

SAP® MaxDB™

Backup & Recovery

Expert Session 11



MaxDB/liveCache Development Support

Roland Mallmann

Junco Ueda

August 3rd, 2010

THE BEST-RUN BUSINESSES RUN SAP™



Agenda



1. Introduction to Backup and Recovery
 - SAP MaxDB Backup Concept
 - Why Backup using SAP MaxDB?
2. Backup
 - Backup Types
 - Backup Templates
 - Snapshots
 - Check Backup
 - Backup Strategy
 - 7.8 Specific Features
3. Recovery
 - Recovery Strategy
 - Recovery Types
4. Monitoring and Troubleshooting
 - Common Errors and their Solutions
 - How to Monitor a Backup / Recovery
5. Questions and Answers

© 2010 SAP AG. All rights reserved. / Page 2

This session is aimed at a wide audience, as the topic is extremely important to any Database Administrator (DBA), regardless if it involves the smallest Open-Source database, or the largest SAP Enterprise sized installations. To protect the contents of the SAP MaxDB database, it is vital to make sure a sound backup strategy is in place, so any loss of data is avoided.

The main goal of this Expert Session on Backup and Recovery is to show what a SAP MaxDB Backup is, why it is important to regularly create, check and ultimately also use it for recoveries. The session will cover the basics of setting up backups, which different backup types exist, which templates (formerly called 'media') exist, which new features are to be expected in the new version 7.8 and how to decide upon a suitable Backup strategy.

Next, the different Recovery types are discussed. Choosing the right Recovery strategy is important, because you want to make sure you don't lose any valuable data, or even time.

Then, you'll learn how to monitor the Backup or Recovery, find out to check what might be a bottleneck, find out how to solve typical issues.

Lastly, there's a round of Q & A.

Agenda

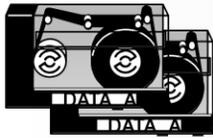


1. **Introduction to Backup and Recovery**
 - SAP MaxDB Backup Concept
 - Why Backup using SAP MaxDB?
2. **Backup**
 - Backup Types
 - Backup Templates
 - Snapshots
 - Check Backup
 - Backup Strategy
 - 7.8 Specific Features
3. **Recovery**
 - Recovery Strategy
 - Recovery Types
4. **Monitoring and Troubleshooting**
 - Common Errors and their Solutions
 - How to Monitor a Backup / Recovery
5. **Questions and Answers**

In this introduction, a first peek into the world of SAP MaxDB Backups and Recoveries is taken.

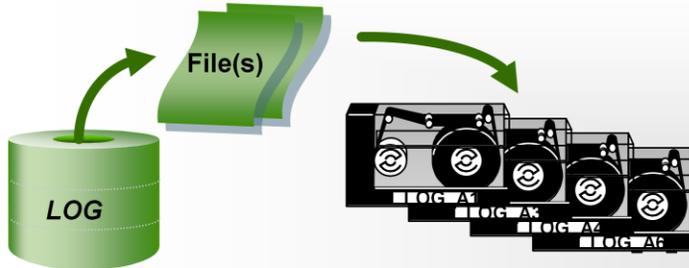


Data Backup



Periodic backup of the dataset to External backup media

Log Backup



Periodic backup of the log area using file backups to external backup media to avoid LOG FULL situations

Important: All backups are performed by the database kernel.

Create file backups of volumes in exceptional cases only.

To prevent data loss in a database used for production operation, it is essential that the data and log areas are backed up regularly.

With Database Studio, MaxDB provides a user-friendly tool for performing these backups. This tool allows you to display the backup history and use external backup tools such as Legato NetWorker, NetVault, and TSM. The widely-used *BACKINT* for Oracle interface has also been offered since early database versions of Max DB (SAP DB). MaxDB features an enhanced version of the standard *BACKINT* interface, namely *BACKINT for SAP MaxDB*. This enhancement also allows you to use pipes for backing up databases using this interface.

A backup of the database instance consists of two elements:

- Periodic backup of the data area
- Backup of the log entries

While backups are being performed, the database remains available without restrictions. Only performance can be limited to some extent. For this reason, we recommend that you perform backups when the SAP dialog workload is low.

Why Backup using SAP MaxDB? Backup vs. Filesystemcopy vs. Snapshot



Type	Description	Pros	Cons
SAP MaxDB Backup	<ul style="list-style-type: none"> • Only database pages marked for backup are backed up • Various pagechecks are performed during backup 	<ul style="list-style-type: none"> • can be processed in parallel (scalability, speed) • result can be checked, verified • can be combined with incremental/log backups for recovery • compression (7.8+) 	<ul style="list-style-type: none"> • can be sluggish, if not configured well • no compression in versions < 7.8
SAP MaxDB Snapshot	<ul style="list-style-type: none"> • Restartpage + Converter are copied 	<ul style="list-style-type: none"> • very fast 	<ul style="list-style-type: none"> • no checking of pages • storage within DB only (i.e. not in DB backups), i.e. no external storage
OS Backup	<ul style="list-style-type: none"> • blunt copy of all relevant files 	<ul style="list-style-type: none"> • (can be) fast • very limited costs 	<ul style="list-style-type: none"> • no MaxDB specific checking of backup • no BackHist integration • error-prone
OS / 3 rd party Filer Snapshot	<ul style="list-style-type: none"> • (example) using NetApp Filer Snapshot tech 	<ul style="list-style-type: none"> • very fast • integration into BackHist (7.8+) 	<ul style="list-style-type: none"> • no MaxDB specific checking of backup • no pagechecks

© 2010 SAP AG. All rights reserved. / Page 5

The benefits and drawbacks of SAP MaxDB backups and its alternatives are shown here, explaining why it makes sense to use SAP MaxDB backups.

Agenda



1. **Introduction to Backup and Recovery**
 - SAP MaxDB Backup Concept
 - Why Backup using SAP MaxDB?
2. **Backup**
 - Backup Types
 - Backup Templates
 - Snapshots
 - Check Backup
 - Backup Strategy
 - 7.8 Specific Features
3. **Recovery**
 - Recovery Strategy
 - Recovery Types
4. **Monitoring and Troubleshooting**
 - Common Errors and their Solutions
 - How to Monitor a Backup / Recovery
5. **Questions and Answers**

© 2010 SAP AG. All rights reserved. / Page 6

In this second chapter, the different Backup Types are discussed in detail. Next, the SAP MaxDB Backup Templates, previously called Backup Media, are shown.

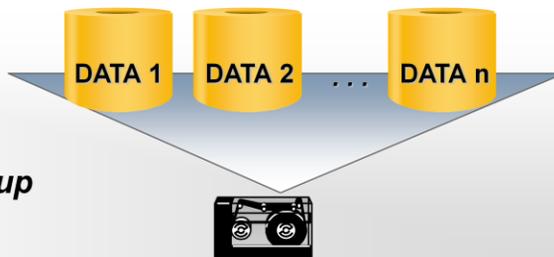
Because creating Backups is so important for any DBA, recommendations regarding Backup Strategies are also on the agenda.

Lastly, a short outlook on the new features to be expected with the next SAP MaxDB release 7.8 is covered.

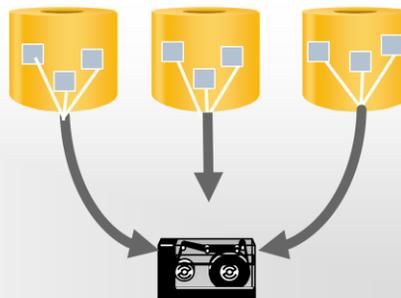
Backup Types I Complete – Incremental DATA



Complete
=
Complete data backup



Incremental
=
**Pages changed
since last complete
backup**



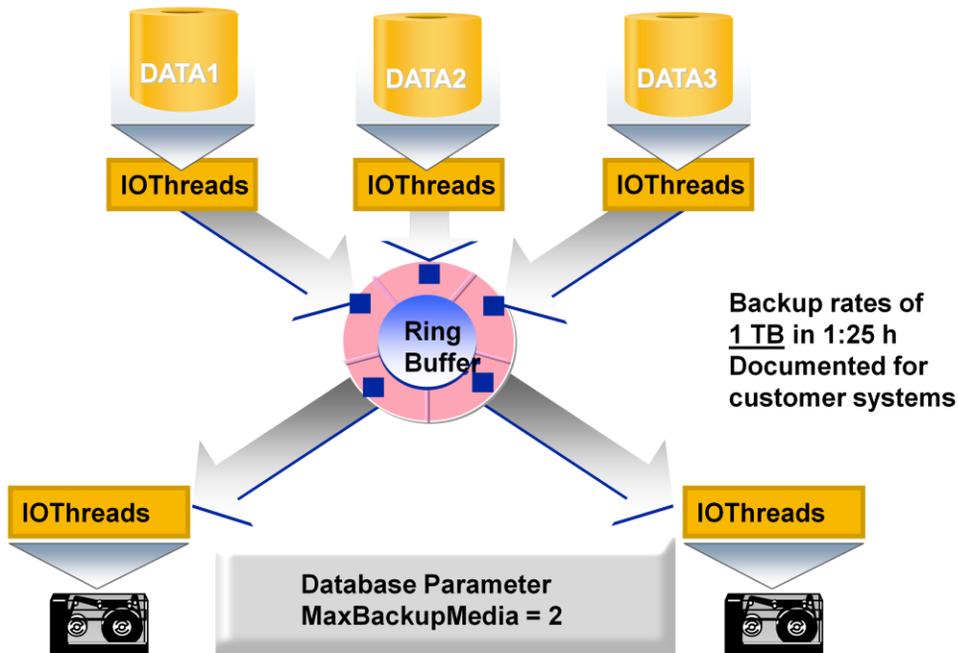
© 2010 SAP AG. All rights reserved. / Page 7

To create a complete backup of all pages used in the data area, choose *Complete Data Backup*. The configuration of the database instance is also backed up. This means that if you perform a recovery, it is also possible to restore the original configuration. Converter pages are not backed up.

To create an incremental backup of the data area, choose *Incremental Data Backup*. This backup contains all pages that have been changed since the last complete backup. Every following incremental data backup still backup only those pages changed since the last full data backup, so the size of the incremental backups will increase over time, depending on how much data was changed.

You access the backup functions in Database Studio by choosing *Backup...* from the context menu of the instance.

Backup Types II Parallel Data Backup



© 2010 SAP AG. All rights reserved. / Page 8

The data pages to be backed up can be saved to several templates (media) at the same time. This speeds up the backup. For example: In an optimal case, the runtime with three tape devices is a third of the runtime when one tape device is used.

You define the maximum number of parallel backup templates with the "MaxBackupMedia" database parameter (1 - 32). Any changes you make to this parameter only take effect after the next database restart.

A memory area where the read and write processes of the server tasks and the I/O thread processes created temporarily for the backup are performed is provided in the system as a (ring) buffer (with a block size of 8 x 8K, or 8 x 64K if page clustering is used).

The limits of the parallel backup process are determined by the access speed of the data volumes, the write performance of the backup devices, or the transport layer (such as a network) between the database server and the backup devices. As long as these limits are not reached, the process scales with every additional backup device in parallel operation.

Parallel backups can be performed only if the following prerequisites are met:

- Multiple tape devices, pipes, or disk areas of equal size are available.
- The value of the "MaxBackupMedia" parameter corresponds to the number of tape devices, backup files, or pipes to be used.
- A parallel backup template group has been created with one backup template for each tape device.

Different backup templates (in terms of size and speed):

- If one of the devices (tape drive, pipe, or hard disk) is slower than the rest, the system writes to this device at the maximum speed possible under the circumstances, regardless of the speed used to write to the other devices. This makes it possible, in emergencies, to combine fast hard disks with slow tape drives. Once the backup has been completed, a proportionally smaller volume of data will have been written to the slower device than to the faster devices.

Supported Backup Media



Data Volumes



Media supported:
tape/ file/pipe



Media groups:
up to 32
tapes/files/pipes
used in parallel

Log Volumes



File(s)



Media supported
for manual log
backup: file/pipe

File.3

File.2

File.1

Automatic log
backup: version
files

© 2010 SAP AG. All rights reserved. / Page 9

One or more media can be used for data backups.

- If multiple media are to be used, these must be organized as a group of parallel backup media ('template group').
- Tapes, files, and pipes can be used as backup media. Pipes are used as an interface to external backup tools, for example.

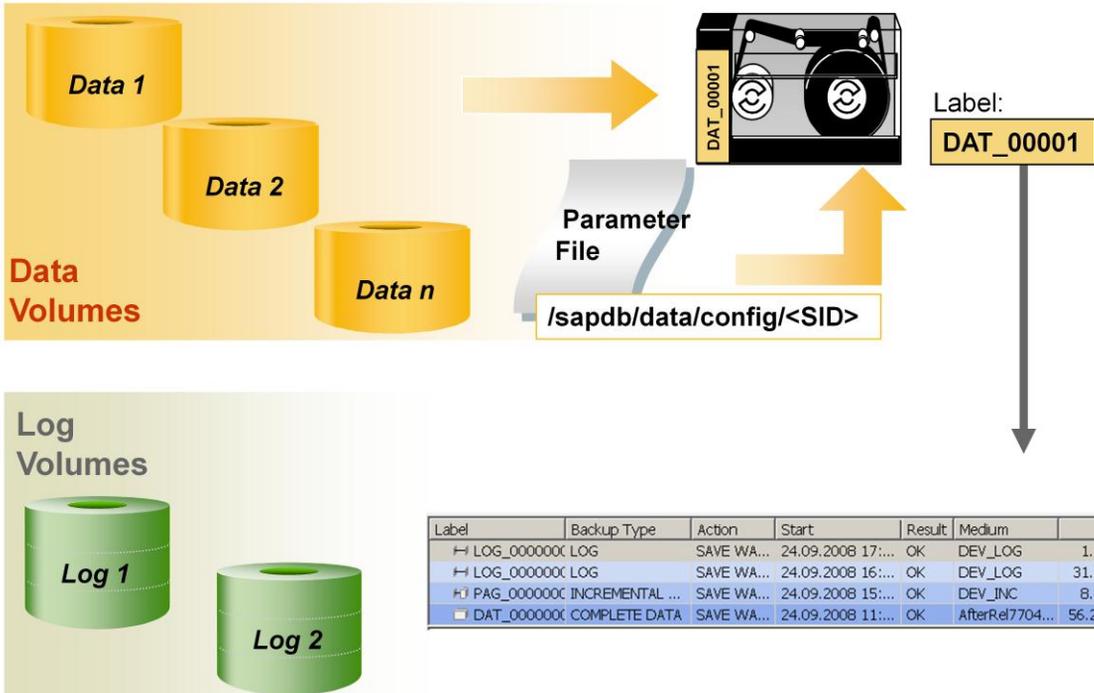
Regular files or pipes can be used for interactive log backups.

- Parallel log backups are not supported.
- Pipes are only supported if backup tools receive the data from the pipe and send a confirmation after the backup is completed.

The *automatic log backup* function can be used only when logs are backed up to files.

- Pipes cannot be used.
- The automatic log backup writes *version files*.
- You can use the "archive_stage" dbmcli command to forward these version files stored in the file system to a backup tool automatically.

