

Why is the initial replication taking more space than expected?

Why is the initial replication taking more space than expected?

Question

How much space is the initial replication going to take?

Often users are surprised to see their initial backup image taking more space than they expected. Meaning the backup image is larger than the amount of disk usage of the partition they have backed up. This is because CDP performs replications at the sector level. While unallocated disk sectors are excluded, the formatting of the file system must still be included, in order to get a true sector level image.

This formatting includes items like super block, block groups, block and inode allocation maps, and directory inodes. This data is not reflected in most disk usage reporting tools...The output of "# df -h" is an example.

Solution

Please see the following example:

```
[root@test-linux ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/sda3 73G 3.6G 66G 6% /
/dev/sda1 101M 9.7M 86M 11% /boot
none 251M 0 2 51M 0% /dev/shm
```

In the example above, "df -h" shows only 3.6G of used data on /dev/sda3. This size value includes only file contents. To get all allocated blocks, we must use a different utility, and make some additional calculations.

```
[root@test-linux ~]# dumpe2fs /dev/sda3
dumpe2fs 1.35 (28-Feb-2004)
Filesystem volume name: /
Last mounted on:
Filesystem UUID: aa613f36-ce53-4158-a512-fef68aabfd3a
Filesystem magic number: 0xEF53
Filesystem revision #: 1 (dynamic)
Filesystem features: has_journal ext_attr filetype needs_recovery sparse_super
Default mount options: (none)
```

```
Filesystem state: clean
Errors behavior: Continue
Filesystem OS type: Linux
Inode count: 9629696
Block count: 77031672
Reserved block count: 3851583
Free blocks: 72007352
Free inodes: 9506828
First block: 1
Block size: 1024
```

In this example, we are interested in Block count, free blocks, and block size.

To get the total number of blocks allocated we take:

block count (77031672) - free blocks (72007352) = 5024320 blocks

To get the amount of data CDP will backup:

5024320 blocks * Block size (1024) = 5144903680 bytes ~4.79 GB

Compare the 4.79 GB to the "df -h" output of 3.6GB. That extra ~1GB is the space used by the ext3 file system.

Related Articles

Page: [Database Disk Image Malformed \(Knowledge Base\)](#) Labels: [database](#), [disk](#), [image](#), [malformed](#), [disksafe](#), [hardware](#), [faults](#)

Page: [The storage cache reported an error reading the native disk \(Knowledge Base\)](#) Labels: [storage](#), [cache](#), [ngetvolume](#), [native](#), [disk](#)

Page: [Why is the initial replication taking more space than expected? \(Knowledge Base\)](#) Labels: [disk](#), [disk_safes](#), [storage](#)

Page: [Why is the initial replication taking more space than expected? \(Knowledge Base\)](#) Labels: [disk](#), [disk_safes](#), [storage](#)

Page: [The storage cache reported an error reading the native disk \(Knowledge Base\)](#) Labels: [storage](#), [cache](#), [ngetvolume](#), [native](#), [disk](#)

Page: [Why is the initial replication taking more space than expected? \(Knowledge Base\)](#) Labels: [disk](#), [disk_safes](#), [storage](#)

Page: [Backup Manager fails to start with INFO - Storage Cache Service starting \(Knowledge Base\)](#) Labels: [sbm](#), [startup](#), [storage](#), [cache](#)
