

Tamino

Backup Guide

Version 9.5 SP1

November 2013

This document applies to Tamino Version 9.5 SP1.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 1999-2013 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, United States of America, and/or their licensors.

The name Software AG, webMethods and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA, Inc. and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://documentation.softwareag.com/legal/>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://documentation.softwareag.com/legal/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices and license terms, please refer to "License Texts, Copyright Notices and Disclaimers of Third-Party Products". This document is part of the product documentation, located at <http://documentation.softwareag.com/legal/> and/or in the root installation directory of the licensed product(s).

Document ID: INS-BACKUP-95SP1-20131030

Table of Contents

Backup Guide	v
1 General Backup Strategies	1
Concepts	2
Requirements for Backup, Logging, and Restore/Recover	4
Recovery from Data Failures	5
Internal and External Backup Solutions	9
Time Considerations	10
2 Internal Backup and Restore in Tamino	11
Backup	12
Restore and Recover	14
Summary	17
3 External Backup in Tamino	19
General Considerations about External Backups (EMC and Network Appliance)	20
Configuring External Backups with EMC or Network Appliance Systems	21
Steps (EMC and Network Appliance)	21
External Backup with Tivoli Storage Manager	25
4 Disaster Recovery	27
Recovering from Internal Configuration Store Corruption	28
Recovering from the Loss of the Host Environment	29
Index	33

Backup Guide

The Tamino Backup Guide gives general information about backup strategies in Tamino and details about how to implement them with or without external backup devices. It is intended for Tamino database administrators with a good knowledge of Tamino administration tasks.

This Guide covers the following topics:

General Backup Strategies

Concepts

Internal Backup and Restore in Tamino

Internal Backup

Restore and Recover

Summary

External Backup in Tamino

External Backup

Configuring External Backups

Steps

Disaster Recovery

Recovering from Internal Configuration Store Corruption

Recovering from the Loss of the Host Environment

1 General Backup Strategies

▪ Concepts	2
▪ Requirements for Backup, Logging, and Restore/Recover	4
▪ Recovery from Data Failures	5
▪ Internal and External Backup Solutions	9
▪ Time Considerations	10

Computers in general are very reliable. You may run your system for months or even years without experiencing any problems that cause you to lose information on your system. But businesses are more and more dependent on computers and the information that is stored in them. The information that is in your computer may not be available anywhere else. So every system needs to back up and restore some or all of its data. There are numerous backup strategies a company can use. In the following, you will find a short introduction to the concepts of backup, restore and recover for databases in general and Tamino in particular. The following topics are covered:

Concepts

In the following, you will find explanations of general notions and terms with regard to backing up databases. Most of them are available in Tamino, unless mentioned otherwise.

Backup

A database is saved to one or more output devices. Note that the term “backup” is used both for the process of saving the data and for the resulting data sets. Making a backup should be possible online, parallel to normal database update activities, and save a transaction-consistent state of the database. Backups should be done at a time when there is a low data load. Several backup concepts are conceivable:

Online backup: A backup during a normal update database session.

Offline backup: A backup when database updates are disabled (the server is down, or in stand-by mode, or in read-only mode).

Complete backup: All data of the whole database or the logical or physical subset of the database is saved.

Incremental backup: Only the data which has been changed since the previous backup is saved. A recovery is only possible if a previous full backup is available, as well as all following incremental backups. Incremental backup and recover operations are fast and efficient.

Full backup: The complete database is saved.

Partial backup: Only a logical or physical subset of the database is saved.



Note: Please note that partial backups are not available in Tamino.

Internal backup: The database system itself saves the information of the database.

External backup: Another system outside of the database saves the content of the database.

Restore

A restore *recreates* the content of the database at backup time from the backup devices.

Full restore: The complete database is restored. A full restore is only possible after a full backup.

Partial restore: Only a logical or physical subset of a database is restored. A partial restore is possible after a full or partial backup (but not available in Tamino).

Recover from Database Logging

A database backup alone allows you only to restore one state of the database as it was at backup time. After a data failure, however, it may not be sufficient to recreate that state of a database, but the state of the database just before the failure occurred. For this reason, all update operations are logged on log spaces.

After the database has been restored, the log spaces are read and the logged update operations are repeated, so that the database is returned to the state that was valid at the time when the last log entry had been created. This process is the recover process.

Full recover: The complete database is recovered. A full recover is only possible after a full restore.

Partial recover: Only a logical subset of the database is recovered. A partial recover is possible after a full or partial restore (but not available in Tamino).

Non-Parallel Backup and Restore/Recover

Normally, a backup is performed in non-parallel mode: The data blocks are written in one stream to the backup devices, or read in one stream from the backup devices.

Parallel Backup and Restore/Recover

A parallel backup writes in parallel to more than one backup device; a parallel restore reads in parallel from more than one backup device. This increases the speed of the backup/restore process, if the backup device is much slower than the disks where the database is stored.

External Backup Based on Mirroring

A copy of the database is generated on a separate volume. For mirroring, there are two possibilities:

1. Mirroring when the backup is started. In this case, it takes some time until the backup finishes. But note that if you do the next backup on the same logical volume, only the blocks modified in the meantime must be modified.
2. Mirroring starts some time before the backup is started, e.g. directly after the previous backup. In this case, the backup must only stop the mirroring, and the time required for the backup is very short.

