

FAQ: OS/DB Migration to Microsoft SQL Server v2.8 Mar 2011

Summary

You are currently running an SAP system on a Unix, Windows or Linux operating system and Oracle, Informix, DB2 or MaxDB database and wish to migrate your SAP system to Microsoft SQL Server.

You may also wish to convert your SAP system to Unicode during the migration to SQL Server.

Background Information

SAP & Microsoft have extended the capabilities of the SAP OS/DB migration tools and procedures to simplify the process of migrating SAP systems to SQL Server. This note contains the latest information regarding the technical capabilities and features for OS/DB Migrations where the target database is SQL Server.

Please review the latest blogs at : <http://blogs.msdn.com/b/saponsqlserver/>

Solution

There are a number of new enhancements that will significantly speed up and simplify the process of migrating a SAP system to SQL Server. In general all of these features are available for all systems based on WAS 6.20 and above (such as SAP R/3 4.7 Extension Set 110 and higher).

The link <http://service.sap.com/systemcopy> contains more information on the OS/DB Migration process. Also review note [82478](#).

Customers should target conversion throughput of around 1-2TB per hour using all of the enhancements contained in this document.

RECOMMENDATIONS

1. Required patch levels for Migration Tools, Windows & SQL Server

You **must** use these patch levels or higher for the following components. It is generally recommended to use the most recent version of these components.

SAPInst & R3SETUP

7.0	Netweaver 7.0 Master SR3 (Date stamp should be >= 09.03.2010)
6.40	Netweaver 04 Master SR1 (compatible with 6.20/R/3 4.7 systems)
4.6D	R3SETUP 46D SR1 (for use on 4.6C - available on request from SAP)

R3LOAD

7.01	Patch Level 34 or higher
7.0	Patch Level 188 or higher
6.40	Patch Level latest
4.6D	Patch Level latest

DBSL

7.01	Patch Level 34 or higher
7.0	Patch Level 188 or higher
6.40	Patch Level latest
4.6D	Patch Level latest

MIGMON

Java based Migration Monitor is downward compatible 7.1, 7.0, 6.40, 4.6C and lower. Use the most recent version. To download Migmon check OSS Note [784118](#)

R3TA

R3TA Table Splitter is only available for 7.1, 7.0 and 6.40. Use the most recent version.

R3LDCTL, loadercli & R3SIZCHK

Use the most recent version.

System Copy OSS Notes

7.1/7.0	Note 888210 and 1120872
6.40	Note 784931 and 771209

4.6D Note [316353](#)
SAP_BASIS Support Pack

Release 4.6C	53
Release 6.20	57
Release 6.40	15
Release 7.00	0 (SQL 2005) 14 (SQL 2008/R2)
Release 7.10	0 (SQL 2005) 0 (SQL 2008/R2)

See note [799058](#) for SQL 2005 and notes [1152240](#) for SQL 2008/R2

Windows & SQL Server

Windows 2003 x64 or IA64 - Use the latest service pack as documented in Note [30478](#) Windows Service Packs

SQL Server 2005, 2008* or 2008 R2** Enterprise Edition x64 or IA64 - [download](#) the latest service pack as documented in Note [62988](#) Service packs for Microsoft SQL Server

*at least CU2 for SQL 2008 must be used. Recommended release is SQL 2008 SP2 (released 24th September 2010)

**Support is documented in note [1076022](#)

Do **not** to use 32bit versions of Windows or SQL. If your system is 4.6C based run 4.6C on 64 bit Windows 2003 and 64 bit SQL 2005.

2. Unsorted Export

An unsorted export is supported and may be imported into a SQL Server database. A sorted export will take much longer to export and is only marginally faster to import into SQL Server. Unicode Conversion customers must export certain cluster tables in sorted mode. This is to allow R3LOAD to read an entire logical cluster record, decompress the entire record (which may be spread over multiple database records) and convert it to Unicode. See Note [954268](#), [1040674](#) and [1066404](#). The content of OSS Note [1054852](#) is out of date and does not account for the fact Intel platforms outperform UNIX platforms dramatically and normal customer DB servers have 128GB to 1TB of RAM. Providing the minimum R3LOAD patch levels are met there are no special actions required and R3LOAD will automatically export the required tables sorted. BW Fact table can also be exported unsorted

Our default recommendation is to export unsorted as in the majority of cases the UNIX/Oracle or DB2 server has only a fraction of the CPU, IO and RAM capacity of a modern Intel commodity server. Even though there is an overhead involved in inserting rows into the clustered index on SQL Server, this overhead is relatively small.

3. Table Splitting

A table split export is fully supported and may be imported into a SQL Server database. Table split packages for the same table may be imported concurrently.

Table splitting is only supported for R3LOAD 6.40 and higher (R3LOAD 6.40 is backwards compatible with Basis 6.20 releases such as R/3 4.7). Review Note [952514](#)

The limitations on SQL Server table splitting listed in the system copy guide are out of date and should be ignored. Customers have successfully split large tables into 20-60 splits and achieved good results on tables that have poor import or export throughput.

