

Getting the Most Out of your 32GFC Storage

Abstract

This paper demonstrates that upgrading your Fibre Channel SAN to 32 Gb/s Fibre Channel (GFC) delivers the best application performance with new 32GFC PowerMax storage arrays.

April 2023

Revisions

Date	Description
May 2020	Initial release
February 2022	Updated information
April 2023	Updated nomenclature

Acknowledgements

Authors:

- Erik Smith, Connectrix Product Engineering
- Dave Barry, Dell Technologies Customer Solution Centers

About Dell Technologies Customer Solution Centers

Dell Technologies Customer Solution Centers help customers explore and validate the Dell Technologies portfolio of hardware, software and solutions. Each of our global Customer Solution Centers is staffed by deep subject matter experts, working in state-of-the-art labs equipped with the latest technologies and showcases.

Services provided by Customer Solution Centers include technical briefings, architectural design workshops and proof-of-concepts that allow customers to validate their chosen solution.

Customers may engage with their account team to take advantage of one of these free services. Learn more at <https://www.dell.com/en-us/dt/what-we-do/customer-engagement-programs/customer-solution-centers.htm>

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 © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [4/17/2023] [White Paper] [2]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
Terminology	6
1 Customer use case	7
Benefits of end-to-end 32Gb/s Fibre Channel (32GFC).....	7
2 Test plan.....	8
Workload generation.....	8
3 Test results	9
3.1 8GFC end-to-end.....	9
3.1.1 Environment description	9
3.1.2 Performance results	9
3.2 32GFC storage introduced to 8GFC environment.....	10
3.2.1 Environment description	10
3.2.2 Performance results	11
3.3 16GFC end-to-end.....	13
3.3.1 Environment description	13
3.3.2 Performance results	13
3.4 32GFC storage introduced to 16GFC environment.....	14
3.4.1 Environment description	14
3.4.2 Performance results	15
3.5 32GFC end-to-end.....	16
3.5.1 Environment description	16
3.5.2 Performance results	16
3.6 32GFC storage in mixed 32GFC and 8GFC environment with Initiator Bandwidth Limits.....	17
3.6.1 Environment description	17
3.6.2 Test procedure.....	18
3.6.3 Initiator bandwidth limits	18
3.6.4 Performance results	19
4 Conclusion.....	22
4.1 The path to 32GFC	22
4.2 Additional benefits of 32GFC SAN Systems	22
A Configuration details.....	23

Table of contents

B	Workload generation script.....	25
C	Technical support and resources	27
C.1	Related resources	27

Executive summary

Customer challenge

As customers deploy Next Generation Storage arrays, special care must be taken to inspect the existing network for health and bandwidth capacity in order to implement a configuration that will optimize performance benefits and avoid issues such as congestion spreading (also known as slow drain). This white paper describes how to avoid the most common configuration issues to enable you to get the most from your 32GFC Storage.

Key findings and recommendations

1. The results in this paper are clear, upgrading your 8GFC or 16GFCb/s Fibre Channel SAN to 32GFC delivers the greatest application performance improvements for your 32GFC Storage.
2. If you must deploy 32G storage into a mixed SAN environment (8G or 16G switches along with 32G switches), the use of Initiator Bandwidth Limits available with PowerMax will allow you to utilize the full bandwidth of the 32G servers, switches and PowerMax paths by limiting the bandwidth for connections to slower devices. Without Initiator Bandwidth Limits throughput could be limited to the bandwidth supported by the slowest host (for example, 8G).
3. Whenever possible, work with your Account Team to perform a SAN Assessment to ensure your environment is configured and performing at the maximum capability.

Terminology

Congestion spreading

Fibre Channel (FC) and other Lossless protocols, such as DCB Ethernet and InfiniBand, prevent buffer overflow at either end of a link by allowing the transmitter to determine when the receiver at the other end of the link is nearing capacity. When a transmitter determines that a corresponding receiver is nearing capacity, the port stops transmitting data until the receiver is ready to receive additional data. While a transmitter is in this state, it is unable to transmit frames and it is said to be experiencing congestion. If a transmitter experiences congestion for a long enough duration, this congestion can propagate backwards towards the source; this phenomenon is known as congestion spreading.

Oversubscription

Oversubscription, in terms of FC SAN, is a condition where demand exceeds supply or capacity of the system, in terms of BW where a host HBA is running slower than the Storage Array port it's accessing, or where the BW ratio between Host and Array is other than 1:1.

Initiator bandwidth (BW) limits

Initiator bandwidth (BW) limits is a PowerMax feature to help users avoid running into congestion spreading in their SAN network. The user can define bandwidth on an Initiator Group (IG) and it will apply to all initiators within the IG, this constrains the amount of data sent from an array to a host. This feature is available on PowerMax arrays running PowerMaxOS 5978 Q2 2019 SR.

Customer use case

This white paper examines a common use case where a customer has an existing 8GFC or 16GFC end-to-end solution (servers, SAN and storage), is running ever more demanding applications on their infrastructure and has run out of storage capacity. To add additional storage capacity, the customer buys a next generation 32GFC PowerMax storage array to alleviate existing capacity issues. The customer has requested guidance on how to best configure the environment to attain the maximum benefits from their new PowerMax.

There are a few elements for the customer to consider that will affect the performance of the PowerMax array. For instance, what is the benefit of transitioning to end-to-end 32G servers, SAN and storage? How will a mixed-bandwidth environment affect PowerMax performance? What are the best options available to introduce the 32GFC PowerMax? The customer needs to maximize the benefits of their new 32GFC storage right from the start.

Risk

If the solution is not configured correctly, congestion spreading will limit the performance of the new storage array and the performance of the new applications running on the new 32G hosts will be roughly the same as the current applications running on the existing 8G servers. As a result, the customer will not experience the full benefits of their new PowerMax system.

Response

The test results clearly show that installing an end-to-end 32GFC environment, including servers, SAN and storage results in the best application performance and greatest benefit from the implementation of the 32GFC PowerMax array. In this way, the entire customer environment can seamlessly provide maximum throughput and response time to the applications on the array.