External Backup Based on Snapshots

Alternatively, you can perform an external backup based on snapshots. A snapshot is not a physical, but a logical copy of the database. When a data block is updated after a snapshot has been created, the block is *not* updated, but copied to a new place. While the snapshot still references

the old block, the original file references the updated block at the new place. Since generating a snapshot does not perform a physical copy of the data, it is very fast.

Replication

Conceptually different from a restore process, a database can be *created* from backup. In this case a new (duplicate) database is created with the content of the database at backup time.

Changes to the database that occur after the creation from backup can then be replicated in a replication database. Unlike a conventional backup that is restored from tape or CD, the replicated database is available to applications as soon as they can be pointed to it. For further information, see the Replication Guide.

Requirements for Backup, Logging, and Restore/Recover

After having mentioned the basic notions of backing up, the question arises why we need backup, restore and recover at all. The simple answer to that is that you want to be able to recreate a previous state of a database after an error has occurred. The reasons for errors are manifold and are dealt with in detail in the next section [Recovery from Data Failures](#). Let us first consider a few requirements for being able to recreate a previous state of a database.

Database Synchronization

When a backup is performed online, it is important to be able to create a consistent state of the database after the corresponding restore. To achieve this, in Tamino a database synchronization is performed at the end of the backup: New transactions are postponed until all open transactions are finished. When all updated blocks have been written to disk, the database is in a consistent state. When this state of all database blocks is contained in the backup, it is possible to restore this (consistent) state of the database.

The restore operation recreates this consistent state of the database. Note that the restored database must be logically identical, but may be physically different. For example, the restore operation can defragment the data.

When the database server is active, it logs all update operations in the database log. After you have restored the database, the recover operation reads the logs and repeats all update operations which have been performed until the required timestamp or until the end of the logs.

Other Requirements

- If you want to be able to perform a recover after a disk error, it is necessary that the database logs and the backups are not on the same disk. This is not necessary if you only want to be able to recover from a software failure, because you have a hardware solution which guarantees that the database is not destroyed because of a hardware failure.

- Disaster Recovery (based on restore/recover) requires that the backup and log spaces be physically copied to a new computer center. For example, if the current log space is copied only after it has been closed, you are not able to reapply the changes that occurred during the current server session.

An alternative to performing a backup is just copying the database spaces to another place. But this has some disadvantages: First, it is only allowed if no update session of the database server is active. Otherwise the saved database spaces are inconsistent. Second, Tamino does not know of these “backups”. This means that old log spaces are not deleted and not released by Tamino. In addition, log spaces cannot be applied after a restore. For this reason, it is not recommended to copy the database spaces to another place instead of performing a normal Tamino backup.

Recovery from Data Failures

One of the most important tasks a database administrator has to accomplish is to define for each database how to handle data failures. There basically are four different kinds of data failure which can occur:

- [Hardware Errors](#)
- [Software Errors](#)
- [Disaster](#)
- [Second Failure](#)

Hardware Errors

A typical hardware error which may destroy a database is a disk failure. In this case, a Tamino restore/recover is a good possibility to handle the situation (see [Internal Backup and Restore in Tamino](#)). If you want to be able to perform a recover after a disk error, it is necessary that the database logs and backups are NOT on the same disk. Normally, a disk error is noticed as soon as it occurs. Hardware errors, that are not immediately recognized, are more problematic, for example if a disk read operation does not display an error, but returns wrong data. This situation is similar to software errors (see next section [Software Errors](#)). Other solutions for handling disk errors are external restore/recover operations with physically separated storage devices (see section [External Backup in Tamino](#) in this Backup Guide) or with saving backups on a tape.

There may be other hardware errors which do not require a restore/recover, but for example a new database start to be performed after repairing the hardware. Note that there are also other concepts of handling disk errors, for example RAID 5 or cluster solutions:

- RAID 5 or disk mirroring: The data is stored redundantly on the disks. If a disk is corrupted, the data is automatically read from other disks. The computer operator must only replace the corrupted disk. The data is automatically recreated on the new disk.

- Replication: After a hardware error has occurred, a replication of the database becomes the master database. It must be made available with the name of the original database. The advantage of this solution is that the database is available without time losses required for a restore/recover process. On the other hand, some transactions may be lost in this process.

The following table compares various possibilities available to recover from a hardware error:

Solution	Special Hardware Requirements	Recovery Time	Loss of Data	Remarks
Internal Restore/Recover in Tamino	None	Long	No	-
External Restore/Recover in Tamino, with physically separated storage devices	Yes	Restore time: short; Recover time: long (same as for internal restore)	No	If you have systems like EMC or Network Appliance, these systems normally use RAID technology, so that the failure of a single disk does not cause problems. There is, however, a small probability that more than one disk or even the complete storage system crash simultaneously. For these rare cases, the database administrator should provide a recovery solution. This could either be a system with physically separated storage devices, or saving the backup to tape.
External Restore/Recover in Tamino, plus saving the backup to tape	Yes	Long, but because of the especially fast and expensive hardware less than with standard hardware	No	(same as above)
RAID 5 or disk mirroring	There are hardware or software based solutions, where the operating system manages the disks	None (the user does not notice that there is a disk failure)	No	If the system is not based on physically separated storage devices, an additional recovery solution should be provided in case the whole storage system fails.
Tamino Replication	None	Short, but in contrast to high availability, the replication database	Yes; because the replication is done	This solution allows also recovery from other hardware errors, for example CPU failures.

Solution	Special Hardware Requirements	Recovery Time	Loss of Data	Remarks
		must be made available as the master database manually.	asynchronously, the last transactions may be lost.	