There are some tables that we always recommend to split due to slow export or import performance: CDCLS, S033, TST03, GLPCA, STXL, CKIT, REPOSRC, APQD, REPOTEXT, INDTEXT

To run R3TA manually use this command line.

```
r3ta -f c:\export\abap\data\



c:\export\abap\data\




```

Using this command in Excel a command line can be built

```
=CONCATENATE("R3TA -f d:\export\abap\data\",A9,".str ","-l ",A9,"_WHR.log"," -o d:\export\abap\data\",A9,".WHR","  
-table ",A9,"%",B9)
```

After generating WHR files with R3TA the WHR splitter must be run to create split packages. Always set the whereLimit parameter to 1, meaning 1 package for each where clause.

```
C:\export>where_splitter.bat -whereDir c:\export\abap\data\ -strDir c:\export\ab  
ap\data -outputDir c:\export\abap\data -whereLimit 1
```

4. Package Splitting

The Java based Package Splitting tool is fully supported in all cases. It is recommended not to use the Perl based splitter.

Note: Exports to SQL Server do not need Extent files and the whole Extent file (*.EXT) file generation process can be skipped to save time. Instead it is recommended to use the following script to determine the largest tables in the Oracle database:

```
spool tablefile.txt
set lines 100 pages 200
col Table format a40
col Owner format a10
col MB format 999,999,999
select owner "Owner", segment_name "Table", bytes/1024/1024 "MB" from dba_segments
where bytes > 100*1024*1024 and segment_type like 'TAB%' order by owner asc, bytes
asc
spool off;
```

Then it is recommended to extract the largest tables (possibly anything more than ~2GB) into their own packages (and also table split if required). The following command can be used. Please note that when using SAPInst EXT files are required. EXT files can be bypassed only when doing a manual Migmon based migration

```
str_splitter.bat -strDirs d:\export\abap\data -outputDir d:\export\abap\data -
tableFile tablefile.txt
```

This command will generate the TPL files and the default STR files (without the EXT files)

```
r31dctl -l logfilename -p D:\exportdirectory
```

5. FASTLOAD

All SAP data types can now be loaded in Bulk Copy mode. It is recommended to set the `-loadprocedure fast` option for all imports to SQL Server. These are the default settings for SAPInst. If migration monitor is used this parameter must be specified. 4.6C/D migrations should use the parameter `-fast` (without the "loadprocedure"). Please also note that to support FastLoad on LOB columns set environment variable `BCP_LOB=1` and review note [1156361](#)

The parameters we recommend for Migmon or SAPInst are `loadArgs=-stop_on_error -merge_bck -loadprocedure fast`

6. Migration Time Analyzer

It is recommended to use MIGTIME with the `-html` option to graphically display the export and/or import time of packages. It is generally recommended to ensure the longest running packages are started at the beginning of the export or import. MIGTIME is available for 4.6C and higher

```
Import_time.bat -installDirs d:\import -html
```

The script below shows the actual status of the SAP Export using SAP MigrationMonitor log files.

The script reloads every 20 seconds and displays

- actual CPU Load
- Actual running Packages
- Actual waiting Packages



MigMonStatus.zip

Before first usage:

- Unzip the MigMonStatus archive in the Migration Monitor directory
- Rename status.txt in status.cmd
- rename queryCPU.txt in queryCPU.vbs
- start the status.cmd

7. Package Order by recommendations

It is recommended to use an OrderBy.txt text file to optimize the export of an Oracle system and the import to SQL. By default a system will export packages in alphabetical order and import packages in size order.

The OrderBy.txt can be used to instruct Migration Monitor to start packages in a specific order. Normally the best order is to start the longest running packages first. It is recommended to perform an export on a test system to determine which tables are likely to run longest.

Note: It is normal for the export and import runtimes of a package to be very different. Some packages may be very slow to export yet very fast to import and vice-versa.

8. Oracle Source System Settings

Please review note [936441](#) - Oracle settings for R3load based system copy

SAP have released [OSS Note 1043380](#) which contains a script that converts the WHERE clause in a WHR file to a ROW ID value. Alternatively the latest versions of SAPInst will automatically generate ROW ID split WHR files if SAPInst is configured for Oracle to Oracle R3LOAD migration. The STR and WHR files generated by SAPInst are independent of OS/DB (as are all aspects of the OS/DB migration process).

The OSS note contains the statement “ROWID table splitting CANNOT be used if the target database is a non Oracle database”. Customers wishing to speed up an export from Oracle may send an OSS message to BC-DB-ORA and request clarification of this restriction. Technically the R3LOAD dump files are completely independent of database and operating system. There is one restriction however, restart of a package during import is not possible on SQL Server. In this scenario the entire table will need to be dropped and all packages for the table restarted.

OS/DB Migrations larger than 1-2TB will benefit from separating the R3LOAD export processes from the Oracle database server.

Note: Windows application servers can be used as R3LOAD export servers even for Unix or mainframe based database servers. Intel based server have far superior performance in SAPS/core than most Unix servers, therefore R3LOAD will run much faster on Windows.

The simplest way to allow Windows R3LOAD to logon to Unix Oracle server is to change the SAP<SID> on schema systems or saps3 on non-schema systems to “saps3” without quotes. This password is hardcoded into R3LOAD. If the password cannot be changed then the user account on the R3LOAD Windows server (normally DOMAIN\<sid>ADM) will need to be added to the SAPUSER table OPS\$<DOMAIN>\<SAPSID>ADM

9. SQL Server Target System Settings

It is recommended to use Windows 2008 SP2 and SQL 2008 SP1 or higher. Only 64bit platforms are recommended. 32bit platforms are now depreciated and **customers are instructed not to use 32bit versions of Windows or SQL Server**. SAP R/3 4.6C offers no native x64 kernel however 4.6C 32 bit kernel can run on Win2003 x64 and is fully supported by MS & SAP.

The SQL Server database should be manually extended so that the SQL Server automatic file growth mechanism is not used as it will slow the import. The transaction log file should be increased to ~10GB for large systems.

