

**NAME**

xfs\_db – debug an XFS filesystem

**SYNOPSIS**

**xfs\_db** [ **-c** *cmd* ] ... [ **-i***r|x|F* ] [ **-f** ] [ **-l** *logdev* ] [ **-p** *progrname* ] *device*  
**xfs\_db** -V

**DESCRIPTION**

**xfs\_db** is used to examine an XFS filesystem. Under rare circumstances it can also be used to modify an XFS filesystem, but that task is normally left to **xfs\_repair**(8) or to scripts such as **xfs\_admin**(8) that run **xfs\_db**.

**OPTIONS**

- c** *cmd* **xfs\_db** commands may be run interactively (the default) or as arguments on the command line. Multiple **-c** arguments may be given. The commands are run in the sequence given, then the program exits.
- f** Specifies that the filesystem image to be processed is stored in a regular file at *device* (see the **mkfs.xfs**(8) **-d** *file* option). This might happen if an image copy of a filesystem has been made into an ordinary file with **xfs\_copy**(8).
- F** Specifies that we want to continue even if the superblock magic is not correct. For use in **xfs\_metadump**.
- i** Allows execution on a mounted filesystem, provided it is mounted read-only. Useful for shell scripts which must only operate on filesystems in a guaranteed consistent state (either unmounted or mounted read-only). These semantics are slightly different to that of the **-r** option.
- l** *logdev* Specifies the device where the filesystems external log resides. Only for those filesystems which use an external log. See the **mkfs.xfs**(8) **-l** option, and refer to **xfs**(5) for a detailed description of the XFS log.
- p** *progrname* Set the program name to *progrname* for prompts and some error messages, the default value is **xfs\_db**.
- r** Open *device* or *filename* read-only. This option is required if the filesystem is mounted. It is only necessary to omit this flag if a command that changes data (**write**, **blocktrash**, **crc**) is to be used.
- x** Specifies expert mode. This enables the (**write**, **blocktrash**, **crc** invalidate/revalidate) commands.
- V** Prints the version number and exits.

**CONCEPTS**

**xfs\_db** commands can be broken up into two classes. Most commands are for the navigation and display of data structures in the filesystem. Other commands are for scanning the filesystem in some way.

Commands which are used to navigate the filesystem structure take arguments which reflect the names of filesystem structure fields. There can be multiple field names separated by dots when the underlying structures are nested, as in C. The field names can be indexed (as an array index) if the underlying field is an array. The array indices can be specified as a range, two numbers separated by a dash.

**xfs\_db** maintains a current address in the filesystem. The granularity of the address is a filesystem structure. This can be a filesystem block, an inode or quota (smaller than a filesystem block), or a directory block (could be larger than a filesystem block). There are a variety of commands to set the current address. Associated with the current address is the current data type, which is the structural type of this data. Commands which follow the structure of the filesystem always set the type as well as the address. Commands which examine pieces of an individual file (inode) need the current inode to be set, this is done with the **inode** command.

The current address/type information is actually maintained in a stack that can be explicitly manipulated with the **push**, **pop**, and **stack** commands. This allows for easy examination of a nested filesystem

structure. Also, the last several locations visited are stored in a ring buffer which can be manipulated with the **forward**, **back**, and **ring** commands.

XFS filesystems are divided into a small number of allocation groups. **xfs\_db** maintains a notion of the current allocation group which is manipulated by some commands. The initial allocation group is 0.

## COMMANDS

Many commands have extensive online help. Use the **help** command for more details on any command.

**a** See the **addr** command.

**ablock** *filoff*

Set current address to the offset *filoff* (a filesystem block number) in the attribute area of the current inode.

**addr** [*field-expression*]

Set current address to the value of the *field-expression*. This is used to "follow" a reference in one structure to the object being referred to. If no argument is given, the current address is printed.

**agf** [*agno*]

Set current address to the AGF block for allocation group *agno*. If no argument is given, use the current allocation group.

**agfl** [*agno*]

Set current address to the AGFL block for allocation group *agno*. If no argument is given, use the current allocation group.

**agi** [*agno*]

Set current address to the AGI block for allocation group *agno*. If no argument is given, use the current allocation group.

**b** See the **back** command.

**back** Move to the previous location in the position ring.

**blockfree**

Free block usage information collected by the last execution of the **blockget** command. This must be done before another **blockget** command can be given, presumably with different arguments than the previous one.

**blockget** [**-npvs**] [**-b** *bno*] ... [**-i** *ino*] ...

Get block usage and check filesystem consistency. The information is saved for use by a subsequent **blockuse**, **ncheck**, or **blocktrash** command.

**-b** is used to specify filesystem block numbers about which verbose information should be printed.

**-i** is used to specify inode numbers about which verbose information should be printed.

**-n** is used to save pathnames for inodes visited, this is used to support the **xfs\_ncheck**(8) command. It also means that pathnames will be printed for inodes that have problems. This option uses a lot of memory so is not enabled by default.

**-p** causes error messages to be prefixed with the filesystem name being processed. This is useful if several copies of **xfs\_db** are run in parallel.

**-s** restricts output to severe errors only. This is useful if the output is too long otherwise.

**-v** enables verbose output. Messages will be printed for every block and inode processed.

**blocktrash** [**-z**] [**-o** *offset*] [**-n** *count*] [**-x** *min*] [**-y** *max*] [**-s** *seed*] [**-0**|**1**|**2**|**3**] [**-t** *type*] ...

Trash randomly selected filesystem metadata blocks. Trashing occurs to randomly selected bits in the chosen blocks. This command is available only in debugging versions of **xfs\_db**. It is useful

for testing **xfs\_repair**(8).

**-0 | -1 | -2 | -3**

These are used to set the operating mode for **blocktrash**. Only one can be used: **-0** changed bits are cleared; **-1** changed bits are set; **-2** changed bits are inverted; **-3** changed bits are randomized.

**-n** supplies the *count* of block-trashings to perform (default 1).

**-o** supplies the bit *offset* at which to start trashing the block. If the value is preceded by a '+', the trashing will start at a randomly chosen offset that is larger than the value supplied. The default is to randomly choose an offset anywhere in the block.

**-s** supplies a *seed* to the random processing.

**-t** gives a *type* of blocks to be selected for trashing. Multiple **-t** options may be given. If no **-t** options are given then all metadata types can be trashed.

**-x** sets the *minimum* size of bit range to be trashed. The default value is 1.

**-y** sets the *maximum* size of bit range to be trashed. The default value is 1024.

**-z** trashes the block at the top of the stack. It is not necessary to run **blockget** if this option is supplied.

**blockuse [-n] [-c count]**

Print usage for current filesystem block(s). For each block, the type and (if any) inode are printed.

**-c** specifies a *count* of blocks to process. The default value is 1 (the current block only).

**-n** specifies that file names should be printed. The prior **blockget** command must have also specified the **-n** option.