Complete Data Backups



© 2010 SAP AG. All rights reserved. / Page 10

A complete data backup stores all used pages of the data volumes on the backup medium using the defined backup template. The current parameter file is also written to every backup medium at the beginning of a backup; in the SAP environment, this is the content of file `/sapdb/data/config/<SID>`. This means it is always possible to restore the instance environment required for this backup after the event.

Every backup is assigned a label in the sequence of backups. The administration tools use these labels to differentiate between the backups.

The backup number is increased by one, even if the backup was not successful.

For every backup, an entry is written to the backup history (dbm.knl file) in the Rrundirectory of the database.

Backups are performed by the database kernel. Do not use operating system tools (such as "dd" or "copy") to perform online backups of database instance disk areas on which the volumes are stored. Such backups are usually of no use since open files cannot be backed up or can only be backed up partially. You should also avoid performing offline backups using operating system tools for similar reasons.

Use the database functions to perform backups.

The advantage of this is that every data volume page that is backed up is checked when it is read (header/trailer comparison), which means that errors can be detected early. This check is comparable to a CHECK DATA (VERIFY) consistency check, although the analysis is not as detailed as that of a regular VERIFY.

Creating a Backup Template



The screenshot shows the SAP Database Studio interface. In the background, the 'twdf1915:DEV' instance is online with 84.69% data and 0.34% log. The 'Backup' tab is active, showing a table of backup history. In the foreground, the 'Backup Template for twdf1915:DEV' dialog box is open, showing configuration for a template named 'AfterRel770427'. The dialog includes fields for Name, Backup Type (COMPLETE DATA), Device Type (FILE), Backup Tool (NONE), Device/File (G:\sapdb\DEV\backup\AfterRel770427.3), Size (0 KB), OS Command, Overwrite, Block Size (64), and Autoloader. Below the dialog, the 'Templates' table is visible, listing several backup templates.

Name	Device/File	Device Type	Backup Type	Backup Tool	Size	Overwrite
AfterRel760509	G:\sapdb\DEV\...	FILE	COMPLETE DATA			
AfterRel770427	G:\sapdb\DEV\...	FILE	COMPLETE DATA			
B4Rel770425	G:\sapdb\DEV\...	FILE	COMPLETE DATA			
B4Rel770427-neu	G:\sapdb\DEV\...	FILE	COMPLETE DATA			

Before you can perform backups, you must define the relevant backup templates. You can create and change backup templates or template groups of parallel backup media in Database Studio by choosing *Backup Templates*.

To be able to create parallel backup media, you must set the value of the "MaxBackupMedia" parameter to match the number of individual media in a parallel medium. For example, if a media group is to comprise 10 individual media, the value of the "MaxBackupMedia" parameter must be "10".

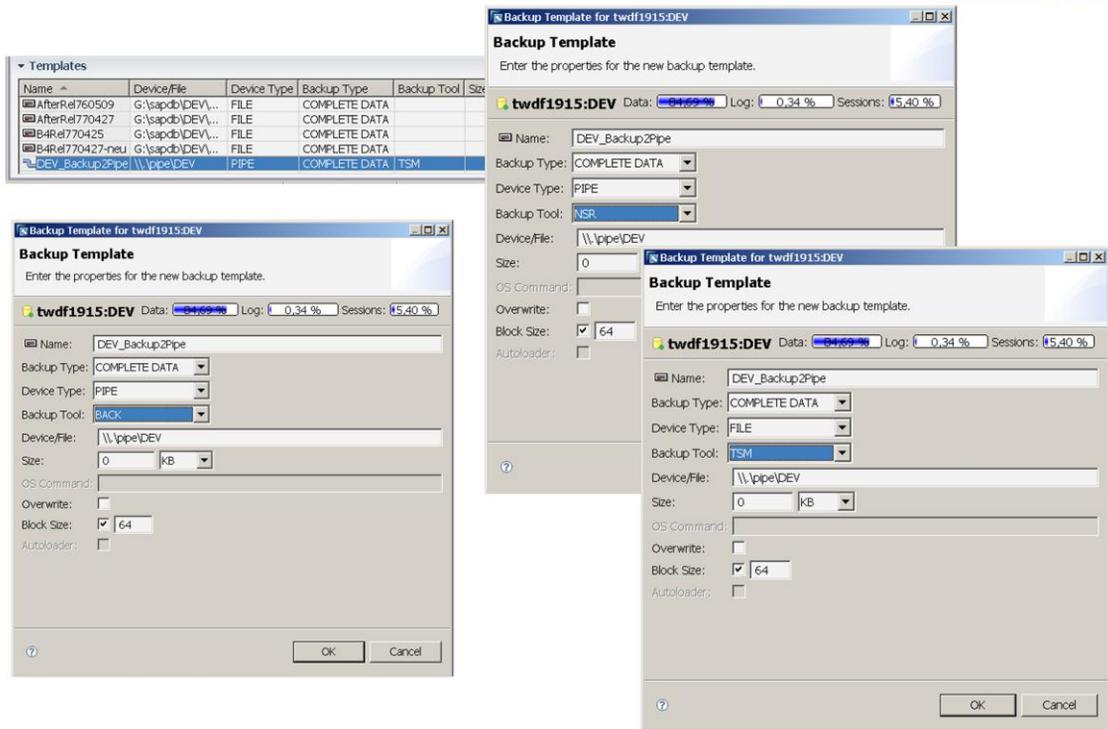
You can specify the following information for the template:

- **Name** of the backup template. This name is freely definable and is not dependent on the storage location used (*Device/File*).
- **Backup Type**: Specify the type of backup for which this template is to be used.
- **Device Type**: Tape, file, or pipe
- **Backup Tool**: Type of external backup tool (if applicable; for more information, see the next page)
- **Device/File**: Path to a device, name of a defined pipe, or name of a file including its path. If you do not specify a path, a file is created in the Rundirectory of the database instance.
- **Size**: Maximum size of the backups that can be created on this medium (if you do not make an entry in this field, files of unlimited size can be created).
- **OS Command**: In this field, you can specify operating system commands for backups

to tape.

- *Overwrite*: This option enables you to perform successive backups to the same file, overwriting the previous backup each time. Use this function carefully since it deletes the backup history and makes it impossible to restore one of the previous backups.
- *Block Size*: The entry in this field defines the size of the data blocks to be written to the medium. If page clustering is used for the instance, the value in this field must be larger than a multiple of the cluster size used (minimum block size, for example, of "64").
- *Autoloader*: Select the *Autoloader* checkbox if you want to use a tape device with automatic tape swapping.

External Backup Tools



MaxDB supports multiple external backup tools and technologies:

- NetWorker (NSR)
- Tivoli Storage Manager (TSM)
- Tools that support the *BACKINT* for SAP MaxDB or *BACKINT* for Oracle interface (BACK), such as:
 - HP Data Protector >6.0 supports *BACKINT* for SAP MaxDB.
 - Comvault QiNetix > 6.1 supports *BACKINT* for SAP MaxDB
 - All other external backup tools on the market must be connected using the *BACKINT* for Oracle interface, and based on our experience, these require additional adapters from their providers.

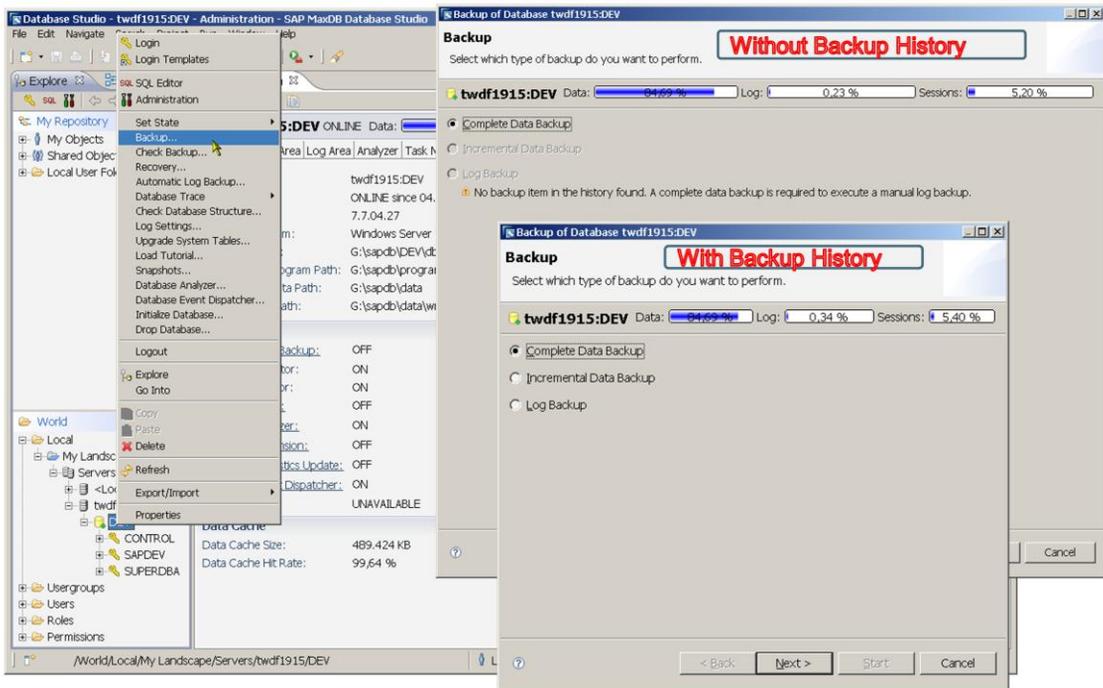
To support one of these tools, the *Device Type* of the backup template must be set to "Pipe".

Additional examples of definitions for templates in Unix and Windows:

- Windows: First tape drive: \\.\tape0
Pipe: \\.\pipe\PipeName
- UNIX: Tape drives, for example: /dev/tape0
Pipes: /tmp/pipe0

Template definitions are stored in the dbm.mmm file in the run directory of the database instance.

Backup Wizard – First Steps



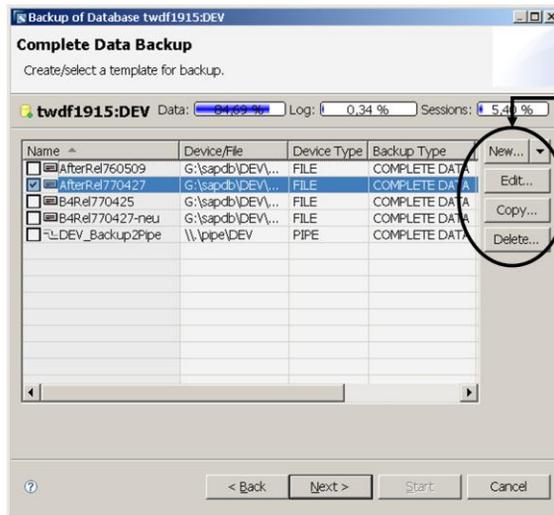
Database Studio includes a Backup Wizard that helps you perform a backup.

To start a complete data backup, choose *Backup...* from the context menu of the instance name.

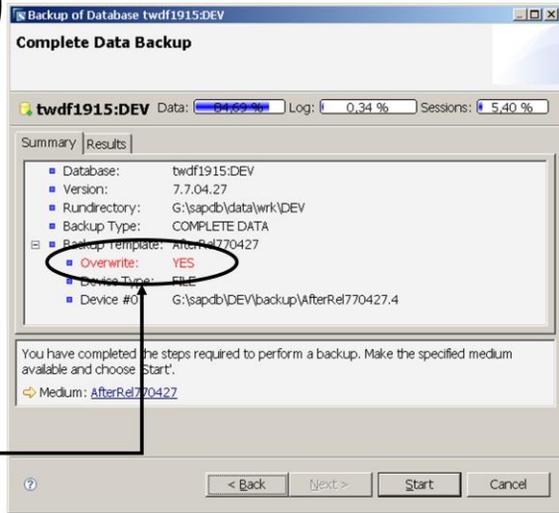
In the first step, you have to specify the type of backup: complete data backup, incremental data backup, or log backup.

If a complete backup has not yet been performed, only this option available when you start the Wizard and a corresponding warning message is displayed.

Backup Wizard – Preparation and Start



Creating and changing a backup template



Be careful when using **OVERWRITE**

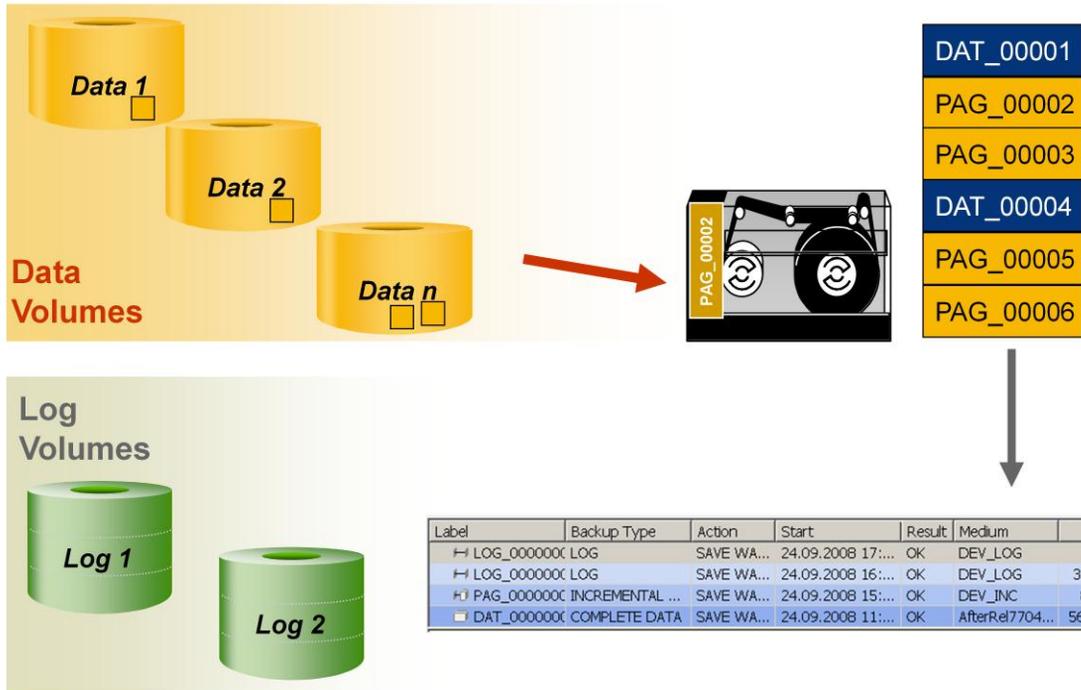
In the next step, you select a medium or template or create a new one.

A summary of the conditions is displayed before the backup is performed.

Using the *Back* or *Next* buttons, you can move back or forward a step in the backup process and change entries.

Be careful when using the "Overwrite" function from the template definition. This function can easily overwrite backup files that are still needed. Instead, create scripts that process the files further (for example, transport them to the downstream backup solution). These scripts delete the original backup files after successful postprocessing so that the database can always work with a valid file name. Check whether it is possible to connect MaxDB directly to the downstream backup solution. For more information, see the "Concepts" slides later in the unit.

Incremental Data Backups



In addition to complete data backups, you can also perform incremental data backups to save individual data pages that have changed. In this case, only the data pages that have been changed since the last complete data backup are saved.

The version number of the label is increased by one for **every** data backup, even if the backup was not successful.