Max Degree of Parallelism should be set to 1 usually. Due to the logic for parallelizing index REBUILD or CREATE statements it is highly likely that most index creation on SAP systems will be single threaded irrespective of what MAXDOP is specified

SQL 2005 CU8 contains minimal logging feature for SAP migrations.

See SAP Note [1241751](#). It is recommended to use SQL 2005 SP3, SP4 or SQL 2008 or higher as these releases have this feature built in.

To activate minimized logging start SQL Server with Trace Flag 610. See note [1482275](#)

If R3LOAD or SQL Server aborts during the import, you need to drop all the tables which were in process at that time. The reason is that there is a small time window where data should be written to disk in a synchronous manner, but the writes are asynchronous. Therefore the consistency of the table cannot be guaranteed and the table should be dropped and the import restarted.

In general we recommend 610, 1117 and 3917. To display trace flags run `DBCC tracestatus`
Remove trace flags 610 and 3917 after the migration. These traceflags should not be used live systems.

10. Setting up a standalone R3LOAD server – SQL and Oracle

OS/DB Migrations larger than 0.5-1TB will benefit from separating the R3LOAD import processes from the database server:

- a. Install SQL Server 2005, 2008 or 2008R2 SNAC (client libraries only)
- b. Apply Service Pack 3 or 4 for SQL Server 2005 or at least SP1 for SQL 2008 or SQL 2008 R2
- c. Install Java SDK on server – either the SAP JVM or the Sun 1.4.2_27 release

- d. Copy the latest versions of R3LOAD.EXE, DBMSSLIB.DLL and MIGMON.SAR (MIGMON.SAR can be found on the SAP installation master DVD)
- e. Set the system environment variables MSSQL_DBNAME=<SID>, MSSQL_SCHEMA=<sid>, MSSQL_SERVER=<hostname> (or MSSQL_SERVER=<hostname>\<inst> named instance) and dbms_type=mss
- f. If the database logins are required please manually create the users Domain\<sid>adm and Domain\SAPService<SID> and then user the script attached to OSS note [551915](#)
- g. Logon as Domain\<sid>adm and run R3LOAD -testconnect

For creating a R3LOAD server for exporting an Oracle system

- a. Install 9.2 or 10g x64 client for Windows (even if connecting to 32 bit system UNIX)
- b. Download the Oracle R3LOAD and DBSL – unzip and place in a directory such as [C:\Export\Oracle\Kernel](#)
- c. Set the following Environment variables (it might be useful to make a small batch file for this):

```

SET DBMS_TYPE=ora
SET dbs_ora_schema=SAPR3 or <SID>SAP for schema systems
SET dbs_ora_tnsname=<SID>
SET NLS_LANG=AMERICAN_AMERICA.WE8DEC (or UTF8 if Unicode)
SET ORACLE_HOME=C:\oracle\product\10.2.0\Winexport
SET ORACLE_SID=<SID>
SET SAPDATA_HOME=C:\oracle\product\10.2.0\Winexport
SET SAPEXE=C:\Export\Oracle\Kernel
SET SAPLOCALHOST=<set to local hostname>
SET SAPSYSTEMNAME=<SID>
SET TNS_ADMIN= C:\oracle\product\10.2.0\Winexport\network\admin

```

- d. Edit the SQLNET.ORA and TNSNAMES.ORA to resemble the below

```

#####
# Filename.....: sqlnet.ora
# Created.....: created by SAP AG, R/3 Rel. >= 6.10
# Name.....:
# Date.....:
# @(#) $Id: //bc/700-
1_REL/src/ins/SAPINST/impl/tpls/ora/ind/SQLNET.ORA#4 $
#####
AUTOMATIC_IPC = ON
TRACE_LEVEL_CLIENT = OFF
NAMES.DEFAULT_DOMAIN = WORLD
SQLNET.EXPIRE_TIME = 10
SQLNET.AUTHENTICATION_SERVICES = (NTS)
DEFAULT_SDU_SIZE=32768
#####
# Filename.....: tnsnames.ora
# Created.....: created by SAP AG, R/3 Rel. >= 6.10
# @(#) $Id: //bc/700-
1_REL/src/ins/SAPINST/impl/tpls/ora/ind/TNSNAMES.ORA#4 $
#####
<SID>.WORLD=
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS =
        (COMMUNITY = SAP.WORLD)
        (PROTOCOL = TCP)
        (HOST = <hostname goes here>)
        (PORT = 1527) or can be 1521 - check each system
      )
    )
  )
  (CONNECT_DATA =
    (SID = <SID>)
    (GLOBAL_NAME = <SID>.WORLD)
  )

```

-)
)
- e. Edit the hosts file on the UNIX server and enter the Windows R3LOAD server ip address and hostname. On the Windows server edit the hosts file and enter the UNIX server ip address and hostname. Test with PING
 - f. Test the Oracle connectivity with TNSPING <SID>.WORLD.
 - g. Run the script attached to OSS Note [50088](#) (userdomain will usually be the local hostname of the R3LOAD server if the server is not a domain member). This script will create the OPS\$ users that are needed for SAP to login to Oracle. :sqlplus /NOLOG @oradbusr.sql <schemaowner> NT <userdomain>
 - h. Try logging into the Oracle database from the Windows server with the following syntax (for schema systems replace SAPR3 with <SID>SAP): sqlplus SAPR3/sap@<SID>.WORLD
 - i. To ensure correct authorizations try running SELECT * FROM T000;
 - j. Try running R3LOAD -testconnect (remember to set the environment first)

It is recommended to review this blog: <http://blogs.msdn.com/b/saponsqlserver/archive/2011/02/04/installation-amp-configuration-of-windows-nfs-client-to-enable-unix-to-windows-to-mount-a-unix-file-system.aspx>

11. Network Settings

Due to the very high volume of traffic it is recommended to configure 1GB or 10GB Ethernet links between a server running R3LOAD and the SQL Server.

