

# 1616401 - Parallelism in the Upgrades, EhPs and Support Packages implementations

<b>Version</b>	21	<b>Type</b>	SAP Knowledge Base Article
<b>Language</b>	English	<b>Master Language</b>	English
<b>Release Status</b>	Released to Customer	<b>Category</b>	How To
<b>Component</b>	BC-UPG-TLS-TLA (Upgrade tools for ABAP)	<b>Released On</b>	02.07.2020

Please find the original document at <https://launchpad.support.sap.com/#/notes/1616401>

## Symptom

Setting up parallel processes dramatically increase the performance of Upgrades and Updates in SAP systems.

This article intends to explain the parallelism import mechanisms and how it works in such procedures.

Here, "parallelism" means:

- 1) the import of more than one package or more than one object of a package at the same time by the transport tools (tp and R3trans);
- 2) the execution of more than one background jobs on the batch phases (both dialog and batch);
- 3) the execution of more than one SQL process on the database level;
- 4) the execution of more than one R3load processes on the EU\_IMPORT phase (loading data from DVDs in release upgrades) and on the migration of data in DMO scenarios (data migration option);
- 5) the execution of more than one phase at same time, during the UPTIME portion of the upgrade.

The information is intended mostly to the upgrade SUM tool (Software Update Manager), but also covers SPAM/SAINT.

## Environment

- SAP Netweaver 7.00 (and higher)
- SUM (Software Update Manager)
- SAINT/SPAM

## Cause

### SUM

When executing the Software Update Manager tool, you're prompted to enter numerical values for several parameters in phase INITSUBST:

Enter the maximum number of ABAP processes (dialogue and batch) during the update:

ABAP PROCESSES (UPTIME)	3
ABAP PROCESSES (DOWNTIME)	4

Enter the maximum number of parallel processes for execution of SQL commands

SQL PROCESSES (UPTIME)	3
SQL PROCESSES (DOWNTIME)	6

Enter the maximum number of parallel import processes:

R3TRANS PROCESSES (UPTIME)	3
R3TRANS PROCESSES (DOWNTIME)	6

Enter the maximum number of parallel R3load or table comparison processes:

R3LOAD PROCESSES (UPTIME)	8
R3LOAD PROCESSES (DOWNTIME)	40

The first parameter, **ABAP PROCESSES**, sets up the number of processes used in **job** phases. This influences all job-run phases, including long runtime ones like DBCLONE (in update scenarios), ACT\_UPG and PARDIST in uptime, and XPRAs in downtime.

During **uptime**, you cannot set up more batch processes than the system has available (seen in SM50 or RZ04), on your source system. If you surpass the number of batch jobs available, the batch system will serialize them, resulting in no performance gain.

For DBCLONE phase there is technical limitation of 10 parallel processes at the most, so a higher number will not help. For ACT\_UPG, the limit is 6 parallel processes, but this phase was benefited from code optimization in SUM SP12, improving dramatically its runtime. For PARDIST phases, 10 processes is the technical limit also.

**ATTENTION!** the parameter `rdisp/wp_no_btc` (Number of work processes of type batch - background processing) holds back the number of parallel processes! Please ensure you have at least an equal number of background processes available, otherwise parallelism will not be obtained. You might also notice the following warning on the logs:

**"2WETQ399 Cannot run 10 batch processes, system has only 1!"**

The second parameter, **SQL PROCESSES**, defines the number of parallel SQL statements to be run on the database at the same time. This setting is specially useful for the DDL-phases (data definition language), like MVNTAB and PARCONV, where the "create table", "alter table", "create index" and other commands take place. Also the phases that create aliases/synonyms of the shadow system use this parameter (SCEXEC\_GRANT + SCEXEC\_ALIAS), but to a less extend. The recommended value is the total number of physical CPUs (using the same logic of R3TRANS PROCESSES below).

