

2753418 - Potential Performance Degradation Due to Timer Fallback

Version	18	Type	SAP Note
Language	English	Master Language	English
Priority	Recommendations / Additional Info	Category	Consulting
Release Status	Released for Customer	Released On	26.10.2020
Component	HAN-DB (SAP HANA Database)		

Please find the original document at <https://launchpad.support.sap.com/#/notes/2753418>

Symptom

You see a trace entry with the following message in the <service> traces after startup:

Fallback to system call for HR timer

Starting with SAP HANA 2 SPS04 you see the alert *Timer fallback*.

Other Terms

VM, virtual machine, hypervisor, Azure, __vdso_gettimeofday

Reason and Prerequisites

Reason:

For optimal performance SAP HANA requires that a machine is configured to provide accurate *Time Stamp Counter (TSC)* timing information. Otherwise significant performance degradation can be expected.

SAP HANA falls back to the system timer implementation if irregularities with the reported timings from processor instruction *RDTSC* are detected. Therefore the actual performance degradation depends on the system timer implementation. If the system timer itself uses *TSC*, the performance degradation could be insignificant.

The timer implementation in SAP HANA 1 (all SPS) and SAP HANA 2 (SPS01 - SPS03) checks whether the *constant_tsc* and *nonstop_tsc* flags are set using */proc/cpuinfo*. Additionally, it checks if the *RDTSC* respectively *RDTSCP* instruction can be used. SAP HANA falls back to the system timer if not all conditions are met.

Starting with SAP HANA 2 SPS04 the check also includes if the *RDTSCP* instruction is supported by the CPU and whether *constant_tsc* and *nonstop_tsc* are set. On Intel platform the system clocksource */sys/devices/system/clocksource/clocksource0/current_clocksource* and */sys/devices/system/clocksource/clocksource0/available_clocksource* are checked. The *current_clocksource* has to be set to *tsc* on Intel. The *available_clocksource* needs to contain *tsc* on Intel. If any of the checks fail, SAP HANA falls back to the system timer.

Affected Releases:

- All Releases of SAP HANA

Prerequisites:

- You run a SAP HANA system on Intel CPUs
- The entry *Fallback to system call for HR timer* is written to the trace
- On SAP HANA 2 SPS04 or later you receive the alert *Timer fallback*

Solution

For systems in Microsoft Azure environments, please follow SAP Note [2791572](#).

On SAP HANA 2 SPS04 if you get a *Timer Fallback* alert, check the info of the alert description to get a detailed description which problem was detected. To do so, you can execute the following select

```
select INFOTEXT from m_events where TYPE like 'TimerFallback'
```

The output can be interpreted as follows:

- **Incorrect system configuration detected.**

- Problem: The cpu flags for the *RDTSCP* instruction or the cpu flags for *constant_tsc* or *nonstop_tsc* are not set or *current_clocksource* and *available_clocksource* are not correctly configured.
- Solution: Make sure *RDTSCP* virtualization is enabled, the cpu flags are set correctly and the clocksource configuration is correct.

You can use the following commands to check:

```
if [ $(cat /proc/cpuinfo | grep nonstop_tsc | wc -l) -gt 0 ]; then echo "CPU flag nonstop_tsc is OK";
else echo "CPU flag nonstop_tsc is not set";fi
if [ $(cat /proc/cpuinfo | grep constant_tsc | wc -l) -gt 0 ]; then echo "CPU flag constant_tsc is
OK"; else echo "CPU flag constant_tsc is not set";fi
if [ "$(cat /sys/devices/system/clocksource/clocksource0/current_clocksource)" == "tsc" ]; then
echo "current_clocksource is OK"; else echo "current_clocksource is not set to 'tsc'";fi
if [ $(cat /sys/devices/system/clocksource/clocksource0/available_clocksource | grep tsc | wc -l) -
gt 0 ]; then echo "available_clocksource is OK"; else echo "available_clocksource does not
contain 'tsc'";fi
```

Please contact your hardware partner or service partner, if you are not able to resolve the issue yourself. Especially cloud providers have their own recommendations.

- **Inconsistent Time Stamp Counter context found during calibration on NUMA node <id>**

- Problem: The time stamp counter values returned during calibration are inconsistent for a certain NUMA node.
- Solution: Please report this as an incident on component HAN-DB.