It is recommended to use network card teaming to increase network bandwidth by combining multiple 1GB LAN connections or using 10Gbps NIC.

It is further recommended to configure Jumbo Frames on both the R3LOAD server and the Database Server both during the export and import. Note that the Jumbo Frame size must be configured **identically** on the Database Server, the Switch ports used by both the DB and R3LOAD Server and the NIC card on the R3LOAD Server. The normal value for Jumbo Frames is 9000, though some network devices may only allow 9004. It is essential that this value is the same on all devices or conversions will occur.

If high kernel times are seen on CPU core 0 or 1, please configure RSS options on the NIC cards. Windows 2008 and higher allows for RSS Ring configuration usually up to 8 CPUs on 1Gbps NIC and up to 16 on 10Gbps cards. Perfmon can be used to monitor "Queued DPC" per CPU. This will indicate how many CPUs are being used for Network DPC traffic and how many RSS Rings are configurable. RSS Ring configuration can be changed under the Advance Network Properties for most NIC drivers.

In some cases the network traffic generated from an import will be so great network errors may cause R3LOAD to fail. If this occurs please review Microsoft KB899599.

It is also recommended to review Note [392892](#) and implement <http://support.microsoft.com/kb/948496>. This is required on Win2003. In all cases use Win2008 SP2 or Win2008 R2 if possible.

Note: Network settings are critical for TCPIP based export/imports.

Note: Advance consultants may wish to setup SOFT NUMA on large NUMA based systems. Testing has shown 20-30% performance boost.

<http://msdn.microsoft.com/en-us/library/ms345346.aspx>

<http://blogs.msdn.com/ddperf/archive/2008/09/09/mainstream-numa-and-the-tcp-ip-stack-part-iv-paralleling-tcp-ip.aspx>

12. Disabling or Deleting Secondary Indexes

Disabling secondary indexes can be done and certain long running indexes built online after the system is restarted and validated. To do this remove the Index definition from the STR structure file. After the system is restarted 10-20 indexes can be built online simultaneously. It is recommended to start the ONLINE index build phase prior to users logging onto the system

13. Disable Hyperthreading (on older Xeon)

Hyperthreading should be disabled on both the R3Load server and the SQL Database server

<http://blogs.msdn.com/slavao/archive/2005/11/12/492119.aspx>

UPDATE Please note: **Intel Nehalem** system our recommendation is to use Hyperthreading (leave hyperthreading enabled on Xeon 55xx, 56xx and 75xx CPU)

14. Purge non-critical tables

Most SAP systems have tables that contain unnecessary data. In many cases these tables can be purged:

Purge SW* tables [738148](#)

Purge Basis tables [706478](#)

Purge SW* tables [702356](#)

Note: The OSS Note [706478](#) contains references to many other OSS notes that contain procedures for purging or archiving many "system" type tables. These tables do not contain business transaction data. In all cases please use the SAP documented procedures for purging or archiving these tables.

15. TCPIP Port Export/Import Procedure

TCPIP Port based export to a SQL Server system is fully supported. In general we recommend this method for advanced migration consultants only.

In such an export procedure R3LOAD will communicate directly with the R3LOAD process on the target server. No dump files will be created as all data is passed via TCPIP. A socket export/import reduces the R3LOAD CPU consumption and may allow slow legacy servers to run a larger total number of R3LOAD processes.

It is not possible to use TCPIP Port based migration procedure when converting from non-Unicode to Unicode.

It is possible to migrate a Unicode SAP system running on an Oracle database to a Unicode SAP system running on SQL Server (even if the source system is running on a big Endian 4102 platform and SQL Server is on a little Endian 4103 platform)

Note: a socket export the OrderBy parameter on the import server must not be set or the import will crash with a Java error (import order is set by the export server).

16. BW Specific Recommendations

Note 991014 discusses Var Decimal storage format. This is particularly important for BW. It is generally recommended to use SQL 2008 compression technologies rather than SQL 2005 Var Decimal.

It is **critical** to fully review

BI 7.0/7.1 [888210](#), [1444413](#) and [1460372](#) (for APO systems [1471193](#))

BW 3.5 [771209](#)

BW 3.1 [777024](#)

before proceeding with a system copy of BW.

Oracle systems with more than 1000 partitions please use SQL 2008 Service Pack 2 or SQL 2008 R2 Service Pack 1. See note [1494789](#) - Enabling 1000+ partitions support on SQL Server

Reference 1010854, 965695, 962124, 961591

This SQL Server script can be used to check the number of table partitions. Review note [1010854](#) when more 700 partitions are observed.

```
select COUNT(partition_id),object_name(object_id),index_id
from sys.partitions
where OBJECTPROPERTY(object_id,'IsUserTable')=1
group by object_id, index_id
order by 2,3 asc
```

Note: Please review this document if a large number of partitions are found -

[http://download.microsoft.com/download/4/C/6/4C611C7F-0906-4604-8DC0-E7076CE2EC5D/SAP_BI_Table_Partitioning II.pdf](http://download.microsoft.com/download/4/C/6/4C611C7F-0906-4604-8DC0-E7076CE2EC5D/SAP_BI_Table_Partitioning_II.pdf)

To check in Oracle:

You can use the following query on your ORACLE database to check in sqlplus if tables with more than 999 partitions exist:

```
select table_name from user_part_tables where partition_count >= 999 and table_name
like '/%';
```

The following two notes are needed when importing onto SQL 2008

Note [1157904](#) and Note [1364683](#)

To repartition systems follow note [1471910](#)

17. Unicode Conversion Specific Recommendation

Please see notes on Unicode conversion, restrictions on unsorted export and socket export. New versions of R3LOAD will export cluster tables sorted always.

