

Intel® QuickAssist Technology (Intel® QAT) SR-IOV driver for VMware ESXi*

Release Notes

Package Version: QAT17_SRIOV.VMW.700.1.1.0-7

August 2020



No license (express or implied, by estoppel or otherwise) to any Intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and noninfringement, as well as any warranty arising from a course of performance, course of dealing, or usage in trade.

This document contains information on products, services, and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications, and roadmaps.


The products and services described may contain defects or errors, known as errata, which may cause deviations from published specifications. Currently characterized errata are available on request.

Copies of documents that have an order number and are referenced in this document may be obtained by calling 1 800 548 4725 or by visiting www.Intel.com/design/literature.htm.

Intel, Xeon, and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

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

Copyright © 2020 Intel Corporation. All rights reserved



Contents

1.0	Description of this Release	5
1.1	Supported Platforms	5
1.2	Supported Operating Systems.....	5
1.3	Audience	5
1.4	Environmental Assumption	5
1.5	Terminology	6
1.6	References and Resources	6
2.0	Getting Started.....	7
2.1	Installing Intel® QAT VMware Driver.....	7
2.2	Pass-through the VF PCI Device to VM	8
2.3	Installing Intel® QAT Software on the Guest	8
2.4	Uninstalling the QAT VMware Driver.....	8
2.5	Recover from Fatal Errors.....	8
2.6	Customer Support	9
3.0	Known Issues	10
3.1	Issues Relating to Violation of Trust Model.....	10
3.1.1	QATE-7495 - An Incorrectly formatted requests to Intel® QAT can render the Intel® QAT endpoint unresponsive.....	10
3.1.2	QATE-30251 - Turning off Bus Master Enable may cause PF hang	10
3.2	Other Limitations	11
3.2.1	QATE-64395 - Usage of DC Session Update API can render the application unresponsive.....	11
3.2.2	VQQ-122 – Intel® QAT HW doesn't support "Number of VFs" SR-IOV configuration	11
3.2.3	Intel® QAT HW doesn't support VF reset functionality.....	12
3.2.4	Intel® QAT HW requires masking some errors in AER register due to HW limitations.....	12
3.2.5	VMware ESXi may require to manually toggle passthrough for Intel® QAT VFs	12

Tables

Table 1.	Terminology.....	6
Table 2.	Additional Documentation	6



Revision History

Revision Number	Description	Date
1.0	Initial Release	August 2020

§



1.0 Description of this Release

This software enables single root I/O virtualization (SR-IOV) for the Intel® QuickAssist Technology (Intel® QAT) on VMware ESXi v7.0.0. SR-IOV enables the creation of Virtual Functions (VF) from a single Intel® QAT device to support acceleration for multiple virtual machines.

For instructions on loading and running the release software, see Section [2.0 Getting Started](#) of this document.

These release notes may also include known issues with third-party or reference platform components that affect the operation of the software.

Note: This driver is fully compatible with Linux* v4.9 release and is the minimum required.

1.1 Supported Platforms

- Intel® Xeon® Scalable Platform with Intel® C62x Chipset
- Intel® QuickAssist Adapter 8960/Intel® QuickAssist Adapter 8970 (formerly known as “Lewis Hill”)

1.2 Supported Operating Systems

The software in this release has been validated against the following Operating Systems:

- VMware ESXi* v7.0.0

1.3 Audience

The user should be familiar with virtualization technologies and has some level of familiarity with the VMware ESXi hypervisor.

Note: This document does not explain how to install VMware ESXi, how to install a virtual machine, and how to administrate ESXi using vSphere Client. For more details on these topics, please refer to VMware’s documentation.

1.4 Environmental Assumption

The following assumptions are made concerning the deployment environment:

- The driver object/executable file on the disk should be protected using the normal file protection mechanisms so that it is writable only by trusted users, for example, a privileged user or an administrator.
- The public key firmware image on the disk should be protected using normal file protection mechanisms so that it is writable only by trusted users, for example, a privileged user or an administrator.
- The Intel® QAT device should not be exposed (via Single-root Input/Output Virtualization (SRIOV)) to untrusted guests.
- The Intel® QAT device should not be exposed (via the “user space direct” deployment model) to untrusted users.

