

# High-Speed Backups without Stopping Business Applications

● Isao Igarashi

*(Manuscript received September 12, 2005)*

The growing demand for 24 × 7 system operations has increased the demands for technologies that back up business volumes online without needing to stop the business applications. Fujitsu provides high-speed backup and recovery solutions for SQL Server, Exchange, Oracle, and DB2 databases. These solutions simplify and reduce the cost of Oracle database backups and recoveries. This paper introduces several techniques for making backups without needing to stop the business applications. It also describes ways to simplify and reduce the cost of system operations.

## 1. Introduction

The growing need to provide 24 × 7 services has increased the demand for technology that can back up business volumes online without stopping the business applications. Fujitsu's ETERNUS disk array series meet this demand with two Advanced Copy features: One Point Copy (OPC)<sup>note 1)</sup> and Equivalent Copy (EC).<sup>note 2)</sup> By using Fujitsu's Advanced Copy Manager (ACM), which performs high-speed backup using an Advanced Copy feature, it has become possible to build reliable, high-speed backup systems.

---

note 1) OPC copies continuous blocks of data on a physical device to a different location. When the hardware receives an instruction from the software via OPC, it creates a state of logical completion of copy and informs the software about the same. The actual copy takes place in the background, and once the copy is complete, the copy source and copy target can be accessed independently.

note 2) EC is a hardware feature that copies continuous blocks of data on a physical device to a separate location and preserves redundancy by reflecting data updates on the copy source to the copy target volume. Starting EC and maintaining redundancy is done via software instructions.

ACM also enables high-speed backups for SQL Server, Exchange, Oracle, and DB2 databases without having to stop the business applications. Moreover, Fujitsu's Recovery Manager for Oracle (RMFO) simplifies and reduces the cost of Oracle database backup and recovery operations.

This paper introduces some techniques for making backups without needing to stop the business applications. It also describes ways to simplify and reduce the cost of system operations.

## 2. Techniques for backing up business volumes without stopping business applications

Business volumes normally contain databases or general files. This section describes some techniques for backing up this type of data without needing to stop the business applications.

When databases are backed up without stopping business applications, the back-up data may become unsynchronized. To avoid this problem, Data Base Management Systems (DBMSs) record changes made to a database in a log file, which can be applied to a restored database to correctly

recover it. However, to use the logs correctly, it is necessary to execute Checkpoint<sup>note 3)</sup> on the log as a pre-process.<sup>note 4)</sup> This pre-process Checkpoint is managed by the database, and the backup copy of the database is processed by ACM. As a result, the database and ACM must cooperate with each other.

The following is a brief description of how ACM cooperates with SQL Server, Exchange, Oracle, and DB2 databases.

1) SQL Server databases: Virtual Device Interface (VDI) is used to communicate with SQL Server. The process flow is shown in **Figure 1**.

- When ACM receives a database backup request for SQL Server, it issues a backup request to SQL Server (① in Figure 1).
- After SQL Server completes the pre-processing, it uses VDI to issue a snapshot request to ACM (②).
- ACM uses the Advanced Copy feature to create a copy of the database (③, ④, ⑤).
- After SQL Server completes the post-processing, it returns control to ACM (⑥).

2) Exchange databases: Volume Shadow Copy Service (VSS) is used to communicate with Exchange. The process flow is shown in **Figure 2**.

- When ACM receives a database backup request for Exchange, it issues a backup request to VSS (① in Figure 2).
- VSS invokes Exchange, which then executes the pre-processing tasks (②).
- After Exchange completes the pre-processing tasks, VSS invokes an ACM provider for ETERNUS (③).
- The ACM provider uses the Advanced Copy feature to create a snapshot (⑤, ⑥).

- After the snapshot has been created, VSS invokes Exchange, which performs the post-processing tasks (⑦, ⑧).
- VSS returns control to ACM (⑨).

