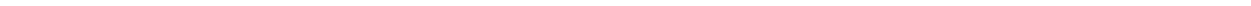




Intel RAID Volume Recovery Procedures



Revision 1.0
12/20/01



Revision History

Date	Revision Number	Modifications
12/20/01	1.0	Release

Disclaimers

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.

The Intel server boards and chassis 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.

Copyright © 2001, Intel Corporation. Intel® is a trademark of Intel Corporation and its subsidiaries. *Other brands and names may be claimed as the property of others.

Table of Contents

Products Affected	4
Description	4
Recovery Procedures	5
Pre-recovery Checklist	6
1.0 Multiple Hard Disk Drives Reported as 'Failed' or 'Offline'	6
2.0 Hot Spare Configured	8
3.0 All RAID Array Drives Failed	10

Products Affected

Product Code	Description	Intel RAID Software
BOXSRCU21	Intel® Server RAID Controller U2-1 (SRCU21)	v3.x and earlier
BOXSRCU31	Intel® Server RAID Controller U3-1 (SRCU31)	v4.x and earlier
BOXSRCU31L	Intel® Server RAID Controller U3-1L (SRCU31L)	v5.x and earlier

Description

This technical procedure is released as a supplement to the Intel Technical Advisory TA-445-2 (ftp://download.intel.com/support/motherboards/server/ta_445_2.pdf). **Read TA-445-2 before attempting any actions noted in this document.**

The RAID volume recovery procedures described in this document apply only to the products noted above. Select the appropriate procedure from the matrix (shown below) and read all of the procedure steps before taking any action. **If you have questions or are not sure which procedure to follow, please contact Intel Customer Support before making any changes to the server. If you do NOT select the correct procedure and/or follow the steps exactly, unrecoverable data loss could result.**

Recovery Procedures

Use the following recovery procedures in order to attempt to recover *Failed* RAID volumes that have been marked *Failed* by the RAID firmware. RAID volumes enter the *Failed* status due to the firmware detecting multiple *Failed* hard disk drives that are members of the *Failed* RAID volume. The failure detected by the firmware can be due to, but not limited to, SCSI bus error conditions, system power spikes, SCSI cabling, or SCSI termination. The RAID firmware will allow the user to recover from a very high percentage of these “false” failures if the correct procedure is followed. The specific recovery procedure used is dependent upon:

1. Configuration of the RAID subsystem just prior to the failure
2. When the failure is detected (while in the operating system or during a reboot/POST)
3. The type of failure (all hard disk drives in the RAID volume reported failed or multiple hard disk drives reported failed)

Table 1. Recovery Procedure Decision Matrix

Failure Condition	Recovery Procedure to be Used
Within the operating system, a RAID volume <i>Fails</i> due to ALL the hard disk drives in that RAID volume being reported as <i>Failed</i> .	'All RAID Array Drives Failed'
A RAID volume <i>Fails</i> due to multiple or all hard disk drives in the array being reported as <i>Failed</i> or <i>Offline</i> when there is a global hot spare configure	'Hot Spare Configured'
Within the operating system, a RAID volume <i>Fails</i> due to multiple (not all) hard disk drives in that RAID volume being reported as <i>Failed</i> or <i>Offline</i> .	'Multiple Hard Disk Drives Reported as <i>Failed</i> or <i>Offline</i> '
During a reboot/POST, a RAID volume <i>Fails</i> due to multiple (not all) hard disk drives in that RAID volume being reported as <i>Failed</i> or <i>Offline</i> .	'Multiple Hard Disk Drives Reported as <i>Failed</i> or <i>Offline</i> '

For all other Failure Conditions, or if you are not clear about your particular failure condition, please call technical support for assistance.