Software Errors

Contrary to recovery from hardware errors, an automatic recovery from software and handling errors is not possible. For the system, a software error is like a normal update operation. The database administrator has several possibilities for handling the problem:

- Perform a restore/recover to a state before the error occurred.
- Tamino software error: In some cases it might help to restore a backup created before the erroneous Tamino version was installed, and to recover all logs with a Tamino server in which the problem had been fixed. (Note that it is possible to restore a backup from a former Tamino version, even if that Tamino version is not installed).
- Try to repair the error, for example by updating the corrupted data or unloading the data that is not corrupt, deleting the corrupted data and reloading the correct data.

The solution depends very much on the individual error situation. Nevertheless, it is useful to perform regular backups, so that backup and restore/recover is a feasible possibility in each situation.

Disaster

It may happen that not only part of the hardware is erroneous, but that the whole hardware system is destroyed, even the complete computer center. In this case, it is necessary to make the data available on another computer, in a different place. This scenario is called disaster recovery. Concepts of disaster recovery are not necessarily based on backup and restore mechanisms. You can also use replications or cluster solutions with physically distributed storage devices. In any case all data required for the disaster recovery must be saved at a remote location. The following table shows the various possibilities you have for disaster recovery in Tamino:

Solution for Disaster Recovery	Special Hardware Requirements	Required Recovery Time	Remarks
Tamino Restore/Recover	None	Long	The updates of the current logs are lost if the log spaces are only copied after they have been finished.

Solution for Disaster Recovery	Special Hardware Requirements	Required Recovery Time	Remarks
Disk mirroring on remote location	Yes, but possibly there are also software-based solutions available	Short, the server needs only to be started on the target machine.	The same precautions as for a cluster solution are necessary.
Tamino Replication	No	Short, but the replication database must be manually made available as the destroyed master database.	Updates may be lost!

For more information about disaster recovery in Tamino, see the section [Disaster Recovery](#) in this guide and the documentation about *High Availability*.

Second Failure

In addition to a single hardware error, the database administrator must be aware of the fact that there is a small risk of a second failure. Standard backup solutions guarantee recovery only if not more than one disk crashes. However, if for example you perform an internal backup, a disk containing a database space has crashed, and the backup is not readable, the complete data is lost.

Depending of the kind of recovery, the following strategies can be provided in case of a second failure:

- If you have a separate solution for disaster recovery, you can use this solution for a second failure. But be aware of the fact that if the only solution for recovery from hardware failures is the disaster recovery solution, this may not be sufficient.
- If you are performing internal backups, you can use a previous backup. In this case, it is important that the different backups are stored on different physical devices. If you want to be able to also perform a recover process in the case of a second failure, you must also save the log spaces to different physical devices. If you restore an older backup, recovery will take longer than usual, because also the logs created between this backup and the backup which is no longer readable must be applied. You can avoid this by copying the backups to another physical location.
- Logs should be copied if you do not want to lose updates because logs are no longer readable.
- If you use RAID technology, you can additionally perform internal or external backups and copy the external backups to another device. You will usually reduce the frequency of the backups, for example once a week instead of once a day.

Internal and External Backup Solutions

Internal Backup

During an internal backup, the database system itself saves the information of the database. Tamino writes or reads all data in the database.

External Backup

When an external backup is performed, another system outside of the database, e.g. software supported special storage devices, saves the content of the database. The initial backup is a full backup. Following backups are incremental. There are two different techniques for external backup:

1. *Mirroring* – the database spaces are mirrored on separate logical volumes. The time required for the external backup depends on when the mirror creation was started. The first backup to a mirror disk must usually copy the entire disk, so this could take some time. However, the hardware permits updates during the copy process, so Tamino can work without interference. Just at the end of the copy the database must be synchronized with the disk and parallel update tasks may be blocked for a short time. All further backups to the same mirror disk will be treated as incremental copies, which means that only the changed data is transferred to the mirror. Hence all subsequent mirror backups could be much faster than the initial one.
2. *Snapshots* – only a logical and not a physical copy of the data is done. The original files and the snapshots reference the same physical blocks. If a block in the original file is updated, it is copied to a new location. After that, the original file references the updated block at its new location, and the snapshot still references the old version of the block at its old location.

Both concepts have advantages and disadvantages:

- Internal backups do not require special hardware, while backup systems required for external backups may be quite expensive.
- An internal backup can be used for recovery from disk failures. If you use external backup systems like for example systems from EMC or Network Appliance, you do not need an external backup for recovery from disk failures. These systems use RAID technology, so no extra external backup is necessary. Using an external backup for recovery from multiple disk errors is only possible if you can save the external backup to another device, for example to tape. This allows you to restore the data if the whole storage system should break.
- While an internal backup is relatively slow, an external backup is usually quite fast. Either only a logical copy (snapshots) exists, or the data has already been copied before (mirroring with a previous initialization of the mirror).

For information on how to back up internally, see [Internal Backup and Restore in Tamino](#). Information on how to back up with external storage devices can be found in the section [External Backup in Tamino](#).

Time Considerations

In the day-to-day administration environment, there normally is a requirement stating that the database should be up after a failure within a given amount of time, for example within two hours. This means that the database administrator must estimate the time necessary for a restore/recover process. Total recovery time is the result of summing up the restore time and the recover time. Use the figures given in the following rule-of-thumb example, to calculate an estimate of the restore/recover time.

