

Dell EMC PowerMax: Service Levels for PowerMaxOS

Abstract

This document describes how Dell EMC™ PowerMax service levels for PowerMaxOS allow users to prioritize applications to provide predictable and consistent performance based on business requirements.

September 2019

Revisions

Date	Description
May 2018	Initial release
September 2019	Content update; new template

œ&} [, |^â*^ { ^} œ

Author: Richard Pace

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2018–2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [8/29/2019] [Technical White Paper] [H17108.1]

Table of contents

Revisions.....	2
Table of contents	3
Executive summary.....	4
Terminology	4
1 Service levels for PowerMaxOS.....	5
1.1 Service levels for hybrid and all flash storage	5
1.2 Service level options.....	5
1.3 How service levels work	6
1.4 Data placement with service levels	7
1.5 Setting service levels	7
1.5.1 Creating a storage group with a service level	8
1.5.2 Modifying a storage group	9
1.6 Alerting and compliance	10
2 Service level functionality	13
2.1 Priority applications	13
2.2 Service providers	14
3 Interoperability.....	15
3.1 Local and remote replication	15
3.1.1 Service levels and SRDF.....	15
3.1.2 Service levels and TimeFinder SnapVX.....	15
3.2 Service levels and host I/O limits.....	15
4 Summary	17
A Technical support and resources	18

Executive summary

Organizations constantly need IT infrastructures that can deliver instant access to largely increasing volumes of data. These organizations spend a great deal of time trying to understand new generations of technology and applications to keep up with the high demand of capacity along with providing high-speed predictable access.

With the cost of flash storage becoming economically viable, organizations are now moving away from hybrid array configurations, such as Dell EMC™ VMAX™ Hybrid arrays with both SSD and spinning disk storage, and on to all-flash storage such as Dell EMC [PowerMax](#) and VMAX All Flash arrays. The capacity and performance of PowerMax and VMAX All Flash arrays provides a single-tier system designed to provide high availability to keep up with increasing demand of data access. Having a single tier of flash storage, users need prioritized data access and have the ability to prioritize critical, high-priority applications while managing lower-priority applications. PowerMaxOS provides this ability with the service levels feature which is described in this paper in detail. Service levels allow users to set expectations for applications in order to provide predictable and consistent performance.

Terminology

The following are explanations of terms that are used throughout the paper:

PowerMaxOS: Beginning with 5978, the operating environment run on PowerMax and VMAX All Flash systems.

Storage group (SG): A logical grouping of thin devices that are provisioned and associated with a particular application.

Response time (RT): The total amount of time it takes to respond to a request for service.

Target response time: The average response time expected for the storage group based on the selected service level.

Upper response time limit: The maximum response time specified by the selected service level.

Lower response time limit: The minimum response time specified by the selected service level.

1 Service levels for PowerMaxOS

Service levels for PowerMaxOS ensure that applications have consistent and predictable performance by allowing users to separate storage groups based on performance requirements and business importance. PowerMaxOS allows you to set specific service levels to ensure the highest-priority application response times are not impacted by lower-priority applications. The available service levels are defined in PowerMaxOS and can be applied at the creation of a storage group or can be modified to an existing storage group at any time.

1.1 Service levels for hybrid and all flash storage

The release of the Dell EMC VMAX3™ hybrid arrays combined both solid-state drives with traditional spinning disk drives. The mix of flash and spinning disk storage was cost effective while providing high-speed performance and capacity. To utilize the capacity more effectively and provide predictable performance, the arrays were equipped with a HYPERMAX OS feature known as service level provisioning, which allowed users to prioritize applications at the storage-group level. Having a mix of flash and spinning disk drives allowed for storage groups with higher service levels to reside on faster flash technology while the lower service levels could be on the lower-tiered spinning disks. Having mixed tiers also allowed for frequently accessed data to be transitioned internally to higher-tiered flash for faster access.

