

Performance Benefits of Deploying Pivotal Greenplum on Dell EMC VxFlex integrated rack

Abstract

This paper highlights the performance results comparing the deployment of Greenplum Database on VxFlex integrated rack vs. Data Computing Appliance

July 2019

Revisions

Date	Description
February 2019	Initial release
July 2019	Updated VxRack FLEX to VxFlex integrated rack as per branding guidelines

Acknowledgements

This paper was produced by the following members of the Dell EMC HCI Solutions engineering team:

Author: Kailas Goliwadekar, Sudhakar Tekkali

Support: Bhanu Jamwal, Sanjay Narahari, Shashikir Chidambara

Other: Sridhar Subramanian

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.

© 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.

Dell 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
Table of contents	3
Executive summary.....	4
1 Introduction.....	5
1.1 Terminology.....	5
2 Technology overview.....	7
2.1 Pivotal Greenplum Database.....	7
2.2 Dell EMC Data Computing Appliance.....	7
2.3 VxFlex integrated rack.....	7
3 Test environment.....	9
3.1 Pivotal Greenplum on VxFlex integrated rack architecture	9
4 Test methodology VxFlex integrated rack vs. DCA.....	11
4.1 TPC-DS like workload configurations	11
4.1.1 Data load for 1 TB	11
4.1.2 99 TPC-DS queries	12
4.1.3 Performance metrics for TPC-DS like workload	12
4.1.4 1 TB multiuser test	13
4.1.5 3 TB multiuser test	13
5 Conclusion.....	14
A Appendix.....	15
A.1 1 TB VxFlex integrated rack vs. DCA data load detailed results.....	15
A.2 1 TB VxFlex integrated rack vs. DCA.....	16
A.3 Performance metrics parameters	20
A.4 Configurations.....	21
A.5 Related resources.....	23
A.6 Additional resources	23

Executive summary

This paper provides performance benefits of deploying the Pivotal Greenplum® Database on VxFlex integrated rack cluster. Pivotal Greenplum provides comprehensive and integrated analytics on multi-structured data.

Dell EMC VxFlex integrated rack (formerly VxRack FLEX) is a rack-scale hyperconverged system that delivers flexibility, scalability and performance for the enterprise data center. VxFlex integrated rack is a flexible architecture that allows multi-OS and multi-hypervisor capabilities and provides the ability to adapt to changing workloads. Scalability comes from starting small and growing incrementally, as well as growing compute and storage independently. It also delivers performance for all workloads in the environment, not just a few, with six-nines of tier one availability. VxFlex integrated rack with Dell EMC VxFlex OS software is a reliable, quick and easy to deploy solution that is ideal for server SAN, heterogeneous virtualized environments, and high-performance databases.

This paper highlights the performance results and compares the deployment of Pivotal Greenplum Database on VxFlex integrated rack vs. Data Computing Appliance (DCA).

These performance tests were carried out using the TPC DS like benchmark. This paper illustrates how Greenplum Database on VxFlex integrated rack performance metrics score* was 50 percent higher than DCA in all the tests that were carried out in a controlled environment.

*Score is one of the TPC-DS like benchmark performance metrics parameters. For more information, see [Table 5](#) in Appendix.

1 Introduction

Dell EMC VxFlex integrated rack is a rack-scale hyperconverged engineered system that delivers flexibility, scalability, performance for the modern data center workloads. The VxFlex integrated rack is powered by software defined storage VxFlex OS, widely adopted enterprise virtualization technology running on enterprise class Dell PowerEdge servers. The VxFlex integrated rack flexible architecture enables not only multi-hypervisor capabilities but also multiple deployment options such as fully hyperconverged, two-layer, hybrid, and bare-metal to become the infrastructure of choice for modern and traditional workload. Scalability comes from starting small and growing incrementally, but also growing compute and storage independently. VxFlex integrated rack also delivers performance for all workloads in the environment.

The TPC Benchmark DS is a decision support benchmark that models several applicable aspects of a decision support system, including queries and data maintenance. Pivotal Greenplum was installed on a VxFlex integrated rack cluster. TPC-DS like queries were run on the VxFlex integrated rack cluster as well as on a Greenplum DCA appliance. This paper compares the results of these tests on both the VxFlex integrated rack cluster and the Greenplum DCA appliance.

