

Release Note TS 20.02

The following copyright notice protects this book under Copyright laws which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Copyright ©Bull SAS 2019

Printed in France

Trademarks and Acknowledgements

We acknowledge the right of proprietors of trademarks mentioned in this manual.

All brand names and software and hardware product names are subject to trademark and/or patent protection.

Quoting of brand and product names is for information purposes only and does not represent trademark and/or patent misuse.

Hardware

March 2019

**Bull Cedoc
357 avenue Patton
BP 20845
49008 Angers Cedex 01
FRANCE**

The information in this document is subject to change without notice. Bull will not be liable for errors contained herein, or for incidental or consequential damages in connection with the use of this material.

Table of Contents

Preface	p-1
Chapter 1. Overview	1-1
1.1. Operating Systems Versions	1-1
1.1.1. VMware ESXi	1-1
1.1.2. Linux	1-2
1.1.3. Windows	1-2
1.2. New Features and Changes	1-3
1.2.1. General	1-3
1.2.2. BIOS_SKL040	1-3
1.2.3. CPLD_P_CP	1-3
1.2.4. EMM_DEFAULT_BIOS_SETTINGS	1-3
1.2.5. EMM33_BMC	1-4
1.2.6. FPGA_CPB	1-4
1.3. Resolved Issues	1-5
Chapter 2. Known Restrictions and Issues	2-1
2.1. Platform Restrictions and Issues	2-1
2.1.1. Intel® DCPMM Hybrid Mode	2-1
2.1.2. Hot Plug of the Broadcom P210tp PCI card	2-1
2.1.3. Serial On LAN (SOL) activation	2-1
2.1.4. Restrictions on Partitioning a BullSequana S800 Server	2-1
2.1.5. LDAP Authentication	2-1
2.1.6. Mounting Virtual Media Files from the Remote Console	2-2
2.1.7. Updating the Server Hardware Console (SHC)	2-2
2.1.8. Locating an FDB Disk	2-2
2.1.9. bsmBiosSettings CLI Command	2-2
2.1.10. Unexpected Server Hardware Console (SHC) Reboot	2-2
2.1.11. Getting FRU Information on Mellanox ConnectX-4 Adapters	2-3
2.1.12. SHC Messages Page Unreachable	2-3
2.1.13. FPGA Update	2-3
2.1.14. Memory Module exclusion	2-3
2.1.15. Locking a 2.5" Disk Storage Unit	2-4
2.2. Software Restrictions and Issues	2-5
2.2.1. PCIe Blade Hot-Plug with RHEL 7.3	2-5
2.2.2. Using SR-IOV	2-5
2.2.3. Powering Off from the Server Hardware Console (SHC)	2-5
2.2.4. Incorrect Allocation of Intel® DCPMM Name Spaces to numa Node	2-6
Chapter 3. Delivery Content	3-1
3.1. Delivered items	3-1
3.2. Documentation	3-1
3.3. Platform Firmware	3-1
3.4. Adapter Firmware	3-3

3.5.	Customer Tools	3-4
3.6.	Management Information Base (MIB)	3-4
3.7.	Bull Admin Tools	3-5
Chapter 4.	Recommendations	4-1
4.1.	Updating from older Technical States	4-1
4.2.	Server Hardware Console (SHC) Firmware Update	4-1
4.3.	FPGA_CPB Update	4-1
4.4.	CVE-2013-4786 IPMI v2.0 vulnerability	4-2
4.5.	Copying the default BIOS settings file	4-2
4.6.	Changing BIOS Settings	4-2
4.7.	Performance Parameters	4-2
4.8.	Mixed Memory Configurations for SAP HANA	4-3
4.9.	Ethernet Ports with VMware	4-3
	4.9.1. Ethernet Ports Limitations	4-3
	4.9.2. Disabling Ethernet Ports	4-3
4.10.	GPUs with VMware	4-4
4.11.	Non Maskable Interrupt	4-4
4.12.	QAT and RedHat	4-4
4.13.	MicroSD cards in URS	4-4
Chapter 5.	Information	5-1
5.1.	Enabling Trusted Platform Module (TPM)	5-1
5.2.	Intel® DCPMM Memory Modules	5-2
	5.2.1. Support by Operating Systems	5-2
	5.2.2. Intel® DCPMM Memory Module Configuration Profiles	5-2
	5.2.3. Reserved Space	5-5
	5.2.4. Memory Mode Configuration: Keep DDR/Intel® DCPMM Ratio ..	5-5
	5.2.5. Replacing an Intel® DCPMM Memory Module	5-5
Chapter 6.	History of Previous Versions	6-1
6.1.	TS 07.02 (November 2018)	6-1
6.2.	TS 06.02 (August 2018)	6-4
6.3.	TS 05.04 (June 2018)	6-6
6.4.	TS 05.03 (May 2018)	6-7
6.5.	TS 05.02 (March 2018)	6-8
6.6.	TS 04.02 (January 2018)	6-10
6.7.	TS 04.01 (December 2017)	6-10

Preface

This document gives information about all changes from the previous version.

It also gives information about restrictions, known problems and the associated workarounds.

Finally it lists the objects delivered in the Technical State and the features of the resources provided on the Resource and Documentation DVDs.

Chapter 1. Overview

Important To fully address the Intel Meltdown/Spectre security alert, it is mandatory to update the Operating System.

1.1. Operating Systems Versions

The following versions are supported on BullSequana S servers. When there are more than one version, (*) indicates a recommended version.