**bmap [-a] [-d] [block [len]]**

Show the block map for the current inode. The map display can be restricted to an area of the file with the *block* and *len* arguments. If *block* is given and *len* is omitted then 1 is assumed for len.

The **-a** and **-d** options are used to select the attribute or data area of the inode, if neither option is given then both areas are shown.

**btdump [-a] [-i]**

If the cursor points to a btree node, dump the btree from that block downward. If instead the cursor points to an inode, dump the data fork block mapping btree if there is one. If the cursor points to a directory or extended attribute btree node, dump that. By default, only records stored in the btree are dumped.

**-a** If the cursor points at an inode, dump the extended attribute block mapping btree, if present.

**-i** Dump all keys and pointers in intermediate btree nodes, and all records in leaf btree nodes.

**check** See the **blockget** command.

**convert type number [type number] ... type**

Convert from one address form to another. The known *types*, with alternate names, are:

**agblock** or **agbno** (filesystem block within an allocation group)

**agino** or **aginode** (inode number within an allocation group)

**agnumber** or **agno** (allocation group number)

**bboff** or **daddroff** (byte offset in a **daddr**)

**blkoff** or **fsboff** or **agboff** (byte offset in a **agblock** or **fsblock**)

**byte** or **fsbyte** (byte address in filesystem)

**daddr** or **bb** (disk address, 512-byte blocks)  
**fsblock** or **fsb** or **fsbno** (filesystem block, see the **fsblock** command)  
**ino** or **inode** (inode number)  
**inoidx** or **offset** (index of inode in filesystem block)  
**inooff** or **inodeoff** (byte offset in inode)

Only conversions that "make sense" are allowed. The compound form (with more than three arguments) is useful for conversions such as **convert agno ag agbno agb fsblock**.

#### **cre [-i|-r|-v]**

Invalidates, revalidates, or validates the CRC (checksum) field of the current structure, if it has one. This command is available only on CRC-enabled filesystems. With no argument, validation is performed. Each command will display the resulting CRC value and state.

- i** Invalidate the structure's CRC value (incrementing it by one), and write it to disk.
- r** Recalculate the current structure's correct CRC value, and write it to disk.
- v** Validate and display the current value and state of the structure's CRC.

#### **daddr [d]**

Set current address to the daddr (512 byte block) given by *d*. If no value for *d* is given, the current address is printed, expressed as a daddr. The type is set to **data** (uninterpreted).

#### **dblock filoff**

Set current address to the offset *filoff* (a filesystem block number) in the data area of the current inode.

#### **debug [flagbits]**

Set debug option bits. These are used for debugging **xfs\_db**. If no value is given for *flagbits*, print the current debug option bits. These are for the use of the implementor.

#### **dquot [-g|-p|-u] id**

Set current address to a group, project or user quota block for the given ID. Defaults to user quota.

#### **echo [arg] ...**

Echo the arguments to the output.

**f** See the **forward** command.

#### **forward**

Move forward to the next entry in the position ring.

#### **frag [-adflqRrv]**

Get file fragmentation data. This prints information about fragmentation of file data in the filesystem (as opposed to fragmentation of freespace, for which see the **freesp** command). Every file in the filesystem is examined to see how far from ideal its extent mappings are. A summary is printed giving the totals.

- v** sets verbosity, every inode has information printed for it. The remaining options select which inodes and extents are examined. If no options are given then all are assumed set, otherwise just those given are enabled.
- a** enables processing of attribute data.
- d** enables processing of directory data.
- f** enables processing of regular file data.
- l** enables processing of symbolic link data.
- q** enables processing of quota file data.
- R** enables processing of realtime control file data.
- r** enables processing of realtime file data.

**freesp** [**-bcds**] [**-A** *alignment*] [**-a** *ag*] ... [**-e** *i*] [**-h** *hl*] ... [**-m** *m*]

Summarize free space for the filesystem. The free blocks are examined and totalled, and displayed in the form of a histogram, with a count of extents in each range of free extent sizes.

- A** reports only free extents with starting blocks aligned to *alignment* blocks.
- a** adds *ag* to the list of allocation groups to be processed. If no **-a** options are given then all allocation groups are processed.
- b** specifies that the histogram buckets are binary-sized, with the starting sizes being the powers of 2.
- c** specifies that **freesp** will search the by-size (cnt) space Btree instead of the default by-block (bno) space Btree.
- d** specifies that every free extent will be displayed.
- e** specifies that the histogram buckets are equal-sized, with the size specified as *i*.
- h** specifies a starting block number for a histogram bucket as *hl*. Multiple **-h**'s are given to specify the complete set of buckets.
- m** specifies that the histogram starting block numbers are powers of *m*. This is the general case of **-b**.
- s** specifies that a final summary of total free extents, free blocks, and the average free extent size is printed.

**fsb** See the **fsblock** command.

**fsblock** [*fsb*]

Set current address to the fsblock value given by *fsb*. If no value for *fsb* is given the current address is printed, expressed as an fsb. The type is set to **data** (uninterpreted). XFS filesystem block numbers are computed  $((agno \ll agshift) | agblock)$  where *agshift* depends on the size of an allocation group. Use the **convert** command to convert to and from this form. Block numbers given for file blocks (for instance from the **bmap** command) are in this form.

**fsmap** [*start*] [*end*]

Prints the mapping of disk blocks used by an XFS filesystem. The map lists each extent used by files, allocation group metadata, journalling logs, and static filesystem metadata, as well as any regions that are unused. All blocks, offsets, and lengths are specified in units of 512-byte blocks, no matter what the filesystem's block size is. **The optional *start* and *end* arguments can be used to constrain the output to a particular range of disk blocks.**

**fuzz** [**-c**] [**-d**] *field action*

Write garbage into a specific structure field on disk. Expert mode must be enabled to use this command. The operation happens immediately; there is no buffering.

The fuzz command can take the following *actions* against a field:

**zeroes**

Clears all bits in the field.

**ones** Sets all bits in the field.

**firstbit**

Flips the first bit in the field. For a scalar value, this is the highest bit.

**middlebit**

Flips the middle bit in the field.

**lastbit**

Flips the last bit in the field. For a scalar value, this is the lowest bit.

**add** Adds a small value to a scalar field.

**sub** Subtracts a small value from a scalar field.

**random**

Randomizes the contents of the field.

The following switches affect the write behavior:

- c** Skip write verifiers and CRC recalculation; allows invalid data to be written to disk.
- d** Skip write verifiers but perform CRC recalculation; allows invalid data to be written to disk to test detection of invalid data.

**hash** *string*

Prints the hash value of *string* using the hash function of the XFS directory and attribute implementation.

**help** [*command*]

Print help for one or all commands.

**info**

Displays selected geometry information about the filesystem. The output will have the same format that **mkfs.xfs**(8) prints when creating a filesystem or **xfs\_info**(8) prints when querying a filesystem.

**inode** [*inode#*]

Set the current inode number. If no *inode#* is given, print the current inode number.