Assume the estimated update time is half an hour and the estimated recover time for the updates of one day is a quarter of an hour, and that you have 5 working days with update activity. Assume the restore/recover time should be no more than 2 hours, after which you should do a weekly backup. If the failure occurs very shortly after the backup, the restore/recover time would be about half an hour. If the failure occurs after one week, the restore time would be about $0.5 \text{ h} + 5 * 0.25 \text{ h} = 1.75 \text{ h}$. In this case, it is recommended to perform a weekly backup. There may be 20% more update activities than usual, and the restore/recover time is still not more than 2 hours.

Note the following rules-of-thumb:

- The restore time is proportional to the backup time. Compare the backup and restore time and use the resulting factor: When the database grows, you can estimate the restore time by multiplying the current backup time by the factor.
- The recover time is proportional to the size of the log files, unless mass loads or index creation operations must be recovered. Also, the recover time for a given amount of log files may vary, depending on the number of update operations.

2 Internal Backup and Restore in Tamino

▪ Backup	12
▪ Restore and Recover	14
▪ Summary	17

The following topics describe how to internally back up and restore your Tamino database in case of any failure that can occur on your system: disk failure, program failure, or human error. When you perform an *internal* backup in Tamino, all data in the database is written block by block. Contrary to an *external* backup, no special hardware is required (see [External Backup in Tamino](#) for further information on external backups). The following topics are covered:

Backup

- [General](#)
- [How to Back Up](#)
- [Backup Generations](#)
- [Disabled Backups](#)

General

An internal Tamino backup can be a full backup of the whole database, or it can be an incremental backup that records the changes since the last full or incremental backup. When developing a backup strategy, apply this principle: the more often your database changes, the more often you should save it. If, for example, your database is very active, you should make daily backup copies. The backup is done online, which means that the server is running, so you do not even have to stop the database when backing up, or notify your users. After a full or incremental backup, a new Tamino log space is started automatically. Backups can also be done in read-only or stand-by mode. This kind of backup that is done while it is not possible to update the database is called *offline backup*.

Incremental backups provide a useful extension to Tamino's backup capabilities. If you decide to include incremental backups in your backup strategy, you should keep a sensible balance between incremental backups and full backups. Regular full backups should still be the central part of your strategy, and the incremental backups can be considered as a quick and safe way of protecting against data loss in the intervals between the full backups.

Generally, internal backups are done to another disk. If you have enough space on a disk other than your working disk, you should back up to a disk as your preferred method. This saves time and effort.



Caution: Please note that when you delete a Tamino database, all backups of the database are also deleted if the option **Keep Backups** has not been specified.

How to Back Up

▶ To create a backup of a database

- Follow the steps described in section Back up a Database in the documentation of the Tamino Manager.

Backup Generations

One important aspect to consider when backing up is the number of backup generations that may co-exist (the default value is 5). If you have enough space on your storage device, a recommendation for the number of backup copies is 7, which means one backup for each day of the week. After the seventh backup, the first one will be deleted, including the log spaces and everything else belonging to the backup.

Why should you have more than one backup? The last backup is always taken as the default, because it takes the least time for the subsequent Recover step. The backup(s) before that should only be used if the most recent backup is not usable. Security increases with the number of backup copies available.

If you use incremental backups as part of your backup strategy, note that incremental backups are not counted as additional backup generations. A backup generation can be thought of as a full backup plus all of the associated incremental backups. A new backup generation begins when a new full backup is made.



Note: Note that newer backups will be removed when an old backup had been restored and the database has been modified.

▶ To define the number of backup generations:

- Follow the steps described in section Using Backup Generations in the documentation of the Tamino Manager.

Disabled Backups

If you perform a database restore without recover, using a backup that is not the most recent backup, all subsequent backups of the database are set automatically to disabled.

Disabled backups are ignored when Tamino counts the number of backup generations for a particular database.

You can use a disabled backup to create a database, using the command `Create Database from Backup`, but otherwise a disabled database cannot be used any more.



Note: There are currently no direct features for deleting a disabled backup. However, a disabled backup will be automatically deleted when there are no more older non-disabled backups.

Restore and Recover

To be prepared for certain types of disk failures, you need to be able to recover everything on your system.

The process of regenerating your data from a backup is done in two steps: First, the restore step, which is the reverse procedure of a backup. As input, you use the backup copy of the database, which automatically restores your data back into the database. The second step, the recover step, repeats database changes since that last backup by using the session logs stored in the log spaces.

A restore is always done for the whole database, not for parts of the database, e.g. for single collections. When you restore from a full backup, the database contents are returned to their exact state at the time the full backup was made. If you restore from an incremental backup, Tamino first restores the most recent full backup that was made before the incremental backup, then restores any other incremental backups that were made after the full backup but before the selected incremental backup, then finally restores the selected incremental backup.

As soon as the restore step successfully finishes, the database is automatically started in standby mode (indicated by a yellow traffic light). Then the recover step is started. During the recover step, all of the completed database transactions that were made since the specified backup are recovered. The input for the recover step are the log spaces; each server session creates a log space that records all completed database transactions that are made during the session. By using the **Recover Until** option, you can restrict the recovery until a certain date and time.

If you restore from a full backup, the subsequent automatic recover step uses only the session logs, even if there are incremental backups available that were made after the full backup.

If you restore from an incremental backup, then Tamino restores first the corresponding full backup and then all incremental backups up to and including the selected incremental backup. The subsequent automatic recover step uses the session logs that were created after the selected incremental backup. Incremental backups that were made after the selected incremental backup are not used in the restore or recover phase.

