Performing a Successful and Cost-Effective Migration of Your SAP System: Understanding the Import Process

Michael A. Moore



Michael A. Moore manages the Information Technology department for a consumer goods manufacturer in the United States. He has 16 years of experience in both application development and system administration. During the past four years, his group has completed two SAP R/3 upgrades, two disaster recovery tests, a disk storage subsystem upgrade, and a successful platform migration.

(complete bio appears on page 74)

When faced with reducing costs, consolidating dissimilar SAP systems, or adding processing power without breaking the bank, you have to consider migrating your SAP R/3 systems. Each vendor's hardware platform has its limitations and associated costs. As companies grow — and shrink — managers face the challenge of "rightsizing" their installed infrastructure and the associated costs to match the company's size. Based on the experience gained from migrating my own organization's SAP R/3 system, the purpose of this article, together with a previously published companion piece,¹ is to help you identify the risks associated with an SAP R/3 migration and to provide insight into how to effectively manage the entire migration process.

An SAP R/3 migration consists of changing either the operating system or the database platform.² During the migration, SAP R/3 data is exported from the source system and imported into a target system. The previous article addressed the requirements you must have in place for SAP to provide you with end-to-end support of the migration process, including how to secure an SAP-certified migration consultant, the keys to a successful migration plan,³ and how to use SAP's migration services. It also detailed how to complete the data export process using the SAP OS/DB Migration Kit,⁴ and pointed out several shortcuts for completing

¹ "Performing a Successful and Cost-Effective Migration of Your SAP System: Understanding the Export Process" (*SAP Professional Journal*, March/April 2004).

In this article series, an SAP migration is defined as the export of an SAP database and the import of that data onto a dissimilar target platform. The operating system or database (sometimes both) will change during the migration process, but the SAP release will remain the same.

³ Migration planning was also covered in detail in the article "A Step-by-Step Guide to Planning for a Successful OS/DB Migration Project" (*SAP Professional Journal*, November/December 2002).

⁴ The SAP OS/DB Migration Kit is an adequate, cost-effective solution for small and medium-sized installations that have less than 1 GB of allocated database space. For all other installations, you will need a customized solution provided by a consultant.

Toolkits Required for the Data Import Process

Two primary SAP software packages will be required during the import process: the SAP OS/DB Migration Kit and the SAP Installation Kit for your chosen target platform and database. The SAP OS/DB Migration Kit includes the tools and documentation that control the overall export and import process. The SAP Installation Kit includes the database and SAP R/3 executables, as well as the initial database required for SAP operation and testing of the target platform.

The SAP Installation Kit and the SAP OS/DB Migration Kit can both be ordered from SAP at **http://service.sap.com/softwarecat**.

and optimizing the export process. This article picks up where the previous one left off by taking you through the steps required to import that exported data into a target system and to validate the results of the import.

The SAP OS/DB Migration Kit provides the software and documentation necessary to complete a full export of the source system data. The same toolkit will be required during the import process, along with an additional set of tools: the SAP Installation Kit for your target platform operating system and database combination, which provides the SAP R/3 and database executables that are installed during the import processing (see the sidebar "Toolkits Required for the Data Import Process" above).

The export process described in the previous article utilized the SAP OS/DB Migration Kit and covered all the steps necessary to complete a data export. Some of the steps performed during the data export, such as installing the OS/DB Migration Kit, will be used during the import as well, but will not be described in this article, since they were covered in detail in the previous article. Instead, I will limit my discussion to the import process itself, concentrating on how to prepare and validate the target system, how to perform the data import, and how to carry out postimport processing. I'll also include some useful tips on optimizing the entire import process to minimize the outage time of your production system. For your reference, **Figure 1** summarizes the migration process as a whole, including the preparation and data export tasks covered in the previous article; a detailed list of tasks required for the entire import process is included in the appendix to this article.

✓ Note!

Remember from the previous article that the migration used to illustrate the discussion involved moving 420 GB of data from one version of the Unix/Informix platform, running SAP R/3 4.6C, to another version of the Unix/Informix platform, also running SAP R/3 4.6C, while using the same database management system in both the source and target systems. We started with a 420 GB database on a three-year-old server with twelve 240 MHz processors and 16 GB of memory. The selected target platform was new equipment with eight 668 MHz processors and 16 GB of memory. The underlying storage architecture did not change — the new server was simply attached to the current Storage Area Network (SAN). The first (untuned) migration required 26 hours to export the data and another 20 hours for the import. The final production migration took 12 hours to export and 8 hours to import.

Eva	luating the Migration	Duration
1.	Identify the potential target platform.	20 days
2.	Ensure all required software is supported on the target platform.	
3.	Evaluate potential migration consultants and migration costs.	
4.	Order the OS/DB Migration Kit.	
5.	Install the OS/DB Migration Kit.	
6.	Download all required patches.	
7.	Complete the first export.	
8.	Make the decision to migrate.	
Pla	nning for the Migration	Duration
1.	Select a migration consultant.	14 days
2.	Sign up for the OS/DB Migration Service.	
3.	Submit a migration plan to SAP.	
4.	Order new hardware.	
5.	Order an SAP Installation Kit for the target platform.	
Pre	paring the Migration Process	Duration
1.	Have the database administrator tune the database for the export.	15 days
2.	Have the operating system engineer tune the operating system for read operations.	
3.	Have the Basis administrator create profiles for the new hardware.	
4.	Have super users and application developers complete the test plan design.	
5.	Receive and install the new hardware.	
6.	Install SAP R/3 on the target hardware.	
7.	Apply patches to the target installation.	
8.	Import the exported data to the target platform.	
9.	Troubleshoot the import process.	
10.	Tune the import process.	
11.	Perform all post-migration system testing.	
Per	forming the Migration	Duration
1.	Receive SAP approval of the migration plan.	21 days
2.	Perform the final development system migration.	
3.	One week later, perform the first production system migration.	
4.	Test the results, including nightly, weekly, and monthly jobs.	
5.	After a minimum of two weeks, perform the final production system migration.	
6.	Perform a homogeneous system copy to create the target test/QA system.	
Tot	al Project Duration	70 days

