Accelerating storage performance in the PowerEdge FX2 converged architecture modular chassis

This white paper highlights the impressive storage performance of the PowerEdge FD332 storage node in the FX2 modular chassis.

Brian Bassett, Solutions Performance Analysis

Dell | Enterprise Solutions Group

November 2015 | Version 1.22

Contents

Executive summary 4
Introduction 4
Key findings 4
Performance 4
Ease of management 4
Configuration 5
FX2 modular chassis 5
FC630 compute node 5
FD332 storage node 5
RAID configuration 5
Iometer software 5
Methodology 6
Performance 7
Single PERC performance 7
Dual PERC (Single Split Host mode) 9
Summary 11

Figures

Figure 1: FX2 modular chassis with two compute nodes and two storage nodes 5
Figure 2: Changing the Storage Mode of the FD332 through the Chassis Management Controller 6
Figure 3: Single PERC mode block diagram 7
Figure 4: Single PERC Iometer performance 7
Figure 5: Single Split Mode block diagram 9
Figure 6: Dual PERC performance in Single Split Host mode 9
Figure 7: Iometer performance summary 11

Executive summary

Introduction

This white paper details the performance characteristics of different modes of operation for the PowerEdge FD332 storage node. Administrators who need maximum performance for storage-heavy workloads, like databases, have three different configurations available to match their needs.

Key findings

Performance

- The PowerEdge FC630 server node combined with the FD332 storage node with a single PERC storage controller installed achieved greater than **204,000 IOPS** (I/O per second) and **1,600 MB/s** (megabytes per second) throughput.
- The FC630 server node combined with the FD332 storage node with dual PERC storage controllers in Single Split Host mode achieved greater than **350,000 IOPS** and **2,737 MB/s** throughput, a **71% improvement** over the single PERC results.

Ease of management

- The PowerEdge FX2 modular chassis features one-stop management of the configuration of the storage nodes, letting administrators easily configure storage nodes to fit their needs from a single administration console.

Test methodology and detailed results are documented in this paper.

Configuration

FX2 modular chassis

The FX2 modular chassis used for this white paper is a four-bay chassis with a FC630 server node and FD332 storage node. This chassis can be configured with two compute nodes and two storage nodes, as seen in Figure 1. In that configuration, each compute node connects to one of the storage nodes. For the testing described in this whitepaper, we used a single FC630 compute node and single FD332 storage node.

Figure 1: FX2 modular chassis with two compute nodes and two storage nodes

FC630 compute node

To ensure compute performance would not be a bottleneck in this testing, the FC630 compute node was configured with two Intel® Xeon® E5-2640 v3 processors and a total of 128 GB of DDR4 memory. Microsoft® Windows Server® 2012 R2 was installed on the internal 300GB SAS hard drives in RAID 1 configuration.

FD332 storage node

In order to push the boundaries of maximum performance, the FD332 storage node was equipped with 16 x 400 GB Mixed Use MLC SAS SSD drives, each connected to the PERC FD33xS controller(s) at their maximum rate of 12 Gbps (gigabits per second).

RAID configuration

For each test, two RAID 10 virtual disks (VDs) were created. Since all drives are SSD, Write Policy for the VDs was set to Write Through and Read Policy was set to No Read Ahead.

Iometer software

The Iometer storage testing software has long been a valuable tool for assessing storage performance. For this testing, Iometer was configured to run a simulated database access pattern consisting of 70% reads, 30% writes, 8k access size, and 100% random accesses. To ensure saturation of the storage, 8 Iometer workers (4 to each VD) were run for each test.

Methodology

In the dual PERC configuration, changing between the available modes is as easy as shutting down the compute node, then accessing the Setup page for the storage node. For the testing detailed in this white paper we used Split Single Host mode, which splits the 16 installed hard drives into two groups of 8, with each group having its own dedicated PERC controller. Also available, but not tested for this white paper, is Split Dual Host mode, which allows two compute nodes to have access to each set of 8 drives in the storage sled.