This information is written to the backup history first.

Snapshots freeze a database state so that it can be restored later, for example. All subsequent changes are made in new (copied) pages.

MaxDB offers multiple snapshots with separate functions within the database.

Snapshots can be created or deleted in the ONLINE operational state.

It is only possible to restore a state from a snapshot in the ADMIN operational state.

Typical Use

- Very fast point-in-time recovery (for example, for SAP upgrades or SP imports)
- Restoring training systems to a defined status



© 2010 SAP AG. All rights reserved. / Page 21

As of version 7.7 you can freeze the data area of a MaxDB installation multiple times using internal snapshots.

A snapshot can be created in the *ONLINE* operational state. You can subsequently revert to the dataset of the snapshot and/or delete the snapshot.

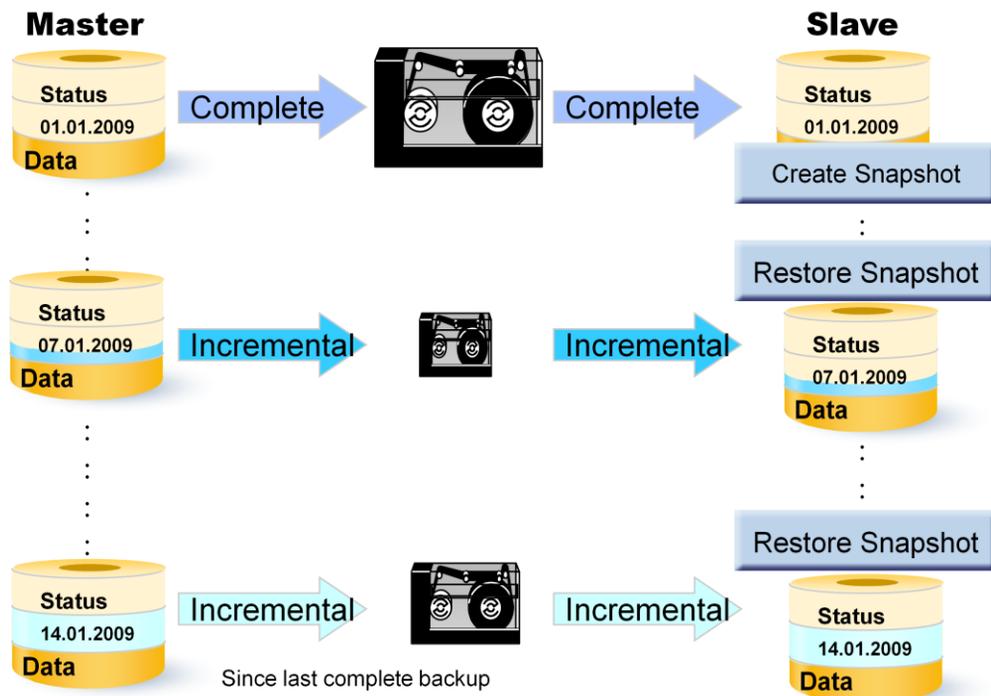
The database kernel uses the `CREATE_SNAPSHOT` command to copy the restart page from the second block of the first data volume to a different position. The complete converter is also copied. The original restart record contains a reference to the restart record relevant for the snapshot.

The `RESTORE_SNAPSHOT` command is used to delete the current converter. All blocks that are no longer needed are marked as free. The transaction information that is no longer needed is deleted from the log so that the `HISTLOST` state is activated. When it is next restarted, the instance works with the data that was frozen using the `CREATE_SNAPSHOT` command and has now been reactivated.

The `DROP_SNAPSHOT` command deletes the restart record and the converter linked to it that is relevant for the snapshot. FreeBlockManagement (FBM) marks all blocks that are no longer needed as free.

MaxDB supports multiple snapshots. Operating the instance with multiple snapshots uses more capacity in the data area.

MaxDB Snapshots: Master / Slave Setup



© 2010 SAP AG. All rights reserved. / Page 22

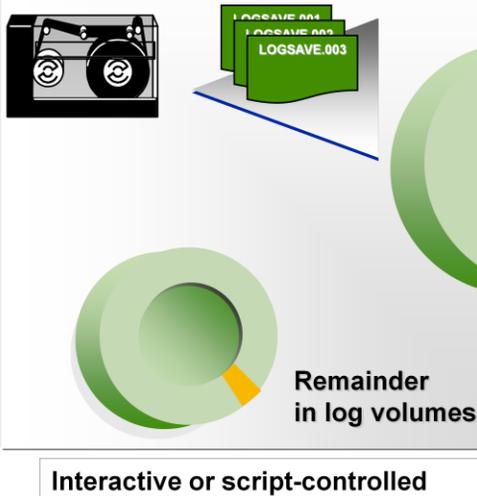
MaxDB gives you the option of using snapshots to synchronize a master and one or more slave instances.

Create the slave instance as a homogeneous system copy using the backup/restore functions. Create a snapshot before the slave instance is first restarted.

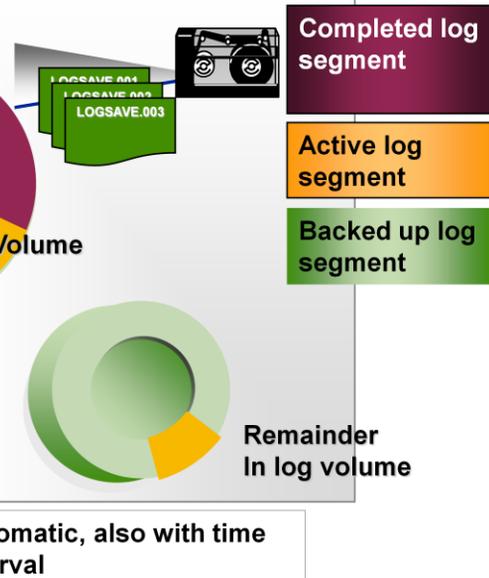
To transfer changes in the master instance to the slave instance, reset the slave instance to the snapshot and then import an incremental backup from the master instance. You can reset the slave instance to the snapshot as often as required and import incremental backups from the master instance.

This procedure works until a complete backup is created in the master instance. New incremental backups then no longer match the snapshot in the slave instance. You can import a complete data backup into the slave instance to synchronize it with the master instance.

Option 1



Option 2



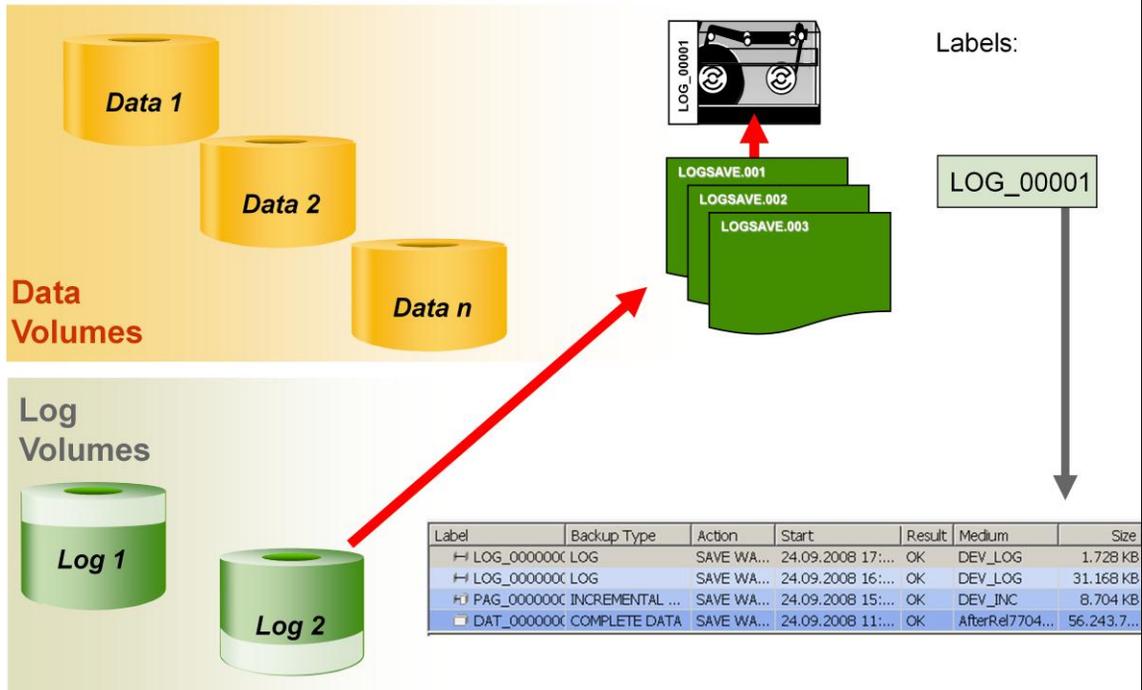
From the context menu of the instance, choose *Backup...* → *Log Backup* to start a log backup that saves all content of the log area that has not yet been saved to a backup medium of your choice. The content of the log area is then released for **overwriting**. Note that the log entries are only released for overwriting and not actively deleted.

- The log backup is performed in sections, the size of which is defined (in pages) by the "AutoLogBackupSize" parameter. By default, the value of the "AutoLogBackupSize" parameter is calculated at the time of installation based on the existing log volumes and set to one third of the total log area.

When you activate automatic log backup, completed log areas are automatically backed up to backup media selected for this purpose.

- We recommend that you use a separate hard disk area for automatic log backups. Only files (*Device Type: File*) can be used as a backup medium at present. To learn how data can be supplied to these files automatically, see the DBMCLI command description for `archive_stage` later in this unit.
- You do not need to deactivate automatic log backup during a data backup or *Verify*. The database kernel monitors the completed log segments.
- If you want to perform an interactive log backup even though the automatic log backup function is activated, you first have to deactivate automatic log backup and then reactivate it after you have performed the interactive log backup.
- You can also specify a specific time interval in which the log is saved automatically

Interactive Log Backups I



© 2010 SAP AG. All rights reserved. / Page 24

Interactive log backups back up all occupied log pages from the log volume that have not yet been backed up and do not belong to an open transaction (without a COMMIT).

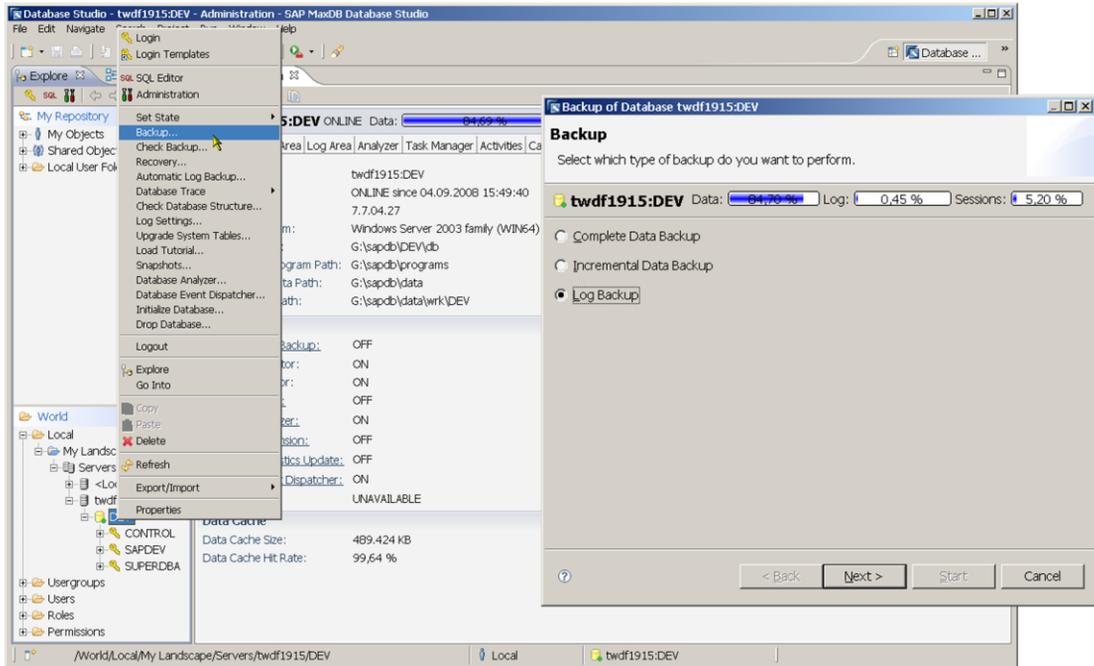
Only *version files* and external backup tools with a confirmation function are accepted as backup media for interactive log backups. For this reason, the log backups should then be stored finally on other backup media.

The system automatically adds a version number (three characters with leading zeros) to the file name defined in the backup medium. Once the number set is exhausted, additional digits are added.

The labels of the log backups are assigned independently of the numbering of the complete and incremental data backups.

All log backups are listed in the backup history in reverse chronological and logical order together with the data backups.

Interactive Log Backups II

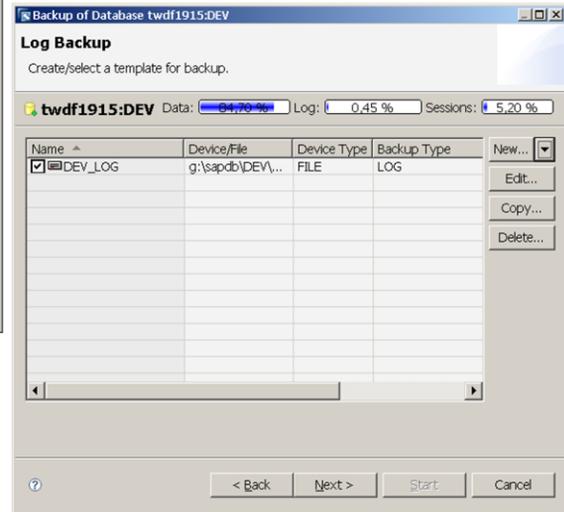
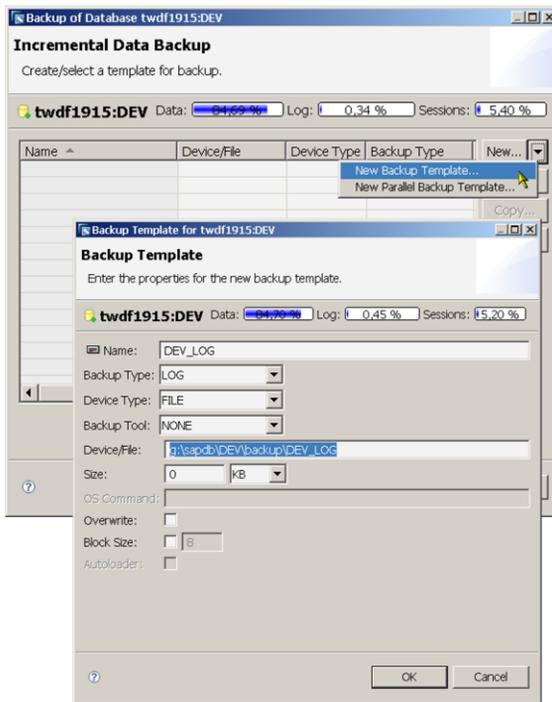


You can start an interactive log backup by choosing *Backup ...*

As always, it is only possible to perform a log backup if a complete data backup has already been carried out. In this context, log backups are delta backups that require a starting point. If this prerequisite is not met, the *Log Backup* option is grayed out.