The test results presented in this white paper also demonstrate how congestion spreading, resulting from mixed bandwidth environments, can negatively affect the application performance on the 32GFC PowerMax throughput. For customers who need to manage mixed-bandwidth environments while transitioning to end-to-end 32GFC the Initiator Bandwidth Limits feature on the 32GFC PowerMax can help customers begin to experience the benefits of 32GFC storage during their transition to a full end-to-end 32GFC SAN solution.

Benefits of end-to-end 32Gb/s Fibre Channel (32GFC)

End-to-end 32GFC would be defined as 32GFC on the server's HBA, connection to a 32GFC switch or director port, ending with a 32GFC connection on the front-end-of the storage array. This configuration would improve your FC-SCSI performance.

Ultimately, if your storage array supports NVMe/FC protocol in the front end and it also uses NVMe protocol in the back end, and you deploy end-to-end 32GFC, you would now have an end-to-end NVMe/FC storage network designed for flash storage.

Test plan

Table 1 describes the tests that were run to demonstrate the increasing performance benefits while transitioning to an end-to-end 32FGC SAN.

Table 1 Test plan

Test	Details	Demonstrates
1. 8GFC end-to-end	8GFC storage, 8G SAN, 8G HBAs	The baseline achievable performance with 8GFC
2. 32GFC storage introduced to 8GFC environment	Combines the end-to-end 8GFC and 32GFC solutions	How congestion spreading impacts performance when 32G hosts and storage are added to an existing 8GFC SAN
3. 16GFC end-to-end	16GFC storage, 16G SAN, 16G HBAs	The baseline achievable performance with 16GFC
4. 32GFC storage introduced to 16GFC environment	Combines the end-to-end 16GFC and 32GFC solutions	How congestion spreading impacts performance when 32GFC storage is added to an existing end-to-end 16GFC SAN
5. 32GFC end-to-end	32GFC storage 32G SAN 32G HBAs	The baseline achievable performance with an end-to-end 32GFC SAN
6. 32GFC storage introduced to 8GFC environment with Bandwidth Limits implemented	Configure initiator-based bandwidth limits on the 32GFC PowerMax	How the introduction of bandwidth limits resolves congestion spreading and increases host performance.

Workload generation

Vdbench was used to generate the workloads for testing that was carried out. Vdbench is a command line utility specifically created to help engineers and customers generate disk I/O workloads to be used for validating storage performance and storage data integrity.

In Vdbench, the Decision Support System (DSS) workload profile was used for all tests.

Table 2 Workload profile

I/O Profile	Block Size	% of Total
Random Read Hit (RRH)	128K	18
Random Read Miss (RRM)	128K	18
Random Write Hit (RWH)	-	
Random Write Miss (RWM)	128K	4
Sequential Read (SR)	256K	48
Sequential Write (SW)	256K	12

Test results

3.1 8GFC end-to-end

This test was carried out to demonstrate the baseline achievable performance with an end-to-end 8GFC solution.

3.1.1 Environment description

Figure 1 illustrates a typical pre-existing 8GFC solution that consists of Dell PowerEdge servers with Fibre Channel Host Bus Adapter (HBA) ports attached to a two-switch Connectrix 8G Fibre Channel SAN. The HBA ports are running at 8Gb/s, and the Connectrix SAN has 8 x 8 Gbs/sec ISL links between the two switches. The storage is provided by a PowerMax with a Fibre Adapter (FA) front-end port running at 8Gb/s.

Note: The storage array used for testing was a PowerMax 2000. For this simulation the front-end ports were hard coded to 8GFC.




Figure 1 8GFC end-to-end solution

3.1.2 Performance results

Running a DSS workload profile, Vdbench was used to produce a maximum throughput of 3 x 8G servers accessing an 8GFC storage array. Refer to Appendix B (Workload Generation Scripts) to review the Vdbench script that was used.

The performance achieved was ~1600 MBs/sec per server or ~800 MBs/sec per port, as demonstrated by the test results captured in Figure 2.

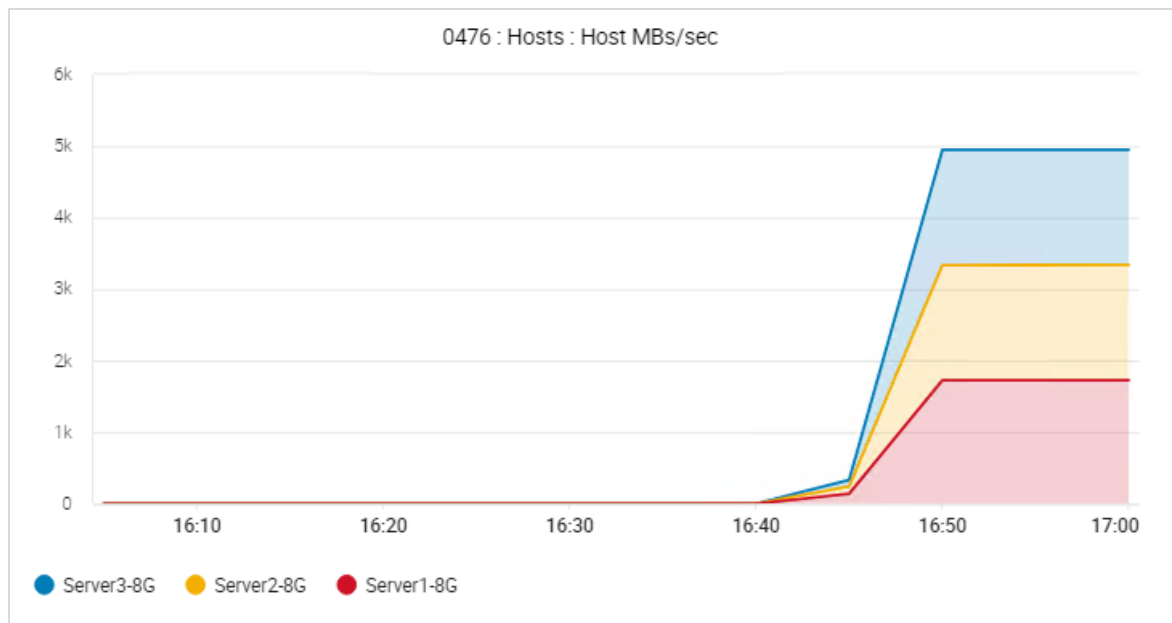


Figure 2 8GFC end-to-end: Host MBs/sec

3.2 32GFC storage introduced to 8GFC environment

This test was carried out to demonstrate how, if a solution is not configured correctly, congestion spreading can limit the performance of a 32GFC storage array added to an existing 8GFC SAN.

