

Information Source:

Kevin Goodman

kevin@colovirt.com

<http://blog.colovirt.com>

Linux/Unix/NAS/Filesystems: Inodes (Part 1) - Checking Availability And High Level Overview

Inodes really tell you how many file handles (files) that can be created on a file system. Most people will never exceed the default setting when the file system is created, nor even know that one is set. I will eventually go into more detail concerning this topic here on the blog. The majority (not all) of file systems that are used on Linux and Unix do not support dynamic inode allocation. What this means is that if you exceed the inode limit of a file system before the storage space, the remainder will be un-usable. That is until some of the current files are removed.

So here is the really bad part. The inodes on ext2 and ext3 (Linux default type) are statically set when the file system is formatted. You can not go back and change the max inode settings. The exceptions to this that I know of are as follows:

- **Riser4**
- **VxFS**
- **XFS**
- **JFS**
- **WAFL (NetApp proprietary)**
- **XZFS**

If you are running one of the above and have max inodes issue, you can correct it.

I have been working with computers for over 15 years and have only ran into this problem once. Luckily, it occurred on a NetApp NAS device that had the ability to increase this value on the live file system. The main killer here are tons of small files. In this case, the file system for that NFS share was 40 gigabytes in size and default was ~1 million inode limit. The quick fix for the issue was to increase this to 3 million.

As far as a ext2 and 3 go, the following shows how to query a file system for relevant inode information

```
root@testbox:~# tune2fs -l /dev/sda1
```

```
tune2fs 1.41.3 (12-Oct-2008)
```

```
Filesystem volume name:      <none>
Last mounted on:            <not available>
Filesystem UUID:            56161dd8-9d1d-4c54-851d-938bb88ce6d4
Filesystem magic number:    0xEF53
Filesystem revision #:      1 (dynamic)
Filesystem features:        has_journal ext_attr resize_inode dir_index filetype needs_recovery
                             sparse_super large_file
Filesystem flags:           signed_directory_hash
Default mount options:     (none)
Filesystem state:          clean
Errors behavior:           Continue
Filesystem OS type:        Linux
Inode count:               4685824
```

```

Block count:          18731782
Reserved block count: 936589
Free blocks:         15534374
Free inodes:         4459463
First block:         0
Block size:          4096
Fragment size:       4096
Reserved GDT blocks: 1019
Blocks per group:    32768
Fragments per group: 32768
Inodes per group:    8192
Inode blocks per group: 256
Filesystem created:  Mon Sep 29 16:25:20 2008
Last mount time:     Fri Jan 23 14:27:02 2009
Last write time:     Fri Jan 23 14:27:02 2009
Mount count:         4
Maximum mount count: 33
Last checked:        Thu Jan 15 09:00:37 2009
Check interval:      15552000 (6 months)
Next check after:    Tue Jul 14 10:00:37 2009
Reserved blocks uid: 0 (user root)
Reserved blocks gid: 0 (group root)
First inode:         11
Inode size:          128
Journal inode:       8
First orphan inode:  2908742
Default directory hash: tea
Directory Hash Seed: a6544c5xxxxxxxxxxxxxxxxxxxx
Journal backup:      inode blocks

```

The above is good to know, especially to check an un-mounted file system. The command below shows a friendlier formatted output

```
root@testbox:~# df -i
```

Filesystem	Inodes	Iused	Ifree	Iuse%	Mounted on
/dev/sda1	4685824	226361	4459463	5%	/
tmpfs	222201	4	222197	1%	/lib/init/rw
varrun	222201	64	222137	1%	/var/run
varlock	222201	5	222196	1%	/var/lock
udev	222201	5142	217059	3%	/dev
tmpfs	222201	5	222196	1%	/dev/shm
/dev/sdb1	61063168	1116	61062052	1%	/media/disk

As you can see, there are no issues to be worried about on this test computer. Most systems administrators perform centralized monitoring of disk usage at a disk space level (capacity). On highly used servers that utilize locally stored and/or direct attached storage, it is a good idea to have a script check and report on the available inodes.

Notes: This information is provided for a high level overview concerning inodes. More in-depth information will be provided in up-coming posts.