Then choose *Next*.

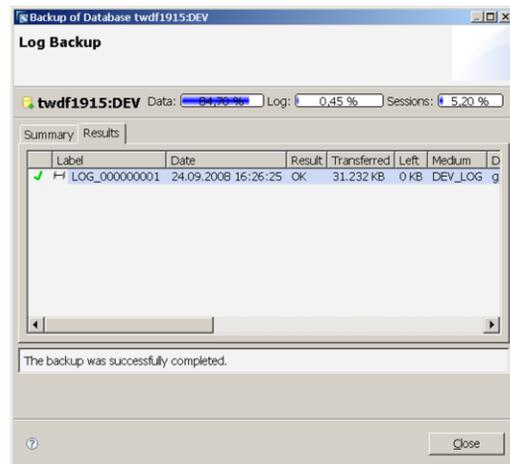
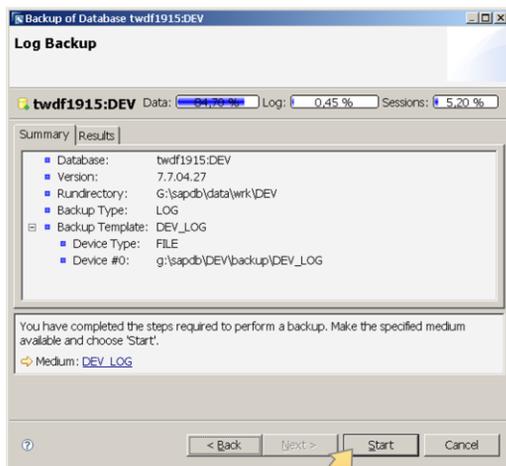
Interactive Log Backups III



As with data backups, you first have to create an appropriate backup medium or select an existing one.

If the automatic log backup function is active, you must deactivate it before starting the interactive log backup and then reactivate it afterwards.

Interactive Log Backup IV



© 2010 SAP AG. All rights reserved. / Page 27

A summary of the settings is displayed before you start the log backup.

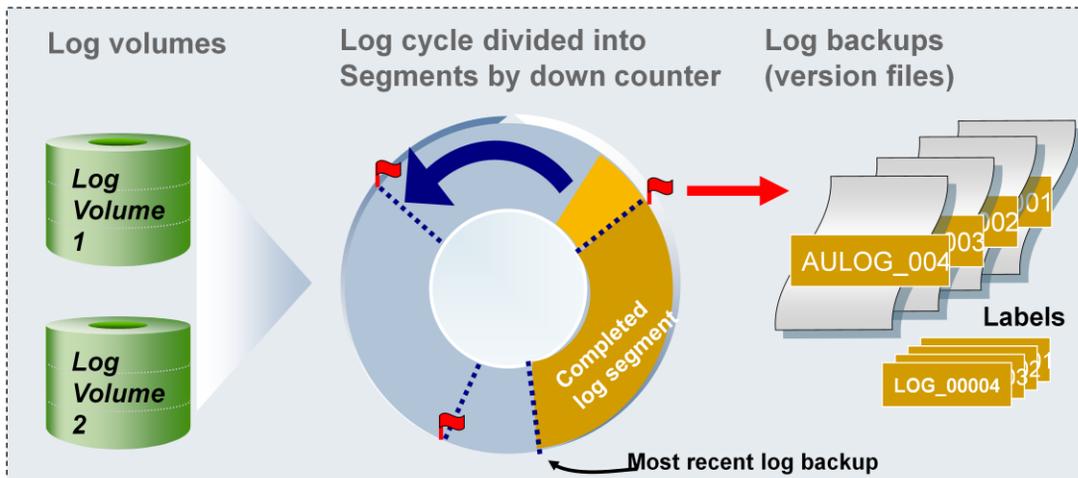
Once the log backup is complete, the system displays the label of the backup.

If the log backup takes a long time, you can move the Wizard to the background and display it again using the *Actions* tab page.

If you are starting the interactive log backup due to a LOG FULL situation, the programs working with the database are stopped until the database can continue processing their transactions. These programs are reactivated automatically when the log backup is complete unless they have been stopped manually or timeout values of the application have been exceeded.

If you do not perform an interactive log backup when a LOG FULL situation occurs and immediately reactivate the automatic log backup function instead, the stopped transactions are reactivated immediately after the first log segment has been backed up (by default, a third of the total log area), while the other segments are still being backed up.

Automatic Log Backups



If automatic log backup is activated, the database instance backs up log information automatically as soon as a predefined amount of data has been added to the log.

This amount is defined by the "AutoLogBackupSize" MaxDB parameter.

You can also activate a time-based control function.

We recommend that you always use this automatic function to avoid LOG FULL situations.

© 2010 SAP AG. All rights reserved. / Page 28

Activate automatic log backup to protect the database against LOG FULL situations.

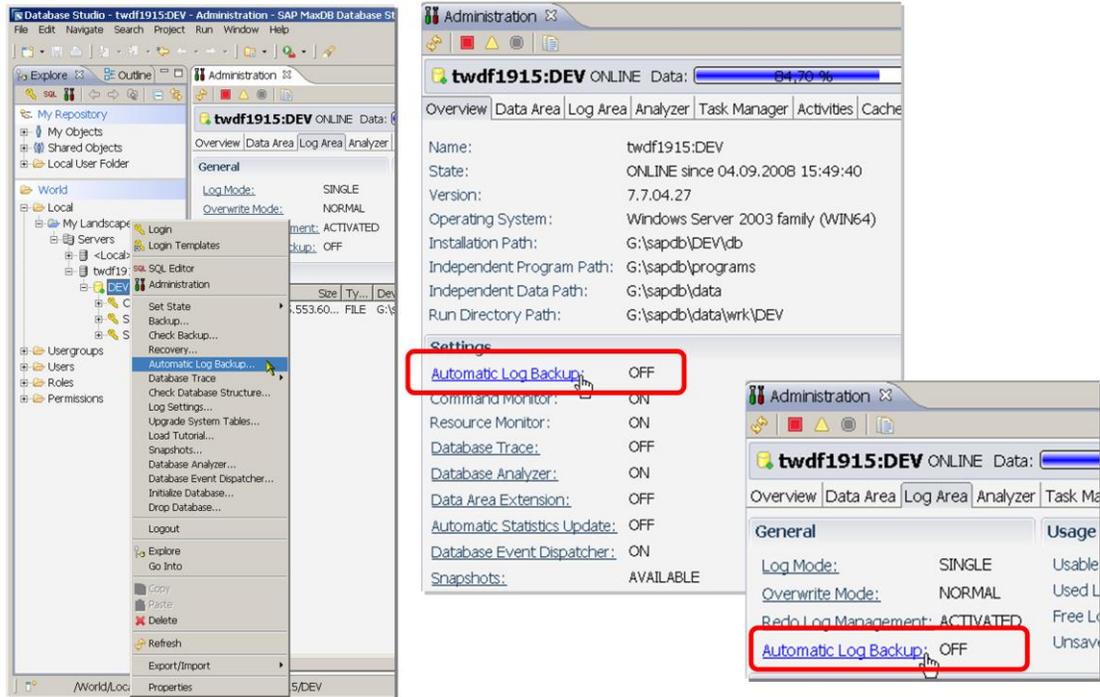
When the number of entries written to the log equals the value of the "AutoLogBackupSize" parameter, this data is written to the defined backup template and the log area is released for overwriting. The log area is **not** deleted by being overwritten with dummy data.

- Each log area is stored in a separate backup file to which a version number is assigned.
- It is only possible to perform automatic log backups to backup files. In turn, these must be backed up to tapes or backup servers with the relevant tools.
- The automatic log backup can also be used alongside script-controlled log backups as a safety net in case the script fails.
- Automatic log backup must be deactivated whenever a script-controlled log backup is performed (for example, with the **dbmcli -U c autolog_off** command).

The longer the duration of data backup actions performed with the utility task, the more important automatic log backups become, since the utility task cannot be used for other purposes during this time, which means that you cannot start any interactive log backups. For example, if data backups take a long time in online mode because backup media are slow, it is possible that online activities carried out during this backup could cause the log area to become full. If automatic log backup is activated, the log area can

still be backed up when the utility task is occupied, thereby preventing a LOG_FULL situation.

Activating the Automatic Log Backup Function



© 2010 SAP AG. All rights reserved. | Page 29

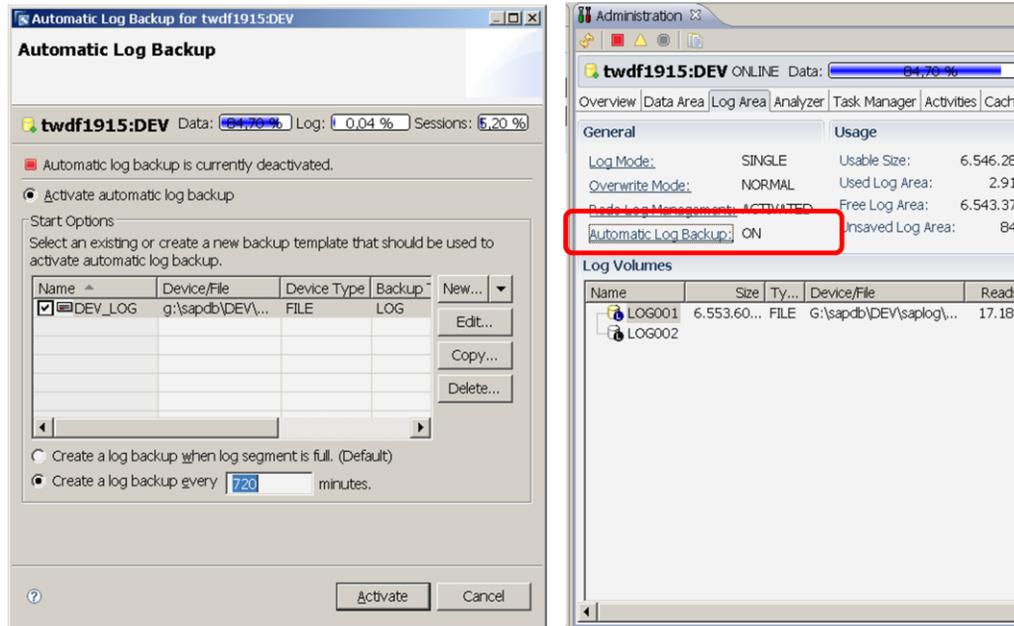
To activate the automatic log backup function, create a corresponding backup template with *Backup Type* "LOG". Next, activate the automatic log backup function by setting *Automatic Log Backup* to "ON".

As always, an initial data backup must exist before you can activate the automatic log backup function. Otherwise, the Wizard does not offer this option.

Automatic log backup remains active until it is deactivated manually or an error occurs during the backup process.

You check the status of the automatic log backup function, for example, on the *Log Area* tab page of the registered database instance.

Activating the Automatic Log Backup Function



© 2010 SAP AG. All rights reserved. / Page 30

Database Studio also gives you the option of activating automatic log backup based on time criteria.

To define the backup interval, you can also use the `autolog_on` or `medium_put` DBCMLI commands.

■ Command overview

`autolog_on ... [INTERVAL <interval>]`

`medium_put ... <interval> (MaxDB Version 7.7.02 or later)`

■ Command syntax (MaxDB Version 7.7.02 or later)

`medium_put <name> <location> <type> <backup_type> <size> <block_size>
<overwrite> <autoloader> <os_command> <tooltype> <interval>`

The default value of the `<interval>` parameter is "0", which means that no interval is to be used.

You can override the value of the `<interval>` parameter for the backup template (backup medium) if the automatic log backup function is activated:

`autolog_on [<medium>] [INTERVAL <interval>]`

*** TODO: check for 'backup_template' dbmcli commands ***

Check the backup to determine whether:

- Backup is complete
- Backup content
(*Page Header*) was read correctly
- The data content is not checked

Using a service database:

- Data not written to the hard disks
- Service database does not occupy any disk space
- Possible to check parallel backups

Call using Database Studio, for example:

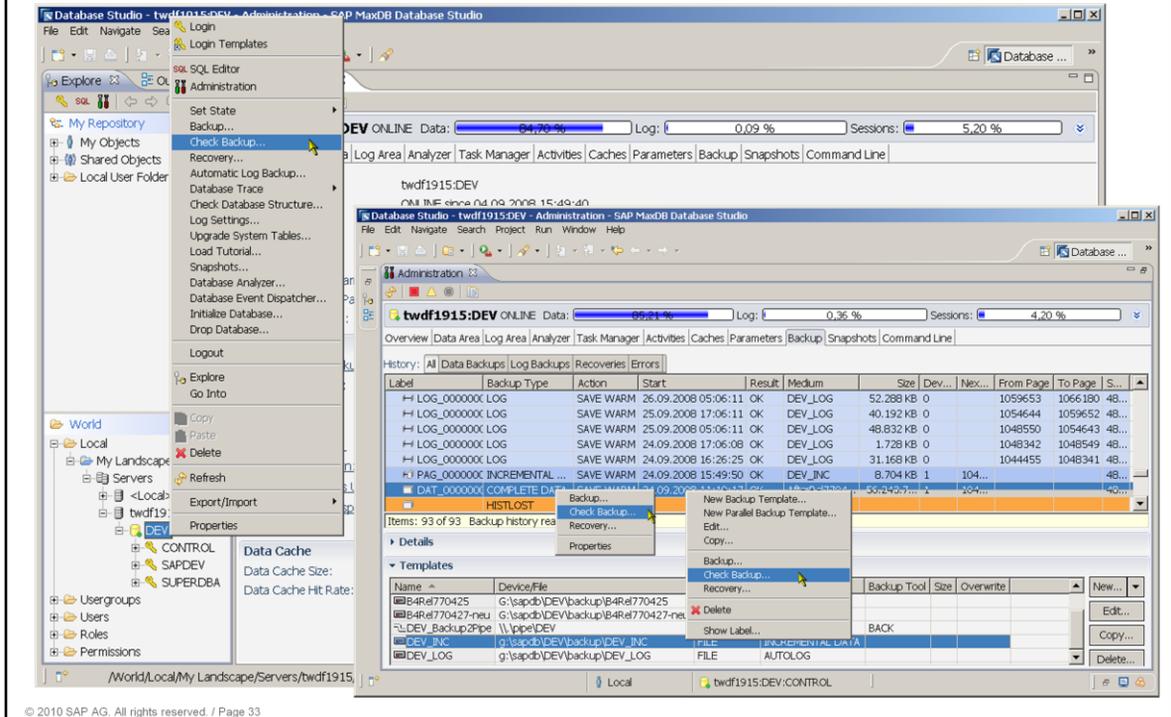
- From the context menu of the database instance → *Check Backup...*
- From the context menu of the backup in the backup history → *Check Backup ...*
- From the context menu of the backup template → *Check Backup ...*

Before you overwrite the backups of a backup generation, check whether an intact backup exists.