With PowerMax and VMAX All Flash systems that have a single tier of flash back-end storage, response time management is used to set performance expectations. Having a single tier allows for all data to be stored on high-speed flash storage without the need to move any data, allowing for better performance with access managed for each storage group with a specified service level.

1.2 Service level options

Service levels are offered with various ranges of performance expectations which are defined by their own characteristics of a target response time. The target response time is the average response time expected for the storage group based on the selected service level. Along with a target response time, service levels also have either an upper response time limit or both an upper and lower response time limit.

The service levels offered are detailed in Table 1. All service levels shown, with the exception of Optimized, have a target response time.

Table 1 Service levels for PowerMaxOS

Service level	Diamond	Platinum	Gold	Silver	Bronze	Optimized
Target response time	✓	✓	✓	✓	✓	
Upper response time limit	✓	✓	✓	✓	✓	
Lower response time limit				✓	✓	



Diamond, Platinum, and Gold: These service levels have the highest priority and performance. Each has an upper response time limit but no lower response time limit which ensures they will be serviced as fast as possible.

Silver and Bronze: These service levels have both an upper and lower limit designed to allow higher-priority service levels to be unaffected. These are managed such that their average response time will be greater than or equal to the lower response time limit.

Optimized: This service level does not have a target response time nor an upper or lower limit. Optimized is designed to use all allowable resources, equal to that of Diamond, and is not managed to assist any other service level. Storage groups set with any other service level will also not be managed to assist storage groups set Optimized. Optimized should be used on systems where application performance and consistency are not of relative importance and should not be mixed on systems with other service levels.

The following table shows the target response time for each service level along with an indication of the lower response time.

Table 2 Service level response times

Service level	Target response time	Lower response time
Diamond	0.6 ms*	None
Platinum	0.8 ms*	None
Gold	1 ms	None
Silver	3.6 ms	~3.6 ms
Bronze	7.2 ms	~7.2 ms
Optimized	N/A	N/A

*Indicates target response time without SCM drives. With SCM, Diamond target response time = 0.4 ms and Platinum target response time = 0.6 ms.

1.3 How service levels work

PowerMaxOS is continually monitoring the system to ensure that any lower-priority applications are minimally disruptive to higher-priority applications. When the higher-priority applications' response times begin to approach the upper limit of the selected service level, the system begins to manage any lower-priority storage groups. The process of monitoring and the management of lower-priority applications both happen in real time.

PowerMaxOS uses real-time machine learning to model workload characteristics. This model provides a predictive function that allows PowerMaxOS to anticipate workload demand for a storage group. With these anticipated workload demands, it can adapt as necessary to changes in block size, write ratio, or I/O load.

A storage group with a higher-priority service level that is affected by any lower-priority storage groups will trigger response-time management to the lower-priority service levels. When the higher-priority storage group reaches its target response time, all lower storage groups will continue to be managed until the lowest-priority storage groups reach their target response time.

The management of any lower-priority service level is imposed by a response-time delay in I/O. The delay gradually increases over time to keep the higher-priority storage group within its respective target response time. The delay gradually decreases to ensure that the higher-priority storage group remains within its response time.

The maximum delay that PowerMaxOS can apply to the response time of a storage group is approximately five times the target response time of the service level.

To view the use cases for service levels, see section 2.

1.4 Data placement with service levels

The PowerMaxOS Q3 2019 release introduces PowerMax support for both storage class memory (SCM) and flash drives. The support for different drive technologies introduces tiered storage called automated data placement (ADP). The functionality for ADP is included as part of the service level options and uses machine learning to place data on the appropriate drive. To improve the response times managed by service levels, ADP complements the management of I/O by placing the most active data on the faster drive technology. ADP movements occur as either a promotion of data or a demotion of data, which are relative to the selected service level as well as the location of data. Promotions are movements of data sent to the SCM drives and demotions are movements out of SCM. PowerMaxOS sets a service level bias in order to apply the priority of promotions and demotions.