Tamino processes incremental backups faster than session logs. Therefore, if you have made a full backup and one or more subsequent incremental backups, then you should select the most recent incremental backup for doing the restore instead of selecting the full backup.

You can view the log spaces in the Tamino Manager. During a mass loading process, the necessary information is written to so called "Utility Recovery Spaces". These spaces are handled like log spaces during the recovery step.

▶ To view the Log Spaces

- Follow the steps described in section View the Log Spaces in the documentation of the Tamino Manager.

An important aspect when restoring and recovering data is the consistency of the database. In Tamino, you can be assured that even if it was not possible to recover all changes, the database is consistent. Transactions are recovered in an “all-or-nothing” fashion. The log spaces carry a timestamp so that you can recover your data up to a certain point in time, e.g. before an application error occurred.

Also, it is possible to restore a backup of a former Tamino version, whether the older Tamino version is installed or not.

How to Restore and Recover

▶ To restore and recover a database from a backup

- Follow the steps described in section Restore a Database from a Backup in the documentation of the Tamino Manager.



Note: During the Restore/Recover step the database is in standby mode, which means that it is not possible to work with it. Hence before and after changing large amounts of data at one point in time, it is recommended to back up the database. Otherwise, the Restore/Recover step might take very long, and the database is not available for that period of time.

How to handle Restore and Recover Error Situations

When you do a Recover, you have the choice between three possibilities: **Recover all**, **Recover until**, and **Don't Recover**. In an error-free situation, you naturally would recover all changes made to the database since the last backup. The other two options are to be used if an application error occurs, or if the Restore/Recover step is not successful.

Restore Error Situations

Restoring data is just the reverse of backing up data. The data saved in the backup copy is put back to where it belongs. Possible error situations for restoring database backup copies are for example:

- The last backup copy was accidentally deleted, or cannot be found.
- The last backup copy is corrupt.

In all these cases, it is important to have a good backup strategy. If you have several backup copies, go back to the most recent one that is not corrupt, and do a Restore/Recover.

Recover Error Situations

For the Recover step, though, other error situations may occur:

- Some or all log spaces are missing for the Recover step.
- The server process is aborted while the Recover process is running.

If there are no or not all log spaces available:

The first case might occur due to a hardware failure or to human error. Maybe someone has deleted or renamed the log spaces accidentally, or the directory no longer exists. If at least some of the log spaces for the Restore/Recover process are available, Tamino processes that information up to a point in time just before information is missing or corrupt. In order to keep the database consistent, transactions that were still open when the error occurred are rolled back. After having recovered the available information, the server is stopped and shut down. The database is consistent, but not as current as it was before the error occurred. Notify your users to update the database manually.

If no log spaces at all can be found, it is not possible to recover any changes. In this case, the only possibility is to check the **Don't Recover** option and to just restore the database, without applying the changes since the last backup. Again, notify your users to update changes manually.

▶ To skip the Recover step

- Select the option **Don't Recover** in the **Restore Data** dialog, as described in section Restore a Database from a Backup in the documentation of the Tamino Manager.

If the server process is aborted while the Recover step is running:

In the second case mentioned above, you have started a Recover step, but it is interrupted because the server is no longer available, e.g. due to a power failure. When you restart the server, several messages are displayed, informing you about the error status of the Recover step. At a certain point in time, you will be asked if you want to continue the server startup:

```
Please reply "yes" or "no" to continue server startup
```

The default is `no`. Reply `no` if you do not expect that the Recover step will be aborted again (for example, if the reason for the abort was a power failure). Redo the Restore/Recover step.

Reply `yes` if an abort will probably reoccur with a second trial, for example if the reason for the abort was a corrupted log space. In this case, you will be informed that restore of existing backups created after the one used for the failed Restore/Recover run is not allowed any more. It is still possible to retry Restore/Recover for the same or older backups. For safety reasons, it is recommended to create a new backup immediately. Alternatively, you could also reply `no` and do a `Recover until`. Specify a point in time just before the server crash occurred, so that the corrupted log space which contains the error is not used.

▶ **To Recover until a certain point in time**

- Select the option **Recover Until** in the **Restore Data** dialog, as described in section **Restore a Database from a Backup** in the documentation of the Tamino Manager. Enter the date and time until which you want the Recover step to run.



Note: You can use the **Recover until** option also if you are sure that an error occurred, e.g. due to a faulty application, or if you want to delete parts of the changes that were made to a database and cannot be deleted otherwise.

Summary

A backup always applies to a whole database. If you back up your data daily and keep your backups in a safe place, you cannot lose more than a day's work. With the Restore/Recover functions you have the possibility to regenerate changes made to the database after the most recent backup. The database will always be restored to a consistent state, but in very rare cases, errors might occur during the process. If an error occurs, there are two possibilities: Either apply changes only up to a certain point in time, or skip the whole recover step and just restore data back to the latest backup copy. In both cases, notify your users to update the database manually with the changes that were lost.

3 External Backup in Tamino

- General Considerations about External Backups (EMC and Network Appliance) 20
- Configuring External Backups with EMC or Network Appliance Systems 21
- Steps (EMC and Network Appliance) 21
- External Backup with Tivoli Storage Manager 25

In Tamino, an external backup is a backup that uses another system outside of the database, e.g. software supported special storage devices, to save the content of a database.

The following topics are covered in this chapter:

General Considerations about External Backups (EMC and Network Appliance)

Tamino supports several external backup methods:

- Mirroring with EMC systems
- Mirroring with Network Appliance systems
- Snapshots with Network Appliance systems

The backup method you choose must be compatible with the used hardware. You may use all backup methods on the same computer, but for a given database it is allowed to use only *either* EMC *or* Network Appliance. If you have stored your database on Network Appliance, you can use Network Appliance mirroring and Network Appliance snapshots for the same database.