**label** [*label*]

Set the filesystem label. The filesystem label can be used by **mount**(8) instead of using a device special file. The maximum length of an XFS label is 12 characters – use of a longer *label* will result in truncation and a warning will be issued. If no *label* is given, the current filesystem label is printed.

**log** [*stop* | *start filename*]

Start logging output to *filename*, stop logging, or print the current logging status.

**logres**

Print transaction reservation size information for each transaction type. This makes it easier to find discrepancies in the reservation calculations between xfsprogs and the kernel, which will help when diagnosing minimum log size calculation errors.

**metadump** [**-egow**] *filename*

Dumps metadata to a file. See **xfs\_metadump**(8) for more information.

**ncheck** [**-s**] [**-i ino**] ...

Print name-inode pairs. A **blockget -n** command must be run first to gather the information.

- i** specifies an inode number to be printed. If no **-i** options are given then all inodes are printed.
- s** specifies that only setuid and setgid files are printed.

**p**

See the **print** command.

**pop**

Pop location from the stack.

**print** [*field-expression*] ...

Print field values. If no argument is given, print all fields in the current structure.

**push** [*command*]

Push location to the stack. If *command* is supplied, set the current location to the results of *command* after pushing the old location.

**q**

See the **quit** command.

**quit**

Exit **xfs\_db**.

**ring** [*index*]

Show position ring (if no *index* argument is given), or move to a specific entry in the position ring given by *index*.

**sb** [*agno*]

Set current address to SB header in allocation group *agno*. If no *agno* is given, use the current allocation group number.

**source** *source-file*

Process commands from *source-file*. **source** commands can be nested.

**stack** View the location stack.**type** [*type*]

Set the current data type to *type*. If no argument is given, show the current data type. The possible data types are: **agf**, **agfl**, **agi**, **attr**, **bmapbta**, **bmapbtd**, **bnobt**, **cntbt**, **data**, **dir**, **dir2**, **dqblk**, **inobt**, **inode**, **log**, **refcntbt**, **rmapbt**, **rtbitmap**, **rtsummary**, **sb**, **symlink** and **text**. See the TYPES section below for more information on these data types.

**uuid** [*uuid* | *generate* | *rewrite* | *restore*]

Set the filesystem universally unique identifier (UUID). The filesystem UUID can be used by **mount**(8) instead of using a device special file. The *uuid* can be set directly to the desired UUID, or it can be automatically generated using the **generate** option. These options will both write the UUID into every copy of the superblock in the filesystem. On a CRC-enabled filesystem, this will set an incompatible superblock flag, and the filesystem will not be mountable with older kernels. This can be reverted with the **restore** option, which will copy the original UUID back into place and clear the incompatible flag as needed. **rewrite** copies the current UUID from the primary superblock to all secondary copies of the superblock. If no argument is given, the current filesystem UUID is printed.

**version** [*feature* | *versionnum* *features2*]

Enable selected features for a filesystem (certain features can be enabled on an unmounted filesystem, after **mkfs.xfs**(8) has created the filesystem). Support for unwritten extents can be enabled using the **extflg** option. Support for version 2 log format can be enabled using the **log2** option. Support for extended attributes can be enabled using the **attr1** or **attr2** option. Once enabled, extended attributes cannot be disabled, but the user may toggle between **attr1** and **attr2** at will (older kernels may not support the newer version).

If no argument is given, the current version and feature bits are printed. With one argument, this command will write the updated version number into every copy of the superblock in the filesystem. If two arguments are given, they will be used as numeric values for the *versionnum* and *features2* bits respectively, and their string equivalent reported (but no modifications are made).

**write** [**-c** | **-d**] [*field value*] ...

Write a value to disk. Specific fields can be set in structures (struct mode), or a block can be set to data values (data mode), or a block can be set to string values (string mode, for symlink blocks). The operation happens immediately: there is no buffering.

Struct mode is in effect when the current type is structural, i.e. not data. For struct mode, the syntax is "**write** *field value*".

Data mode is in effect when the current type is data. In this case the contents of the block can be shifted or rotated left or right, or filled with a sequence, a constant value, or a random value. In this mode **write** with no arguments gives more information on the allowed commands.

- c** Skip write verifiers and CRC recalculation; allows invalid data to be written to disk.
- d** Skip write verifiers but perform CRC recalculation. This allows invalid data to be written to disk to test detection of invalid data. (This is not possible for some types.)

## TYPES

This section gives the fields in each structure type and their meanings. Note that some types of block cover multiple actual structures, for instance directory blocks.