1.1.1. VMware ESXi

ESXi 6.7u1 build 10302608 (*)

ESXi 6.5u2 patch 09 build 10175896 (*)

ESXi 6.5u2 build 8294253

ESXi 6.5u1 patch 02 build 7388607

Intel® Xeon® 1 st generation processors	Intel® Xeon® 2 nd generation processors
<ul style="list-style-type: none">• ESXi 6.7u1 build 10302608 (*)• ESXi 6.5 patch 11 build 10719125• ESXi 6.5u2 build 8294253• ESXi 6.5u1 patch 02 build 7388607	<ul style="list-style-type: none">• ESXi 6.7u1 build 10302608 (*)• ESXi 6.5 patch 11 build 10719125

Note Available on the Bull Support Website: <https://support.bull.com>

For BullSequana S200 server certification details check:

<https://www.vmware.com/resources/compatibility/detail.php?deviceCategory=server&productid=44639>

For BullSequana S400 server certification details check:

<https://www.vmware.com/resources/compatibility/detail.php?deviceCategory=server&productid=44854>

For BullSequana S800 server certification details check:

<https://www.vmware.com/resources/compatibility/detail.php?deviceCategory=server&productid=44745>

1.1.2. Linux

1.1.2.1. Red Hat

Intel® Xeon® 1st generation processors	Intel® Xeon® 2nd generation processors	Intel® DCPMM
<ul style="list-style-type: none">• RHEL 7.6: All models (*)• RHEL 7.5: All models (*)• RHEL 7.4: All models• RHEL 7.3: All models	<ul style="list-style-type: none">• RHEL 7.6: All models• RHEL 7.5: All models	<ul style="list-style-type: none">• RHEL 7.6: All models (*)• RHEL 7.5: All models

1.1.2.2. Suse

Intel® Xeon® 1st generation processors	Intel® Xeon® 2nd generation processors	Intel® DCPMM
<ul style="list-style-type: none">• SLES 12 SP3	<ul style="list-style-type: none">• SLES 12 SP4	<ul style="list-style-type: none">• SLES 12 SP4

1.1.3. Windows

Windows Server 2016 (with iaStorA.free.win8.64bit.4.3.0.1198 driver)

1.2. New Features and Changes

1.2.1. General

This version supports the following main new features:

- Second generation Intel® Xeon® Scalable processors
- Intel® DCPMM memory modules

1.2.2. BIOS_SKL040

- Supports display of ATOS logo on the BIOS access screen.
- Add 2 corrections for BIOS settings that could return to default values and cause loss of customer settings:
 - Preventing reset BIOS settings to default in case of CMOS issue (battery or checksum).
 - Avoid copying defaultbiossetup file to current biossetup file (then resetting configuration) prior to the transfer of the current BIOS settings back to EMM, at the end of BIOS phase.
- RAS: supports Partial Memory Mirroring, adapt the event sent to BMC in consequence.
- Enable data poisoning but let Viral disabled, used for error containment.
- Update Intel® Server Platform Services manageability engine firmware version to SPS_E5_04.00.04.393.0.
- First generation Intel® Xeon® Scalable processor microcode has evolved to MB750654_02000050 in Intel's Reference Code

1.2.3. CPLD_P_CP

- PEB Phy reset and SPI Mux selection logic changed to take care of corner cases of PEB flash corruption during AC Power ON/OFF operation.

1.2.4. EMM_DEFAULT_BIOS_SETTINGS

- Configure volMemMode to AUTO instead of 1LM to natively support MemoryMode (2LM) required by Intel® DCPMM memory modules.
- Enable packet poisoning by default to allow a better errors containment.

See 4.6. Changing BIOS Settings for more information on how to propagate the setting in your current BIOS settings file.

1.2.5. EMM33_BMC

- Support of Intel® DCPMM memory modules
- New Atos branding support
- An informative message is added in the message log during boot if the memory size is different from the one stored during previous boot
- An informative message is added in the message log during boot if the number of CPUs is different from the one stored during previous boot
- User can add free text in BMC message log through web interface: Maintenance > Maintenance Operations > Add User Message
- New RAS event: Partial Memory Mirroring activated
- A new option is available to User in Maintenance > Remote Console Setting > User Specific > "Launch Remote Console in Java WEB Start" to enable/disable the use of Java Web Start to launch the Remote Console.
- Improvement in security, some deprecated ciphers have been disabled:
 - SEED-SHA, RC4-SHA, RC4-MD5, DES-CBC3-SHA, DES-CBC-SHA, EXP-DES-CBC-SHA, EXP-RC2-CBC-MD5, EXP-RC4-MD5
 - No possibility of enabling enable-weak-ssl-ciphers - support have been removed in openssl
 - SSL v3 is disabled

1.2.6. FPGA_CPB

- Implemented BMC watchdog time out reset counter.
- Support to reset PCH and PEB PHY using ASRAM register.

1.3. Resolved Issues

Note The four first issues listed below are actually resolved by the TS 07.02 but were missing from its release note.

PEB Ethernet Activity LEDs

PEB Ethernet LEDs behave now correctly.

See Description Guide for more information on LED behavior

Fan Messages at Power On

No more inconsistent fan status messages are issued at Power On.

Power Supply Unit (PSU) Redundancy Sensor

The Power Redundancy sensor is now reliable to check power supply.