1.1 Terminology

The following table defines acronyms and terms that are used throughout this document:

Table 1 Terminology

Term	Definition	Description
DAS	Direct Attached Storage	Storage device/devices that is/are attached directly to a computer or a server.
MDM	Meta Data Manager	A VxFlex OS component that maintains storage cluster meta data information.
SDC	Storage Data Client	A VxFlex OS component that consumes storage from the software defined storage cluster.
SDS	Storage Data Server	A VxFlex OS component that contributes its DAS to the software defined storage cluster.
SVM	Storage Virtual Machine	A VM in ESXi environment that runs SDC, SDS, and MDM components of VXFflex operating system.
MDW	Master Host	The master is the entry point to the Greenplum Database system and the database instance to which users connect and submit SQL statements. For more information, see Introduction to Greenplum ETL tool – Overview .
SDW	Segment	The segment nodes that handle data processing and storage. For more information, see Introduction to Greenplum ETL tool – Overview .
RHV	Red Hat Virtualization	Refers a complete open-source virtualization solution, which is derived from the Red Hat Enterprise Linux kernel, Kernel-based Virtual Machine (KVM) technology, and oVirt virtualization management projects

Term	Definition	Description
MPP	Massively Parallel Processing	Refers to the use of large number of processors to perform a set of coordination computations in parallel.
TPC-DS (like)	NA	TPC-DS benchmark scripts for Greenplum database from https://github.com/pivotalguru/TPC-DS .

2 Technology overview

2.1 Pivotal Greenplum Database

The Pivotal Greenplum Database is an MPP database server that supports next generation data warehousing and large-scale analytics processing. By automatically partitioning data and running parallel queries, it enables a cluster of servers to operate as a single database supercomputer performing tens or hundreds of times faster than a traditional database. It supports SQL, MapReduce parallel processing, and data volumes ranging from hundreds of gigabytes, to hundreds of terabytes.

2.2 Dell EMC Data Computing Appliance

Dell EMC DCA is an integrated analytics platform that accelerates Big Data analyses within a single appliance, delivering faster time to value and lower integration risks and total cost. EMC DCA maximizes security, availability, and performance for an analytic data warehouse—without the complexity and constraints of proprietary hardware. DCA provides a converged infrastructure - servers, networks, and storage for the distributed database systems. The EMC DCA is purpose built to be both modular and fully integrated. It includes independent modules for Pivotal Greenplum and Pivotal HD. EMC DCAs enhance flexibility and can be scaled at any time by adding new modules. The solution deploys multiple DCAs dedicated for database/Hadoop and functionality across multiple sites.

2.3 VxFlex integrated rack

VxFlex integrated rack is a modular hyperconverged platform that enables extreme scalability and flexibility for next generation cloud applications and mixed workloads. VxFlex integrated rack is designed for deployments involving large numbers of VMs and users. It provides the following features:

- Low provisioning complexity with simple-to-use functionality
- Flexible and variable configuration options
- Horizontal scaling by adding, moving, and removing nodes on-the-fly to extend DAS capacity exponentially
- Step-sized building blocks for future data center environments
- Pre-configured, pre-loaded, pre-tested, and fully optimized IT stack that is delivered as a fully assembled and supported system

