

CrossCheck Travel 3.0

Database Backup and Recovery Procedures
for Agency Owned Equipment





Introduction

The development of a CrossCheck Travel 3.0 Database backup and recovery strategy is an exercise carried out in phases. In order to work through these phases we need to understand some database terminology such as, backup and recovery,

In this document we outline the phases in developing a backup and recovery plan and cite an example. We show how time becomes an important factor in database recovery. We also explain what to do when the backed up database fails to restore from your backup medium, and discuss what database health checks should be performed to ensure the database and log files are valid.

An example of database backup and recovery commands is provided to illustrate the statements used to backup and recovery a database. Finally, we discuss key points in developing a disaster recovery strategy in the event the physical machine running your database is no longer available.

Backup, Recovery, and Disaster Recovery

We tend to think of 'backup and recovery' as one topic and 'disaster recovery' as another.

'DBBackup' is a utility program used to make a copy of the contents of database files and log files. The database files consist of a database root file and a log file.

'Recovery' is a sequence of tasks performed to restore a database to some point-in-time. Recovery is performed when either a hardware or media failure occurs. Hardware failure is a physical component failure in your machine, such as, a disk drive, controller card, or power supply. Media failure is the result of unexpected database error when processing data.

Before you begin recovery, it is a good practice to back up the failing database. Backing up the failing database preserves the situation, provides a safe location so files are not accidentally overridden, and if unexpected errors occur during the recovery process Galileo Southern Cross Support may request these files be forwarded to them.

'Disaster recovery' differs from a database recovery scenario because the operating system and all related software must be recovered before any database recovery can begin.





Files that make up a CrossCheck Travel 3.0 Database

Adaptive Server Anywhere CrossCheck Travel 3.0 databases consist of disk files that store data. The default database file name is **cct.db**. This main database file contains database tables, system tables, and indexes.

The transaction log, **cct.log**, is a file that records database modifications. Database modifications consist of inserts, updates, deletes, commits, rollbacks, and database schema changes. A transaction log is not required but is highly recommended. The database engine uses a transaction log to apply any changes made between the most recent checkpoint and the system failure. The checkpoint ensures that all committed transactions are written to disk. During recovery the database engine must find the log file at its default location. When the transaction log file is not specifically identified then the database engine presumes that the log file is in the same directory as the database file.

Online CCT 3.0 Database Backup

CrossCheck Travel 3.0 database backups can be performed while the database is being actively accessed (**online**).

An online database backup is performed by executing DBBACKUP from the command-line, from the Sybase Central 'Backup Database' utility or, in our case, by a batch file scheduled to run at a certain time referenced later in this document.

When an online backup process begins the database engine externalizes all cached data pages kept in memory to the database file(s) on disk. This process is called a checkpoint. The database engine continues recording activity in the transaction log file while the database is being backed up. The log file is backed up after the backup utility finishes backing up the database. The log file contains all of the transactions recorded since the last database backup. For this reason the log file from an online full backup must be 'applied' to the database during recovery.





Developing a CrossCheck Travel 3 Database backup and recovery strategy

The steps suggested in the development of a CCT 3.0 db backup and recovery strategy consist of the following:

- Understand what backup and recovery means to your business
- Management commits time and resources for the project
- Develop, test, time, document, health check, deploy, and monitor
- Beware of any external factors that affect recovery
- Address secondary backup issues.

Understand what backup and recovery means to your business

How long can my business survive without access to the corporate data?
Express your answer in terms of minutes, hours, or days.

If your recovery time is in minutes then CCT 3.0 database backup and recovery is critical to your business needs and it is paramount that you implement some kind of backup and recovery strategy. If recovery can take hours then you have more time to perform the tasks. If recovery can be expressed in terms of days then the urgency to recover the database still exists, but time appears to be less of a factor.

Management commits time and resources for the project

Management must decide to commit financial resources towards the development and implementation of a CCT 3.0 database backup and recovery strategy. The strategy can be basic or quite extensive depending upon the business needs of the company. After developing a backup and recovery strategy management should be informed of the expected backup and recovery times. Anticipate management countering the timings by preparing alternative solutions. These alternative solutions could be requesting additional hardware, improved backup medium, altering backup schedule, accepting a longer recovery time versus backup time. Then it will be up to management to decide what solution fits their corporate needs.





