



Target queue depth customization

ONTAP FlexArray

NetApp
September 20, 2021

Table of Contents

- Target queue depth customization 1
 - Guidelines for specifying the appropriate target queue depth 1
 - Setting the target queue depth (ONTAP prior to 8.2) 1
 - Setting the target queue depth 2
 - Displaying target queue depth statistics 2
 - Target port utilization policy settings 4

Target queue depth customization

The target queue depth defines the number of ONTAP commands that can be queued (outstanding) on a storage array target port. ONTAP supplies a default value. For most deployments, the default target queue depth is appropriate; however, you can change it to correct performance issues.

The default target queue depth differs with different releases of ONTAP:

- For ONTAP, the default is 512.
- For all releases prior to Data ONTAP 8.2, the default is 256.

When a storage array is configured with multiple initiators sharing target ports, you do not want the outstanding commands in the queue buffer from all initiators together to exceed what the storage array can handle. Otherwise, the performance of all hosts can suffer. Storage arrays differ in the number of commands they can handle in the queue buffer.



Target queue depth might also be referred to as “target queue length,” “Q-Depth,” or “Max Throttle.”

Guidelines for specifying the appropriate target queue depth

You need to consider the impact of all the initiators accessing the storage array port when you are planning the configuration for a specific ONTAP system or a specific host that does not run ONTAP.

If your deployment includes more than one initiator on a target port, you need to consider the total number of commands sent to a target port by all initiators when setting the target queue depth.

Guidelines for specifying the appropriate target queue depth are as follows:

- Do not configure a value of 0 (zero).

A value of 0 means that there is no limit on the outstanding commands.

- Consider the volume of commands that specific initiators would be likely to send to the target port.

You could then configure higher values for initiators likely to send a greater number of requests and a lower value for initiators likely to send a lesser number of requests.

- Configure hosts that do not run ONTAP according to the guidelines provided for those hosts.
- Consider setting the target queue depth on a per-target-port basis when workloads differ across ports.

Setting the target queue depth (ONTAP prior to 8.2)

The default target queue depth is acceptable for most implementations, but you can change the default value if required.

This setting is per ONTAP system, and it applies to all target ports on all storage arrays. For ONTAP systems running version prior to 8.2, you can use this option.

Step

1. Use the following option to set the target queue depth: `options disk.target_port.cmd_queue_depth value`

Setting the target queue depth

The default target queue depth is acceptable for most implementations, but can be changed if performance issues are encountered.

You can set the target queue depth by storage array or on a per-target port basis.

Step

1. Use one of the following commands to set the target port queue depth on all target ports or on a specific target port of a storage array.

If you want to...	Use this command sequence...
Set the target port queue depth on all target ports for a storage array	<code>set advanced storage array port modify -name array_name -max-queue-depth value</code>
Set the target port queue depth on a specific target port on a storage array	<code>set advanced storage array port modify -name array_name -wwnn value -wwpn value -max-queue-depth value</code>

For more information about these commands, see the man pages.

Displaying target queue depth statistics

If you suspect that a target queue depth setting is causing performance issues on your storage array, you should check the value that is set for the queue depth and check the state of requests on the FC initiator ports.

There are different levels of detail that you can access to determine whether there are problems processing requests on target ports. The following steps describe how to determine the current setting for the target port queue depth, determine whether there are requests waiting on the ports, and display detailed port statistics to help you understand the workload on the port.

Steps

1. Use the `storage array show` command with the `-instance` parameter to show the current value of the target port queue depth.

```

> set advanced
> storage array show -instance

Name: HP2
      Prefix: HP-2
      Vendor: HP
      Model: HSV300
      options:
        Serial Number: 50014380025d1500
Target Port Queue Depth: 512
      LUN Queue Depth: 32
      Upgrade Pending: false
      Optimization Policy: eALUA
      Affinity: aaa
      Error Text: -

```

2. Use the `storage array port show -fields max-queue-depth` command to show the queue depth setting for each port on the storage array.

```

> set advanced
> storage array port show -fields max-queue-depth

name                wwnn                wwpn                max-queue-depth
-----
EMC_SYMMETRIX_1    50060480000001a0    50060480000001a0    -
EMC_SYMMETRIX_1    50060480000001a1    50060480000001a1    -
EMC_SYMMETRIX_1    50060480000001b0    50060480000001b0    -
EMC_SYMMETRIX_1    50060480000001b1    50060480000001b1    256

```

A value of "-" for Max Queue Depth indicates that the port does not have a specific max queue depth setting and is using the value set at the storage array level.

3. Use the `storage array port show` command to display performance information about storage array target ports.

The results of this command help you determine whether there are performance problems related to the ports. The `%busy` and `%waiting` values provide a high-level view of the performance on a port. If these values show a high percentage of requests waiting to be processed or show that the port is busy for a great percentage of time, then you might want to investigate further into the state of the port.

```
vgv3070f51::*> storage array port show
```

```
Array Name: HP2  
WWNN: 50014380025d1500  
WWPN: 50014380025d1508  
Connection Type: fabric  
Switch Port: vgbr300s70:9  
Link Speed: 4 GB/s  
Max Queue Depth: -
```

Node	Initiator	Count	LUN IOPS	KB/s	%busy	%waiting	Link Errs
vgv51-02	0a	21	2	53	0	0	0
vgv51-01	0a	21	2	48	1	0	0

4. You can obtain more detailed information about ports by using the `storage array port show -fields` command with the `average-latency-per-iop`, `average-pending`, `average-waiting`, `max-pending`, or `max-waiting` fields.

Target port utilization policy settings

ONTAP can detect events of resource contention, such as I/O queue being full, commands being timed out or HBA resource getting exhausted, on a target port.

You can set the target port utilization policies by using the `storage array port modify` command if you detect such events on a given array target port.

The following table describes the two utilization policies associated with a target port:

Policy	Description
normal	When ONTAP detects target port resource contention on a given array target port, it reduces the target port's queue depth and throttles I/O to the target port. In this mode, the reduction of the target port queue depth is lesser than the defer policy for each target port resource contention event. The subsequent rise in the target port queue depth is faster than the defer policy. normal is the default policy.
defer	When ONTAP detects target port resource contention on a given array target port, it reduces the target port's queue depth and throttles I/O to the target port. In this mode, the reduction of the target port queue depth is greater than the normal policy for each target port resource contention event. The subsequent rise in the target port queue depth is slower than the normal utilization policy.

Output examples to view and modify array target port utilization policies

The following command displays the target port utilization policy associated with an array target port:

```
vgv3170_jon::> storage array port show -wwnn 2703750270235
    Array Name: HITACHI_DF600F_1
        WWNN: 2703750270235
        WWPN: 2703750270235
    Connection Type: fabric
        Switch Port: vgbr300s89:9
        Link Speed: 4 GB/s
    Max Queue Depth: 1024
    Utilization Policy: defer

                                LUN
Link
Node                               Initiator  Count  IOPS  KB/s  %busy  %waiting
Errs
-----
-----
0          vgv3170f54a             0a      2     50   1956   85     0
0          vgv3170f54b             0a      2    350  15366  100    40
```

By default, the I/O policy for a given array target port is **normal**. You can modify the I/O policy associated with the port by running the following command:

```
vgv3070f50ab::> storage array port modify -wwpn 50014380025d1509
-utilization-policy ?

normal      This policy aggressively competes for target port resources,
in effect competing with other hosts.
            (normal)

defer       This policy does not aggressively compete for target port
resources, in effect deferring to other hosts.

vgv3070f50ab::> storage array port modify -wwpn 50014380025d1509
-utilization-policy defer
1 record updated.
```

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.