Dismounting and Mounting Back a Module from a Partition

Partitioning from the SHC after having dismounted and mounted back a module from a partition on a multi-module with a partition made of two modules is now possible.

Memory Module exclusion

Excluding a memory module from the SHC is now possible.

BIOS Update

Using the Preserve NVRAM option when updating the BIOS firmware does not lead to PEB/PEBS issues anymore.

Note To avoid any issues with firmware update, it is strongly recommended to use the global firmware update feature available through iCare or the BSM CLI commands.

Chapter 2. Known Restrictions and Issues

2.1. Platform Restrictions and Issues

2.1.1. Intel® DCPMM Hybrid Mode

Restriction

Intel® DCPMM Hybrid Mode is not yet supported.

2.1.2. Hot Plug of the Broadcom P210tp PCI card

Restriction

Hot-plugging the Broadcom PCI card BCM 957416A4160C is not possible. Insert or remove the card only when the operating system is stopped.

2.1.3. Serial On LAN (SOL) activation

Issue

When using the ipmi command "SOL activate" for Serial On LAN, there are issues with the keyboard.

Workaround

Open a ssh session on the SHC and use the terminal command.

2.1.4. Restrictions on Partitioning a BullSequana S800 Server

Intel restrictions on how processors can be linked in a glue-less configuration lead to restrictions on how a BullSequana S800 server can be partitioned.

The available partitioning schemes are:

- Two partitions:
 - Three modules and one module: all schemes are available
 - Two modules and two modules: only the (module 0, module 1) and (module 2 , module 3) scheme is available
- Four partitions

2.1.5. LDAP Authentication

Issue

When the DNS server configured from the Network Settings page the Server Hardware Console (SHC) belongs to the Active Directory domain, the LDAP authentication of the embedded controller fails without any error notification.

Workaround

Do not configure the DNS server before performing the LDAP authentication.

2.1.6. Mounting Virtual Media Files from the Remote Console

Issue

Installing software from a very large file via the Remote Console may fail with several medium errors reported.

Workaround

Use smaller files.

2.1.7. Updating the Server Hardware Console (SHC)

Issue

Updating the SHC firmware results in the Alert Setting Policies page being cleared.

Workaround

If you have modified the Alert Policies default settings, collect settings before the SHC update and restore them afterward.

2.1.8. Locating an FDB Disk

Issue

The command designed to locate a failed FDB disk fails to switch on the disk's LED, making it impossible to locate it.

Workaround

See [BullSequana S Description Guide, 86 A1 13FR](#), to locate FDB disks.

2.1.9. bsmBiosSettings CLI Command

Issue

In rare cases, the bsmBiosSettings.sh CLI command may hang, displaying infinite lines on screen.

Workaround

Kill and restart the command.

2.1.10. Unexpected Server Hardware Console (SHC) Reboot

Issue

In some cases, the SHC reboots without any prompting from the user.

Workaround

Connection is lost during reboot. Wait for a few minutes for the SHC to be available again.

2.1.11. Getting FRU Information on Mellanox ConnectX-4 Adapters

Issue

On servers running RHEL 7.5 and equipped with Mellanox ConnectX-4 adapters, there may be some discrepancies between the FRU page of the Server Hardware Console (SHC) and the results of the `lspciconf_m3.pl` script.

Workaround

The correct information is obtained from the SHC.

2.1.12. SHC Messages Page Unreachable

Issue

With Internet Explorer web browser, when adding a message from the **Add User Message** page, using the `”` character will make the **Messages** page consequently unreachable. The following message is displayed in a pop-up window:

Connection lost: your session has expired or Web server is not responding.

Workaround

Avoid using the `”` character in added messages.

2.1.13. FPGA Update

Issue

In rare cases, using either the SHC or the CLI commands, the update of the FPGA firmware appears unfinished but a message in the SHC Messages page indicates that the update has been completed successfully.

Workaround

Check firmware versions to be sure that the update is successful.

2.1.14. Memory Module exclusion

Issue

Excluding a memory module using the Hardware Management CLIs has no effect.

Workaround

Use the SHC to exclude a memory module.


2.1.15. Locking a 2.5" Disk Storage Unit

Issue

In some cases the latch does not insert properly into the locking bracket

Workaround

1. Check that the latch (A) is parallel to the top of the side wall (B) as shown below.



2. Press the latch (A) down to secure the HDD cage in place.

2.2. Software Restrictions and Issues

2.2.1. PCIe Blade Hot-Plug with RHEL 7.3

Issue

When hot-plugging a PCIe blade on a server running RHEL 7.3, the blade's Power LED blinks for a while then shuts down though it should stay on.

Workaround

This is an OS issue. Update to RHEL 7.4 or 7.5.

2.2.2. Using SR-IOV

Issue

On a BullSequana S400 or S800 server, attempting to assign the SR-IOV passthrough to a virtual machine fails, resulting in the following error message: *unsupported configuration: host does not support passthrough of host PCI devices*

Restriction

SR-IOV is not supported on Virtual Machines running SLES 12 SP2.

2.2.3. Powering Off from the Server Hardware Console (SHC)

Issue

On servers running RHEL, clicking the Power Off button available in the Power Management page of the SHC does not result in the complete shutdown of the system. Clicking the Force Power Off button is necessary to get a complete shutdown.

Workaround

The Operating System must be configured to accept the power off request.