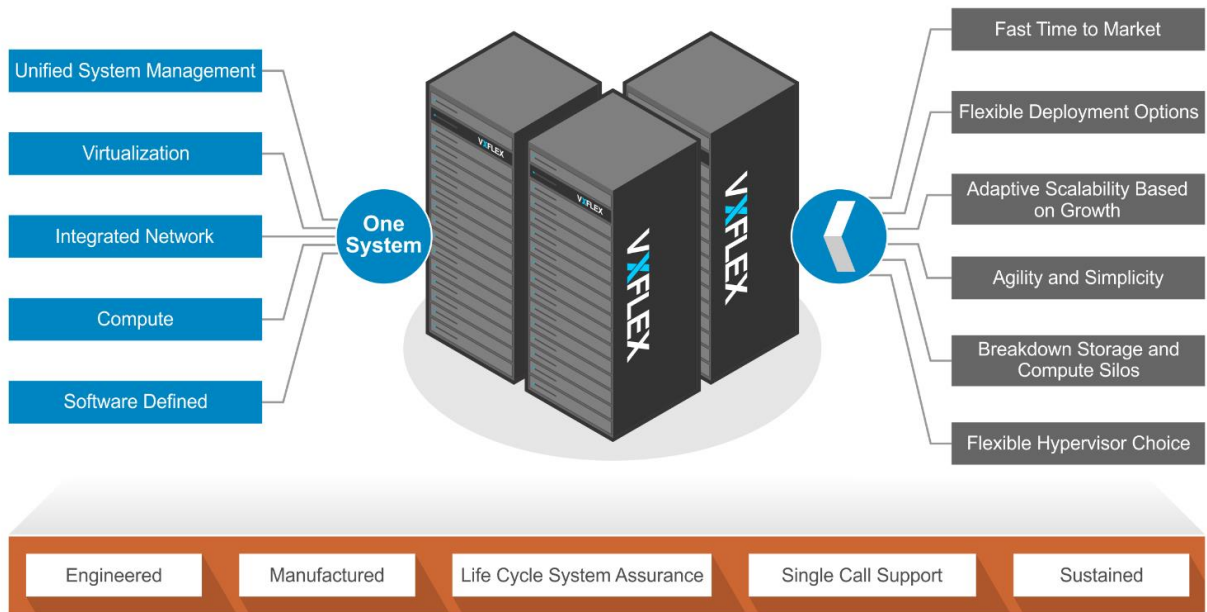


Figure 1 VxFlex integrated rack benefits

For more information about the product, see [Dell EMC Technical Resource Center](#).

3 Test environment

This section shows the test environment that was set up for Pivotal Greenplum on VxFlex integrated rack.

3.1 Pivotal Greenplum on VxFlex integrated rack architecture

For this solution, a VxFlex integrated rack cluster comprised of three VxFlex integrated rack R740xd nodes were used. The SDS and SDC were installed on the nodes. Virtual machines using an RHV hypervisor were created on the nodes and the VxFlex OS volumes were mapped to these virtual machines. Two virtual machines were created on one node and one virtual machine were created on each of the other two nodes.

Greenplum Database was installed on the virtual machines. One of the virtual machines was configured as the MDW and the other three virtual machines were configured as the SDW. There was one master host and three segment hosts.