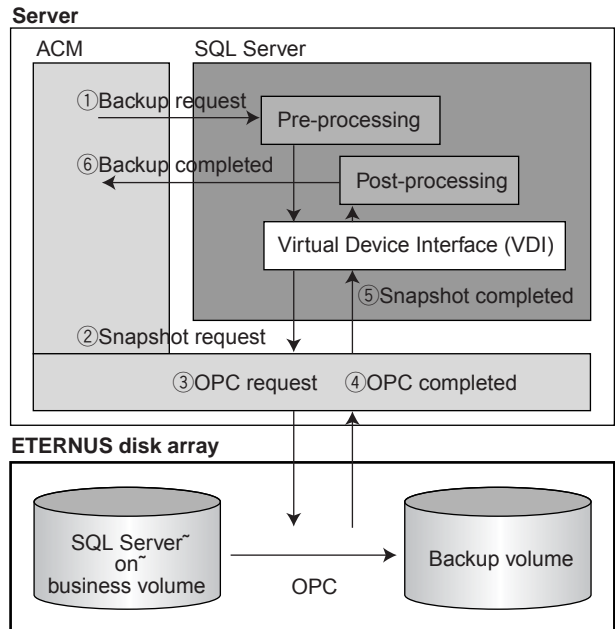


Figure 1 Backup process of SQL Server via Virtual Device Interface.

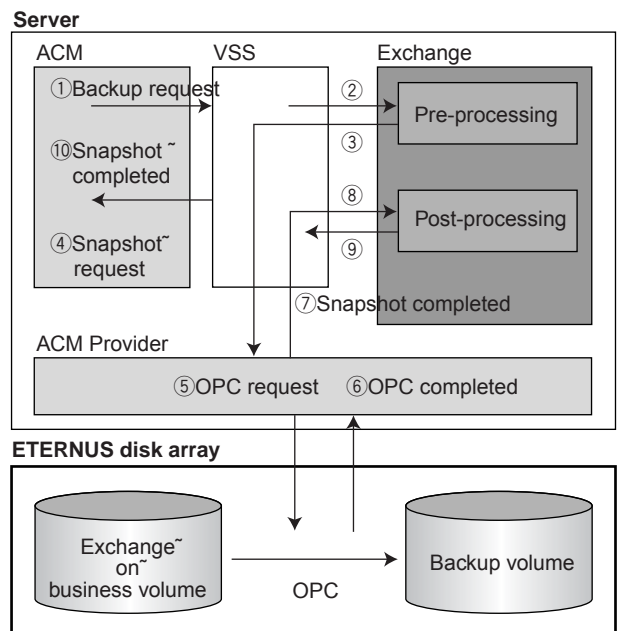


Figure 2 Backup process of Exchange via VSS.

note 3) This helps to distinguish which information is needed for recovery.

note 4) This pre-processing depends on the DBMS being used. Some DBMSs allow copying of data once data updates from the business applications are stopped.

ACM manages the backed up data and log files of the SQL Server and Exchange databases, and these can be restored as follows.

- Point-in-time Restore  
Restore to the point when the backup was performed.
  - Roll forward Restore  
After restoring the backed up data, apply the log file contained on the business volume and restore the database to the latest state.
- 3) Oracle databases: Using Oracle's begin backup command, the database is placed in backup mode (①, ② in **Figure 3**) and the copy process is started (③, ④, ⑤). When the copy process is completed (⑥), the end backup command is used to release the database from backup mode (⑦, ⑧, ⑨).  
The details of Oracle database backup and recovery are covered in a later section.
- 4) DB2 databases: As for an Oracle database, the database is placed in backup mode and the copy process is started. When the copy process is completed, the database is released from backup mode.

