DP2 FCHECK User's Guide

1. Introduction

FCHECK is a Guardian program that checks the internal consistency of Disk Process 2 (DP2) structured files. It *does not* correct any of the structural problems it finds.

You can use FCHECK to verify the consistency of standard (Format 1) and large (Format 2) files.

Note: The older DP2 tool, FILCHECK, does not work on Format 2 files.

After you install G06.24, you can also use FCHECK to:

- Validate mirrored drive consistency. You can use the –SCAN option to verify the checksum of the source and target drives and to ensure that the data on both drives of the mirrored pair is the same. If one of the drives has a valid checksum and the other drive has an invalid checksum, the –SCAN option automatically copies the data with the valid checksum to the drive with the invalid checksum.

- Determine the space requirement for migrating data from internal, 514-sector volumes to 512-sector volumes. Before initiating a migratory revive, use the –SIZE option to report the current space used on the volume and the amount of space required to migrate the data to ensure that a disk drive in an Enterprise Storage System (ESS) or Fibre Channel disk drive enclosure is properly sized. If space is not sufficient, the migratory revive will abort.

FCHECK validation is actually performed in DP2. FCHECK drives the validation and reports the errors.

You can run FCHECK concurrently with any application. The DP2 volume process performs data validation. As a result, application activity does not interfere with consistency checks.

This user guide is organized as follows:

Section	Title	Describes...
1	Introduction	When to use FCHECK
2	Data Scrubbing	FCHECK data validation
3	Disk Scrubbing	FCHECK disk validation
4	FCHECK Size Option	FCHECK capacity checks for migration of data from 514-sector disks to a disk drive in an ESS or Fibre Channel disk drive enclosure.
5	FCHECK Syntax	How to run FCHECK
6	Disk Scan Checks	Disk scrubbing validation checks
7	Structured Data Validation	Data scrubbing validation checks
8	FCHECK Output	How to use the information returned by FCHECK
9	Detailed Information for Broken Files	How FCHECK reports problems with a file
10	Physical Disk Integrity Validation	FCHECK –SCAN and Automatic Sector Reallocation
	Appendix	Outstanding issues

You can acquire and use the FFIX utility to rebuild the B-tree structure in place for a broken file or table partition. You must alter the partition to a nonaudited state before you run the FFIX utility, and exclusive access is required. For assistance with using the FFIX utility, contact your service provider.

2. Data Scrubbing

As a general guideline, run FCHECK periodically to ensure that data is valid and consistent. You can initiate FCHECK routinely during nonpeak periods to cover entire volumes or limit it to critical data sets.

When you run the FCHECK utility without a --SCAN option, and without the --SIZE option, a full validation of the structures is performed for any target structured files. The validation is typically performed in response to certain application errors, especially error 59 and includes key-sequenced B-tree consistency and block header and trailer validation. Data scrub can be an intense operation, requiring significant CPU and I/O cycles to validate the data. Consider disk scrub for routine scrubbing requirements, with ad-hoc or less routine data scrub to validate key data sets.

FCHECK also allows you to validate unstructured files with data scrub to ensure that such files are readable. Use disk scrub instead if the only files in the set are unstructured files.

3. Disk Scrubbing

For mirrored volumes, HP recommends using the –SCAN option for periodic disk scrubbing. You can use NETBATCH or a similar product that automatically initiates one or more FCHECK validation scan operations during nonpeak periods. For example, you might execute this command for key disk volumes once each night or weekend:

```
FCHECK –SCAN –FINDBAD –VOL $DATA
```

HP strongly recommends that you perform the scan before you remove (SCF STOP) one-half of a mirrored pair. In addition to the –SCAN option, FCHECK supports the –SCANPRIMARY and –SCANMIRROR options:

```
FCHECK –SCANMIRROR –FINDBAD –VOL $DATA
```

If a checksum error is detected on the specified drive, the driver attempts to read the data on the nonscanned drive. If that data has a valid checksum, the driver automatically rewrites the data to the drive with the defective data. Otherwise, the scan does not read the data on the drive that was not specified.

This scan is particularly useful when a drive is suspect or the subject of a smart trip warning. By scanning only the good drive, the scan is less likely to result in a complete failure of the suspect or smart-tripped drive.

Always consider a disk scrub of the mirrored pair before you stop any drive associated with the mirrored pair. These conditions include:

- Drive replacement
- Stopping one drive of each mirrored pair as a fallback position
- Drive upgrade
- Migratory revive to a disk drive in an ESS or Fibre Channel disk drive enclosure.

Performing a disk scrub will reduce the possibility of data loss associated with the loss of the mirrored protection.

4. FCHECK –SIZE Option

Before you perform a migratory revive from a 514-byte sector disk to a disk drive in an ESS or Fibre Channel disk drive enclosure, run FCHECK –SIZE to ensure that the drive drive has sufficient capacity to handle any possible increased space requirement.