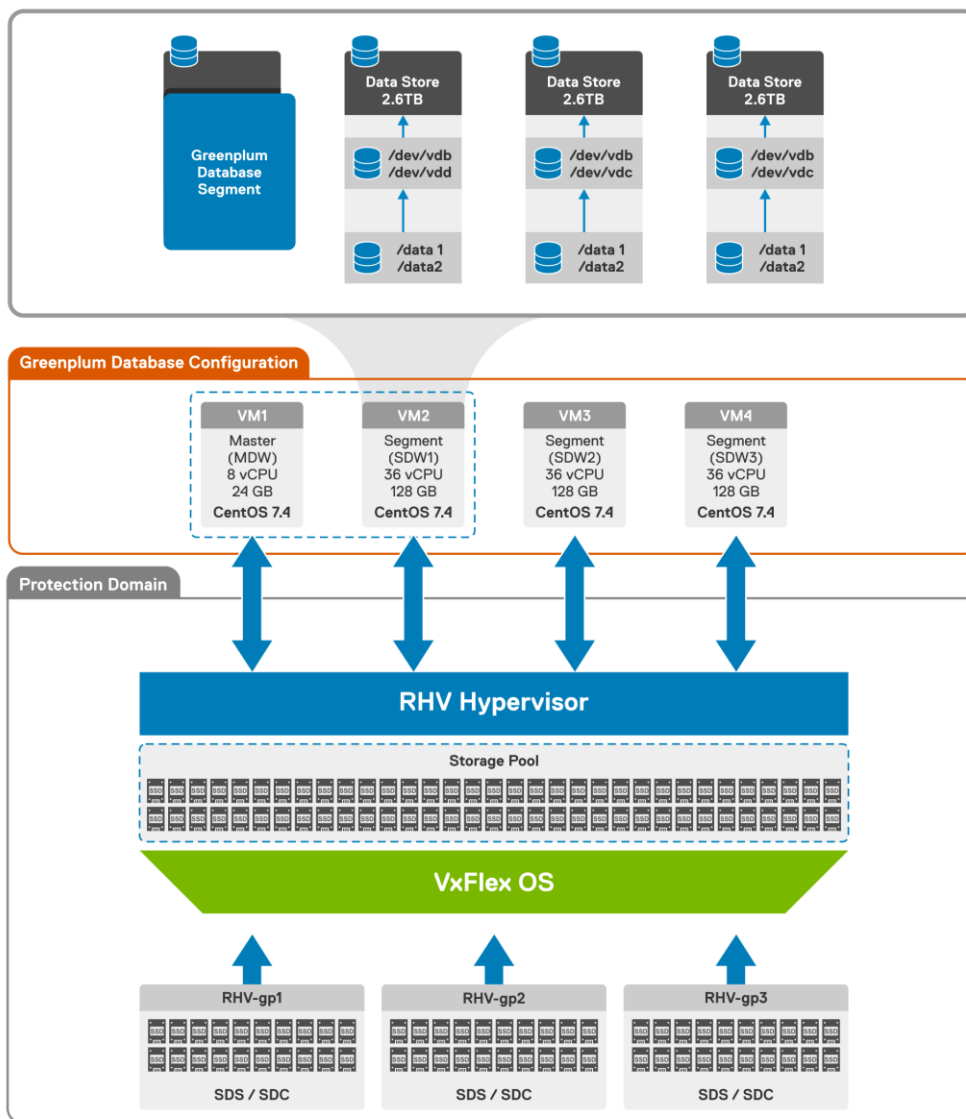


Figure 2 Pivotal Greenplum architecture diagram

There were three nodes in the VxFlex integrated rack cluster. All the nodes had a similar configuration.

A similar configuration was set up on the DCA. For a detailed description of the DCA nodes, see [Table 9](#) in Appendix.

4 Test methodology VxFlex integrated rack vs. DCA

There were two types of test approaches that are followed in Pivotal Greenplum performance testing, for example, baseline hardware performance and TPC-DS like performance benchmark.

- Baseline hardware performance was done using the **GPCHECKPERF** utility from Greenplum. The baseline hardware performance measures write, read, network performance, and other parameters. For more information about this utility, see [gpcheckperf](#).
- Different TPC-DS like datasets at different scale factors such as 1 TB and 3 TB were generated for VxFlex integrated rack and DCA. The data load time and execution of the 99 queries were captured. The TPC-DS like toolkit is available at <https://github.com/pivotalguru/TPC-DS>.

Note: TPC-DS like transactions for Greenplum database were simulated using the tool from the above URL.

4.1 TPC-DS like workload configurations

The purpose of TPC-DS like benchmarks is to provide relevant, objective performance data to industry users.

The testing scope for the study includes TPC-DS like workload scenarios that are listed below:

- Time for data load
- Running the 99 queries on Greenplum for single user
- Calculation of performance metrics
- Running the 99 queries on Greenplum for multi users
 - 1 TB
 - 3 TB

4.1.1 Data load for 1 TB

The time to load for 1-TB of data for VxFlex integrated rack was 30.76 minutes, whereas in DCA it was 40.9 minutes. This shows the data load on VxFlex integrated rack was 30 percent faster than DCA.

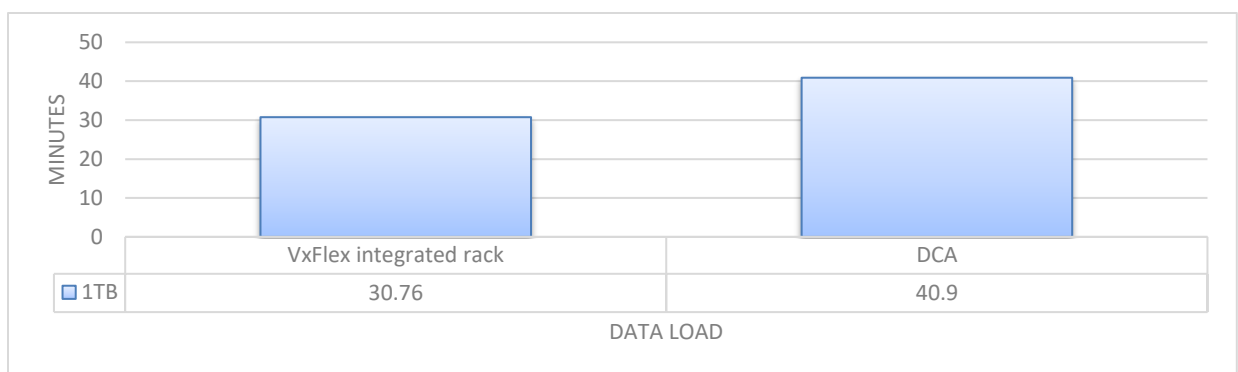


Figure 3 1 TB data load

For more information on data load, see [Table 2](#) in Appendix.