If the solution is not configured correctly, congestion spreading will limit the performance of the new storage array and the performance of the new applications running on the new 32G hosts will be roughly the same as the current applications running on the existing 8G servers. As a result, the customer will not experience the full benefits of their new PowerMax system.

3.2.1 Environment description

New 32GFC storage arrays are often not acquired for use in green field scenarios. A typical use case involves the addition of new storage to an existing solution where applications are running out of storage. Therefore, consideration must be given to how a new 32GFC storage array is added to an existing solution.

To expand the original end-to-end 8GFC solution, new 8G servers were added that access data on both the 8GFC and 32GFC storage arrays giving the mixed solution described in Figure 8. Six servers with 8G HBAs connected to an 8G edge switch, ISL'd to an 8GFC core switch and connected to 8GFC storage. In addition, there is an 8G edge switch ISL'd to a 32GFC core switch that is connected to 32GFC storage.

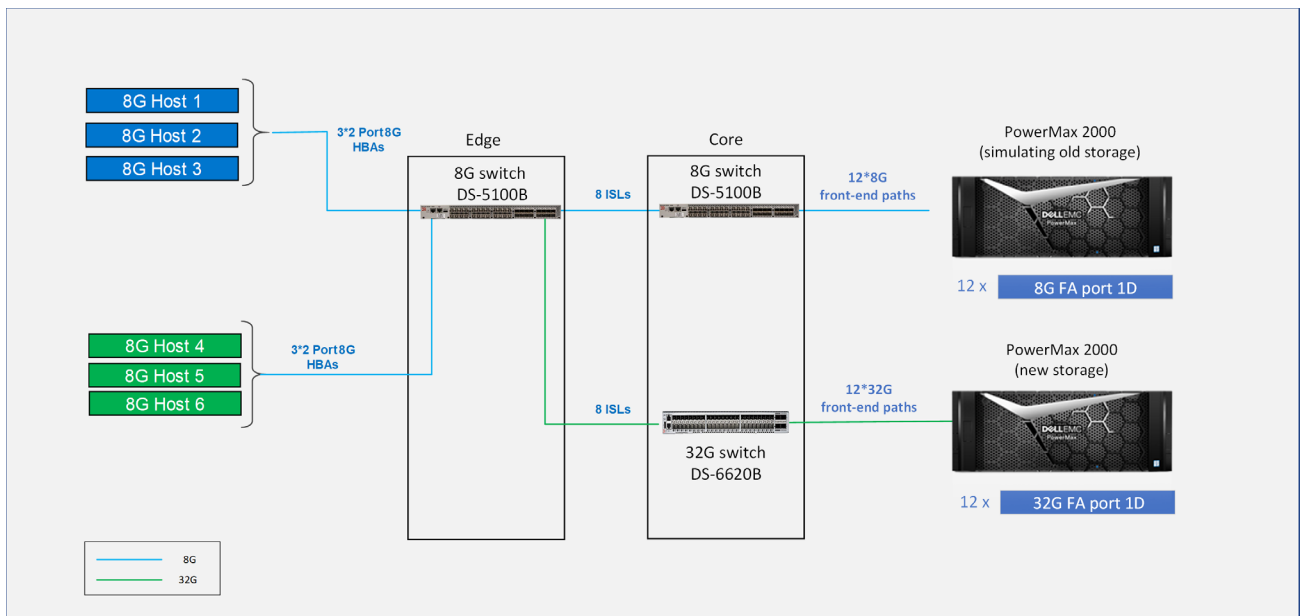


Figure 3 Mixed solution: 8G servers access both 8GFC and 32GFC storage

3.2.2 Performance results

The Vdbench script was started on the new 8G hosts (hosts 4, 5, and 6) using the same DSS profile as before.

Note: Each 8GFC server port is attempting to access an 8GFC interface on the 8GFC storage and a 32GFC interface on the 32GFC storage. As a result, the total bandwidth available on the arrays is 4000MB/s per HBA interface.

The throughput on each of the three new 8G servers (hosts 4, 5, 6) was ~800MB/s, as shown in Figure 4.




Figure 4 Servers 4, 5, 6 during script run: Host MBs/sec

However, there was a drop in the throughput on the original 8G servers (hosts 1, 2, 3) as shown in Figure 5. This drop occurred because the new 8G servers (800MB/s) were receiving data from both the 8GFC and

32GFC arrays (4000 MB/s total) and the hosts are only capable of consuming 800MB/s. The difference in the rate of consumption versus the rate of delivery (that is, 3200MB/s) had to be buffered by the SAN.




Figure 5 Servers 1, 2, 3 during script run: Host MBs/sec

The original servers (hosts 1, 2, 3) return to their maximum throughput when the Vdbench script was halted on the 3 new servers which are accessing both 8G and 32G storage. IO on hosts 4, 5 and 6 was started at time stamp 14:45 and stopped at time stamp 15:10, as illustrated in Figure 6.




Figure 6 Servers 4, 5, 6 script halted: Host MBs/sec

Figure 7 illustrates the impact of the IO on hosts 4, 5 and 6 on the performance of hosts 1, 2 and 3. At time stamp 14:45, when the IO was started on hosts 4, 5 and 6 a drop in performance was recorded on hosts 1, 2 and 3. Performance returns to the original throughput after the IO on hosts 4, 5 and 6 is stopped.

Since most FC switches have a modest number of buffers available, no single switch in the SAN will be able to buffer all the data. When a switch reaches its buffer capacity, it will no longer accept additional frames. When this happens on an ISL, if any two flows happen to share the same Virtual Channel (VC) they will both be impacted by congestion spreading.