Target System Preparation and Validation

Prior to the actual import of the source system data, the target platform must be prepared to receive it. This process consists of four major steps:

- 1. Install and configure the target operating system.
- 2. Maintain user IDs (i.e., *sapr3*, *<sid>adm*, etc.) for the target operating system.
- 3. Define the target file system.
- Install SAP software⁵ on the target operating system (the SAP software installation process includes installing a basic, functional SAP instance, which is included as part of the SAP Installation Kit, for validation purposes).

Step 1: Install and Configure the Target Operating System

To run both SAP and your database platform on the target server, all operating system prerequisites must be identified and completed. Review the documentation included in the SAP Installation Kit for your target platform. This documentation is unique to your combination of operating system and database software. As an illustrative example, the details associated with the successful SAP R/3 installation we performed on our target platform are described in the sidebar on the next page.

In addition to general operating system considerations, the SAP Installation Kit includes references to several SAP Notes that deal with special requirements for your particular operating system and database configuration, such as Unix kernel tuning, Microsoft Windows registry settings, and installing the German language file sets. For a smooth installation, be sure to locate and read the designated SAP Notes for your specific OS/DB combination.

Recall that when designing an SAP system, SAP central instance services can be installed either on the database server or on a separate server. For the purposes of this article, we will assume that the database and SAP central instance software will be installed on the same server in a Unix environment.

AIX/Informix Operating System Requirements

Each platform has its own set of specific requirements for the operating system installation. Individual installation guides can be found in the SAP Service Marketplace at http://service.sap.com/instguides.

The following example is taken from the 4.6C "R/3 Installation on UNIX: Informix Database" guide:

- ✓ Verify that all required operating system software is installed. The AIX command *Islpp -I* will produce a listing of all installed software. The output should include the following file sets:
 - bos.rte -- Base operating system runtime
 - bos.adt Base application development
 - bos.data Base operating system data
 - bos.sysmgt System management
 - bos.diag.rte Hardware diagnostics database
 - bos.msg.en_US Base operating system runtime messages (US English)
 - bos.net.nfs Network file system
 - bos.net.tcp TCP/IP
 - perfagent Performance agent
 - bos.loc.iso.en_US Base system locale code set (US English)
 - bos.loc.iso.de_DE Base system locale code set (German)
 - bos.iconv.de_DE Base level file set (required for locale code set)
 - bos.iconv.en_US -- Base level file set (required for locale code set)
 - devices.* Device drivers for all installed hardware
 - printers.rte Printer backend (if printer installed)

Step 2: Maintain User IDs for the Target Operating System

Once the operating system has been installed, the user IDs must be configured. In order to minimize the amount of change to the migrated data, you should maintain consistent user and group properties between the source and target platforms.

On Unix systems, this requires identical SAP R/3

UIDs and GIDs (the unique number assigned to user and group names, respectively). If the target operating system already has a particular source SAP R/3 ID assigned elsewhere by default, you can reassign the ID and all associated files (the sidebar to the right shows you how). Verify that additional user properties such as CPU time, file size creation, data and stack allocation, and dump file creation are configured as specified in the installation documentation.

- X11.base AIX/Windows runtime
- X11.apps AIX/Windows applications
- X11.motif AIX/Windows motif
- X11.fnt.iso1 AIX/Windows Latin 1 fonts
- X11.loc.en_US AIX/Windows locale (US English)
- X11.msg.en_US AIX/Windows messages (US English)
- X11.Dt AIX/Windows desktop
- *xlC.rte* C++ set for AIX application runtime (version 3.1.4.8 or higher)
- Add the following command to the standard Unix */etc/inittab* directory to enable the Portable Stream Environment (PSE) at system initialization:

strload:2:once:/usr/sbin/strload

Add kernel-specific parameters for asynchronous I/O, including the maximum number of processes per CPU, the process priority, and the queue length. Your database documentation will indicate initial settings; once the system is in production, however, these parameters will need to be tuned.

The most overlooked step in the process (at least in the US) is the installation of the German language file sets (*bos.loc.iso.de_DE* and *bos.iconv.de_DE* in the example above). The SAP software installation will fail if these file sets are not installed.

Also keep in mind that time synchronization is configured during the target operating system construction process, and that not all time synchronization programs are compatible. For instance, there is a known problem with synchronizing some Unix variants with Microsoft Windows time servers due to the differences between Network Time Protocol (NTP) and Simple Network Time Protocol (SNTP). Make no assumptions when changing platforms — verify that all requirements can be met in the new environment. For example, we had to react to the NTP/SNTP issue by creating a new time server for the new SAP environment, and then scheduling a firewall upgrade.

