


DELL EMC ISILON PERFORMANCE ANALYSIS - T-ANALYZER



Nidhi Shree N

Specialist 2, Inside Product
Dell Technologies
Nidhishreen.shreen@dell.com

Sai Ganesh

Specialist 2, Inside Product
Dell Technologies
Saiganesh27.is@gmail.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at www.dell.com/certification](http://www.dell.com/certification)

Table of Contents

1. Introduction	4
2. How is it different from other tools?	4
3. Target Audience	4
4. Post-Sales Performance Analysis	5
• Field Scenarios	5
5. Performance Indicators Review	5
a) Collecting Statistics	6
b) T-Analyzer Analysis	7
c) Review Report	7
6. Understanding the Performance Metrics	8
7. Future Enhancement	13
8. Summary	13
9. References	14

Abstract

Every year the volume of data and the speed at which it is created grows exponentially. Due to this rapid growth, storing data has become an increasingly difficult challenge. Larger volumes of data which are now much more important to an organization's operations require storage infrastructure that has high performance, is available and reliable. But how do you know if your storage is performing as it should? What performance can you realistically expect from your current infrastructure? Will a new storage solution be able to support your existing workloads? To answer these questions - it is important to gather storage metrics and perform a level of analysis on the storage infrastructure.

Dell EMC Isilon is a scale-out network-attached storage platform for high-volume storage, backup, and archiving of unstructured data. It provides a cluster-based storage array based on industry-standard hardware and is scalable to sixty-eight petabytes in a single filesystem using its FreeBSD-derived OneFS file system. An Isilon clustered storage system is composed of three or more nodes. Each node is a server integrated with proprietary operating system software called OneFS (based on FreeBSD[4]), which unifies a cluster of nodes into a single shared resource. Since the system is highly scalable it becomes necessary for the admins to understand the cluster's performance statistics.

T-Analyzer is one such tool that helps visualize the Dell EMC Isilon's cluster performance metrics with the help of the logs generated using the T-Analyzer (TA) scripts.

Introduction

T-Analyzer is a program written in Java by **Anders Bray** and **Denzel Holmes** to help with analyzing performance on an Isilon cluster by automating the tell analytics process and providing an interactive interface for reviewing the data.

The statistical data is collected first - either manually, or by deploying from within the program; then the tool parses the data and produces an interactive user interface. The current version is still in beta and thus has basic formatting and features that will help in most Tell Analytic cases.

Though this type of analysis is already being done, the existing procedure involves running long commands to parse the data as well as manually producing graphs for use in displaying the results that can take several days/weeks to complete.

A customer-facing conversation is needed to obtain approval to run the script on their cluster for the metric gathering. Expectations are that we will gather the log details and compile a report for review and analysis.

How T-Analyzer differs from other tools?

T-Analyzer differs from other tools such as isiperf and InsightIQ because it uses weighted averages to help determine where a bottleneck is, which is not easily found over several samples with live stats.

The tool also operates on the principle of taking a mathematical limit, because the more data used to measure the average, the more irrefutable the findings.

In addition to this, data insights from InsightIQ/DataIQ and other tools can be combined for more powerful storytelling, analysis, and introspection.

Audience

This article is ideal for those who specialize in troubleshooting Isilon performance problems or want to correlate several data points to determine the possible bottleneck.

While other – (Technical Support Engineers) (TSEs) may find the information helpful, if not taken with the correct context some of the data could lead to false positives. Hence, you should always consult with an L2 if using this tool.

1. Post-Sales Performance Analysis

- **Field Scenarios**

Below are some of the field scenarios where T-Analyzer can help.

- When a sales campaign needs a performance review leading to a customer facing discussion that will focus on their performance objectives, expectations and addressing concerns.
- When a performance issue arises, what steps can the SE take to proactively accelerate the troubleshooting process and collaborate with support.
- Proactively perform a performance assessment and capture performance data points and baselines to share with the customer to help with their future planning.
- Customer may be adding a new workload; gather before-after metrics to make recommendations.
- A post-incident assessment can provide immediate recuperative benefits.
- If there is a perceived performance concern, how can the SE be proactive and assist?