Figure 2: Changing the Storage Mode of the FD332 through the Chassis Management Controller

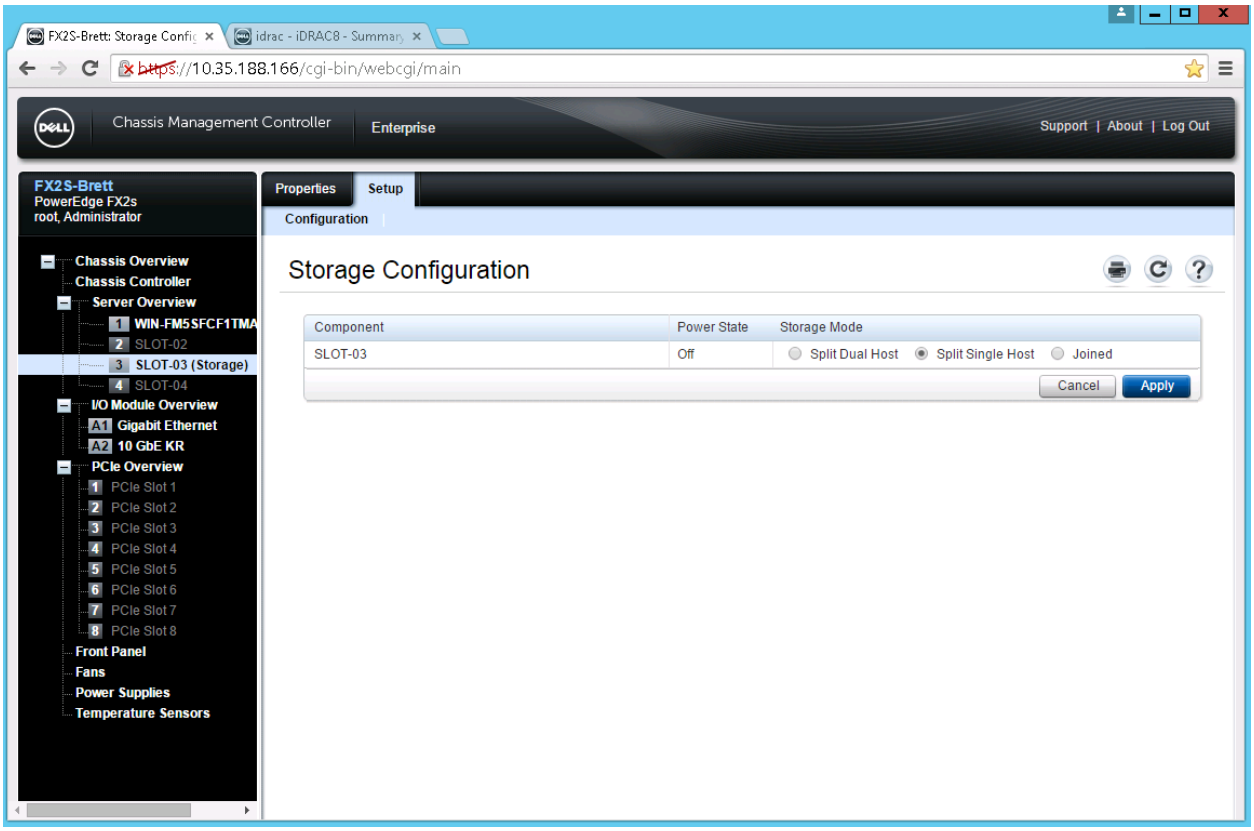

Performance

Single PERC performance

First we tested in the base configuration with a single PERC controller connected to all 16 SSDs in the FD332 storage node. In this configuration, each compute node has access to all 16 SSDs in the storage node installed below it, as seen in Figure 3. The FX2 4-bay chassis could also host an additional compute and storage node combination in the rightmost slots, but for our testing we used a single compute node and a single storage node.

Figure 3: Single PERC mode block diagram

As Figure 4 shows, performance in this configuration is already impressive, with this completely random access pattern topping out at over 200k IOPS and more than 1.6 GB/s of throughput.