It is not possible to have more than one database on a logical volume. The reason is that the external backup with EMC or Network Appliance saves complete logical volumes, which requires a database synchronization to be done beforehand. These two activities are performed when you invoke the external backup. If a second database was stored on the same volume, a database synchronization (plus the necessary update of the internal configuration store) for that second database would not be done.

The internal configuration store and the log spaces must not be stored on the same logical volume as the database spaces. Otherwise the information on the log spaces or the log spaces themselves would be lost after a restore, without a recovery being possible.

If you restore an old Network Appliance snapshot, all newer snapshots are lost.

If you use snapshots, you can get a database space overflow although you do not add new data to the database, but perform updates only. This is due to the fact that the updates are not performed in the same place, but are copied to wherever there is some free space available. As a consequence, you should monitor the free space on the logical volumes which contain database spaces.

After the first full backup has been done with an external backup system, the backups that follow are done incrementally - only the changes are backed up. This leads to a much better performance of backup and restore operations.

Configuring External Backups with EMC or Network Appliance Systems

If you run a database on a Network Appliance Filer or EMC Symmetrix, it seems that the database files are part of the file system. On a Windows system, the device has to be mapped to a device letter. On a UNIX machine, the device has to be mounted. In the case of Network Appliance, the filer is divided into volumes. In the case of EMC Timefinder, the Symmetrix with the related disks is divided into device groups. Generally, it is not allowed to run more than one database on a volume or device group. The database should not be on the same volume as \$SAG. Additionally, another volume/device group has to be configured for the log location and log archive location. This is very important if Tamino is to run error-free.

As described in the following section [Steps](#), the configuration of the respective external backup method is done within the Tamino Manager. It is possible to configure more than one external backup method. Also, it is possible to use a mix of internal and external backups.

Please note that Tamino does not check if the Database Administrator has configured the external backup correctly. Forgetting a logical volume containing database spaces for the database to be saved would be fatal, since after a restore the database would be inconsistent, due to the fact that some database spaces were not restored.

Steps (EMC and Network Appliance)

This section describes the steps required for doing external backups on a Network Appliance Filer or an EMC Symmetrix System. The following topics are covered:

- [External Backup with Network Appliance Filer](#)
- [External Backup with EMC Timefinder](#)

To enable one of the methods, you first need to select the option **Add Backup Method** under **External Backup Environment** in the Tamino Manager.

External Backup with Network Appliance Filer

For an external backup with Network Appliance Filer, the database administrator has to allow access to the file system via the command `exportfs`. The Network Appliance Filer concept allows adding and removing disks without stopping any application. If there is not enough disk space, an additional disk can be plugged in without shutting down the server or waiting for transactions to finish. This is very important for 24/7 availability. All Network Appliance specific commands produce information that is logged in an error file. If an error occurs, the error file can be found in the database specific log location.

There are two backup methods with a Network Appliance Filer: Snapshots and Mirroring.

Network Appliance Snapshot

For this method, you first have to configure a volume for the database files, with a log location that points to a different location. Then, in the Tamino Manager, choose the database you want to back up. Under **External Backup Environment**, choose **Network Appliance Snapshot**. To set the filer and volume name, select the option **Set Method Parameters**. This opens a dialog that lets you specify the filer and volume name as parameters for the backup method.

After having modified the method parameters, select the object **Backups**. Then choose **Backup** from the context menu. In the resulting dialog, choose the desired backup method **Network Appliance Snapshot**.

Snapshots are available via their name, so the next step is to define it. Choose **Next** to display the page **Create Backup - Specify Backup Configuration Name** and enter a Backup Configuration Name.

Choose **Finish** to start the backup. During the backup process, the database suspends update requests.

The restore functionality is available as for internal backups (see section [Restore and Recover](#)). There are no special requirements for restoring snapshots. Note that it is possible to take more than one snapshot. If you do not restore the latest snapshot, newer backups are also not available any longer, since the snapshots are stored on the same volume. If the whole volume is restored, the snapshots are lost.

Network Appliance Mirroring

For this method, you have to configure a volume for the database, and one volume for each mirrored database. The mirror volumes do not have to be on the same filer, but must have exactly the same sizes as the database volume. The mirror volume must be initialized by the filer administrator. To do so, select the object **Network Appliance Mirror** under **External Backup Environment** and choose **Set Method Parameters** from the context menu. In the resulting dialog, enter the names of the filer and volume.

The next step is to define a configuration name. To do so, choose **Add Configuration** and enter a name in the field **Configuration Name**.

After having configured the parameters for the Network Appliance Mirror, select the object **Backups**. Then choose **Backup** from the context menu. In the dialog box that appears, choose the desired backup method **Network Appliance Mirror**. Then choose **Next** and select the backup configuration:

If you choose a mirror of an existing backup, it is overwritten. Consequently, only one backup with a mirror name is managed. Before the system performs the backup, a test is done to find out whether the destination filer for the mirror exists. During the backup process, the database suspends update requests. For the restore process, all other mirror backups are still available. If Network

Appliance snapshots are combined with mirrors in the backup sequence, all snapshots later than the restored mirror are no longer available.