The third parameter, **R3TRANS PROCESSES**, is used for tp and R3trans transport tools in parallel, on the phases where import of support packages take place (DDIC\_UPG, SHADOW\_IMPORT\* and TABIM\_UPG). This type of import can be performed by the transport tools since SAP\_BASIS release 700, according SAP Note [1127194](#) "**R3trans import with parallel processes**". This is the setting that brings the most benefits in terms of overall performance in upgrades because of the number and size of the packages imported at different stages. The recommended value is the same number of physical CPUs available on the server, to follow the "1 parallel import process per CPU". You can view the number of CPUs available in transaction

ST06, on "CPU count". If your CPU is a multi-core one, SAP System will consider them as a physical, available CPU. For example: if you have 4 quad-core CPUs, they are seen by SAP System as 16 CPUs in ST06. On this case, you can enter '16' for both UPTIME and DOWNTIME R3TRANS PROCESSES fields (remember that ordinary transport requests import might still take place during uptime, and they share the same tp/R3trans resources. So avoid importing requests during upgrades).

During such phases, the number of parallel processes on the OS will vary greatly, due to the many packages/objects reading and writing into the database. Please refer to the technical details explained on the "Resolution" part below.

Depending of the upgrade strategy, the phases can be run during the uptime or in downtime portion:

PHASES in SUM tool	DOWNTIME-MINIMIZED (Standard or Advanced mode)	RESOURCE-MINIMIZED (Single System mode)
SHADOW_IMPORT_UPG1	uptime	downtime
DDIC_UPG	uptime	downtime
SHADOW_IMPORT_UPG2	uptime	downtime
SHADOW_IMPORT_INC	uptime	downtime
TABIM_UPG	downtime	downtime

Please take into consideration the database parallelism information below before setting up a definitive number for **R3TRANS PROCESSES**!

The fourth parameter, **R3LOAD PROCESSES**, tells the upgrade tool how many R3load processes will be triggered during phases **EU\_IMPORT\*** (in system releases upgrade) and in **DMO (data migration option)** scenarios, to move the data from the old database into the new HANA one. Since R3load utilizes small CPU processing, you can safely enter a significant number for this parameter: 3 to 5 times the number of CPUs, being "5" the theoretical top limit.

### The SUM Guide

The Upgrade Guide also provides information about parallelism. Specially in chapter **A.3.2: Configuring Parallel Processes During the Runtime**, where it instructs how to modify the number of processes after they were entered in phase INITSUBST:

This procedure is also described in KBA #[1904239](#): **SUM options to change parameters entered in INPUT and INITSUBST phases**, and can be used in case you need to modify the number of processes **before** starting the phase you want to have the changes in effect.

( Refer to the relevant [SUM Guide](#) for specific commands )

### SPAM and SAINT

Parallelism can be obtained on the ABAP tools Support Package Manager and Add-On Installation tool. Please find below a comparison between the tools:

PHASES (SUM)	PHASES (SAINT/SPAM)	PARALLELISM?
ACT_UPG	DD_ACTIVATION	BATCH
DDIC_UPG	DD_IMPORT	R3TRANS
PARCONV_UPG	MOVE NAMETAB/CONVERSION	DDL
PARDIST_SHD	DISTRIBUTION	BATCH
SHADOW_IMPORT_INC		R3TRANS
SHADOW_IMPORT_UPG1		R3TRANS
SHADOW_IMPORT_UPG2		R3TRANS
TABIM_UPG	MAIN_IMPORT	R3TRANS
XPRAS_AIMMRG	XPRAS	BATCH