OSS Note [1139642](#) quotes DB growth of 60% when moving to Unicode on SQL Server. This figure has never been observed on any real customer system. Even without using SQL Server 2008/R2 compression technologies only a small increase has been observed. The estimate of 60% is factually wrong.

18. SQL 2008 ROW Compression

NW 7.0 SR3 and above installations and migrations can load data in ROW compressed format. To enable this functionality a new R3LOAD must be used. Please obtain the latest R3LOAD and DBSL. For all SQL 2008 migrations it is recommended to import in ROW compressed format. See note [1143246](#)

19. SQL 2008 PAGE Compression

SQL Server includes powerful and efficient DB compression technologies. Please see blogs on <http://blogs.msdn.com/b/saponsqlserver/> for further information.

The compression of secondary indexes is supported as of September 2010. See note [1488135](#) - Database compression for SQL Server. Also see Note SAP Note [1505884](#) for information on the latest R3LOAD versions

To force PAGE compression edit the DDLMSS.TPL file to read

```
cretab: CREATE TABLE &tab_name&  
        ( /{ &fld_name& &fld_desc& /-, /} ) &compression_page&
```

```
drptab: DROP TABLE &tab_name&
```

```
crepky: ALTER TABLE &tab_name&  
        ADD CONSTRAINT &pri_key&  
        PRIMARY KEY &clustered& ( /{ &key_fld& /-, /} ) &compression_page&
```

```
drppky: ALTER TABLE &tab_name&  
        DROP CONSTRAINT &pri_key&
```

```
creind: CREATE &unique& INDEX &ind_name&  
        ON &tab_name&  
        ( /{ &fld_name& /-, /} ) &compression_page&
```

This procedure is fully supported for all customer systems as of Sept 2010.

The TPL file that needs to be modified is in the `c:\program files\SAPInst` directory (do a search of this directory) or the TPL file that is referenced in the package import log

SAP also fully support compression of Secondary Indexes. See Note [1459005](#) - Enabling index compression for SQL Server

During the Pilot phase of Table & Index compression customers with typical SAP on Oracle system (Unicode or Non-Unicode) running on SQL 2008 R2 (which offers Unicode compression) and secondary index compression will compress to about 20% of original size. This means a 10TB Oracle system will become ~2TB on SQL Server. This allows customers running on multi-million dollar UNIX hardware to migrate to inexpensive highly performing Intel hardware easily. Testing around performance has shown very significant improvements in performance on servers such as the HP DL580 G7 compared to UNIX servers that cost ~\$3m USD.

To check the compression properties of a particular table run the following in SQL Management Studio

```
select OBJECT_NAME(object_id), index_id, data_compression, data_compression_desc from  
sys.partitions where object_id = OBJECT_ID('<TABLENAME>');
```


20. Overview of SAP Tools & Releases

Every version of including SAP 4.6C and higher can run on SQL 2005. Some older releases of SAP R/3 do **not** have R3SETUP/SAPInst for SQL 2005 – a major source of confusion.

Basis Release	3.1I	4.0B	4.5B	4.6x	6.20	6.40	7.00	and higher
R3SETUP	4.6D	4.6D	4.6D	4.6D	-	-	-	
SAPINST	-	-	-	-	6.40	6.40	7.00	
R3LDCTL	3.1I	4.0B	4.5B	4.6D	6.40	6.40	7.00	
R3SZCHK	-	-	4.5B	4.6D	6.40	6.40	7.00	
R3LOAD	3.1I	4.0B	4.5B	4.6D	6.40	6.40	7.00	
MIGMON	yes	yes	yes	yes	yes	yes	yes	
R3TA	no	no	no	no	yes	yes	yes	
DISTMON	no	no	no	no	yes	yes	yes	
Package Splitter:								
-Perl (don't use)	no	yes	yes	yes	yes	yes	yes	
-Java	yes	yes	yes	yes	yes	yes	yes	
SQL 2000	no	no	yes	yes	yes	yes	yes	(not recommended)
SQL 2005 SP2	no	no	no	yes	yes	yes	yes	
SQL 2008*	no	no	no	no	no	no	yes	(64 bit only)
SQL 2008 R2**	no	no	no	no	no	no	yes	(64 bit only)
SQL 11***	no	no	no	no	no	no	yes	(64 bit only)
Win 2003 SP2 x64	no	no	yes	yes	yes	yes	yes	(not recommended)
Win 2008	no	no	no	no	no	no	yes	(64 bit only)
Win 2008 R2**	no	no	no	no	no	no	yes	(64 bit only)

*at least SQL 2008 SP1 (or CU2) must be used **SQL 2008 R2 supported 09/2010 and Win 2008 R2 supported as of November 2010 ***SQL 11 in beta testing as at Dec 2010.
64bit means both x64 and IA64 (Itanium)

21. Common Problems & Errors

The system copy procedure must be followed exactly or some of the errors below may occur.

