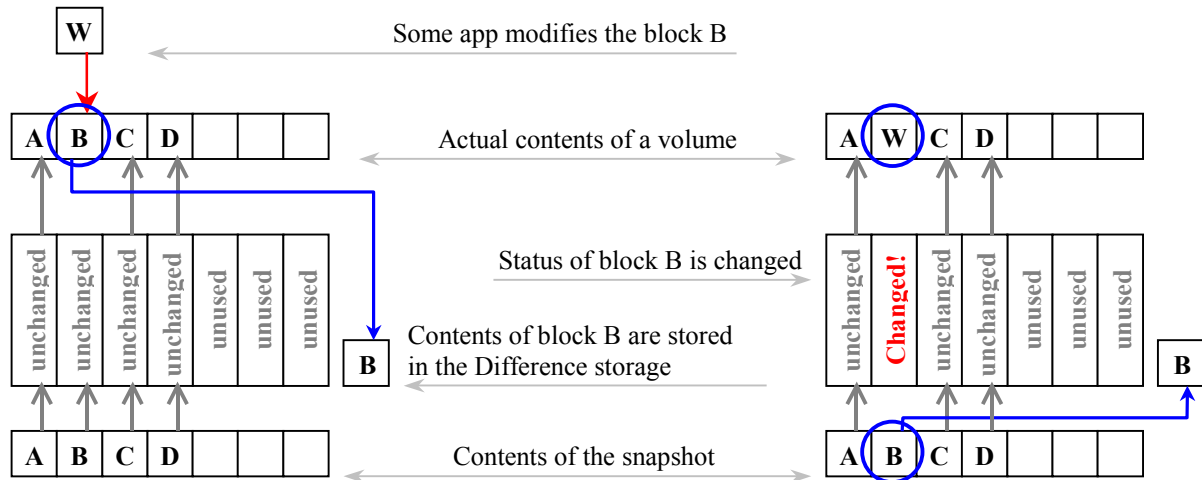
The *Copy-On-Write (COW)* method predominates among software snapshot techniques. It is based on the idea of preserving original contents of modified blocks before modifications are written to a disk, in special *difference storage*. COW maintains an original volume in the actual state and stores its original content at the moment of snapshot creation in the difference storage.

A system-level agent, which is responsible for snapshot maintaining, must keep track of all changes being made on a volume. On first attempt to overwrite a block, the snapshot agent copies original block's content to the difference storage and then allows modifying that block. When a backup utility acquires some block from the snapshot, the agent determines its status and takes a "changed" block from the difference storage while "unchanged" block is taken from the original volume. The following illustrates COW basics:

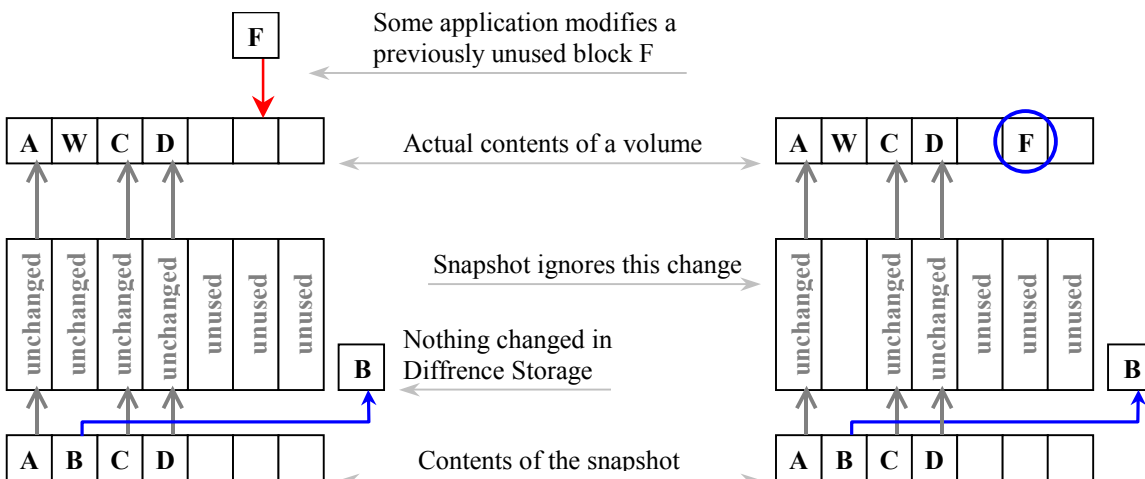
1. A volume was partially in use. Blocks A, B, C, and D contained data. A COW based snapshot was created for this volume. The snapshot provider watches only the used blocks of a volume. Initially they are marked as "unchanged".



- Some application tried to modify the block B. Before modifications are made, the snapshot provider copies contents of the block B to the difference storage. The block is now marked as "changed". Then the block B is updated. The snapshot will redirect all queries to the block B to data stored in the difference storage. Queries to blocks A, C and D will be directed to appropriate blocks of the volume.



- Some application tried to modify the block F, which was not originally in use. This block was not included in the snapshot. The snapshot provider does not take care of this change. Nothing is changed in the snapshot data.



1.2.6 Data Synchronization

Another problem for online backup functionality is that applications may temporarily hold open files in logically inconsistent state. The snapshot technique does not solve that problem as it concerns solely to business application operation. The true reason of the problem is that applications are unaware about a backup routine running and do not synchronize their data.

Microsoft has made a great attempt to solve the problem. The snapshot backup framework referred to as Volume Shadow Copy Service (VSS) has been built in latest versions of Windows, exactly to Windows XP, Windows Server 2003 and Windows Vista. VSS includes mechanisms for notifying applications about a backup, interchanging of related information between VSS participants and synchronizing execution of software components involved to the process.

VSS has several significant limitations:

- Only VSS compliant applications can benefit from VSS framework.

- VSS is a local solution that works within a single host.
Remote applications and distributed data systems aren't controlled by VSS.
- VSS currently works to its full capacity on Windows 2003 only.

1.2.7 Write Inactivity Paradigm

There are other (partial) solutions for synchronization problem that are applicable for non-VSS compliant software. There is a popular solution based on the paradigm of *Write Inactivity Period* (WIP) that was introduced by St. Bernard Software Company in the middle of 1990-th.

It is supposed that business applications for intensive information processing use transactions mechanisms. A *transaction* collects I/O writes into a compact group in order to reduce a chance of incomplete transaction committing if a failure of any type occurred. Data are in consistent state between transactions, so that a period when application(s) do not write to a disk is the best moment for a snapshot capture. This period is referred to as WIP – *Write Inactivity Period*.

2 Technology Overview

2.1 Paragon HotBackup Description

Paragon HotBackup is an online backup technology for Win'NT+ family operating systems. It was developed in 2001 and integrated to all company's backup solutions in 2002-2003. Currently it supports all versions of Windows NT4, 2000, XP and 2003 (including x86, IA64 and AMD64 versions).

2.1.1 Hotbackup Concepts

HotBackup is not a snapshot technology. However its concepts appear to be similar to ones used in software snapshot technologies. In particular, a sort of WIP observation and COW scheme are implemented.

During an online backup, Drive Backup uses the kernel mode driver HOTCORE.SYS in order to monitor and control write activity of applications and an operating system. The driver intercepts disk I/O requests and implements the most time-critical part of COW scheme while the Drive Backup utility includes disk data analysis and archiving functions.

The HOTCORE driver is installed during the standard Drive Backup setup procedure, and it is the reason why the system restart is required in order to complete the setup procedure. HOTCORE does nothing until it is activated by the Drive Backup. In the idle mode the driver bypasses any calls, makes no impact to the disk subsystem performance and only takes few kilobytes of system memory.

2.1.2 How it Works

HOTCORE driver is activated only in case the online backup is performed. In an offline backup mode, the driver is not involved to the process.

- Drive Backup activates the HOTCORE driver in the beginning of the "physical" backup.
- HOTCORE waits for a pause between IO writes on the targeted volume.
- When the pause is observed, the driver takes a "snapshot".

Within the framework of HotBackup, a "snapshot" is a map of blocks to be protected by the COW scheme against losing their original contents. The process requires the close cooperation between the driver and the utility. Finally, the "protection area" of the snapshot includes only used blocks with optional exception of *excluded files* (e.g. PAGEFILE.SYS and HIBERFIL.SYS). The embedded module for filesystem analysis allows reducing this step down to few seconds or less.

- During the snapshot capture, the driver watches the volume against I/O writes. If a write occurred before "snapshot" is created, the step is repeated.
- Upon successful snapshot creation, the driver applies the COW scheme to the "protection area".
- The utility starts the backup process. Archived blocks are immediately excluded from the protection area, so that the area is shrinking during the backup.
- If a foreign application tried to modify a "protected" block, the driver preserves its original contents in a buffer. Initially, blocks are stored in a memory buffer. When it becomes near full, the utility moves blocks to temporary files (named like "X:\hb_nn.tmp", where X: is a drive defined in the Settings).

- Normally Drive Backup performs the fastest streamlined backup of used blocks. However, if temporary files grew very fast or became very large, the utility pauses the streamlined mode and begins emergency backup of buffered blocks.

In online backup the following time-related restrictions are used:

- A snapshot must be created within 60 seconds.
- Buffered data must be processed within 10 seconds.

These restrictions are hard-coded and cannot be changed. If any of these restrictions were not satisfied, the online backup session fails. If a user aborted the backup operation or the utility failed, the driver automatically switches to the idle mode in 10 seconds.

2.2 MS VSS Basics

Volume Shadow Copy Service (VSS) is an open system-level framework for snapshot backup solutions. It was developed by Microsoft in close cooperation with leading vendors of backup solutions. VSS is included in Windows XP, Windows Server 2003 and Windows Vista.

VSS provides mechanisms for notifying applications about a backup, interchanging of related information between VSS participants and synchronizing execution of software involved to the process. These mechanisms ensure consistent backup of online data for VSS compliant applications.

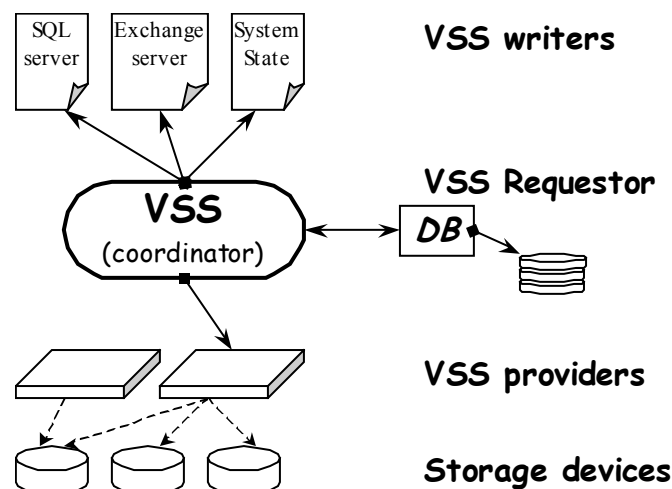
2.2.1 VSS Concepts

VSS is based on concepts of a snapshot and a volume shadow copy. Being invoked by a VSS aware backup utility, VSS creates snapshots for selected volumes and represents them as virtual read-only devices, which are referred to as *shadow copies of volumes*. Once shadow copies are created, a backup utility stores data from shadow copies while business applications continue writing to original volumes.

There are three kinds of software involved to the VSS framework:

- *Providers* (snapshot providers) – tools that create and maintain hardware or software snapshots.
- *Requestors* – utilities that acquire shadow copies, usually backup utilities.
- *Writers* – VSS compliant applications that hold open files on volumes, actual backup objectives.

Within the VSS framework, VSS writers are able to inform other VSS participants about files being in use, file grouping and restoration conditions. A group of files that constitute a whole entity in a business application and should be backed up together is named a *writer's component*. For example, all files that constitute a mailbox store in MS Exchange Server are represented as a single writer's component. VSS writer can be an application itself or a special agent, which provides VSS-to-application interaction.



VSS itself only coordinates activity of providers, writers and requestors. A standard Windows XP/2003 distribution includes the VSS coordinator, the universal software provider (VOLSNAP), several VSS writers for system components and the universal VSS writer for MS Desktop Engine (MSDEwriter), which is responsible for integration of MS SQL Server to VSS framework.

Microsoft Exchange Server 2003 is the first VSS compliant version of Exchange. The VSS writer for Exchange, which is named "Microsoft Exchange Writer", is installed to the system at the normal Exchange 2003 setup procedure in Windows Server 2003 environment.

2.2.2 How it Works

Here is only a brief description of MS VSS operation. The comprehensive description of VSS can be found on appropriate Microsoft TechNet pages (e.g. see the descriptive topic "How Volume Shadow Copy Service Works", <http://www.microsoft.com/technet/prodtechnol/windowsserver2003/library/TechRef/2b0d2457-b7d8-42c3-b6c9-59c145b7765f.mspx>).

1. A VSS requestor (backup utility) invokes VSS to initialize the backup routine.
2. VSS acquires information from all VSS writers. As a result the *Writers Metadata Document* (WMD) is created. It contains distinctive names of applications and components, restoration parameters, list of files and other information.
3. A requestor receives WMD and makes a decision what to backup. It creates the *Backup Components Document* (BCD), which defines volumes, writers and components involved to the process and sends it to VSS. As an option, preferred VSS providers can be selected.
4. VSS commands to all involved VSS writers to finish running transactions and then "freeze". VSS waits until all involved writers complete that step.
5. Then VSS commands to appropriate VSS providers to create snapshots for selected volumes.
6. After snapshots are created, VSS allows writers to "thaw" and to write to disk. In addition, VSS asks writers if write operations were indeed suspended during the snapshot creation. If it was not the case, the whole backup procedure fails, snapshots are removed and the VSS requestor is notified.
7. Upon successful completion of above steps (within predefined time intervals), VSS creates *shadow copies* from snapshots and provides VSS requestor with appropriate references.
8. VSS requestor starts copying information from shadow copies...
9. Upon the process completion, the VSS requestor informs VSS about the backup completion status.
10. VSS informs all involved VSS writers about the backup completion status. VSS writers may use this notification to perform some specific actions (for example, Exchange truncates log files).

A shadow copy can be deleted immediately after the backup completes, or it can persist in the system. In the last case it can be mounted like an ordinary volume (this feature is available in Windows Server 2003 and Vista).

2.2.3 MS VSS Limitations

MS VSS has several significant limitations:

- Only VSS compliant applications can benefit from VSS framework.
- VSS is a local solution working within a single host. Remote applications aren't controlled by VSS.
- VSS currently works to its full capacity on Windows 2003 only.

2.3 How Drive Backup Integrates with MS VSS

Storage management frameworks obviously provide benefits and can simplify storage management activity. Now Paragon Drive Backup Enterprise is adapted to participate in VSS framework operation as a requestor.

2.3.1 Enabling Backup via VSS

To perform backup operations via VSS services, select the "Microsoft Volume Shadow Copy Service " option:

```
(menu:) Tools
      ↳ Settings...
          ↳ Hot processing options
              ↳ Hot processing technology
                  ↳ Microsoft Volume Shadow Copy Service
```

Note that MS VSS service is available only in Windows XP, 2003 and Vista. In other operating systems, only Hotbackup technology is available for online backup.

Drive Backup provides a simplified VSS management, so that not all VSS options are controllable by a user. The program internally chooses only most reliable VSS operation modes.

2.3.2 How it Works

A VSS aware backup is performed in the following manner:

- A user chooses volume(s) to be backed up in the program's interface.
- When the "physical" backup operation begins, the program receives the compound WMD document from VSS (see #2 in the VSS description topic).
- Drive Backup determines VSS writers having components located on chosen volumes. These VSS writers are included to the BCD document (see #3). In other words, all VSS writers containing files on targeted volumes should participate in VSS operation at the snapshot creation (i.e. should be "frozen" and "thawed" by VSS).
- Drive Backup commands VSS to use the default order of VSS providers invocation: a hardware provider first (if available), a third-party software provider next (if available), the last is Microsoft's universal system provider VOLSnap (always available).
- The program performs the volume-level backup of the created shadow copy set.
- After the operation completes, the shadow copy set is deleted.

Currently the program does not support restoration via VSS writers. In fact, VSS based restoration is generally just a file copying. Applications should be stopped, or appropriate components should be detached/dismounted manually in order to be restored.

2.4 Choosing between Online and Offline Backup

The reasons to prefer offline backup are:

- Online backup via Hotbackup or VSS option is slower than offline backup.
- VSS initialization is long and generally unstable under high IO traffic on a targeted volume.
- A resulting image produced by Hotbackup is slightly larger than one produced in offline mode because of non-sequential image structure (see Hotbackup description).
- Neither of online backup options totally eliminates problems that are naturally inherent for online backup technologies in general.

