



Datadobi DobiMiner vs. Robocopy & Rsync N2N Migration Benchmark Testing

April 2016

Document: DataDobi N2N Migration Benchmark Testing (April 2016)
Authors: M. Baquiran, D. Wren
Company: PassMark Software
Date: 26 April 2016
File: Datadobi_N2N_Migration_Benchmark_Testing_April_2016_Edition_1.docx
Edition: 1

Table of Contents

TABLE OF CONTENTS	2
REVISION HISTORY	3
EXECUTIVE SUMMARY	4
PRODUCTS AND VERSIONS	5
PERFORMANCE METRICS SUMMARY	6
BENCHMARK 1: INITIAL MIGRATION.....	6
BENCHMARK 2: DELTA ITERATION.....	6
TEST RESULTS	7
BENCHMARK 1A – INITIAL MIGRATION (NFS).....	7
BENCHMARK 1B – INITIAL MIGRATION (SMB).....	7
BENCHMARK 2A – DELTA ITERATION (NFS).....	8
BENCHMARK 2B – DELTA ITERATION (SMB).....	8
CONTACT DETAILS	9
APPENDIX 1: TEST ENVIRONMENT	10
APPENDIX 2: METHODOLOGY DESCRIPTION	13
TESTING NOTES.....	13
SETTINGS AND PREPARATION.....	13
BENCHMARK 1: INITIAL MIGRATION (FIRST SCAN AND COPY).....	13
BENCHMARK 2: DELTA ITERATION.....	15

Revision History

Rev	Revision History	Date
Edition 1	Initial version of this report, includes results for Datadobi vs. Robocopy and Rsync.	26 April 2016

Executive Summary

PassMark Software® conducted objective performance testing to compare three tools suitable for NAS to NAS migrations on both Windows and Linux in April 2016. An initial as well as a subsequent delta migration was performed using each tool. This report presents our results and findings as a result of performance benchmark testing conducted on these tools.

Products and Versions

For all products, we have tested the most current and publicly available version of each tool. The names and versions of products are given below:

Manufacturer	Product Name	Release Year	Product Version	Date Tested
Datadobi	DobiMiner	2016	3.5.29	April 2016
Microsoft	Robocopy	2016	XP027 5.1.10.1027	April 2016
N/A (Linux Utility)	Rsync	2016	3.0.6	April 2016

Performance Metrics Summary

To compare the performance of each product, two separate benchmarks were measured. Firstly, the time it takes to scan and copy over a large number of files (13.8GB) from one NAS to another NAS was measured. Once this was complete, files were added (1.38GB) to the original data set on the source NAS. A subsequent copy between the same source and destination folders was then run and the time it takes to reconcile this change on the target NAS was measured.

Datadobi's DobiMiner offers a single solution for both SMB and NFS migrations. These respective capabilities are comparable to that of Robocopy (a Windows tool) and Rsync (a Unix/Linux tool). Thus, we have made the following comparisons:

Benchmark 1: Initial Migration

This test measures the time it takes for the tool to run an initial scan and copy of all files in the dataset from the source NAS to the target NAS. It was carried out on all three tools to make the following comparisons:

- a) DobiMiner using an SMB proxy vs. Robocopy
- b) DobiMiner using an NFS proxy vs. Rsync

Benchmark 2: Delta Iteration

This test is taken after Benchmark 1 has been completed. It measures the time it takes for the tool to copy over changes made to the original data set (a 10% addition) from the source NAS to the target NAS. This was carried out on all three tools to make the following comparisons:


- a) DobiMiner using an SMB proxy vs. Robocopy
- b) DobiMiner using an NFS proxy vs. Rsync

Test Results

In the following charts, Dobiminer's results have been highlighted in green.


Benchmark 1a – Initial Migration (NFS)

The following chart compares the time to run an initial scan and copy of the dataset from the source NAS to the target NAS using NFS. The times measured are displayed in seconds.




Benchmark 1b – Initial Migration (SMB)

The following chart compares the time to run an initial scan and copy of the dataset from the source NAS to the target NAS using SMB. The times measured are displayed in seconds.




Benchmark 2a – Delta Iteration (NFS)

The following chart compares the time to run a subsequent iteration in which an addition to the dataset is scanned for on the source NAS, and then copied over to the target NAS using NFS. The times measured are displayed in seconds.



Benchmark 2b – Delta Iteration (SMB)

The following chart compares the time to run a subsequent iteration in which an addition to the dataset is scanned for on the source NAS, and then copied over to the target NAS using SMB. The times measured are displayed in seconds.



Disclaimer and Disclosure

This report only covers versions of products that were available at the time of testing. The tested versions are as noted in the “Products and Versions” section of this report. The products we have tested are not an exhaustive list of all products available in these categories.