Changing a Default User or Group ID on the Target System That Conflicts with a Source SAP R/3 ID

During our migration, the default installation for the target Unix environment used both the *<sid>adm* UID (200) and the *sapsys* GID (200) to control access to system documentation. In order to avoid changing the user and group properties on all the SAP files to be copied from the source system, we chose instead to change the UID and GID for the documentation application delivered with the operating system.

(continued on next page)

Here are the steps we followed to complete this change, which can be easily modified to fit your particular scenario:

- 1. Stop the documentation application and associated services.
- 2. Change the UID for the user docadm.
- 3. Change the GID for the group docadm.
- 4. Run the following script to change user properties on all files:

5. Run the following script to change group properties on all files:

6. Restart the system documentation application.

Step 3: Define the Target File System

As with user IDs, the SAP file systems should have consistent properties across the source and target landscapes. When building the required file systems on the target system, you will need to diverge from the standard SAP installation documentation and base your initial file system sizes on the existing source system instead. The reason for this is that SAP documentation indicates minimum required file system

Name	Size	Ownership	Permission
/usr/sap/ <sid></sid>	1 GB	<sid>adm:sapsys</sid>	775
/usr/sap/trans	500 MB	<sid>adm:sapsys</sid>	775
/informix/ <sid></sid>	1 GB	informix:informix	777
/sapmnt/ <sid></sid>	1 GB	<sid>adm:sapsys</sid>	777
/usr/sap/put	250 MB	<sid>adm:sapsys</sid>	777
/usr/sap/put2	250 MB	<sid>adm:sapsys</sid>	777
/sap_cd	12 GB	<sid>adm:sapsys</sid>	777

Figure 2 Example File System Sizes and Permission/Ownership Values (Unix/Informix)

sizes and will not account for the growth of your system over time. Be certain to change the ownership and permission values for the target file system to the values specified by the installation documentation; otherwise, the configuration process will fail. **Figure 2** shows the permissions and file system sizes required for our target file systems. Note that along with moving the database contents, you will need to move several associated transaction files from the source to the target system (the standard Unix /usr/sap/trans directories, for example) during the migration process. This will not happen automatically — you will need to manually back up and restore the contents.

After allocating the file system used for SAP files, allocate the database table space using the source sys-

tem sizes as a guideline. If the database installation requires raw disk (i.e., Informix), be certain to set the permissions and ownership appropriately (see the sidebar below).

Network resources will need to be defined in preparation for the SAP import to the target platform. Configure the target system with the network-defined printers and any network file systems that are in use on the source SAP instance. You may want to script the printer definition process, as the process will need to be completed on each spool server and system (development, quality assurance, and production). Complete the operating system installation by configuring the remaining maintenance processes — including time synchronization, backups, and error reporting.

Defining Raw Disk

Most current database management systems configure table spaces in terms of either defined file systems or raw disk. File systems have the benefit of being more flexible with regard to growth, but their use also incurs additional operating system overhead. For large databases where performance is a requirement, raw disk configuration is often the optimal choice. In AIX, raw disk is carved into logical volumes and then assigned to the Informix database as "chunks." For our migration, we selected a standard size of 2 GB per chunk in order to make the task of allocating additional database space more manageable.

The R3SETUP installation program locates disk for the Informix database in the /Informix/<SID>/sapdata directory. Each subdirectory should point to a separate physical volume, such as:

- /informix/<SID>/sapdata/physdev01
- /informix/<SID>/sapdata/physdev02
- /informix/<SID>/sapdata/physdev03

Physical volumes on Unix systems, however, are defined as special device files in the /dev directory:

- /dev/rhdisk1
- /dev/rhdisk2
- /dev/rhdisk3

Therefore, links are required to allow SAP to allocate disk space from the directory where it expects to locate the disk definition (*/informix/<SID>/sapdata/*).

(continued on next page)

The following command ensures that the physical device is available to R3SETUP when it begins to assign disk to the database spaces:

ln -s /dev/rhdisk1 /informix/<SID>/sapdata/physdev1

Ensure that the user and group ownership permissions are correct for the /dev/rhdisk1 device! These raw devices should have read and write permissions and be owned by the *<sid>adm* user and the *sapsys* group. Otherwise, the R3SETUP program will terminate when it tries to allocate space.

🖌 Tip

It is best to document in one place all the steps required to build out the operating system. Script the user, group, and file system creation processes since these will be completed several times — on the development, quality assurance, and production platforms, for example.

Step 4: Install SAP Software

At this point, the operating system, required user and group IDs, and file systems must be configured and defined. The next step is to install the base SAP R/3 database as delivered in the SAP Installation Kit for your target platform.

This basic SAP R/3 installation will be used to validate the operating system configuration and to perform database backup and restore testing. The primary purpose of this basic installation is to enable you to complete a test installation on the target platform, log in to an SAP instance, and perform backup and restore testing in a controlled environment using a small database. This will be much more efficient than using your production instance, and should allow you to quickly identify any potential problems before you perform the actual import.

✓ Note!

A larger database size will likely change the database configuration parameters used for the SAP-delivered installation, and the backup will begin to span tapes, create more log files, etc. Consequently, the restore process will almost assuredly be changed as a result of the larger database size. For this reason, backup and restore testing on the SAP-delivered database does <u>not</u> replace the final backup and restore testing of the migrated database.