Before proceeding with Tell Analytics ensure the following has been considered

- Determine which Dell and Customer stakeholders will be involved. Resources will vary based on whether this is being driven from a sales or support standpoint.
- Is there a perceived performance concern or is this a routine check-up?
- Determine if an SR should be opened or if one is open, can this performance analysis help?
- Is the issue repeatable and reproduceable? (If so, captures can be taken during the perceived issue)
- Are any of these items relevant to the performance review? (Top Talkers by IP and hostname, Client Node Balance, Protocol Ops by Class, Protocol by Operations, Protocol ops by Class Table, Disk Heat by Class, Heat statistics by path [Shows deadlock, contended, blocked by path])

2. Performance Indicators Review

Performance Indicators Review (PIR) is a deliverable by Presales Solution Services team which adopts the T-Analyzer tool.

PIR flow is classified into three major categories:



It is important to have a clear understanding of the workflow used in this cluster before starting, such as the protocols used, file size, peak workload times, etc.

Reminder: A customer-facing conversation is needed to obtain approval to run the script on their cluster for the metric gathering. Expectations are that we will gather the log details and compile a report for review and analysis.

Collecting Statistics

Review the objectives and TA script with the customer. Schedule time to run the script in the environment on the cluster and gather Logs “

In any case, the aim is to gather data, if possible, for the script during the time they have problems or workload to be evaluated while minimizing the amount of time it runs at other times.

Using the information collected on the first two items identify the ideal time the script can be collected. (The other data requested on those items will be essential when the analysis is performed)

Note: The TA Script is available on the confluence page referenced in this article’s Reference section.

Steps to run TA Script

1. Copy the ta#.sh onto the Isilon cluster in the **/ifs/data/Isilon Support** folder.
2. Prior to launching ensure that no files exist within **/ifs/data/Isilon Support/TA** as this Script will overwrite them.
 - We recommended to have a copy of the data in **/ifs/data/Isilon Support/TA** if the folder already exists to avoid overwriting.
3. Run the Script for one iteration and determine the time taken for a single iteration using the below:

/bin/bash ta#.\$sh 1

4. Divide the total time you plan to collect data by the time per iteration to identify how many iterations to start the Script for. (We recommend always running with at least five iterations)

/bin/bash ta#.\$sh <number of iterations to run>

T-Analyzer Analysis:

“Run the gathered logs through the T-analyzer tool. Obtain immediate results on the User Interface.”

The prerequisite for gathering graphs via the T-Analyzer tool is to install JRE on windows and download the T-Analyzer tool (from the confluence web page).

Analysis Method

Open T-Analyzer.jar → Browse the Required Log Folder → Run

Further, for each iteration, run the collected logs through the T-Analyzer tool and obtain the Graphical representations of Performance Metrics as output. Several iterations may be used to better define the bottleneck at different intervals.

An image of every graph will be saved automatically using the default view to a folder within your home drive called TA-Reports.

For windows this is: **C:\users\%username%\TA-Reports**

For Unix/Linux this is: **~/TA-Reports**

a) Report Review:

“Process and interpolate Graphs with the desired PIR Template. Analyze the data and review internally as needed.

Schedule time to review findings with the customer for feedback, recommendations, and next steps to discuss the outcomes with customers.

If an SR is in play, collaborate with support and let them drive the troubleshooting plan and engagement.

The following Performance Metrics are available in the report:

1. Top Talkers by IP and Hostname

2. Client Node Balance
3. Protocol Ops by Class
4. Protocol by Operations
5. Protocol ops by Class Table
6. Disk Heat by Class
7. Heat statistics by path (Shows deadlock, contended, blocked by path)

