



Foreign LUN Import Performance

ONTAP

NetApp
September 20, 2021

Table of Contents

- Foreign LUN Import Performance 1
 - Performance enhancements in ONTAP 8.3.1 1
 - Variables that affect Foreign LUN Import migration performance 2
 - Benchmarks for estimating migration durations 2
 - Best practices for Foreign LUN Import migration 2
 - ESXi CAW/ATS remediation 3
 - Host remediation 4
 - Clearing SCSI-3 persistent reservations 5
 - Creating the host to destination zones 7

Foreign LUN Import Performance

You should be aware of certain things that affect the performance of your Foreign LUN Import, such as certain variables, best practices and host remediation.

Performance enhancements in ONTAP 8.3.1

There have been some enhancements to FLI to better manage performance and stop workload starvation from occurring. FLI enhancements in ONTAP 8.3.1 include a new throttle command and LUN import show enhancements showing throughput and QoS policy groups.

The LUN import throttle command is used to limit the maximum speed at which an import can run.

```
cluster::*> lun import throttle -vserver fli_72C -path /vol/flivol/72Clun1
-max-throughput-limit

{<integer>[KB|MB|GB|TB|PB]} Maximum Throughput Limit (per sec)
```

Use the instance switch with `lun import show` to display extended LUN import information, including throttle and QoS information.

```
cluster::*> lun import show -instance

Vserver Name: fli_72C
LUN Path: /vol/flivol/72Clun1
Foreign Disk Serial Number: D0i1E+G8Wg6m
Import Home Node: ontaptme-fc-cluster-01
Import Current Node: ontaptme-fc-cluster-01
Operation In Progress: import
Admin State: stopped
Operational State: stopped
Percent Complete: 0
Blocks Imported: -
Blocks Compared: -
Total Blocks: 6297480
Estimated Remaining Duration: -
Failure Reason: -
Maximum Throughput Limit(per sec): -
Current Throughput (per sec): -
QoS Policy Group: -
```

The values for the `current` throughput show the current rate of throughput of import or verify operations. Users should check this before setting a throttle value. It is empty when not running. The `QoS` policy group

shows the QoS group if the LUN import throttle was used.

Variables that affect Foreign LUN Import migration performance

There are a number of variables that affect how quickly a given migration completes.

These variables include:

- How many concurrent migrations are running between a given source and destination
- Source array capabilities
- Source array load
- Destination array capabilities
- Destination array load
- How much I/O is being generated to the LUN during migration
- The type, bandwidth, and fan-ins/fan-outs on front-end fabrics

For best performance, use no more than 16 concurrent FLI migrations per node.

Given the number of variables that affect migration performance, it is recommended that a number of test migrations be performed. Generally, the larger the test sample, the better the characterization will be. Therefore, we recommend that a number of different sized test migrations be performed in order to get an accurate sampling of throughput performance. Performance data from those tests can then be used to extrapolate timing and durations of the planned production migrations.

Benchmarks for estimating migration durations

For planning purposes certain assumptions can be used to estimate the level of effort and duration of data migrations.

To get an accurate estimate of your actual performance, you should run a number of test migrations of different sizes in order to get accurate performance numbers for your specific environments.



The following benchmarks are strictly for planning purposes and are unlikely to be particularly accurate for specific environments.

Assumptions: Five hours per host migration based on a host with 8 LUNs with a total of 2 TB of data. These parameters provide a planning number of approximately 400 GB per hour.

Best practices for Foreign LUN Import migration

NetApp strongly recommends a professional services or partner professional services, engagement to scope and plan the migration as well as to train customer personnel on how to perform data migrations using Foreign LUN Import (FLI) 7-Mode to ONTAP.

- Perform one or more test migrations at least a week prior to your migration project in order to verify configuration, connectivity, and throughput; uncover any other issues; and validate your methodology.

- For maximum throughput, don't run more than 16 migrations concurrently per node.
- Verification is not required, but we encourage you to verify a subset of the LUNs imported/migrated to validate the import process.
- Use the throughput observed in your test migrations to plan production migration durations.
- For best performance, migrate LUNs during non-peak demand periods.