Develop, test, time, document, health check, deploy, and monitor

These phases are the core in developing a backup and recovery strategy:

Build the CCT 3.0 backup batch file CCT_BACK.CMD (script details are referenced later in this document). Verify this batch file works as designed. Does your full online backup work? Verify that the CCT_BACK.CMD produces the desired results.

Time estimates from executing backup and recover the CCT 3.0 database help to get a feel for how long will these tasks take. Use this information to identify what batch file will be executed and when.

Document the CCT_BACK.CMD backup batch file and create written procedures outlining where your backups are kept and identify the naming convention used as well as the kind of backups performed. This information can be very important when an individual must check the backups or perform a database recovery and the data base administrator (DBA) is not available.

Incorporate health checks into the CCT 3.0 backup procedures. You should check the database to ensure the database is not corrupted. You can perform a database health check prior to backing up a database or on a copy of the database from your backup.

Deployment of your backup and recovery consists of setting up your backup procedures on the CCT 3.0 production server. Verify the necessary hardware is in place, including a reliable UPS, and any other supporting software necessary to perform these tasks.

Monitor backup procedures to avoid unexpected errors. Make sure any changes in the process are reflected in the documentation.

Beware of external factors that affect recovery

External factors that effect database recovery are time, hardware, and software. Allow additional recovery time for entering miscellaneous tasks that must be performed. These tasks could be as simple as retrieving and loading tapes. Factors that influence time are the size of database files, recovery medium, disk space, and unexpected errors. The more files you add into the recovery scenario increase the places where recovery can fail. As the backup and recovery strategy develops it may be necessary to check the performance of the equipment and software ensuring it meets your expectations.





Protect CrossCheck Travel 3 database backups by performing health checks

Database health checks are run against the database and log files to ensure they are not corrupt. The database validity utility called DBVALID is used to scan every record in every table and looks up each record in each index on the table. If the database file is corrupt, you need to recovery from your previous database backup. A database can be validated before being backed up or against a copy of the database from your backup.

To run a CCT 3.0 database validation before backing it up, simply access the “Database Tools” under the Galileo Southern Cross\CrossCheck Travel program group on the production Server and select “Validate Database”. This will automatically run the DBVALID command against the CCT 3.0 production database.

DBVALID can take time to check the database. You will need to decide when to run DBVALID. As a general rule, a database validity check done before running a full database backup would ensure the database is free of database pointer errors before being backed up.

When you run DBVALID against a restored copy from a back up medium you will be accomplishing the following checks:

You will verify the database can be successfully restored from the backup medium


You will be able to run you database validity check against the database without impacting your production environment

Run DB VALIDATION after hours or on a non-production machine to avoid impacting your production environment.

IMPORTANT:

The database used in the validation process cannot participate in recovery since log offsets will be changed once the database is started.





CrossCheck Travel 3 Backup and recovery strategy example

Let's develop a CCT 3.0 backup and recovery strategy for a fictitious Travel Agent called 'Galactic Travel'.

Galactic Travel requires access to corporate data stored in the database during business hours. Business hours are defined as a time period from Monday through Friday between the hours of 8AM to 8PM. During business hours Galactic Travel cannot be without this information for more than 24 hours. Access to the corporate data outside of business hours is not critical.

The database is not being backed up. The size of the **cct.db** database is 600 MB consisting of a main database file containing several tables and indexes. The database is using a transaction log. Daily business activities will cause the log files to grow at a daily rate of 80MB. We will now be able to develop a base line for the time it will take to backup and recover the database. The timings are fictitious since they will vary depending upon the machine, size of database, and additional software running on the machine.

The name of the database backup utility is DBBACKUP referenced in the CCT_BACK.COMD batch file provided here below. It can be used to backup database and log file.

A full database backup using CCT_BACK.COMD will make a copy of the main database file and transaction log file.

Time the basic CCT 3.0 Database backup and recovery operations

After running several full database backup tests, we calculated the average backup time to be 30 minutes.

Galactic Travel has a third party software tool to backup disk files tape. All database backups will be copied onto tape. We calculate the average database and log backup and restore from tape to be 15 minutes and 5 minutes respectively for a full database backup. Now we can use these timings to develop our recovery strategy.



Estimate a total recovery time from failure

The times estimated in the previous section give us a way to estimate what our recovery time will be given the following conditions.

If we need to recover the database from a backup it will take 15 minutes to restore the database from tape to disk

We presume the current database transaction log file is not corrupt and can be used with the database. Recovery process of the log file takes approximately 5 minutes.