Note: VxFlex integrated rack loads 1 TB data 30 percent faster than DCA.

4.1.2 99 TPC-DS queries

The TPC-DS like result comparison between DCA and VxFlex integrated rack are shown below. For more information on the timings for all the 99 queries, see [Table 3](#) in Appendix.

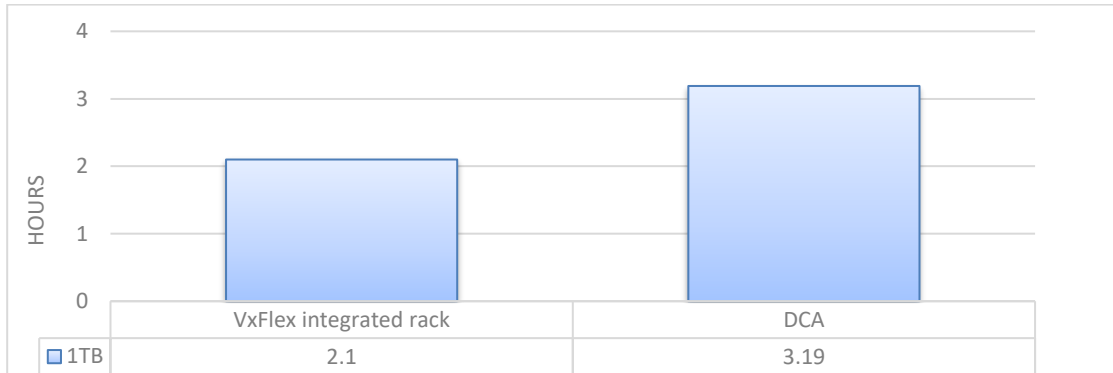


Figure 4 1 TB TPC-DS like single user total time query execution

Note: VxFlex integrated rack runs 1 TB single user queries 35 percent faster than DCA.

4.1.3 Performance metrics for TPC-DS like workload

The performance metrics that were captured for 1 TB of data on VxFlex integrated rack cluster and on DCA are shown below. These performance metrics are the standard set of parameters for TPC-DS like benchmark. For more information about these parameters, see https://github.com/pivotalguru/TPC-DS/blob/master/09_score/rollout.sh. Score is one of the TPC-DS like benchmark performance metrics parameters. The higher the number, the better the performance.

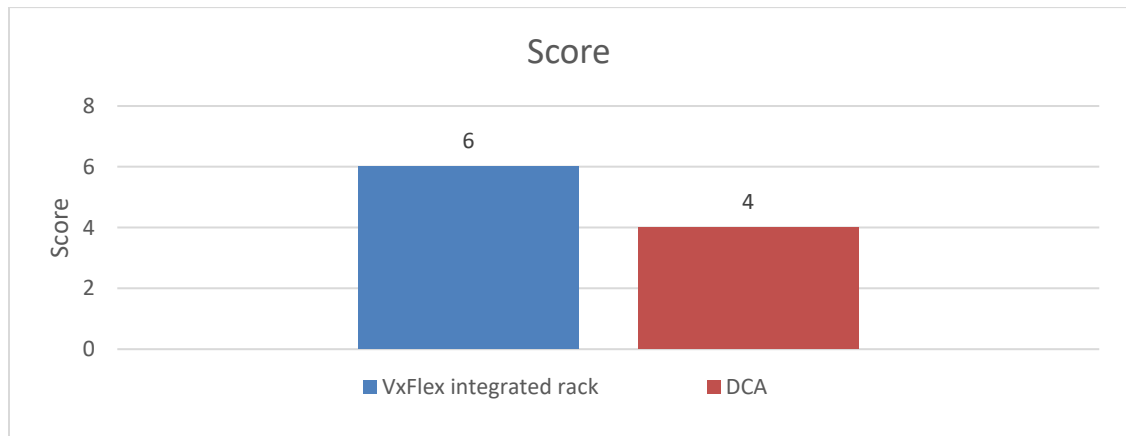


Figure 5 Performance metrics Score for TPC-DS like benchmark

Note: VxFlex integrated rack has a performance metrics score 50 percent higher than DCA.