Promotion:

- Diamond:
 - Highest promotion priority
 - During optimal utilization, PowerMaxOS attempts to put all data labeled Diamond on SCM drives
- Platinum, Gold, Optimized: All data has equal priority
- Silver, Bronze: Excluded from promotion

Demotion:

- Silver, Bronze: Highest demotion priority
- Diamond:
 - Data is demoted when there is a need for more active Diamond data to be promoted
 - Can also be demoted if the SCM usage exceeds pool reserved capacity
- Platinum, Gold, Optimized:
 - Equal demotion priority
 - Demotions occur when there is a need to create available space in SCM for higher priority data or more active data with the same priority.

1.5 Setting service levels

Service levels can be applied to a storage group when either creating a new storage group or by modifying an existing storage group. Users also have the ability to change service levels at any time to apply the desired response time performance expectation.

Note: When creating a cascading storage group, service levels are applied to the child storage group and the parent storage group remains set to none.

Applying and modifying service levels are done through PowerMax management tools:

- Unisphere™ for PowerMax
- SymCLI
- RestAPI

1.5.1 Creating a storage group with a service level

Setting service levels in Unisphere for PowerMax can be accomplished through either the Storage Group dashboard or the Service Level dashboard.

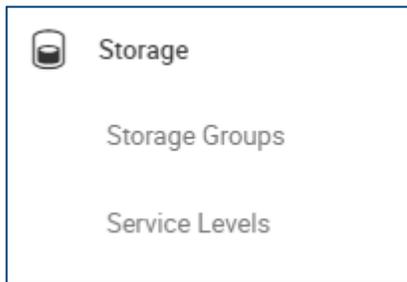


Figure 1 Storage Groups and Service Levels dashboard selection

The following example creates a storage group and sets a service level using the Service Levels dashboard. In the Service Levels dashboard, the available service levels are displayed along with information such as the target response time. From this view, select the desired service level and click **Provision** in the top-left corner. This opens the provision wizard.

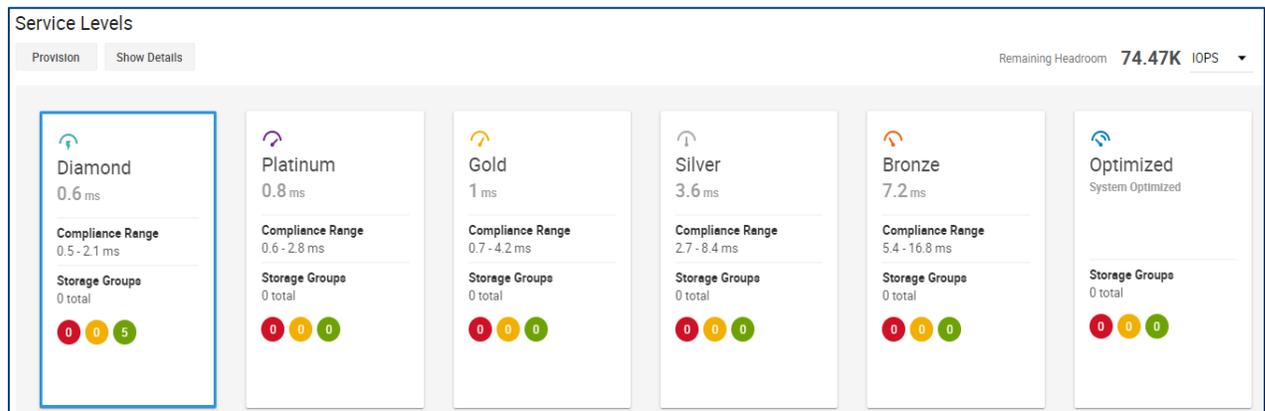


Figure 2 Service level dashboard

In the provisioning wizard, select the desired service level from the drop-down menu and enter a name for the storage groups. Select the storage resource pool, number of volumes, and capacity to create devices.