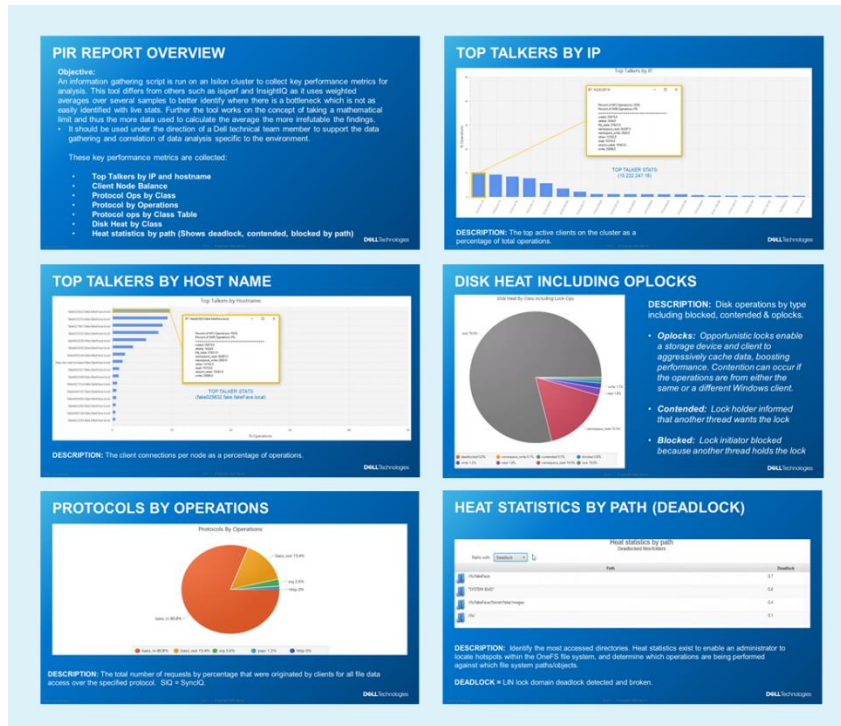


Figure: Sample PIR Report created by PreSales Solution Services team

PIR is available along with PSS - generated State of the Cluster (SOTC) which can be requested via PSS Dash. (Link in the reference section)

Understanding the Performance Metrics

- **Top Talkers by IP**

Describes the percentage of total operations from 0 to 100 of the top active clients on the cluster indicated by their IP address. The percentage of operations include NFS & SMB. It also indicates the active clients that are consuming the most system resources.

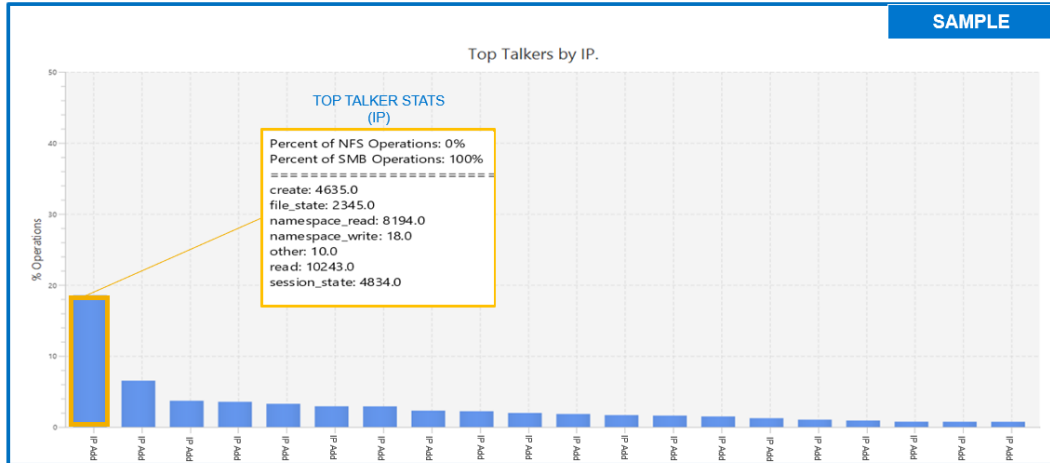


Image Source Sample PIR Report created by PreSales Solution Services team

- **Top Talkers by Hostname**

Describes the percentage of total operations of all the top active clients on the cluster indicated by their Host name. Like that of Top Talkers by IP, this indicates the active clients that are consuming the most system resources by Host name.