VSS provide data consistency for VSS compliant applications only. Hotbackup does not guarantee 100% data consistency in any case, but only a very high probability of that. In fact, it provides perfect results for any applications that use transactions (such as MS Exchange and SQL servers).

Drive Backup provides some flexibility of choosing between offline and online backup mode. A user can choose between three modes: (a) "always offline backup", (b) "switch to online if a volume was in use" and (c) "always online backup". The differences between these modes are the following:

(a) Always offline mode:

- The program does not backup volumes being in use. It suggests to reboot the computer in order to complete the operation in the "Startup Bluescreen" mode.
- If a volume wasn't in use, the program switches to exclusive use of the volume. No applications can access any files on the volume until the backup procedure is finished.

(b) Switch to online if a volume was in use:

- If a volume was in use, the program performs the online backup. Other applications are allowed to access the volume during the operation.
- If a volume wasn't in use, the program performs the offline backup. As it was described above, no applications can access any files on the volume until the backup procedure is finished.

(c) "Always online backup":

- The program unconditionally performs the online backup. Other applications are allowed to access the volume during the operation.

3 Protecting Microsoft Exchange Server

Microsoft Exchange Server is a high-availability enterprise-class mail server for Windows platform, which can scale from hosting simple messaging systems to mission critical business applications. The more your business depends on Exchange Server, the more important it is to protect it.

This version of Paragon Drive Backup Enterprise has no Exchange specific functionality. For this reason, some considerations should be taken into account in order to provide reliable backup of Exchange data.

The best practices presented in this guide are general principles, not guidelines for specific environments. This chapter discusses general requirements for Exchange configurations.

3.1 Microsoft Exchange Server Storage Management Overview

Protecting Exchange data is critical for business. Understanding the structure of information storage is essential for optimal use of the Exchange Server and effective use of backup utilities for data protection.

Exchange uses a transaction-based logging informational system to securely store messaging content under conditions of high load. A hierarchical structure of Exchange databases is dedicated for providing high level of data availability. The core of this system is named Extensible Storage Engine (ESE, also known as JET engine). ESE is optimized for fast data retrieval since this is the main function that the database performs.

Knowledge of the underlying database technology can help to better understand the backup and restore processes and avoid hidden obstacles. A detailed description of Exchange Information Store is available on Microsoft TechNet resources. This section briefly discusses aspects of the Exchange databases layout related to Drive Backup operation.

The structure of Exchange Information Store has been noticeably changed since moving from Exchange 5.5 to Exchange 2000 and 2003. However, most common features in the underlying databases operation change insignificantly in relation to Drive Backup function.

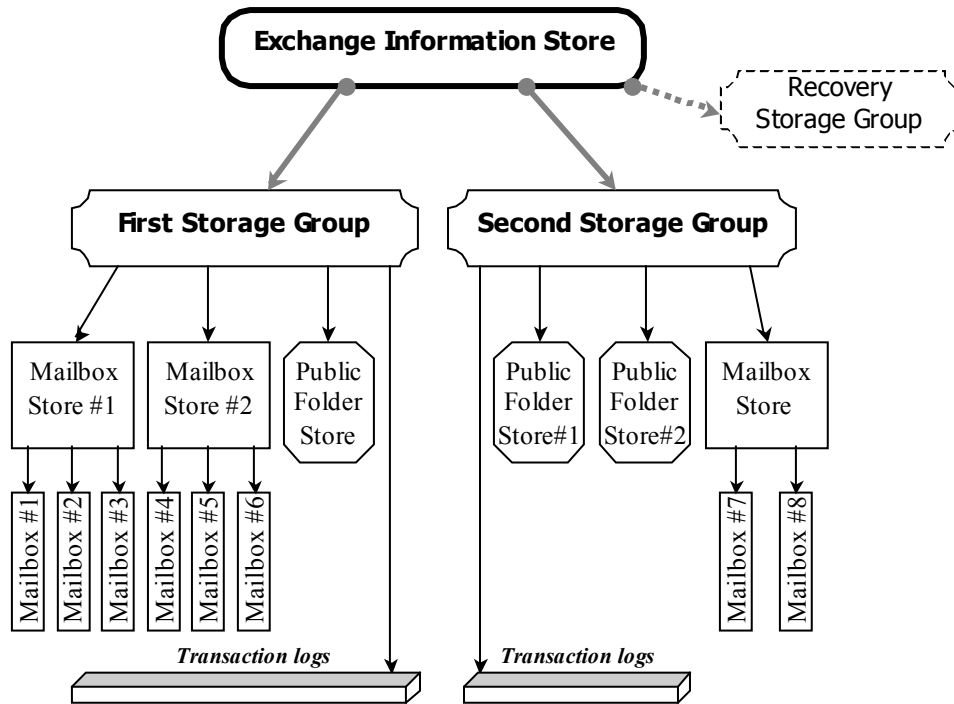
3.1.1 Exchange Storage Components

Exchange Server has two main databases for user information – the Directory Store and Message Store.

The *Directory Store* is the database of Exchange users. In Exchange 5.5 and earlier versions, the Directory Store was placed in a dedicated Exchange database. With Exchange 2000 and 2003, the Directory Store is integrated to the Active Directory. Advantages of that approach are centralized administration and transient data availability. The side effect is that Exchange 2000+ must run on Windows in an environment where Active Directory is available. Although the Directory Store data change infrequently in comparison to the Message Store, it is critical that the backup schedule for the Directory Store is coordinated with the backup schedule for the Information Store in order to maintain consistency between users and messaging data.

The *Message Store* is where messaging data are stored. The Message Store is logically divided into two types of databases – "public" and "private". The first type contains Public Folder data, the second one keeps user Mailbox Stores. For better scalability and data protection, Exchange Server 2000 and 2003 allow the Message Store to be split into several Storage Groups of databases.

Each *Storage Group* contains several databases and shares *transaction logs* between databases within the Group. Each *database* can be either a *Mailbox Store* (contain private user mailboxes) or a *Public Folder Store* (contain public data). The hierarchical storage layout provides a very flexible data protection scheme as well as a possibility for data processing parallelism.



Difference between Exchange versions in relation to Information Store capacity:

Feature	Exchange 5.5 and earlier	Exchange 2000/2003 Standard edition	Exchange 2000/2003 Enterprise edition
Amount of Groups	N/A	1 SG + 1 RSG	4 SG + 1 RSG
Amount of Stores	1 Mailbox + 1 Public Folder	1 Mailbox + 1 Public Folder	5 Stores total (e.g. 3 Mailbox + 2 Public Folder)
Store Size max.	16 GB total	16 GB per Store (75 GB for SP2)	16 TB per Store

Here the **SG** abbreviates "Store Group" and the **RSG** abbreviates "Recovery Storage Group".

Recovery Storage Group is used as a temporal replacement for ordinary Storage Group(s) in case of some types maintenance activity for avoiding blackouts for Storage Group users.

3.1.2 Exchange Files

Exchange databases consist of following several type of files, most important of which are the following:

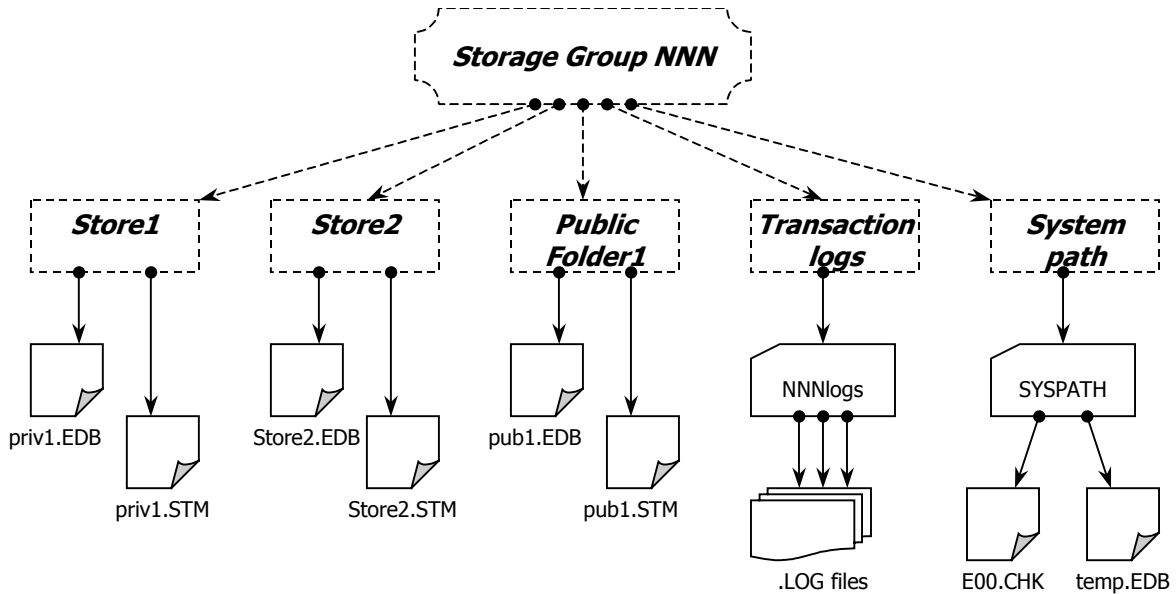
.EDB	"Property data", contain a textual part of messages in MS Database Encapsulated Format (MDBEF).
.STM	"Streaming data", contain binary content of messages in MIME format.
.LOG	Transaction logs files.
.CHK	The Checkpoint file. It stores a separator between flushed and un-flushed parts of transaction logs.

Each database such as a Mailbox Store or a Public Folder Store consist of two files: **xxx.EDB** and **xxx.STM**, where "xxx" denotes some same name.

Default file names of Exchange files:

- A Mailbox Store consists of **PRIV1.EDB** and **PRIV1.STM** files,
- A Public Folder Store consists of **PUB1.EDB** and **PUB1.STM** files,
- Exchange 5.5 Directory Store is located in the **DIR.EDB** file,
- A temporary store is located in the **TEMP.EDB** file (it holds intermediate transactions data),
- The most recent part of the transaction log is placed in the **E\$\$\$.LOG** file,
- The older parts of the transaction log are located in the **E\$\$\$###.LOG** files,
- Two reserved fragments of the transaction log in **RES1.LOG** and **RES2.LOG** files,
- The checkpoint information is placed in the **E\$\$\$CHK** file.

Exchange provides a high flexibility for Information Store files allocation. It is the root of high-level scalability of Exchange Information Store. Database files for each Store, transaction log files for every Storage Group as well as mail queues for every virtual SMTP server can be placed in different locations. The below picture demonstrates relationship between Storage Group components and physical files and directories:



Dashed frames and lines designate logical components and logical relationship, while solid frames and lines designate physical directories and files and dependencies. The picture demonstrates that database files for every Store within same Storage Group can be individually located. Transaction log files are shared between all Stores within the Storage Group; they are always placed in same directory, which can be individually located. The specific files (named "system files" within the Exchange framework, **E00.CHK** and **TEMP.EDB**) are always placed in same directory, which can be individually located.

By default, Exchange places log files and system files to a same directory while all individual Stores within the Storage Group are placed to another directory on same volume. The default locations for Storage Groups are separate directories on the volume where Exchange binaries are installed, which in turn is the Windows system volume (by default again).

You should take into consideration that such distribution of Exchange Information Store components probably is easy to maintain but is contrary to all recommendations for Exchange reliability and performance improvements from all sources including Microsoft itself. It is strongly recommended to rearrange newly added components of the Exchange Information Store (i.e. individual Stores and Storage Groups) immediately after its creation, or at least as soon as it possible. This task can be completed from Exchange System Management console. An empty Store is moved to another location within few seconds while for a large Store the operation may take a long while.

To choose a better configuration for Exchange Information Store, read and follow recommendations for improving Exchange reliability and performance. In addition, you can use the Exchange Best Practices Analyzer Tool from Microsoft.

3.1.3 Transaction Logs

Transaction logs are essential part of the Exchange Information Store. It is the key component of the disaster recovery procedure for Exchange. The most exciting feature of Exchange transaction logs is that it is able to compensate the data loss between the moments of the last full backup and a disaster.

Exchange keeps logging transactions for every database. Transaction log files are shared between all databases within each Storage Group. For this reason backup, restore, log re-play and truncation procedures are coordinated for databases within a Storage Group.

Each transaction log is divided into small portions of 5MB in size. The most recent part of the log is placed in the **\$\$\$LOG** file, where "\$\$" denotes the logical number of the Storage Group. For the Exchange Enterprise edition, it can be **E00.LOG** through **E03.LOG** (because up to four regular Storage Groups are allowed). For the Standard edition, there is only one regular Storage Group, so that there is only one transaction log and one the most recent log-file named **E00.LOG**.

As soon as the most recent part of the log becomes full, it is renamed to an **E\$\$#####.LOG** file, and a new **E\$\$LOG** file is created. An "E\$\$#####" filename consists of the Storage Group's *system prefix* "E\$\$" (e.g. "E00" for the First Storage Group) followed by the 5-digit hexadecimal sequential number of a log file.

Exchange supports two modes of transactions logging. In the normal mode, the transaction log grows unrestrictedly and log files are sequentially indexed, as it was described in previous sections. The transaction log is truncated only in case of a full backup of entire Storage Group, i.e. all Stores nested in the Storage Group, its transaction log and system files, which are shared between all databases within the Group. The normal mode of transactions logging enables logs re-play at database restoration.

There is another mode named "circular logging". In the circular logging mode, Exchange uses circularly only existing log files without creating new ones. This mode allows to avoid wasting disk space. However the log re-play option is basically unavailable in circular logging mode. Microsoft does not recommend to use circular logging for business applications.

3.1.4 Clean Shutdown State

When Exchange is under load, transactions are intensively stored to the log-files and Exchange lazily flushes logged data to a main database (EDB and STM files) in a background mode. This allows avoiding Exchange productivity degradation, however the transaction log usually contains large amounts of un-flushed operational information.

Exchange forcedly flushes all currently logged data to database files only in particular cases:

- Clean dismount of a Store or a clean shutdown of Exchange Server,
- Maintenance tasks that require dismounting a Store (e.g. for Store migration).
- Backup of a Store, Storage Group or a whole Information Store.
- Performing a recovery operation on a dismounted Store (by using the ESEUTIL utility).

The concept of "clean shutdown" state of a Store is essential within the framework of Exchange because Exchange server can mount only a Store being in the "clean shutdown" state. A Store that is in the "dirty shutdown" state must be *recovered* first (by using the ESEUTIL /R utility). The Store Recovery operation flushes transactions to the main database and places a Store to a clean shutdown state.

3.1.5 Log Truncation

Thanks to keeping transaction log files, Exchange is able to compensate the data loss beyond the moment of a full backup in case of a disaster occurred. However under conditions of intensive loading, transaction logs may grow very large. Exchange does not truncate it and allows the log to grow unrestrictedly. Even flushed parts of the log are left on the disk.

A *log truncation* is really equal to removal of numerous unnecessary .LOG files. It is performed only upon successful completion of a full backup of a Storage Group. For this reason, a *backup* is the essential procedure within the frame of Exchange operation. It is strongly recommended to backup every Storage Group of Exchange Information Store on a base of regular schedule.

3.1.6 Data Restoration and Rolling Forward

When a disaster occurred, and a Store is restored from an archive, "old" parts of the transaction log are restored together with the main database files. The "old" parts can be combined together with undamaged "new" parts of the log, which were created after a moment of a full backup. The combined transaction log can be re-played as a whole, and this will *roll forward* databases beyond the moment of a backup. As a result the Store will come to a "clean shutdown" state that corresponds to a point in time very close to the moment of a disaster.