The syntax for this operation is:

```
FCHECK –SIZE –VOL $DATA
```

FCHECK examines the volume directory and reports the current usage for volume metadata, structured files, and unstructured files. In addition, FCHECK provides the amount of space required to migrate the entire volume to a disk drive in an ESS or Fibre Channel disk drive enclosure. The first section of this report shows current allocated space, and the second section shows the required allocated space for the target disk.

```
Disk space usages on volume (in MB):
  Maximum space:           4963030
    Meta data:             376
    Structured data:       1897
    Unstructured data:     4960757
  Allocated space:         11043
    Meta data:             2
    Structured data:       23
    Unstructured data:     11018

To revive data to an ESS/JBOD disk,
the target disk needs disk space for:
  Maximum space:           5693004
    Meta data:             378
    Structured data:       1897
    Unstructured data:     5690729
  Allocated space:         12671
    Meta data:             4
    Structured data:       23
    Unstructured data:     12644
```

The maximum amount of space reported is based on full allocation of all files on the disk volume. There is typically no reason to size based on this value, but the information is provided to ensure that the disk has sufficient capacity for any expected increase in space requirements.

5. FCHECK Syntax

```
FCHECK  [ -RATE nnn ] [options] { [-FILE file]
                                  [-VOL volume]
                                  [-SUBVOL subvolume]
                                  [ * ] }
            [ -HELP ]
```

–RATE controls the rate at which the file checking is performed. The minimum value is 1 (slowest) and the maximum value is 100 (fastest). The larger the specified value, the greater the impact on system performance.

options = any combination of the following options, separated by spaces:

–Q	Quiet mode
–S	Summary mode
–SIZE	Provides information on space usage
–FINDBAD	Combines –Q and –S
–FIXEOF	Reduces EOF to exclude blocks with invalid checksums
–NOBTREE	Skips B-tree checks for key-sequenced files or tables
–SCAN	Performs scan check
–FASTSCAN	Performs scan check without checksum validation
–FULLSCAN	Performs scan check followed by data validation
–SCANPRIMARY	Performs scan check of only the primary drive
–SCANMIRROR	Performs scan check of only the mirror drive

Use one of the following options to specify the file set(s) to be checked by FCHECK:

–FILE	[*system*.][$*volume*.][*subvolume*.]{*filename*}
–VOL	[*system*.]{$*volume*}
–SUBVOL	[*system*.]{$*volume.subvolume*}
*	Check all files in the current $*volume.subvolume*

–HELP = print this information

You can specify **–FILE $volume.SYS00.DIRECTRY** to check the directory.

To limit impact to higher priority workloads, initiate FCHECK at a priority less than the higher priority workloads.

6. Disk Scan Checks

When you use the –SCAN option, the validation examines the data on both drives of the mirrored pair and ensures that the blocks have valid checksums. If both the primary and mirror drives contain invalid checksums, an error is reported for the file in question. If one drive has a valid checksum and the other drive has an invalid checksum, the good data is rewritten to the drive with the invalid data and the checking continues. If neither drive contains invalid checksums, the data is compared and any mismatch results in an error reported to the user, which identifies the location of the mismatch.

To limit the scan to the primary or the mirror drive, use the –PRIMARYSCAN or –MIRRORSCAN option in the place of the –SCAN option. When you specify either of these options, the scan validates checksums on the specified drive. If the other drive is available, it is used for checksum error recovery if the checksum is valid on that drive.

When you use the –FASTSCAN option, the validation only compares the data on the two drives of the mirrored pair. Any mismatch is reported to the user.

You can specify –FULLSCAN to perform both scan validation and data validation. When you use –FULLSCAN, the scan validation is performed first, followed by the structure validation.

Note: Neither the –SCAN nor the –FASTSCAN option performs the structured data validation.

Scan validation is performed for all file types, including unstructured files.

FCHECK errors displayed during a scan validation include:

```
Checksum   300
Medium     311
Compare    313
```

FCHECK might also report other errors, but they are typically related to path loss or a down drive. Check Event Management Service (EMS) logs for disk subsystem events related to such errors.

7. Structured Data Validation

When the –SCAN or –FASTSCAN options are not specified, FCHECK performs structured data validation.

FCHECK will prompt DP2 to examine each block to ensure a proper structure, which must be consistent for the file or table type. If the –NOBTREE option is not specified, DP2 examines key-sequenced B-tree structures to ensure that the B-tree is consistent.

DP2 examines the file or table sequentially, and preempts the data validation if higher priority work is present. FCHECK uses the –RATE value to pause between preemptions to limit the impact on active production.

When DP2 encounters errors, EMS events are generated, and FCHECK reports the details of the errors.

8. FCHECK Output

FCHECK provides a visual report for broken key-sequenced B-tree structures, based on the errors that DP2 detects. This data shows the inconsistence in the b-tree.

When no errors are encountered, FCHECK reports the following information:

```
 FileCheck : $SYSTEM.SYS20.TEMPLATE                    2003/03/31
07:19
   Version :  File format version Standard
   Blocklen:       4096
     Reclen:        510
     Keylen:         30

    Disk I/Os    :          87
    Data blocks  :        1293
    Index blocks :          22
    Bitmap blocks:           1
    Elapsed time :        0.66  seconds

=========================================================================
* * FileCheck did not find errors * * *
```

Note: The last line is displayed after all files are checked.