The installation process is started by installing the R3SETUP installation tools from the SAP Installation Kit into the */usr/sap/put* directory.⁶ Next, as the root user, start the installation process from the */usr/sap/put* directory using the following command:

R3SETUP -f CENTRDB.R3S

This scripted installation process will validate the base configuration of the target system, install the

This is the same process used to install the SAP OS/DB Migration Kit, as described in "Performing a Successful and Cost-Effective Migration of Your SAP System: Understanding the Export Process" (*SAP Professional Journal*, March/April 2004).

🖌 Tip

The installation CDs are necessary for both the test installation of the SAP-delivered database and the actual import process. To speed up the installation process, create a file system /sap_cd with about 12 GB of space and define the following subdirectories:

- export1
- export2
- export3
- export4
- informix
- kernel

Then copy the contents of the associated CDs to the corresponding subdirectory. On certain Unix systems, the files will be copied, but the case of the file name will be lost — all files will be named in lowercase. Unfortunately, the file names are required to be in uppercase for the installation program to run. The following script will convert lowercase characters to uppercase for all files in a given directory (see also SAP Note 48550).

for i in 'ls'
 do
 mv \$i echo \$i | tr "[a-z]" "[A-Z]"
done

SAP-delivered database and SAP R/3 software, and then import a minimal base SAP R/3 system. During this process, the user ID and file system permissions are validated, and the SAP R/3 and database software executables are then installed and configured. Finally, the data import of the SAP-delivered database will start. Following that import, database indexes will be created, an update statistics process will run, and the SAP R/3 application will be started.

The detailed installation process flow, as scripted by the *CENTRDB.R3S* command file, is shown in **Figure 3**.

Figure 3 The Installation Process Flow

Installation Details

- 1. Verify the SAP system:
 - SAP system ID
 - SAP instance number
 - Database and central instance host names
 - Location of the sapmount directory
 - Location of the installation CDs
 - Database installation details (i.e., 32 or 64 bit, etc.)
- 2. Verify the operating system:
 - Memory allocation details (i.e., allocation of RAM to SAP instance vs. database instance)
 - Network parameters (i.e., TCP/IP port information)
 - User and group configuration (both SAP and database)
 - Installation of German file sets
 - Kernel/registry settings
 - Whether LDAP will be used
- 3. Verify the database software:
 - License information
 - Basic configuration parameters
 - Migration key
 - Number of parallel processes to use when loading the database
 - Password for the SAPconnect (sapr3) user

Installation Process

- 4. Install the SAP R/3 software.
- 5. Create the SAP R/3 default instance profiles.
- 6. Install the backend software (database management system).
- 7. Install the frontend software (client access interface).
- 8. Build database spaces.
- 9. Extend database spaces.
- 10. Create default database configuration.
- 11. Start database software.
- 12. Grant permissions to the SAPconnect (sapr3) user.
- Import of the SAP-Delivered Database
- 13. Complete the SAP installation using the SAPdelivered data.

Post-Import Processing

- 14. Create database indexes and views.
- 15. Update database statistics.
- 16. Start SAP R/3.

If you encounter an error at any time during this installation process, the process will exit and the automatically created *CENTRDB.log* file can be inspected to determine the reason behind the failure. In addition, several other log files (named *.*log*) will be created in the installation directory. List the log files in the */usr/sap/put* directory by date/time last modified files (use, for example, the Unix command *ls -lt* to list the files). Once the problem has been identified and resolved, the process can be restarted from the */usr/sap/put* directory by entering the following command as the root user:

R3SETUP -f CENTRDB.R3S

🖌 Tip

The main reason for failure at this point will be related to inadequate database space or log space. If the failure occurs during a load process, determine which database space is full using SAPDBA. If the process fails during the index creation process, add space to the TEMP database space and restart the R3SETUP -f CENTRDB.R3S process.

This procedure constitutes a relatively safe process; it affords minimal chances for problems and will create a sense of confidence in the target environment by showing an actual, live SAP system running on the new operating system or database platform. Complete the process by logging on to the new instance using SAPGUI as user DDIC in client 000 with the default password 19920706.

After celebrating the successful installation, request the new SAP hardware key and a permanent SAP license:

1. As the root user, execute the command:

/sapmnt/<SID>/exe/saplicense -get"

- Request a license key at http://service.sap.com/licensekey.
- 3. Once you have received the license key, install it using the command:

saplicense -install

At this point, you can rest assured that your target system can successfully run SAP R/3. But before you can truly consider the import complete, you need to install the actual, exported SAP R/3 production data on the new platform. We'll look at this next.

Importing the Exported Data into the Target System

The actual migration of the exported data into the target system takes place in two steps:

- 1. Move a copy of the exported data to the target system.
- 2. Import the production database tables and load them with the exported data.

There are a few differences between the SAPdelivered installation and the actual imported system that may create problems during the migration process — these are mainly related to actual space requirements vs. estimations. The most common problems occur because the physical space in the database is inadequate for holding the tables to be imported.

🖌 Tip

If you cannot overcome an import problem, go ahead and contact SAP support. In rare cases, the problems will require SAP support to deliver modified executables, such as R3load and R3ldctl.