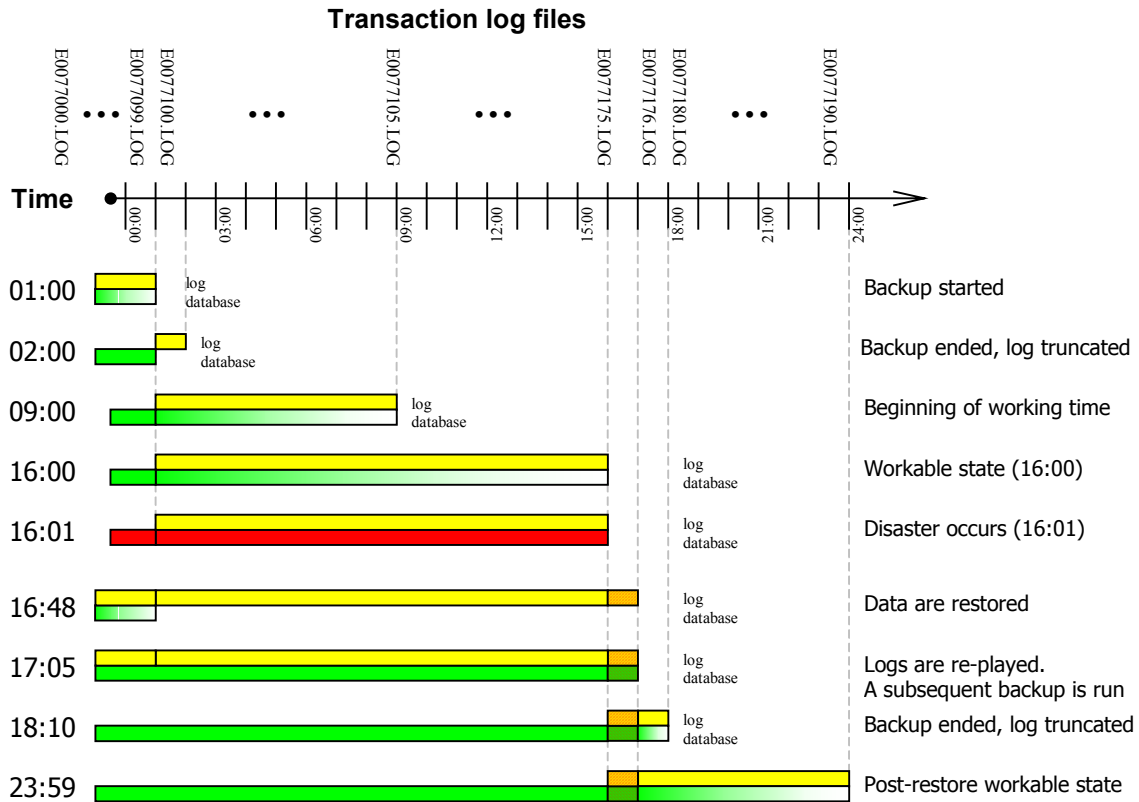
This can nearly nullify the amount of messages being lost due to a disaster. It is the key feature for Exchange data high recoverability but it requires careful management of transaction log files.

Let's demonstrate the behavior described above. Suppose there is an Exchange based messaging system. There is the "Sales" Mailbox Store. The Store is scheduled to back up every night at 1:00 AM, and the backup procedure takes ~1 hour. We will have the following picture:

1. At 01:00 the database consisted of the Sales.EDB, Sales.STM, E00.LOG and one hundred "old" log files E0077000.LOG to E0077099.LOG that keep mail traffic for the previous day.

2. At 01:00 the backup routine is started.
It backs up all files mentioned above. Also the new log file E0077100.LOG is created. It begins accumulating messages received while the backup routine runs.
3. At 02:00 the backup routine has been finished. Exchange flushes data from log files E0077000.LOG to E0077099.LOG to the main database and then deletes these log files. Now the Store consists of the Sales.EDB, Sales.STM, E00.LOG and E0077100.LOG files.
4. By 09:00 (the beginning of an operational day) the Store consists of the Sales.EDB, Sales.STM, E00.LOG and E0077100.LOG to E0077105.LOG files. Some part (the most part) of logged information has been flushed to the main database.
5. The operational day begins and the log grows very quickly. At 16:00 it consists of E0077100.LOG to E0077175.LOG files.
6. At 16:01 a failure occurred and a main database of the Store becomes corrupted (for example, Sales.STM become unreadable).
7. An administrator invokes the restore routine. Suppose it has been finished at 16:48. The procedure restores the 15-hours-old version of Sales.EDB and Sales.STM and "old" log files E0077000.LOG to E0077099.LOG. These files keep database state at the moment of 01:00 AM. If the data retrieval process finished here, all operational data for this day will be lost.
8. In case the "new" log files E0077100.LOG to E0077175.LOG, which are created since 01:00, are not corrupted, the restore routine combines "old" and "new" log files. Now the transaction log consists of E0077000.LOG to E0077175.LOG.
9. At the end of the files restoration, the Exchange database restore procedure is invoked. It re-plays accumulated transaction log files (from E0077000.LOG to E0077175.LOG). Upon completion of that process, the database will become at the state that corresponds to 16:00.
10. Suppose this phase is ended at 17:05. At this moment an administrator can mount the recovered Store and Store users can continue working.
11. As it recommended by Microsoft, a successfully restored Exchange system should be backed up in order to (a) save the restored system state and (b) force Exchange to truncate log files. The backup procedure is started at 17:05 and is ended at 18:10. As a result, the recently unpacked log files are surely flushed to a main database and transaction logs are truncated. Only few log files will remain (e.g. from E0077176.LOG to E0077180.LOG)
12. At 23:59 the Store's transaction log includes E0077180.LOG to E0077190.LOG. The tomorrow backup will truncate the log and drop unnecessary log files.

Below is the illustration of data restoration and log re-play:



Legend:

- Data logged in LOG files
- Data stored in the database files
- Data partially flushed to the database
- Corrupted data / files
- } Store blackout period
- } Store blackout period

Restore from a backup image retrieves data processed to the moment of the last full backup (up to today's 01:00 in our example). The log re-play allows retrieving data processed since the full backup to the moment of failure (today's from 01:00 to 16:00 in our example). The only trouble period is the period of the damaged Store blackout (16:01-17:05 in this example).

In Exchange 2000 and 2003, a Recovery Storage Group can be used for providing partial operation for users of the damaged Store. Exchange administrator can quickly reconfigure Exchange so that users of damaged database(s) will use the Recovery Storage Group to send and receive messages while the damaged Storage is under repair. When the recovery process completes, administrator can invoke merge procedure in order to add messages from the Recovery Store Group to the fixed ordinary Storage.

3.1.7 Performance

There are multiple factors that affect Exchange databases performance. Most significant and common of them are available memory and disk subsystem throughput. Exchange Information Store service (STORE.EXE) consumes vast amount of memory for effective indexing and logged data caching.

High throughput of disk subsystem generally enhances performance of Exchange Information Store as well. Most effective ways to increase disks throughput are using faster disks and using RAIDs. Striped RAID sets provide highest read/write performance that is found nearly a multiple of single disk performance.

RAID 0	Striped Disk Array (! RAID-0 is not fault-tolerant!)
RAID 0+1	Mirroring of striped segments (RAID 1 over RAID 0)
RAID 3	Striped Disk Array with Isolated Parity
RAID 5	Striped Disk Array with Distributed Parity
RAID 10	Striping of mirrored segments (RAID 0 over RAID 1)

It is a good practice to place your production databases on a high-performance RAID set.

For Exchange, placing its components onto separate storage devices can accelerate the Information Store – an additional performance gain is achieved by parallel execution of I/O requests.

3.1.8 Reliability

While full backup of databases provides a disaster recovery option, there are common ways to reduce the chance of a failure. A hardware-based method is to use fault-tolerant RAID configurations (for example, RAID-10 or RAID-5).

Isolating its components from an operating system, from each other and from other intensively used file resources can enhance an overall reliability of the Exchange Information Store. It is a good practice to place your production databases on an isolated non-system volume, or better to a dedicated disk device.

3.1.9 Backup

Exchange databases must be regularly backed up. It is not only needed to protect business applications from disaster but is the obligatory activity for normal Exchange operability.

There is an exception from this rule. Mail queues of Exchange virtual SMTP servers (**Queue**, **Pickup** and **Badmail** subdirectories of the virtual SMTP root directory, which is by default is located at "C:\Program Files\Exchsrvr\Mailroot\vs1" directory). Their contents are extremely volatile and cannot be re-used after a post-disaster restoration. During a working session the mail queues may grow large, so that it is reasonable to surely exclude their contents from backups. A good practice is to isolate these components on dedicated volume(s) or disk devices.

Exchange supports several backup modes of its data. These modes differ in (a) set of data files being backed up and (b) post-backup data processing, or actions being made after a successful completion of a backup procedure. The table below describes the difference between backup types:

Backup Type	What is backed up	Post-backup actions
Normal (=Full)	Transaction logs and database files	Logged data are flushed to main database files and transaction logs are truncated
Copy	Transaction logs and database files	Nothing
Incremental	Transaction logs only	Logged data are flushed to main database files and transaction logs are truncated
Differential	Transaction logs only	Nothing

Full backup and Incremental backup require exchange to be aware of backup routine operation and results for their successful completion. Copy backup and Differential backup do not require that – backup routines of these types can be performed by non Exchange aware backup utilities. The only condition is that these utilities should somehow verify that files being backed up are in consistent state, i.e. that Exchange will interpret these files as consistent at post-restore attempt to mount databases.

3.2 Backing up Exchange Databases with DBE

Some concerns should be taken into account in order to use Paragon Drive Backup Enterprise for successful backing up of Microsoft Exchange Server data. The following circumstances should be considered:

- Is a Store (or Storage Group) located on a mapped network drive(s)?
- Is it spread over multiple volumes?
- Should it be backed up online, or it is acceptable to temporarily hold it offline?
- Which operating system is running on the host?

The governing factors that limit Drive Backup adaptability for SQL Server backup are the following:

- This version of Drive Backup does backup network drives.

- This version of Drive Backup does not include Exchange Server specific agents. Exchange databases are backed up like ordinary files. The WIP detection mechanism is used for database's data synchronization.
- Only backup via MS VSS option provides Exchange log truncation, and only for Exchange 2003 that is the only VSS compliant version of Exchange. VSS is available in Windows server 2003 and Windows XP only.
- Databases are always backed up in a mode that is effectively equal to the *full backup*.

This version of Drive Backup is unaware of logical structure of Exchange databases. Drive Backup's "full" and "incremental" modes refer to volume backup modes at block level, not to database backup modes.

Concerning to abilities of this version of Drive Backup Enterprise, the minimum component of the Exchange Information Store that can be backed up is Storage Group. To backup some Storage Group, you should include all volumes that contain components of that Storage Group, to the set of volumes being backed up simultaneously. In this case some components of that Storage Group are not included to the set of volumes, Exchange will not truncate the transaction log of that Storage Group. In addition, an incomplete backup of data may not be restored correctly.

Drive Backup supports both volume-level and file-level restoration of data. In case of volume-level data restoration, a whole Storage Group will be rolled back to the pre-backup state including database files and transaction log. In case of file-level restoration, there is an ability to implement the log re-play mechanism in order to roll forward the restored databases.

3.2.1 Recommended Backup Options

This version of Drive Backup limited support of Exchange versions that precede Exchange Server 2003. It is recommended to use Drive Backup Enterprise with Exchange Server 2003 installed on Windows Server 2003 host operating system. In addition, it is recommended to update Exchange 2003 to Service Pack 2 and Windows 2003 to Service Pack 1.

For any version of Exchange Server, Exchange data will be successfully backed up in offline mode, by using an offline backup as well as any online backup option (HotBackup, MS VSS). However it requires to temporarily dismount Exchange databases, in Exchange 2000 and 2003 a whole Storage Group should be dismounted. So that during offline backup some Exchange resources are unavailable.

In case Exchange data blackouts during backup are unacceptable, you should choose between online backup options. It is strongly recommended to use only the "Microsoft Volume Shadow Copy" online backup option to backup Exchange Server 2003 databases.

3.2.2 Recommended Information Store Layouts

General recommendations:

1. Place Exchange production databases on local volumes.
2. Isolate Exchange production databases from an operating system.
3. Place Exchange production databases on high-performance RAID set(s).
4. Place Exchange databases on a dedicated volume(s).
5. Avoid distributing a databases belonging to same Storage Group over multiple volumes whenever it is possible. Place all database files (.EDB and .STM files and transaction log) on a single volume.
6. Isolate mail queues (Queue, Pickup and Badmail) of virtual SMTP server(s) from production databases. It is preferred to place them on a dedicated volume or a dedicated disk device.

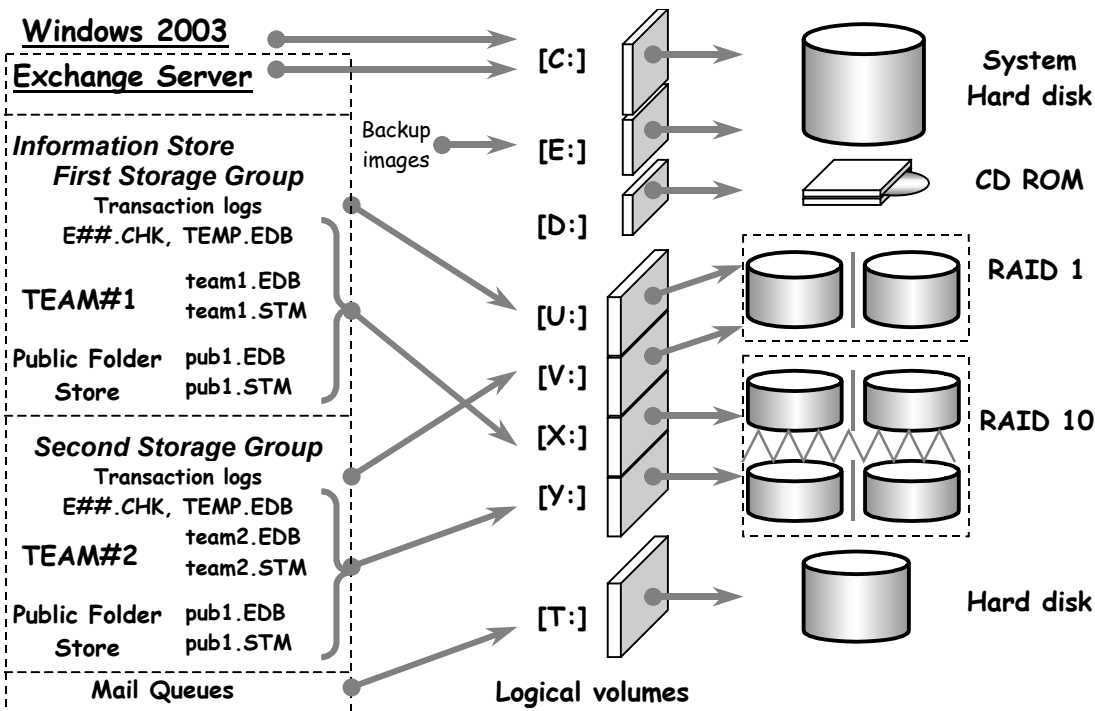
System-dependent recommendations:

7. Use only MS VSS option for online backup of Exchange 2003 databases in Windows Server 2003 and Vista environment. Choose the "Microsoft Volume Shadow Copy" item in the "Hot processing technology" pull-down list in the program settings for that purpose.
8. Avoid online backup of Exchange databases for versions that precede Exchange 2003.
9. Avoid online backup of Exchange databases in non Windows Server 2003/Vista environment.
10. Exchange databases can be successfully backed up in offline mode, for any version of Exchange and Windows. Dismount a Store or stop the Exchange services before backing up volumes having Exchange data. For offline databases, both online and offline backup modes provide successful results.

A preferred system configuration could be like this:

- Exchange production databases should be placed on a dedicated fault-tolerant RAID set. For example, a RAID 10 provides both high reliability and performance levels.
- The operating system and Exchange binaries are placed on the dedicated system hard disk. In case this disk is big enough, it can be partitioned in several volumes. OS and Exchange binaries should be placed on the dedicated system volume. Other volume(s) can be used for infrequently accessed data, e.g. for a temporal or a permanent storing of backup images.
- In case the budget allows this, add an extra hard disk dedicated for storing backup images and recovery tool set.
- The volatile components of the Exchange Information Store, such as the temporary database and mail queues, should be isolated from system files and other Exchange data onto a dedicated storage device.
- It is preferred to isolate Exchange Storage Groups from each other, and isolate database files (.EDB and .STM) from transaction logs within each Storage Group.
- Exchange is placed on the dedicated server.
- It is preferred that Exchange dedicated server is not the primary domain controller.

Placing Exchange on a dedicated server enhances application reliability and performance but it increases the cost of Exchange based solution. Exchange 2000 and 2003 require Active Directory to operate, which in turn requires a domain based network and domain controllers, which usually are dedicated computers. Domain controller hosts with Active Directory installed are not a best Exchange placement because of potential reliability reduction and performance degradation (<http://support.microsoft.com/kb/321543/?sd=RMVP&fr=1> or <http://support.microsoft.com/?kbid=888794> for details).