Figure 4: Single PERC Iometer performance

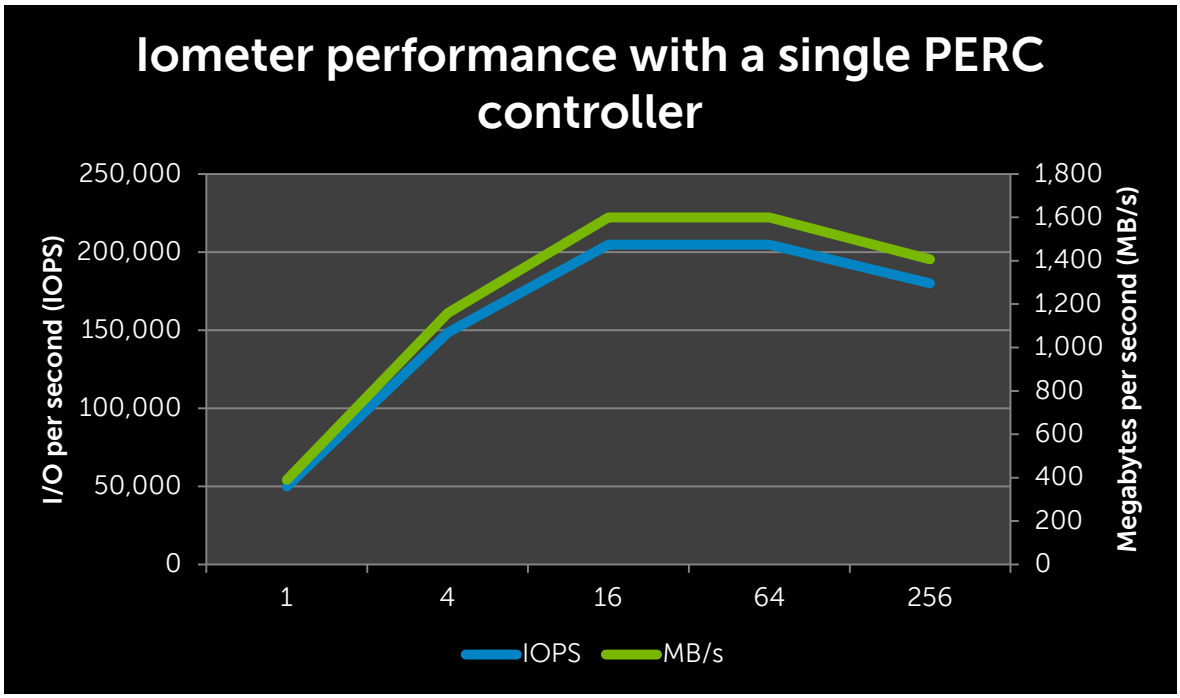

Table 1: Storage performance summary with a single PERC installed

Single PERC (2 VD, 8 workers) - RAID 10					
Queue Depth	1	4	16	64	256
IOPS	49,882	148,337	204,703	204,751	179,984
Read IOps	33,386	99,393	137,194	137,138	120,560
Write IOps	16,495	48,944	67,508	67,613	59,423
MB/s	390	1,159	1,599	1,600	1,406
Read MBps	261	777	1,072	1,071	942
Write MBps	129	382	527	528	464

Dual PERC (Single Split Host mode)

When dual PERCs are present, the FD332 storage node can be configured in Single Split Host mode. This mode allows the FC630 compute node to see both PERC controllers independently, with each controlling 8 of the 16 hard drives, as seen in Figure 5. In this mode, it is necessary to create at least 2 virtual disks in order to use all the storage present in the FD332.

Figure 5: Single Split Mode block diagram

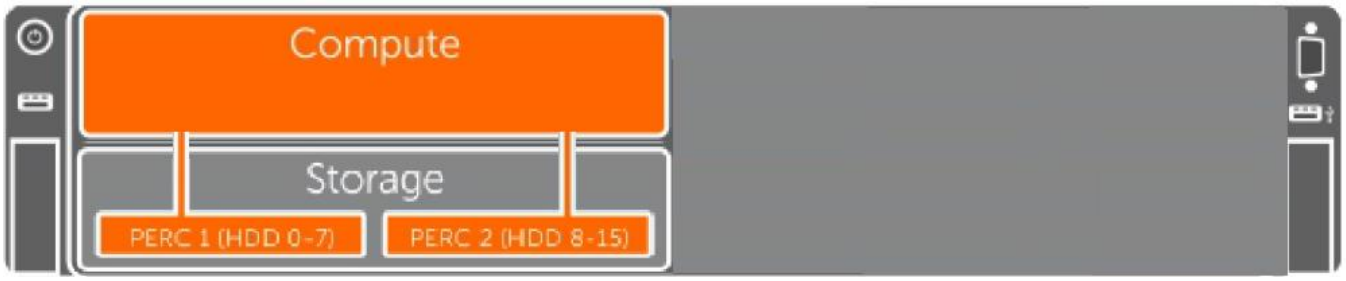

Iometer was run against both of these virtual disks simultaneously and the IOPS and MB/s result aggregated. Figure 6 shows the results of this testing, with IOPS topping 350k and throughput of over 2.7 GB/s.

