

Dell EMC Storage Manager Scalability Solutions

Dell Storage Engineering
April 2017

Revisions

Date	Description
April 2017	Initial release

Acknowledgements

Author: Darin Schmitz

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 © 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [4/17/2017] [Best Practices] [3395-BP-SC]

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	4
1 Deployment guidance.....	5
1.1 Configuration tips.....	5
1.2 Storage Center arrays per Data Collector	5
1.3 Management domains	6
1.4 Geographic distance.....	7
1.5 Failure domains	7
2 Conclusion.....	8
A Technical support and resources	8

Executive summary

Dell EMC™ Storage Manager (DSM) contains features on the Data Collector server that are a crucial piece to managing the SAN infrastructure. With that said, it is important to keep the Data Collector running efficiently so that transactions complete in a timely manner and prevent administrative delays. This guide provides high-level techniques and strategies for deploying the DSM Data Collector servers in a way that promotes acceptable transactions times to execute storage management tasks.

1 Deployment guidance

With the release of DSM 2016 R3, the release notes communicate that to maintain acceptable performance, a Data Collector should manage no more than ten Dell EMC SC Series arrays. This recommendation was due to detailed data analysis indicating that limiting a Data Collector to ten arrays should offer a responsive client experience for most customers.

However, ensuring adequate performance of the Data Collector is more complicated than simply limiting the number of arrays it manages. While the recommendation of ten arrays is a sound guideline for most customers, it does not reveal the entire picture. The Data Collector can actually scale beyond ten arrays, but exactly how many is dependent on a large number of factors. The most prominent factors being network latency, number of volumes per array, and the number of arrays managed.

The following sections will provide additional guidance and tips for maintaining acceptable performance of DSM Data Collectors.

1.1 Configuration tips

Implementing the following suggestions will help ensure improved performance of the Data Collector.

- When running the Data Collector within a virtual machine, be sure to manage it as a tier 1 application. This means taking snapshots, and using hypervisor resource regulation features to ensure that the virtual machine does not become resource starved.
- Be sure to dedicate enough RAM and CPU resources to the machine per the advice of the most current *Dell Storage Manager Installation Guide*.
- If available, store the Data Collector and the corresponding SQL database on high performance storage. Reducing disk transaction times helps to improve operations within the Data Collector and database.
- Be mindful of SQL server RAM consumption. In configurations where SQL server is installed on the same machine as the Data Collector service, SQL server can consume the entire free RAM in the host, subsequently starving the Data Collector service.

1.2 Storage Center arrays per Data Collector

Engineering testing has revealed that the DSM client achieves the best performance when the storage arrays and the Data Collector reside within the same data center. The sub-millisecond network latencies allow for timely statistics collection and task transaction times. When trying to decide how many arrays each Data Collector should manage for intra-data-center scenarios, administrators can make rough estimates by means of volumes per array (see Figure 1).