Pre-recovery Checklist

1. **Before proceeding with any of the following procedures you must ensure that, if a Hot Spare was configured prior to the failure, the Hot Spare WAS NOT placed into the RAID array by the RAID firmware as an attempt to Auto-rebuild the 'Failed' RAID volume. In this situation, even when the attempt to recover the 'Failed' RAID array is successful, it has been shown that the data that existed on the RAID volume prior to the failure may be lost. Follow the procedures in section 2.0, 'Hot Spare Configured', to decrease the possibility of lost of data.**
2. **While in the operating system, in the event that ALL the hard disk drives in a RAID array are reported as 'Failed' and no other RAID array is configured on the RAID controller (the OS is not on the 'Failed' RAID array), follow the procedures in section 3.0, 'All RAID Array Drives Failed', to decrease the possibility of data loss.**
3. **If a RAID volume is in the 'Failed' status due to multiple failed hard disk drives (where ALL the hard disk drives in the failed array ARE NOT failed; i.e. you have at least one hard disk drive in the 'Failed' RAID array has its status as 'OK'). During POST/Reboot the firmware will report the error condition and the RAID volume will be reported as 'Failed' on the monitor. Use the recovery procedures in section 1.0, 'Multiple Hard Disk Drives Reported as Failed or Offline'.**

1.0 Multiple Hard Disk Drives Reported as 'Failed' or 'Offline'

- Case 1: If the failure is encountered while in the operating system **DO NOT** reboot. Use the RAID Storage Console to carry out the following procedures. Replace the keystrokes with mouse clicks and ARCU with RAID Storage Console. Continue with step I.
- Case 2: If the failure is encountered during POST/Reboot, power down the server and boot to the Advanced RAID Configuration Utility (ARCU) located on the RAID distribution CD or the ARCU web package. Continue with step I.

I. Return hard disk drives that are in an error condition to a normal condition

1. From the Main menu, <↓> to Physical Disks+View/Actions link and press <Enter>

Note: upon a reboot, the RAID controller should detect any hard disk drives that were reported as *Missing*. If any drives in the *Failed* array are reported as *Missing* or do not appear in the disk list at all, they are not being seen when the RAID controller scans the SCSI bus. This is an indication that there may be something wrong with the cabling or termination on that SCSI bus. Also the hard disk drive(s) may not be properly seated into the SCSI connector. Power-off the server and re-check all cabling, SCSI termination, and hard disk drives. Reboot the server to detect the hard disk drives that were reported as *Missing*. If drives are detected and displayed continue.

DO NOT attempt further recovery in the case of actual hard disk drive failure, Call technical support.

2. <↓> to the Action column of the disk drive that is showing status of *Failed* or *Offline* and press <Enter>
3. <↓> to "Mark as Normal" (or "Mark Online") and press <Enter>

4. Repeat the previous two steps (2 and 3) for all disk drives that are members of the *Failed* RAID volume.
5. <↓> to Submit and press <Enter>
6. <↓> to Yes() and press <Enter>
7. <↓> to Submit and press <Enter>
8. Return to the Physical Disk List page and confirm that all the drives that were just marked back to Normal or Online still indicates *OK* status.

If any drives are still indicating a status of *Failed* or *Offline*, this may indicate that the hard disk drive(s) has actually failed. Repeat the steps in this procedure to attempt to bring the drive(s) back to status of *OK*. If after several attempts the drive(s) remain *Failed* or *Offline*, The drive(s) will probably have to be replaced or reformatted. Successful recovery of the RAID volume and its data may not be possible. **Stop and call Technical Support!**

9. Return to Main Menu

II. Mark the *Failed* RAID volume as *Normal*

1. From the Main menu, <↓> to RAID Volumes+View/Actions and press <Enter>
2. <↓> to the Action column of the RAID volume with status *Failed* (the volume in which the *Failed* or *Offline* hard disk drives were just returned to status of *OK*) and press <Enter>
3. <↓> to "Mark as Normal" and press <Enter>
4. Repeat for all *Failed* RAID volumes being recovered
5. <↓> to Submit and press <Enter>; you will see a Warning! Message
6. <↓> to Yes() and press <Enter>
7. <↓> to Submit and press <Enter>
8. Confirm the action
9. Return to Main Menu

The RAID volume should have status of *OK*. RAID 5 volumes will have a status of *Uninitialized* and will need to have parity re-initialized using the following steps.