You can also use the following to check a data backup from the commandline:

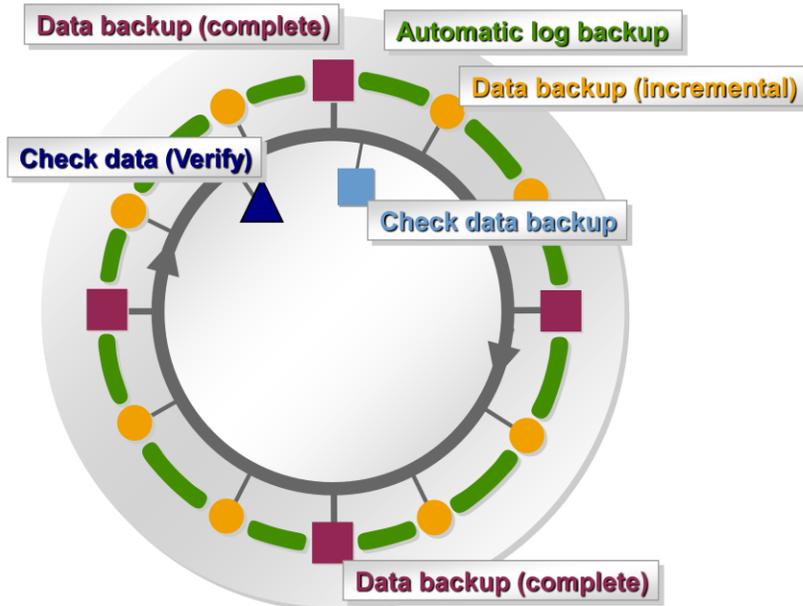
```
dbmcli -U c -uSRV recover_check <template name> DATA
```

Checking Backups



In Database Studio, you can check backups by choosing *Check Backup...* from the context menu of the instance and using the *backup check Wizard* that then appears. You can check both log and data backups. Select the relevant backup medium and choose *Next*.

Exemplary Backup Strategy



© 2010 SAP AG. All rights reserved. / Page 36

To ensure data security, it is necessary to perform data and log backups at appropriate intervals. We recommend that you perform:

- A complete data backup (of all data) at least once a week or, if possible, daily
- An incremental data backup (of individual pages) after more extensive system activities during the week
- Log backups at least daily, ideally using the automatic log backup function

This creates four backup generations in a 28-day cycle. The backup media can then be reused.

- Before reusing backup media, you must perform a consistency check (using *Check Data (VERIFY)*) of the database at least once within the backup cycle. This ensures that the database is physically consistent and it is safe to overwrite these tapes.
- At the beginning of each backup cycle, you must also check whether the complete data backups can be read (see "Checking Backups") to ensure that the backup concept is working as expected.
- If you want to use snapshots of the file system as a substitute for a backup of the database, you must check more often whether the system is consistent using the *Check Data* function. Do this at least once a week (based on the complete backup in the example above, once a week).

Weekly Planner for Backups – DBA Cockpit



The screenshot displays the SAP DBA Cockpit interface. The main window is the 'Jobs: DBA Planning Calendar' for November 2007, showing a weekly grid with scheduled actions like 'PrepUpdStat...' and 'DataBackup'. Overlaid windows include 'Schedule a New Action' for 'Mark tables requiring statistics update', 'Complete data backup', and a table for 'Restrict Value Range' with columns for backup media, location, save type, and size.

Name of Backup Media	Location of Backup Media	Save Ty	Type of Backup Media	Size of Media	Block Si	Overwrite Opt
COPY_GTP	\\PIPE\copy_gtp	DATA	PIPE	0	8	NO
NSRGT	\\PIPE\NSRGT	DATA	PIPE	0	8	NO

© 2010 SAP AG. All rights reserved. / Page 37

With the DBA Cockpit and its embedded DBA Planning Calendar, the Computing Center Management System (CCMS) provides you with an easy way to schedule the most important administrative tasks.

- The scheduled actions run automatically and are controlled by SAP NetWeaver. The system administrator is responsible for ensuring that the correct backup media are available.
- You can check whether the scheduled actions were successful in the DBA Cockpit directly.
- You can schedule actions to run at one-week or multi-week intervals, depending on the backup strategy you require.

The Planning Calendar accesses the same information as Database Studio.

- It displays the backup media that you defined using Database Studio.
- You can view the backup history in the DBA action logs (transaction **DB12**).

In addition to performing backup actions, you can also use the Planning Calendar to refresh the optimizer statistics. This is described in the *Performance Tuning and Troubleshooting* unit.

Important: Planning entries take on the current status of the database environment. If actions are no longer performed on the original database server after switch processes in failsafe system landscapes, this can lead to errors. For many database-related scheduling entries, the current database server is determined and this information

included in the planning data. After a database server failure, the CCMS can no longer access this. If this happens, reschedule the actions and revert to the normal state of the system landscape. It is not possible to ensure with complete certainty that the same backup media are defined on the second server.

Integration of external backups via snapshot clones or split mirror into the MaxDB backup history

- External backups are supported for I/O consistent snapshots and mirrors.
- MaxDB provides an integration into the backup history for those backups.

Integrated support of compressed backups

- Automatic compression and decompression of backups
- Useful for full/incremental data backups and log backups

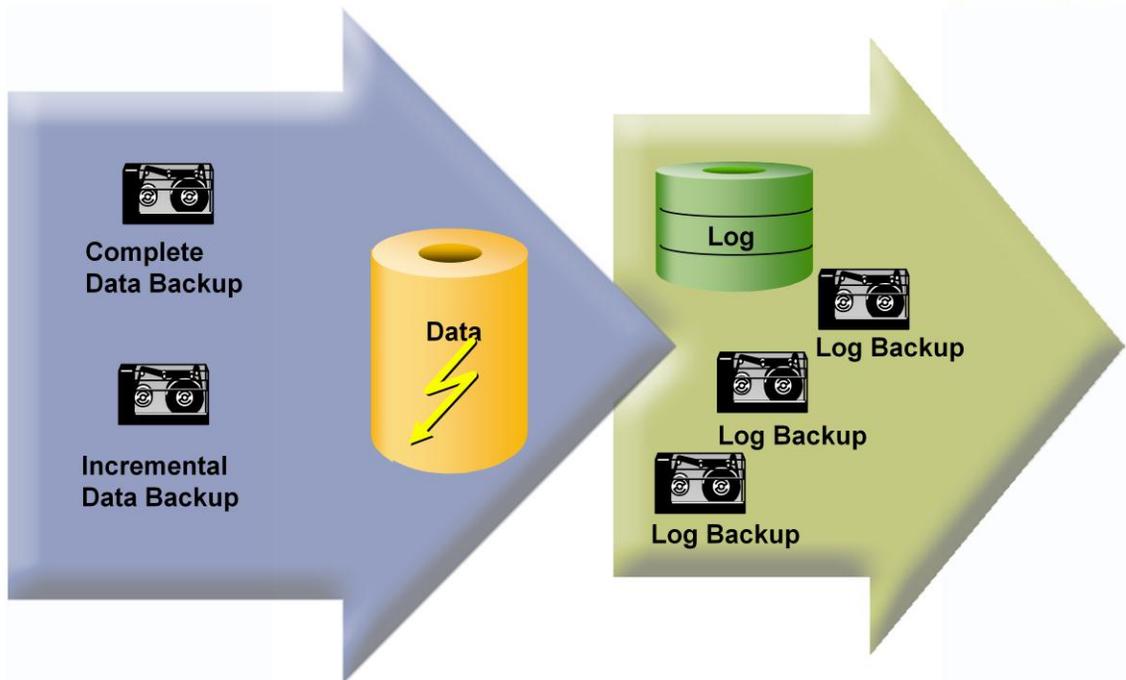
New *Backup Template* definition supporting all of the above backup/restore scenarios

Agenda



1. **Introduction to Backup and Recovery**
 - SAP MaxDB Backup Concept
 - Why Backup using SAP MaxDB?
2. **Backup**
 - Backup Types
 - Backup Templates
 - Snapshots
 - Check Backup
 - Backup Strategy
 - 7.8 Specific Features
3. **Recovery**
 - Recovery Strategy
 - Recovery Types
4. **Monitoring and Troubleshooting**
 - Common Errors and their Solutions
 - How to Monitor a Backup / Recovery
5. **Questions and Answers**

In this 3rd chapter, we take a look into the Recovery topic. How to decide upon a Recovery strategy is discussed here, as well as an example of an actual Recovery using Database Studio.



© 2010 SAP AG. All rights reserved. / Page 40

If you follow our recommendations for the disk configuration of your database instances and the backup strategy, the current log entries and at least four backup generations are always available for you to restore the content of the database instance if problems occur. It is then very unlikely that you will lose any data.

If a data volume sustains physical damage, a complete database recovery needs to be performed. This basis for this type of recovery is normally the complete and incremental data backups as well as log backups of the latest backup generation.

If a logical error occurs in the SAP system, making it necessary to reset the system to a previous state, you also do this by performing a database recovery using a complete data backup and then importing incremental data and log backups. The administrator can specify whether all available log information is to be recovered up to the most recent point in time possible, or only up to a specific time in the past without the most recent transactions.

To ensure you are well prepared for a recovery, we recommend that the DBAs regularly test a complete database recovery using the backups from the production system. For these tests, you require a test server comparable to the database server. This could, for example, be your quality assurance system.

Diagnosis in Database Studio



Runtime Environment

knlMsg
knlMsg.old */sapdb/data/wrk/<SID>*
knlMsg.err

Backup History

dbm.knl
dbm.mdf */sapdb/data/wrk/<SID>*

Actions of the Utility Task

dbm.utl */sapdb/data/wrk/<SID>*

Log of the Client Commands

dbm.prt */sapdb/data/wrk/<SID>*

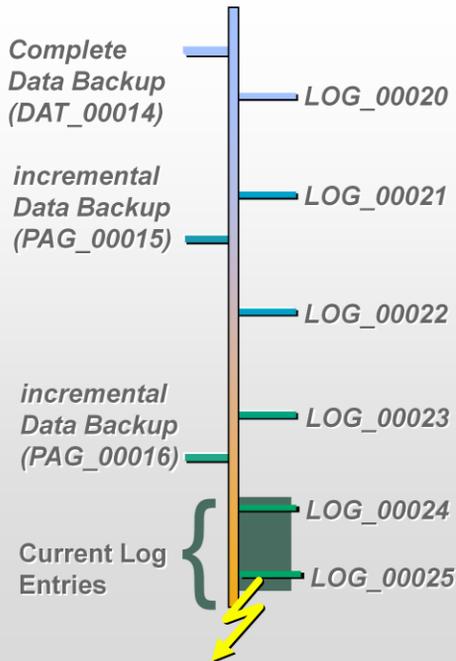
The following files contain information for identifying the cause of database problems:

- The knlMsg diagnosis file is the most important source of information if database problems occur. This file logs the messages issued during communication between the runtime environment of MaxDB and the database kernel. Every time the database kernel is started, a new file is created and the existing file saved as knlMsg.old. The knlMsg file is also backed up immediately after a database crash and is kept until two further database crashes have occurred. These files are stored in the directory defined by the "DiagnoseHistoryPath" parameter.
- The error messages in knlMsg are also logged in knlMsg.err. This file is not overwritten and can be archived and then deleted if it becomes very large. A new version of the file is then created automatically.
- The dbm.knl and dbm.mdf files contain information about the backup history and the use of backup media and labels. Based on this information, the database instance can provide assistance during recovery processes. If these files are lost, you have to determine the actions required for the recovery yourself.
- In contrast, the "dbm.prt" and "dbm.utl" files contain information about the commands sent to the database server (dbm.prt) and the resulting actions performed using the utility task (dbm.utl).

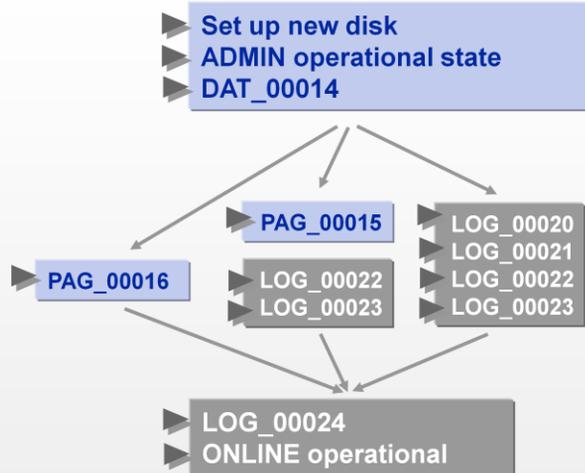
Recovery Strategy



Backup History



Recovery Processes



To restore the database instance after a structure or disk error, you first have to make new hard disk space available.

You can use the backup history to obtain an overview of which actions are required. Perform the recovery in the ADMIN operational state.

The first step of the recovery process is to import the most recent complete data backup.

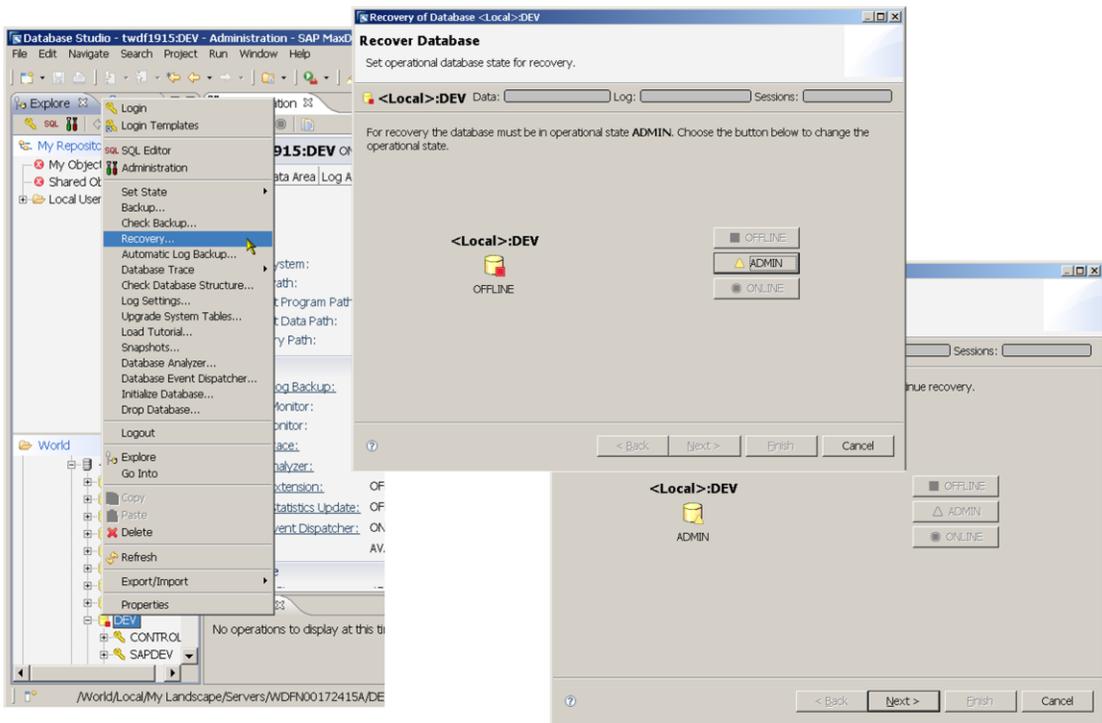
If the information in the log area was created before the start of this data backup, the database instance can restore the most recent database state using the information from the log area immediately after the data has been imported successfully.