<b>agf</b>	<p>The AGF block is the header for block allocation information; it is in the second 512-byte block of each allocation group. The following fields are defined:</p> <table> <tr> <td><b>magicnum</b></td><td>AGF block magic number, 0x58414746 ('XAGF').</td></tr> <tr> <td><b>versionnum</b></td><td>version number, currently 1.</td></tr> <tr> <td><b>seqno</b></td><td>sequence number starting from 0.</td></tr> <tr> <td><b>length</b></td><td>size in filesystem blocks of the allocation group. All allocation groups except the last one of the filesystem have the superblock's <b>agblocks</b> value here.</td></tr> <tr> <td><b>bnoroot</b></td><td>block number of the root of the Btree holding free space information sorted by block number.</td></tr> <tr> <td><b>cntroot</b></td><td>block number of the root of the Btree holding free space information sorted by block count.</td></tr> <tr> <td><b>bnolevel</b></td><td>number of levels in the by-block-number Btree.</td></tr> <tr> <td><b>cntlevel</b></td><td>number of levels in the by-block-count Btree.</td></tr> <tr> <td><b>ffirst</b></td><td>index into the AGFL block of the first active entry.</td></tr> <tr> <td><b>flast</b></td><td>index into the AGFL block of the last active entry.</td></tr> <tr> <td><b>flcount</b></td><td>count of active entries in the AGFL block.</td></tr> <tr> <td><b>freeblks</b></td><td>count of blocks represented in the freespace Btrees.</td></tr> <tr> <td><b>longest</b></td><td>longest free space represented in the freespace Btrees.</td></tr> <tr> <td><b>btreeblks</b></td><td>number of blocks held in the AGF Btrees.</td></tr> </table>	<b>magicnum</b>	AGF block magic number, 0x58414746 ('XAGF').	<b>versionnum</b>	version number, currently 1.	<b>seqno</b>	sequence number starting from 0.	<b>length</b>	size in filesystem blocks of the allocation group. All allocation groups except the last one of the filesystem have the superblock's <b>agblocks</b> value here.	<b>bnoroot</b>	block number of the root of the Btree holding free space information sorted by block number.	<b>cntroot</b>	block number of the root of the Btree holding free space information sorted by block count.	<b>bnolevel</b>	number of levels in the by-block-number Btree.	<b>cntlevel</b>	number of levels in the by-block-count Btree.	<b>ffirst</b>	index into the AGFL block of the first active entry.	<b>flast</b>	index into the AGFL block of the last active entry.	<b>flcount</b>	count of active entries in the AGFL block.	<b>freeblks</b>	count of blocks represented in the freespace Btrees.	<b>longest</b>	longest free space represented in the freespace Btrees.	<b>btreeblks</b>	number of blocks held in the AGF Btrees.
<b>magicnum</b>	AGF block magic number, 0x58414746 ('XAGF').																												
<b>versionnum</b>	version number, currently 1.																												
<b>seqno</b>	sequence number starting from 0.																												
<b>length</b>	size in filesystem blocks of the allocation group. All allocation groups except the last one of the filesystem have the superblock's <b>agblocks</b> value here.																												
<b>bnoroot</b>	block number of the root of the Btree holding free space information sorted by block number.																												
<b>cntroot</b>	block number of the root of the Btree holding free space information sorted by block count.																												
<b>bnolevel</b>	number of levels in the by-block-number Btree.																												
<b>cntlevel</b>	number of levels in the by-block-count Btree.																												
<b>ffirst</b>	index into the AGFL block of the first active entry.																												
<b>flast</b>	index into the AGFL block of the last active entry.																												
<b>flcount</b>	count of active entries in the AGFL block.																												
<b>freeblks</b>	count of blocks represented in the freespace Btrees.																												
<b>longest</b>	longest free space represented in the freespace Btrees.																												
<b>btreeblks</b>	number of blocks held in the AGF Btrees.																												
<b>agfl</b>	<p>The AGFL block contains block numbers for use of the block allocator; it is in the fourth 512-byte block of each allocation group. Each entry in the active list is a block number within the allocation group that can be used for any purpose if space runs low. The AGF block fields <b>ffirst</b>, <b>flast</b>, and <b>flcount</b> designate which entries are currently active. Entry space is allocated in a circular manner within the AGFL block. Fields defined:</p> <table> <tr> <td><b>bno</b></td><td>array of all block numbers. Even those which are not active are printed.</td></tr> </table>	<b>bno</b>	array of all block numbers. Even those which are not active are printed.																										
<b>bno</b>	array of all block numbers. Even those which are not active are printed.																												
<b>agi</b>	<p>The AGI block is the header for inode allocation information; it is in the third 512-byte block of each allocation group. Fields defined:</p> <table> <tr> <td><b>magicnum</b></td><td>AGI block magic number, 0x58414749 ('XAGI').</td></tr> <tr> <td><b>versionnum</b></td><td>version number, currently 1.</td></tr> <tr> <td><b>seqno</b></td><td>sequence number starting from 0.</td></tr> <tr> <td><b>length</b></td><td>size in filesystem blocks of the allocation group.</td></tr> <tr> <td><b>count</b></td><td>count of inodes allocated.</td></tr> <tr> <td><b>root</b></td><td>block number of the root of the Btree holding inode allocation information.</td></tr> <tr> <td><b>level</b></td><td>number of levels in the inode allocation Btree.</td></tr> <tr> <td><b>freecount</b></td><td>count of allocated inodes that are not in use.</td></tr> <tr> <td><b>newino</b></td><td>last inode number allocated.</td></tr> <tr> <td><b>dirino</b></td><td>unused.</td></tr> <tr> <td><b>unlinked</b></td><td>an array of inode numbers within the allocation group. The entries in the AGI block are the heads of lists which run through the inode <b>next_unlinked</b> field. These inodes are to be unlinked the next time the filesystem is mounted.</td></tr> </table>	<b>magicnum</b>	AGI block magic number, 0x58414749 ('XAGI').	<b>versionnum</b>	version number, currently 1.	<b>seqno</b>	sequence number starting from 0.	<b>length</b>	size in filesystem blocks of the allocation group.	<b>count</b>	count of inodes allocated.	<b>root</b>	block number of the root of the Btree holding inode allocation information.	<b>level</b>	number of levels in the inode allocation Btree.	<b>freecount</b>	count of allocated inodes that are not in use.	<b>newino</b>	last inode number allocated.	<b>dirino</b>	unused.	<b>unlinked</b>	an array of inode numbers within the allocation group. The entries in the AGI block are the heads of lists which run through the inode <b>next_unlinked</b> field. These inodes are to be unlinked the next time the filesystem is mounted.						
<b>magicnum</b>	AGI block magic number, 0x58414749 ('XAGI').																												
<b>versionnum</b>	version number, currently 1.																												
<b>seqno</b>	sequence number starting from 0.																												
<b>length</b>	size in filesystem blocks of the allocation group.																												
<b>count</b>	count of inodes allocated.																												
<b>root</b>	block number of the root of the Btree holding inode allocation information.																												
<b>level</b>	number of levels in the inode allocation Btree.																												
<b>freecount</b>	count of allocated inodes that are not in use.																												
<b>newino</b>	last inode number allocated.																												
<b>dirino</b>	unused.																												
<b>unlinked</b>	an array of inode numbers within the allocation group. The entries in the AGI block are the heads of lists which run through the inode <b>next_unlinked</b> field. These inodes are to be unlinked the next time the filesystem is mounted.																												
<b>attr</b>	<p>An attribute fork is organized as a Btree with the actual data embedded in the leaf blocks. The root of the Btree is found in block 0 of the fork. The index (sort order) of the Btree is the hash value of the attribute name. All the blocks contain a <b>blkinfo</b> structure at the beginning, see type <b>dir</b> for a description. Nonleaf blocks are identical in format to those for version 1 and version 2 directories, see type <b>dir</b> for a description. Leaf blocks can refer to "local" or "remote" attribute values. Local values are stored directly in the leaf</p>																												



block. Leaf blocks contain the following fields:

<b>hdr</b>	header containing a <b>blkinfo</b> structure <b>info</b> (magic number 0xfbee), a <b>count</b> of active entries, <b>usedbytes</b> total bytes of names and values, the <b>firstused</b> byte in the name area, <b>holes</b> set if the block needs compaction, and array <b>freemap</b> as for <b>dir</b> leaf blocks.
<b>entries</b>	array of structures containing a <b>hashval</b> , <b>nameidx</b> (index into the block of the name), and flags <b>incomplete</b> , <b>root</b> , and <b>local</b> .
<b>nvlist</b>	array of structures describing the attribute names and values. Fields always present: <b>valuelen</b> (length of value in bytes), <b>namelen</b> , and <b>name</b> . Fields present for local values: <b>value</b> (value string). Fields present for remote values: <b>valueblk</b> (fork block number of containing the value).

Remote values are stored in an independent block in the attribute fork. Prior to v5, value blocks had no structure, but in v5 they acquired a header structure with the following fields:

<b>magic</b>	attr3 remote block magic number, 0x5841524d ('XARM').
<b>offset</b>	Byte offset of this data block within the overall attribute value.
<b>bytes</b>	Number of bytes stored in this block.
<b>crc</b>	Checksum of the attribute block contents.
<b>uuid</b>	Filesystem UUID.
<b>owner</b>	Inode that owns this attribute value.
<b>bno</b>	Block offset of this block within the inode's attribute fork.
<b>lsn</b>	Log serial number of the last time this block was logged.
<b>data</b>	The attribute value data.