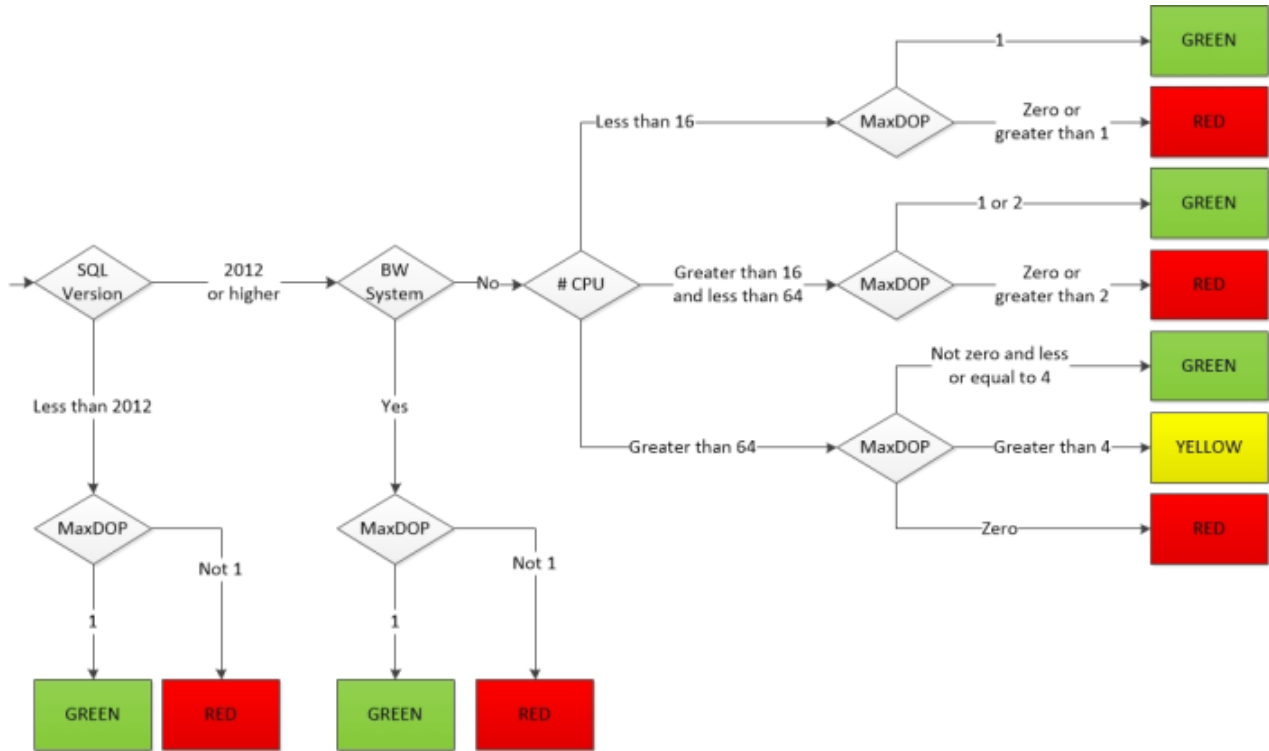
In some occasions parallelism might not work in SPAM and SAINT, even if the corresponding settings are correctly configured and the transport tools are up-to-date. On such cases please check of the **ACC\_IMPORT** parameter is set in the transport profile. This parameter deactivates the parallelism for SPAM/SAINT imports!

The parameter ACC\_IMPORT is useful for large number of small transports. In this mode, tp and R3trans process one import after another, using interprocess communication to avoid starting/stopping of processes. For large transports (like Support Packages), this parameter is counterproductive, because it forbids parallelism. Please refer to note #[1223360](#) for more information about it.

### Additional Database-specific information

Each database has its own internal parallel mechanics, which work apart from the SAP Upgrade tools. You have to consider the resources usage on the database side alongside the upgrade tool parallel operations to avoid exhausting the Database, **specially in case you run the database on the same server of the Primary Application Server (PAS, the old Central Instance Server)**! In case of questions please contact your database vendor directly.

1. **MSSQL**: please refer to the following article: [Max degree of parallelism \(MaxDOP\)](#). You have to set up the MaxDOP setting accordingly, paying attention to not go overboard. The following picture illustrates the configuration quite simple:



2. **Oracle:** there are no special instructions other than having the parameters correctly configured according to note #[1888485](#) (Oracle 12), #[1431798](#) (Oracle 11) and #[724716](#) (Oracle 10). An Oracle parallelism FAQ do exist in SAP note #[651060](#).
3. **Sybase (ASE):** Sybase's Adaptive Server Enterprise has its own Upgrade parameters, available as a PDF attachment in note #[1680803](#). The section "IMPORT" is of particular interest since it deals with the DDL statements on the database (page 28. Also page 49, on the section "Reconfigure Engines and Parallel Processing"). The normal operations parameters can be found in note #[1539124](#). Also please refer to Sybase's online DB manual here, specially to chapter 7 "Parallel Query Processing".
4. **HANA:** HANA has a parallelism mechanism designed by concept, so there is not special adjustments **for Upgrades**. [This document](#) provides a good higher level overview of its mechanics. Anyway you can tweak parallel workers as per notes #[2100040](#) and #[2000003](#) for daily use.
5. **MaxDB:** no special instructions related directly to upgrade. As a general performance recommendation, please refer to the [MaxDB's Database Parameter page](#) (about parameter MaxUserTasks, MaxCPUs and MaxUserTasks) and to the following notes:
  - #[814704](#): MaxDB Version 7.6 parameter settings for OLTP/BW
  - #[1004886](#): SAP MaxDB 7.7 database parameter recommendations
  - #[1308217](#): SAP MaxDB 7.8 database parameter recommendations
  - #[1346964](#): SAP MaxDB 7.9 database parameter recommendations
6. **DB2 for i:** please refer to the white paper which is available on IBM's techdocs library: [IBM Techdocs White Paper: SAP Software Update Manager: Behaviors and Insights](#). It contains optimizations and platform specific advises for different stages of the upgrade. This document has been completely overworked to cover aspects of SAP's current standard upgrade tool SUM on platform IBM i.