Figure 3 Provision Storage wizard

At this point, either complete provisioning storage to a host, or create the storage group and provision to a host at some other time.

1.5.2 Modifying a storage group

Use the Storage Groups dashboard to modify service levels for a storage group. The dashboard contains a list of existing storage groups and their options.

	Name ▲	Compliance	SRP	Service Level	Capacity (GB)
<input type="checkbox"/>	Prod_DB				
<input checked="" type="checkbox"/>	Prod_DB	✓	SRP_1	Diamond	0.00

Figure 4 Storage Groups dashboard

Note: You can choose to create a storage group from this dashboard. Click **Create** to open the provisioning wizard described in section 1.5.1, Creating a storage group with a service level.

Once the desired storage group is selected, click **Modify** to change the service level. The Modify Storage Group wizard opens. From the drop-down menu, select a new service level.

Modify Storage Group | Prod_DB

1 Storage Group

Storage Group Name * Prod_DB Storage Resource Pool SRP_1

Optimized

Diamond

Platinum

Gold

Silver

Bronze

None

Volume Capacity 5 TB

500 GB Total Service Levels 1

Enable Compression Run Suitability Check

Figure 5 Modifying a storage group

1.6 Alerting and compliance

Alerts are offered with service levels, allowing users to be notified when a specific service level has reached its respective upper response time limit. Alerting is available on all service levels with the exception of Optimized, which does not have a target response time or upper response time limit.

Alerts are only available through Unisphere for PowerMax and are enabled by default. As with other alerts in Unisphere, users can set up email notifications for these alerts.

When a storage group with a service level sets an alert, it is viewable from either the Unisphere dashboard > **Alerts** tab or from the Storage Groups dashboard > **Compliance** column.

The following example demonstrates how to view a compliance alert for service levels and displays the details within the alert. An alert is shown in the Storage Groups dashboard as indicated by the yellow triangle under Compliance (see Figure 6).

Storage Groups						
<input type="button" value="Create"/> <input type="button" value="Modify"/> <input type="button" value="Provision"/> <input type="button" value="Protect"/> <input type="button" value="Set Host I/O Limits"/> ⋮						
<input type="checkbox"/>	Name	Compliance ▲	SRP	Service Level	Capacity (GB)	Emulation
<input type="checkbox"/>	Win146_WL_	<input type="text"/>				
<input type="checkbox"/>	WIN146_WL	—	NONE	NONE	1,610.01	FBA
<input checked="" type="checkbox"/>	WIN146_WL_D		SRP_1	Diamond	230.00	FBA

Figure 6 Alert in the Storage Groups dashboard

In Figure 7, users can further examine the storage group to gather more information on the alert. Navigate to the **Compliance** tab under the **Details** tab of the storage group.

Storage Groups > WIN146_WL_D

< Details Compliance Volumes Performance

Properties

Symmetrix ID	000197800130
Compliance	
Service Level	Diamond

Figure 7 Storage Groups details

Under the Compliance tab, there are details of the timeframe and response time management. Figure 8 shows the response time of the storage group exceeds the upper response time limit briefly on Thursday and then again between Thursday and Friday.