Next, we explain a technique for backing up general file systems without stopping the business applications. Unlike databases, file systems do not generally have a log file mechanism. Therefore, to ensure that data synchronization is achieved through correct backup, it becomes necessary to stop any I/O from the business applications and to ensure the data in the OS memory cache is written to disk. Therefore, the backup command of ACM executes the unmounting, copy, and mounting processes automatically.<sup>note 5)</sup> However, there are cases when, based on the customer's requirements, it is not possible to unmount the file system. In such cases, instead of unmounting the file system, the I/O operations on this file system are stopped and the data from the cache is forcibly flushed. Once the copy process is completed, these I/O operations are resumed.<sup>note 6)</sup>

The file system can be unmounted by executing the unmount command as part of a backup command in a pre-copy script. The file system can be mounted by executing the mount command in a post-copy script. These scripts can be completely customized to meet the system operation requirements.

The process flow for backing up a file system without stopping the business applications is shown in **Figure 4**.

In the remaining sections, we look at some of the problems associated with Oracle database backup and recovery and how Fujitsu overcomes them.

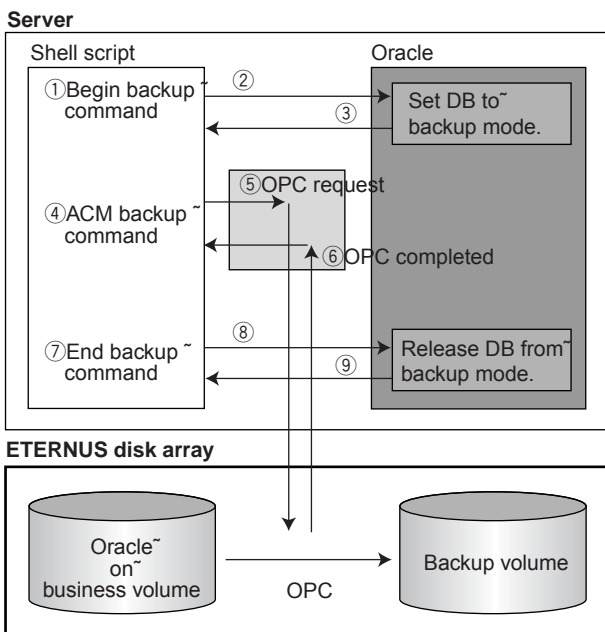


Figure 3  
Backup process of Oracle DB.

note 5) For UNIX only. For Windows, a process of locking the file system, copying, and unlocking the file system is followed.

note 6) Some operating systems cannot stop I/O. Moreover, when I/O is being stopped, some I/Os may still be in progress. This must be carefully considered on the operation side because copying should not be done during an I/O. Therefore, there must be some way to guarantee data synchronization. Generally, taking file system backups without stopping the business applications requires features such as Solaris' fssnap or Windows' VSS. This problem will be tackled in the near future.

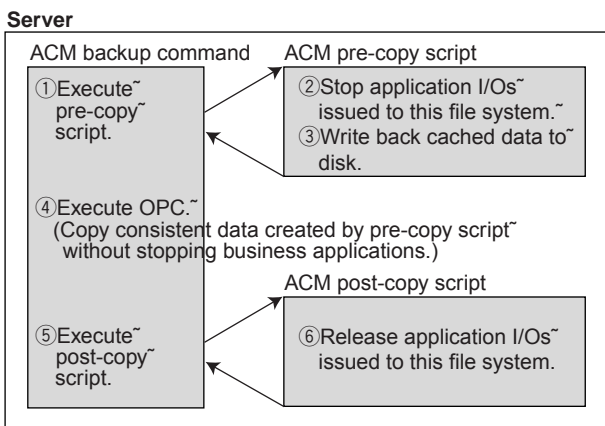


Figure 4  
Backing up a file system without stopping business applications.