External Backup with EMC Timefinder

Tamino supports database backups in conjunction with EMC's Timefinder software. The concept is to split a BCV (Business Continuance Volume) from the so called standard device(s) where the database is located for usage as a backup save set. The following configurations are supported:

- Single BCV on the same Symmetrix
- Multi BCV on the same Symmetrix
- Remote BCV on a different Symmetrix

Configuration

The first step you need to take is to configure the external backup environment. Specify the devices where the database is located.

To do so, select the object **EMC TimeFinder** under **External Backup Environment** and choose **Set Method Parameters** from the context menu. In the resulting dialog, enter the following information:

- The target (the Symmetrix ID of the database).
- The Symmetrix devices where the database is located.
- The library path of the EMC Symmetrix API.

As the database files are located on file systems, you have to determine on which physical devices the file system is mounted. A file system may spread over multiple physical disks, particularly if logical volume managers are involved. The SYMCLI call `symhostfs` may help to collect that information.

The following rules must be followed when configuring your database for external backup with EMC Timefinder:

- All index and data locations must be taken into account (including the Reserved Location, because a database may open container files at this location).
- It is not necessary to take journal and temporary files into account (which means that these files may be located at the same place as index and data). To speed up the backup/restore process, however, it may be more efficient to put the journal in a different location.
- The log location and log archive location must be at a different place, which is not part of the external backup process.

You can choose between two kinds of device specifications. Either specify the physical device name (`/dev/rdisk...`), or the Symmetrix device name (3-hexdigits value):

1. Physical:

```
Target = pd  
Devices=/def/rdsk/... ..
```

2. Symmetrix:

```
Target = <SymmID>  
Devices=XXX...
```

The Symmetrix ID (SymmId) can be abbreviated from right to left until it is unique for a host. However, it must have at least 4 digits.

To connect to EMC's software component SYMAPI, you have to define the full path name of the symapi library within the associated EMC installation.

After having defined the standard devices, choose **Add Configuration** from the context menu to configure the BCV. Note that the number and size of the BCV devices must match the ones of the database. Each BCV receives a freely selectable group name, which is to be used when the backup is done. For a Multi-BCV environment, each BCV set must be entered in the same way.

Backup

After having finished the configuration, the database can be backed up to the configured BCV(s). The configuration is checked and the devices are inserted into a device group which receives the name of the database. Thus you can view this device group with the appropriate SYMCLI functions.

There are two phases, the establish and the split phase. During the establish phase, normal database activity is possible without interference. During the split phase, which is less time consuming, the database suspends update requests. Only retrieval is possible.

The first backup after configuration is a full backup. After this first backup, all following backups are incremental. Thus the complete backup and restore process is much faster.

Restore

The restore operation restores the BCV and then plays back the redo log. As it is an incremental restore (only changed data is restored), it is a relatively fast action. Because the disks are changed below the mounted file system, it must be remounted (on UNIX) to get rid of cached data. During this restore procedure, all involved file systems will be unmounted. After the restore is completed, they will be mounted again. Thus no other processes must be active on the file systems involved (the mount and unmount processes are managed by Tamino). Note that the BCV-restore plays back the entire file system(s), not only the database files. Therefore you must be aware of the fact that this affects all other files located in such file system(s).

Logging

Some basic functions (e.g. completion of a backup or a restore) are logged with EMC's standard logging facility (which is the logfile in the log directory of the EMC installation). If an error occurs, Tamino displays an error message. If a SYMAPI function fails, a detailed error code is displayed. The errors are also logged in the standard EMC log file.

External Backup with Tivoli Storage Manager

Tivoli Storage Manager provides backup facilities for networks of computers. If you want to include Tamino databases into your Tivoli Storage Manager based backup strategy, you have the following alternatives:

- Shut down the Tamino database servers before the backup, back up the data and index spaces using Tivoli Storage Manager, then restart the database servers. In this case, it is recommended to have just one database per location, and to back up all locations assigned to index or data spaces to be sure to include all relevant files.
- Switch the database to read-only by changing the server property, then stop and restart the database server. Back up the data and index spaces using Tivoli Storage Manager, switch the read-only property off, then stop and restart the database servers. In this case, it is recommended to have just one database per location, and to back up all locations assigned to index or data spaces to be sure to include all relevant files.
- Do a Tamino backup as usual. Back up the resulting backup space with Tivoli Storage Manager.
- On UNIX, do a Tamino backup to a named pipe location, using the Tamino database as usual. Have Tivoli Storage Manager read from the named pipe.

The file names of the backup spaces can be displayed with the `argbatch` command `show backupfiles`.

4 Disaster Recovery

- Recovering from Internal Configuration Store Corruption 28
- Recovering from the Loss of the Host Environment 29

In Tamino, disaster recovery stands for the restoration of a Tamino database when everything is lost except for a backup of the database and recovery files. This may be the case after a total loss of the computing environment, in which backup and recovery files are still available on an archive. Unlike the normal restore and recover process, no information about the locations of the recovery files exists within the internal configuration store. Nevertheless, Tamino offers the possibility to recover from this disaster scenario if the following conditions are met:

- The log spaces since the last backup used to create a database must be complete (no log space must be missing).
- All log spaces must be accessible.
- The log spaces must not have been used for a **Recover until** or **Skip recover** process before.

In order to meet these conditions, you should have a good archiving concept. Please contact Software AG Support, if you require support for this task.

The following two disaster recovery scenarios are described in this chapter:

Recovering from Internal Configuration Store Corruption

If the content of your internal configuration store (registry) is damaged for any reason (e.g. someone has accidentally deleted the information, or a disk error has occurred), you will not be able to access the Tamino databases. However, the file *Registry.log* keeps track of all internal configuration store entries that are relevant to Tamino and have occurred since the last installation of Tamino.

