



Monthly Specification Update

Intel® Visual Compute Accelerator VCA1283LVV

Feb 2017

Intel Server Boards and Systems

Revision History

Date	Modifications
Dec. 2015	Initial release.
May 2016	Update for 1.2 Software Release.
Feb 2017	<ul style="list-style-type: none"> • Updated Erratum 2, 4, 6, 7, 8, 9, 14, 15 • Add Erratum 16 • Updated Product Scope Table

Disclaimers

The specified Server System may contain design defects or errors known as errata that may cause the product to deviate from the published specifications. Current characterized errata are documented in this Specification Update.

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation.

*Other brands and names may be claimed as the property of others.

Copyright © Intel Corporation 2016-2017

Contents

Preface	1
Product Scope	2
Summary Tables of Changes.....	3
Errata.....	5
Media Server Studio Errata Items impacting VCA.....	11
Documentation Updates	12

This page intentionally left blank

Preface

This document is intended to communicate product errata, published specification changes, published specification clarifications, and published document changes for the following Intel Server products:

- *Intel® Visual Compute Accelerator VCA1283LVV*

It is intended for system integrators and software developers of applications, operating systems, or tools.

Nomenclature

1. **Specification Changes** are modifications to the current published specifications for Intel server boards. These changes will be incorporated in the next release of the specifications.
2. **Specification Clarifications** describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.
3. **Documentation Changes** include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.
4. **Errata** are design defects or errors. Errata may cause the server board or system behavior to deviate from published specifications. Hardware and software designed to be used with any given processor stepping must assume that all errata documented for that processor stepping are present on all devices.

Product Scope

The following Intel product codes and Intel factory installed System Software Stack are covered by this update:

Product Code	Server PBA	TA	BIOS Revision	BMC Firmware Revision	ME Firmware Revision	FRU/SDR Revision
VCA1283LVV	H57152-404	H90713-001	VCA-BIOS_0ACGC100.0000	N/A	N/A	N/A
VCA1283LVV	H57152-405	H90713-002	VCA-BIOS_0ACGC104.0000	N/A	N/A	N/A
VCA1283LVV	H57152-500	H90713-003	VCA-BIOS_0ACGC106.0000	N/A	N/A	N/A

Summary Tables of Changes

The following tables indicate the errata and the document changes that apply to the Monthly Specification Update. Intel intends to fix some of the errata in a future stepping of components, and to account for the other outstanding issues through documentation or specification changes as noted. The tables use the following notations:

Doc: Intel intends to update the appropriate documentation in a future revision.

Fix: Intel intends to fix this item in the future.

Fixed: This item has been corrected.

No Fix: There are no plans to fix this item.

Shaded: This item is either new or has been modified from the previous specification update.

Table 1. VCA Errata Summary Table

No.	Plans	Description of Errata
1	No Fix	Simultaneous condition of active transfer of network packets over the PCI network to the VCA nodes and hard failure of one of the nodes sometimes results in DMA engine error.
2	Fixed	Network on the VCA node hangs after several minutes of low memory.
3	No Fix	High loss of network packets when small UDP packets are sent and throughput greater than 2Gbps.
4	Fixed	Xeon E3 system is not recognized as valid host for VCA upon RPM installation.
5	Fix	PLX 8713 device name displayed incorrectly
6	Fixed	vcactl reboot: may fail if path to image is > 256 characters long.
7	Fixed	User in "vcausers" group cannot run all "vcactl" commands
8	Fixed	Autoboot may take nearly 3 minutes to complete with NFS image
9	Fixed	Autoboot fails if vcausers group does not have read access to the image directory
10	No Fix	DomU to DomU transfers are significantly slower than DomU – DomX or DomU – bare metal
11	No Fix	(Informational) Why does the Linux kernel indicate it is "tainted"?
12	No Fix	Time Zone is not automatically synchronized between host and nodes.
13	Fix	Node becomes unresponsive when using NFS and heavy memory load
14	No Fix	CATERR on node requires AC power cycle of host to clear.
15	Fixed	Hosts running Intel® Xeon® E5-2600 v4 processor family must add "nosmap" to kernel boot line
16	Fix	An Intel Hyper-Threading Technology Enabled Processor may exhibit internal parity errors or unpredictable system behavior.

Table 2. Intel® MSS Errata Summary Table

No.	Plans	Description of Errata
1	Fix	OpenCL BufferReadWrite certification test fails if running from the virtual machine (GuestOS) under Xen hypervisor
2	Fix	If the VCA node runs of system memory then GPU crashes upon decode that creates resulting file larger than available memory in the system.