- With RHEL 7.3:
 - a. In the RHEL Graphical User Interface, go to Applications > Utilities > Tweak Tool > Power > Power button action.
 - b. Choose the Shutdown option.
- With RHEL 7.4, 7.5 and 7.6:
 - a. Install the acpid package.
 - b. Replace `/etc/acpi/actions/power.sh` content with the following content:

```
#!/bin/sh
PATH=/usr/sbin:/usr/bin
shutdown -h now
```

2.2.4. Incorrect Allocation of Intel® DCPMM Name Spaces to numa Node

Issue

On a server running a RedHat prior to 7.6 with Intel® DCPMM memory modules configured in Application Direct mode, the `numactl` command returns an incorrect answer.

Workaround

Update to RHEL 7.6.

Chapter 3. Delivery Content

3.1. Delivered items

- Documentation, firmware and customer tools are delivered on the Resource and Documentation DVDs
- BSMHW_NG and iCare are delivered on the Resource and Documentation DVDs
- VMware ESXi Installer is delivered, if ordered, on a bootable USB key

3.2. Documentation

Note (*) indicates a new version, (**) indicates a new item.

Name	Description	Version
BullSequana S Customer Documentation Portfolio	Complete documentation dedicated to the customer	08 (*)
BullSequana S Field Documentation Portfolio	Complete documentation dedicated to the field	07 (*)

3.3. Platform Firmware

Notes • (*) indicates a new version, (**) indicates a new item

- The FPGA_W3WEO firmware version always reads as (0.0.0) even if the firmware has been successfully updated to a more advanced version.
- There are now two different images of each BIOS firmware: one compatible with the PEB board and the other with the PEBS board. Their versions are numbered as follows:
 - x=0 for PEBS
 - x=1 for PEB

The managing tools are configured to automatically select the adequate BIOS image.

Name	Description	Version
ApachePAss	Firmware for Apache Pass memory modules	01.2.0 build 5321 (**)
BIOS_CCL041	BIOS firmware for second generation Intel® Xeon® Scalable processors	41.27.01 build x04 (**)
BIOS_SKL040	BIOS firmware for first generation Intel® Xeon® Scalable processors	40.61.01 build x06 (*)
CPLD_IO_CPB	Flash image for the IO CPLD component on the CPB board	2.6.0.0

Name	Description	Version
CPLD_NBB	Flash image for the CPLD component on the NBB board	2.2.0
CPLD_P_CPB	Flash image for the CPLD component on the CPB board	2.4.9.0 (*)
EMM33_BMC	Server Hardware Console (SHC) firmware	33.32.02 build 0003 (*)
EMM_DEFAULT_BIOS_SETTINGS	Default BIOS settings file	1.6 (*)
ESXi_6.5_BullSequana_S	VMware supervisor	6.5u2p9 build 10175896
ESXi_6.7_BullSequana_S	VMware supervisor	6.7u1 build 10302608
FPGA_CPB	FPGA firmware for theCPB board	2.7.0.0 (*)
FPGA_FLASH_M3WEO	Flash image for the embedded firmware of the sWitch Ethernet One Gigabit (WEO).	1.0.0
FPGA_W3WEO	FPGA image for the sWitch Ethernet One Gigabit (WEO)	2.0.0 (appears as 0.0.0)
FW_PEB	Flash image for the SPI 4Mbit 85MHz 8SOIC 256Byte per page	2.B.9
FW_PHY_PEBS	Flash image for the EEPROM 256KBIT 400KHZ 8SOIC	1.0.0
FW_URS	Flash for the SPI 4Mbit 75MHz 8SO.	0.0.1

3.4. Adapter Firmware

Note (*) indicates a new version, (**) indicates a new item.

Name	Version
Broadcom_PCIe_BCM957416A4160C	212.0.147.0
Emulex_PCIe_LPe12002-M8	fw202a3 - UniversalBootCode1140a6 OneCommandManagerCLI 11.4.204.12
Emulex_PCIe_LPe31002-M6	11.4.204.20 OneCommandManagerCLI 11.4.204.12
Emulex_PCIe_LPe31004-M6	11.4.204.25 OneCommandManagerCLI 11.4.204.12
Ethernet_Intel_I350-X520	22.9
INTEL_PCIe_OPA_HFI	10.4.2.0.7-110
LSI_MegaRaid_SAS_9361-8i	Package 24.21.0-0012 Firmware 4.680.00-8249
LSI_MegaRaid_SAS_9361-16i	Package 24.21.0-0012 Firmware 4.680.00-8249
LSI_SAS_9300_8e	P16 (**)
LSI_SAS_9300-8i	15.00.02.00/8.35.04.00
LSI_SAS_9305-16i	15.00.00/8.35.00
Mellanox_ConnectX-4Lx	14.23.1020
Mellanox_ConnectX-5	16.23.1020

3.5. Customer Tools

Note (*) indicates a new version, (**) indicates a new item.

Name	Description	Version
EMM_REGS_DUMP	This set of files gives the list of registers to dump in CPU and FPGA devices in case of CATERR detection or of IPMI dump command.	1.0
EMM33_BMC_Bckp	The backup image of the Server Hardware Console (SHC) firmware.	33.14.00 build 2
mc-setup	A Linux Utility used to discover the embedded management board's MAC address and to change the embedded management board's IP address.	1.2.1 build 2
MceLog	A Linux tool dedicated to collect MCE logs on servers. This version is updated to fully support errors reported by the Intel® Xeon® Scalable processors.	158 build 50
psetup	A Windows Utility used to discover the embedded management board's MAC address and to change the embedded management board's IP-address.	1.2.4

3.6. Management Information Base (MIB)

Note (*) indicates a new version, (**) indicates a new item.