ESXi CAW/ATS remediation

Online FLI does not support VMware Atomic Test and Set (ATS)/SCSI Compare and Write (CAW). This is important if you are using VMFS5 and your source array supports CAW. To remediate the host you must follow the process outlined in this section.

FLI online LUN relationships don't support ATS/CAW commands, and the VMFS5 file system would fail mounting on the destination ESXi 5.x host. This is the result of a VMware maintaining an ATS bit on the VMFS5 header, which enforces CAW/ATS and will not allow the header to work on a host or array without ATS. The ATS bit is carried in the VMFS header, which is part of the first LUN listed in the *Partitions spanned*. This is the only LUN, if there are multiple extents listed, that needs to be remediated.

If the LUN is shared by more than one host, then updating it on one of the hosts is sufficient. All other hosts automatically update after a rescan. Disabling ATS/CAW will fail if any VM or ESXi active I/O from any of the sharing hosts is running on the LUN. We recommend shutting down the VMs and other host machines sharing the LUN while making the necessary ATS/CAW changes. This action can be performed at the beginning of disruptive parts of the host repoint/cutover listed in the *Disruptive Cutover* section of the appropriate FLI workflow.

If the LUN is shared by more than one host, all hosts will need to be offline while the ATS bit is enabled or disabled. After enabling or disabling ATS, you will need to refresh LUNs. After you have completed any remapping, you can then bring the hosts back up and verify that you are able to access the LUNs.

If you are running a previous version of VMFS or upgraded from a previous version, then you shouldn't have to perform any remediation. If you do need to either enable or disable ATS/CAW, you can use the commands listed below. However, neither will work if the VM is active and there is any I/O running to the VMFS5 datastore. We recommend shutting down the host machine, making the necessary ATS/CAW changes, and performing the rest of the disruptive parts of the host repoint/cutover listed in the *Disruptive Cutover* section of the appropriate FLI workflow.

You can check ATS/CAW status by running the following command:

```

~ # vmkfstools -Ph -v 1 /vmfs/volumes/fli-orig-3
VMFS-5.58 file system spanning 1 partitions.
File system label (if any): fli-orig-3
Mode: public ATS-only
Capacity 99.8 GB, 58.8 GB available, file block size 1 MB, max file size
62.9 TB
Volume Creation Time: Wed Jun 10 13:56:05 2015
Files (max/free): 130000/129979
Ptr Blocks (max/free): 64512/64456
Sub Blocks (max/free): 32000/31995
Secondary Ptr Blocks (max/free): 256/256
File Blocks (overcommit/used/overcommit %): 0/41931/0
Ptr Blocks (overcommit/used/overcommit %): 0/56/0
Sub Blocks (overcommit/used/overcommit %): 0/5/0
Volume Metadata size: 804159488
UUID: 557841f5-145136df-8de6-0025b501a002
Partitions spanned (on "lvm"):
naa.60080e50001f83d4000003075576b218:1
Is Native Snapshot Capable: YES
OBJLIB-LIB: ObjLib cleanup done.
~ # vmkfstools -Ph -v 1 /vmfs/volumes/fli-orig-3
~ # vmkfstools --help

```

If the mode had listed the word public only, no remediation would be necessary. In the case above public ATS-only means that the ATS is enabled and needs to be disabled until the import is completed, at which time it can be reenabled.

To disable ATS/CAW on a LUN, use the following command:

```
# vmkfstools --configATSOOnly 0 /vmfs/devices/disks/naa.aaaaaaaaaaaaaaaa
```

To reenable ATS/CAW, after the migration is complete, use:

```
# vmkfstools --configATSOOnly 1 /vmfs/devices/disks/naa.aaaaaaaaaaaaaaaa
```

Host remediation

Depending on the type of migration, host remediation may take place inline to the migration (Foreign LUN Import online, and 7-Mode to ONTAP) or might it occur after the migration is complete (Foreign LUN Import offline).

Use the for remediation steps for different host OSs. Consult your gap analysis, put together during the planning and analysis phases, and appropriate NetApp and vendor documentation for steps specific to your migration.



FLI uses the same remediation procedures that would be used with the 7MTT. Therefore, it makes sense to leverage the same remediation document rather than document those procedures multiple times in different places.



For CAW remediation, use the ESXi CAW/ATS remediation process.