Table 3. Documentation Updates

No.	Doc Name	Description of Documentation Change
-----	----------	-------------------------------------

Following are in-depth descriptions of each erratum / documentation change indicated in the tables above. The errata and documentation change numbers below correspond to the numbers in the tables.

Errata

1. Simultaneous condition of active transfer of network packets over the PCI network to the VCA nodes and hard failure of one of the nodes sometimes results in DMA engine error.

Problem	Simultaneous condition of active transfer of network packets over the PCI network to the VCA nodes and hard failure of one of the nodes (such as kernel panic that causes CPU self reset or hardware thermal trip CPU shutdown) sometimes results in DMA engine error and this particular engine is not usable any more.
Implication	After individually re-booting this VCA node network interface is still broken - no data is transferred between the VCA node and the VCA host.
Status	Will not fix
Workaround	Reboot of the entire host system is required to recover DMA engine from hung state.

2. Network on the VCA node hangs after several minutes of low memory.

Problem	Network on the VCA node hangs after several minutes of low memory. Usually caused by NFS file caching when streaming several files larger than system memory.
Implication	Lost network connection to the node
Status	Fixed
Workaround	If the available memory is less in the system than 180MB then do housekeeping of the NFS cache when transferring files exceeding in size the system memory.

```
echo 4 > /proc/sys/vm/drop_caches
while sleep 10; do echo 3 > /proc/sys/vm/drop_caches; done &
```

It is fixed by software release 1.2 and later versions. It is required to use SWAP partition for each OS based on E3 node.

3. High loss of network packets when small UDP packets are sent and throughput is greater than 2Gbps.

Problem	High loss of network packets when small UDP packets are sent and throughput is greater than 2Gbps. The PCI network interfaces have MTU of 65535 bytes so can accept larger UDP packets and scale performance this way up to 12Gbps.
Implication	Large network packet loss
Status	Will not fix
Workaround	Use close to MTU packet sizes when sending UDP from the VCA nodes.

4. Xeon E3 system is not recognized as valid host for VCA upon RPM installation.

Problem Installation script (`vca_config.sh`) run during modules installation (`vcass-modules-3.10.0-1.*.rpm`) doesn't recognize properly a host with Xeon E3 or pre-production E5 CPU. In such a case, the following warning during installation is displayed:
`Unknown VCA side; configuration skipped`

Implication Soft dependencies are not set in the system and `vop` module isn't loaded, so virtual network to VCA card cannot be established.

Status Fixed

Workaround After failed installation type the following: `vca_setup.sh host`. It is fixed by software release 1.1 and later versions.

5. PLX 8713 device name displayed incorrectly.

Problem PLX 8713 PCIe device names change after a cold (AC or DC) reset (device 8086:2952 (ca) changes to device 10b5:87d0 (rev ca) in `lspci` utility output).

Implication No functional implication; cosmetic only unless a user script is written with hard-coded values.

Status Fix

Workaround Warm reset (reboot) host system.

6. `vcactl` reboot: may fail if path to image is > 256 characters long

Problem If configuration file `.xml` contains path longer than 256 characters for entry "last-os-image", and autoboot is configured, nodes may fail to boot due to the long path.

Implication Nodes may not boot, reducing functionality for end user.

Status Fixed

Workaround Move boot images to a path with fewer total characters in the path (e.g., `/media/vca_images/`). It is fixed by software release 1.3 and later versions

7. User in "vcausers" group cannot run all "vcactl" commands

Problem	command “vcactl config-use” fails for non-root user, even if user is in “vcausers” group, which should allow execution of all vcactl commands without root elevation.
Implication	Cannot instruct VCA to use updated configuration XML file without root escalation (e.g., sudo or su -).
Status	Fixed
Workaround	It is fixed by software release 1.3 and later versions.

8. Autoboot may take nearly 3 minutes to complete with NFS image

Problem	When using NFS based (“persistent”) images and “Autoboot” feature is enabled, automatic boot has been observed to take as long as 160 seconds, due to timing of NFS service startup.
Implication	Booting persistent NFS based images at host boot time may cause “wait timeout” messages. Nodes still boot, but require longer than the defined timeout value in the configuration file.
Status	Fixed
Workaround	It is fixed by software release 1.3 and later versions.