Moving Files via File Transfer Protocol (FTP)

File Transfer Protocol (FTP) can be used to move files between the source and target systems. FTP supports two transfer modes: ASCII (a.k.a. TEXT) and BINARY. In a Unix ASCII file, each line of text ends with a single line-feed character, while a Microsoft Windows ASCII (TEXT) file uses a pair of characters — the carriage return and line-feed (in that order) — to signify end of line. In order to process the information correctly, the R3SETUP installation program requires that files be correctly formatted after they are moved to the target system. To further complicate matters, all non-ASCII files need to be moved as is using the BINARY transfer mode.

The consequence of this restriction is that if you are moving files from one operating system to another, you must be cognizant of differences in how those systems treat various file modes; otherwise, the data will be unusable on the target system. The table to the right details representative file types and the correct transfer mode to use to move the files between dissimilar operating systems.

If you are changing database management system software and *not* operating system software during the migration, just move all files using the BINARY mode.

File Extension	Mode
.STR	ASCII
.TOC	ASCII
. <i>nnn</i> (<i>n</i> = 0-9)	BINARY
.TPL	ASCII
.EXT	ASCII

Similar problems exist when the import process fails during index creation due to inadequate space in the temporary database spaces. Other problems are related to unpatched SAP R/3 and database software that is installed during the import process. The good news is that these issues can be easily overcome.

Step 1: Move a Copy of the Exported Data to the Target System

The first step in importing the data is to move a copy of the exported data to the target environment. Remember from the previous article that the database export process created data files and control files; however, several other files will also need to be moved to perform adequate system testing — primarily the transport history data in the standard Unix */usr/sap/trans* file system and any data or file system associated with bolt-on applications.

Move the data via File Transfer Protocol (FTP). Do not attempt to import data from a network file system volume, since your performance will be severely degraded due to the overhead associated with network file system processing. In addition, consider the physical network path when you start moving this data. The ideal situation will be to use a dedicated switched network circuit without sharing any bandwidth with other processes — for example, use a high-speed backup network while it is idle.

Refer to the sidebar above for instructions on moving the exported data to the target system using FTP. The target system directory you move the exported data to (in our case, a directory called /dbexport) must have an identical underlying file structure to that of the source system export data directory. Ensure that the file permissions and ownership are identical on both the source and target systems. Next, move all associated data files (those in */usr/sap/trans* as well as those required for bolt-on applications).

✓ Note!

Keep in mind that bolt-on applications may need to be reinstalled and configured in order to function on the target system. Bolt-ons are beyond the scope of this article and must be addressed on a case-by-case basis with your migration consultant and bolt-on vendor support.

If you haven't done so already, request the migration key required when migrating from one database or operating system environment to another. The migration key is generated automatically when you access **http://service.sap.com/migrationkey** and enter the source and target operating system and database specifics (refer to the previous article on the export process for more details). You should already have the appropriate bolt-on software products for the target platform. In addition, you need to determine if any software keys that are hardware-dependent will be required during the target system installation (i.e., for database or bolt-on applications).

Step 2: Import the Production Database Tables and Load Them with the Exported Data

The next step is to import the production database tables and load them with the export data you copied to the target system, but first you need to perform some cleanup. Stop both the SAP R/3 and database software (for Unix, don't forget to stop the SAPOSCOL process). Delete all files from the following directories, as they will be re-created during the migration import process:

- /informix/<SID>/
- /usr/sap/<SID>/
- /sapmnt/<SID>/

• database user home directory (*/home/Informix*, for example)

✓ Note!

The directory names listed previously are particular to the use of Informix as a database if you are using a different database system, you will need to adjust the names accordingly.

●^{*} Caution!

Be careful to ensure that the database software will still be able to access the disk space that is assigned for database table spaces. In our target environment of Unix and Informix, the database volumes are raw logical volumes that are made available to the database via links in the directory /informix/<SID>/sapdata. If this directory and its contents are deleted, they must be re-created before the migration process can successfully allocate space to the database spaces.

Now's the time to install the OS/DB Migration Kit on the target system. This is the same installation process used during the export except that the target directory will be */usr/sap/put2*. Prior to running the INSTTOOL command, remember to change to the directory in which you want to install the kit. As we did on the export system, patch the tool with the most recent software available from SAP's software distribution center (be sure to include the latest R3ldctl, R3load, and R3szchk tools). The same files are installed as during the source system migration kit installation — the difference is that when the migration process is started, the R3SETUP control file used is *CEDBMIG.R3S*.

Starting and Stopping the Export Process

To optimize export and import processes, you must be able to manipulate the .R3S control files to stop and start the export and import processes as needed. The example here has been coded to stop the .R3S control file execution just before the *SPLITSTR.PL* program runs during the data export process. The same syntax is used to stop processing during the import phase.

Begin by editing the DBEXPORT.R3S file and locating the [EXE] header, which is the control section of the file that will execute all the other sections in the correct order. Under the [EXE] header, locate the string 120=SPLITSTRFILES_IND_IND and insert above it the command 115=EXE_STOP:

```
[EXE]
...
100=R3LDCTL_IND_INF
110=R3SZCHK_IND_INF
115=EXE_STOP
120=SPLITSTRFILES_IND_IND
130=DBEXPCOPYR3LDCTLFILES_IND_IND
...
```

Immediately following the *[EXE]* section, create a new section called *[EXE_STOP]* with the instruction *CExitStep*, which tells the R3SETUP interpreter to stop all command processing:

```
[EXE_STOP]
CLASS=CExitStep
```

(continued on next page)