- ERROR: ExeFastLoad: rc = 2**
Please review note [942540](#). It is probable that the DFACT.SQL file has not been generated by the SMIGR_CREATE_DDL report or the file is not in the <export dir>DBMSS directory. If the problem continues try setting the NO_BCP=1 to stop FASTLOAD. This will allow R3LOAD to output a more specific error message. Also check the SQL Server Error Log.
- SQL Stack Dump LATCH TIMEOUT**
It is likely that the SAPDATAx files or SAPLOG1 file was not created large enough and SQL Server has tried to extend this file. Under extremely heavy load this error may be seen. Expand the database to the expected final size prior to beginning the import. Ensure the log file is at least 10GB for larger systems.
- Deadlock error in package log file**
If message : Transaction was deadlocked on lock resources with another process and has been chosen as the deadlock victim. This message is extremely rare but can theoretically occur on large split tables. In the majority of cases the fastest resolution will be to drop the table and reset the status of the TSK files and import all packages of the split table again.

```
(IMP) INFO: EndFastLoad failed with <2: Bulk-copy commit unsuccessful:[208]
Invalid object name '<sid>.MSSDEADLCK'.
[1205] Transaction (Process ID xxx) was deadlocked on lock resources with
another process and has been chosen as the deadlock victim. Rerun the
transaction.
[208] Invalid object nam>
(IMP) ERROR: EndFastload: rc = 2
```

If you are using table splitting you may wish to reduce the default BCP batch size from 10000 to 5000. Note that larger batch sizes produce faster import times. For non-split tables you may wish to consider increasing the batch size after careful testing. The MIGMON/R3LOAD parameter BCP_BATCH_SIZE=50000 or set the environment variable set BCP_BATCH_SIZE=50000.

Note: SQL 2008 R2 contains an enhancement that will greatly reduce the occurrence of this problem. It is expected that the batchsize can be increased significantly on SQL 2008 R2. Testing with values of 100,000 have shown good results. It is recommended to monitor memory consumption with large batch sizes

- d. Various Error Messages Installing 4.6x/4.7 on x64 Servers
All SAP systems should run on 64bit versions of Windows and SQL Server. The installation routines for 4.6C/D were created before x64 versions of Windows were available. Because of this there may be errors. Please review Notes [862789](#), [960769](#), [899111](#) & [814834](#).
- e. BW/BI or other Netweaver Systems
BW or Netweaver systems have special restrictions and special requirements.
In general we recommend that for BW or Netweaver systems the following support packs are implemented:

Netweaver 7.0 BI 7.0

Support Pack Stack 9

Review Notes: 888210, 962124, 1010854

Netweaver '04 BW 3.5

Support Pack Stack 18

Review Notes: 771209, 996263, 960504, 1010854

SAP BW 3.0b/3.1c

SAP_BASIS >=58

BW 3.0B/3.1C >=25/19

Review Notes: 777024

OS/DB Migrations on support packs lower than the above recommendations are still possible and carefully reviewing all related notes and implementing these corrections in the source system. In some cases new functionality such as table partitioning will not be activated unless certain prerequisites are met.

In all cases customers must carefully implement note [875986](#) before implementing any notes with SNOTE.

- f. 4.6C Error (BEK) ERROR: SlicGetInstallationNo() failed
The system environment variable SAPSYSTEMNAME = <SID> is not set. Set this variable for the user <sid>adm
- g. 4.6C error in dev_w* - Long Datatype Conversion not performed" please see Note [126973](#) - SICK messages with MS SQL Server
- h. R3SETUP and possibly very old SAPInst may attempt to create a SAP database with code page 850BIN prior to the import of the dump files. Note [799058](#) and [600027](#) strictly forbid the use of code page 850BIN and require conversion to **850BIN2**.
Also note that the utility for converting codepage 850BIN to 850BIN2 does **not** work on SQL 2005 or higher (the fast conversion feature was dropped from SQL 2005). Therefore care should be taken to avoid the case where R3SETUP creates a 850BIN database on SQL 2005 and then MIGMON is used to import the system into this database. Clearly this will result in an unsupported system running code page 850BIN on SQL 2005. Conversion will be impossible and the import will need to be repeated after dropping and then manually creating the database.

The following commands display the server (default) and database collations:

```
SELECT SERVERPROPERTY('Collation')
SQL_Latin1_General_CP850_BIN2
SELECT DATABASEPROPERTYEX('<SID>', 'Collation')
SQL_Latin1_General_CP850_BIN2
```

- i. ABAP Shortdump & SM21 error max. marker count = 2090
>B *** ERROR => dbtran ERROR (set_input_da_spec): statement too big
> marker count = 2576 > max. marker count = 2090

This is because the limit on the number of parameters on a stored procedure is 2100 on SQL. It is higher on other databases

<http://technet.microsoft.com/en-us/library/ms191132.aspx>

- j. In very rare cases a JOIN on Oracle may not work on SQL Server. This can happen on systems such as CRM where GUIDs are stored in RAW datatypes and a JOIN is attempted on a CHAR datatype. Please review Note [1294101](#)
- k. A simple and easy way to suspend and release all batch jobs on a system is to run these reports in SE38
Suspend: BTCTRNS1
Release: BTCTRNS2