## Resolution

Technical aspects of tp/R3trans parallelism in SUM

The SUM tool takes into account the information entered in field "**R3TRANS PROCESSES**" and optimizes

the import stages according to the number of packages to be imported and according to their ordenance on the import buffer file, where they are split by the **SYNCKMARK** entries.

SUM then calculates the number of parallel **tp** and **R3trans** processes for each group, where each tp will read the packages, one by one, and each R3trans will write objects from that package into the database, also one by one. Depending on the distribution of the packages, the tool will balance the number of parallel **tps** and **R3trans**, looking for the most optimal import procedure for that particular batch of packages.

**SYNCKMARK:** are entries present on the system buffer, that divide the number of packages that can be imported at same time by the upgrade/update tools. These entries are only needed when TP is allowed to import several transport requests (or packages) in parallel, and it's used to avoid that two transport requests that contain the same object are imported in parallel. So only the group of transport request not separated by a "SYNCKMARK" will be imported in parallel. For example: a TP process starts new parallel R3trans jobs for every transport request in the buffer up to a SYNCKMARK entry; at this point, TP will wait until the last of the started R3trans jobs has finished before starting the next R3trans (for a transport request after the current SYNCKMARK).

Please consider the example below:

#### •Analyze buffer structure in blocks

```

/0/SYNCKMARK
/0/<TR1>
/0/<TR2>
/0/<TR3>
/0/<TR4>
/0/<TR5>
/0/<TR6>
/0/SYNCKMARK
/0/<TR7>
/0/<TR8>
/0/SYNCKMARK
/0/<TR9>
/0/SYNCKMARK
/0/<TR10>
/0/SYNCKMARK
/0/<TR11>
/0/SYNCKMARK
  
```

} mainimp\_proc = 6  
 } parallel = 1  
  
 } mainimp\_proc = 2  
 } parallel = 3  
  
 } mainimp\_proc = 1  
 } parallel = 6

#### 1. Set stopmark at <TR6>

```
tp put <SID> -Dmainimp_proc=6 -Dparallel=no
```

#### 2. Set stopmark at <TR8>

```
tp put <SID> -Dmainimp_proc=2 -Dparallel=3
```

#### 3. Set stopmark at <TR11>

```
tp put <SID> -Dmainimp_proc=1 -Dparallel=6
```

On this example, the import buffer has divided the packages into groups that vary from 1 to 6 in between SYNCKMARK entries. For this example a value of 6 was given on the field "**R3TRANS PROCESSES**", so SUM will calculate dynamically the best combination for each block. On the first block, composed of 6 packages, SUM set the parameter `mainimp_proc` to 6 and `parallel` to 1, which means 6 tp processes and 1 R3trans process (probably because the packages are small and contain a reduced number of objects. So each tp + R3trans import will finish quickly and there is no sense in having more R3trans working in parallel. The overall execution is optimized by having a higher number of packages processed at same time, each one by one tp process).

Now for the second group, there are 2 packages in between SYNCKMARK entries. SUM sets `mainimp_proc` to 2 and `parallel` to 3, meaning 2 tps processe at the most, at same time, and up to 3 R3trans processe per tp. Probably those 2 packages are bigger and contain more objects, so having a high number of R3trans processes speed up the processing of each package.

For the last 3 groups, which are made of only 1 package in between SYNCKMARK entries, there is no need to have more than 1 tp, so all resources are allocated to R3trans via the "parallel" parameter.

The value of parameter `mainimp_proc` sets the number of parallel tps, responsible for handling the packages.

The value of parameter **parallel** sets the number of parallel R3trans, responsible for handling the objects inside each package

Now what if the size of the packages vary in each block? The below example illustrate what happens (using the first block of packages from the example above, but) with a bigger package 1 (represented below by <TR1>) and smaller packages from 2 - 6:

(the horizontal green bar means the time to import each package)



Can be done equally well as:



Should be faster by a factor of 3 due to parameter „parallel“!