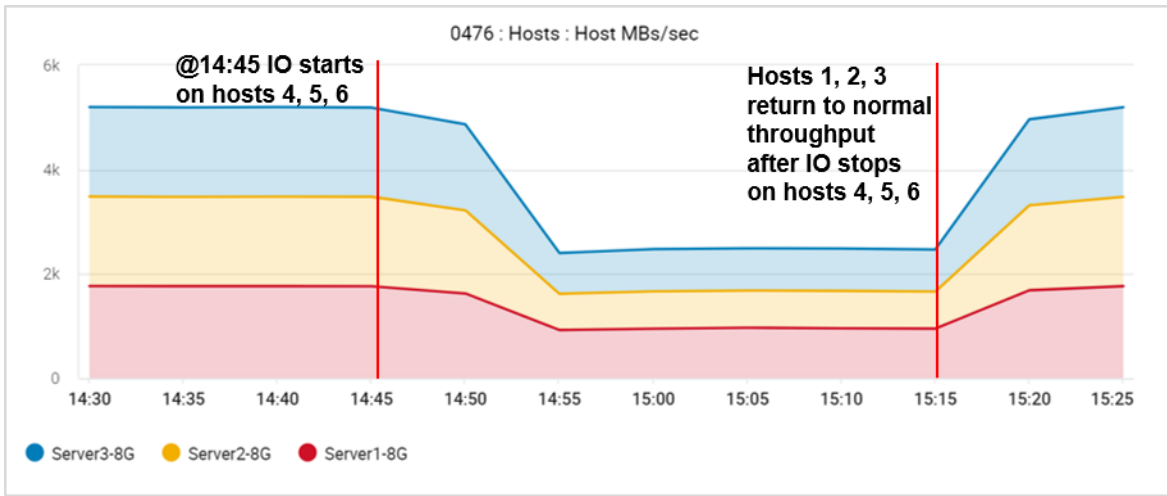


Figure 7 Servers 1, 2, 3 script halted: Host MBs/sec

3.3 16GFC end-to-end

This test was carried out to demonstrate the baseline achievable performance with an end-to-end 16GFC solution.

3.3.1 Environment description

Illustrates a typical pre-existing 8GFC solution that consists of Dell PowerEdge servers with Fibre Channel Host Bus Adapter (HBA) ports attached to a two-switch Connectrix 16G Fibre Channel SAN. The HBA ports are running at 16Gb/s, and the Connectrix SAN has 8 x 16 Gbs/sec ISL links between the two switches. The storage is provided by a Dell PowerMax with a Fibre Adapter (FA) front-end port running at 16Gb/s.

Note: The storage array used for testing was a PowerMax 2000. For this simulation the front-end ports were hard coded to 16GFC.

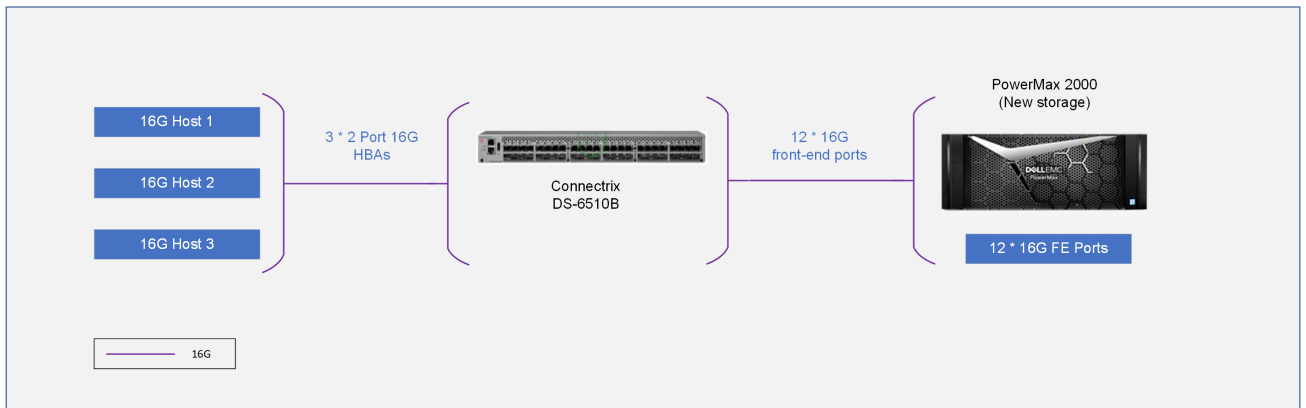


Figure 8 16GFC end-to-end solution

3.3.2 Performance results

Running a DSS workload profile, Vdbench was used to produce a maximum throughput of 3 x 16G servers accessing an 16GFC storage array. Refer to Appendix B (Workload Generation Scripts) to review the Vdbench script that was used.

The performance achieved was ~ 3.3GBs/sec per server as demonstrated by the test results captured in Figure 9.




Figure 9 16GFC end-to-end: Host MBs/sec

3.4 32GFC storage introduced to 16GFC environment

This test demonstrates that even if you have 32GFC storage, you will not get the full benefits unless you have a 32G SAN (+ compute).

3.4.1 Environment description

Figure 10 illustrates a typical pre-existing 16GFC compute that consists of PowerEdge servers with Fibre Channel Host Bus Adapter (HBA) ports attached to a single-switch Connectrix 32G Fibre Channel SAN. The storage is provided by a PowerMax with a Fibre Adapter (FA) front-end port running at 32Gb/s. Due to the 16GFC compute the end-to-end speed is negotiated to 16FC.




Figure 10 Mixed solution: 32GFC storage with 16G Hosts and 16GFC SAN switch

3.4.2 Performance results

Running a DSS workload profile, Vdbench was used to produce a maximum throughput of 3 x 16G servers accessing a 32GFC storage array. Refer to Appendix B (Workload Generation Scripts) to review the Vdbench script that was used.

The performance achieved was ~ 3.3GBs/sec per server as demonstrated by the test results captured in Figure 11.




Figure 11 16G hosts access 32GFC storage over a 16G SAN: Hosts MBs/sec

3.5 32GFC end-to-end

The test was carried out to demonstrate the baseline achievable performance with an end-to-end 32GFC solution.

