



# Intel® Cache Acceleration Software for Linux\*

---

*Release Notes*

*February 2018*

*Release 3.5.1 (GA)*



**Intel® CAS for Linux\* v3.5.1 (GA)**

**Notice:** This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Intel, the Intel logo, Intel® Solid State Drive Data Center S3700 Series, Intel® Solid State Drive Data Center P3700 Series, are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2018 Intel Corporation. All rights reserved.



# Contents

---

<b>1</b>	<b>Release Description.....</b>	<b>5</b>
1.1	Features.....	5
1.2	Supported SSD Drives.....	5
1.3	Limitations.....	6
1.4	Supported Platforms.....	7
<b>2</b>	<b>Known Issues.....</b>	<b>8</b>
<b>3</b>	<b>Resolved Issues.....</b>	<b>9</b>



## Revision History

Document Revision	Software Release Description	Software Release Date
001B	Initial Beta release of document.	January 2013
001	Initial public release of document.	February 2013
002	Documentation Updates for v2.1	May 2013
003	Documentation Updates for v2.5	August 2013
004	Documentation Updates for v2.6	December 2013
005	Documentation Updates for v2.6.1	April 2014
006	Documentation Updates for v2.7	May 2014
007	Documentation Updates for v2.7 GA	July 2014
008	Documentation Updates for v2.7.1 GA	October 2014
009	Documentation Updates for v2.8 GA	December 2014
010	Documentation Updates for v2.9 GA	June 2015
011	Documentation Updates for v3.0 GA	December 2015
012	Documentation Updates for v3.1 GA	May 2016
013	Documentation Updates for v3.5 GA (QEMU)	April 2017
014	Documentation Updates for v3.5 Beta	June 2017
015	Documentation Updates for v3.5 GA	July 2017
016	Documentation Updates for v3.5.1 Beta	December 2017
017	Documentation Updates for v3.5.1 GA	January 2018
018	Documentation Updates for v3.5.1 GA (Solid Driver/RPM Update)	February 2018



# 1 Release Description

---

This document describes extensions and deviations from the release functionality described in the *Intel® Cache Acceleration Software (Intel® CAS) for Linux\* v3.5.1 (GA) Administrator Guide*.

## 1.1 Features

3.5.1:

**CAS Startup and Udev rules interaction** – The intelcas service was retired in favor of using udev rules to handle startup of caching. CAS is now enabled by default. Udev rules survey block devices in the system; if one or more of those devices is configured for caching, Udev rules perform actions to activate those devices. Also, handling of broken/failed devices has been improved.

3.5:

**TRIM Support** - TRIM is enabled by default and supported by CAS for Intel® Solid State Drives (SSDs). This helps improve performance especially in WB (write-back) mode because CAS does not have to flush the data that is no longer needed to the backend storage.

**Atomic Writes** - This feature is available only for Intel Data Center SSDs, such as the DC P3600 or DC P3700 Series. It allows CAS to write user data and cache metadata in one write request, reducing number of writes to cache devices. This has a positive impact on performance and drive endurance.

**Warm Cache** – This feature provides faster recovery after an un-planned shutdown. In previous versions, the user had to recover the cache and then start the cache without any data cached.

## 1.2 Supported SSD Drives

Any Linux flash device (SAS, SATA, PCIe\*, Fibre Channel, RAID) is supported, and can be direct attached, expander attached, or attached via SAN (with a single worker).

Additionally, the following Intel® SSDs have been fully qualified:

- Intel® SSD DC S3700 Series
- Intel® SSD DC P3700 Series



## 1.3 Limitations

This release has the following known limitations:

- An "ext3" file system is limited to 16 TiB max core drive capacity (this is an intrinsic limitation of the "ext3" file system, not an Intel CAS limit.).
- For VMs, paravirtualized drivers are not supported.
- System sleep (S3) and hibernation (S4) power states and resumption from these states is not supported.
- Hot-plugging of either core or cache drive is not supported.
- For best performance, it is recommended to use the *noop* IO scheduler on the cache device (SSD).
- NUMA node IO transfers can result in performance degradation on multi-CPU platforms with high number of worker threads, if the threads send IO to a CPU that is not directly attached to the PCIe\* device that is the target of the IO. This is a known limitation of multi-CPU platform architecture and not of Intel CAS.
- Linux file systems with less than 4096 byte (4KiB) block sizes may encounter less than ideal performance when workload mix has a high percentage of writes in write-back mode.
- Sub-partitioning of an intelcas virtual block device (e.g. /dev/intelcas1-1) that is built on an existing partition of the core device (eg. /dev/sdc1) is not supported.
- Stack overrun can occur in kernels with 8KiB stack size that are caching to an NVMe\* SSD, when IO payloads have gaps. This is due to Linux implementation of NVMe\* IO using Physical Region Page (PRP) lists, which require there to be no gaps in the IO payload. When a payload is received with IO gaps, the kernel splits it into many sub IOs that meet PRP list requirements. As these IOs complete, it is possible to exceed the 8KiB stack size limit, which would result in a stack overrun. This is a known limitation of Linux NVMe implementation and not of Intel® CAS. To avoid this issue, it is recommended to use a kernel with 16KiB stack size, such as RHEL 6.7 or newer.
- With older versions of LVM, you must first create a partition on the core device (for example, if the core device is /dev/sdc, you must create /dev/sdc1) prior to accelerating that device with Intel CAS device and creating the physical volume or creation of the logical volume will fail. If you see the following warning: *"WARNING: Ignoring duplicate config node: types (seeking types); Found duplicate PV"*, then you must use this workaround.
- When utilizing the Atomic Writes feature, you must reboot the servers after formatting the NVMe devices and **before any further changes are made to the system.**
- When using large RAID volume starting cache on such device could take longer time than usual. This is not a CAS issue, but is related to how RAID device handles DISCARD/TRIM requests.
- In SysV based Linux distributions (ex. CentOS/RHEL 6.x) if the filesystem created on the CAS device is being used by a user process, the CAS service is not able to umount that filesystem and CAS cannot be stopped. In these distributions there are no methods to provide dependencies between various services for proper ordering of startup/shutdown. To prevent startup issues after this situation we recommended using the load option in intelcas config file. This issue does not affect SystemD based distros.
- It is not recommended to create a file system on an intelcas device.
- Certain methods for creating partitions on devices (CentoOS/RHEL 6.x) may result in a partition alignment message concerning possible performance degradation. This message may be disregarded.



## 1.4 Supported Platforms

The following table lists the supported platforms for 64-bit processors. The supported QEMU version is 2.1.2 for CentOS\* 7.2.

**Table 1-1 Supported Platforms**

Operating System	Kernel
Red Hat* Enterprise Linux* (RHEL*) 6.7	x86_64, Kernel 2.6.32-573
Red Hat* Enterprise Linux* (RHEL*) 6.8	x86_64, Kernel 2.6.32-642
Red Hat* Enterprise Linux* (RHEL) 7.2	x86_64, Kernel 3.10.0-327
Red Hat* Enterprise Linux* (RHEL) 7.3	X86_64, kernel 3.10.0-514
CentOS* 6.7	x86_64, Kernel 2.6.32-573
CentOS* 6.8	x86_64, Kernel 2.6.32-642
CentOS* 7.2	x86_64, Kernel 3.10.0-327
CentOS* 7.3	x86_64, Kernel 3.10.0-514
SUSE* Linux* Enterprise Server (SLES*) Version 12 SP2	x86_64, Kernel 4.4.74-92.38
Ubuntu Server 14.04.5	x86_64, Kernel 4.4.0-59
Ubuntu Server 16.04	x86_64, Kernel 4.4.0-28
Other distros – Intel CAS will install and compile from source on other distros and other kernels, but the user may be required to reproduce any issues on a validated distro & kernel to receive support.	Other kernels



## 2 Known Issues

The following are current known issues:

**Table 2-1 Known Issues**

Reference	Description	Impact/Status/Workaround
1805729730	Report generation when multiple errors occur	If there are multiple conditions preventing casadm from starting cache on device, only the first one is reported. For example, if the device contains partitions or filesystem and there is not enough memory to start the cache, only the first one is printed, and information about low memory is missing.
1805170798	Soft lockup in metadata flushing	Sporadically, in some specific I/O traffic (very small I/O requests interleaved with larger requests), CPU soft lockup may be reported in system logs but metadata flushing will proceed correctly.  This is a minor issue with no work-around.
1805941826	Incorrect installer output when upgrading in-flight from 3.5.1 RPM to enterprise installer	User is presented with an unusable error message rather than an advisory to clean the CAS configuration prior to in-flight upgrade.



## 3 Resolved Issues

---

The following table lists issues Intel has resolved or closed since the previous release.

**Table 3-1 Resolved Issues**

Reference	Description	Impact/Status/Workaround
98942	Config statistics for core display wrong size, when partition size is not multiplication of cache line size.	Resolved