III. Initialize the parity on the RAID 5 volume

1. From the Main menu, <↓> to RAID Volumes+View/Actions and press <Enter>
2. <↓> to the Action column of the RAID 5 volume with status *Uninitialized* and press <Enter>
3. <↓> to "Initialize Parity" and press <Enter>
4. Repeat for all *Failed* RAID volumes being recovered
5. <↓> to Submit and press <Enter>

The RAID 5 volume(s) will now begin to re-initialize the parity. This could take some time depending upon equipment and size of RAID volumes.

6. Press <Enter> to go to the View Volume Status page to confirm initialization. If the initialization fails this may indicate that a hard disk drive has actually failed.

DO NOT interrupt the parity initialization! DO NOT reboot the server until the initialization has completed!!

When initialization completes, the RAID 5 volume should now be accessible.

2.0 Hot Spare Configured

Note: Before proceeding, this procedure is for RAID volumes that have encountered multiple hard disk drive failures. The assumption is that at least one of two things is known: (1) the identity of the original disk drives for each RAID array or (2) the identity of all configured Hot Spare drives.

The RAID volume is in a Failed status and is not accessible to the host. This is an attempt to recover the RAID volume and its data. If currently in the operating system, DO NOT reboot. Use the RAID Storage Console to complete the following procedures.

If failure occurs during reboot, cancel the reboot and boot to the ARCU located on the RAID CD or the ARCU web package.

- I. **Determine if the hot spare hard disk drive has been added to the *Failed* RAID volume's array by the RAID firmware.**
 1. From the main menu
 2. <↓> to Physical Disks+View/Actions and press <Enter>
 3. Determine if one of the original RAID member disks has changed its Usage from *In Use* to *Available*. If it is confirmed that the hot spare drive was not placed into the RAID volume to be recovered go to section 1.0 procedures. If the hot spare's Usage is *In Use*, continue with **step II** of this section because this indicates that this disk has been replaced in the RAID array by the hot spare disk.

- II. **Ensure that the *Failed* RAID volume recovery is done with all the original hard disk drives (and only the original hard disk drives)**
 1. From the Physical Disks Lists page
 2. Locate the original Hot Spare drive that was placed into the *Failed* array (its state should be *In Use + OK*) and <↓> to the Action column of the disk drive, press <Enter>
 3. <↓> to "Mark **Offline**" and press <Enter>
 4. <↓> to Submit and press <Enter>
 5. <↓> to Yes() and press <Enter>
 6. <↓> to Submit and press <Enter>
 7. Return to the Physical Disks List page
 8. Locate the *Failed* original member drives of the RAID volume being recovered and <↓> to the Action column of the first hard disk drive press <Enter>
 9. <↓> to "Mark as Normal" or "Mark Online" and press <Enter>
 10. For all remaining *Failed* (or *Offline*) hard disk drives that are members of the RAID volume being recovered, perform the "Mark as Normal" or "Mark Online" action.
 11. <↓> to Submit and press <Enter>
 12. <↓> to Yes() and press <Enter>
 13. <↓> to Submit and press <Enter>
 14. Return to the Physical Disks List page
 15. **Ensure that:**
 - a. **The original Hot Spare is of status Offline**
 - b. **The original member hard disk drives of the RAID volume being recovered are of status OK and usage In Use (with exception of the one that was replaced in the array by the original hot spare, its usage will be Available).**
 16. When the configuration in step 15 is verified, Return to the Main Menu