Related information

[SAN host transition and remediation](#)

Clearing SCSI-3 persistent reservations

If you have a Windows cluster, you need to remove SCSI-3 reservations for the quorum disk, even if all clustered hosts are offline.

If you attempt to label the source LUN as a foreign disk, the following error message will appear:

```
Error: command failed: The specified foreign disk has SCSI persistent reservations. Disk serial number: "6006016021402700787BAC217B44E411". Clear the reservation using the "storage disk remove-reservation" command before creating the import relationship.
```

You can remove SCSI-3 reservations for the quorum disk on the NetApp controller using the `storage disk remove-reservation` command:

```
storage disk remove-reservation -disk disk_name
```

Here is a snippet showing this error and the remediation for it:

```

cluster-4b:*> lun offline -vserver fli_cluster -path
/vol/fli_volume/cluster_CVS
cluster-4b:*> lun import create -vserver fli_cluster -path
/vol/fli_volume/cluster_CVS -foreign-disk 6006016021402700787BAC217B44E411
Error: command failed: The specified foreign disk is not marked as
foreign. Disk serial number: "6006016021402700787BAC217B44E411".

cluster-4b:*> sto disk show -disk DGC-1.6 -fields serial-number,is-
foreign
  (storage disk show)
disk is-foreign serial-number
-----
DGC-1.6 true 6006016021402700787BAC217B44E411

cluster-4b:*> lun import create -vserver fli_cluster -path
/vol/fli_volume/cluster_CVS -foreign-disk 6006016021402700787BAC217B44E411

Error: command failed: The specified foreign disk has SCSI persistent
reservations. Disk serial number: "6006016021402700787BAC217B44E411".
Clear the reservation using the "storage disk remove-reservation" command
before creating the import relationship.

cluster-4b:*> storage disk remove-reservation -disk DGC-1.6
cluster-4b:*> lun import create -vserver fli_cluster -path
/vol/fli_volume/cluster_CVS -foreign-disk 6006016021402700787BAC217B44E411
cluster-4b:*> lun online -vserver fli_cluster -path
/vol/fli_volume/cluster_CVS
cluster-4b:*> lun import show

vserver foreign-disk path operation admin operational percent in progress
state state complete
-----
-----
fli_cluster 6006016021402700787BAC217B44E411 /vol/fli_volume/cluster_CVS
import stopped stopped 0

cluster-4b:*> lun import start -vserver fli_cluster -path
/vol/fli_volume/cluster_CVS
cluster-4b:*> lun import show

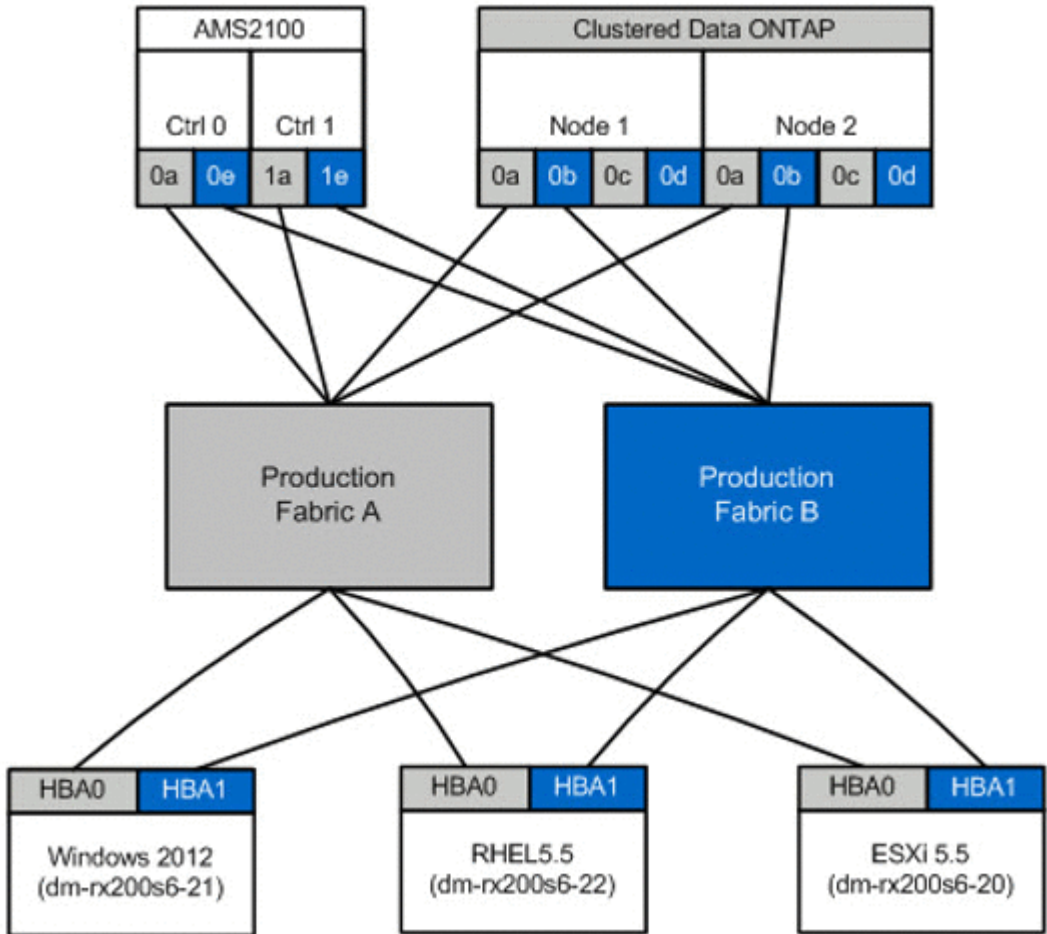
vserver foreign-disk path operation admin operational percent in progress
state state complete
-----
-----
fli_cluster 6006016021402700787BAC217B44E411 /vol/fli_volume/cluster_CVS
import started in_progress 7

