

# Tape I/O Horizontal Scalability with Oracle Hierarchical Storage Manager 6.0

Oracle Hierarchical Storage Manager (Oracle HSM) is now horizontally scalable. This capability enables users to increase the total throughput of the archive by 10x compared to older releases of Oracle HSM and archive multiple petabytes of data per day. I/O from Oracle's StorageTek QFS Client servers can be used to archive and stage files to and from the disk cache to tape, preventing the metadata server (MDS) from becoming a bottleneck. Up to nine StorageTek QFS Clients running Oracle Solaris, in addition to the MDS, can be used to archive data to tape, versus a single server in the previous architecture.

## Hardware and Software Engineered to Work Together

### Benefits

Improve data ingest rates:

- » 10x increase in the amount of data per day that can be archived (use 10 servers instead of 1 to move data to tape).
- » Multiple PBs of data can be archived per day.
- » Faster retrieval of data from tape.

Migrate data faster to new storage technology within the archive:

- » Enabling more servers to read data from existing media and writing it to new media shortens the time needed to migrate to new tape drive technology.

Enable use of commodity servers:

- » Distributing the workload across many servers enables use of a smaller and less costly metadata server.
- » If a backup MDS is used, that server can be smaller as well.

Increase utilization of existing server assets:

- » Client servers are often underutilized. By enabling them to archive data to tape, users get more value from them even when they are not being used to read and write data to the archive.

- Archive up to 10x more data per day and retrieve data faster.
- Increase utilization of existing server assets.
- Migrate data to new tape drive technology faster.

### How It Works

Previously the I/O from StorageTek QFS Client servers could be used only to write data to the disk cache of an Oracle HSM archive. To make archive copies of data on tape, all of the I/O had to be driven by the MDS. For large archives with high data ingest rates, this means that a large and expensive MDS is required. Most archives employ a failover MDS as well, meaning that two expensive servers are required.

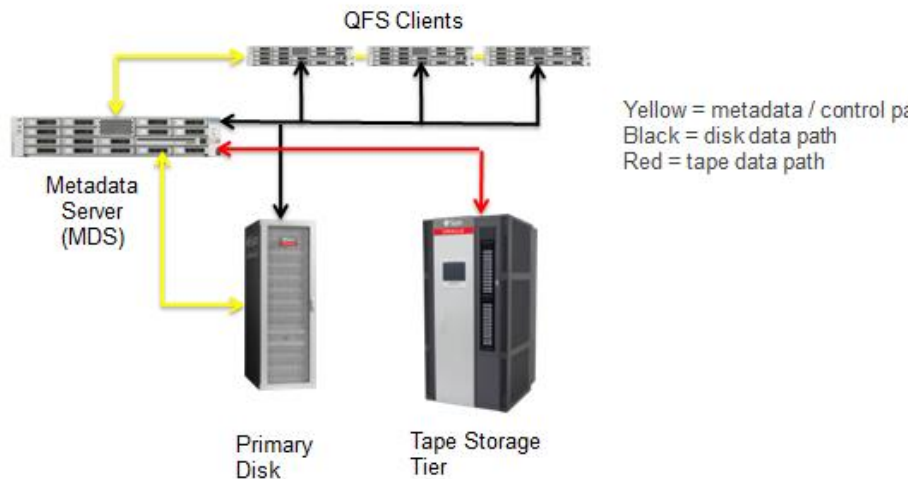


Figure 1: Oracle HSM 5.3 and previous versions, without tape I/O horizontal scalability

Now the I/O from StorageTek QFS Client servers also can be used to read/write data to tape. Up to 9 StorageTek QFS Clients can run on Oracle Solaris and act as data movers. (Hundreds of additional client servers can be added to the shared StorageTek QFS file system, but they cannot act as data movers.) I/O from the metadata server will continue to be used to read/write to tape. Because 10 servers (1 MDS + 9 clients) can be used instead of only the MDS, data ingest rates can be increased by 10x. All metadata operations are still handled by the MDS. When a data mover client reads or writes file from tape, it notifies the MDS so the file metadata can be updated.

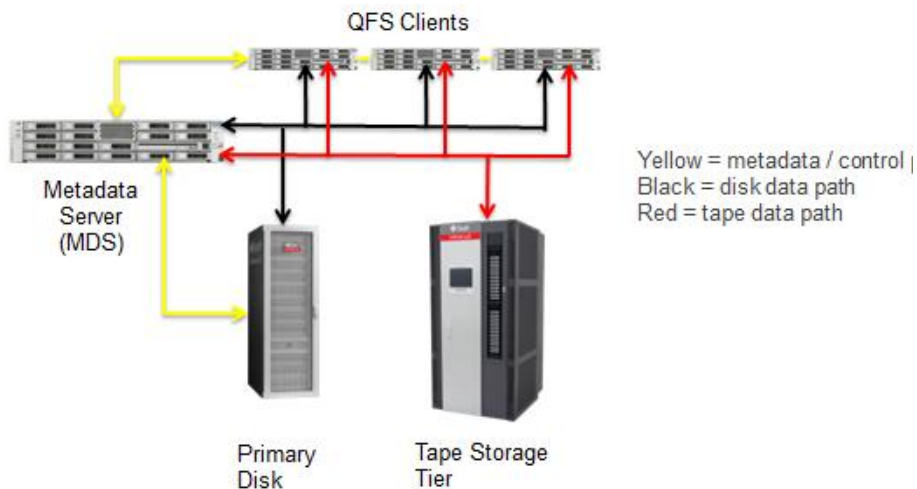


Figure 2: Oracle HSM 6.0, with tape I/O horizontal scalability

CONNECT WITH US

-  [blogs.oracle.com/oracle](http://blogs.oracle.com/oracle)
-  [facebook.com/oracle](http://facebook.com/oracle)
-  [twitter.com/oracle](http://twitter.com/oracle)
-  [oracle.com](http://oracle.com)

FOR MORE INFORMATION  
Contact: 1.800.ORACLE1



**Hardware and Software, Engineered to Work Together**

Copyright © 2015, Oracle and/or its affiliates. All rights reserved. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. 0115

Data can be ingested into the archive and written to the disk cache by any of the StorageTek QFS Clients, whether they are acting as data movers or traditional StorageTek QFS Clients. Only the StorageTek QFS Clients that are acting as data movers and the MDS can read and write data to tape. Which server supplies the I/O to create the tape copies is determined using a round-robin method. In other words, I/O for the first tape archive copy of a file is provided by the MDS, the second from data mover client 1, the third from data client 2, and so on, up to the last data mover client, and then back to the MDS.

In addition to selecting which StorageTek QFS Client servers can act as data movers, users also can select which tape drives will be used by the data movers. Users have the option to select whether the tape drives will archive data, stage data, or do both when using the I/O from the data movers. This provides flexibility in how the archive system can be configured.

Users also can select which files or archive sets will take advantage of the data movers. Users may take advantage of this functionality to optimize the utilization of their tape drives by having the data movers write large files to tape but have smaller files archived using only the MDS.

## CONTACT US

For more information about Oracle's storage solutions for digital archiving, visit [oracle.com](http://oracle.com) call +1.800.ORACLE1 to speak to an Oracle representative.

### CONNECT WITH US

-  [blogs.oracle.com/oracle](http://blogs.oracle.com/oracle)
-  [facebook.com/oracle](http://facebook.com/oracle)
-  [twitter.com/oracle](http://twitter.com/oracle)
-  [oracle.com](http://oracle.com)

FOR MORE INFORMATION  
Contact: 1.800.ORACLE1