Name	Description	Version
MIB_bull_PlatformManagement	Defines Platform Management SNMP interfaces of Bull servers.	201807171200Z
MIB_PlatformEventTraps	The Platform Event Trap definition file. This MIB (Management Information Base) file is used by SNMP (Simple Network Management Protocol) managers to receive server hardware events.	2.3.6

3.7. Bull Admin Tools

Note (*) indicates a new version, (**) indicates a new item.

Name	Description	Version
BSMHW_NG	A set of prompt commands, based on free IPMI open source commands, used to manage server or device hardware. These commands can be used to return information and status and/ or to remotely control and configure server hardware.	1.5.8 (*)
Bull_Admin_Tools_VM_Appliance	An appliance that delivers Bull Administration tools on CentOS system.	2.4.0 (*)
iCare	A WEB application used for hardware unit maintenance. Both Linux and Windows versions are provided.	2.5.4 (*)

Chapter 4. Recommendations

4.1. Updating from older Technical States

Important Due to a change in the EMM33_BMC firmware's size between TS 05.04 and TS 06.02, its mandatory to update to TS 06.02 before updating to any later TS.

If the server is running a TS preceding 06.02, perform the following steps:

1. Update FULLY to TS 06.02.
2. Perform an AC cycle an ALL modules. For each module, perform the following steps:
 - a. Power off the module.
 - b. Unplug the power cords.
 - c. Wait until the power LEDs are off.
 - d. Plug in the power cords.
 - e. Power on the module.
3. Clear the Internet browser's cache before using the Server Hardware Console (SHC) for the first time.
4. Update to any later TS.

4.2. Server Hardware Console (SHC) Firmware Update

- It is strongly recommended to power off the system before updating the SHC firmware. Otherwise, some slave modules may be lost.
- If the PCIe slot 0 is not visible after updating the SHC, do an AC/Off - AC/On to see the slot.

Note To avoid any issues with firmware update, it is strongly recommended to use the global firmware update feature available through iCare or the BSM CLI commands.

4.3. FPGA_CPB Update

It is mandatory to update the SHC firmware before updating the FPGA_CPB firmware.

Note To avoid any issues with firmware update, it is strongly recommended to use the global firmware update feature available through iCare or the BSM CLI commands.

4.4. CVE-2013-4786 IPMI v2.0 vulnerability

To address this vulnerability, it is strongly recommended to change the super default account username.

4.5. Copying the default BIOS settings file

The EMM_DEFAULT_BIOS_SETTINGS must be present on the server.

See CLI Reference Guide for more information on how to copy the default BIOS settings file

4.6. Changing BIOS Settings

W083 WARNING

W083:

Do not change BIOS setup settings unless directed to do so by the support team.

W082 WARNING

W082:

These procedures are for advanced users only. Risk of system damage.

- To configure volMemMode to AUTO, use the following command:

```
bsmBiosSettings.sh -H <IP address> -u <user> -p <pwd> -a set -n 'MEM.volMemMode 2'
```

- To enable packet poisoning by default, use the following command:

```
bsmBiosSettings.sh -H <IP address> -u <user> -p <pwd> -a set -n 'SETUP.PoisonEn 1'
```

4.7. Performance Parameters

1. It is recommended to update the defaultbiossetup file to its latest version and to manually set the StaleAtoSEn value at least once:

- a. Check the StaleAtoSEn value.

```
BsmBiosSettings.sh -H ip -u user -p pwd -a get -n 'UPI.StaleAtoSOptEn'
```

- b. Set the StaleAtoSEn value.

```
BsmBiosSettings.sh -H ip -u user -p pwd -a set -n 'UPI.StaleAtoSOptEn 1'
```

2. For systems that are running SAP Hana/SAP BW, except BullSequana S200 servers, some BIOS settings may be tuned to improve performance with Intel® Xeon® Scalable processors by disabling HW prefetchers and adjusting IRQ/RRQ threshold.

```
bsmBiosSettings.sh -H ip -u user -p pwd -a set -n 'CPU.DCUStreamerPrefetcherEnable 0'  
bsmBiosSettings.sh -H ip -u user -p pwd -a set -n 'CPU.DCUIPPrefetcherEnable 0'  
bsmBiosSettings.sh -H ip -u user -p pwd -a set -n 'CPU.MlcSpatialPrefetcherEnable 0'  
bsmBiosSettings.sh -H ip -u user -p pwd -a set -n 'CPU.MlcStreamerPrefetcherEnable 0'  
bsmBiosSettings.sh -H ip -u user -p pwd -a set -n 'UPI.IrqThreshold 3'
```

4.8. Mixed Memory Configurations for SAP HANA

On servers equipped with first generation Intel® Xeon® Scalable processors, two specific configurations mixing memory modules are allowed under the following conditions:

- For SAP HANA only
- Using only authorized parts:
 - RDIMM 3DS 128 GB (Samsung M393AAK40B42-CWD) with RDIMM 3DS 64 GB (Samsung M393A8K40B22-CWD)
 - LRDIMM 64 GB (Samsung M386A8K40BM2-CTD) with LRDIMM 32 GB (Smasung M386A4K40BB0-CRC)
- In the configuration recommended: the larger size memory module in slot 0 and the smaller size one in slot 1
- All memory modules in slot 0 are the same and all memory modules in slot 1 are the same
- With processors with M suffix if 128 GB and 64 GB memory modules are being mixed

Note With a mix of 64 GB and 32 GB memory modules, the speed is 2 400 instead of 2 666.