Otherwise, the most recent incremental data backup is imported followed by the missing log backups. You can specify the point in time up to which the log entries are to be recovered.

If the most recent incremental data backup is not available, it is also possible to use any recent incremental data backup instead as a workaround. This correspondingly increases the effort required to recover the log.

When the database instance is started in the ONLINE operational state, the most recent log information from the online log volume is used and the recovery is complete.

Performing Recoveries I

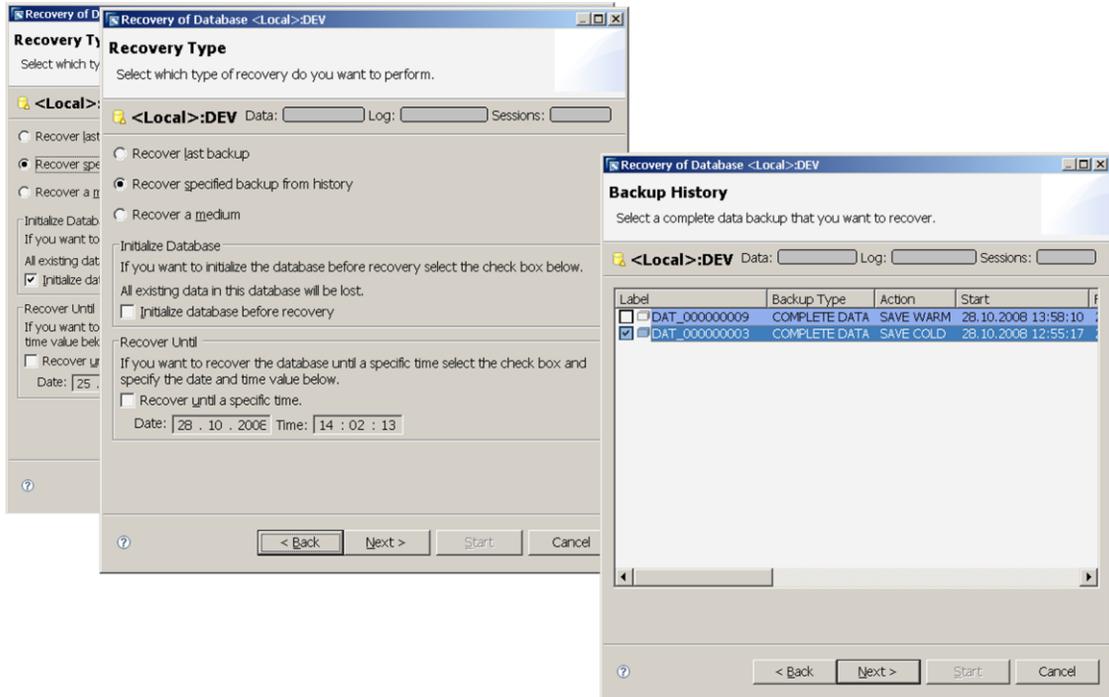


© 2010 SAP AG. All rights reserved. / Page 43

The following slides explain the recovery process in Database Studio.

Choose → *Recovery...* from the context menu of the instance to display the recovery Wizard.

Now switch the database instance to the ADMIN operational state.



© 2010 SAP AG. All rights reserved. / Page 44

Choose the recovery option required.

- **Recover last backup:** Resets the database to the state it was in before the crash.
- **Recover specified backup from history:** Resets the database to a certain previous state stored in the backup history.
- **Recover a medium:** Restore without a backup history, for example, to import a backup of a different database.

Using the *Initialize database before recovery* option, you can reinitialize the database and the log in particular.

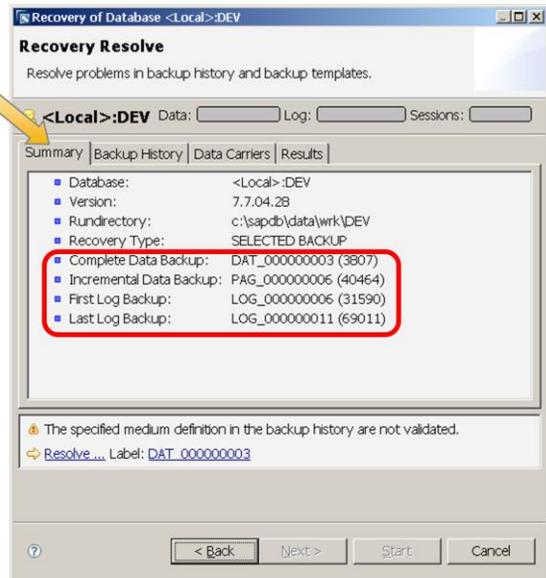
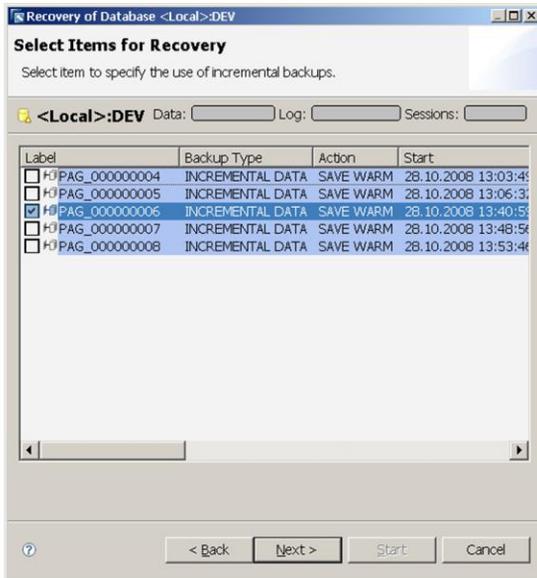
- If you select this option, all log information that has not already been backed up is overwritten irrevocably. This means that transaction data may be lost.
- If the log volume is defective, it is important to choose this option so you can repair the log. You must perform another log backup beforehand, however, to extract as much of the log from the defective log volume as possible.
- If you are creating a system copy of another database for this instance, you must select this option so that the log volume is deleted and the old log information does not conflict with the different dataset.

You can specify a time in the past up to which data is to be recovered. This affects only the recovery of log entries, however; all information already contained in the data backups is recovered completely. A recovery time can only be taken into account when log entries are recovered from additional log backups or the log volume.

The subsequent steps depend on your selection. In this example, the *Recover specified backup from history* option is selected.

All complete data backups are then listed as the data basis (starting point) for the recovery.

Performing Recoveries III



You select the complete backup as the starting point and now continue with the recovery process.

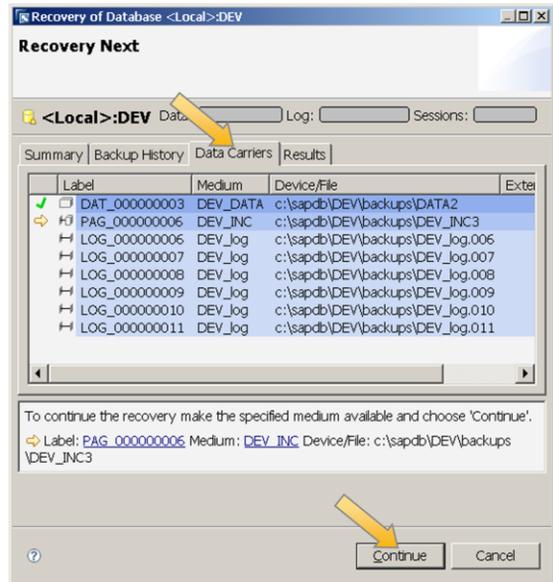
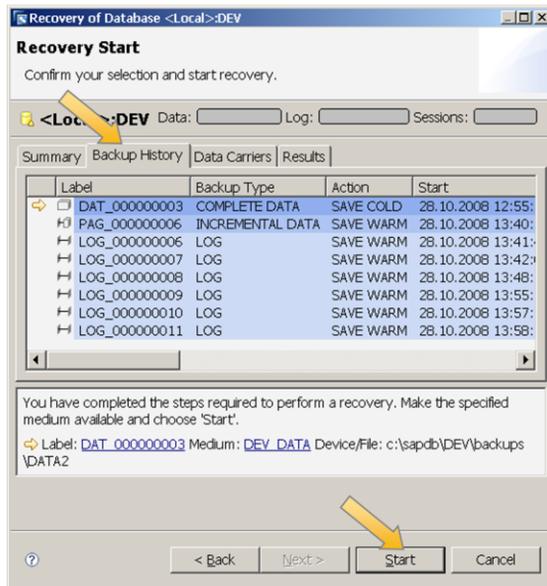
You now have to select the incremental backup, if available.

The database then fills any transactional gaps with log backups and/or the online log volume.

The *Summary* tab page displays this information.

Since recovering transactional log information is often a lengthy process, we recommend that you use the most recent incremental backup available. This directly speeds up the recovery process.

Performing Recoveries VI



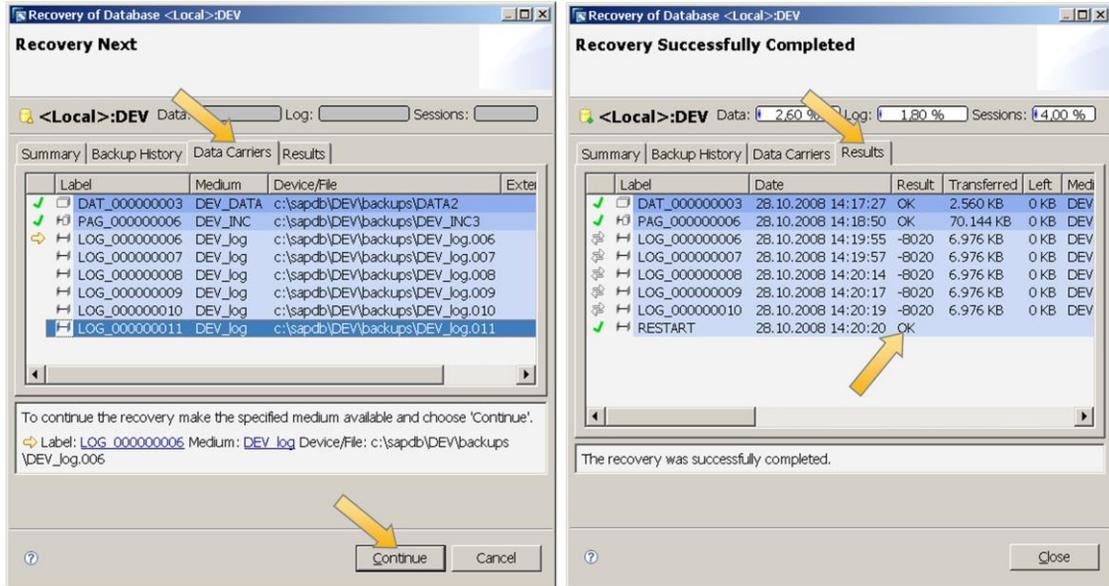
The "Start" button now becomes active and you can start the recovery.

The process stops after the complete data backup has been recovered, and you can continue the process by recovering the incremental backup.

Alternatively, you can stop the recovery at this point by choosing *Cancel* and restart it later as long as the database instance has not been switched from the ADMIN to ONLINE operational state in the meantime.

- This makes it possible to build a shadow instance manually. The next unit provides more information about these downtime security solutions with MaxDB.

Performing Recoveries VII

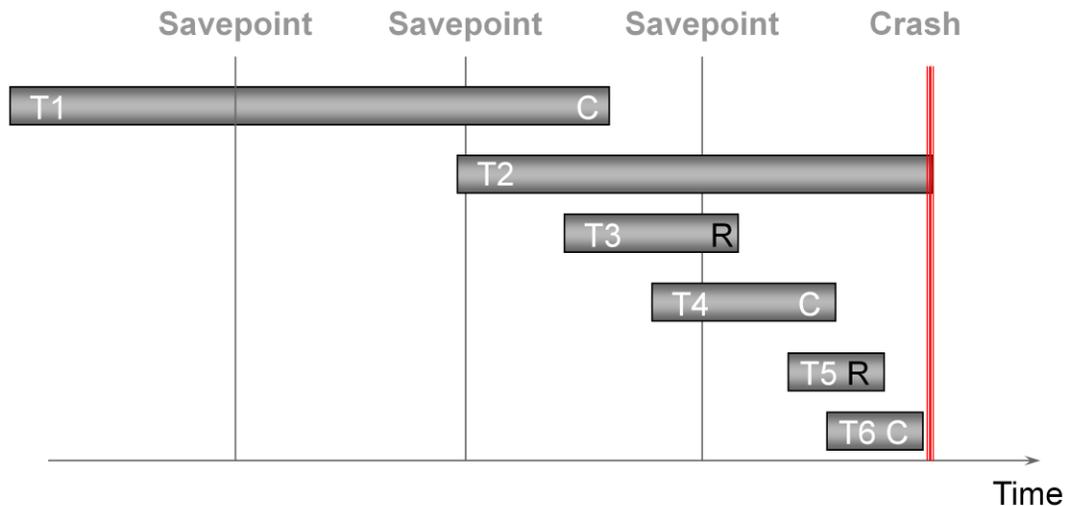


The system behaves in the same way after the incremental backup has been recovered: The Wizard stops and gives you the option of interrupting the recovery process.

At the end of the recovery process, the "Results" tab page displays a summary of the process including the return values of the various steps.

- Return value -8020 displayed here is not an error, but an indication that this log backup has been completed successfully and the next log backup has been requested.

Handling Transactions During Restarts



Transactions T1 and T5 not recovered

Transactions T4 and T6 recovered

Transactions T2 and T3 rolled back

© 2010 SAP AG. All rights reserved. / Page 50

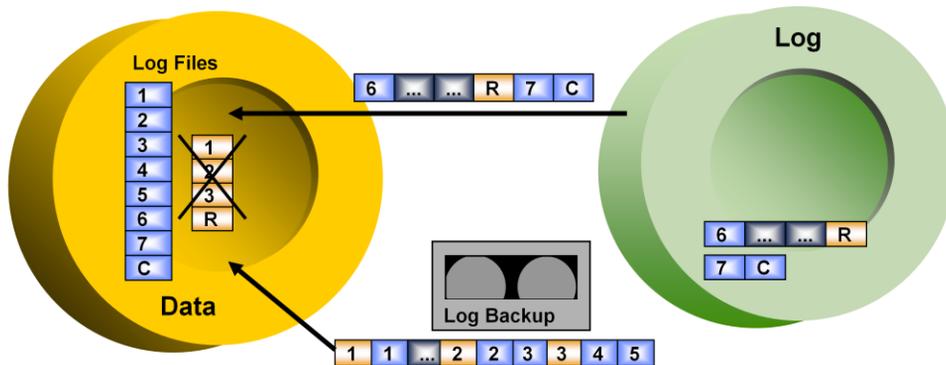
Whenever the database is restarted (for example, after an error), the system searches for open transactions or changes made after the last savepoint and starts an automatic recovery if necessary.

Only the data that was valid at the time of the last savepoint was written to the data volumes and is still available now. If the database crashes, it is only possible to trace the changes made between the savepoints by viewing the log entries.

The last savepoint is now used as the starting point for the automatic recovery.

