



Managing Performance Service Levels

Active IQ Unified Manager 9.14

NetApp
March 05, 2024

Table of Contents

- Performance Service Levels 1
 - Managing workloads by assigning PSLs 2
 - Creating and editing Performance Service Levels 6

Performance Service Levels

A Performance Service Level (PSL) enables you to define the performance and storage objectives for a workload. You can assign a PSL to a workload when initially creating the workload, or afterwards by editing the workload.

The management and monitoring of storage resources are based on Service Level Objectives (SLOs). SLOs are defined by service level agreements that are based on required performance and capacity. In Unified Manager, SLOs refer to the PSL definitions of the applications that are running on NetApp storage. Storage services are differentiated based on the performance and utilization of the underlying resources. A PSL is a description of the storage service objectives. A PSL enables the storage provider to specify the performance and capacity objectives for the workload. When you assign a PSL on a workload, the corresponding workload on ONTAP is managed by its performance and capacity objectives. Each PSL is governed by peak, expected, and absolute minimum IOPs, and expected latency.

Unified Manager has the following types of PSLs:

- **System-defined:** Unified Manager provides a few canned policies that cannot be changed. These predefined PSLs are:

- Extreme Performance
- Performance
- Value

The Extreme Performance, Performance, and Value PSLs are applicable for most of the common storage workloads in a data center.

Unified Manager also offers three Performance Service Levels for database applications. These are extremely high-performance PSLs that support bursty IOPS and are appropriate for database applications with the highest throughput demand.

- Extreme for Database Logs
- Extreme for Database Shared Data
- Extreme for Database Data

- **User-defined:** If the predefined Performance Service Levels do not meet your requirements, then you can create new PSLs to meet your needs. For information, see [Creating and editing Performance Service Levels](#).
- **Beyond Extreme:** The Beyond Extreme PSLs are the system-recommended PSLs that are suggested for workloads that demand IOPs higher than Extreme. The workloads are internally analyzed based on their IOPS, capacity, and latency, and a Beyond Extreme PSL is recommended for each of these workloads on the **Storage > Workloads > All Workloads** screen. You can apply the PSLs to the workloads to ensure an optimum performance.

The IOPs parameters for the workloads are dynamically generated, depending on the workload behavior, and appended to the name of the Beyond Extreme PSL in the format `Beyond Extreme <number-(peak IOPS/TB)> <number(expected IOPS/TB)>`. For example, if the system determines a workload to have the peak and expected IOPs as 106345 and 37929 respectively, the Beyond Extreme PSL that is generated for the workload is named as `Beyond Extreme 106345 37929`. Though these PSLs are recommended by the system, when you assign them to workloads, these PSLs are labeled as User-defined in type.

Managing workloads by assigning PSLs

You can access PSLs from the **Policies > Performance Service Levels** page and by using the storage provider APIs. Managing storage workloads by assigning PSLs to them is convenient as you do not have to individually manage the storage workloads. Any modifications can also be managed by reassigning another PSL rather than managing them individually. Unified Manager helps you to assign PSLs on your workloads based on internal assessment and recommendations.

For information on assigning system-recommended PSLs to workloads, see [Assigning system-recommended PSLs to workloads](#)

The Performance Service Levels page lists the available PSL policies and enables you to add, edit, and delete them.



You cannot modify a PSL that is system-defined or that is currently assigned to a workload. You cannot delete a PSL that is assigned to a workload, or if it is the only available PSL.

This page displays the following information:

Field	Description
Name	Name of the PSL.
Type	Whether the policy is system-defined or user-defined.
Expected IOPS/TB	Minimum number of IOPS that an application is expected to perform on a LUN or file share. Expected IOPS specifies the minimum expected IOPS allocated, based on the storage object allocated size.
Peak IOPS/TB	<p>Maximum number of IOPS that an application can perform on a LUN or file share. Peak IOPS specifies the maximum possible IOPS allocated, based on the storage object allocated size or the storage object used size.</p> <p>Peak IOPS are based on an allocation policy. The allocation policy is either allocated-space or used-space. When the allocation policy is set to allocated-space, the peak IOPS is calculated based on the size of the storage object. When the allocation policy is set to used-space, the peak IOPS is calculated based on the amount of data stored in the storage object, taking into account storage efficiencies. By default, the allocation policy is set to used-space.</p>