#### **bmapbt**

Files with many extents in their data or attribute fork will have the extents described by the contents of a Btree for that fork, instead of being stored directly in the inode. Each bmap Btree starts with a root block contained within the inode. The other levels of the Btree are stored in filesystem blocks. The blocks are linked to sibling left and right blocks at each level, as well as by pointers from parent to child blocks. Each block contains the following fields:

<b>magic</b>	bmap Btree block magic number, 0x424d4150 ('BMAP').
<b>level</b>	level of this block above the leaf level.
<b>numrecs</b>	number of records or keys in the block.
<b>leftsib</b>	left (logically lower) sibling block, 0 if none.
<b>rightsib</b>	right (logically higher) sibling block, 0 if none.
<b>recs</b>	[leaf blocks only] array of extent records. Each record contains <b>startoff</b> , <b>startblock</b> , <b>blockcount</b> , and <b>extentflag</b> (1 if the extent is unwritten).
<b>keys</b>	[non-leaf blocks only] array of key records. These are the first key value of each block in the level below this one. Each record contains <b>startoff</b> .
<b>ptrs</b>	[non-leaf blocks only] array of child block pointers. Each pointer is a filesystem block number to the next level in the Btree.

#### **bnobt**

There is one set of filesystem blocks forming the by-block-number allocation Btree for each allocation group. The root block of this Btree is designated by the **bnoroot** field in the corresponding AGF block. The blocks are linked to sibling left and right blocks at each level, as well as by pointers from parent to child blocks. Each block has the following fields:

	<b>magic</b> BNOBT block magic number, 0x41425442 ('ABTB'). <b>level</b> level number of this block, 0 is a leaf. <b>numrecs</b> number of data entries in the block. <b>leftsib</b> left (logically lower) sibling block, 0 if none. <b>rightsib</b> right (logically higher) sibling block, 0 if none. <b>recs</b> [leaf blocks only] array of freespace records. Each record contains <b>startblock</b> and <b>blockcount</b> . <b>keys</b> [non-leaf blocks only] array of key records. These are the first value of each block in the level below this one. Each record contains <b>startblock</b> and <b>blockcount</b> . <b>ptrs</b> [non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.
<b>cntbt</b>	<p>There is one set of filesystem blocks forming the by-block-count allocation Btree for each allocation group. The root block of this Btree is designated by the <b>cntroot</b> field in the corresponding AGF block. The blocks are linked to sibling left and right blocks at each level, as well as by pointers from parent to child blocks. Each block has the following fields:</p> <b>magic</b> CNTBT block magic number, 0x41425443 ('ABTC'). <b>level</b> level number of this block, 0 is a leaf. <b>numrecs</b> number of data entries in the block. <b>leftsib</b> left (logically lower) sibling block, 0 if none. <b>rightsib</b> right (logically higher) sibling block, 0 if none. <b>recs</b> [leaf blocks only] array of freespace records. Each record contains <b>startblock</b> and <b>blockcount</b> . <b>keys</b> [non-leaf blocks only] array of key records. These are the first value of each block in the level below this one. Each record contains <b>blockcount</b> and <b>startblock</b> . <b>ptrs</b> [non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.
<b>data</b>	User file blocks, and other blocks whose type is unknown, have this type for display purposes in <b>xfs_db</b> . The block data is displayed in hexadecimal format.
<b>dir</b>	<p>A version 1 directory is organized as a Btree with the directory data embedded in the leaf blocks. The root of the Btree is found in block 0 of the file. The index (sort order) of the Btree is the hash value of the entry name. All the blocks contain a <b>blkinfo</b> structure at the beginning with the following fields:</p> <b>forw</b> next sibling block. <b>back</b> previous sibling block. <b>magic</b> magic number for this block type. <p>The non-leaf (node) blocks have the following fields:</p> <b>hdr</b> header containing a <b>blkinfo</b> structure <b>info</b> (magic number 0xfebe), the <b>count</b> of active entries, and the <b>level</b> of this block above the leaves. <b>btree</b> array of entries containing <b>hashval</b> and <b>before</b> fields. The <b>before</b> value is a block number within the directory file to the child block, the <b>hashval</b> is the last hash value in that block. <p>The leaf blocks have the following fields:</p> <b>hdr</b> header containing a <b>blkinfo</b> structure <b>info</b> (magic number 0xfeeb), the <b>count</b> of active entries, <b>namebytes</b> (total name string bytes), <b>holes</b> flag (block needs compaction), and <b>freemap</b> (array of <b>base</b> , <b>size</b> entries for free regions).

**entries** array of structures containing **hashval**, **nameidx** (byte index into the block of the name string), and **namelen**.  
**namelist** array of structures containing **inumber** and **name**.

**dir2**

A version 2 directory has four kinds of blocks. Data blocks start at offset 0 in the file. There are two kinds of data blocks: single-block directories have the leaf information embedded at the end of the block, data blocks in multi-block directories do not. Node and leaf blocks start at offset 32GiB (with either a single leaf block or the root node block). Freespace blocks start at offset 64GiB. The node and leaf blocks form a Btree, with references to the data in the data blocks. The freespace blocks form an index of longest free spaces within the data blocks.

A single-block directory block contains the following fields:

**bhdr** header containing **magic** number 0x58443242 ('XD2B') and an array **bestfree** of the longest 3 free spaces in the block (**offset**, **length**).  
**bu** array of union structures. Each element is either an entry or a freespace. For entries, there are the following fields: **inumber**, **namelen**, **name**, and **tag**. For freespace, there are the following fields: **freetag** (0xffff), **length**, and **tag**. The **tag** value is the byte offset in the block of the start of the entry it is contained in.  
**bleaf** array of leaf entries containing **hashval** and **address**. The **address** is a 64-bit word offset into the file.  
**btail** tail structure containing the total **count** of leaf entries and **stale** count of unused leaf entries.

A data block contains the following fields:

**dhdr** header containing **magic** number 0x58443244 ('XD2D') and an array **bestfree** of the longest 3 free spaces in the block (**offset**, **length**).  
**du** array of union structures as for **bu**.

Leaf blocks have two possible forms. If the Btree consists of a single leaf then the freespace information is in the leaf block, otherwise it is in separate blocks and the root of the Btree is a node block. A leaf block contains the following fields:

**lhdr** header containing a **blkinfo** structure **info** (magic number 0xd2f1 for the single leaf case, 0xd2ff for the true Btree case), the total **count** of leaf entries, and **stale** count of unused leaf entries.  
**lents** leaf entries, as for **bleaf**.  
**lbests** [single leaf only] array of values which represent the longest freespace in each data block in the directory.  
**ltail** [single leaf only] tail structure containing **bestcount** count of **lbests**.

A node block is identical to that for types **attr** and **dir**.