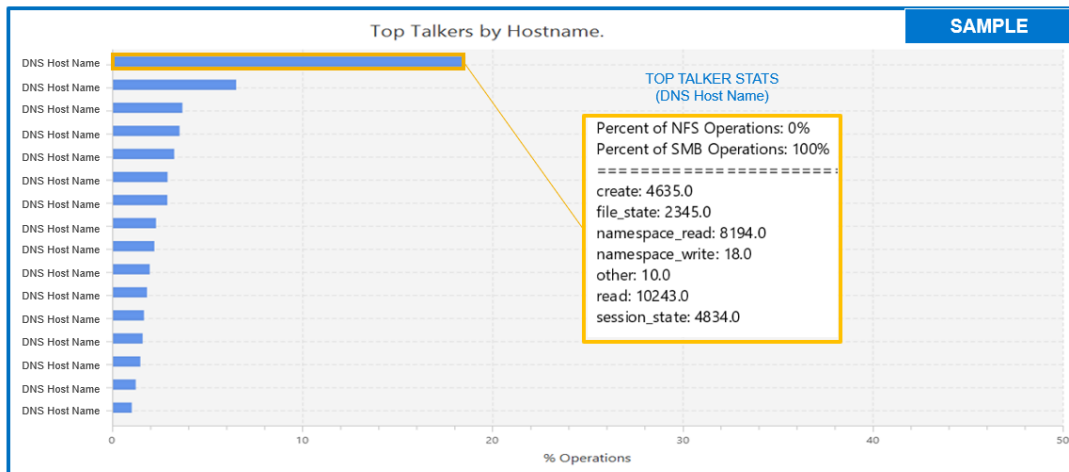


Image Source: Sample PIR Report created by PreSales Solution Services team

- **Client Node Balance**

Client node balance represents the number of nodes currently connected to the cluster and the total number of operations indicated on each node by percentage.

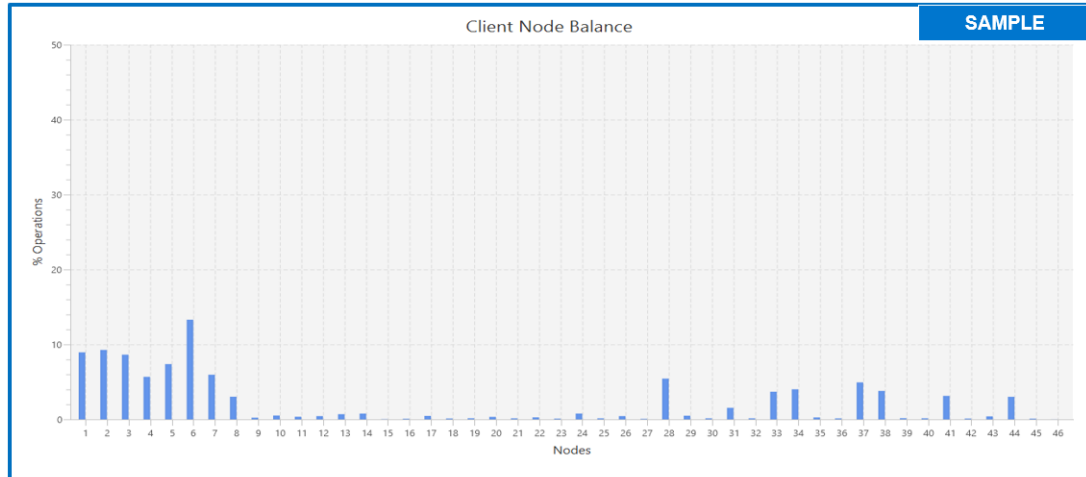


Image Source: Sample PIR Report created by PreSales Solution Services team

- **Protocol Ops by Class**

The Pie chart displays the total number of times the Clients request access to the file data over a specified protocol by class as percentage.

Example: Percentage of read, write, etc.