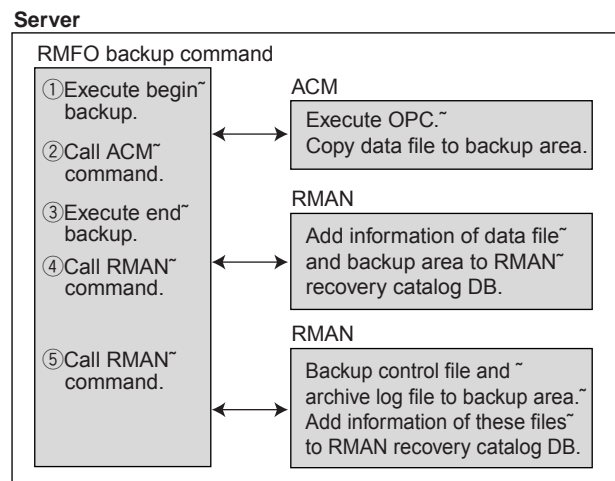


Figure 5  
Summary of RMFO backup.

### 3. Problems associated with database backup and recovery

In this section, using Oracle as an example, we describe some of the problems associated with database backup and recovery.

#### 1) Complexity of online backup procedure

The data files (tablespace area), control files, and archive log files (update log) must be backed up. However, a simple copy does not ensure synchronization, and it becomes necessary to use the following complex procedure.

- Do the following for each tablespace area in order to backup the data files (tablespace areas):
  - ① Issue the online backup start command.
  - ② Copy the data files.
  - ③ Issue the online backup end command.
- Back up the control files using an SQL command.
- Back up the archive log files (update log) as follows:
  - ① Forcibly output the archive log file.
  - ② Monitor the completion of file creation.
  - ③ Copy the archive log file.

If there is a mistake in any of the above steps, the backed up files become useless and recovery becomes impossible. Moreover, which file was backed up when and where needs to

be suitably administered, which leads to increased administration costs.

#### 2) Complexity of recovery procedure

When a database fails, the database administrator recovers it using the following complex procedure.

- Locate the database failure.
- Formulate a detailed recovery procedure.
- Select the files to be restored.
- Restore the logs required for recovery.
- Use a complex SQL command to recover the database.

The above procedure contains many operations, and if a mistake is made in any one of them, the downtime will be further increased. Moreover, the time spent by the database administrator to make a recovery increases the operation costs.

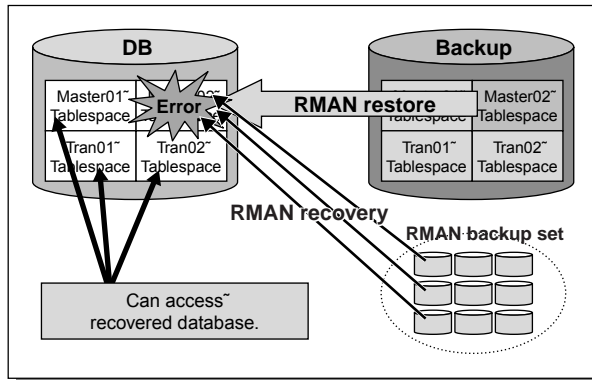
### 4. Ways to simplify and reduce costs of backup and recovery

The problems described in the previous section can be resolved by combining RMFO with ACM and Oracle Recovery Manager (RMAN).

#### 1) Simplifying and automating the administration of online backups

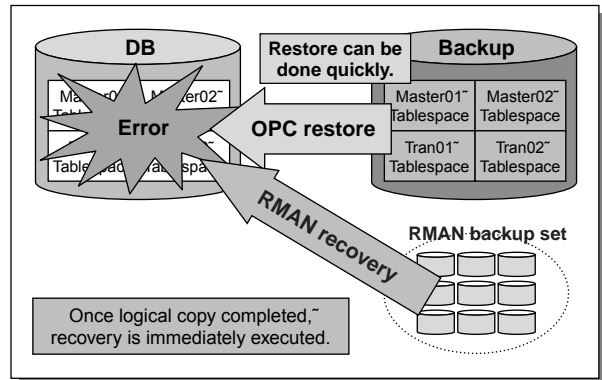
Backup scripts can be created automatically for backing up data files (tablespace area), control files, and archive log files (update log).

**Partial online recovery**



(a) If damage extent is small, partial recovery by RMAN restore is performed without stopping DB operation.

**OPC fast restore and recovery**



(b) If damage extent is big, OPC fast restore is quickly performed regardless of data size.