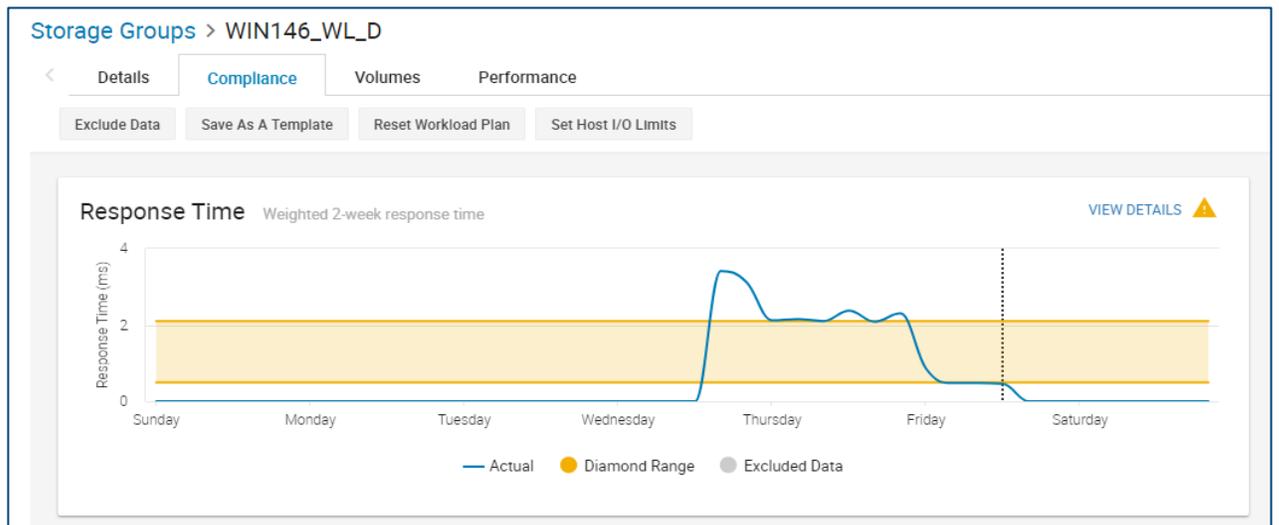


Figure 8 Compliance tab

Near the two timeframe buckets, click **View Details** to see more information on the response time (Figure 9). The first four-hour interval is the average response time within those four hours, and the second interval is the average of two weeks.

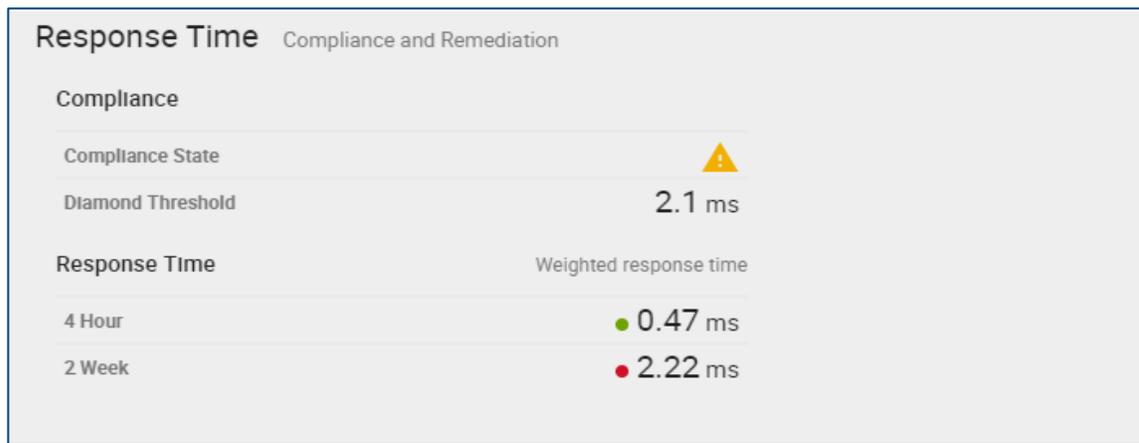


Figure 9 Response Time compliance

2 Service level functionality

Service levels for PowerMaxOS are designed to put relative importance on application storage groups. This function allows users to manage applications based on predictability and consistency in response time. With the ability to manage the performance of the array, users can determine whether to utilize as much of the resources the array allows or as little of the resources regardless of the system capabilities. With this ability to manage performance, users assign importance on applications based on specific needs.