On this case, SUM could have improved the import of this groups of packages by setting up 2 tp parallel process (mainimp\_proc = 2) and using up to 3 parallel R3trans process (parallel = 3), speeding up the total import time of the block by a factor of 3.

**IMPORTANT INFORMATION:** The formula that SUM uses to calculate the parallel process is the following:  
 $nrosprocs = 1 + mainimp\_proc + (mainimp\_proc * parallel)$  where nrosprocs --> number of total parallel process running on the OS level (tps + R3trans) mainimp\_proc --> number of tps running in parallel process on the OS level parallel --> number of R3trans running in parallel process on the OS level

As per the formula, there is always 1 tp process starting to process the buffer. It then starts the "mainimp\_proc" processes (depending on each block of packages), which, by their turn, starts the "parallel" R3trans managed processes.

This also leads, in some cases, to a slightly different total number of parallel processes, as seen on the below example:

(example taken from the TABIM\_UPG phase, where "**R3TRANS PROCESSES**" was set to **8**)

```
>> 2013/04/22 18:29:13 START OF PHASE MAIN_NEWBAS/TABIM_UPG
Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/DEFAULT.TPP checkimpdp T90
Importing for group 1/44: Setting stop mark at 'SAPK-701DRINSAPBASIS'.
Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -
Dmainimp_proc=1 -Dparallel=8
Importing for group 2/44: Setting stop mark at 'SAPK-701DFINPIBASIS'.
Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -
```

*Dmainimp\_proc=2 -Dparallel=4*

*Importing for group 3/44: Setting stop mark at 'SAPKITLRD5'.*

*Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -*

*Dmainimp\_proc=3 -Dparallel=3*

*Importing for group 4/44: Setting stop mark at 'SAPK-701DRINSAPBW'.*

*Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -*

*Dmainimp\_proc=1 -Dparallel=8*

*Importing for group 5/44: Setting stop mark at 'SAPKE60028'.*

*Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -*

*Dmainimp\_proc=2 -Dparallel=4*

*Importing for group 6/44: Setting stop mark at 'SAPK-603DFINEAAPPL'.*

*Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -*

*Dmainimp\_proc=1 -Dparallel=8*

*Importing for group 7/44: Setting stop mark at 'SAPKGPFD07'.*

*Running /usr/sap/SAP/T95/exe/tp pf=/usr/sap/T95/SUM/abap/var/TABIMUPG.TPP put T90 -*

*Dmainimp\_proc=3 -Dparallel=3*

Note that, for some groups, there is 9 parallel processes (3 mainimp\_proc x 3 parallel). Such minor differences are expected as SUM is always looking for the most optimal combination.

- How many R3trans process are forked during the import?

Looking in the correspondent import log files, it is possible to identify how many R3trans processes were forked during the import. For this, search for the message for the phrase "child processes forked". If the message is not found, no parallelism was done.

<pre>.../abap/log/SAPHB70105.CPT 1 ETP153 DICTIONARY IMPORT 4 ETW000 &gt; parallel 3 &lt;-- calculated dynamically 4 ETW000 3 child processes forked &lt;--number of R3trans child executed in parallel</pre>	<pre>.../abap/log/SAPKB70105.CPT 1 ETP187 R3TRANS SHADOW IMPORT 4 ETW000 &gt; parallel 3 &lt;-- no parallelism done if there is no message "child processes forked".</pre>
---	--

For troubleshooting about parallelism, please refer to note [#1597414](#): "SHADOW\_IMPORT phases are not using R3trans parallelism"

- Tips & tricks:

In order to avoid the error message "**====> HALT: creation of child process for parallel import failed. Please contact the SAP support.**" during the parallel import, make sure the UNIX system is set as **unlimited** and the parameter MAXUPROC (Maximum users processes) is large enough, it is normal during the parallelism the user <SID>adm starts a very high number of processes;

Before starting the upgrade, make sure the **database is optimized**, updating statistics and parameters. In additional use the **latest version** of **TP** and **R3trans**;

Avoid to use **NFS** in the upgrade directory, it should be a **local filesystem**, due the high I/O volume;

## See Also

SAP Note [1223360](#): Composite SAP Note: Performance optimization during import