Figure 6: Dual PERC performance in Single Split Host mode

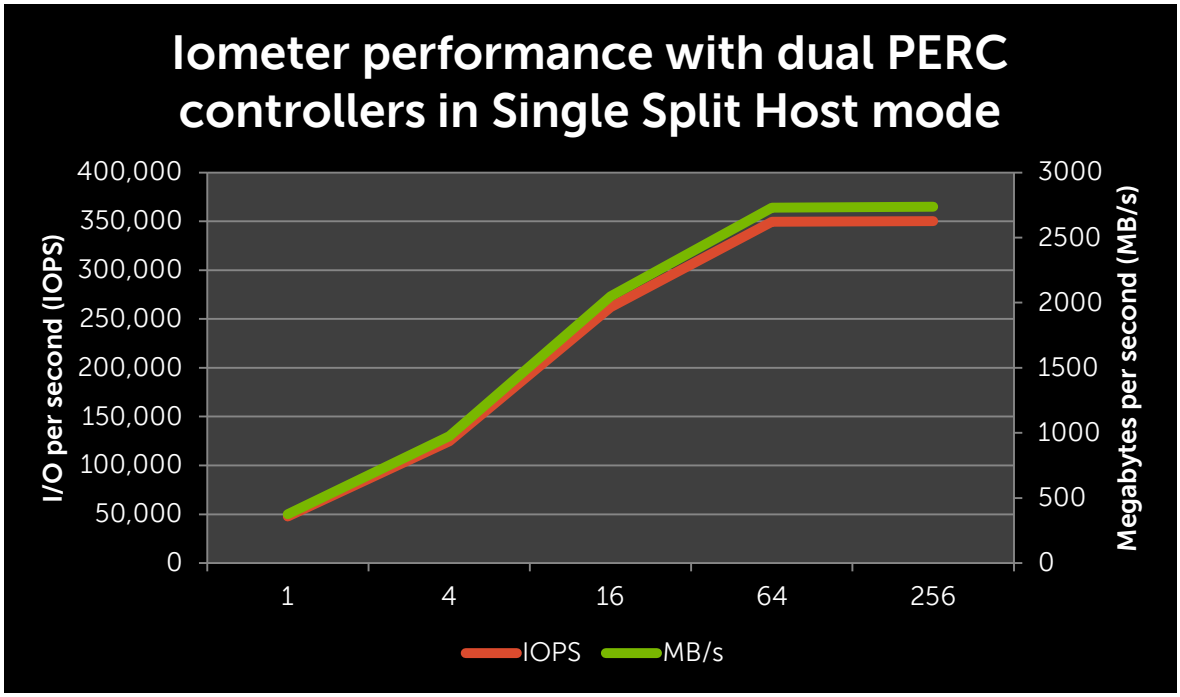

Table 2: Storage performance summary with dual PERC

Single Split Host (Dual PERC, 2 VD, 8 workers) - RAID 10						
Queue Depth	1	4	16	64	256	**% improvement over Single PERC**
IOps	47,823	124,835	262,111	349,542	350,356	71%
Read IOps	33,470	87,357	183,438	244,703	245,240	79%
Write IOps	14,352	37,478	78,673	104,838	105,116	55%
MBps	374	975	2,048	2,731	2,737	71%
Read MBps	261	682	1,433	1,912	1,916	79%
Write MBps	112	293	615	819	821	55%

Summary

For some workloads, the single PERC configuration may perform well, providing excellent performance at more than 204K IOPS. If splitting the drives into two groups of 8 is an option, configuring the FD332 with an additional PERC in Split Single Host mode provides more than 350K IOPS, almost twice the performance of single PERC mode.

Administrators looking for a high-performing storage solution have an array of options available in the FX2 converged architecture modular chassis.

Figure 7: Iometer performance summary