III. Place the RAID volume to status of *Degraded* to Auto-rebuild it

1. From the main menu
2. <↓> to RAID Volumes+View/Actions and press <Enter>
3. <↓> to the Action column of the RAID volume with status of Failed and press <Enter>
4. <↓> to “Mark as Normal” and press <Enter>
5. <↓> to Submit and press <Enter>, You will see a Warning message.
6. <↓> to Yes() and press <Enter>
7. <↓> to Submit and press <Enter>
8. Confirm Yes and Submit
9. Return to the RAID Volume List page and confirm that all the RAID volumes that you are attempting recovery are in the *Degraded* status
10. Return to Main Menu (The original array member disk must now be marked to *Spare* in order to properly rebuild the *Degraded* RAID volumes to the *OK* status with data intact)
11. <↓> to Physical Disks+View/Actions and press <Enter>
12. Locate the original member drive that was replaced by the hot spare (it should have Usage as *Available* and Status *OK*), <↓> to the Action column of the disk drive press <Enter>
13. Perform the “Mark As Spare” action and submit to complete the action

At this point the *Degraded* RAID volume(s) being recovered should auto-rebuild to the original drives in the RAID volume. **DO NOT interrupt the rebuild! DO NOT reboot the server until the rebuild has completed!**

3.0 All RAID Array Drives Failed

This procedure is required in the event that all the hard disk drives that are members of RAID volumes go to status of *Fail* while the operating system is running and the operating system is not on located on the RAID volume that has *Failed*.

I. Determine failed or error condition and correct hard disk drive error status.

1. Launch the RAID Storage Console main menu
Note: If a hot spare was configured prior to the hard disks being reported as failed, determine if the hot spare replaced one of the original array member hard disks and was inserted into the Failed RAID array by the RAID firmware (the exact configuration of all affected RAID volumes, RAID arrays, and hard disk drives must be known).
 - **In the Physical Disk List, locate the hot spare drive and determine its Usage. If it has changed from 'Spare' to 'In Use', then it has been placed into the array in an attempt to auto-rebuild the failed array. In this case go [section 2.0, 'Hot Spare Configured'](#). **DO NOT Reboot the computer! Stay in the OS and run the procedure using the RAID Storage Console.****
2. Go to the Physical Disks List page (in left menu, click the Physical Disks+View/Actions link) and locate the *Failed* hard disk drives that are the member disks of the RAID volume that is to be recovered (**the recovery must be done using all the hard disks that were members of the RAID volume at the time of the failure, otherwise, data recovery will fail even though the RAID volume can be repaired**)
3. Go to the Disk Actions column and click on the box to access the drop down list.
4. From the list select "Mark as Normal", repeat for all *Failed* hard disks.
5. Click the [Submit] button; a Warning Message will appear
6. Click the [Yes] button to complete the action. When the physical disk list reappears, all of the *Failed* disk should now be indicating a status of *OK*. If some or all disks are not of status *OK*, repeat steps 3 through 6. After repeating, if all disks are still not of status of *OK*, this may be an indication of an actual hard disk failure (**it may not be possible to recover the data on this RAID volume, Do Not continue! Call Technical Support!**)
7. Return to the main menu

II. Mark the RAID volume "Normal"

1. Go to the RAID Volume List page (in left menu, click the RAID Volumes+View/Action link)
2. Locate the RAID volume to be recovered (status of *Failed*) and go to the Action column
3. Click on the list box to access the drop down list
4. Select "Mark as Normal" then click on [Submit]
5. At the Warning Message, click [Yes] and confirm by clicking [Yes] again
6. When don you will be the text: "Requested action(s) completed"

The RAID volume should now have status of *OK*. RAID 5 volumes will have a status of *Uninitialized* and will need to have parity re-initialized using the following steps.

III. Re-initialize parity for RAID 5 volumes

1. Go to the RAID Volume List page (in left menu, click the RAID Volumes+View/Action link)
2. Go to Actions column of the RAID 5 volume being recovered and click on the list box to access the drop down list
3. Select "Initialize Parity" and click the [Submit] button to begin the parity initialization

DO NOT interrupt the parity initialization or reboot the server! Let the initialization complete before attempting any other actions.

Please contact Intel Customer Support if you require more specific information about this issue.