Description of this Release



- DRAM is considered to be inside the trust boundary. The typical memory-protection schemes provided by the Intel architecture processor and memory controller, and by the operating system, prevent unauthorized access to these memory regions.
- Persistent keys were not considered, but the storage media are considered inside the cryptographic boundary

1.5 Terminology

Table 1. Terminology

Term	Definition
API	Application Programming Interface
BDF	Bus Device Function
Intel® QAT	Intel® QuickAssist Technology
Intel® VT	Intel® Virtualization Technology
IPS	Intel® Premier Support
PVSP	Partner Verified and Supported Products
PCI	Peripheral Connect Interface
PF	Physical Function
SSH	Secure Shell
SR-IOV	single root I/O virtualization
UI	User Interface
VF	Virtual Functions
VM	Virtual Machine


1.6 References and Resources

Refer to the following open-source website URL for the associated software and collateral:

<https://01.org/intel-quickassist-technology>

Table 2. Additional Documentation

Title	Location/Document Number
<i>Using Intel® Virtualization Technology (Intel® VT) with Intel® QAT Application Note</i>	330689
<i>Intel® QuickAssist Technology Software for Linux Getting Started Guide</i>	336212
<i>Intel® QuickAssist Technology Software for Linux Software Programmer's Guide</i>	336210
<i>Intel® QuickAssist Technology Cryptographic API Reference Manual</i>	330685



2.0 Getting Started

The following section details the instructions on how to install and uninstall the Intel® QAT VMware driver and the Intel® QAT software guest as well as a guide as to how to recover from potential errors.

2.1 Installing Intel® QAT VMware Driver

1. Open a Secure Shell (SSH) connection to the target ESXi host (ESXi shell and SSH needs to be enabled to perform this operation).
2. Copy the component bundle to the ESXi server.

Technically, you can place the file anywhere that is accessible to the ESXi console shell, but for these instructions, we'll assume the location is in '/tmp.'

Here's an example of using the Linux '> scp' utility to copy the file from downloaded location to the remote/target ESXi server located at 10.10.10.10:

```
> scp QAT17_SRIOV.VMW.700.1.1.0-*.tar.gz root@10.10.10.10:/tmp
```

3. Extract the package (assuming it has been copied to /tmp folder)

```
> cd /tmp
> tar -xzf QAT17_SRIOV.VMW.700.1.1.0-*.tar.gz
```

4. Set host acceptance level to "PartnerSupported":

```
> esxcli software acceptance set --level PartnerSupported
```

5. There are two ways to install the driver:

- a. From component via the following command (recommended approach):

```
> esxcli software component apply -d /tmp/quickassist/DRIVER/ESXi7.0/icp-qat-pf-drv_1.1.0.7-15843807_16673147.zip
```

- b. From the VIB package extracted from component:

```
> unzip /tmp/quickassist/DRIVER/ESXi7.0/icp-qat-pf-drv_1.1.0.7-15843807_16673147.zip
> esxcli software vib install --maintenance-mode -v tmp/vib20/icp-qat-pf/INT_bootbank_icp-qat-pf_1.1.0.7-1OEM.700.1.0.15843807.vib
```

6. Reboot the system:

```
> reboot
```

7. If the Intel® QAT driver has been loaded without errors, you should see the `icp_qat_pf` module showing in the list of system modules:

```
> esxcfg-module -l | grep icp_qat_pf
> icp_qat_pf          3      1972
```

8. Now the Intel® QAT VFs are enabled in the system. You can verify this by running the `lspci` command. For example:

```
> lspci -vn | grep 37c9
0000:85:01.0 Class 0b40: 8086:37c9 [PF_0.133.0_VF_0]
0000:85:01.1 Class 0b40: 8086:37c9 [PF_0.133.0_VF_1]
0000:85:01.2 Class 0b40: 8086:37c9 [PF_0.133.0_VF_2]
0000:85:01.3 Class 0b40: 8086:37c9 [PF_0.133.0_VF_3]
0000:85:01.4 Class 0b40: 8086:37c9 [PF_0.133.0_VF_4]
0000:85:01.5 Class 0b40: 8086:37c9 [PF_0.133.0_VF_5]
0000:85:01.6 Class 0b40: 8086:37c9 [PF_0.133.0_VF_6]
0000:85:01.7 Class 0b40: 8086:37c9 [PF_0.133.0_VF_7]
0000:85:02.0 Class 0b40: 8086:37c9 [PF_0.133.0_VF_8]
0000:85:02.1 Class 0b40: 8086:37c9 [PF_0.133.0_VF_9]
```



```
0000:85:02.2 Class 0b40: 8086:37c9 [PF_0.133.0_VF_10]
0000:85:02.3 Class 0b40: 8086:37c9 [PF_0.133.0_VF_11]
0000:85:02.4 Class 0b40: 8086:37c9 [PF_0.133.0_VF_12]
0000:85:02.5 Class 0b40: 8086:37c9 [PF_0.133.0_VF_13]
0000:85:02.6 Class 0b40: 8086:37c9 [PF_0.133.0_VF_14]
0000:85:02.7 Class 0b40: 8086:37c9 [PF_0.133.0_VF_15]
```

At this point the Intel® QAT VFs can be attached to a guest Virtual Machine (VM).