An estimated recovery time is calculated to be a total of 20 minutes, as shown below.

Recovery from backup medium phase

Operation	Time
Full database backup	15 minutes
Database log file backup	5 minutes
Subtotal time for backup medium	20 minutes

This recovery example indicates that it takes approximately 20 minutes to recover the database from tape. A full database backup from the night before will recover the database to the end of the last business day and meet our business objective of database recovery within 24 hours.

We suggest that the database backups be retained on tape for 5 weeks.

We now have a backup and recovery strategy to present to our management.

What do you do when you encounter errors during recovery?

During the restore process, the full database backup fails to restore from tape. This backup copy of the database is now invalid. Identify the next available full database backup needed to recover the database. Use the previous day's full database backup to begin the recovery process. Recovery time would add an additional 20 minutes to restore the previous day's db files. Then you will have to apply the last db backup log file to the previous full db backup to recover the database to the last 24 hours. Failing that, you will lose 48 hours worth of data as you will have to go back at least another 24 hours to recover the next database backup.





Recovery from previous day's backup

Recovery from medium phase: first attempt fails to successfully restore from tape a full database backup

Operation	Time
Full database backup	15 minutes
Database log backup	5 minutes
Subtotal time for backup medium	20 minutes

Second attempt successfully restoring from tape the previous day's full database backup

Operation	Time
Full database backup	15 minutes
Database log backup	5 minutes
Subtotal time for backup medium	20 minutes

Database recovery phase (applying database log to recovered database)

Operation	Time
Apply log to the recovered database	30 minutes
Issuing commands	30 minutes

Total recovery time: 1 hour and 40 minutes

This second scenario illustrates the need to be ready for the unexpected. Every database recovery will be different from the previous one. If one of the database files from a full database backup is bad verify the log files are valid. If so, then at least you can recover from a previous day's full database backup and recover past this point of failure by applying the log.

Additional Information regarding recovery

Before doing any kind of database recovery you should backup your failing database. If you feel uncomfortable or unsure of what to do, then it's time to contact Galileo Southern Cross Support. Even if you have come up with a way to recover the database, it is worth getting a second opinion and to have your recovery plans verified by Galileo Technical Support.





Secondary backup issues

Define a backup retention timeframe for your backups.

If your backup medium is tape then decide the number of times they will be used before discarding.

Test your backup procedures by restoring from them to ensure they can be restored.

Make up different recovery scenarios and try to recover the database from them.

Verify your recovery procedures are kept current.

Backup and recovery commands

The CrossCheck Travel 3 database server used in this section is Adaptive Server Anywhere 8.0.3. The executables for standalone and network servers are **dbeng8/dbsrv8** respectively. The path in the following examples is not indicated since you will either use the installation default or change the path accordingly.

Test your commands

These commands are to be copied as is into a backup script. Always, test, test, and test your scripts before deploying them.