Field	Description
Absolute minimum IOPS	<p>The absolute minimum IOPS is used as an override, when the expected IOPS is less than this value. The default values of the system-defined PSLs are the following:</p> <ul style="list-style-type: none"> • Extreme Performance: If expected IOPS \geq 6144/TB, then absolute minimum IOPS = 1000 • Performance: If expected IOPS \geq 2048/TB and $<$ 6144/TB, then absolute minimum IOPS = 500 • Value: If expected IOPS \geq 128/TB and $<$ 2048/TB, then absolute minimum IOPS = 75 <p>The default values of the system-defined database PSLs are the following:</p> <ul style="list-style-type: none"> • Extreme for Database Logs: If expected IOPS \geq 22528, then absolute minimum IOPS = 4000 • Extreme for Database Shared Data: If expected IOPS \geq 16384, then absolute minimum IOPS = 2000 • Extreme for Database Data: If expected IOPS \geq 12288, then absolute minimum IOPS = 2000 <p>The higher value of the absolute minimum IOPS for custom PSLs can be a maximum of 75000. The lower value is calculated as the following:</p> <p>$1000/\text{expected latency}$</p>
Expected latency	Expected latency for storage IOPS in milliseconds per operation (ms/op).
Capacity	Total available and used capacity in the clusters.
Workloads	Number of storage workloads that have been assigned the PSL.

For information about how the peak IOPS and expected IOPS help in achieving consistent differentiated performance on ONTAP clusters, see the following KB article: [What is Performance Budgeting?](#)

Events generated for workloads breaching the threshold defined by PSLs

Note that if workloads exceed the expected latency value for 30% of the time during the previous hour, Unified Manager generates one of the following events to notify you of a potential performance issue:

- Workload Volume Latency Threshold Breached as defined by Performance Service Level Policy
- Workload LUN Latency Threshold Breached as defined by Performance Service Level Policy.

You may want to analyze the workload to see what may be causing the higher latency values.

For more information, see the following links:

- [Volume events](#)
- [What happens when a performance threshold policy is breached](#)
- [How Unified Manager uses workload latency to identify performance issues](#)
- [What performance events are](#)

System-defined PSLs

The following table provides information about the system-defined PSLs:

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme Performance	Provides extremely high throughput at a very low latency Ideal for latency-sensitive applications	1	12288	6144	1000
Performance	Provides high throughput at a low latency Ideal for database and virtualized applications	2	4096	2048	500
Value	Provides high storage capacity and moderate latency Ideal for high-capacity applications such as email, web content, file shares, and backup targets	17	512	128	75

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme for Database Logs	<p>Provides maximum throughput at the lowest latency.</p> <p>Ideal for database applications supporting database logs. This PSL provides the highest throughput because database logs are extremely bursty and logging is constantly in demand.</p>	1	45056	22528	4000
Extreme for Database Shared Data	<p>Provides very high throughput at the lowest latency.</p> <p>Ideal for database applications data that is stored in a common data store, but is shared across databases.</p>	1	32768	16384	2000
Extreme for Database Data	<p>Provides high throughput at the lowest latency.</p> <p>Ideal for database applications data, such as database table information and metadata.</p>	1	24576	12288	2000

Creating and editing Performance Service Levels

When the system-defined Performance Service Levels do not match your workload requirements, you can create your own Performance Service Levels that are optimized for your workloads.

What you'll need

- You must have the Application Administrator role.
- The Performance Service Level name must be unique, and you cannot use the following reserved keywords:

Prime, Extreme, Performance, Value, Unassigned, Learning, Idle, Default, and None.

You create and edit custom Performance Service Levels from the Performance Service Levels page by defining the service level objectives you require for the applications that will access storage.



You cannot modify a Performance Service Level if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies > Performance Service Levels**.
2. In the **Performance Service Levels** page, click the appropriate button depending on whether you want to create a new Performance Service Level or if you want to edit an existing Performance Service Level.

To...	Follow these steps...
Create a new Performance Service Level	Click Add .
Edit an existing Performance Service Level	Select an existing Performance Service Level, and then click Edit .

The page to add or edit a Performance Service Level is displayed.

3. Customize the Performance Service Level by specifying the performance objectives, and then click **Submit** to save the Performance Service Level.

You can apply the new or changed Performance Service Level to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.