3.2.3 Unsupported Database Layouts

This version of Drive Backup has some limitations in relation to backing up Exchange data:

1. A Store or a Storage Group that is entirely located on a mapped network drive cannot be backed up by using Drive Backup, because this version of the program does not backup network drives.
2. A Store or a Storage Group that is partially located on a mapped network drive cannot be backed up by using Drive Backup, because not all database files can be stored.

Generally, a Store or a Storage Group partially located on a mapped network drive may be an obstacle for accurate data restoration for the whole volume. As Drive Backup is unable to backup completely such databases, it is unable to restore them correctly, too. To avoid possible problems caused by incomplete database restoration, do not apply a volume-level restoration of volumes containing distributed databases. Instead, a file-level restoration should be applied for such volumes. Use the "Image Explorer" utility for the file-level restoration from Drive Backup's backup images.

3. If Exchange (any version) was installed in Windows 2000 or NT 4.0 environments, only the "Hotbackup" online backup option is available. In this case, distributed Exchange databases cannot be successfully backed up in online mode (because of Hotbackup synchronization limitations). Such databases can be accurately backed up in offline mode only. You should dismount a Store (a whole Storage Group in Exchange 2000 and 2003) before backing up volumes containing distributed databases.

3.3 Restoring Exchange Databases with DBE

This version of Drive Backup does not apply online restoration for Exchange Server databases. Offline restoration is the only available method. A database migration-at-restoration is not available. In fact these limitations are common for snapshot-based backup solutions for Exchange. Before starting restoration of Exchange data, ensure that a Store or Storage Group of interest is really dismounted.

A database can be restored in two ways – as a part of the volume restoration procedure or by manual extraction of files from a volume's image and placing them to the original location. Because of specific abilities of restoration of Exchange databases (log re-play abilities), the primary method is the file-based restoration of data. The volume-based restoration should be used only in cases of severe damage of storage devices that leave no chances for retrieval of full or partial set of transaction log files.

The difference between the volume-based and file-based restoration methods is the following:

- At *volume-level restoration*, all contents of a targeted volume are rolled back. File system metadata are also restored and put into consistent state. The volume-level restoration is very fast and is able to recover volumes from scratch.
- At *volume-level restoration*, in case a restored volume contains transaction log files, they will be deleted, and only the set of "old" log files will be retrieved. With the resulting set of log files, the log re-play will only restore Exchange databases to the moment of time preceding to the backup. In addition, any other databases and objects, which are located on that volume, will be returned to their previous state.
- At *file-level restoration*, one can extract only required files from an image. It provides selective data retrieval. Other volume's contents are not affected. As concerns to Exchange data recovery, it allows to implement the log re-play technique.
- However the *file-level restoration* is rather slow and requires a targeted volume to be healthy. File-level restoration cannot be used for recovering damaged volumes.

To restore an Exchange database, one should first dismount all Stores that share same Storage Group with the restored database, or ensure that all Stores are indeed unavailable (e.g. because of failure). After a database is restored, it should be mounted again as well as all Stores that share same Storage Group with it.

In specific cases the Exchange Server may fail to mount a restored database, with generating errors. In this case a situation investigation should be performed, the problem should be isolated and an appropriate recovery procedure should be performed. There are several dedicated Internet resources containing large amounts of Exchange troubleshooting related information:

- Microsoft TechNet: <http://www.microsoft.com/exchange/techinfo/default.mspx>
- MExchange.ORG: <http://www.msexchange.org/>

It is impossible to describe any possible issues regarding to data restoration. Few most typical situations will be discussed in the next sections.

3.4 Choosing between Hotbackup and VSS Online Backup Options

Drive Backup provides two options for snapshot-based online backup in Windows XP, 2003 and Vista. In Windows 2000 and Windows NT4.0, only Hotbackup option is available.

This version of Drive Backup limited support of Exchange versions that precede Exchange Server 2003. It is recommended to use Drive Backup Enterprise with Exchange Server 2003 installed on Windows Server 2003 host operating system. In addition, it is recommended to update Exchange 2003 to Service Pack 2 and Windows 2003 to Service Pack 1.

For any version of Exchange Server, Exchange data will be successfully backed up in offline mode, by using an offline backup as well as any online backup option (HotBackup, MS VSS). However it requires to

temporary dismount Exchange databases, in Exchange 2000 and 2003 a whole Storage Group should be dismounted. So that during offline backup some Exchange resources are unavailable.

To avoid Exchange blackouts, online backup options should be used. Underlying technologies (Hotbackup and MS VSS) are different by their concepts and features. As regards to backing up Exchange Server databases, these options exhibit different levels of operation stability and resulting data consistency. Following considerations can be taken into account when choosing between options:

- Exchange Server 2003 has obvious benefits from VSS based backup. In particular, it automatically performs transaction log truncation after a successful completion of full backup of a Storage Groups. It is strongly recommended to use the "Microsoft Volume Shadow Copy" online backup option for backing up Exchange Server 2003 databases.
- Only Exchange Server 2003 is VSS compliant while earlier versions are not. For earlier versions of Microsoft Exchange Server, the difference between online backup options is generally unimportant.
- VSS option allows creating coherent backup of multiple volumes. With Hotbackup option, Drive Backup always performs an asynchronous backup of multiple volumes. This feature can be determinative for backing up databases distributed over multiple volumes.
- Hotbackup requires less memory and disk resources to operate. For example, latest updates of VSS (for Windows XP and 2003) need at least 300MB of disk space per every created shadow copy (at the moment of snapshot initialization), while Hotbackup will consume disk space only under noticeable disk I/O traffic. It will take 300MB only under high disk load.
- Both VSS and Hotbackup may fail to initialize or fail to maintain a created snapshot under conditions of high I/O traffic on a volume being backed up. VSS requires large amount of free space to maintain a snapshot, while Hotbackup mostly needs high backup performance compared to Exchange load. The backup performance depends on (a) throughput of archived volumes and backup storage and (b) compression throughput, which in turn depends mostly on CPU and memory speed.
- Hotbackup initialization is more stable in comparison to VSS. As concerns to Exchange, Hotbackup practically always initializes within predefined time intervals, while VSS may fail under high load.
- Only latest versions of VSS are stable enough. It is highly recommended to install the latest service pack available and also check for latest hot fixes for Windows XP/2003 related to VSS.

The tests have revealed a general instability of VSS initialization under a "stressed load" of the Exchange Server. To solve such a problem, either use the Hotbackup option or revise your maintenance plan in order to schedule the backup jobs at less stressed periods.

3.5 How to Use DBE for Protecting Exchange Data

3.5.1 How to Redistribute Exchange Data

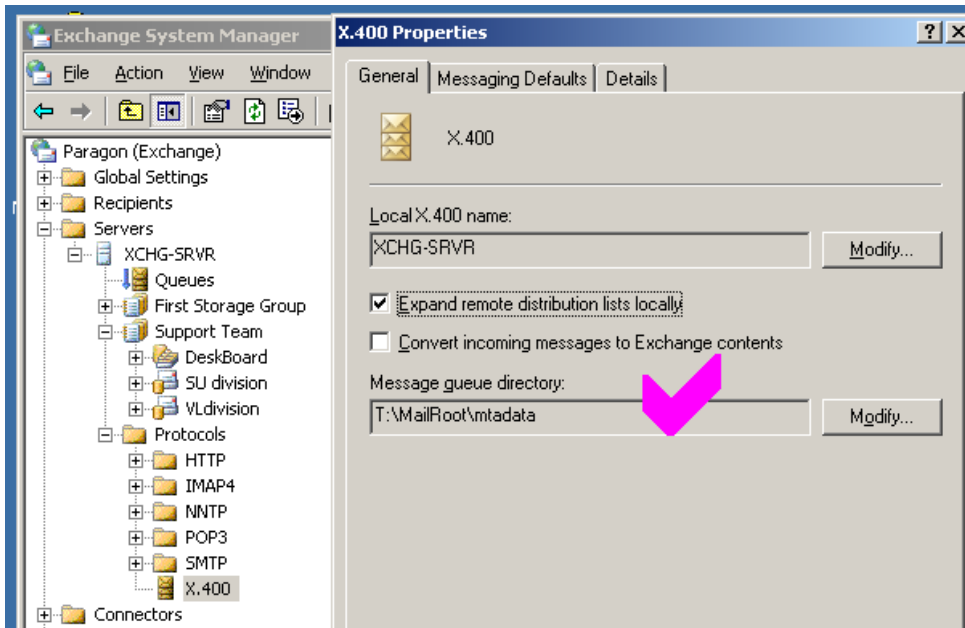
This section demonstrates how to change location of a Storage Group files. Such procedure can be applied to a newly created Store or Storage Group in order to move it on a dedicated production storage device (e.g. RAID or a dedicated hard disk).

Example 1: Moving mail queue directories.

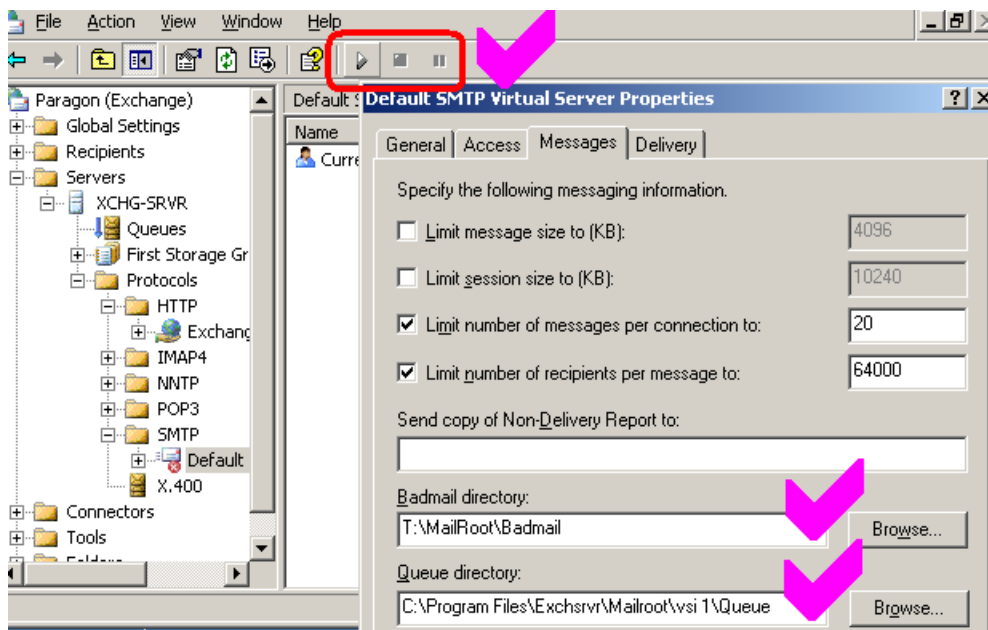
It is recommended to relocate mail queues to a dedicated volume.

1. Run Exchange System Manager (Start → Programs → Microsoft Exchange → System Manager)
2. Expand the "Servers → <Exchange_server_name> → Protocols" node in the tree view.
3. Select the "X.400" node, invoke the pull-down menu and select the "Properties" item.

- Change the "Message Queue directory" value in the X.400 Properties dialog. Press the "Modify..." button; enter the new directory name and press OK button to start the process. It may take a while.



- Then expand the "SMTP" node. You will see the list of virtual SMTP servers supported by Exchange. Select every desired virtual SMTP server, invoke the pull-down menu and select the "Properties" item.
- Stop the selected virtual SMTP server for enabling mail queues re-location. Invoke the popup menu and select the "Stop" item, or press the small black square button on the Toolbar [■].



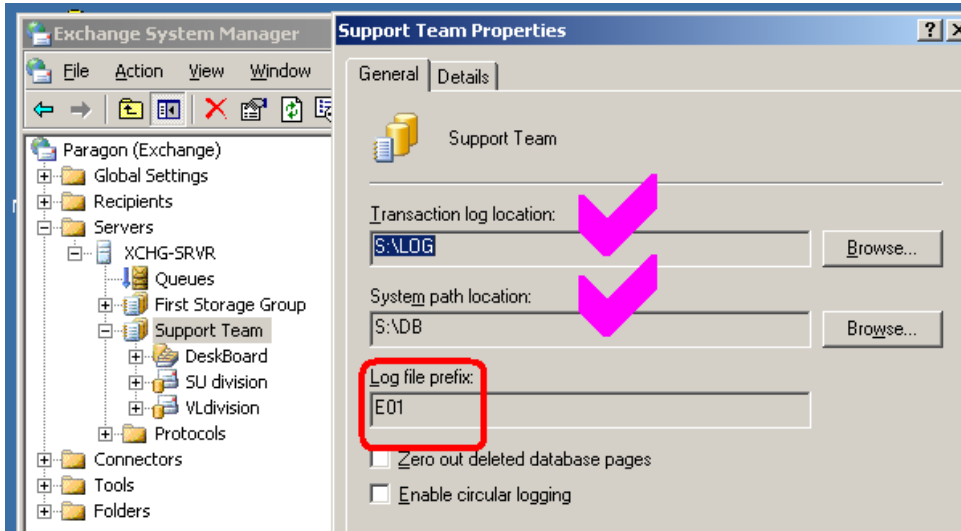
- Change "Badmail directory" and "Queue directory" values in the "SMTP Properties" dialog. For each of them, press the "Modify..." button and select the new directory. Press OK button to start the data movement process. The operation may take a while, especially in case Badmail directory contains a lot of undeliverable mail files.
- Re-start the virtual SMTP server. Invoke the popup menu and select the "Start" item, or press the small black triangle button on the Toolbar [▶].

Example 2: Moving transaction log files and E##.CHK of a Storage Group.

It is recommended to locate transaction logs on a separate volume on a dedicated fault-tolerant storage (RAID-1, RAID-10) for protecting transaction log from damage. The volume capacity should be at least 2-3 times larger than the planned maximum log capacity – this space is required for the log re-play at


post-disaster database recovery procedure. Anyway, it is strongly recommended at least to isolate the log files from the Windows system volume.

1. Run Exchange System Manager (Start → Programs → Microsoft Exchange → System Manager)
2. Expand the "Servers" → <Exchange_server_name> node in the tree view.
3. Choose the node that corresponds to the Storage Group of interest, invoke the pull-down menu and select the "Properties" item.
4. In the Storage Group Properties dialog, change the "Transaction log location" and the "System path location" directories (this one holds E\$.CHK and TEMP.EDB).



The so-called *system prefix* is displayed on this page, which presents in the beginning of filenames for log files and checkpoint file of that Storage Group.

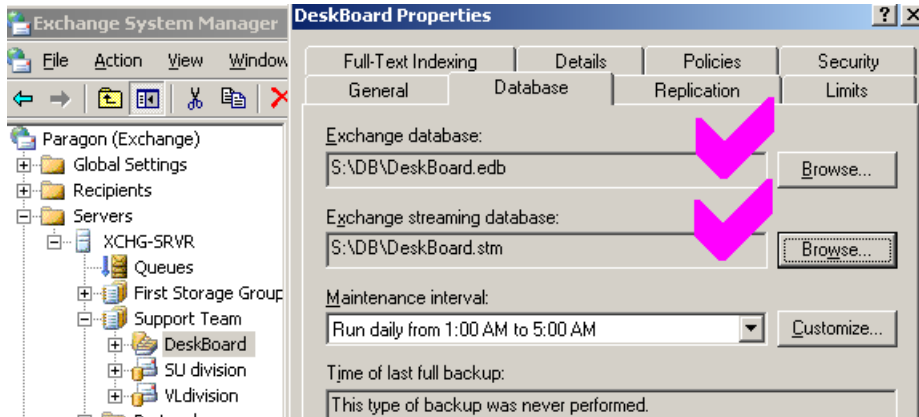
5. After new directory names have been selected, press OK button in order to start data re-location.
6. Exchange will display the warning message that the Storage Group will be temporarily dismounted (unavailable) to complete data re-location. Press "Yes" button to confirm.
7. Exchange will display the data movement progress. Time required for moving these files depends on amount of log files. For a newly created Storage Group, it takes only few seconds.



Example 3: Moving database files of a Store:

It is recommended to locate database files on a high production and fault-tolerant storage (RAID-5, RAID-10) for protecting data from disaster. The volume should be large enough to maintain planned databases growth. Anyway, it is strongly recommended to isolate production databases from the Windows system volume.

8. Run Exchange System Manager.
 1. Find the appropriate Storage Group in the tree view:
 "Servers" → <Exchange_server_name> → <Storage_Group_name>
 2. Expand the Group node, select the subnode that corresponds to the Store of interest, invoke the pull-down menu and select the "Properties" item.
 3. Go to the "Databases" tab in the Store Properties dialog. Change "Exchange database" and "Exchange streaming database" values.