Directories used

C:\Galileo\Database\backup

Creating backups

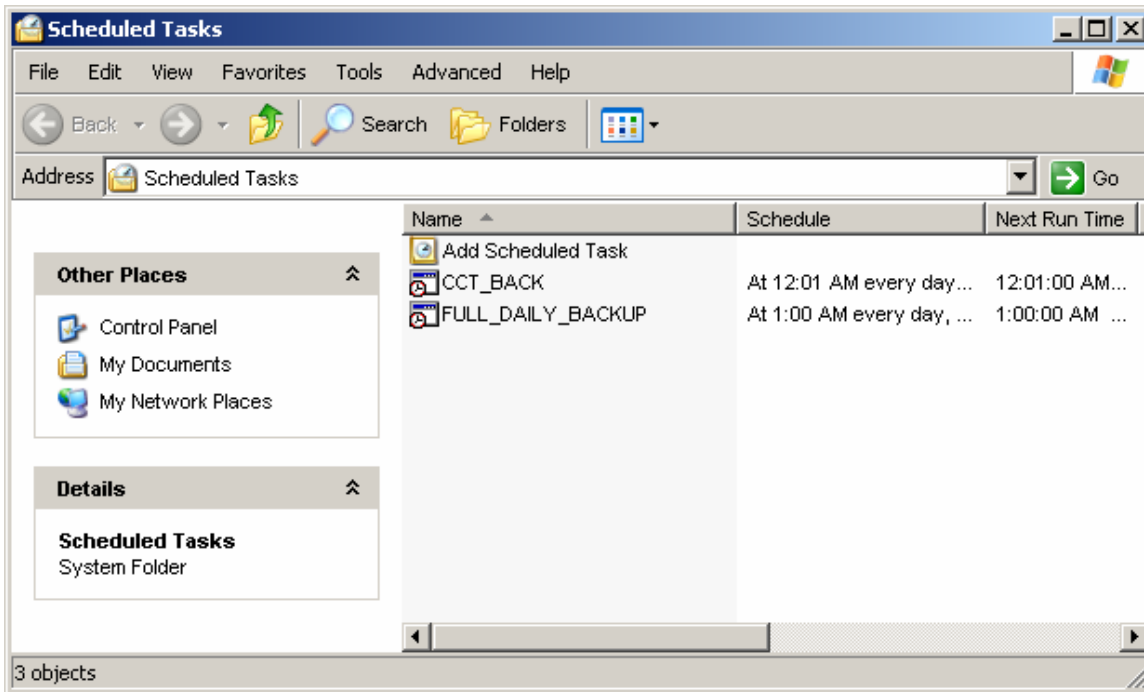
Create a full online backup of the database and log file by running the below commands in a scheduled batch file named **CCT_BACK.CMD**:

```
@echo off
CD \Progra~1\GalileoSouthernCross\CrossCheckTravel\Server\
dbbackup -y -t -c "UID=dev;PWD=dev;ENG=Galileo" C:\Galileo\Database\backup
exit
```



Scheduling CCT_BACK.CMD to run automatically

To make sure that the CCT_BACK.CMD batch file is scheduled to run automatically, access the “Scheduled Tasks” applet: Start\All Programs\Accessories\System Tools. Follow the wizard to set it up and make sure the scripts are scheduled to run as per example below:



This system will allow first for the CCT 3.0 database files to be backed up on HDD on C:\Galileo\Database\backup. Then the FULL_DAILY_BACKUP will backup the CCT 3.0 database files from that location together with system state and other relevant files to tape according to your site full backup strategy. Please, test thoroughly both backup scripts and the scheduler to make sure there are no bad surprises in case a database recovery is necessary.