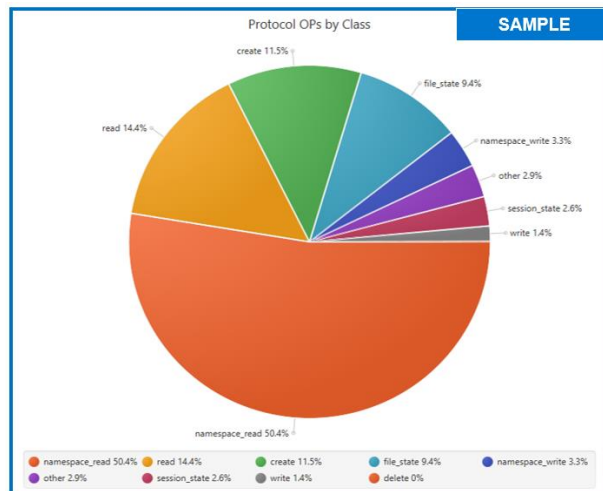


Image Source: Sample PIR Report created by PreSales Solution Services team

- **Protocols by Operations**

Protocols by Operations indicates the percentage of total number of requests by the clients to access the file data over a specified protocol.

Siq indicates = SyncIQ. **IRP** = The filter manager provides a plug-in framework for pre - and post-input/output request packet (IRP). The IRP provides the mechanism to encode a protocol request handled by Likewise Input Output (LWIO) and encodes the request handled by the file system drivers.

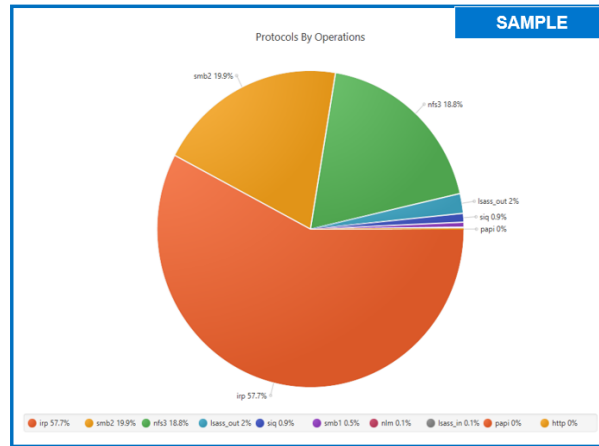


Image Source: Sample PIR Report created by PreSales Solution Services team

- **Disk heat by Class**

Indicates the percentage of total number of disk operations by type.

When there is an increase in namespace read values that means there is a significant metadata work is being performed, which may result in an issue.

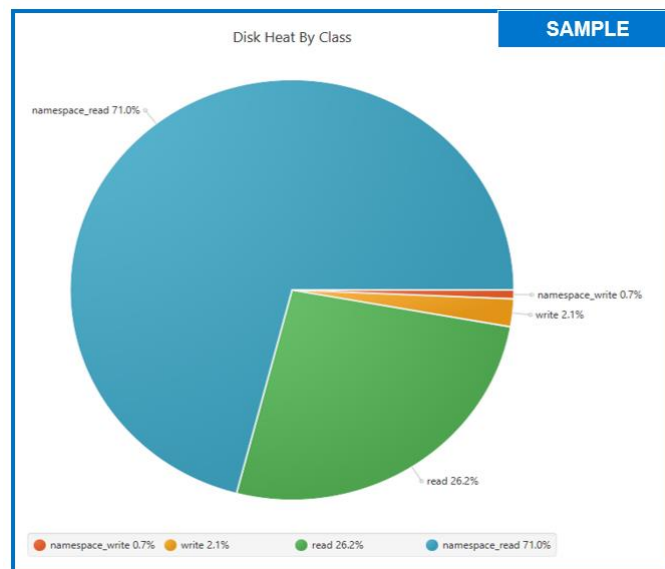


Image Source: Sample PIR Report created by PreSales Solution Services team

- **Disk heat including Oplocks**

Indicates the disk operations by type which includes oplocks, contended & Blocked.

Opportunistic lock (OpLock): A form of file locking used to facilitate caching, access control, and improve performance. It enables the storage device and client to aggressively perform caching.

Contended Locks: There is at least one more T2 thread that wants to access the same one that the previous T1 thread has already locked.

Blocked: Occurs when two processes require simultaneous access to the same piece of data so that the data is locked by one process and the other must wait for the other to complete and unlock the lock. The blocked process resumes service as soon as the first process is done.