Figure 6

RMFO recovery operation. Fujitsu's RMFO can perform partial online recovery or entire restore and recovery based on extent of database damage.

**Server**

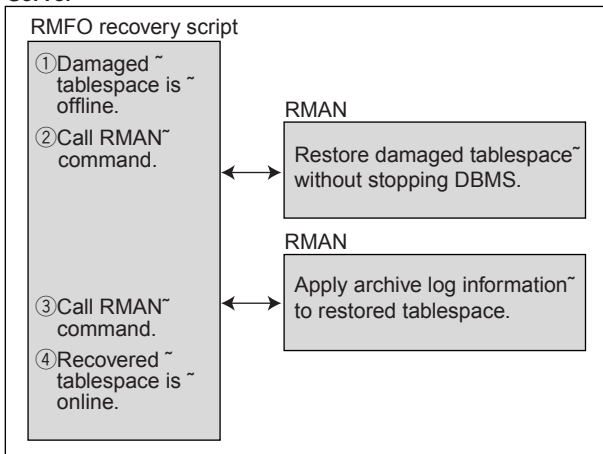


Figure 7

Summary of RMFO partial online recovery.

**Server**

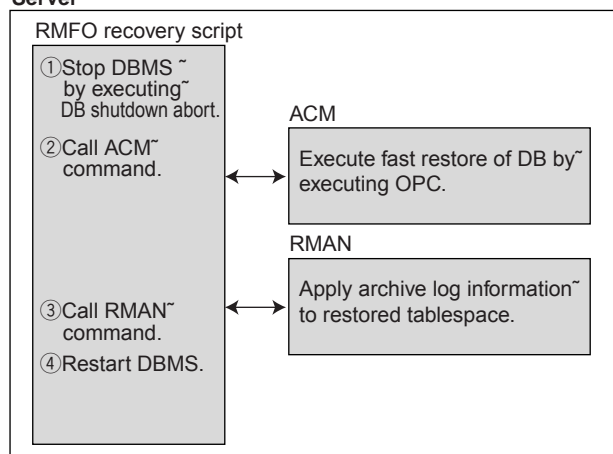


Figure 8

Summary of OPC fast restore and recovery.

Administrative tasks can be simplified and automated by using RMAN to administer the backed up data.

**Figure 5** summarizes the RMFO backup operation.

2) Simplifying the recovery procedure

By using the simple, user-friendly menu provided by RMFO, it is no longer necessary to use complicated commands, and recovery can be performed using simple operations. Also, RMFO provides information in the menu so that recovery commands can be automatically created and

executed. Moreover, in conjunction with RMAN, specific portions of a database can be recovered without stopping the DBMS (RMAN partial online recovery). Further, because an entire database can be recovered using OPC, which performs high-speed copying, the downtime can be minimized (OPC fast restore and recovery). These two types of recovery are shown in **Figure 6**.

**Figures 7** and **8** summarize the RMFO partial online recovery and OPC fast restore and recovery operations, respectively.

The RMFO scripts can be scheduled using

scheduling software such as Systemwalker OperationMGR. By combining ACM and ETERNUS's Advanced Copy feature, it is possible to build a simple backup system for an Oracle database and achieve high-speed backup and automation without special database skills.

## 5. Conclusion

There are several database backup solutions that can be used without needing to stop the database. However, the distinguishing feature of

Fujitsu's ACM backup and recovery solution is that it performs high-speed backups without interrupting business applications. In addition, it has a user-friendly menu that enables users to automatically create and execute backup scripts and rapidly recover a database without needing special database skills. We will continue to automate and simplify backup and recovery procedures by supporting more DBMS types, supporting disk additions, and preventing problems due to configuration changes and other causes.



**Isao Igarashi** received the B.S. degree in Communication Engineering from Tohoku University, Sendai, Japan in 1977. He joined Fujitsu Ltd., Tokyo, Japan in 1977, where he has been working in the Software Group.