The transactions shown are handled as follows:

- Transactions T1 and T5 are not recovered:
T1 was completely backed up at the last savepoint before the database crash.
T5 was started after the last savepoint but manually interrupted and rolled back before the crash, making it superfluous.
- Transactions T4 and T6 are recovered:
T4 was active at the time of the last savepoint and completed successfully.
T6 was started after the savepoint and completed successfully.
- Transactions T2 and T3 are rolled back:
T2 was active during the crash and must, therefore, be rolled back.
T3 was active at the last savepoint. This means the changes made by this transaction are included in the savepoint and must, therefore, be removed since the transaction was rolled back after the savepoint.



The log entries are imported in parallel.

No changes are made to the current online log volume.

All transaction-relevant actions are collected in log files for each transaction within the data volumes.

The log information can be taken from a log backup or log volume.

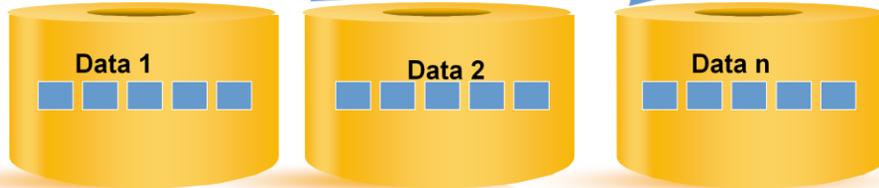
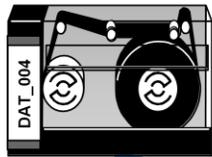
The transaction to be recovered is processed as soon as all its data has been collected.

Recovering the log entries, therefore, requires space in the data volume and can, in theory, fill the data area completely depending on the quantity of data in the log.

Performing Recoveries – Complete Data Backups

SAP

Restore
DAT_004



© 2010 SAP AG. All rights reserved. / Page 52

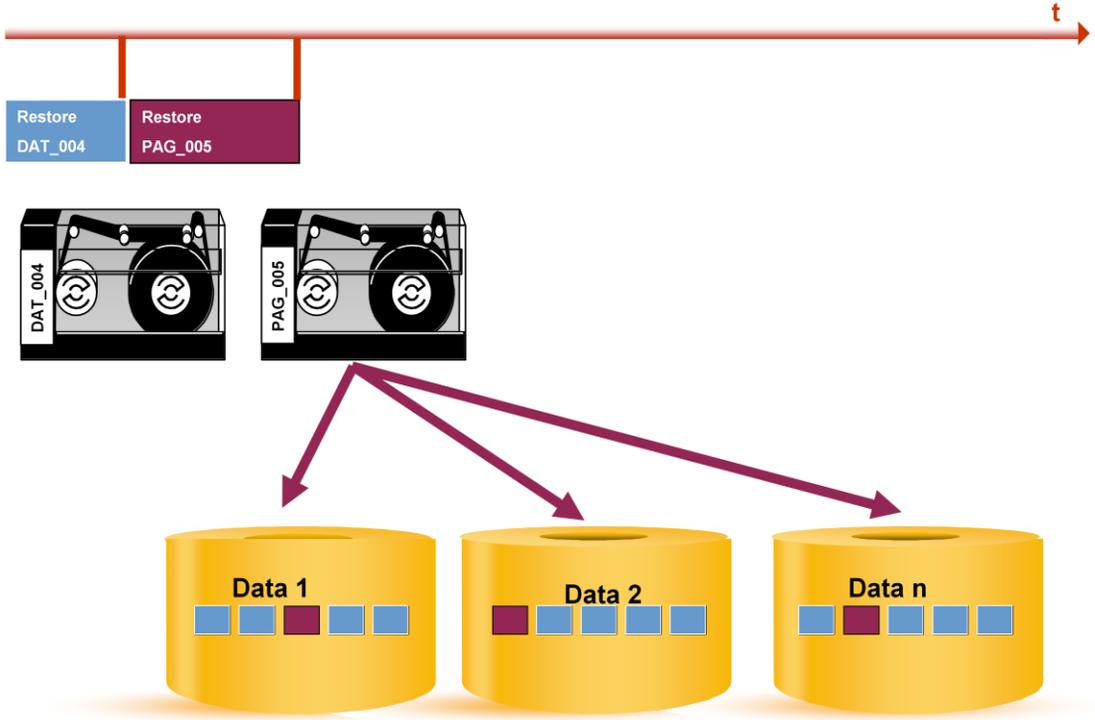
This figure shows a summary of the database processes carried out when the different backups (data and log) are recovered.

Step 1:

Switch the database to the ADMIN operational state.

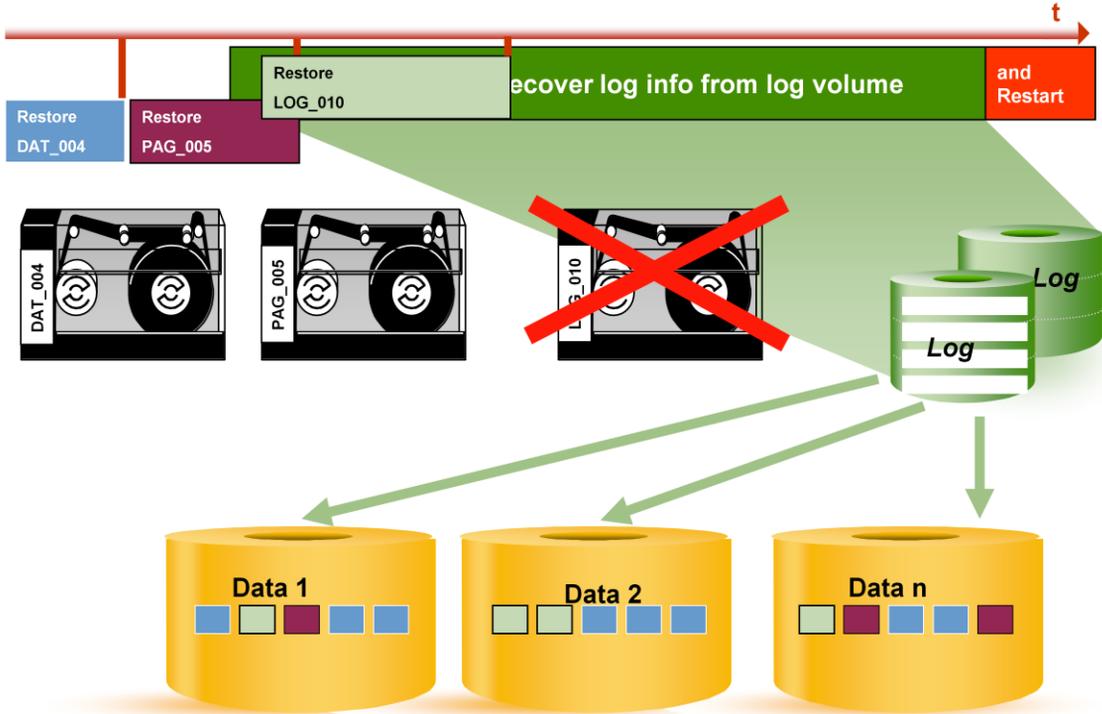
The recovery of the complete data backup is started and the pages are transferred to the data volumes.

Performing Recoveries – Incremental Backups



Step 2: The existing pages from the complete data backup are overwritten with the pages from the incremental data backup.

Performing Recoveries – Log backups

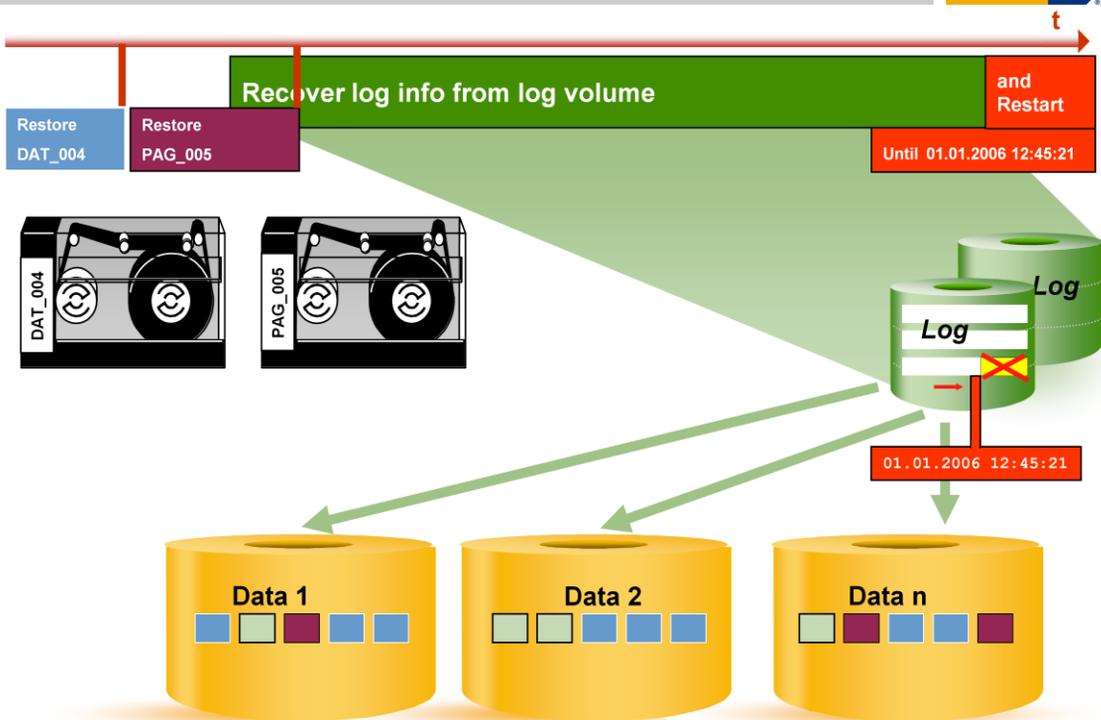


© 2010 SAP AG. All rights reserved. / Page 54

Step 3: Restoring changes after data backups

- Option 1: All the required data still exists in the log volume, even if one or more log backups have already been created. All log entries are taken from the log volume and not from the log backup.
- If you switch the database instance to the ONLINE operational state, the log entries are recovered.

Performing Recoveries – Recovering Log Entries

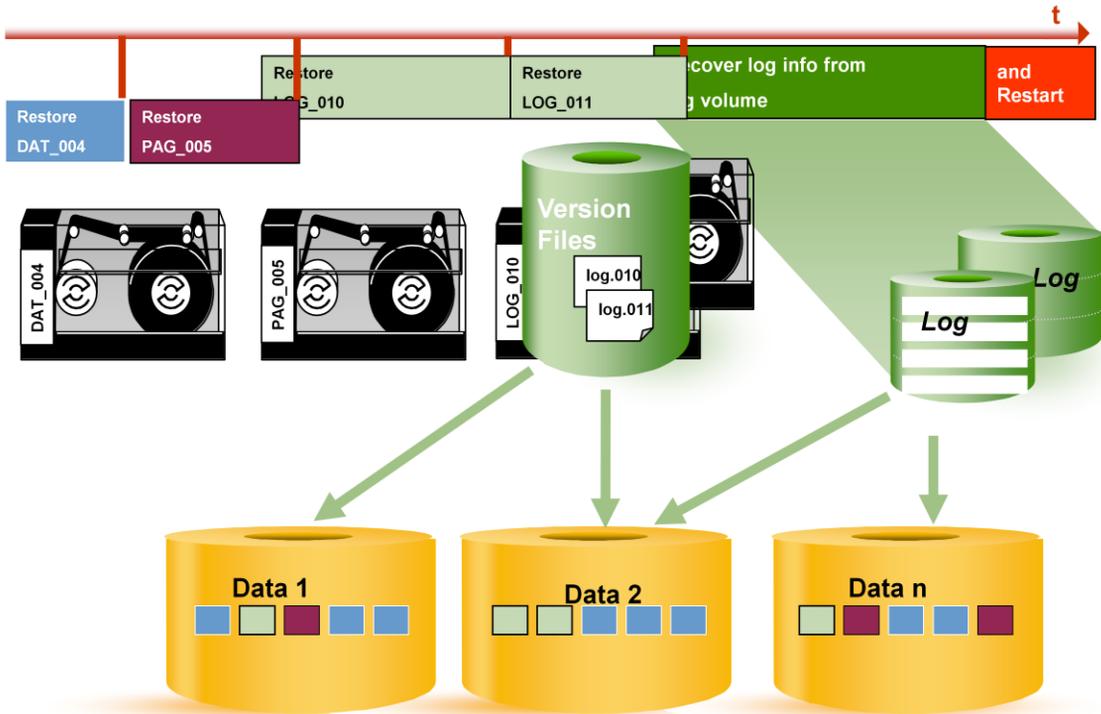


© 2010 SAP AG. All rights reserved. / Page 55

Step 3: Restoring changes after data backups

- Option 2: The database instance is only to be recovered up to a specific time in the past to undo an input error and return to the original state of the data.
- You specify a *Recover Until* time for starting the database instance and recovering the log entries. Log entries created after this time are ignored and deleted.
- To ensure that this log information is available for subsequent restore operations (such as another "Restore Until" with a later time specified), we strongly recommend that you create a log backup **before** the *Recover Until* time.

Performing Recoveries – Log Version Files

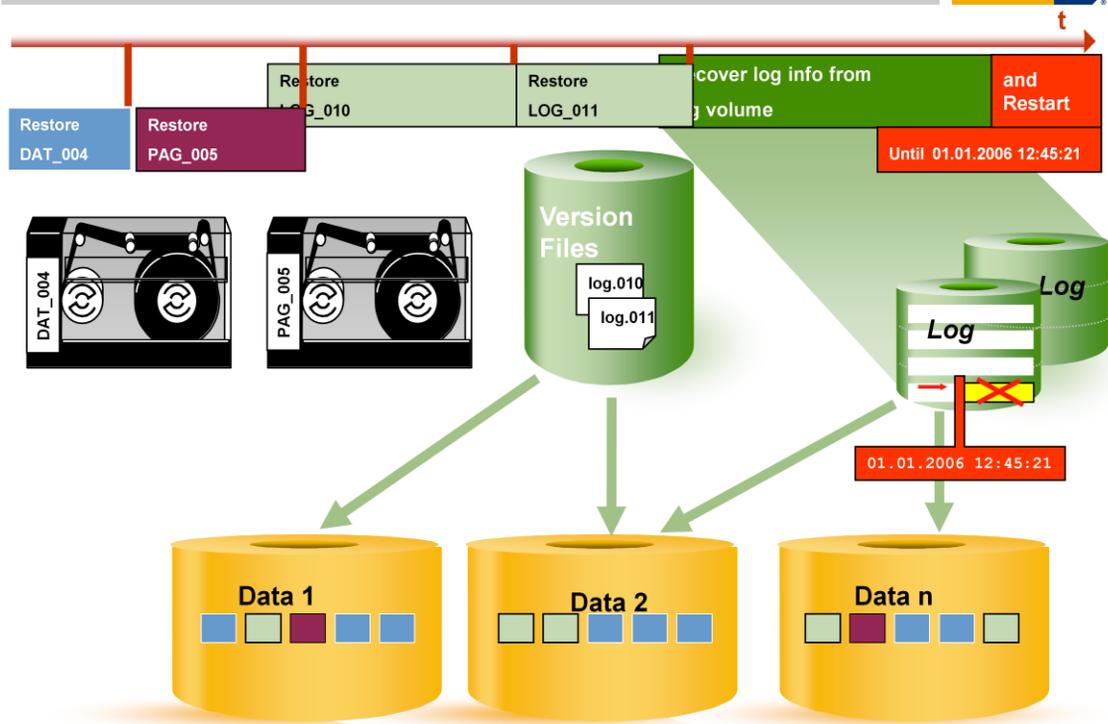


© 2010 SAP AG. All rights reserved. / Page 56

Step 3:

- Option 3: Since the last log backup, so many log entries have been created due to the operation of the database instance that the log area released for overwriting has been overwritten with newer log entries.
- The log information that no longer exists in the log volume must be recovered from the log backups.
- This is done using the data volumes so that the log entries already in the log volume are not overwritten.
- If the log backup files have already been saved to tape, they must be made available again on the disks under the same name.
- The log entries are recovered when the database instance is switched to the ONLINE operational state.
- The log entries from the backups are read up to the first entry that exists in the log volume.