For more information about performance metrics parameters on VxFlex integrated rack vs. DCA, see [Table 4](#) in Appendix.

4.1.4 1 TB multiuser test

The TPC-DS like 1 TB multiuser results comparison between VxFlex integrated rack and DCA for all 99 queries run time is shown below. In this case, five users running 1-TB TPC DS queries were evaluated.

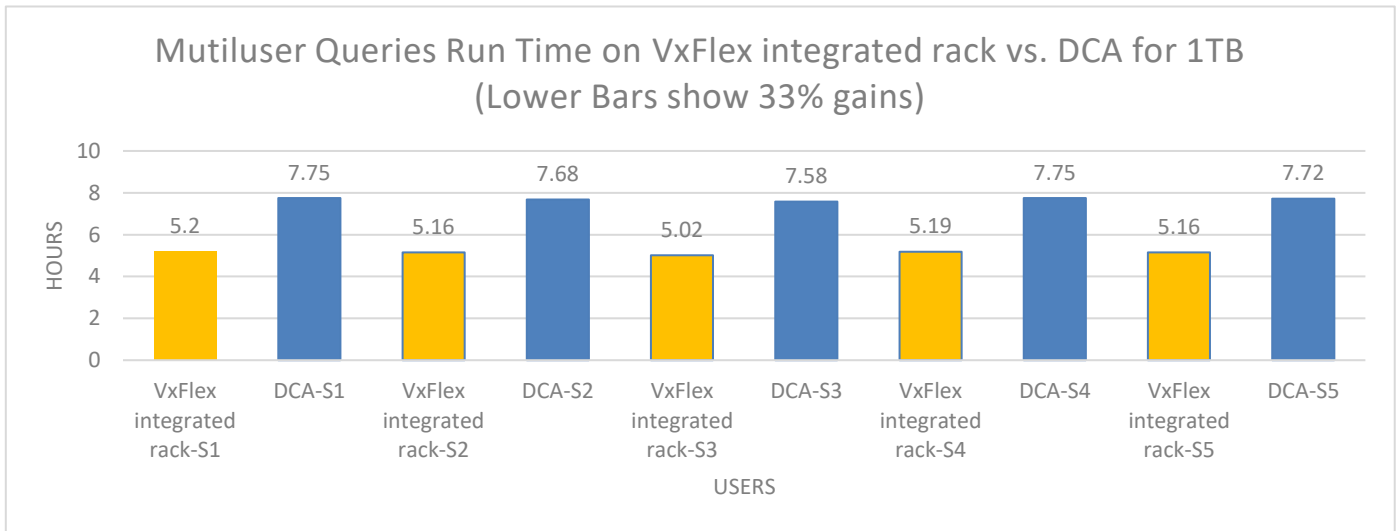


Figure 6 1 TB multiuser queries

Note: VxFlex integrated rack runs 1 TB multiuser user queries 33 percent faster than DCA.

4.1.5 3 TB multiuser test

The TPC-DS like 3 TB multiuser results comparison between VxFlex integrated rack and DCA for all the 99 queries run time is shown below. DCA could not complete the 3 TB multiuser test as the appliance was running the queries more than three days and became unresponsive.




Figure 7 Multiuser queries run time on VxFlex integrated rack 3 TB

Note: DCA could not complete the 3 TB multiuser test. VxFlex integrated rack took about 31 hours for each user.

5 Conclusion

This paper shows the performance results for Pivotal Greenplum on VxFlex integrated rack vs. DCA. The TPC-DS like tests were performed on VxFlex integrated rack and DCA for Pivotal Greenplum. Initially, baseline hardware performance tests were carried out, which was followed by other TPC-DS like tests.

- For the baseline hardware performance test, VxFlex integrated rack was 32 percent better than DCA for write operations, and 63 percent higher for read operations.
- Data load on VxFlex integrated rack was 35 percent faster than DCA.
- VxFlex integrated rack outperformed DCA for 1 TB TPC-DS like single user test by 35 percent.
- For the 1 TB multiuser test, VxFlex integrated rack ran 99 queries 32 percent faster than DCA.
- For the 3 TB multiuser test, DCA could not complete the test, whereas VxFlex integrated rack could run all the queries successfully.