2.2 Pass-through the VF PCI Device to VM

1. Connect to the target ESXi host via Web User Interface (UI).
2. In the left pane, click on **VMs**.
3. In the center pane, click on the desired **VM**.

Note: Ensure that the VM is powered off.

4. Click on the **"Edit"** button to edit the VM settings.

A pop-up window with the VM settings will appear.

5. Click on **"Add other Device"** and select **"PCI device"**.

The new PCI device will be added. By default, it selects the first Intel® QAT VF. To select a specific VF, click the drop-down list. The Bus Device Functions (BDFs) listed here will match with the output of the `"lspci -vn | grep 37c9"` command. Additional VFs can be added by repeating this step.

6. Click **"Save"**.

Now you have one or more VFs attached to your guest.

2.3 Installing Intel® QAT Software on the Guest

For instructions on how to install the VF driver on a Linux guest, refer to [Table 2, Using Intel® Virtualization Technology \(Intel® VT\) with Intel® QAT Application Note](#).

2.4 Uninstalling the QAT VMware Driver

1. Open an SSH connection to the target ESXi host.
2. Actual process depends on the way how the driver was installed on system:

- a. If it was installed as component (recommended) - run next command to remove driver component:

```
> esxcli software component remove -n icp-qat-pf-driv
```

- b. If it was installed as VIB package – run next command to remove VIB:

```
> esxcli software vib remove -n icp-qat-pf --maintenance-mode
```

3. Reboot system to complete removal:


```
> reboot
```

2.5 Recover from Fatal Errors

In the event of a persistent device error state that cannot be recovered by software, it is recommended to manually reload the PF driver on the ESXi host or reset the host itself.

The driver resets and recovers the Physical Function (PF) device during driver reloading.

The following are instructions to reload the driver:



1. Power off all the VMs that are using Intel® QAT hardware.
2. Execute the next commands to reload the PF driver and `devmgr`:

```
> esxcfg-module -u icp_qat_pf  
> kill -HUP $(cat /var/run/vmware/vmkdevmgr.pid)
```

2.6 Customer Support

The Intel® QAT driver supporting ESXi is certified under VMware's Partner Verified and Supported Products (PVSP) certification program. Under the PVSP program, Intel acts as the primary support team for all Intel® QAT related issues. To submit an Intel® Premier Support (IPS) ticket, please go here:

<https://www.intel.com/content/www/us/en/secure/my-intel/dashboard.html>

VMware forwards all issues they suspect to be related to Intel® QAT to Intel to help triage and resolve with the customer directly.

§



3.0 Known Issues

3.1 Issues Relating to Violation of Trust Model

The second generation of Intel® QAT was designed with performance as the primary objective. To attain the best possible performance, applications are exposed directly to the hardware with no bounds checking. This approach implies a trusted programming model, wherein an application is expected to supply correctly formatted addresses and arguments at the Application Programming Interface (API).


An application failing to follow the programming conventions runs the risk of negatively impacting the platform.

3.1.1 QATE-7495 - An Incorrectly formatted requests to Intel® QAT can render the Intel® QAT endpoint unresponsive

Title	An Incorrectly formatted request to Intel® QAT can render the Intel® QAT endpoint unresponsive
Reference #	QATE-7495
Description	This version of the Intel® QAT hardware does not perform exhaustive request address and parameter checking. It follows that a malicious application could submit requests that can bring down an entire Intel® QAT endpoint or platform itself, which can impact other Intel® QAT jobs associated with the hardware. This presents a risk that must be managed by the host and guest operating systems and other system policies. The exposure can extend to other guest operating systems or applications outside of the typical access boundary of the malicious guest or application. Conditions like invalid address, address range that crosses the SecureRAM boundaries, and invalid request can cause HW to hang or system reset.
Implication	All guest operating systems and applications using QAT must be trusted, and/or other steps must be taken to ensure that an untrusted application or guest cannot submit incorrectly formatted requests.
Resolution	There is no workaround available. However, system policies (including limiting specific operating system permissions) can help to mitigate this issue.
Affected OS	All
Driver/Module	CPM IA - Common
Related Issues	<p>QATE-14706 - Partial descriptor submission may cause hang QATE-14287 - IOMMU page fault provokes device hang QATE-15430 - Malformed NULL descriptor may cause hang QATE-30895 - Crossing SecureRAM boundaries may cause device hang QATE-39377 - Continuous submitting malformed requests in VM may cause the platform to hang or reboot</p> <p>Root Cause Analysis: <i>When a non-posted transaction is initiated to an invalid target (bad memory address), a UR is returned. Later, a Completion Time Out happens. A tag is issued to the non-posted transaction, and a tag is returned by the UR and a tag is returned by the CTO. This means that for every one tag issued, two are returned. This causes unexpected overflows in counters, too many outstanding transactions, and eventually leads to system instability and a crash.</i></p>