3.5.1 Environment description

This solution consisted of three servers with 32G HBAs connected to a 32GFC edge switch, ISLed to a 32GFC core switch and connected to 32GFC PowerMax storage.

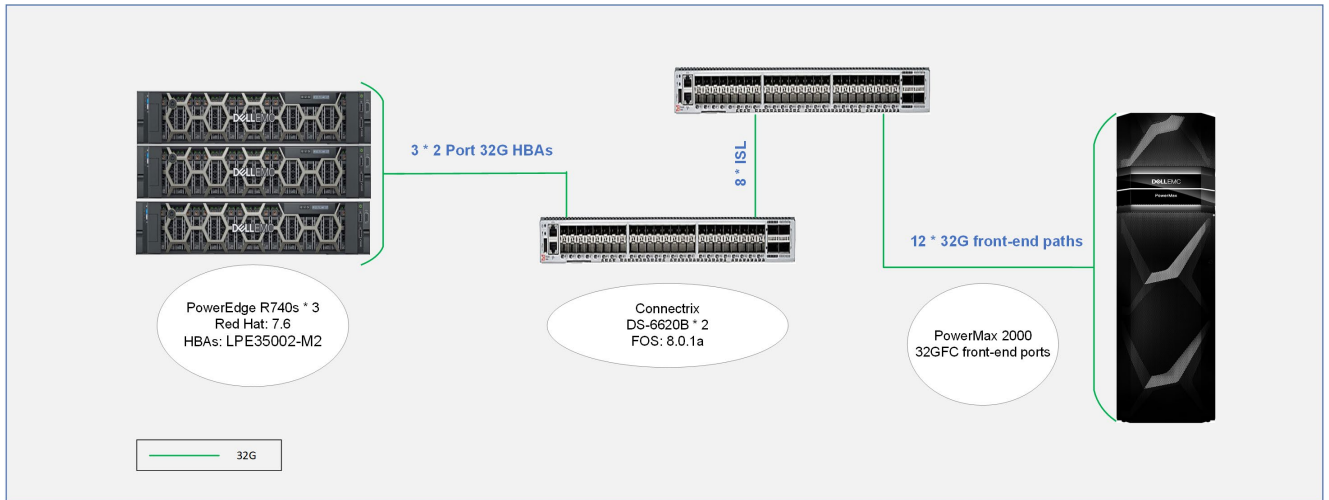


Figure 12 32GFC end-to-end solution

3.5.2 Performance results

To compare the performance profile of the 8GFC end-to-end solution with a 32GFC end-to-end solution the same Vdbench script (see Appendix B Workload generation script) was run against the configuration shown in Figure 12.

As expected, the throughput was much higher achieving ~7000MBs/sec per server as shown in Figure 13.




Figure 13 32G end-to-end: Host MBs/sec

3.6 32GFC storage in mixed 32GFC and 8GFC environment with Initiator Bandwidth Limits

This test was carried out to demonstrate how the introduction of initiator bandwidth limits resolves congestion spreading and maximizes the benefits of the 32G host, switch and PowerMax performance in mixed 32GFC and 8GFC environments, by limiting the bandwidth available to the slower connections.

3.6.1 Environment description

To demonstrate how the performance of 32GFC storage can be optimized in a mixed 8GFC and 32GFC environment, the solution shown in Figure 14 was built and tested.

- 3 x 8G servers (hosts 1, 2, 3) access storage on the 8GFC array
- 3 x 8G servers (hosts 4, 5, 6) access storage on the 32GFC array
- 3 x 32G servers (hosts xyz, zxy, yzx) access storage on the 32GFC array