Before executing the import process for the first time, the *CEDBMIGR3S* file needs to be modified to stop after the SAP R/3 and database software has been installed, but prior to beginning the data import, in order to apply patches to these applications (see the sidebar above). Insert a breakpoint just after installation of the database software (for Informix installations, this is after the *INSTALLDBBE_IND_INF* and *INSTALLDBFE_IND_INF* steps):

```
390=INSTALLDBBE_IND_INF
400=INSTALLDBFE_IND_INF
405=EXE_STOP
```

As the root user, start the import process by

changing to the */usr/sap/put2* directory and issuing the command:

R3SETUP -f CEDBMIG.R3S

Answer the questions exactly the same as you did when installing the SAP-delivered database with one exception: when prompted to *Enter/confirm the path to the EXPORT1 CD*, enter the directory to which you copied the exported data (i.e., */dbexport*). Once the breakpoint has been reached, the SAP R/3 and database software will have been installed and the R3SETUP process will stop. Download all available kernel and database patches from **http://service.sap.com**. Apply these patches to the target system prior to continuing the migration, in

Next, begin execution of the R3SETUP program using the following command:

R3SETUP -f DBEXPORT.R3S

When the R3SETUP program reaches this step, a message will ask you to confirm that you wish to exit:

Please enter the parameter or confirm default [EXIT]:

Press *Enter* and the R3SETUP program will exit and return to the operating system.

You can also bypass a step without running it and without stopping the export process, by simply inserting the line *STATUS=OK* in the appropriate section. In the example below, the update statistics step is bypassed without stopping the export process, which becomes very handy if you complete the update statistics prior to beginning the export outage:

```
[DBCOMPUTESTAT4MIG_IND_INF]
CLASS=CExecUpd
...
RUN_AS_USER=@INFUSER@
STATUS=OK
STEP_ENV=DB_ENV
```

All of these steps are explained in greater detail in SAP Note 118059.

order to avoid potential problems that have been fixed with the patched software (for example, we ran into a problem with buffer overruns that could have easily been prevented with a patch).

Edit the *CEDBMIG.R3S* file and remove the breakpoint. Restart the import (as the root user, change to the */usr/sap/put2* directory and issue the R3SETUP command). The steps already completed will be bypassed and import processing will start. During the actual import, several R3load processes will start and data loading will begin in parallel. In the event that a single process fails, the remaining processes will continue to run and additional import processes will start automatically. Once the final import control file has been processed, either the index creation step will begin or the R3SETUP process will terminate with error conditions.

The primary reason for errors at this point in the

import is that the process used to calculate the space required for the data did not size the database space large enough to hold all the required data. The good news is that these errors can be easily corrected and the import process can be restarted without any manual cleanup or repetition of successfully completed steps. Using the SAPDBA tool, determine which database spaces are full and add space to them. Document the database spaces that require additional space and how much space was added. Then, restart the import process and repeat these steps until all the data has been successfully imported.

Once the data is loaded, the next step is to re-create table views and all the primary and secondary indexes. During the migration process, the temporary database space that is created on the target system is identical to the space allocated on the source system. However, additional temporary database space is required in order to create all the necessary indexes. During our migration, for instance, the database to be migrated was 420 GB with an initial temporary database space size of 12 GB. In order to successfully complete the index creation, we had to triple the size of the temporary database space available to 36 GB just to be able to successfully create the indexes.

🗸 Tip

If you have split the control files manually, it is very possible that the SAPVIEW.STR file may not be present in the /dbexport directory. The SPLITSTR.PL process does not create a SAPVIEW.STR file in the target directory because a view does not have physical data associated to it, as does a table proper. When the files in a target directory that were created as a result of the SPLITSTR.PL program are copied back into the working directory during the export process, ensure that the SAPVIEW.STR file is not removed inadvertently. For more information on this subject, see SAP Note 709389.

At this point the data and the indexes have been created. Next, the R3SETUP program will create the views and start updating the database statistics. The import processing continues by installing a temporary SAP license and starting the SAP R/3 central instance. Once the SAP instance is started, the final steps of the migration initialize the Workbench Organizer (transaction *SE06*) and schedule a minimal set of maintenance jobs.

At this point in the migration, a working SAP system is now present on the target hardware/software platform with the source SAP system data. Before turning the system over to your end users, however, you need to perform several cleanup steps to return the environment to a truly productive status. The checklist in **Figure 4** is provided by SAP in the *R/3 Homogeneous System Copy Release 4.6D* document and has been augmented with the minimal additional steps required for heterogeneous migrations.

Figure 4 Readying the Migrated System for Productive Use

Final Cleanup Tasks
Reinstall, configure, and move associated data for all bolt-on applications.
Reconfigure the transport layer.
✓ Optimize the SAP profiles for the target system.
Optimize the database configuration for the target system.
✓ Verify the database logging mechanism.
Run the installation check tool (transaction <i>SICK</i>).
✓ Verify the configuration of the Workbench Organizer (transaction <i>SE06</i>).
Configure jobs on the target system.
Check the consistency of the TEMSE objects (transaction <i>SP12</i> ; see SAP Note 16875).
Verify the printer configuration in both SAP and the operating system (transaction <i>SPAD</i>).
✓ Verify all RFC destinations (transaction <i>SM59</i>). Be careful to ensure that the RFC definitions point to the systems in the target landscape and not back to the source systems!
Configure the operation mode timetables that were disabled prior to the export (transaction <i>RZ04</i>).
\checkmark Maintain the instance profile records (transaction <i>RZ04</i>).
Change the default SAP system user passwords (transaction <i>SU01</i>).
Configure CCMS according to your specific requirements.
✓ Install the permanent SAP license (see SAP Note 566263).
Review the system log on all target servers for errors (transaction <i>SM21</i>).
Perform a consistency check of the database (transaction <i>DB02</i>).
Verify the application and central instance server processes (transaction <i>SM51</i>).
Perform additional checks as identified by your SAP-certified migration consultant.