3.1.2 QATE-30251 - Turning off Bus Master Enable may cause PF hang

Title	Turning off Bus Master Enable may cause PF hang
Reference #	QATE-30251
Description	Specific guest's operations to rings with disabled BME bit may cause PF to hang
Implication	-



Title	Turning off Bus Master Enable may cause PF hang
Resolution	If PF hangs, the system administrator should shut down all the VMs and manually reload the driver or restart the whole system.
Affected OS	ESXi
Driver/Module	HW

3.2 Other Limitations

3.2.1 QATE-64395 - Usage of DC Session Update API can render the application unresponsive

Title	Usage of DC Session Update API can render the application unresponsive
Reference #	QATE-64395
Description	In case of using Linux driver 4.10 with DC Session Update API user may get a timeout and fatal errors in Guest OS <code>dmesg</code> : <pre>c6xxvf 0000:ff:00:0: Fatal error received from PF 0x6ac20013</pre>
Implication	The application which is using the mentioned API may get a timeout-related error or stuck on waiting for a response from HW.
Resolution	Use smaller chunks for submission or increase timeout values in the application that is using Intel® QAT. Also, it's possible to increase a Heartbeat and Quiesce timeouts for PF driver itself via the next steps: <ol style="list-style-type: none"> 1. Power off all VMs that is using Intel® QAT hardware 2. Unload driver: <pre>> vmkload_mod -u icp_qat_pf</pre> 3. Load driver module via next command where <code>hb_interval</code> is ranged from 500 to 5000 ms and <code>quiesce_timeout</code> is from 350 to 20000 ms: <pre>> vmkload_mod icp_qat_pf hb_interval=5000 quiesce_timeout=20000</pre> 4. Reset device manager: <pre>> kill -HUP \$(cat /var/run/vmware/vmkdevmgr.pid)</pre> <p>To reset <code>hb_interval</code> and <code>quiesce_timeout</code>, just repeat all steps omitting setting time values on step 3.</p>
Affected OS	All
Driver/Module	<code>icp_qat_pf</code>

3.2.2 VQQ-122 – Intel® QAT HW doesn't support "Number of VFs" SR-IOV configuration

Title	Intel® QAT HW doesn't support "Number of VFs" SR-IOV configuration
Reference #	VQQ-122
Description	If the system administrator configures the number of VFs less than the total number of VFs supported by PF, the driver will fail to attach the device.
Implication	The system administrator can't configure the VF number less than the total VFs.
Resolution	Go with the default configuration, and it will enable all VFs or configure the number of VFs equal to the total number of supported VFs per PF.



Title	Intel® QAT HW doesn't support "Number of VFs" SR-IOV configuration
Affected OS	All
Driver/Module	HW

3.2.3 Intel® QAT HW doesn't support VF reset functionality


Title	Intel® QAT HW doesn't support VF reset functionality
Reference #	-
Description	The Intel® QAT HW doesn't implement the SR-IOV specification section, which requires VFs to support Function Level Reset (FLR).
Implication	-
Resolution	ESXi PF SR-IOV driver emulates VFRLR.
Affected OS	All
Driver/Module	HW

3.2.4 Intel® QAT HW requires masking some errors in AER register due to HW limitations

Title	Intel® QAT HW requires masking some errors in AER register due to HW limitations
Reference #	-
Description	QAT HW does not process Completion Timeout, Unsupported Request, and Uncorrectable Internal Errors correctly, and the associated bits should be masked in the AER mask register to prevent NMI failures which may lead to platform crash.
Implication	-
Resolution	ESXi PF SR-IOV driver masking appropriate errors to mitigate platform crashes.
Affected OS	All
Driver/Module	HW

3.2.5 VMware ESXi may require to manually toggle passthrough for Intel® QAT VFs

Title	Intel® QAT HW requires masking some errors in AER register due to HW limitations
Reference #	-
Description	Due to limitations in VMware ESXi v7.0 Intel® QAT VFs could be not marked for passthrough and require to enable passthrough for VFs in vSphere UI manually.
Implication	System administrator need to manually toggle passthrough for VFs before assigning to VMs
Resolution	Update to Update 1 or newer release for v7.0.0 or follow next instruction to toggle passthrough manually: <ol style="list-style-type: none"> 1. Connect to the target ESXi host via Web User Interface (UI) 2. In the left pane, click on Manage. 3. Choose Hardware tab. 4. Using checkboxes select Intel® Co-processor devices that have Disabled passthrough state. 5. Click on Toggle passthrough button to enable passthrough for disabled devices.
Affected OS	VMware ESXi v7.0.0 without Update 1 or any future update.



Title	Intel® QAT HW requires masking some errors in AER register due to HW limitations
Driver/Module	-

§