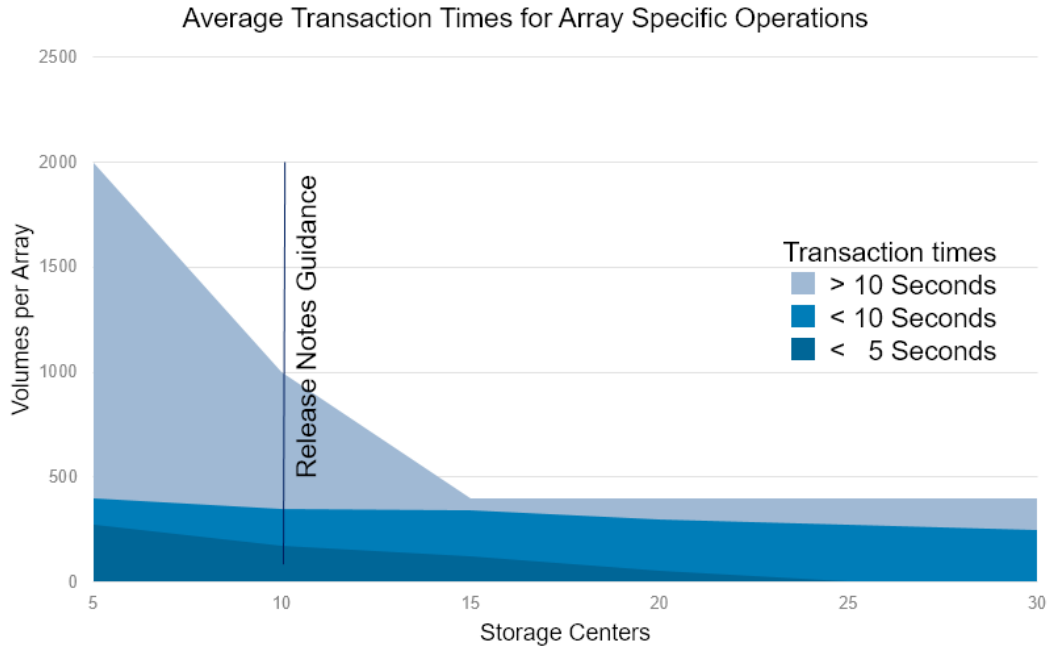


Figure 1 Graph showing approximate transaction times with various configurations

While the development teams work to improve scalability with each software release, the current data indicates that with a high quantity of volumes per array, the more likely that transaction times exceed the transaction time testing goal of <10 seconds. The darkest blue areas of the graph are where admins can expect the most responsiveness from the client. In contrast, it is best to avoid configurations outside of the blue < 10 seconds area, as the extended response times can become unpredictable and operations could potentially time-out.

Effectively, what this means is that while a single Data Collector under ideal conditions may be able to actually manage dozens of arrays with low volume counts, a higher volume count translates to each data collector managing fewer arrays.

1.3 Management domains

When approaching the scalability limits of a single Data Collector, the load can be spread out by installing additional primary Data Collectors to separate the storage arrays into functional management domains. For example, one Data Collector to manage arrays in building A, while a second Data Collector to manage arrays in building B. In larger data centers, the management domains can be broken down to racks, rows, floors, replication groups, or any other logical grouping.

Note: During periods of slower than normal client performance, remember that if the array is running SCOS 7.0 or later, the DSM client can connect directly to the array to perform the array specific management operations. In addition, if the array is running SCOS 6.7 or earlier, it has a built in web interface for management.

1.4 Geographic distance

When the Data Collector is separated from the arrays by distance, the added network latency can possibly magnify the transaction times, reducing the total number of volumes and arrays that can be managed by a single Data Collector. While certain client transactions can tolerate up to 75 ms of network latency before crossing the 10-second threshold testing goal, other types of transactions can exceed the goal with just 30 ms of latency separating them. While individual results may vary, benchmarking shows that the DSM client operations routinely finish under the 10-second goal where network latencies are less than 50 ms. Of course, every environment is different, so administrators must adjust the number of arrays each Data Collector manages to suit their specific environment.

1.5 Failure domains

When supporting multi-data-center replicated environments, if more than one Data Collector is required, Dell recommends separating arrays into failure domains. In other words, arrays that use each other for disaster recovery, fail-over, or business continuance, must be added to the same Data Collector. In hub and spoke replication architectures, this means dedicating at least one pair of primary and secondary Data Collectors per spoke. Keep in mind that based on latency between sites, multiple Data Collector pairs may be required between each site for larger environments.

2 Conclusion

For more information about the various replication scenarios specific to each application, refer to the best practices guides on Dell TechCenter at [Dell.com/StorageResources](https://www.dell.com/storage/resources).

A Technical support and resources

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

[Dell TechCenter](https://www.dell.com/techcenter) is an online technical community where IT professionals have access to numerous resources for Dell software, hardware and services.

[Storage Solutions Technical Documents](https://www.dell.com/techcenter/storage-solutions-technical-documents) on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.