In all cases, VxFlex integrated rack out performed DCA by at least 30 percent. Some of the other key features of VxFlex integrated rack engineered system are:

- Delivers unmatched **performance**
- Unprecedented **scale** (1000 nodes and 100s PB storage capacity)
- Build in redundant hardware components and VxFlex OS mesh mirror architecture delivers unparalleled **resiliency**

A Appendix

A.1 1 TB VxFlex integrated rack vs. DCA data load detailed results

Table 2 1 TB VxFlex integrated rack vs. DCA data load detailed results

Table name	VxFlex integrated rack (seconds)	DCA (seconds)
call_center	0.38	0.14
catalog_page	0.30	0.18
catalog_returns	101.10	105.11
catalog_sales	418.42	614.61
customer	2.25	3.34
customer_address	1.14	1.14
customer_demographics	0.44	0.38
date_dim	0.32	0.42
household_demographics	0.12	0.11
income_band	0.49	0.10
inventory	37.37	69.69
item	1.11	1.18
promotion	0.34	0.15
reason	0.07	0.10
ship_mode	0.46	0.12
store	0.37	0.15
store_returns	58.59	93.93
store_sales	977.98	1207.12
time_dim	0.28	0.34
warehouse	0.11	0.12
web_page	0.54	0.14
web_returns	18.18	26.26
web_sales	225.23	329.33

Table name	VxFlex integrated rack (seconds)	DCA (seconds)
web_site	0.37	0.14

A.2 1 TB VxFlex integrated rack vs. DCA




Figure 8 99 queries of VxFlex integrated rack vs. DCA

Note: VxFlex integrated rack runs 1 TB single users queries 35 percent faster than DCA.

Table 3 99 queries of VxFlex integrated rack vs. DCA

Queries	VxFlex integrated rack (seconds)	DCA (seconds)
1	9.99	13.13
2	127.13	194.19
3	7.78	14.14
4	281.28	472.47
5	83.83	132.13

Queries	VxFlex integrated rack (seconds)	DCA (seconds)
6	2.26	3.40
7	13.14	25.25
8	3.35	6.64
9	123.12	149.15
10	12.12	18.18
11	172.17	279.28
12	1.14	1.18
13	24.25	39.40
14	213.21	334.33
15	4.43	6.61
16	49.50	86.87
17	23.23	32.33
18	22.23	32.32
19	3.37	5.60
20	1.16	2.22
21	2.21	2.26
22	27.28	35.36
23	1266.13	1942.19
24	616.62	958.96
25	28.28	37.38
26	7.73	11.11
27	12.13	20.21
28	63.63	90.90
29	53.54	113.11
30	5.50	5.59
31	25.26	42.42
32	4.42	6.62

Queries	VxFlex integrated rack (seconds)	DCA (seconds)
33	4.46	6.66
34	18.19	31.31
35	33.33	48.49
36	44.44	61.61
37	14.14	26.27
38	62.63	106.11
39	23.24	34.35
40	8.81	16.17
41	0.20	0.34
42	2.22	3.31
43	26.27	40.40
44	26.26	47.48
45	4.42	5.60
46	52.52	66.66
47	73.73	101.10
48	25.26	33.34
49	7.76	12.13
50	55.56	179.18
51	62.62	104.10
52	2.22	3.31
53	9.97	15.15
54	4.44	5.59
55	2.28	4.42
56	5.52	6.62
57	41.41	53.54
58	2.26	3.38
59	171.17	252.25

Queries	VxFlex integrated rack (seconds)	DCA (seconds)
60	6.62	9.92
61	8.90	13.14
62	38.38	63.64
63	8.88	13.13
64	107.11	186.19
65	127.13	194.19
66	17.17	28.28
67	1177.12	1431.14
68	14.14	25.25
69	10.11	15.15
70	51.51	78.79
71	5.57	8.90
72	234.23	342.34
73	9.99	15.16
74	134.13	198.20
75	89.90	149.15
76	38.39	66.66
77	5.53	8.81
78	285.29	409.41
79	54.55	91.91
80	87.87	84.85
81	8.84	10.11
82	46.47	132.13
83	2.27	3.37
84	4.43	7.79
85	14.15	20.20
86	15.15	19.19