4.9. Ethernet Ports with VMware

4.9.1. Ethernet Ports Limitations

There is a number of configurations rules to follow when using VMware:

<https://configmax.vmware.com/home>

In particular, the number of Ethernet ports supported by VMware is limited:

- Combination of 1GB and 10 GB ports: four 1GB ports and sixteen 10GB ports
- i40en 10GB (Intel) ports: eight ports

4.9.2. Disabling Ethernet Ports

1. Set up the server in Single Point of Management network configuration.

See BullSequana S Configuration Guide, 86 A1 27FR, for more information

2. Disable Ethernet ports using BSM CLI commands.

Module 1

```
./bsmBiosSettings.sh -H <module IP address> -u super -p pass -a set -n 'PCI.PciePortDisable_47 0'
```

Module 2

```
./bsmBiosSettings.sh -H <module IP address> -u super -p pass -a set -n 'PCI.PciePortDisable_89 0'
```

Module 3

```
./bsmBiosSettings.sh -H <module IP address> -u super -p pass -a set -n 'PCI.PciePortDisable_131 0'
```

See BullSequana S Remote Hardware Management CLI Reference Guide, 86 A1 19FR, for more information

4.10. GPUs with VMware

It is mandatory to set the MMIOH BIOS parameter to 4 using the following command:

```
./bsmBiosSettings.sh -H <IP address> -u <user> -p <pwd> -a set -n 'RC.MmiohBase 4'
```

Note With Nvidia Tesla boards and if global memory size is bigger than 4 TB, set the parameter to 6.

4.11. Non Maskable Interrupt

On specific configurations with servers running Linux, it is not always possible to get a dump when using NMI.

Set the Linux boot `crashkernel` parameter value to 2048.

4.12. QAT and RedHat

It is strongly recommended to disable the QAT (Quick Assist Technology) option under RedHat as it does not work correctly.

This option is disabled by default in the BIOS settings.

4.13. MicroSD cards in URS

In order to work properly in the Internal Dual RAID board, the microSDs must be formatted correctly. Please use only those provided by Atos representatives.

Chapter 5. Information

5.1. Enabling Trusted Platform Module (TPM)

Important Before enabling the TPM feature, it is mandatory to verify that its usage complies with local laws, regulations and policies and get approvals or licenses where applicable. Bull SAS will not be responsible for any related liabilities to any compliance issues arising from your usage of TPM violating the above mentioned requirements.

Important 在启用TPM功能之前，您必须确认其使用符合当地法律，法规和政，并在适用情况下获得批准或许可。违反上述要求，Bull SAS将不承担由于您使用TPM而导致的任何合规问题的相关责任。

5.2. Intel® DCPMM Memory Modules

5.2.1. Support by Operating Systems


Check that the Intel® DCPMM mode that is going to be used is correctly supported by the server's OS.

5.2.2. Intel® DCPMM Memory Module Configuration Profiles

The following profiles are recommended per CPU socket . For a multi-module server, each CPU socket should be populated identically.


For each profiles, the three modes (Memory mode, Application Direct mode and Hybrid mode) can be used.

Profile 2-1-1 (6 DDR4 + 2 Intel® DCPMM memory modules)




Mark	Memory module type	iMC	Channel	DIMM
C0	DDR4	0	2	0
B0	DDR4	0	1	0
A0	DDR4	0	0	0
A1	Intel® DCPMM	0	0	1
D1	Intel® DCPMM	1	3	1
D0	DDR4	1	3	0
E0	DDR4	1	4	0
F0	DDR4	1	5	0

Profile 2-2-1 (6 DDR4 + 4 Intel® DCPMM memory modules)



Mark	Memory module type	iMC	Channel	DIMM
C0	DDR4	0	2	0
B0	DDR4	0	1	0
B1	Intel® DCPMM	0	1	1
A0	DDR4	0	0	0
A1	Intel® DCPMM	0	0	1
D1	Intel® DCPMM	1	3	1
D0	DDR4	1	3	0
E1	Intel® DCPMM	1	4	1
E0	DDR4	1	4	0
F0	DDR4	1	5	0

Profile 2-2-2 (6 DDR4 + 6 Intel® DCPMM memory modules)



Mark	Memory module type	iMC	Channel	DIMM
C0	DDR4	0	2	0
C1	Intel® DCPMM	0	2	1
B0	DDR4	0	1	0
B1	Intel® DCPMM	0	1	1
A0	DDR4	0	0	0
A1	Intel® DCPMM	0	0	1
D1	Intel® DCPMM	1	3	1
D0	DDR4	1	3	0
E1	Intel® DCPMM	1	4	1
E0	DDR4	1	4	0
F1	Intel® DCPMM	1	5	1
F0	DDR4	1	5	0

5.2.3. Reserved Space

Some memory space is reserved for internal Intel® DCPMM usage and is not available for the configuration modes.

5.2.4. Memory Mode Configuration: Keep DDR/Intel® DCPMM Ratio

When using Intel® DCPMM in memory mode configuration, the DDR memory modules are used as a cache to speed up the accesses to the Intel® DCPMM.