Using the Unix diff Program to Compare Report Outputs

As the test plan is executed, several reports are produced that you can use to verify that the migrated data was moved without change. Compare these outputs (from both source and target) to ensure that they are equivalent. Unfortunately, manually looking for unexpected differences between voluminous reports can be a tedious and error-prone activity. Furthermore, be aware that even though some of their lines may differ, two report outputs can be considered successfully equivalent. (This happens most commonly when report output contains lines with date and/or time stamps, or perhaps the name of the user who ran the report.) Fortunately, Unix systems contain utility programs (*awk* and *diff*) that can be used to automate effective comparison of two text files.*

The actual report data can be written to a spool file and then exported as a text file using the SAP menu path $System \rightarrow List \rightarrow Save \rightarrow Local$ file. Place all the source system list files in one directory (e.g., *source_reports*) and the target system files in a separate directory (e.g., *target_reports*). In Unix, you can do a line-by-line comparison of the report outputs using the *diff* utility. To compare all files in the directories that have the same name, use the command:

diff source_reports target_reports

The output will detail the differences between the two files, which should include only expected variances (such as date or time stamps), if the results from the two systems are equivalent.**

Both Unix and Microsoft Windows users can also take advantage of *awk* scripts to remove all date and time stamp information from the files before file comparisons take place, thus eliminating from the comparison output the "noise" of expected or insignificant differences.

- * While Microsoft Windows doesn't sport a native *diff*, there are multiple Windows diff-like implementations available (as well as *awk* implementations). A quick Google for "Windows diff" and/or "Windows awk" will lead you to sources for useful OEM implementations of this functionality.
- ** Similar Microsoft Windows commands can be used their syntax will vary depending on the particular file compare utility being used.

Completing this checklist will validate the target system integrity prior to allowing users access to the migrated system.

Once the checklist is complete, the technical team can release the system for super-user validation. This process needs to be defined by the ABAP support staff in conjunction with the super users. A thorough test plan should be written by the application developers and super users that includes daily, weekly, monthly, quarterly, and year-end processing activities — both dialog and batch — and executed on both the source and target platforms. Super users should then compare the results and question any discrepancies (refer to the sidebar above for details on using the Unix *diff* program to compare report outputs). Once the super users are satisfied, initiate the nightly batch processing. Once you've verified that the nightly batch process has successfully completed, the system can be considered ready for end-user/production processing.

Helpful Hints for Optimizing the Import Process

The import process can be optimized in a number of ways in order to minimize outage time for your production system.

- ✓ The system import process can be started before the actual production outage begins. Complete the import process up to and including the point that the target system database and SAP software is patched. The target system is now staged and ready for the exported data to be copied to local disk and for the import to begin.
- ✓ The number of parallel import processes can be adjusted to minimize the time required to complete the database load. Be certain to use all the operating system resources, without causing any undue context switching or system paging. To optimize the process, import the database several times and monitor the resources in use during each import run. Use this method to determine the optimum number of sequential import processes that should be used during the production import.
- ✓ The default database configuration as defined by R3SETUP will not be optimal. R3SETUP creates an online transaction processing environment configuration, which is not optimized for massive data loads. By reviewing the database-specific literature for Business Intelligence processing (a.k.a. data warehousing), you can learn how to tune the database configuration for data load and index creation processing. (This subject is discussed in more detail in the previous article about the data export process.)
- ✓ Once the first database import has been successfully completed, the target database space size requirements will be defined. The additional space requirements can be made with one stop of the migration import process using the SAPDBA tool.
- ✓ Define a parallel update statistics process to use after the import completes. The update statistics process that R3SETUP will begin is single-streamed, and hence will only update statistics for one table at a time. This is a very inefficient process, and a custom parallel update statistics process can eliminate hours from the outage requirements. This is the exact same process that was used to optimize the export process (refer to the export article for additional details).

Post-Import Processing

At this point, the initial "test" migration can be considered complete. The data has been exported from the source platform and imported onto the target platform and validated. As this is the first time the process has been completed, both the source and target environments will contain identical data. SAP requires that you maintain a parallel environment for a minimum of two weeks. During that time, users should access both systems — i.e., enter production data into the source and target systems and then verify that the target system functions identically to the source system. Once the two weeks have passed and all parties are comfortable that the migration was successful, the actual final production (go-live) migration can be completed.

As part of the OS/DB Migration Check, SAP will complete two post-migration system checks following

the go-live migration. The first system health check will focus on tuning the SAP system profile and database parameters in order to optimize performance on the target system. The second health check will focus on the response times of the target system and compare the results to the historical times from the source system. SAP will provide reports that summarize the findings and provide recommendations to improve performance and reliability.