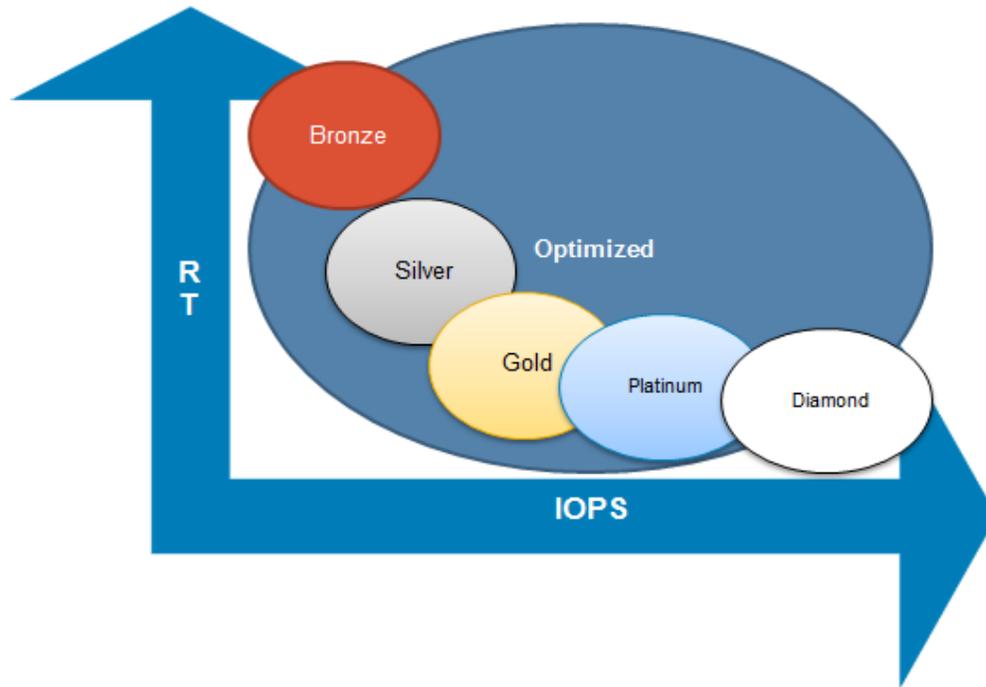


Figure 10 Relative response time compared to IOPS

2.1 Priority applications

Service levels allow users to insulate specific storage groups from the performance impact of other noisy-neighbor applications. The user can assign critical applications to higher service levels such as Diamond, Platinum, or Gold which allow for these storage groups to utilize all available resources at all times. These critical applications are not managed unless the system exhibits resource constraints causing the applications to fail to maintain desired performance levels.

In Figure 11, the lower-priority storage groups begin to impede on the response-time boundary of a storage group with a higher service level. The lower-priority storage group is then managed by PowerMaxOS. The management of lower-priority storage groups subsides once the higher-priority storage group is within its respective target service level.