To be efficient, a ratio of 1/16 to 1/4 must be maintained between the DDR and Intel® DCPMM capacity.

The standard recommended value is 1/8.

Example

Intel® DCPMM memory module capacity	DDR4 memory module capacity: Ratio 1/4	DDR4 memory module capacity: Ratio 1/16
6 TB	1,5 TB	375 MB

For a 6 TB of Intel® DCPMM capacity, use between 375 MB and 1,5 TB of DDR4.

5.2.5. Replacing an Intel® DCPMM Memory Module

Before replacing an Intel® DCPMM memory module configured in Application Direct Mode, back up the data present on the memory module, as you would with any other disk.

Chapter 6. History of Previous Versions

6.1. TS 07.02 (November 2018)

New features and changes

BIOS_SKL040

- Add new entry "12TB" for MMIOH_Base setting in BIOS setup to solve a failure with VMware ESXi and Tesla GPU cards
- Update Skylake microcode for security issues Spectre_NG and L1TF (SA-00115 and SA-00161) to version MB750654_0200004D
- Fix the ACPI SLIT table when SNC is enabled, now the table is getting the correct distances
- Implement Page Retire mechanism for VMware (remove memory pages when too much corrected errors occur) by using CMCI interface. This mechanism is controlled with 3 new settings and enabled by default:
 - MEM.PageRetireEn: Activate or not the Page Retire (0=Disable/1=Enable)
 - MEM.PageRetireErrThreshold: Num of errors in a timeframe (default:10)
 - MEM.PageRetireThresholdWindow: Timeframe in hours (default:24)
- Remove "Lacking IO resources" warning because it is only useful in Legacy mode not UEFI mode.
- Workaround for UPI Topology issue, rerun several times the process instead of aborting immediately (max 4 times).
- Send additional information to EMM when error with DIMMs to know if memory is excluded or not.
- Use BIOS Code drop 59

CPLD_IO_CPB

- Filter removed from BMC hang status signal to decrease action delay during the BMC hang event Filter added for BMC hang status signal

EMM33_BMC

- New version of OpenSSL 1.0.2k, to increase the security level (support of TLS 1.2)
- New version of OpenLDAP 2.4.46

EMM_DEFAULT_BIOS_SETTINGS

- Revert UPI.StaleAtoSOptEn to 1, mistakenly switched to 0
- Add a revision number to the defaultbiossetup file, name will now have the following naming : "defaultbiossetup.X.Y"
 - X= Major revision (example: adding or removing settings)
 - Y= Minor revision (example: changing settings values)

Important The revision number must be removed from the name before the file is uploaded to BMC.

- To avoid any compatibility issues with USB devices, all the USB ports located at the front of the server are configured as USB2 ports. Default value are switched to 0 instead of 1.
 - PCH.PchUsbSsPort_3 : control topmost connector
 - PCH.PchUsbSsPort_4 : control bottommost
 - PCH.PchUsbSsPort_5 : control middle
- Add 3 new BIOS settings for VMware Page Retirement
 - MEM.PageRetireEn : Activate or not the Page Retire (0=Disable/1=Enable)
 - MEM.PageRetireErrThreshold : Num of errors in a Timeframe (default:20)
 - MEM.PageRetireThresholdWindow : Timeframe value (default:24 hours)

To modify the settings, use the following command:

```
./bsmBiosSettings.sh -H <IP address> -u <user> -p <pwd> -a set -n '<parameter> <value>'
```

FPGA_CPB

- ASRAM operating frequency changed to 100Mhz from 200Mhz
- PEB buffer was in Flip Flops. To do timing closure, this is changed to RAM
- FPGA_CS_N signal is double synced before using in the counter
- Latches were there in CAT Error and CPU F/U/C Error Timers and this are removed
- Clock enable signals were used as clocks in few places so these are changed too
- The 6.25MHz clock to shifty bus logic were output of counter earlier. Now this clock is generated from PLL itself
- Some dangling logics are removed
- Clock and other timing constraints were updated to make sure there are no internal timing issue
- Latches were there in Sync block and these are removed now

FPGA_M3WEO

- Fixed major revision ID
- The following registers were wrongly mapped for module 1 and module 2:
 - SPI1_register_addr_B16
 - o_SPI1_register_addr_B17
 - o_SPI2_register_addr_B16
 - o_SPI2_register_addr_B17.

This been corrected: the following operations will now function correctly:

- FRU Read/Write
- BCM register Read/Write
- BCM Flash Read/Write
- M3WEO register access for module 1 and module 2

Resolved issues

FAN Regulation Messages

Fan speed is now suitably regulated: there are no longer multiple alarming fan speed sensor statuses or messages in the System Event Log (SEL).

Mounting Drives as Virtual Media

Virtual Media now works correctly with two drives.

Be aware that clicking on Connect or Disconnect in the Virtual Media dialog box causes the existing virtual media to be disconnected and a new USB device to be connected with the updated virtual media configuration.

Also, clicking on Virtual Media Connect or Disconnect buttons while an installation is running from virtual media is likely to interrupt the installation.

PXE Boot with a Mellanox_ConnectX-4Lx Adapter

With the adapter firmware provided with this TS, the UEFI firmware is now installed by default. Consequently, booting with PXE is no longer an issue.

If this adapter is part of the server, be sure to update its firmware to the latest version.

IO Port Resource Message

The irrelevant *Lacking IO port resource* message no longer appears when booting.