Deadlock: In the operating system, a deadlock occurs when two or more processes need a resource retained by the other process to complete their execution.

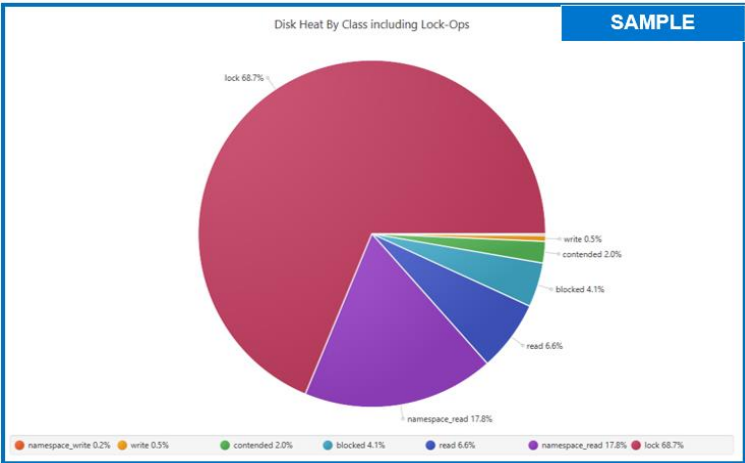


Image Source: Sample PIR Report created by PreSales Solution Services team

- **Top 10 Busiest Path**

As the name indicates it gives the details of the all the busiest paths of the cluster out of which we will identify the most accessed directories by percentage of operations.

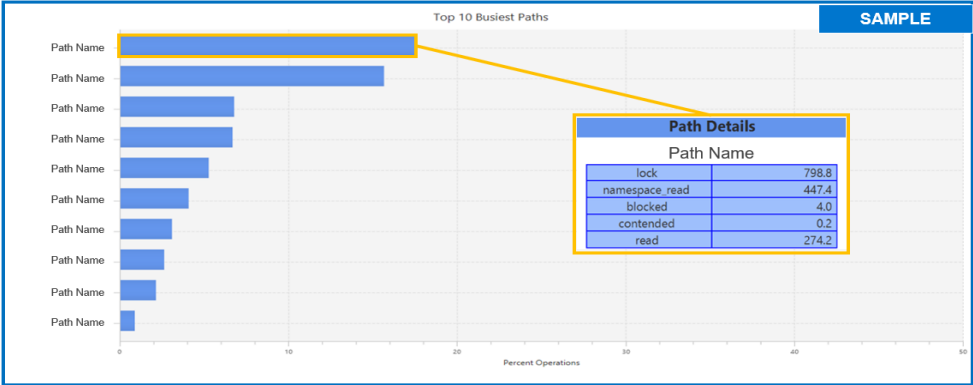


Image Source: Sample PIR Report created by PreSales Solution Services team

Future Enhancements

The aim is to create a web-based statistic viewer that contains all the T-Analyzer’s report and is available directly in Elvis as performance viewer.

elvis.emc.com/perf_viewer – Application to investigate cluster performance issues and perform proactive workflow analysis.

During a performance issue, perf_view.sh can be run on a live cluster to collect performance data. Performance Viewer processes perf_view data to provide a webpage with graphs, tables, and automated performance checks.

It can be used to investigate high protocol response times and overall cluster workflow.

Summary

For an organization to prevail in the advanced digital future, they may adopt a variety of measures relying upon the business verticals and their kind of customers.

A better understanding of how OneFS performs will lead to a better understanding of workloads, how they communicate with an Isilon system, and how they can be made more efficient.

Proactively perform a performance assessment like PIR using the T-Analyzer approach and collect performance data points as - baselines to share with the customer to assist with their future planning.

The final objective is to monitor the performance of a cluster effectively and granularly.

References

- **ISG-Storage Confluence**
<https://confluence.cec.lab.emc.com/display/ISILON/T-Analyzer+-+Automating+Statistic+Analysis>
- **PSS Dash**
<https://dash-dellemc.tfaforms.net/67>
- **UDS Specialty - Presales Solution Services**
<https://inside.dell.com/docs/DOC-367289>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” DELL TECHNOLOGIES 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 Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.