4. After new directory names have been selected, press OK button in order to start data re-location.
5. Exchange will display the warning message that the Storage Group will be temporarily dismantled (unavailable) to complete data re-location. Press "Yes" button to confirm.
6. Exchange will display the data movement progress. Time required for moving these files depends on amount of stored data. For a newly created Store, it takes only few seconds. For a long-played database, the operation may take a long while.

It is recommended to isolate database files (.EDB and .STM) from transaction log files. The *checkpoint file* (E\$.CHK, stored in the "System path" directory for every Storage Group) needs not to be isolated from other components. The "System path" directory can be assigned to a directory containing log files or to a common directory for database files, for example. The checkpoint plays important role at the log re-play process. It usually is restored together with "old" log and database files from a backup image.

3.5.2 How to Backup Exchange Data

As was it mentioned above, the current version of Drive Backup is unaware of Exchange Information Store structure and only allows to backup volumes. This section demonstrates how to correctly choose a set of volumes to be backed up in order to protect a Storage Group of interest.

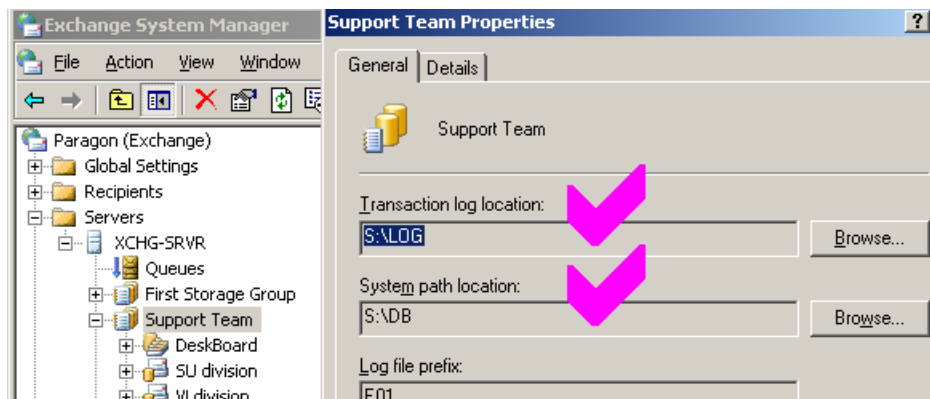
Conditions: The Exchange Information Store has been already configured and existing Storage Groups will not be changed for a long period. There is the "Support Team" Storage Group that should be backed up regularly. The "Support Team" Storage Group is completely located on the volume [S:].

Such Storage Group layout is not recommended for Exchange data because of low performance and low degree of reliability; it can be used only in case of tight budget. We will discuss it only for illustrating basic techniques.

Purpose: Set up correct parameters for the task of a full backup of the "Support Team" Storage Group. Then create an automated scheduled task that will make a full backup of the Storage Group.

Stage 1: Gathering information:

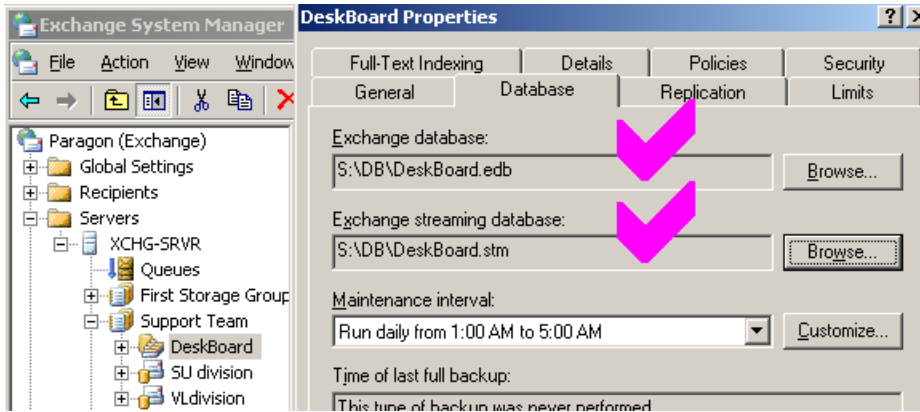
1. Run Exchange System Manager and find the required Storage Group in the tree view: "Servers" → <Exchange_server_name> → "Support Team"
2. Invoke the popup menu for that tree node and open the Properties dialog.



Remember volumes where "Transaction log" and "System path" directories are located. Then close the Properties dialog.

In this example, the "Support Team" Storage Group keeps transaction logs in the "S:\LOG". The "System path" directory is "S:\DB".

- 3. Expand the Storage Group node.
- 4. Select the first Store subnode and open Properties dialog for it.



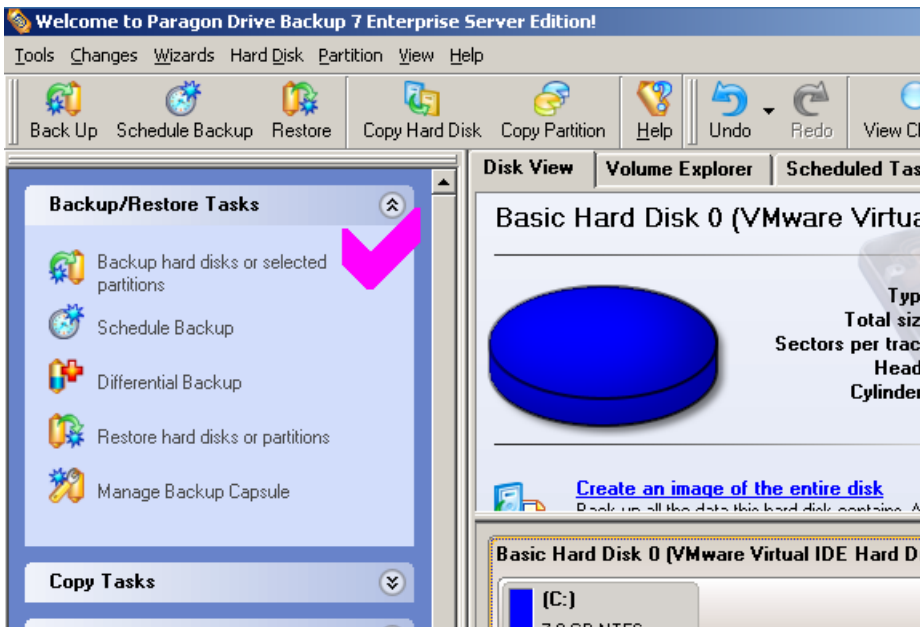
Remember volumes where the "properties data" and the "streaming data" files are located. Then close the Properties dialog.

- 5. Repeat the previous step for every Store in the "Support Team" Storage Group.

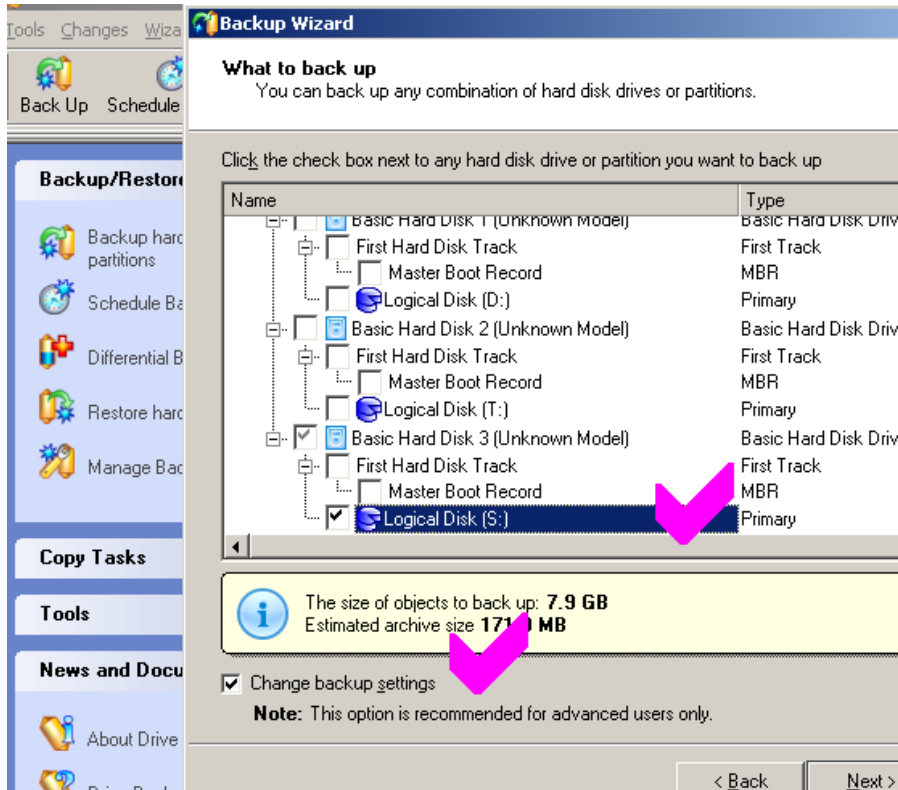
In this example, the "Support Team" Storage Group consists of three Stores ("DeskBoard", "SU division" and "VLdivision"), which all are placed in the "S:\DB". Correspondent database files are: DeskBoard.EDB, DeskBoard.STM, SUdivision.EDB, SUdivision.STM, VLdivision.EDB and VLdivision.STM.

Stage 2: Setting up backup properties:

- 6. Run Drive Backup and invoke the Backup Wizard.

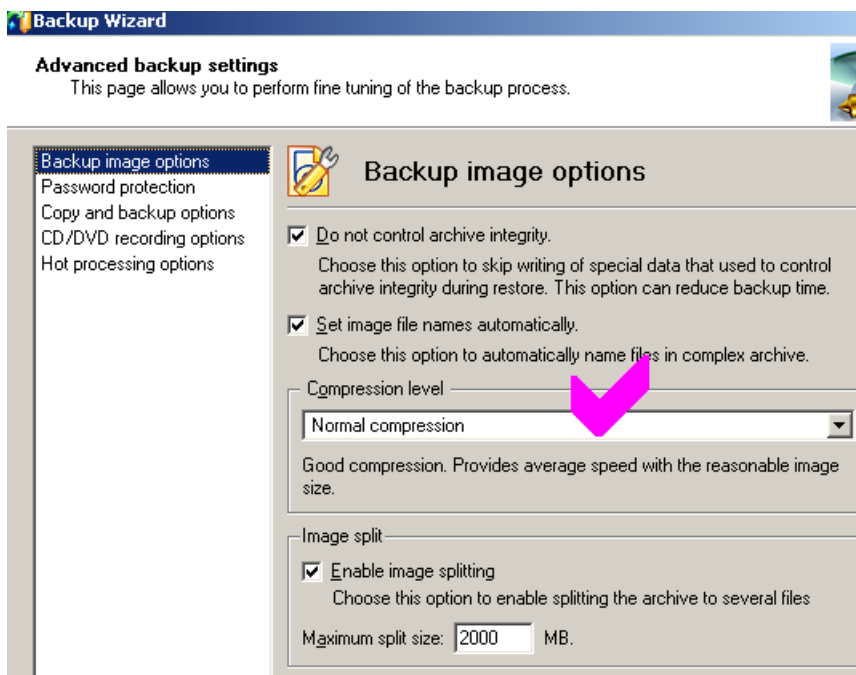


- 7. In the tree view of disks and volumes, select all volumes that were selected in the Stage 1. **Important!** There is the "Change backup settings" checkbox on the bottom of the screen. Set this checkmark in order to control the backup settings for that task. Then press "Next" button.

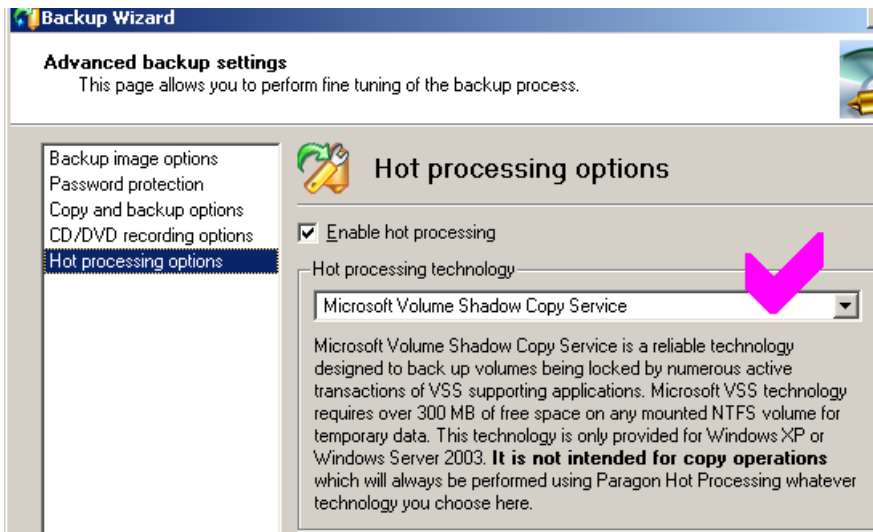


In this example, the whole "Support Team" Storage Group is located on the dedicated volume [S:]. So that at least the volume S: must be selected in the tree view.

- 8. On the "Backup Settings" screen, set the backup parameters. For the absolute most of them default values are optimal. We will change only two of them in this example:
- 9. Go to the "Backup image options" page and set the "Normal" compression level. It usually provides good compression rate for Exchange data (3:1 – 6:1) with harmless speed deceleration.



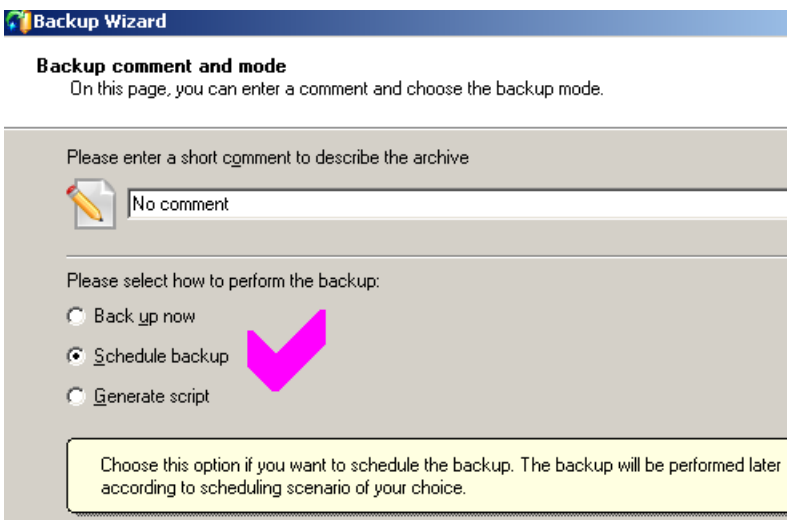
- 10. Go to the "Hot processing options" page and select the "Microsoft Volume Shadow Copy Service" item in the "Hot processing technology" pull-down list.



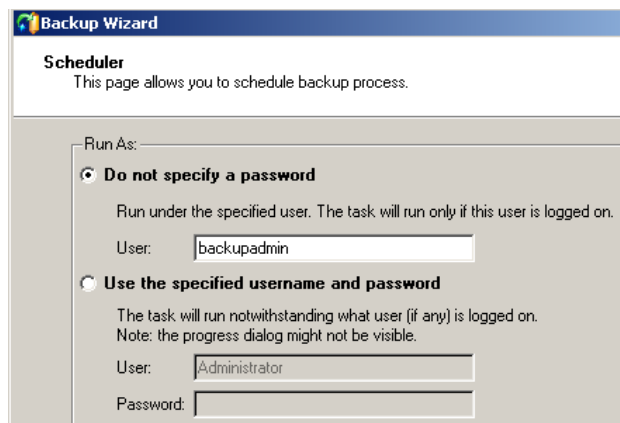
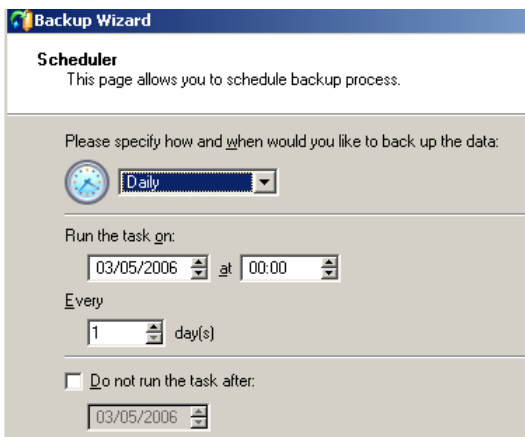
- 11. On the two next screens, choose the "Save data on the local/network disk" and select location for storing backup images.

