Data Integrity Validation and Splitting Mirrors: Best Practice

(Version 3.0, Revised 2005-08-18)

Overview

Splitting mirrors is the practice of bringing down one-half of the mirrored pair of disk drives. This operation may be performed for a variety of reasons, including disk upgrade, hardware repair, and data protection. The volume is protected from invalid data due to checksum and media failures only when both drives in the mirrored pair are available. One or more distinct blocks of data may be invalid on each of the drives, but when both drives are fully available, the I/O subsystem can automatically re-route read requests to the other drive of the mirror in order to overcome such failures.

Good practice should require that all files and tables be validated prior to splitting mirrors, to detect and correct latent failures.

This document describes the best practice for data validation prior to splitting mirrors. It covers both G-series starting with G06.24 and all H-series. While the validation process is meant to reduce the risk of splitting mirrored drives, there is no method available to eliminate this risk**. Splitting mirrored drives creates a single point of failure. Although disk hardware failures are relatively rare, they do occur, and mirrored drives should be split only when absolutely necessary.** Routine data backup procedures should be followed to avoid single point of failure data loss.

Prior to bringing down a drive in a mirrored pair, the NSK FCHECK utility should be used to validate that the data is the same on both the primary and mirrored drives. The FCHECK –SCAN option will detect checksum errors and correct them when possible, and validate that the data is consistent across the mirrored pair.

When any read detects a checksum error on either drive, the DP2 driver will automatically attempt to read the data from the other drive of the mirrored pair. If the data is valid, the driver will automatically rewrite the data to the drive with the invalid checksum.

When the FCHECK –SCAN utility is executed, all target data is read on both drives and compared. If the data on both drives incurs an invalid checksum, the data cannot be repaired. EMS events will be generated for this case reporting the disk sector and FCHECK will report the file relative location of the checksum error.

If the volume is mirrored, and both drives of the mirrored pair incur an irrecoverable checksum error at the same location, this may be a symptom of various failures:

1. Path problems – in this case, it is typical for there to be multiple reports of unrecoverable checksum errors in a variety of locations across a variety of files and tables.
2. Unwritten data – in this case, there may have been some interrupted data load or restore operation that left the data set incomplete. It is typical for there to be multiple reports of unrecoverable checksum errors for contiguous sets of data within the same file.
3. Software problems – in rare cases, software issues may result in a variety of failure modes that can result in unrecoverable checksum errors.

Data Validation Timing

The data validation can require anywhere from minutes to hours, depending upon the system load, FCHECK process priority, the RATE used to perform the validation, and the size of the data set that must be validated. The validation can be performed in parallel across multiple volumes for each volume subject to a mirror split. The rate at which data is processed for each volume can range up to 50mbytes/second or more, depending upon the file set. The best way to determine the data validation timing is with actual experience. The risk of undetected latent failures is reduced when data is validated prior to splitting the mirrors. The validation should be performed as close as possible to the mirror split to ensure the lowest possible risk.

Recovery from Data Integrity Loss

Any inconsistencies reported by FCHECK -SCAN require special handling to determine the correct course of action to properly recover the data set. As a general rule, an FCHECK without the –SCAN option should be performed if the reported file or table is structured.

The first step is to identify if one of the disks in the mirrored pair is defective or contains all the inconsistencies. If this can be determined, the defective or inconsistent drive should be the target of a revive operation to correct the inconsistencies.

WARNING: If there is any doubt as to the location of the inconsistent or mismatched data, contact your HP support representative for assistance. Removing the wrong drive from service, or removing either drive from service when both contain inconsistencies, may result in an inability to recover the suspect data.

If the file or table cannot be otherwise recovered via standard data recovery mechanisms such as RESTORE or TMF RECOVER FILES, contact your HP support representative for assistance.

FCHECK -SCAN Instructions

The FCHECK -SCAN operation for data validation associated with splitting mirrors is performed by running the FCHECK process:

```
FCHECK [/ PRI nnn /] –SCAN [-RATE <nnn>] –VOL <$volume>
                                         -SUBVOL <subvol>
                                         -FILE <filename>
```

Every file and table on the volume will be scanned when the –VOL option is specified. The user can specify –FILE for a specific file or –SUBVOL for the set of files contained within a sub-volume.

A priority should be specified to avoid impact to production activity. The specified priority should be less than the priority of the production workload, and should also be less than 151.

The –RATE option should also be considered for environments when production response time increases cannot be tolerated. The default rate is 100. It can be reduced to a value between 1 and 99. The lower the rate, the smaller the impact will be to production activity.

A reduced RATE and a reduced PRIORITY will not typically slow the pace of the scrub when there is no concurrent activity.