A MICRON
TECHNICAL
MARKETING
BRIEF

Getting the Most Out of Your Micron M500DC SSD + Microsoft SQL Server 2014 Investment

Host RAID Optimizations for OLTP

Overview

When investing in enterprise-grade solid state drives (SSDs), like Micron's M500DC SATA SSDs, for use in Microsoft SQL Server 2014 OLTP deployments, it is critical to understand the benefits—and potential tradeoffs—of different RAID configuration options in order to get the most value from your investment.

Microsoft's SQL Server 2014 is one of the most widely deployed relational database management systems (RDBMS) in enterprise data centers today. It is primarily used to manage transaction-based applications—ranging from order entry and fulfillment to real-time data acquisition, management, and analysis—as well as some commercial processes. In the vast majority of Microsoft SQL Server 2014 online transaction processing (OLTP) deployments, application responsiveness (latency) and overall performance—measured in transactions per minute (TPM) and new orders per minute (NOPM)—are the driving factors of success.

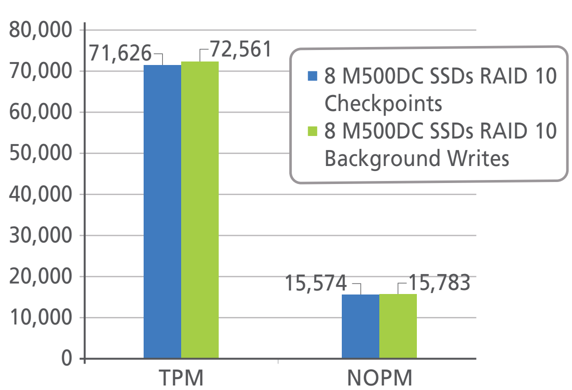

Figure 1: Performance of 8 M500DC SSDs With Checkpoints vs. Background Writes

Typical OLTP Application Requirements

- » **Immediate access** to mission-critical data
- » **Uninterrupted service** to avoid costly delays in access to data during order entry and fulfillment and data acquisition, analysis, and management.
- » **Overall consistent performance** for smoother workflows and more predictable process execution

This technical brief looks at several Microsoft SQL Server 2014 configuration options and identifies those that can take full advantage of Micron's M500DC to improve the overall responsiveness of SQL 2014-based OLTP applications and those that enable simpler platform design.

Checkpoints or Background Writes?

Following an unexpected failure, Microsoft SQL Server 2014 provides known good recovery starting points (database checkpoints) in the log as indicators of where to start applying changes during the recovery process. Checkpoints can cause periodic burst write traffic to storage devices, with corresponding periodic increases in latency. Alternatively, background writes via indirect checkpoints (of dirty buffer pages) can be set to trigger more frequently, which results in more I/O to the storage but limits I/O spikes, yielding more consistent latency.

As shown in Figure 1, a 49-user workload using eight M500DC SSDs enabled with checkpoints generated approximately the same performance (TPM and NOPM) as the same configuration

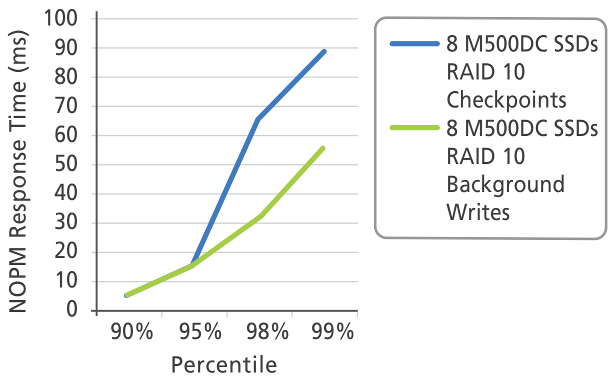

Figure 2: Latency of 8 M500DC SSDs With Checkpoints vs. Background Writes

enabled with background writes. However, enabling background writes instead of checkpoints showed consistently lower latency (same workload, 99th percentile), as shown in Figure 2.

RAID and Database Layout

For SQL 2014 platforms using conventional RAID implementations, the RAID level can have a dramatic impact on application responsiveness and data protection. When choosing a RAID level, the right balance between usable capacity, data protection, and performance must be determined.

For mission-critical applications, unprotected RAID 0 should not be used; nested RAID levels (such as, RAID 50 or RAID 10) are most common. It can also be beneficial to place portions of the database on separate RAID sets (for example, the log files on one RAID set and the database files and tempdb on another).

To determine the optimal RAID configuration, tests were conducted with eight M500DC drives in three different configurations:

1. Six M500DC SSDs at RAID 10 with database and tempdb on a logical disk 1; two M500DC SSDs at RAID 1 with logs on a logical disk
2. Eight M500DC SSDs at RAID 10 with database, logs, and tempdb on a single logical disk
3. Eight M500DC SSDs at RAID 50 with database, logs and tempdb on a single logical disk

As shown in Figure 3, TPM and NOPM performance was highest when the log files were placed on a dedicated RAID 1 set and the database and tempdb were placed on a separate dedicated RAID 10 set. When the logs were placed on the same RAID set as the database and tempdb, both TPM and NOPM decreased slightly. The lowest TPM and NOPM were seen on the RAID 50 configuration.