Stage 3: Creating the scheduled task:

- 12. On the next screen, choose "schedule backup" to create a scheduled backup task.



- 13. Customize the schedule parameters.



At this moment, the program generates a script file containing appropriate commands for back up selected volumes and stores it in the "Scripts" subdirectory of the Drive Backup's installation directory. The script is intended for executing by the so-called Script Interpreter utility (SCRIPTS.EXE – for Windows environment).

Few notes:

- The standard Scheduled Tasks service is used to run the Script Interpreter. Scheduled Tasks should be enabled in order to execute the automated backup task.
- For smoothly running a scheduled backup in a system where no one has been logged on, Drive Backup creates a special local user account, which is a member of the local Administrators group. This account should be created and enabled to log on locally otherwise scheduled backup tasks may fail to run.

The automatically generated script always uses the same filename(s) for storing a backup image. In other words, every time a scheduled backup task is run, it overwrites the previously created backup image. Either you must take care of regularly saving successfully created backup images or you should enhance the script.

Paragon Scripting Language allows to program sophisticated scripts with smart behavior. As concerns to this example, the so-called "cyclic backup" script is more convenient for this purpose. The "cyclic backup" script uses a limited-length stack of last backup images. As soon as the new image is created, the oldest image in the stack is deleted. So that the "cyclic backup" script guarantees availability of the last good backup image and implicitly controls the total capacity of backup images stored on the disk.

3.5.3 How to Restore Exchange Data

This section demonstrates how to correctly restore a Storage Group from a previously created backup image.

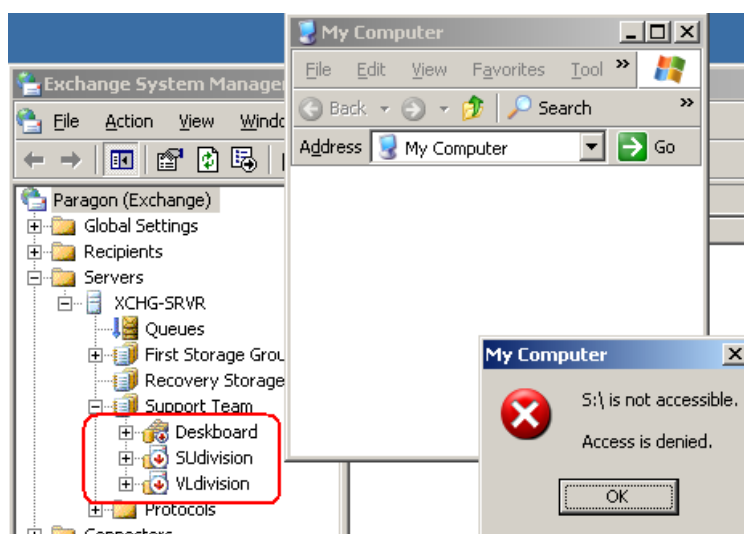
Conditions: The "Support Team" Storage Group was backed up by using the MS VSS online backup option. Later, the production volume was corrupted and databases became unavailable. Transaction logs of the Group are not damaged. The host operating system is still running. The production volume has been fixed against filesystem errors but some databases were irreversibly corrupted.

Purpose: Restore the "Support Team" Storage Group from the full backup image and bring the whole Group to a workable state.

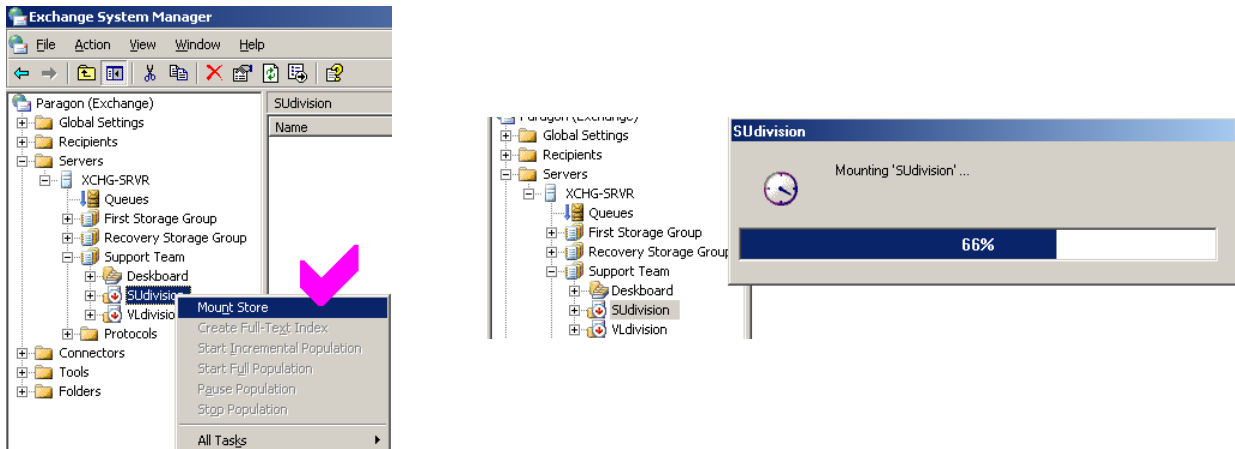
The "Support Team" is the second Storage Group in the Exchange Information Store. Its system prefix is E01. The Storage Group is entirely isolated on the drive [S:]. Transaction logs are located in "S:\LOG" and database files (.EDB and .STM) were located in "S:\DB" directory.

Stage 1: Try the soft recovery the databases:

1. Run Exchange System Manager and open the required Storage Group in the tree view:
"Servers" → <Exchange_server_name> → "Support Team"
2. Expand the Group node. Dismounted or corrupted Stores are marked by a small circle with the red down-arrow on the icon.



3. Try to mount dismounted databases normally. Select a Store subnode, invoke the popup menu and select the "Mount Store" menu item. Wait until Exchange mounts the Store or displays an error message. The mount routine may take a long while.



4. In case **all** dismounted databases are started normally, you fortunately need not a recovery. Just let Exchange working and make a backup of the Storage Group.
5. If some Stores were not started, Exchange will display the error message. In this case, you really need to restore data from the backup.

Stage 2: Dismount the Storage Group:

Dismount all Stores nested in the selected Storage Group. Exchange allows restoring only offline databases. You must dismount all Stores because transaction log is shared between all Stores within the Storage Group.

6. Select the Store node, invoke the popup menu and select the "Dismount Store" menu item. Wait until Exchange dismounts the Store. The process may take a while.
7. Repeat the previous step for all Stores within the Storage Group.

Remember which Stores were not affected by a failure. These Stores do not necessary to be restored from a backup image, and you can reduce the recovery time by skipping their files at restore.

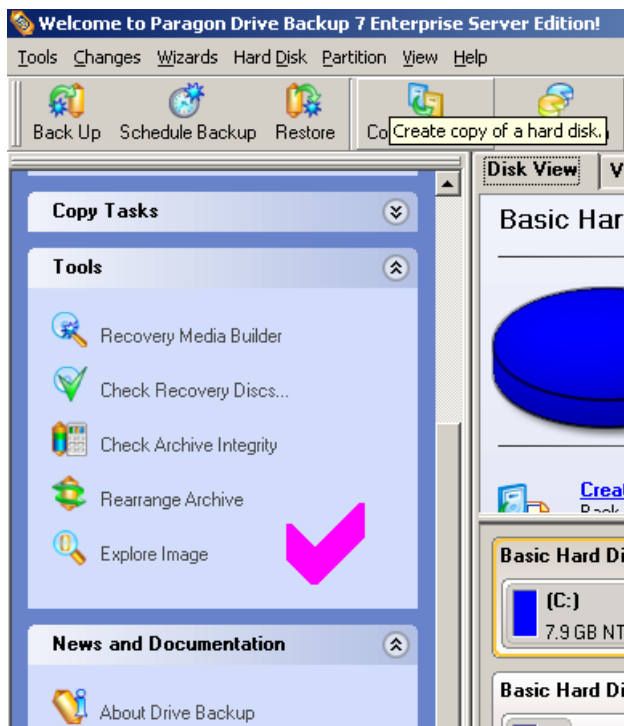
Stage 3: Extract required files from the backup image. Generally you need to extract damaged EDB and STM files, stored transaction log files and the stored .CHK file.

8. It is strongly recommended to make a copy of all currently available log files prior to starting next steps.
9. Perform the filesystem check on volumes where transaction logs and database files are placed.

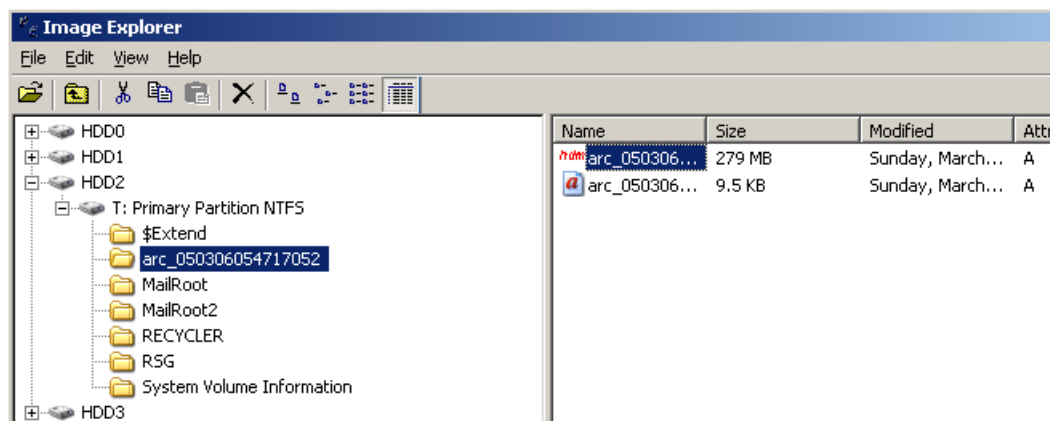
Data recovery requires volumes to be in healthy state. In this example, the "Support Team" Storage Group is located on the volume [S:]. Run the "CHKDSK S:" command to check filesystem or "CHKDSK S: /F /R" to check and fix filesystem state.

10. Find the last backup image and place it on a local volume or on a shared network resource. It is the required step in case the image was stored on removable media.

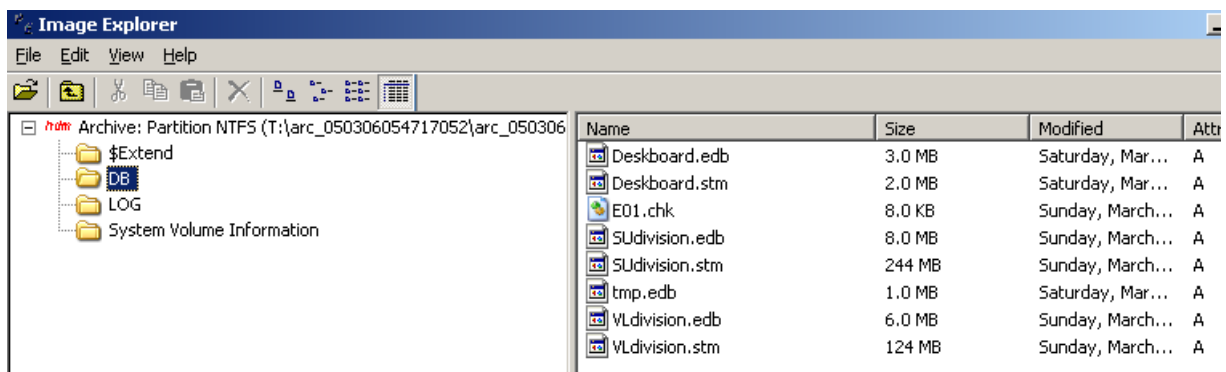
11. Run Drive Backup and start the Image Explorer tool.



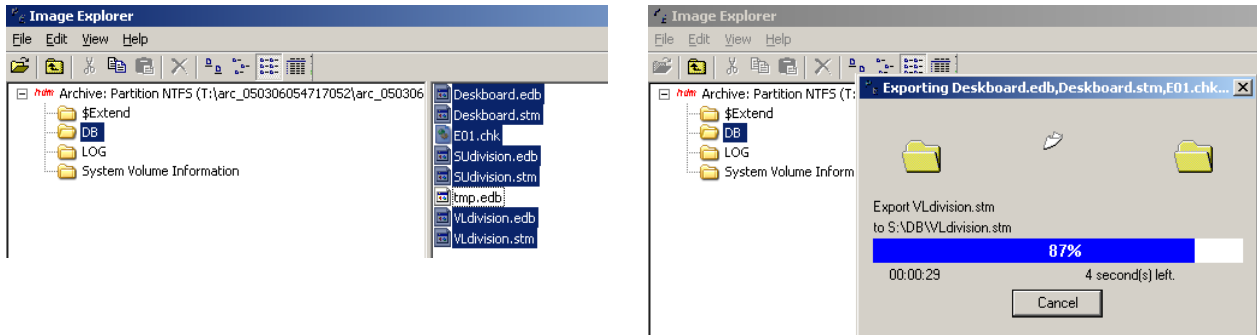
12. Open the backup image in Image Explorer:
(menu:) File → Load Archive



13. Browse the image and find database files (.EDB and .STM) of damaged Stores.



14. Select all .EDB and .STM files, invoke the popup menu and select the "Export" item and extract them to their original location. An export of multiple files and directories is supported.



In case all databases were corrupted, simply export the whole "DB" directory. If most databases were not damaged, select only files that were corrupted (to reduce recovery time).

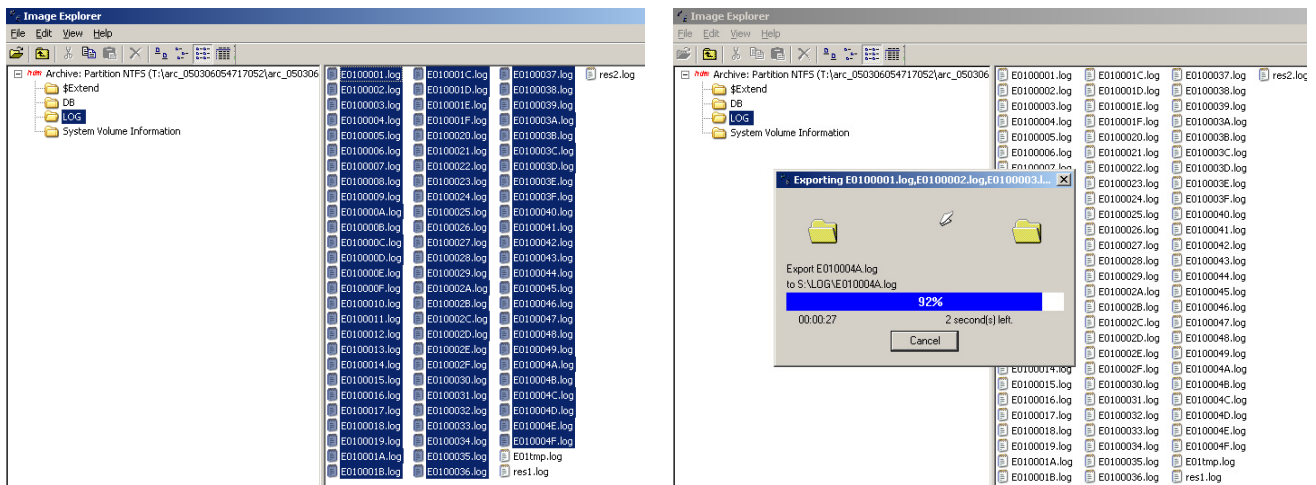
To correctly export a directory from an image:

- o Choose a directory in the backup image.
- o Invoke the popup menu and select the "Export" item.
- o In the "Select directory" dialog, choose the directory that is the parent for the exported one.

For example, to export the whole "S:\DB" directory from the backup image to its original location, you should select the drive "S:" in the "Select directory" dialog.

15. Browse the image for stored transaction log files. Extract stored log files to their original directory:

- o All E\$\$\$###.LOG files *must* be extracted (E01###.LOG in this example).
- o E\$.LOG *should not* be extracted (E01.LOG in this example).
- o RES1.LOG, RES2.LOG and E\$tmp.LOG need not to be extracted.



16. Inspect the resulting set of transaction log files.

The hexadecimal numbers of log files, which positions in log filenames are designated here by the "#" symbol, must constitute the consecutive set of hexadecimal numbers (example: 00123...00129, 0012A, 0012B...0012F, 00130...00139, 0013A...0013F, 00140...0014A).