9. Autoboot fails if vcausers group does not have read access to the image directory

Problem	Autoboot function will fail to boot nodes if the directory tree containing the image configured for boot is not readable by “vcausers” group, or alternatively, world-readable.
Implication	User must change group ownership of directory tree containing boot images, or much make the directory tree world-readable.
Status	Fixed.
Workaround	Issue the command as a root-enabled user: <pre>chgrp -R vcausers <directoryname></pre> (e.g., “chgrp -R vcausers /media/vca images”) Add a note to VCA_SoftwareUserGuide, that for auto-boot to work, the directory containing the boot image and all parent directories of that directory must either be world-readable, or be owned by the “vcausers” group.

10. DomU to DomU transfers are significantly slower than DomU – DomX or DomU – bare metal

Problem	Significant (80-85%) performance degradation occurs when transferring data between Dom U virtualized systems compared to bare metal configurations.
Implication	Transfers will take significantly longer in this use case than in anticipated use cases.

Status	No Fix
Workaround	Do not transfer between Dom U instances; instead, transfer between Dom U and bare metal, or between Dom U and Dom X.

11. (Informational) Why does the Linux kernel indicate it is “tainted”?

Problem	Linux kernel indicates (on screen, or in a log file) that the kernel is “tainted”.
Implication	<p>No functional implication; informational notice to the user that the Linux distribution is running in a “non-standard” configuration. The following tables provide additional information:</p> <p>(source: https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/tree/kernel/panic.c retrieved 15 May 2016, lines 289-310)</p> <pre>/** * print_tainted - return a string to represent the kernel taint * state. * * 'P' - Proprietary module has been loaded. * 'F' - Module has been forcibly loaded. * 'S' - SMP with CPUs not designed for SMP. * 'R' - User forced a module unload. * 'M' - System experienced a machine check exception. * 'B' - System has hit bad_page. * 'U' - Userspace-defined naughtiness. * 'D' - Kernel has oopsed before * 'A' - ACPI table overridden. * 'W' - Taint on warning. * 'C' - modules from drivers/staging are loaded. * 'I' - Working around severe firmware bug. * 'O' - Out-of-tree module has been loaded. * 'E' - Unsigned module has been loaded. * 'L' - A soft lockup has previously occurred. * 'K' - Kernel has been live patched. * * The string is overwritten by the next call to print_tainted(). */</pre> <p>(source: https://www.kernel.org/doc/Documentation/sysctl/kernel.txt retrieved 15 May 2016)</p> <p>tainted:</p>

Non-zero if the kernel has been tainted. Numeric values, which can be ORed together:

- 1 - A module with a non-GPL license has been loaded, this includes modules with no license.
Set by modutils >= 2.4.9 and module-init-tools.
- 2 - A module was force loaded by insmod -f.
Set by modutils >= 2.4.9 and module-init-tools.
- 4 - Unsafe SMP processors: SMP with CPUs not designed for SMP.
- 8 - A module was forcibly unloaded from the system by rmmod -f.

- 16 - A hardware machine check error occurred on the system.
- 32 - A bad page was discovered on the system.
- 64 - The user has asked that the system be marked "tainted". This could be because they are running software that directly modifies the hardware, or for other reasons.
- 128 - The system has died.
- 256 - The ACPI DSDT has been overridden with one supplied by the user instead of using the one provided by the hardware.
- 512 - A kernel warning has occurred.
- 1024 - A module from drivers/staging was loaded.
- 2048 - The system is working around a severe firmware bug.
- 4096 - An out-of-tree module has been loaded.
- 8192 - An unsigned module has been loaded in a kernel supporting module signature.
- 16384 - A soft lockup has previously occurred on the system.
- 32768 - The kernel has been live patched.

Status No Fix

Workaround Not required.

12. Time Zone is not automatically synchronized between host and nodes.

Problem If time synchronization is important for tools used by the end user, the lack of time zone synchronization may cause these tools to not function correctly or return errors.

Implication Time zones not in synch between the host, Intel® Visual Compute Accelerator nodes, and (if present) virtual machines running on the Intel® VCA nodes may cause application/utility/tool failure.

Status Fix

Workaround If time zone synchronization is important in installed environment, script time zone update to be run after Intel® VCA node OS boot.

13. Node becomes unresponsive when using NFS and heavy memory load.

Problem Intel® Visual Compute Accelerator stops responding to external input, appears hung. Issue root caused to memory starvation and un-cleared caches.

Implication One or more nodes may stop processing workloads, and appear to be offline, significantly reducing overall throughput performance of the card.