```


Creating the host to destination zones

You will need to create the host to destination zones. There are two types of production fabric, fabric A and fabric B.

The following is an illustration of the host and destination storage zoning.



Production zones available in production fabric A.

Zone	WWPN	Zone Members
Zone: rx21_flicDOT	21:00:00:24:ff:30:14:c5	RX21 HBA 0
	20:01:00:a0:98:2f:94:d1	flicDOT lif1
	20:03:00:a0:98:2f:94:d1	flicDOT lif3
Zone: rx22_flicDOT	21:00:00:24:ff:30:04:85	RX22 HBA 0
	20:01:00:a0:98:2f:94:d1	flicDOT lif1
	20:03:00:a0:98:2f:94:d1	flicDOT lif3

Zone	WWPN	Zone Members
Zone: rx20_flicDOT	21:00:00:24:ff:30:03:ea	RX20 HBA 0
	20:01:00:a0:98:2f:94:d1	flicDOT lif1
	20:03:00:a0:98:2f:94:d1	flicDOT lif3

Production zones available in production fabric B.

Zone	WWPN	Zone Members
Zone: rx21_flicDOT	21:00:00:24:ff:30:14:c4	RX21 HBA 1
	20:02:00:a0:98:2f:94:d1	flicDOT lif2
	20:04:00:a0:98:2f:94:d1	flicDOT lif4
Zone: rx22_flicDOT	21:00:00:24:ff:30:04:84	RX22 HBA 1
	20:02:00:a0:98:2f:94:d1	flicDOT lif2
	20:04:00:a0:98:2f:94:d1	flicDOT lif4
Zone: rx20_flicDOT	21:00:00:24:ff:30:03:eb	RX20 HBA 1
	20:02:00:a0:98:2f:94:d1	flicDOT lif2
	20:04:00:a0:98:2f:94:d1	flicDOT lif4

Brocade fabric in production fabric A example

The following is an example of a Brocade fabric in production fabric A.