A freespace block contains the following fields:

**fhdr** header containing **magic** number 0x58443246 ('XD2F'), **firstdb** first data block number covered by this freespace block, **nvalid** number of valid entries, and **nused** number of entries representing real data blocks.  
**fbests** array of values as for **lbests**.

**dqblk**

The quota information is stored in files referred to by the superblock **uquotino** and **pquotino** fields. Each filesystem block in a quota file contains a constant number of quota entries. The quota entry size is currently 136 bytes, so with a 4KiB filesystem block size there are 30 quota entries per block. The **dquot** command is used to locate these entries

in the filesystem. The file entries are indexed by the user or project identifier to determine the block and offset. Each quota entry has the following fields:

<b>magic</b>	magic number, 0x4451 ('DQ').
<b>version</b>	version number, currently 1.
<b>flags</b>	flags, values include 0x01 for user quota, 0x02 for project quota.
<b>id</b>	user or project identifier.
<b>blk_hardlimit</b>	absolute limit on blocks in use.
<b>blk_softlimit</b>	preferred limit on blocks in use.
<b>ino_hardlimit</b>	absolute limit on inodes in use.
<b>ino_softlimit</b>	preferred limit on inodes in use.
<b>bcount</b>	blocks actually in use.
<b>icount</b>	inodes actually in use.
<b>itimer</b>	time when service will be refused if soft limit is violated for inodes.
<b>btimer</b>	time when service will be refused if soft limit is violated for blocks.
<b>iwarns</b>	number of warnings issued about inode limit violations.
<b>bwarns</b>	number of warnings issued about block limit violations.
<b>rtb_hardlimit</b>	absolute limit on realtime blocks in use.
<b>rtb_softlimit</b>	preferred limit on realtime blocks in use.
<b>rtbcount</b>	realtime blocks actually in use.
<b>rtbtimer</b>	time when service will be refused if soft limit is violated for realtime blocks.
<b>rtbwarns</b>	number of warnings issued about realtime block limit violations.

### **inobt**

There is one set of filesystem blocks forming the inode allocation Btree for each allocation group. The root block of this Btree is designated by the **root** field in the corresponding AGI block. The blocks are linked to sibling left and right blocks at each level, as well as by pointers from parent to child blocks. Each block has the following fields:

<b>magic</b>	INOBT block magic number, 0x49414254 ('IABT').
<b>level</b>	level number of this block, 0 is a leaf.
<b>numrecs</b>	number of data entries in the block.
<b>leftsib</b>	left (logically lower) sibling block, 0 if none.
<b>rightsib</b>	right (logically higher) sibling block, 0 if none.
<b>recs</b>	[leaf blocks only] array of inode records. Each record contains <b>startino</b> allocation-group relative inode number, <b>freecount</b> count of free inodes in this chunk, and <b>free</b> bit-map, LSB corresponds to inode 0.
<b>keys</b>	[non-leaf blocks only] array of key records. These are the first value of each block in the level below this one. Each record contains <b>startino</b> .
<b>ptrs</b>	[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.

### **inode**

Inodes are allocated in "chunks" of 64 inodes each. Usually a chunk is multiple filesystem blocks, although there are cases with large filesystem blocks where a chunk is less than one block. The inode Btree (see **inobt** above) refers to the inode numbers per allocation group. The inode numbers directly reflect the location of the inode block on disk. Use the **inode** command to point **xfs\_db** to a specific inode. Each inode contains four regions: **core**, **next\_unlinked**, **u**, and **a**. **core** contains the fixed information. **next\_unlinked** is separated from the core due to journaling considerations, see type **agi** field

**unlinked.** **u** is a union structure that is different in size and format depending on the type and representation of the file data ("data fork"). **a** is an optional union structure to describe attribute data, that is different in size, format, and location depending on the presence and representation of attribute data, and the size of the **u** data ("attribute fork"). **xfs\_db** automatically selects the proper union members based on information in the inode.

The following are fields in the inode core:

<b>magic</b>	inode magic number, 0x494e ('IN').
<b>mode</b>	mode and type of file, as described in <b>chmod(2)</b> , <b>mknod(2)</b> , and <b>stat(2)</b> .
<b>version</b>	inode version, 1 or 2.
<b>format</b>	format of <b>u</b> union data (0: xfs_dev_t, 1: local file – in-inode directory or symlink, 2: extent list, 3: Btree root, 4: unique id [unused]).
<b>nlinkv1</b>	number of links to the file in a version 1 inode.
<b>nlinkv2</b>	number of links to the file in a version 2 inode.
<b>projid_lo</b>	owner's project id (low word; version 2 inode only). <b>projid_hi</b> owner's project id (high word; version 2 inode only).
<b>uid</b>	owner's user id.
<b>gid</b>	owner's group id.
<b>atime</b>	time last accessed (seconds and nanoseconds).
<b>mtime</b>	time last modified.
<b>ctime</b>	time created or inode last modified.
<b>size</b>	number of bytes in the file.
<b>nblocks</b>	total number of blocks in the file including indirect and attribute.
<b>extsize</b>	basic/minimum extent size for the file.
<b>nextents</b>	number of extents in the data fork.
<b>naextents</b>	number of extents in the attribute fork.
<b>forkoff</b>	attribute fork offset in the inode, in 64-bit words from the start of <b>u</b> .
<b>aformat</b>	format of <b>a</b> data (1: local attribute data, 2: extent list, 3: Btree root).
<b>dmevmask</b>	DMAPI event mask.
<b>dmstate</b>	DMAPI state information.
<b>newrtbm</b>	file is the realtime bitmap and is "new" format.
<b>prealloc</b>	file has preallocated data space after EOF.
<b>realtime</b>	file data is in the realtime subvolume.
<b>gen</b>	inode generation number.

The following fields are in the **u** data fork union:

<b>bmbt</b>	bmap Btree root. This looks like a <b>bmapbtd</b> block with redundant information removed.
<b>bm</b>	array of extent descriptors.
<b>dev</b>	dev_t for the block or character device.
<b>sfd</b>	shortform (in-inode) version 1 directory. This consists of a <b>hdr</b> containing the <b>parent</b> inode number and a <b>count</b> of active entries in the directory, followed by an array <b>list</b> of <b>hdr.count</b> entries. Each such entry contains <b>inumber</b> , <b>namelen</b> , and <b>name</b> string.
<b>sfd2</b>	shortform (in-inode) version 2 directory. This consists of a <b>hdr</b> containing a <b>count</b> of active entries in the directory, an <b>iscount</b> of entries with inumbers that don't fit in a 32-bit value, and the <b>parent</b> inode number, followed by an array <b>list</b> of <b>hdr.count</b> entries. Each such entry contains <b>namelen</b> ,