- **Non-monotonic Time Stamp Counter update found on NUMA node <id>**

- Problem: The values read from the time stamp counter weren't monotonic. This hints at a hardware issue or a live-migration to a different node that doesn't hold this invariant.
- Solution: Make sure *RDTSCP* instruction works correctly on your machine. Please contact your hardware partner or service partner, if you are not able to resolve the issue yourself.

- **Bad Time Stamp Counter synchronization between cores on NUMA node <id>**

- Problem: The result from the *RDTSCP* instruction returned during re-calibration is inconsistent or diverged too much from what was expected for a certain NUMA node. This hints at a hardware issue.
- Solution: Make sure *RDTSCP* instruction works correctly on your machine. Please contact your hardware partner or service partner, if you are not able to resolve the issue yourself.

- **Could not calculate current time on NUMA node <id>**

- Problem: The result from the *RDTSCP* instruction returned is inconsistent or diverged too much

from what was expected. This hints at a hardware issue.

- Solution: Make sure *RDTS* instruction works correctly on your machine. A possible cause is that system time was adjusted by a large difference in one step. If you use *ntp* to synchronize the system time make sure it is configured to *slew* rather than *step*. Please contact your hardware partner or service partner, if you are not able to resolve the issue yourself.

On earlier SAP HANA versions you have to check all of the above mentioned configurations and make sure *RDTS* instruction works correctly on your machine.

Remarks:

- You have to restart the SAP HANA system in order to make the changes effective, as the timer is only checked during startup.
- If everything is configured correctly the SAP HANA service traces will contain the following message: *Using RDTS for HR timer*
- The trace message "*Using RDTS for HR timer*" is not traced on HANA 2 Revision 040.00 due to a programming error.
- As several popular hypervisors (Xen, Azure, KVM) have optimized system timer implementation, using the system timer by default can be sufficient and will prevent the alert from appearing, if the hypervisor is not able to propagate the flags correctly. You can enable the system timer by following one of below two options on all hosts of your HANA instance:
 - a. Setting environment variable *HDB_TIMER=system* as described in SAP Note [1827566](#) or
 - b. Setting parameter:
`[daemon]`
`environment = HDB_TIMER=system`

in the *daemon.ini* on HOST level. If you choose to set the parameter in the HOST specific *daemon.ini*, please consider that it will be applied to all HANA services started by the daemon, but not for the daemon itself.

This is still a valid measure as the daemon itself is not affected by the performance degradation.

- You have to restart your HANA instance afterwards such that the parameter change takes effect.
- If you encounter the alert after changing system time, please refer to SAP Note [2819623](#)
- If you encounter timer related performance degradation after live-migration of the virtual machine refer to SAP Note [2820725](#)
- The NTP configurations mean the following
 - *slew*: system time is shifted at most 0.5ms / second. e.g. If shift is 4ms, it will take 8 seconds to adjust to new time
 - *step*: system time is shift all at once.

Software Components

Software Component	Release
HDB	1.00 - 1.00
HDB	2.00 - 2.00

This document refers to

SAP Note/KBA	Title
1827566	How to set environment variables for SAP system?
2820725	Performance Degradation After VM Live Migration
2819623	Alert 7 "Timer fallback" Reported Due to Switched Back System Time
2791572	Performance Degradation Because of Missing VDSO Support in Azure VM
2380229	SAP HANA Platform 2.0 - Central Note
1523337	SAP HANA Database 1.00 - Central Note

This document is referenced by

SAP Note/KBA	Title
2915232	HANA timer fallback alert on AWS due to incorrect configuration
2843506	100% CPU usage on Time Stamp Counter (TSC)
2823760	HANA Timer Fallback on VMware based servers
2380176	FAQ: SAP HANA Database Trace
2100040	FAQ: SAP HANA CPU
2798857	Changing System Time in Operating System Can Cause HANA Crash
2859234	SAP HANA Timer Fallback alert on SLES 12 SP3
1656099	SAP Applications on AWS: Supported DB/OS and AWS EC2 products
2820725	Performance Degradation After VM Live Migration
2819623	Alert 7 "Timer fallback" Reported Due to Switched Back System Time
2791572	Performance Degradation Because of Missing VDSO Support in Azure VM

[Terms of use](#) | [Copyright](#) | [Trademark](#) | [Legal Disclosure](#) | [Privacy](#)