Note: Due to the inherent capacity loss with this configuration, it may be preferable to place the logs on the same RAID set as the database and tempdb to save costs. In this case, logs can be protected using periodic backups, database replication, or mirroring methods.

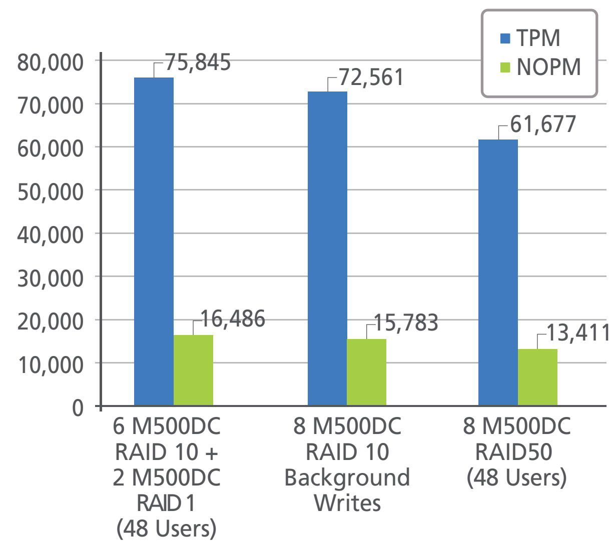

Figure 3: Effect of RAID Configuration on M500DC Performance

Benefits of Dedicated RAID Sets

» **Write Latency:** Because transaction WRITE latency is directly dependent on log WRITE latency, placing the log files on a dedicated RAID 1 set improves overall write responsiveness. The corresponding 99th percentile of latency for new order transactions (write transactions) decreases when the log files are placed on a dedicated RAID 1 set, and because this configuration dedicates two M500DC SSDs for logs, the remaining number of M500DC SSDs available within the given set is six. This reduction in RAID 10 array drive count causes the data disk to incur slightly higher latency because it has to absorb all background writes for dirty pages spread across only six M500DCs, as shown in Figure 4.

» **Log Functions:** Dedicating a RAID set to log functions is a general SQL best practice, providing better protection against failures. For example, when the database RAID set fails, the log disk data can still be used to recover the database to the latest point in time by replaying logs on a backup set. If the database is protected via AlwaysOn Availability Group using synchronous commit mode, then a single RAID set hosting database and logs can survive RAID set failures. The replica database on another server instance will have the latest log information to replay.

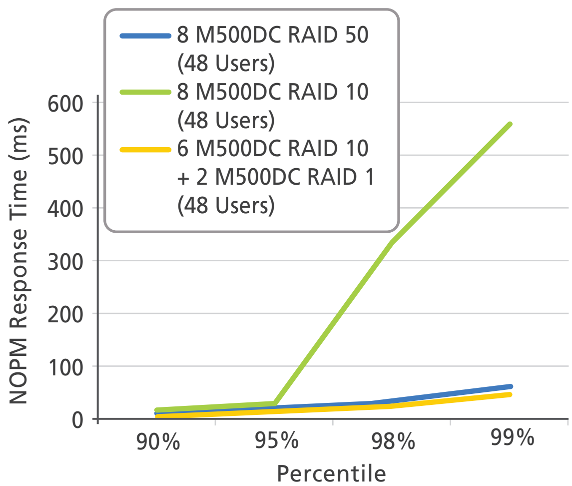

Figure 4: Times With Different RAID Configurations

Quick Summary

Background Writes Instead of Checkpoints: Configure the platform to use background writes, which can significantly reduce latency (especially around the 99th percentile) while still supporting comparable TPM and NOPM performance rates.

Dedicated RAID Sets: For best performance and lowest latency, configure the platform to place the database and tempdb on a dedicated RAID10 set and the log on a dedicated RAID 1 set. This configuration yields better performance and improves the corresponding 99th percentile latency for new order transactions (write transactions).

- A single RAID 50 set can be used to provide additional capacity, but it reduces latency consistency (approximately 9%).
- If the log files must be placed on the same RAID set as the database and tempdb to save costs, protect the logs using periodic backups, database replication, or mirroring methods.

The Bottom Line

Micron's M500DC enterprise SSD is an ideal match for OLTP workloads running on Microsoft SQL Server 2014. It offers extremely low latency and consistent responsiveness, exceptional mixed workload performance, leading endurance (in its class), and an approachable price point. In order to take full advantage of the M500DC, additional care must be taken when configuring the host and application.

micron.com

This technical marketing brief is published by Micron and has not been authorized, sponsored, or otherwise approved by Microsoft. Products are warranted only to meet Micron's production data sheet specifications. Products and specifications are subject to change without notice.