Conclusion

As information technology professionals, we are constantly being asked to deliver more for less — or face the threat of outsourcing. In the case of my organization, we were originally asked to reduce the overall system cost by 50% without affecting performance. The only way we could accomplish that aggressive goal was by changing hardware platforms. The SAP OS/DB Migration Kit not only simplified the migration process for us, it also helped to control the associated risks. In this two-part article series, I detailed what it means to perform a migration using the SAP OS/DB Migration Kit, and I provided you with some insight into where additional performance gains are possible.

The export process requires a stable, static environment to ensure data quality. The target system must be configured correctly to receive the exported data with minimal change. The testing that follows the migration will ensure that the data is complete and valid. This is clearly a complex and involved process, but it can be highly profitable. For example, our migration was considered a huge success for the company — we reduced systems costs by 60%, cut response times in half, and the baseline job's runtime dropped from 20 hours on the source system to 4 hours on the target. There were no visible negative effects — all data was migrated, all external systems worked, and there were no surprises. The best part was that all of this was accomplished using internal staff and a single migration consultant.

Now that you understand what the process requires and the rewards that it can deliver, you can make an educated decision about whether this is the right approach to cost control for your organization. Talk to your current vendor and then their competitors. Run the numbers and have the chosen vendor loan you a system (or time in their lab) to test the migration process. Try the migration, and leverage the vendor's SAP Competency Center for assistance. If successful, you're already on your way. Good luck!

Michael A. Moore manages the Information Technology department for a consumer goods manufacturer in the United States. He has 16 years of experience in both applications development and systems administration. Mike started working with SAP in January 2000, and during the past four years, his group has completed two SAP R/3 upgrades, two disaster recovery tests, a disk storage subsystem upgrade, and a successful platform migration.

Prior to his SAP experience, Mike was a senior operating system engineer for a national insurance company, where he was responsible for Unix systems located all over the East Coast of the United States. While there, Mike specialized in installing and managing Unix servers and applications requiring high availability, including data warehousing and imaging. Mike also has eight years of application development and software quality assurance experience gained while employed at various federal nuclear installations. Mike is a co-holder of a software patent (along with his former employer) for an application design that manages complex storage arrays. He can be reached at MikeAMoore@aol.com.

Appendix: The Complete Import Process Task List

Target System Preparation

- 1. Install the operating system:
 - Configure time synchronization.
 - Change the *docadm* user and group ID.
 - Verify all required software file sets have been installed.
 - Add the Portable Stream Environment (PSE) command to the standard Unix /etc/inittab directory.
 - Create users and groups.
 - Configure kernel asynchronous I/O on the database server only.
 - Create raw logical volumes.
 - Change ownership of the raw disk.
 - Create links to raw disk from /informix/<SID>/sapdata.
- 2. Install the SAP Installation Kit into the /usr/sap/put directory:
 - Download and update the files when newer versions are available.

3. Copy installation CDs to disk:

- Create file system.
- Create subdirectories.
- Copy contents of installation CDs.
- Change lowercase file names to uppercase.
- 4. Complete the SAP-delivered database installation:
 - Execute the command R3SETUP -f CENTRDB.R3S.
 - Stop the database and SAP R/3.
 - Patch both SAP and Informix.
 - Restart the application.
 - Verify logon.
- 5. Complete a test backup and restore:
- Resolve all issues before continuing.
- 6. Clean up the installation directories.
- 7. Insert a breakpoint into the CEDBMIG.R3S file after the database software is installed.
- 8. Start the SAP-delivered installation for a second time until the breakpoint is reached:
 - Patch Informix and SAP.

Production Outage Begins

- 9. Configure and complete the export.
- 10. Copy the following data to the target system:
 - · The data exported from the source system
 - The transport data in the standard Unix /usr/sap/trans directory

(continued on next page)

- 11. Remove the breakpoint from CEDBMIG.R3S.
- 12. Insert a new breakpoint before the database load.
- 13. Restart the import process until the breakpoint is reached:
 - Add space for TEMP and other database spaces.
- 14. Remove the breakpoint from CEDBMIG.R3S.
- 15. Insert a new breakpoint before the update statistics run.
- 16. Restart the import process:
 - Complete a parallel update statistics run.
- 17. Remove breakpoint from CEDBMIG.R3S.
- 18. Restart the import process.
- 19. Wait for the import process to run to completion.
- 20. Complete post-import processing:
 - Reinstall, configure, and move associated data for all bolt-on applications.
 - Reconfigure the transport layer.
 - Optimize the SAP profiles for the target system.
 - Optimize the database configuration for the target system.
 - Verify the database logging mechanism.
 - Run the installation check tool (transaction SICK).
 - Verify the configuration of the Workbench Organizer (transaction SE06).
 - Configure jobs on the target system.
 - Check the consistency of the TEMSE objects (transaction SP12; see SAP Note 16875).
 - Verify the printer configuration in both SAP and the operating system (transaction SPAD).
 - Verify all RFC destinations (transaction *SM59*). Be careful to ensure that the RFC definitions point to the systems in the target landscape and not back to the source systems!
 - Configure the operation mode timetables that were disabled prior to the export (transaction RZ04).
 - Maintain the instance profile records (transaction RZ04).
 - Change the default SAP system user passwords (transaction SU01).
 - Configure CCMS according to your specific requirements.
 - Install the permanent SAP license (see SAP Note 566263).
 - Review the system log on all target servers for errors (transaction SM21).
 - Perform a consistency check of the database (transaction *DB02*).
 - Verify the application and central instance server processes (transaction SM51).
 - Perform additional checks as identified by your certified migration consultant.

Migration Complete