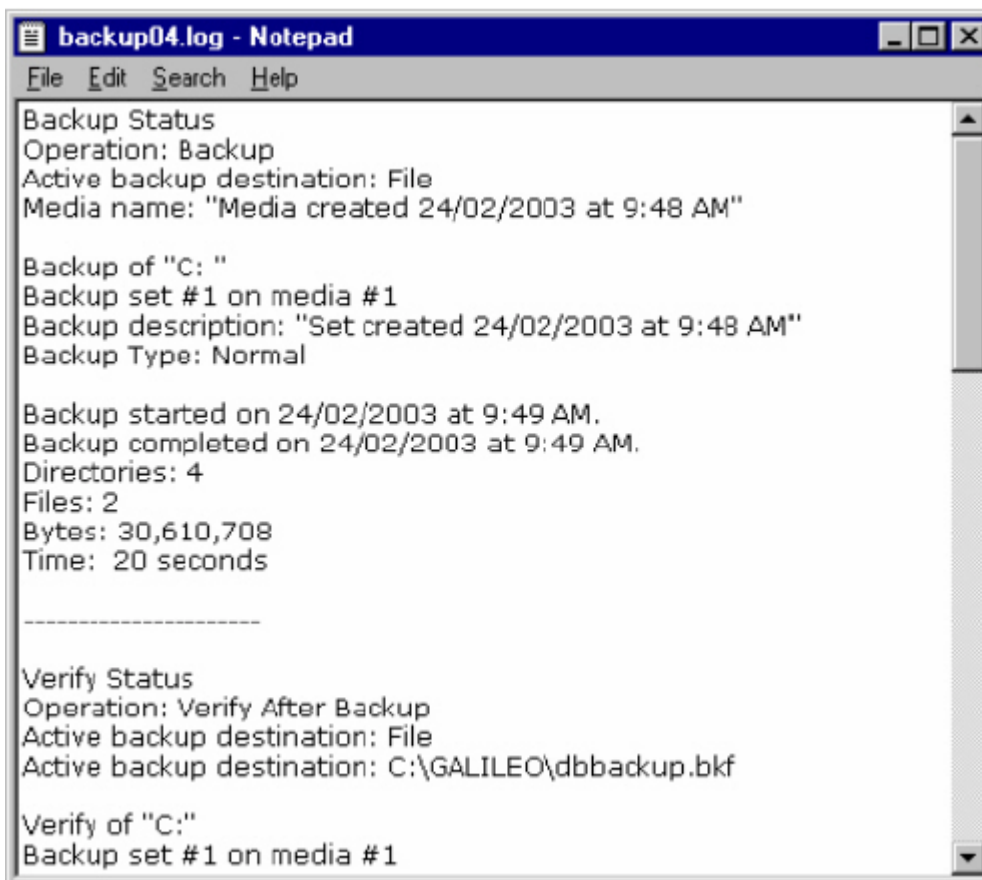
Full System Backup Verification

The following procedure should be performed first thing in the morning, prior to any transaction being entered into CrossCheck Travel 3.

It is assumed that the full system backup (NTBACKUP if using the Windows Tool) is run on the production CCT 3.0 Database Server.

Access the backup.log through the NTBACKUP applet. Go to Tools\Report to access the Full System Backup log file.

A Notepad test editor will open:



```
backup04.log - Notepad
File Edit Search Help
Backup Status
Operation: Backup
Active backup destination: File
Media name: "Media created 24/02/2003 at 9:48 AM"

Backup of "C: "
Backup set #1 on media #1
Backup description: "Set created 24/02/2003 at 9:48 AM"
Backup Type: Normal

Backup started on 24/02/2003 at 9:49 AM.
Backup completed on 24/02/2003 at 9:49 AM.
Directories: 4
Files: 2
Bytes: 30,610,708
Time: 20 seconds

-----

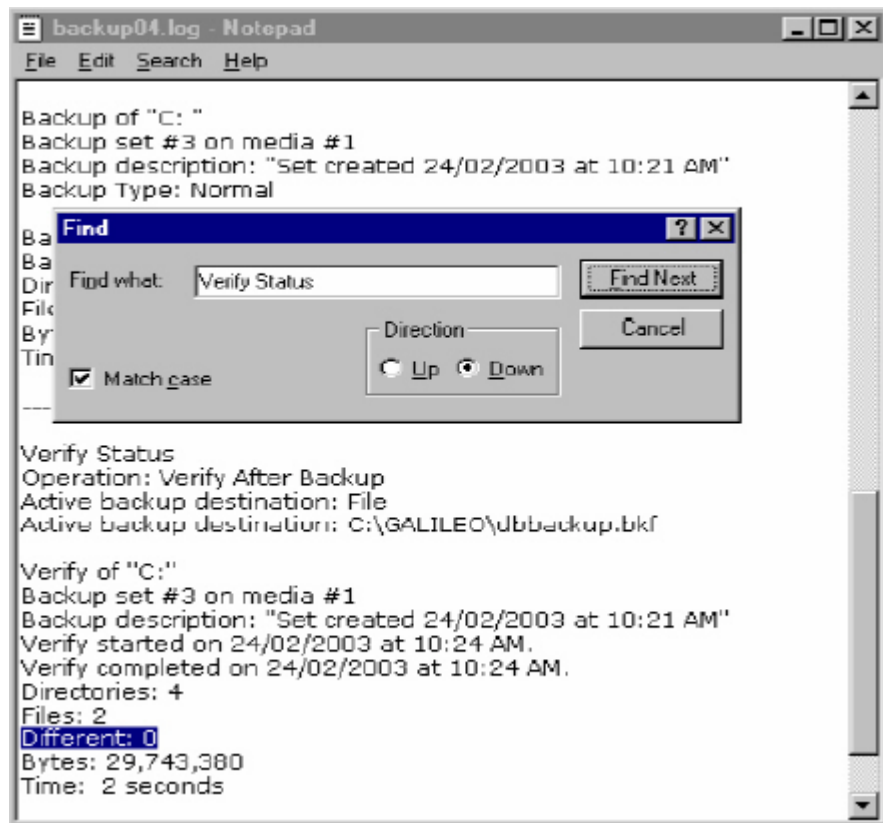
Verify Status
Operation: Verify After Backup
Active backup destination: File
Active backup destination: C:\GALILEO\dbbackup.bkf

Verify of "C:"
Backup set #1 on media #1
```

Each time NTBACKUP (or other third party backup software) is executed, the details of the backup should be appended to this log file.