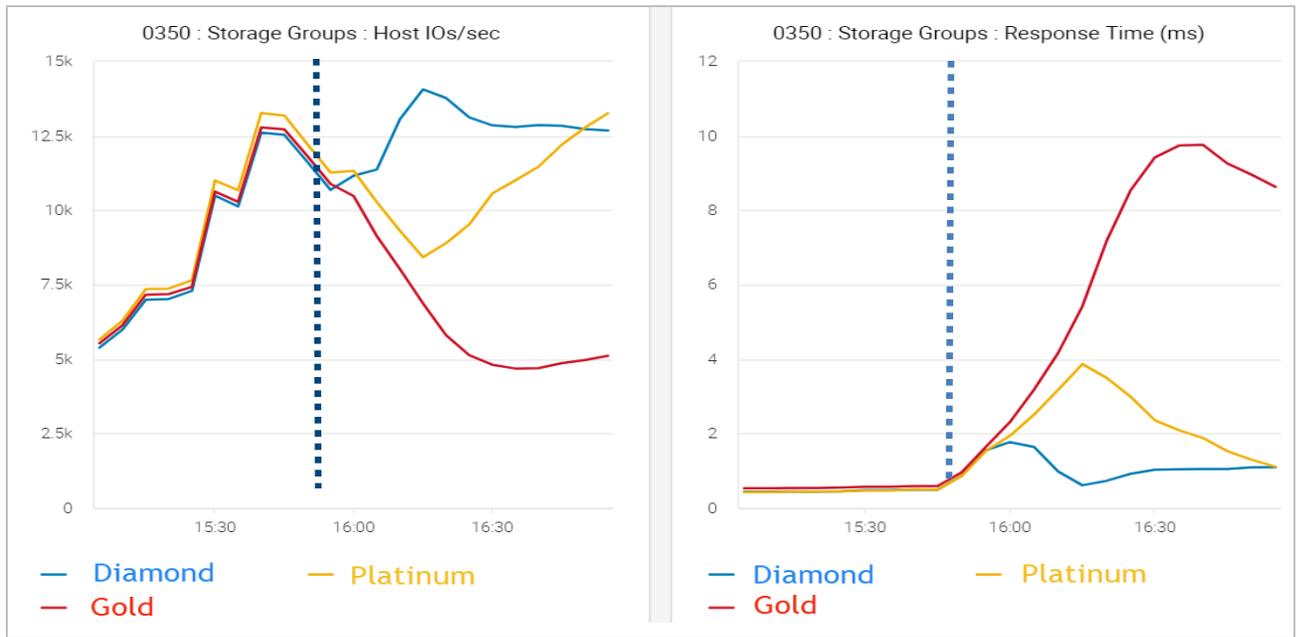


Figure 11 Example of noisy neighbor

2.2 Service providers

Service levels can be used to implement a chargeback model where service providers assign a higher cost to higher-performing storage. Silver and Bronze service levels give the provider the ability to impose differentiated performance levels regardless of system resource constraints or the ability of higher-priority service levels to reach target performance levels. Users can also set service levels while applications are running, which allows providers to adjust the response times of storage groups based on revolving price points.

Figure 12 shows the lower-tier Bronze and Silver storage groups being managed by PowerMaxOS, while all other storage groups remain at a lower response time.

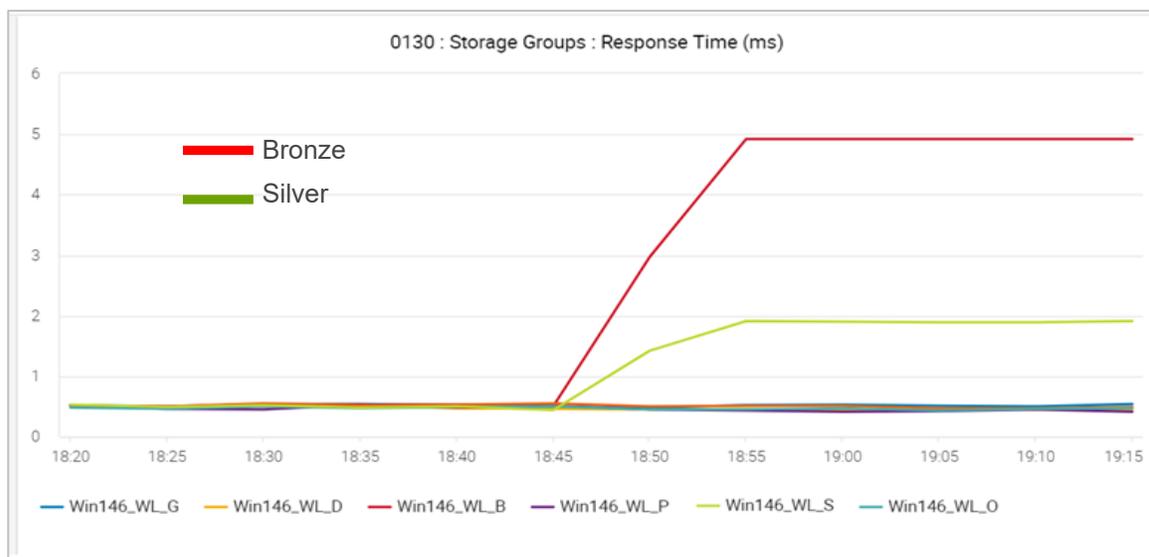


Figure 12 Service provider response time management

3 Interoperability

Service levels operate seamlessly alongside all PowerMaxOS data services and features. Service levels are not supported in a mainframe environment.