22. Troubleshooting Tips

- a. R3LOAD Connection Problems
Review note [98678](#). The system environment variable MSSQL_DBSLPROFILE=1 will write a trace file dbsl_w0 to the current directory. This file will become very large and seriously reduce the performance of a system. In some cases it may be necessary to set the SAPSYSTEMNAME=<SID> system environment variable.
- b. R3LOAD Cannot Find DFACT.SQL, STR or Dumpfiles
The system environment variable R3LOAD_WL=1 will output extra information in the <package>.LOG file
- c. Scan log files with Windows FINDSTR (Windows version of grep)
The command line below will output all the error lines from the export or import directory
Findstr /C:ERROR: <path to log files>*.log
- d. A large number of R3LOAD processes are configured and Oracle issues this error

The system Error message returned by DbSl:

```
ORA-00018: maximum number of sessions exceeded
(DB) INFO: disconnected from DB
```

Solution:

Increase the parameters in

```
unix: $ORACLE_HOME/dbs/init<DBSID>.ora
windows:$ORACLE_HOME/database/<initDBSID>.ora
PROCESSES=1000
SESSIONS=1105
```

- e. Sorting some BW or other large tables can consume massive amounts of PSAPTEMP. If this occurs there are two options: (1) switch to Unsorted export (see earlier section in this document or (2) run the commands below to increase PSAPTEMP

```
(EXP) ERROR: DbSlExeRead failed
rc = 99, table "/BIC/B0000585000"
(SQL error 1652)
error message returned by DbSl:
ORA-01652: unable to extend temp segment by 128 in tablespace PSAPTEMP
(DB) INFO: disconnected from DB
```

```
Sqlplus /nolog
Connect / as sysdba
SQL> ALTER TABLESPACE PSAPTEMP ADD TEMPFILE
'E:\oracle\BWP\sapdata1\temp_1\TEMP.DATA2' SIZE 20000M
SELECT * FROM V$TEMP_SPACE_HEADER;
```

- f. FASTLOAD Errors
The system environment variable NO_BCP=1 will override the -loadprocedure -fast option and force R3LOAD to use the normal DBSL interface for import
- g. Special characters are corrupted
Please review this note [1279882](#)
- h. To Enable fastload on LOB columns in 6.40 & 7.00 set BCP_LOB=1 and review note [1156361](#)
- i. If this error occurs during a MDMP Unicode conversion review [992956](#)
(DB) INFO: UMGPMDDII~WRD created

```
(DB) INFO: UMGPMDDIT created
(DB) INFO: UMGPMDDIT~0 created
(IMP) INFO: ExeFastLoad failed with <2: BCP Commit failed:[2627] Violation of
PRIMARY KEY constraint 'UMGPMDDIT~0'. Cannot insert duplicate key in object
'dbo.UMGPMDDIT'.
[3621] The statement has been terminated.>
(IMP) ERROR: ExeFastload: rc = 2
(DB) INFO: disconnected from DB
```

- j. ASSERTION_FAILED during generation of DFACT.SQL. Please cross reference [984396](#) first. If this is unsuccessful please run RSDDS_CHANGERUN_TMPTABLS_DEL
- k. Very slow import times on systems with a large amount of RAM
SQL 2005 may be very slow to import during the data load phase if inserts are overflowing” the page flush buffer. This is because the overflow is stored in the SQL datacache. The process to store and retrieve overflow rows is single threaded. The larger the amount of memory available for SQL Server data cache the worse the problem becomes. **This problem is fixed in SQL 2008 and page flush manager queue has been increased to 1024 pages (under trace flag 3917 it is 8192 pages)**. It is recommended to contact Microsoft support if this problem is encountered on SQL 2005. If a customer is using a SAP Kernel 7.0 or higher component it is strongly recommended to use SQL 2008 or higher
- l. Below is a useful script to run if an Import fails and the entire SAP database needs to be purged of all tables. Thanks to Amit for providing this. **WARNING: Running this script will drop all tables in the current database**

```
Use <SID>;
Exec sp_MSforeachtable 'drop table ?';
```

- m. Towards the end of an import there may be many “suspended” SQL processes. These can be viewed with SQL Management Studio Activity Monitor. Clicking on the suspended process may show that a process is performing a CREATE INDEX. Towards the end of an import most of the table data import is complete and SQL Server will be building secondary indexes. The primary clustered index has been built simultaneously as the table data is loaded. Often these secondary indexes are non-standard Z indexes or sometimes unused SAP standard indexes. These indexes may be deleted in the source system before export or created after the system has been restarted and the downtime period is over. SQL 2005 and higher supports online index creation. The memory consumption during index creation can be substantial, especially if many indexes are being built simultaneously. This script is useful to detect situations when SQL is suspending index creation due to insufficient memory

```
-- current memory grants per query/session
select
session_id, request_time, grant_time ,
requested_memory_kb / (1024.0 * 1024 ) as requested_memory_gb ,
granted_memory_kb / ( 1024.0 * 1024 ) as used_memory_gb ,
used_memory_kb / ( 1024.0 * 1024 ) as used_memory_gb ,
st.text
from
sys.dm_exec_query_memory_grants g cross apply
sys.dm_exec_sql_text(sql_handle) as st
-- uncomment the where conditions as needed
-- where grant_time is not null -- these sessions are using memory
allocations
-- where grant_time is null -- these sessions are waiting for memory
allocations

-- overall server status

select * from sys.dm_exec_query_resource_semaphores
```

If many R3LOAD BCP or CREATE INDEX Processes are in status SUSPENDED with RESOURCE_SEMAPHORE wait type in the DMV below:

```
select session_id, request_id, start_time, status ,
command, wait_type, wait_type from sys.dm_exec_requests where session_id >49;
```