Place the cursor at the top of the page (Ctrl + home); click on "Search" and then "Find". Type: "Verify Status" in the text box. Select: "Find Next".



Search until you find the backup description of the set created at today's date.
For example: "Set Created 18/01/06 at 01:00AM".

Make sure "data verification was selected, completed and no differences were detected. See lower section of the above graphic.

Once satisfied that the Full Backup was successful, select file/exit to quit notepad.

Recovering backups

Retrieve the CCT 3.0 database files (cct.db and cct.log) from tape.
Copy both files in C:\Program Files\GalileoSouthernCross\CrossCheckTravel\Database directory





Perform health checks

Now start the database using the personal server (dbeng6) or stand alone database engine (dbeng8.exe). We do this is to ensure the users will not be able to connect to the database until we are certain it is valid.

The database is mounted to a personal server by accessing the command line and issuing the following command:

```
DBENG8-n Galileo C:\Progra~1\GalileoSouthernCross\CrossCheckTravel\Database\cct.db
```

Once connected, verify that the database is not corrupt by running DBVALID.

Go to Start\All Programs\GalileoSouthernCross\Database Tools\Validate Database

This tool will automatically run a validation on your recovered database.

If “no errors” are returned, the validation was successful. If any error is returned, please call Galileo Southern Cross Support for further troubleshooting.

Create a post-recovery backup

Clear out old copies of the database in the backup directory. You may want to backup the old files to some removable backup medium such as a CD or a DVD.

Create a post-recovery backup of the database and log file

Go to Start\All Programs\GalileoSouthernCross\Database Tools\Manual DBBackup


This tool will perform a one off database files backup in the default location: C:\Galileo\Database\backup






Allow user access the database

At this time we now want to shutdown the database because we used the stand alone or personal server database engine to perform health checks, verify the database was valid, and post-recovery of the database.

Stop the temporary engine by locating the  icon on the system tray, next to the system clock, right-clicking it and selecting "Stop".

Start the database engine in the normal manner prior to the recovery. Locate and select the CCT 3.0 DBMS (Sybase Central) icon on the CCT 3.0 production server desktop. Once able to access Sybase Central, click on "Services" on the explorer like window to display the "Galileo" Database Service; right-click on the Galileo service and select "Start".

Please, note the difference between the standalone engine icon here above and the network server icon .

The network server will allow users to connect to the database through the agency network from the CCT 3.0 client application installed on their PC.

Monitor the database server engine and user activity against the database for unexpected error messages. Call the Galileo Service Centre if any error is reported.





Disaster recovery

Disaster recovery is different from a database recovery because it usually implies the production server machine is no longer available. This may be the result of a flood, natural disaster, or an inoperable machine. This type of situation requires the operating system, system software, database software, and application software recovery onto a physically different machine. The machines may be similar or identical. If the machine is different this could impact your disaster recovery. After system and database software is installed then CrossCheck Travel 3 database recovery procedures can be started.

Additional information regarding disaster recovery planning

The following list is an example of what information is needed in developing a disaster recovery strategy. You should review your strategy and update it accordingly if you have omitted any of these points of interest. This list is only representative of may be required and can vary based upon each site's configuration of hardware, software, location, and personnel.

Establish an offsite location with compatible system and network hardware.

Establish procedures for retrieving offsite backups.

Disaster recovery procedures should be kept in multiple places. At the offsite location with the backups for the operating system, system software, database, and application software as well as with individuals involved with disaster recovery.

Schedule regular disaster recovery tests and after each test review the outcome.

Verify the disaster recovery procedures contain copies of all installation software and their passwords for the operating system, system software, database, and application software.

Record the total time to complete your disaster recovery.

Keep the disaster recovery procedures current and update the copies kept offsite.





Summary

The material presented here provides detailed information in your development and implementation of a database backup and recovery strategy. We have identified the files used in a database and what is an online backup.

We stressed the importance of incorporating health checks into the backup and recovery strategy to safeguard your investment. A backup and recovery example is given showing factors involved in the development of a backup and recovery strategy. We discuss what is involved with disaster recovery. Disaster recovery should not be overlooked. An accident to your equipment can happen causing your server and database to become inoperable. What we also want to stress is to test, test, and test your backup and recovery procedures. Assume nothing and verify everything works as expected. We hope this discussion provides you with the necessary information to assist in your development and implementation of a CrossCheck Travel 3 Database backup and recovery strategy.