6.2. TS 06.02 (August 2018)

New Features and Changes

BIOS_SKL040

- Enable StaleAtoSEn BIOS setting by default to improve performance

See Chapter 4. Recommendations for more information on performance improvement

- Fix S800 USB booting timeout for VMware
- UPI warning message sent to BMC only when failure in fast mode
- Suppress UPI warning for non-existing UPI link
- Avoid errors on Intel® Xeon® Scalable processors with only two UPI links
- Fix in DMAR table avoiding error messages with RHEL
- Fix Bootdev issues with bootable USB or VMware
- Provide relevant memory module location information in case of memory failure or warning (module/socket/iMC/channel/dim/rank)
- Fix SRAT APIC and X2APIC affinity structures

EMM33_BMC

- In Messages log, BIOS messages are not displayed as BMC messages anymore
- Support of OEM model 38 in SNMP traps

FPGA_CPB

- Logic used to run RPL_FAN at full speed changed
- ID LED turning ON or OFF logic moved to IOCPLD

Resolved Issues

CPU Power Consumption Sensors

The CPU power consumption sensors now reports correct values.

WEO fault Message

A WEO fault message is no longer issued when the WEO sensor has no reading.

Boot Manager Entries

When there are more than 15 entries in the boot manager, each entry is now assigned a unique EFI network number.

Memory Module Messages during BIOS Initialization

Inconsistent warning messages about the memory modules are no longer issued during BIOS initialization.

Updating Firmware from the Server Hardware Console (SHC)

When a firmware update is successful, the following message is no longer displayed:

Please wait for the connection to be established.

Missing Processors When Booting the server

There no longer processors missing from the configuration with the following message in the SEL:

```
2018-05-14 18:14:01 BMC Message BIOS Init Warning Message on Module: 0 DIMM:  
([Major-code:58h; Minor-code:02h])
```

6.3. TS 05.04 (June 2018)

New features and changes

EMM33_BMC

New release fixing the following issues:

- Incorrect system name displayed by the NFC tag
- DFM LEDs turning on red randomly

FPGA_CPB

New release fixing the following issue:

DFM fans always running at full speed

Resolved Issues

Incorrect system name displayed by the NFC tag

There are no longer errors in the system name displayed by the NFC tag.

DFM LEDs turning on red randomly

The DFM LEDs do not become red randomly anymore.

DFM fans always running at full speed

The DFM fans are now running at suitable speed.

6.4. TS 05.03 (May 2018)

New features and changes

EMM33_BMC

- New release fixing the following issue:
DFM fans randomly unavailable with TS 05.02.

Resolved Issues

DFM fans randomly unavailable with TS 05.02

With the present release of the EMM33_BMC firmware, the fans are running normally, without random faults.

6.5. TS 05.02 (March 2018)

New features and changes

BIOS_SKL040

- Intel fix for Spectre and Meltdown issues
- Memory SddcPlusOne RAS feature enabled by default.
- Fixed excluded dimm display in setup memory topology.
- Improved PatrolScrubbing logging messages on error.
- The integrated Gbe controller is now reported to the Server Hardware Console (SHC).
- Improved dmidecode type9 display for PCIe slots information.
- The Press Esc line is now displayed at 60% of window height for small screens.
- Added Rank Sparing RAS feature (1 or 2 spare ranks).
- Improved RAS messages sent to SHC for SDDC, ADDDC, RankSparing, Leaky Bucket RAS features.

EMM33_BMC

- Changed display of identification LED for better understanding of actions.
- SEL events can be displayed in multiple or single web pages.
- Added the SEL binary file to Collect Log files.
- Partitioning is now available from the SHC, including from a slave console.
- Boot device and instance can be selected from the SHC. This is used to set parameters that direct the system boot to a particular option after a system power up or reset. This feature is the same as the IPMI boot device option.
- PCIe hot plug is available under Red Hat and Suse only.
- On the Power Management web page, Force Power Off, Force Power Cycle, Hard Reset and Diagnostic Dump commands need to be confirmed.
- The "super" user name can be modified from the SHC.
- Implemented reset to default function.

FPGA_CPB

- Fans run at FULL SPEED when the SHC hangs in power on state.

Resolved Issues

Simultaneous power on of different partitions

Powering on two modules of different partitions simultaneously is now possible.

FPGA Update on a BullSequana S800 Server

Inconsistent messages are no longer issued at power on after updating the FPGA on a BullSequana S800 server.

BullSequana S200 Server BIOS Update with Error in SEL

Inconsistent messages are no longer issued when the BIOS update is successful.

Unable to Update Bios with the Preserved Nvram Option

Updating the BIOS firmware from the SHC with the preserved Nvram option is now possible. On a multi-module server, every module is updated successfully.

ESXi 6.5 Installation Failure on USB Raid SD Card (URS)

Installing ESXi 6.5 on a USB Raid SD Card (URS) with Virtual Media is now possible without failure.

Updating the SHC firmware on a multi-module server

The SHC will not show the firmware update as completed if it is not completed on all modules.

6.6. TS 04.02 (January 2018)

This Technical State 04.02 is a patched one compared to the Technical State 04.01. It addresses the Intel Meltdown/Spectre patch.

See The Technical Support Bulletin 400-18-02 for more details, available on the Bull Support Website: <https://support.bull.com>

6.7. TS 04.01 (December 2017)

First delivery

Bull Cedoc
357 avenue Patton
BP 20845
49008 Angers Cedex 01
FRANCE