This report only covers results obtained from one particular test environment configuration. In real life there are many permutations of hardware, software & networking setups. There is also great variation in the file sets that one might encounter during migrations between NASs. Different computing environments will give different performance than those presented here.

Datadobi funded the production of this report. The metrics included in the report were selected by Datadobi.

Disclaimer of Liability

While every effort has been made to ensure that the information presented in this report is accurate, PassMark Software Pty Ltd assumes no responsibility for errors, omissions, or out-of-date information and shall not be liable in any manner whatsoever for direct, indirect, incidental, consequential, or punitive damages resulting from the availability of, use of, access of, or inability to use this information.

Trademarks

All trademarks are the property of their respective owners.

Contact Details

PassMark Software Pty Ltd

Level 5

63 Foveaux St.

Surry Hills, 2010

Sydney, Australia

Phone + 61 (2) 9690 0444

Fax + 61 (2) 9690 0445

Web www.passmark.com

Appendix 1: Test Environment

The following test environment was used:

Physical Machine to Host Virtual Machines

This physical machine running VMWare ESXi was used to host the three virtual machines listed below.

Model: Generic hardware
O/S: VMWare ESXi Version 2016.03.02
CPU: Intel Xeon E3-1220v2 CPU
Motherboard: Intel S1200BTL Server Motherboard
RAM: 24GB (2 x 4GB + 2x8GB) ECC RAM
Video: 1GB nVIDIA GeForce GT 620M
SSDs (spanned together to make a single volume)
SSD #1: Crucial m4 CT256M4SSD2 256 GB
SSD #2: Corsair Force 3 90GB
Network Cards (one assigned per Virtual Machine)
#1: Intel Gigabit CT Desktop Adaptor
#2: 1 Gbit/sec on board
#3: 1 Gbit/sec on board



Virtual Machine for DobiMiner SMB Proxy

O/S: Windows Server 2008 R2 Standard
RAM: 7GB
Storage: 55 GB



The DobiMiner SMB Proxy service was installed onto a virtual machine running Windows Server 2008.

Virtual Machine for DobiMiner NFS Proxy

O/S: Linux CentOS 6.6
RAM: 8GB
Storage: 24 GB



The DobiMiner NFS Proxy was delivered as an Open Virtualization Archive (.ova) which provided a complete specification of the virtual machine.

Virtual Machine for Controller (DobiMiner Core)

O/S: Linux CentOS 6.6

RAM: 8GB

Storage: 74 GB



The DobiMiner Core Controller was delivered as an Open Virtualization Archive (.ova) which provided a complete specification of the virtual machine.

NAS (Two to be setup – 1 Source and 1 Destination NAS)

Model:

QNAP TS-453U-RP

SSD (x4):

SAMSUNG SSD 850 PRO 2.5" SATA III 512GB configured as a single volume in RAID 10 on both the source and destination NAS.




Network Switch

D-Link Gigabit Switch DGS-1016D (16-Port). Cat6 cables.



Setup Diagram



Note: During testing, the DobiMiner NFS Proxy Virtual Machine and SMB Proxy Virtual Machines were not run simultaneously.

Appendix 2: Methodology Description

Testing Notes

The following acronyms apply throughout this document:

NFS = Network File System. This is a protocol used to allow file access on a remote computer or device over a network. It was created by Sun Microsystems (now owned by Oracle) and is used mostly on Unix & Linux based systems.

SMB = Server Message Block. This is a protocol from IBM / Microsoft for remote file access. Thus it is found mostly on machines running Windows. The SMB implementation in Windows Server 2008 R2 will be used (Version 2.1).

DobiMiner supports both NFS and SMB migrations. Rsync, a Linux/Unix command line tool uses the NFS protocol to copy files. Robocopy, a Windows command line tool, uses the SMB protocol to copy files.

Thus, the following test rounds were carried out:

- Rsync on Linux (NFS only)
- Robocopy on Windows (SMB only)
- DobiMiner on Windows and Linux (SMB and NFS)

Settings and Preparation

- Unless otherwise specified, the default settings for each tool were used.
- All antivirus / malware scanning functions were disabled in the O/S.
- The destination hard drives were empty before the start of the test.
- A predefined set of data files was placed on the source NAS (on 4 x SSDs in RAID 10).
- The necessary shares and read/write permissions were configured on both the source and destination NASs. This includes the following:
 - A user with read and write privileges was added to each NAS. This user is used as the login required by each migration tool to connect to the source and target NAS.
 - NFS access rights were configured to have read and write permissions.

Benchmark 1: Initial Migration (First Scan and Copy)

This metric measures the time taken by the tool to copy all files in the dataset from the source NAS to the target NAS. The copy time also includes any preliminary stages (e.g. scanning) that are required to complete the copy. Timing will begin as soon as the operation is manually initiated. Timing will complete as soon as the operation is indicated to be complete and the changes are reflected in the destination NAS.