SAP Note [1597414](#): SHADOW\_IMPORT phases are not using R3trans parallelism

SAP Note [1127194](#): R3trans import with parallel processes

SAP Note [1309506](#): SPAM/SAINT with parallel R3trans

SAP Note [2655872](#): Performance analysis for SPAM / SAINT run

SAP Note [2615119](#): R3trans import with statistical information



[Support Package Manager Settings](#) (to configure parallelism for SPAM)

[Defining the Settings in Add-On Installation Tool](#) (to configure parallelism for SAINT)

## Keywords

R3trans, parallel, mainimp\_proc, MaxMainImpProc, child process, SHADOW\_IMPORT\_UPG1, SHADOW\_IMPORT, SHADOW\_IMPORT, DDIC\_UPG, SHADOW\_IMPORT\_UPG2, SHADOW\_IMPORT\_INC, TABIM\_UPG, TABIM, tpp, tp0, config file, speed up, hang, slow, forever, dump, error, stop, windows, upgrade, enhancement package, r3up, optimize, tune up, turbo, speed, mock, copy, alternate, alternative, parconv, virtual, machine

## Other Components

Component	Description
BC-EHP-INS-TLA	Please use BC-UPG-TLS-TLA
BC-UPG-OCS-SPA	Support package tools for ABAP

## This document refers to

SAP Note/KBA	Title
2655872	Performance analysis for SPAM / SAINT run
2100040	FAQ: SAP HANA CPU
2000003	FAQ: SAP HANA
1904239	SUM options to change parameters entered in INITPUT and INITSUBST phases
1597414	SHADOW_IMPORT phases are not using R3trans parallelism
814704	MaxDB Version 7.6 parameter settings for OLTP/BW
724716	IS-H: Function 'Next Appointment' is not executed
651060	FAQ: Oracle Parallel Execution
2615119	R3trans import with statistical information
1888485	Database Parameter for 12.1.0.2
1680803	SYB: SAP Adaptive Server Enterprise - Best Practice for SAP Business Suite and SAP BW
1539124	SYB: Database Configuration for SAP applications on SAP ASE
1431798	Oracle 11.2.0: Database Parameter Settings

1346964	SAP MaxDB 7.9 database parameter recommendations
1308217	SAP MaxDB 7.8 database parameter recommendations
1223360	Composite SAP Note: Performance optimization during import
1127194	R3trans import with parallel processes
1004886	SAP MaxDB 7.7 database parameter recommendations
	<a href="http://help.sap.com/saphelp_nw70/helpdata/EN/49/3e920c29002221e10000000a42189d/content.htm">http://help.sap.com/saphelp_nw70/helpdata/EN/49/3e920c29002221e10000000a42189d/content.htm</a>
	<a href="http://help.sap.com/saphelp_nw70/helpdata/EN/b1/8c8838754ebb10e10000009b38f8cf/content.htm">http://help.sap.com/saphelp_nw70/helpdata/EN/b1/8c8838754ebb10e10000009b38f8cf/content.htm</a>
	<a href="http://help.sap.com/Download/Multimedia/ASE_16.0/ptquery.pdf">http://help.sap.com/Download/Multimedia/ASE_16.0/ptquery.pdf</a>
	<a href="http://www-01.ibm.com/support/docview.wss?uid=tss1wp101412">http://www-01.ibm.com/support/docview.wss?uid=tss1wp101412</a>
	<a href="https://blogs.saphana.com/2013/04/22/massively-parallel-processing-on-hana/">https://blogs.saphana.com/2013/04/22/massively-parallel-processing-on-hana/</a>
	<a href="http://maxdb.sap.com/doc/7_8/44/bd1ec6a5d51388e1000000a155369/content.htm">http://maxdb.sap.com/doc/7_8/44/bd1ec6a5d51388e1000000a155369/content.htm</a>
	<a href="http://support.sap.com/sltoolset">http://support.sap.com/sltoolset</a>
	<a href="http://blogs.msdn.com/b/saponsqlserver/archive/2015/04/29/max-degree-of-parallelism-maxdop.aspx">http://blogs.msdn.com/b/saponsqlserver/archive/2015/04/29/max-degree-of-parallelism-maxdop.aspx</a>
	IBM Techdocs White Paper: SAP Software Update Manager: Behaviors and Insights

Terms of use | Copyright | Trademark | Legal Disclosure | Privacy