If this is the case, it may be useful to cap the amount of memory that a particular secondary index build task can consume. This will force the Secondary Index Build to use TEMPDB. The way to cap memory is to Active Resource Governor (by right clicking on it in SSMS). Adjust the memory percentage value as needed. By default SQL Server can easily consume 10-40GB RAM per Index Build if no limit is set – the actual value depends on the amount of RAM in the server. This substantially improves Index build speed, however if too many secondary indexes are built at one time this will consume all available memory, thereby blocking other resources. It is recommended to monitor TempDB utilization when setting this option

```
USE master;
BEGIN TRAN;
-- Create 1 workload group for SAP R3Load
-- Workload group is getting assigned to default pool automatically
CREATE WORKLOAD GROUP R3load;
GO
COMMIT TRAN;
go
-- Create a classification function.
CREATE FUNCTION dbo.classify_r3load() RETURNS sysname
WITH SCHEMABINDING AS
BEGIN
    DECLARE @grp_name sysname
    IF (APP_NAME() LIKE 'R3 00%')
        SET @grp_name = 'R3load'
    RETURN @grp_name
END;
GO
-- Register the classifier function with Resource Governor
ALTER RESOURCE GOVERNOR WITH (CLASSIFIER_FUNCTION= dbo.classify_r3load);
GO
--change maximum memory grant a query can get. Default = 25%
ALTER WORKLOAD GROUP R3load with (REQUEST_MAX_MEMORY_GRANT_PERCENT=5);
go
-- Start Resource Governor
ALTER RESOURCE GOVERNOR RECONFIGURE;
GO
```

- n. To transfer all objects from the dbo schema (or any other schema) into the <sid> schema run the scripts attached to OSS Note 551915 and then copy the output of this script and paste into a new query window and execute:

```
DECLARE @sourceschema VARCHAR(100)
DECLARE @destinationschema VARCHAR(100)
SET @sourceschema = 'sourceschema'
SET @destinationschema = 'destinationschema'
SELECT 'ALTER SCHEMA ' + @destinationschema + ' TRANSFER ' + b.name + '.' + '
[' + a.name + ']'
FROM sys.objects a JOIN sys.schemas b ON a.schema_id=b.schema_id WHERE a.type
IN ('U','P','FN','V')
AND b.name=@sourceschema;
```

After running this script, re-run MSSPROCS in SE38

23. Migration for 4.6C or lower based systems : High level process:

- Raise an OSS message requesting a copy of the 4.6D SAP R3SETUP. (R3SETUP is no longer available for download)
- Prepare system according to 4.6D system copy guide

- c. Install R3SETUP on the source system and update the DBMSSLIB.DLL, R3LOAD.EXE & R3SZCHK.EXE
- d. Modify R3SETUP DBEXPORT.R3S to force R3SETUP to exit just before starting the export

```

<xx>=R3SZCHK_IND_IND
<xx>=DBEXPCOPYEXTFILES_NT_IND
<xx>=DBR3LOADEXECDUMMY_IND_IND    ***delete***
<xx>=CUSTOMER_EXIT_FOR_EXPORT    ***add***
<xx>=DBEXPR3LOADEXEC_NT_IND      ***delete***
<xx>=DBGETDATABASESIZE_IND_IND

[CUSTOMER_EXIT_FOR_EXPORT]        ***add***
CLASS=CExitStep                   ***add***
EXIT=YES                           ***add***

```

- e. Run R3SETUP and open DBEXPORT.R3S. Do **not** select the Perl based package splitter. Exit at the customer stop point
- f. Copy the Java based splitter to the R3SETUP install directory. Copy *.EXT and *.STR files from <export dir>DATA to the installation directory. Configure and run the Java based package splitter tool. The package splitter will process the EXT and STR files and rename them to *.OLD and create new EXT and STR files.
- g. Copy Migration Monitor to the installation directory and run Migration Monitor to export the system
- h. R3SETUP and open DBEXPORT.R3S to continue the export steps. These steps will generate the DBSIZE.TPL
- i. Run Migration Time Analyzer to check which packages run the longest. Try to optimize the export by starting these packages first using the OrderBy.txt file
- j. Start a CMD.EXE session from the \Windows\syswow64 directory and run SETUP.BAT to install R3SETUP on target server. Immediately update the DBMSSLIB.DLL and R3LOAD.EXE
- k. Modify DBMIG.R3S with exit point

```

190=DBDBSLTESTCONNECT_NT_IND
200=MIGRATIONKEY_IND_IND
<xx>=CUSTOMER_EXIT_FOR_IMPORT      ***add***
210=DBR3LOADEXECDUMMY_IND_IND      ***delete***
220=DBR3LOADEXEC_NT_MSS             ***delete***
230=DBR3LOADVIEWDUMMY_IND_IND      ***delete***
240=DBR3LOADVIEW_NT_IND            ***delete***
250=DBPOSTLOAD_NT_MSS
260=DBCONFIGURATION_NT_MSS

[CUSTOMER_EXIT_FOR_IMPORT]        ***add***
CLASS=CExitStep                   ***add***
EXIT=YES                           ***add***

```

- l. Run R3SETUP and open DBMIG.R3S. Exit at the customer stop point
- m. Copy the <export dir> to the target system and run Migration Monitor to import the system
- n. Run R3SETUP to continue the installation. If R3SETUP fails review note [965145](#)
- o. Run Migration Time Analyzer and review OrderBy.txt
- p. Perform the post system copy steps as per the 4.6D system copy guide