If you specify the –Q (quiet) option, only the last line of the report appears:

```
* * FileCheck did not find errors * * *
```

If you specify the –S (summary) option, each checked file is reported as either correct (Ok) or broken:

```
File Ok : $SYSTEM.SYS20.TEMPLATE
 * * * FileCheck did not find errors * * *
```

If any errors are found, the last line of the report is:

```
* * FileCheck found broken files  * * *
```

If you specify the –FINDBAD option, any bad files are reported as:

```
File Bad: $DATA22.GDUMPS.TEMPLATE
* * FileCheck found broken files  * * *
```

An FCHECK summary display includes the following information:

```
FCHECK -findbad -vol $data01
DP2 File Check Server -- FCHECK -- T9208G11 -- (15MAR2004)


||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
||

   File check summary:
    Files checked:            2876
    Files skipped:            0
    Open error 12:      8
    Enscribe files:           2882
    SQL tables:         2
    Broken files:       0
    Structured xsums :  0

 * * * Files not checked, open error * * *
```

FCHECK might encounter various failures when attempting to open a file for validation. While the summary shows the number of instances of open error 12, all open errors are included in this count, including error 12. In some cases, open errors might be transient. In other cases, the open error might be persistent. For example, if an audited partition requires HP NonStop Transaction Management Facility (TMF) volume recovery, the error persists until the recovery is performed. If the error occurs because a redo or an undo operation is needed, you can use the SNOOP utility to reset the recovery flags, which will allow FCHECK to open the partition. Use this procedure only if volume or file recovery has failed to recover the partition.

Other possible causes for open errors include:

- Reserved file name
- Conflicting open
- Active NonStop SQL/MP or NonStop SQL/MP Structured Query Language (SQL) Data Definition Language (DDL) operation

9. Detailed Information for Broken Files

The FCHECK utility displays information about three major file-structure problems:

- Broken block
- Block checksum error
- Broken key-sequenced B-tree structure

For a broken block header or block checksum error, FCHECK displays:

```
>>>    Bad Block code:     BLK^19
   Block RSN:           0          Expected RSN:           0
       Flags:     %100000        Expected Flags:     %006400
     Records:       31309          Error Recnum: ----------
   Blk Level:           0           File Levels:           0
    Prev RSN:          -1              Next RSN:          -1
 >>>>>>>    KS: wrong block type
```

The block code equates to an internal symptom string generated by the disk process and is translated in the last line of the preceding display. In this case, the block flags do not match what is expected.

For a broken B-tree structure, FCHECK displays:

```
>>>   KS B-tree error
   Block RSN:          24           Expected RSN:          24
       Flags:    %006000         Expected Flags:    %006000
     Records:          33           Error Recnum:          32
   Blk Level:           0            File Levels:           3
    Prev RSN:          16               Next RSN:          64
-------------------------------------------------------------------
Level    Prev RSN           RSN      Recnum      Next RSN
-------------------------------------------------------------------
 02                           8           0
 01                        1040           1
       ---------------------------------------------------------
      Index rec -1            |             Index rec +1
                              |              0<-- <     32>
                              |
                              |
                             RSN
                              |
               16  <---->    24  *---->          64
                                                  |
             8832  <----------------------------------
-------------------------------------------------------------------
```

This display shows the index paths that lead to the data block (block levels), and the data block chains (previous and next). This information can aid manual recovery of the broken B-tree and provide some diagnostic information to help manage the problem.

If you cannot recover the broken data with some form of file-recovery method (such as restore or TMF recovery), you might be able to correct the broken structure manually or use the FFIX utility to rebuild the structure in place. If the structure needs to be corrected manually, contact your service provider for assistance.

10. Physical Disk Integrity Validation

You can use FCHECK –SCAN as the key part of a regularly scheduled disk data-integrity validation. As described earlier, FCHECK –SCAN finds all sector checksum problems and automatically these problems when one of the drives of the mirrored pair contains valid and checksum error-free data. In addition, if T8109 Automatic Sector Reallocation is enabled, running FCHECK –SCAN causes every in-use sector on both halves of a logical disk to be read. The DP2 driver automatically reports any media errors encountered during this scan, and the T8109 reallocation recovers the data from the good copy. As a result, media errors are reported and automatically corrected, if possible, whenever you run FCHECK –SCAN.

FCHECK –SCAN results are also reported to and cause an alarm on the HP NonStop Open System Mangement (OSM) interface.

You can use FCHECK –SCAN inside an HP Tandem Advanced Command Language (TACL) macro or other automatic mechanism. Because OSM alarms show any problems that could not be fixed, you can then use these alarms to help you correct the problems. When the problem is corrected, you can delete the alarm or rerun FCHECK –SCAN to confirm the repair and update the OSM alarm status.