

Technical Advisory

Intel® RSTe Problem Report —

Intel® Rapid Storage Technology enterprise: RAID 5 Data Loss Exposure

October 2, 2015

Description

Table 1: Affected Intel Products

Product	Product (Driver) Versions
Intel® Rapid Storage Technology enterprise Driver	All Windows* based production driver releases of the following baselines: <ul style="list-style-type: none">• All currently available up to and including RSTe 4.2.0.1143 (Please contact your platform OEM or ODM for details on obtaining this version)

Problem Statement

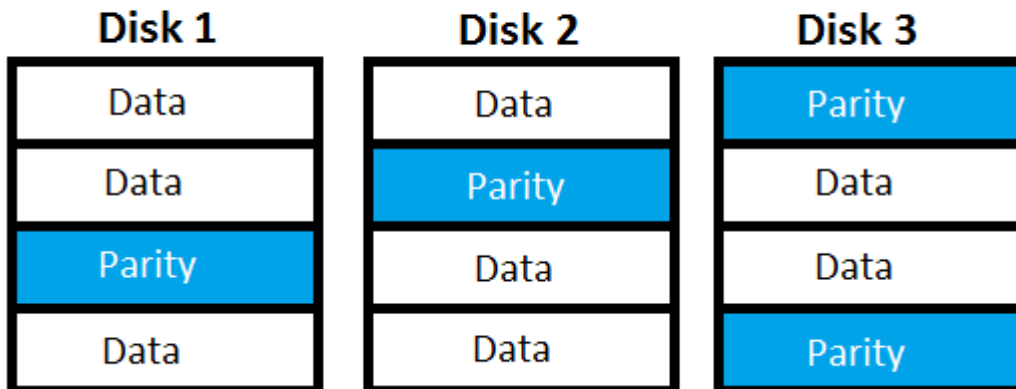
A potential data loss condition exists when an Intel® Rapid Storage Technology enterprise (Intel® RSTe) managed RAID 5 volume becomes degraded. The data loss exposure only exists after a RAID 5 volume becomes degraded and there is sustained heavy I/O (with I/O buffering disabled) for more than approximately 60 minutes.

Root Cause

When an RSTe RAID5 volume resides on an N-disk matrix in normal mode, the RAID manager writes data blocks to N-1 disks and then computes a parity (XOR from data blocks) based off of the data written. The parity is then written drive N. The following is a simple diagram showing this on a 3 driver RAID 5 volume.




Raid 5 volume "Normal" mode



When one of the RAID5 volume member disks fails, the RAID5 volume goes to degraded state. When processing 'read data' commands the RSTe RAID manager is able to recover the lost data from the healthy disks (data and/or parity) and is able to respond with the complete data blocks back up to the OS.

When processing 'write data' commands the RSTe RAID manager uses the RAID5 logic to write the data and parity blocks, skipping writes to failed drive.

Raid 5 volume "Degraded" mode



In the failing scenario, a user runs a heavy IO traffic with buffering disabled, a race condition may occur when accessing the pointer to data blocks processing 'write data' commands. When the error occurs the correct pointers may be inadvertently overwritten and as a result incorrect data may be written to a disk.

This exposure to data loss does occur in any other RAID5 states, only degraded.

Impact of Defect

Windows has buffering enabled by default. In order to expose the data loss condition, this buffer option must be manually disabled. Another condition that can expose this condition is using an application that provides an option to configure buffering and that option is set to not use buffering. Under these conditions, if a RAID 5 volume becomes degraded, after 60-plus minutes of heavy I/O, the platform could encounter data loss.



Workaround

1. Ensure that buffering is always enabled.
2. When a RAID5 volume becomes degraded, replace the failed drive immediately and start the rebuild process as quickly as possible.
3. Add another drive to the system and mark that drive as a Hot Spare. This will result in the degraded RAID 5 transitioning to a rebuild state in a very short period of time (thereby closing the exposure window).

Corrective Action/Resolution

Intel has resolved this issue and has released the following packages (Please contact your platform OEM or ODM for details on obtaining these driver versions):

1. For Windows Server 2003 and Server 2008
 - a. **Driver version number:** RSTe 4.0.2.1019
2. For Windows Vista and XP (64 bit)
 - a. **Driver version number:** RSTe 4.1.2.1011
3. For Windows 8 (32bit and 64bit)
 - a. **RSTe Kit version number:** RSTe 4.2.2.1005
4. For Windows 7, 8.1, Server 2008 R2, Server 2012, Server 2012 R2 and Windows 10
 - a. **RSTe Kit version number:** RSTe 4.3.0.1223



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.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

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

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting [Intel's Web Site](#).

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries

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

Copyright © 2015, Intel Corporation. All rights reserved.