Performing Recoveries – Log Version Files



© 2010 SAP AG. All rights reserved. / Page 57

Step 3:

- Option 3: This variant also allows you to specify a *Recover Until* time.
- This time can lie in the period during which log backups were performed.

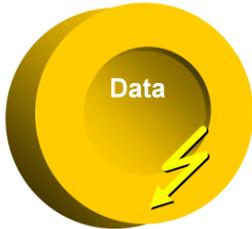
Agenda



1. **Introduction to Backup and Recovery**
 - SAP MaxDB Backup Concept
 - Why Backup using SAP MaxDB?
2. **Backup**
 - Backup Types
 - Backup Templates
 - Snapshots
 - Check Backup
 - Backup Strategy
 - 7.8 Specific Features
3. **Recovery**
 - Recovery Strategy
 - Recovery Types
4. **Monitoring and Troubleshooting**
 - Common Errors and their Solutions
 - How to Monitor a Backup / Recovery
5. **Questions and Answers**

In this 4th and last chapter, we give anybody interested in the topic an introduction into the world of monitoring and troubleshooting typical Backup and Recovery questions and issues. It is meant as a first step towards using more advanced SAP MaxDB analysis tools.

Failure of a
Data Volume

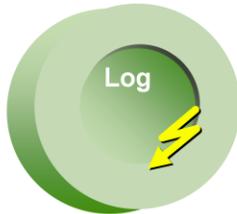


Normal recovery of the database instance

- The existing log volume is used to restore the most recent state of the database instance.
- If the log entries required are no longer in the log volume (because they have been backed up and then overwritten in the log volume), you must also recover log backups.

This slide shows the typical action to be undertaken when a data volume fails, for example when the hardware on which it is located, can no longer be accessed. In this simple representation, we assume the log volume is unaffected, you can simply recover the last full data backup in ADMIN mode and the database should reach state ONLINE after the ReDo has finished on the logvolume.

Failure of a log volume



Procedure is similar to a system copy (see SAP note 129352)

If DB can reach the state ADMIN:

- To keep the loss of data as small as possible, execute a data backup in the state ADMIN. This saves all changes that were stored in the data area with the last savepoint.
- To recover the database, choose the recovery with initialization and use this backup.

If DB can not reach the state ADMIN:

- If a data backup in the state ADMIN is no longer possible because the data volumes are also affected by the disk error, you must execute a recovery with initialization using the last data backup and all of the following log backups.

If DB can reach admin mode:

The loss of data includes the period as of the last savepoint, which you can determine using the detailed information for this backup under 'Last Save Point' (Database Studio -> Administration -> Backup -> Details or DBMCLI command `backup_history_list`).

We always recommend that you back up the most recent log information from the online log volume one more time, even if parts of the log volume may be defective. The backup process then saves all intact log pages until the first page error occurs.

If DB can not reach admin mode:

In this case, the loss of data includes the period as of the last log backup.

Additional useful info:

If you are using the mirroring log mode (mirrored volumes) and one of the two log volumes has failed, refer to SAP note 1177052: Mirrored log: Recovery of deficient log volume



Problem: Backup fails due to operating system parameters that are inadequately configured.

```
(edited) KnlMsg excerpt -

2009-09-14 10:11:08 ERR RTEIO
1: The opening of path '/sapdb/PBI/sapdata/DISKD0015'
failed,DESCRIPTION=open(READ,EXISTING),ERRORTEXT=Too many open
files,RETURNCODE=24,_FILE=RTEIO_StreamMedium+noPIC.cpp,_LINE=5144

DESCRIPTION:

An attempt to open the path '/sapdb/PBI/sapdata/DISKD0015' failed. The
operating system returned the error code 24 and the error text 'Too
many open files'. The complete command line reads 'open
READ,EXISTING)'.

...
Thread 0x69C Task      - 2009-09-06 18:23:52
RTEKernel   36: Resource limits
RTEKernel   37: Maximum CPU time: unlimited
RTEKernel   38: Maximum number of processes: 522240
RTEKernel   39: Maximum core file size: 0 MB
RTEKernel   40: Maximum number of open files: 1024
RTEKernel   41: Maximum file size: unlimited
RTEKernel   42: Maximum number of threads: unlimited
RTEKernel   43: Maximum stack size: 10 MB
RTEKernel   44: Lockable memory size: 0 MB
RTEKernel   45: Maximum virtual memory size: unlimited
RTEKernel   46: Resident set size: unlimited
RTEKernel   47: Maximum heap size: unlimited
Thread 0x69C Task      - 2009-09-06 18:23:52
RunTime     9: Program environment dump
```

© 2010 SAP AG. All rights reserved. / Page 61

SAP note 628131: SAP DB/MaxDB operating system parameters on Unix

2. General parameters

a) User limits

For this example, Maximum number of open files (maxfiles,nofiles) is relevant.

- maxfiles, nofiles >= 2048 (for SUN OS: >= 4096)

Generally, a value of 2048 is sufficient for MaxDB OLTP systems. For very large database instances, check and increase the value using the formula in the note.

How to Monitor the Recovery Activities I



Using the Database Console (x_cons) tool to check if Recovery is progressing:

```
C:\sapdb\programs\bin> x_cons <database_name> show active
```

Recovery of Data Backup:

```
SERUERDB: MAX76510

ID   UKT  Win  TASK          APPL Current      Timeout Region      Wait
   tid type          pid state          priority cnt try      item
T77   7  0xB48 User          4260 JobWait BckRec         0 0         449(s)
T220  4  0xEB0 BUPvol        Ususpend (204)  0 0         435672(s)
T221  4  0xEB0 BUPvol        Ususpend (204)  0 0         435672(s)
T222  4  0xEB0 BUPmed        AsynWaitRead    0 0         435672(s)

Console command finished (2010-07-28 16:49:12).
```

'show active' command -

provides an overview of the states of all active database sessions (user tasks) and special tasks. Only those tasks that are currently processed by the database kernel are displayed. If a user task is waiting for a new command with status, 'Command Wait', then it is not displayed. Special tasks are displayed only when they are not in a sleep or suspend state (vsleep, vsuspend). You can also show the statistics for specific task groups only: DW (data writer), SV (server task), US (user task), GC (garbage collector).

Reference material: x_cons documentation on SAP MaxDB SDN Wiki pages.

How to Monitor the Recovery Activities II



Using the Database Console (x_cons) tool to check if Recovery is progressing:

```
C:\sapdb\programs\bin> x_cons <database_name> show active
```

Recovery of Log Backup:

```
SERVERDB: MAX76510

ID   UKT   Win   TASK      APPL Current      Timeout Region   Wait
   tid type      pid state      priority cnt try   item
T73   7   0xB48 User      4260 JobWait BckRec      0 0      696(s)
T89   4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T90   4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T91   4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T92   4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T93   4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T94   4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
...
T151  4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T152  4   0xEB0 RedoExe   NoRedoJob(231) 0 0      850073(s)
T153  4   0xEB0 RedoRea   Ususpend (204) 0 0      850073(s)
T221  4   0xEB0 BUPvol   JobWait Redo   0 0      850073(s)
T222  4   0xEB0 BUPmed   AsynWaitRead   0 0      850073(s)

Console command finished (2010-07-28 17:40:28).
```

Task that started the recovery but will not do the actual work. It will wait for backup or recovery to complete:

```
T73   7   0xB48 User      4260 JobWait BckRec      0 0
696(s)
```

Reading data from the log backup:

```
T221  4   0xEB0 BUPvol   JobWait Redo   0 0
850073(s)
```

```
T222  4   0xEB0 BUPmed   AsynWaitRead   0 0
850073(s)
```

How to Monitor the Recovery Activities III



Using the Database Console (x_cons) tool to check if Recovery is progressing:

```
C:\sapdb\programs\bin> x_cons <database_name> show io
```

Recovery of Data Backup:

```
SERUERDB: LU1

Devspace          No.      Read(s)    RPages    Write(s)    WPages
-----
knltrace          0         0          0         1           1
/sapdb/LU1/sapdata/DISKD0001  1      93171      93171      0           0
/sapdb/LU1/sapdata/DISKD0002  2     94031      94031      0           0
/sapdb/LU1/sapdata/DISKD0003  3     93147      93147      0           0
/sapdb/LU1/sapdata/DISKD0004  4     93290      93290      0           0
/sapdb/LU1/sapdata/DISKD0005  5     93821      93821      0           0
/sapdb/LU1/sapdata/DISKD0006  6     93965      93965      0           0
/sapdb/LU1/sapdata/DISKD0007  7     93995      93995      0           0
/sapdb/LU1/sapdata/DISKD0008  8     93385      93385      0           0
/sapdb/LU1/sapdata/DISKD0009  9     93219      93219      0           0
/sapdb/LU1/sapdata/DISKD0010 10     93881      93881      0           0
/sapdb/LU1/saplog/DISKL001   11         1          1          1           1
total I/O                935906    935906      2           2

Console command finished (2010-06-07 19:54:06).
```

‘show io’ command -

This command displays the number of read and write operations per volume as well as the number of 8KB page reads. These numbers are independent of whether the I/O is synchronous or asynchronous.

How to Monitor the Recovery Activities IV



Using knldiag/knlmsg to check how Recovery is progressing:

```
RESTORE 52024: 20200 pages <- "sapdb\backup\sc1_log.003"
Log      20012: redo-read#150000:TR11423019(0)[586588]@NIL.2704'InsertRecord':20091114:171203
...
RESTORE 52024: 25600 pages <- "sapdb\backup\sc1_log.003"
Log      20012: redo-read#250000:TR11455362(0)[620667]@NIL.2064'UpdateRecord':20091120:232718
RESTORE 52024: 26800 pages <- "sapdb\backup\sc1_log.003"
RESTORE 52024: 27800 pages <- "sapdb\backup\sc1_log.003"
...
IO       19626: Async I/O thread stopped, "\\PALD00412980A\sapdb\backup\sc1_log.003"
RESTORE 52024: 28656 pages <- "sapdb\backup\sc1_log.003"
...
Log      20044: normal end of log found at off NIL lastseq 634496.
Log      20012: last-redo-read empty errlist#289254:TR11468216(1)[634496]@NIL.2344'Commit':20091123:110947
Savepoint 1: Savepoint (Distance) started by T153
Pager    20004: SUP(1) Start Write Data
...
SAUPOINT 53071: B20SUP_COMPLETED: 2496
Rst      5: RedoControl: End of redo processing
Rst      4: RedoControl: RedoTasks: 64 , MaxRedoTrans: 168,LOG_IO_SEQUENCE=536919
Rst     12: RedoControl: redo reader T0 wait reason: isNotWaiting, waiting redo executor: 0
Rst      8: RedoControl: Transactions processed: 88484 (88484)
Savepoint 1: Savepoint (Distance) started by T221
...
SAUPOINT 53071: B20SUP_COMPLETED: 2498
Admin    20039: Kernel state: 'REDO LOG' finished
...
RunTime  3: State changed from ADMIN to ONLINE
Admin    20041: Kernel state: 'ONLINE' reached
```

© 2010 SAP AG. All rights reserved. / Page 65

While a Recovery is active, you can also check upon it's status in the KnIMsg file (versions 7.6+, otherwise in 'knldiag').

In the above example, you can see how many pages have been recovered already, that the end of the Redo Log has been reached and a Savepoint is written.

You can even see that the state of the DB is then changed from ADMIN to ONLINE, i.e. the Recovery was successful!

Agenda



1. **Introduction to Backup and Recovery**
 - SAP MaxDB Backup Concept
 - Why Backup using SAP MaxDB?
2. **Backup**
 - Backup Types
 - Backup Templates
 - Backup Strategy
 - 7.8 Specific Features
3. **Recovery**
 - Recovery Strategy
 - Recovery Types
4. **Monitoring and Troubleshooting**
 - How long will the Backup / Recovery take?
 - How to Monitor a Backup / Recovery
 - Common Errors and their Solutions
5. **Questions and Answers**

© 2010 SAP AG. All rights reserved. / Page 66

Now, any remaining questions can be answered.

Any other remaining questions, that might require longer answer, or if time no longer permits answering, please direct your questions at our SAP MaxDB / liveCache forum on sdn.sap.com

Questions and Answers



SAP MaxDB resources In need of more information?



Home	maxdb.sap.com
SAP Education	The place for SAP Education
SDN Forum	SAP MaxDB / livecache
SDN Wiki	SAP MaxDB Wiki

© 2010 SAP AG. All rights reserved. / Page 68

In case you're ever in need of more information on any kind of subject on SAP MaxDB (or liveCache), please direct your search towards:

- The SAP MaxDB site: maxdb.sap.com. This site contains a lot of information, ranging from our official documentation to the recordings of previous Expert Sessions
- Next is the official SAP Education site: it contains MANY offers for all kind of courses on any SAP topics, including for example, the ADM515 administration course on SAP MaxDB and the UMEW60, concentrating on SAP MaxDB monitoring and optimization.
- Then, we have the heavily use SAP MaxDB forum. In case of questions on SAP MaxDB products, please register and join the Community!
- Lastly, we have our also equally well visited Wiki pages. We've added a lot of information here that might interest any SAP MaxDB DBA, including a documentation on tools like `x_cons` and a Support Guide.

Reference material used:

- `x_cons` documentation on SAP MaxDB SDN Wiki pages.

- AMD515 (altered version)
- Expert Session 8: New Features in SAP MaxDB Version 7.8

Thank You!
Bye, Bye – And Stay Tuned for Further Sessions



Sept 7, 2010

Session 11: Analysis of SQL Locking Situations

Feedback and further information:
<http://www.sdn.sap.com/irj/sdn/maxdb>

Thanks a lot for your interest in this session!

The recording of this session will be available on maxdb.sap.com/training soon. There you'll find also the recordings of all previously held Expert Sessions about SAP MaxDB.

Further Expert Sessions about SAP MaxDB are currently planned. It will be posted on our SDN page <http://www.sdn.sap.com/irj/sdn/maxdb> if and when these sessions will take place.

Thank you!



Thanks a lot for your interest in this session.

The recording of this session will be available on maxdb.sap.com/training soon. There you'll find also the recordings of all previously held Expert Sessions about SAP MaxDB.

Further Expert Sessions about SAP MaxDB are currently planned. It will be posted on our SDN page <http://www.sdn.sap.com/irj/sdn/maxdb> if and when these sessions will take place.