Status Fix

Workaround Run the following scripts on the host and nodes after booting nodes, but before starting workloads on the nodes if it is suspected that the nodes may enter a low memory condition, and NFS is being used at all by the nodes:

host.sh

```
#!/bin/bash
```

```
echo 4 > /proc/sys/vm/drop_caches
while sleep 10; do echo 3 > /proc/sys/vm/drop_caches; done &
```

node.sh

```
#!/bin/bash
echo 4 > /proc/sys/vm/drop_caches
while sleep 1; do echo 3 > /proc/sys/vm/drop_caches; done &
```

14. CATERR on node requires full power cycle of host to clear.

Problem If a node on an Intel® Visual Compute Accelerator experiences a catastrophic error (CATERR), it can only be cleared by performing an AC power cycle of the host.

Implication Because more than one Intel® VCA may be installed in a host, performing an AC power cycle on the host may bring offline nodes that were still in functional state, reducing overall data processing until they are brought back online.

Status No Fix

Workaround Perform full power cycle of host (remove power from power supplies, wait 30 seconds, restore power to power supplies, power on host).

15. Hosts running Intel® Xeon® E5-2600 v4 processor family must add “nosmap” to kernel boot line.

Problem If the host system is running the latest Intel® Xeon® E5-2600 v4 processor family and has current Intel® Visual Compute Accelerator card(s) installed, the Linux* kernel will fail to boot during PCIe enumeration.

Implication Host system is unusable if customer has Intel® Visual Compute Accelerator card(s) installed and upgrades host to Intel® Xeon® E5-2600 v4 processor family processors.

Status Fixed.

Workaround Update GRUB boot loader configuration to include kernel option “nosmap”. Process is outlined as follows:

- 1) Log in as sudo enabled user
- 2) Issue the following commands:
 - a. `cd /boot/grub2`
 - b. `mv grub.cfg grub.cfg.orig`
 - `# edit /etc/default/grub`
 - `# Find the line that begins with "GRUB_CMDLINE_LINUX", and add "nosmap" immediately before the final " character on the line`
 - `#Save and exit`
 - c. `grub2-mkconfig -o /boot/grub2/grub.cfg`
 - d. inspection of the new /boot/grub2/grub.cfg file should show "nosmap" added to the boot line (which begins with "linux16")

It is fixed by software release 1.2 and later versions.

16. An Intel Hyper-Threading Technology Enabled Processor may exhibit internal parity errors or unpredictable system behavior.

Problem Under a complex series of micro architectural events while running Hyper-Threading Technology, a correctable internal parity error or unpredictable system behavior may occur.

Implication A correctable error (IA32_MC0_STATUS.MCACOD=0005H and IA32_MC0_STATUS.MSCOD=0001H) may be logged.

Status Fix.

Workaround Need to update VCA BIOS to 107 if use Media Server Studio version 16.5 and future release

Run the command below as root

```
ssh <node IP> dmidecode -t bios|grep ACGC
```

If the number after string "ACGC" is lower than 106, then it's recommend to update the bios to 107.

Follow the steps below to upgrade the BIOS.

Note: The VCA1283LVV board (all nodes) will reboot during this process

- Download the latest BIOS (<https://downloadcenter.intel.com/download/25590>)
- Unzip Intel®VisualComputeAccelerator_BIOS_<version>.zip into a temporary location on the host.
- Make sure BIOS image selection jumpers are correctly placed (refer to Section BIOS Image Selection Jumpers in VCA_Spec_HW_Users_Guide)
- Run "vcactl reset"
- Run "vcactl wait-BIOS"
- Run "vcactl update-BIOS VCA-BIOS_0ACGC<version>.img"
- Run "vcactl wait-BIOS"
- After running the command above, the screen should show that the BIOS is "UP"
- To verify if the BIOS was successfully updated, run "ssh <node_IP> dmidecode -t bios|grep ACGC"
- The string after "ACGC" should now show a number greater than 106.

On success, BIOS upgrade is finished. Host reboot is NOT required with this upgrade.

The BIOS is available for download at <https://downloadcenter.intel.com/download/25590>

Media Server Studio Errata Items impacting VCA

Review the release notes of the Intel® Media Server Studio for a full list of errata items. The following are errata items not yet documented in the Intel® MSS release notes and impacting VCA.

1. OpenCL BufferReadWrite certification test fails if running from the virtual machine (GuestOS) under Xen hypervisor.

Problem	OpenCL BufferReadWrite certification test fails if running from the virtual machine (GuestOS) under Xen hypervisor.
Implication	OpenCL certification test fails
Status	Fix
Workaround	None

2. If the VCA node runs of system memory then GPU crashes upon decode that creates resulting file larger than available memory in the system.

Problem	If the VCA node runs of system memory then GPU crashes upon decode that creates resulting file larger than available memory in the system.
Implication	GPU will crash
Status	Fix
Workaround	None

Documentation Updates