	a saved <b>offset</b> used when the directory is converted to a larger form, a <b>name</b> string, and the <b>inumber</b> .																
<b>symlink</b>	symbolic link string value.																
The following fields are in the <b>a</b> attribute fork union if it exists:																	
<b>bmbt</b>	bmap Btree root, as above.																
<b>bm</b>	array of extent descriptors.																
<b>sfattr</b>	shortform (in-inode) attribute values. This consists of a <b>hdr</b> containing a <b>totsize</b> (total size in bytes) and a <b>count</b> of active entries, followed by an array <b>list</b> of <b>hdr.count</b> entries. Each such entry contains <b>namelen</b> , <b>valuelen</b> , <b>root</b> flag, <b>name</b> , and <b>value</b> .																
<b>log</b>	Log blocks contain the journal entries for XFS. It's not useful to examine these with <b>xfs_db</b> , use <b>xfs_logprint(8)</b> instead.																
<b>refcntbt</b>	There is one set of filesystem blocks forming the reference count Btree for each allocation group. The root block of this Btree is designated by the <b>refcntroot</b> field in the corresponding AGF block. The blocks are linked to sibling left and right blocks at each level, as well as by pointers from parent to child blocks. Each block has the following fields: <table> <tr> <td><b>magic</b></td><td>REFC block magic number, 0x52334643 ('R3FC').</td></tr> <tr> <td><b>level</b></td><td>level number of this block, 0 is a leaf.</td></tr> <tr> <td><b>numrecs</b></td><td>number of data entries in the block.</td></tr> <tr> <td><b>leftsib</b></td><td>left (logically lower) sibling block, 0 if none.</td></tr> <tr> <td><b>rightsib</b></td><td>right (logically higher) sibling block, 0 if none.</td></tr> <tr> <td><b>recs</b></td><td>[leaf blocks only] array of reference count records. Each record contains <b>startblock</b>, <b>blockcount</b>, and <b>refcount</b>.</td></tr> <tr> <td><b>keys</b></td><td>[non-leaf blocks only] array of key records. These are the first value of each block in the level below this one. Each record contains <b>startblock</b>.</td></tr> <tr> <td><b>ptrs</b></td><td>[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.</td></tr> </table>	<b>magic</b>	REFC block magic number, 0x52334643 ('R3FC').	<b>level</b>	level number of this block, 0 is a leaf.	<b>numrecs</b>	number of data entries in the block.	<b>leftsib</b>	left (logically lower) sibling block, 0 if none.	<b>rightsib</b>	right (logically higher) sibling block, 0 if none.	<b>recs</b>	[leaf blocks only] array of reference count records. Each record contains <b>startblock</b> , <b>blockcount</b> , and <b>refcount</b> .	<b>keys</b>	[non-leaf blocks only] array of key records. These are the first value of each block in the level below this one. Each record contains <b>startblock</b> .	<b>ptrs</b>	[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.
<b>magic</b>	REFC block magic number, 0x52334643 ('R3FC').																
<b>level</b>	level number of this block, 0 is a leaf.																
<b>numrecs</b>	number of data entries in the block.																
<b>leftsib</b>	left (logically lower) sibling block, 0 if none.																
<b>rightsib</b>	right (logically higher) sibling block, 0 if none.																
<b>recs</b>	[leaf blocks only] array of reference count records. Each record contains <b>startblock</b> , <b>blockcount</b> , and <b>refcount</b> .																
<b>keys</b>	[non-leaf blocks only] array of key records. These are the first value of each block in the level below this one. Each record contains <b>startblock</b> .																
<b>ptrs</b>	[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.																
<b>rmapbt</b>	There is one set of filesystem blocks forming the reverse mapping Btree for each allocation group. The root block of this Btree is designated by the <b>rmaproot</b> field in the corresponding AGF block. The blocks are linked to sibling left and right blocks at each level, as well as by pointers from parent to child blocks. Each block has the following fields: <table> <tr> <td><b>magic</b></td><td>RMAP block magic number, 0x524d4233 ('RMB3').</td></tr> <tr> <td><b>level</b></td><td>level number of this block, 0 is a leaf.</td></tr> <tr> <td><b>numrecs</b></td><td>number of data entries in the block.</td></tr> <tr> <td><b>leftsib</b></td><td>left (logically lower) sibling block, 0 if none.</td></tr> <tr> <td><b>rightsib</b></td><td>right (logically higher) sibling block, 0 if none.</td></tr> <tr> <td><b>recs</b></td><td>[leaf blocks only] array of reference count records. Each record contains <b>startblock</b>, <b>blockcount</b>, <b>owner</b>, <b>offset</b>, <b>attr_fork</b>, <b>bmbt_block</b>, and <b>unwritten</b>.</td></tr> <tr> <td><b>keys</b></td><td>[non-leaf blocks only] array of double-key records. The first ("low") key contains the first value of each block in the level below this one. The second ("high") key contains the largest key that can be used to identify any record in the subtree. Each record contains <b>startblock</b>, <b>owner</b>, <b>offset</b>, <b>attr_fork</b>, and <b>bmbt_block</b>.</td></tr> <tr> <td><b>ptrs</b></td><td>[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.</td></tr> </table>	<b>magic</b>	RMAP block magic number, 0x524d4233 ('RMB3').	<b>level</b>	level number of this block, 0 is a leaf.	<b>numrecs</b>	number of data entries in the block.	<b>leftsib</b>	left (logically lower) sibling block, 0 if none.	<b>rightsib</b>	right (logically higher) sibling block, 0 if none.	<b>recs</b>	[leaf blocks only] array of reference count records. Each record contains <b>startblock</b> , <b>blockcount</b> , <b>owner</b> , <b>offset</b> , <b>attr_fork</b> , <b>bmbt_block</b> , and <b>unwritten</b> .	<b>keys</b>	[non-leaf blocks only] array of double-key records. The first ("low") key contains the first value of each block in the level below this one. The second ("high") key contains the largest key that can be used to identify any record in the subtree. Each record contains <b>startblock</b> , <b>owner</b> , <b>offset</b> , <b>attr_fork</b> , and <b>bmbt_block</b> .	<b>ptrs</b>	[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.
<b>magic</b>	RMAP block magic number, 0x524d4233 ('RMB3').																
<b>level</b>	level number of this block, 0 is a leaf.																
<b>numrecs</b>	number of data entries in the block.																
<b>leftsib</b>	left (logically lower) sibling block, 0 if none.																
<b>rightsib</b>	right (logically higher) sibling block, 0 if none.																
<b>recs</b>	[leaf blocks only] array of reference count records. Each record contains <b>startblock</b> , <b>blockcount</b> , <b>owner</b> , <b>offset</b> , <b>attr_fork</b> , <b>bmbt_block</b> , and <b>unwritten</b> .																
<b>keys</b>	[non-leaf blocks only] array of double-key records. The first ("low") key contains the first value of each block in the level below this one. The second ("high") key contains the largest key that can be used to identify any record in the subtree. Each record contains <b>startblock</b> , <b>owner</b> , <b>offset</b> , <b>attr_fork</b> , and <b>bmbt_block</b> .																
<b>ptrs</b>	[non-leaf blocks only] array of child block pointers. Each pointer is a block number within the allocation group to the next level in the Btree.																
<b>rtbitmap</b>	If the filesystem has a realtime subvolume, then the <b>rbmino</b> field in the superblock refers to a file that contains the realtime bitmap. Each bit in the bitmap file controls the																

allocation of a single realtime extent (set == free). The bitmap is processed in 32-bit words, the LSB of a word is used for the first extent controlled by that bitmap word. The **atime** field of the realtime bitmap inode contains a counter that is used to control where the next new realtime file will start.