1. Create the zone in production fabric A.

```
zoneCreate "rx21_flicDOT", "21:00:00:24:ff:30:14:c5"
zoneAdd "rx21_flicDOT", "20:01:00:a0:98:2f:94:d1"
zoneAdd "rx21_flicDOT", "20:03:00:a0:98:2f:94:d1"
zoneCreate "rx22_flicDOT", "21:00:00:24:ff:30:04:85"
zoneAdd "rx22_flicDOT", "20:01:00:a0:98:2f:94:d1"
zoneAdd "rx22_flicDOT", "20:03:00:a0:98:2f:94:d1"
zoneCreate "rx20_flicDOT", "21:00:00:24:ff:30:03:ea"
zoneAdd "rx20_flicDOT", "20:01:00:a0:98:2f:94:d1"
zoneAdd "rx20_flicDOT", "20:03:00:a0:98:2f:94:d1"
```

2. Activate the zone in production fabric A.

```
cfgAdd "PROD_LEFT", "rx21_flicDOT"  
cfgAdd "PROD_LEFT", "rx22_flicDOT"  
cfgAdd "PROD_LEFT", "rx20_flicDOT"  
cfgEnable "PROD_LEFT"  
cfgSave
```

Brocade fabric in production fabric B example

The following is an example of a Brocade fabric in production fabric B

1. Create the zone in production fabric B.

```
zoneCreate "rx21_flicDOT", "21:00:00:24:ff:30:14:c4"  
zoneAdd "rx21_flicDOT", "20:02:00:a0:98:2f:94:d1"  
zoneAdd "rx21_flicDOT", "20:04:00:a0:98:2f:94:d1"  
zoneCreate "rx22_flicDOT", "21:00:00:24:ff:30:04:84"  
zoneAdd "rx22_flicDOT", "20:02:00:a0:98:2f:94:d1"  
zoneAdd "rx22_flicDOT", "20:04:00:a0:98:2f:94:d1"  
zoneCreate "rx20_flicDOT", "21:00:00:24:ff:30:03:eb"  
zoneAdd "rx20_flicDOT", "20:02:00:a0:98:2f:94:d1"  
zoneAdd "rx20_flicDOT", "20:04:00:a0:98:2f:94:d1"
```

2. Activate the zone in production fabric B.

```
cfgAdd "PROD_RIGHT", "rx21_flicDOT"  
cfgAdd "PROD_RIGHT", "rx22_flicDOT"  
cfgAdd "PROD_RIGHT", "rx20_flicDOT"  
cfgEnable "PROD_RIGHT"  
cfgSave
```

Cisco fabric in production fabric A example

The following is an example of a Cisco fabric in production fabric A.

Steps

1. Create the zone in production fabric A.

```
conf t
zone name rx21_flicDOT vsan 10
member pwn 21:00:00:24:ff:30:14:c5
member pwn 20:01:00:a0:98:2f:94:d1
member pwn 20:03:00:a0:98:2f:94:d1
zone name rx22_flicDOT vsan 10
member pwn 21:00:00:24:ff:30:04:85
member pwn 20:01:00:a0:98:2f:94:d1
member pwn 20:03:00:a0:98:2f:94:d1
zone name rx20_flicDOT vsan 10
member pwn 21:00:00:24:ff:30:03:ea
member pwn 20:01:00:a0:98:2f:94:d1
member pwn 20:03:00:a0:98:2f:94:d1
exit
end
```

2. Activate the zone in production fabric A.

```
conf t
zoneset name PROD_LEFT vsan 10
member rx21_flicDOT
member rx22_flicDOT
member rx20_flicDOT
exit
zoneset activate name PROD_LEFT vsan 10
end
copy running-config startup-config
```

Cisco fabric in production fabric B example

The following is an example of a Cisco fabric in production fabric B.

Steps

1. Create the zone in production fabric B.

```
conf t
zone name rx21_flicDOT vsan 10
member pwn 21:00:00:24:ff:30:14:c4
member pwn 20:02:00:a0:98:2f:94:d1
member pwn 20:04:00:a0:98:2f:94:d1
zone name rx22_flicDOT vsan 10
member pwn 21:00:00:24:ff:30:04:84
member pwn 20:02:00:a0:98:2f:94:d1
member pwn 20:04:00:a0:98:2f:94:d1
zone name rx20_flicDOT vsan 10
member pwn 21:00:00:24:ff:30:03:eb
member pwn 20:02:00:a0:98:2f:94:d1
member pwn 20:04:00:a0:98:2f:94:d1
exit
end
```

2. Activate the zone in production fabric B.

```
conf t
zoneset name PROD_RIGHT vsan 10
member rx21_flicDOT
member rx22_flicDOT
member rx20_flicDOT
exit
zoneset activate name PROD_RIGHT vsan 10
end
copy running-config startup-config
```

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.