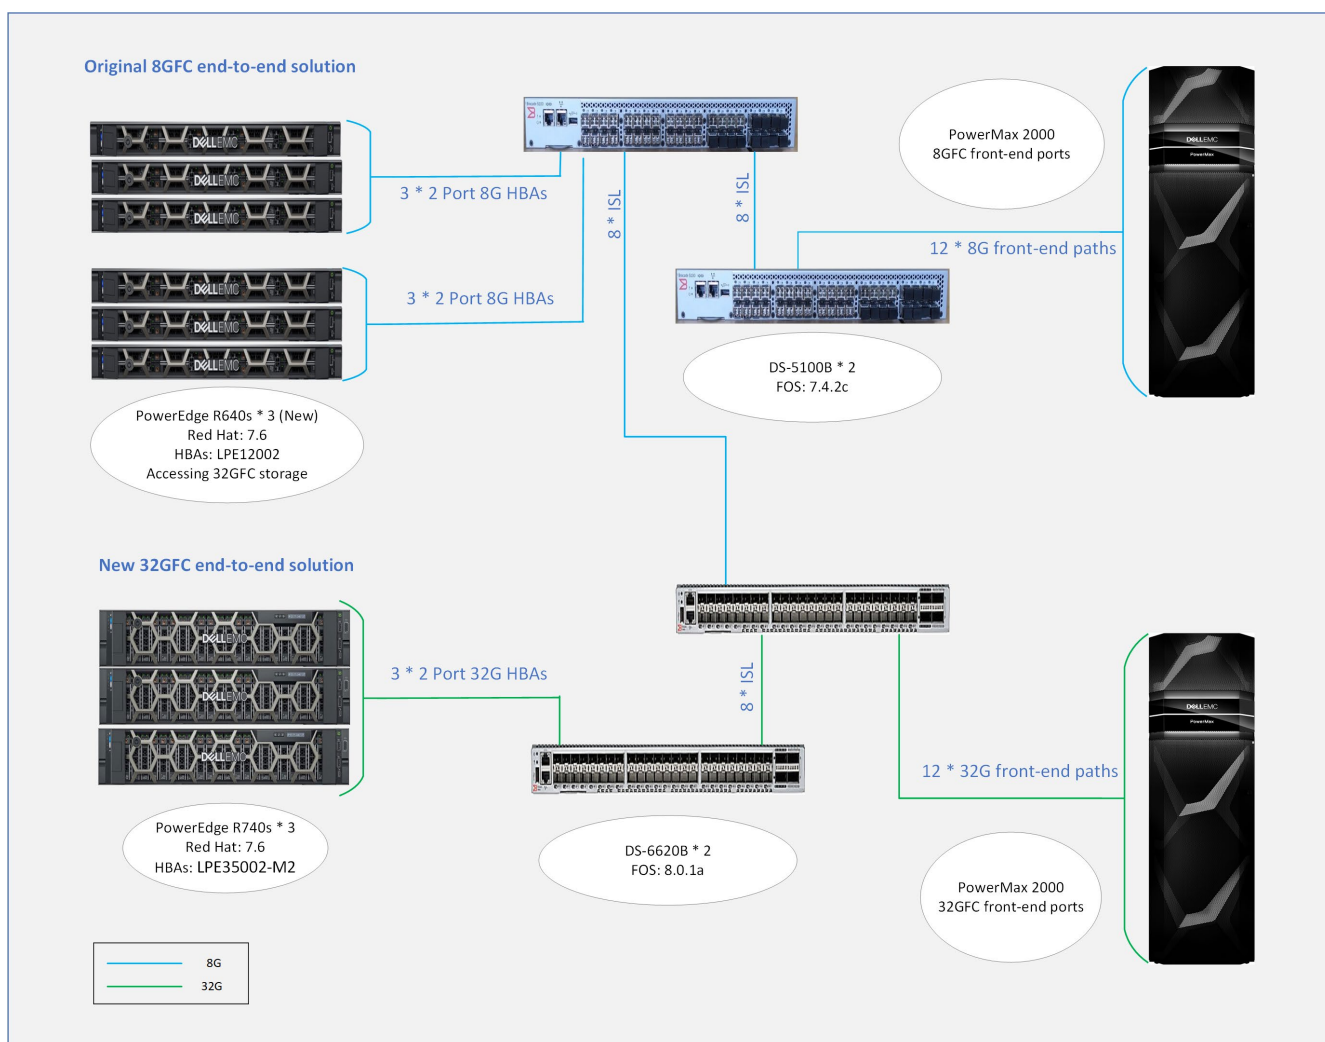


Figure 14 Combined solution: 8GFC and 32GFC servers and storage

3.6.2 Test procedure

1. The Vdbench workload script was run on the three 8G servers (hosts 4, 5, 6) that accessed storage on the 32GFC array. Performance data was captured.
2. While the 8G load was running, the same Vdbench workload script was run on the three 32G servers (hosts xyz, zxy, yzx) that accessed storage on the 32GFC array. Performance data was captured.
3. Initiator bandwidth limits were set on the 32G PowerMax array to constrain the amount of data being sent from the storage array to hosts 4, 5, 6 (the three 8G servers that accessed storage on the 32GFC array). Performance data was captured.

3.6.3 Initiator bandwidth limits

In the latest PowerMaxOS release (5978 Q2 2019 SR) for PowerMax, the ability to define initiator bandwidth limits was introduced.

To configure bandwidth limits for a host:

1. Launch Unisphere and select the appropriate PowerMax.
2. Select the **Hosts** tab.
3. Select the appropriate hosts.
4. Click the vertical three dots icon, as shown in Figure 15.
5. Choose **Set Bandwidth Limit**.
6. In the **Set Host BW Limit** dialog box, type a value for the Bandwidth Limit, and click **Apply**.

Figure 15 shows the steps required to set bandwidth limits on 32GFC PowerMax.