3.1 Local and remote replication

Service levels are interoperable with replication features including TimeFinder™ SnapVX and SRDF™.

3.1.1 Service levels and SRDF

When creating remote replication of a storage group with SRDF, the user can initially set the desired service level of the remote storage group through the SRDF wizard in the local array. Once SRDF protection of a storage group is set, the user can modify either the local or remote storage group service level with any of the PowerMaxOS management tools.

3.1.2 Service levels and TimeFinder SnapVX

When linking a new target storage group to a snapshot using TimeFinder SnapVX, the target storage group inherits the same service level of the source storage group. If linking an existing storage group to a snapshot, the existing storage group service level remains the same as it was when created. The service level of the target storage group can be modified at any time to adjust a desired service level depending on the user requirements.

3.2 Service levels and host I/O limits

Service levels ensure storage groups have an expectation of performance in terms of response time while host I/O limits provide a function to limit the amount of front-end port performance. Host I/O limits do this by allowing users to set a maximum front-end throughput on either IOPS, MB/s, or a combination of both. When a host I/O limit is applied to a storage group that has a service level set, the storage group will still be managed by any higher-priority storage group.

Both service levels and host I/O limits are set per storage group and can work together for more predictable and consistent performance. Host I/O limits can be set on a storage group that has a specified service level to manage the front-end throughput if the desired response-time performance of the storage group is continually being exceeded or if the storage group is being impeded upon by other storage groups. Host I/O limits are not a method to maintain response time but allow users to control how much data is being driven to the array, which help service levels maintain consistent response-time performance.

Figure 13 illustrates service levels as they apply to the response-time upper and lower limits and shows how host I/O limits relate to setting the maximum allowable throughput IOPS.

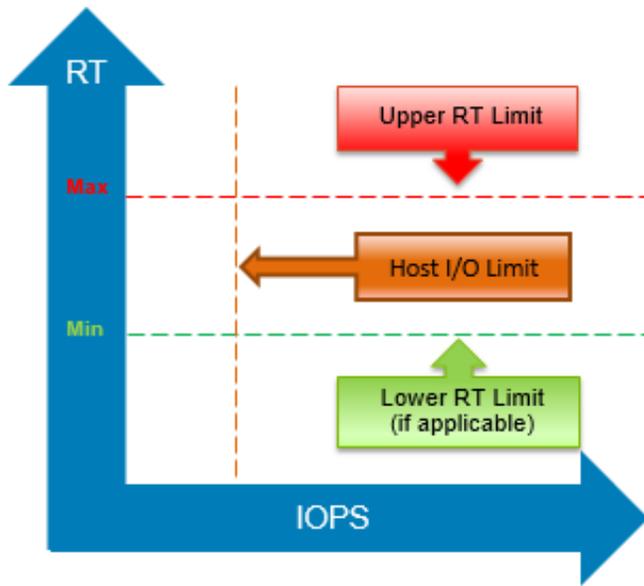


Figure 13 Service level management relative to host I/O limits

4 Summary

Service levels for PowerMaxOS give organizations the ability to prioritize the importance of business-critical applications while consolidating lower-priority applications. PowerMaxOS ensures that higher-priority applications are not impacted by any other application with a lower service level. Users can set priority based on a range service levels that each provide varying performance expectations and relative target response times. This allows users to set their desired expectations while providing predictable and consistent performance.

A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.