



Figure 6.1. Accessing Snapshot

### 6.3.7. Tuning Performance

This section provides details regarding improving the system performance in an SMB environment. The various enhancements tasks can be classified into:

- ✧ Enabling Metadata Caching to improve the performance of SMB access of Red Hat Gluster Storage volumes.
- ✧ Enhancing Directory Listing Performance
- ✧ Enhancing File/Directory Create Performance

More detailed information for each of this is provided in the sections ahead.

#### 6.3.7.1. Enabling Metadata Caching

Enable metadata caching to improve the performance of directory operations. Execute the following commands from any one of the nodes on the trusted storage pool in the order mentioned below.



### Note

If majority of the workload is modifying the same set of files and directories simultaneously from multiple clients, then enabling metadata caching might not provide the desired performance improvement.

1. Execute the following command to enable metadata caching and cache invalidation:

```
# gluster volume set <volname> group metadata-cache
```

This is group set option which sets multiple volume options in a single command.

2. To increase the number of files that can be cached, execute the following command:

```
# gluster volume set <VOLNAME> network.inode-lru-limit <n>
```

*n*, is set to 50000. It can be increased if the number of active files in the volume is very high. Increasing this number increases the memory footprint of the brick processes.

### 6.3.7.2. Enhancing Directory Listing Performance

The directory listing gets slower as the number of bricks/nodes increases in a volume, though the file/directory numbers remain unchanged. By enabling the parallel readdir volume option, the performance of directory listing is made independent of the number of nodes/bricks in the volume. Thus, the increase in the scale of the volume does not reduce the directory listing performance.



### Note

You can expect an increase in performance only if the distribute count of the volume is 2 or greater and the size of the directory is small (< 3000 entries). The larger the volume (distribute count) greater is the performance benefit.

To enable parallel readdir execute the following commands:

1. Verify if the **performance.readdir-ahead** option is enabled by executing the following command:

```
# gluster volume get <VOLNAME> performance.readdir-ahead
```

If the **performance.readdir-ahead** is not enabled then execute the following command:

```
# gluster volume set <VOLNAME> performance.readdir-ahead on
```

2. Execute the following command to enable **parallel-readdir** option:

```
# gluster volume set <VOLNAME> performance.parallel-readdir on
```

The **cluster.lookup-optimize** configuration option enables DHT lookup optimization. To enable this option run the following command:

```
# gluster volume set VOLNAME cluster.lookup-optimize <on/off>\
```



### Note

The configuration takes effect for newly created directories immediately post setting the above option. For existing directories, a rebalance is required to ensure the volume is in balance before DHT applies the optimization on older directories.

## 21.6. Replication

If a system is configured for two ways, active-active replication, write throughput will generally be half of what it would be in a non-replicated configuration. However, read throughput is generally improved by replication, as reads can be delivered from either storage node.

## 21.7. Directory Operations

In order to improve the performance of directory operations of Red Hat Gluster Storage volumes, the maximum metadata (stat, xattr) caching time on the client side is increased to 10 minutes, without compromising on the consistency of the cache.

Significant performance improvements can be achieved in the following workloads by enabling metadata caching:

- ✧ Listing of directories (recursive)
- ✧ Creating files
- ✧ Deleting files
- ✧ Renaming files

### 21.7.1. Enabling Metadata Caching

Enable metadata caching to improve the performance of directory operations. Execute the following commands from any one of the nodes on the trusted storage pool in the order mentioned below.



### Note

If majority of the workload is modifying the same set of files and directories simultaneously from multiple clients, then enabling metadata caching might not provide the desired performance improvement.

1. Execute the following command to enable metadata caching and cache invalidation:

```
# gluster volume set <volname> group metadata-cache
```

This is group set option which sets multiple volume options in a single command.

2. To increase the number of files that can be cached, execute the following command:

```
# gluster volume set <VOLNAME> network.inode-lru-limit <n>
```

*n*, is set to 50000. It can be increased if the number of active files in the volume is very high. Increasing this number increases the memory footprint of the brick processes.