Queries	VxFlex integrated rack (seconds)	DCA (seconds)
87	60.60	105.11
88	91.91	145.15
89	12.12	18.18
90	11.12	20.21
91	1.16	2.26
92	2.29	3.38
93	74.74	126.13
94	25.26	48.48
95	390.39	639.64
96	38.39	69.69
97	126.13	187.19
98	2.23	3.37
99	74.74	126.13

A.3 Performance metrics parameters

Table 4 TPC-DS like benchmark performance metrics parameters on VxFlex integrated rack vs. DCA

Performance Metrics Parameters	VxFlex integrated rack	DCA	Percent Difference
Scale factor	1000	1000	
Load	1845.96	2454.30	25%
Analyze	234.23	451.45	48%
1 user queries	7591.68	11505.71	34%
Concurrent queries	92677.05	138573.23	33%
Q	1485	1485	0%
TPT	37958.39	57528.56	34%
TTT	92677.05	138573.23	33%
TLD	92.30	122.715	25%
Score	6	4	50%

Table 5 Performance metrics parameters

Performance Metrics Parameters	Comments
Scale factor	Total data volume tested
Load	Time taken to load
Analyze	Time that is taken for performing analyze
1 user queries	Time that is taken for single queries running
Concurrent queries	Time taken while running concurrent queries
Q	Total number of weighted queries
TPT	$T_{Power} * S_q$, where T_{Power} is the total elapsed time to complete the Power Test, and S_q is the number of streams that are run in a Throughput Test.
TTT	$TTT1 + TTT2$, where $TTT1$ is the total elapsed time of Throughput Test 1 and $TTT2$ is the total elapsed time of Throughput Test 2.
TLD	TLD is the load factor that is computed as $TLD = 0.01 * S_q * T_{Load}$, where S_q is the number of streams that are run in a Throughput Test and T_{Load} is the time to finish the load.
Score	Higher the number, better the performance.

A.4 Configurations

Table 6 VxFlex integrated rack node configuration

Component	Definition
Server	Dell EMC VxFlex integrated rack R740xd
CPU	2 socket Intel(R) Xeon(R) Gold 6126 CPU @ 2.60GHz. Each has 12 physical cores (24 logical cores)
Memory	192 GB
Network	4 x 25-Gb NIC ports (Mellanox)
Disks	20 x Dell Express Flash PM1725a 800 GB SFF
Operating system	CentOS Linux release 7.4.1708 (Core)

Component	Definition
Hypervisor	KVM
VxFlex OS	R2_6.10002.101

Table 7 Platform details

Vendor	Name	Version	Description
Dell EMC	VxFlex OS	R2_6.10002.101	SDS
VMware	KVM	6.5	Hypervisor
VMware	Virtual Manager	6.5	Management
RedHat / CentOS	CentOS Linux release	7.4.1708 (Core)	Operating System (for Greenplum VM)

Table 8 The VxFlex OS and Greenplum segment configuration test configuration

Component	Definition
Nodes	3 SDS and 3 SDC nodes in HCI configuration
VxFlex OS volumes	8 volumes in total
Greenplum Database	Six of 2.6 TB volumes are used for the database
Segment nodes	Two of 2.6 TB volumes were mapped to each segment node One 16 GB volume and one 2.6 TB volume was mapped to the master
vCPUs	36 vCPUs to master and to each segment

Table 9 Test environment DCA node configuration for Greenplum database

Component	Definition
Total CPU cores	32
Total RAM	62 GB
Total drives per node	24
Drive type	SAS
Number of nodes	4
Total storage per nodes	21.6 TB

A.5 Related resources

For more information related to this solution, see the following links:

Note: The links below are open to customers although some may require registration for access.

- [VxFlex OS blog](#)
- [VxFlex integrated rack Datasheet](#)
- [Pivotal Greenplum Database features](#)
- [Pivotal Greenplum 5.10.2 Release Notes](#)
- [TPS Benchmark DC](#)
- [Pivotal Network](#)
- [EMC Greenplum Database collateral](#)
- [Github links](#)
- [Big Data Support Site](#)
- [Pivotal Greenplum Best Practices Summary](#)
- [Introduction to Greenplum ETL tool – Overview](#)

A.6 Additional resources

[Dell.com](#) is focused on meeting customer needs with proven services and support.

[Dell EMC Technical Resource Center](#) provides expertise that helps to ensure customer success on Dell EMC VxFlex integrated rack.