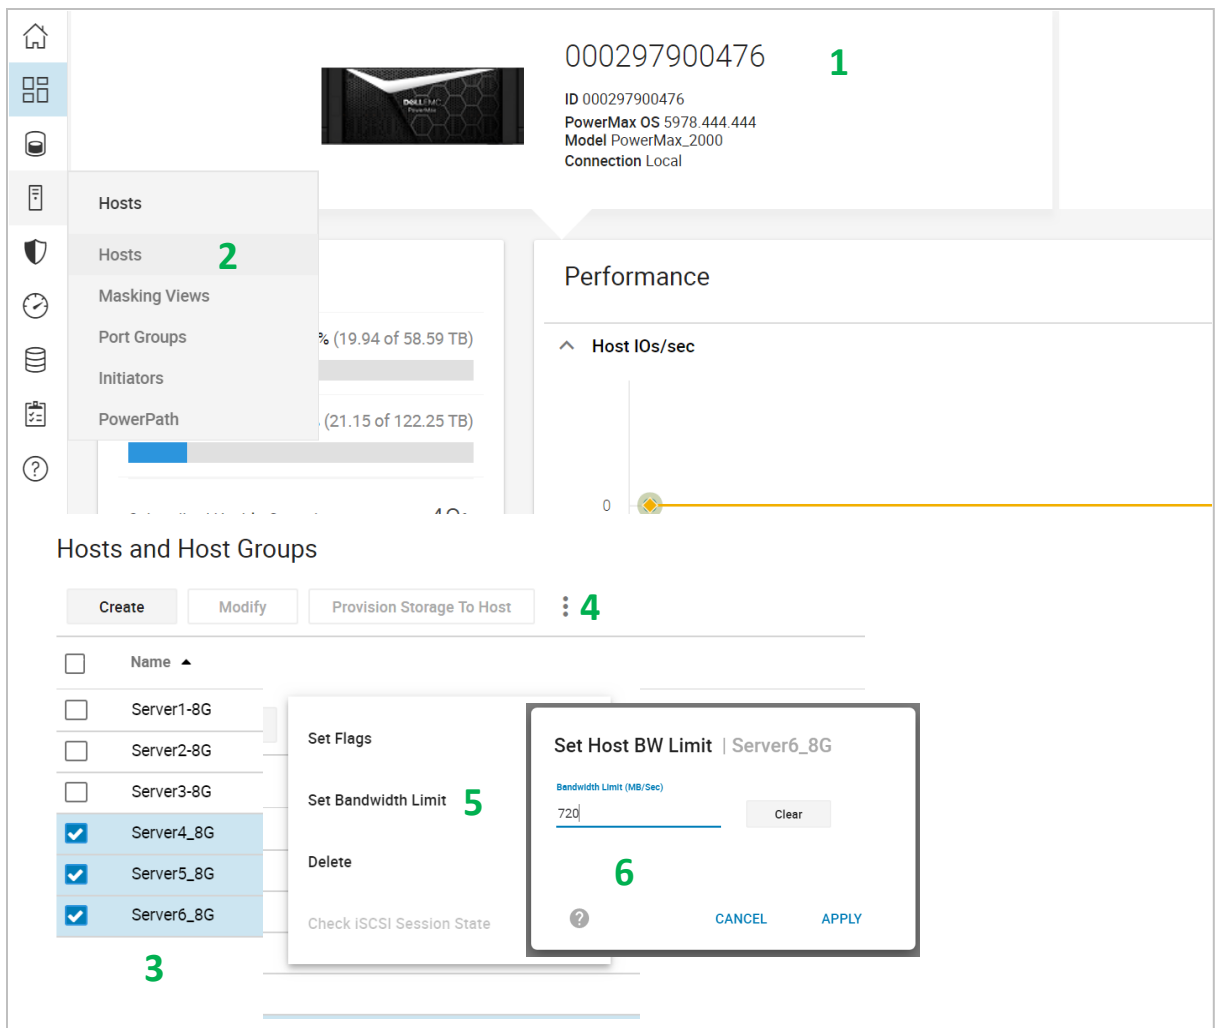


Figure 15 Configuration of bandwidth limits on 32GFC PowerMax

3.6.4 Performance results

Step 1: 8G hosts accessing the 32GFC array

Figure 16 records the performance throughput of the 8G servers (hosts 4, 5, 6) as they accessed the 32GFC storage while the Vdbench DSS script was run. The 8G servers achieved ~1600MBs/sec per server which is roughly the maximum throughput as it is going through an 8GFC SAN.




Figure 16 Combined solution 8GFC servers: Host MBs/sec

Step 2: 32G hosts accessing the 32GFC array

Figure 17 records the performance throughput of the 32G servers (hosts xyz, zxy, yzx) as they accessed the 32GFC storage while the Vdbench DSS script was run. As you can see, the 32G servers are not close to the maximum throughput of ~7000MBs/sec that we recorded for 32G servers in the 32GFC end-to-end solution. The performance data recorded in Figure 17 show the 32G server throughput at ~1700MBs/sec which indicates the presence of congestion spreading in the SAN.




Figure 17 Combined solution 32G servers: Host MBs/sec

Step 3: 32G hosts accessing the 32G array after introduction of initiator bandwidth limits

Once the initiator bandwidth limits were introduced (720 MB/s per port), the 32G hosts increased throughput to ~6000MBs/sec as illustrated in Figure 18.

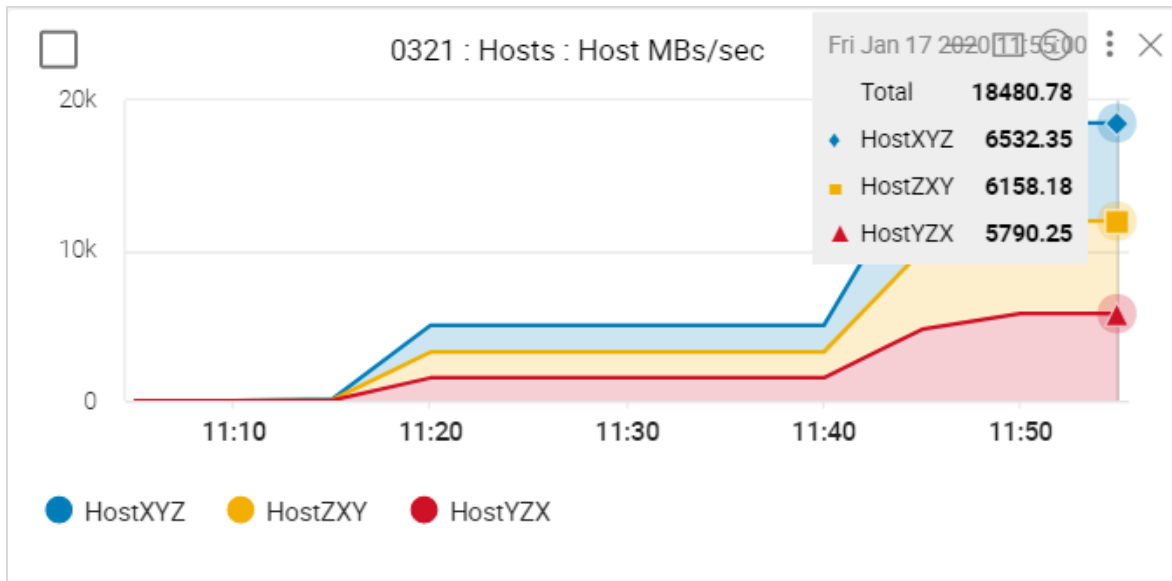


Figure 18 32G hosts after introduction of bandwidth limitats : Host MBs/sec

Conclusion

4.1 The path to 32GFC

The test results described in this paper demonstrate that matching the speed of your SAN to the speed of the storage array interface that is being attached to, allows you to maximize the value you get from your next generation storage system. If you do not match the speed of your SAN to the speed of your storage array, for example, connect a 32GFC storage interface to an 8G or 16G switch, 50-75% of the array's available bandwidth will be unusable.

Furthermore, Dell PowerMax provides Initiator Bandwidth Limits that will allow your 32GFC storage to achieve the maximum performance of the 32GFC connections in a mixed speed environment.