For each tool, the following steps were carried out:

- **DobiMiner (SMB and NFS)**

After DobiMiner Core and the NFS and SMB proxies were setup, the migrations were configured via the web console. Migrations were configured to be run using 64 threads on NFS and 48 threads on SMB. Also, copying security descriptors and ACLs was deselected in the SMB migration settings.

- The migration was initiated manually and left to run up until the “Steady State” stage was reached. At this point the files have completed scanned and copied over from the source NAS to the target NAS. The time to complete this process was taken from the detailed results tab viewed from the DobiMiner web console.

- **Rsync (on Linux)**

This test was run from the same (virtual) machine that runs the DobiMiner NFS proxy. Since Rsync does not support a direct file copy between two remote servers, the required folders on the source and destination NAS were shared and then mounted as local folders. This was done using the following commands:

```
sudo mount -t nfs 192.168.2.10:/Migration /mnt/sourcefolder
```

and

```
sudo mount -t nfs 192.168.2.9:/Migration /mnt/targetfolder
```

The copy was then carried out using the command:

```
time rsync -avzh /mnt/sourcefolder/TD1 /mnt/targetfolder
```

The “time” command was prepended to the command to time the execution automatically.

Note: 192.168.2.10 is the source NAS IP address and 192.168.2.9 is the destination NAS IP address.

- **Robocopy (on Windows)**

This test was run from the same machine that runs the DobiMiner SMB proxy. First, both the source and target NAS were logged into via Windows Explorer. The copy was then executed and logged using the following command:

```
Robocopy.exe \\192.168.2.10\Migration\TD1 \\192.168.2.9\Migration\TD1 /E  
/PURGE /COPY:DAT /R:1 /W:1 /V /LOG:sTD.log
```

PassMark’s tool CommandTimer.exe was used to measure and log the copy time, so that the full command was:

```
CommandTimer.exe “Robocopy.exe \\192.168.2.10\Migration\TD1  
\\192.168.2.9\Migration\TD1 /E /PURGE /COPY:DAT /R:1 /W:1 /V /LOG:sTD.log”
```

Note: 192.168.2.10 is the source NAS IP address and 192.168.2.9 is the destination NAS IP address.

Benchmark 2: Delta Iteration

Once Benchmark 1 was completed on each tool, a set of additional files (roughly 10% of the original data set) was added to the original data set on the source NAS. This test measures the time it took for the tool to scan and copy over this additional data from the source NAS to the target NAS. Timing began as soon as the delta iteration was initiated, and timing was complete as soon as the new files were copied over and the changes were reflected in the target NAS.

For each tool, the following steps were carried out:

- **DobiMiner (SMB and NFS)**

This test was run as a subsequent iteration following the Initial Scan & Copy (i.e. Benchmark 1).

Before the delta iteration was initiated, the initial migration was confirmed to be in the “Steady State”. The delta iteration was then started by checking the migration in the **Scheduling > Steady State** tab and pressing **Start**. The time to complete this process was recorded from the detailed results tab viewed from the DobiMiner web console.

- **Rsync (on Linux)**

This test was carried out on the same virtual machine that runs the DobiMiner NFS proxy.

The delta iteration was run by adding the `-u` flag to the copy command, making the command:

```
time rsync -avzh -u /mnt/sourcefolder/TD1 /mnt/targetfolder
```

The `time` command was prepended to the command to time the execution.

- **Robocopy (on Windows)**

By default, Robocopy skips any files that are already on the source and destination, making the command the same as that used in Benchmark 1:

```
Robocopy.exe \\192.168.2.10\Migration\TD1 \\192.168.2.9\Migration\TD1  
/E /PURGE /COPY:DAT /R:1 /W:1 /V /LOG:sTD.log
```

PassMark’s tool CommandTimer.exe was used to measure and log the copy time, so that the full command becomes:

```
CommandTimer.exe “Robocopy.exe \\192.168.2.10\Migration\TD1  
\\192.168.2.9\Migration\TD1 /E /PURGE /COPY:DAT /R:1 /W:1 /V /XO  
/LOG:sTD.log”
```

Initial Migration Data Set

Total Size: 13.8 GB

Number of Files: 1914773

Number of Directories: 40903

The data set consists of four file sizes spread over the source drives in the following proportions:

File Size	Percentage
2KB	96.7465%
200KB	0.9675%
1MB	0.1935%
100MB	0.0019%
Folders	2.0905%

Delta Iteration Data Set

Total Size: 1.38GB (Roughly 10% of the above data set)

Number of Files: 210454

Number of Directories: 3