It is an optional step. When trying to soft recover and mount a Store, Exchange will perform all checks by self and will report about detected errors. However, this step allows ensuring that you chose a right backup image to restore and immediately recognize situations when the full log re-play cannot be completed.

- 17. Delete the checkpoint file of the corrupted Storage Group (S:\DB\E01.CHK in this example).
- 18. Extract the old E\$.CHK file from the backup image, to its original location (S:\DB\E01.CHK in this example). The TEMP.EDB file needs not to be extracted.

The checkpoint file stores the separator position between the flushed and not-flushed-yet logged data. As concerns to the restore routine:

- The "new" checkpoint reports to the Exchange database engine (ESE) that most data are flushed to the database, but it is not true for the restored old database files.
- The old checkpoint file contains correct information about flushed and logged data.
- Finally, in case there is no checkpoint file at all, ESE will re-play all available log files against the databases. It will enlarge recovery time, however in some cases it may be the only way to revive Exchange databases.

19. Close Image Explorer and Drive Backup. Close any file managers and windows of the Windows Explorer that browse directories with just restored Exchange databases and log files as well.

The thing is that Exchange must have the full access rights to directories containing transaction logs and databases for successful completion of the recovery routine (see MS KB articles ID 823022 and 896143). In practice, opening these directories by other applications can impede Exchange to complete the recovery. Closing applications eliminates possible impact to the Exchange data recovery process.

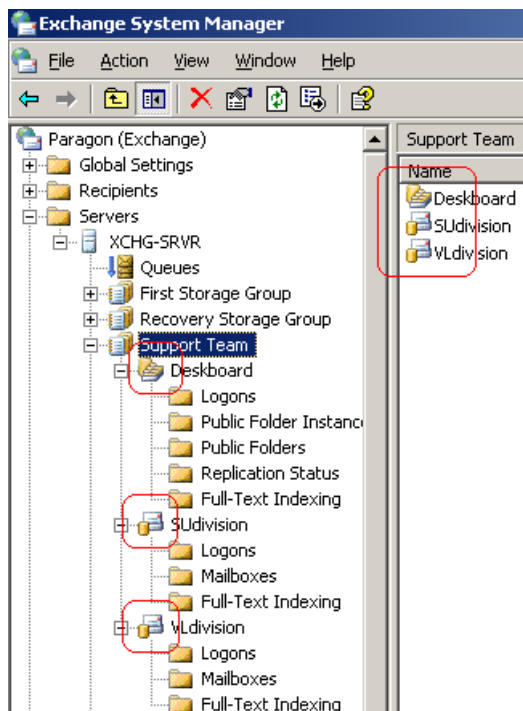
Stage 4: Invoke the soft recovery routine of Exchange databases.

20. Run Exchange System Manager and open the restored Storage Group in the tree view:
"Servers" → <Exchange_server_name> → "Support Team"

21. Expand the Storage Group node. You will see the list of all Stores nested in this Storage Group. Try to mount every Store by invoking the popup menu and selecting the "Mount Store" item.

For every mounted database Exchange will perform the *soft recovery* procedure, which is the selective log re-play for the selected Store. The process usually takes a long while (see the step #3).

22. Restoration is successful only in case all Stores within the Storage Group are mounted smoothly:



In case all Stores in the Storage Group were mounted, it can be involved to the normal operation.

If this was not the case, the further troubleshoot activity should be performed in order to fix damaged Exchange databases.

Stage 5: Perform the online backup of the successfully restored Storage Group.

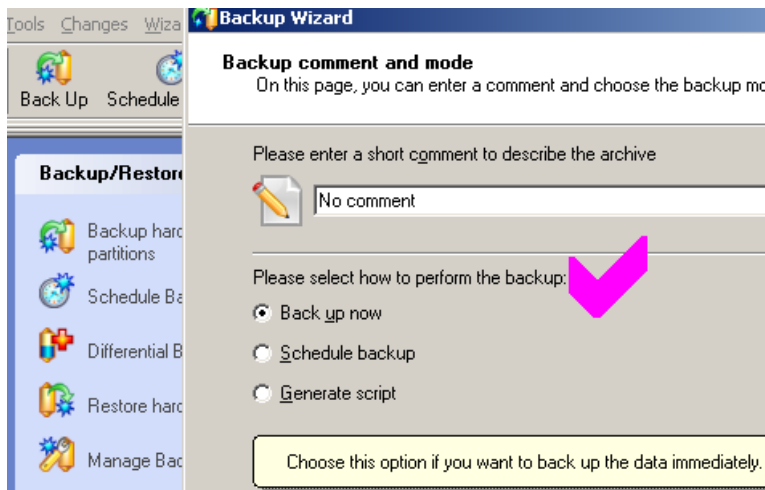
This step is needed to delete unnecessary transaction log files of the "Support Team" Storage Group.

23. Run Drive Backup and invoke the Backup Wizard (see step#6 in the previous section).

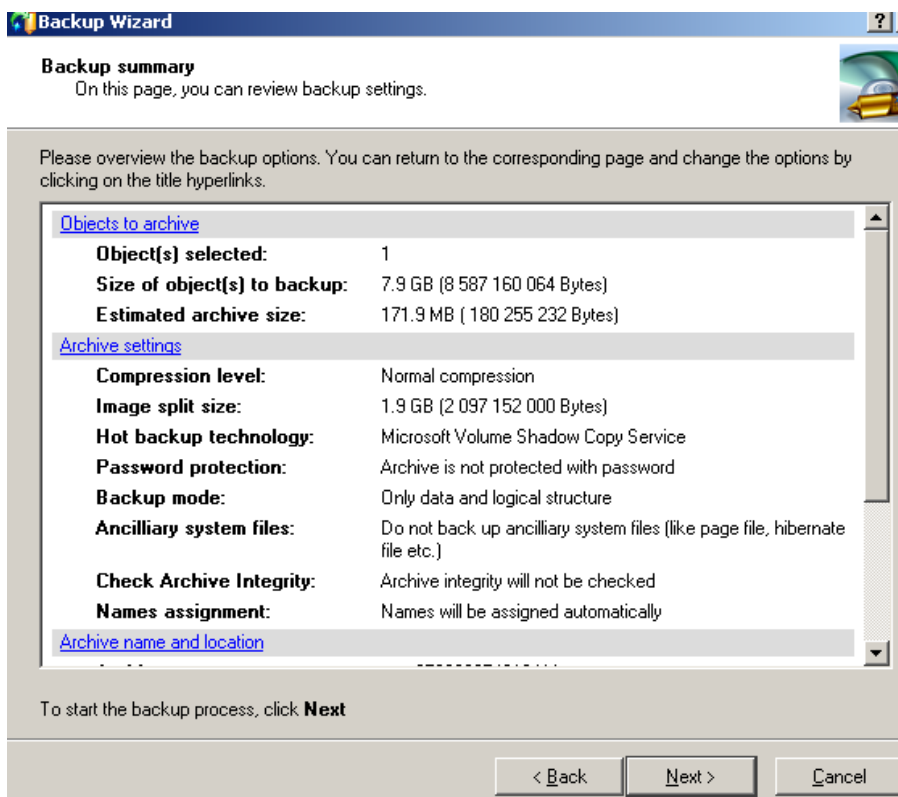
24. Select all volumes that contain databases and logs of the "Support Team" Storage Group, in the tree view of disks and volumes (see step#7 in the previous section). Set "Change backup settings" checkmark in order to control the backup settings for that task. Then press "Next" button.

In this example, the whole "Support Team" Storage Group is located on the dedicated volume [S:]. So that at least the volume S: must be selected in the tree view.


- 25. On the "Backup Settings" screen, set the backup parameters (see steps #8-10 in the previous section). Go to the "Backup image options" page and set the "Normal" compression level. Then go to the "Hot processing options" page and select the "Microsoft Volume Shadow Copy Service" item in the "Hot processing technology" pull-down list.
- 26. On the two next screens, choose the "Save data on the local/network disk" and select location for storing backup images.
- 27. On the next screen, choose "Back up now" to create a run-once backup task.



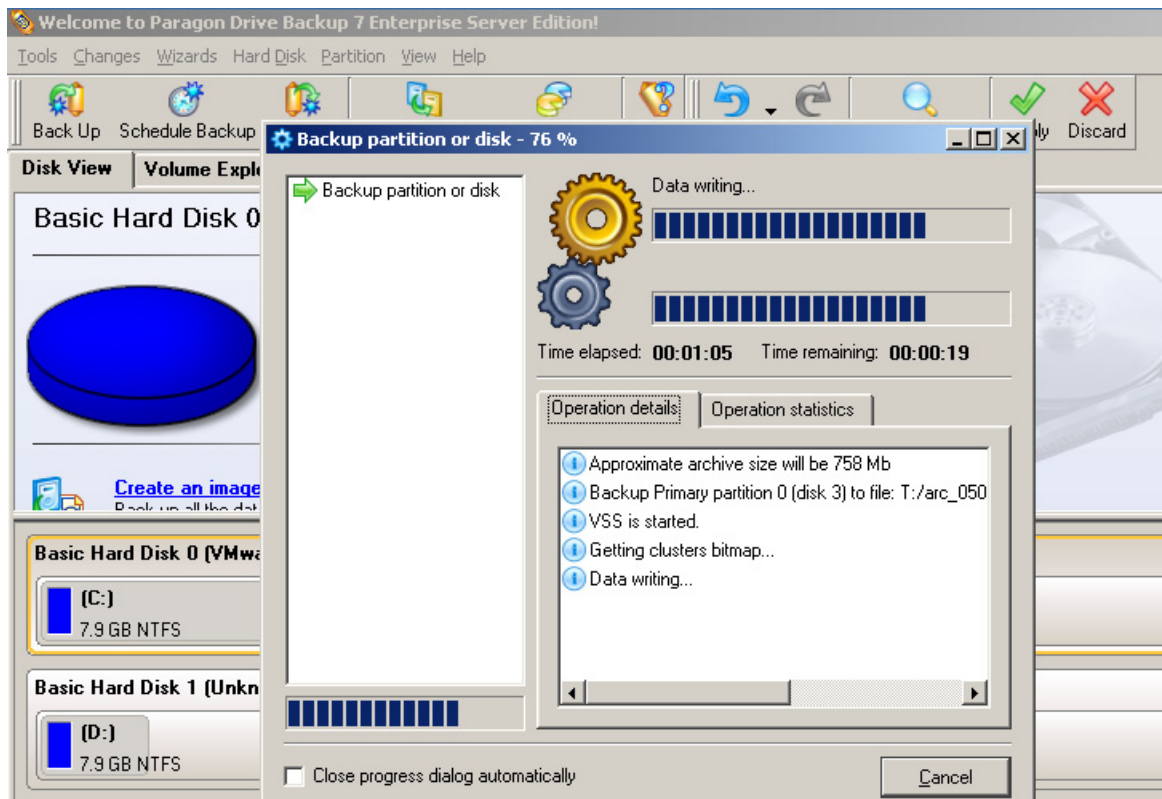
- 28. On the next screen, you can inspect parameters of the created backup task.



29. Drive Backup creates the virtual backup task but does not starts it immediately. The program allows you accumulating multiple backup tasks and disk management operations in a single batch. To play the job batch, either select the "Changes → Apply Changes" item in main menu or press the "Apply" button on the Operations toolbar.



30. The program will back up the selected data in the interactive mode. After the backup operation completes, Exchange transaction log will be truncated.



3.5.4 How to backup and restore Exchange databases distributed across multiple volumes

Exchange databases are usually not fit within a single volume but spread across multiple volumes. As it was mentioned in previous sections, this allows to improve the Information Store reliability and performance.

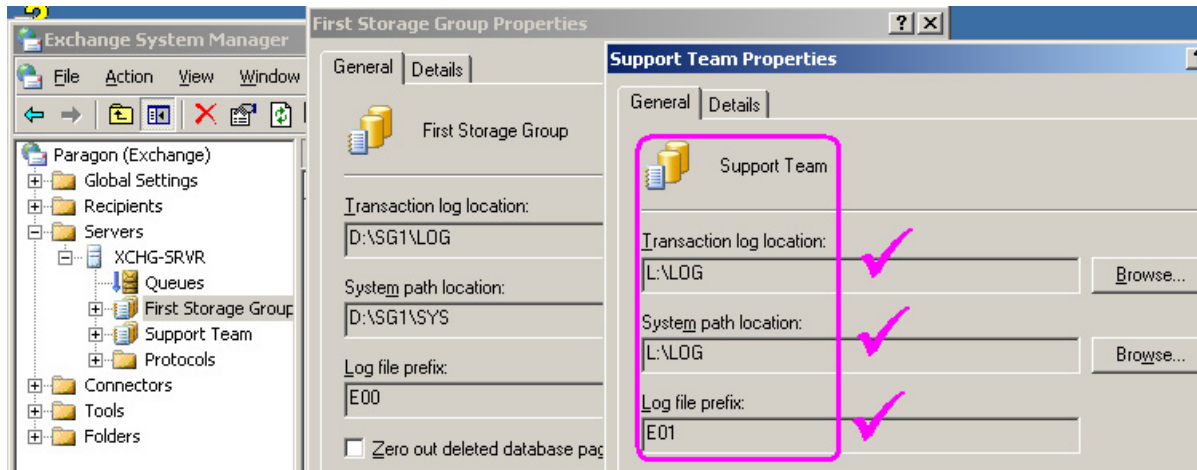
This section demonstrates how to deal with distributed Exchange databases.

Conditions: The "Support Team" Storage Group is spread across three volumes (L:, S: and V:).

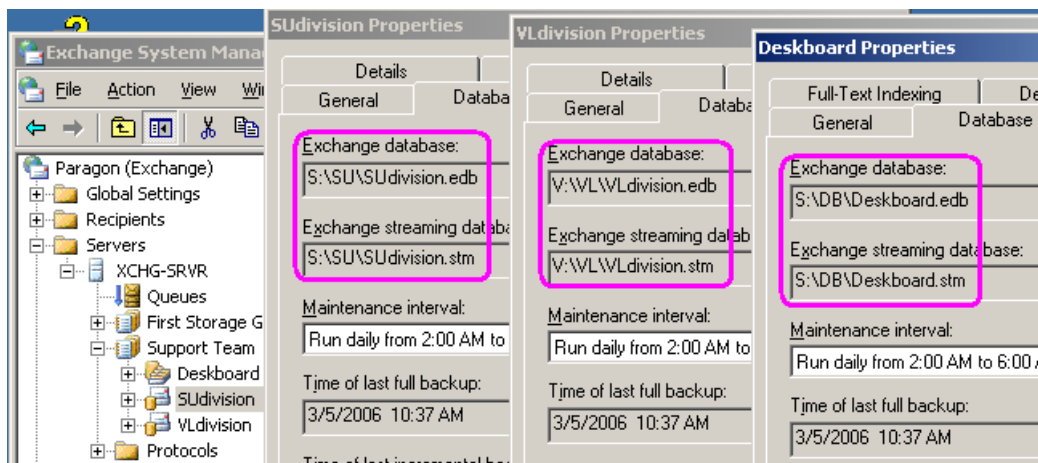
Purpose: Perform a correct backup of the "Support Team" Storage Group. Then, perform a correct restore of it.