### rtsummary

If the filesystem has a realtime subvolume, then the **rsumino** field in the superblock refers to a file that contains the realtime summary data. The summary file contains a two-dimensional array of 16-bit values. Each value counts the number of free extent runs (consecutive free realtime extents) of a given range of sizes that starts in a given bitmap block. The size ranges are binary buckets (low size in the bucket is a power of 2). There are as many size ranges as are necessary given the size of the realtime subvolume. The first dimension is the size range, the second dimension is the starting bitmap block number (adjacent entries are for the same size, adjacent bitmap blocks).

### sb

There is one sb (superblock) structure per allocation group. It is the first disk block in the allocation group. Only the first one (block 0 of the filesystem) is actually used; the other blocks are redundant information for **xfs\_repair**(8) to use if the first superblock is damaged. Fields defined:

<b>magicnum</b>	superblock magic number, 0x58465342 ('XFSB').
<b>blocksize</b>	filesystem block size in bytes.
<b>dblocks</b>	number of filesystem blocks present in the data subvolume.
<b>rblocks</b>	number of filesystem blocks present in the realtime subvolume.
<b>rextents</b>	number of realtime extents that <b>rblocks</b> contain.
<b>uuid</b>	unique identifier of the filesystem.
<b>logstart</b>	starting filesystem block number of the log (journal). If this value is 0 the log is "external".
<b>rootino</b>	root inode number.
<b>rbmino</b>	realtime bitmap inode number.
<b>rsumino</b>	realtime summary data inode number.
<b>rextsize</b>	realtime extent size in filesystem blocks.
<b>agblocks</b>	size of an allocation group in filesystem blocks.
<b>agcount</b>	number of allocation groups.
<b>rbmblocks</b>	number of realtime bitmap blocks.
<b>logblocks</b>	number of log blocks (filesystem blocks).
<b>versionnum</b>	filesystem version information. This value is currently 1, 2, 3, or 4 in the low 4 bits. If the low bits are 4 then the other bits have additional meanings. 1 is the original value. 2 means that attributes were used. 3 means that version 2 inodes (large link counts) were used. 4 is the bitmask version of the version number. In this case, the other bits are used as flags (0x0010: attributes were used, 0x0020: version 2 inodes were used, 0x0040: quotas were used, 0x0080: inode cluster alignment is in force, 0x0100: data stripe alignment is in force, 0x0200: the <b>shared_vn</b> field is used, 0x1000: unwritten extent tracking is on, 0x2000: version 2 directories are in use).
<b>sectsize</b>	sector size in bytes, currently always 512. This is the size of the superblock and the other header blocks.
<b>inodesize</b>	inode size in bytes.
<b>inopblock</b>	number of inodes per filesystem block.
<b>fname</b>	obsolete, filesystem name.
<b>fpack</b>	obsolete, filesystem pack name.
<b>blocklog</b>	log2 of <b>blocksize</b> .
<b>sectlog</b>	log2 of <b>sectsize</b> .

	<b>inodelog</b>	log2 of <b>inodesize</b> .
	<b>inopblog</b>	log2 of <b>inopblock</b> .
	<b>agblklog</b>	log2 of <b>agblocks</b> (rounded up).
	<b>rextslog</b>	log2 of <b>rextents</b> .
	<b>inprogress</b>	<b>mkfs.xfs(8)</b> or <b>xfs_copy(8)</b> aborted before completing this filesystem.
	<b>imax_pct</b>	maximum percentage of filesystem space used for inode blocks.
	<b>icount</b>	number of allocated inodes.
	<b>ifree</b>	number of allocated inodes that are not in use.
	<b>fdblocks</b>	number of free data blocks.
	<b>frextents</b>	number of free realtime extents.
	<b>uquotino</b>	user quota inode number.
	<b>pquotino</b>	project quota inode number; this is currently unused.
	<b>qflags</b>	quota status flags (0x01: user quota accounting is on, 0x02: user quota limits are enforced, 0x04: quotacheck has been run on user quotas, 0x08: project quota accounting is on, 0x10: project quota limits are enforced, 0x20: quotacheck has been run on project quotas).
	<b>flags</b>	random flags. 0x01: only read-only mounts are allowed.
	<b>shared_vn</b>	shared version number (shared readonly filesystems).
	<b>inoalignmt</b>	inode chunk alignment in filesystem blocks.
	<b>unit</b>	stripe or RAID unit.
	<b>width</b>	stripe or RAID width.
	<b>dirblklog</b>	log2 of directory block size (filesystem blocks).
<b>symlink</b>	Symbolic link blocks are used only when the symbolic link value does not fit inside the inode. The block content is just the string value. Bytes past the logical end of the symbolic link value have arbitrary values.	
<b>text</b>	User file blocks, and other blocks whose type is unknown, have this type for display purposes in <b>xfs_db</b> . The block data is displayed in two columns: Hexadecimal format and printable ASCII chars.	

## DIAGNOSTICS

Many messages can come from the **check (blockget)** command. If the filesystem is completely corrupt, a core dump might be produced instead of the message

*device is not a valid filesystem*

If the filesystem is very large (has many files) then **check** might run out of memory. In this case the message

**out of memory**

is printed.

The following is a description of the most likely problems and the associated messages. Most of the diagnostics produced are only meaningful with an understanding of the structure of the filesystem.

**agf\_freeblks *n*, counted *m* in ag *a***

The freeblocks count in the allocation group header for allocation group *a* doesn't match the number of blocks counted free.

**agf\_longest *n*, counted *m* in ag *a***

The longest free extent in the allocation group header for allocation group *a* doesn't match the longest free extent found in the allocation group.

**agi\_count *n*, counted *m* in ag *a***

The allocated inode count in the allocation group header for allocation group *a* doesn't match the number of inodes counted in the allocation group.



**agi\_freecount *n*, counted *m* in ag *a***

The free inode count in the allocation group header for allocation group *a* doesn't match the number of inodes counted free in the allocation group.

**block *a/b* expected inum 0 got *i***

The block number is specified as a pair (allocation group number, block in the allocation group). The block is used multiple times (shared), between multiple inodes. This message usually follows a message of the next type.

**block *a/b* expected type unknown got *y***

The block is used multiple times (shared).

**block *a/b* type unknown not expected****SEE ALSO**

**mkfs.xfs(8), xfs\_admin(8), xfs\_copy(8), xfs\_logprint(8), xfs\_metadump(8), xfs\_ncheck(8), xfs\_repair(8), mount(8), chmod(2), mknod(2), stat(2), xfs(5).**