Table 3 Test comparisons: 8GFC, 16GFC and 32GFC

	8GFC SAN (uncongested)	8GFC SAN (congested)	16GFC SAN	32GFC SAN
Server 1 (MBs/sec)	1660	890	3535	7212
Server 2 (MBs/sec)	1560	792	3284	7031
Server 3 (MBs/sec)	1560	810	3283	7025
<i>Average (MBs/sec)</i>	<i>1593</i>	<i>830</i>	<i>3367</i>	<i>7089</i>

4.2 Additional benefits of 32GFC SAN Systems

There are also other valid reasons to upgrade the SAN from 8GFC and 16GFC to 32GFC. The technology included with the 32Gb/s Fibre Channel allows for the following enhancements:

- SAN-based Analytics
- Reporting for NVMe/FC
- Automatic detection of degraded SAN performance
- Increased scale for VMware environments
- Increased IOPs
- Increased Buffer Credits

Other pragmatic reasons to upgrade to 32GFC include:

- The 32GFC PowerMax IO Module provides increased performance for critical business workloads and it supports NVMe/FC for end-to-end NVMe performance.
- 8GFC and 16GFC products are transitioning to end-of-sales status, which means end-of-service status will soon follow.

A Configuration details

Table 4 Dell PowerMax - 01 SN 476 (8GFC front end)

Component	Details
PowerMax	2000
vBricks	2
Cache	2TB usable
Storage	34 * 2.84TB SSD
PowerMaxOS	5978.444.444
Front-end ports	16 * 8G ports (Note: all front-end ports were hardcoded to 8G via the binfile)

Table 5 Dell PowerMax - 02 SN 321 (32GFC front end)

Component	Details
PowerMax	2000
vBricks	2
Cache	2TB usable
Storage	34 * 2.84TB SSD
PowerMaxOS	5978.444.444
Front-end ports	16 * 32G Ports

Table 6 Dell Connectrix DS-5100B switch

Component	Details
Quantity	2 (core and edge)
Fabric OS	7.4.2c
Ports	40 * 8G ports

Table 7 Dell Connectrix DS-6620B switch

Component	Details
Quantity	2 (core and edge)
Fabric OS	8.0.1a
Ports	48 * 32G ports

Table 8 Dell PowerEdge R640 Servers

Component	Details
Purpose	8GFC Load testing
Quantity	6
Operating System	Red Hat 7.6
Applications	Vdbench
Host Bus Adapters	Emulex LPE12002 2 Port HBA (8G)

Table 9 Dell PowerEdge R740 Servers

Component	Details
Purpose	32GFC Load testing
Quantity	3
Operating System	Red Hat 7.6
Applications	Vdbench
Host Bus Adapters	LPE35002-M2 2Port HBA (32G)

B Workload generation script

Using the Vdworkbench DSS workload profile, the following script was used to generate workload for each test.

```
#List of servers
hd=default,shell=vdbench,user=root
hd=vm0,system=10.102.10.61
hd=vm1,system=10.102.10.62
hd=vm2,system=10.102.10.63
hd=vm3,system=10.102.10.64
hd=vm4,system=10.102.10.65
hd=vm5,system=10.102.10.66

#storage devices / List of devices used
sd=sd1,lun=/dev/emcpowera,size=450G,openflags=o_direct
sd=sd2,lun=/dev/emcpowerb,size=450G,openflags=o_direct
sd=sd3,lun=/dev/emcpowerc,size=450G,openflags=o_direct
sd=sd4,lun=/dev/emcpowerd,size=450G,openflags=o_direct
sd=sd5,lun=/dev/emcpowere,size=450G,openflags=o_direct
sd=sd6,lun=/dev/emcpowerf,size=450G,openflags=o_direct
sd=sd7,lun=/dev/emcpowerg,size=450G,openflags=o_direct
sd=sd8,lun=/dev/emcpowerh,size=450G,openflags=o_direct
sd=sd9,lun=/dev/emcpoweri,size=450G,openflags=o_direct
sd=sd10,lun=/dev/emcpowerj,size=450G,openflags=o_direct
sd=sd11,lun=/dev/emcpowerk,size=450G,openflags=o_direct
sd=sd12,lun=/dev/emcpowerl,size=450G,openflags=o_direct
sd=sd13,lun=/dev/emcpowerm,size=450G,openflags=o_direct
sd=sd14,lun=/dev/emcpowern,size=450G,openflags=o_direct
sd=sd15,lun=/dev/emcpowero,size=450G,openflags=o_direct
sd=sd16,lun=/dev/emcpowerp,size=450G,openflags=o_direct
sd=sd17,lun=/dev/emcpowerq,size=450G,openflags=o_direct
sd=sd18,lun=/dev/emcpowerr,size=450G,openflags=o_direct
sd=sd19,lun=/dev/emcpowers,size=450G,openflags=o_direct
sd=sd20,lun=/dev/emcpowert,size=450G,openflags=o_direct

#Workload : breakdown of profile using a workload skew of 90/10
wd=wd_RRH1,sd=sd*,seekpct=100,rhpct=100,rdpct=100,skew=16,range=(0,10),xfersize=128k
wd=wd_RRH2,sd=sd*,seekpct=100,rhpct=100,rdpct=100,skew=2,range=(10,100),xfersize=128k
wd=wd_RRM1,sd=sd*,seekpct=100,rdpct=100,skew=16,range=(0,10),xfersize=128k
wd=wd_RRM2,sd=sd*,seekpct=100,rdpct=100,skew=2,range=(10,100),xfersize=128k
wd=wd_RWM1,sd=sd*,seekpct=100,rdpct=0,skew=3,range=(0,10),xfersize=128k
wd=wd_RWM2,sd=sd*,seekpct=100,rdpct=0,skew=1,range=(10,100),xfersize=128k
wd=wd_SR1,sd=sd*,seekpct=0,rdpct=100,skew=43,range=(0,10),xfersize=256k
wd=wd_SR2,sd=sd*,seekpct=0,rdpct=100,skew=5,range=(10,100),xfersize=256k
wd=wd_SW1,sd=sd*,seekpct=0,rdpct=0,skew=11,range=(0,10),xfersize=256k
wd=wd_SW2,sd=sd*,seekpct=0,rdpct=0,skew=1,range=(10,100),xfersize=256k
```

Workload generation script

```
#run : 3000 IOPS a server with 4 threads  
rd=DSS,wd=wd_*,iorate=max, interval=1,elapsed=6000,threads=8,
```

C Technical support and resources

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

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell Technologies storage platforms.

C.1 Related resources

[Congestion Spreading and How to Avoid it \(White Paper\)](#)