Stage 1: Gathering information:

1. Run Exchange System Manager and find the required Storage Group in the tree view: "Servers" → <Exchange_server_name> → "Support Team"
2. Inspect properties of the "Support Team" Storage Group:



3. Inspect properties of databases included to the "Support Team" Storage Group:

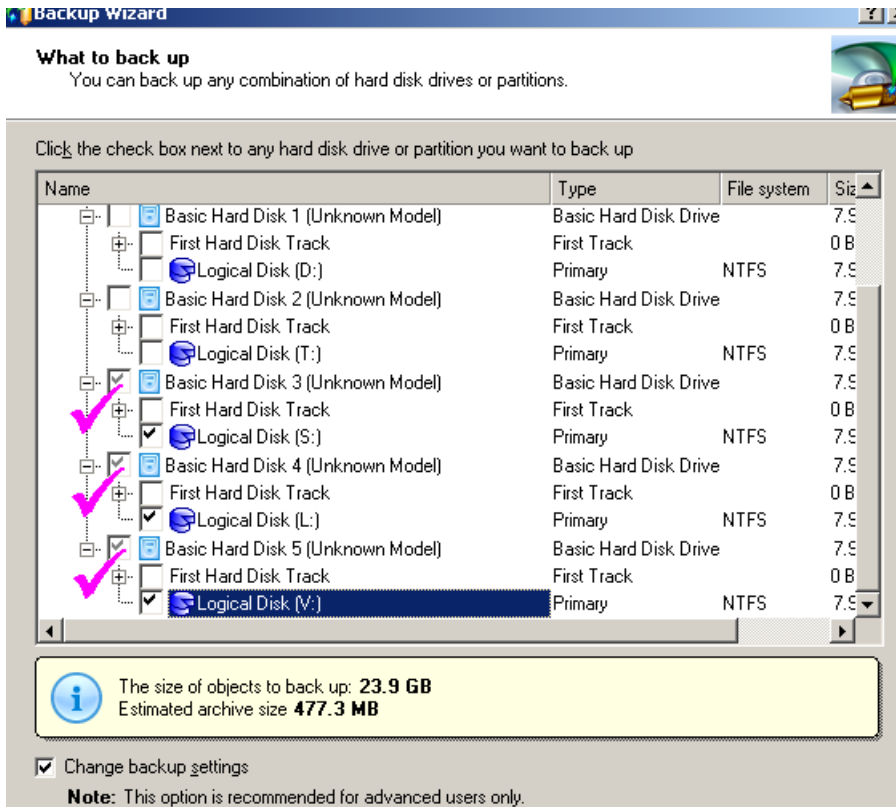


In our example, the "Support Team" Storage Group is spread across three volumes (L:, S: and V:):

- The "VLdivision" database is located on the [V:] volume.
- The "Deskboard" and "SUdivision" are located on the [S:] volume.
- The transaction log files of the Storage Group are located on the [L:] volume.

Stage 2: Setting up backup properties:

4. Run Drive Backup
5. Set up Drive Backup to use the "Microsoft Volume Shadow Copy" online backup option:
 - (menu:) Tools
 - ↳ Settings...
 - ↳ Hot processing options
 - ↳ Hot processing technology
 - ↳ Microsoft Volume Shadow Copy Service
6. Invoke the Backup Wizard. In the tree view of disks and volumes, select all volumes that were selected in the Stage 1. In this example, [L:], [S:] and [V:] volumes should be selected.



There is the "Change backup settings" checkbox on the bottom of the screen. You can set this checkmark in order to control the backup settings for that task.

7. Set up or re-define other backup parameters, if required. See the "How to Backup Exchange Data" section, steps #8-11 for more details.

Stage 3: Create the scheduled task, or run the backup operation once:

8. Set up the scheduled backup task. See the "How to Backup Exchange Data" section, stage 3 for more details.

For test purposes, you can choose a "Back up now" option instead of "Scheduled backup". In this case, the program will accumulate the backup operation in the queue of pending operations. You can add more operations to the queue. To execute all accumulated pending operations, press the "Apply" button on the toolbar, or select the "Changes → Apply Changes" menu item.

Now suppose the "Support Team" Storage Group becomes corrupted but an operating system and Exchange are still running. We need to restore data and bring the Storage Group to a workable state.

Stage 4: Try the soft recovery the databases.

The first you should try is to let Exchange to fix databases by self. See the "How to Restore Exchange Data" section, stage 1, how to force the soft recovery routine.

Stage 5: Dismount the Storage Group.

In case the soft recovery fails, dismount the damaged Storage Group before starting the data restoration from a backup image. See the "How to Restore Exchange Data" section, stage 2 for more details.

Stage 6: Restore data from the backup image.

9. At first, determine the actual state of transaction logs and databases. Results of this check allow to select the best data restoration mode for Exchange data.

In case transaction logs are not corrupted, you are able to restore and roll forward Exchange databases to a state very close to the moment of a disaster. This really means an almost complete elimination of data loss.

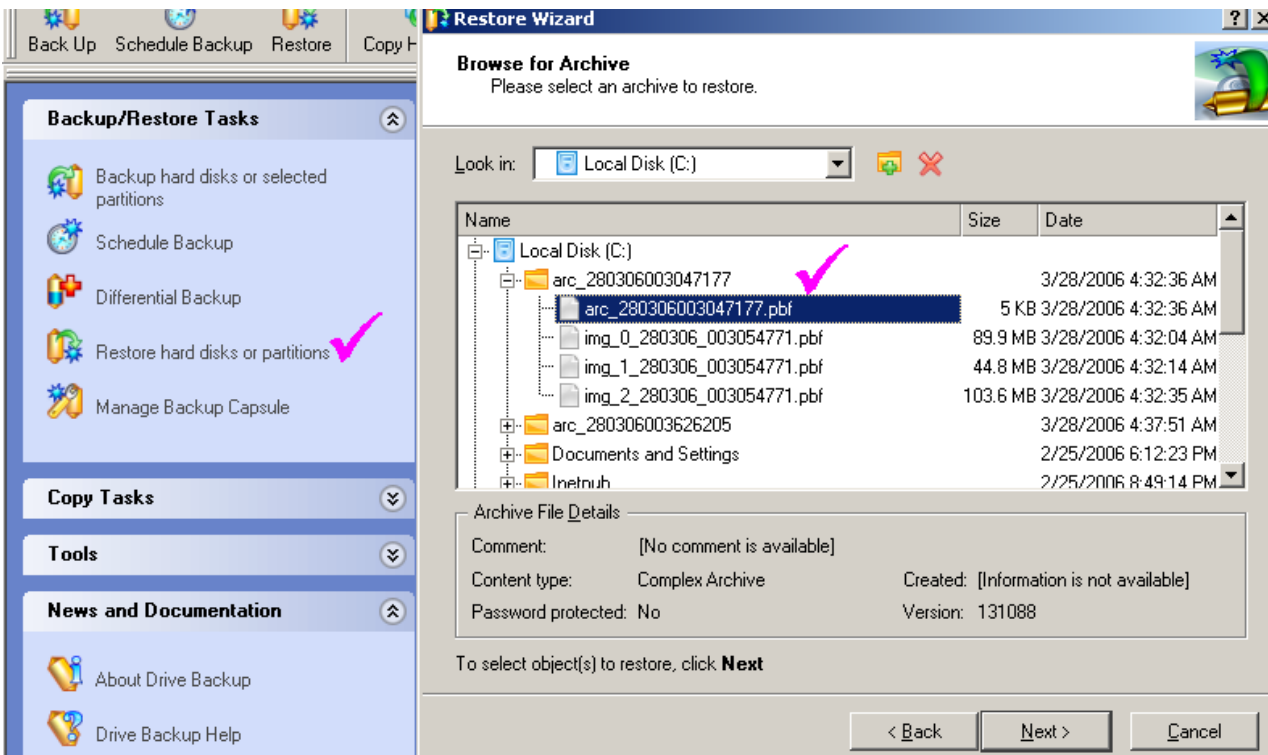
Otherwise, you can only restore Exchange databases to a state at the moment of the full backup; data modifications made between moments of the full backup and the disaster will be lost.

- 10. It is strongly recommended to make a copy of all currently available log files prior to starting next steps.
- 11. Perform the filesystem check on volumes where transaction logs and database files are placed. In this example, run the following commands:
 - o CHKDSK L: /F /R
 - o CHKDSK S: /F /R
 - o CHKDSK V: /F /R

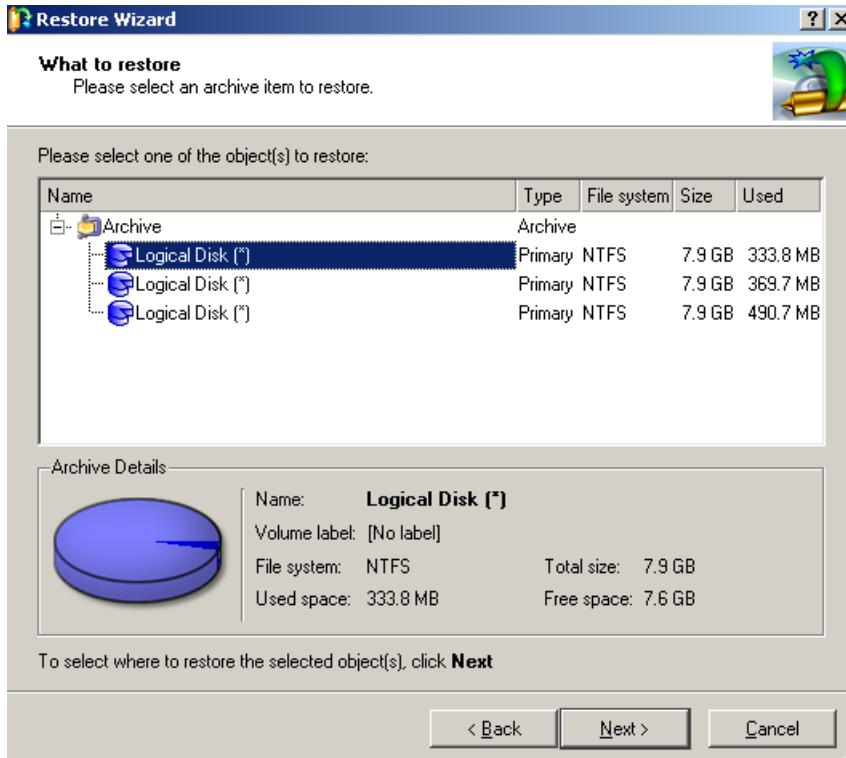
Remember, which components and volumes are not damaged. In case transaction logs are not corrupted, you can implement the "restore + roll forward" routine in order to eliminate data loss. In this case, you should apply file-level restoration to transaction log files. If transaction logs are corrupted, you will lose some amount of most recent data and data modifications. In this situation, a better way of data restoration depends on the actual state of database files.

We assume that transaction logs of the Storage Group were not damaged. In addition, we assume that there are no other databases and files on the volumes [S:] and [V:] than the databases, which is to be restored. In this case, the volumes [S:] and [V:] can be restored in volume-level restore mode.

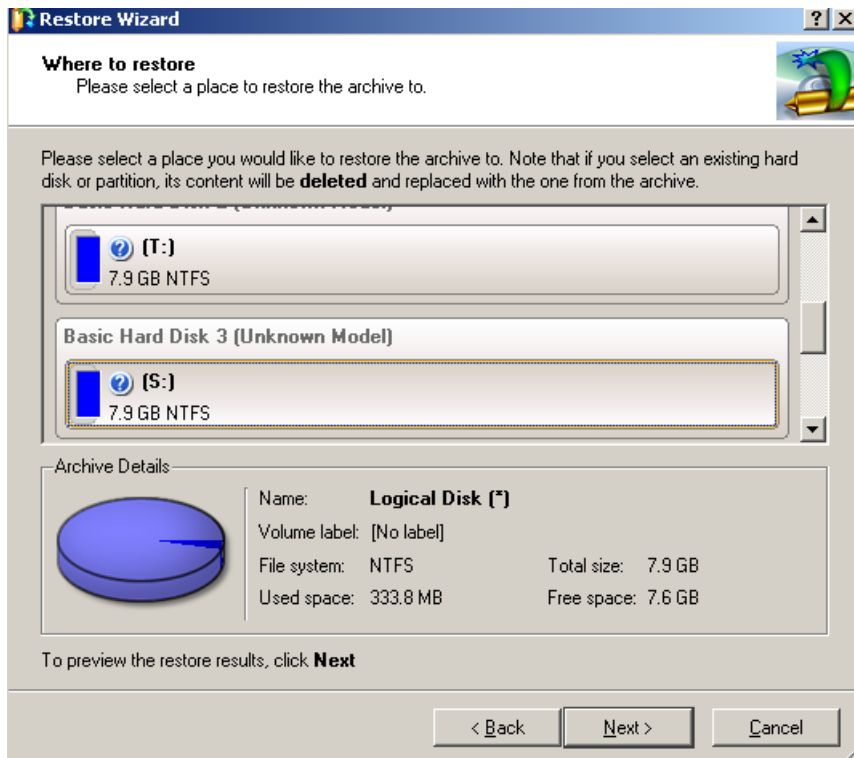
- 12. Find the last backup image and place it on a local volume or on a shared network resource.
It is the required step in case the image was stored on removable media.
- 13. Run Drive Backup, invoke the Restore Wizard and open the appropriate backup image:




- 14. Select a volume stored in the image:



- 15. Select a volume on a disk where data should be restored. By default, the program automatically suggests to restore an image over an original volume:



- 16. The program will accumulate the restore operation in the queue of pending operations. To execute pending operation(s), press the "Apply" button on the toolbar.



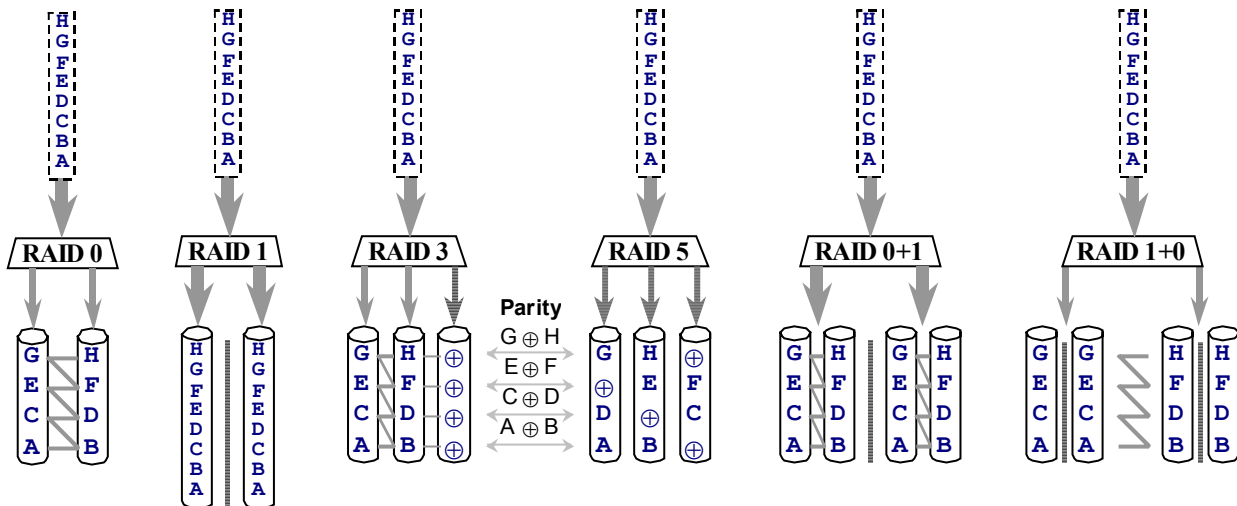
17. After restoration of volumes containing databases of the "Support Team" Storage Group, you should restore the transaction log files in the file-level mode. See the "How to Restore Exchange Data" section, stage 3 for more details.
18. Finally, Invoke the soft recovery routine of Exchange databases. See the "How to Restore Exchange Data" section, stage 4 for more details.

4 Appendix

4.1 RAID levels

Use of RAIDs is the most effective way to increase throughput and provide a hardware-level fault tolerance of a disk subsystem. Striped RAID sets provide highest read/write performance, which is found nearly a multiple of a single disk performance.

RAID level	Brief description	Amount of disks	Size factor	Performance factor	Fault tolerance
0	Striped Disk Array	$N \geq 2$	$\times N$	$\times N$	No
1	Mirrored Disk Array	$M \geq 2$	$\times 1$	$\times 1$	Yes
0+1	Mirroring of striped segments (RAID 1 over RAID 0)	$M \cdot N \geq 4$ $M \geq 2, N \geq 2$	$\times N$	$\times N$	Yes
3	Striped Disk Array with Isolated Parity	$N+1 \geq 3$	$\times N$	$\times N$ (for reads) $\times 1$ (for writes)	Yes
5	Striped Disk Array with Distributed Parity	$N+1 \geq 3$	$\times N$	$\times N$ (for reads) $\times 1$ (for writes)	Yes
10	Striping of mirrored segments (RAID 0 over RAID 1)	$N \cdot M \geq 4$ $N \geq 2, M \geq 2$	$\times N$	$\times N$	Yes



A more detailed characteristic of various RAID levels can be found in the Internet.

4.2 Exchange Disaster Recovery and Troubleshooting Resources

For more in-depth information on Exchange Server disaster recovery, you can read and bookmark:

- Exchange resources on Microsoft TechNet:
<http://www.microsoft.com/exchange/techinfo/default.mspx>
- Exchange dedicated resources MExchange.ORG:
<http://www.msexchange.org/>
- Microsoft's Exchange Server 2003 Disaster Recovery Operations Guide
<http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/disrecopgde.mspx>
- MS TechNet article "What to Do When an Exchange Store Won't Mount"
<http://www.microsoft.com/technet/prodtechnol/exchange/2003/wontmount.mspx>