On Windows

▶ To recover from registry corruption

- 1 Use the import file provided by the installation of the latest Tamino version

```
regutil delete "HKLM\Software AG\Tamino"  
  
regutil import "<DRIVE_LETTER>:\Documents and Settings\All Users\Application  
Data\Software AG\Tamino\Registry.export"
```

- 2 Apply the internal configuration store entries from the saved log with

```
regutil -f "<DRIVE_LETTER>:\Documents and Settings\All Users\Application  
Data\Software AG\Tamino \Registry.log"
```

On UNIX

▶ Case 1: A backup of `$SAG/common/rgs/REGFILE` is available

- 1 Restore the backup of the internal configuration store
`$SAG/common/rgs/REGFILE`
- 2 Apply the internal configuration store entries from the saved log with
`regutil -f "$SAG/ino/Registry.log"`

▶ Case 2: A backup of `$SAG/common/rgs/REGFILE` is not available

- 1 Save the file "`$SAG/ino/Registry.log`"
- 2 Reinstall Tamino, including the internal products.
- 3 Apply the internal configuration store entries from the saved log with
`regutil -f <saved Registry.log>`



Note: Note that the backup file of the *Registry.log* must be a backup of the current Tamino installation, and not of a former version of Tamino.

Recovering from the Loss of the Host Environment

If you have lost your total computing environment due to a disaster, it is still possible to recover if a backup and all subsequent recovery files are available. However, the recovery files must not have been used for a restore with no full recover (recover with Skip Recover or Recover Until) process. The steps for disaster recovery in this case are:

▶ To recover from the loss of the host environment

- 1 In a new computing environment, install Tamino.
- 2 Perform a **Create from Backup** with the given backup file.



Note: The backup file you specify must be a full backup file. You cannot use an incremental backup file to perform this step.

If no recovery files are to be applied or available, this is all that needs to be done. This may be the case if you are doing frequent backups.

If recovery files are available, apply them by following the next steps:

3 Add the internal configuration store entries for the recovery files.

To do so, create a script using the `regutil` utility to add the entries for the recovery files. An entry for a log space (with the extension `.1LO`) has the following structure:

```
HKEY_LOCAL_MACHINE
\SOFTWARE
  \Software AG
    \Tamino
      \servers
        \<database name>
          \savepoints
            \000000000000
              \log space
                \<session number>
                  \<extent number>
```

The session number must be extended by leading zeroes as a 10 digit number, the extent number as a 5 digit number.

Here is an example: Two log spaces, `AAF000010000000025.1LO` and `AAF000010000000026.1LO` on the directory `D:\archie`, are to be processed.

```
create "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space"
create "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00025"
create "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00025\00001"
setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00025\00001" type "file"
setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00025\00001" path "D:\archie\AAF000010000000025.1LO"
create "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00026"
create "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00026\00001"
setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00026\00001" type "file"
setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space\00026\00001"
path "D:\archie\AAF000010000000026.1LO"
```

Other recovery files than log spaces (extension .1L0) only exists when the Tamino Data Loader was used without the `concurrentWrite` option. These are files with an extension of .1C0, .1C1 or .1C2. In case you want to do a disaster recovery with Tamino Data Loader recovery files, please contact Software AG Support to ask how to add registry entries.

- 4 Use the Restore/Recover function.



Note: Note that the server must not be started between **Create from Backup** and performing the recover function.

Index

B

- backup
 - complete, 2
 - create, 13
 - definition, 2
 - external, 2, 9, 19
 - configuration, 21
 - EMC Timefinder, 23
 - mirroring, 3, 9, 20
 - snapshot, 3, 9, 20
 - full, 2
 - incremental, 2
 - internal, 2, 9, 11
 - non-parallel, 3
 - offline, 2
 - online, 2
 - parallel, 3
 - partial, 2
 - requirements, 4
 - time, 10
- backup generations
 - number of, 13
- backup strategy, 1
- BCV, 23
 - multi, 23
 - remote, 23
 - single, 23
- Business Continence Volume (see see BCV)

C

- change
 - number of backup generations, 13
- complete backup, 12
- corrupted data, 7

D

- data corruption, 7
- data failure
 - recovery, 5
- database
 - synchronization, 4
- disaster
 - recovery, 5, 7
- disaster recovery, 27
- disk failure, 14

- recovery, 5
- disk mirroring, 5

F

- failure
 - recovery, 8

H

- hardware error
 - recovery, 5
 - solutions, 6
- host environment
 - loss of, 29

I

- incremental backup, 2, 20, 24
- internal configuration store
 - disaster recovery, 28

L

- log space, 15

M

- mirroring
 - with Network Appliance, 22
- multi BCV, 23

N

- Network Appliance Filer, 21

R

- RAID 5, 5
- recover, 14
 - additional failure, 8
 - definition, 3
 - disaster, 7
 - error situations, 15
 - from data failure, 5
 - from disaster, 27
 - full, 3
 - hardware error, 5
 - partial, 3

- skip, 16
- software error, 7
- until point in time, 17

registry corruption, 28

remote BCV, 23

replication, 4, 6

restore, 14

- definition, 2
- error situations, 15
- full, 2
- partial, 2

S

single BCV, 23

snapshot

- with Network Appliance, 22

software error

- recovery, 7

synchronization

- of a database, 4

T

time

- for backup, 10
- for recover, 10
- for restore, 10